



MicroHH 1.0: a computational fluid dynamics code for direct numerical simulation and large-eddy simulation of atmospheric boundary layer flows

Chiel C. van Heerwaarden^{1,2}, Bart J. H. van Stratum^{1,2}, Thijs Heus³, Jeremy A. Gibbs⁴, Evgeni Fedorovich⁵, and Juan Pedro Mellado²

¹Meteorology and Air Quality Group, Wageningen University, Wageningen, the Netherlands

²Max Planck Institute for Meteorology, Hamburg, Germany

³Cleveland State University, Cleveland, OH, USA

⁴Department of Mechanical Engineering, University of Utah, Salt Lake City, UT, USA

⁵University of Oklahoma, Norman, OK, USA

Correspondence to: Chiel C. van Heerwaarden (chiel.vanheerwaarden@wur.nl)

Received: 17 February 2017 – Discussion started: 21 March 2017

Revised: 19 June 2017 – Accepted: 26 June 2017 – Published: 28 August 2017

Abstract. This paper describes MicroHH 1.0, a new and open-source (www.microhh.org) computational fluid dynamics code for the simulation of turbulent flows in the atmosphere. It is primarily made for direct numerical simulation but also supports large-eddy simulation (LES). The paper covers the description of the governing equations, their numerical implementation, and the parameterizations included in the code. Furthermore, the paper presents the validation of the dynamical core in the form of convergence and conservation tests, and comparison of simulations of channel flows and slope flows against well-established test cases. The full numerical model, including the associated parameterizations for LES, has been tested for a set of cases under stable and unstable conditions, under the Boussinesq and anelastic approximations, and with dry and moist convection under stationary and time-varying boundary conditions. The paper presents performance tests showing good scaling from 256 to 32 768 processes. The graphical processing unit (GPU)-enabled version of the code can reach a speedup of more than an order of magnitude for simulations that fit in the memory of a single GPU.

1 Introduction

In this paper, we present a description of MicroHH 1.0, a new computational fluid dynamics code for the simulation of turbulent flows in doubly periodic domains, with a focus on those in the atmosphere. MicroHH is designed for the direct numerical simulation (DNS) technique but also supports the large-eddy simulation (LES) technique. Its applications range from neutral channel flows to cloudy atmospheric boundary layers in large domains. MicroHH is written in C++ and the graphical processing unit (GPU)-enabled parts of the code in NVIDIA's CUDA. The simulation algorithms have been designed and are written from scratch with the goal to create a fast and highly parallel code that is able to run on machines with more than 10 000 cores. It is a key requirement for the code to be able to perform DNS at very high Reynolds numbers or to conduct LES at very fine grids (grid spacing less than 1 m), or in domains that approach the synoptic scales (beyond 1000 km). We decided to start from scratch, in order to be able to use C++ and its extensive possibilities in object-oriented and metaprogramming. Furthermore, the implementation of a dynamical core that is fully fourth order in space, which is very beneficial for DNS, but to retain the option to switch to second-order accuracy for LES, required a new code design.

Even though we started from scratch, many of the ideas are the results of our experiences with other codes. Here, DALES

(Heus et al., 2010), UCLA-LES (Stevens et al., 2005), and PALM (Maronga et al., 2015), deserve a reference as MicroHH could not have been possible without those.

This paper is built up as follows: in Sect. 2, we provide a full description of the governing equations of the dynamical core, and their numerical implementation is discussed in Sect. 3. Subsequently, in Sect. 4, we present the parameterizations and their underlying assumptions. Section 5 discusses the technical details of the code, and Sects. 6 and 7 explain how to run the model and which output is generated. This is followed by a series of model tests on the validity and accuracy of the dynamical core in Sect. 8, and a series of more applied atmospheric flow cases based on previous studies (Sect. 9). Hereafter, the parallel performance is evaluated (Sect. 10). Then, an overview of published work with MicroHH is presented (Sect. 11), followed by the future plans (Sect. 12) and the concluding remarks (Sect. 13). Finally, there is a short description of where to get MicroHH, and where to find its tutorials and a selection of visualizations (see code availability section).

2 Dynamical core: governing equations

The dynamical core of MicroHH solves the conservation equations of mass, momentum, and energy under the anelastic approximation (Bannon, 1996). Under this approximation, the state variables density, pressure, and temperature are described as small fluctuations (denoted with a prime in this paper) from corresponding vertical reference profiles (denoted with subscript zero) that are functions of height only. This form of the approximation directly simplifies to the Boussinesq approximation if the reference density $\rho_0(z)$ is taken to be constant with height z . Consequently, MicroHH does not need separate implementations of Boussinesq and anelastic approximations. To facilitate the subsequent discussion of the conservation equations, we define the scale height for density H_ρ based on the reference density profile:

$$H_\rho \equiv \left(\frac{1}{\rho_0} \frac{d\rho_0}{dz} \right)^{-1}. \quad (1)$$

2.1 Conservation of mass

The conservation of mass is formulated using Einstein summation as

$$\frac{\partial \rho_0 u_i}{\partial x_i} = \rho_0 \frac{\partial u_i}{\partial x_i} + \rho_0 w H_\rho^{-1} = 0, \quad (2)$$

where u_i represents the components of the velocity vector (u, v, w) and x_i represents the components of the position vector (x, y, z) . This formulation conserves the reference mass, as density perturbations are ignored in the equation (Lilly, 1996).

Under the Boussinesq approximation ($H_\rho \rightarrow \infty$), Eq. (2) simplifies to conservation of volume:

$$\frac{\partial u_i}{\partial x_i} = 0. \quad (3)$$

2.2 Thermodynamic relations and conservation of momentum

The thermodynamic relation between the fluctuations of virtual potential temperature, pressure, and density under the anelastic approximation is (see Bannon, 1996 for its derivation)

$$\frac{\theta'_v}{\theta_{v0}} = \frac{p'}{\rho_0 g H_\rho} - \frac{\rho'}{\rho_0}, \quad (4)$$

where θ'_v is the perturbation virtual potential temperature, θ_{v0} the reference virtual potential temperature, p' is the perturbation pressure, g is the gravity acceleration, and ρ' is the perturbation density.

The corresponding momentum equation is written in the flux form in order to assure momentum conservation. The hydrostatic balance $dp_0/dz = -\rho_0 g$ has been subtracted, and Eq. (4) has been used to introduce potential temperature as the buoyancy variable to formulate the conservation of momentum as

$$\begin{aligned} \frac{\partial u_i}{\partial t} = & -\frac{1}{\rho_0} \frac{\partial \rho_0 u_i u_j}{\partial x_j} - \frac{\partial}{\partial x_i} \left(\frac{p'}{\rho_0} \right) \\ & + \delta_{i3} g \frac{\theta'_v}{\theta_{v0}} + \nu \frac{\partial^2 u_i}{\partial x_j^2} + F_i, \end{aligned} \quad (5)$$

where δ is the Kronecker delta, ν is the kinematic viscosity, and vector F_i represents external forces resulting from parameterizations or large-scale forcings. As Bannon (1996) showed, this formulation is energy-conserving in the sense that there is a consistent transfer between kinetic and potential energy.

Under the Boussinesq approximation, the two equations simplify to

$$\frac{\theta'_v}{\theta_{v0}} = -\frac{\rho'}{\rho_0}, \quad (6)$$

$$\frac{\partial u_i}{\partial t} = -\frac{\partial u_i u_j}{\partial x_j} - \frac{1}{\rho_0} \frac{\partial p'}{\partial x_i} + \delta_{i3} g \frac{\theta'_v}{\theta_{v0}} + \nu \frac{\partial^2 u_i}{\partial x_j^2} + F_i. \quad (7)$$

2.3 Pressure equation

The equation to acquire the pressure is diagnostic because density fluctuations are neglected in the mass conservation equation under the anelastic approximation (Eq. 2). To simplify the notation, we define a function $f(u_i)$ that contains all right-hand-side terms of Eq. (5), except the pressure gradient. To arrive at the equation that allows us to solve for the pressure, we multiply the equation with the base density ρ_0

and take its divergence. Conservation of mass ensures that the tendency term vanishes, and an elliptic equation for pressure remains:

$$\frac{\partial}{\partial x_i} \left[\rho_0 \frac{\partial}{\partial x_i} \left(\frac{p'}{\rho_0} \right) \right] = \frac{\partial \rho_0 f(u_i)}{\partial x_i}. \quad (8)$$

Under the Boussinesq approximation, the equation simplifies to

$$\frac{\partial^2}{\partial x_i^2} \left(\frac{p'}{\rho_0} \right) = \frac{\partial f(u_i)}{\partial x_i}. \quad (9)$$

In Sect. 3, we explain how these equations are solved numerically.

2.4 Conservation of an arbitrary scalar

The conservation equation of an arbitrary scalar ϕ is written in flux form:

$$\frac{\partial \phi}{\partial t} = -\frac{1}{\rho_0} \frac{\partial \rho_0 u_j \phi}{\partial x_j} + \kappa_\phi \frac{\partial^2 \phi}{\partial x_j^2} + S_\phi, \quad (10)$$

where κ_ϕ is the diffusivity of the scalar and S_ϕ represents sources and sinks of the variable.

2.5 Conservation of energy

MicroHH provides multiple options for the energy conservation equation. The conservation equation for potential temperature for dry dynamics θ can be written as

$$\frac{\partial \theta}{\partial t} = -\frac{1}{\rho_0} \frac{\partial \rho_0 u_j \theta}{\partial x_j} + \kappa_\theta \frac{\partial^2 \theta}{\partial x_j^2} + \frac{\theta_0}{\rho_0 c_p T_0} Q, \quad (11)$$

where κ_θ is the thermal diffusivity for heat, and Q represents external sources and sinks of heat. A second option for moist dynamics is available. This has an identical conservation equation, but with liquid water potential temperature θ_l rather than θ as the conserved variable (see Sect. 3.9 for details).

A third, more simplified mode, is available for dry dynamics under the Boussinesq approximation. Here, the equation of state (Eq. 6) can be eliminated and the conservation of momentum and energy can be written in terms of buoyancy $b \equiv (g/\theta_{v0})\theta'_v$ as

$$\frac{\partial u_i}{\partial t} + \frac{\partial u_i u_j}{\partial x_j} = -\frac{1}{\rho_0} \frac{\partial p'}{\partial x_i} + \delta_{i3} b + v \frac{\partial^2 u_i}{\partial x_j^2}, \quad (12)$$

$$\frac{\partial b}{\partial t} + \frac{\partial b u_j}{\partial x_j} = \kappa_b \frac{\partial^2 b}{\partial x_j^2} + Q_b, \quad (13)$$

with κ_b as the diffusivity for buoyancy and Q_b as an external buoyancy source. By using buoyancy, length and time remain as the only two dimensions, which proves convenient

for dimensional analysis. In this formulation, θ'_v is the fluctuation of the virtual potential temperature with respect to the surface value θ_{v0} . The consequence is that the buoyancy increases with height in a stratified atmosphere, analogously to the virtual potential temperature (see Garcia and Mellado, 2014, their Fig. B1 and van Heerwaarden and Mellado, 2016, their Fig. 7a)

With a slight modification to the definition of θ'_v , it is possible to study slope flows in periodic domains. We define θ'_v as the fluctuation with respect to a linearly stratified background profile $\theta_{v0} + (d\theta_v/dz)_0 z$. The background stratification in units of buoyancy is $N^2 \equiv (g/\theta_{v0})(d\theta_v/dz)_0$. If we work out the governing equations again and introduce a slope α (x axis pointing upslope; see Fedorovich and Shapiro, 2009, their Fig. 1) in the x direction, we find

$$\frac{\partial u}{\partial t} + \frac{\partial u_j u}{\partial x_j} = -\frac{1}{\rho_0} \frac{\partial p'}{\partial x} + \sin(\alpha) b + v \frac{\partial^2 u}{\partial x_j^2}, \quad (14)$$

$$\frac{\partial w}{\partial t} + \frac{\partial u_j w}{\partial x_j} = -\frac{1}{\rho_0} \frac{\partial p'}{\partial z} + \cos(\alpha) b + v \frac{\partial^2 w}{\partial x_j^2}, \quad (15)$$

$$\begin{aligned} \frac{\partial b}{\partial t} + \frac{\partial b u_j}{\partial x_j} &= \kappa_b \frac{\partial^2 b}{\partial x_j^2} - (u \sin(\alpha) \\ &+ w \cos(\alpha)) N^2 + Q_b, \end{aligned} \quad (16)$$

where the evolution equation of v is omitted, as it contains no changes.

