

A Nested Multi-Scale System Implemented in the Large-Eddy Simulation Model PALM model system 6.0

Antti Hellsten¹, Klaus Ketelsen², Matthias Sühling³, Mikko Auvinen¹, Björn Maronga^{3,6}, Christoph Knigge⁴, Fotios Barmpas⁵, Georgios Tsegas⁵, Nicolas Moussiopoulos⁵, and Siegfried Raasch³

¹Finnish Meteorological Institute, P.O.Box 503, FI-00101, Helsinki, Finland

²Software Consultant, Beethovenstr. 29A, 12247 Berlin, Germany

³Leibniz University Hannover, Institute of Meteorology and Climatology, Herrenhäuser Strasse 2, 30419 Hannover, Germany

⁴Deutscher Wetterdienst, Frankfurter Straße 135, D-63067 Offenbach, Germany

⁵Aristotle University Thessaloniki, P.O.Box 483, GR-54124, Thessaloniki, Greece

⁶University of Bergen, Geophysical Institute, Postboks 7803, 5020 Bergen, Norway

Correspondence to: A. Hellsten (antti.hellsten@fmi.fi)

Abstract.

Large-eddy simulation provides a physically sound approach to study complex turbulent processes within the atmospheric boundary layer including urban boundary layer flows. However, such flow problems often involve a large separation of turbulent scales, requiring a large computational domain and very high grid resolution near the surface features, leading to prohibitive computational costs. To overcome this problem, an online LES-LES nesting scheme is implemented into the PALM model system 6.0. The hereby documented and evaluated nesting method is capable of supporting multiple child domains which can be nested within their parent domain either in a parallel or recursively cascading configuration. The nesting system is evaluated by simulating first a purely convective boundary layer flow system and then three different neutrally-stratified flow scenarios with increasing order of topographic complexity. The results of the nested runs are compared with corresponding non-nested high- and low-resolution results. The results reveal that the solution accuracy within the high-resolution nest domain is clearly improved as the solutions approach the non-nested high-resolution reference results. In obstacle-resolving LES, the two-way coupling becomes problematic as antlerpolation introduces a regional discrepancy within the obstacle canopy of the parent domain. This is remedied by introducing canopy-restricted antlerpolation where the operation is only performed above the obstacle canopy. The test simulations make evident that this approach is the most suitable coupling strategy for obstacle-resolving LES. The performed simulations testify that nesting can reduce the CPU time up to 80% compared to the fine-resolution reference runs while the computational overhead from the nesting operations remained below 16% for the two-way coupling approach and significantly less for the one-way alternative.

1 Introduction

Large-Eddy Simulation (LES) has been used for basic research of atmospheric boundary layer (ABL) phenomena using idealized model setups for decades. Now, it is becoming an important method in applied research on realistic, very detailed and complicated flow systems such as urban ABL problems, e.g. (Bitter and Hanna, 2003; Tseng et al., 2006; Bou-Zeid et al.,

2009; Tominaga and Stathopoulos, 2013; Giometto et al., 2016; Buccolieri and Hang, 2019). Until the recent years, there were no ABL LES models capable of modelling detailed surface structures, such as buildings or steep complex terrain shapes in ABL. Nowadays, it is possible to carry out LES for complex built areas (e.g., Letzel et al., 2008), but this is still limited to relatively small areas because of the high spatial resolution requirement. Concerning urban LES, Xie and Castro (2006) have
5 shown that at least from 15 to 20 grid nodes are needed across street canyons to satisfactorily resolve the most important turbulent structures within the canyons. This requirement typically leads to grid spacings on the order of 1 m. However, the vertical extent of the LES domain should scale with the ABL height, and the horizontal size should span over several ABL heights in order to capture the ABL-scale turbulent structures (de Roode et al., 2004; Fishpool et al., 2009; Chung and McKeon, 2010; Auvinen et al., 2020a). To adequately capture processes on the street-scale and to simultaneously capture large ABL-scale tur-
10 bulence, sufficiently large model domains at small grid sizes are required, posing high demands on the computational resources in terms of CPU time and memory. Moreover, the uncertainty related to the lateral boundary conditions usually decreases as the domain is made larger.

Many conventional continuum-based numerical solution methods (e.g. finite-element and finite-volume methods) allow variable resolution so that the resolution can be concentrated to the area of principal interest and relaxed elsewhere. However, only
15 unstructured grid systems allow to take full advantage of spatially variable resolution. Many general-purpose computational fluid dynamics packages offer unstructured grid systems, but according to our experience such solvers are usually computationally decidedly less efficient than ABL-tailored LES models, such as PALM (Raasch and Schröter, 2001; Maronga et al., 2015, 2020), the Weather Research and Forecasting Model (WRF) (Skamarock et al., 2008) with its LES option and the Dutch Atmospheric Large-Eddy Simulation (DALES) (Heus et al., 2010) that are based on structured orthogonal grid system with
20 constant horizontal resolution. The model nesting approach can be exploited to further speed up ABL LES models or to allow larger domain sizes without compromising the resolution in the area of primary interest. The authors acknowledge that alternative solution approaches, such as Lattice Boltzmann Method (e.g., Aidun and Clausen, 2010; Ahmad et al., 2017), also offer strategies for computational efficiency improvements. The presented nesting approach is primarily relevant for the family of structured, finite-difference Navier-Stokes solvers.

25 The idea of grid nesting is to simultaneously run a series of two or more LES model domains with different spatial extents and grid resolutions. In this implementation the outermost and coarsest-resolution LES domain (termed *root* domain henceforth), which acts as a *parent* to its *child* domains, obtains its boundary conditions in a conventional manner, whereas the nested LES domain (child) always obtains its boundary condition from its respective parent domain through interpolation. In one-way coupled nesting only the children obtain information from their parents. In such a coupling strategy, the instantaneous
30 child and parent solutions can deviate within the volume of the nest. If a stronger binding between the solutions is desired, the child solution needs to be incorporated into the parent solution. This is achieved in two-way coupled nesting, where the parent solutions are influenced by their children through so-called anterpolation (Clark and Farley, 1984; Clark and Hall, 1991; Sullivan et al., 1996). The term anterpolation was coined by Sullivan et al. (1996) although the concept is older.

The child-to-parent anterpolation can be implemented using, for instance, the post insertion (PI) approach (Clark and Hall,
35 1991) where the parent solution is replaced by the child solution within the volume occupied by both domains. In practice, some

buffer zones where interpolation is omitted are necessary near the child boundaries to avoid growth of unphysical perturbations near the child boundaries (Moeng et al., 2007). An example of a two-way coupled nesting implemented in the WRF-LES model is given by Moeng et al. (2007), and later successfully applied to a stratocumulus study by Zhu et al. (2010). ~~However, this WRF-LES nesting system is limited to horizontal directions, i.e. all the domains have equal height which may lead to computational inefficiency.~~ The WRF-LES nesting system can also be used in one-way coupled mode (Mirocha et al., 2013), and this way it has been applied, e.g. to a complex-terrain study (e.g. Nunalee et al., 2014; Muñoz-Esparza et al., 2017) and to an offshore convective boundary layer study (Muñoz-Esparza et al., 2014). ~~More recently~~ Subsequently, Daniels et al. (2016) introduced a vertical interpolation into the WRF model, but this method is restricted to one-way coupled nesting. ~~Moreover, according to our knowledge, Recent implementation of an immersed boundary method has made WRF-LES is not applicable to blunt-obstacle-resolving LES required for urban turbulence studies, with one-way coupled nesting approach applicable to obstacle-resolving LES (Wiersema et al., 2020).~~ In addition to WRF-LES, the numerical weather prediction model ICON (Zängl et al., 2015) features an LES mode and includes an online nesting capability (Heinze et al., 2017). However, due to their terrain following coordinate system, neither WRF-LES nor ICON-LES can resolve sharp obstacle structures, hence these models cannot be applied to urban studies in obstacle-resolving fashion.

15 Recently, Huq et al. (2019) implemented a purely vertical nesting system into PALM in which the child and parent domains are required to have the same horizontal extent. Although this approach is useful, e.g. when the grid resolution near the surface needs to be refined to better capture the atmosphere-surface exchange, the requirement of equal horizontal domain extensions poses high demands on the computational resources, limiting this approach to only academic studies. This implementation is also limited to have a single child domain only. For these reasons, we decided to develop the present, more general and fully

20 three-dimensional nesting system in PALM. It can also be run in a pure vertical nesting mode.

One-way coupled obstacle-resolving LES has been applied to a built environment by Nakayama et al. (2016) and by Vonlanthen et al. (2016, 2017). Also the present PALM implementation has already been demonstrated by Maronga et al. (2019, 2020) and applied to obstacle-resolving urban studies (Kurppa et al., 2019; Auvinen et al., 2020a; Karttunen et al., 2020; Kurppa et al., 2020) using the one-way coupling. At current stage, we are not aware of any research on obstacle-resolving LES

25 in the ABL context employing two-way coupled nesting approach. Through our studies, we have observed that the application of two-way coupling in obstacle-resolving LES can become problematic. Therefore, in addition to documenting and evaluating the newly implemented nesting method in the PALM model, this paper addresses the applicability of the two-way coupled nesting approach in obstacle-resolving LES.

The paper is organized as follows: Sec. 2 gives a brief description of the LES mode of the PALM model system 6.0. Sec. 3 presents the technical, algorithmic and numerical aspects of the implemented nesting. In Sec. 4 the implemented nesting is evaluated for a series of test cases featuring different kinds of boundary-layer flow. Finally, Sec. 5 summarizes the results and gives and outlook of future developments.

2 The PALM model system 6.0 (LES mode)

The PALM model system (Raasch and Schröter, 2001; Maronga et al., 2015, 2020) is based on the non-hydrostatic, filtered, Navier-Stokes-equations in the Boussinesq-approximated or anelastic form. It solves the prognostic equations for the conservation of momentum, mass, energy, and moisture on a staggered Cartesian Arakawa-C grid. Subgrid-scale turbulence is parameterized using a 1.5-order closure after Deardorff (1980) in the formulation of Saiki et al. (2000). In its standard configuration PALM has thus seven prognostic quantities: the velocity components u_i (where $u_1 = u, u_2 = v, u_3 = w$), the potential temperature θ , specific humidity q_v , a passive scalar s , and the subgrid-scale (SGS) turbulent kinetic energy e . By default, discretization in time and space is achieved using a ~~3rd-order~~ third-order Runge-Kutta scheme after Williamson (1980) and a ~~5th-order~~ fifth-order advection scheme after Wicker and Skamarock (2002). The horizontal grid spacing is always equidistant, whereas it is possible to use variable grid spacing in the vertical direction. Often, the vertical grid spacing is set equidistant within the boundary layer, and stretching is applied above the boundary layer to save computational time in the non-turbulent free atmosphere. At the lateral boundaries cyclic conditions or more advanced in- and outflow conditions can be employed.

Both the Boussinesq and the anelastic approximation require incompressibility of the flow. To provide this feature a predictor-corrector method is used where an equation is solved for the modified perturbation pressure after every Runge-Kutta sub-time step (e.g. Patrinos and Kistler, 1977). The method involves the calculation of a preliminary prognostic velocity. Divergences in the flow field are then attributed solely to the pressure term, leading to a Poisson equation for the perturbation pressure. In case of cyclic lateral boundaries, the Poisson equation is solved by using a direct fast Fourier transform (FFT) method. However, in case of non-cyclic boundary conditions, an iterative multigrid scheme is used (e.g. Hackbusch, 1985).

Parallelization of PALM is achieved by using the Message Passing Interface (MPI, e.g. Gropp et al., 1999) and a two-dimensional (horizontal) domain decomposition.

PALM offers several embedded models to simulate physical processes within the urban environment. Namely, without intention to provide an exhaustive list, this embraces a land-surface (Gehrke et al., 2020) and a building surface model (Resler et al., 2017) to consider the surface-atmosphere exchange of heat and moisture; a radiative-transfer model (Krč et al., 2020) to include complex three-dimensional mutual radiative interactions between surfaces and plants; an indoor and building-energy demand model; an aerosol (Kurppa et al., 2019) and an air-chemistry model (Khan et al., 2020); as well as an embedded Lagrangian particle model for dispersion. For a complete overview we refer to Maronga et al. (2020).

3 Nesting system

3.1 General concept

The nesting system we have developed is based on the concept of parent and child domains. Each parent domain can enfold multiple child domains but a child domain can, naturally, only have one parent domain. The top-level domain, also called root domain, acts as a parent domain to child domains at the first nesting level. The child domains at first nesting level might have subsequent child domains for which they then act as parent domains (cascading arrangement), see Fig. 1. Our nesting

system allows for up to 63 nested domains plus the root domain. The implementation requires that all child domains are always completely located inside their respective parent domain. Also, the grid spacings of a child domain naturally have to be smaller than the grid spacings of its parent domain. The grid-spacing ratios $\Delta X_i / \Delta x_i$ must always be integer valued although different ratios may be used in different directions. Therefore in nested runs the grid stretching is only allowed in the root domain and

- 5 only above the top boundary of the highest nested domain. There may be multiple child domains at the same nesting levels, but overlapping child domains at the same nesting level are not permitted. Finally, all the nest domains have to be surface-bound so that elevated child domains are not allowed. Time synchronization is taken care by simply selecting the minimum of the time steps determined by each model independently and broadcasting this time-step value for all models. Each model inputs and outputs in the same way.

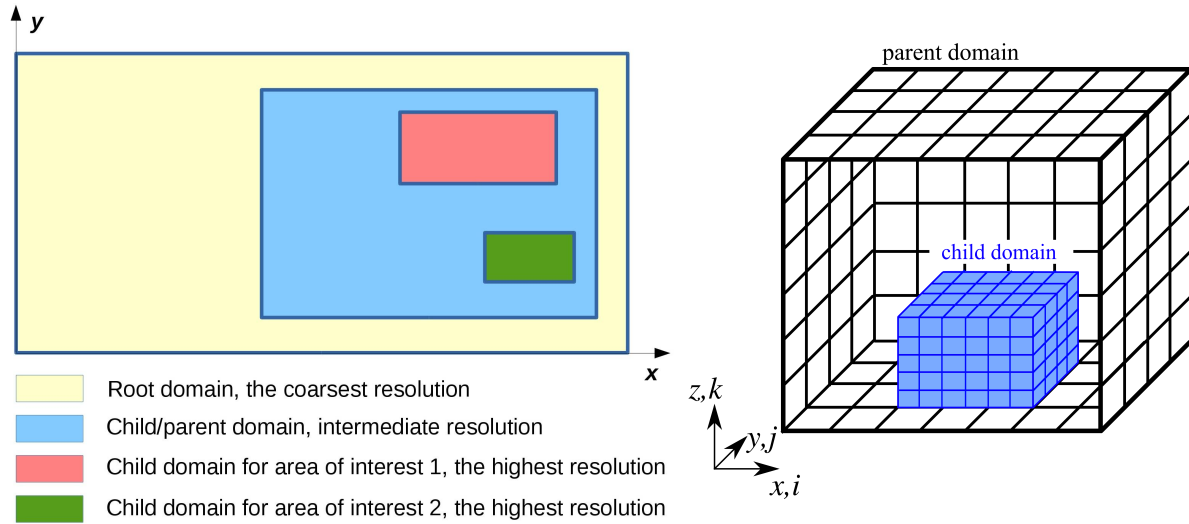


Figure 1. A schematic example of a nested configuration involving both cascading and parallel child domains is shown on x, y -plane on the left-hand side. On the right-hand side, a three-dimensional view of a nested child domain inside its parent domain is shown.

- 10 In general, the system is designed as two-way coupled nesting, in which a child domain can affect its parent domain and vice versa. It is possible, however, to run the system in a one-way coupled mode where no feedback from the child domain is incorporated in its parent domain. Moreover, it is possible to use the system as a pure vertical one-dimensional nesting, where the lowest part of the model (e.g. the atmospheric surface layer where the dominant turbulent eddies are usually very small) can be run as a child domain with finer grid spacing than its parent domain that compasses the entire boundary layer. In the
- 15 case of pure vertical nesting, cyclic boundary conditions must be set on all the lateral boundaries. Unlike the method proposed by Huq et al. (2019), the present method allows a cascade of more than one child domain also in the pure vertical nesting cases.

The present nesting approach is a variant of the PI method, in which the communication between each parent-child couple is realized via interpolations (from parent to child) and antinterpolations (from child to parent) after each Runge-Kutta sub-step

and just before the pressure solver. The latter then ensures that mass conservation is enforced in the antepolated solution in the parent domain.

