

TOSCA – an open-source, finite-volume, large-eddy simulation (LES) environment for wind farm flows

Stipa, S.; Ajay, Arjun ; Allaerts, D.J.N.; Brinkerhoff, Joshua

DOI

[10.5194/wes-9-297-2024](https://doi.org/10.5194/wes-9-297-2024)

Publication date

2024

Document Version

Final published version

Published in

Wind Energy Science

Citation (APA)

Stipa, S., Ajay, A., Allaerts, D. J. N., & Brinkerhoff, J. (2024). TOSCA – an open-source, finite-volume, large-eddy simulation (LES) environment for wind farm flows. *Wind Energy Science*, 9(2), 297–320. <https://doi.org/10.5194/wes-9-297-2024>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.



TOSCA – an open-source, finite-volume, large-eddy simulation (LES) environment for wind farm flows

Sebastiano Stipa¹, Arjun Ajay¹, Dries Allaerts², and Joshua Brinkerhoff¹

¹School of Engineering, University of British Columbia–Okanagan, Kelowna, Canada

²Aerospace Engineering, TU Delft, Delft, the Netherlands

Correspondence: Sebastiano Stipa (sebstipa@mail.ubc.ca)

Received: 18 April 2023 – Discussion started: 11 May 2023

Revised: 15 November 2023 – Accepted: 5 December 2023 – Published: 5 February 2024

Abstract. The growing number and growing size of wind energy projects coupled with the rapid growth in high-performance computing technology are driving researchers toward conducting large-scale simulations of the flow field surrounding entire wind farms. This requires highly parallel-efficient tools, given the large number of degrees of freedom involved in such simulations, and yields valuable insights into farm-scale physical phenomena, such as gravity wave interaction with the wind farm and farm–farm wake interactions. In the current study, we introduce the open-source, finite-volume, large-eddy simulation (LES) code TOSCA (Toolbox fOr Stratified Convective Atmospheres) and demonstrate its capabilities by simulating the flow around a finite-size wind farm immersed in a shallow, conventionally neutral boundary layer (CNBL), ultimately assessing gravity-wave-induced blockage effects. Turbulent inflow conditions are generated using a new hybrid off-line–concurrent–precursor method. Velocity is forced with a novel pressure controller that allows us to prescribe a desired average hub-height wind speed while avoiding inertial oscillations above the atmospheric boundary layer (ABL) caused by the Coriolis force, a known problem in wind farm LES studies. Moreover, to eliminate the dependency of the potential-temperature profile evolution on the code architecture observed in previous studies, we introduce a method that allows us to maintain the mean potential-temperature profile constant throughout the precursor simulation. Furthermore, we highlight that different codes do not predict the same velocity inside the boundary layer under geostrophic forcing owing to their intrinsically different numerical dissipation. The proposed methodology allows us to reduce such spread by ensuring that inflow conditions produced from different codes feature the same hub wind and thermal stratification, regardless of the adopted precursor run time. Finally, validation of actuator line and disk models, CNBL evolution, and velocity profiles inside a periodic wind farm is also presented to assess TOSCA’s ability to model large-scale wind farm flows accurately and with high parallel efficiency.

1 Introduction

In 2018, Ørsted, a leading company in developing, constructing and operating offshore and onshore wind farms, concluded a project aimed at understanding the limits of models and processes used for wind energy forecasts. The investigation pointed out that blockage and wake effects are currently neglected and underestimated, respectively, when performing wind power predictions (Ørsted, 2019). Blockage, also referred to as turbine or farm induction (Bleeg et al., 2018), is defined as the wind slowdown approaching the wind farm.

On the other hand, wake losses are characterized by a power production deficit by waked turbines and are claimed to be underestimated both inside and especially between neighboring sites (Pedersen et al., 2022). While wind farm losses arising from individual turbine wakes have been the subject of extensive research, farm–farm wake effects have gained importance only recently (Lundquist et al., 2019; Ahsbabs et al., 2020; Schneemann et al., 2020). Specifically, as more plants are constructed in the proximity of pre-existing ones, the evolution of neighboring farm wakes is an increasingly

important aspect to account for and model (Nygaard et al., 2020).

Turbine-level induction has been researched for many years (Trolborg and Meyer Forsting, 2017; Gribben and Hawkes, 2019; Branlard and Gaunaa, 2014; Branlard et al., 2020), and extensions to wind turbine clusters have been attempted using a linear superposition of individual effects (Branlard and Meyer Forsting, 2020; Segalini, 2021). However, recent studies suggest that this could underestimate – if not totally misrepresent – wind-farm-level blockage, which is heavily influenced by the mutual interaction between the wind farm and the density-stratified atmospheric boundary layer (ABL) (Smith, 2010; Wu and Porté-Agel, 2017; Allaerts and Meyers, 2017, 2018, 2019; Centurelli et al., 2021). In fact, the flow deceleration in the wind farm displaces the capping inversion layer, and interfacial waves are formed. Subsequently, their energy is transported vertically and horizontally by atmospheric internal gravity waves. This mechanism triggers pressure disturbances inside the boundary layer, altering the velocity field around the wind farm.

In industry, annual energy captures are made using low-cost but fast, often analytical, reduced-order wake models (Jensen, 1983; Ainslie, 1988; Larsen, 1988; Bastankhah and Porté-Agel, 2014; Niayifar and Porté-Agel, 2016), aimed at capturing the gross aerodynamic processes within the farm. While they have been used effectively for hundreds of wind energy projects, the majority of these models currently struggle in accurately reproducing wind farm blockage and farm-farm wake interactions (Nygaard et al., 2022). This is classified as an industry-wide issue, as over-predicting annual energy production can have a negative impact on all companies' financial estimates.

Reduced-order models need to be thoroughly validated, but comprehensive observation datasets are difficult to obtain. Numerical analyses, in particular large-eddy simulations (LESs), are able to provide such data, together with valuable insight into the physical processes. LES resolves the largest and most energetic turbulent eddies, while the smallest ones are modeled. Nevertheless, LES of large wind farms in the atmospheric boundary layer (ABL) is extremely challenging, given the breadth of scales involved, spanning from resolved turbulence eddies of a few meters to gravity waves characterized by wavelengths of several kilometers. To successfully tackle these problems, LES solvers must possess good parallel efficiency and an optimized code input–output (I/O), allowing them to operate on modern high-performance computing (HPC) architectures. Additionally, many numerical aspects have to be carefully treated, such as the coupling with mesoscale models (Haupt et al., 2023), handling wave reflections produced by the domain boundaries or the generation of a suitable time-resolved turbulent inflow (Lanzilao and Meyers, 2022a). The last two tasks can be achieved simultaneously using the concurrent-precursor method, where a simulation without wind turbines (precursor) is advanced in sync with the wind farm simulation (successor). The latter

features a fringe region, where body forces are used to damp gravity wave reflections and to restore the desired turbulent inflow. At each time step, such body forces are calculated based on the concurrent-precursor instantaneous fields, leading to the precursor and successor solutions matching at the fringe region exit. More details on precursor techniques are given in Sect. 2.4, where our new hybrid method is also described.

Several LES codes have been developed by the research community so far (see Breton et al., 2017, for a review), among which only a few can effectively tackle the above-mentioned application. The KU Leuven code SP-Wind, for example, has been successfully used for finite wind farm simulations capturing gravity wave effects (Lanzilao and Meyers, 2022b) but unfortunately is not open-source. Conversely, open-source tools, such as the Parallelized Large-Eddy Simulation Model (PALM; Maronga et al., 2015), developed by the Institute of Meteorology and Climatology at Leibniz Universität Hannover (Germany), or SOWFA (the Simulator fOr Wind Farm Applications), maintained by the National Renewable Energy Laboratory (NREL), do not implement the concurrent-precursor method, making it difficult to properly simulate gravity wave effects at the same time as avoiding inlet–outlet reflections. In addition, although SOWFA has been used in several research studies in the last decade (Churchfield et al., 2012b, a; Fleming et al., 2014; Johlas et al., 2021, to name a few), it is based on OpenFOAM 6 (OpenCFD, 2018), a general-purpose set of libraries that are not specifically designed to efficiently run at scale. SOWFA's greatest limitation is its massive generation of output files. In particular, a directory containing all simulation fields is generated for each processor at run time. When dealing with thousands of processors, the number of produced files can easily saturate, within a few checkpoint iterations, the maximum file count in many HPC architectures. To address these shortcomings, NREL has started the ExaWind project (Min et al., 2022), but the latter is not yet at a final production stage, and it does not feature the concurrent-precursor technique or a method to model complex terrains.

For the aforementioned reasons, we have developed an open-source, finite-volume framework that is tailored for large-scale studies of wind-farm-induced gravity waves and cluster wake–atmosphere interaction, with the objective of *gaining sufficient understanding of the physics of atmospheric flow within and around wind plants*, a grand challenge of modern wind energy according to Shaw et al. (2022). The new framework is called TOSCA (Toolbox fOr Stratified Convective Atmospheres) and exploits state-of-the-art parallel libraries, such as Open MPI (Gabriel et al., 2004), PETSc (Balay et al., 2022), HYPRE (Falgout and Yang, 2002) and HDF5 (The HDF Group, 2006), for the parallel solution of partial differential equations and handling of intense I/O operations. TOSCA is specifically designed to enable LESs of large finite wind farms. Wind turbines can be modeled using the actuator line (Sørensen and Shen, 2002)

and the actuator disk (Jimenez et al., 2007, 2008) models. As inlet–outlet boundary conditions produce a consistent and undesirable reflection of atmospheric gravity waves, we introduce a hybrid off-line precursor–concurrent-precursor methodology which, coupled with periodic boundary conditions, limits artificial wave reflections while simultaneously reducing the computational cost associated with initializing the turbulent precursor. The concurrent-precursor method (Inoue et al., 2014) is to our knowledge not available in other finite-volume solvers, though it is extensively used in pseudo-spectral methods (Wu and Porté-Agel, 2017; Allaerts and Meyers, 2017, 2018). For this reason, gravity wave studies to date have been only performed using the latter discretization technique, which does not allow for grid refinement in the pseudo-spectral directions. This forces a uniform grid resolution, leading to high cell counts. Conversely, the finite-volume method allows for grid stretching, enabling us to resolve larger domains with the same number of degrees of freedom while providing greater geometrical flexibility. Finally, TOSCA also features a sharp-interface immersed boundary method (IBM) based on Haji Mohammadi et al. (2019) that allows us to simulate moving objects and complex terrain features, but its validation will be covered in a follow-up paper.

The present paper is organized as follows. First, in Sect. 2, we describe the developed LES framework. Next, Sect. 3 presents comparisons with existing numerical and experimental studies to validate TOSCA’s actuator models, the evolution of thermally stratified ABLs and wake interactions inside a periodic wind farm in neutral conditions. In Sect. 4, we compare results obtained from CNBL simulations using the newly developed velocity and temperature controlling techniques against the commonly used geostrophic forcing combined with a wind angle controller. In Sect. 5, we present the simulated flow field around a reference 100-turbine finite wind farm immersed in a turbulent CNBL, highlighting TOSCA’s ability to accurately predict gravity wave blockage effects. Finally, conclusions are outlined in Sect. 6.

2 Methodology

TOSCA is a finite-volume code, formulated in generalized curvilinear coordinates, allowing it to also take as input non-Cartesian structured meshes. The present section is organized as follows. The governing equations in Cartesian coordinates are reported in Sect. 2.1, while actuator models used to represent wind turbines in the domain are described in Sect. 2.2.

To provide a better flow of the paper, the numerical method, the governing equations in curvilinear coordinates actually solved in TOSCA and a brief overview of generalized curvilinear coordinates are reported in Appendix A, while the LES turbulence model in the curvilinear frame is detailed in Appendix B.

An overview of TOSCA’s parallel efficiency is given in Appendix D, where we analyze the time per iteration with an increasing number of nodes and mesh elements on the Niagara (Niagara, 2018; Loken et al., 2010; Ponce et al., 2019) high-performance computer at the SciNet HPC Consortium. In addition, TOSCA has been used to run finite wind farm simulations on the whole Niagara cluster (2024 nodes, 40 cores per node) and on all Cascade nodes of the UBC ARC (Sockeye, 2023; 170 nodes, 40 cores per node) cluster, demonstrating its capability to handle massively parallel computations.

In order to run ABL simulations, we developed a novel methodology, described in Sect. 2.3, that enforces a desired hub-height wind speed while simultaneously avoiding inertial oscillations produced by the Coriolis force above the boundary layer. In addition, we show that disagreement exists between different computational fluid dynamics (CFD) codes in predicting the final mean potential-temperature profile inside the boundary layer. In this regard, we propose the use of a mean temperature controller which maintains a prescribed average potential-temperature profile, harmonizing the comparison of simulation results in future studies. Finally, Sect. 2.4 details our hybrid off-line–concurrent-precursor methodology, which saves computational resources when performing the turbulence initialization in the precursor phase.

2.1 Governing equations

Governing equations correspond to mass and momentum conservation for an incompressible flow with Coriolis forces and Boussinesq approximation for the buoyancy term. The latter is calculated using the modified density ρ_k , evaluated by solving a transport equation for the potential temperature. These equations, expressed in Cartesian coordinates using tensor notation, read

$$\frac{\partial u_i}{\partial x_i} = 0, \quad (1)$$

$$\frac{\partial u_i}{\partial t} + \frac{\partial}{\partial x_j} (u_j u_i) = -\frac{1}{\rho_0} \frac{\partial p}{\partial x_i} + \frac{\partial}{\partial x_j} \left[v_{\text{eff}} \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) \right] - \frac{1}{\rho_0} \frac{\partial p_\infty}{\partial x_i} + \frac{\rho_k}{\rho_0} g_i - 2\epsilon_{ijk} \Omega_j u_k + f_i + s_i^v + s_i^h, \quad (2)$$

$$\frac{\partial \theta}{\partial t} + \frac{\partial}{\partial x_j} (u_j \theta) = \frac{\partial}{\partial x_j} \left(\kappa_{\text{eff}} \frac{\partial \theta}{\partial x_j} \right), \quad (3)$$

where u_i is the Cartesian velocity; p/ρ_0 is the kinematic pressure; θ is the potential temperature, defined as $\theta = T(p_0/p)^{R/c_p}$ (T is the absolute temperature, R is the specific gas constant for dry air, c_p is the specific heat at constant pressure and p_0 is the reference pressure); g_i is the gravitational acceleration vector; and Ω_j is the rotation rate vector at an arbitrary location on the planetary surface (defined as $\omega \cos \phi \hat{y} + \omega \sin \phi \hat{z}$, where ϕ is the latitude, in a local reference frame having \hat{z} aligned and opposite to the gravitational

acceleration vector, \hat{x} tangent to Earth's parallels, and \hat{y} such that the frame is right-handed). Source terms f_i , s_i^v and s_i^h are body forces introduced by turbines and by vertical and horizontal damping regions, respectively. Moreover, the modified density ρ_k is defined as

$$\frac{\rho_k}{\rho_0} = 1 - \left(\frac{\theta - \theta_0}{\theta_0} \right), \quad (4)$$

where θ_0 is a reference potential temperature, chosen as the ground temperature. Parameters ν_{eff} and κ_{eff} are the effective viscosity and thermal diffusivity, respectively. The former is the sum of the kinematic viscosity ν and the sub-grid-scale viscosity ν_t , while the latter is sum between the thermal diffusivity $\kappa = \nu/Pr$ and the turbulent thermal diffusivity κ_t . Both ν_t and κ_t are defined in Appendix B, while the Prandtl number Pr is set to 0.7 in all simulations. The third term on the right-hand side of Eq. (2) is a uniform horizontal pressure gradient that balances turbulent stresses and the Coriolis force, allowing the boundary layer to reach a statistically steady state. This term is commonly referred to as the velocity controller, and it is explained in Sect. 2.3.1.

