

University of British Columbia

UBCO-UL NSERC Alliance Grant “Reduced-Order Models of Wind
Farm Induction and Far-Field Wake Recovery”

Response to Reviewer 1

Exec. S. Stipa - November 14, 2023

We would like to thank the reviewer for the time dedicated to revising the paper. We proceed with answering and clarifying, where possible, the proposed comments.

Our response, denoted in black, is shown below. Modified text of the paper is shown between quotes in *italic*, while the reviewers' comments are denoted in blue. Please refer to the track changes section at the end of this document for a detailed overview of the changes made to the manuscript.

I very much appreciate the very extensive responses that have been provided. As previously indicated, I strongly support the publication of the manuscript. Overall, this is an excellent and very important study for the community and certainly qualifies for publication in *Wind Energy Science*. Nevertheless, I feel that some of the answers should be better reflected in the updated manuscript, as indicated below.

While I agree that some of the suggested controllers are useful, as explained in the manuscript, I believe the statement that these are "better" (line 641) is too strong. While the temperature profile controller is useful in the context of reproducibility, it does affect physics as the way buoyancy is handled is changed. This aspect should be better reflected in the used formulation. In the abstract and conclusions, the reason for reproducibility should be better reflected instead of claiming it is simply better. As an aside, as the manuscript shows that results from different codes can vary from each other for different reasons, the proposed technique may not even assure full reproducibility between different codes.

The comment has been introduced in the manuscript. In particular abstract, result section and conclusions have been modified.

I thank the authors for updating Table 1. While the difference between TOSCA and OpenFOAM is much smaller than in the previous version, a difference of 6% is still significant, considering that a uniform inflow of cases is considered. The explanation below in Table 1 seems rather speculative, and I believe it should be replaced by a statement that the reason for the difference is unknown. Unless the authors could prove that their statement line 501 is indeed correct.


Updated accordingly.

Figure 7: There is no equation 36.


Thank you very much for pointing it out. The reference to Eq. 11 in Fig. 7 has now been updated after moving the LES section in the Appendix.

Figures 9 and 10 have a high-frequency oscillation in the spanwise direction. → In the manuscript, no clarification has been made on this point. Some of the statements provided in answer to the referee could be elaborated, in particular, why these waves already seem to be present in the inflow, so these waves also travel upstream. Is the statement "free of spurious (i.e. non-physical) waves" based on a calculation, or it this interpretation of the simulation result?

An explanation has been added in the paper after figures mentioned by the reviewer. However, the phenomenon was already explained at the end of Sec. 2.4. In addition, 15,000 s may not be sufficient to obtain perfectly clean average fields, and the artefact of inflow turbulence may still be observed. As a reference, we report below a comparison between two hub-height averages obtained from two similar cases where one has been run for 40,000 s without a spanwise shift in the periodized inflow turbulence, while the second has been run for 20,000 s with a spanwise shift of the inflow condition. As can be seen, while high-frequency oscillations can be detected in the first case, the second simulation depicts almost a perfectly uniform average velocity despite running for half the time. We do not believe this artefact poses any problem apart from yielding visualizations of lower quality.



(a)



(b)

Figure 1. (a) Hub-height wind speed for the simulation that ran for 40,000 s without spanwise inflow shift; (b) Hub-height wind speed for the simulation that ran for 20,000 s with spanwise inflow shift.

Line 677 " If averages are gathered for a sufficient amount of time, these streaks do not alter the simulation results, from the wind farm performance point of view" → The statement is too strong as turbulent characteristics of streaks would affect coherence between different downstream rows. The present statement mostly refers to power production. Please formulate this more carefully.

Rephrased according to the reviewer comment.

Why is case S2 in Figure 6d not zero for $h/H \leq 0.8$ as is the case for case P?