3.2 Restrictions

The current implementation poses a few restrictions for the nested setups. Moreover, the interpolation and antepolation methods, which are discussed in the following sections, are based on certain assumptions, e.g. on the grid-line matching between parent and child domains leading to a few more restrictions. Altogether these restrictions are:

- the child domain must always be completely inside its parent domain and there must be a margin of four parent-grid cells between the boundaries of child and parent domains
- parallel child domains must not overlap each other and there must be a margin of at least four child-grid cells between two parallel child domains
- the domain decomposition of all child domains must be such that the sub-domain size is never smaller than the parent grid-spacing in the respective direction
- buildings or any other topography must not reach the child domain top
- all the grid-spacing ratios must be integer valued
- the outer boundaries of child domains must match with grid planes in its parent domain
- no grid stretching is allowed in the child domains and in root domain it is allowed only above the top boundary of the highest child domain

3.3 Structure of the nesting algorithm

Ideally, the coupling actions, i.e data transfers between the domains, antepolation and interpolation, would be performed after the pressure-correction step using the final divergence free velocity field on both parent and child. To achieve this in the context of the pressure-correction method employed in PALM requires a staged arrangement of the coupling actions such that a child first sends data to its parent and after receiving the data the parent antepolates and performs the pressure correction step. After the pressure-correction step the parent sends data to the child which interpolates new boundary conditions from the received data and performs the pressure-correction step. The purely vertical nesting method implemented in PALM by Huq et al. (2019) features this kind of staged structure. However, Huq et al.’s method may lead to excessive waiting times as the child has to wait until the parent performs the pressure-correction step and vice versa. Moreover, the staged coupling approach becomes more complicated and more inefficient when a cascade of several nested domains is used. Therefore, Huq’s implementation allows for only one child domain. The possibility to employ cascades of child domains was an initial requirement for the present system design and therefore the staged coupling arrangement had to be abandoned. In principle, it would be possible to perform the pressure-correction step twice, first time before the coupling actions for all domains and second time for all

parent domains after the coupling to make the antepolated fields divergence free. However, this would be computationally very expensive severely compromising the benefits from the nesting. This is because the pressure solution is typically the most time consuming part of the solution process. To avoid this extra penalty, the coupling is based on the preliminary prognostic velocity fields \mathbf{u}_{pre} in the present implementation. The sequence of the coupling actions is illustrated in Fig. 2. This choice

5 has the consequence that the interpolated velocity boundary conditions for a child domain may violate the global mass balance over the child domain such that

$$\int_S \rho \tilde{\mathbf{u}}_{\text{pre}} \cdot \mathbf{n} dS \neq 0, \quad (1)$$

where the tilde symbol is the interpolation operator and \mathbf{n} is the unit inner surface normal vector of the child domain boundary S excluding the bottom boundary. This mass-conservation error, though typically small, is eliminated in an integral sense by

10 adding a constant velocity correction $\Delta \mathbf{u}_{\text{pre}}$

$$\Delta u_{l,\text{pre}} = -n_l \frac{\int_S \rho \tilde{\mathbf{u}}_{\text{pre}} \cdot \mathbf{n} dS}{\int_S \rho dS} \quad (2)$$

to the interpolated child boundary values to exactly eliminate the global mass-balance error in Eq. (1). Here $l \in \{1, 1, 2, 2, 3\}$ corresponding to the left (1), right (1), south (2), north (2), and top (3) boundaries [following the boundary nomenclature used in PALM](#). In case of purely vertical nesting mode, the correction is applied only on the top boundary and S spans only over

15 it. This correction is made for all child domains right before the pressure-correction step. According to our tests, $\Delta \mathbf{u}_{\text{pre}}$ is typically three or four orders of magnitude smaller than the dominant velocity scales of the flow.

Huq et al. (2019) showed results for a zero mean-wind CBL case. In this case, especially if the nest-top boundary is set on a relatively low level, unphysical overestimation of horizontal velocity component variances easily develop if the coupling is based on \mathbf{u}_{pre} . Huq et al. (2019) showed that using the staged sequence of coupling actions, allowing the coupling based

20 on the final velocity field \mathbf{u} , mostly removes the overestimation of the horizontal-velocity variances. We have confirmed this by temporarily modifying the current implementation to adhere to Huq et al.'s staged arrangement and simulating a vertically nested zero mean wind CBL case similar to Huq et al.'s test case.

In the present method, the overestimation of the horizontal velocity component variances can be mostly avoided by using the integral mass-balance forcing (Eq. 2) and further by setting a narrow buffer zone below the top boundary in which the

25 antepolation is not performed. This is described in more detail in Sec. 3.5.

In addition to the velocity field, also all other prognostic variables are coupled except the SGS turbulent kinetic energy e , as it depends on the resolution by definition and therefore it is not straightforward to couple between parent and child domains having different resolutions. The antepolated values should be increased by some unknown resolution dependent factor and the interpolated values should be reduced accordingly. e strongly follows the velocity-gradient field and therefore it tends to

30 adapt to the antepolated velocity field on the parent side during the next Runge-Kutta step without being antepolated itself. Relying on this reasoning, we omit the antepolation of e . Moreover, we assume that the local generation of e often dominates its advection implying that replacing the interpolation of its child-boundary values by simple zero-gradient conditions may be

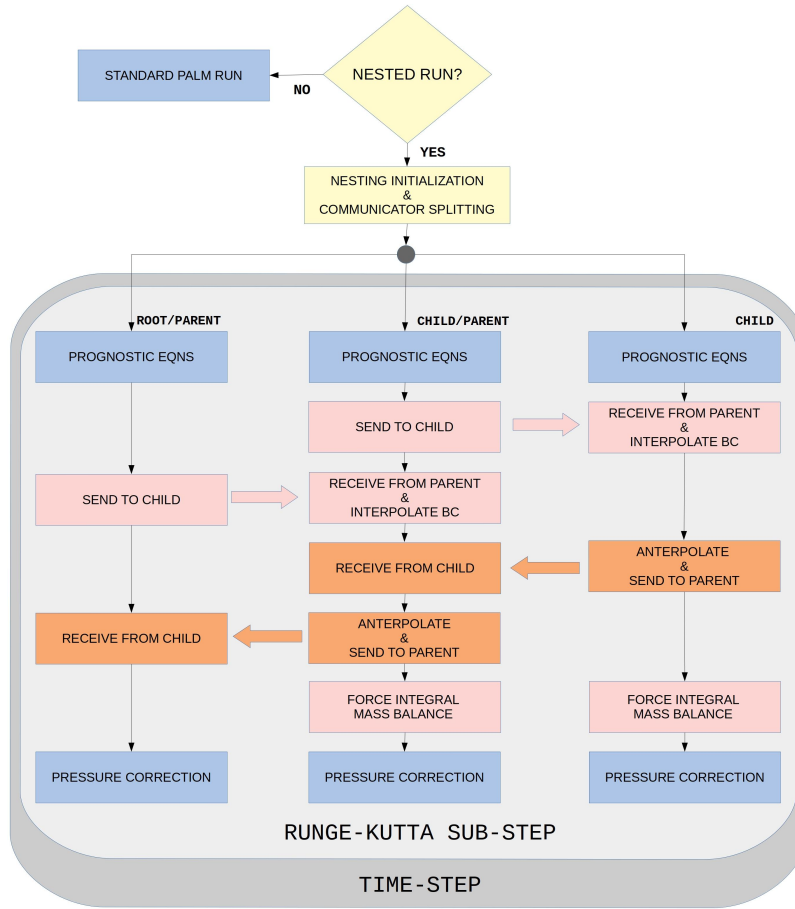


Figure 2. Flowchart illustrating the nesting actions in case of three domains in cascading order. In case of more than three levels of domains, more branches similar to the current middle branch would be added. Blue boxes represent baseline PALM actions while the other colors indicate nesting-specific actions. In one-way coupling only the actions indicated by pink color are invoked.

acceptable. In our numerical tests we compared the zero-gradient conditions with interpolated boundary values reduced by an estimated resolution-difference dependent factor. The comparisons revealed ~~no significant~~ only negligible differences in the results.

Further technical implementation issues are discussed in Appendix A.

5 3.4 Interpolation (parent to child)

Boundary conditions for the child domain lateral and top boundaries are given using data from its parent domain. This data is interpolated and set to the boundary grid points right behind the outer boundary of the respective child sub-domain. As mentioned in Sec. 3.3, all prognostic variables are interpolated except the SGS turbulent kinetic energy e . It is not interpolated since it depends on the grid resolution by definition and therefore the parent e interpolated to the child-domain boundaries

would be inconsistent with the child grid. Instead, a simple Neumann condition (zero-gradient condition) is used on child-domain boundaries.

It is very important that the interpolation method conserves the mass (volume) flow rate through the boundaries. If the mass conservation is violated in a two-way coupled run, a non-physical secondary circulation usually develops. Clark and Farley (1984) introduced a quadratic interpolation scheme that forms a reversible pair with their antinterpolation scheme which we also employ. This reversibility guarantees the mass conservation. However, as recently noted by Zhou et al. (2018), the conservation of fluxes other than the mass flux is violated if both advective velocity component and advected variable are interpolated using the Clark and Farley (1984) scheme or in fact any interpolation scheme of higher than zeroth order. They selected to use the Clark and Farley (1984) method only for the advective velocity component and the simple zeroth-order method, which in one dimension is

$$\phi_i = \Phi_I \quad \text{for } i \text{ within the parent-grid cell } I, \quad (3)$$

for all advected variables. ~~The~~ (here the child variables and indices are denoted by lowercase letters and those of parent by uppercase). This choice satisfies the flux conservation requirement for all variables if the grid-spacing ratio is odd valued. Conservation of fluxes through the nest boundaries is an important condition for a nesting algorithm. According to our tests the lack of flux conservation for example for the velocity component v on the left nest boundary may lead to a wrong mean-flow angle in the whole system of domains, and such an error can be remarkably large. Therefore, we design the interpolation method such a way that the flux conservation errors on the nested boundaries are minimized. In PALM, the interpolation algorithm has to cope with complex topography. Therefore we do not select the quadratic scheme of Clark and Farley (1984) at all. Instead, we use a zeroth-order interpolation for all variables. Our approach is to use Eq. (3) for the boundary normal velocity component u^N and all scalar variables, and another zeroth-order interpolation for the staggered velocity components u_i^S as follows (in one dimension)

$$u_i^S = \begin{cases} U_I^S & \text{for grid points } i \text{ co-located with a parent-grid point } I \text{ in the direction of the interpolation} \\ \frac{1}{2}(U_{I-1}^S + U_I^S) & \text{for all grid points } i \text{ between parent-grid points } I-1 \text{ and } I. \end{cases} \quad (4)$$

The reason behind this choice is explained and discussed in what follows.

In principle, the most straightforward way to satisfy the flux conservation would be to directly use the flux on the parent-grid cell face on the boundary and to distribute it onto the underlying child-grid cell faces in the proportion of the cell-face areas in the fashion of the finite volume method in which the fluxes are typically stored. However, PALM is formulated as a finite difference method and thence its architecture does not support this method. Therefore the interpolation procedure should be at least approximately flux conservative.

Zhou et al. (2018) require separately conservation of a prognostic variable ϕ and its resolved-scale turbulent flux $\langle u^{n'} \phi' \rangle$ through the boundary, where u^N is the boundary-normal velocity component, as

$$\langle \phi \rangle_b = \langle \Phi \rangle_b, \quad (5)$$

$$\langle u^{N'} \phi' \rangle_b = \langle U^{N'} \Phi' \rangle_b, \quad (6)$$

where ϕ and u^N are variables of the child grid and Φ and U^N are variables on the parent grid, and $\langle \cdot \rangle_b$ denote spatial averaging over the child domain boundary. The first condition Eq. (5) has originally been stated by Kurihara et al. (1979). Later Clark and Farley (1984) stated a stronger local form of this condition called reversibility condition. Reversibility of interpolation and anterpolation means that the following holds locally, i.e. not only in the sense of spatial averaging over the whole boundary

$$5 \quad \widehat{\widetilde{\phi}}(\Phi) = \Phi, \quad (7)$$

where the tilde is the interpolation operator and the hat is the anterpolation operator. Naturally, the reversibility guarantees that the conservation condition Eq. (5) is fulfilled. Obviously the zeroth-order interpolation satisfies the reversibility condition Eq. (7) in addition to Eq. (5). It is straightforward to show that if Eq. (5) holds, the flux conservation condition Eq. (6) can be also be written for the total flux as

$$10 \quad \langle u^N \phi \rangle_b = \langle U^N \Phi \rangle_b. \quad (8)$$

We shall study the flux conservation using this condition instead of Eq. (6).

Zhou et al. (2018) use the quadratic interpolation scheme by Clark and Farley (1984) for the boundary-normal velocity component, i.e. the advective component. For all other variables they selected to use the simple zeroth-order interpolation Eq. (3). This selection was made in order to satisfy the flux conservation condition Eq. (6) which is equivalent to the condition Eq. (8). This choice readily satisfies the flux conservation for all variables which are non-staggered relative to u^N and U^N . It also satisfies the flux conservation condition for the velocity components u^S and U^S staggered relative to u^N and U^N in the boundary plane, but only if the grid-spacing ratio is odd-valued. The difference of the odd- and even-valued grid-spacing ratio cases is illustrated in Fig. 3. The method by Zhou et al. (2018) is, indeed, strictly limited to odd-valued grid-spacing ratios. This is a strong restriction and therefore we decided to allow also even valued grid-spacing ratios. In such cases, the flux conservation condition is still readily satisfied for the non-staggered variables, but for staggered velocity components the situation becomes more complicated and it is not necessary possible to satisfy it exactly. By using Eq. (4) for the staggered velocity components the flux conservation condition can be satisfied approximately for both odd and even grid-spacing ratios. As an example and for the sake of clarity but without any loss of generality, we compute the spatially averaged fluxes of v - and V -components for a nested domain face and assume that the boundary-normal direction is x . The velocity components in the x -direction are u and U . ~~The~~ Within the advection scheme, the advective u velocity for the flux is interpolated linearly to the flux point of v as $(u_{j-1} + u_j)/2$. The chosen interpolation technique Eq. (4) leads to the following averaged resolved-scale advection flux of v through the child boundary (the SGS-fluxes are assumed small and omitted here)

$$\langle uv \rangle_b = \frac{1}{[R_y(N_y + 1) - 1]N_z} \left[R_y \sum_{J=J_s}^{J_n} \sum_{K=K_b}^{K_t} \frac{U_{J-1,K} + U_{J,K}}{2} V_{J,K} + (R_y - 1) \sum_{K=K_b}^{K_t} \left(\frac{U_{0,K}}{2} V_{0,K} + \frac{U_{N_y,K}}{2} V_{N_y+1,K} \right) \right], \quad (9)$$

while on the parent grid, the averaged advection flux of V is (as expanded by R_y for easier comparison)

$$30 \quad \langle UV \rangle_b = \frac{1}{R_y N_y N_z} \sum_{J=J_s}^{J_n} \sum_{K=K_b}^{K_t} R_y \frac{U_{J-1,K} + U_{J,K}}{2} V_{J,K}. \quad (10)$$




Figure 3. Staggered velocity-component nodes in cases of odd (3) a) and even (4) b) grid-spacing ratios. The staggered velocity-component nodes are shown as arrows, thick arrows are for the parent grid and thin ones for the child grid. The parent scalar grid-cell faces are drawn with solid lines and the corresponding child-grid cells with dotted lines. Locations of the corresponding parent-grid scalar nodes are shown as black dots. The blue color indicates the left-hand parent-grid node and red the right-hand node. The violet colored child grid nodes receive the averaged values according to Eq. (4).

Here, J and K are the parent-grid indices, and the child-domain boundary covers the parent-grid node range where $J_s \leq J \leq J_n$ and $K_b \leq K \leq K_t$ with $N_y = J_n - J_s + 1$ and $N_z = K_t - K_b + 1$. R_y is the integer-valued grid-spacing ratio in the y -direction. Clearly Eq. (9) and Eq. (10) do not exactly equal each other because of the edge terms depending on e.g. V_0 and V_{N_y+1} , and because the denominator of Eq. (9) slightly deviates from $R_y N_y N_z$. It should be noted that these two differences usually have opposite effects, and that $\langle uv \rangle - \langle UV \rangle$ tends towards zero as N_y becomes large. In typical applications, the order of magnitude of N_y is hundreds making the flux conservation error negligibly small. N_z is usually smaller than N_y , maybe even one order of magnitude smaller, making the flux conservation error of the vertical velocity component w possibly somewhat larger than that of the horizontal components. If $N_z = 32$ and the grid-spacing ratio $R_z = 3$ for example, $\langle uv \rangle - \langle UV \rangle$ can be expected to be of the order of 1%. Moreover, a small flux conservation error does not distort w as easily as the horizontal components, because w unlike the horizontal components is relatively strongly controlled by its surface boundary condition.

In the above considerations, it is assumed that the advected parent-grid variable values on the child boundary are equal to the values used for the parent-grid flux in the advection-term computation. This is not the case in reality since on the parent side the advected variable values used for calculating the advection fluxes are interpolated using the 5th-order-fifth-order interpolation scheme by Wicker and Skamarock (2002) while on the child boundary this is not the case. Here, it is important to understand, that the above mentioned interpolation onto the flux point in the advection scheme is a separate procedure from the interpolation from parent to child, and it is performed in a different phase of the time-step cycle.