2.2 Actuator models

To represent wind turbines, different models have been implemented. In TOSCA, they are referred to as the actuator line model (ALM), actuator disk model (ADM) and uniform actuator disk (UAD) model. Following Sørensen and Shen (2002), Sørensen et al. (2015), Porté-Agel et al. (2010), and Jimenez et al. (2007), the first two models require detailed blade information (i.e., airfoils, twist and chord), while the UAD only requires the turbine thrust coefficient and general rotor information such as diameter and hub height (Jimenez et al., 2007, 2008). The idea behind actuator models is to represent the wind turbine as a distribution of points, each associated with a Lagrangian force. For the UAD model, the sum of forces from all points must be equal to the total wind turbine thrust, while the ALM and ADM in TOSCA additionally include rotor torque, as they also model blade rotation. Once the Lagrangian force at each point has been calculated, it is distributed to the surrounding mesh cells through a projection function. In TOSCA, a classical isotropic Gaussian projection is used, namely

$$g(x, y, z) = \frac{1}{\epsilon \pi^{3/2}} \exp\left(-\frac{(x-x_0)^2 + (y-y_0)^2 + (z-z_0)^2}{\epsilon^2}\right), \quad (5)$$

where (x_0, y_0, z_0) is the position of the actuator point and ϵ is a tunable parameter, corresponding to the standard deviation of the Gaussian projection function. Note that while the projection function should integrate to unity to preserve each point force, this is never exactly possible, and the projection distance is cut when 99% of the Gaussian volume

has been taken into account. Moreover, the Gaussian width-to-grid size ratio should be larger than 2 in order to avoid large projection errors and numerical instabilities (Martínez-Tossas et al., 2015).

The definition of the turbine point mesh and the evaluation of the point force are different depending on the specific model. In the ALM, each rotor blade is represented by a line of points, which are physically rotated at each iteration, making it an unsteady model. In the ADM and UAD model, the number of points in the azimuthal (tangential) direction is not equal to the number of blades, and it is usually set to a high value. For both the ALM and ADM, the point force is calculated exploiting the blade element momentum theory (BEM; see Glauert, 1935). First, the radially varying velocity is estimated at each point, using information from the CFD mesh and the wind turbine angular velocity. This is known as velocity sampling, and different methods have been proposed (Churchfield et al., 2017). TOSCA samples the velocity at the actuator point, using nearest-neighbor interpolation from the closest mesh cell. Next, the velocity magnitude and angle of attack are given as input to appropriate airfoil tables, which return lift and drag at the point location. Various airfoil tables are used along the blade span, as does the operating Reynolds number. Lift and drag at each actuator point are then distributed to the surrounding CFD cells by convolution with the projection function. Conversely, for the UAD model, the blade loading is uniform, and the force is calculated by dimensionalizing the turbine thrust coefficient with the freestream velocity and the portion of rotor area belonging to each actuator point. In waked conditions, the concept of freestream velocity is not well defined. Hence, a common practice is to first average the wind velocity on the rotor disk and then use the momentum theory to infer the corresponding freestream velocity (Meyers and Meneveau, 2010). In our framework, since the ADM and ALM also account for blade rotation, they can be coupled with a rotor inertia and control system dynamics solver (pitch and angular velocity controllers), while nacelle yaw can be applied to any of the three models.

2.3 Controllers

This section reviews the current state of the art for velocity controllers in precursor simulations and presents a novel technique, which we refer to as geostrophic damping, which allows control of the hub-height velocity while avoiding inertial oscillations generated by the Coriolis force. Moreover, a simple temperature controller is also presented that maintains a constant average potential-temperature profile throughout the precursor simulation.

2.3.1 Velocity controller

In the precursor simulation, the flow is usually driven by a uniform horizontal pressure gradient, which is related to the geostrophic wind components by the geostrophic balance at equilibrium:

$$\frac{1}{\rho_0} \frac{\partial p_\infty}{\partial x} = f_c V_G \frac{1}{\rho_0} \frac{\partial p_\infty}{\partial y} = -f_c U_G, \quad (6)$$

where $f_c = 2\Omega_z$ is known as the Coriolis parameter. Using the above equations to prescribe the driving pressure gradient does not give any control over velocity magnitude and direction at the wind turbine hub height. In fact, the latter will be a result of the turbulent stresses inside the boundary layer, which are not known a priori. However, being able to control these parameters is convenient in wind farm simulations as it allows the operation point of the turbines to be easily prescribed. To this end, Sescu and Meneveau (2014) and Allaerts and Meyers (2015) developed and tuned an algorithm that slowly rotates the flow in the domain, allowing us to control the wind direction at a specified height. Later, Stieren et al. (2021) used the same approach to impose dynamic wind direction changes. Besides the driving pressure gradient, evaluated using Eq. (6), the additional cross product $-\epsilon_{ijk}\omega_i u_j \hat{x}_k$ is added to the momentum equation's right-hand side, where the angular frequency ω_i is calculated based on the angle difference at the reference height (see Allaerts and Meyers, 2015, for details on this procedure). Such a method, which we will refer to as the geostrophic controller, does not entirely solve the issue, as velocity magnitude at the hub height is still unknown a priori. Nevertheless, a different approach exists, available for example in SOWFA, which allows us to prescribe both magnitude and direction at a specified height h_{ref} . In particular, given a desired velocity $u_{\text{ref},i}$, which should be maintained at h_{ref} , an error vector can be defined as the difference between the reference wind and the velocity sampled at the reference height, averaged over the homogeneous directions. At this point, a proportional–integral controller can be used to evaluate the driving pressure gradient (i.e., the third term on the right-hand-side of Eq. (2) such that the desired speed and angle are maintained at h_{ref} . This approach will be referred to as the pressure controller. In TOSCA both controller methods are implemented, and the driving pressure gradient in the second type of controller is evaluated as

$$\frac{1}{\rho_0} \frac{\partial p_\infty}{\partial x_i} = r (\alpha e_{P,i} + (1 - \alpha) e_{I,i}^n), \quad (7)$$

$$e_{P,i} = (u_{\text{ref},i} - \langle u_i(h_{\text{ref}}) \rangle_{xy}) / \Delta t, \quad (8)$$

$$e_{I,i}^n = (1 - \Delta t/T) e_{I,i}^{n-1} + (\Delta t/T) e_{P,i}, \quad (9)$$

where subscript i refers to the i th component, $e_{P,i}$ is the proportional error, $e_{I,i}^n$ is the integral error evaluated at time step n , r is a relaxation factor, α is the proportional fraction of the controlling action, T is the time filter for the integral error,

Δt is the time step size, and $\langle \cdot \rangle_{xy}$ denotes a spatial average along the homogeneous directions x and y . In the present study, we set $r = 0.7$, $\alpha = 0.8$ and $T = 2$ h.

On one hand, the pressure controller is more convenient for wind turbine simulations, as hub wind and direction can be directly specified. However, unlike the geostrophic controller, it does not provide knowledge of the geostrophic wind a priori, making it impossible to initialize the flow such that Eq. (6) is satisfied. An inconsistency in the initial condition produces inertial oscillations above the boundary layer, as the initial wind speed aloft differs from its equilibrium geostrophic value. This can be easily verified by noting that the unsteady form of Eq. (6) (see for example Stull, 2016), namely

$$\begin{cases} \frac{\partial u}{\partial t} + f_c(V_G - v) = 0 \\ \frac{\partial v}{\partial t} - f_c(U_G - u) = 0, \end{cases} \quad (10)$$

represents an undamped linear oscillator with angular frequency f_c . In particular, if $v \neq V_G$ or $u \neq U_G$, at any point during the simulation, inertial oscillations will be produced. In some wind energy applications, for example, when studying the formation of atmospheric gravity waves above the boundary layer, the physics of the problem strongly depends on the magnitude of the geostrophic wind. In such cases, results would be negatively impacted by these inertial oscillations, whose amplitude depends on the initial condition.

Nevertheless, being able to exactly define the wind speed and direction at a specified height is a desirable property of the pressure controller. For this reason, we developed a new methodology that allows us to remove these inertial oscillations, also enabling the use of the pressure controller in those cases where a steady-state geostrophic wind is preferred. First, we note that the system of Eq. (10) can be damped by introducing an additional term as follows:

$$\begin{cases} \frac{\partial u}{\partial t} + 2\alpha f_c(u - U_G) + f_c(V_G - v) = 0 \\ \frac{\partial v}{\partial t} + 2\alpha f_c(v - V_G) - f_c(U_G - u) = 0, \end{cases} \quad (11)$$

where the coefficient α determines if the system is over-damped ($\alpha > 1$), under-damped ($\alpha < 1$) or critically damped $\alpha = 1$. With some manipulation, Eq. (11) can be rewritten as

$$\begin{cases} \frac{\partial^2 u}{\partial t^2} + 2\alpha f_c \frac{\partial u}{\partial t} + 2\alpha f_c^2(v - V_G) + f_c^2(u - U_G) = 0 \\ \frac{\partial^2 v}{\partial t^2} + 2\alpha f_c \frac{\partial v}{\partial t} - 2\alpha f_c^2(u - U_G) + f_c^2(v - V_G) = 0. \end{cases} \quad (12)$$

These equations slightly differ from a conventional spring–mass–damper system in the additional coupling terms $2\alpha f_c^2(v - V_G)$ and $2\alpha f_c^2(u - U_G)$.

We observed that the presence of these terms enhances the damping action, halving the exponent of the decay rate, characterized by an e -folding time of $1/(2\alpha f_c)$. In order for the oscillation amplitude to reach less than 3% of the initial value, a damping time $T_{3\%} = \ln(100/3)/(2\alpha f_c)$ is necessary.

Note that, in order for the damping term to be evaluated, knowledge about the geostrophic wind components is still required. We deduce U_G and V_G from the driving pressure gradients imposed by the pressure controller (Eqs. 7, 8, 9) by means of the definition of the geostrophic wind speed (i.e., the geostrophic balance given by Eq. 6). In addition, we filter the obtained geostrophic components using a filter constant of $0.2\pi/f_c$ (corresponding to 1/10 of the inertial oscillation period). Finally, we highlight that the damping action should start after the boundary layer has become fully developed, as the pressure gradient prescribed by the controller depends on the turbulent stresses if h_{ref} is inside the ABL. In our simulations, we start the damping action after almost one inertial period ($T_D \approx 2\pi/f_c$), and we maintain the damping active for at least a time equal to $T_3\%$. In order for this geostrophic damping method not to affect the velocity inside the ABL, we smoothly bring the damping to zero below a certain height by multiplying the damping terms with the following function:

$$f_d = \frac{1}{2} \left[1 + \tanh \left(\frac{7(h - H_d)}{\Delta_d} \right) \right], \quad (13)$$

where H_d is the height where the damping has halved its strength. If the simulation models a capping inversion layer, we set $H_d = H$, where H is the capping inversion center, and $\Delta_d = \Delta$, where Δ is set as the capping inversion width.

2.3.2 Temperature controller

When running precursor CNBL simulations, the predicted ABL height, as well as the final value of potential temperature at the ground, depends on the mixing history experienced inside the boundary layer. This is in turn affected by the specific LES setup and the type of discretization used. Moreover, the impact of such code details is made even more noticeable by the fact that these simulations usually run for a very long time (of the order of $2\pi/f_c$). For example, SP-Wind, which employs a pseudo-spectral discretization in the horizontal directions and an energy-conservative fourth-order advection scheme in the vertical, predicts less mixing than other pseudo-spectral codes, such as NCAR-LES (Pedersen et al., 2014) or WiRE-LES (Abkar, M. and Porté-Agel, 2013), which use for example a second-order central scheme in the vertical direction (this can be appreciated in Fig. 3). Besides, in finite-volume codes, like TOSCA or SOWFA, an upwind-biased advection scheme is usually preferred, as it stabilizes the numerical method but does not allow us to conserve mechanical energy. These considerations pose comparison issues among different codes, as their differences will have an impact for example on the final ABL height, inversion thickness and potential-temperature jump and, in general, on the heating history of the boundary layer, ultimately affecting the successor solution, in which wind turbines are present. In particular, if a CNBL precursor is run with a certain initial potential-temperature profile, this will evolve differently based on both the adopted simulation framework and

the length of the precursor run, determining a discrepancy in the initial condition of the wind farm simulation.

In the present work, we propose to apply a potential-temperature controller in the precursor simulation so that the successor can be exactly run with the intended temperature profile. This could be beneficial for example when making comparisons between different codes, as it ensures that the successor background potential-temperature profile matches the precursor initial condition. In particular, we apply the following height-dependent source term on the right-hand side of Eq. (3):

$$s_\theta(h) = r \frac{\bar{\theta}(h) - \langle \theta(h) \rangle_{xy}}{\Delta t}, \quad (14)$$

where $\bar{\theta}(h)$ is the desired vertical potential-temperature profile, taken as the initial value of $\langle \theta(h) \rangle_{xy}$, which is average of the potential temperature along the homogeneous directions at a given height. The parameter r is a relaxation coefficient that we set to 0.7. Note that a very similar method was used by Allaerts et al. (2020, 2023) to drive LESs using mesoscale models or observations.

2.4 Hybrid off-line–concurrent precursor

Wind turbine wake recovery and thus power production are greatly influenced by background atmospheric turbulence. As a consequence, prescribing a physical turbulent inflow is necessary if real wind turbine operation is to be simulated. A commonly used approach is the so-called precursor–successor method (Churchfield et al., 2012b, a), where a first simulation of the sole ABL, without wind turbines, is run until turbulence reaches steady-state statistics. After this first phase, the latter is further progressed, and velocity and potential temperature are saved on a plane parallel to the inlet boundary, at each iteration, forming the inflow database. At this point, the simulation with the wind turbines (successor) is started and, at each time step, the inflow boundary condition is interpolated from the saved slices at the two closest times in the database so that precursor and successor time steps can take different values. The above methodology implies no periodicity of the domain in the streamwise direction and has proved to work extremely well for isolated turbine simulations, cases where the ABL height is not perturbed by objects located below it or in the absence of thermal stratification. On the contrary, when thermal stratification is present, atmospheric gravity waves can be triggered above the ABL, and a careful design of the simulation should prevent such waves from being reflected by the physical boundaries. In particular, the LES setup should be equipped with damping regions at the top, inlet and outlet, or if streamwise periodic boundary conditions are used, only one damping region in the streamwise direction is then required (Calaf et al., 2010; Inoue et al., 2014; Allaerts and Meyers, 2017). The latter, also known as the fringe region,

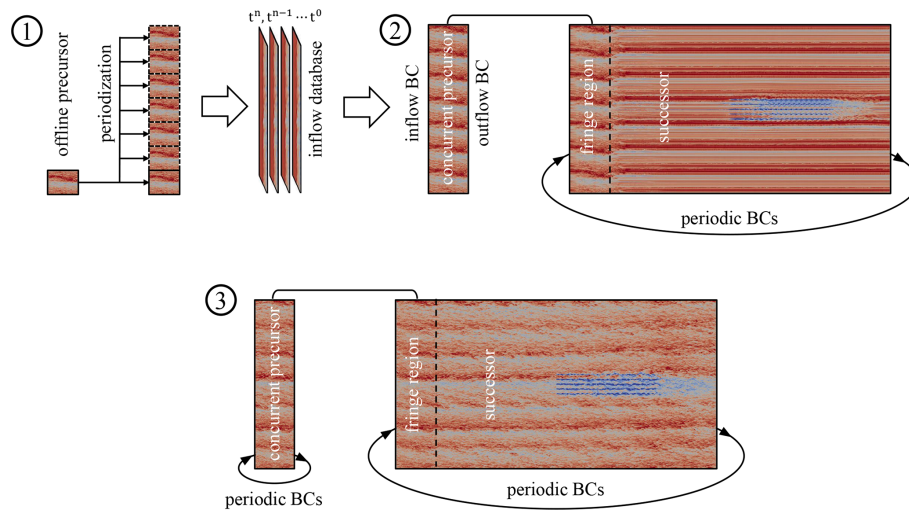


Figure 1. Sketch of the hybrid off-line–concurrent-precursor method, including definition of boundary conditions (BCs).

must ensure that turbine wakes are prevented from being re-introduced into the domain by periodic boundaries and that an unperturbed turbulent inflow is reached again at the fringe exit. To achieve this, the desired flow that is used to compute the damping term should ideally contain time-resolved turbulent structures at every cell located in the fringe. The precursor is then advanced in sync with the successor, in a domain larger than or equal to the fringe region, so that velocity and temperature fields are available, at each time step and spatial location, to compute damping sources.