The case S2, described in Allaerts and Meyers (2017), does not feature potential temperature control. As a consequence, the average θ inside the boundary layer increases over time due to the fact that turbulent mixing slowly “eats up” the free atmosphere. This is explained in the paper and it is exactly the reason why we suggest that potential temperature control should be used in these type of simulations. Case P, which also does not feature potential temperature control, is actually superimposed to case G on the plot and, for both, the higher numerical dissipation of TOSCA leads to faster increase of the average potential temperature inside the CNBL (these profiles are also took at a later time, as explained in the paper, so the difference is magnified). Finally, case PT is the only one that can reach a truly statistically steady state, as the mean potential temperature profile is maintained constant with the proposed controller. The reviewer is also referred to Fig. 3, where final potential temperature profiles – after their evolution – predicted by different codes are displayed. This test allows, among other things, to understand how dissipative a code is based on how much the ground temperature has increased over time. SP-Wind (the same code that

produced the S2 profile in Fig. 6) is much less dissipative than NCAR-LES and Wire-LES, despite them also being pseudo-spectral codes. Conversely, TOSCA's dissipation is comparable with the latter two even if it employs a finite-volume discretization.

References

Allaerts, D. and Meyers, J.: Boundary-layer development and gravity waves in conventionally neutral wind farms, *Journal of Fluid Mechanics*, 814, 95–130, <https://doi.org/10.1017/jfm.2017.11>, 2017.

University of British Columbia

UBCO-UL NSERC Alliance Grant “Reduced-Order Models of Wind
Farm Induction and Far-Field Wake Recovery”

Response to Reviewer 2

Exec. S. Stipa - November 14, 2023

We would like to thank the reviewer for the time dedicated to revising the paper. We proceed with answering and clarifying, where possible, the proposed comments. Our response, denoted in black, is shown below. Modified text of the paper is shown between quotes in *italic*, while the reviewers' comments are denoted in blue. Please refer to the track changes section at the end of this document for a detailed overview of the changes made to the manuscript.

The authors have provided lengthy responses to my comments and suggestions in addition to an addendum presenting a simple analytical model to explain the origin of global wind farm blockage. However, the question is not about wind plant blockage, but about the reality of gravity wave effects induced by a wind plant. Gravity waves may be induced, but the question is if their effects as pronounced as in simulations that utilize incompressible solvers with Boussinesq approximation. The same question can be extended to the simple analytical mode presented in the addendum.

At the end of the addendum the authors state: “As a concluding remark, wind farm simulations—for example LES—which do not feature a sufficient vertical resolution of the free atmosphere, inevitably fail to capture the full evolution of gravity waves aloft, as there is not enough domain to actually resolve them.” The sufficient height to resolve gravity waves implies that gravity waves are of significant vertical wavelength, i.e., a wavelength that is some large fraction of the depth of the troposphere. The question is if Boussinesq approximation can accurately capture such gravity waves.

We acknowledge that the reviewer remains skeptical about the use of an incompressible code employing the Boussinesq approximation to capture wind farm-induced gravity waves. We will try to allay this skepticism in the present answer.

The Boussinesq (1903) approximation was re-examined by Spiegel and Veronis (1960) for thermal convection problems. In particular, given a generic thermodynamic variable f , the authors introduce the following nomenclature

$$f(x, y, z, t) = f_m + f_0(z) + f'(x, y, z, t) \quad (1)$$

where f_m is the constant space-average of f , f_0 is the variation in the absence of motion and f' are the motion-induced fluctuations. Applying this nomenclature to the potential temperature field θ of a conventionally neutral atmosphere where a generic wind farm is immersed, we identify θ_m as the mean value of θ at the ground, $\theta_0(z)$ as the horizontally averaged vertical excess in temperature w.r.t. θ_m , and $\theta'(x, y, z)$ as the spatial fluctuation in the mean temperature produced by the wind turbines. In the context of the finite wind farm simulations presented in the paper, $\theta_0(z)$ coincides with the horizontally-averaged potential temperature profile in the precursor simulation, while $\theta'(x, y, z)$ is the field obtained by removing θ_m and $\theta_0(z)$ from θ in the successor simulation after it has reached a statistically steady state. In terms of the time-averaged results, the precursor simulation does not induce any variation in $\theta_0(z)$ because the mean velocity field has a null vertical component. For this reason, the wind farm is the only cause of the perturbation field $\theta'(x, y, z)$, and equals zero if no turbines are present. (Here, our reasoning excludes the turbulent fluctuations, as we only focus our attention on steady-state spatial perturbations induced by gravity waves.)