3 Dynamical core: numerical implementation

3.1 Grid

MicroHH is discretized on a staggered Arakawa C-grid, where the scalars are located in the center of a grid cell and the three velocity components are at the faces. The code can work with stretched grids in the vertical dimension. The grid is initialized from a vertical profile that contains the heights of the cell centers. The locations of the faces are determined consistently with the spatial order of the interpolations that are described in Sect. 3.4. All spatial operators in the model, such as the advection and diffusion, default to the same order as the grid and can be overridden according to the user's wishes (see Sect. 6).

There is the option to apply a uniform translation velocity to the grid and thus to let the grid move with the flow. This so-called Galilean transformation is allowed as the Navier–Stokes equations are invariant under translation. It has the potential to allow for larger time steps and to increase the accuracy of simulations.

3.2 Three-dimensional fields

In order to solve the governing equations, MicroHH generates at initialization three-dimensional fields of the prognostic variables. These are the three velocity components (Eqs. 5

or 7) and the thermodynamic variables (Eqs. 11, 13, or 16). Furthermore, the user has the option to define additional passive scalars (Eq. 10). Each of the prognostic fields has an additional three-dimensional field assigned to store its tendency (see Sect. 3.3). Furthermore, a diagnostic field is assigned for the pressure, as well as three or four additional ones for intermediate computations. Newly implemented physical parameterizations have the option to request additional three-dimensional fields at initialization of the specific parameterization.

The generation of turbulence requires perturbations to the initial fields. MicroHH has two options to superimpose perturbations on any of the prognostic variables. These perturbations can be random noise of which the amplitude and location can be controlled, as well as two-dimensional rotating vortices with an axis aligned with the x or y dimension. The former option is the most commonly used method to start convective turbulence, whereas the latter is the default for neutral or stably stratified flows, which develop turbulence more easily from larger perturbations.

3.3 Time integration

The prognostic equations are solved using low-storage Runge–Kutta time integration schemes. Such schemes require two fields per variable: one that contains the actual value, which we denote with ϕ in this section, and one that represents the tendencies, denoted with $\delta\phi$. The code provides two options: a three-stage, third-order scheme (Williamson, 1980) and a five-stage, fourth-order scheme (Carpenter and Kennedy, 1994). Both can be written in the same generic form in semi-discrete formulation as

$$(\delta\phi)_n = f(\phi_n) + a_n(\delta\phi)_{n-1} \quad (17)$$

$$\phi_{n+1} = \phi_n + b_n \Delta t (\delta\phi)_n, \quad (18)$$

where f is a function that represents the computation of all right-hand-side terms, a_n and b_n are the coefficients for the Runge–Kutta method at stage n , and Δt is the time step. Expression $f(\phi_n)$ thus represents the actual tendency calculated using, for instance, Eqs. (5) or (10), whereas $(\delta\phi)_n$ is a composite of the actual tendency and those from the previous stages. In low-storage form, the tendencies of the previous stage $(\delta\phi)_{n-1}$ are retained and multiplied with a_n at the beginning of a stage, except for the first stage, where $a_1 = 0$.

For the third-order scheme, the vectors a_n and b_n are

$$a_n = \left\{ 0, -\frac{5}{9}, -\frac{153}{128} \right\}, \quad (19)$$

$$b_n = \left\{ \frac{1}{3}, \frac{15}{16}, \frac{8}{15} \right\}. \quad (20)$$

For the fourth-order scheme, the vectors a and b are

$$a_n = \left\{ 0, -\frac{567301805773}{1357537059087}, -\frac{2404267990393}{2016746695238}, \right. \\ \left. -\frac{3550918686646}{2091501179385}, -\frac{1275806237668}{842570457699} \right\} \quad (21)$$

$$b_n = \left\{ \frac{1432997174477}{9575080441755}, \frac{5161836677717}{13612068292357}, \frac{1720146321549}{2090206949498}, \right. \\ \left. \frac{3134564353537}{4481467310338}, \frac{2277821191437}{14882151754819} \right\}. \quad (22)$$

The reduced truncation error of the fourth-order scheme makes the scheme preferable over the third-order scheme under many conditions (see Sect. 8.2). The code can be run with a fixed Δt , as well as an adaptive time step based on the local flow velocities.

3.4 Building blocks of the spatial discretization

The spatial operators are based on finite differences. The code supports second-order and fourth-order accurate discretizations following Morinishi et al. (1998) and Vasilyev (2000). From Taylor series, spatial operators can be derived that constitute the building blocks of more advanced operators, such as the advection and diffusion operators. In the following subsections, we describe the elementary operators and the composite operators that can be derived from them. We have selected a set of examples that cover the relevant operators.

We define two second-order interpolation operators, one with a small stencil and one with a wide stencil, as

$$\phi_{i,j,k} \approx \bar{\phi}_{i,j,k}^{2x} \equiv \frac{\phi_{i-\frac{1}{2},j,k} + \phi_{i+\frac{1}{2},j,k}}{2}, \quad (23)$$

$$\phi_{i,j,k} \approx \bar{\phi}_{i,j,k}^{2xL} \equiv \frac{\phi_{i-\frac{3}{2},j,k} + \phi_{i+\frac{3}{2},j,k}}{2}. \quad (24)$$

Interpolations are marked with a bar. The superscript indicates the spatial order (2) and the direction (x), and has an extra qualifier L when it is taken using the wide stencil. The subscript indicates the position on the grid (i, j).

The gradient operators, denoted with letter δ , are defined in a similar way:

$$\frac{\partial\phi}{\partial x} \Big|_{i,j,k} \approx \delta^{2x}(\phi)_{i,j,k} \equiv \frac{\phi_{i+\frac{1}{2},j,k} - \phi_{i-\frac{1}{2},j,k}}{x_{i+\frac{1}{2}} - x_{i-\frac{1}{2}}} \quad (25)$$

$$\frac{\partial\phi}{\partial x} \Big|_{i,j,k} \approx \delta^{2xL}(\phi)_{i,j,k} \equiv \frac{\phi_{i+\frac{3}{2},j,k} - \phi_{i-\frac{3}{2},j,k}}{x_{i+\frac{3}{2}} - x_{i-\frac{3}{2}}}. \quad (26)$$

We use the Einstein summation in the operators. For instance, the divergence of vector $u_i|_{i,j,k}$ can be written as $\delta^{2x_i}(u_i)_{i,j,k}$.

The fourth-order operators, written down in the same notation, are defined as

$$\begin{aligned} \phi_{i,j,k} &\approx \overline{\phi}_{i,j,k}^{-4x} \\ &\equiv \frac{-\phi_{i-\frac{3}{2},j,k} + 9\phi_{i-\frac{1}{2},j,k} + 9\phi_{i+\frac{1}{2},j,k} - \phi_{i+\frac{3}{2},j,k}}{16}. \end{aligned} \quad (27)$$

The biased version of this operator (suffix b in the superscript) can be applied in the vicinity of the boundaries at the bottom and top. Here, we show the biased stencil that can be applied for vertical interpolation near the bottom:

$$\begin{aligned} \phi_{i,j,k} &\approx \overline{\phi}_{i,j,k}^{-4zb} \\ &\equiv \frac{5\phi_{i,j,k-\frac{1}{2}} + 15\phi_{i,j,k+\frac{1}{2}} - 5\phi_{i,j,k+\frac{3}{2}} + \phi_{i,j,k+\frac{5}{2}}}{16}. \end{aligned} \quad (28)$$

Note that we only write down the bottom boundary for brevity.

The centered and biased fourth-order gradient operators are

$$\begin{aligned} \left. \frac{\partial \phi}{\partial x} \right|_{i,j,k} &\approx \delta^{4x}(\phi)_{i,j,k} \\ &\equiv \frac{\phi_{i-\frac{3}{2},j,k} - 27\phi_{i-\frac{1}{2},j,k} + 27\phi_{i+\frac{1}{2},j,k} - \phi_{i+\frac{3}{2},j,k}}{x_{i-\frac{3}{2}} - 27x_{i-\frac{1}{2}} + 27x_{i+\frac{1}{2}} - x_{i+\frac{3}{2}}}, \end{aligned} \quad (29)$$

and

$$\begin{aligned} \left. \frac{\partial \phi}{\partial z} \right|_{i,j,k} &\approx \delta^{4zb}(\phi)_{i,j,k} \\ &\equiv \frac{-23\phi_{i,j,k-\frac{1}{2}} + 21\phi_{i,j,k+\frac{1}{2}} + 3\phi_{i,j,k+\frac{3}{2}} - \phi_{i,j,k+\frac{5}{2}}}{-23z_{k-\frac{1}{2}} + 21z_{k+\frac{1}{2}} + 3z_{k+\frac{3}{2}} - z_{k+\frac{5}{2}}}. \end{aligned} \quad (30)$$

3.5 Boundary conditions

The lateral boundaries in MicroHH are periodic. The bottom and top boundary conditions can be formulated in their most general form as the Robin boundary condition:

$$a\phi_s + b \left. \frac{\partial \phi}{\partial z} \right|_s = c, \quad (31)$$

with a , b , and c as constants. This gives the Dirichlet boundary condition when $a = 1$, $b = 0$, and the Neumann boundary condition when $a = 0$, $b = 1$.

MicroHH makes use of ghost cells in order to avoid the need of biased schemes for single interpolation or gradient operators near the wall. The values at the ghost cells are derived making use of the boundary conditions following Morinishi et al. (1998). The ghost cells for the Dirichlet boundary conditions in the second-order accurate discretization are

$$\phi_{-\frac{1}{2}} = 2c - \phi_{\frac{1}{2}}, \quad (32)$$

whereas those for the Neumann boundary condition are

$$\phi_{-\frac{1}{2}} = -c \left(-z_{-\frac{1}{2}} + z_{\frac{1}{2}} \right) + \phi_{\frac{1}{2}}. \quad (33)$$

In the case of the fourth-order scheme, we have two ghost cells, and therefore a second boundary condition is required. Here, we set the third derivative equal to zero following Morinishi et al. (1998). For the Dirichlet boundary condition we then acquire the following expressions for the ghost cells:

$$\phi_{-\frac{1}{2}} = \frac{8c - 6\phi_{\frac{1}{2}} + \phi_{\frac{3}{2}}}{3}, \quad (34)$$

$$\phi_{-\frac{3}{2}} = 8c - 6\phi_{\frac{1}{2}} + \phi_{\frac{3}{2}}, \quad (35)$$

whereas in the case of a Neumann boundary condition, we find

$$\phi_{-\frac{1}{2}} = -c \frac{z_{-\frac{3}{2}} - 27z_{-\frac{1}{2}} + 27z_{\frac{1}{2}} - z_{\frac{3}{2}}}{24} + \phi_{\frac{1}{2}}, \quad (36)$$

$$\phi_{-\frac{3}{2}} = -3c \frac{z_{-\frac{3}{2}} - 27z_{-\frac{1}{2}} + 27z_{\frac{1}{2}} - z_{\frac{3}{2}}}{24} + \phi_{\frac{3}{2}}. \quad (37)$$

3.6 Advection

We use the previously introduced notation to describe the more complex operators and expand them for illustration. The advection term is discretized in the flux form, where ϕ is an arbitrary scalar located in the center of the grid cell. In the second-order case, this gives the following discretization:

$$\begin{aligned} \left. \frac{\partial u\phi}{\partial x} \right|_{i,j,k} + \left. \frac{\partial v\phi}{\partial y} \right|_{i,j,k} &\approx \delta^{2x} \left(u\overline{\phi}^{2x} \right)_{i,j,k} + \delta^{2y} \left(v\overline{\phi}^{2y} \right)_{i,j,k} \\ &= \frac{u_{i+\frac{1}{2},j,k} \overline{\phi}_{i+\frac{1}{2},j,k}^{2x} - u_{i-\frac{1}{2},j,k} \overline{\phi}_{i-\frac{1}{2},j,k}^{2x}}{x_{i+\frac{1}{2}} - x_{i-\frac{1}{2}}} \\ &\quad + \frac{v_{i,j+\frac{1}{2},k} \overline{\phi}_{i,j+\frac{1}{2},k}^{2y} - v_{i,j-\frac{1}{2},k} \overline{\phi}_{i,j-\frac{1}{2},k}^{2y}}{y_{j+\frac{1}{2}} - y_{j-\frac{1}{2}}}. \end{aligned} \quad (38)$$