In order to spin up the wind farm simulation, we developed a three-step procedure (see Fig. 1), where both off-line and concurrent-precursor methodologies are adopted. We believe that the latter method is necessary when dealing with wind-farm-induced gravity waves, as it allows us to avoid wave reflections while prescribing a time-resolved turbulent inflow at the same time. Since the concurrent-precursor domain should coincide with the successor – whose size is dictated by the gravity waves and wind farm – in both the spanwise and the vertical directions, it is usually oversized from the turbulence generation point of view. This leads to a considerable quantity of computational resources being consumed when starting up the unperturbed ABL. In fact, domains of a much smaller size are used in the literature when the sole ABL is of interest or when the inflow data are generated using the off-line precursor technique.

In TOSCA, we exploit the flexibility of the finite-volume formulation by combining the two techniques. In particular, we initialize the ABL on what we refer to as the off-line precursor domain, which can be arbitrarily defined in both the streamwise and the vertical directions. The only requirement is that the spanwise size of the successor is an exact multiple of the off-line precursor domain. When turbulence has reached a statistically steady state, we save flow slices of velocity and potential temperature from this domain into a so-

called inflow database. After this first phase, the concurrent-precursor and successor simulations are started, and streamwise inflow–outflow boundary conditions are used in the former for one flow-through time. Inflow slices from the inflow database are periodized in the spanwise direction, while extrapolation is performed in the vertical direction. We note that it is extremely important that the flow above the inversion layer does not contain any periodic variations in time, as this would be noticed in the successor, at streamwise intervals equal to the concurrent-precursor domain length. For this reason, we average the off-line precursor data at the 10 highest cells and slowly merge the instantaneous data to this average across such an interval. This removes even the smallest periodic content in the flow above the boundary layer, which is now characterized by a truly constant geostrophic wind. For instance, we weight the average and instantaneous velocities using two hyperbolic weighting functions (Eq. 13 is used for both, but a minus sign after unity is applied for the instantaneous velocity). H_d is set to the height of the fifth cell center from the off-line precursor top boundary, while Δ_d is equal to the width of the 10 highest cells.

After the concurrent precursor has run for one flow-through time, the turbulent inflow has reached the outlet, and streamwise boundary conditions are switched to periodic. At this point, the simulation is self-sustained, and we run the precursor and successor simultaneously for one successor flow-through time so that gravity waves and wind turbine wakes are formed. For the simulation presented in Sect. 5, we ran the off-line precursor for 10^5 s, the concurrent-precursor spin-up phase where we used inflow–outflow boundary conditions for 600 s, and the overall successor spin-up for 5000 s (note that this phase can start in parallel with the previous one). Data were gathered from 105 000 to 120 000 s.

The drawback of such a method is that a spanwise periodicity is introduced into the concurrent-precursor and succes-

sor domains. If the Coriolis force is active, this can be broken everywhere except at hub height, where the flow is aligned with the x axis, by setting the off-line precursor and the concurrent-precursor streamwise domain length to a different value. At the hub height, the larger turbulence structures might be locked in position if they span the whole domain length. Although they will eventually disappear, they result in slow convergence of flow averages at the hub height. This issue has been already observed in the past for example by Munters et al. (2016), who proposed to use shifted periodic boundary conditions in the concurrent-precursor domain.

Nevertheless, the proposed hybrid method, sketched in Fig. 1, is very convenient as it allows us to reduce the overall computational cost of the ABL spin-up phase, where wind-farm-induced gravity waves are not yet present. In fact, this initial phase is run on a domain whose size is dictated by the current flow physics, rather than on quantities that will only become relevant at later simulation stages.

3 Validation

In this section, we validate the developed solver using three different benchmark cases. In Sect. 3.1, we simulate an NREL 5 MW reference wind turbine, operating in a uniform inflow equal to 8 m s^{-1} , and compare our results to those of Martínez-Tossas et al. (2015). In Sect. 3.2 we validate the ability of TOSCA to simulate conventionally neutral boundary layer (CNBL) evolution, comparing our results against data from different LES codes reported by Allaerts (2016). Finally, in Sect. 3.3, an infinite wind farm in a turbulent boundary layer without thermal stratification is compared to experimental and numerical data collected by Chamorro and Porté-Agel (2011) and Stevens et al. (2018), respectively.

3.1 Isolated rotor in uniform inflow

In this validation case, we perform two simulations of the NREL 5 MW reference wind turbine (Jonkman et al., 2009) using the ADM and the ALM techniques, with a uniform inflow velocity of 8 m s^{-1} . Periodic boundary conditions are applied at the upper, lower and spanwise boundaries. At the outlet, a zero normal gradient on velocity outflow is specified. The domain is 10 rotor diameters in all directions, with the turbine rotor placed in the geometric center of the domain. The mesh is graded in all directions from a resolution of 16.8 m next to all boundaries to 2.1 m near the wind turbine. In particular, this fine region, where the mesh is uniform, extends 1 diameter upstream of the turbine and 5 diameters downstream. In the other two directions, it extends beyond the edge of the rotor for 1 diameter.

For this case, the standard Smagorinsky model was used, where we set the C_s coefficient of Eq. (B6) to 0.028224 (corresponding to the value of $c_s = \sqrt{C_s} = 0.168$ used in Martínez-Tossas et al., 2015). The ALM has 63 points in the radial direction, while the ADM has 63 and 72 points

Table 1. Wind turbine power as predicted by TOSCA and OpenFOAM for the ADM and ALM model. OpenFOAM data correspond to Martínez-Tossas et al. (2015) with a mesh resolution of 2.1 m and projection width equal to 4.2 m.

	ALM [MW]	ADM [MW]
TOSCA	2.14	2.04
OpenFOAM	2.01	2.08

in the radial and azimuthal directions, respectively. The rotational speed of the wind turbine is set to a constant value of 9.1552 rpm in all cases. The projection ϵ in Eq. (5) is set to 4.2 m. Both simulations are advanced for 300 s, after which data are averaged for the next 300 s. Figure 2 shows the normalized wind speed deficit at 1 and 4 downstream rotor diameters for ALM and ADM simulations performed both with TOSCA and by Martínez-Tossas et al. (2015). An excellent match can be observed at 1 diameter for both models, while at 4 diameters TOSCA predicts slightly higher deficits, especially for the ALM. This difference is due to an earlier breakdown of the blade-tip vortices in the simulations of Martínez-Tossas et al. (2015), which were performed with OpenFOAM. As OpenFOAM is an unstructured code, we believe that non-hexahedral elements arising from the three successive refinement regions produce some small oscillations in the velocity, which is seen by the simulation as added turbulence intensity, determining an earlier breakdown of the blade-tip vortices. This effect is not present in TOSCA, as the mesh is fully structured and smoothly graded from 16.8 to 2.1 m. In van der Laan et al. (2014), the same case is run without a turbulence model using EllipSys3D and SnS, and a higher maximum deficit than both TOSCA and Martínez-Tossas et al. (2015) has been observed at 2.5 diameters. Such a discrepancy between different codes in predicting turbine wake recovery is not observed when a precursor is used to prescribe the inflow, as wake mixing is guided by ABL turbulence instead of numerical oscillations. In Table 1 we report the aerodynamic power produced by the wind turbine as predicted by TOSCA with the two actuator models and that obtained by Martínez-Tossas et al. (2015).

The ADM matches well with results from Martínez-Tossas et al. (2015), while TOSCA's ALM predicts a slightly higher power. The reason for such a difference is presently unknown to the authors.

3.2 CNBL evolution

In this section, we validate TOSCA's ability to perform CNBL simulations by running two cases of a neutral boundary layer, developing against a stable background stratification with lapse rate of 1 and 10 K km^{-1} . In both validation cases, the geostrophic wind is $G = 10 \text{ m s}^{-1}$, the Coriolis parameter is set to $f_c = 10^{-4} \text{ s}^{-1}$, the surface rough-




Figure 2. Normalized wind speed deficit 1 D and 4 D behind a wind turbine represented by (a) ADM and (b) ALM.

ness is $z_0 = 0.01$ m and the reference temperature is 290 K. The numerical domain size is $3 \text{ km} \times 3 \text{ km} \times 2 \text{ km}$, with 256^3 grid points. This corresponds to a grid resolution of approximately $11.7 \text{ m} \times 11.7 \text{ m} \times 7.8 \text{ m}$. Results are compared with results from SP-Wind (Allaerts, 2016), WiRE-LES (Abkar, M. and Porté-Agel, 2013) and NCAR-LES (Pedersen et al., 2014), which all use a pseudo-spectral horizontal discretization.

Velocity and temperature fields are initialized with a constant and linear profile equal to the geostrophic velocity and background stratification, respectively. Furthermore, sinusoidal perturbations are added to the velocity profile, below 100 m, with an amplitude of $0.1G$ and 12 periods in the x and y directions to trigger turbulent fluctuations. Simulations are advanced in time for 24 h, and results are averaged over the last hour. Periodic boundary conditions are applied in the horizontal directions, while a slip boundary condition is applied at the upper boundary. At the ground, the wall shear stress is prescribed through classic Monin–Obukhov similarity laws (Monin and Obukhov, 1954; Paulson, 1970; Etling, 1996), following the approach of Yang et al. (2017) to address the log-layer mismatch. The flow is driven using pressure gradients obtained from Eq. (6). We do not apply any velocity controller in these simulations so that the wind direction inside the boundary layer is free to change while the geostrophic wind remains aligned with the x axis. Figure 3 compares vertical profiles of velocity magnitude, horizontal wind direction, potential temperature and kinematic heat flux obtained from TOSCA, with profiles reported by Allaerts (2016), Abkar, M. and Porté-Agel (2013) and Pedersen et al. (2014). Very good agreement is found in the horizontal wind direction and magnitude. Regarding temperature profiles, TOSCA is more aligned with NCAR-LES and WiRE-LES results, while SP-Wind predicts a slightly lower inversion layer and potential temperature at the ground. This highlights how turbulent mixing is predicted differently by the four codes. Heat flux profiles agree well below the inversion layer, with the exception that TOSCA predicts a more diffused kinematic heat flux profile above the inversion layer

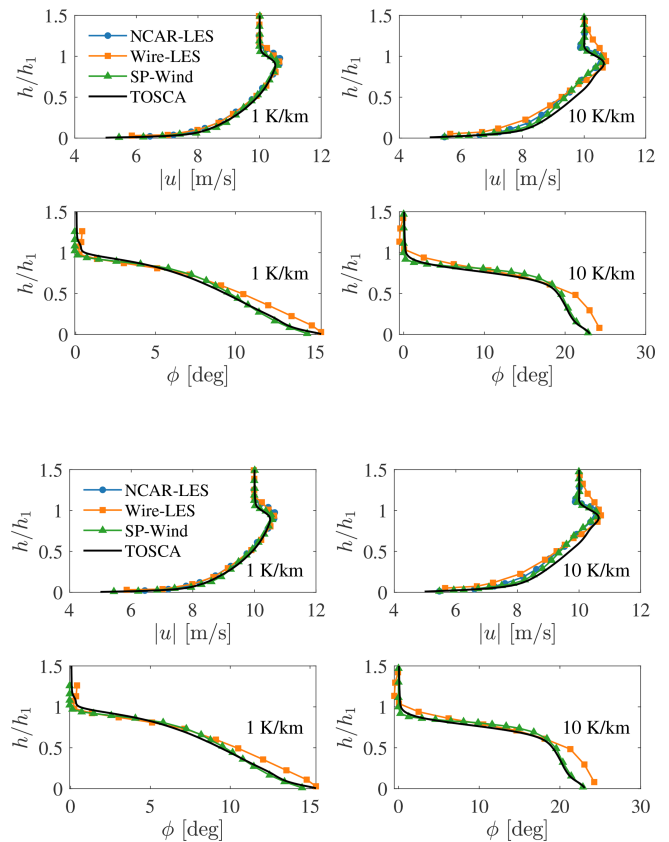


Figure 3. Vertical profiles averaged over the last simulation hour of, from top to bottom, velocity magnitude, wind angle, potential temperature and heat flux. The initial lapse rate is (left) 1 K km^{-1} and (right) 10 K km^{-1} .

for the 10 K km^{-1} case. We do not have a clear explanation for such behavior.

In Table 2 quantitative parameters of the resulting ABLs are reported for the different cases and LES codes. The reference temperature θ_0 , the capping inversion strength $\Delta\theta$ and the inversion width Δh are evaluated by a least-squares fit of the resulting temperature profiles with the model proposed

by Rampanelli and Zardi (2004). The ABL height is taken as the center of the capping inversion layer. Both vertical profiles of Fig. 3 and quantitative ABL parameters reported in Table 2 demonstrate that TOSCA is well aligned with results from Pedersen et al. (2014) and Abkar, M. and Porté-Agel (2013) and is thus capable of conducting CNBL simulations.

In addition, we note that simulations performed using SP-Wind consistently predict less mixing than other codes, which is confirmed by the lower absolute value of the minimum kinematic heat flux q_{\min} . This could be due to the fourth-order energy-conservative scheme, which is adopted in SP-Wind simulations, while other codes employ second- or third-order non-conservative advection schemes.

3.3 Infinite wind farm in neutral conditions

In this section, we run the same infinite wind farm simulation as has been conducted in Stevens et al. (2018), corresponding to the wind tunnel experiments performed by Chamorro and Porté-Agel (2011). The scaled wind farm consists of 30 wind turbines, arranged in an aligned configuration with 3 columns and 10 rows. Spanwise and streamwise spacings are set to $S_y = 4D$ and $S_x = 5D$, respectively, where $D = 0.15$ m is the turbine diameter. The wind farm is made periodic in the spanwise direction by placing turbine columns 1 and 3 at a distance of $S_y/2$ from the lateral boundaries, where periodic boundary conditions are applied. In Stevens et al. (2018), simulations are run with both the ALM and the non-rotating uniform ADM model, which we refer to as the uniform actuator disk model (UADM), and each wind farm row has a different C_t value. For this validation case, we did not attempt to use rotating actuator models (ADM or ALM), as C_t coefficients applied in Stevens et al. (2018) at some rows are higher than the value of $C_{t,\max}$ from their reported BEM calculations. Therefore, since it would not have been possible to match their exact angular velocity for some of the rows, we decided to opt for the UADM, where turbine-specific thrust calculation at the p th disk element is solely based on the thrust coefficient C_t and the freestream velocity U_∞ as

$$\mathbf{f}_p = \frac{1}{2} U_\infty^2 dA_p C_t \hat{\mathbf{e}}_t. \quad (15)$$

In the above expression, dA_p is the disk area associated with the p th actuator disk point; $\hat{\mathbf{e}}_t$ is a vector normal to the rotor disk, pointing in the upstream direction; and U_∞ is evaluated from the average disk velocity U_{disk} exploiting the momentum theory. For instance, $U_\infty = U_{\text{disk}}/(1-a)$, where a is the induction factor, related to the thrust coefficient as $C_t = 4a(1-a)$. Note that Eq. (15) can be rewritten in an equivalent form by using U_{disk} and the disk-based thrust coefficient $C'_t = C_t/(1-a)^2$ in place of U_∞ and C_t . We use the latter formulation in the present case, as Stevens et al. (2018) reported the value of C'_t at each wind farm row.

