Digital Comprehensive Summaries of Uppsala Dissertations
from the Faculty of Science and Technology 2596

Of Lattices and Lines

Building A Highly Efficient Wind Farm Flow Solver With The Lattice Boltzmann Method

HENRY KORB

ACTA UNIVERSITATIS
UPSALIENSIS
2025

ISSN 1651-6214
ISBN 978-91-513-2605-4
urn:nbn:se:uu:diva-568037

UPPSALA
UNIVERSITET

Dissertation presented at Uppsala University to be publicly examined in E22, Cramérgatan 5, Visby, Friday, 21 November 2025 at 10:00 for the degree of Doctor of Philosophy. The examination will be conducted in English. Faculty examiner: Assistant Professor Richard Stevens (University of Twente).

Abstract

Korb, H. 2025. Of Lattices and Lines. Building A Highly Efficient Wind Farm Flow Solver With The Lattice Boltzmann Method. *Digital Comprehensive Summaries of Uppsala Dissertations from the Faculty of Science and Technology* 2596. 92 pp. Uppsala: Uppsala University. ISBN 978-91-513-2605-4.

High-fidelity Large Eddy Simulation (LES) is the most common method for researchers to understand the highly complex interactions of wind farms and the atmospheric boundary layer (ABL), but is extremely expensive in terms of computational cost.

The lattice Boltzmann method (LBM) can reduce computational cost of wind farm simulations by several orders of magnitude by efficiently leveraging graphics processing units (GPUs). This thesis develops a software framework to simulate wind farms in the ABL using an GPU-resident LBM solver that is both accurate and computationally efficient.

First, theoretical and runtime requirements for such a framework are established. It is found that most of the required methodology is available and runtime requirements for many industry applications can be met with available methodology and hardware.

Then, a framework consisting of the GPU-resident LBM solver VirtualFluids and the turbine modelling tool WiFI is developed. The framework is validated against benchmark measurements and other LES solvers. Furthermore, the effect of resolution on accuracy and correction methods for rotor modelling at coarse resolution are studied. Finally, the model is extended to simulate thermally stratified boundary layers. The framework exhibits excellent computational efficiency and is as accurate as other LES solvers. By enabling simulations at coarse resolutions, computational performance can be increased even further.

A predecessor of the developed framework is used to demonstrate how the LBM enables novel and accelerates traditional research approaches. The ability to quickly generate large amounts of training data enables the use of data-driven methods, for example, a convolutional neural network to predict average flow quantities. Furthermore, larger parameter spaces can be examined, here, for example, of the helix approach for wind-farm flow control.

This thesis matures the LBM for wind energy from academic cases towards real-world applicability, enabling industrial use of LES and demonstrating its benefits for research. Further expansion of modelling capabilities to include complex terrain and forested conditions is needed as well as more extensive validation.

Keywords: Lattice Boltzmann Method, Wind Energy, Wind Farm Control

Henry Korb, Department of Earth Sciences, Wind Energy, Villav. 16, Uppsala University, SE-75236 Uppsala, Sweden. Department of Earth Sciences, LUVAL, Villav. 16, Uppsala University, SE-75236 Uppsala, Sweden.

© Henry Korb 2025

ISSN 1651-6214
ISBN 978-91-513-2605-4
URN urn:nbn:se:uu:diva-568037 (http://urn.kb.se/resolve?urn=urn:nbn:se:uu:diva-568037)

Technology is the name we give to stuff
that doesn't work properly yet
— Douglas Adams

List of papers

This thesis is based on the following papers, which are referred to in the text by their roman numerals.

I **H. Korb**, J. Bastin, H. Asmuth, and S. Ivanell. "The Lattice Boltzmann Method for Wind Farm Simulations: A Review". In: *Wind Energy Science Discussions* (2025), pp. 1–41. In Review.

II H. Asmuth, **H. Korb**, and S. Ivanell. "How Fast Is Fast Enough? Industry Perspectives on the Use of Large-eddy Simulation in Wind Energy". In: *Journal of Physics: Conference Series* 2505.1 (2023), p. 012001

III **H. Korb**, H. Asmuth, and S. Ivanell. "Validation of a Lattice Boltzmann Solver Against Wind Turbine Response and Wake Measurements". In: *Journal of Physics: Conference Series* 2505.1 (2023), p. 012008

IV **H. Korb** and E. Taschner. "Coarse, Fast, and Still Accurate? Comparing Corrections for the Actuator Line Model". In: *Journal of Physics: Conference Series* 3016.1 (2025), p. 012051

V **H. Korb**, H. Asmuth, M. Schönherr, M. Geier, and S. Ivanell. "Large Eddy Simulation of Thermally Stratified Atmospheric Boundary Layers with a Lattice Boltzmann Method". In: *Wind Energy Science Discussions* (2025). In Review.

VI H. Asmuth and **H. Korb**. "WakeNet 0.1 - A Simple Three-dimensional Wake Model Based on Convolutional Neural Networks". In: *Journal of Physics: Conference Series* 2265.2 (2022), p. 022066

VII **H. Korb**, H. Asmuth, and S. Ivanell. "The Characteristics of Helically Deflected Wind Turbine Wakes". In: *Journal of Fluid Mechanics* 965 (2023), A2

Reprints were made with permission from the publishers.

The work was divided among authors as follows:

Paper I – **Henry Korb (HK)** performed the literature review with support from Jean Bastin (JB) and input from Henrik Asmuth (HA). **HK** wrote the manuscript with inputs from JB, HA, and Stefan Ivanell (SI).

Paper II – HA conceptualised the survey with input from **HK** and SI. HA performed the processing of the results with input from **HK** and SI. **HK** setup and performed the comparison simulation. The manuscript was written by HA with input from **HK** and SI.

Paper III – **HK** performed the simulation and analysed the results. HA acquired the data from the benchmark. The manuscript was written by **HK** with input from HA and SI.

Paper IV – **HK** and Emanuel Taschner (ET) conceptualised the work. **HK** implemented the correction methods and performed the simulations. **HK** performed the post-processing and analysis with input from ET. The manuscript was written by **HK** with input from ET.

Paper V – Development of the boundary condition was performed by **HK** with input from HA and Martin Geier (MG). The implementation was performed by **HK** and HA with input from Martin Schönherr (MS). Simulations and analysis was performed by **HK** with input by HA. The manuscript was written by **HK** with input from HA, MG, MS and SI.

Paper VI – HA and **HK** conceptualised the work and implemented the model to equal parts. **HK** performed the simulations to generate training data. HA and **HK** analysed the data and build the processing pipeline. The manuscript was written by HA with inputs from **HK**.

Paper VII – The theoretical analysis was done by **HK**. The numerical study was designed by **HK** with input from HA and SI. **HK** performed the simulations and analysed the results with input from HA. The mauscript was written by **HK** with input from HA and SI.

Papers II, III, VI and VII were also included in my licentiate thesis "*The Lattice Boltzmann Method for Wind Farm Simulations: Validation and Application*".

Contents

1 Introduction 9
1.1 Current Challenges in Wind Energy Research 10
1.1.1 Wind-Farm Flows 10
1.1.2 Turbine Modelling 11
1.1.3 Wind-Farm Control 12
1.2 The Lattice Boltzmann Method and Wind Energy 12
1.2.1 The Lattice Boltzmann Method 12
1.2.2 Simulating Atmospheric Boundary Layers with the LBM 13
1.2.3 Wind Turbine Simulations with the LBM 14
1.2.4 Computational Performance 14
1.2.5 A LBM Framework for Wind Energy Simulations 14
1.3 Aims of This Work 16
1.4 Outline 16

2 Methodology 17
2.1 Requirements 17
2.2 Governing Equations 18
2.3 The Lattice Boltzmann Method 18
2.3.1 The Collision Operator 20
2.3.2 Refinement 22
2.3.3 Subgrid-Scale Model 22
2.3.4 Boundary Conditions 22
2.3.5 Advection-Diffusion 25
2.4 Turbine Modelling with the Actuator Line Method 25
2.5 Wind Turbine and Wind Farm Control 28
2.6 Implementation 29

3 Requirements for Industrial Applications 31
3.1 Survey Design 31
3.2 Survey Results 31
3.2.1 Background of Respondents 31
3.2.2 Current Modelling Practices 32
3.2.3 Requirements for the Use of LES 32
3.2.4 Comparison to GPU-resident LBM 33
3.3 Conclusion 34

4 Development and Validation of a Framework for Wind Energy Research Based on the LBM 35
4.1 Validation of the Framework for a Single Turbine 35
4.1.1 Setup 35
4.1.2 Results 36
4.1.3 Conclusion 39
4.2 Improving Rotor Modelling 40
4.2.1 Corrections for the ALM 40
4.2.2 LES Study 40
4.2.3 Conclusion 44
4.3 Simulating Thermally Stratified Boundary Layers 45
4.3.1 Extensions for a DDF-LBM Solver 45
4.3.2 Validation in a Conventionally Neutral Boundary Layer 47
4.4 Results of a Stably Stratified Boundary Layer 50
4.4.1 Conclusion 51

5 Applications 53
5.1 A Convolutional Neural Network for Wake Prediction 53
5.1.1 A Convolutional Neural Network to Predict Wind Turbine Wakes 54
5.1.2 An Open-Source Dataset for Training of Wind-Turbine Wakes 54
5.1.3 Results 55
5.1.4 Conclusion 58
5.2 Characterizing the Wake of the Helix Approach 59
5.2.1 Methodology 59
5.2.2 The Helical Wake of Single Turbine 60
5.2.3 Multiple Helices 64
5.2.4 Conclusion 67

6 Conclusion 68

7 Sammanfattning (Swedish Summary) 71

Acknowledgements 74

1. Introduction

As of 2024, wind energy supplies approximately 8 % of all electricity worldwide and over 17 % in the European Union, where it is the largest supplier of renewable energy [42]. However, in 2024, even a record growth in renewable energy production did not match the increase in electricity demand. There are many factors limiting the expansion of wind energy, however, one of the principal technical challenges is the high uncertainty of power and load predictions [168].

One of the reasons for the high uncertainty is that the interaction of large wind farms with the atmospheric boundary layer (ABL) remains insufficiently understood and difficult to model. The interactions occur on a wide range of scales, from large-scale weather patterns spanning hundreds of kilometres down to millimetres at the turbine blade, as illustrated schematically in Figure 1.1 [164]. Furthermore, wind turbines and wind farms are growing rapidly in size, resulting in new phenomena, such as wind-farm wakes and gravity waves, as well as new opportunities to improve wind-farm efficiency, such as wind-farm flow control. To develop a thorough understanding of wind-farm physics requires the use of computational fluid dynamics (CFD), with the highest fidelity achievable today being Large Eddy Simulation (LES). LES numerically solves the largest structures of the flow, requiring little parametrisation and giving a detailed, transient picture of the flow [125]. However, the computational cost of doing LES spanning the scales required for wind-farm flows is exceptionally large. Hence, its application in industry has been extremely limited, despite being a standard tool in research for over 20 years.

For the most part, reductions in the computational cost of LES have been driven by an increase in the computational power of central processing units (CPUs). However, the increase in CPU power has slowed down over the past two and a half decades [143]. Thus, researchers have opted for parallelisation to larger high performance computing (HPC) systems. Nevertheless, most solvers are based on discretisations of the incompressible Navier-Stokes equations, which are structurally ill-suited to parallelisation [89]. The lattice Boltzmann method (LBM) is an alternative approach to LES that is inherently well-suited to massive parallelisation and has been introduced to wind energy research in recent years [124, 16, 130]. The advantages of the LBM are even further amplified by the widespread introduction of graphics processing units (GPUs), which have been used for the LBM for over two decades [36]. However, applications of the LBM to wind energy research have so far been confined to relatively simplistic academic cases.

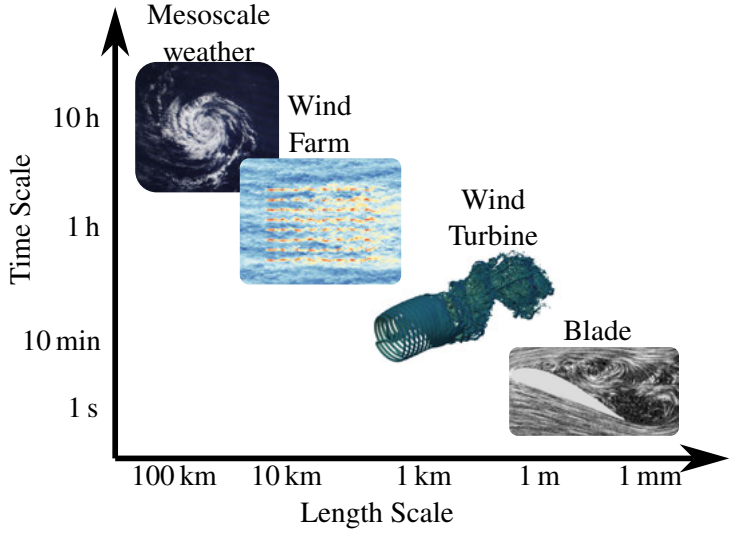

Figure 1.1. Different time and length scales of physical phenomena in wind-farm flows.

1.1 Current Challenges in Wind Energy Research

Several works have summarised the current challenges in wind energy, such as the *Long-Term Research Agenda* of the European Academy of Wind Energy (EAWE) by van Kuik and Peinke [149], and more recently, a series of publications entitled the *Grand Challenges* by Veers et al. [151], including a review of the impact of turbulence in the ABL on wind turbine performance and loads [77].

1.1.1 Wind-Farm Flows

Kosović et al. [77] note that as a result of the rapid increase in wind energy deployment, wind turbines are increasingly installed in locations where conditions are not well described by simple formulations. Such conditions include high or low levels of turbulence, thermal stratification, as well as wind shear and veer. Furthermore, due to their increase in size, wind turbines are now often situated outside of the surface layer, which typically spans the lowest 100 m of the atmosphere. Indeed, in highly stable boundary layers, wind turbines heights may even extend to the top of the ABL. Other modelling challenges mentioned by Kosović et al. include topography, low-level jets and gravity waves.

As the wind passes a wind turbine, it loses some of its kinetic energy and becomes more turbulent, creating the wake. Wakes have been one of the main areas of research in wind energy and much is known for simple conditions.

However, effects of thermal stratification, shear, veer and the interaction of multiple wakes are still active areas of research [119].

Another challenge that arises with increasing size and number of wind farms is the phenomenon of wind-farm wakes. Wind-farm wakes behave fundamentally differently from wind turbine wakes [84] and can span up to 60 km downstream of a wind farm, impinging on downstream wind farms, resulting in so-called wind theft [39].

All these challenges warrant the use of high-fidelity simulation tools, capable of including complex boundary conditions and efficiently simulating very large domains, see, for example, Ouro et al. [113] for a discussion. But even the highest fidelity simulation tools have uncertainties because the inputs to the model themselves are uncertain and the modelled system is chaotic. Understanding model sensitivities has been attempted in a number of studies, for example, Andersen et al. [4] and Andersen et al. [5], yet, a proper uncertainty quantification has been limited to lower fidelity models, for example Zhang and Zhao [169].

1.1.2 Turbine Modelling

As the flow becomes more complex, so does its effect on turbine loads. Furthermore, blades are becoming longer and more flexible, exhibiting considerable twisting. Accurate prediction of damage equivalent loads thus requires not only high-fidelity flow models but also the coupling to accurate models of the turbine structure, as was done for example in Hodgson et al. [58], Taschner et al. [142], and Trigaux et al. [145]. Trigaux et al. highlight the need to include blade twist for accurate prediction of loading at the tip, which had been disregarded in many structural models in the past.

In LES, turbines are usually modelled using actuator methods that impose body forces on the flow . The actuator disc and line are the most commonly used, either representing the turbine as a disc or using lines. The actuator line method (ALM) requires a smaller time-step size and higher grid resolution, but also includes more realistic physics. Hence, the ALM is required for a wide range of applications, such as aeroelastic simulations mentioned above or many wind-farm control approaches discussed in subsection 1.1.3. Since its introduction by Sørensen and Shen [133], many improvements and corrections have been proposed. Work has been conducted examining the velocity sampling location and averaging [166], velocity correction [146], and force spreading [64, 63]. In particular, the width of the force smearing has been widely debated, since it affects a variety of metrics. The most common used recommendation follows Troldborg [147], while Meyer Forsting and Troldborg [103] and Martínez-Tossas et al. [95] give essentially contradicting recommendations based on the quantity of interest to optimise. Furthermore, it was recognised by Dağ [32] and later proven by Caprace et al. [28], Meyer Forsting

et al. [101], and Martínez-Tossas and Meneveau [98], that the ALM leads to reduced induction in areas of high vorticity gradients, i.e. at the root and tip of the blade. Numerous corrections for this error have been proposed, including Meyer Forsting et al. [101], Martínez-Tossas et al. [99], Kleine et al. [68], and Dağ and Sørensen [33].

1.1.3 Wind-Farm Control

The turbines and the wind in a wind farm interact. As laid out by Meyers et al. [104], controlling turbines to delibaretely modify the flow can allow to improve a variety of metrics of a wind farm, first and foremost power production [161]. In recent years, there has been considerable attention on various wind-farm flow control paradigms, most of them either focusing on wake steering or individual-pitch control, see for example Brown et al. [26] and Taschner et al. [141]. However, assessing the effect on power production and loads by these approaches also requires high-fidelity simulation and most simulations of these approaches have been limited to one or two turbines.

Another area of increasing attention is the provision of ancillary services, that is, using wind farms to stabilise the electrical grid. The ability of wind energy to provide such ancillary services, such as frequency stabilisation, becomes critical as wind and solar energy gain a large share in electrical grids [110]. Safely providing such services, such as a power reserve, also requires an accurate estimation of available power, as described, for example, in Göçmen et al. [51].

A key metric for applying LES in industrial control applications is the real-time factor, measuring the ratio of simulated to real time. A real-time factor above 1 implies a simulation progressing faster than the real world, thus making it applicable to short-term forecasting. Traditional LES methods can perform real-time simulations only at extremely coarse resolutions, as demonstrated in Bauweraerts and Meyers [20].

1.2 The Lattice Boltzmann Method and Wind Energy

Perhaps the largest barrier to a wider application of LES is its immense computational cost. Yet, LES is required for a wide variety of applications in academia and industry, as the previous section demonstrates. One of the most promising candidates to reducing that computational cost is the LBM utilizing GPUs.

1.2.1 The Lattice Boltzmann Method

The LBM was developed from Lattice Gas Automata in the late 1980s [100, 57]. A comprehensive introduction can be found in Krüger et al. [79]. The

LBM consists of a two-step algorithm, a collision and a streaming step. The streaming step is non-local but linear, while the collision is potentially non-linear but local. It is this separation of non-locality and non-linearity that makes the LBM highly parallelisable. Much research on the LBM has focused on the collision step, since it is crucial for the accuracy and stability of the method. Early collision operators were unstable at high Reynolds numbers, but a new generation of operators, such as the hybrid regularised recursive [59] or the cumulant operator [50] have mitigated that issue. A comparison can be found in Coreixas et al. [30].

Other important improvements of the LBM include the development of refinement algorithms [38, 45, 81] and boundary conditions for curved boundaries [25, 50]. Additionally, many algorithmic improvements specifically for the use of GPUs have been developed, such as the improved streaming patterns *esoteric twist* [49], *esoteric pull*, and *esoteric push* [86] or indirect addressing [128, ch. 6]. Furthermore, the use of multiple GPUs has been improved by the development of communication hiding techniques [165].

1.2.2 Simulating Atmospheric Boundary Layers with the LBM

Simulations of the ABL feature very high Reynolds numbers and typically cover large domains. The high Reynolds numbers make direct numerical simulation (DNS) unattainable for the foreseeable future. Instead LES is currently the highest fidelity possible. Malaspinas and Sagaut [92] demonstrated that eddy viscosity models can readily be used with the LBM, enabling LES. That the LBM is suitable for the simulation of very large domains was demonstrated in various cases, many of them from research on urban boundary layers, such as the simulation of $10\,\mathrm{km} \times 10\,\mathrm{km}$ of Tokyo's inner city at $1\,\mathrm{m}$ resolution presented by Onodera et al. [112]. Other studies featuring large domains include the study Watanabe et al. [158] of flows above forests, Lenz et al. [87] of the flow through Basel and Jacob et al. [59].

However, these studies do not include wall models, which are commonly used in ABL flows to allow coarser resolutions near the wall. Malaspinas and Sagaut [93] proposed a first approach to include wall models in the LBM, but the approach requires the modelling of turbulent viscosity at the wall in addition to the velocity profile. An improved version of it was employed in the solver ProLB and used to simulate atmospheric boundary layers in Feng et al. [37]. Another approach was developed by Asmuth et al. [9], which does not require specification of a viscosity and agreed very well with reference data in simulations of pressure driven atmospheric boundary layers.

1.2.3 Wind Turbine Simulations with the LBM

Simulations of wind turbines were first focused on rotor resolved simulations, such as Perot et al. [115] and Xu [167]. The first ALM simulation using the LBM was presented in Rullaud et al. [124], albeit a 2D simulation of a vertical axis wind turbine. Asmuth et al. [16] features the first 3D simulation of a horizontal axis wind turbine, simulating a single turbine using the ALM and the GPU-resident LBM solver ELBE in laminar inflow, demonstrating the high accuracy and computational efficiency of the method. The follow-up studies examined many fundamental aspects of LBM-ALM simulations and featured turbulent inflows [15, 10]. Again, results showed excellent agreement with reference data. The methodology was compared to traditional LES solvers and measurements in Asmuth et al. [14] and was shown to be as accurate as the other solvers. Schottenhamml et al. [130] also presented an implementation of the ALM using the solver WALBERLA and compared results and computational performance to SOWFA, a widely used LES tool and also found very good agreement with great reductions in computational cost. Further validation and scaling measurements were presented in Schottenhamml et al. [129]. Watanabe and Hu [157] presented simulations of a wind farm of 8 turbines at 80 nodes per diameter using up to 40 GPUs, again demonstrating the excellent scalability of the LBM to large simulations domains. Very recently, Wang et al. [156] presented simulations of a farm of model turbines in a thermally stratified wind tunnel using PROLB [37].

1.2.4 Computational Performance

One of the main motivations for the LBM is its supposed computational efficiency. In the LBM, computational performance is usually measured in million node updates per second (MNUPS).

The evolution of the computational performance is shown in Figure 1.2. The graph shows an exponential increase, with performance doubling every 2.5 years. Currently, the highest performance reported in the literature is about 3000 MNUPS. The LBM is also especially well suited for very large domains, as the memory available per GPU is increasing drastically. Comparisons to Navier-Stokes based solvers come with many caveats, but a number of studies have tried to make comparisons in terms of time-to-solution, including Asmuth et al. [15] and Schottenhamml et al. [130]. Both studies find reductions in cost by three orders of magnitude.

1.2.5 A LBM Framework for Wind Energy Simulations

Most of the aforementioned studies only study simplified cases, either lacking wall models, turbine control, or thermal stratification. Combining all available models in a single framework requires much development and validation ef-

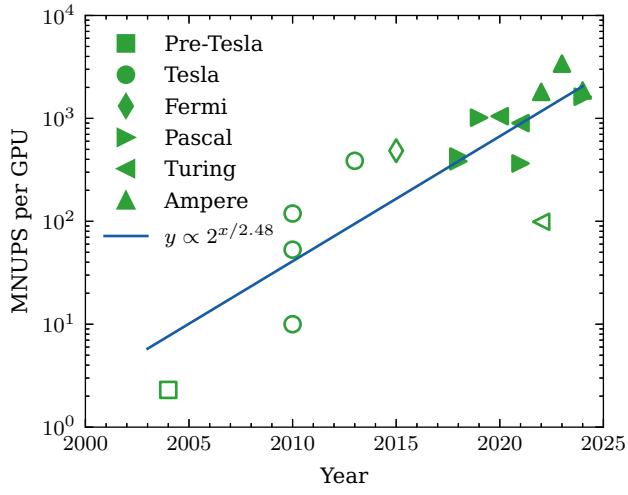

Figure 1.2. Computational performance per GPU over time. Unfilled markers use D3Q19 lattice, filled markers D3Q27. Different markers represent different generations of GPU architecture. Figure adapted from paper I.

fort. The wind energy community has long relied on only a few open-source frameworks for simulations, with the most commonly used being SOWFA [108] and PALM [94], while closed-source in-house codes are still the norm. Only a single widely applied framework exists for aeroelastic simulations, the OpenFAST toolbox [66]. However, the wide adoption of SOWFA despite its limitations makes a strong argument for the willingness of the wind energy community to use open-source solvers. There are also numerous other arguments for using open-source research software, outlined in Barker et al. [18], such as ensuring reproducibility, improved productivity, and higher software quality. Recently, a number of new open-source solvers have been presented in the literature, among them Sauer and Muñoz-Esparza [127], Stipa et al. [138], and Kuhn et al. [80]. However, no open-source framework for LES based on the LBM for wind energy is available yet.

1.3 Aims of This Work

The objective of this thesis to develop a framework for conducting LES of large wind farms based on the LBM utilizing GPUs. This framework should be able to conduct simulations at very high computational efficiency, making LES suitable for industrial application; model turbines with high fidelity; and conduct simulations of large numbers of turbine and in complex conditions such as thermally stratified boundary layers. Furthermore, the framework will be applied to varying test cases for validation and further understanding of wind-farm flows with particular attention to wind-farm control and machine learning. This work is divided into three main components:

1. Establish requirements for industrial application and review available methodology
2. Develop, implement, and validate the framework
3. Apply the framework to wind-farm control and training of machine learning models

1.4 Outline

This thesis is focused on the development of the lattice Boltzmann method for wind-energy applications. It first develops a more thorough understanding of requirements and already existing capabilities. Furthermore, this work develops and validates a framework for LES of wind turbines with the LBM, that is then improved and extended to more realistic cases. Finally, the method is applied to current challenges in wind energy research.

Chapter 2 is largely based on paper I and introduces the methodology common to this work, that is, the lattice Boltzmann method and the actuator line method. Furthermore, it discusses some aspects of the developed framework. Chapter 3 summarizes paper II, establishing criteria for industrial use. Chapter 4 presents the developments of the framework, including validation of the basic model (paper III), improvements to rotor modelling (paper IV) and developing the methodology for simulation thermally stratified boundary layers (paper V). The applications to wind-farm control and machine learning from papers paper VI and paper VII are presented in chapter 5. Finally, general conclusions are presented in chapter 6.

2. Methodology

This chapter introduces the common methodology used throughout this work. It also includes a general review of the LBM for wind-farm flows that was performed in paper I.