The discretization of the advection of the velocity components (see Eqs. 5 and 7) involves extra interpolations as the following example illustrates:

$$\begin{aligned} \left. \frac{\partial vu}{\partial x} \right|_{i,j,k} &= \delta^{2x} \left(\overline{v}^{2y} \overline{u}^{2x} \right)_{i,j,k} \\ &= \frac{\overline{v}_{i+\frac{1}{2},j,k}^{2y} \overline{u}_{i+\frac{1}{2},j,k}^{2x} - \overline{v}_{i-\frac{1}{2},j,k}^{2y} \overline{u}_{i-\frac{1}{2},j,k}^{2x}}{x_{i+\frac{1}{2}} - x_{i-\frac{1}{2}}}. \end{aligned} \quad (39)$$

In the standard fourth-order scheme, the scalar advection in flux form is represented by

$$\begin{aligned} \left. \frac{\partial u\phi}{\partial x} \right|_{i,j,k} &\approx \delta^{4x} \left(u\overline{\phi}^{4x} \right)_{i,j,k} \\ &= \left(u_{i-\frac{3}{2},j,k} \overline{\phi}_{i-\frac{3}{2},j,k}^{4x} - 27u_{i-\frac{1}{2},j,k} \overline{\phi}_{i-\frac{1}{2},j,k}^{4x} \right. \\ &\quad \left. + 27u_{i+\frac{1}{2},j,k} \overline{\phi}_{i+\frac{1}{2},j,k}^{4x} - u_{i+\frac{3}{2},j,k} \overline{\phi}_{i+\frac{3}{2},j,k}^{4x} \right) / \\ &\quad \left(x_{i-\frac{3}{2}} - 27x_{i-\frac{1}{2}} + 27x_{i+\frac{1}{2}} - x_{i+\frac{3}{2}} \right). \end{aligned} \quad (40)$$

Hereafter, we assume that operator notation is clear and only expand it where necessary.

MicroHH has a fully kinetic energy-conserving fourth-order advection scheme (Morinishi et al., 1998) available. The scheme is constructed by interpolation of two kinetic energy-conserving second-order discretizations to eliminate the second-order error (as illustrated below)

$$\frac{\partial u\phi}{\partial x} \Big|_{i,j,k} \approx \frac{9}{8}\delta^{2x}\left(u\bar{\phi}^{2x}\right)_{i,j,k} - \frac{1}{8}\delta^{2xL}\left(u\bar{\phi}^{2xL}\right)_{i,j,k}, \quad (41)$$

to ensure that velocity variances are conserved under advection.

Velocity interpolations, such as those in Eq. (39), still need to be performed with fourth-order accuracy (Eq. 27) in order to be fourth-order accurate (see Morinishi et al., 1998 for details). The expression

$$\frac{\partial vu}{\partial x} \Big|_{i,j,k} \approx \frac{9}{8}\delta^{2x}\left(\bar{v}^4\bar{u}^{2x}\right)_{i,j,k} - \frac{1}{8}\delta^{2xL}\left(\bar{v}^4\bar{u}^{2xL}\right)_{i,j,k} \quad (42)$$

includes, for instance, a combination of second- and fourth-order interpolations.

To increase the overall accuracy of the second-order advection operator, there is an option available to only increase the interpolation part to fourth order:

$$\frac{\partial u\phi}{\partial x} \Big|_{i,j,k} \approx \delta^{2x}\left(u\bar{\phi}^{4x}\right)_{i,j,k}. \quad (43)$$

3.7 Diffusion

We apply a discretization for diffusion that can be written as the divergence of a gradient, using the building blocks defined earlier in this section. As this operator is identical in all directions, we present it in one direction only:

$$\kappa_\phi \frac{\partial^2 \phi}{\partial x^2} \Big|_{i,j,k} \approx \kappa_\phi \delta^{2x}\left(\delta^{2x}(\phi)\right)_{i,j,k}, \quad (44)$$

$$\kappa_\phi \frac{\partial^2 \phi}{\partial x^2} \Big|_{i,j,k} \approx \kappa_\phi \delta^{4x}\left(\delta^{4x}(\phi)\right)_{i,j,k}. \quad (45)$$

On an equidistant grid, this provides the well-known second-order accurate operator for the second derivative:

$$\kappa_\phi \delta^{2x}\left(\delta^{2x}(\phi)\right)_{i,j,k} = \kappa_\phi \frac{\phi_{i-1,j,k} - 2\phi_{i,j,k} + \phi_{i+1,j,k}}{(\Delta x)^2}, \quad (46)$$

where Δx is the uniform grid spacing.

For the fourth-order accurate operator, a seven-point stencil is used:

$$\begin{aligned} & \kappa_\phi \delta^{4x}\left(\delta^{4x}(\phi)\right)_{i,j,k} \\ &= \frac{\kappa_\phi}{576(\Delta x)^2} (\phi_{i-3,j,k} - 54\phi_{i-2,j,k} + 783\phi_{i-1,j,k} \\ & \quad - 1460\phi_{i,j,k} + 783\phi_{i+1,j,k} - 54\phi_{i+2,j,k} + \phi_{i+3,j,k}). \end{aligned} \quad (47)$$

Whereas diffusion can be computed with fourth-order accuracy using a five-point stencil, we use a seven-point stencil, as it extends naturally to non-uniform grids as explained




Figure 1. Schematic of the diffusion discretization near the wall. The green node is the evaluation point at the center of the first cell above the wall, the red nodes are the stencil of the divergence operator, and yellow nodes show the stencils of the four gradient operators over which the divergence is evaluated. White nodes indicate the extent of the stencil.

in Castillo et al. (1995). The usage of a seven-point stencil requires special care near the walls. In Fig. 1, we show an example of how the second derivative in the vertical direction is computed for a scalar at the first model level (green node in Fig. 1). The calculation of the divergence (Fig. 1, red stencil) requires the gradient located at the first face below the wall (lowest red node in Fig. 1), which can only be acquired using the biased gradient operator (Eq. 30 and yellow stencil connected to lowest red node in Fig. 1). The extent of the complete stencil near the wall (white nodes; Fig. 1) is thus six points, rather than seven.

3.8 Pressure

Equations (8) and (9) are solved following the method of Chorin (1968). This is a fractional step method that first computes intermediate values of the velocity components for the next time step, based on all right-hand-side terms of the momentum conservation equation (Eq. 5):

$$u_i^*|_{i,j,k}^{t+1} = u_i|_{i,j,k}^t + \Delta t f_i|_{i,j,k}^t, \quad (48)$$

with the intermediate velocity components denoted with an asterisk.

The velocity values at the next time step can be computed as soon as the pressure is known, using

$$u_i|_{i,j,k}^{t+1} = u_i^*|_{i,j,k}^{t+1} - \Delta t \delta^{nx_i} \left(\frac{p}{\rho_0} \right) \Big|_{i,j,k}^t. \quad (49)$$

In order to compute the pressure, we multiply the previous equation with the reference density and take its gradient, ar-

riving at

$$\delta^{n x_i} (\rho_0 u_i) \Big|_{i,j,k}^{t+1} = \delta^{n x_i} (\rho_0 u_i^*) \Big|_{i,j,k}^{t+1} - \Delta t \delta^{n x_i} \left[\rho_0 \delta^{n x_i} \left(\frac{p}{\rho_0} \right) \right] \Big|_{i,j,k}^t, \quad (50)$$

where n indicates the spatial order, and the subscript i in superscript x_i indicates that $\delta^{n x_i}$ is a divergence operator. The left-hand side equals zero due to mass conservation at the next time step (Eq. 2). The resulting equation is the Poisson equation that is the discrete equivalent of Eq. (8). Rewriting this equation leads to

$$\frac{\delta^{n x_i} (\rho_0 u_i^*) \Big|_{i,j,k}^{t+1}}{\Delta t} = \delta^{n x_i} \left[\rho_0 \delta^{n x_i} \left(\frac{p}{\rho_0} \right) \right] \Big|_{i,j,k}^t. \quad (51)$$

To simplify the notation, we denote the left-hand-side term as ψ and the p/ρ_0 term on the right-hand side as π . Solving a Poisson equation is a global operation. Because the computed fields are periodic in the horizontal directions on an equidistant grid, and a Poisson equation is linear, we can perform a Fourier transform in the two horizontal directions:

$$\widehat{\psi}_{l,m,k} = -k_{*n}^2 \widehat{\pi}_{l,m,k} - l_{*n}^2 \widehat{\pi}_{l,m,k} + \delta^{n z} [\rho_0 \delta^{n z} (\widehat{\pi})]_{l,m,k}, \quad (52)$$

where Fourier-transformed variables are denoted with a hat, the spatial order of the operation with n , and the wavenumbers in the two horizontal dimensions x and y are l and m , respectively. Variables k_*^2 and l_*^2 are the squares of the modified wavenumbers:

$$-k_{*2}^2 \equiv 2 \frac{\cos(k \Delta x)}{(\Delta x)^2} - \frac{2}{(\Delta x)^2} \quad (53)$$

and

$$-k_{*4}^2 \equiv 2 \frac{\cos(3k \Delta x) - 54 \cos(2k \Delta x) + 783 \cos(k \Delta x)}{576(\Delta x)^2} - \frac{1460}{576(\Delta x)^2}, \quad (54)$$

where the former is the modified wavenumber for the second-order accurate solver and the latter is the wavenumber for the fourth-order one. Note that the coefficients correspond to those in Eqs. (46) and (47). Both expressions satisfy the limit $\lim_{\Delta x \rightarrow 0} k_{*n}^2 = k^2$, where n is the order of the scheme.

Solving Eq. (52) for $\widehat{\pi}$ requires solving a banded matrix for the vertical direction in which the walls are located. This matrix is tridiagonal for the second-order solver and heptadiagonal for the fourth-order solver. For this, a standard Thomas algorithm (Thomas, 1949) is used. After the pressure is acquired, inverse Fourier transforms are applied and subsequently the pressure gradient term (see Eqs. 5 and 7) is computed for all three components of the velocity tendency. Note that the computation of the corrected velocity components does not require a boundary condition for pressure (see Vreman, 2014 for details).

3.9 Thermodynamics

MicroHH supports the potential (θ) and liquid water potential (θ_l) temperature as thermodynamic variables (Sect. 2.5). The dry (θ) and moist (θ_l) thermodynamics are related through the use of a total specific humidity q_t , which is defined as the sum of the water vapor specific humidity (q_v) and the cloud liquid water specific humidity (q_l). In the absence of liquid water, $\theta_l = \theta$; in the presence of liquid water, the liquid water potential temperature is approximated as (Betts, 1973)

$$\theta_l \approx \theta - \frac{L_v}{c_p \Pi} q_l, \quad (55)$$

where L_v is the latent heat of vaporization, c_p the specific heat of dry air at constant pressure, and Π is the Exner function:

$$\Pi = \left(\frac{p}{p_{00}} \right)^{R_d/c_p}, \quad (56)$$

where p is the hydrostatic pressure, p_{00} a constant reference pressure, and R_d the gas constant for dry air. The cloud liquid water content is calculated as

$$q_l = \max(0, q_t - q_s), \quad (57)$$

where q_s is the saturation specific humidity:

$$q_s = \frac{\epsilon e_s}{p - (1 - \epsilon) e_s}, \quad (58)$$

with ϵ the ratio between the gas constant for dry air and the gas constant for water vapor (R_d/R_v), and e_s the saturation vapor pressure. The latter is approximated using a 10th order Taylor expansion at $T = 0^\circ\text{C}$ of the Arden Buck equation (Buck, 1981). q_l is adjusted iteratively to arrive at a consistent state where $q_v = q_s$. Finally, the virtual potential temperature (Eq. 5) is defined in MicroHH as

$$\theta_v \equiv \theta \left(1 - \left[1 - \frac{R_v}{R_d} \right] q_t - \frac{R_v}{R_d} q_l \right). \quad (59)$$

The base state pressure and density are calculated assuming a hydrostatic equilibrium: $dp_0 = -\rho_0 g dz$, with the density defined as $\rho_0 = p_0 / (R_d \Pi \theta_{v0})$. Integration with height results in

$$p_{0;k+1} = p_{0;k} \exp \left(\frac{-g(z_{k+1} - z_k)}{R_d \Pi \theta_{v0}} \right), \quad (60)$$

where θ_{v0} is the average virtual potential temperature between z_k and z_{k+1} . This equation is applied from a given surface pressure to the model top, alternating the calculations at the full and half model levels. That is, given the full thermodynamic state (pressure and density) at a full level k , the thermodynamic state can be advanced from the half level $k - \frac{1}{2}$

to $k + \frac{1}{2}$. Using the newly calculated state at $k + \frac{1}{2}$, pressure and density at $k + 1$ can be calculated.