The wind tunnel model used in the experiment by Chamorro and Porté-Agel (2011) is the GWS/EP-6030 tur-

bine; it has a hub height of 0.125 m, an overhang of 0.03 m, a hub radius of 0.0075 m and a tower diameter of 0.01 m. The tower and nacelle have been modeled following the approach of Stevens et al. (2018), except for the projection function, which is given by Eq. (5). The value of ϵ has been set to 0.02452 for the tower and nacelle and to 0.03515625 for the rotor in order to closely match their approach. The tower is represented by 50 actuator points and is characterized by a drag coefficient of 0.68, while the nacelle consists of a single point where the force is calculated by dimensionalizing a drag coefficient of 4. Moreover, the rotor has been discretized using 20 radial points and 50 azimuthal points. To prescribe a turbulent inflow, Stevens et al. (2018) used the concurrent-precursor technique, as their code is pseudo-spectral. Since TOSCA is a finite-volume code and inflow-outflow conditions can be applied, we opted for the computationally cheaper off-line precursor technique described in Sect. 2.4. The precursor domain is $1.8 \text{ m} \times 1.8 \text{ m} \times 0.675 \text{ m}$, with $129 \times 129 \times 145$ cells in each direction in order to match their cell size. The flow is driven by the pressure controller described in Sect. 2.3.1, with a desired flow velocity u_{ref} of 3 m s^{-1} at the hub height. Potential-temperature stratification is turned off so that the boundary layer height coincides with the domain size in the vertical direction z . We used periodic boundary conditions in the horizontal directions, while at the upper boundary a slip condition is applied. At the ground, we used the same similarity laws as Stevens et al. (2018), with an equivalent roughness height of 0.03 mm.

The precursor is run for 100 s (corresponding to ≈ 160 flow-through times), after which we saved the inflow field at each time step for 300 s. In the successor, we apply the pre-calculated inflow and source terms from the precursor, linearly interpolating from the two closest available times. At the outlet, we use a zero-normal-gradient condition applied to the velocity. The remaining boundaries are treated in the same manner as the precursor. The successor domain is $8.25 \text{ m} \times 1.8 \text{ m} \times 0.675 \text{ m}$, with $588 \times 129 \times 145$ cells in each direction. The first row of the wind farm is located 5 diameters from the inlet boundary, matching the setup of Stevens et al. (2018). The successor is advanced in time for 300 s, and we start gathering data after one flow-through time (≈ 3 s).

In Fig. 4 we show the velocity profiles, averaged across the wind farm columns, for rows 1, 4, 7 and 10, together with experimental data from Chamorro and Porté-Agel (2011) and numerical results from Stevens et al. (2018). As can be noticed, TOSCA matches very well with both numerical and experimental data. In the upper portion of the velocity profile, for increasing wind turbine row, both TOSCA and results from Stevens et al. (2018) predict higher velocities than the experiment. This effect is given by wind farm area blockage in the numerical domain, which causes the flow to accelerate close to the upper boundary in order to conserve mass. In the

Table 2. Quantitative ABL results from the two ABL simulations performed with different codes.

	γ [K km ⁻¹]	θ_0 [K]	$\Delta\theta$ [K]	Δh [m]	H [m]	u^* [m s ⁻¹]	$q_{\min}/10^{-4}$ [K m s ⁻¹]
TOSCA	1	290.38	0.50	213	783	0.34	-7.0
NCAR LES	1	290.36	1.18	342	800	0.37	-5.8
Wire LES	1	290.36	0.54	229	717	0.36	-
SP-Wind	1	290.34	0.41	148	687	0.34	-4.2
TOSCA	10	292.08	2.85	160	429	0.34	-22.5
NCAR LES	10	292.17	2.92	119	439	0.37	-25.5
Wire LES	10	292.28	3.03	210	425	0.35	-
SP-Wind	10	291.72	2.08	97	356	0.34	-13.8

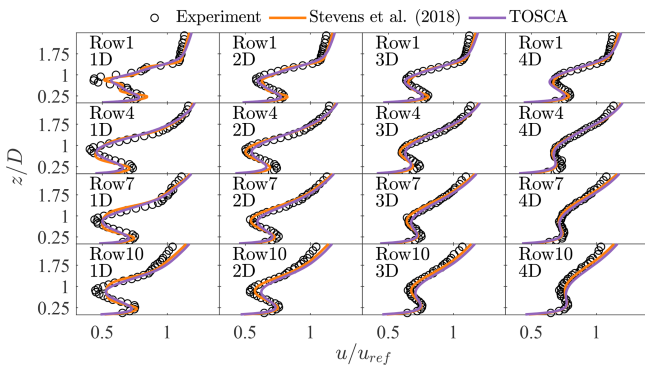


Figure 4. Velocity profiles made non-dimensional with the hub-height velocity for rows 1, 4, 7 and 10, averaged over the three columns. For each row, the wake evolution is reported at 1, 2, 3 and 4 diameters downstream.

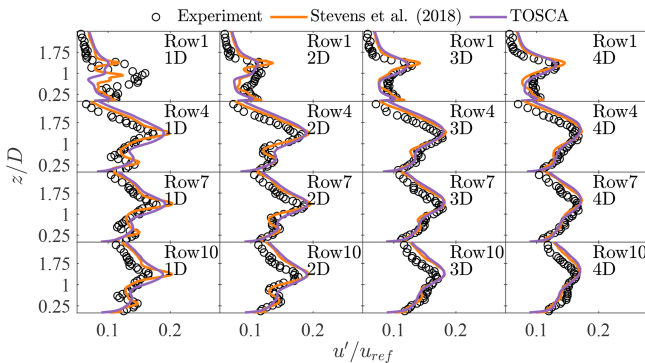


Figure 5. Profiles of streamwise fluctuations non-dimensionalized with hub-height velocity for rows 1, 4, 7 and 10, averaged over the three columns. For each row, the wake evolution is reported at 1, 2, 3 and 4 diameters downstream.

experimental data, this is not observed, as the height of the wind tunnel test section was 1.7 m.

In Fig. 5, profiles of u'/u_{ref} are reported for the same location as Fig. 4. Given the velocity time history at a point, we first evaluate $u'u'$ by averaging the square of the fluctuation history, obtained as the difference between the velocity

signal and its average. Then u' is obtained as the square root of $u'u'$. Results show that TOSCA is well aligned with results from Stevens et al. (2018), both predicting higher fluctuations than experiments in the topmost downwind part of the wind farm for the reason mentioned above. These results demonstrate that TOSCA accurately predicts turbine–wake interactions inside a wind farm, in terms of both the mean and the fluctuations, making it suitable for the simulation of wind turbines immersed in a turbulent boundary layer.

4 CNBL simulations with different controllers

In this section, we present CNBL results obtained using the different velocity and temperature controllers described in Sect. 2.3. In particular, we compare case S2 from Allaerts and Meyers (2017) against results obtained from TOSCA using both pressure and temperature controllers at the same time (case PT), and pressure and geostrophic controllers with no temperature forcing (case P and G, respectively). A summary of the different cases with the relative controllers is given in Table 3. The simulations employ periodic boundary conditions in the horizontal directions and a slip boundary condition at the upper boundary. Classic Monin–Obukhov similarity theory is enforced at the ground.

Following Allaerts and Meyers (2017), the domain size is 9.6 km × 4.8 km × 1.5 km in the streamwise, spanwise and vertical directions, respectively, discretized using 320 × 320 × 300 cells in each direction. Case G is forced with a geostrophic wind speed of 12 m s⁻¹, matching the setup used by Allaerts and Meyers (2017) in case S2. Conversely, in the P and PT cases the pressure controller aims to maintain a wind speed of 10.871 m s⁻¹ at $h_{\text{ref}} = 100$ m, thus matching the hub-height wind speed obtained from case G. The Coriolis parameter f_c is set to 10⁻⁴. In all cases, potential temperature has been initialized using the Rampanelli and Zardi (2004) model, the inputs of which are reported in Table 4.

Table 3. Velocity and potential-temperature controlling strategies for the cases presented in this section. Case S2 corresponds to Allaerts and Meyers (2017) and has the same setup as case G.

	Velocity	Temperature	Geo. damping
G	geostrophic forcing + hub wind angle	off	off
P	pressure forcing based on hub wind	off	on
PT	pressure forcing based on hub wind	mean temperature forcing	on
S2	geostrophic forcing + hub wind angle	off	off

Table 4. Inputs for the Rampanelli and Zardi (2004) model used to initialize the CNBL simulations described in the present section. In the present work, H identifies the capping inversion center, while it refers to the capping inversion base in Allaerts and Meyers (2017).

γ [K km ⁻¹]	θ_0 [K]	$\Delta\theta$ [K]	Δh [m]	H [m]
1	288.15	2.0	100	550




Figure 6. Comparison of results extracted from cases G, P and PT against data from Allaerts and Meyers (2017). Flow statistics from cases G, P and PT are averaged from 92 800 to 100 000 s, while in Allaerts and Meyers (2017) (case S2) data are averaged from 54 000 to 72 000 s.

The model parameter c which determines potential-temperature smearing across the capping inversion is set to 0.33. For the velocity, we use a uniform log law to prescribe the initial condition; namely

$$\begin{cases} u(z) = \frac{u^*}{\kappa} \ln\left(\frac{z}{z_0}\right) & z < H \\ u(z) = \frac{u^*}{\kappa} \ln\left(\frac{H}{z_0}\right) & z \geq H. \end{cases} \quad (16)$$

In case G, where geostrophic damping is not applied, care must be paid in prescribing an initial geostrophic wind consistent with geostrophic forcing in order not to trigger inertial oscillations above the capping inversion.

For this reason, u^* is set to $\kappa G / \ln(H/z_0)$ in case G, while it is calculated as $\kappa u_{\text{ref}} / \ln(h_{\text{ref}}/z_0)$ for cases P and PT. Note that, when the pressure controller is used, inertial oscillations

cannot be avoided since geostrophic wind is not known a priori.

All three simulations are carried out for 10^5 s (≈ 27.8 h), while Allaerts and Meyers (2017) run case S2 for 20 h, gathering statistics from 54 000 to 72 000 s (over the last 5 h of simulated time). Geostrophic damping in cases P and PT starts at $T_D = 6 \times 10^4$ s (this value is close to the oscillation period of the geostrophic wind), and it will be later shown that the wind angle controller of case G stabilizes the wind angle at around 6.5×10^4 s. Hence, we average flow statistics from cases G, P and PT from 92 800 to 100 000 s, while in Allaerts and Meyers (2017) results are averaged from 54 000 to 72 000 s. In Fig. 6, we report vertical profiles of velocity magnitude, direction, shear stress and potential temperature obtained from the four different cases.

As shown before in Sect. 3.2, TOSCA predicts more mixing than SP-Wind, used by Allaerts and Meyers (2017). This results in a higher inversion height for a given set of ABL parameters and can be observed by comparing cases G and S2, which feature the same wind angle controller but which differ in the profile of potential temperature obtained. This leads to an increased surface temperature predicted by TOSCA and a different wind veer profile between the two codes. Although we note that such differences are accentuated by the fact that statistics from SP-Wind are collected at an earlier time, i.e., the CNBL has grown by a lower extent, case S2 seems to be more aligned to case PT, where the average potential-temperature profile is kept constant by the controller. The difference in mixing between the two codes also affects the average hub-height velocity, which differs by 0.33 m s^{-1} between case G and case S2. For cases P and PT such a parameter is an input, and it has been set according to results from case G. In Table 5, output quantities extracted from the four different simulations are reported, averaging flow statistics in the abovementioned time intervals. The capping inversion center H , ground temperature θ_0 , inversion strength $\Delta\theta$ and inversion width Δh are calculated by fitting the Rampanelli and Zardi (2004) model in a least-squares sense.

Figure 6, together with quantitative data reported in Table 5, demonstrates how the pressure controller with geostrophic damping (case P) almost exactly matches results obtained using the geostrophic controller (case G), predicting a geostrophic wind that only differs by 0.25 % with respect to $G = 12 \text{ m s}^{-1}$. Figure 6 also highlights how sensitive the

Table 5. ABL parameters obtained by fitting the Rampanelli and Zardi (2004) model for the CNBL cases presented in this section, together with the resulting friction velocity, minimum heat flux and geostrophic wind angle. Data from cases G, P, and PT are averaged from 92 800 to 100 000 s, while in Allaerts and Meyers (2017) (case S2) data are averaged from 54 000 to 72 000 s.

	u_{ref} [m s ⁻¹]	G [m s ⁻¹]	θ_0 [K]	$\Delta\theta$ [K]	H [m]	Δh [m]	u^* [m s ⁻¹]	$q_{\text{min}}/10^{-4}$ [K m s ⁻¹]	ϕ_G [°]
G	10.871	12.00	288.33	1.95	612	113.5	0.323	-1.04	-12.23
P	10.871	11.97	288.33	1.95	612	113.5	0.323	-1.04	-12.23
PT	10.871	11.69	288.15	2.0	550	84.5	0.323	-1.33	-11.48
S2	11.200	12.00	288.19	1.99	585	93.8	0.315	-	-

ABL is to its heating history, since case PT – where the average θ profile is kept constant – predicts a lower geostrophic wind than cases G and P. In fact, it can be noticed from Fig. 7b how the inversion height is kept constant in case PT, while it grows in time during simulations P and G. For this reason, in the latter cases the ABL will experience a slightly higher amount of dissipation, which results in a small increase in the geostrophic wind if compared to case PT. Therefore, in simulations P, G and S2, the boundary layer is developing against a potential-temperature profile that is slowly evolving, in turn affecting the mean velocity profile. This mechanism, which is of course physical, does not reproduce what happens in real life, where the boundary layer stability evolves following the timescale of the diurnal cycle instead. As a consequence, since such temperature drift is physical but arises from an idealization, we believe that fixing the average potential-temperature profile would represent a more consistent method for conducting these types of idealized simulations. In addition, we also suggest driving the ABL with a pressure controller, which allows specifying the hub-height velocity, as the issue related to geostrophic inertial oscillations can be addressed using the proposed geostrophic damping method. This would allow precursor simulations to reach a truly statistically steady state, leading to better agreement between different codes applied to inflow conditions used for successor simulations.

Figure 7a shows the evolution of the geostrophic wind (components and magnitude), calculated as the spatial average in the homogeneous directions and at those cells where $h > H + \Delta$, produced by cases P and G. It can be seen how the developed damping technique is able to stop inertial oscillations after a time $T_D + T_3\%$, reaching a geostrophic wind that only differs by 0.25 % from the simulation where the geostrophic controller has been applied ($T_3\% = 17500$; see Sect. 2.3.1 for definition). Moreover, in Fig. 7b we report wind angle and velocity magnitude horizontally averaged at the reference height, together with the height of the inversion center over time, evaluated by fitting the Rampanelli and Zardi (2004) model at each time step. It is evident that the pressure controller precisely maintains the wind at the desired speed and in the desired direction. Interestingly, it can be also noticed that the geostrophic controller produces

small oscillations in the hub-height wind speed. These are inertial oscillations as well, but they are naturally damped by turbulence as they happen inside the boundary layer. Finally, looking at the evolution of the inversion layer height in Fig. 7b and at the final potential-temperature profile in Fig. 6, it is clear that controlling the mean potential temperature prevents the boundary layer from growing indefinitely, preserving the initial capping inversion height and the initial value of potential temperature at the ground.