The default advection scheme in PALM is the fifth-order Wicker and Skamarock (2002) scheme which employs a stencil of seven nodes, i.e. three boundary ghost-point values would be needed behind the nest boundaries. However, in PALM, all boundaries except cyclic boundaries are treated in a special way. Next to a boundary, the advection scheme is degraded such that only one layer of boundary ghost nodes is actually needed. This is achieved by using the third-order Wicker and Skamarock (2002) scheme for the second layer inside the domain and the first-order upwind scheme for the first layer of nodes next to the boundary. This has to be taken into account in designing of the interpolation procedure since the use of the first-order upwind scheme may increase the flux-conservation error and compromise the accuracy around the nest boundaries. However, we can reduce this additional flux-conservation error by using the following 'trick'. Instead of substituting the original parent-grid values in the interpolation formulae (3) and (4), we replace them by values interpolated onto the flux points using a scheme of higher than first order and substitute these values instead. As a result, the formally first-order upwind advection scheme becomes the selected higher order scheme if the local flow direction is into the domain. The main purpose of this ~~'trick'~~-substitution, which we refer to as Transfer to Boundary Plane (TBP), is to reduce the flux-conservation error. Ideally the fifth order scheme would be used for the ~~flux-point~~-TBP interpolation. However, in practice it is not possible to use any interpolation scheme using more than one grid point behind child-domain boundaries. The reason for this is that the child has no information about the parent domain topography outside the first parent-grid layer behind a child-domain boundary. An interpolation stencil reaching further away than this could penetrate a vertical wall leading to erroneous interpolation. Therefore, the only alternative to the ~~1st-order~~-~~first-order~~ upwind interpolation is to use the simple average of the parent-grid values both sides of the child-domain boundary. This leads to the ~~2nd-order~~-~~second-order~~ central advection scheme. Obviously it is different from the ~~5th-order~~-~~fifth-order~~ scheme, but we argue that the difference between the fluxes computed using the ~~5th-order and 2nd-order~~-~~fifth-order and second-order~~ schemes is usually smaller than the difference between the fluxes computed using the ~~5th-order and 1st-order~~-~~fifth-order and first-order~~ schemes. Our numerical tests support this argument. On the top boundary, there is no topography and hence we can use wider interpolation stencil there. We ended up using the ~~3rd-order~~-~~Wicker and Skamarock (2002) scheme~~-~~third-order~~ Wicker and Skamarock (2002) scheme there because in our numerical tests it yielded similar results as the more complicated ~~and~~ more communication-intensive ~~5th-order~~-~~fifth-order~~ scheme. It should also be noted, that this ~~trick~~-TBP-substitution improves the accuracy only on those boundary regions where the flux is into the child domain. The sequence of interpolation operations is illustrated in Fig. 4 using the left child boundary as an example. Phase 1 operations belong to the above explained TBP-substitution and phase 2 operations to the actual interpolation using formulae (3) and (4).

According to our experience, the conservation properties of an interpolation method are more important than its local accuracy. Increasing the interpolation accuracy on the child boundary planes does not yield any evident benefit as the solution requires a development distance (i.e. a border zone) as it adapts to the changed resolution within the child domain. Therefore the zeroth-order method has turned out to be fully sufficient and remains the only interpolation method implemented in PALM. However, we have considered also an alternative interpolation approach for the advected variable based on tri-linear interpolation with a specific reversibility correction. Although it is not implemented a short discussion is provided in Appendix B.




Figure 4. A schematic illustration of the interpolation operations on the left child boundary as an example. Phase 1 operations belong to the TBP-substitution and phase 2 operations to the actual interpolation using formulae (3) and (4). Filled blue symbols denote intermediate values on the boundary plane (flux plane) resulting from the TBP. Filled red symbols denote the final boundary-condition values for the child boundary.

3.5 Anterpolation (child to parent)

Anterpolation is used to feed the child domain solution back to its parent domain. Generally, anterpolation consists of filtering the fine-grid child solution $\phi_{i,j,k}$ and mapping it to the parent domain grid. We select to employ the anterpolation scheme proposed by Clark and Farley (1984), which consists of simple averaging over one parent-domain grid volume around the parent grid node of the variable in question corresponding to top-hat filtering, viz.

$$\hat{\phi}_{I,J,K} = \frac{1}{N_{I,J,K}} \sum_{i_1(I)} \sum_{j_1(J)} \sum_{k_1(K)} \phi_{i,j,k} . \quad (11)$$

The original parent solution $\Phi_{I,J,K}$ is replaced by the ant interpolated solution in the domain of overlap. Here, i, j, k and I, J, K are the child- and parent-grid indices, respectively, and the hat is the anterpolation operator. The summation index limits, i.e.

the span of the ant interpolation cell $i_1(I)$, $i_2(I)$, $j_1(J)$, $j_2(J)$, $k_1(K)$ and $k_2(K)$ are pre-computed during the initialization and they depend on the grid configuration and the variable in question, i.e. the staggered velocity components have different index limits than the grid-cell centered scalars. Note that for the velocity components, the ant interpolation volume is reduced to the grid cell face on which the velocity component is defined. This means that the upper index limit in the direction of the velocity component is reduced to the lower one, for instance $i_2 = i_1$ for u , because the coordinates of the velocity component node in the respective direction in the parent and the child readily coincide, thus there is no need for ant interpolation in this direction. $N_{I,J,K}$ is the number of child domain values used for ant interpolation at a given parent-grid location, and is pre-computed during the initialization as

$$N_{I,J,K} = [i_2(I) - i_1(I) + 1][j_2(J) - j_1(J) + 1][k_2(K) - k_1(K) + 1] . \quad (12)$$

- Note that due to the staggered grid, four sets of the index limits and $N_{I,J,K}$ are pre-computed and stored: one for each velocity component and one for all scalars. Generally, the ant interpolation cells can be spanned in more than one way. We define the ant interpolation cells similarly to Clark and Farley (1984). For scalar variables (non-staggered variables) the ant interpolation cell spans $X_i \pm \Delta X_i/2$ where X_i ($i = 1, 2, 3$) are the coordinates of the scalar node in the parent grid. For the velocity components (staggered variables), for example for u , the ant interpolation cell spans X_1 , $X_2 \pm \Delta X_2/2$, $X_3 \pm \Delta X_3/2$, where X_i are the coordinates of the staggered u -node in the parent grid.

- Buffer zones where the ant interpolation is omitted are applied next to the child-domain boundaries except the bottom boundary. The main purpose of the buffer zones is to avoid an unstable feedback loop between the ant interpolation and interpolation. The default width of these buffer zones is two prognostic grid nodes. The user may choose a different value for the buffer width, but the minimum allowed width is one parent-grid spacing. This is because the layer of nodes nearest to the child boundary is directly used in the interpolation, and using an ant interpolated value for interpolation leads to a strongly unstable behaviour. The buffer zones are comparable to the relaxation zones applied in the nesting system of the WRF-LES model (Moeng et al., 2007). In the WRF-LES nesting system the ant interpolation is under-relaxed within these zones such that the under-relaxation coefficient varies linearly across the relaxation zones which are five grid spacings wide. As mentioned in Sec. 3.3 the buffer zone below the top boundary also reduces the overestimation of the horizontal velocity variances observed in zero mean-wind CBL tests in purely vertical nesting mode. According to these tests in purely vertical nesting mode, simulation results are not particularly sensitive to the extent of the vertical downward shift of the upper edge of the ant interpolation domain.

Canopy-restricted ant interpolation

- The ant interpolation algorithm is implemented in the PALM model with a feature that enables its application in a spatially selective manner such that the operation is only performed within the computational domain that is above a user-defined vertical threshold. This practice is discovered to resolve complications that arise when two-way coupled nesting is applied in obstacle-resolved LES simulations where the ant interpolated solution within the obstacle canopy introduces discrepancies in the coarser parent solution. Thus we label this approach *canopy-restricted* (CR) ant interpolation and the coupling is referred to as

two-way CR, for short. The necessity of this anterpolation strategy is motivated and its effectiveness demonstrated in Sec. 4.2.3 where nesting is applied to obstacle-resolved LES test case.

4 Numerical experiments

In order to evaluate the nesting strategy, to show its benefits and point out its limits, we performed a series of nested model simulations for different grid-spacing ratios and respective non-nested reference simulations for different atmospheric situations. ~~For this purpose, we~~ The idea is not to mainly validate the PALM model against experimental data, but to systematically compare the nested-domain results to corresponding non-nested fine- and coarse-grid reference results, and to show that the nested-domain solutions are closer to the fine-grid reference solutions than the coarse-grid reference solutions are. PALM has been already evaluated for various ABL flows against measurement data (Letzel et al., 2008). Nevertheless, we show one comparison against wind-tunnel data in Sec. 4.2.2.

Furthermore, the idea is not to present grid convergence studies since the grid convergence of the PALM model has been demonstrated previously e.g. for convective boundary layer by Hellsten and Zilitinkevich (2013). Instead, our idea is to systematically compare the nested-domain results to corresponding non-nested fine- and coarse-grid reference results, and to show that the nested-domain solution is closer to the fine-grid reference solution than the coarse-grid reference solution is.

We simulated a homogeneously-heated flat-terrain convective boundary layer as well as a purely shear-driven flat-terrain boundary layer. Further, to investigate the performance of the grid nesting in more complex situations where non-flat topography is present, we performed two-staged nested simulations for a neutrally-stratified flow over a smooth three-dimensional hill and will compare the results against wind-tunnel data. Second, we simulated a neutrally-stratified urban boundary-layer flow over a regular staggered arrangement of building cubes using one- and two-staged nesting, and will compare the nested simulation results to corresponding non-nested fine- and coarse-grid simulation results. ~~Finally, we will demonstrate the applicability of a two-stage nesting for the same flow over the cube array.~~ Details concerning the different simulation setups are given in their respective sections. ~~Please note,~~ Note that for the sake of simplicity velocity components will hereafter be addressed by lower case variable names only, no matter if it refers to the flow in the parent or the child domain.

4.1 Convective boundary layer

The nesting method is first evaluated for a pure convective boundary layer (CBL) with zero mean wind. We set up one child domain that is centered within the parent domain. For the root domain, cyclic lateral boundary conditions were set. A homogeneous and time-constant surface sensible heat flux of 0.1 K m s^{-1} was prescribed. The simulation was initialized with a potential temperature profile that increases linearly with height at a lapse-rate of $0.3 \text{ K}/100 \text{ m}$. The root-model domain size is $10.2 \text{ km} \times 10.2 \text{ km} \times 3.0 \text{ km}$ in the x -, y - and z -directions, respectively, with an isotropic grid spacing of 20 m . The top of the child domain is set to be within the middle part of the CBL, and the domain size is $2.5 \text{ km} \times 2.5 \text{ km} \times 0.48 \text{ km}$ in the x -, y - and z -directions, respectively, with an isotropic grid spacing of 10 m , resulting in a grid-spacing ratio of 2. In order to examine how turbulence statistics behave for different grid-spacing ratios between parent and child in the CBL, we additionally run

nested simulations with grid-spacing ratio of 3 and 4 by increasing the isotropic grid spacing in the parent domain to 30 m and 40 m, respectively. Non-nested coarse- and fine-grid reference simulations were carried out corresponding to the nested simulations with different grid-spacing ratios. The simulated time was 4 four hours for all convective cases. Data analysis started after 2 hours of simulated time when model spin-up effects are not present any more and the simulations reached steady-state conditions. In order to perform a spectral analysis of time-series data, the time step was held constant at 1.0 s in all convective simulations during the data analysis period.

Figure 5a) shows an instantaneous horizontal cross-section of the w -component at a height of 40 m for the parent and child (overlaid) domains for the grid-spacing ratio 2. A hexagonal pattern of convective cells with strong updrafts and weaker downdrafts is visible, as it can be typically observed in LES. The transition between parent and child appears smooth and the flow structures are continuous in terms of shape and amplitude, while within the inner part of the child domain more fine-scale structures can be observed with slightly stronger up- and downdrafts, as also reported by Moeng et al. (2007). Furthermore, Fig. 5b), showing an instantaneous vertical cross-section for the w -component, also depicts how the up- and downdrafts are consistently maintained across the child boundary without any obvious impact on the turbulent structures.

Figure 6 shows horizontally- and time-averaged vertical profiles of potential temperature θ , vertical turbulent heat flux $\langle w'\theta' \rangle$, variances of horizontal and vertical velocity components, as well as the skewness of the vertical velocity component w , being one of the most grid sensitive quantities (Sullivan and Patton, 2011). The profiles of $\langle \theta \rangle$ indicate a well mixed CBL. With increasing grid-spacing ratio the corresponding parent and coarse grid simulations deviate from the fine grid reference, particularly near the surface and within the inversion layer, while the child results adhere well with the non-nested fine-reference simulation, indicating that the profiles of $\langle \theta \rangle$ in the child domains are rather independent of the parent grid for the employed grid spacings.

The heat flux profiles in the child and parent simulations decrease linearly with height within the CBL and are in good agreement with the fine reference simulation. For the parent simulation we note the near-surface kink in the heat flux (see Fig. 6c for a close-up view). Moeng et al. (2007) observed a similar kink in the heat flux and attributed it to inaccuracies in the statistical evaluation of the heat flux, more precisely, to errors that arise from interpolation from a mass- to a height-coordinate system. However, to evaluate fluxes PALM does not apply any interpolations but uses directly the resolved- and subgrid-scale fluxes as calculated in the advection scheme and the subgrid model, respectively, so that interpolation errors cannot explain the kink in this case. Instead, we attribute the kink in the parent domain to the anterpolation from the fine child solution. In simulations with different vertical grid spacing, the vertical gradients of $\langle \theta \rangle$ within the unstable near-surface layer are differently resolved, resulting in slightly different near-surface temperatures, as it can, e.g., be observed between the fine and coarse reference simulations in Fig. 6a. This indicates that the parent simulation will yield slightly different $\langle \theta \rangle$ -profiles than the child simulation. After the anterpolation is performed, the parent solution is replaced by the underlying child solution, where the near-surface vertical gradients of $\langle \theta \rangle$ in the parent domain partly deviate from the ones the model would create without feedback from the child domain, i.e. the near-surface $\langle \theta \rangle$ -profile in the parent is not in equilibrium with the applied surface boundary condition any more. In the following time step the parent model tries to re-adjust the post-inserted $\langle \theta \rangle$ to the vertical gradients as being present without feedback from the child, altering the heating rates and thus the near surface vertical




Figure 5. Instantaneous horizontal a) and vertical b) cross section of w after 4 h of simulated time for the grid-spacing ratio case 2. The horizontal cross-section is given at a height of 40 m. The black box indicates the lateral and top boundaries of the child domain. The white line indicates the y -position of the vertical cross section of w shown in (b). The vertical axis in (b) is normalized with the horizontal mean boundary-layer depth z_i . Note that only part of the parent domain is shown for the sake of visibility.

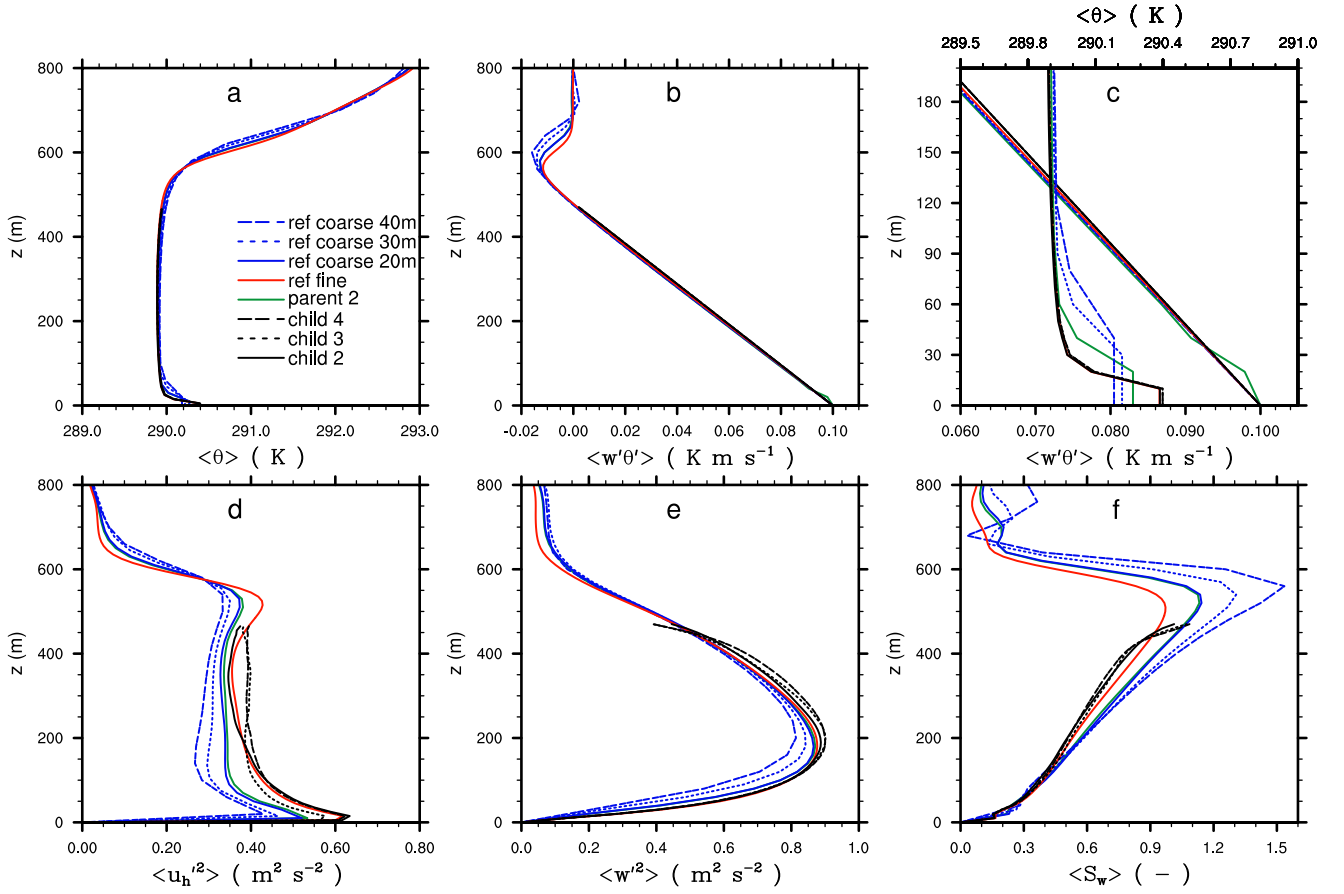


Figure 6. 30-min time and horizontally averaged profiles of a) $\langle \theta \rangle$, b) $\langle w' \theta' \rangle$, c) a close-up view of $\langle \theta \rangle$ and $\langle w' \theta' \rangle$, d) variance of the horizontal velocity components, e) variance of the vertical velocity component, and f) skewness of the vertical velocity component after 4 four hours of simulated time. Please note the second upper abscissa in c). Profiles are shown for the grid-spacing ratio of 2, 3, and 4 for the respective child domains, indicated by the respective numbers. All these child domains have the same resolution as the fine-grid reference simulation. The corresponding profiles from the coarse-grid reference simulations for the 20 m, 30 m, and 40 m grid spacing are indicated the same. For the sake of clarity, the resulting profiles for the parent domain are only shown for the grid-spacing ratio of 2. Squared Angle brackets indicate a horizontal domain-average over the domain.

gradients of the heat flux, which in turn becomes visible as near-surface kink. In fact, we verified this hypothesis in a test case by using identical vertical grid spacing in parent and child. In this case, no kink in the vertical heat flux was visible any more (not shown).