The base state density ρ_0 that is used in the dynamical core (Sect. 2) is calculated using the initial virtual potential temperature profile and is not updated during the experiment. The density and hydrostatic pressure used in the moist thermodynamics can optionally be updated every time step, following the same procedure as explained in Boing (2014).

3.10 Rotation

The effects of a rotating reference frame on an f plane can be included through the Coriolis force. The acceleration due to the Coriolis force $F_{i,\text{cor}}$ is computed for the two horizontal velocity components (indices 1 and 2 in Eqs. 5 and 7) as

$$F_{1,\text{cor}}|_{i,j,k} = f_0 v_{i,j,k}, \quad (61)$$

$$F_{2,\text{cor}}|_{i,j,k} = -f_0 u_{i,j,k}, \quad (62)$$

with f_0 as Coriolis parameter specified by the user.

4 Physical parameterizations

4.1 Subfilter-scale model for large-eddy simulation

With the governing equations described in Sect. 2 it is possible to resolve the flow down to the scales where molecular viscosity acts. In many applications, however, such simulations are too costly. In that case, one may opt for large-eddy simulation (LES), where filtered equations are used to describe the largest scales of the flow, and the subfilter-scale motions are modeled. The LES implementation in MicroHH assumes very high Reynolds numbers in which the molecular viscosity is neglected. Filtering of the anelastic conservation of momentum equation (Eq. 5), with a tilde applied to denote filtered variables, leads to

$$\begin{aligned} \frac{\partial \tilde{u}_i}{\partial t} = & -\frac{1}{\rho_0} \frac{\partial \rho_0 \tilde{u}_i \tilde{u}_j}{\partial x_j} - \frac{\partial \pi}{\partial x_i} - \frac{1}{\rho_0} \frac{\partial \rho_0 \tau_{ij}}{\partial x_j} \\ & + \delta_{i3} g \frac{\tilde{\theta}'_v}{\theta_{v0}} + F_i. \end{aligned} \quad (63)$$

In this equation, a tensor τ_{ij} is defined as

$$\tau_{ij} \equiv \widetilde{u_i u_j} - \tilde{u}_i \tilde{u}_j - \frac{1}{3} (\widetilde{u_i u_i} - \tilde{u}_i \tilde{u}_i). \quad (64)$$

This is the anisotropic subfilter-scale kinematic momentum flux tensor. The isotropic part of the full momentum flux tensor has been added to the pressure, providing the modified pressure:

$$\pi \equiv \frac{\tilde{p}'}{\rho_0} + \frac{1}{3} (\widetilde{u_i u_i} - \tilde{u}_i \tilde{u}_i). \quad (65)$$

As τ_{ij} contains the filtered product of unfiltered velocity components, this quantity needs to be parameterized. MicroHH uses the Smagorinsky–Lilly (Lilly, 1968) model, in

which τ_{ij} is modeled as

$$\tau_{ij} = -K_m \left(\frac{\partial \tilde{u}_i}{\partial x_j} + \frac{\partial \tilde{u}_j}{\partial x_i} \right), \quad (66)$$

with K_m interpreted as the subfilter eddy diffusivity. This quantity is modeled as

$$K_m = \lambda^2 S \left(1 - \frac{g}{\theta_{v0}} \frac{\partial \tilde{\theta}_v}{\partial z} \right)^{\frac{1}{2}}, \quad (67)$$

and is proportional to the magnitude $S \equiv (2S_{ij}S_{ij})^{\frac{1}{2}}$ of the strain tensor S_{ij} , which is defined as

$$S_{ij} \equiv \frac{1}{2} \left(\frac{\partial \tilde{u}_i}{\partial x_j} + \frac{\partial \tilde{u}_j}{\partial x_i} \right). \quad (68)$$

The subfilter eddy diffusivity thus takes into account the local stratification and the turbulent Prandtl number Pr_t . The latter is set to one-third by default but can be overridden in the settings. The length scale λ is the mixing length defined following Mason and Thomson (1992), as

$$\frac{1}{\lambda^n} = \frac{1}{[\kappa(z+z_0)]^n} + \frac{1}{(c_s \Delta)^n}, \quad (69)$$

which is an arbitrary matching function (n is a free parameter, set to 2 in MicroHH) between the mixing length following wall scaling to the subfilter length scale (filter size) $\Delta \equiv (\Delta x \Delta y \Delta z)^{1/3}$, related to the grid spacing. The grid scale is used as an implicit filter; thus, no explicit filtering is applied. In the case of a high Reynolds number atmospheric LES with an unresolved near-wall flow, the vertical gradients of the horizontal velocity components $\partial \tilde{u}_{i,j} / \partial z$ in the strain tensor are replaced with the theoretical gradients predicted from Monin–Obukhov similarity theory. Evaluation of these gradients is explained in detail in Sect. 4.2.

The same approach is followed for all scalars, including the thermodynamic variables discussed in Sect. 2.5:

$$\frac{\partial \tilde{\phi}}{\partial t} = -\frac{1}{\rho_0} \frac{\partial \rho_0 \tilde{u}_j \tilde{\phi}}{\partial x_j} - \frac{1}{\rho_0} \frac{\partial \rho_0 R_{\phi,j}}{\partial x_j} + \tilde{S}_\phi. \quad (70)$$

The term $R_{\phi,j}$ refers to the subfilter flux of $\tilde{\phi}$ and is defined as

$$R_{\phi,j} = \widetilde{u_j \phi} - \tilde{u}_j \tilde{\phi}. \quad (71)$$

The subfilter-scale flux is parameterized in terms of the gradient

$$R_{\phi,j} = -\frac{K_m}{Pr_t} \frac{\partial \tilde{\phi}}{\partial x_j}. \quad (72)$$

4.2 Surface model

The LES implementation of MicroHH uses a surface model that is constrained to rough surfaces and high Reynolds numbers, which is a typical configuration for atmospheric flows. This model computes the surface fluxes of the horizontal momentum components and the scalars (including thermodynamic variables) using Monin–Obukhov similarity theory (MOST) (Wyngaard, 2010, his Sect. 10.2). MOST relates surface fluxes of variables to their near-surface gradients using empirical functions that depend on the height of the first model level z_1 divided by the Obukhov length L as an argument. Length L is defined as

$$L \equiv -\frac{u_*^3}{\kappa B_0}, \quad (73)$$

where u_* is the friction velocity, κ is the Von Karman constant, and B_0 is the surface kinematic buoyancy flux. L represents the height at which the buoyancy production/destruction of turbulence kinetic energy equals the shear production. In MicroHH, we use a local implementation of MOST, i.e., each grid point has its own value of L . This choice can lead to an overestimation of near-surface wind due to violation of the MOST assumption of horizontal homogeneity (Bou-Zeid et al., 2005, their Fig. 18), but it allows for a more straightforward extension to heterogeneous land surfaces.

Following MOST, the friction velocity u_* and the momentum fluxes may be related to the near-surface wind gradient as

$$\frac{\kappa z_1}{u_*} \frac{\partial U}{\partial z} \approx -\frac{\kappa z_1 u_*}{u'w'} \frac{\partial \tilde{u}}{\partial z} \approx -\frac{\kappa z_1 u_*}{v'w'} \frac{\partial \tilde{v}}{\partial z} \approx \phi_m \left(\frac{z_1}{L} \right), \quad (74)$$

where U is defined as $\sqrt{\tilde{u}^2 + \tilde{v}^2}$, and $\overline{u'w'}$ and $\overline{v'w'}$ as the surface momentum fluxes for the two wind components. These relationships can be integrated from the roughness length z_{0m} to z_1 resulting in

$$u_* = f_m (U_1 - U_0), \quad (75)$$

$$\overline{u'w'} = -u_* f_m (\tilde{u}_1 - \tilde{u}_0), \quad (76)$$

$$\overline{v'w'} = -u_* f_m (\tilde{v}_1 - \tilde{v}_0), \quad (77)$$

with f_m defined as

$$f_m \equiv \frac{\kappa}{\ln \left(\frac{z_1}{z_{0m}} \right) - \Psi_m \left(\frac{z_1}{L} \right) + \Psi_m \left(\frac{z_{0m}}{L} \right)}, \quad (78)$$

with Ψ_m described in Eqs. (83) and (85).

The same procedure for scalars is followed, with

$$\frac{\kappa z_1 u_*}{\phi'w'} \frac{\partial \tilde{\phi}}{\partial z} = \phi_h \left(\frac{z_1}{L} \right), \quad (79)$$

and in integrated form:

$$\overline{\phi'w'} = u_* f_h (\tilde{\phi}_1 - \tilde{\phi}_0), \quad (80)$$

with

$$f_h \equiv \frac{\kappa}{\ln \left(\frac{z_1}{z_{0h}} \right) - \Psi_h \left(\frac{z_1}{L} \right) + \Psi_h \left(\frac{z_{0h}}{L} \right)}, \quad (81)$$

with Ψ_h described in Eqs. (83) and (85).

The functions ϕ_m , ϕ_h , Ψ_m , and Ψ_h are empirical and depend on the static stability of the atmosphere. Under unstable conditions, we follow (Wilson, 2001; Wyngaard, 2010):

$$\phi_{m,h} = \left(1 + \gamma_{m,h} |\zeta|^{2/3} \right)^{-1/2}, \quad (82)$$

$$\Psi_{m,h} = 3 \ln \left(\frac{1 + \phi_{m,h}^{-1}}{2} \right), \quad (83)$$

where ζ is the ratio of a height and the Obukhov length L , $\gamma_m = 3.6$, and $\gamma_h = 7.9$. Under stable conditions, we use (Högström, 1988; Wyngaard, 2010)

$$\phi_{m,h} = 1 + \lambda_{m,h} \zeta, \quad (84)$$

$$\Psi_{m,h} = -\lambda_{m,h} \zeta, \quad (85)$$

where $\lambda_m = 4.8$ and $\lambda_h = 7.8$.

With the equations above, the surface fluxes, surface values, and near-surface gradients can be computed but only if the Obukhov length L is known. The surface model calculates the Obukhov length by relating the dimensionless parameter z_1/L to a Richardson number. The employed formulation of the Richardson number depends on the chosen boundary condition in the model. Three possible options are available:

- The first option is fixed momentum fluxes and a fixed surface buoyancy flux. Both the friction velocity u_* and the surface buoyancy flux B_0 are specified. Under these conditions, we define the Richardson number Ri_a equal to z_1/L ; L can be computed directly from the expression

$$Ri_a \equiv \frac{z_1}{L} = -\frac{\kappa z_1 B_0}{u_*^3}. \quad (86)$$

- The second option is a fixed horizontal velocity U_0 at the surface and a fixed surface buoyancy flux B_0 . The friction velocity u_* is unknown. Now, L needs to be retrieved from the implicit relationship:

$$Ri_b \equiv \frac{z_1}{L} f_m^3 = -\frac{\kappa z_1 B_0}{(U_1 - U_0)^3}. \quad (87)$$

- The third option is a fixed surface velocity U_0 and a fixed surface buoyancy b_0 . With this boundary condition, the surface values of the horizontal velocities and

the buoyancy are given, and both u_* and the surface buoyancy flux B_0 are unknown. L is then retrieved from

$$Ri_c \equiv \frac{z_1}{L} \frac{f_m^2}{f_h} = \frac{\kappa z_1 (\tilde{b}_1 - \tilde{b}_0)}{(U_1 - U_0)^2}. \quad (88)$$

In the event of the latter two options, a solver is required to find the value of L that satisfies the equation, as f_m (Eq. 78) and f_h (Eq. 81) both depend on L as well. For performance reasons, we have created a lookup-table-based approach that relates L to the Richardson number. The lookup table has 10^4 entries, of which 90% is spaced uniformly between $z_1/L = -5$ to 5. The remaining 10% are used to stretch the negative range up to $z_1/L = -10^4$ to allow for the correct free convection limit.