We proceed to assessing that the Boussinesq approximation can be used with incompressible codes to study wind farm generated gravity waves, and even terrain-generated waves, which induce larger perturbations. To do so, we create a map on a G vs γ space, where G is the geostrophic wind speed and γ is the free-atmosphere lapse rate. In particular, for each point on the map, we evaluate the criteria introduced by Spiegel and Veronis (1960) and see if they hold. As more stringent criteria were outlined by Miesen et al. (1988), we will consider the Boussinesq approximation to be valid only if all criteria are verified at the

same time. Specifically, Spiegel and Veronis (1960) require that

$$\frac{\Delta\theta_0}{\theta_m} \ll 1 \quad (2)$$

where $\Delta\theta_0$ is the maximum variation in $\theta_0(z)$ across the domain height. Moreover, for non-linear problems where the flow departs from the initial condition (e.g., convection problems), Spiegel and Veronis (1960) require that

$$\alpha \ll \frac{\Delta\theta_0}{\theta_m} \ll 1 \rightarrow \frac{\alpha\theta_m}{\Delta\theta_0} \ll 1 \quad (3)$$

where $\alpha = \theta'/\theta_m$ is a measure of the wave amplitude. Eq. 3 states that motion-induced fluctuations cannot exceed, in order of magnitude, the static variations. In this context, we consider $\theta'(x,y,z)$ to be the maximum observed value of $|\theta(x,y,z) - \theta_m - \theta_0(z)|$, so that all perturbations are taken with the positive sign. Specifically referring to gravity wave problems, Miesen et al. (1988) further adds that

$$\frac{\lambda_z}{H} \ll 1 \quad (4)$$

where $\lambda_z = 2\pi G / \sqrt{g\gamma/\theta_m}$ is the vertical wavelength of gravity waves and $H = \theta_m\gamma^{-1}$ is the scale height of the problem (Miesen et al., 1988). Moreover, the same authors add a further requirement that the gravity waves' amplitude should satisfy for the approximation to hold, namely

$$\alpha \ll \frac{\lambda_z}{H} \rightarrow \frac{\alpha H}{\lambda_z} \ll 1. \quad (5)$$

As the value of α depends on the solution, for each (G, γ) state we evaluate θ' using a simple 2D case corresponding to the uniform flow over a Gaussian hill that has a height equal to 250 m and half width equal to 500 m. The θ' corresponding to such a scenario is easily obtained using linear theory, and is well documented in Nappo (2012). The domain height is chosen as $2\lambda_z$, i.e. the minimum required to resolve gravity waves within numerical simulations, as one wavelength is allocated to the physical portion of the domain, while the second wavelength corresponds to the Rayleigh damping layer. The use of linear theory to estimate the maximum value of θ' is the only feasible option within the time constraints of the paper review process. We acknowledge that the linear theory itself employs the Boussinesq approximation, and so the solution may not be representative of the real physics when the validity bounds expressed by Eqs. 2 to 5 are approached. Nevertheless, when the left hand sides of Eqs. 2 to 5 are much less than unity, the Boussinesq approximation is justified and results from linear theory should be trustworthy.

Fig. 1 shows the left hand side of Eqs. 2 to 5 on a G, γ space that is representative of atmospheric conditions for wind farm simulations. First, it can be noticed that conditions involving the perturbation amplitude α are critical at very low geostrophic wind speeds (low wavelengths), and mildly depend on the lapse rate, as both θ' and λ_z decrease as γ increases. Conversely, the left hand side expressed by Eqs. 2 and 4 increases with both G and γ , with Spiegel and Veronis (1960) being more restrictive than Miesen et al. (1988).




Figure 1. Value of the left hand side of Eqs. 2 to 5 on the G, γ space. (a) maximum value; (b) individual contributions associated with each condition. Dashed line indicates contour level 0.3.

As a consequence, the combined diagram shows that prohibited regions are towards high values of both G and γ , or toward high values of γ at low geostrophic speeds. Fig. 2 shows the gravity waves patterns for $G = 1, 10, 20$ m/s and $\gamma = 10$ K/km. While both the vertical and horizontal wavelengths increase with increasing geostrophic wind ($\lambda_z \approx 350, 3500$ and 7000 , respectively), the maximum perturbation amplitude is more constant, close to ≈ 2 K. According to Fig. 1, the solution corresponding to $G = 1$ m/s lies outside of the validity bounds for the Boussinesq approximation, as the left-hand side of Eq. 5 is greater than unity. The case corresponding to $G = 20$ m/s could still be treated using the Boussinesq approximation, as the maximum value among the left hand sides of Eqs. 2 to 5 is around 0.4. Finally, the case corresponding to $G = 10$ m/s is well inside the validity bounds.