2.1 Requirements

Wind turbines are situated in the atmospheric boundary layer. We can formulate criteria for a wind-farm flow solver based on the following considerations. The flow regime of the ABL is characterised by velocities on the order of $V_0 = \mathrm{O}(10\,\mathrm{m/s})$, while the speed of sound of air $c_\mathrm{s} \approx 343\,\mathrm{m/s}$, hence the Mach number $\mathrm{Ma} = \frac{V_0}{c_\mathrm{s}} \ll 0.1$ and no compressibility effects need to be taken into account. However, the relevant length scales are $L = \mathrm{O}(10\,\mathrm{m}) - \mathrm{O}\left(10^3\,\mathrm{m}\right)$ and the kinematic viscosity of air is $\nu = \mathrm{O}\left(10^{-5}\,\mathrm{m^2/s}\right)$, resulting in Reynolds numbers $\mathrm{Re} = \frac{V_0 L}{\nu} \approx 10^7$–$10^9$. Thus, the model must remain stable at such high Reynolds numbers. The importance of subgrid-scale (SGS) models to correctly model the ABL is highlighted by, for example, Gadde et al. [43]. Sufficiently resolving the length scales near the wall renders ABL simulation too computationally expensive since length scales decrease with proximity to the wall. Employing a wall model significantly reduces computational cost. Similarly, resolving the boundary layer on the blades also requires very high resolution. Modelling the turbines via actuator models significantly reduces computational cost and complexity of the simulation, as actuator models eliminate the need to simulate moving solid bodies. Lastly, the conditions in which turbines are placed need to be accounted for. Turbines are often placed in hilly or forested area, which significantly effects the flow conditions [7]. Furthermore, thermal stratification has a variety of effects on wind-farm flows, as outlined in subsection 1.1.1.

Hence, we identify the following criteria:

- Low Mach number
- Stability at high Reynolds number
- Advanced subgrid models
- Wall modelling
- Turbine modelling
- Complex terrain and forested areas
- Thermal stratification

2.2 Governing Equations

The flow in the ABL is adequately described by the incompressible Navier-Stokes equations, optionally including Coriolis, buoyancy and turbine forces [140]. This system of equations has to be solved numerically, in this work using LBM-LES. In LES, the filtered equations are solved [139]:

$$\frac{\partial u_i}{\partial x_i} = 0 \tag{2.1}$$

$$\frac{\partial u_i}{\partial t} + u_j \frac{\partial u_i}{\partial\, x_j} = -\frac{1}{\rho}\frac{\partial p}{\partial x_i} + \frac{\partial}{\partial x_i}\left(\nu_\mathrm{e}\frac{\partial u_j}{\partial x_j}\right) + F_i^\mathrm{D} + F_i^\mathrm{B} + F_i^\mathrm{T}. \tag{2.2}$$

The spatial coordinates are $\boldsymbol{x} = [x, y, z]^T = [x_1, x_2, x_3]^T$. The density and filtered velocity of the fluid are denoted as ρ and $\boldsymbol{u} = [u_1, u_2, u_3]^T = [u, v, w]^T$, respectively. The deviation from background pressure is p, the effective kinematic viscosity is ν_e, F_i^B and F_i^T are buoyancy and turbine forces. Einstein summation applies to all Latin, but not Greek, indices. F_i^D is the driving force, either equal to a constant pressure gradient

$$F_i^\mathrm{D} = \frac{\partial p_\infty}{\partial x_i} \tag{2.3}$$

or the Coriolis force in geostrophic balance

$$F_i^\mathrm{D} = -\epsilon_{ij3}(G_j - u_j) f_\mathrm{C}, \tag{2.4}$$

where f_C is the Coriolis parameter that depends on the latitude, $\boldsymbol{G}$ is the geostrophic wind, and the Levi-Civita symbol is denoted as ϵ_{ijk}.

If thermal stratification is considered, an advection diffusion equation for the potential temperature θ also has to be solved and coupled to the Navier-Stokes equations via the Boussinesq approximation:

$$\frac{\partial \theta}{\partial t} + u_j \frac{\partial \theta}{\partial x_j} = \frac{\partial}{\partial x_j}\left(D_\mathrm{e}\frac{\partial \theta}{\partial x_j}\right) \tag{2.5}$$

$$F_i^\mathrm{B} = g\frac{\theta - \langle\theta\rangle(z)}{\theta_\mathrm{r}}\delta_{i3}. \tag{2.6}$$

The effective diffusivity is denoted as D_e, g is the gravitational acceleration, θ_r is a reference temperature, $\langle\cdot\rangle$ denotes horizontal averaging, and δ_{ij} is the Kronecker delta. In this work, the turbine forces are computed from the actuator line method, which will be discussed further in section 2.4.

2.3 The Lattice Boltzmann Method

The LBM solves the Navier-Stokes equations numerically by solving a transport equation for populations $f_{\iota\kappa\lambda}$ of particles with discrete velocities $\boldsymbol{c}_{\iota\kappa\lambda}$ on a cubic

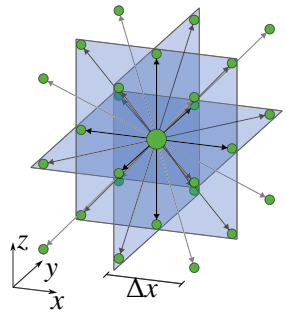

Figure 2.1. Schematic of a D3Q27 lattice. Green dots represent grid nodes, while black arrows represent velocity directions.

grid of nodes with spacing Δx and constant time-step size Δt [79]. Velocities, spacing and time-step size are related by $\boldsymbol{c}_{\iota\kappa\lambda} = c[\iota,\kappa,\lambda]^T$ with $c = \frac{\Delta x}{\Delta t}$. The set of velocities $c_{\iota\kappa\lambda}$ forms the lattice giving the lattice Boltzmann method its name. Throughout this work we use a so-called D3Q27 lattice, meaning a lattice of 27 velocities in 3 dimensions. Kang and Hassan [67] and White and Chong [159] find that the D3Q19 lattice is sufficient for turbulent flows if the flow is aligned with the lattice directions, in other cases, a D3Q27 lattice is recommended. This was confirmed for the case of wind turbine wakes in Asmuth et al. [15]. The D3Q27 lattice is constructed from the set of $\iota,\kappa,\lambda \in \{-1,0,1\}$. Figure 2.1 gives a visualization of the lattice. In addition to the velocity set, the lattice also defines a set of weights $w_{\iota\kappa\lambda}$ and a lattice speed of sound c_s, that is $\frac{1}{\sqrt{3}}$ for most lattices, including the D3Q27 lattice. The LBM can be derived by discretizing the Boltzmann equation for the particle density function, hence its name. However, the LBM does not solve the Boltzmann equation, only the Navier-Stokes equations. The LBM consists of a two-step algorithm:

$$f^*_{\iota\kappa\lambda}(\boldsymbol{x},t) = f_{\iota\kappa\lambda}(\boldsymbol{x},t) + \Omega_{\iota\kappa\lambda}(\boldsymbol{x},t) \tag{2.7}$$

$$f_{\iota\kappa\lambda}(\boldsymbol{x}+\boldsymbol{c}_{\iota\kappa\lambda}\Delta t, t+\Delta t) = f^*_{\iota\kappa\lambda}(\boldsymbol{x},t). \tag{2.8}$$

In the collision step (2.7), the collision operator Ω and the post-collision populations $f^*_{\iota\kappa\lambda}(\boldsymbol{x},t)$ are computed. The collision operator can be very complex and highly non-linear, however, it only relies on the local populations. In the second step, the streaming step, the post-collision populations are streamed to the neighbouring nodes, making this step linear but non-local. The moments of the populations are defined as

$$m_{\alpha\beta\gamma} = \sum_{\iota\kappa\lambda} (c\iota)^\alpha (c\kappa)^\beta (c\lambda)^\gamma f_{\iota\kappa\lambda}, \tag{2.9}$$

where we call the sum of α, β and γ the order of the moment. The moments are related to macroscopic quantities, such as density and velocity by

$$\rho = m_{000} = \sum f_{\iota\kappa\lambda} \tag{2.10}$$

$$u = \frac{m_{100}}{\rho} + \frac{F_x}{2}, v = \frac{m_{010}}{\rho} + \frac{F_y}{2}, w = \frac{m_{001}}{\rho} + \frac{F_z}{2}, \tag{2.11}$$

where $\boldsymbol{F}$ is the total force density. Higher order moments are related to the stress tensor and the internal energy. The work of Dubois [35] shows that Taylor expanding (2.7) and (2.8) to second order and taking the zeroth and first order moments yields the incompressible Navier-Stokes equations, given a small number of requirements. A third order error is proportional to the cube of the Mach number

$$\mathrm{Ma} = \frac{V_0}{c_\mathrm{s}} = \frac{\sqrt{3}V_0\Delta t}{\Delta x}. \tag{2.12}$$

That is why the LBM is often said to be weakly compressible, since the error is small only in the low Mach number regime, usually $\mathrm{Ma} < 0.1$. On the other hand, artificially increasing the Mach number increases the time-step size and thus reduces computational cost.

2.3.1 The Collision Operator

The accuracy of the LBM is governed by the choice of collision operator Ω. A collision operator for the Navier-Stokes equations needs to satisfy certain requirements, as detailed, for example, in Krüger et al. [79], such as conserving the zeroth and first order moments. Therefore, a large number of collision operators has been proposed and it is still an active area of research.

The simplest collision operator is the Bhatnagar-Gross-Krook (BGK) or single-relaxation-time operator. The BGK operator relaxes all populations towards their equilibrium with a common rate and actually predates the LBM as it was developed to analyze the Boltzmann equation in Bhatnagar et al. [23]. Building on the idea that collision is a relaxation process, three branches of collision operators have been developed. The first branch is that of multi-relaxation time (MRT) operators, that relax moments of the populations at different rates. MRT operators usually work in three steps:

1. transform populations to moments
2. relax moments
3. transform moments to populations

Popular methods include the original MRT [31], the cascaded or central moment [46] and the cumulant [50] operators. Early MRT operators suffered from low numerical stability, violation of Galilean invariance and many other issues, see for example the results from Freitas et al. [41]. Geier et al. argue that relaxing statistically dependent quantities at different rates is invalid. Cu-

mulants are by construction statistically independent, furthermore the Galilean invariance of the collision process is improved significantly.

Another widely used class are the regularized operators. By utilizing a relationship between the stress tensor and the populations, populations are regularized before collision. The approach was first introduced by Latt and Chopard [85] and further developed in Malaspinas [91]. Jacob et al. [59] proposed to combine the regularization with a finite-difference method to compute the stress tensor, yielding the hybrid regularized recursive (HRR) method.

Finally, the entropic collision operators enforce the H-theorem and are widely applied to supersonic hydrodynamic flows but in general more complex and slower than the MRT and regularized operators [30].

Spinelli et al. [136] compare different collision operators (MRT, HRR and cumulant) and SGS models and find that the cumulant method consistently yields the best results, even without the use of a SGS model. However, Gehrke et al. [44] find that, at coarse resolution, an explicit SGS model is required also with the cumulant operator.

Throughout this thesis we employ the cumulant collision operator, therefore we describe it in more detail here. Cumulants are defined from the cumulant generating function

$$c_{\alpha\beta\gamma} = c^{-\alpha-\beta-\gamma} \frac{\partial^\alpha \partial^\beta \partial^\gamma}{\partial \Xi^\alpha \partial \Upsilon^\beta \partial Z^\gamma} \ln\left(F\left(\boldsymbol{\Xi}\right)\right)\Big|_{\Xi=\Upsilon=Z=0}, \tag{2.13}$$

where $F(\boldsymbol{\Xi})$ is the Laplace transformed population

$$F(\boldsymbol{\Xi}) = \mathcal{L}\left\{\sum_{\iota\kappa\lambda} f_{\iota\kappa\lambda}\delta(\iota c-\xi)\delta(\kappa c-\upsilon)\delta(\lambda c-\zeta)\right\}, \tag{2.14}$$

with $\delta(x)$ being the Dirac delta function. Cumulants of second order and higher are relaxed individually by

$$c^*_{\alpha\beta\gamma} = \omega_{\alpha\beta\gamma} c^{\mathrm{eq}}_{\alpha\beta\gamma} + (1+\omega_{\alpha\beta\gamma}) c_{\alpha\beta\gamma}, \tag{2.15}$$

where $\omega_{\alpha\beta\gamma}$ and $c^{\mathrm{eq}}_{\alpha\beta\gamma}$ are the relaxation rate and equilibrium of the respective cumulant. See Geier et al. [50] for further details. The relaxation rate of the second order cumulants is related to the viscosity by

$$\frac{1}{\omega_{(110)}} = \frac{\nu}{c_{\mathrm{s}}^2 \Delta t} + \frac{1}{2}, \tag{2.16}$$

where parentheses denote all permutations of the indices. The accuracy of the method can be improved by choosing specific relaxation rates, as described in Geier et al. [48], eliminating leading order errors in the diffusion terms, leading to fourth order accuracy in diffusion.

2.3.2 Refinement

The LBM is restricted to cubic grids with constant grid size due to the coupling of velocity, time and space discretization. However, in many flows some areas are of higher interest or require higher resolution due to smaller dominating scales, such as near the ground or in the wind-turbine wake, in the case of wind-farm simulation. To remedy this shortcoming, multi-level refinement was introduced to the LBM by Filippova and Hänel [38]. A more accurate approach was introduced by Geier et al. [45], making use of the entries of the stress tensor available locally. A refined version is presented in Kutscher et al. [81], which is employed in a number of studies in this thesis.

2.3.3 Subgrid-Scale Model

As LES does not resolve all scales, the effect of sub-grid scales needs to be accounted for. Malaspinas and Sagaut [92] showed that SGS models can be applied in the LBM using (2.16). The original SGS model by Smagorinsky [131] is still widely used, despite well-known drawbacks. The QR model, on the other hand, is the first of a new class of SGS models based on providing the appropriate amount of dissipation [152]. The QR model only relies on invariants of the stress-tensor, which can be computed directly from the populations, making the QR model ideally suited for the LBM. It was generalized to the anistropic minimum dissipation (AMD) for anisotropic grids by Rozema et al. [123], note, however, that the LBM employs isotropic grids. The AMD model was also extended to account for thermal stratification by Abkar et al. [1]. In this thesis, most studies employ the QR model, while papers VI and VII employ the standard AMD model and paper V employs the extended AMD model.

2.3.4 Boundary Conditions

Since the primary variable of the LBM are the populations, boundary conditions also have to be formulated in terms of $f_{\iota\kappa\lambda}$. This opens up many ways to prescribe the same macroscopic boundary condition.

Bounce-Back Boundary Conditions

This work only employs one class of boundary conditions, bounce-back conditions. Here, the wall is assumed to lie halfway between two nodes and populations reflect (i.e. bounce back). The bounce-back approach can be used to define no-slip and velocity boundary conditions [79, ch. 5]:

$$f_{\overline{\iota\kappa\lambda}}(\boldsymbol{x}, t+\Delta t) = f^{*}_{\iota\kappa\lambda}(\boldsymbol{x}, t) - 2\rho^{\mathrm{w}} \frac{w_{\iota\kappa\lambda}}{c_{\mathrm{s}}^2} u_i^{\mathrm{w}} c_{\iota\kappa\lambda,i}, \tag{2.17}$$

where quantities at the wall are denoted with the superscript w. By setting $\boldsymbol{u}^{\mathrm{w}} = \boldsymbol{u}(\boldsymbol{x}, t)$ it is also possible impose a zero-stress boundary condition.

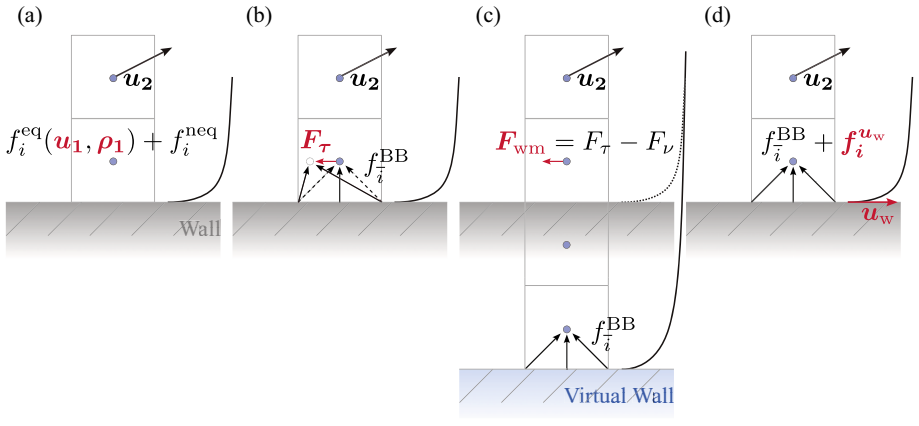

Figure 2.2. Illustration of different approaches of wall modelling. Red symbols indicate inputs to the model. (a) Population reconstruction approach, (b) wall function bounce, (c) immersed virtual wall method, (d) wall velocity approaches

Wall Modelling

Near the ground, turbulent scales decrease in size and sufficiently resolving the flow becomes prohibitively expensive. To that end, wall models, in the case of the ABL Monin-Obukhov similarity theory (MOST)[111], can be employed to predict a wall shear stress that has to be applied at the boundary. Wall modelling was only introduced rather recently in the LBM by Malaspinas and Sagaut [93]. We classify 4 approaches for wall modelling, which are schematically shown in Figure 2.2.

Direct population reconstruction methods, like the model by Malaspinas and Sagaut, compute the entire set of populations at the wall from macroscopic quantities. The model was extended to curved boundaries in Wilhelm et al. [160]. A similar approach is proposed in Haussmann et al. [56].

The *wall-function bounce* approach by Han et al. [54] alters the classical bounce-forward approach for slip boundary conditions. The effect of the wall model is taken into account via a drag term at the first node off the wall. Since the drag term is evaluated directly from the shear stress, no viscosity is required, however, the model is originally defined in D3Q19 and an extension to D3Q27 or curved boundaries is not straightforward, if possible at all.

The *immersed virtual wall* approach was introduced by Kuwata and Suga [82] and also applies a drag force on the node. However, this force is applied at the first node at an virtual wall location. A virtual fluid layer is located below the virtual wall, thus extending the domain. This makes the model easily applicable at curved boundaries. However, the wall shear stress is applied at the node, not at the location of the virtual wall.

We termed the final class of approaches *wall velocity approaches*, as they compute a wall velocity to match the shear stress of the wall model and then apply that wall velocity via the bounce back approach. Pasquali et al. [114]

uses the information about the stress tensor included in cumulants to compute a skin friction coefficient and a wall velocity. However, this method requires a relation of skin friction and bulk velocity and a viscosity to relate cumulants and stress tensor. However, it is applicable to curved boundaries. The inverse Momentum Exchange Method (iMEM) by Asmuth et al. [9] combines the bounce-back method and the momentum Exchange Method [25] to compute the wall velocity required such that the momentum exchanged between wall and fluid matches the shear stress from the wall model. It is only based on the bounce back approach and can therefore be applied with all collision operators and it does not require a viscosity. However, although possible, no generalization to curved boundaries has been presented as of yet.

There has been no consistent comparison of wall models taking into account different wall models, boundary conditions, collision operators and SGS models. The most thorough comparison can be found in Spinelli et al. [135], testing different wall models, SGS models and collision operators, but only one boundary condition and at a Reynolds number much too low for atmospheric flows. The only wall modelling approaches that have been applied to atmospheric flows are the reconstruction approach in Feng et al. [37] and the iMEM approach by Asmuth et al. [9]. Feng et al. [37] yields decent agreement in simulations of a neutral boundary layer, while the results of Asmuth et al. [9] yield excellent agreement in a simulation of a pressure driven boundary layer and compare to results from a pseudo-spectral solver presented in Gadde et al. [43].

In this thesis we employ the iMEM approach and present an extension to thermally stratified flows in paper V.

Precursor-Successor Simulations

In the course of this thesis we also developed an approach to conduct precursor-successor simulations, which are commonly used for wind-farm simulations. A precursor simulation is a bi-periodic empty boundary layer used to generate turbulent inflow conditions for the successor simulation. The precursor simulation can often have significantly smaller streamwise extent than the successor simulation containing the wind farm.

The successor contains the wind farm. It is only periodic in the lateral direction and needs to prescribe inflow and outflow conditions in the streamwise direction. In the LBM, the inflow boundary condition can be either set to a velocity boundary condition or we can prescribe the missing populations directly. The first approach has the advantage of requiring less memory, while the latter requires less computation and is more accurate as the populations also contain higher order moments related to the stress tensor. By linearly interpolating to every 10th time step we were able to reduce the memory overhead of the populations approach with negligible effect on accuracy, yielding a very efficient and accurate boundary condition.

Complex Geometries

Another consequence of the cubic grid is that cells can not be deformed to follow geometries, as with finite volume methods. However, the geometry can be approximated by using interpolated bounce back methods, such as the method proposed by Bouzidi et al. [25] or in the appendix of Geier et al. [50]. A more simple method for incorporating complex geometries is the Immersed Boundary Method pioneered by Peskin [116]. The IBM can be applied in both Navier-Stokes and LBM solvers. Most IBM approaches aim at imposing a no-slip condition. However, for an ABL flow framework, IBM in conjunction with a wall model is far more interesting. This has been done in Navier-Stokes solvers, however, we could not find an example in an LBM solver.

2.3.5 Advection-Diffusion

The LBM can also be used to solve the advection diffusion equation of a scalar, such as the potential temperature. Introducing a second set of populations $g_{\iota\kappa\lambda}$ and a modified collision operator Ω^{AD}, the LBM for advection diffusion is

$$g_{\iota\kappa\lambda}(\boldsymbol{x} + \Delta t \boldsymbol{c}_{\iota\kappa\lambda}, t + \Delta t) = g_{\iota\kappa\lambda} + \Omega^{\mathrm{AD}}_{\iota\kappa\lambda} \tag{2.18}$$

and

$$\theta = \sum_{\iota\kappa\lambda} g_{\iota\kappa\lambda}. \tag{2.19}$$

The main difference between an LBM solving the advection-diffusion equation and the Navier-Stokes equations is that the LBM for advection-diffusion only conserves the zeroth order moment during collision. This enables some simplifications in the method. In this work, in particular, paper V, we employ a cascaded LBM to solve the advection-diffusion equation alongside the Navier-Stokes equations to solve the full system of equations described in the beginning of this chapter.

2.4 Turbine Modelling with the Actuator Line Method

Simulating wind-farm flows requires simulation of both the flow and the turbines. Fully resolving the geometry of a wind turbine or even just a blade is extremely costly. However, the exact details of the flow along the blade surface only have a negligible effect on the flow in the whole wind farm. Hence, the effect of the blade on the flow can be modelled via actuator methods that impose a body force on the flow. The two main methods are the actuator disc and actuator line method, modelling the rotor either as a rotor disc or a line per blade, respectively. There exist also methods in-between, such as the actuator sector method [106].

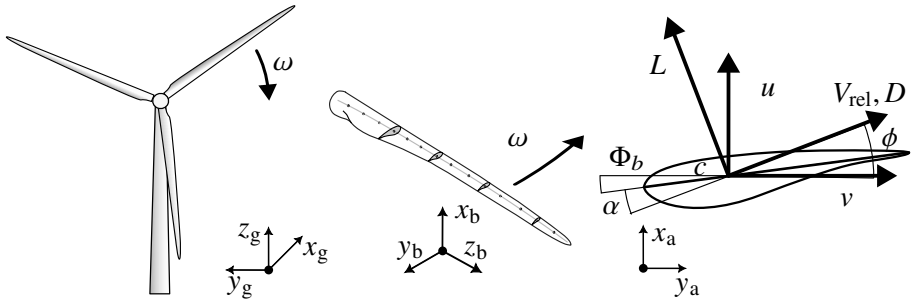

Figure 2.3. Illustration of global, blade and airfoil coordinate system and angles, velocities and forces in the actuator line method.

Since the ALM represents the turbine with more geometrical features, it can also resolve more physical flow phenomena, for example the tip vortices. It was first proposed in Sørensen and Shen [134] and has since been applied widely, see for example Sørensen et al. [132] or Martínez-Tossas et al. [97] for an overview. The ALM is based on blade element theory. The blade is divided into sections and lift and drag forces are computed for each section, typically based on a look-up table. We describe here the most simplistic form of the ALM, with collocated velocity sampling and force projection, but many alterations have been proposed over the years. An overview over the setup is given in Figure 2.3. There are three frames of reference, the global frame of reference $\boldsymbol{x}_\mathrm{g}$, the blade frame of reference $\boldsymbol{x}_\mathrm{b}$ and the airfoil frame of reference $\boldsymbol{x}_\mathrm{a}$. First, a velocity sampled from the flow domain at each actuator point i has to be transformed into the airfoil frame of reference. Defining Ψ_b as the azimuth and Φ_b as the pitch of blade b, and γ as the yaw of the turbine, the velocity $\boldsymbol{u}0_i^\mathrm{B}$ in the rotating blade coordinate system can be computed from the sampled velocity $\boldsymbol{u}_i^\mathrm{LES}$ via

$$\boldsymbol{u}_i^\mathrm{B} = \begin{pmatrix} 0 \\ \omega r_i \\ 0 \end{pmatrix} + \boldsymbol{R}_x(\Psi_b) \cdot \boldsymbol{R}_z(\gamma) \cdot \boldsymbol{u}_i^\mathrm{LES}, \tag{2.20}$$

where ω is the rotor speed and r_i is the radius of the blade section and $\boldsymbol{R}_x$ and $\boldsymbol{R}_z$ are the rotation matrices about x- and z-axis, respectively. Dropping the superscript signifying the frame of reference and the section index i, the angle of attack α and relative velocity V_rel can be extracted from the local velocity:

$$\phi = \arctan\frac{u}{v} \tag{2.21}$$

$$\alpha = \phi - \Phi_b \tag{2.22}$$

$$V_\mathrm{rel} = \sqrt{u^2 + v^2} \tag{2.23}$$

Lift and drag force L and D can then simply computed from

$$L = \frac{1}{2}\rho c V_{\text{rel}}^2 C_{\text{L}}(\alpha) \tag{2.24}$$

$$D = \frac{1}{2}\rho c V_{\text{rel}}^2 C_{\text{D}}(\alpha), \tag{2.25}$$

where the lift and drag coefficients C_{L} and C_{D} are obtained from a look-up table and c is the chord length of the blade section. Finally, the line force on the blade section in the global frame of reference can be computed from rotating lift and drag via

$$\boldsymbol{f} = R_z(-\gamma) \cdot \boldsymbol{R}_x(-\Psi_b) \cdot \boldsymbol{R}_z(\phi) \cdot \begin{pmatrix} L \\ D \\ 0 \end{pmatrix}. \tag{2.26}$$

The reaction force to the line force that is imposed on the flow has to be converted into a body force and is smeared via convolution with a kernel function to improve numerical stability:

$$\boldsymbol{F}(\boldsymbol{x}) = -\frac{1}{\Delta x^3} \sum_i \boldsymbol{f}_i \eta_\epsilon(|\boldsymbol{x} - \boldsymbol{x}_i|)\Delta r_i, \tag{2.27}$$

where

$$\eta_\epsilon(x) = \frac{1}{\left(\sqrt{\pi}\epsilon\right)^3} \exp\left(-\left(\frac{x}{\epsilon}\right)^2\right) \tag{2.28}$$

is the isotropic Gaussian kernel with smearing parameter ϵ, which was proposed in the original formulation and is adopted in this work. Other options for smearing kernels are presented, for example, in Schottenhamml et al. [130] or Jha and Schmitz [64]. A common guideline for the smearing width is to relate it to the grid size via $\epsilon = 2\Delta x$ based on recommendation in Troldborg [147]. However, the smearing also introduces an error, leading to an overprediction of forces near the root and tip of the blade, which has garnered some interest in the scientific community in recent years. We examine the effect of the smearing width and compare some of the correction methods in detail in paper IV.