4.3 Large-scale forcings

4.3.1 Pressure force

MicroHH provides two options to introduce a large-scale pressure force into the model. The first is to enforce a constant mass flux through the domain in the streamwise direction. In this approach, the desired large-scale velocity U_f is set, and the corresponding pressure gradient is computed. We follow here the approach of van Reeuwijk (2007). In this approach, the u component of the horizontal momentum equation (Eq. 5) is volume averaged to acquire

$$\frac{\langle u \rangle^{n+1} - \langle u \rangle^n}{\Delta t} = \langle f_1 \rangle - \left\langle \frac{\partial}{\partial x} \left(\frac{p}{\rho_0} \right) \right\rangle + F_{p;ls}, \quad (89)$$

where brackets indicate a volume average, f_1 contains all the right-hand-side terms of the u component of the conservation of momentum, except for the dynamic pressure, which is contained in the second term. The large-scale pressure force $F_{p;ls}$, which is to be computed, is the last term. We can now set $\langle u \rangle^{n+1} = U_f$ to explicitly set the volume-averaged velocity in the next time step. Furthermore, the volume-averaged horizontal pressure gradient vanishes, because of the periodic boundary condition, which makes $F_{p;ls}$ the only unknown. The acquired pressure force $F_{p;ls}$ will be added as an external force in the equation of zonal velocity (F_1 in Eqs. 5 and 7).

The second option is to enforce a large-scale pressure force through the geostrophic wind components u_g and v_g , in combination with rotation, with the accelerations of the two horizontal velocity components $F_{i,p;ls}$ calculated as

$$F_{1,p;ls}|_{i,j,k} = -f_0 v_{g;k}, \quad (90)$$

$$F_{2,p;ls}|_{i,j,k} = f_0 u_{g;k}, \quad (91)$$

where $u_{g;k}$ and $v_{g;k}$ are user-specified vertical profiles of geostrophic wind components.

4.3.2 Large-scale sources and sinks

Large-scale sources and sinks, representing, for instance, large-scale advection or radiative cooling, can be prescribed for each variable separately. The user has to provide vertical profiles of large-scale sources and sinks $S_{\phi;ls}$ that are added to the total tendencies.

4.3.3 Large-scale vertical velocity

A second method of introducing large-scale thermodynamic effects is through the inclusion of a large-scale vertical velocity. In this case, each scalar gets an additional source term $S_{\phi,w,ls}$ of the form

$$S_{\phi,w,ls}|_{i,j,k} = -w_{ls;k} \delta^{2x} (\langle \phi \rangle_k), \quad (92)$$

where $w_{ls;k}$ is a user-specified vertical profile of large-scale vertical velocity and $\langle \phi \rangle_k$ is the horizontally averaged vertical profile at height z_k for scalar ϕ . The tendency term is not applied to the momentum variables.

4.4 Buffer layer

MicroHH has the option to damp gravity waves in the top of the simulation domain in a so-called buffer layer. The source term $S_{\phi,buf}$ associated with the damping at grid cell i, j, k is calculated for an arbitrary variable ϕ as

$$S_{\phi,buf}|_{i,j,k} = \frac{\phi_{i,j,k} - \phi_{0;k}}{\tau_{d;k}}, \quad (93)$$

where ϕ_0 is taken from a user-specified vertical reference profile, and timescale τ_d is a measure for the strength of the damping. It varies with height and is calculated at height z_k following

$$\tau_{d;k}^{-1} = \sigma \left(\frac{z_k - z_{b;bot}}{z_{b;top} - z_{b;bot}} \right)^\beta, \quad (94)$$

where σ is the damping frequency chosen by the user and β an exponent that describes the shape of the vertical profile of the damping frequency, which is always zero at the bottom ($z_{b;bot}$) and σ at the top ($z_{b;top}$) of the layer.

5 Technical details of the code

5.1 Code structure

MicroHH is written in C++ and uses elements of object-oriented programming and template metaprogramming. The code components are written in classes that define the interface. Inheritance is used to allow for specializations of classes. This way of organizing the code has two advantages: it minimizes switches and it allows code components and their extensions to reside in their own file, which increases code clarity and facilitates the merging of new code

extensions. High performance of computational kernels is achieved by executing kernels in their own function, with explicit inclusions of the `restrict` keyword to notify the compiler that fields do not overlap in memory. Furthermore, compiler-specific pragmas are used to aid vectorization on Intel compilers.

5.2 GPU

MicroHH is enabled to run on fast and energy-efficient graphical processing units (GPUs). This promising technique has been pioneered in atmospheric flows by Schalkwijk et al. (2012) and has shown its potential in weather forecasting (Schalkwijk et al., 2015). The implementation is based on NVIDIA's CUDA. The CPU and GPU version reside in the same code base, where the GPU implementation is activated with the help of precompiler statements. The philosophy is that the entire model is initialized on the CPU and that the GPU implementation is only activated right before starting the main time loop. At that moment, the required fields are copied in double precision accuracy to the GPU, and the entire time integration is done there. Synchronization only happens when statistics have to be computed or when restart files or cross sections of flow fields are saved to disk, to ensure high performance. The design choice to do the entire initialization on the CPU minimizes the amount of GPU code and therefore allows for maintaining a single code base for the CPU and GPU code.

5.3 Parallelization

The code uses the Message Passing Interface (MPI) in order to run on a large number of cores. The three-dimensional simulation domain is split into vertically oriented columns standing on a two-dimensional grid. The code assigns one MPI task to each grid column using the `MPI_Cart_create` function and uses this grid to detect the IDs of neighboring processes. In order to avoid complex packing routines, we make use of MPI datatypes wherever possible. The MPI calls are hidden in an interface to avoid the need to manually write MPI calls.

The input/output (IO) is entirely based on MPI-IO, the parallel IO framework that comes with MPI, to ensure that three-dimensional fields and two-dimensional cross sections are stored as single files. We have chosen MPI-IO in order to limit the number of files resulting from simulations on a large number of processes and to allow for restarts on a different number of processes. In order to keep complexity of the IO as low as possible, we make use of the `MPI_Sub_array` function in combination with `MPI_File_write_all` in order to write the fields.

5.4 External dependencies

MicroHH depends on several external software tools or libraries. It uses the CMake (<https://cmake.org>) build system

for the generation of Makefiles. CMake allows for parallel builds, which minimizes the compilation time, and it is easy to add configurations for different machines. Furthermore, the FFTW3 library (Frigo and Johnson, 2005) is used for the computation of fast Fourier transforms. The statistical routines make use of netCDF software developed by UCAR/Unidata (<http://doi.org/10.5065/D6H70CW6>). In order to run the provided test cases and their output scripts, a Python (<https://www.python.org>) installation including the NumPy (van der Walt et al., 2011; <http://www.numpy.org>) and Matplotlib (Hunter, 2007; <https://matplotlib.org>) modules is required. Automatic documentation generation can be done using Doxygen (<http://doxygen.org>), but this is optional.

6 Running simulations

In order to run a simulation with MicroHH, a sequence of steps needs to be taken. Each simulation has an `.ini` file that contains the settings of the simulation, a `.prof` file that contains the (initial) vertical profiles of all required variables, and, if time-varying boundary conditions are desired, a file with the prescribed time evolution for all time-varying boundary conditions. MicroHH provides a document (`doc/input.pdf`) that contains an overview of all possible options that can be specified in the `.ini` file.

To prepare a simulation with the name `test_simulation`, MicroHH needs to be run in the following way:

```
./microhh init test_simulation
```

where it is assumed that `test_simulation.ini` and `test_simulation.prof` are available in the directory where the command is triggered. This procedure will create the initial fields of all prognostic variables and save the required fields for those model components that need to save their state to guarantee bitwise identical restarts.

After the previously described initialization phase, the execution of

```
./microhh run test_simulation
```

will start the actual simulation. Depending on the chosen output intervals, the simulation will create restart files, statistics, cross sections, and field dumps. MicroHH can restart from any time at which the restart files are saved.

The last mode in which the code can run is the post-processing mode. By running

```
./microhh post test_simulation
```

the code will generate statistics from saved restart files at a specified time interval. This mode allows the user to create new statistics and calculate those from saved data without having to rerun the simulation.




Figure 2. Convergence of the spatial discretization error in the two-dimensional Taylor–Green vortex. Subscript 2 indicates the second-order scheme, subscript 4 the most accurate fourth-order scheme, and subscript 4M the fully energy-conserving fourth-order scheme. The dashed black line is the reference for second-order convergence; the dotted black lines indicate fourth-order convergence.

7 Model output

7.1 Statistics

A large set of output statistics can be computed during runtime at user-specified time intervals. The statistics module provides vertical profiles of means, second-, third- and fourth-order moments of all prognostic variables, as well as advective and diffusive fluxes. Furthermore, there are multiple diagnostic variables, such as the pressure, the pressure variance, and its vertical flux. The thermodynamics generate their own statistics based on the chosen option. The moist thermodynamics provides a large set of cloud statistics. There is a separate module for budget statistics, which provides the budgets of all components of the Reynolds stress tensor, and those of the variance and vertical flux of the thermodynamic variables.

One of the key features of the MicroHH statistics routine is that an arbitrary mask can be passed into the routine over which the statistics are calculated. This allows, for instance, for a very simple way of computing conditional statistics in updrafts or clouds, which is demonstrated later in Sect. 9.2.

7.2 Two- and three-dimensional output

It is possible to save two-dimensional cross sections and three-dimensional fields of any of the prognostic and diagnostic variables of the model, as well as of derived variables. This output can be made at user-specified time intervals. Cross sections can be made in any chosen xy , xz , and yz planes. Derived variables (any arbitrary function of existing model variables), can be easily added to the code by the user.




Figure 3. Time evolution of the kinetic energy change ΔKE during 1000 time units of random noise advection for the RK3 and RK4 time integration schemes with three different time steps (a). Kinetic energy change convergence of the temporal discretization for the RK3 and RK4 schemes (b).

8 Validation of the dynamical core

In this section, we present a series of cases intended to validate MicroHH under a wide range of settings. Each of these test cases is available in the `cases/` directory in the MicroHH repository, where all detailed settings can be found (see code availability section). Below, we present only the most relevant information per case.

8.1 Taylor–Green vortex

The two-dimensional Taylor–Green vortex (`cases/taylorgreen`) presents an ideal test case for a dynamical core, as it has an analytical solution even though it is nonlinear. This flow is composed of two rotating vortices whose evolution in a domain $[0, 1; 0, 0.5]$ is described with

$$u(x, z, t) = \sin(2\pi x) \cos(\pi z) f(t), \quad (95)$$

$$w(x, z, t) = \cos(2\pi x) \sin(\pi z) f(t), \quad (96)$$

$$p(x, z, t) = \frac{1}{4} (\sin(4\pi x) + \sin(4\pi z)) f(t)^2, \quad (97)$$

where $f(t) = 8\pi^2 \nu t$.

We use the analytical form at $t = 0$ as the initial condition and run this case for one vortex rotation ($t = 1$), with




Figure 4. Normalized numerical Prandtl model solutions for velocity u (a) and buoyancy b (b) compared to their analytical counterparts.

$\nu = (800\pi^2)^{-1}$. We compare the result against the analytical solution for a set of grid spacings and with the second-order and fourth-order dynamical cores; for the latter we compare the most accurate advection scheme and the fully energy-conserving one.