The variances of the horizontal and vertical velocity components, as well the skewness of the vertical component, depend strongly on the grid spacing as the coarse- and fine-grid reference simulations show, where the variances (skewness) become smaller (larger) for increasing grid spacing. The parent simulation agrees well with the coarse-resolution simulation, indicating

that the interpolation changes the parent flow field only marginally. The variances and skewness in the child simulations agree with the fine-reference profiles, except for the upper regions of the child domain where the variances are slightly overestimated. The child profiles are almost independent of grid-spacing ratio, and are close to the reference simulation profile. This indicates that the child solutions are almost independent on the chosen grid-spacing ratio in the studied cases.

5 Although there can be no mean horizontal advection in the zero-mean wind CBLs, spatially and temporally local horizontal advection always takes place and therefore flow structures are advected locally from parent to child (and vice versa). Therefore, advected flow structures may need a certain fetch to adjust to the changed grid spacing. In order to get an idea of how ~~much~~
~~distance apart~~ long distance from the lateral child boundaries is required to observe similar turbulence properties as in a
non-nested fine resolution reference simulation, we performed a spectral analysis. Therefore, we sampled time series of TKE
10 and θ at different locations ~~within~~ inside the child domain and calculated frequency spectra from the sampled time series. Subsequently, we averaged spectra over all sampling locations with the same distance from the lateral child boundaries. Overall,
we calculated spectra inside the child domain at locations 75, 100, 200, 300, and 500 m away from the lateral boundary. It
should be noted that transforming frequency spectra into wave number spectra using Taylor’s hypothesis in order to directly
link spectral information and grid spacing is not strictly correct in this case where we have no background wind; nevertheless
15 we will assume that frequency and wave number space are connected, i.e. large frequencies belong to small spatial scales and
vice versa. Fig. 7 shows the resulting frequency spectra, as well as corresponding spectra from fine- and coarse-grid reference
simulation. As expected, the coarse-resolution spectra exhibit less spectral energy at larger frequencies, compared to the child-
and fine-resolution spectra. This is due to the larger filter length assumed for the subgrid model, removing more energy at larger
spatial scales and thus also affect smaller frequencies. The child-spectra agree well with the fine-reference spectra, especially
20 for the grid-spacing ratio case 2 where even locations close to the lateral boundaries show good agreement with the reference.
We attribute this to the nature of the CBL, where turbulence is mostly produced locally by buoyancy and horizontal advection
is almost negligible, so that turbulence is almost not affected by any transport from the boundaries. Also for the grid-spacing
ratios of 3 and 4 the differences are generally small, but locations close to the lateral boundaries show slightly smaller spectral
densities at larger frequencies, indicating that for larger grid-spacing ratios the flow needs a fetch of a few tens of meters in a
25 purely buoyancy-driven boundary layer to adjust to the finer grid spacing.

Even though the child simulations yield turbulence profiles, spectra and instantaneous flow patterns similar to the fine-grid
reference simulation for a pure-buoyancy driven flow, the nested simulation nevertheless creates side effects on the flow which
appear as a secondary circulation (SC). This SC is not caused by a violation of mass conservation that has been discussed
in Section 3.4. Figure 8 shows the 5-hour time-averaged w -component at the middle part of the CBL in a homogeneously-
30 heated nested simulation. In order to compute the 5-hour time average, we continued the simulation with grid-spacing ratio
2 for further 3 hours. Within the region of the nested child domain a mean updraft can be observed, which is in the range of
 $0.4 - 0.9 \text{ m s}^{-1}$ and extends throughout the entire depth of the CBL (not shown). At the child domain boundaries and outside
the child-domain region the flow subsides in average, and horizontally directed branches at the upper and lower parts of the
CBL occur, giving the overall picture of a SC. The strength of this SC, indicated by the amplitude of the mean updraft, is in the

order of the strength of SCs observed in previous simulations over idealized stripe-like surface heterogeneities (Sühling et al., 2014) and even exceeds the strength of SCs observed in simulations over realistic surface forms (Maronga and Raasch, 2013).

SCs develop above surface heterogeneities mainly due to differential surface heating of the air, resulting in mean updrafts and downdrafts over the stronger- and less-heated patches, respectively. However, since we prescribe the same surface sensible
5 heat flux in the parent as well as in the child simulation, differential surface heating cannot be the reason of the SC in the nested simulation. Moeng et al. (2007) observed a temperature bias in their child domain that led to mean vertical motion to compensate the temperature bias. They observed temperature biases that go either way, i.e. a too cold or a too warm child domain, which they attributed to a nested child domain of too small horizontal extent. If only a few up- or downdrafts are resolved in the child domain, the vertical transport is dominated by these up- or downdrafts and thus a warmer or cooler
10 CBL can be quickly produced in the child domain, respectively. They showed that for larger horizontal child domain size the temperature bias and thus the associated vertical motion vanished. However, they only considered instantaneous differences between parent and child, meaning that the temperature bias is a result of insufficient sampling of the large up- and downdrafts rather than an inherent feature of the nesting which can only be observed after time-averaging. In our case the SC becomes visible only after considerable time-averaging. The updraft branch of the secondary circulations is always located within the
15 child domain also for larger child domain extensions. We hypothesize that this SC is triggered by a slightly different divergence of the vertical heat flux between the region occupied by the child domain and the remaining parent domain due to different grid spacing. It might be impossible to eliminate, because higher resolution better represents the turbulent mixing, so differences between the parent and the child solutions are to be expected in general.

Even though this inherent artificially-induced SC only appears when the flow is averaged over a longer time under quasi-
20 stationary conditions (no daily cycle, no change in the mean wind, etc.), nested simulation results should be interpreted carefully in terms of SCs. In particular, since the strength of the artificial SC is in the order of 'real-world' circulations over heterogeneous terrain, these ~~may superimpose each other~~two may become superimposed, altering the pattern of the vertical transport of sensible and latent heat. Although we did not succeed to proof our hypothesis, we encourage other researchers to look for the existence of such SCs in any nested models by analyzing the time averaged results.

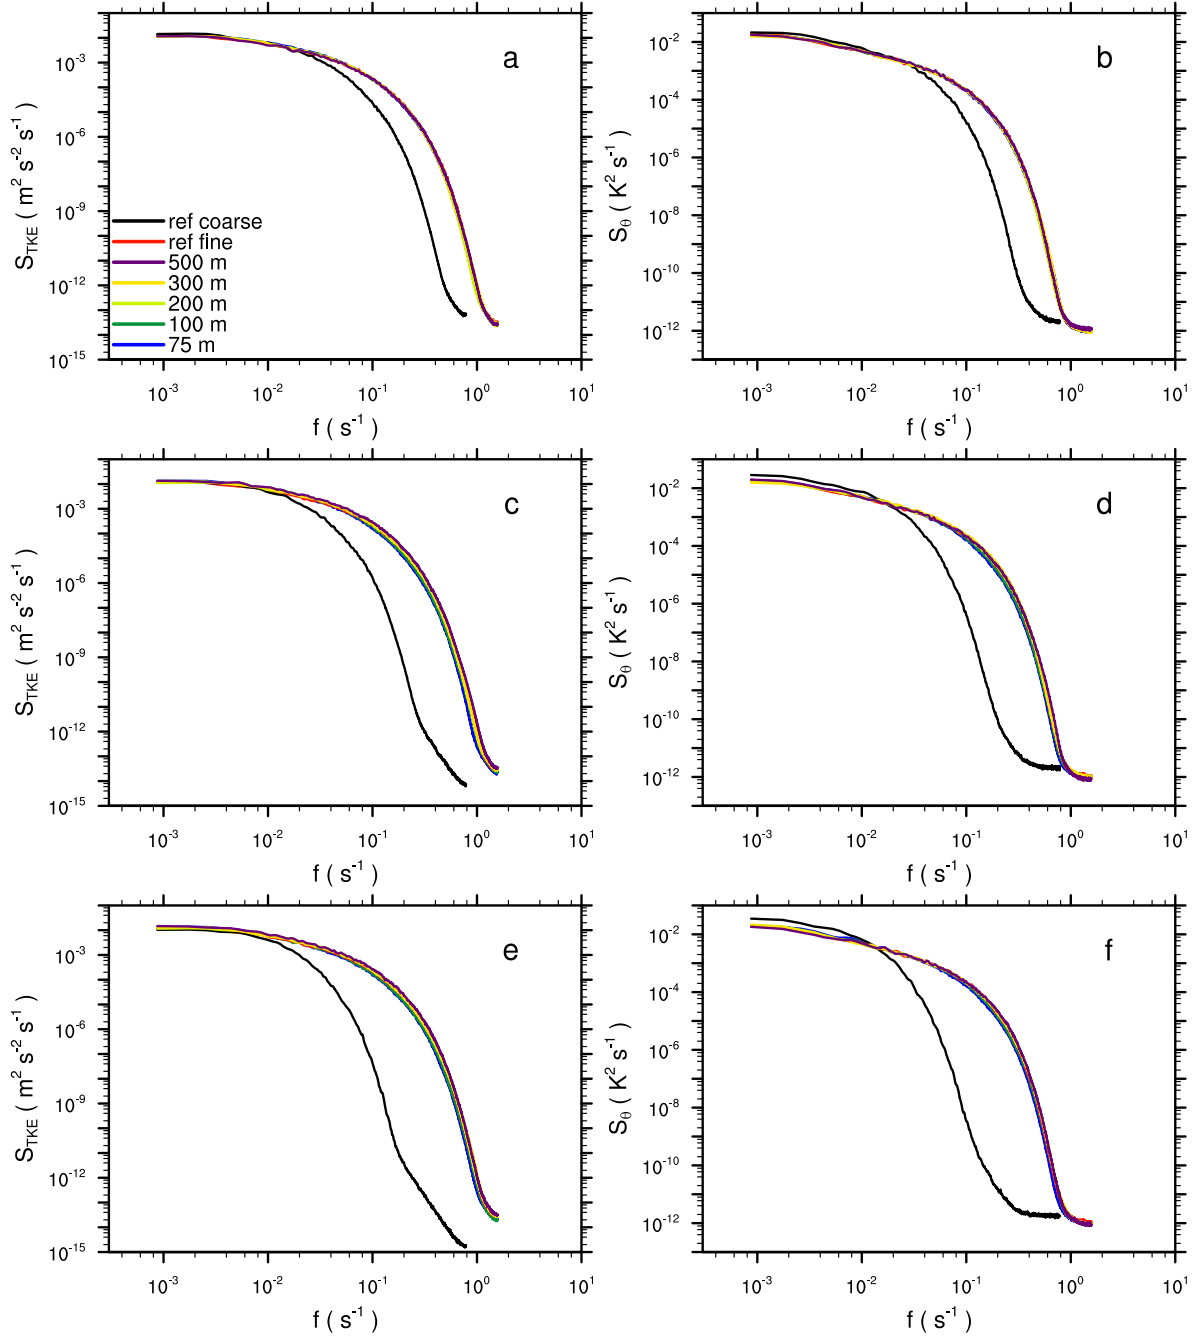


Figure 7. Frequency spectra of the TKE (left column) and θ variances (right column) ~~at different distances from inside~~ the ~~lateral-child~~ ~~boundaries-domain~~ for the grid-spacing ratio of 2 (a,b), 3 (c,d), and 4 (e,f). The coloured curves represent spectra taken at different distances away from the lateral child boundaries. Furthermore, spectra for fine- and corresponding coarse-grid references simulations are displayed. TKE and θ were sampled at $z = 120$ m.




Figure 8. Horizontal cross-section of 5-h time-averaged vertical velocity at $z = 400$ m in a nested simulation with grid ratio of 2. The black box indicates the location of the child domain.

4.2 Neutrally stratified boundary layer tests

Initialization and inflow conditions

As further test cases, we set up boundary layer flow simulations with increasing order of complexity. First, to evaluate the performance of grid nesting in shear-driven boundary-layer flows, we simulated a flow over a homogeneous flat surface in order to compare first and ~~second-order~~ second-order moments from a nested simulation against reference simulations. In a second step, we simulated a flow over a smooth three-dimensional hill for comparison of nested simulation results against wind-tunnel data. Finally, in order to illustrate the advantages of the grid nesting in more complex setups, we simulated a flow over a staggered arrangement of cubes mounted on a flat surface.

The parent domain size for all neutrally-stratified simulations was $L_x \times L_y \times L_z = 5.1 \times 1.5 \times 0.32 \text{ km}^3$ in the x -, y -, and z -directions, respectively. In all neutral simulations we prescribed a homogeneous roughness length of $z_0 = 0.01 \text{ m}$. At the top boundary we applied a free-slip condition for the horizontal wind components and zero vertical motion. At the spanwise lateral boundaries (north and south boundary) we applied cyclic conditions. At the western lateral boundary (hereafter referred to as inflow boundary) we prescribed mean inflow profiles for the u and v component, obtained from a cyclic precursor run. Two different precursor simulations were employed for the subsequent test cases. The one used for the flat surface and for the smooth hill featured a geostrophic wind of $u_g = 4.8 \text{ m s}^{-1}$ and $v_g = -1.3 \text{ m s}^{-1}$ at a latitude of 55 degrees, adjusted such that the surface-layer mean flow became parallel with the x -axis. This precursor simulation ran for 36 hours to reach a stationary state. The second precursor simulation, used for the cuboid case, was driven by a fixed pressure gradient angled to result in a mean flow of $u = 10 \text{ m s}^{-1}$ at $z = L_z$ with a 3 degree angle from the x -axis.

In order to obtain a turbulent inflow, we applied a turbulence recycling method according to Kataoka and Mizuno (2002), where the inflow mean vertical profiles of u and v are superimposed by turbulent fluctuations sampled at a recycling plane, which is placed at $x_{rc} = 1.5 \text{ km}$ downstream the inflow boundary. The recycling plane is placed sufficiently far apart from the inflow boundary to allow for statistically-independent turbulence, but also sufficiently far apart from the location of the child domain to avoid any feedback between the grid nesting and the inflow conditions. For further details on the implementation of the turbulence recycling method see Maronga et al. (2015).

Further, in order to avoid persistent streaks in the u -component, which may develop in neutrally-stratified flows and will be recurrently recycled in case of vanishing v -component, we shifted the recycled turbulent signals along the y -direction at the inflow boundary, following Munters et al. (2016). At the eastern outflow boundary we set a radiation boundary condition (Miller and Thorpe, 1981). The root domain was initialized with three-dimensional data recursively copied from the precursor run, while the child domain was initialized with data obtained from the parent. We used an isotropic grid spacing of 4 m and 2 m within the root and the nested child domain, respectively. The cuboid ~~case also encompasses~~ and the smooth-hill case also encompass a third domain with 1 m resolution (two-stage nesting).

In order to evaluate the effect of the nesting, we performed additional non-nested reference simulations with 4 m and 2 m grid spacing. However, due to its high computational demands, a 1 m non-nested reference simulation was not performed. The simulated time of the neutrally stratified simulations ranged from 4 to 7 hours. Data analysis started after 2 hours of

simulated time. When spectral analysis was performed ([homogeneous flat case](#)), the time step was held constant at 1.0 s for that simulation.

4.2.1 Neutrally stratified boundary-layer flow over flat terrain

Figure 9 shows an instantaneous horizontal cross-section of the u -component for the nested simulation. As typical for a neutrally stratified boundary layer, elongated streak-like structures can be observed (Hutchins and Marusic, 2007; Hutchins et al., 2012). These elongated structures preserve their size and amplitude when entering the child domain from the left and exiting to the right.




Figure 9. Instantaneous horizontal cross-section of the u -component in m s^{-1} at $z = 40$ m. The black box indicates the location of the child domain. The white solid lines indicate the x -locations where the profiles shown in Fig. 11 are averaged over the y -direction.