One of the major drawbacks of the actuator line model is that it limits the time-step size by

$$\Delta t < \frac{\Delta x}{\lambda V_0}, \tag{2.29}$$

where $\lambda = \frac{\omega R}{V_0}$ is the tip-speed ratio and V_0 is the inflow velocity. Nevertheless, a comparison to (2.12) shows, that the low Mach number requirement of the LBM is stricter than the time step limit of the ALM for $\lambda < \frac{\sqrt{3}}{\text{Ma}} \approx 17$, assuming $\text{Ma} = 0.1$. Thus, the ALM does not restrict the time-step size, when used with the LBM. However, the ALM also requires higher spatial resolution than the actuator disc, typically $\Delta x < \frac{D}{24}$, as we also discuss in paper IV. The actuator

disc can be used with just $\Delta x < \frac{D}{7}$ [163]. Hence, the actuator disc can be attractive when studying very coarse resolution cases, e.g., wind-farm wakes. Yet, there also exists a minimum resolution to properly resolve the boundary layer, in the neutral case typically around $\Delta x = 10\,\mathrm{m}$, in thermally stratified simulations often even $\Delta x = 5\,\mathrm{m}$ [3]. Thus, the boundary layer will usually be resolved enough for the ALM if $D \geq 240\,\mathrm{m}$ in neutral and $D \geq 120\,\mathrm{m}$ in thermally stratified conditions.

2.5 Wind Turbine and Wind Farm Control

Wind turbines have three main controllable elements: the yaw angle, the pitch angle and the generator torque M_G. The power output P of a turbine is computed by

$$P = \omega M_\mathrm{G}. \tag{2.30}$$

Rotor speed, generator torque and aerodynamic torque M_aero are further related by Euler's second law:

$$I\frac{\partial \omega}{\partial t} = M_\mathrm{aero}(\gamma, \boldsymbol{\Phi}, \omega) - M_\mathrm{G}, \tag{2.31}$$

where I is the moment of inertia of the rotor. The aerodynamic torque can simply computed from the force along the blades,

$$M_\mathrm{aero} = \sum_b \int (x_i(r_b) - x_{\mathrm{h},i}) f_j \epsilon_{ijk} t_k \mathrm{d}r_b, \tag{2.32}$$

where $\boldsymbol{x}_\mathrm{h}$ and $\boldsymbol{t}$ are location of the hub and rotation axis of the rotor, respectively, and $\boldsymbol{x}(r_b)$ follows the center line of blade b. It can be shown that controlling the generator with $M_\mathrm{G} = k\omega^2$ yields a constant rotor speed and hence power [55]. Thus, a single turbine can be optimally controlled (meaning with maximal power output) by tuning k and setting yaw and pitch to $0°$. This control paradigm is also referred to as greedy control as each turbine maximizes its power output.

However, when wind turbines are gathered in a wind farm, the settings of one turbine affect the surrounding turbines. Determining the optimal settings is therefore non-trivial, since controlling one turbine sub-optimally can lead to improvements at another turbine, typically in the form of power gains. This is the field of wind-farm control. There are two main ways of doing wind-farm control, induction and wake-steering. In wake-steering control, the wake of the turbine is steered away from downwind turbines, thus decreasing wake losses at the cost of decreased power production of the upwind turbine. In induction control the induction factor of the upwind turbine is reduced with the objective of reducing wake losses and improving power production of a downwind turbine. The induction can be reduced either by changing k or by pitching the

blades. It is now well established that static induction control does not lead to overall power gains. However, dynamically altering the induction can excite instability modes in the wake of the turbine and thus lead to rapidly recovering wakes and overall power gains. The effectiveness of dynamically changing the induction was first observed in Munters and Meyers [107], where an optimal control framework lead to oscillating induction control, later named the *pulse* approach. Instead of altering the induction, Frederik et al. [40] proposed to pitch the blades individually, creating a helical wake in an approach named the *helix* approach. Introducing the excitation frequency ω_e and the pitch amplitude A, the helix approach consists of prescribing the blade pitch as

$$\Phi_b(t) = A \sin(\Psi_b + \omega_e) . \tag{2.33}$$

We examined many aspects of the helix approach in paper VII.

2.6 Implementation

Most methods necessary to build a LES framework that meets the requirements postulated in subsection 1.2.5 are available. However, they have not been combined into a single framework. In this thesis, we develop such a framework based on two components: the GPU-resident version of the general-purpose LBM solver VirtualFluids, developed over 20 years at the Institute for Computational Modelling in Civil Engineering at TU Braunschweig [47]; and the Wind Farm Interface (WiFI), developed during the course of this thesis [70]. The basic idea is similar to that of SOWFA, that is, build an extension for wind energy on top of an already established solver and thus benefit from developments in that solver. The general division of work between VirtualFluids and WiFI is shown in Figure 2.4. For optimal computation speed, VirtualFluids is implemented in C/C++ and uses the NVIDIA CUDA Toolkit [109]. Compared to a single time step in the LBM solver, the computations required by WiFI are much simpler, while complex processing of configuration files, airfoil data, etc. dominate. Therefore, WiFI was primarily implemented in Python and only performance critical calls, such as computing the actuator line, were implemented in CUDA. To couple the two components we decided to extend the python bindings of VirtualFluids using pybind11 [60] to the GPU code as well. Another motivation for the design of WiFI was the consideration of coupling to an aeroelastic tool. However, we found that the currently available tools, i.e. OpenFAST, were not sufficient to meet our requirements for computational performance. The ALM includes implementations of the smearing corrections by Meyer Forsting et al. [102] and Martínez-Tossas et al. [99].

In addition to computing forces via the ALM, WiFI also contains implementations of a number of controllers, such as a simple constant speed controller, the reference controller described in Jonkman et al. [65] and a pitch controller for the helix approach. Furthermore, it was designed for application to large

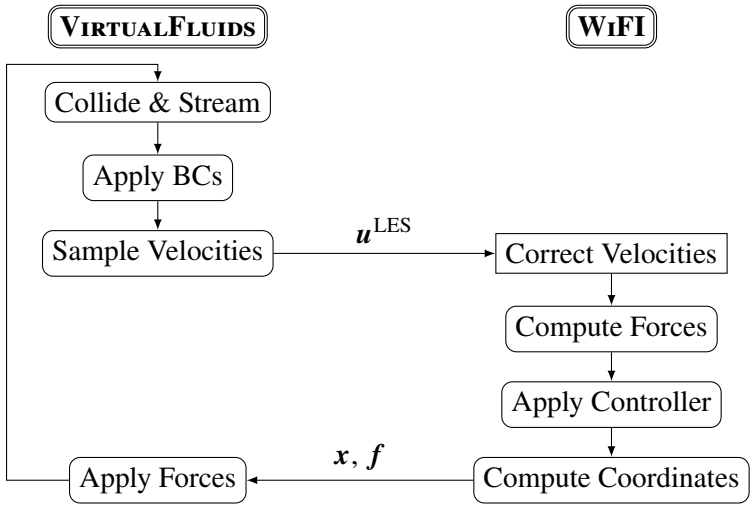

Figure 2.4. General schematic of the coupling between VirtualFluids and WiFI. Steps in rectangles without rounded corners are optional.

wind farms, so that scaling to a large number of turbines only incurs small overhead. The framework is used throughout this thesis with the exception of paper VI and paper VII.

3. Requirements for Industrial Applications

The growing size of wind turbines and wind farms brings many new modelling challenges, for both academia and industry, as outlined in subsection 1.1.1. Yet, the adoption of LES in industry is generally assumed to remain marginal and considerable academic effort is spent on improving and extending engineering wake models. However, there is little knowledge on actual usage of models in the industry beyond anecdotal data from individual conversations with industry representatives. To better establish which models are being used for which task and what the barriers are for industrial application, we conducted a survey among industry experts in paper II. The focus of the survey was on establishing requirements for computational performance but also to understand what other factors limit the adoption of more sophisticated models.

3.1 Survey Design

The survey was conducted in the form of an anonymous online questionnaire of 11 questions, distributed via email and social media. A total of 48 participants answered the survey over the span of two months. It is difficult to establish, whether this represents a significant fraction of wind industry employees working in this area, however, a random subsample of 50 % of responses exhibits similar trends to the whole dataset, indicating converged data.

3.2 Survey Results

3.2.1 Background of Respondents

First, we established the background of the respondents. Most stakeholder groups in the wind energy industry are present, with 40 % of respondents working as consultants, 32 % working in wind-farm development and 11 % in original equipment manufacturing. Furthermore, more than half of respondents had 10 years or more of experience in the wind energy industry and only 10 % less than 2 years. Furthermore, more than half of correspondents reported advanced or expert level knowledge of CFD. Thus, we conclude that the background and level of expertise of the respondents was overall high enough to provide valuable insights.

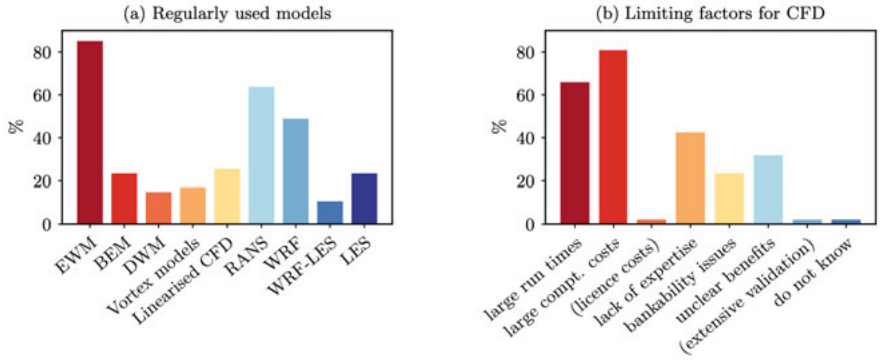

Figure 3.1. Survey results on (a) used models and (b) limiting factors of using CFD.

3.2.2 Current Modelling Practices

Figure 3.1 (a) shows which kinds of models are currently applied in the industry on a regular basis. The vast majority of respondents used engineering wake models (EWMs), but also Reynolds Averaged Navier Stokes (RANS) and the Weather Research and Forecasting (WRF) model. Almost a quarter of participants used LES and another 13 % used WRF-LES, highlighting that LES exists in industrial use.

When asked about factors limiting the use of CFD, computational cost and large run times were cited as the most common reasons, which is in accordance with the general understanding. However, other factors such as lack of expertise, bankability issues and unclear benefits were also indicated by at least 20 % each. All three factors can be read as a need to improve collaboration between industry and academia, and as a call to academia to offer training in these models.

3.2.3 Requirements for the Use of LES

When asked which improvements would be necessary to use LES, respondents again identified "faster run times" and "lower computational cost". The next most important point is "more validation". This is another aspect that can be addressed in academia, by providing more validation studies, although the limited availability of validation data is a major limiting factor in that. Several individual comments also pointed to the need for validation and quantification of uncertainties. Participants were also asked identify the areas with the largest potential benefits for the use of LES. Respondents anticipated that all modelling challenges could benefit from LES in all four areas, complex terrain, stratification, blockage and gravity waves, and wind-farm wakes, but respondents considered wind-farm wakes to be most relevant, followed by complex terrain.

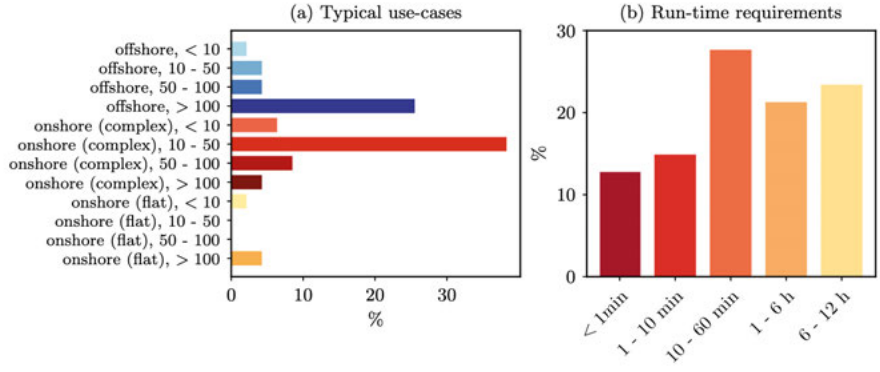

Figure 3.2. Survey results on (a) typical use case for LES and (b) corresponding runtime requirements.

In order to quantify runtime requirements, respondents were asked to, first, describe their typical use case in terms of number of turbines and terrain, and secondly, give an acceptable runtime for a simulation of 1 h. The results are shown in Figure 3.2. By far the most common cases are medium sized wind farms in complex terrain and large wind farms in offshore conditions. Responses on runtime requirements were much more evenly distributed, likely linked to the variety of application for which the models are used. For example, layout optimization has significantly lower maximal runtime than calibration of a model for annual expected production. About 45 % of the responses correspond to a real-time factor of below or equal to 1, while another 27 % require real-time factors of 1 to 6.

3.2.4 Comparison to GPU-resident LBM

To compare the requirements indicated by the correspondents to the currently available technology, we conducted a simulation of a wind farm comprising 64 DTU 10 MW reference turbines [17]. The turbine diameter is $D = 198\,\mathrm{m}$ and the turbines are arranged in an aligned grid with $7\,D$ streamwise and $5\,D$ lateral spacing. The domain has a size of $20\,\mathrm{km} \times 15\,\mathrm{km} \times 1.6\,\mathrm{km}$ is discretized with a grid spacing of $\Delta x = 10\,\mathrm{m}$ in the lowest 500 m. The total grid comprises 187×10^6 nodes. Subgrid stresses are modelled using the QR model and the bottom boundary employs a log-law as wall model, with friction velocity $u_* = 0.425\,\mathrm{m/s}$ and roughness length $z_0 = 0.01\,\mathrm{m}$. The flow is driven by a constant pressure gradient, resulting in velocity $u_0 = 11.2\,\mathrm{m/s}$ and turbulence intensity $Ti = 6.1\,\%$ at hub height. A snapshot of the instantaneous streamwise velocity is shown in Figure 3.3. The simulation is run for a total runtime of 8000 s on a single Nvidia A100 GPU and is completed after 5112 s, corresponding to a real-time factor of 1.56. Thus, the capabilities at

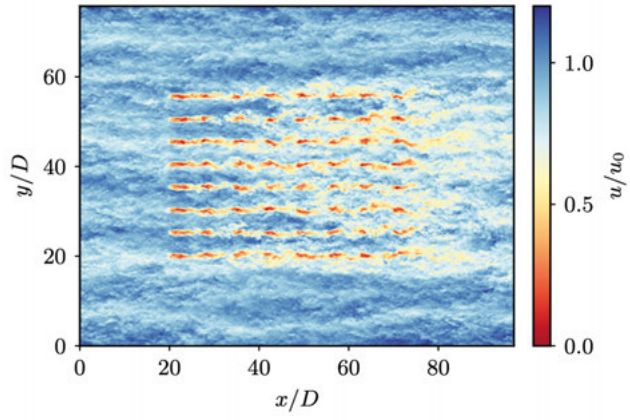

Figure 3.3. Instantaneous visualization of the flow field at hub height using streamwise velocity.

time of writing paper II fulfil the runtime requirements of about 45 % of all participants. Further improvement in computational efficiency since 2023 has put the computational requirements of another 12 % in reach within the next few years.

3.3 Conclusion

The survey confirms the widely held belief that one of the major barriers to adoption of LES is the large computational cost / long runtimes. However, it also highlights that other aspects play an important role as well, such as lacking expertise / training and validation. Issues such as the lack of bankability furthermore indicate that adoption needs to happen across different sectors of the industry, underlining the importance of flexible frameworks that can be adopted to different use cases. At the same time, the results showed that current technology can provide sufficiently fast runtimes for about half of the respondents. However, adoption will not take place unless industrial users also know about these results. Relating the findings to the aim of this thesis, that is, developing a framework for performing LES that is applicable in the industry, it highlights the needs for thorough validation and accessibility of the software and training material.

4. Development and Validation of a Framework for Wind Energy Research Based on the LBM

The following chapter comprises papers III, IV and V. First, the framework for simulating wind turbines in atmospheric boundary layers is validated against benchmark data in section 4.1. In section 4.2, parameters and correction methods for the turbine modelling are examined. Finally, the framework's capabilities are extended to simulating thermally stratified boundary layers in section 4.3.

4.1 Validation of the Framework for a Single Turbine

As the last two papers showed, the LBM is suitable for wind energy applications and promises to be computationally efficient enough to meet many industrial users' requirements. However, paper II also showed that another criterion for acceptance in the industry is validation. Validation against full scale measurements is rare for numerous reasons. Prescribing exact boundary conditions for simulations that match those of the measurements is difficult, since the measurements required for boundary conditions often do not exist. Furthermore, even cross-code comparisons have shown not to yield the same results due to different models, configuration and even small implementation details, see, for example, Martínez-Tossas et al. [97]. Additionally, publicly available full-scale measurements are extremely rare and even then, comparison is difficult, as highlighted in Asmuth et al. [14].

The SWiFT Benchmark by Doubrawa et al. [34] compares simulation tools of varying fidelities to results from a measurement campaign at the Scaled Wind Technology Facility (SWiFT). The comparison comprises measurements of the turbine response and velocities in the turbine wake. In paper III, we repeat the benchmark case using the framework described in section 2.6 and compare with numerical and experimental results. Furthermore, we compare computational cost and accuracy of the different models.

4.1.1 Setup

The benchmark consists of measurements conducted at the SWiFT facility in Lubbock, Texas. The turbine is a Vestas V27 turbine with $D = 27\,\mathrm{m}$ and hub

Table 4.1. *Parameters of the benchmark cases as described in Doubrawa et al. [34], including mean and standard deviation of velocity at hub-height, turbulence intensity at hub height and friction velocity.*

U / m/s	σ_u / m/s	σ_v / m/s	σ_w / m/s	Ti_{hub} / %	u_* / m/s
8.7	0.93	0.77	0.50	10.7	0.45

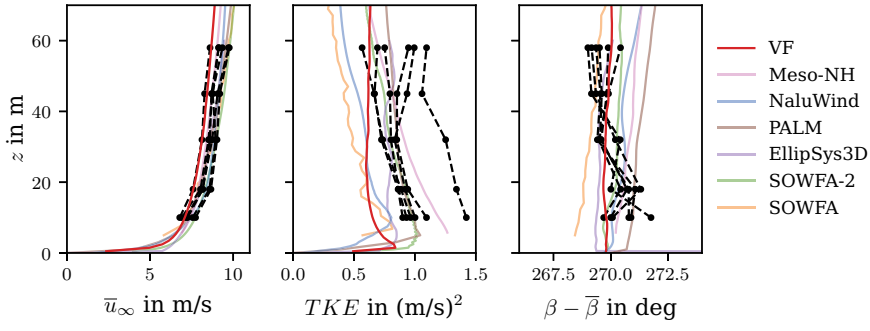

Figure 4.1. Inflow conditions, including velocity (left), *TKE* (center) and veer (right) measured 2.5D upstream of the turbine.

height $z_{\text{hub}} = 32.1\,\text{m}$. The inflow is captured by a meteorological mast 2.5 D upstream of the turbine, while the wake between 1 D and 5 D is measured using a DTU SpinnerLidar. Additionally, rotor speed, generator power, and torque of the turbine are measured. We only consider the neutral case of the three benchmarks defined in the original publication. Six ten-minute periods from the data were chosen and an ensemble average was computed to define the flow conditions given in Table 4.1.

We employ a precursor to generate realistic inflow conditions with a size of $6\,\text{km} \times 6\,\text{km} \times 1\,\text{km}$ and an isotropic grid of $\Delta x = 7.8125\,\text{m}$. We set periodic conditions in streamwise and lateral directions and employ the wall model with $z_0 = 1.4 \times 10^{-2}\,\text{m}$ to match the inflow conditions. The flow is driven with a constant pressure gradient. The successor simulation has the same outer dimensions as the precursor with successive refinement towards the turbine which is placed at $(3000\,\text{m}, 2500\,\text{m}, 32\,\text{m})$. At the finest grid the resolution is $\Delta x \approx 1\,\text{m}$. The turbine model includes the smearing correction presented in Meyer Forsting et al. [102].

4.1.2 Results

We compare our results to the results from the other LES solvers from Doubrawa et al. [34] and the results presented in Jézéquel et al. [62].

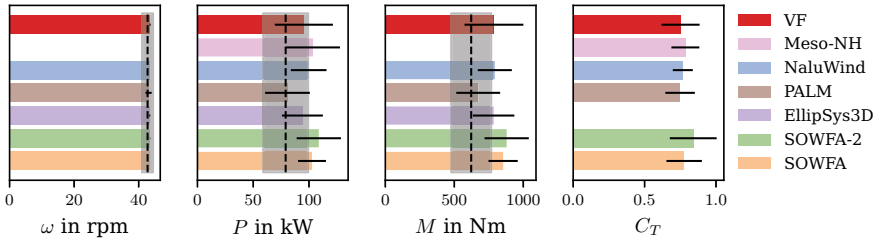

Figure 4.2. Turbine Response including average rotor speed, turbine power, generator torque and thrust coefficient.

Inflow

The vertical mean profiles of inflow velocity u_∞, turbulence kinetic energy *TKE* and veer β are shown in Figure 4.1. We find that the velocity profile is in close agreement with the measurement results, albeit somewhat lower than most other solvers. The profile of *TKE* is more constant with height than for most other solvers, however, this agrees quite closely to the average of the measurements. The turbulence intensity from the simulation matches the measured turbulence intensity well, since both *TKE* and u_∞ are on the lower end of the ensemble distribution. Our results don't exhibit any veer as there is no Coriolis force featured in our simulation.

Turbine Response

The average turbine response is shown in Figure 4.2. All solvers agree closely in the rotor speed, including ours. However, when it comes to power, the differences are significantly larger. The results of our solver match the measurements well, but PALM yields the closest agreement. Since all match the rotor speed well, the torque shows a similar picture as the power. Finally, we also compare the thrust coefficient C_T. There are no measurements available. Again, the solvers differ quite widely, with PALM and VirtualFluids reporting the lowest coefficient and SOWFA-2 the highest. We attribute the close agreement of our results with the measurements to the use of the smearing correction, as it is known that the ALM leads to overprediction of the forces at the tip, resulting in overprediction of power. We will analyze this phenomenon in more detail in paper IV.

Wake

To compare the flow field in the wake, we portray the measured flow field in the wake alongside the simulation results and their difference in Figure 4.3. The wake is slightly shifted to the left in the experiments, but otherwise agrees closely with the simulation results in shape and strength. Inside the wake the differences are small as Figure 4.3 (c) illustrates.

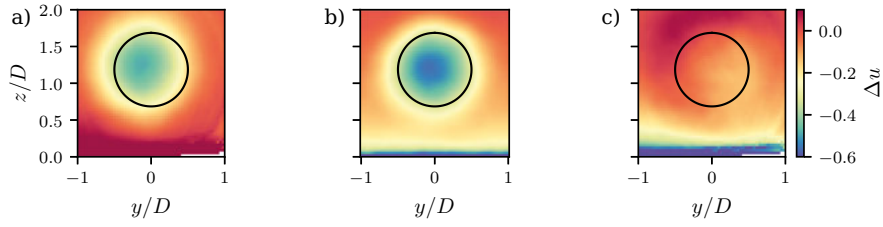

Figure 4.3. Contour of velocity deficit in the wake of the turbine at $x = 3D$. (a) Measurements, (b) simulation with VirtualFluids, (c) difference between measurements and simulation.

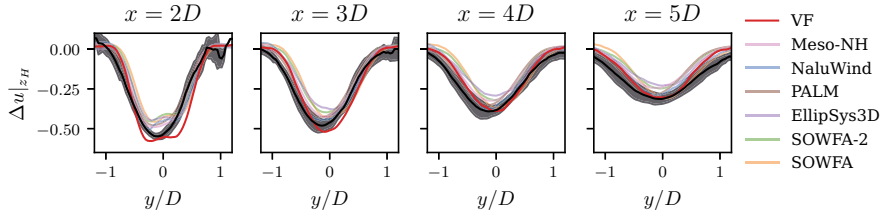

Figure 4.4. Horizontal profiles of velocity deficit at various downstream distance.

A more detailed view of the wake is given in Figure 4.4. We observe that most solvers underestimate the velocity deficit throughout the wake, while our results agree well further down the wake. In the near wake at $x = 2D$, our solver overestimates the velocity deficit. Furthermore we observe that the wake has not yet transitioned at that position, while the measurements indicate an already transitioned wake as shown by the single Gaussian shape.

Computational Cost

While LES can offer the highest fidelity of results, a user can also choose a lower fidelity tool if it results in acceptable accuracy for the particular use case. To be able to choose the appropriate tool, the user must usually balance

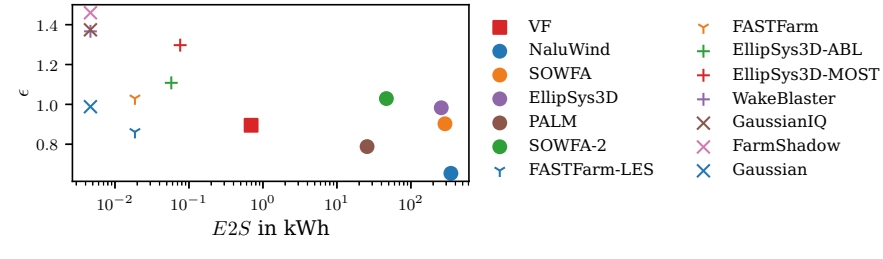

Figure 4.5. Comparison of energy-to-solution and error for participants in the benchmark.

computational cost and accuracy. The quantification of computational cost is difficult, as runtime alone is not a good indicator. Results depend on available hardware, scale of parallelization and numerous other factors. The use of GPUs exacerbates the difficulties, as a single GPU has many more cores than a CPU but the cores are slower and less independent. A metric that balances some of these effects is the energy-to-solution (*E2S*), that is, all electricity used by the hardware to arrive at the solution of the given problem. Based on data provided by Doubrawa et al. [34] and assumptions on the consumption of a CPU we compute the energy to solution for all participants of the benchmark. The accuracy is measured as the average mean squared difference between simulation and measurement along horizontal and vertical line through the hub. The results of that comparison are shown in Figure 4.5. We clearly find that VirtualFluids is one to two orders of magnitude more efficient than all other LES solvers, while being as accurate. On the other hand, lower fidelity models are still significantly more efficient than the LBM solver, however, they are also consistently less accurate.