Figure 2 shows the error convergence of the simulations. The error for a variable ϕ is computed as $\sum \Delta x \Delta z |\phi_{i,k} - \phi_{\text{ref},i,k}|$ over the two-dimensional domain, where Δx and Δz are the uniform grid spacings used in this case and ϕ_{ref} is the analytical solution. All variables converge according to the order of the numerical scheme. The fourth-order dynamical core loses accuracy at fine grid spacings. This is due to the boundary condition for the vertical velocity that sacrifices an order of accuracy to ensure global mass conservation (Morinishi et al., 1998).

8.2 Kinetic energy conservation and time accuracy

The second test of the dynamical core consists of combined evaluation of kinetic energy ($\text{KE} \equiv \frac{1}{2}(u^2 + v^2 + w^2)$) conservation and time accuracy (`cases/conservation`). In this experiment, we run the model with only the advection and pressure solver enabled and advect random noise through the domain for 1000 s. These tests have been conducted with the third- and fourth-order Runge–Kutta schemes. We have chosen the fourth-order spatial discretization in order to demonstrate its energy conservation.

The loss of kinetic energy as a function of time is shown in Fig. 3a. The fourth-order scheme results in a smaller energy




Figure 5. Velocity means (a) and variances (b) for Moser et al. (1999) channel flow case at a Reynolds τ of 590. The dashed vertical lines mark the spectra locations. Height z is normalized with u_τ/ν ; velocities with u_τ^{-1} .

loss for the same time step and a faster convergence. The error-convergence plot (Fig. 3b) shows that the error convergence is in accordance with the order of the respective scheme. Furthermore, it illustrates the fact that, if high accuracy in time is desired, the five-stage, fourth-order scheme is less expensive than the three-stage, third-order scheme. For instance, at a Δt of 10, the error of the fourth-order scheme is approximately equal to the error of the third-order scheme at a Δt of 2.5. To reach this accuracy, the fourth-order scheme needs only 5 steps per 10 time units, whereas the third-order scheme needs 12 steps.

8.3 Laminar anabatic flow

To test the buoyancy routine and the option to put the domain on a slope, a Prandtl-type anabatic slope flow (Prandtl, 1942) has been simulated (`cases/prandtlslope`). In this test case, the surface is inclined at an angle of 30° and a linearly stratified atmosphere ($N = 1 \text{ s}^{-1}$) is heated from below with a fixed surface buoyancy flux of $0.005 \text{ m}^2 \text{ s}^{-3}$.

The fluid, which was initially at rest, goes through a series of decaying oscillations after the buoyancy flux is applied at the surface. Eventually, it reaches the steady state corresponding to the Prandtl model solution. Numerical integration was performed sufficiently long for the oscillation amplitude to become a small fraction of the amplitude of the first oscillation. Comparison of horizontal wind u and buoyancy b of analytical and numerical solutions is shown in Fig. 4. For both variables, the solutions closely agree with each other.

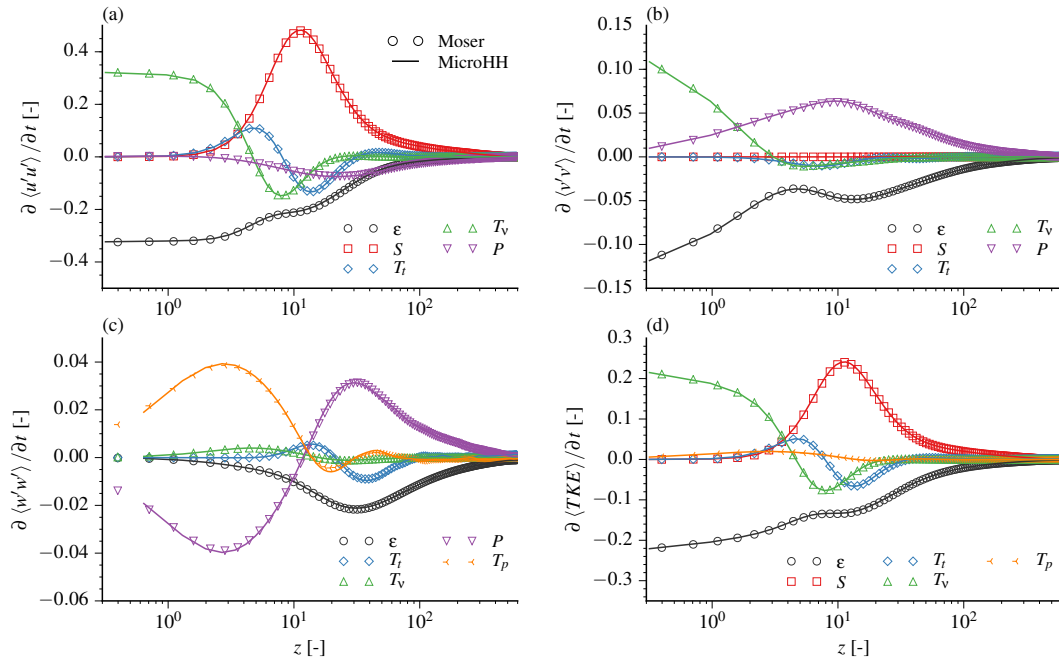


Figure 6. Budgets of variances and turbulence kinetic energy (TKE, defined as $(u'^2 + v'^2 + w'^2)/2$) compared against Moser et al. (1999)'s reference data at a Reynolds τ of 590. Height z is normalized with u_τ/ν ; the variances and TKE budget with ν/u_τ^4 .

8.4 Turbulent channel flow

For fully turbulent flows, the numerical solutions cannot be compared against any analytical test cases. Therefore, we validate our results against a channel flow at a Reynolds τ number of 590 (Moser et al., 1999) for means, variances, spectra, and second-order budget statistics (*cases/moser590*). The case is run at a resolution of $768 \times 384 \times 256$ grid points. The original numerical simulation data of Moser et al. (1999) have been produced on a $384 \times 384 \times 256$ grid with spectral schemes in the horizontal dimensions and Chebyshev polynomials in the vertical.

Figure 5a shows the normalized horizontally averaged streamwise velocity. The normalized rms values of all three velocity components are presented in Fig. 5b. All plotted variables show a perfect match with the data and are indistinguishable from Moser's data. In order to further assess the accuracy of the data, we show the second-order budgets of the variances in Fig. 6. Also here, the match with the reference data is excellent, which indicates that the whole range of spatial scales in the flow is represented well and that the fourth-order scheme is well able to pick up the small-scale details of the flow.

The findings in the previous paragraph are further corroborated by the spectra shown in Fig. 7. Over the whole range of scales, the match between our simulation and that of Moser et al. (1999) is excellent. Note that the spectra from Moser's simulation display an increase in pressure variance at the highest wavenumbers. This increase is the result of aliasing

errors at high wavenumbers that are typical for the spectral schemes that Moser et al. (1999) used.

8.5 Turbulent katabatic flow

The final evaluation of the dynamical core without parameterizations enabled is based on the direct numerical simulation of a turbulent katabatic flow. Here, a buoyancy-driven slope flow is simulated following the setup of Fedorovich and Shapiro (2009) (*cases/drycblslope*). A flow over a slope inclined at an angle α of 60° is simulated with a fixed uniform surface buoyancy flux of $-0.5 \text{ m}^2 \text{ s}^{-3}$. The simulation is performed in a domain of $0.64 \text{ m} \times 0.64 \text{ m} \times 1.6 \text{ m}$ using a uniform grid of $256 \times 256 \times 640$ points. The initial state is a fluid at rest with a linear buoyancy stratification N of 1 s^{-1} . No-slip boundary conditions are applied at the bottom; free-slip at the top.

Turbulent motion starts quickly after the buoyancy flux is applied at the surface. It is characterized by random, large-amplitude fluctuations of velocity and buoyancy fields in the near-slope core region and shows quasi-periodic oscillatory behavior at larger distances from the slope. Mean profiles of the along-slope velocity component and buoyancy, as well as profiles of second-order turbulence statistics, such as kinematic turbulent fluxes of momentum and buoyancy, and velocity component and buoyancy fluctuation variances, were evaluated by averaging the simulated flow fields spatially over the along-slope planes and temporally over five oscillation periods beyond the transition stage.

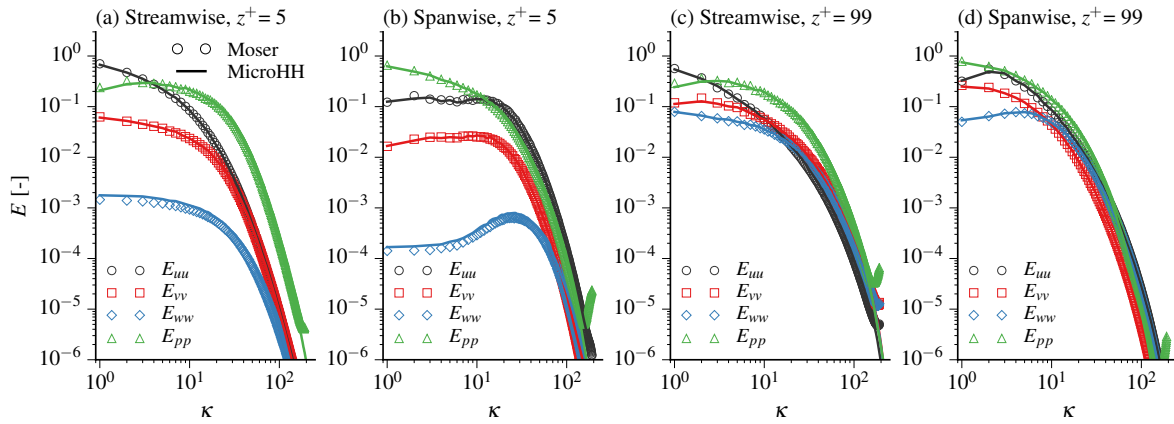


Figure 7. Spectra of all velocity components and pressure compared against Moser et al. (1999)'s reference data at a Reynolds τ of 590. The velocity spectra are normalized with u_{τ}^{-2} ; the pressure spectra with u_{τ}^{-4} .

For comparison, the same katabatic flow case was reproduced using the numerical code (hereafter referred to as FS09) that was employed to simulate turbulent slope flows in Shapiro and Fedorovich (2008) and Fedorovich and Shapiro (2009). In that code, the time advancement was performed with an Asselin-filtered second-order leapfrog scheme (Durran, 2013). The spatial discretizations are identical to the second-order accurate ones of MicroHH.

Numerical results obtained with both numerical codes testify that stable environmental stratification in combination with negative surface buoyancy forcing in the katabatic flow leads to an effective suppression of vertical turbulent exchange in the flow region adjacent to the slope. This suppression results in a shallow near-surface sublayer with strong buoyancy gradients (Fig. 8a) capped by a narrow jet with peak velocity located very close to the ground (Fig. 8b). Further comparison has been performed on the slope-normal fluxes of momentum and buoyancy (not shown), where a nearly perfect match has been reproduced as well.

9 Validation of atmospheric large-eddy simulations

9.1 Dry convective boundary layer with strong inversion

As a first test case of MicroHH in LES mode, we present that of Sullivan and Patton (2011) (*cases/sullivan2011*). This is a dry clear convective boundary layer that grows into a linearly stratified atmosphere with a very strong capping inversion (see Fig. 9a). The system is heated from the bottom by applying a constant kinematic heat flux of 0.24 K m s^{-1} . The domain size is $5120 \text{ m} \times 5120 \text{ m} \times 2048 \text{ m}$. Gravity wave damping has been applied in the top 25 % of the domain. Simulations have been run for 3 h with three spatial resolutions. The time-averaged profiles have been calculated over the last hour.

The results show the formation of a well-mixed layer with an overlying capping inversion (see Fig. 9a) and the associated linear heat flux profile with negative flux values in the entrainment zone (see Fig. 9b). The change of both quantities with resolution highlights the intrinsic challenge in resolving a boundary layer with an inversion layer that is stronger than the numerical schemes can accurately resolve. At coarse resolution, the strong inversion leads to an unphysical overshoot of the potential temperature flux above the boundary layer top (see Fig. 9b). This overshoot leads to a numerical mixed layer on top of the entrainment zone that vanishes with increasing resolution.