Figure 10 shows horizontal profiles of the time- and y -averaged friction velocity u_* within the child domain and the corresponding coarse- and fine-grid reference cases. In the coarse and fine reference cases u_* is constant along the x -axis indicating that the flow is in equilibrium with the surface friction. In the coarse-grid simulations u_* shows slightly higher values compared to the fine-grid reference simulation, even though the prescribed surface roughness is identical in all simulations. This suggests that the flow in the coarse-grid simulations sees a slightly rougher surface, which we attribute to the less accurate representation of the vertical near-surface gradients of the wind profile compared to the fine-grid simulation. When the flow enters the child domain, the coarse-grid inflow wind profile is not in equilibrium with the surface friction any more and the near-surface flow decelerates, indicated by the higher values of u_* near the child inflow boundaries. With increasing distance to the inflow boundary, u_* rapidly decreases and reaches a minimum with lower values compared to the reference cases, until it increases again reaching a secondary maximum and then asymptotically approaches a constant value, which is similar to the value of the fine-grid reference case in the grid-ratio case 2 and 3. However, at least for the given model domain size, u_* does not approach the fine-grid solution in grid-ratio case 4 but still exhibits higher values. This kind of spatial oscillation of u_* , which indicates

an alternating deceleration and acceleration of the near-surface flow along the x -direction, shows that the surface-momentum exchange in the child domain needs a sufficiently large development length. For grid-ratio case 2 the required fetch length is of at least 1 km to adjust to the fine-grid resolution. With increasing grid ratio the amplitude of the spatial oscillation increases and the fetch length becomes longer, and, as in grid-ratio case 4 even exceeds the model domain size.

5 Figure 10 shows that u_* gradually adjusts to the fine-reference value, at least for grid ratio case 2 and 3. This is in contrast to Moeng et al. (2007), who revealed a friction velocity bias between parent and child in their neutrally-stratified simulation when employing grid-size dependent SGS model as it is also used in this study.

Figure 11 shows time- and y -averaged profiles of the horizontal wind speed within the child domain for different grid ratios, taken at different distances downstream of the inflow boundary, indicated by the white solid lines in Fig. 9. At a distance of
10 100 m the profiles in the child domains agree well with the fine-reference profile within the lowest 10 m. Even though the surface-momentum exchange is still not in equilibrium at that position (see Fig. 10), one could already conclude from the near-surface wind profiles that the flow has already been adapted to the finer grid resolution. However, further above, the wind profiles of the child model still deviate from the fine-reference solution and are closer to the coarse reference profiles. This is especially obvious for the grid-spacing ratio case 4, where the wind profile shows a discontinuity at a height of $z = 20$ m. With
15 increasing distance from the child inflow boundary, the child-profiles gradually adjust to the fine-reference simulation, while at a distance of 2000 m the child profiles agree with the fine-reference solution, except for the grid-ratio case 4 which still deviates from the fine-reference solution.

In order to further analyze the flow adjustment within the child domain, we computed resolved-scale turbulent kinetic energy (TKE) spectra at different distances from the child inflow boundary. The spectra were calculated from time-series of the three
20 velocity components that were sampled at different locations within the domain. The final spectra were then obtain by averaging individual spectra over all locations with identical distance to the inflow boundary, assuming that the flow is parallel to the x -axis. Figure 12 shows TKE spectra obtained from the child domain and for the corresponding reference simulations. At low frequencies (large wave numbers), the spectra look quite similar and no obvious differences to the fine- and coarse-reference spectra can be observed, indicating the grid nesting does not induce any larger-scale oscillations which propagate through the
25 model domain. At higher frequencies, however, especially the near-inflow boundary child spectra differ to the fine-reference spectra and resemble more the corresponding coarse reference spectra. With increasing distance from the inflow boundary, the spectral properties gradually adjusts to those of the fine-reference case, while at a fetch of 500–1000 m almost no differences can be observed any more at that height level.

In contrast to a buoyancy-driven boundary layer, the flow in a purely shear-driven boundary layer requires a sufficiently
30 large development distance to adjust to the finer grid resolution in terms of spectrally similar conclusion. However, a purely shear-driven flow over a flat homogeneous surface can certainly be considered as an extreme case in terms of flow adjustment, as the vertical turbulent exchange, which is primarily driven by surface-roughness induced shear, is rather low compared to less idealized flows over non-flat terrain or with obstacles included. Hence, we expect that the required fetch length may decrease for rougher surfaces and more complex surface geometries.




Figure 10. Two-hour time- and y -averaged horizontal profiles of the friction velocity for the grid-spacing ratio of 2, 3, and 4, as well as the corresponding coarse- and fine-grid reference simulations.




Figure 11. Two-hour time- and y -averaged profiles of the horizontal wind speed for the grid-spacing ratio of 2, 3, and 4, taken at different distances downstream of the inflow boundary, indicated by the white solid lines in Fig. 9. Also, corresponding time- and y -averaged profiles from the fine and coarse reference simulations taken at the same locations are shown.




Figure 12. Frequency spectra of the resolved-scale TKE taken at different sampling locations downstream of the inflow boundary for the neutrally-stratified boundary layer at $z = 40$ m, for grid-spacing ratio a) 2, b) 3, and c) 4. TKE-spectra for the fine- and the corresponding coarse-grid reference simulations are also shown.

4.2.2 Neutrally stratified boundary-layer flow over a smooth three-dimensional hill

The hill case ~~was setup~~ is studied to compare flow statistics against the wind-tunnel observations conducted by Ishihara et al. (1999), who sampled data at different up- and downstream locations along the centre hill axis. This flow is simulated using a two-stage nesting configuration in which a second child domain with 1 m resolution is placed within the first child domain with 2 m resolution.

The terrain height of the smooth three-dimensional hill is given by

$$z(x, y) = H \cos^2 \left(\frac{\pi \sqrt{(x - x_0)^2 + (y - y_0)^2}}{2l} \right), \quad (13)$$

with the hill height $H = 40$ m and a hill radius $l = 100$ m, while x and y ~~indicated~~ indicate the location on the discrete grid and ~~x_0 and y_0~~ $x_0 = 2024$ m and $y_0 = 384$ m the location of the hill top. ~~Please note,~~ with respect to the ~~wind-tunnel experiment,~~ parent domain dimensions. Note that we up-scaled the hill dimension by a factor of ~~1000-~~ 1000 with respect to the wind-tunnel model. Again, the parent domain size is $L_x \times L_y \times L_z = 5.1 \times 1.5 \times 0.32$ km³ in the x -, y -, and z -directions, respectively. The child domain sizes are $L_x \times L_y \times L_z = 0.768 \times 0.384 \times 0.16$ km³ and $L_x \times L_y \times L_z = 0.576 \times 0.288 \times 0.12$ km³ for the first and second child domains, respectively. The upstream boundaries of the first and second child domains are placed about $9.7 H$ and $6.3 H$ upstream of the hill top, respectively. The smooth hill geometry is approximated with the grid-following stair-step geometry in PALM due to its orthogonal grid arrangement and its topography description system, see (Maronga et al., 2015).

Figure 13 shows the mean flow field along the centerline of the three-dimensional hill for the nested child simulations as well as ~~fine and coarse reference simulation~~ the fine- and coarse-grid reference simulations. Upwind of the hill the mean flow in the nested ~~simulation agrees~~ simulations agree with the one in the fine and coarse reference simulation. ~~Also in the lee~~ In the ~~coarse-grid reference simulation the re-circulation extends further downstream up to about $4.1 H$ on the lee side of the hill~~ the ~~flow field in the~~ compared to the 2 m nested and fine reference simulation ~~agree, both showing~~ which both show a re-circulation that extends up to about $3.75 H$ downstream of the hill top, ~~while in the coarse reference simulation,~~ With further increasing the grid resolution to 1 m the re-circulation ~~extends farther downstream up to about $4.1 H$.~~ Figure ~~zone~~ zone further shortens to about $3 H$. Figures 14 and 15 show the corresponding standard deviations of the u - and w -~~component~~ components sampled at different locations along the centerline of the hill. Upwind of the hill the standard deviations of the u - and w -~~component~~ components agree with the observations and show no significant difference between the simulations. However, leeward of the hill at $1.25 H$, the LES underestimates the standard ~~deviation~~ deviations of the u - and w -~~component~~ components in the 2 m and 4 m simulations, which is most pronounced in the ~~coarse reference simulation~~ coarse-grid reference simulation, while the profiles in the 1 m child domain of the nested simulation agree fairly well with the measurements. Further downstream, the ~~coarse-coarse-grid~~ reference run still slightly underestimates the observed standard deviations, while the ~~nested and fine reference simulation~~ 2 m child domain result and fine-grid reference result slightly overestimate the standard deviations, which is in agreement with results from the EPFL-LES model presented in Diebold et al. (2013) who employed a similar ~~grid resolution of the hill.~~ The fine grid resolution for this hill flow. With 1 m grid spacing the standard deviations further downstream are still slightly overestimated, though their vertical shape and amplitude is captured quite well. In order to

Table 1. Root-mean square deviation (RMSD) of the simulated u - and w -standard deviations evaluated against the corresponding wind-tunnel measurements at the positions shown in Fig. 14 and 15. The RMSD is evaluated at all respective heights of the observations and is averaged over all heights and profile locations.

Simulation	RMSD (σ_u)	RMSD (σ_w)
REF 4 m	0.043	0.078
REF 2 m	0.022	0.054
Child 2 m	0.026	0.055
Child 1 m	0.015	0.039

provide a more quantitative measure, Tab. 1 provides the root-mean square deviation (RMSD) of the profiles shown in Fig. 14 and 15. RMSD is defined as

$$RMSD(\psi) = \sqrt{\langle (\bar{\psi} - \bar{\psi}_{\text{Ref}})^2 \rangle}. \quad (14)$$

where ψ is any prognostic variable from the considered solution and ψ_{Ref} refers to the corresponding measurement value.

- 5 The RMSD-values indicate that the profiles converge towards the observations with increasing grid resolutions. Furthermore, the standard deviations from the ~~nested and fine reference simulation agree, showing 2 m child domain and from the fine-grid reference simulation show~~ only marginal differences among each other, ~~indicating that the nesting method and the fine-grid reference simulations lead to almost identical results~~. Considering that the hilltop is placed only about ~~300 m apart~~ 390 m and 250 m downstream of the child domain inflow boundary, this indicates that in more complex setups where topography is
- 10 present, the adjustment fetch can become significantly smaller compared to purely flat terrain as discussed in ~~Seet~~ Sec. 4.2.1.

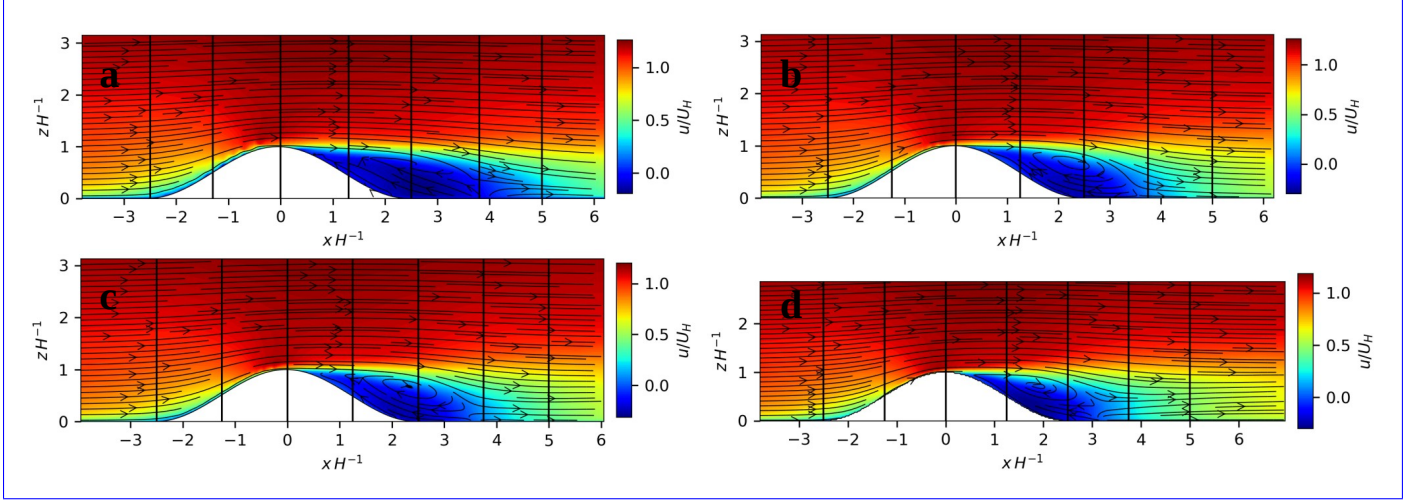


Figure 13. 2-hour Two-hour time-averaged vertical cross-section of the u -component field (colored contours) as well as the $u-w$ mean flow field (vector arrows) displayed along the centerline of the three-dimensional hill for a) the nested child domain 4 m reference simulation, b) the 2 m reference simulation, and c) the 4 m reference-2 m nested child simulation, and d) the 1 m nested child simulation. Vector arrows as well as the u -component are normalized with the reference wind speed taken at $z = H$ upwind of the hill. The ordinate and the abscissa are scaled with the hill height H . Note, the abscissa is centered at the hill top. The black vertical lines indicate the positions of the profiles displayed in Fig. 14 and 15. Note that the ordinate in d) is constrained by the vertical dimension of the 1 m child domain. Moreover, note that the x -dimension in c) and d) do not show the entire child domain extents but cut out minor parts at the left, right and top.

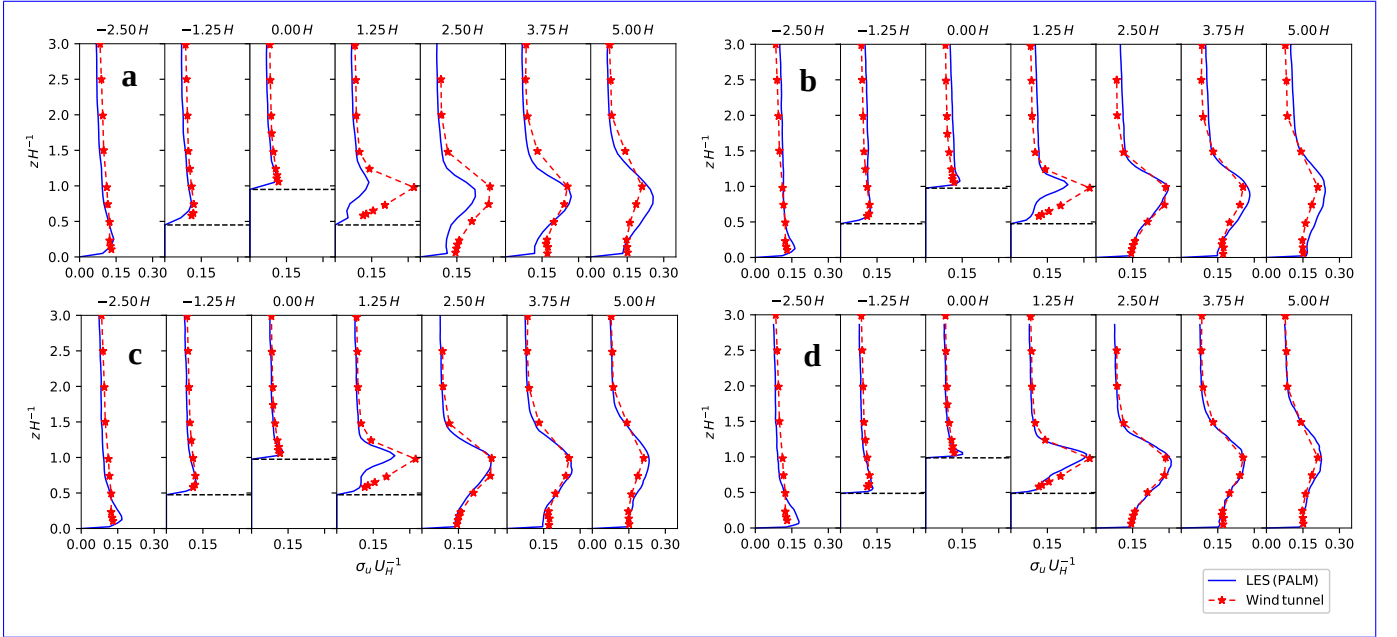


Figure 14. Two-hour time-averaged vertical profiles of the standard deviation of the u -component for the LES and well as the observed wind tunnel flow, for a) the nested-child-domain 4 m reference simulation, b) the 2 m reference simulation, and c) the 4-m-reference-2 m nested child simulation, and d) the 1 m nested child simulation. The ordinate is scaled with the hill height H . The standard deviation is normalized with the reference wind speed taken at $z = H$ upwind of the hill. The black dashed horizontal lines indicate the discrete height of the surface at the sampling location. Note that the vertical dimension of the 1 m nested child domain does not cover the entire vertical range of the measurements so that some data points from the LES are missing on the normalized height coordinate in d).