5 Finite wind farm with thermal effects

In this section, we present results from the simulation of a finite-size wind farm consisting of 100 NREL 5 MW wind turbines, aligned in 20 rows and 5 columns, with stream-wise and spanwise spacing of 5 and 4.76 rotor diameters, respectively. We include thermal stratification to assess the effects of gravity wave blockage for a lapse rate of 1 K km⁻¹, a capping inversion centered at 500 m with a strength of 7.312 K. Given the large scale of the gravity waves, the numerical domain is set to 40 × 21 × 28 km in the streamwise, spanwise and vertical directions, respectively. All directions are graded to reach a mesh resolution of 30 × 12.5 × 10 m around the wind turbines. The hybrid off-line-concurrent-precursor technique described in Sect. 2.4 has been used to spin up turbulence in the precursor, providing a time-resolved CNBL inflow for the successor simulation. This technique is combined with a Rayleigh damping layer and the advection damping technique (Lanzilao and Meyers, 2022a) to ensure low reflectivity of gravity waves from the top boundary and the fringe region exit. Further details on the successor-precursor meshes and simulations; CNBL parameters; and tuning of the fringe, Rayleigh and advection damping region coefficients are given in Appendix C.

Figure 8 shows hub-height instantaneous velocity and pressure contours around the wind farm. The gravity wave footprint inside the ABL can be clearly noticed in the pressure field, together with the small-scale pressure increase in front of each rotor. This effect is superimposed on the much larger pressure variation due to atmospheric gravity waves, which take place from the farm entrance to the exit.




Figure 7. (a) Time evolution of average geostrophic wind (streamwise, spanwise components and magnitude from top to bottom) from cases P and G, where pressure and geostrophic controllers are used, respectively. Predictions using Eq. (11) are also shown. (b) Time evolution of the hub-height wind angle, hub-height velocity magnitude and capping inversion center from cases P and G (no potential-temperature control) and PT (with potential-temperature controller described in Sect. 2.3.2). Time is non-dimensionalized with the start time of the geostrophic damping action T_D .




Figure 8. Contours of instantaneous velocity (a) and pressure (b) at the wind turbine hub height.




Figure 9. (a) Perturbation pressure inside the ABL. (b) Perturbation potential temperature at the capping inversion height.

Regarding the instantaneous velocity field, streamwise streaks generated by elongated turbulence structures can be appreciated. The large size of these structures is related to the high value of the prescribed equivalent roughness height z_0 and to the fact that periodic boundary conditions artificially increase their length when they span the entire domain in the streamwise direction. If averages are gathered for a sufficient amount of time, these streaks are not expected to alter the simulation results from the wind farm power production standpoint.

At any given location, we define the perturbation value of a quantity as the difference between its successor time average and the precursor time average, evaluated at the same height. Figure 9 shows horizontal contours of pressure and temperature perturbations at the hub and inversion heights, respectively. An interesting aspect is that, due to the presence of the Coriolis force, the direction of propagation of interfacial waves in the inversion layer is not aligned with the wind farm streamwise symmetry axis.

For instance, the two trains of waves generated by the positive and negative inversion layer displacements, at the wind farm entrance and exit, respectively, have a spanwise offset, resulting in a much more complex interaction. Moreover, spurious wave interactions with their periodic images




Figure 10. (a) Hub-height wind speed; (b) hub-height perturbation velocity upstream of the first wind farm row. Each upstream location is averaged along the spanwise direction within the wind farm envelope.

can also be noticed, but the spanwise size of the domain ensures that they happen far from and downstream of the wind farm. Nevertheless, we are developing a lateral fringe region, which is aimed at removing this effect, where the instantaneous desired flow is reconstructed from the concurrent precursor, allowing for a smaller spanwise domain size. At this time, we do not believe wave interactions from periodic images alter the gravity wave pattern in the region near and up-

stream of the wind farm, which is of primary interest in order to assess wind farm blockage.

Figure 10a shows the average hub-height velocity field, from which the effect of gravity waves on velocity can be assessed. In particular, positive perturbations are observed where negative pressure gradients are experienced and vice versa. The perturbation, averaged along the spanwise direction and within the region enclosed by the wind farm spanwise limits, is shown quantitatively in Fig. 10b. In particular, we record velocity reductions as high as 5.80 %, 2.15 % and 1.02 % at 2.5, 20 and 40 diameters upstream of the first turbine row. Another aspect that is noticeable from the average fields of Figs. 9 and 10 is the presence of high-frequency oscillations in the spanwise direction. These are caused by the turbulence streaks mentioned earlier in the present section, and, while we plan on eliminating this shortcoming by applying a spanwise offset in the streamwise periodicity to artificially break the locking in position of such structures (Munters et al., 2016), we note that their manifestation could be readily eliminated by averaging for a time period longer than 15000 s.

Overall, these results indicate that the developed framework and methodology allow for conducting finite-size wind farm simulations, capturing gravity wave effects unaltered by spurious wave reflections from the fringe region or interaction from periodic images.

6 Conclusions

In the present paper, we introduced TOSCA, a new open-source LES framework for the simulation of large wind farms interacting with thermally stratified boundary layers. We validated TOSCA's wind turbine models and its ability to simulate the evolution of conventionally neutral boundary layers and to accurately predict the flow around infinite wind farms in neutral conditions. We presented a new controlling methodology for ABL precursors that allows us to prescribe a desired wind speed at a reference height – located inside the boundary layer – while at the same time avoiding velocity oscillations produced by the Coriolis force in the geostrophic region above the inversion height. This approach, if combined with a potential-temperature controller, allows different codes to obtain CNBL inflow profiles that are characterized by the same potential-temperature profile and hub-height wind speed. Conversely, using geostrophic forcing makes the hub-height velocity dependent on the amount of numerical dissipation specific to the adopted code, while the final temperature profile depends on both numerical dissipation and precursor simulated time. We also described a new methodology for simulating finite-size wind farms under atmospheric gravity wave effects. In particular, we introduced the hybrid off-line precursor–concurrent-precursor method, where the off-line technique is used on a small domain, in order to spin up ABL turbulence, while the concur-

rent method is adopted for the turbine simulation. In fact, we found that the concurrent precursor, combined with a fringe region, is a crucial element to simultaneously avoid spurious gravity wave reflections and provide a time-resolved turbulent inflow. The off-line precursor data are used to start up the flow field in the concurrent precursor by means of spanwise periodization. The concurrent-precursor domain is usually bigger than required, as its size is determined by the successor domain that runs concurrently. Hence, being able to reach steady-state turbulent statistics on a smaller domain is indeed convenient, as it makes finite wind farm simulations less computationally intensive.

Finally, we demonstrated that TOSCA is able to simulate wind farm gravity wave interactions and large-scale blockage effects. Specifically, for the CNBL simulated herein, we measured a velocity reduction of 5.80 % at 2.5 diameters upstream of the first row.

In the future, we will implement shifted periodic boundary conditions to obtain field statistics which are less dependent on the spanwise location, and we will address the heat flux mismatch above the inversion layer.

Appendix A: Numerical procedure

The adoption of generalized curvilinear coordinates allows for the computational mesh to follow terrain coordinates if required or to be stretched and deformed with the only condition that the indexing remains structured. We denote a set of generalized curvilinear coordinates as l_i , with $i = 1, 2, 3$, by which points in a three-dimensional Euclidean space E^3 may be defined. Cartesian coordinates are a special case of such a generalization and will be denoted as x_i , with $i = 1, 2, 3$. When using explicit notation, the three curvilinear directions will be identified by Greek symbols as ξ , η and ζ . With these definitions, and given the position vector \mathbf{r} of a point P in Cartesian space, the covariant base vectors can be expressed as $\mathbf{g}_i = \partial \mathbf{r} / \partial l_i$ (with Cartesian components $(\mathbf{g}_i)_j = \partial x_j / \partial l_i$), while contravariant base vectors are given by $\mathbf{g}^i = \nabla l_i$ (with Cartesian components $(\mathbf{g}^i)_j = \partial l_i / \partial x_j$). As a result, the following relation holds between covariant and contravariant base vectors:

$$\mathbf{g}^i = \frac{\mathbf{g}_{i+1} \times \mathbf{g}_{i+2}}{\mathbf{g}_i \cdot (\mathbf{g}_{i+1} \times \mathbf{g}_{i+2})} = J (\mathbf{g}_{i+1} \times \mathbf{g}_{i+2}), \quad (\text{A1})$$

where J is the Jacobian of the transformation defining l_i in terms of x_j , i.e., the determinant of the matrix of partial derivatives $\partial l_i / \partial x_j$. It is required that $J \neq 0$, which is equivalent to asking that covariant base vectors are not co-planar. Note that they are usually neither unit vectors nor orthogonal to each other. Given a set of curvilinear coordinates l_i , with covariant base vectors \mathbf{g}_i and contravariant base vectors \mathbf{g}^i , it is possible to define the covariant and contravariant metric tensors through the scalar products:

$$g_{ij} = \mathbf{g}_i \cdot \mathbf{g}_j = \frac{\partial x_k}{\partial l_i} \frac{\partial x_k}{\partial l_j}, g^{ij} = \mathbf{g}^i \cdot \mathbf{g}^j = \frac{\partial l_i}{\partial x_k} \frac{\partial l_j}{\partial x_k}, \quad (\text{A2})$$

where the repeated index implies summation. Metric tensors satisfy $J = \sqrt{\det(g_{ij})}$ and $J^{-1} = \sqrt{\det(g^{ij})}$.

The use of generalized curvilinear coordinates allows differential operators on any structured mesh to be expressed using a Cartesian-like discretization along the curvilinear directions, which are chosen to be the local structured grid lines. Moreover, the quantities

$$\begin{aligned} S^\xi &= \frac{1}{J} \left(\frac{\partial \xi}{\partial x} \hat{x} + \frac{\partial \xi}{\partial y} \hat{y} + \frac{\partial \xi}{\partial z} \hat{z} \right), \\ S^\eta &= \frac{1}{J} \left(\frac{\partial \eta}{\partial x} \hat{x} + \frac{\partial \eta}{\partial y} \hat{y} + \frac{\partial \eta}{\partial z} \hat{z} \right), \\ S^\zeta &= \frac{1}{J} \left(\frac{\partial \zeta}{\partial x} \hat{x} + \frac{\partial \zeta}{\partial y} \hat{y} + \frac{\partial \zeta}{\partial z} \hat{z} \right) \end{aligned} \tag{A3}$$

are equal to face area vectors if evaluated at cell faces. Contravariant base vectors components $\partial l_i / \partial x_j$ in Eq. (A3) are evaluated using Eq. (A1), i.e., from the covariant base vectors, which are easily obtained exploiting finite differences.

Contravariant components of any vector field \mathbf{u} , a function of position \mathbf{r} , can be expressed in terms of its Cartesian components as $u^i = \mathbf{u} \cdot \mathbf{g}^i$. If one instead uses \mathbf{g}^i / J (namely the face area vector along the i th direction S^i) and the relation is again evaluated at cell faces, contravariant fluxes are obtained as $V^i = \mathbf{u} \cdot S^i$. In TOSCA, only the independent variables (positions) are transformed in curvilinear coordinates using the chain rule and integration by parts, while dependent variables are retained in Cartesian coordinates. This partial transformation avoids computing the Christoffel symbols of the second kind, which are cumbersome to evaluate numerically. Moreover, they would increase the requirements of smoothness of the computational mesh, as they involve second-order derivatives of the transformation metrics. The momentum equation is finally dotted with the face area vectors, so it can be partially written in terms of contravariant fluxes as

$$\frac{\partial V^q}{\partial l_q} = 0, \tag{A4}$$

$$\begin{aligned} \frac{\partial V^q}{\partial t} + \frac{\partial l_q}{\partial x_i} \frac{\partial}{\partial l_r} (V^r u_i) &= -\frac{1}{\rho_0} \frac{\partial p}{\partial l_r} g^{rq} \\ &+ \frac{\partial l_q}{\partial x_i} \frac{\partial}{\partial l_r} \left[\frac{v_{eff}}{J} \left(\frac{\partial u_i}{\partial l_k} g^{rk} + \frac{\partial u_j}{\partial l_k} \frac{\partial l_r}{\partial x_j} \frac{\partial l_k}{\partial x_i} \right) \right] \\ &- \frac{1}{\rho_0} \frac{\partial p_\infty}{\partial l_r} g^{rq} + \frac{1}{J} \frac{\partial l_q}{\partial x_i} \frac{\rho_k}{\rho_0} g_i \\ &- \frac{2}{J} \frac{\partial l_q}{\partial x_i} \epsilon_{ijk} \Omega_j u_k + \frac{1}{J} \frac{\partial l_q}{\partial x_i} (f_i + s_i^v + s_i^h), \end{aligned} \tag{A5}$$

$$\frac{\partial \theta}{\partial t} + \frac{\partial}{\partial l_r} (V^r \theta) = J \frac{\partial}{\partial l_r} \left(\frac{\kappa_{eff}}{J} \frac{\partial \theta}{\partial l_k} g^{rk} \right). \tag{A6}$$

Equation (A5) is used to solve for contravariant fluxes, which are staggered at cell faces, while pressure is located at cell centers. In contrast to a staggered formulation using a full transformation, where Cartesian velocity does not appear in

the equations, in a partial transformation all Cartesian velocity components are required at each face center in order to discretize Eq. (A5) with the same accuracy. One alternative would be to solve all components of the momentum equation at each face center in order for the Cartesian velocity to be attainable without interpolation. Although this approach has been adopted in the literature (Maliska and Raithby, 1984), it triples the computational cost. In TOSCA, we follow the approach of Ge and Sotiropoulos (2007), where the momentum equation is first discretized at cell centers and then interpolated and solved at face centers in a staggered fashion. Cartesian velocity is subsequently reconstructed at cell centers by interpolating contravariant fluxes at the same location. With respect to a standard staggered formulation (e.g., in Cartesian coordinates), this procedure encompasses additional steps for interpolating the discretized momentum equation at face centers and for transforming the interpolated fluxes into the Cartesian velocity at cell centers (flux interpolation is required in either case). It should be noted that the overhead in computational cost is minimal because it only involves 1D interpolations along grid lines, for which a second-order central scheme is used. Another slightly different approach (Rosenfeld et al., 1992) is to discretize the momentum equation in a staggered manner. This avoids interpolating the whole momentum right-hand side at cell faces, but it requires interpolating contravariant fluxes instead. In either case, methods based on partially transformed equations involve an additional interpolation step (as contravariant fluxes and Cartesian velocity are defined at different locations). This imposes slightly tighter constraints on the time step value in order to keep the method stable. For this reason, we opted for an implicit treatment of advection and viscous terms in Eq. (A5). Specifically, we use the matrix-free Newton–Krylov solver implemented in PETSc (Balay et al., 2022), where the iterative Krylov subspace generalized minimum residual (GMRES) method (Saad and Schultz, 1986) is used to solve the linear system associated with each inner iteration. (See Knoll and Keyes, 2004, for a comprehensive review and application of matrix-free Newton–Krylov methods.) In addition, such a hybrid staggered–non-staggered formulation facilitates the application of boundary conditions, which are prescribed for the Cartesian velocity using ghost cells.