Figure 2. Gravity wave patterns for (a) $G = 1$ m/s; (b) $G = 10$ m/s; (c) $G = 20$ m/s; $\gamma = 10$ K/km in all cases. Horizontal coordinate normalized with the half width of the hill, b . Vertical coordinate normalized with λ_z .

Focusing now on the finite wind farm simulation presented in the paper, $G = 10.815$ m/s, $\gamma = 1$ K/km ($\lambda_z \approx 11.8$ km). As can be noticed, this lies in a region of Fig. 1a where the Boussinesq approximation is fully applicable, as the maximum among the left hand sides of Eqs. 2 to 5 is less than 0.1. We additionally note that the Boussinesq validity plots reported above are likely to be conservative if applied to wind

farms instead of terrain features. In Fig. 3, we report a comparison between the potential temperature perturbations observed in the finite wind farm simulation presented in the paper and around the hill.




Figure 3. Gravity wave patterns with $G = 10.815$ m/s and $\gamma = 1$ K/km obtained for (a) the hill using linear wave theory and (b) the wind farm using LES. The horizontal dashed line in the Fig. (b) indicates the height of the Rayleigh damping region, i.e. waves are damped moving above this line (the analytical solution does not require vertical damping).

Only focusing on the free atmosphere, the maximum perturbation values of potential temperature are ≈ 0.09 K for the wind farm and ≈ 0.25 K (both negative) for the hill used in the previous part of this analysis (the wind farm perturbs the free atmosphere almost three times less). Finally, one last aspect to mention is that the wind farm simulation also features a capping inversion layer. Here, higher fluctuations in potential temperature are observed (≈ 3 K) but, since the inversion layer is extremely thin (100 m in our study), these do not represent a threat in violating the Boussinesq approximation.

In conclusion, we argue from the above analysis that gravity waves can be studied using incompressible codes following the Boussinesq approximation. Indeed, most of the G , γ variable space can be covered by this approximation. Limiting conditions are represented by very low geostrophic winds (roughly below the cut-in wind speed for modern large wind turbines) and high geostrophic winds combined with a strong free atmosphere lapse rate.

As a last remark, the authors of this paper are aware that attempts of comparing WRF-LES (Weather Research and Forecasting Model in LES mode) results against incompressible codes such as TOSCA or SOWFA are underway at the University of Indiana (Matt Churchfield, Hrishikesh Sivanandan, Mehtab Khan, personal communication). In particular, WRF-LES has been tested on a hill case setup similar to the one presented in this answer, and similar gravity waves patterns were observed (preliminary results of this work have been recently presented as a poster at the NAWEA 2023 Conference in Boulder, Colorado). However, as WRF does not feature anti-reflection measures apart from the vertical Rayleigh damping layer, wave patterns were distorted by likely high wave reflectivity. Additionally, for those WRF studies that use nested grids, we argue that their use could impede gravity waves to propagate from inner to outer grids. Although this might not always occur, it should be considered that the relaxation regions used to drive the boundary conditions of inner grids from outer ones might act as a damping layers, distorting or damping these waves.

The authors have followed my suggestion and moved some of the background material about Numerical Procedure in Appendix A. They also made necessary corrections related to some of the minor comments. However, in a number of cases they did not consider my comments carefully and did not make any modifications to the manuscript. In some cases they misinterpreted my comments. For example, in my review I stated “Wind Energy Science journal may not be appropriate journal for publication of the work primarily focused on model development and idealized simulations.” The key word is “may.” Considering that the manuscript presents a new model and its validation, I do think that the work is worth publishing, the question is whether Wind Energy Science it the right journal and that decision rests with the editor.

On the other hand, the manuscript can be published “only if the numerous issues regarding inaccurate statements about the state of the science addressed and modeling choices better justified.” The revised manuscript still does not meet this threshold.