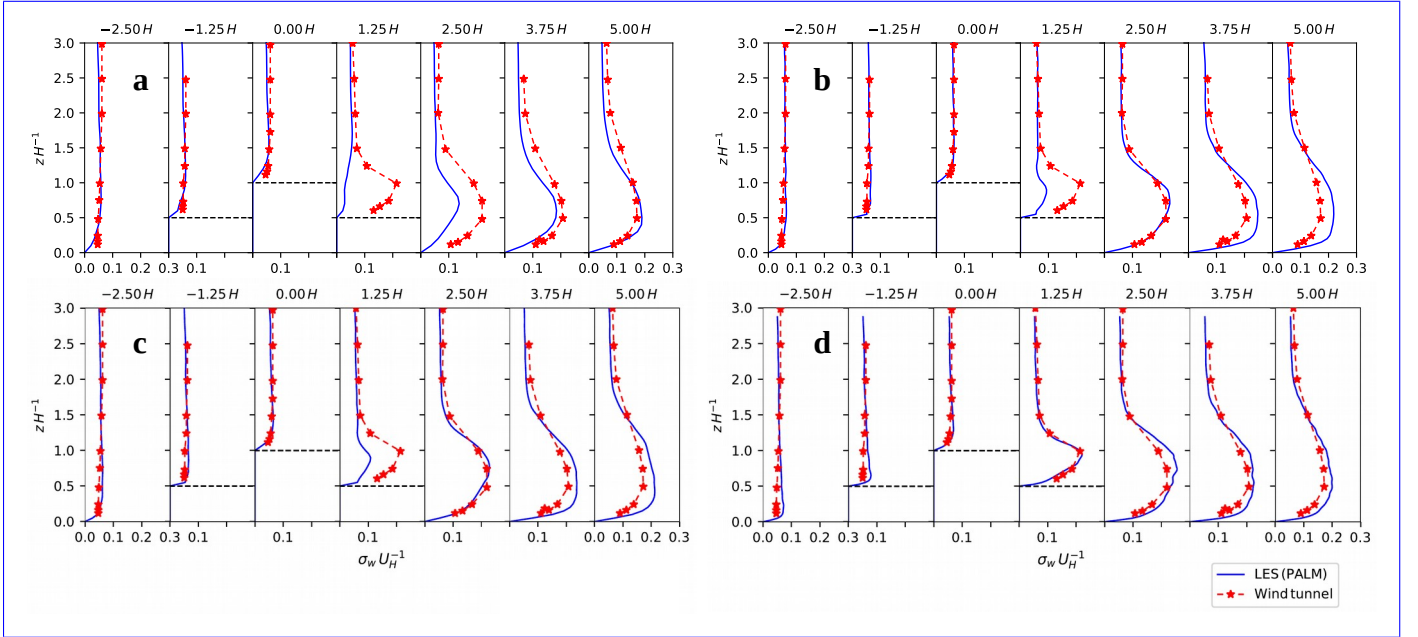


Figure 15. Two-hour time-averaged vertical profiles of the standard deviation of the w -component for the LES and well-as the observed wind tunnel flow, for a) the ~~nested-child domain~~ 4 m reference simulation, b) the 2 m reference simulation, and c) the ~~4-m-reference-2 m nested child simulation~~, and d) the 1 m nested child simulation. The ordinate is scaled with the hill height H . The standard deviation is normalized with the reference wind speed taken at $z = H$ upwind of the hill. The black dashed horizontal lines indicate the discrete height of the surface at the sampling location. Note that the vertical dimension of the 1 m nested child domain does not cover the entire vertical range of the measurements so that some data points from the LES are missing on the normalized height coordinate in d).

4.2.3 Neutrally stratified boundary layer over a regular array of cubes

The final test case features a neutral atmospheric boundary layer flow over flat terrain which becomes incident with a staggered pattern of cubical obstacles. The resulting flow scheme resembles urban canopy turbulence where the interaction between roughness elements and ABL turbulence is primarily resolved. Here, the cubical obstacle height is $H = 40$ m. The distance between the obstacles is $3H$ in the x -directions and $1H$ in the y -direction.

To demonstrate the flexibility of the nesting implementation, we carried out simulations with two different nested configurations illustrated in Fig. 16. The first (v1) case features a single child domain while the second (v2) case contains a two-stage nesting system where a second child domain is nested within the first. In the latter configuration, the first child acts as a parent for the second child domain. The isotropic grid spacing is 4 m in the root domain, 2 m in the second level nest (first child) and 1 m in the third level nest (second child). (Note, that the implementation does allow child locations to be selected such that their domain boundaries intersect with the obstacles.) The two example configurations represent nesting applications designed to meet different levels of accuracy demands. The v2 configuration is set to resolve the transition effect at the leading edge of the cube canopy and to capture the blunt-body wake interactions in sufficient detail within the center region of the cuboid canopy.

First, in the context of obstacle-resolved LES, we motivate the employment of an optional *canopy-restricted* (CR) interpolation strategy introduced in Sec. 3.5. For this purpose, consider Fig. 17 showing an instantaneous horizontal cross-section of vorticity vector magnitude at $z = 0.9H$ height for configuration v2. The image is focusing on a region where all domains with different resolutions are visible.

The visualization indicates the strength and spatial structure of the resolved turbulent eddies and how they are affected by grid resolution. The differences are significant. In such obstacle-resolving LES, the increased grid resolution has the ability alter the flow solution to such a degree that the anteropulation introduces details to the coarser parent which are inconsistent with the rest of the parent's flow solution. Particularly with blunt-body obstacle canopy flows, this discrepancy is clearly manifested as a locally changing resultant pressure drag (caused by the obstacles) within the anteropulated domain. To inspect this, we compute the resultant pressure drag coefficient

$$C_{Fp} = \frac{2}{\rho U_{ref}^2 S_{ref}} (F_{p,x}^2 + F_{p,y}^2)^{1/2} \quad (15)$$

for the differently coupled simulations. In Eq. 15 $u_{ref} = \langle \bar{u} \rangle|_{z=1.25H}$ is the reference wind speed, F_p is the resultant pressure force exerted on the cubes obtained by integrating the pressure over vertical walls, ρ is the density of air and S_{ref} is the accumulated frontal area of the cubes. The results are listed in Table 2, which makes evident the drastic difference between the values for the coarse reference and the two-way coupled parent (C_{Fp} [Coarse] vs. C_{Fp} [Root]: two-way). This large difference arises as the anteropulated solution within the obstacle canopy introduces a large-scale disturbance to the parent solution giving rise to unphysical secondary effects. These effects, in turn, lead to complicated feedback systems in the two-way coupled solutions whose realizations become ~~dependent~~dependent on the chosen nesting configuration.




Figure 16. An overview of the cubical-obstacle case layout. The obstacles are cubes with 40 m sides. The figure displays two nested arrangements: version 1 (v1) featuring a root domain and a secondary nest domain whereas version 2 (v2) also includes a tertiary nest domain embedded within a larger secondary nest. The root and nested domains are indicated with (\star) , $(\star\star)$, and $(\star\star\star)$ respectively in the upper left-hand-corner of each domain. The first child domain is displayed with white background for better visualization.

This problematic behavior is significantly abated by adopting the CR antepolation strategy setting here the vertical threshold at $1.25H$ via experimenting. This CR antepolation allows the parent and child flow fields to become strongly coupled while minimizing global inconsistencies in the parent solution. While all the child domain solutions over-predict the pressure drag, the two-way CR solution yield $C_{Fp}[\text{Child 2}]$ values that are ~~eloses~~ closest to the fine reference.

- 5 To further evaluate the nesting performance, we exploit root (normalized) mean square difference ($RNMSD$ or $RMSD$) and fractional bias (FB) as comparison metrics (see, Britter and Hanna, 2003) evaluated over successive xy -planes to assess the effectiveness of the nesting approach in obstacle-resolving LES cases. $RNMSD$ and $RMSD$ provide a measure of mean difference that is composed of random scatter and systematic bias whereas the fractional bias (FB) yields a specific measure for the systematic bias between the two solutions. The ~~metries~~ $RMSD$ metric is defined by Eq. (14) while $RNMSD$ and FB are




Figure 17. Instantaneous close-up view of vorticity magnitude (s^{-1}) on xy -plane at elevation $z = 0.9H$ for the v2 case with two nested domains. Black lines indicate the bounds of the first and the second nest. Note, only parts of the domain extents are displayed.

Table 2. Resultant pressure force coefficients C_{Fp} evaluated over Child 1 (**) domain shown in Fig. 16 (v1). Results for parent and child solutions are reported for one-way, two-way and two-way canopy-restricted (CR) methods, where the latter is a modified two-way coupling approach where the anteropulation is restricted (i.e. not allowed) within the obstacle canopy.

Version 1	Version 2		
	one-way	two-way	two-way CR
$C_{Fp}[\text{Root}]$	0.592	0.735	0.549
$C_{Fp}[\text{Child 1}]$	0.602	0.599	0.594
$C_{Fp}[\text{Fine}]$		0.583	
$C_{Fp}[\text{Coarse}]$		0.592	

defined as

$$RNMSD(\psi) = \sqrt{\frac{\langle (\bar{\psi} - \bar{\psi}_{\text{Ref}})^2 \rangle}{\langle \bar{\psi} \rangle \langle \bar{\psi}_{\text{Ref}} \rangle}} \quad \text{or } \underline{RNMSD\psi} = \sqrt{\langle (\bar{\psi} - \bar{\psi}_{\text{Ref}})^2 \rangle} \quad (16)$$

$$FB(\psi) = \frac{\langle \bar{\psi} \rangle - \langle \bar{\psi}_{\text{Ref}} \rangle}{0.5 (\langle \bar{\psi} \rangle + \langle \bar{\psi}_{\text{Ref}} \rangle)}. \quad (17)$$

where ψ is a generic prognostic variable from the considered case, while [here](#) ψ_{Ref} refers to the value from the reference simulation with ~~2 m~~ 2 m resolution. $RMSD$ is used instead of $RNMSD$ in cases where the product of double-averaged quantities used for normalization approaches zero. Similarly, FB is only evaluated for the streamwise velocity component because other components yield a near-zero denominator which contaminates the metric. The evaluations are performed for 15 xy -planes within the child domain (zone (**) in Fig. 16) which are equally spaced over the range $0 \leq z/H \leq 1.5$. We have excluded 128 m and 64 m wide development zones at the boundaries in the x and y -directions respectively. When the coarse (4 m resolution) reference solution is compared to the fine (2 m resolution) reference, the coarse solution is interpolated onto the fine grid before the comparison metrics are evaluated.

Both model variants v1 and v2 are included in the analysis to demonstrate how the size and placement of the child domains effect the metrics and also to illustrate the possibility to employ a cascade of nested domains. Although no comparison metrics are presented for the second child solution featuring 1 m resolution, its influence is embedded in the solution of the first child.

The $RNMSD$ and $RMSD$ profiles for the velocity components and their variances depicted in Figs. 18 and 19 lay bare the effectiveness of the presented nesting system and reveal the added benefit of the CR antepolation. While all the coupling approaches succeed in significantly reducing the discrepancy compared to the fine reference, the conventional two-way coupling exhibits the most pronounced level of deviation. The FB results in Fig. 20 indicate that the two-way coupled solution also contains the most systematic deviation, which is conform with the pressure drag results.

The one-way coupling approach performs consistently better than the unmodified two-way coupled in all metrics, but it is also associated with a systematic bias that is larger than the value by coarse reference. However, if the modest systematic shift in streamwise velocity can be accepted, the one-way coupling offers a cost-effective nesting coupling approach (see Sec. 4.3 for performance measures). Nonetheless, the results conclude that the introduced CR antepolation approach presents the most recommended coupling strategy for obstacle-resolving LES as it provided the best metrics in every category.

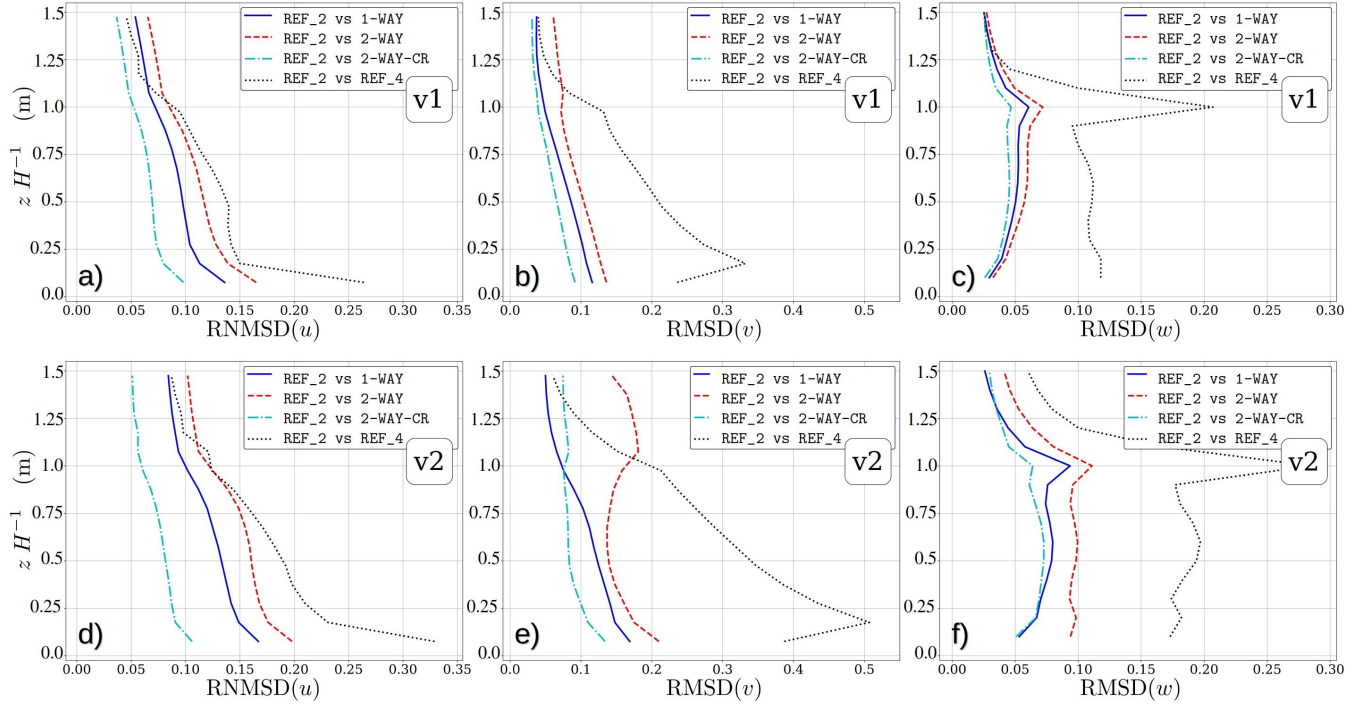


Figure 18. Vertical distributions of root (normalized) mean square difference ($RNMSD$ or $RMSD$) of velocity components for configuration v1 a-c) and configuration v2 d-f). The metrics are evaluated for 15 horizontal planes between the range $0 \leq z/H \leq 1.5$.




Figure 19. Vertical distribution of $RNMSD$ of the horizontal velocity variances for configuration v1 a-c) and configuration v2 d-f). The metrics are evaluated for 15 horizontal planes between the range $0 \leq z/H \leq 1.5$.




Figure 20. Fractional bias (FB) values evaluated over 15 xy -planes between the range $0 \leq z/H \leq 1.5$, for a) v1 and b) v2.

4.3 Performance issues

- Table 3 gives an overview of the consumed CPU time T_{CPU} , wall-clock time T_{wall} and some performance measures in the nested as well as fine- and coarse-grid reference simulations for the hill and the cube case simulations. T_{CPU} is here defined simply as $n_{\text{procs}} T_{\text{wall}}$ and n_{procs} is the number of parallel processes. As a rule of thumb, doubling the resolution leads to an increase in CPU time by approximately a factor of 16 (when the numerical time step is determined according to the CFL criterion). This can be observed comparing the coarse- and fine-grid reference simulations. Compared to the fine-grid reference simulations, the nested simulations consumed significantly less CPU time (up to 80% reduction) while increasing the computational cost by factors of 3.8 and 3.4 in hill and cube canopy cases compared to the ~~coarse-reference-coarse-grid~~ reference. The direct overhead due to the nesting, i.e. the fraction of the CPU-time used by all the nesting-related operations during the time-stepping is reasonably small in these tests, only from 2% to 16%. Also the CPU time per grid point per time step increases only quite moderately due to the nesting. These figures differ surprisingly much between the hill and the cube-array cases. This is most likely due to the fact that these cases were run on different computer systems. Although these factors depend on the child domain size, these ~~test-tests~~ make evident that the nesting technique can significantly reduce the computational cost, while yielding results that closely adhere to the non-nested fine-resolution simulation.
- Due to the inter/interpolation and the accompanied inter-model data transfer, the nesting itself consumes CPU time. In our tests the workload with respect to the number of grid points treated by a processor element was equal among the parent and the child simulation. With this optimal configuration, the two-way nesting consumed about 10-16% of the CPU time in our tests, while it consumes only about 2% in the one-way nesting. This suggests that most of the CPU time taken by two-way nesting is consumed in the interpolation and the associated child to parent data transfer.
- ~~Please note, It is important to bear in mind that~~ if the workload between child and parent processes is not well balanced, the faster processes need to wait before the data-transfer can start until the slower processes reach that point, reducing the computational efficiency of the nesting.

Table 3. ~~Required~~ Total number of grid points N_{gp} , number of parallel processes N_{procs} , required CPU time and wall-clock times and other performance measures for the ~~neutrally-stratified-hill~~ and the ~~convective-cube array v1 test case~~ cases. Here $N_{\Delta t}$ is the number of time steps taken during the run.

heightCase	CPU time $(\times 10^2 \text{ h}) N_{\text{gp}}$	Overhead nesting N_{procs}	$T_{\text{CPU, h}}$	Nesting overhead	$T_{\text{wall, h}}$	$T_{\text{CPU}}/(N_{\Delta t} N_{\text{gp}})$
Hill coarse	2.2 39 321 600	512	220	–	0.4297	0.658
Hill fine	39.1 314 572 800	2048	3910	–	1.909	0.692
Hill nest	8.31 45 219 840	768	831	16%	1.082	1.085
Cubes coarse	8.55 39 321 600	120	855	–	7.130	1.290
Cubes fine	129.0 314 572 800	480	12 900	–	26.88	1.166
Cubes v1 nest one-way	25.8 56 098 816	184	2580	2.1%	14.02	1.427
Cubes v1 nest two-way	29.1 56 098 816	184	2910	12%	15.82	1.636

5 Conclusions and future outlook

This article documents and evaluates an online LES-LES nesting scheme implemented into the PALM model system 6.0. The nesting system relies on the post-insertion approach and features both one-way and two-way coupling approaches. We give a detailed description of the model's relevant technical, algorithmic and numerical aspects and provide evidence for the accuracy gains the method introduces with a dramatically reduced computational cost compared to globally refined grid resolution. Particularly in urban boundary layer studies requiring obstacle-resolving LES, the nesting approach has proven essential.