4.1.3 Conclusion

The goal of paper III is to validate the methodology of the LES framework developed in this thesis and compare the computational cost of the method. To that end we compare to the results from the SWiFT benchmark. We find that the precursor-successor setup yields good inflow conditions. The turbine response, including rotor speed and power match very well with the measurements, which we also attribute to the use of the smearing correction by Meyer Forsting et al. [102]. In the wake we find that our results also agree well with the measurements, but slightly overpredict the wake strength. Conversely, the other participants in the benchmark consequently underestimated the wake strength. Comparing the energy-to-solution and a global error measure we find that our method yields a similar accuracy as the other LES solvers and consistently higher accuracy than lower-fidelity models. Meanwhile,our model is at least one order of magnitude more efficient than all other LES solvers. A drawback of the benchmark is the size and age of the turbine, which is not representative of turbines being built today, therefore many aspects of modern turbines can not be assessed. Nevertheless, the study represents an important step for the framework and validation of LBM based solvers in general.

4.2 Improving Rotor Modelling

The comparison of rotor response in paper III showed that LES solvers tended to overpredict the power, with the exception of PALM. This is a known issue of the actuator line method that all of the participants use, with the exception of PALM (forces are imposed as lines but sampled upstream [78]). The ALM is a widely used method, and often its capabilities to capture blade effects are required, for example in the study of individual pitch control or fatigue load analysis. However, this accuracy comes at increased requirements in terms of grid and time-step size. Furthermore, it is well documented that the forces computed with the ALM become inaccurate at coarser resolutions due to the smearing of body forces in the LES domain, see section 2.4. In paper IV, we study the effect of the grid size, the smearing width, and the induction corrections on the ALM and establish criteria for users to choose the level of accuracy and computational cost they desire.

4.2.1 Corrections for the ALM

Of the many induction corrections proposed, we decided to examine the filtered lifting line correction (FLLC) by Martínez-Tossas et al. [99] and the vortex-based smearing correction (VBSC) by Meyer Forsting et al. [101], including the improvements from Meyer Forsting et al. [102]. Both corrections follow the same reasoning: Vortices are shed at the edges of the sections of the actuator line due to a change in circulation. These vortices, in turn, induce a velocity at the blade. The induced velocity decreases the angle of attack and thus reduces circulation and lift. However, the vortices shed by the actuator line are weakened by the smearing, resulting in an increased lift and loading. This effect is especially pronounced in regions of high lift gradients, which is near the root and tip of the blade. The correction methods compared here aim to compute that missing induced velocity $u_{\text{ind}} = u^{\text{opt}} - u^{\text{LES}}$, where u^{opt} is the velocity induced by an optimal actuator line, and add it back to the sampled velocity u^{LES}. The correction methods differ in three points. First, the FLLC considers an actuator line with $\epsilon = 0.25c$ as optimal, whereas the VBSC assumes an ideal lifting line. Second, the methods assume different shapes of the vortices: In the FLLC, all vortices are straight, since it relies on linearization, whereas the VBSC includes the near-wake model by Pirrung et al. [118] and Pirrung et al. [117] to track the shed vortices. And third, the resulting equations require different solution methods.

4.2.2 LES Study

We conduct a parameter study to examine the effect of grid resolution and smearing width by simulating a single NREL 5 MW reference turbine [65] with $D = 126\,\text{m}$. We again use VirtualFluids coupled to WiFI, using the same

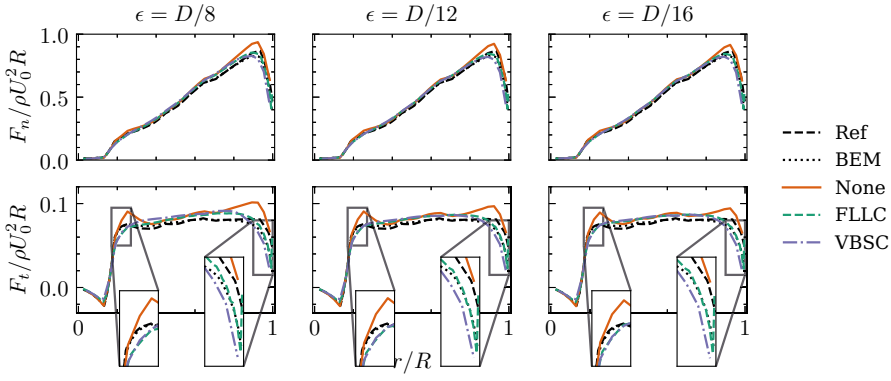

Figure 4.6. Averaged normal and tangential force along the blade for different smearing widths at $\Delta x = D/32$.

method as in paper III. The domain is $24\,D \times 8\,D \times 8\,D$ with periodic lateral boundaries. The inflow is a velocity boundary condition. We vary grid resolution $\Delta x \in \{D/16, D/24, D/32\}$ and smearing width $\epsilon \in \{D/8, D/12, D/16\}$. Note that we examine smearing width and grid resolution independently from each other. We simulate all possible combinations of the aforementioned parameters in uniform inflow. In addition we conduct simulations with 5 % and 10 % turbulence intensity of the highest resolution and smallest smearing width. Each simulation is repeated without and with both corrections. Furthermore, we simulate one high-resolution reference case at $\Delta x = D/196$. The turbines all operate at constant rotor speed with $\lambda = 7.55$.

An additional simulation of a wind farm in a pressure driven boundary layer is conducted, again without correction and for both corrections. The farm comprises 6 rows of 5 NREL 5MW turbines in a domain of $60\,D \times 40\,D \times 10\,D$. The domain has a resolution of $\Delta x = D/12$ and is refined around the turbine by a factor of 2. The inflow is generated with a precursor. At the bottom, we apply the iMEM boundary condition with $z_0 = 1 \times 10^{-3}$ m. The flow is driven by a constant pressure gradient resulting in a hub-height velocity $u_\mathrm{h} \approx 8.4\,\mathrm{m/s}$ and $Ti = 5.7\,\%$.

Accuracy and Computational Cost of the Corrections

We begin with examining the effect the smearing corrections on the forces computed by the actuator line, shown in Figure 4.6. We compare the results of the uniform inflow case at $\Delta x = D/32$, with the reference simulation and computations using Blade Element Momentum theory (BEM) [55]. For brevity's sake, we only show a single resolution, but the results at other resolutions follow a similar trend. We find that the reference case and BEM results agree closely up to the tip, where they begin to diverge. This can be due to differing assumptions of the methods but also due to the fact that even in this very highly resolved case $\epsilon > c_\mathrm{Tip}$, underlining the need for corrections. The normal

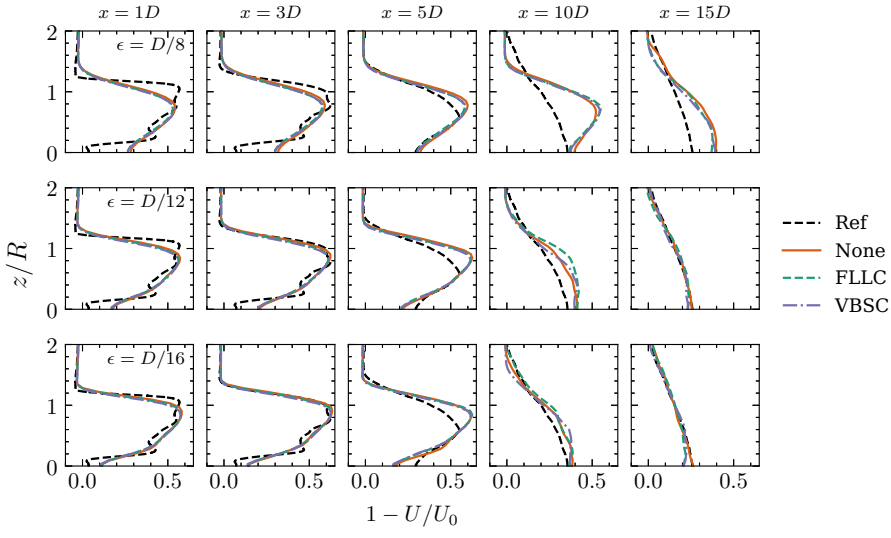

Figure 4.7. Averaged velocity deficit profiles in the wake of the turbine at $\Delta x = D/32$.

and tangential forces F_n and F_t in the middle section of the blade agree well, regardless of smearing width. However, towards the root and tip, we see the effect of the smearing width and the corrections. Without correction, we find an overprediction of forces, which increases with smearing width. The corrections are largely able to offset that overprediction. At the tip, the FLLC matches the BEM results very closely, while the VBSC predicts slightly smaller forces.

In Figure 4.7, we can observe the effect of the smearing width and the correction on the flow field by analysing velocity deficit profiles at various downstream distances. We observe a large effect of the smearing width on wake shape and transition, whereas the corrections only have a very limited effect. This can be expected, as the smearing directly effects the forcing in the flow, whereas the corrections only act indirectly by lowering the force at the tip. We also observe that at the smallest smearing width, the results generally agree well with the reference case with the exception of $x = 5D$. At large smearing widths, however, we find a considerable delay in wake transition due to a weaker wake shear layer [96].

To condense all simulations into a single picture, we constructed a heat map of four metrics: the error in power and thrust, the error in rotor equivalent wind speed [154] at $x = 10D$, and the real-time factor. We observe that power and thrust prediction are sensitive to smearing width, but much less to grid resolution. Meanwhile, the wake becomes more accurate with higher resolution and smaller smearing width. On the other hand, higher resolution leads to much higher computational cost. Finally, we find that either correction leads to significant improvements in power and thrust prediction while having only marginal effects on the other metrics. The VBSC consistently outperforms the

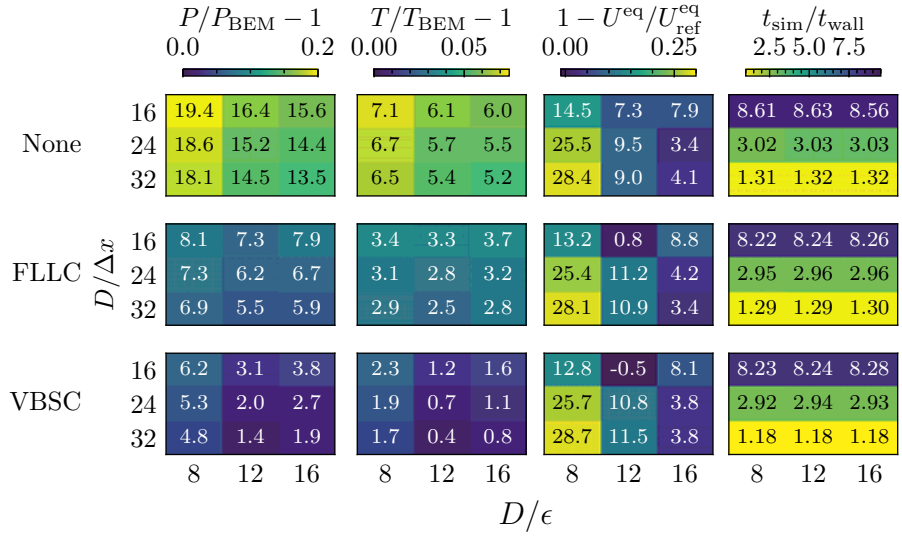

Figure 4.8. Heatmap of results from the parameter study. (Left to right) Error in turbine power relative to BEM; error in thrust relative to BEM; error in rotor equivalent wind speed at $x = 10D$ relative to reference simulation; and real-time factor. Numbers in the first three columns of the heatmaps are given in percent.

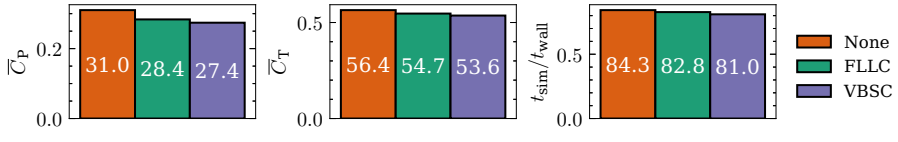

Figure 4.9. Average power and thrust coefficient of the wind farm based on the hub height velocity and real-time factor.

FLLC, but comes with a higher cost. The effect of turbulent inflow is discussed in more depth in paper IV, but the results are generally in line with previous findings.

Finally, we analyze the results from the simulation of a wind farm in a pressure driven boundary layer. Farm average power and thrust coefficient based on hub height velocity, as well as real-time factor are shown in Figure 4.9. Again, we find a significant decrease in power when using the corrections, with the VBSC reducing power prediction by 12 % compared to the uncorrected case. The effect on the thrust coefficient is smaller but still noticeable. All simulations run slightly below real-time, with the uncorrected case being the most efficient, but even the costlier VBSC only incurs an overhead of 4 %. These findings clearly show how important a correction is.

However, Figure 4.10 shows that the effect of the corrections is more complex than simply reducing power. Here, the reduction in power and thrust relative to the uncorrected case is shown per turbine. We find a surprisingly

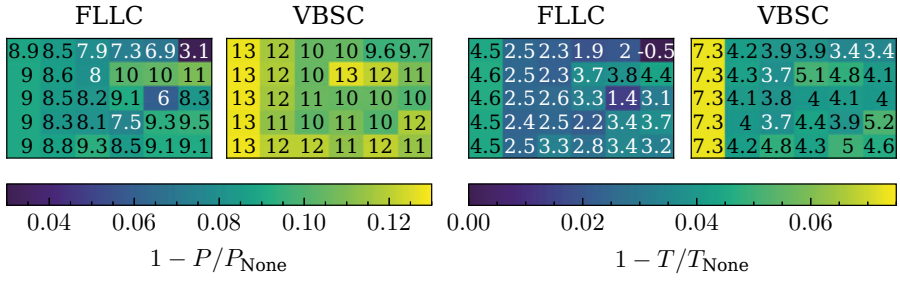

Figure 4.10. Difference in power and thrust when applying the smearing corrections relative to the uncorrected wind farm.

inconsistent effect of the corrections. Both have the strongest effect on the first column of turbines, however, deep in the wind farm we find another row of turbines strongly affected. Furthermore, the corrections differ in the strength of reduction.

4.2.3 Conclusion

This work highlights the effect that the smearing width and the grid resolution have on power, thrust, and wake of the ALM. We also demonstrated that smearing width and grid resolution should be considered two separate parameters, contrary to the common notion that $\epsilon = 2\Delta x$. We showed that at $\Delta x = D/24$, the actuator line can accurately predict the flow, if a sufficiently small smearing width is used. However, the smearing width is limited by numerical stability. The LBM is able to sustain lower smearing widths than the recommended value of $\epsilon = 2\Delta x$ and remains stable even at $\epsilon = \Delta x$. In addition, we compared two corrections for the missing induction and found that both perform well. The VBSC had a stronger effect, while the FLLC incurs slightly lower overhead. Nevertheless, we showed how important the use of a smearing correction is, reducing average power coefficient of a wind farm by up to 12 %. Furthermore, we showed that the corrections have a complex, turbine-specific effect and thus cannot be replaced by a simple correction factor. Improving the accuracy of these basic metrics is critical for industrial application. Hence, further research on this topic is warranted.

4.3 Simulating Thermally Stratified Boundary Layers

Many phenomena observed for large wind farms are affected, or even caused, by thermal stratification, such as wind farm wakes [90] and hydrostatic blockage [83]. Wind farm wakes and blockage also require extremely large domains, see Ouro et al. [113], thus making them a good candidate for the application of the LBM framework.

To consider thermal stratification, we need to solve an additional transport equation for potential temperature, see chapter 2. Generally, there are two possibilities to solve this equation within an LBM framework: discretizing the equation using finite differences or the finite volume method in a hybrid approach; or solving it with another LBM using the double-distribution function (DDF) approach. The hybrid approach has the advantage of requiring less memory, as it only uses of a single value per node, while the DDF approach requires 7 to 27 values, depending on the chosen lattice. On the other hand, in the DDF approach, momentum and temperature solver have the same numerical structure.

Very few works on thermally stratified atmospheric flows using the LBM have been presented in the literature, most notably using the solver ProLB, which uses a hybrid approach [37]. ProLB was also very recently employed by Wang et al. [156] to simulate a model wind farm in stratified conditions. However, it should be noted that these simulations did not include a capping inversion.

Application of the LBM to solve the advection-diffusion equation was pioneered by Wolf-Gladrow [162] and has since been applied in many cases. However, most applications have been restricted to two dimensions. An overview over LBM solvers for advection-diffusion is given in Gruszczyński and Łaniewski-Wołłk [52]. The authors compare a number of advanced collision operators and find that the most accurate method is the cascaded LBM using two relaxation times. Wang et al. [155] employ a DDF approach to simulate flow over a ridge with two MRT collision operators but the results are not satisfactory. We present a model based on the DDF approach using the cumulant and a cascaded LBM for momentum and temperature, respectively, in paper V. We develop a novel combined boundary condition to use Monin-Obukhov similarity theory following Moeng [105] and we implement Coriolis and buoyancy forces in VirtualFluids. The novel model is validated for simulations of conventionally neutral and stable boundary layers by comparing to reference data.

4.3.1 Extensions for a DDF-LBM Solver

As briefly outlined in section 2.3, solving the advection-diffusion equation with LBM is very similar to solving the Navier-Stokes equations. The only difference is in the collision operator and the interpretation of moments of the

populations. Recall that

$$\theta = \sum g_{\iota\kappa\lambda}, \tag{4.1}$$

where $g_{\iota\kappa\lambda}$ is the second set of populations.

Collision Operator

The cascaded collision operator we employ in our model was introduced in Alihussein et al. [2] for simulation of dissolution in porous media. It is also a MRT collision operator and based on the model presented in Geier et al. [46]. Populations are transformed to central moments,

$$k_{\alpha\beta\gamma} = \sum \left(\iota - \frac{u}{c}\right)^{\alpha} \left(\kappa - \frac{v}{c}\right)^{\beta} \left(\lambda - \frac{w}{c}\right)^{\gamma} g_{\iota\kappa\lambda}, \tag{4.2}$$

orthogonalized, and relaxed towards their equilibrium $k^{\mathrm{eq}}_{\alpha\beta\gamma}$

$$k^{*}_{\alpha\beta\gamma} = \omega_{\alpha\beta\gamma} k^{\mathrm{eq}}_{\alpha\beta\gamma} + \left(1 - \omega_{\alpha\beta\gamma}\right) k_{\alpha\beta\gamma}. \tag{4.3}$$

Contrary to the LBM for the Navier-Stokes equations, only the zeroth order moment, i.e. the temperature, is conserved. The relaxation rates of the first order moments are related to the diffusivity by

$$\frac{1}{\omega_{(100)}} = \frac{D}{c_{\mathrm{s}}^2 \Delta t} + \frac{1}{2}, \tag{4.4}$$

where the parentheses imply all permutations of the indices. We set all other relaxation rates to one.

Coupling

The two systems of equations, i.e. the Navier-Stokes equations and the temperature equation need to be coupled to interact. The momentum equations act on the temperature equations through $\boldsymbol{u}$, which appears in transformation of populations to central moments. The effect of the temperature on the flow is modelled via the Boussinesq approximation as temperature changes are assumed to be small.

Boundary Conditions

Boundary conditions differ to the regular LBM because the quantity of interest is the zeroth order moment (analogue to the density of the regular LBM). To apply a wall model, a heat flux $q_{\mathrm{w}} = -D \frac{\partial \theta}{\partial x_i} n_i$, where $\boldsymbol{n}$ is the normal vector to the wall, must be set. The flux at the node can be computed as

$$j_i = j_i^{\mathrm{D}} + j_i^{\mathrm{A}} = -D \frac{\partial \theta}{\partial x_i} + u_i \theta = \sum_{\iota\kappa\lambda} c_{\iota\kappa\lambda,i} g_{\iota\kappa\lambda}, \tag{4.5}$$

where $\boldsymbol{j}$ is the total flux as sum of diffusive and advective flux j_i^{D} and j_i^{A}, respectively. We prescribe the flux at the wall j^{W} to be equal to the diffusive

flux in wall-tangential direction and equal to the heat flux in normal direction by

$$j^{\mathrm{W}} = j_i^{\mathrm{D}} - (j_j^{\mathrm{D}} n_j) n_i + q_{\mathrm{w}} n_i. \tag{4.6}$$

Finally, the new flux can be prescribed via the bounce-back rule:

$$g_{\overline{\iota\kappa\lambda}} = g_{\iota\kappa\lambda} - 2 w_{\iota\kappa\lambda} \frac{c_{\iota\kappa\lambda,i} j_i^{\mathrm{W}}}{c_{\mathrm{s}}^2}. \tag{4.7}$$

The same boundary condition can also be used to prescribe a Neumann type boundary condition.

Monin Obukhov Similarity Theory Wall Model

At the bottom wall, we determine wall shear stress τ_{w} and heat flux q_{w} using Monin-Obukhov similarity theory:

$$\zeta(z) = \frac{z}{L} = -\frac{z \kappa g q_{\mathrm{w}}}{u_*^3 \theta_{\mathrm{r}}} \tag{4.8}$$

$$u(z) = \frac{u_*}{\kappa} \left(\ln \frac{z}{z_0} - \psi_{\mathrm{M}}(\zeta) \right) \tag{4.9}$$

$$\theta(z) - \theta_0 = -\frac{q_{\mathrm{w}}}{u_* \kappa} \left(\ln \frac{z}{z_{0,\mathrm{H}}} - \psi_{\mathrm{H}}(\zeta) \right), \tag{4.10}$$

where ζ is a stability parameter based on the Obukhov length L, $u_* = \sqrt{\frac{\tau_{\mathrm{w}}}{\rho}}$ is the friction velocity, κ is the von-Kármán constant, z_0 and $z_{0,\mathrm{H}}$ are roughness lengths for momentum and temperature, respectively, and θ_0 is the surface temperature [8]. We use the classical similarity functions $\psi_{\mathrm{M}}(\zeta)$ and $\psi_{\mathrm{H}}(\zeta)$ determined by Businger et al. [27]. The combined wall model and boundary condition is hereafter referred to as the surface layer boundary condition.

4.3.2 Validation in a Conventionally Neutral Boundary Layer

The first validation case is a conventionally neutral boundary later (CNBL) following Berg et al. [22]. This testcase was also used for validation of AMR-Wind [80].

Case Description

A bi-periodic domain of size $2560\,\mathrm{m} \times 2560\,\mathrm{m} \times 896\,\mathrm{m}$ is initialized with uniform geostrophic wind $\boldsymbol{G} = [5, 0, 0]^T\,\mathrm{m/s}$ and a constant temperature gradient of $\frac{\partial \theta}{\partial z} = 3\,\mathrm{K/km}$. The top and bottom boundary conditions are slip and surface layer conditions with prescribed heat flux of $q_{\mathrm{w}} = 0\,\mathrm{K/(m\,s)}$ and $z_0 = 0.05\,\mathrm{m}$, respectively. The Coriolis parameter is set to $f^{\mathrm{C}} = 10^{-4}\,1/\mathrm{s}$ and the reference temperature is $\theta_{\mathrm{r}} = 290\,\mathrm{K}$. The simulation is run for 65 T_{E} with eddy turnover time assumed to be $T_{\mathrm{E}} = 1700\,\mathrm{s}$, with results averaged over the last 10 T_{E}. We

Table 4.2. *Inversion height and friction velocity in the CNBL case compared to references.*

	VF		Berg		AMR-Wind	
Quantity	B	C	C	D	C	D
z_i / m	381.5	362.25	372	350	352.1	337.0
u_* / m/s	0.207	0.205	0.225	0.221	0.208	0.203

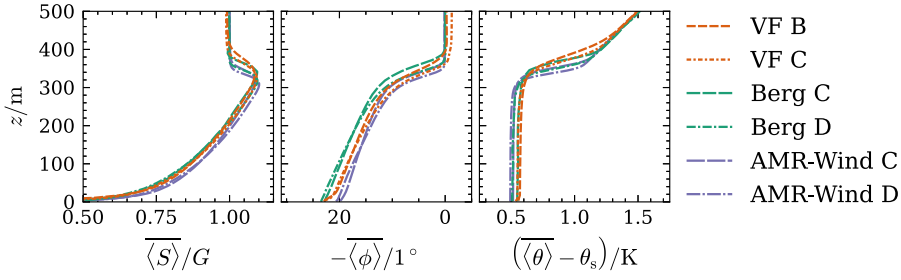

Figure 4.11. Horizontally and temporally averaged profiles of wind speed, veer and temperature of the conventionally neutral boundary layer test case after $65T_\mathrm{E}$

conduct simulations at two resolutions, corresponding to grids B and C in Berg et al. [22] with $\Delta x = 7\,\mathrm{m}$ and $\Delta x = 3.5\,\mathrm{m}$, respectively. We compare to the results obtained with grids C and D ($\Delta x = 1.75\,\mathrm{m}$).

Results

The resulting friction velocity and inversion height z_i, computed as the height with the maximum temperature gradient, alongside results from the references are shown in Table 4.2. We find that for all solvers, an increase in resolution from grid B to C or C to D leads to a decrease in inversion height and friction velocity. However, the solvers differ in inversion height and friction velocity for the same resolution by 3 % to 6 %, with VirtualFluids being between the reference results for inversion height and having the lowest friction velocity. We attribute these differences to the different convergence orders of the solvers. Nevertheless, the results are reasonably close to each other.