Pressure–velocity coupling is provided using a second-order fractional step method similar to that of van Kan (1986), where velocity is first guessed by solving for the contravariant fluxes, which are then projected into a divergence-free space by means of a pressure correction ϕ obtained by solving a Poisson equation. Potential temperature is subsequently solved using the new velocity field, with an implicit treatment of the right-hand side. Time discretization uses a second-order implicit scheme for both momentum and temperature equations. All derivatives are discretized using the second-order central scheme, while the advection term in Eq. (A5) is discretized using a blend between central and

QUICK (Leonard, 1979) schemes. The blending is such that QUICK is used in regions of almost uniform or slowly varying velocity, avoiding the oscillations produced by the central scheme in such regions.

Appendix B: LES modeling

To model the sub-grid stresses, TOSCA uses the dynamic Smagorinsky model (Lilly, 1992; Germano et al., 1991), with Lagrangian averaging of the model coefficient C_s (Meneveau et al., 1996). The model has been recast in generalized curvilinear coordinates, similar to what was presented in Armenio and Piomelli (2000). The effect of unresolved scales in the momentum equation, after the filtering operation, appears in Cartesian coordinates through the equation

$$\frac{\partial}{\partial x_j} (\overline{u_i u_j}) = \frac{\partial}{\partial x_j} (\overline{u_i} \overline{u_j}) + \frac{\partial}{\partial x_j} (\tau_{ij}^D), \tag{B1}$$

where the overbar indicates the filtering operation applied e.g. to a variable q as

$$\overline{q} = \int_{\mathcal{D}} q(\mathbf{x}', t) G(\mathbf{x} - \mathbf{x}') d\mathbf{x}', \tag{B2}$$

where the integration is performed over the entire flow domain \mathcal{D} (Moeng, 1984). In curvilinear coordinates, Eq. (B1) reads

$$\frac{\partial}{\partial l_k} (\overline{v^k u_j}) = \frac{\partial}{\partial l_k} (\overline{v^k} \overline{u_j}) + \frac{\partial}{\partial l_k} (\sigma_j^k), \tag{B3}$$

where

$$\sigma_j^k = \overline{v^k u_j} - \overline{v^k} \overline{u_j}, \tag{B4}$$

and assuming a linear eddy viscosity model,

$$\sigma_i^k = -2\nu_t S_j^k S_{ij}, \tag{B5}$$

$$\nu_t = C_s \Delta^2 \sqrt{2 S_{ij} S_{ij}}, \tag{B6}$$

where $S_j^k = 1/J \partial l_k / \partial x_j$ comprises the face area vectors, $S_{ij} = \frac{1}{2}(\partial u_i / \partial x_j + \partial u_j / \partial x_i)$ is the symmetric part of the velocity gradient tensor and Δ is the cubic root of the local cell volume. Using the idea of Germano et al. (1991), a second filter, denoted as $\widetilde{\cdot}$, can be applied which has $\widetilde{\Delta} = 3\Delta$ in TOSCA, leading to the tensor

$$T_j^k = \widetilde{\overline{v^k u_j}} - \widetilde{\overline{v^k}} \widetilde{\overline{u_j}}, \tag{B7}$$

which accounts for the effect of the unresolved plus the smallest resolved scales. The Germano tensor, i.e., the contribution to the resolved stresses from the largest unresolved motions, is defined in generalized curvilinear coordinates by subtracting the tilde-filtered Eq. (B3) from Eq. (B7):

$$G_j^k = T_j^k - \widetilde{\sigma_j^k} = \widetilde{\overline{v^k u_j}} - \widetilde{\overline{v^k}} \widetilde{\overline{u_j}}. \tag{B8}$$

Using Eqs. (B5) and (B6) to express $\widetilde{\sigma_j^k}$ and T_j^k gives

$$\widetilde{\sigma_j^k} = -2C_s \widetilde{\Delta}^2 \widetilde{|\overline{S}|} \widetilde{S_j^k} S_{ij}, \tag{B9}$$

$$T_j^k = -2C_s \widetilde{\Delta}^2 \widetilde{|\overline{S}|} \widetilde{S_j^k} \widetilde{S_{ij}}, \tag{B10}$$

where in Eq. (B10) the approximation $\widetilde{S_j^k} \widetilde{S_{ij}} = \widetilde{S_j^k} \widetilde{S_{ij}}$ has been used. In fact, as the LES filter has a size of three mesh cells in each direction, S_j^k (the face area vectors at the central cell) and $\widetilde{S_j^k}$ (the filtered face area vectors within the box) are almost identical provided that mesh grading is smooth enough. Conversely, the equality holds exactly if a homogeneous filter and a uniform mesh are considered. Inserting Eqs. (B9) and (B10) into Eq. (B8) leads to

$$G_j^k = -2C_s \widetilde{\Delta}^2 \widetilde{|\overline{S}|} \widetilde{S_j^k} \widetilde{S_{ij}} + 2C_s \Delta^2 \widetilde{|\overline{S}|} S_j^k S_{ij} = C_s M_j^k, \tag{B11}$$

where

$$M_j^k = 2 \left[\Delta^2 \widetilde{|\overline{S}|} S_j^k S_{ij} - \widetilde{\Delta}^2 \widetilde{|\overline{S}|} \widetilde{S_j^k} \widetilde{S_{ij}} \right]. \tag{B12}$$

It is now possible to find C_s in a least-squares sense as

$$C_s(\mathbf{x}, t) = \frac{M_j^k G_j^k}{M_m^n M_m^n}.$$

Note that the above relation is not invariant with respect to rotation of the reference frame because it implicitly contains the face area vectors; hence tensors are no longer symmetric. Variables must then be transformed into physical space to find C_s as

$$C_s(\mathbf{x}, t) = \frac{G_i^k M_i^q g_{kq}}{M_n^m M_n^l g_{ml}}, \tag{B13}$$

where g_{ij} is the covariant metric tensor. Since the C_s coefficient oscillates in space, some sort of average is required. TOSCA follows the approach presented in Meneveau et al. (1996), where the numerator and denominator of Eq. (B13) are averaged along streamlines as

$$\langle G_i^k M_i^q g_{kq} \rangle = I_{GM} = \int_{-\infty}^t G_i^k(t') M_i^q(t') g_{kq} W(t-t') dt', \tag{B14}$$

$$\langle M_n^m M_n^l g_{ml} \rangle = I_{MM} = \int_{-\infty}^t M_n^m(t') M_n^l(t') g_{ml} W(t-t') dt', \tag{B15}$$

where $W(t) = 1/T_s \exp(-t/T)$ is a weighting function and T_s is a timescale defined as

$$T_s = 1.5 \Delta \left[8 G_i^k M_i^q g_{kq} M_n^m M_n^l g_{ml} \right]^{-1/8}. \tag{B16}$$

The integrals of Eqs. (B14) and (B15) can be evaluated as

$$I_{GM}^n(\mathbf{x}) = \epsilon(G_i^k M_i^q g_{kq}) + (1 - \epsilon)(I_{GM}^{n-1}(\mathbf{x} - \mathbf{u}\Delta t)), \quad (\text{B17})$$

$$I_{MM}^n(\mathbf{x}) = \epsilon(M_n^m M_n^l g_{ml}) + (1 - \epsilon)(I_{MM}^{n-1}(\mathbf{x} - \mathbf{u}\Delta t)), \quad (\text{B18})$$

where $\epsilon = (\Delta t/T_s)/(1 + \Delta t/T_s)$. We use tri-linear interpolation formulas to evaluate the integrals I_{GM} and I_{MM} at the $\mathbf{x} - \mathbf{u}\Delta t$ position, and all quantities are evaluated at cell centers, including contravariant fluxes, which are linearly interpolated from the faces.

Regarding the potential-temperature equation, sub-grid fluxes are evaluated following the approach of Moeng (1984), i.e., through the definition of a thermal eddy diffusivity $\kappa_t = \nu_t/Pr_t$, where Pr_t is the turbulent Prandtl number, which depends on stability as

$$Pr = \frac{1}{1 + 2l/\Delta},$$

$$l = \begin{cases} \min\left(\frac{7.6\nu_t}{\Delta}, \sqrt{\theta_0/|s|}, \Delta\right) & \text{if } s < 0 \\ \Delta & \text{if } s \geq 0, \end{cases}$$

$$s = g_i \frac{\partial \theta}{\partial x_i}. \quad (\text{B19})$$

Note that, if the potential-temperature gradient is locally stable, $Pr_t \rightarrow 1$, while for neutral or unstable cases, $Pr_t = 1/3$. This reflects the decrease in the mixing length scale under stable conditions (Schumann, 1991).

Appendix C: Finite wind farm setup

In this section, we describe in detail the setup of the finite wind farm case presented in Sect. 5. To avoid wave reflections from inflow–outflow boundaries, we adopt periodic boundary conditions and the concurrent-precursor technique. This also provides a suitable turbulent inflow, eliminating the wind farm wake re-advected at the inlet by the periodic boundaries. To avoid wave reflections from the upper boundary, we use a Rayleigh damping layer, while lateral boundaries are periodic. Spanwise periodicity implies that gravity waves induced by the wind farm will interact with their periodic images, requiring the domain to be sufficiently large for these interactions to happen far from and downstream of the wind turbines. Moreover, we use the advection damping technique developed by Lanzilao and Meyers (2022a) to ensure that interactions between fringe-generated and physical waves are not advected downstream but remain trapped inside the advection damping region.

The size of the successor domain is $40 \text{ km} \times 21 \text{ km} \times 28 \text{ km}$ in the streamwise, spanwise and vertical directions, respectively, discretized with $1554 \times 1194 \times 345$ cells. All directions are graded to reach a mesh resolution of $30 \text{ m} \times 12.5 \text{ m} \times 10 \text{ m}$ around the wind farm.

The wind farm has a rectangular planform, with 20 rows and 5 columns. The first row is located at $x = 0$, and extends from 300 to 2700 m. This determines a lateral spacing of 600 m (4.76 D), while streamwise spacing is set to 630 m (5 D). Wind turbines are equipped with angular velocity and pitch controllers described in Jonkman et al. (2009). A very simple yaw controller is also added, which rotates each wind turbine independently with a uniform speed of 0.5° s^{-1} when flow misalignment exceeds 1° . The flow angle is calculated by filtering the wind velocity at a sampling point located 1 D upstream of the rotor center, using a time constant of 600 s. Turbines are modeled using the ADM, while tower and nacelle are not accounted for. The projection width ϵ is set to 18.75 m.

The concurrent-precursor mesh coincides with the portion of the successor domain which is located inside the fringe region. In particular, it is $5 \text{ km} \times 21 \text{ km} \times 28 \text{ km}$. The mesh resolution in the streamwise direction is 15 m, while in the spanwise and vertical directions, it is the same as the successor.

In order to save computational resources, we do not run the whole precursor simulation on the concurrent-precursor mesh, whose size is determined by the wind farm and gravity wave parameters. Instead, we perform the spin-up phase on a $6 \text{ km} \times 3 \text{ km} \times 1 \text{ km}$ domain, characterized by a resolution of $15 \text{ m} \times 15 \text{ m}$ in the streamwise and spanwise directions. The vertical direction is discretized in the same manner as the successor in order to increase the resolution inside the capping inversion layer. This spin-up phase is carried out for 10^5 s , after which an inflow database is collected. The generated inflow database is then used to start up the solution in the concurrent-precursor and successor domains. This technique, which we refer to as the hybrid off-line–concurrent precursor, is explained in Sect. 2.4. In the successor, after a spin-up of 5000 s, corresponding to slightly more than one flow-through time, data are gathered for 15 000 s.

The off-line precursor simulation uses the pressure and temperature controllers described in Sect. 2.3, while in both the concurrent-precursor and the successor simulations, velocity is controlled using a constant source term, obtained by averaging the off-line precursor source from 100 000 to 120 000 s. The temperature controller is retained in the concurrent-precursor simulation, but it is switched off in the successor so that the inversion height is free to be perturbed by the wind farm.

CNBL parameters used for the off-line precursor are summarized in Table C2. They are calculated based on the sensitivity analysis performed in Allaerts and Meyers (2019). In particular, our objective is to choose a set of non-dimensional parameters such that the capping inversion layer is strongly perturbed by the wind farm. This results in a capping inversion Froude number of $Fr = 0.94$ and an internal wave parameter of $P_N = 3.02$. These non-dimensional groups are related to the physics and magnitude of interfacial waves inside the inversion layer and internal gravity waves above the

Table C1. Mesh information for the finite wind farm case.

(a) Vertical discretization parameters				
z_s [km]	z_e [km]	Δz [m]	N [-]	f [-]
0	0.4	10	40	1
0.4	0.5	10–4.85	14	0.94591
0.5	0.6	4.59–10	15	1.05125
0.6	1	10	40	1
1	3	10–100	51	1.04698
3	17	100	140	1
17	28	100–500	44	1.03818

(b) Streamwise (left) and spanwise (right) discretization parameters.									
x_s [km]	x_e [km]	Δx [m]	N [-]	f [-]	y_s [km]	y_e [km]	Δy [m]	N [-]	f [-]
–20	–15.005	15	333	1	–9	–1.5	20	375	1
–15.005	–13	15–30	94	1.00748	–1.5	–0.5	20–12.5	62	0.99269
–13	18.02	30	1035	1	–0.5	3.5	12.5	320	1
18.02	19.97	30–15	90	0.9923	3.5	4.5	12.5–20	62	1.00805
19.97	20	15	2	1	4.5	12	20	375	1

ABL, respectively. The temperature profile is initialized according to Rampanelli and Zardi (2004).

In Fig. C1 we show vertical profiles of the wind speed magnitude, inflow angle, non-dimensional shear stress and potential temperature, averaged over the last 15000 s from the concurrent-precursor domain. It can be noticed how the pressure controller accurately maintains the desired wind speed and direction at h_{ref} and how the temperature controller removes the ground temperature shift observed in the previous sections by keeping the average profile constant in time. As a consequence, inversion height and strength are maintained equal to their initial values (Table C2), while the resulting friction velocity corresponds to 0.432 m s^{-1} .

Regarding the successor case, we followed the approach of Lanzilao and Meyers (2022a) to choose the damping layer and fringe coefficients. In particular, after a reflectivity study that employed a computationally cheap canopy model (not shown here), we found that a Rayleigh damping coefficient of $\nu_{RDL} = 0.05$ and a fringe damping coefficient of $\nu_{FR} = 0.03$ yielded minimal gravity wave reflectivity. The dominant vertical wavelength of the gravity waves is estimated as $\lambda_z = 2\pi G/N \approx 11.8 \text{ km}$ (Allaerts and Meyers, 2017), where N is the Brunt–Väisälä frequency ($N = 5.72 \times 10^{-3} \text{ s}^{-1}$ based on parameters listed in Table C2) and $G = 10.815 \text{ m s}^{-1}$. We ensure that at least one λ_z is contained in the Rayleigh damping layer by setting its width to 12 km. Regarding the advection damping technique developed by Lanzilao and Meyers (2022a), we observed that their guidelines in how to choose the length of the advection damping region did not apply to our case, which is characterized by a very strong inversion layer and a shallow boundary layer. In fact, we believe that a




Figure C1. Precursor inflow data, averaged from the concurrent-precursor domain from 105000 s onward. We only show a subsection of the domain in the vertical direction corresponding to $10 \times h_{ref}$. All profiles are uniform above what is plotted, except for the potential-temperature profile, which exhibits a linear increase equal to the lapse rate γ .

key parameter that needs to be tuned in order to avoid spurious gravity wave interactions is the length of the region where advection damping is applied after the fringe. This holds in particular for sub-critical ($Fr < 1$) cases, where waves inside the capping inversion can propagate against the flow. Here, perturbations would be propagated upstream of the wind farm to the fringe exit, being suddenly forced to obey the precursor inflow inside the fringe region. Such a sharp change in the boundary layer displacement at the fringe exit induces spurious gravity waves which remain trapped at their streamwise location if horizontal advection of vertical

Table C2. ABL parameters used for the finite wind farm simulation presented in this section.

u_{ref} [m s ⁻¹]	h_{ref} [m]	θ_0 [K]	$\Delta\theta$ [K]	Δh [m]	γ [K km ⁻¹]	H [m]	f_c [s ⁻¹]	z_0 [m]
9.0	90	300	7.312	100	1	500	9.6057×10^{-5}	0.05

Table C3. Fringe and advection damping region information.