The implementation of this three-dimensional nesting system is based on two-level parallelism involving inter-model and intra-model parallelization using MPI. This enables our nesting implementation to flexibly support multiple child domains which can be nested within their parent domain either in a parallel or recursively cascading configuration. All solutions involved within the nested simulation are advanced using a globally synchronized time step whereas the coupling between each parent/child pair is performed with interpolation (parent to child) and antinterpolation (child to parent) operations.

The nesting method is evaluated by performing a series of numerical experiments with an objective to demonstrate that the refined child solution (nested within a coarser parent) approaches the non-nested reference solution obtained by employing fine resolution globally.

The first test case features horizontally homogeneous convective boundary layer (CBL) with no mean mean wind. In this case, first and second order boundary-layer statistics are well captured in the child domain and are closely comparable to non-nested high-resolution reference statistics. Further, due to the local nature of turbulence production and the weak advection from parent into the child, the flow statistics show almost no dependence on the distance to the child boundaries. However, in case of several hours long averaging times we found that a nonphysical secondary circulation develops although the surface heating is homogeneous. We hypothesize that this secondary circulation is an inherent consequence of the spatially changing description of flow physics in the parent and child solutions. ~~This~~ Even though we demonstrated this issue with a rather idealized setup using unrealistically long averaging time and such nonphysical circulation are probably minimized in less idealized simulations, e.g. with wind, a daily cycle and realistic terrain surfaces, we believe that this should be kept in mind when applying the nesting system to CBL-problems.

The second test case simulated neutrally stratified boundary layer flow over flat terrain. The nested simulations reveal that the flow solution within the child domain must undergo a development phase, as the flow solution adjusts to the higher resolution, before reaching equilibrium state again. The required development length depends on the grid-spacing ratio between parent and child. However, a purely shear-driven flow over a homogeneous flat terrain can be considered as an extreme scenario with respect to the development length of turbulence, while in cases with more complex surface geometry the flow adapts within shorter development distances. Beyond the development distance, the child solution for grid-spacing ratios of two and three agree well with the non-nested fine-reference solution, but in case of grid-spacing ratio of four the results clearly deviate from the fine-reference solution.

The third numerical experiment featured boundary layer flow (similar to second test case) over a smooth three-dimensional hill. This test case also exploits wind tunnel measurements to strengthen the nesting model evaluation. In this case, the flow

statistics in the windward and the leeward part of the hill are almost the same as in a fine-reference simulation and agree well with wind-tunnel observations presented in Ishihara et al. (1999).

The final test case examines a flow system where a fully developed boundary layer flow becomes incident with a staggered arrangement of cube-shaped obstacles. This flow scenario closely resembles an obstacle-resolving urban boundary layer flow situation. The case revealed that in two-way coupled simulations, the interpolated child solution introduces discrepancies within the parent domain which is manifested as elevated pressure drag within the interpolated zone. This complication is remedied by introducing *canopy-restricted* interpolation approach, where interpolation is omitted within the obstacle canopy. By computing comparison metrics, root-normalized mean square difference and fractional bias, to quantify the difference between the fine reference and the nested solutions, the canopy-restricted two-way coupling is shown to be the best coupling strategy for obstacle-resolving LES studies.

Future outlook

Future development is planned to include the following tasks. Incorporation of PALM's Lagrangian particle model in the nesting system in order to enable Lagrangian dispersion studies in urban environments in such a way that particles can be transferred between parent and child domains depending on their position. Thus, the long-distance transport of e.g. pollutants, can be simulated in a coarse-resolution parent grid, while dispersion on the street-scale for specific locations can be simulated in a fine-resolution child domain. We note that this has been already implemented into PALM and is available to users, but further sensitivity tests with respect to the treatment of stochastic subgrid-scale particle speeds (Weil et al., 2004) are still pending. A thorough description and verification of the particle nesting will be published in a follow-up article.

Further, we note that the PALM model system 6.0 includes also a RANS (Reynolds-averaged Navier-Stokes) mode offering two different turbulence closures to calculate the eddy diffusivity, that are a TKE- l and a TKE- ϵ closure according to Mellor and Yamada (1974, 1982). Besides the LES-LES nesting the nesting system is being extended to handle RANS-LES and RANS-RANS nesting, which require coupling of additional RANS-variables. Moreover, in a companion paper in this special issue we present a pure one-way off-line mesoscale nesting method in which the PALM model system 6.0 is nested into meso-scale models such as COSMO or WRF. This will allow modelling of meso-scale processes on a much larger coarse-grid domain as e.g. shown by Muñoz-Esparza et al. (2017), while concurrently focusing on fine-scale processes within certain areas using the present LES nesting approach.

Furthermore, to date, the timestep in all parent and child models is synchronized and restricted to the minimum of the time steps determined by each model independently using the CFL criterion. To our experience, the global timestep is often restricted by the flow around building edges where high wind speeds occur within the fine-grid child domains. Hence, we plan to implement a time-splitting into PALM where the parent and child models will be coupled only at the end of the parent timesteps. This would allow to run coarser-scale parent domains with larger time steps. Thus, computational time could be saved in the time-integration of the parent simulation as well as in the inter-model communication between parent and child.

Appendix A: Technical realization

A1 General

The nested model system is implemented using two levels of MPI communicators. The inter-model communication (communication between model domains) is handled by a global communicator using the one-sided communication pattern (Remote Memory Access, RMA). The intra-model communication (communication between subdomains within each model domain) is two-sided and it is handled using a 2-D communicator that has different color for each model. The intra-model communication system is the baseline parallelization of PALM (Maronga et al., 2015).

Data transferred from parent to child and from child to parent is always stored in the coarser parent-model grid in order to minimize the amount of data transfer. This means that the interpolations and antinterpolations are always performed by the child. For these purposes, children contain auxiliary arrays which follow the parent-grid spacings and indexing for each prognostic variable to be coupled covering the overlap domain plus necessary number of ghost-node layers.

A2 Initialization

Mapping between each parent and child model domain decompositions as well as all the necessary index mappings are determined in the initialization phase and stored so that the coupling actions during the time-stepping are straightforward and efficient.

Initial conditions for the root are set similarly to non-nested runs. The root then sends initial field data to its children which interpolate their own initial conditions from the data received from the root. Next the first-level children send their data to their children, if any, and so on. The basic interpolation subroutines for child boundary conditions operate only on the ghost nodes behind the child-model boundaries. Therefore a separate three-dimensional interpolation subroutine is implemented to generate initial fields for all the nest domains from their parent-model fields. The same interpolation algorithm is used here as in the interpolations for child boundary conditions.

A3 ~~Time-synchronization~~

~~Time-synchronization is taken care by simply selecting the minimum of the time steps determined by each model independently and broadcasting this time-step value for all models. Each model inputs and outputs in the same way.~~

25 A3 Modularization

The data transfer between parents and children is conducted by code contained by five specific fortran modules forming a module set called PALM Model Coupler (PMC). Calls to the PMC-subroutines are mostly made in PMC-interface module (pmc_interface_mod.f90) such that only a small number of calls to the PMC-interface subroutines are needed within the baseline PALM code. This way, the changes to the baseline code were kept minimal. The PMC-interface module also contains subroutines for the nesting-related initialization actions, interpolation, antinterpolation, child mass-balance forcing, etc.

A4 MPI implementation

While reading the input namelists, the PALM root process checks if a namelist called "&nesting_parameters" is given in the parameter input file PARIN. If not, subroutine called `pmc_init_model` resets all nesting-related parameters (`coupling_layout` etc.) and sets `MPI_COMM_WORLD` as the base global MPI-communicator `comm_palm`. The run then continues in standard

5 way without nesting. If the namelist "&nesting_parameters" is found and correctly input, the root process of the root model distributes this information to all other processes via `MPI_COMM_WORLD`. Then, all the necessary nesting related parameters are determined and the base communicator is split into different colors ¹ for each model based on the model id number. This is done by calling `MPI_COMM_SPLIT`. Now each model has its own process group and associated individual base commu-

10 nicator color such that each model's internal communication is not visible to other models. After this, the mappings between models are determined. Each model, except the root model, identifies its parent model and creates an inter-communicator between the process groups of itself and its parent model. This is realized by calling `MPI_INTERCOMM_CREATE`. In the same way, each model identifies its all children if any, and creates inter communicators between the process groups of itself and all of its children. These inter communicators are only used to transfer setup data between the root processes of the parent and child models. For 3D model data transfers between parent and child, specific ~~intra-communicators are~~

15 intra-communicator is created by merging inter-communicators of each processes within the remote process groups. This is made after `pmc_init_model` separately for child and parent models (note that a model may be both parent and child) in subroutines `pmc_childdinit` and `pmc_parentinit` by calling `MPI_INTERCOMM_MERGE`. After the `pmc`-initialization, the run of each model goes as usual. Cartesian topology-based communicator `comm_2d` is created by each model from its color of the base communicator `comm_palm` using `MPI_CART_CREATE`.

20 The model internal communication is done in the usual way, e.g. by calling the boundary exchange routines. All data transfer between parent and child models is done within the PMC interface. For this communication MPI one sided communication (RMA) is used. An RMA window is opened on the parent side. To transfer data from parent to child, the parent fills the RMA window via local copy. After synchronization via `MPI_WIN_FENCE`, the child processes can fetch the data across the network with `MPI_GET`. While transferring data in the opposite direction, the child first transfers the data via `MPI_PUT`. After another

25 `MPI_WIN_FENCE` call, the parent copies the data out of the RMA window into the local model data area.

Appendix B: Thoughts on an alternative interpolation method

Should higher interpolation accuracy across the boundaries be sought, the following considerations are relevant. As stated by Zhou et al. (2018), to satisfy the global flux-conservation requirement, one of the flux factors, either the advective velocity component or the advected variable, must be constant within the anterpolation cell. This implies zeroth-order interpolation.

30 The other factor must be interpolated using any reversible interpolation scheme.

¹The term color means here that the communicator has the same name for all models (process groups), but they are, however, individual communicators guaranteeing that communication of one model is not interfered by the others.

As stated above, the quadratic Clark and Farley (1984) scheme should not be used because it employs a stencil wider than the parent-grid cell which leads to problems with complex geometries. On the other hand, tri-linear interpolation has a favorable stencil width, but it is not suitable for the scheme as such is not reversible. However, linearly interpolated values $\tilde{\phi}_{i,j,k}$ can be made reversible by introducing an additional correction $\phi_{i,j,k} = \tilde{\phi}_{i,j,k} + \Delta\phi_{i,j,k}$ which guarantees the reversibility. The

5 reversibility correction $\Delta\phi_{i,j,k}$ depends on the difference $\Delta\Phi_{I,J,K}$ between the original parent-grid value $\Phi_{I,J,K}$ and the value obtained by interpolating the linearly interpolated values to the parent-grid node I, J, K as

$$\Delta\Phi = \Phi - \hat{\phi}. \quad (\text{B1})$$

$\Delta\Phi_{I,J,K}$ is a constant value within the parent-grid cell, hence a question arises: how to distribute the correction to the child-grid nodes i, j, k such that $\hat{\Delta\phi}_{I,J,K} = \Delta\Phi_{I,J,K}$? The simplest choice is $\Delta\phi_{i,j,k} = \Delta\Phi_{I,J,K}$, but this choice is not recommendable

10 in the cases of positive definite scalar variables as it could lead to negative values when Φ is close to zero. In principle this problem could be avoided by weighting the local corrections in proportion to the local differences $\Phi_{I,J,K} - \tilde{\phi}_{i,j,k}$ but this simply reduces the method back to the zeroth-order baseline method. To make this approach useful, a more advanced technique to distribute the correction ought to be developed. However, this is beyond the scope of the present work as stated in Sec. 3.4.

Acknowledgements. Test runs with PALM have been performed at the supercomputers of the North-German Super-computing Alliance

15 (HLRN), Germany, and CSC – IT Center for Science, Finland.

This research was funded by Academy of Finland grant number 277664. BM and MS were supported by the Federal German Ministry of Education and Research (BMBF) under grant 01LP1601 within the framework of *Research for Sustainable Development* (FONA; <https://www.fona.de>). KK was supported by the Federal German Ministry of Economy and Energy (BMU) under grant 0325719C.

Author contributions. Coordination of the study: AH, SR. Design and implementation of the inter-model communication: KK. Theoretical

20 considerations and implementation of the nesting interface: mainly AH with contributions from SR, MA, MS and BM. Simulations, post-processing and analysis of model results: AH, MS, MA, BM, CK, FB, GT, NM. Drafting of the manuscript: AH, MA, MS, BM. Revision of the manuscript: all authors.

Appendix C: Code availability

The PALM model system is freely available at <http://palm-model.org> and distributed under the GNU General Public Licence v3 (<http://www.gnu.org/copyleft/gpl.html>). However, the simulations presented in this document were performed using a slightly modified code based on revision 4295. This modified source code (4295M) as well as the input files for the test runs are available at <https://doi.org/10.25835/0090593> (Hellsten et al., 2020). Numerous pre- and post-processing scripts are available at <http://doi.org/10.5281/zenodo.4005687> (Auvinen et al., 2020b).

25

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

References

- Ahmad, N. H., Inagaki, A., Kanda, M., Onodera, N., and Aoki, T.: Large-eddy simulation of the gust index in an urban area using the lattice Boltzmann method, *Boundary-Layer Meteorology*, 163, 447–467, 2017.
- Aidun, C. K. and Clausen, J. R.: Lattice-Boltzmann method for complex flows, *Annual review of fluid mechanics*, 42, 439–472, 2010.
- 5 Auvinen, M., Boi, S., Hellsten, A., Tanhuanpää, T., and Järvi, L.: Study of Realistic Urban Boundary Layer Turbulence with High-Resolution Large-Eddy Simulation, *Atmosphere*, 11, 201, doi:10.3390/atmos11020201, <http://dx.doi.org/10.3390/atmos11020201>, 2020a.
- Auvinen, M., Karttunen, S., and Kurppa, M.: P4UL: Pre- and Post-Processing Python Library for Urban LES Simulations, doi:10.5281/zenodo.4005687, <https://doi.org/10.5281/zenodo.804282>, 2020b.
- Bou-Zeid, E., Overney, J., Rogers, B. D., and Parlange, M. B.: The Effects of Building Representation and Clustering in Large-Eddy Simulations of Flows in Urban Canopies, *Boundary-Layer Meteorology*, 132, 415–436, doi:10.1007/s10546-009-9410-6, <https://doi.org/10.1007/s10546-009-9410-6>, 2009.
- 10 Britter, R. E. and Hanna, S. R.: FLOW AND DISPERSION IN URBAN AREAS, *Annual Review of Fluid Mechanics*, 35, 469–496, doi:10.1146/annurev.fluid.35.101101.161147, <https://doi.org/10.1146/annurev.fluid.35.101101.161147>, 2003.
- Buccolieri, R. and Hang, J.: Recent Advances in Urban Ventilation Assessment and Flow Modelling, *Atmosphere*, 10, 144, doi:10.3390/atmos10030144, <http://dx.doi.org/10.3390/atmos10030144>, 2019.
- 15 Chung, D. and McKeon, B. J.: Large-eddy simulation of large-scale structures in long channel flow, *Journal of Fluid Mechanics*, 661, 341–364, doi:10.1017/S0022112010002995, 2010.
- Clark, T. and Farley, R.: Severe downslope windstorm calculations in two and three spatial dimensions using anelastic interactive grid nesting: A possible mechanism for gustiness, *Journal of Atmospheric Sciences*, 41, 329–350, 1984.
- 20 Clark, T. and Hall, W.: Multi-domain simulations of the time dependent Navier-Stokes equations: benchmark error analysis of some nesting procedures, *Journal of Computational Physics*, 92, 456–481, 1991.
- Daniels, M., Lundquist, K., Mirocha, J., Wiersema, D., and Chow, F.: A new vertical grid nesting capability in the Weather Research and Forecasting WRF Model, *Monthly Weather Review*, 144, 3725–3747, 2016.
- de Roode, S. R., Duijkerke, P. G., and Jonker, H. J.: Large-eddy simulation: How large is large enough?, *Journal of the atmospheric sciences*, 61, 403–421, 2004.
- 25 Deardorff, J.: Stratoculumus-capped mixed layers derived from a three-dimensional model, *Boundary-Layer Meteorology*, 18, 495–527, 1980.
- Diebold, M., Higgins, C., Fang, J., Bechmann, A., and Parlange, M. B.: Flow over Hills: A Large-Eddy Simulation of the Bolund Case, *Boundary-Layer Meteorology*, 148, 177–194, doi:10.1007/s10546-013-9807-0, 2013.
- 30 Fishpool, G. M., Lardeau, S., and Leschziner, M. A.: Persistent Non-Homogeneous Features in Periodic Channel-Flow Simulations, *Flow, Turbulence and Combustion*, 83, 323–342, doi:10.1007/s10494-009-9209-z, <https://doi.org/10.1007/s10494-009-9209-z>, 2009.
- Gehrke, K. F., Sührling, M., and Maronga, B.: Modeling of land-surface interactions in the PALM model system 6.0: Land surface model description, first evaluation, and sensitivity to model parameters, *Geoscientific Model Development Discussions*, 2020, 1–34, doi:10.5194/gmd-2020-197, <https://gmd.copernicus.org/preprints/gmd-2020-197/>, 2020.
- 35 Giometto, M., Christen, A., Meneveau, C., Fang, J., Krafczyk, M., and Parlange, M.: Spatial Characteristics of Roughness Sublayer Mean Flow and Turbulence Over a Realistic Urban Surface, *Boundary-Layer Meteorology*, pp. 1–28, 2016.