9.2 BOMEX

The BOMEX shallow cumulus case (Siebesma et al., 2003) (*cases/bomex*), S03 hereafter, provides the opportunity to evaluate the moist thermodynamics (see Sect. 3.9) and large-scale forcings. We have repeated the case as described in S03 at the original resolution of the study ($100 \text{ m} \times 100 \text{ m} \times 40 \text{ m}$) and at a higher resolution ($10 \text{ m} \times 10 \text{ m} \times 9.375 \text{ m}$).

This case produces non-precipitating shallow cumulus. It has a large-scale cooling applied that represents radiation, as well as a large-scale drying to allow the atmosphere to relax to a steady state. In addition, a large-scale vertical velocity is applied over a certain height range to reproduce the appropriate synoptic conditions.

The simulation is run for 6 h. Statistics are recorded during the final hour, including conditional statistics for the cloud-covered fields ($q_l > 0$) and for the cloud cores ($q_l > 0$ and $\theta_v > 0$). The vertical profile of area coverage and profiles of horizontally averaged liquid water potential temperature θ_l , total water q_t , and vertical velocity w are shown in Fig. 10. All mean and conditionally sampled statistics at the original resolution are predominantly within 1 standard deviation of the ensemble mean of data from all models that participated




Figure 8. Profile of the mean along-slope velocity (a) and buoyancy (b) as predicted by MicroHH and FS09.

in S03. The horizontally averaged vertical velocities in the cloud and cloud core decrease considerably with an increase in resolution.

9.3 GABLS1

To evaluate the LES mode for stable atmospheric conditions, the GABLS1 LES intercomparison case (Beare et al., 2006) (*cases/gabls1*) was reproduced. The boundary layer develops in this case from a shallow, well-mixed layer into a weakly stable boundary layer, driven by a prescribed negative tendency of the surface temperature over a total integration time of 9 h. The Boussinesq approximation is used, the advection scheme uses fourth-order accurate interpolations (Eq. 27), and the Smagorinsky subgrid turbulence scheme is set up with a Smagorinsky constant equal to 0.12 and a subgrid turbulent Prandtl number of unity. The experiments are performed at two different resolutions with grid cells of 2 and 6.25 m, and compared to the models which participated in the study of Beare et al. (2006).

Figure 11 shows the domain and time-averaged (over a period from 28 800 to 32 400 s) vertical profiles of potential temperature ($\langle\theta\rangle$) and the velocity component ($\langle u \rangle$), and also time series of the boundary layer depth (z_{ABL}) and friction velocity (u_*). At the largest grid spacing of 6.25 m, it takes approximately 2 h for the flow to become turbulent, as is evident from the delayed boundary layer growth and abrupt changes in u_* . Nonetheless, typical features like the low-level jet (Fig. 11b) are well reproduced, and all statistics are predominantly within the range of results from Beare et al. (2006). With the grid spacing reduced to 2 m, the flow




Figure 9. Vertical profiles of horizontally averaged potential temperature (a) and normalized kinematic heat flux (b). The boundary layer depth z_i is the location of the maximum vertical gradient in the potential temperature profile shown in (a).

becomes turbulent nearly instantaneously, but the resulting boundary layer depth and surface friction velocity are on the low side compared to the five models from Beare et al. (2006) which were run at this resolution.

10 Performance

10.1 CPU

The parallel performance of MicroHH has been evaluated in strong- (*cases/strongscaling*) and weak-scaling (*cases/weaksaling*) experiments. The case used is direct numerical simulation of a buoyancy-driven atmospheric boundary layer based on van Heerwaarden et al. (2014). For each simulation in the scaling experiments, a series of time steps is performed, and the mean cost per step is calculated over the series. The strong-scaling experiment has been performed on the Leibniz-Rechenzentrum der Bayerischen Akademie der Wissenschaften (LRZ)'s SuperMUC¹ machine (phase 1 thin node eight-core Sandy Bridge-EP Xeon E5-2680 8C, two processors per node, Infiniband FDR10 interconnect). In this experiment, simulations were performed on $1024 \times 1024 \times 1024$ and $2048 \times 2048 \times 1024$ grid points, with the number of processes increased throughout the scaling experiment. The weak-scaling experiment has been performed on Forschungszentrum Jülich's Juqueen² machine (IBM PowerPC A2, 1.6 GHz, 16 cores per node, 5D Torus network, 40 GBps). In this experiment, a fixed $64 \times 32 \times 1024$ grid is assigned to each processor, and throughout the experiment the domain size is increased. The results of both experiments are shown in Fig. 12.

¹<https://www.lrz.de/services/compute/supermuc/>

²http://www.fz-juelich.de/ias/jsc/EN/Expertise/Supercomputers/JUQUEEN/JUQUEEN_node.html




Figure 10. BOMEX LES intercomparison (S03). Shown are the domain mean, and conditionally sampled cloud ($q_l > 0$) and cloud core ($q_l > 0$ and $b - \langle b \rangle > 0$) vertical profiles of (a) area coverage of cloud and cloud core, (b) liquid water potential temperature, (c) total specific humidity, and (d) vertical velocity. The results are averaged over $t = 18\,000\text{--}21\,600$ s. The shaded area denotes the mean ± 1 standard deviation of the participating models from S03; the solid and dashed lines the results from MicroHH, using the original (solid) and a higher-resolution (dashed) setup.




Figure 11. GABLS1 LES intercomparison (Beare et al., 2006). Shown are the vertical profiles of (a) potential temperature and (b) u component of the velocity, and time series of the (c) boundary layer depth and (d) surface friction velocity. The shaded areas show the range in the results from the models that participated in the Beare et al. (2006) study. The dotted black lines show the initial conditions.

The strong-scaling experiment shows that increasing the number of processors leads to faster simulations. The speedup is initially close to linear, but each consecutive increase in the number of cores makes the model less efficient. Based on these results, we conclude that for the chosen test case and for the used supercomputers, a work load of approximately 2×10^6 grid points per core is the best balance between speed and computational efficiency.

The weak scaling shows almost 90 % efficiency going from 512 to 8192 cores; beyond that, the scaling falls off to 80 %. This can be explained by physical properties of the machine; beyond 8192 cores, a simulation no longer fits on one midplane (a physical unit consisting of 8192 cores), leading to slower communication.

10.2 Performance GPU (CUDA) implementation

The GPU implementation of MicroHH allows for fast simulations on a single device. The current state-of-the-art GPUs

feature 12 GB of memory; thus, simulations of maximally $512 \times 512 \times 512$ grid points of a flow with three velocity components, pressure, two scratch fields for temporary storage, and a single scalar fit in memory. Within this experiment, we compare thus GPU simulations that do not need communication against CPU simulations that require communication between cores and nodes. The reason for doing so is that nearly all of the simulations of the presented results in Sects. 8 and 9 fit within the memory of a single GPU.

To test the performance of such simulations, the performance of MicroHH on an NVIDIA Quadro K6000 (using CUDA 6.5) has been compared against the Max Planck Institute for Meteorology’s cluster Thunder (two Intel Xeon E5-2670 CPUs per node, 16 cores per node). Three benchmark cases have been chosen: the BOMEX moist convection case on grids of 64^3 , 128^3 , 256^3 , and $512^2 \times 384$, and the channel flow cases of Moser et al. (1999) at Reynolds τ numbers of 180 and 590.

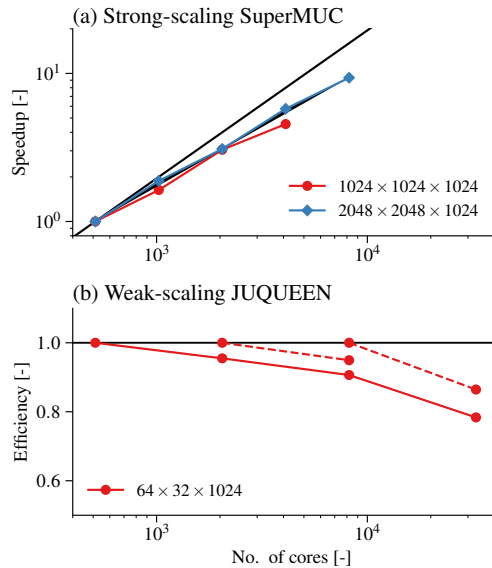


Figure 12. Speedup from a strong-scaling experiment (a). Efficiency from a weak-scaling experiments (b). Black lines indicate perfect speedup and efficiency. The dashed red lines show the efficiency change relative to the previous measurement.

The results shown in Table 1 point to the great potential of GPU computing. For the considered cases, which all fit on a single GPU, it takes at least 32 cores to reach equal performance. Only at 64 cores, the CPU simulations are notably faster. Therefore, for simulations that fit into its memory, the GPU provides an excellent alternative for the CPU, especially because the GPU is very energy efficient.

11 Published studies

To date, several studies have been published that make use of MicroHH or data generated with MicroHH. The work of van Heerwaarden et al. (2014) focused on the scaling of flow over heterogeneously heated land surfaces using DNS and LES, Gentine et al. (2015) used LES to study the structure of the inversion of a convective boundary layer, van Heerwaarden and Mellado (2016) developed scaling laws for the convective boundary layer over a surface with a constant temperature from DNS data, McColl et al. (2017) improved surface-layer similarity under mildly convective conditions with the help of DNS data, and Umphrey et al. (2017) used DNS data produced with MicroHH as a reference for their simulations of slope flow.

12 Future plans

There are several ongoing projects to extend the model. Currently, a parameterization for microphysics has been developed, and an interactive land surface model is under devel-

Table 1. Speedup of GPU simulation compared to respective CPU simulation performed on n cores.

Case	$n = 1$	$n = 16$	$n = 32$	$n = 64$
B64	18.49	1.93	1.14	0.95
B128	28.01	2.98	1.51	0.92
B256	27.76	3.02	1.59	0.91
B512	29.88	3.03	1.56	0.86
M180	21.57	2.17	1.13	0.69
M600	22.55	2.25	1.06	0.60

opment. In addition, the immersed boundary method following Tseng and Ferziger (2003) is being implemented to allow for simulations of flow over obstacles, urban canopy, and hills. Furthermore, preliminary experiments have been performed to include a domain-specific language (DSL) to enable the expression of complex finite difference operators in simple and compact syntax (<https://github.com/microhh/stencilbuilder/>). This development has shown great potential for two reasons. First, the DSL prevents implementation errors, as the explicit indexing in computational kernels with spatial operators can be omitted. Second, the DSL allows for simple implementation of system-specific tuning, such as loop tiling or OpenMP.

13 Conclusions

This paper has presented a full description of MicroHH, a new computational fluid dynamics code for simulations of turbulent flows in the atmospheric boundary layer. The governing equations and their implementation have been presented, and a broad validation under a wide range of settings has been shown. MicroHH delivers the expected error convergence of the spatial and temporal schemes, and has proven to be mass, momentum, and energy conserving. The code delivers good performance in weak- and strong-scaling experiments. Its current limitations are the absence of horizontal boundary conditions other than periodic, and the limited set of available physical parameterizations. Both limitations will be addressed in future versions of the code.

Code availability. MicroHH has its own website at <http://microhh.org>. The code is hosted at GitHub and can be accessed either via the website or directly from <https://github.com/microhh/microhh>. The GitHub website includes a wiki with several tutorials, including one to compile and run the code. The GitHub repository is coupled to Zenodo, which provides DOIs for released software. The release on which the reference paper is based is found at <http://dx.doi.org/10.5281/zenodo.822842> (van Heerwaarden et al., 2017). A selection of visualizations can be viewed at the MicroHH channel on Vimeo (<https://vimeo.com/channels/microhh/>).

Appendix A: Physical constants

Table A1 presents an overview of the chosen values for physical constants in the code.

Table A1. Overview of used constants.