(a) Fringe region parameters.			
x_s [km]	x_e [km]	Δ_s [km]	Δ_e [km]
-20	-15	1	1
(b) Advection damping region parameters.			
x_s [km]	x_e [km]	Δ_s [km]	Δ_e [km]
-18	-11	1	1

velocity is turned off. Nevertheless, these waves would interact with physical waves from inside the domain, resulting in more spurious interactions. As a consequence, it is crucial to ensure that all spurious interactions generated by this mechanism are fully contained within the advection damping region and are not advected downstream.

We used the same damping functions as Lanzilao and Meyers (2022a), and in Table C3 their parameters are reported for our finite wind farm simulation.

Appendix D: TOSCA parallel scaling

In this section, we show TOSCA's strong and weak parallel performance by running CNBL simulations with an increasing number of Niagara nodes for three different mesh sizes. The simulation setup corresponds to the off-line precursor described in Appendix C. The different meshes are evaluated by systematically doubling the number of elements in each direction, starting from $300 \times 300 \times 100$ cells. As a consequence, they consist of 9 million, 72 million and 576 million elements in total. Table D1 reports the number of nodes for each run, which only consisted of 2 h of wall-clock time.

Tests have been performed on Compute Canada's Niagara (2018) cluster, which consists of 2024 nodes, each with 40 Intel "Skylake" cores at 2.4 GHz or 40 Intel "Cascade Lake" cores at 2.5 GHz. Node interconnection consists of an EDR InfiniBand network, organized in a "Dragonfly+" topology with five dragonfly wings. Figure D1 shows the time per iteration as the node count increases for each of the CNBL meshes. TOSCA's strong scaling performance remains close to linear until roughly 25 000 cells per core are reached, which we identify as a reasonable trade-off between effi-

Table D1. Scaling tests performed on the Niagara Compute Canada cluster, each node consists of 40 CPUs.

	CNBL 9×10^6 cells	CNBL 72×10^6 cells	CNBL 576×10^6 cells
Number of Niagara nodes	1	5	100
	2	10	200
	4	20	400
	8	40	800
	16	80	-

ciency and speed. TOSCA was also successfully run at scale on the entire Niagara cluster to simulate a finite wind farm on a mesh exceeding 1 billion elements, proving TOSCA's suitability for massively parallel computations.

We also highlight that the time per iteration does not reflect the actual speed at which the simulation advances in time, as the time step size depends on the numerical method. In implicit methods like the Newton–Krylov solvers employed by TOSCA, the computational cost of each time step depends on the time step size, whereas these quantities are unrelated in explicit methods. Nevertheless, implicit methods are able to advance in time with a Courant–Friedrichs–Lewy (CFL) number greater than 1 (we used 0.9 for these analyses), while explicit methods are usually limited to a value close to 0.5.

Regarding actuator models, their solution and I/O operations are also parallelized in TOSCA. Specifically, we define a sphere of cells, for each wind turbine, that contains all cells that the rotor can possibly intersect when yawing. Processors owning mesh cells belonging to the sphere are then grouped into turbine-specific sub-communicators, which are used to solve wind turbines simultaneously. Hence, provided that a sufficiently high core count is used, the wind turbine update time in TOSCA is independent of the number of wind turbines in the simulation, and each communicator can write turbine data to file simultaneously.

For the finite wind farm simulation presented in Sect. 5 (100 wind turbines), the turbine update time was less than 0.1 s. The individual turbine update time depends on the cell size and on the actuator point-processor ownership search (which processor in the communicator controls which actuator point). The latter is only triggered by a change in yaw for the ADM and uniform ADM models, while it has to be performed at every iteration for the ALM, as actuator points are physically rotating.




Figure D1. TOSCA strong scaling performance on Compute Canada's Niagara cluster. All simulations share the same setup and only differ in the number of mesh elements.

Code availability. TOSCA is available at <https://doi.org/10.17605/OSF.IO/Q4VAF> (Stipa et al., 2023).

Data availability. The dataset can be made available from the authors upon request.

Author contributions. Conceptualization: SS, DA, JB; methodology: SS; software: SS, AA; validation: SS; formal analysis: SS; investigation: SS; computational resources: JB; data curation: SS; writing – original draft preparation: SS; writing – review and editing: JB, DA; visualization: SS; supervision: JB, DA; project administration: JB; funding acquisition: JB. All authors have read and agreed to the published version of the paper.

Competing interests. The contact author has declared that none of the authors has any competing interests.

Disclaimer. Publisher's note: Copernicus Publications remains neutral with regard to jurisdictional claims made in the text, published maps, institutional affiliations, or any other geographical representation in this paper. While Copernicus Publications makes every effort to include appropriate place names, the final responsibility lies with the authors.

Acknowledgements. Computational resources provided by the Digital Research Alliance of Canada and Advanced Research Computing at the University of British Columbia are gratefully acknowledged.

Financial support. This research has been supported by the Natural Sciences and Engineering Research Council of Canada (grant no. 556326).

Review statement. This paper was edited by Sukanta Basu and reviewed by two anonymous referees.

References

- Abkar, M. and Porté-Agel, F.: The effect of free-atmosphere stratification on boundary-layer flow and power output from very large wind farms, *Energies*, 6, 2338–2361, <https://doi.org/10.3390/en6052338>, 2013.
- Ahsbahs, T., Nygaard, N. G., Newcombe, A., and Badger, M.: Wind Farm Wakes from SAR and Doppler Radar, *Remote Sensing*, 12, 462, <https://doi.org/10.3390/rs12030462>, 2020.
- Ainslie, J.: Calculating the flowfield in the wake of wind turbines, *J. Wind Eng. Ind. Aerod.*, 27, 213–224, [https://doi.org/10.1016/0167-6105\(88\)90037-2](https://doi.org/10.1016/0167-6105(88)90037-2), 1988.
- Allaerts, D.: PhD Thesis: Large-eddy Simulation of Wind Farms in Conventionally Neutral and Stable Atmospheric Boundary Layers, Tech. rep., KU Leuven, <https://lirias.kuleuven.be/1731391&lang=en> (last access: 1 November 2022), 2016.
- Allaerts, D. and Meyers, J.: Large eddy simulation of a large wind-turbine array in a conventionally neutral atmospheric boundary layer, *Phys. Fluids*, 27, 065108, <https://doi.org/10.1063/1.4922339>, 2015.
- Allaerts, D. and Meyers, J.: Boundary-layer development and gravity waves in conventionally neutral wind farms, *J. Fluid Mech.*, 814, 95–130, <https://doi.org/10.1017/jfm.2017.11>, 2017.
- Allaerts, D. and Meyers, J.: Gravity Waves and Wind-Farm Efficiency in Neutral and Stable Conditions, *Bound.-Lay. Meteorol.*, 166, 269–299, <https://doi.org/10.1007/s10546-017-0307-5>, 2018.
- Allaerts, D. and Meyers, J.: Sensitivity and feedback of wind-farm-induced gravity waves, *J. Fluid Mech.*, 862, 990–1028, <https://doi.org/10.1017/jfm.2018.969>, 2019.
- Allaerts, D., Quon, E., Draxl, C., and Churchfield, M.: Development of a Time–Height Profile Assimilation Technique for Large-Eddy Simulation, *Bound.-Lay. Meteorol.*, 176, 329–348, <https://doi.org/10.1007/s10546-020-00538-5>, 2020.
- Allaerts, D., Quon, E., and Churchfield, M.: Using observational mean-flow data to drive large-eddy simulations of a diurnal cycle at the SWiFT site, *Wind Energy*, 26, 469–492, <https://doi.org/10.1002/we.2811>, 2023.
- Armenio, V. and Piomelli, U.: A Lagrangian Mixed Subgrid-Scale Model in Generalized Coordinates, *Flow Turbul. Combust.*, 65, 51–81, <https://doi.org/10.1023/A:1009998919233>, 2000.
- Balay, S., Abhyankar, S., Adams, M. F., Benson, S., Brown, J., Brune, P., Buschelman, K., Constantinescu, E., Dalcin, L., Dener, A., Eijkhout, V., Faibussovitch, J., Gropp, W. D., Hapla, V., Isaac, T., Jolivet, P., Karpeev, D., Kaushik, D., Knepley, M. G., Kong, F., Kruger, S., May, D. A., McInnes, L. C., Mills, R. T., Mitchell, L., Munson, T., Roman, J. E., Rupp, K., Sanan, P., Sarich, J., Smith, B. F., Zampini, S., Zhang, H., Zhang, H., and Zhang, J.: PETSc/TAO Users Manual, Tech. Rep. ANL-21/39 – Revision 3.18, Argonne National Laboratory, <https://doi.org/10.2172/1893326>, 2022.
- Bastankhah, M. and Porté-Agel, F.: A new analytical model for wind-turbine wakes, *Renew. Energ.*, 70, 116–123, <https://doi.org/10.1016/j.renene.2014.01.002>, 2014.
- Bleeg, J., Purcell, M., Ruisi, R., and Traiger, E.: Wind Farm Blockage and the Consequences of Neglecting

- Its Impact on Energy Production, *Energies*, 11, 1609, <https://doi.org/10.3390/en11061609>, 2018.
- Branlard, E. and Gaunaa, M.: Cylindrical vortex wake model: right cylinder, *Wind Energy*, 524, 1973–1987, 2014.
- Branlard, E. and Meyer Forsting, A. R.: Assessing the blockage effect of wind turbines and wind farms using an analytical vortex model, *Wind Energy*, 23, 2068–2086, <https://doi.org/10.1002/we.2546>, 2020.
- Branlard, E., Quon, E., Forsting, A. R. M., King, J., and Moriarty, P.: Wind farm blockage effects: comparison of different engineering models, *J. Phys. Conf. Ser.*, 1618, 062036, <https://doi.org/10.1088/1742-6596/1618/6/062036>, 2020.
- Breton, S., Sumner, J., Sørensen, J., Hansen, K., Sarmast, S., and Ivanell, S.: A survey of modelling methods for high-fidelity wind farm simulations using large eddy simulation, *Philos. T. R. Soc. A*, 375, 20160097, <https://doi.org/10.1098/rsta.2016.0097>, 2017.
- Calaf, M., Meneveau, C., and Meyers, J.: Large eddy simulations of fully developed wind-turbine array boundary layers, *Phys. Fluids*, 22, 015110, <https://doi.org/10.1063/1.3291077>, 2010.
- Centurelli, G., Vollmer, L., Schmidt, J., Dörenkämper, M., Schröder, M., Lukassen, L. J., and Peinke, J.: Evaluating Global Blockage engineering parametrizations with LES, *J. Phys. Conf. Ser.*, 1934, 012021, <https://doi.org/10.1088/1742-6596/1934/1/012021>, 2021.
- Chamorro, L. and Porté-Agel, F.: Turbulent Flow Inside and Above a Wind Farm: A Wind-Tunnel Study, *Energies*, 4, 1916–1936, <https://doi.org/10.3390/en4111916>, 2011.
- Churchfield, M., Lee, S., Michalakes, J., and Moriarty, P.: Numerical Study of the Effects of Atmospheric and Wake Turbulence on Wind Turbine Dynamics, *J. Turbul.*, 13, N14, <https://doi.org/10.1080/14685248.2012.668191>, 2012a.
- Churchfield, M., Lee, S., Moriarty, P., Martínez Tossas, L., Leonardi, S., Vijayakumar, G., and Brasseur, J.: A Large-Eddy Simulation of Wind-Plant Aerodynamics, <https://doi.org/10.2514/6.2012-537>, ISBN 978-1-60086-936-5, 2012b.
- Churchfield, M., Schreck, S., Martínez Tossas, L., Meneveau, C., and Spalart, P.: An Advanced Actuator Line Method for Wind Energy Applications and Beyond, in: AIAA SciTech 35th Wind Energy Symposium, 9–13 January 2017, Grapevine, Texas, <https://doi.org/10.2514/6.2017-1998>, 2017.
- Etling, D.: Modelling the vertical ABL structure, in: collection: Modelling of Atmospheric Flow Fields, edited by: Lalas, D. P. and Ratto, C. F., World Scientific, 45–86, https://doi.org/10.1142/9789814447164_0003, 1996.
- Falgout, R. D. and Yang, U. M.: hypre: A Library of High Performance Preconditioners, in: Computational Science – ICCS 2002, edited by Sloot, P. M. A., Hoekstra, A. G., Tan, C. J. K., and Dongarra, J. J., Springer Berlin Heidelberg, Berlin, Heidelberg, 632–641, ISBN 978-3-540-47789-1, 2002.
- Fleming, P. A., Geeraad, P. M., Lee, S., van Wingerden, J.-W., Johnson, K., Churchfield, M., Michalakes, J., Spalart, P., and Moriarty, P.: Evaluating techniques for redirecting turbine wakes using SOWFA, *Renew. Energ.*, 70, 211–218, <https://doi.org/10.1016/j.renene.2014.02.015>, 2014.
- Gabriel, E., Fagg, G. E., Bosilca, G., Angskun, T., Dongarra, J. J., Squyres, J. M., Sahay, V., Kambadur, P., Barrett, B., Lumsdaine, A., Castain, R. H., Daniel, D. J., Graham, R. L., and Woodall, T. S.: Open MPI: Goals, Concept, and Design of a Next Generation MPI Implementation, in: Proceedings, 11th European PVM/MPI Users' Group Meeting, Budapest, Hungary, 97–104, https://doi.org/10.1007/978-3-540-30218-6_19, 2004.
- Ge, L. and Sotiropoulos, F.: A numerical method for solving the 3D unsteady incompressible Navier–Stokes equations in curvilinear domains with complex immersed boundaries, *J. Comput. Phys.*, 225, 1782–1809, <https://doi.org/10.1016/j.jcp.2007.02.017>, 2007.
- Germano, M., Piomelli, U., Moin, P., and Cabot, W. H.: A dynamic subgrid-scale eddy viscosity model, *Phys. Fluids A-Fluid*, 3, 1760–1765, <https://doi.org/10.1063/1.857955>, 1991.
- Glauert, H.: *Airplane Propellers*, Springer Berlin Heidelberg, Berlin, Heidelberg, 169–360, https://doi.org/10.1007/978-3-642-91487-4_3, ISBN 978-3-642-91487-4, 1935.
- Gribben, B. J. and Hawkes, G. S.: A potential flow model for wind turbine induction and wind farm blockage, *Systems and Engineering Technology*, <https://www.fnc.co.uk/wind-farm-blockage/> (last access: 1 April 2023), 2019.
- Haji Mohammadi, M., Sotiropoulos, F., and Brinkerhoff, J.: Moving least squares reconstruction for sharp interface immersed boundary methods, *International Journal for Numerical Methods in Fluids*, 90, 57–80, <https://doi.org/10.1002/flid.4711>, 2019.
- Haupt, S. E., Kosović, B., Berg, L. K., Kaul, C. M., Churchfield, M., Mirocha, J., Allaerts, D., Brummet, T., Davis, S., DeCastro, A., Dettling, S., Draxl, C., Gagne, D. J., Hawbecker, P., Jha, P., Juliano, T., Lassman, W., Quon, E., Rai, R. K., Robinson, M., Shaw, W., and Thedin, R.: Lessons learned in coupling atmospheric models across scales for onshore and offshore wind energy, *Wind Energ. Sci.*, 8, 1251–1275, <https://doi.org/10.5194/wes-8-1251-2023>, 2023.
- Inoue, M., Matheou, G., and Teixeira, J.: LES of a Spatially Developing Atmospheric Boundary Layer: Application of a Fringe Method for the Stratocumulus to Shallow Cumulus Cloud Transition, *Mon. Weather Rev.*, 142, 3418–3424, <https://doi.org/10.1175/MWR-D-13-00400.1>, 2014.
- Jensen, N.: A note on wind generator interaction, *Tech. Rep. Risø M-2411*, <https://orbit.dtu.dk/en/publications/a-note-on-wind-generator-interaction> (last access: 1 April 2023), 1983.
- Jimenez, A., Crespo, A., Migoya, E., and Garcia, J.: Advances in large-eddy simulation of a wind turbine wake, *J. Phys. Conf. Ser.*, 75, 012041, <https://doi.org/10.1088/1742-6596/75/1/012041>, 2007.
- Jimenez, A., Crespo, A., Migoya, E., and Garcia, J.: Large-eddy simulation of spectral coherence in a wind turbine wake, *Environ. Res. Lett.*, 3, 015004, <https://doi.org/10.1088/1748-9326/3/1/015004>, 2008.
- Johlas, H. M., Martínez-Tossas, L. A., Churchfield, M. J., Lackner, M. A., and Schmidt, D. P.: Floating platform effects on power generation in spar and semisubmersible wind turbines, *Wind Energy*, 24, 901–916, <https://doi.org/10.1002/we.2608>, 2021.
- Jonkman, J., Butterfield, S., Musial, W., and Scott, G.: Definition of a 5MW Reference Wind Turbine for Offshore System Development, National Renewable Energy Laboratory (NREL), <https://doi.org/10.2172/947422>, 2009.
- Knoll, D. and Keyes, D.: Jacobian-free Newton–Krylov methods: a survey of approaches and applications, *J. Computat. Phys.*, 193, 357–397, <https://doi.org/10.1016/j.jcp.2003.08.010>, 2004.