Mean profiles of wind speed S, veer ϕ and temperature are shown in Figure 4.11. All profiles agree well with the reference data below the inversion layer. The coarser simulation yields a weaker gradient at the top of the inversion layer in both wind speed and temperature, in line with results from Berg et al. [22] and observations by van Heerwaarden et al. [148] for a similar case. The higher resolution case presents a slightly askew veer profile in the free atmosphere. This is due to the smallness of the Coriolis force in the free atmosphere and the use of single precision floating point numbers, exacerbated by the normalization procedure required by the LBM, see Krüger et al. [79,

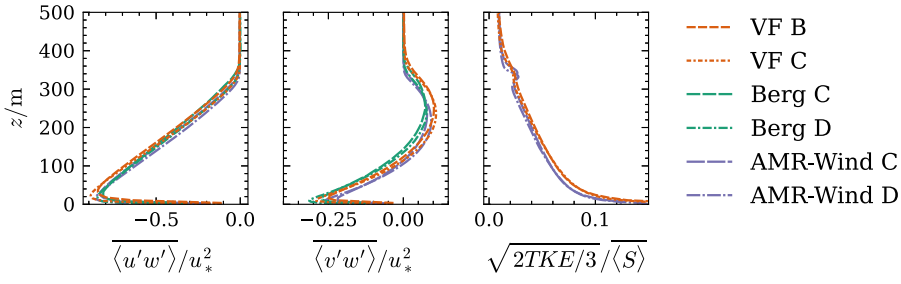

Figure 4.12. Horizontally and temporally averaged profiles of Reynolds stresses and a surrogate of turbulence intensity of the conventionally neutral boundary layer test case after $65T_\mathrm{E}$

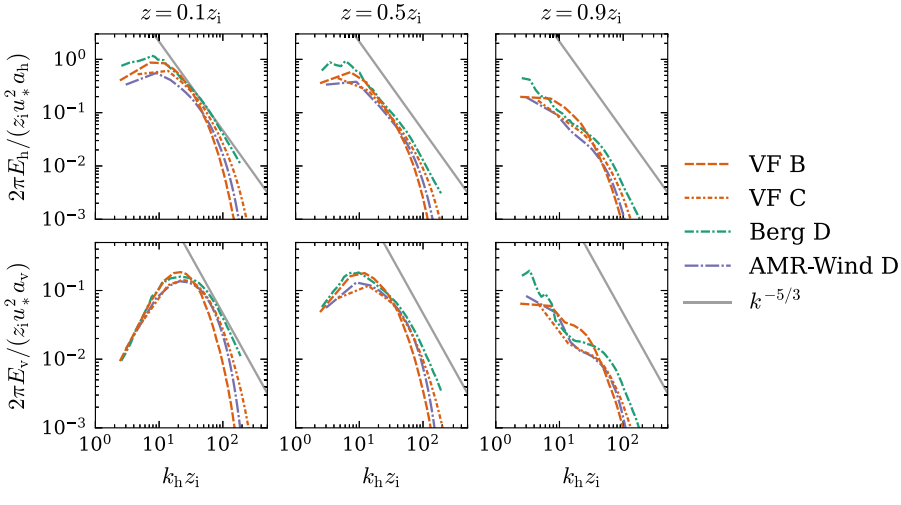

Figure 4.13. Horizontal (top) and vertical (bottom) velocity spectra at $z = 37\,\mathrm{m}$ (left), $z = 186\,\mathrm{m}$ (center), and $z = 333\,\mathrm{m}$ (right)

p. 273ff] which leads to smaller force magnitudes relative to velocity with increased resolution.

We show Reynolds stresses and a surrogate for turbulence intensity, see [6], in Figure 4.12. Reynolds stresses agree very closely among all compared solvers. The surrogate turbulence intensity obtained from our results exhibits a less pronounced bump at the inversion layer.

Figure 4.13 shows velocity spectra obtained at three heights in the domain, see paper V for a detailed description of the computation procedure. The spectra agree well amongst all solvers at low wave numbers. At higher wave numbers the difference in the numerical methods become visible. Berg et al. employ a pseudo-spectral solver, resulting in the lowest decay at high wave numbers. AMR-Wind employs a second order accurate scheme, while the cumulant LBM is fourth order accurate in diffusion. Therefore, we find that

Table 4.3. *Boundary layer height, friction velocity, and buoyancy flux of GABLS1 test case compared to references.*

	VF		Beare et al.		Gadde et al.	
Quantity	2 m	1 m	2 m	1 m	Smag.	LASD
h/m	182	158	162 – 197	149 – 164	166	166
$u_*/\mathrm{m/s}$	0.28	0.26	0.24 – 0.28	—	0.265	0.253
$-\frac{g\langle w\theta\rangle}{\theta_0 \cdot 10^4}/\mathrm{m^2/s^3}$	4.08	3.59	3.5 – 4.7	—	4.1	3.8

VirtualFluids yields a more accurate decay at high wave numbers than AMR-Wind, despite lower resolution.

The simulations are conducted on a single NVIDIA RTX 6000 GPU. Simulation of grid B and C require 3.4×10^4 s and 5.1×10^5 s of wall time respectively, resulting in a realtime factor of 3.6 and 0.25 respectively.

4.4 Results of a Stably Stratified Boundary Layer

We further validate the model against the commonly used stably stratified GABLS1 benchmark case presented by Beare et al. [21]. We compare to the results from the original benchmark and additional results from Gadde et al. [43] using two different SGS models, the Smagorinsky model (denoted Smag.) and the Lagrangian Average Scale Dependent (LASD) model [24].

Case Description

The domain is bi-periodic with dimensions $400\,\mathrm{m} \times 400\,\mathrm{m} \times 400\,\mathrm{m}$, geostrophic wind in streamwise direction is set to $8\,\mathrm{m/s}$ and the Coriolis parameter is $f^{\mathrm{C}} = 1.39 \times 10^{-4}\,1/\mathrm{s}$. The domain is initialized with uniform geostrophic wind and a two-layered temperature profile, the lower 100 m are initialized with $\theta_0 = 263.5\,\mathrm{K}$, while the upper part of the domain is initialized with a constant temperature gradient of $\frac{\partial \theta}{\partial z} = 10\,\mathrm{K/km}$. At the top boundary a slip boundary is applied, while a constant cooling rate of $0.25\,\mathrm{K/h}$ and a roughness length of $z_0 = 0.1\,\mathrm{m}$ are applied at the bottom wall using the surface layer boundary condition. The reference temperature is $\theta_\mathrm{r} = 265\,\mathrm{K}$. A total of 9 h is simulated with statistics collected over the last hour. We conduct two simulations with $\Delta x = 2\,\mathrm{m}$ and $\Delta x = 1\,\mathrm{m}$.

Results

Quantities of interest of the simulation are gathered in Table 4.3. Here the boundary layer height h is computed as $h = \frac{h_{95}}{0.95}$ with h_{95} being the height at which the shear stress is less than 95 % of the wall shear stress. Again we find that the boundary layer height decreases with increasing resolution and that VirtualFluids agrees reasonably close with the references. The friction

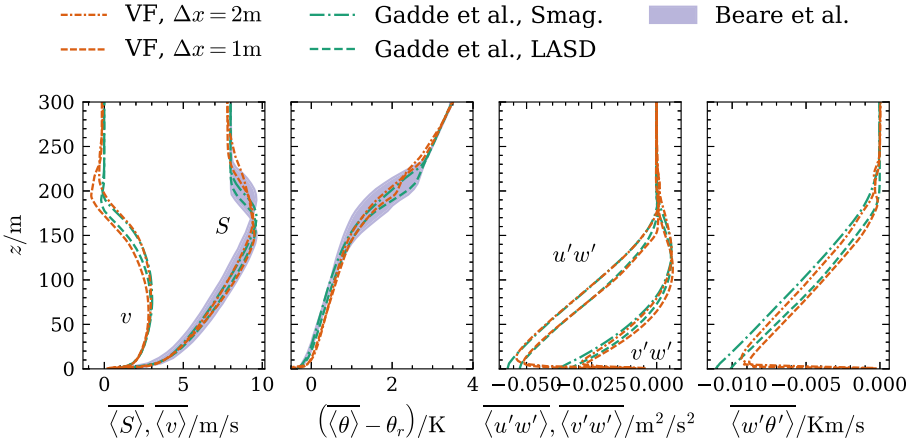

Figure 4.14. Profiles of wind speed, lateral velocity and fluxes of the GABLS1 case

velocity and buoyancy flux also agree well with the references, our results are always within the span of the different results from the original benchmark.

Profiles of mean wind speed, lateral velocity, and temperature, as well as vertical momentum and temperature fluxes, are shown in Figure 4.14. Within the boundary layer, the profiles agree well with the references. At the top of the inversion layer our results show a stronger lateral velocity and higher wind speed, and the temperature exhibits a double inversion layer profile. The double inversion layer leads to weaker suppression of vertical motion in the inversion layer, resulting in weaker velocity gradients above the inversion layer. This could possibly be addressed by enhancing the accuracy of the collision operator in a similar manner as was done by Gruszczyński and Łaniewski-Wołłk [52] for the two-relaxation time cascaded model. However, it is noteworthy that there is already a large variation in the reference results and the results of our model fall within the range of the references in most cases, therefore we expect little effect of these inaccuracies.

4.4.1 Conclusion

Many phenomena important in wind energy are related to thermal stratification. In this work, we develop, implement and validate a method for simulating thermally stratified boundary layers using a double-distribution function LBM. We develop a novel combined wall model boundary condition to apply Monin-Obukhov similarity theory to momentum and temperature. The model is validated in simulations of conventionally neutral and stable boundary layers. The results agree well with the reference cases in the boundary layer. The excellent spectral properties of the LBM in the neutral case are particularly noteworthy. Small differences are evident at the top of the inversion layer and the free atmosphere. This work represents an important step towards

comprehensive simulations of boundary layers and wind farms using the LBM. Although the model in this work has not yet been optimized for computational performance, it already achieves above real-time performance for the CNBL case at $\Delta x = 7\,\mathrm{m}$ and a real-time factor of around 0.25 at a resolution of $\Delta x = 3.5\,\mathrm{m}$. Further improvements to the collision operator and computational performance are expected to increase accuracy and speed even further.

5. Applications

We applied the LBM to two different cases of wind energy research to demonstrate its benefits. The high computational efficiency of the LBM allows researchers to simulate a larger parameter space, which we make use of in two different ways. In paper VI, we create a set of training data that we then use to train a convolutional neural network to predict average flow quantities. However, higher computational efficiency can also be advantageous to traditional research question, which we address in paper VII. Here, we want to characterize the flow in the wake of a wind turbine using the helix approach and understand the interaction of multiple turbines.

5.1 A Convolutional Neural Network for Wake Prediction

In certain applications in wind energy, such as layout optimization, hundreds or thousands of simulations of wind farms have to be conducted, as indicated by the results of paper II. With currently available technology, it is unlikely that LES will be efficient enough to perform such computations. However, recent years have shown that machine learning algorithms can provide extremely fast running models for a wide range of applications in fluid dynamics [153]. In contrast to classical engineering models, machine learning models do not rely on the fine tuning of single parameters. Most machine learning models rely on a large volume of training data, which must be generated by a higher-fidelity method. Furthermore, the larger the parameter space covered by the training data, the broader the applicability of the learned model. A promising class of machine learning models are convolutional neural networks (CNNs). A CNN maps a multi-layered input grid to a multi-layered output grid. The traditional encoder-decoder architecture first encodes a large input map to a smaller map containing more layers, encoding larger features. The decoder then maps these large features back to a larger grid with fewer layers. CNNs usually require to learn far fewer parameters than fully connected networks, which makes them well suited to large, gridded input data, such as images, videos, or fluid flows. We develop a model based on convolutional neural networks called WakeNet and generate a set of training data of a single turbine wake in paper VI.

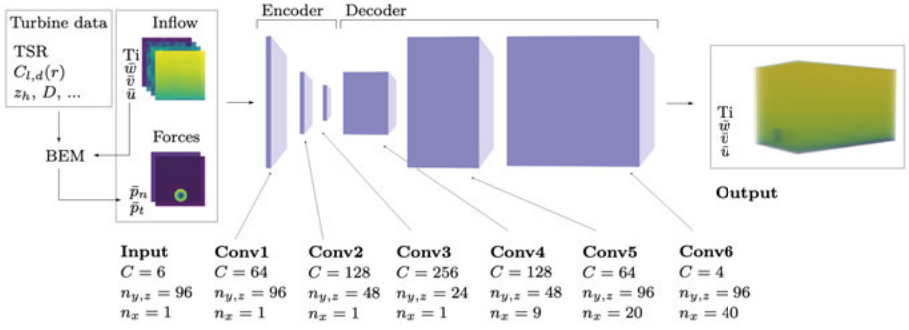

Figure 5.1. Schematic of WakeNet including inputs, architecture of the convolutional neural network and output.

5.1.1 A Convolutional Neural Network to Predict Wind Turbine Wakes

The goal of the network is to predict the time-averaged three-dimensional flow field in the wake of a single turbine, based on a plane of inflow and turbine data, as well as operational parameters, such as rotor speed.

The input of the model consists of a 2D plane of averaged velocity, $\overline{u}, \overline{v}, \overline{w}$, and Ti. Additionally, we use a BEM model to compute normal and tangential force densities p_n and p_t based on turbine input data comprising airfoil data, hub height and tip-speed ratio (TSR). Velocities and forces are normalized with u_h and $\rho u_\mathrm{h}^2 D$, respectively.

The combined input data is then fed to a fully convolutional neural network with an encoder-decoder architecture, mapping the 2D input field of size $n_y = n_z = 96$ and six feature layers or channels to a three dimensional output with size $n_x \times n_y \times n_z = 40 \times 96 \times 96$ and four output channels $\overline{u}, \overline{v}, \overline{w}$ and Ti. The output spans $10\,D \times 4\,D \times 4\,D$. An overview of the model is shown in Figure 5.1. Note that the dimensions of the input and output plane are the same, so that instances of the model can be chained together to simulate a whole wind farm. Further details on the network architecture can be found in the original publication and the model is openly available at Asmuth and Korb [11].

5.1.2 An Open-Source Dataset for Training of Wind-Turbine Wakes

We generate a dataset for training based on simulations of a single NREL 5 MW turbine using the GPU-resident LBM-LES solver elbe [61]. The solver is also based on the cumulant LBM and uses the AMD SGS model and an ALM implementation very similar to that of WiFI. We simulate a domain of $14\,D \times 6\,D \times 6\,D$ with $\Delta x = D/24$. The inflow boundary condition is a velocity boundary with logarithmic inflow profile superimposed with synthetically

Table 5.1. *Simulation parameters that are varied in the training dataset.*

Quantity	Unit	Values
λ		5, 6, 7, 7.5, 8
z_h	D	0.71, 0.95, 1.19
u_h	m/s	5, 7.5, 10, 12.5, 15
z_0	m	1×10^{-3}, 1×10^{-2}, 1×10^{-1}
Ti	%	1, 5, 10, 15

generated turbulence using the Mann model. The bottom boundary employs the iMEM boundary with varying roughness length and lateral boundaries are periodic. The turbine is located at $x = 3D$. The inflow plane used as input to the network is sampled 2 diameters upstream of the turbine. Five parameters vary between the cases: turbine hub height z_h, TSR λ, mean hub-height velocity u_h, roughness length z_0 and turbulence intensity Ti. The values are summarized in Table 5.1 and simulations are conducted for all possible combinations of values, leading to a total of 900 simulations. The time-averaged data is available at Korb and Asmuth [71].

5.1.3 Results

Training

We first analyze the volume of training data required, to determine whether we generated enough training data and whether all 900 cases are necessary. The available dataset is split into a training and a validation set to evaluate the trained model on data that was not used during the training. We compare training with 10 % to 80 % of the dataset and find that especially prediction of $\overline{u}$ requires the full dataset to reduce the error. We measure the accuracy of the network as

$$L_2 = \sqrt{\frac{\sum_{i=1}^{N}\left(Y_j^{\mathrm{WN}}(x_i) - Y_j^{\mathrm{LES}}(x_i)\right)^2}{\sum_{i=1}^{N} Y_j^{\mathrm{LES}}(x_i)}} \tag{5.1}$$

for each case j, where N is the number of points in the output, $Y^{\mathrm{WN}}(x)$ is the output of the network at x and $Y^{\mathrm{LES}}(x)$ is the result from the simulation. The lowest L_2 achieved is 0.56 % for $\overline{u}$, while the error in $\overline{v}$ and $\overline{w}$ is 27.95 % and 27.41 %, respectively. The error for Ti is 17.01 %. We attribute this large difference in error to a better normalization of $\overline{u}$, as the other errors are about the same size in absolute terms. Thus, the results could potentially be improved by a more appropriate normalisation scale. However, the accuracy of prediction of the mean streamwise velocity is noteworthy.

The difference between solution of the LES and the prediction for an individual case of streamwise velocity and turbulence intensity at three locations

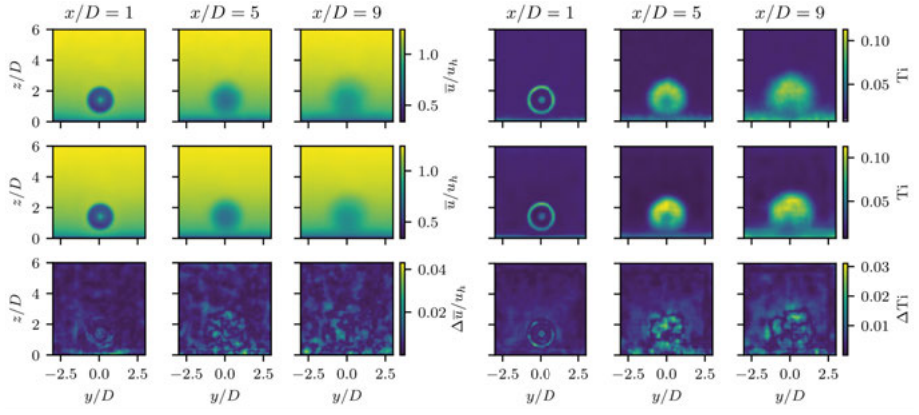

Figure 5.2. Contour plots of the LES data (top rows), model prediction (mid rows) and respective difference (bottom row) of $\overline{u}/u_\mathrm{h}$ (left) and Ti (right) with $\lambda = 7$, $z_\mathrm{h} = 0.95\,D$, $u_\mathrm{h} = 7.5\,\mathrm{m/s}$, $z_0 = 0.1\,\mathrm{m}$ and $Ti = 5\,\%$.

is shown in Figure 5.2. The differences in streamwise velocity are at most 5 % of hub height velocity and are scattered throughout the flow field. The characteristics of the turbine wake are visible at all shown locations. The turbulence intensity is also captured well in its general characteristic, however the relative error is significantly larger and more concentrated around the area of the wake.

To see how well the model predicts in different operating conditions, we show profiles of all predicted quantities at different λ in Figure 5.3, comparing the predictions by WakeNet and the LES. The characteristics of the mean streamwise velocity are excellently captured by the model, such as transition from near to far wake and dissipation of the wake, regardless of TSR. Errors in the other velocity components are of a similar magnitude as the streamwise velocity. Turbulence intensity is slightly overpredicted in the far wake in all cases shown. However, the shape and distinctive changes from near to far wake are well-captured.

To analyze the spatial variation of the model prediction, we compute the local mean absolute error

$$MAE(\boldsymbol{x}) = \frac{1}{N_\mathrm{test}} \sum_{i}^{N_\mathrm{test}} \left| Y_i^\mathrm{WN}(\boldsymbol{x}) - Y_i^\mathrm{LES}(\boldsymbol{x}) \right|, \tag{5.2}$$

where N_test is the number of test cases, and show it in Figure 5.4. The highest error in all velocity components is found in the near wake at $x = 1\,D$ in the area of the wake. Nevertheless, the largest errors remain below 2 %, and they decrease with increasing distance to the turbine. The turbulence intensity shows the highest errors around $x = 5\,D$. We attribute this to the fact that this is the area with highest variability, since the location of wake transition, which occurs around $x = 5\,D$, depends on thrust of the turbine and turbulence intensity.

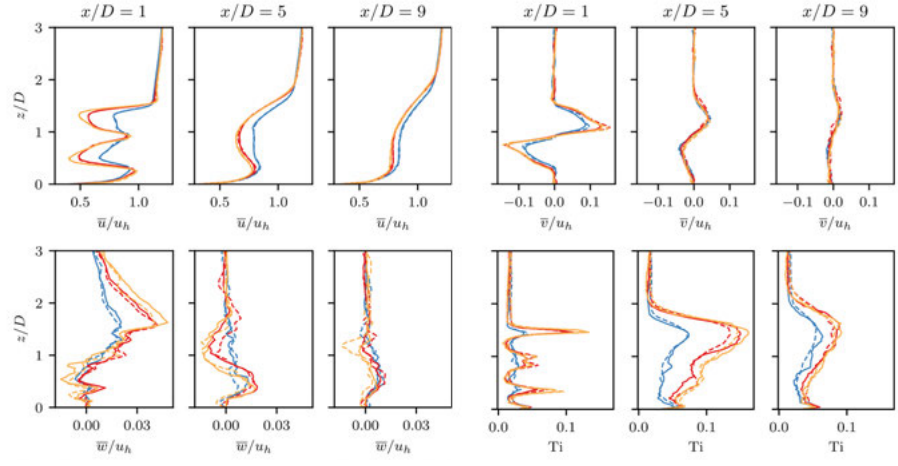

Figure 5.3. Vertical profiles of all predicted quantities. Predictions by WakeNet (full lines) compared against the corresponding LES results (dashed lines). All cases refer to the same inflow conditions as shown in Figure 5.2. TSR varies between 5 (blue), 7 (red) and 8 (orange).

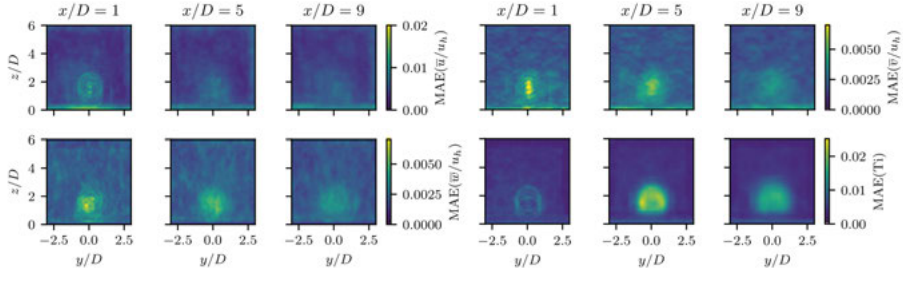

Figure 5.4. Mean absolute error of velocity and turbulence intensity prediction across the test data set.

5.1.4 Conclusion

The machine learning model developed in this work excellently predicts average quantities of the flow behind a single turbine. Furthermore, its execution speed on a single GPU is on a similar scale as that of a wake model while containing significantly more information. However, further work will be required to extend this model for the prediction of wind-farm flows. Particularly, new training data with waked inflows will have to be generated which drastically increases the number of parameters. Nevertheless, it shows how promising this type of architecture is for flow field prediction. This work can be taken as an example of how higher computational efficiency of LES enables higher accuracy of lower order models. By creating larger datasets, larger parameter spaces can be covered, and in turn accuracy of the lower order models can be improved. A similar approach could also be applied for engineering models, where the models are calibrated on a number of LES simulations of a specific wind farm, turbine or flow condition and then applied to, e.g., layout optimization.

5.2 Characterizing the Wake of the Helix Approach

The application of the lattice Boltzmann method for wind energy is valuable not only because it enables industrial application, but also because it accelerates research of "classical" questions. One of the oldest questions in wind energy is how to operate wind turbines to reduce wake effects. Wind turbines are typically clustered in wind farms, where the wake of one turbine impinges upon downstream turbines. The wake is an area of reduced velocity and increased turbulence intensity, decreasing power production and increasing fatigue loads on downstream turbines. Consequently, reducing wake effects by altering wind turbine control has been of interest for nearly 30 years [137]. Such an approach is called wind-farm flow control, as outlined in subsection 1.1.3. The novel class of individual pitch control has attracted considerable attention in recent years. We examine some of the fundamental aspects of the fluid dynamics of the *helix* approach for the wake of a single turbine as well as the interaction of multiple helical wakes in paper VII.

5.2.1 Methodology

To examine the fluid dynamic properties of the wake of turbines applying the helix approach, we conduct a range of LES simulations. We perform two highly resolved LES simulations to examine the effect of the helix approach on the tip vortices, and eight medium-resolved simulations to examine deflection, entrainment, and wake shape under varying inflow conditions. Additionally, we examine the interactions of three helical wakes in five simulations of three turbines. We simulate the NREL 5MW turbine using the LBM framework ELBE. The domain for simulations of a single turbine has dimensions $15\,D \times 8\,D \times 8\,D$ and grid spacing $\Delta x = D/32$, with a refined zone around the turbine of size $13\,D \times 6\,D \times 6\,D$ and a spacing of $\Delta x = D/32$. The grid spacings are the same for simulations of three turbines, but coarse and refined domains are enlarged to $18\,D$ and $16\,D$ in the streamwise direction. The lateral boundaries are treated as slip conditions and the inlet prescribed as a constant uniform inflow superimposed with synthetic turbulence. We examine the effect on a single wake of three inflow parameters, the inflow velocity V_0, turbulence intensity Ti, and turbulent length scale L by varying each parameter. Furthermore, each simulation is conducted with and without the helix. The simulations are summarized in Table 5.2.

By simulating three turbines in a row we can investigate the effect of the phase difference between the helices, denoted as Φ. To that end, we conduct five simulations, one without the helix, serving as a baseline case, and four where the helix is applied at all turbines but with differing phase differences. The cases are listed in Table 5.3.

Details on the control of these turbines can be found in paper VII.

Table 5.2. *Simulation parameters of single turbine cases*

Name	Helix	V_0 / m/s	Ti / %	L / m	λ
G1V9 / G1I5 / G1L40	No	9	5	40	7.74
G1V13	No	13	5	40	6.14
G1I10	No	9	10	40	7.74
G1L120	No	9	5	120	7.74
H1V9 / H1I5 / H1L40	Yes	9	5	40	7.74
H1V13	Yes	13	5	40	6.14
H1I10	Yes	9	10	40	7.74
H1L120	Yes	9	5	120	7.74

Table 5.3. *Simulation parameters of three-turbine cases*

Name	Helix	Φ	V_0 / m/s	Ti / %	L / m	λ_0	λ_1	λ_2
G3	No	-	9	5	40	7.70	6.15	6.11
H3PS0	Yes	0°	9	5	40	7.70	6.49	6.50
H3PS90	Yes	90°	9	5	40	7.70	6.49	6.35
H3PS180	Yes	180°	9	5	40	7.70	6.48	6.18
H3PS270	Yes	270°	9	5	40	7.70	6.48	6.36

5.2.2 The Helical Wake of Single Turbine

The Origin of the Helical Wake

The helix approach can be derived from the definition of a global pitch vector $\mathbf{\Phi}$ that is prescribed to rotate with a frequency ω_e and amplitude A

$$\mathbf{\Phi} = A \begin{bmatrix} 0 \\ \sin \omega_e t \\ \cos \omega_e t \end{bmatrix} \tag{5.3}$$

Transforming this global pitch to the blade local frame yields the pitch angle ϕ_b of the blade b

$$\phi_b = A \sin (\Psi_b + \omega_e t). \tag{5.4}$$

We can define a corresponding Strouhal number $\mathrm{St} = \frac{\omega_e D}{2\pi V_0}$. From the blade pitch and a linearization of the lift coefficient we then compute the force of the rotor on the flow to be

$$\boldsymbol{F} \propto -N_\mathrm{b} \sum_{i=1}^{N_\mathrm{b}} \begin{bmatrix} \cos(\phi)\, C_\mathrm{L}(\alpha_0) \\ C_1 \sin(\omega_e t) \\ C_1 \cos(\omega_e t) \end{bmatrix}, \tag{5.5}$$

where $C_1 = A(\partial C_\mathrm{L}/\partial\alpha)(sin(\phi)/2)$, ϕ is the inflow angle, α is the angle of attack, and N_b is the number of blades. We then postulate that the helical shape of the wake is caused by the radial component of that force due to a rotating

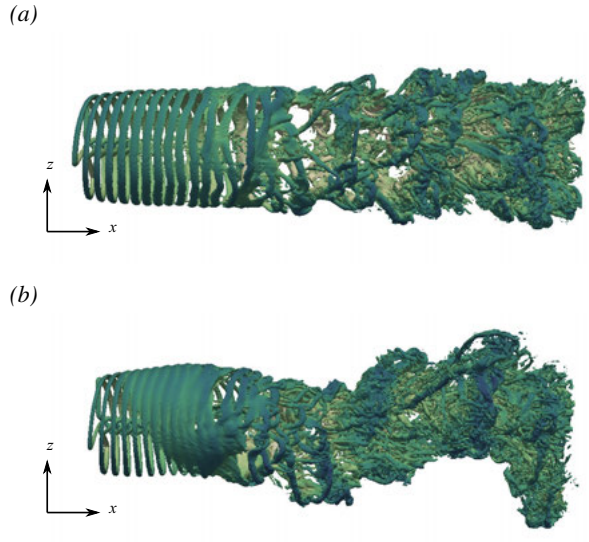

Figure 5.5. Tip spirals in the wake of a turbine with greedy (a) and helix (b) control visualized as vorticity magnitude iso-contours.

deflection due to conservation of momentum. Later research by Coquelet et al. [29] has challenged this assumption, noting that the radial force only has a magnitude of 1 % of the thrust force. Instead, they attributed the wake deflection to the system of vortices caused by imbalanced forces. Gutknecht et al. [53] also highlights the importance of the imbalance of forces, in line with Coquelet et al. Nevertheless, the observations in the following still remain valid.