Symbol	Description	Value	Units
κ	Von Karman constant	0.4	–
g	Gravitational acceleration	9.81	m s^{-2}
c_p	Specific heat of dry air at constant pressure	1005	$\text{J kg}^{-1} \text{K}^{-1}$
p_{00}	Reference pressure	1×10^5	Pa
R_d	Gas constant for dry air	287.04	$\text{J K}^{-1} \text{kg}^{-1}$
R_v	Gas constant for water vapor	461.5	$\text{J K}^{-1} \text{kg}^{-1}$
L_v	Latent heat of vaporization	2.5×10^6	J kg^{-1}

Appendix B: Nomenclature

Symbol	Description	Units
ρ_0	Reference density	kg m^{-3}
H_ρ	Scale height for density	m
u_i	Components of velocity vector (u, v, w)	m s^{-1}
x_i	Components of position vector (x, y, z)	m
p'	Perturbation pressure	Pa
p_0	Reference pressure	Pa
ρ'	Perturbation density	kg m^{-3}
p'	Perturbation pressure	Pa
p_0	Reference pressure	Pa
θ'_v	Perturbation virtual potential temperature	K
θ_{v0}	Reference virtual potential temperature	K
g	Gravity acceleration	m s^{-2}
ν	Kinematic viscosity	$\text{m}^2 \text{s}^{-1}$
κ	Scalar diffusivity	$\text{m}^2 \text{s}^{-1}$
F_i	External acceleration vector	m s^{-2}
S	External sources and sinks	variable dependent
θ	Potential temperature	K
θ_l	Liquid water potential temperature	K
b	Buoyancy	m s^{-2}
T_0	Reference absolute temperature	K
Q	Heat input	$\text{J m}^{-3} \text{s}^{-1}$
α	Slope of surface	rad
N	Brunt–Väisälä frequency	s^{-1}

Competing interests. The authors declare that they have no conflict of interest.

Acknowledgements. Version 1.0 of MicroHH was completed thanks to extensive discussions and exchange of code with many of the developers of DALES (Huug Ouwersloot, Stephan de Roode, Arnold Moene, and Jordi Vilà-Guerau de Arellano), PALM (Björn Maronga and Siegfried Raasch), UCLA-LES (Linda Schlemmer and Bjorn Stevens), and ICON (Leonidas Linares-Daquis). We thank Alan Shapiro for the discussions on the bottom boundary condition for pressure for buoyancy-driven flows and for his comments on the manuscript. We thank Elie Bou-Zeid and an anonymous reviewer for their constructive comments on the manuscript. The authors gratefully acknowledge the Gauss Centre for Supercomputing (GCS) for providing computing time for a GCS Large-Scale Project on the GCS share of the supercomputer JUQUEEN at the Jülich Supercomputing Centre (JSC).

The article processing charges for this open-access publication were covered by the Max Planck Society.

Edited by: Paul Ullrich

Reviewed by: Elie Bou-Zeid and one anonymous referee

References

- Bannon, P. R.: On the anelastic approximation for a compressible atmosphere, *J. Atmos. Sci.*, 53, 3618–3628, 1996.
- Beare, R. J., Macvean, M. K., Holtslag, A. A., Cuxart, J., Esau, I., Golaz, J. C., Jimenez, M. A., Khairoutdinov, M., Kosovic, B., Lewellen, D., and Lund, T. S.: An intercomparison of large-eddy simulations of the stable boundary layer, *Bound.-Lay. Meteorol.*, 118, 247–272, 2006.
- Betts, A. K.: Non-precipitating cumulus convection and its parameterization, *Q. J. Roy. Meteor. Soc.*, 99, 178–196, 1973.
- Boing, S. J.: The interaction of deep convective clouds and their environment, TU Delft, Delft University of Technology, the Netherlands, 2014.
- Bou-Zeid, E., Meneveau, C., and Parlange, M.: A scale-dependent Lagrangian dynamic model for large eddy simulation of complex turbulent flows, *Phys. Fluids*, 17, 025105, <https://doi.org/10.1063/1.1839152>, 2005.
- Buck, A. L.: New equations for computing vapor pressure and enhancement factor, *J. Appl. Meteorol.*, 20, 1527–1532, 1981.
- Carpenter, M. H. and Kennedy, C. A.: Fourth-order 2N-storage Runge-Kutta schemes, Tech. Rep. TM-109112, NASA Langley Research Center, 1994.
- Castillo, J. E., Hyman, J. M., Shashkov, M. J., and Steinberg, S.: The sensitivity and accuracy of fourth order finite-difference schemes on nonuniform grids in one dimension, *Comput. Math. Appl.*, 30, 41–55, 1995.
- Chorin, A. J.: Numerical solution of the Navier-Stokes equations, *Math. Comput.*, 22, 745–762, 1968.
- Durran, D. R.: Numerical methods for wave equations in geophysical fluid dynamics, Vol. 32, Springer Science & Business Media, 2013.
- Fedorovich, E. and Shapiro, A.: Structure of numerically simulated katabatic and anabatic flows along steep slopes, *Acta Geophys.*, 57, 981–1010, 2009.
- Frigo, M. and Johnson, S. G.: The design and implementation of FFTW3, *Proc. IEEE*, 93, 216–231, 2005.
- Garcia, J. R. and Mellado, J. P.: The two-layer structure of the entrainment zone in the convective boundary layer, *J. Atmos. Sci.*, 71, 1935–1955, <https://doi.org/10.1175/JAS-D-13-0148.1>, 2014.
- Gentine, P., Bellon, G., and van Heerwaarden, C. C.: A Closer Look at Boundary Layer Inversion in Large-Eddy Simulations and Bulk Models: Buoyancy-Driven Case, *J. Atmos. Sci.*, 72, 728–749, <https://doi.org/10.1175/JAS-D-13-0377.1>, 2015.
- Heus, T., van Heerwaarden, C. C., Jonker, H. J. J., Pier Siebesma, A., Axelsen, S., van den Dries, K., Geoffroy, O., Moene, A. F., Pino, D., de Roode, S. R., and Vilà-Guerau de Arellano, J.: Formulation of the Dutch Atmospheric Large-Eddy Simulation (DALES) and overview of its applications, *Geosci. Model Dev.*, 3, 415–444, <https://doi.org/10.5194/gmd-3-415-2010>, 2010.
- Högström, U.: Non-dimensional wind and temperature profiles in the atmospheric surface layer: a re-evaluation, *Bound.-Lay. Meteorol.*, 42, 55–78, 1988.
- Hunter, J. D.: Matplotlib: A 2D Graphics Environment, *Comput. Sci. Eng.*, 9, 90–95, <https://doi.org/10.1109/MCSE.2007.55>, 2007.
- Lilly, D.: A comparison of incompressible, anelastic and Boussinesq dynamics, *Atmos. Res.*, 40, 143–151, 1996.
- Lilly, D. K.: Models of cloud-topped mixed layers under a strong inversion, *Q. J. Roy. Meteor. Soc.*, 94, 292–309, 1968.
- Maronga, B., Gryschka, M., Heinze, R., Hoffmann, F., Kanani-Sühring, F., Keck, M., Ketelsen, K., Letzel, M. O., Sühring, M., and Raasch, S.: The Parallelized Large-Eddy Simulation Model (PALM) version 4.0 for atmospheric and oceanic flows: model formulation, recent developments, and future perspectives, *Geosci. Model Dev.*, 8, 2515–2551, <https://doi.org/10.5194/gmd-8-2515-2015>, 2015.
- Mason, P. J. and Thomson, D. J.: Stochastic backscatter in large-eddy simulation of boundary layers, *J. Fluid Mech.*, 242, 51–78, 1992.
- McColl, K. A., van Heerwaarden, C. C., Katul, G. G., Gentine, P., and Entekhabi, D.: Role of large eddies in the breakdown of the Reynolds analogy in an idealized mildly unstable atmospheric surface layer, *Q. J. Roy. Meteor. Soc.*, <https://doi.org/10.1002/qj.3077>, online first, 2017.
- Morinishi, Y., Lund, T. S., Vasilyev, O. V., and Moin, P.: Fully conservative higher order finite difference schemes for incompressible flow, *J. Comput. Phys.*, 143, 90–124, 1998.
- Moser, R. D., Kim, J., and Mansour, N. N.: Direct numerical simulation of turbulent channel flow up to $Re = 590$, *Phys. Fluids*, 11, 943–945, 1999.
- Prandtl, L.: Führer durch die Strömungslehre, Vieweg and Sohn, Braunschweig, 1942.
- Schalkwijk, J., Griffith, E. J., Post, F. H., and Jonker, H. J.: High-performance simulations of turbulent clouds on a desktop PC: Exploiting the GPU, *B. Am. Meteorol. Soc.*, 93, 307–314, 2012.
- Schalkwijk, J., Jonker, H. J., Siebesma, A. P., and Van Meijgaard, E.: Weather forecasting using GPU-based large-eddy simulations, *B. Am. Meteorol. Soc.*, 96, 715–723, 2015.

- Shapiro, A. and Fedorovich, E.: Coriolis effects in homogeneous and inhomogeneous katabatic flows, *Q. J. Roy. Meteor. Soc.*, 134, 353–370, 2008.
- Siebesma, A. P., Bretherton, C. S., Brown, A., Chlond, A., Cuxart, J., Duynkerke, P. G., Jiang, H., Khairoutdinov, M., Lewellen, D., Moeng, C. H., and Sanchez, E.: A large eddy simulation inter-comparison study of shallow cumulus convection, *J. Atmos. Sci.*, 60, 1201–1219, 2003.
- Stevens, B., Moeng, C.-H., Ackerman, A. S., Bretherton, C. S., Chlond, A., de Roode, S., Edwards, J., Golaz, J.-C., Jiang, H., Khairoutdinov, M., Kirkpatrick, M. P., Lewellen, D. C., Lock, A., Müller, F., Stevens, D. E., Whelan, E., and Zhu, P.: Evaluation of Large-Eddy Simulations via Observations of Nocturnal Marine Stratocumulus, *Mon. Weather Rev.*, 133, 1443–1462, 2005.
- Sullivan, P. P. and Patton, E. G.: The effect of mesh resolution on convective boundary layer statistics and structures generated by large-eddy simulation, *J. Atmos. Sci.*, 68, 2395–2415, 2011.
- Thomas, L. H.: Elliptic problems in linear difference equations over a network, *Watson Sci. Comput. Lab. Rept.*, Columbia University, New York, 1, 1949.
- Tseng, Y.-H. and Ferziger, J. H.: A ghost-cell immersed boundary method for flow in complex geometry, *J. Comput. Phys.*, 192, 593–623, 2003.
- Umphrey, C., DeLeon, R., and Senocak, I.: Direct Numerical Simulation of Turbulent Katabatic Slope Flows with an Immersed-Boundary Method, *Bound.-Lay. Meteorol.*, 164, 367–382, <https://doi.org/10.1007/s10546-017-0252-3>, 2017.
- van der Walt, S., Colbert, S. C., and Varoquaux, G.: The NumPy Array: A Structure for Efficient Numerical Computation, *Comput. Sci. Eng.*, 13, 22–30, <https://doi.org/10.1109/MCSE.2011.37>, 2011.
- van Heerwaarden, C. C. and Mellado, J. P.: Growth and decay of a convective boundary layer over a surface with a constant temperature, *J. Atmos. Sci.*, 73, 2165–2177, 2016.
- van Heerwaarden, C. C., Mellado, J. P., and De Lozar, A.: Scaling laws for the heterogeneously heated free convective boundary layer, *J. Atmos. Sci.*, 71, 3975–4000, 2014.
- van Heerwaarden, C. C., van Stratum, B. J. H., and Heus, T.: *microhh/microhh: 1.0.0*, Zenodo, <https://doi.org/10.5281/zenodo.822842>, 2017.
- van Reeuwijk, M.: Direct simulation and regularization modeling of turbulent thermal convection, TU Delft, Delft University of Technology, Delft, the Netherlands, 2007.
- Vasilyev, O. V.: High order finite difference schemes on non-uniform meshes with good conservation properties, *J. Comput. Phys.*, 157, 746–761, 2000.
- Vreman, A. W.: The projection method for the incompressible Navier–Stokes equations: The pressure near a no-slip wall, *J. Comput. Phys.*, 263, 353–374, 2014.
- Williamson, J. H.: Low-Storage Runge-Kutta schemes, *J. Comput. Phys.*, 35, 48–56, 1980.
- Wilson, D. K.: An alternative function for the wind and temperature gradients in unstable surface layers, *Bound.-Lay. Meteorol.*, 99, 151–158, 2001.
- Wyngaard, J. C.: *Turbulence in the Atmosphere*, Cambridge University Press, Cambridge, UK, 2010.