- Lanzilao, L. and Meyers, J.: An Improved Fringe-Region Technique for the Representation of Gravity Waves in Large Eddy Simulation with Application to Wind Farms, *Bound.-Lay. Meteorol.*, 186, 567–593, <https://doi.org/10.1007/s10546-022-00772-z>, 2022a.
- Lanzilao, L. and Meyers, J.: Effects of self-induced gravity waves on finite wind-farm operations using a large-eddy simulation framework, *J. Phys.: Conf. Ser.*, 2265, 022043, <https://doi.org/10.1088/1742-6596/2265/2/022043>, 2022b.
- Larsen, G.: A simple wake calculation procedure. Tech. Rep. Risø-M-2760 Risø., Tech. rep., <https://orbit.dtu.dk/en/publications/a-simple-wake-calculation-procedure> (last access: 1 April 2023) 1988.
- Leonard, B.: A stable and accurate convective modelling procedure based on quadratic upstream interpolation, *Comput. Meth. Appl. M.*, 19, 59–98, [https://doi.org/10.1016/0045-7825\(79\)90034-3](https://doi.org/10.1016/0045-7825(79)90034-3), 1979.
- Lilly, D. K.: A proposed modification of the Germano subgrid-scale closure method, *Phys. Fluids A-Fluid*, 4, 633–635, <https://doi.org/10.1063/1.858280>, 1992.
- Loken, C., Gruner, D., Groer, L., Peltier, R., Bunn, N., Craig, M., Henriques, T., Dempsey, J., Yu, C.-H., Chen, J., Dursi, L. J., Chong, J., Northrup, S., Pinto, J., Knecht, N., and Zon, R. V.: SciNet: Lessons Learned from Building a Power-efficient Top-20 System and Data Centre, *J. Phys. Conf. Ser.*, 256, 012026, <https://doi.org/10.1088/1742-6596/256/1/012026>, 2010.
- Lundquist, J., Duvivier, K., Kaffine, D., and Tomaszewski, J.: Costs and consequences of wind turbine wake effects arising from uncoordinated wind energy development, *Nature Energy*, 4, 25–34, <https://doi.org/10.1038/s41560-018-0281-2>, 2019.
- Maliska, C. R. and Raithby, G. D.: A method for computing three dimensional flows using non-orthogonal boundary-fitted coordinates, *International Journal for Numerical Methods in Fluids*, 4, 519–537, <https://doi.org/10.1002/flid.1650040606>, 1984.
- Maronga, B., Gryscha, M., Heinze, R., Hoffmann, F., Kanani-Sühring, F., Keck, M., Ketelsen, K., Letzel, M. O., Sühring, M., and Raasch, S.: The Parallelized Large-Eddy Simulation Model (PALM) version 4.0 for atmospheric and oceanic flows: model formulation, recent developments, and future perspectives, *Geosci. Model Dev.*, 8, 2515–2551, <https://doi.org/10.5194/gmd-8-2515-2015>, 2015.
- Martínez-Tossas, L. A., Churchfield, M. J., and Leonardi, S.: Large eddy simulations of the flow past wind turbines: actuator line and disk modeling, *Wind Energy*, 18, 1047–1060, <https://doi.org/10.1002/we.1747>, 2015.
- Meneveau, C., Lund, T. S., and Cabot, W. H.: A Lagrangian dynamic subgrid-scale model of turbulence, *J. Fluid Mech.*, 319, 353–385, <https://doi.org/10.1017/S0022112096007379>, 1996.
- Meyers, J. and Meneveau, C.: Large Eddy Simulations of Large Wind-Turbine Arrays in the Atmospheric Boundary Layer, in: 48th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition, 4–7 January 2010, Orlando, FL, <https://doi.org/10.2514/6.2010-827>, 2010.
- Min, M., Brazell, M., Tomboulides, A., Churchfield, M., Fischer, P., and Sprague, M.: Towards Exascale for Wind Energy Simulations, *arXiv [preprint]*, <https://doi.org/10.48550/ARXIV.2210.00904>, 2022.
- Moeng, C.-H.: A Large-Eddy-Simulation Model for the Study of Planetary Boundary-Layer Turbulence, *J. Atmos. Sci.*, 41, 2052–2062, [https://doi.org/10.1175/1520-0469\(1984\)041<2052:ALESMF>2.0.CO;2](https://doi.org/10.1175/1520-0469(1984)041<2052:ALESMF>2.0.CO;2), 1984.
- Monin, A. and Obukhov, A.: Basic laws of turbulent mixing in the surface layer of the atmosphere, *Tr. Akad. Nauk SSSR Geophys. Inst.*, 151, 163–187, 1954.
- Munters, W., Meneveau, C., and Meyers, J.: Shifted periodic boundary conditions for simulations of wall-bounded turbulent flows, *Phys. Fluids*, 28, 025112, <https://doi.org/10.1063/1.4941912>, 2016.
- Niagara: Niagara, <https://docs.alliancecan.ca/wiki/Niagara> (last access: 1 April 2023), 2018.
- Niayifar, A. and Porté-Agel, F.: Analytical Modeling of Wind Farms: A New Approach for Power Prediction, *Energies*, 9, 741, <https://doi.org/10.3390/en9090741>, 2016.
- Nygaard, N. G., Steen, S., Poulsen, L., and Pedersen, J. G.: Modelling cluster wakes and wind farm blockage, *J. Phys. Conf. Ser.*, 1618, 062072, <https://doi.org/10.1088/1742-6596/1618/6/062072>, 2020.
- Nygaard, N. G., Poulsen, L., Svensson, E., and Grønnegaard Pedersen, J.: Large-scale benchmarking of wake models for offshore wind farms, *J. Phys.: Conf. Ser.*, 2265, 022008, <https://doi.org/10.1088/1742-6596/2265/2/022008>, 2022.
- OpenCFD: OpenFOAM – The Open Source CFD Toolbox – User’s Guide – Version 6, OpenCFD Ltd., United Kingdom, <https://www.openfoam.com/> (last access: 1 April 2023) 2018.
- Ørsted: Ørsted presents update on its long-term financial targets, <https://orsted.com/en/company-announcement-list/2019/10/1937002> (last access: 5 March 2023), 2019.
- Paulson, C. A.: The Mathematical Representation of Wind Speed and Temperature Profiles in the Unstable Atmospheric Surface Layer, *J. Appl. Meteorol. Climatol.*, 9, 857–861, [https://doi.org/10.1175/1520-0450\(1970\)009<0857:TMROWS>2.0.CO;2](https://doi.org/10.1175/1520-0450(1970)009<0857:TMROWS>2.0.CO;2), 1970.
- Pedersen, J. G., Gryning, S.-E., and Kelly, M.: On the Structure and Adjustment of Inversion-Capped Neutral Atmospheric Boundary-Layer Flows: Large-Eddy Simulation Study, *Bound.-Lay. Meteorol.*, 153, 43–62, <https://doi.org/10.1007/s10546-014-9937-z>, 2014.
- Pedersen, J. G., Svensson, E., Poulsen, L., and Nygaard, N. G.: Turbulence Optimized Park model with Gaussian wake profile, *J. Phys.: Conf. Ser.*, 2265, 022063, <https://doi.org/10.1088/1742-6596/2265/2/022063>, 2022.
- Ponce, M., van Zon, R., Northrup, S., Gruner, D., Chen, J., Ertinaz, F., Fedoseev, A., Groer, L., Mao, F., Mundim, B. C., Nolta, M., Pinto, J., Saldarriaga, M., Slavnic, V., Spence, E., Yu, C.-H., and Peltier, W. R.: Deploying a Top-100 Supercomputer for Large Parallel Workloads: The Niagara Supercomputer, in: Proceedings of the Practice and Experience in Advanced Research Computing on Rise of the Machines (Learning), PEARC ’19, Association for Computing Machinery, New York, NY, USA, <https://doi.org/10.1145/3332186.3332195>, ISBN 9781450372275, 2019.
- Porté-Agel, F., Lu, H., and Wu, Y.-T.: A large-eddy simulation framework for wind energy applications, in: Fifth International Symposium on Computational Wind Engineering (CWE2010), 23–27 May 2010, Chapel Hill, North Carolina, USA, 2010.
- Rampanelli, G. and Zardi, D.: A Method to Determine the Capping Inversion of the Convective Boundary Layer, *J.*

- Appl. Meteorol., 43, 925–933, [https://doi.org/10.1175/1520-0450\(2004\)043<0925:AMTDTC>2.0.CO;2](https://doi.org/10.1175/1520-0450(2004)043<0925:AMTDTC>2.0.CO;2), 2004.
- Rosenfeld, M., Kwak, D., and Vinokur, M.: Development of a fractional-step method for the unsteady incompressible Navier-Stokes equations in generalized coordinate systems, NASA Technical Memorandum no. 103912, NASA, 1992.
- Saad, Y. and Schultz, M. H.: GMRES: A Generalized Minimal Residual Algorithm for Solving Nonsymmetric Linear Systems, *SIAM J. Sci. Stat. Comp.*, 7, 856–869, <https://doi.org/10.1137/0907058>, 1986.
- Schneemann, J., Rott, A., Dörenkämper, M., Steinfeld, G., and Kühn, M.: Cluster wakes impact on a far-distant offshore wind farm's power, *Wind Energy Science*, 5, 29–49, <https://doi.org/10.5194/wes-5-29-2020>, 2020.
- Schumann, U.: Subgrid length-scales for large-eddy simulation of stratified turbulence, *Theor. Comput. Fluid Dyn.*, 2, 279–290, 1991.
- Segalini, A.: An analytical model of wind-farm blockage, *J. Renew. Sustain. Ener.*, 13, 033307, <https://doi.org/10.1063/5.0046680>, 2021.
- Sescu, A. and Meneveau, C.: A control algorithm for statistically stationary Large Eddy Simulations of thermally stratified boundary layers, *Q. J. Roy. Meteor. Soc.*, 140, 2017–2022, <https://doi.org/10.1002/qj.2266>, 2014.
- Shaw, W. J., Berg, L. K., Debnath, M., Deskos, G., Draxl, C., Ghate, V. P., Hasager, C. B., Kotamarthi, R., Mirocha, J. D., Muradyan, P., Pringle, W. J., Turner, D. D., and Wilczak, J. M.: Scientific challenges to characterizing the wind resource in the marine atmospheric boundary layer, *Wind Energ. Sci.*, 7, 2307–2334, <https://doi.org/10.5194/wes-7-2307-2022>, 2022.
- Smith, R. B.: Gravity wave effects on wind farm efficiency, *Wind Energy*, 13, 449–458, <https://doi.org/10.1002/we.366>, 2010.
- Sockeye: UBC ARC: Sockeye – Detailed Technical Specifications, <https://arc.ubc.ca/sockeye-techspecs> (last access: 1 April 2023), 2023.
- Sørensen, J. N. and Shen, W. Z.: Numerical Modeling of Wind Turbine Wakes, *J. Fluids Eng.*, 124, 393–399, <https://doi.org/10.1115/1.1471361>, 2002.
- Sørensen, J. N., Mikkelsen, R. F., Henningson, D. S., Ivanell, S., Sarmast, S., and Andersen, S. J.: Simulation of wind turbine wakes using the actuator line technique, *Philos. T. R. Soc. A*, 373, 20140071, <https://doi.org/10.1098/rsta.2014.0071>, 2015.
- Stevens, R. J., Martínez-Tossas, L. A., and Meneveau, C.: Comparison of wind farm large eddy simulations using actuator disk and actuator line models with wind tunnel experiments, *Renew. Energ.*, 116, 470–478, <https://doi.org/10.1016/j.renene.2017.08.072>, 2018.
- Stieren, A., Gadde, S. N., and Stevens, R. J.: Modeling dynamic wind direction changes in large eddy simulations of wind farms, *Renew. Energ.*, 170, 1342–1352, <https://doi.org/10.1016/j.renene.2021.02.018>, 2021.
- Stipa, S., Ajay, A., and Brinkerhoff, J.: Toolbox fOr Stratified Convective Atmospheres (TOSCA), OSF [code], <https://doi.org/10.17605/OSF.IO/Q4VAF>, 2023.
- Stull, R.: Practical Meteorology: An Algebra-based Survey of Atmospheric Science, BC Open Textbook Collection, AVP International, University of British Columbia, ISBN 9780888651761, <https://books.google.nl/books?id=xP2sDAEACAAJ> (last access: 1 April 2023), 2016.
- The HDF Group: Hierarchical data format version 5, <http://www.hdfgroup.org/HDF5> (last access: 1 April 2023), 2006.
- Troldborg, N. and Meyer Forsting, A.: A simple model of the wind turbine induction zone derived from numerical simulations, *Wind Energy*, 20, 2011–2020, <https://doi.org/10.1002/we.2137>, 2017.
- van der Laan, M. P., Storey, R., Sørensen, N., Norris, S., and Cater, J.: A CFD code comparison of wind turbine wakes, *J. Phys. Conf. Ser.*, 524, 012140, <https://doi.org/10.1088/1742-6596/524/1/012140>, 2014.
- van Kan, J.: A Second-Order Accurate Pressure-Correction Scheme for Viscous Incompressible Flow, *SIAM J. Sci. Stat. Comp.*, 7, 870–891, <https://doi.org/10.1137/0907059>, 1986.
- Wu, K. L. and Porté-Agel, F.: Flow Adjustment Inside and Around Large Finite-Size Wind Farms, *Energies*, 10, 2164, <https://doi.org/10.3390/en10122164>, 2017.
- Yang, X. I. A., Park, G. I., and Moin, P.: Log-layer mismatch and modeling of the fluctuating wall stress in wall-modeled large-eddy simulations, *Phys. Rev. Fluids*, 2, 104601, <https://doi.org/10.1103/PhysRevFluids.2.104601>, 2017.