Tip Vortices

We first examine the effect of the helix approach on the tip vortices via a qualitative assessment of the high-resolution cases. The tip vortices are identified via iso-contours of the vorticity magnitude in Figure 5.5. Tip vortices play a crucial role in wake recovery as they shield the wake from entraining kinetic energy, as discussed, for example by Lignarolo et al. [88]. In the greedy case, we observe the characteristic sequence of pairing, leapfrogging and break up, as reported by Tophøj and Aref [144] and Sarmast et al. [126]. The breakdown of the helix wake looks markedly different. Due to the varying pitch angle, we observe tip vortices of different strength. In regions of strong vorticity, the tip vortices fuse, followed by a more violent breakdown. Furthermore, we can clearly see the helix-shaped wake.

Wake Deflection

As is evident from Figure 5.5, the helix approach effects the position of the centreline of the wake. We define three terms to discuss the geometry of

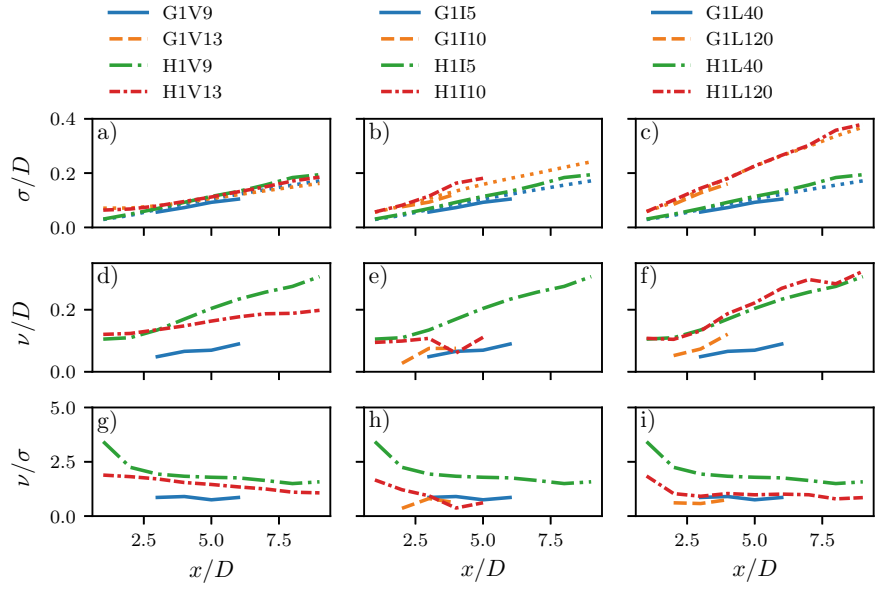

Figure 5.6. Estimated parameters of the Rice distribution of wake centre distances.

the wake: wake deflection, wake meandering and wake deformation. Wake deflection refers to the displacement of the mean wake centre due to forces exerted by the turbine, while wake meandering denotes the fluctuation of the wake centre around the mean center. Wake deformation refers to a change in shape of the wake without altering the mean wake center. Importantly, here we define the mean as the mean in the frame of reference rotating with the helix. We identify the wake centre using the SAMWICH library [120]. The wake centre distance d is the distance of the centre of the rotor plane to the wake center. Due to the combined effect of meandering and deflection, the wake centre distance follows a Rice distribution

$$f(d|\nu,\sigma) = \frac{d}{\sigma^2}\exp\left(\frac{-\left(d^2+\nu^2\right)}{2\sigma^2}\right) I_0\left(\frac{d\nu}{\sigma^2},\right) \tag{5.6}$$

where f is the probability density function with the two parameters ν and σ and I_0 denotes the modified Bessel function of the first kind of order zero [121, 122]. The parameters ν and σ characterize deflection and meandering, respectively. We determine these parameters from the sampled wake distances using the method of Koay and Basser [69]. The results are shown in Figure 5.6.

We find that the effect of the helix on wake meandering is small in most cases, while turbulence intensity and turbulence length scale have a much larger effect. However, we can clearly observe a strong deflection of the wake at low turbulence intensities. At high turbulence intensities, the algorithm for

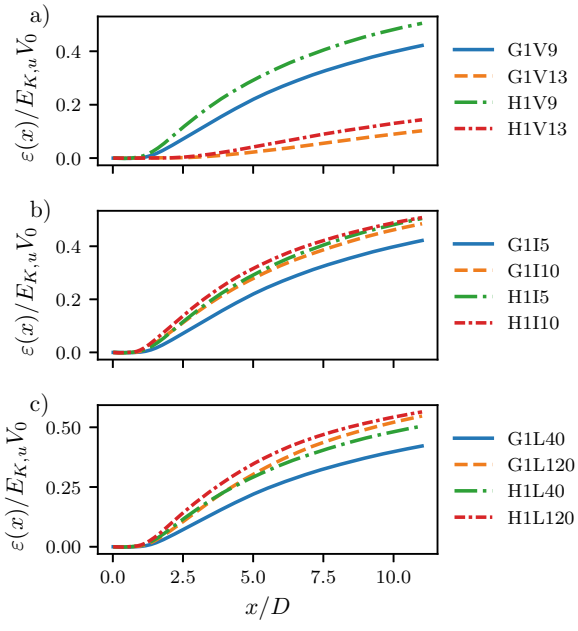

Figure 5.7. Cumulative entrainment into the rotor area. Normalized with streamwise transport of mean kinetic energy.

determining the parameters did not always converge, which occurs when the meandering is much stronger than the deflection. At higher inflow velocity, the thrust coefficient of the turbine is lower and so is the deflection of the wake center. Nevertheless, in the baseline case, the wake is deflected almost half a radius, thus reducing overlap with downstream turbines and reducing the effect of the wake.

Entrainment

In addition to deflection, the helix approach also improves wake mixing, thereby reducing of the wake deficit. Indeed, this is usually regarded as the main effect of the helix. We quantify the effect of improved wake mixing with the total transport of kinetic energy into the wake, the entrainment

$$\epsilon(x) = -\int_{x_0}^{x} \iint_{A_R} \frac{\partial}{\partial r} \left[\overline{u}_i \overline{u_i' u_r'} \right] \mathrm{d}A\mathrm{d}x, \tag{5.7}$$

where A_R is the rotor area. The results are shown in Figure 5.7, normalized with the kinetic energy of the undisturbed inflow $E_{K,u} = 1/2V_0^2 A_R$ and the inflow velocity.

The helix approach significantly increases entrainment in every case, however, the strength of the effect depends on the inflow conditions. It is most

pronounced in the baseline case, while higher turbulence intensity and turbulent length scale lead to higher entrainment regardless of the control strategy. The overall highest entrainment is found at large length scales and with the helix. At higher inflow velocity, we once again find a decreased effect. The helix influences the breakdown of tip vortices, leading to an earlier transition and, consequently, and earlier increase in entrainment. This especially effects the area $1\,D$ to $6\,D$ downstream of the turbine, crucially, where a downstream turbine would be typically placed.

Wake Shape
We examine the wake shape in the helix-averaged frame of reference and find that the wake has a kidney shape, similar to that observed in wake steering, see Bastankhah and Porté-Agel [19]. The shape arises from two vortices in the wake, however, unlike in the wake steered case, these vortices have the same polarity. While only mentioned briefly in this work, it was later found that this vortex pair is of importance for the efficiency of the helix approach by Coquelet et al. [29].

Kinematic Description of a Single Helical Wake
We develop a simple description of the angle of the direction in which the helix deflects the wake, that will be used to understand the interactions of multiple helices. The direction of the helix is $\varphi(x,t)$. Under the assumption that the helix angle is transported downstream with a constant transport velocity $\overline{u}_{helix}$, we find that

$$\varphi(x,t) - \varphi(x_0,t) = 2\pi \mathrm{St} \frac{x - x_0}{D} \frac{V_0}{\overline{u}_{helix}} = \Delta\varphi(x), \tag{5.8}$$

which is constant with time and only depends on the ratio of transport to inflow velocity. We fit this analytical expression to a measured phase difference and find that $\overline{u}_{helix} = 0.7V_0$ is a good fit for the angle at $x = 5D$ for all cases except that of higher inflow velocity.

5.2.3 Multiple Helices

The helix has usually only been applied to a single turbine, see for example Taschner et al. [141], Coquelet et al. [29], and Brown et al. [26]. However, the interactions of multiple helices can significantly effect the efficacy of the helix of the downstream turbine. Van Vondelen et al. [150] develop a Kalman filter-based controller to synchronize the helix phases, however, they do not assess whether this approach actually leads to the highest power gains. We aim to determine the phase difference a priori based on (5.8). We find that by introducing a controllable phase difference Φ_1, and adjusting (5.4) to

$$\phi_1 = A \sin\left(\Psi_b + \omega_e t - (\Phi_1 + \Delta\varphi_0(S_x))\right), \tag{5.9}$$

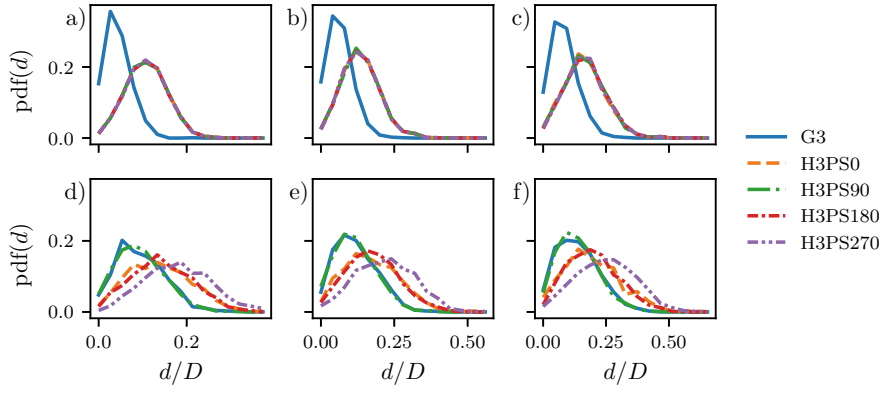

Figure 5.8. Distribution of wake centre distances at 2 D, 3 D, and 4 D downstream of turbines 0 (a-c) and 1 (d-f).

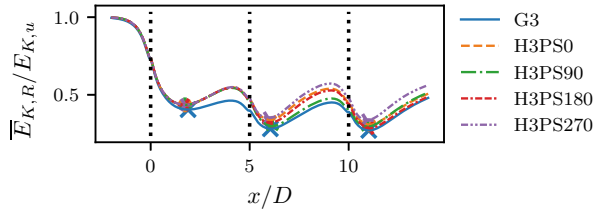

Figure 5.9. Mean kinetic energy in the rotor area. Local minima and turbine locations marked by markers and dotted black line, respectively

where the subscripts denote the index of the turbine starting with the upstream turbine at index 0, we can a-priori set the phase difference between the helices of turbines 0 and 1.

Wake deflection

We examine the effect the phase shift has on the distribution of wake centre distances in Figure 5.8. We see that, compared to the greedy case, the wake centres are deflected in the wake of the first turbine and the phase difference has no influence as expected. However, at the second turbine, we observe a strong dependency of the distribution on the phase difference. In the case H3PS90, the distribution closely resembles the greedy case. Cases H3PS0 and H3PS180 are very similar, with a medium mean deflection, while case H3PS270 shows the strongest deflection. Overall, we find that the spread of the wake centre distances is larger, due to the increased incoming turbulence.

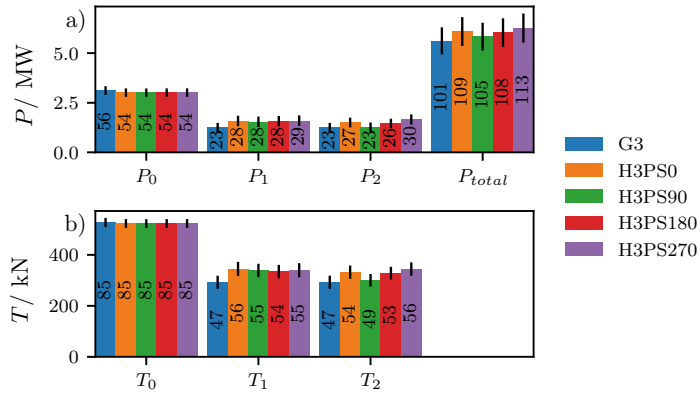

Figure 5.10. Average power and thrust of all turbines. Numbers in bars are mechanical power relative to mean power of the undisturbed inflow and thrust relative to mean force on the rotor disk in the undisturbed inflow, in per cent.

Kinetic Energy in the Wake

The differences in wake deflection are also observable in the mean kinetic energy in the rotor area $\overline{E}_{K,R}$ behind each turbine, shown in Figure 5.9. In the wake of the first turbine, the helix approach has the same effect on all turbines, with an earlier onset of recovery and significantly higher mean kinetic energy at the second turbine compared to the greedy case. However, in the wake of the second turbine we observe the same trends as for Figure 5.8. Case H3PS90 is almost equal to the greedy case, while the highest kinetic energy can be found for case H3PS270. In the wake of the third turbine, case H3PS180 yields the lowest improvements, while H3PS270 is still the best performing case. Nevertheless, the difference between H3PS2670 and G3 is significantly lower in the wake of turbine 2, indicating that the helix approach is less effective deeper in the farm, which is to be expected due to its decreased efficiency in highly turbulent inflows.

Total Power and Thrust

The increase in mean kinetic energy corresponds directly to increases in power and thrust, as shown in Figure 5.10. While there are small losses at the first turbine, the power production at the first downstream turbine increases compared to the greedy cases for all phase differences. However, already at this turbine small differences occur, with H3PS270 yielding slightly more than the other cases. The differences are more pronounced at the second downstream turbine, where H3PS90 produces the same as the greedy case, while H3PS270 produces about 30 % more than the greedy case. Overall, power production increases in all cases. However, while the case H3PS90 produces 5 % more

than the greedy case, the best case of H3PS270 produces 12 % more. The effects on the thrust follow the same trends as the power production.

5.2.4 Conclusion

In paper VII, we examine various aspects of the wake produced by the helix approach, including wake deflection, wake mixing and the application of the helix to multiple turbines. We find that wake deflection is a significant contributor to the power gains observed for the helix approach, in addition to the added wake mixing. We develop a simple model for the direction of deflection of the wake and apply that model to understand the interactions of multiple helices. The findings of the model are augmented with simulations of a row of three turbines and find that in contrast to the model, the highest power gains are found for a phase difference of 270°.

A number of papers on the helix approach were published after this work, some questioning the findings we have made. In particular, our model for the origin of the helical wake has been questioned and subsequently improved in later work by Coquelet et al. [29] and Gutknecht et al. [53]. Furthermore, a more in depth analysis in paper VII of phase differences comes to the conclusion that the direction of the force relative to the direction of the incoming helix is the cause for the difference in sustaining the helix. In light of the findings by Coquelet et al. [29] it is questionable, whether this explanation is a correlation or really the cause. Further research in this direction is required for fully understanding the interaction of multiple helices.

Although not directly focused on the LBM, the work in this paper highlights its potential as a conventional CFD tool for wind energy research. The reduction in computational cost enables researchers to work more efficiently, conduct larger or more numerous simulations, and shorten development time for new models. It also enables more researchers to do cutting edge research, as the hardware costs of a consumer-grade GPU are much lower than that of a HPC cluster.

6. Conclusion

The lattice Boltzmann method (LBM) is a promising approach for enabling cost-efficient Large Eddy Simulation (LES) of wind-farm flows at high fidelity, rendering it attractive to both research and industry. This thesis expands the applicability of the LBM by identifying knowledge gaps in the literature and establishing requirements for industrial use; validating, improving, and extending an open-source software framework for high-fidelity wind-farm flow simulations; and demonstrating its relevance for low-cost model development and wind-farm flow control research.

Many of the current and urgent questions in wind energy research depend on high-fidelity flow simulations. Meanwhile, industry relies on simplified models that are incapable of reproducing the complex interactions between wind farms and the atmosphere. To reduce the computational cost, new solvers have been developed that use graphics processing units (GPUs) for massively parallel computations. However, traditional solvers that discretise the Navier-Stokes equations are ill-suited to such hardware. By contrast, the LBM, is excellently parallelisable and has been employed on GPUs for more than 20 years. It was recently introduced to wind energy research, exhibiting considerable potential for decreasing computational cost and time-to-solution [16, 129, 14]. However, previous studies have focused on highly idealised cases. In this work, an open-source framework is developed, comprising two main components: the open-source GPU-resident LBM solver VirtualFluids and the wind-farm simulation tool Wind Farm Interface (WiFI).

The available literature relevant to LBM for wind-farm flows is reviewed in paper I. The vast majority of methods required for a general-purpose wind-farm flow solvers are available now; however, no currently available framework combines all of them. Many of the required methods have been applied in numerous studies and the LBM has repeatedly been shown to be accurate and suitable for large-scale applications. Previous applications indicate that computational performance doubles roughly every 2.5 years, and performance exceeding real-time is already attainable with current hardware.

In paper II, a survey is conducted to understand barriers to the adoption of LES for industrial use. It is found that computational cost and runtime are the primary factors deterring industrial users, while lack of expertise and validation are also significant. Responses regarding acceptable runtime for a representative case varied from seconds to days, with around half of the respondents satisfied with real-time performance. For comparison, a simulation

of a wind farm comprising 64 turbines is conducted using the framework developed in this thesis, demonstrating that real-time performance can be achieved on currently available hardware.

Paper III presents a validation of the newly developed framework against measurement data and other solvers from the SWiFT benchmark [34]. The LBM-based model is at least as accurate as other methods for predicting turbine response and wind turbine wakes. A comparison of energy-to-solution demonstrates that the solver is more than an order of magnitude more efficient than the next most efficient LES solver, and even more than two orders of magnitude more efficient than the most expensive solver, while achieving comparable accuracy. However, it is also one to two orders of magnitude more expensive than lower fidelity models.

The actuator line method used to model the rotor has a well-known shortcoming of overpredicting the forces near the root and the tip due to the smearing of forces in the fluid domain. Several models for alleviating this shortcoming have been proposed. In paper IV, the models proposed by Martínez-Tossas et al. [99] and Meyer Forsting et al. [101] are compared. The effects of grid resolution and smearing width on rotor response, wake and computational time are assessed, as well as the performance of the corrections for the error introduced by the smearing. Computational time is drastically reduced with reduced resolution, while accuracy in rotor response can be increased by using either of the two compared corrections. However, the correction by Meyer Forsting et al. performs better but also incurs greater computational overhead. The wake is affected primarily by resolution and smearing width, while the corrections have a limited effect. Finally, the impact of the corrections on a wind farm is examined. The corrections affect turbines differently, meaning power predictions cannot be corrected by a constant factor.

In paper V, VirtualFluids is extended to include thermal stratification by solving a second transport equation for temperature. A new wall boundary condition to apply Monin-Obukhov similarity theory (MOST) is developed, as well as Coriolis and buoyancy force and damping layers are implemented. The model is validated in conventionally neutral and stably stratified boundary layers. The model agrees closely in most examined quantities, however, marginal disagreement in the veer is found in the inversion layer. The newly developed boundary condition performs well in both conditions. While the computational cost increases two-fold compared to isothermal simulations, the model is still significantly more computationally efficient than traditional solvers.

Papers VI and VII apply the LBM to typical questions of wind energy research. In paper VI, a predecessor of the model developed in this thesis is used to generate a large dataset to train a machine learning model. The model is trained to predict the flow field in the wake of a single turbine and agrees excellently with test data. Predictions can be generated in a few hundred microseconds. The properties of wakes of turbines employing the helix approach to mitigate wake effects are examined in paper VII. Analytical

descriptions of the wake are given, and the effect of the helix approach on wake deflection, wake mixing, and wake shape is characterised. Furthermore, the interactions of multiple wakes is studied. The phase difference between helices is found to be a crucial parameter for efficiently applying the helix to a row of turbines.

Further validation of the model is required, particularly for larger wind turbines and farms. However, as of today, no validation data for such a case is publicly available. Additionally, while the current framework possesses more capabilities than other LBM solvers, there are still features missing. This includes the ability to model complex terrain and forests, couple to mesoscale models, and simulate wind farms with thermal stratification.

In summary, this work advances the LBM for wind energy from a promising method to being ready for wide-spread use in both industry and research. The developed framework enables highly efficient and accurate computations of flows in large wind farms in, given some additional implementations, stratified atmospheric boundary layers. By reducing computational cost and improving high-fidelity modelling capabilities, this work enables researchers and industrial users to address complex research questions and reduce uncertainties in the planning and operation of wind farms.

7. Sammanfattning (Swedish Summary)

Lattice-Boltzmann-metoden (LBM) är en lovande metod för att möjliggöra kostnadseffektiv Large Eddy Simulering (LES) av vindkraftsflöden med en god fysikalisk representation, vilket gör den attraktiv för både forskning och industriella tillämpningar. Denna avhandling utökar tillämpbarheten av LBM genom att utveckla och validera ett förbättrat (open-source) ramverk för högupplösta flödessimuleringar av vindkraftsparker och demonstrera dess relevans för utveckling av lågkostnadsmodeller och forskning om kontroll av vindkraftsparker. Med utgångspunkt i definierade krav för industriell användning och dess kunskapsluckor har ramverket utvecklats, validerats, förbättrats och utökats.

Många av de aktuella och brådskande frågorna inom vindkraftsforskning är beroende av högupplösta flödessimuleringar. Samtidigt förlitar sig industrin på enkla modeller som inte kan reproducera de komplexa interaktionerna mellan vindkraftsparker och atmosfären. För att minska beräkningskostnaden har nya numeriska metoder utvecklats som använder grafikprocessorer (GPUs) för parallella beräkningar. Traditionella lösare som diskretiserar Navier-Stokes-ekvationer är dock mindre lämpade för sådan hårdvara. LBM är däremot utmärkt parallelliserbar och har använts på GPU:er i mer än 20 år. De introducerades nyligen inom vindkraftsforskning och uppvisar stor potential för att minska beräkningskostnader och beräkningstid, se till exempel, Asmuth m. fl. [16], Schottenhamml m. fl. [130] och Asmuth m. fl. [15]. Tidigare studier fokuserade dock på förenklade fall. I detta arbete utvecklar vi ett open-source ramverk som består av två huvudkomponenter: den open-source GPU-baserade LBM-lösaren VirtualFluids och vindkraftssimuleringsverktyget Wind Farm Interface (WiFI).

I artikel I granskar och sammanfattar vi tillgänglig relevant litteratur kring LBM för vindkraftsflöden. Majoriteten av metoderna som krävs för att skapa generella flödeslösare för vindkraftsparker finns tillgängliga idag men de saknas sammanställt ramverk som kombinerar dem. Dessutom har många av de nödvändiga metoderna tillämpats i ett antal studier och LBM har upprepade gånger visats vara korrekt och lämplig för storskaliga tillämpningar. Tidigare tillämpningar indikerar att beräkningsprestanda fördubblas ungefär vart 2,5:e år och att beräkningstid motsvarande realtid redan är uppnåelig med nuvarande hårdvara. I artikel II genomför vi en undersökning för att förstå hinder för införandet av LES för industriellt bruk. Vi finner att beräknings-kostnad och tid är de största faktorerna som avskräcker industriella användare, men även att brist på expertis och validering är viktiga skäl. Svaren angående acceptabel beräkningstid för ett representativt fall varierade från sekunder till daga med ungefär

hälften av respondenterna i undersökningen var nöjda med realtidsprestanda. Som jämförelse genomför vi en simulering av en vindkraftspark bestående av 64 turbiner med hjälp av det ramverk som utvecklats i denna avhandling och finner att den kan uppnå realtidsprestanda på den för närvarande tillgängliga hårdvaran.

Artikel III presenterar en validering av det nyutvecklade ramverket mot mätdata och andra lösare från SWiFT-benchmarken [34]. Den LBM-baserade modellen är minst lika noggrann som andra metoder för att förutsäga turbinernas respons och dess vakar. En jämförelse av energi-till-lösning visar att lösaren är mer än en storleksordning effektivare än den näst mest effektiva LES-lösaren och till och med mer än två storleksordningar effektivare än den dyraste lösaren, samtidigt som den uppnår liknande noggrannhetsnivåer. Den är dock också en till två storleksordningar dyrare än modeller med lägre fysikalisk beskrivning. Den s.k. actuatorline-metoden som används för att modellera rotorn har en välkänd brist att överpredicera krafterna nära roten och spetsen på grund av utsmetning av krafter i flödesdomänen. Ett antal modeller för att minska denna brist har föreslagits. Ett antal modeller för att lindra denna brist har föreslagits. I artikel IV jämför vi modellerna som föreslagits av Martínez-Tossas m. fl. [99] och Meyer Forsting m. fl. [101]. Vi utvärderar effekten av beräkningsnäts upplösning och utdragning på rotorns respons, vak och beräkningstid samt huruvida korrigeringarna tillräckligt korrigerar för fel som introduceras av utdragningen. Beräkningstiden minskar drastiskt med reducerad upplösning medan noggrannheten i rotorresponsen kan ökas genom att använda endera korrigeringen. Vaken påverkas mestadels av upplösning och utdragningens storlek. Slutligen undersöker vi effekten av korrigeringarna på en vindkraftspark och finner att korrigeringarna påverkar turbiner på olika sätt, vilket innebär att effektprognoser inte kan korrigeras med en konstant faktor. I artikel V utökar vi VirtualFluids till att inkludera termisk stratifiering genom att lösa en andra transportekvation för temperatur. Vi utvecklar ett nytt väggrandvillkor för att tillämpa Monin-Obukhov likhetsteori (MOST), implementerar Coriolis och flytkraft och dämpningslager. Modellen valideras i konventionellt neutrala och stabilt stratifierade gränsskikt. Vi finner mycket god överensstämmelse i de flesta undersökta kvantiteter, men i inversionslagret finner vi marginell skillnad i vridningen. De nyutvecklade randvillkoren presterar bra under båda förhållandena. Medan beräkningskostnaden ökar tvåfaldigt är modellen fortfarande mycket mer beräkningseffektiv än jämförbara beräkningsmetoder.