- Gropp, W., Lusk, E., and Skjellum, A.: Using MPI: Portable parallel programming with the Message Passing Interface, 2nd edition, MIT Press, Cambridge, MA, 1999.
- Hackbusch, W.: Multigrid methods and applications, Springer, Berlin, Heidelberg, New York, 378 pp., 1985.
- Heinze, R., Dipankar, A., Henken, C. C., Moseley, C., Sourdeval, O., Trömel, S., Xie, X., Adamidis, P., Ament, F., Baars, H., Barthlott, C., Behrendt, A., Blahak, U., Bley, S., Brdar, S., Brueck, M., Crewell, S., Deneke, H., Di Girolamo, P., Evaristo, R., Fischer, J., Frank, C., Friederichs, P., Göcke, T., Gorges, K., Hande, L., Hanke, M., Hansen, A., Hege, H.-C., Hoose, C., Jahns, T., Kalthoff, N., Klocke, D., Kneifel, S., Knippertz, P., Kuhn, A., van Laar, T., Macke, A., Maurer, V., Mayer, B., Meyer, C. I., Muppa, S. K., Neggers, R. A. J., Orlandi, E., Pantillon, F., Pospichal, B., Röber, N., Scheck, L., Seifert, A., Seifert, P., Senf, F., Siligam, P., Simmer, C., Steinke, S., Stevens, B., Wapler, K., Weniger, M., Wulfmeyer, V., Zängl, G., Zhang, D., and Quaas, J.: Large-eddy simulations over Germany using
- 10 ICON: a comprehensive evaluation, *Quarterly Journal of the Royal Meteorological Society*, 143, 69–100, doi:10.1002/qj.2947, <https://rmets.onlinelibrary.wiley.com/doi/abs/10.1002/qj.2947>, 2017.
- Hellsten, A. and Zilitinkevich, S.: Role of convective structures and background turbulence in the dry convective boundary layer, *Boundary-Layer Meteorology*, 149, 323–353, 2013.
- Hellsten, A., Ketelsen, K., Sührling, M., Auvinen, M., Maronga, B., Knigge, C., Barmpas, F., Tsegas, G., Moussiopoulos, N., and
- 15 Raasch, S.: Dataset: A Nested Multi-Scale System Implemented in the Large-Eddy Simulation Model PALM model system 6.0, doi:10.25835/0090593, <https://doi.org/10.25835/0090593>, 2020.
- Heus, T., van Heerwaarden, C., Jonker, H., Siebesma, A. P., Axelsen, S., van den Dries, K., Geoffroy, O., Moene, A., Pino, D., de Roode, S., et al.: Formulation of and numerical studies with the Dutch Atmospheric Large-Eddy Simulation (DALES), *Geosci. Model Dev*, 3, 415–444, 2010.
- 20 Huq, S., De Roo, F., Raasch, S., and Mauder, M.: Vertically nested LES for high-resolution simulation of the surface layer in PALM (version 5.0), *Geoscientific Model Development*, 12, 2523–2538, 2019.
- Hutchins, N. and Marusic, I.: Evidence of very long meandering features in the logarithmic region of turbulent boundary layers, *Journal of Fluid Mechanics*, 579, 1–28, 2007.
- Hutchins, N., Chauhan, K., Marusic, I., Monty, J., and Klewicki, J.: Towards reconciling the large-scale structure of turbulent boundary layers
- 25 in the atmosphere and laboratory, *Boundary-layer meteorology*, 145, 273–306, 2012.
- Ishihara, T., Hibi, K., and Oikawa, S.: A wind tunnel study of turbulent flow over a three-dimensional steep hill, *Journal of Wind Engineering and Industrial Aerodynamics*, 83, 95 – 107, doi:[https://doi.org/10.1016/S0167-6105\(99\)00064-1](https://doi.org/10.1016/S0167-6105(99)00064-1), 1999.
- Karttunen, S., Kurppa, M., Auvinen, M., Hellsten, A., and Järvi, L.: Large-eddy simulation of the optimal street-tree layout for pedestrian-level aerosol particle concentrations – A case study from a city-boulevard, *Atmospheric Environment: X*, 6, 100073, doi:<https://doi.org/10.1016/j.aeaoa.2020.100073>, <http://www.sciencedirect.com/science/article/pii/S2590162120300125>, 2020.
- 30 Kataoka, H. and Mizuno, M.: Numerical flow computation around aeroelastic 3D square cylinder using inflow turbulence, 2002.
- Khan, B., Banzhaf, S., Chan, E. C., Forkel, R., Kanani-Sührling, F., Ketelsen, K., Kurppa, M., Maronga, B., Mauder, M., Raasch, S., Russo, E., Schaap, M., and Sührling, M.: Development of an atmospheric chemistry model coupled to the PALM model system 6.0: Implementation and first applications, *Geoscientific Model Development Discussions*, 2020, 1–34, doi:10.5194/gmd-2020-286, <https://gmd.copernicus.org/preprints/gmd-2020-286/>, 2020.
- 35 Krč, P., Resler, J., Sührling, M., Schubert, S., Salim, M. H., and Fuka, V.: Radiative Transfer Model 3.0 integrated into the PALM model system 6.0, *Geoscientific Model Development Discussions*, 2020, 1–41, doi:10.5194/gmd-2020-168, <https://gmd.copernicus.org/preprints/gmd-2020-168/>, 2020.

- Kurihara, Y., Tripoli, G., and Bender, M.: Design of a Movable Nested-Mesh Primitive Equation Model, *Monthly Weather Review*, 107, 239–249, 1979.
- Kurppa, M., Hellsten, A., Roldin, P., Kokkola, H., Tonttila, J., Auvinen, M., Kent, C., Kumar, P., Maronga, B., and Järvi, L.: Implementation of the sectional aerosol module SALSA2.0 into the PALM model system 6.0: model development and first evaluation, *Geoscientific Model Development*, 12, 1403–1422, doi:10.5194/gmd-12-1403-2019, <https://www.geosci-model-dev.net/12/1403/2019/>, 2019.
- 5 Kurppa, M., Roldin, P., Strömberg, J., Balling, A., Karttunen, S., Kuuluvainen, H., Niemi, J. V., Pirjola, L., Rönkkö, T., Timonen, H., Hellsten, A., and Järvi, L.: Sensitivity of spatial aerosol particle distributions to the boundary conditions in the PALM model system 6.0, *Geoscientific Model Development Discussions*, 2020, 1–33, doi:10.5194/gmd-2020-163, <https://www.geosci-model-dev-discuss.net/gmd-2020-163/>, 2020.
- 10 Letzel, M., Krane, M., and Raasch, S.: High resolution urban large-eddy simulation studies from street canyon to neighbourhood scale, *Atmospheric Environment*, 42, 8770–8784, 2008.
- Maronga, B. and Raasch, S.: Large-Eddy Simulations of Surface Heterogeneity Effects on the Convective Boundary Layer During the LITFASS-2003 Experiment, *Boundary-Layer Meteorology*, 146, 17–44, doi:10.1007/s10546-012-9748-z, <https://doi.org/10.1007/s10546-012-9748-z>, 2013.
- 15 Maronga, B., Gryschka, M., Heinze, R., Hoffmann, F., Kanani-Sühring, F., Keck, M., Ketelsen, K., Letzel, M., 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, *Geoscientific Model Development*, 8, 1539–1637, 2015.
- Maronga, B., Gross, G., Raasch, S., Banzhaf, S., Forkel, R., Heldens, W., Kanani-Sühring, F., Matzarakis, A., Mauder, M., Pavlik, D., Pfafferott, J., Schubert, S., Seckmeyer, G., Sieker, H., and Winderlich, K.: Development of a new urban climate model based on the model
- 20 PALM - Project overview, planned work, and first achievements, *Meteorol. Z.*, 28, 105–119, doi:10.1127/metz/2019/0909, 2019.
- Maronga, B., Banzhaf, S., Burmeister, C., Esch, T., Forkel, R., Fröhlich, D., Fuka, V., Gehrke, K. F., Geletič, J., Giersch, S., Gronemeier, T., Groß, G., Heldens, W., Hellsten, A., Hoffmann, F., Inagaki, A., Kadasch, E., Kanani-Sühring, F., Ketelsen, K., Khan, B. A., Knigge, C., Knoop, H., Krč, P., Kurppa, M., Maamari, H., Matzarakis, A., Mauder, M., Pallasch, M., Pavlik, D., Pfafferott, J., Resler, J., Rissmann, S., Russo, E., Salim, M., Schrempf, M., Schwenkel, J., Seckmeyer, G., Schubert, S., Sühring, M., von Tils, R., Vollmer, L., Ward, S., Witha,
- 25 B., Wurps, H., Zeidler, J., and Raasch, S.: Overview of the PALM model system 6.0, *Geoscientific Model Development*, 13, 1335–1372, doi:10.5194/gmd-13-1335-2020, <https://www.geosci-model-dev.net/13/1335/2020/>, 2020.
- Mellor, G. L. and Yamada, T.: A Hierarchy of Turbulence Closure Models for Planetary Boundary Layers, *Journal of the Atmospheric Sciences*, 31, 1791–1806, doi:10.1175/1520-0469(1974)031<1791:AHOTCM>2.0.CO;2, 1974.
- Mellor, G. L. and Yamada, T.: Development of a turbulence closure model for geophysical fluid problems, *Reviews of Geophysics*, 20,
- 30 851–875, doi:10.1029/RG020i004p00851, 1982.
- Miller, M. J. and Thorpe, A. J.: Radiation conditions for the lateral boundaries of limited-area numerical models, *Quarterly Journal of the Royal Meteorological Society*, 107, 615–628, doi:10.1002/qj.49710745310, <https://rmets.onlinelibrary.wiley.com/doi/abs/10.1002/qj.49710745310>, 1981.
- Mirocha, J., Kirkil, G., Bou-Zeid, E., Chow, F. K., and Kosović, B.: Transition and equilibration of neutral atmospheric boundary layer flow in one-way nested large-eddy simulations using the Weather Research and Forecasting model, *Monthly Weather Review*, 141, 918–940, doi:10.1175/MWR-D-11-00263.1, <https://doi.org/10.1175/MWR-D-11-00263.1>, 2013.
- Moeng, C., Dudhia, J., Klemp, J., and Sullivan, P.: Examining two-way grid nesting for large-eddy simulation of the PBL using the WRF model, *Monthly Weather Review*, 135, 2295–2311, 2007.

- Muñoz-Esparza, D., Kosović, B., García-Sánchez, C., and van Beeck, J.: Nesting turbulence in an offshore convective boundary layer using large-eddy simulations, *Boundary-layer meteorology*, 151, 453–478, 2014.
- Muñoz-Esparza, D., Lundquist, J. K., Sauer, J. A., Kosović, B., and Linn, R. R.: Coupled mesoscale-LES modeling of a diurnal cycle during the CWEX-13 field campaign: From weather to boundary-layer eddies, 2017.
- 5 Munters, W., Meneveau, C., and Meyers, J.: Shifted periodic boundary conditions for simulations of wall-bounded turbulent flows, *Physics of Fluids*, 28, 025 112, doi:10.1063/1.4941912, <https://doi.org/10.1063/1.4941912>, 2016.
- Nakayama, H., Takemi, T., and Nagai, H.: Development of LOcal-scale High-resolution atmospheric DIspersion Model using Large-Eddy Simulation. Part 5: detailed simulation of turbulent flows and plume dispersion in an actual urban area under real meteorological conditions., *Journal Nuclear Science and Technology*, 53, 887–908, doi:10.1080/00223131.2015.1078262, <https://doi.org/10.1080/00223131.2015.1078262>, 2016.
- 10 Nunalee, C. G., Kosović, B., and Bieringer, P. E.: Development of LOcal-scale High-resolution atmospheric DIspersion Model using Large-Eddy Simulation. Part 5: detailed simulation of turbulent flows and plume dispersion in an actual urban area under real meteorological conditions., *Atmospheric Environment*, 99, 571–581, doi:10.1016/j.atmosenv.2014.09.070, <https://doi.org/10.1016/j.atmosenv.2014.09.070>, 2014.
- 15 Patrinos, A. N. A. and Kistler, A. L.: A numerical study of the Chicago lake breeze, *Boundary-Layer Meteorol.*, 12, 93–123, 1977.
- Raasch, S. and Schröter, M.: PALM – A large-eddy simulation model performing on massively parallel computers, *Meteorologische Zeitschrift*, 10, 363–372, 2001.
- Resler, J., Krč, P., Belda, M., Juruš, P., Benešová, N., Lopata, J., Vlček, O., Damašková, D., Eben, K., Derbek, P., Maronga, B., and Kanani-Sühring, F.: PALM-USM v1.0: A new urban surface model integrated into the PALM large-eddy simulation model, *Geoscientific Model Development*, 10, 3635–3659, doi:10.5194/gmd-10-3635-2017, 2017.
- 20 Saiki, E. M., Moeng, C.-H., and Sullivan, P. P.: Large-eddy simulation of the stably stratified planetary boundary layer, *Boundary Layer Meteorol.*, 95, 1–30, 2000.
- Skamarock, W. C., Klemp, J. B., Dudhia, J., Gill, D. O., Barker, D. M., Duda, M. G., Huang, X.-Y., Wang, W., and Powers, J. G.: G.: A description of the Advanced Research WRF version 3, Tech. Rep. NCAR/TN-475+ STR, NCAR, 2008.
- 25 Sühring, M., Maronga, B., Herbort, F., and Raasch, S.: On the Effect of Surface Heat-Flux Heterogeneities on the Mixed-Layer-Top Entrainment, *Boundary-Layer Meteorology*, 151, 531–556, doi:10.1007/s10546-014-9913-7, <https://doi.org/10.1007/s10546-014-9913-7>, 2014.
- Sullivan, P. and Patton, E.: The effect of mesh resolution on convective boundary layer statistics and structures generated by large-eddy simulation, *Journal of the Atmospheric Sciences*, 68, 2395–2415, 2011.
- Sullivan, P., McWilliams, J., and Moeng, C.-H.: A grid nesting method for large-eddy simulation of planetary boundary-layer flow, *Boundary-Layer Meteorology*, 80, 167–202, 1996.
- 30 Tominaga, Y. and Stathopoulos, T.: CFD simulation of near-field pollutant dispersion in the urban environment: A review of current modeling techniques, *Atmospheric Environment*, 79, 716 – 730, doi:<https://doi.org/10.1016/j.atmosenv.2013.07.028>, <http://www.sciencedirect.com/science/article/pii/S1352231013005499>, 2013.
- Tseng, Y.-H., Meneveau, C., and Parlange, M. B.: Modeling flow around bluff bodies and predicting urban dispersion using large eddy simulation, *Environmental science & technology*, 40, 2653–2662, 2006.
- 35 Vonlanthen, M., Allegrini, J., and Carmeliet, J.: Assessment of a one-way nesting procedure for obstacle resolved large eddy simulation of the ABL, *Computers and Fluids*, 140, 136–147, doi:10.1016/j.compfluid.2016.09.016, <https://doi.org/10.1016/j.compfluid.2016.09.016>, 2016.

- Vonlanthen, M., Allegrini, J., and Carmeliet, J.: Multiscale interaction between a cluster of buildings and the abl developing over a real terrain, *Urban Climate*, 20, 1–19, doi:10.1016/j.uclim.2017.02.009, <https://doi.org/10.1016/j.uclim.2017.02.009>, 2017.
- Weil, J., Sullivan, P., and Moeng, C.: The use of large-eddy simulations in Lagrangian particle dispersion models, *Journal of the Atmospheric Sciences*, 61, 2877–2887, 2004.
- 5 Wicker, L. J. and Skamarock, W. C.: Time-splitting methods for elastic models using forward time schemes, *Mon. Wea. Rev.*, 130, 2088–2097, 2002.
- Wiersema, D. J., Lundquist, K. A., and Chow, F. K.: Mesoscale to microscale simulations over complex terrain with the immersed boundary method in the Weather Research and Forecasting Model, *Monthly Weather Review*, 148, 577–595, 2020.
- Williamson, J. H.: Low-storage Runge-Kutta schemes, *J Comput. Phys.*, 35, 48–56, 1980.
- 10 Xie, Z. and Castro, I.: LES and RANS for turbulent flows over arrays of wall-mounted obstacles, *Flow, Turbulence and Combustion*, 76, 291–312, 2006.
- Zhou, B., Xue, M., and Zhu, K.: A grid-refinement-based approach for modelling the convective boundary layer in the gray zone: Algorithm implementation and testing, *Journal of the Atmospheric Sciences*, 75, 1143–1161, 2018.
- Zhu, P., Albrecht, B. A., Ghate, V. P., and Zhu, Z.: Multiple-scale simulations of stratocumulus clouds, *Journal of Geophysical Research: Atmospheres*, 115, 2010.
- 15 Zängl, G., Reinert, D., Rípodas, P., and Baldauf, M.: The ICON (ICOsahedral Non-hydrostatic) modelling framework of DWD and MPI-M: Description of the non-hydrostatic dynamical core, *Quarterly Journal of the Royal Meteorological Society*, 141, 563–579, doi:10.1002/qj.2378, 2015.