Artiklar VI och VII tillämpar LBM på typiska frågor inom vindkraftsforskning. I artikel VI använder vi en föregångare till modellen som utvecklats i denna avhandling för att generera en stor datamängd för att träna e maskininlärningsmodell att förutsäga flödesfältet i vaken av en enda turbin. Modellen överensstämmer utmärkt med testdata och kan generera förutsägelser på några hundra mikrosekunder. I artikel VII undersöker vi egenskaperna hos turbiners vak genom att använda helixmetoden för att mildra vakeffekter. Vi ger ana-

lytiska beskrivningar av vaken och karakteriserar effekten på vakavböjning, vakmixning och vakens form. Dessutom studerar vi interaktionerna mellan flera vakar och finner att fasskillnaden mellan helixar är en avgörande parameter för att effektivt applicera helixmetoden på en rad turbiner.

Ytterligare validering av modellen krävs, särskilt för större vindkraftverk och parker. För närvarande finns dock inga valideringsdata för ett sådant fall offentligt tillgängliga. Dessutom, även om det nuvarande ramverket har fler funktioner än andra LBM-lösare, saknas det fortfarande funktioner. Detta inkluderar förmågan att modellera komplex terräng och skogar, koppling till mesoskaliga modeller och simulera vindkraftparker med termisk stratifiering. Sammanfattningsvis för detta arbete framåt LBM för vindenergi från en lovande metod till att vara redo för bred användning inom både industri och forskning.

Det utvecklade ramverket möjliggör mycket effektiva och noggranna beräkningar av flöden i stora vindkraftparker i, givet några ytterligare implementeringar, stratifierade atmosfäriska gränsskikt. Genom att minska beräkningskostnaderna och förbättra högkvalitativa modelleringsmöjligheter gör detta arbete det möjligt för forskare och industriella användare att besvara komplex forskningsfrågor och minska osäkerheter i planering och drift av vindkraftparker.

Acknowledgments

This work has been filled with many highs and lows. Luckily, the highs far outweighed the lows, thanks to the many kind, smart and wonderful people that have accompanied me on this journey. At first, I would like to thank my supervisor Stefan Ivanell, for the trust he put in me to do what I thought was important, and for supporting me in my ideas. I would also like to thank Joakim Widén, for being there when needed. Furthermore, I want to thank Hugo Olivares-Espinosa for keeping a critical eye and continuing asking questions, but also giving advice where to go.

My PhD was part of the Graduate School on Sustainability Studies (GRASS) of Uppsala University Campus Gotland. We set out to do things differently, but first had to understand how do anything at all amidst a global pandemic. But once we had found our way, GRASS has been a wonderful environment to do a PhD, filled with many kind and generous people. I want to thank my fellow GRASS students for all the fun and interesting time we have shared and teaching me a wider perspective on the world and research. I especially want to thank Marije Poort for annoying the hell out of me, always ripping me out of my concentration, and being a great friend. I want to thank the director of GRASS, Jenny Helin, for always having an ear to listen to our troubles and all the effort she has put into making GRASS and now also Campus Gotland a place where people and research thrive. To my collaborators on the GoSET project, Gubb Marit Stigson and Lukas Dahlström, I want to thank you for all the mutual support and for at least trying to understand each others projects and perspectives.

At the same time as I was a GRASS student, I was also a member of the Wind Energy division. I want to thank Annie Nyström for her never-ending kindness and support with everything administrative and your help with the PhD Seminar. We all know that nothing here would work without you. Thank you to everyone at the wind energy division for fun discussions over lunch and coffee breaks, sometimes strenuous planning days or early Monday morning meetings. I especially want to thank my fellow PhD students. Memo (Mohammed Mehdi) Mohammadi, thank you for fun travels, discussions and showing me what it means to be effective. To Jean Bastin, Roko Topić, Adrián Silva Caballero, Damián Lopéz, Vikranth Nayankanthi, Karthikeyan Prabakaran, and Zahojie Qin, I wish you the best of luck for the rest of your journey.

I also want to thank the many other people at Campus Gotland that have made every lunch break an opportunity for an interesting discussion. In particular,

I want to thank Nicholas Hoad for having been a great roommate and good friend, and Lotte van den Heuvel for showing me how healthy it can be to have a more distant approach to your work. I also want to thank Antonio Segalini for showing me that experiments can also be fun and Johan Arnqvist for actually teaching me some meteorology.

During my time as a PhD I had the fortune of meeting many researchers in the wind energy community, particularly the other PhDs at the PhD Seminars, which we had the hair-costing honour to host one year. I like to think I have made many friends along the years from various places, so that going to a conference has become more like a meeting of old friends than a work trip. Special thanks to the gang from KU Leuven, a remarkable confluence of scientific and partying excellence. I would like to thank my colleagues from DTU, Alexander Meyer-Forsting, Niels Troldborg and Nils Horneff, for being the first ones to start believing in this thing as well and for many nice trips to Copenhagen. I want to thank my collaborator on paper IV, Emanuel Taschner, for the shared enthusiasm on doing fluid dynamics and LES and being curious and for endless interesting discussions between and often during conference sessions.

I also had the opportunity to visit the Instituto de Mecánica de los Fluidos e Ingenería Ambiental at the University of the Republic Uruguay for a research stay. I want to thank everyone there for welcoming me and putting up with my broken Spanish. I especially want to thank Bruno Lopéz for helping with the organisation and for being a good friend and sharing late nights at the office, Paula Gomez for letting me share the office with her and her great artwork that is now decorating my home and Martín Draper, muchas gracias por tu hospidalidad y tu asado delicioso de Malin.

The work in this thesis was, of course, only possible because I could use what so many smart and hard working people have build and invented before me. I would like to thank everyone at the IRMB of TU Braunschweig for letting us use and abuse VirtualFluids and sharing so generously. Thanks to Martin Geier for knowing the answer to every question about the LBM I could imagine and some more and for also willingly sharing that answer. I want to thank Anna Wellman for her tremendous efforts in trying to teach an a little older pony some new coding tricks. Lastly, I want thank Martin Schönherr, for all the time he has given to us, discussing code, reading merge requests, and talking about everything else in the world that may or may not be LBM-related.

The journey of this thesis began not with a PhD, but actually at an internship I was doing at the ultraFluidX team, where I had the great pleasure of learning from, among others, Christoph Niedermeier and Christian Janssen, giving me the first glimpses into LBM and GPU-programming. Funnily, the questions of inclined boundary conditions that plagued me then still plague me today. Through the keen awareness of Merten Stender, I was introduced to this researcher on an island I had never heard of, of the name of Henrik Asmuth, who happened to do exactly what I had in mind for my Diploma thesis. This was the

start to a great friendship that also happened to be academically fruitful. I have learnt so much from you, from camping at the sea to hacking away at a new code to breaking up all my spaghetti sentences into comprehensible English. You have laid the foundation so much of this work and I am so grateful for your guidance and continued mentorship.

I would like to thank my parents, for doing everything in their possibilities to support me on my way, even if they don't always like exactly everything I do. It seems to have not gone that poorly. I want to thank Peter, who did exactly the same. To my little sister Katha, for being a smart and kind person that supports me exactly then when our parents don't think it is a good idea. I also want to thank all my friends back in Germany that have also been like a family to me, Helene, Lisa, Kersten, the boat crew, my former roommates Tom and Martin, you all have made the last five and many more years before that so much fun and helped me endure on my little island with visits, care packages and listening to my self-made problems.

Lastly, I would like to thank Malin, my partner, for loving me for who I am, even though that includes long monologues on the physics of wind turbines and distributions of particles. And for accepting that I am sometimes late at the office but also not accepting that I am sometimes late at the office. You have made the lows not so low and the highs so much higher, and for that I am thankful.

Bibliography

[1] M. Abkar, H. J. Bae, and P. Moin. "Minimum-Dissipation Scalar Transport Model for Large-Eddy Simulation of Turbulent Flows". In: *Physical Review Fluids* 1.4 (2016), p. 041701.

[2] H. Alihussein, M. Geier, and M. Krafczyk. "A Parallel Coupled Lattice Boltzmann-Volume of Fluid Framework for Modeling Porous Media Evolution". In: *Materials* 14.10 (10 2021), p. 2510.

[3] D. Allaerts and J. Meyers. "Large Eddy Simulation of a Large Wind-Turbine Array in a Conventionally Neutral Atmospheric Boundary Layer". In: *Physics of Fluids* 27.6 (2015), p. 065108.

[4] S. J. Andersen, B. Witha, S.-P. Breton, J. N. Sørensen, R. F. Mikkelsen, and S. Ivanell. "Quantifying Variability of Large Eddy Simulations of Very Large Wind Farms". In: *Journal of Physics: Conference Series* 625.1 (2015), p. 012027.

[5] S. J. Andersen, S.-P. Breton, B. Witha, S. Ivanell, and J. N. Sørensen. "Global Trends in the Performance of Large Wind Farms Based on High-Fidelity Simulations". In: *Wind Energy Science* 5.4 (2020), pp. 1689–1703.

[6] C. L. Archer. "Brief Communication: A Note on the Variance of Wind Speed and Turbulence Intensity". In: *Wind Energy Science* 10.7 (2025), pp. 1433–1438.

[7] J. Arnqvist. "Mean Wind and Turbulence Conditions in the Boundary Layer above Forests". Uppsala: Uppsala University, 2015.

[8] S. P. Arya. *Introduction to Micrometeorology*. 2nd ed. Vol. 79. International Geophysics Series. San Diego: Academic Press, 2001. 420 pp.

[9] H. Asmuth, C. F. Janßen, H. Olivares-Espinosa, and S. Ivanell. "Wall-Modeled Lattice Boltzmann Large-Eddy Simulation of Neutral Atmospheric Boundary Layers". In: *Physics of Fluids* 33.10 (2021), p. 105111.

[10] H. Asmuth, C. F. Janßen, H. Olivares-Espinosa, K. Nilsson, and S. Ivanell. "Assessment of Weak Compressibility in Actuator Line Simulations of Wind Turbine Wakes". In: *Journal of Physics: Conference Series* 1618 (2020), p. 062057.

[11] H. Asmuth and H. Korb. *WakeNet*. 2024. URL: `https://source.coderefinery.org/wind_energy_uu/wakenet` (visited on 09/11/2025).

[12] H. Asmuth and H. Korb. “WakeNet 0.1 - A Simple Three-dimensional Wake Model Based on Convolutional Neural Networks”. In: *Journal of Physics: Conference Series* 2265.2 (2022), p. 022066.

[13] H. Asmuth, H. Korb, and S. Ivanell. “How Fast Is Fast Enough? Industry Perspectives on the Use of Large-eddy Simulation in Wind Energy”. In: *Journal of Physics: Conference Series* 2505.1 (2023), p. 012001.

[14] H. Asmuth, G. P. Navarro Diaz, H. A. Madsen, E. Branlard, A. R. Meyer Forsting, K. Nilsson, J. Jonkman, and S. Ivanell. “Wind Turbine Response in Waked Inflow: A Modelling Benchmark against Full-Scale Measurements”. In: *Renewable Energy* 191 (2022), pp. 868–887.

[15] H. Asmuth, H. Olivares-Espinosa, and S. Ivanell. “Actuator Line Simulations of Wind Turbine Wakes Using the Lattice Boltzmann Method”. In: *Wind Energy Science* 5.2 (2020), pp. 623–645.

[16] H. Asmuth, H. Olivares-Espinosa, K. Nilsson, and S. Ivanell. “The Actuator Line Model in Lattice Boltzmann Frameworks: Numerical Sensitivity and Computational Performance”. In: *Journal of Physics: Conference Series* 1256 (2019), p. 012022.

[17] C. Bak, R. Bitsche, A. Yde, T. Kim, M. H. Hansen, F. Zahle, M. Gaunaa, J. P. A. A. Blasques, M. Døssing, J.-J. W. Heinen, et al. “Light Rotor: The 10-MW Reference Wind Turbine”. In: *EWEA 2012-European Wind Energy Conference & Exhibition*. European Wind Energy Association (EWEA), 2012.

[18] M. Barker, N. P. Chue Hong, D. S. Katz, A.-L. Lamprecht, C. Martinez-Ortiz, F. Psomopoulos, J. Harrow, L. J. Castro, M. Gruenpeter, P. A. Martinez, and T. Honeyman. “Introducing the FAIR Principles for Research Software”. In: *Scientific Data* 9.1 (2022), p. 622.

[19] M. Bastankhah and F. Porté-Agel. “Experimental and Theoretical Study of Wind Turbine Wakes in Yawed Conditions”. In: *Journal of Fluid Mechanics* 806 (2016), pp. 506–541.

[20] P. Bauweraerts and J. Meyers. “On the Feasibility of Using Large-Eddy Simulations for Real-Time Turbulent-Flow Forecasting in the Atmospheric Boundary Layer”. In: *Boundary-Layer Meteorology* 171.2 (2019), pp. 213–235.

[21] R. J. Beare, M. K. Macvean, A. A. M. Holtslag, J. Cuxart, I. Esau, J.-C. Golaz, M. A. Jimenez, M. Khairoutdinov, B. Kosovic, D. Lewellen, T. S. Lund, J. K. Lundquist, A. Mccabe, A. F. Moene, Y. Noh, S. Raasch, and P. Sullivan. “An Intercomparison of Large-Eddy Simulations of the Stable Boundary Layer”. In: *Boundary-Layer Meteorology* 118.2 (2006), pp. 247–272.

[22] J. Berg, E. G. Patton, and P. P. Sullivan. “Large-Eddy Simulation of Conditionally Neutral Boundary Layers: A Mesh Resolution Sensitivity Study”. In: *Journal of the Atmospheric Sciences* 77.6 (2020), pp. 1969–1991.

[23] P. L. Bhatnagar, E. P. Gross, and M. Krook. “A Model for Collision Processes in Gases. I. Small Amplitude Processes in Charged and Neutral One-Component Systems”. In: *Physical Review* 94.3 (1954), pp. 511–525.

[24] E. Bou-Zeid, C. Meneveau, and M. Parlange. “A Scale-Dependent Lagrangian Dynamic Model for Large Eddy Simulation of Complex Turbulent Flows”. In: *Physics of Fluids* 17.2 (2005), p. 025105.

[25] M. Bouzidi, M. Firdaouss, and P. Lallemand. “Momentum Transfer of a Boltzmann-lattice Fluid with Boundaries”. In: *Physics of Fluids* 13.11 (2001), pp. 3452–3459.

[26] K. Brown, G. Yalla, L. Cheung, J. Frederik, D. Houck, N. deVelder, E. Simley, and P. Fleming. “Comparison of Wind-Farm Control Strategies under Realistic Offshore Wind Conditions: Wake Quantities of Interest”. In: *Wind Energy Science Discussions* (2025), pp. 1–39.

[27] J. A. Businger, J. C. Wyngaard, Y. Izumi, and E. F. Bradley. “Flux-Profile Relationships in the Atmospheric Surface Layer”. In: *Journal of the Atmospheric Sciences* 28.2 (1971), pp. 181–189.

[28] D.-G. Caprace, P. Chatelain, and G. Winckelmans. “Lifting Line with Various Mollifications: Theory and Application to an Elliptical Wing”. In: *AIAA Journal* 57.1 (2019), pp. 17–28.

[29] M. Coquelet, J. Gutknecht, J. V. Wingerden, M. Duponcheel, and P. Chatelain. “Dynamic Individual Pitch Control for Wake Mitigation: Why Does the Helix Handedness in the Wake Matter?” In: *Journal of Physics: Conference Series* 2767.9 (2024), p. 092084.

[30] C. Coreixas, B. Chopard, and J. Latt. “Comprehensive Comparison of Collision Models in the Lattice Boltzmann Framework: Theoretical Investigations”. In: *Physical Review E* 100.3 (2019), p. 033305.

[31] D. d’Humières. “Generalized Lattice-Boltzmann Equations”. In: *Rarefied Gas Dynamics: Theory and Simulations*. Progress in Astronautics and Aeronautics. American Institute of Aeronautics and Astronautics, 1994, pp. 450–458.

[32] K. O. Dağ. “Combined Pseudo-Spectral / Actuator Line Model for Wind Turbine Applications”. 2017.

[33] K. O. Dağ and J. N. Sørensen. “A New Tip Correction for Actuator Line Computations”. In: *Wind Energy* 23.2 (2020), pp. 148–160.

[34] P. Doubrawa, E. W. Quon, L. A. Martinez-Tossas, K. Shaler, M. Debnath, N. Hamilton, T. G. Herges, D. Maniaci, C. L. Kelley, A. S. Hsieh, M. L. Blaylock, P. van der Laan, S. J. Andersen, S. Krueger, M. Cathelain, W. Schlez, J. Jonkman, E. Branlard, G. Steinfeld, S. Schmidt, F. Blondel, L. J. Lukassen, and P. Moriarty. "Multimodel Validation of Single Wakes in Neutral and Stratified Atmospheric Conditions". In: *Wind Energy* 23.11 (2020), pp. 2027–2055.

[35] F. Dubois. "Equivalent Partial Differential Equations of a Lattice Boltzmann Scheme". In: *Computers & Mathematics with Applications*. Mesoscopic Methods in Engineering and Science 55.7 (2008), pp. 1441–1449.

[36] Z. Fan, F. Qiu, A. Kaufman, and S. Yoakum-Stover. "GPU Cluster for High Performance Computing". In: *SC '04: Proceedings of the 2004 ACM/IEEE Conference on Supercomputing*. SC '04: Proceedings of the 2004 ACM/IEEE Conference on Supercomputing. 2004, pp. 47–47.

[37] Y. Feng, J. Miranda-Fuentes, S. Guo, J. Jacob, and P. Sagaut. "ProLB: A Lattice Boltzmann Solver of Large-Eddy Simulation for Atmospheric Boundary Layer Flows". In: *Journal of Advances in Modeling Earth Systems* 13.3 (2021), e2020MS002107.

[38] O. Filippova and D. Hänel. "Grid Refinement for Lattice-BGK Models". In: *Journal of Computational Physics* 147.1 (1998), pp. 219–228.

[39] E. Finserås, I. Herrera Anchustegui, E. Cheynet, C. G. Gebhardt, and J. Reuder. "Gone with the Wind? Wind Farm-Induced Wakes and Regulatory Gaps". In: *Marine Policy* 159 (2024), p. 105897.

[40] J. A. Frederik, R. Weber, S. Cacciola, F. Campagnolo, A. Croce, C. Bottasso, and J.-W. van Wingerden. "Periodic Dynamic Induction Control of Wind Farms: Proving the Potential in Simulations and Wind Tunnel Experiments". In: *Wind Energy Science* 5.1 (2020), pp. 245–257.

[41] R. K. Freitas, A. Henze, M. Meinke, and W. Schröder. "Analysis of Lattice-Boltzmann Methods for Internal Flows". In: *Computers & Fluids* 47.1 (2011), pp. 115–121.

[42] N. Fulghum and E. Graham. *Global Electricity Review*. Ember, 2025.

[43] S. N. Gadde, A. Stieren, and R. J. A. M. Stevens. "Large-Eddy Simulations of Stratified Atmospheric Boundary Layers: Comparison of Different Subgrid Models". In: *Boundary-Layer Meteorology* 178.3 (2021), pp. 363–382.

[44] M. Gehrke, C. F. Janßen, and T. Rung. "Scrutinizing Lattice Boltzmann Methods for Direct Numerical Simulations of Turbulent Channel Flows". In: *Computers & Fluids*. Ninth International Conference on Computational Fluid Dynamics (ICCFD9) 156 (2017), pp. 247–263.

[45] M. Geier, A. Greiner, and J. G. Korvink. “Bubble Functions for the Lattice Boltzmann Method and Their Application to Grid Refinement”. In: *The European Physical Journal Special Topics* 171.1 (2009), pp. 173–179.

[46] M. Geier, A. Greiner, and J. G. Korvink. “Cascaded Digital Lattice Boltzmann Automata for High Reynolds Number Flow”. In: *Physical Review E* 73.6 (2006), p. 066705.

[47] M. Geier, K. Kutscher, M. Schönherr, A. Wellmann, S. Peters, H. Alihussein, J. Linxweiler, and M. Krafczyk. “VirtualFluids – Open Source Parallel LBM Solver”. In: *Computer Physics Communications* (2025), p. 109810.

[48] M. Geier, A. Pasquali, and M. Schönherr. “Parametrization of the Cumulant Lattice Boltzmann Method for Fourth Order Accurate Diffusion Part I: Derivation and Validation”. In: *Journal of Computational Physics* 348 (2017), pp. 862–888.

[49] M. Geier and M. Schönherr. “Esoteric Twist: An Efficient in-Place Streaming Algorithmus for the Lattice Boltzmann Method on Massively Parallel Hardware”. In: *Computation* 5.2 (2017), p. 19.

[50] M. Geier, M. Schönherr, A. Pasquali, and M. Krafczyk. “The Cumulant Lattice Boltzmann Equation in Three Dimensions: Theory and Validation”. In: *Computers & Mathematics with Applications* 70.4 (2015), pp. 507–547.

[51] T. Göçmen, G. Giebel, N. K. Poulsen, and P. E. Sørensen. “Possible Power of Down-Regulated Offshore Wind Power Plants: The PossPOW Algorithm”. In: *Wind Energy* 22.2 (2019), pp. 205–218.

[52] G. Gruszczyński and Ł. Łaniewski-Wołłk. “A Comparative Study of 3D Cumulant and Central Moments Lattice Boltzmann Schemes with Interpolated Boundary Conditions for the Simulation of Thermal Flows in High Prandtl Number Regime”. In: *International Journal of Heat and Mass Transfer* 197 (2022), p. 123259.

[53] J. Gutknecht, E. Taschner, M. Coquelet, A. Viré, and J. W. van Wingerden. “The Impact of Coherent Large-Scale Vortices Generated by Helix Active Wake Control on the Recovery Process of Wind Turbine Wakes”. In: *Physics of Fluids* 37.6 (2025), p. 067162.

[54] M. Han, R. Ooka, and H. Kikumoto. “Effects of Wall Function Model in Lattice Boltzmann Method-Based Large-Eddy Simulation on Built Environment Flows”. In: *Building and Environment* 195 (2021), p. 107764.

[55] M. O. L. Hansen. *Aerodynamics of Wind Turbines*. 2nd ed. London ; Sterling, VA: Earthscan, 2008. 181 pp.

[56] M. Haussmann, F. Ries, J. B. Jeppener-Haltenhoff, Y. Li, M. Schmidt, C. Welch, L. Illmann, B. Böhm, H. Nirschl, M. J. Krause, and A. Sadiki. "Evaluation of a Near-Wall-Modeled Large Eddy Lattice Boltzmann Method for the Analysis of Complex Flows Relevant to IC Engines". In: *Computation* 8.2 (2 2020), p. 43.

[57] F. J. Higuera and J. Jiménez. "Boltzmann Approach to Lattice Gas Simulations". In: *Europhysics Letters* 9.7 (1989), p. 663.

[58] E. L. Hodgson, C. Grinderslev, A. R. Meyer Forsting, N. Troldborg, N. N. Sørensen, J. N. Sørensen, and S. J. Andersen. "Validation of Aeroelastic Actuator Line for Wind Turbine Modelling in Complex Flows". In: *Frontiers in Energy Research* 10 (2022).

[59] J. Jacob, O. Malaspinas, and P. Sagaut. "A New Hybrid Recursive Regularised Bhatnagar–Gross–Krook Collision Model for Lattice Boltzmann Method-Based Large Eddy Simulation". In: *Journal of Turbulence* 19.11–12 (2018), pp. 1051–1076.

[60] W. Jakob. *Pybind11*. 2025. URL: `https://github.com/pybind/pybind11` (visited on 08/20/2025).

[61] C. F. Janßen, D. Mierke, M. Überrück, S. Gralher, and T. Rung. "Validation of the GPU-Accelerated CFD Solver ELBE for Free Surface Flow Problems in Civil and Environmental Engineering". In: *Computation* 3.3 (2015), pp. 354–385.

[62] E. Jézéquel, M. Cathelain, V. Masson, and F. Blondel. "Validation of Wind Turbine Wakes Modelled by the Meso-NH LES Solver under Different Cases of Stability". In: *Journal of Physics: Conference Series* 1934.1 (2021), p. 012003.

[63] P. K. Jha, M. J. Churchfield, P. J. Moriarty, and S. Schmitz. "Guidelines for Volume Force Distributions Within Actuator Line Modeling of Wind Turbines on Large-Eddy Simulation-Type Grids". In: *Journal of Solar Energy Engineering* 136.031003 (2014).

[64] P. K. Jha and S. Schmitz. "Actuator Curve Embedding – an Advanced Actuator Line Model". In: *Journal of Fluid Mechanics* 834 (2018).

[65] J. Jonkman, S. Butterfield, W. Musial, and G. Scott. *Definition of a 5-MW Reference Wind Turbine for Offshore System Development*. NREL/TP-500-38060. National Renewable Energy Laboratory, 2009.

[66] J. Jonkman. "The New Modularization Framework for the FAST Wind Turbine CAE Tool". In: *51st AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition*. American Institute of Aeronautics and Astronautics, 2013.

[67] S. K. Kang and Y. A. Hassan. “The Effect of Lattice Models within the Lattice Boltzmann Method in the Simulation of Wall-Bounded Turbulent Flows”. In: *Journal of Computational Physics* 232.1 (2013), pp. 100–117.

[68] V. G. Kleine, A. Hanifi, and D. S. Henningson. “Non-Iterative Vortex-Based Smearing Correction for the Actuator Line Method”. In: *Journal of Fluid Mechanics* 961 (2023), A29.

[69] C. G. Koay and P. J. Basser. “Analytically Exact Correction Scheme for Signal Extraction from Noisy Magnitude MR Signals”. In: *Journal of Magnetic Resonance* 179.2 (2006), pp. 317–322.

[70] H. Korb. *WiFI*. 2025. URL: `https://source.coderefinery.org/Hkorb/wifi` (visited on 08/20/2025).

[71] H. Korb and H. Asmuth. *WakeNet Dataset*. Zenodo, 2022. DOI: `10.5281/zenodo.5911015`.

[72] H. Korb, H. Asmuth, and S. Ivanell. “The Characteristics of Helically Deflected Wind Turbine Wakes”. In: *Journal of Fluid Mechanics* 965 (2023), A2.

[73] H. Korb, H. Asmuth, and S. Ivanell. “Validation of a Lattice Boltzmann Solver Against Wind Turbine Response and Wake Measurements”. In: *Journal of Physics: Conference Series* 2505.1 (2023), p. 012008.

[74] H. Korb, H. Asmuth, M. Schönherr, M. Geier, and S. Ivanell. “Large Eddy Simulation of Thermally Stratified Atmospheric Boundary Layers with a Lattice Boltzmann Method”. In: *Wind Energy Science Discussions* (2025). In Review.

[75] H. Korb, J. Bastin, H. Asmuth, and S. Ivanell. “The Lattice Boltzmann Method for Wind Farm Simulations: A Review”. In: *Wind Energy Science Discussions* (2025), pp. 1–41. In Review.

[76] H. Korb and E. Taschner. “Coarse, Fast, and Still Accurate? Comparing Corrections for the Actuator Line Model”. In: *Journal of Physics: Conference Series* 3016.1 (2025), p. 012051.

[77] B. Kosović, S. Basu, J. Berg, L. K. Berg, S. E. Haupt, X. G. Larsén, J. Peinke, R. J. A. M. Stevens, P. Veers, and S. Watson. “Impact of Atmospheric Turbulence on Performance and Loads of Wind Turbines: Knowledge Gaps and Research Challenges”. In: *Wind Energy Science Discussions* (2025), pp. 1–67.

[78] S. Krüger, G. Steinfeld, M. Kraft, and L. J. Lukassen. “Validation of a Coupled Atmospheric–Aeroelastic Model System for Wind Turbine Power and Load Calculations”. In: *Wind Energy Science* 7.1 (2022), pp. 323–344.

[79] T. Krüger, H. Kusumaatmaja, A. Kuzmin, O. Shardt, G. Silva, and E. M. Viggen. *The Lattice Boltzmann Method: Principles and Practice*. 1st ed. Graduate Texts in Physics. Springer International Publishing, 2017. 694 pp.

[80] M. B. Kuhn, M. T. Henry de Frahan, P. Mohan, G. Deskos, M. Churchfield, L. Cheung, A. Sharma, A. Almgren, S. Ananthan, M. J. Brazell, L. A. Martínez-Tossas, R. Thedin, J. Rood, P. Sakievich, G. Vijayakumar, W. Zhang, and M. Sprague. "AMR-Wind: A Performance-Portable, High-Fidelity Flow Solver for Wind Farm Simulations". In: *Wind Energy* 28.5 (2025), e70010.

[81] K. Kutscher, M. Geier, and M. Krafczyk. "Multiscale Simulation of Turbulent Flow Interacting with Porous Media Based on a Massively Parallel Implementation of the Cumulant Lattice Boltzmann Method". In: *Computers & Fluids* 193 (2019), p. 103733.

[82] Y. Kuwata and K. Suga. "Wall-Modeled Large Eddy Simulation of Turbulent Heat Transfer by the Lattice Boltzmann Method". In: *Journal of Computational Physics* 433 (2021), p. 110186.

[83] L. Lanzilao and J. Meyers. "A Parametric Large-Eddy Simulation Study of Wind-Farm Blockage and Gravity Waves in Conventionally Neutral Boundary Layers". In: *Journal of Fluid Mechanics* 979 (2024), A54.

[84] L. Lanzilao and J. Meyers. "Wind-Farm Wake Recovery Mechanisms in Conventionally Neutral Boundary Layers". In: *Journal of Fluid Mechanics* 1015 (2025), A5.

[85] J. Latt and B. Chopard. "Lattice Boltzmann Method with Regularized Pre-Collision Distribution Functions". In: *Mathematics and Computers in Simulation*. Discrete Simulation of Fluid Dynamics in Complex Systems 72.2 (2006), pp. 165–168.

[86] M. Lehmann. "Esoteric Pull and Esoteric Push: Two Simple In-Place Streaming Schemes for the Lattice Boltzmann Method on GPUs". In: *Computation* 10.6 (2022), p. 92.

[87] S. Lenz, M. Schönherr, M. Geier, M. Krafczyk, A. Pasquali, A. Christen, and M. Giometto. "Towards Real-Time Simulation of Turbulent Air Flow over a Resolved Urban Canopy Using the Cumulant Lattice Boltzmann Method on a GPGPU". In: *Journal of Wind Engineering and Industrial Aerodynamics* 189 (2019), pp. 151–162.

[88] L. Lignarolo, D. Ragni, C. Krishnaswami, Q. Chen, C. Simão Ferreira, and G. van Bussel. "Experimental Analysis of the Wake of a Horizontal-Axis Wind-Turbine Model". In: *Renewable Energy* 70 (2014), pp. 31–46.

[89] R. Löhner. "Towards Overcoming the LES Crisis". In: *International Journal of Computational Fluid Dynamics* 33.3 (2019), pp. 87–97.

[90] O. Maas and S. Raasch. “Wake Properties and Power Output of Very Large Wind Farms for Different Meteorological Conditions and Turbine Spacings: A Large-Eddy Simulation Case Study for the German Bight”. In: *Wind Energy Science* 7.2 (2022), pp. 715–739.

[91] O. Malaspinas. “Increasing Stability and Accuracy of the Lattice Boltzmann Scheme: Recursivity and Regularization”. 2015. arXiv: 1505.06900 [physics].

[92] O. Malaspinas and P. Sagaut. “Consistent Subgrid Scale Modelling for Lattice Boltzmann Methods”. In: *Journal of Fluid Mechanics* 700 (2012), pp. 514–542.

[93] O. Malaspinas and P. Sagaut. “Wall Model for Large-Eddy Simulation Based on the Lattice Boltzmann Method”. In: *Journal of Computational Physics* 275 (2014), pp. 25–40.

[94] B. Maronga, S. Banzhaf, C. Burmeister, T. Esch, R. Forkel, D. Fröhlich, V. Fuka, K. F. Gehrke, J. Geletič, S. Giersch, T. Gronemeier, G. Groß, W. Heldens, A. Hellsten, F. Hoffmann, A. Inagaki, E. Kadasch, F. Kanani-Sühring, K. Ketelsen, B. A. Khan, C. Knigge, H. Knoop, P. Krč, M. Kurppa, H. Maamari, A. Matzarakis, M. Mauder, M. Pallasch, D. Pavlik, J. Pfafferott, J. Resler, S. Rissmann, E. Russo, M. Salim, M. Schrempf, J. Schwenkel, G. Seckmeyer, S. Schubert, M. Sühring, R. von Tils, L. Vollmer, S. Ward, B. Witha, H. Wurps, J. Zeidler, and S. Raasch. “Overview of the PALM Model System 6.0”. In: *Geoscientific Model Development* 13.3 (2020), pp. 1335–1372.

[95] L. A. Martínez-Tossas, M. J. Churchfield, and C. Meneveau. “Optimal Smoothing Length Scale for Actuator Line Models of Wind Turbine Blades Based on Gaussian Body Force Distribution”. In: *Wind Energy* 20.6 (2017), pp. 1083–1096.

[96] L. A. Martínez-Tossas, M. J. Churchfield, and S. Leonardi. “Large Eddy Simulations of the Flow Past Wind Turbines: Actuator Line and Disk Modeling”. In: *Wind Energy* 18.6 (2015), pp. 1047–1060.

[97] L. A. Martínez-Tossas, M. J. Churchfield, A. E. Yilmaz, H. Sarlak, P. L. Johnson, J. N. Sørensen, J. Meyers, and C. Meneveau. “Comparison of Four Large-Eddy Simulation Research Codes and Effects of Model Coefficient and Inflow Turbulence in Actuator-Line-Based Wind Turbine Modeling”. In: *Journal of Renewable and Sustainable Energy* 10.3 (2018), p. 033301.

[98] L. A. Martínez-Tossas and C. Meneveau. “Filtered Lifting Line Theory and Application to the Actuator Line Model”. In: *Journal of Fluid Mechanics* 863 (2019), pp. 269–292.

[99] L. A. Martínez-Tossas, P. Sakievich, M. J. Churchfield, and C. Meneveau. “Generalized Filtered Lifting Line Theory for Arbitrary Chord Lengths and Application to Wind Turbine Blades”. In: *Wind Energy* 27.1 (2024), pp. 101–106.

[100] G. R. McNamara and G. Zanetti. “Use of the Boltzmann Equation to Simulate Lattice-Gas Automata”. In: *Physical Review Letters* 61.20 (1988), pp. 2332–2335.

[101] A. R. Meyer Forsting, G. R. Pirrung, and N. Ramos-García. “A Vortex-Based Tip/Smearing Correction for the Actuator Line”. In: *Wind Energy Science* 4.2 (2019), pp. 369–383.

[102] A. R. Meyer Forsting, G. R. Pirrung, and N. Ramos-García. “Brief Communication: A Fast Vortex-Based Smearing Correction for the Actuator Line”. In: *Wind Energy Science* 5.1 (2020), pp. 349–353.

[103] A. R. Meyer Forsting and N. Troldborg. “Generalised Grid Requirements Minimizing the Actuator Line Angle-of-Attack Error”. In: *Journal of Physics: Conference Series* 1618.5 (2020), p. 052001.

[104] J. Meyers, C. Bottasso, K. Dykes, P. Fleming, P. Gebraad, G. Giebel, T. Göçmen, and J.-W. van Wingerden. “Wind Farm Flow Control: Prospects and Challenges”. In: *Wind Energy Science* 7.6 (2022), pp. 2271–2306.

[105] C.-H. Moeng. “A Large-Eddy-Simulation Model for the Study of Planetary Boundary-Layer Turbulence”. In: *Journal of the Atmospheric Sciences* 41.13 (1984), pp. 2052–2062.

[106] M. M. Mohammadi, H. Olivares-Espinosa, G. P. Navarro Diaz, and S. Ivanell. “An Actuator Sector Model for Wind Power Applications: A Parametric Study”. In: *Wind Energy Science Discussions* (2023), pp. 1–24.

[107] W. Munters and J. Meyers. “An Optimal Control Framework for Dynamic Induction Control of Wind Farms and Their Interaction with the Atmospheric Boundary Layer”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 375.2091 (2017), p. 20160100.

[108] National Renewable Energy Laboratory. *SOWFA: Simulator fOr Wind Farm Applications*. URL: `https://www.nrel.gov/wind/nwtc/sowfa` (visited on 09/04/2025).

[109] NVIDIA. *CUDA Toolkit*. NVIDIA Developer. URL: `https://developer.nvidia.com/cuda-toolkit` (visited on 09/04/2025).

[110] M. O'Malley, H. Holttinen, N. Cutululis, T. K. Vrana, J. King, V. Gevorgian, X. Wang, F. Rajaei-Najafabadi, and A. Hadjileonidas. "Grand Challenges of Wind Energy Science – Meeting the Needs and Services of the Power System". In: *Wind Energy Science* 9.11 (2024), pp. 2087–2112.

[111] A. M. Obukhov. "Turbulence in an Atmosphere with a Non-Uniform Temperature". In: *Boundary-Layer Meteorology* 2.1 (1971), pp. 7–29.

[112] N. Onodera, T. Aoki, T. Shimokawabe, and H. Kobayashi. "Large-Scale LES Wind Simulation Using Lattice Boltzmann Method for a 10 Km x 10 Km Area in Metropolitan Tokyo". In: *TSUBAME ESJ* 9 (2013).

[113] P. Ouro, M. Ghobrial, K. Ali, and T. Stallard. "Numerical Modelling of Offshore Wind-Farm Cluster Wakes". In: *Renewable and Sustainable Energy Reviews* 215 (2025), p. 115526.

[114] A. Pasquali, M. Geier, and M. Krafczyk. "Near-Wall Treatment for the Simulation of Turbulent Flow by the Cumulant Lattice Boltzmann Method". In: *Computers & Mathematics with Applications*. Mesoscopic Methods in Engineering and Science 79.1 (2020), pp. 195–212.

[115] F. Perot, M.-S. Kim, and M. Meskine. "NREL Wind Turbine Aerodynamics Validation and Noise Predictions Using a Lattice Boltzmann Method". In: *18th AIAA/CEAS Aeroacoustics Conference (33rd AIAA Aeroacoustics Conference)*. American Institute of Aeronautics and Astronautics, 2012.

[116] C. S. Peskin. "The Immersed Boundary Method". In: *Acta Numerica* 11 (2002), pp. 479–517.

[117] G. R. Pirrung, H. A. Madsen, and S. Schreck. "Trailed Vorticity Modeling for Aeroelastic Wind Turbine Simulations in Standstill". In: *Wind Energy Science* 2.2 (2017), pp. 521–532.

[118] G. R. Pirrung, H. Madsen, and S. Schreck. "Trailed Vorticity Modeling for Aeroelastic Wind Turbine Simulations in Stand Still". In: *Journal of Physics: Conference Series* 753 (2016), p. 042007.

[119] F. Porté-Agel, M. Bastankhah, and S. Shamsoddin. "Wind-Turbine and Wind-Farm Flows: A Review". In: *Boundary-Layer Meteorology* 174.1 (2020), pp. 1–59.

[120] E. W. Quon, P. Doubrawa, and M. Debnath. "Comparison of Rotor Wake Identification and Characterization Methods for the Analysis of Wake Dynamics and Evolution". In: *Journal of Physics: Conference Series* 1452.1 (2020), p. 012070.

[121] S. O. Rice. "Mathematical Analysis of Random Noise". In: *The Bell System Technical Journal* 23.3 (1944), pp. 282–332.

[122] S. O. Rice. "Mathematical Analysis of Random Noise". In: *The Bell System Technical Journal* 24.1 (1945), pp. 46–156.

[123] W. Rozema, H. J. Bae, P. Moin, and R. Verstappen. “Minimum-Dissipation Models for Large-Eddy Simulation”. In: *Physics of Fluids* 27.8 (2015), p. 085107.

[124] S. Rullaud, F. Blondel, and M. Cathelain. “Actuator-Line Model in a Lattice Boltzmann Framework for Wind Turbine Simulations”. In: *Journal of Physics: Conference Series* 1037 (2018), p. 022023.

[125] P. Sagaut. *Large Eddy Simulation for Incompressible Flows*. Scientific Computation. Berlin/Heidelberg: Springer-Verlag, 2006.

[126] S. Sarmast, R. Dadfar, R. F. Mikkelsen, P. Schlatter, S. Ivanell, J. N. Sørensen, and D. S. Henningson. “Mutual Inductance Instability of the Tip Vortices behind a Wind Turbine”. In: *Journal of Fluid Mechanics* 755 (2014), pp. 705–731.

[127] J. A. Sauer and D. Muñoz-Esparza. “The FastEddy® Resident-GPU Accelerated Large-Eddy Simulation Framework: Model Formulation, Dynamical-Core Validation and Performance Benchmarks”. In: *Journal of Advances in Modeling Earth Systems* 12.11 (2020), e2020MS002100.

[128] M. Schönherr. “Towards Reliable LES-CFD Computations Based on Advanced LBM Models Utilizing (Multi-) GPGPU Hardware”. PhD thesis. TU Braunschweig, Institut für rechnergestützte Modellierung im Bauingenieurswesen, 2015.

[129] H. Schottenhamml, A. Anciaux Sedrakian, F. Blondel, H. Köstler, and U. Rüde. “waLBerla-wind: A Lattice-Boltzmann-based High-Performance Flow Solver for Wind Energy Applications”. In: *Concurrency and Computation: Practice and Experience* 36.16 (2024), e8117.

[130] H. Schottenhamml, A. Anciaux-Sedrakian, F. Blondel, A. Borras-Nadal, P.-A. Joulin, and U. Rüde. “Evaluation of a Lattice Boltzmann-based Wind-Turbine Actuator Line Model against a Navier-Stokes Approach”. In: *Journal of Physics: Conference Series* 2265.2 (2022), p. 022027.

[131] J. Smagorinsky. “GENERAL CIRCULATION EXPERIMENTS WITH THE PRIMITIVE EQUATIONS: I. THE BASIC EXPERIMENT”. In: *Monthly Weather Review* 91.3 (1963), pp. 99–164.

[132] J. N. Sørensen, R. F. Mikkelsen, D. S. Henningson, S. Ivanell, S. Sarmast, and S. J. Andersen. “Simulation of Wind Turbine Wakes Using the Actuator Line Technique”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 373.2035 (2015), p. 20140071.

[133] J. N. Sørensen and W. Z. Shen. “Computation of Wind Turbine Wakes Using Combined Navier-Stokes/Actuator-line Methodology”. In: *Proceedings of European Wind Energy Conference EWEC ’99, Nice*. 1999 European Wind Energy Conference and Exhibition. 1999, pp. 156–159.

[134] J. N. Sørensen and W. Z. Shen. “Numerical Modeling of Wind Turbine Wakes”. In: *Journal of Fluids Engineering* 124.2 (2002), pp. 393–399.

[135] G. G. Spinelli, J. Gericke, K. Masilamani, and H. G. Klimach. “Key Ingredients for Wall-Modeled LES with the Lattice Boltzmann Method: Systematic Comparison of Collision Schemes, SGS Models, and Wall Functions on Simulation Accuracy and Efficiency for Turbulent Channel Flow”. In: *Discrete and Continuous Dynamical Systems - S* 17.11 (2024), pp. 3224–3251.

[136] G. G. Spinelli, T. Horstmann, K. Masilamani, M. M. Soni, H. Klimach, A. Stück, and S. Roller. “HPC Performance Study of Different Collision Models Using the Lattice Boltzmann Solver Musubi”. In: *Computers & Fluids* 255 (2023), p. 105833.

[137] M. Steinbuch, W. W. de Boer, O. H. Bosgra, S. A. W. M. Peters, and J. Ploeg. “Optimal Control of Wind Power Plants”. In: *Journal of Wind Engineering and Industrial Aerodynamics* 27.1 (1988), pp. 237–246.

[138] S. Stipa, A. Ajay, D. Allaerts, and J. Brinkerhoff. “TOSCA – an Open-Source, Finite-Volume, Large-Eddy Simulation (LES) Environment for Wind Farm Flows”. In: *Wind Energy Science* 9.2 (2024), pp. 297–320.

[139] R. Stoll, J. A. Gibbs, S. T. Salesky, W. Anderson, and M. Calaf. “Large-Eddy Simulation of the Atmospheric Boundary Layer”. In: *Boundary-Layer Meteorology* 177.2 (2020), pp. 541–581.

[140] R. B. Stull, ed. *An Introduction to Boundary Layer Meteorology*. Dordrecht: Springer Netherlands, 1988.

[141] E. Taschner, M. Becker, R. Verzijlbergh, and J. Van Wingerden. “Comparison of Helix and Wake Steering Control for Varying Turbine Spacing and Wind Direction”. In: *Journal of Physics: Conference Series* 2767.3 (2024), p. 032023.

[142] E. Taschner, M. Folkersma, L. A Martínez-Tossas, R. Verzijlbergh, and J.-W. van Wingerden. “A New Coupling of a GPU-resident Large-Eddy Simulation Code with a Multiphysics Wind Turbine Simulation Tool”. In: *Wind Energy* 27.11 (2023), pp. 1152–1172.

[143] T. N. Theis and H.-S. P. Wong. “The End of Moore’s Law: A New Beginning for Information Technology”. In: *Computing in Science & Engineering* 19.2 (2017), pp. 41–50.

[144] L. Tophøj and H. Aref. “Instability of Vortex Pair Leapfrogging”. In: *Physics of Fluids* 25.1 (2013), p. 014107.

[145] F. Trigaux, P. Chatelain, and G. Winckelmans. “Investigation of Blade Flexibility Effects on the Loads and Wake of a 15-MW Wind Turbine Using a Flexible Actuator Line Method”. In: *Wind Energy Science* 9.8 (2024), pp. 1765–1789.

[146] F. Trigaux, T. Villeneuve, G. Dumas, and G. Winckelmans. “Near-Tip Correction Functions for the Actuator Line Method to Improve the Predicted Lift and Drag Distributions”. In: *Journal of Fluid Mechanics* 989 (2024), A1.

[147] N. Troldborg. “Actuator Line Modeling of Wind Turbine Wakes”. 2009.

[148] C. C. van Heerwaarden, B. J. H. van Stratum, T. Heus, J. A. Gibbs, E. Fedorovich, and J. P. Mellado. “MicroHH 1.0: A Computational Fluid Dynamics Code for Direct Numerical Simulation and Large-Eddy Simulation of Atmospheric Boundary Layer Flows”. In: *Geoscientific Model Development* 10.8 (2017), pp. 3145–3165.

[149] G. van Kuik and J. Peinke, eds. *Long-Term Research Challenges in Wind Energy - A Research Agenda by the European Academy of Wind Energy*. Vol. 6. Research Topics in Wind Energy. Cham: Springer International Publishing, 2016.

[150] A. A. W. van Vondelen, J. Ottenheym, A. K. Pamososuryo, S. T. Navalkar, and J.-W. van Wingerden. “Phase Synchronization for Helix Enhanced Wake Mixing in Downstream Wind Turbines”. In: *IFAC-PapersOnLine*. 22nd IFAC World Congress 56.2 (2023), pp. 8426–8431.

[151] P. Veers, K. Dykes, S. Basu, A. Bianchini, A. Clifton, P. Green, H. Holttinen, L. Kitzing, B. Kosovic, J. K. Lundquist, J. Meyers, M. O’Malley, W. J. Shaw, and B. Straw. “Grand Challenges: Wind Energy Research Needs for a Global Energy Transition”. In: *Wind Energy Science* 7.6 (2022), pp. 2491–2496.

[152] R. Verstappen. “When Does Eddy Viscosity Damp Subfilter Scales Sufficiently?” In: *Journal of Scientific Computing* 49.1 (2011), p. 94.

[153] R. Vinuesa and S. L. Brunton. “Enhancing Computational Fluid Dynamics with Machine Learning”. In: *Nature Computational Science* 2.6 (2022), pp. 358–366.

[154] R. Wagner, M. Courtney, J. Gottschall, and P. Lindelöw-Marsden. “Accounting for the Speed Shear in Wind Turbine Power Performance Measurement”. In: *Wind Energy* 14.8 (2011), pp. 993–1004.

[155] Y. Wang, B. T. MacCall, C. M. Hocut, X. Zeng, and H. J. S. Fernando. “Simulation of Stratified Flows over a Ridge Using a Lattice Boltzmann Model”. In: *Environmental Fluid Mechanics* 20.5 (2020), pp. 1333–1355.

[156] Z. Wang, J. Jacob, F. Marlow, and P. Sagaut. “Lattice Boltzmann Method-Based Large Eddy Simulation of the Influence of Atmospheric Stability on Wind Farm Wake”. In: *Physics of Fluids* 37.6 (2025), p. 065131.

[157] S. Watanabe and C. Hu. “Lattice Boltzmann Simulation for Wake Interactions of Aligned Wind Turbines Using Actuator Line Model with Turbine Control”. In: *Journal of Physics: Conference Series* 2767.5 (2024), p. 052020.

[158] T. Watanabe, K. Shimoyama, M. Kawashima, Y. Mizoguchi, and A. Inagaki. “Large-Eddy Simulation of Neutrally-Stratified Turbulent Flow Within and Above Plant Canopy Using the Central-Moments-Based Lattice Boltzmann Method”. In: *Boundary-Layer Meteorology* 176.1 (2020), pp. 35–60.

[159] A. T. White and C. K. Chong. “Rotational Invariance in the Three-Dimensional Lattice Boltzmann Method Is Dependent on the Choice of Lattice”. In: *Journal of Computational Physics* 230.16 (2011), pp. 6367–6378.

[160] S. Wilhelm, J. Jacob, and P. Sagaut. “An Explicit Power-Law-Based Wall Model for Lattice Boltzmann Method–Reynolds-averaged Numerical Simulations of the Flow around Airfoils”. In: *Physics of Fluids* 30.6 (2018), p. 065111.

[161] J. W. van Wingerden, P. A. Fleming, T. Göçmen, I. Eguinoa, B. M. Doekemeijer, K. Dykes, M. Lawson, E. Simley, J. King, D. Astrain, M. Iribas, C. L. Bottasso, J. Meyers, S. Raach, K. Kölle, and G. Giebel. “Expert Elicitation on Wind Farm Control”. In: *Journal of Physics: Conference Series* 1618 (2020), p. 022025.

[162] D. Wolf-Gladrow. “A Lattice Boltzmann Equation for Diffusion”. In: *Journal of Statistical Physics* 79.5 (1995), pp. 1023–1032.

[163] Y.-T. Wu and F. Porté-Agel. “Simulation of Turbulent Flow Inside and Above Wind Farms: Model Validation and Layout Effects”. In: *Boundary-Layer Meteorology* 146.2 (2013), pp. 181–205.

[164] J. C. Wyngaard. *Turbulence in the Atmosphere*. Cambridge University Press, 2010. 407 pp.

[165] W. Xian and A. Takayuki. “Multi-GPU Performance of Incompressible Flow Computation by Lattice Boltzmann Method on GPU Cluster”. In: *Parallel Computing*. Emerging Programming Paradigms for Large-Scale Scientific Computing 37.9 (2011), pp. 521–535.

[166] S. Xie. “An Actuator-line Model with Lagrangian-averaged Velocity Sampling and Piecewise Projection for Wind Turbine Simulations”. In: *Wind Energy* (2021), we.2619.

[167] J. Xu. “Wake Interaction of NREL Wind Turbines Using a Lattice Boltzmann Method”. In: *Sustainable Energy* 4.1 (2016), pp. 1–6.

[168] J. Yan, C. Möhrlen, T. Göçmen, M. Kelly, A. Wessel, and G. Giebel. "Uncovering Wind Power Forecasting Uncertainty Sources and Their Propagation through the Whole Modelling Chain". In: *Renewable and Sustainable Energy Reviews* 165 (2022), p. 112519.

[169] J. Zhang and X. Zhao. "Quantification of Parameter Uncertainty in Wind Farm Wake Modeling". In: *Energy* 196 (2020), p. 117065.

Acta Universitatis Upsaliensis

Digital Comprehensive Summaries of Uppsala Dissertations from the Faculty of Science and Technology 2596

Editor: The Dean of the Faculty of Science and Technology

A doctoral dissertation from the Faculty of Science and Technology, Uppsala University, is usually a summary of a number of papers. A few copies of the complete dissertation are kept at major Swedish research libraries, while the summary alone is distributed internationally through the series Digital Comprehensive Summaries of Uppsala Dissertations from the Faculty of Science and Technology. (Prior to January, 2005, the series was published under the title "Comprehensive Summaries of Uppsala Dissertations from the Faculty of Science and Technology".)

Distribution: publications.uu.se
urn:nbn:se:uu:diva-568037

ACTA UNIVERSITATIS
UPSALIENSIS
2025