The papers cited as justification that the manuscript devoted mainly to model development should be published in *Wind Energy Science* are all focused on wind turbine wake models and not on the large-eddy simulation models (i.e. simulation systems).

We agree that the decision about whether or not *Wind Energy Science* is the right journal for publishing the present work rests with the editor. We argue that TOSCA is more than merely a large eddy simulation model, but is a bespoke software specifically developed for wind energy applications, and which we believe will be useful for stimulating research within the wind research community. Moreover, it contains several novel features aimed at addressing research questions that have garnered significant attention in the past few years. We are making TOSCA available for the benefit of stimulating further research and scientific discussion within the community served by *Wind Energy Science*. Some of these beneficial features include

- A new open-source high-fidelity framework tailored to simulate large wind farms with realistic inflow conditions, i.e. with detailed turbulence resolution. The code can potentially simulate large-scale variation in hub-height wind magnitude and potential temperature structure as well. Advanced actuator models with pitch, yaw and angular velocity controller allow to realistically capture the wind farm response to the incoming wind and its interaction with the atmospheric flow.
- A new driving pressure gradient that allows to remove inertial oscillations that are generated from inconsistent initial conditions. There is no way, as explained in the paper, to choose a consistent set of initial conditions if the wind is controlled at the hub height. Hence, these oscillations, although physical in their nature, are just a spurious effect in this class of LES simulations.
- A new hybrid off-line/concurrent precursor method.
- A simulation of a finite-size wind farm (finite both in the spanwise and streamwise direction) under truly statistically steady-state CNBL. The CNBL profile includes wind veer and thermal stratification, while the wind turbine models include angular velocity, pitch and yaw controllers.
- An assessment of large scale blockage effects for a specific CNBL case, showcasing the capabilities of the developed open-source code.

We would submit that the level of dialogue that has been stimulated with this reviewer is evidence that the scientific debate surrounding the relevance of gravity waves to wind farms and the proper methodology for studying them remains important to the field. Tools that contribute to this debate – in spite of their necessary limitations – will be relevant to *Wind Energy Science*. Moreover, we counter the reviewer’s argument that “idealized simulations” are less capable in advancing our understanding of physical processes occurring within and around wind farms. For example, wake models are fast, allowing a broader exploration of the variable space at the price of lower accuracy. LES, depending on the specific case, may offer a detailed picture that is capable of isolating the dominant physical processes. Weather models allow to model a broader physics covering a bigger spatial scale, sacrificing high resolution at the wind turbine scale. Each of these models has its strengths and weaknesses, but advancements in all, we believe, are worth publishing in this journal, especially if tightly related to the wind energy field, which we believe our paper to be.

The authors insist on keeping the statement that SOWFA is not “sufficiently parallel-efficient,” should not be made without providing evidence that TOSCA is comparatively more “parallel-efficient,” however, no comparison in performance between TOSCA and other models is provided. Comparison of the scaling plot presented in Appendix D, with the scaling plot presented in Min et al. (2022, arXiv:2210.00904v1), Figure

5, shows that on 4 Summit nodes (44 cores) CPU-based NekRS with needs about 7 s per time step for 512^3 grid cell simulation and AMR-Wind needs about 5 s per time step, while TOSCA needs more than 20 s per time step for $600 \times 600 \times 200$ grid cell simulation (1.8 times smaller number of cells) on 5 40-core nodes. Based on this comparison it would seem that TOSCA is more than factor of 3 to 4 slower than the NekRS and AMR-Wind. In their reply they state that “We do not want to include a scaling performance comparison in our paper, as a fair comparison would require the same HPC platform and case.” If that is the case then statements, like the one quoted above, “sufficiently parallel-efficient,” that is included in the revised manuscript cannot be supported without direct comparison.

The statement regarding SOWFA not being sufficiently parallel-efficient has been modified in this second revision of the paper. We instead state that OpenFOAM-6, which SOWFA is based on, is “...a general-purpose set of libraries that are not specifically designed to efficiently run at scale.”

Furthermore, we believe that the reviewer is conflating scalability and solver performance. Parallel efficiency is not a measure of absolute solver speed, but rather the relative gain in speed given a certain increase in the number of processors used for the calculation (strong scalability) or the relative loss of performance produced by an increase in job size with a constant number of cells per processor (weak scalability). As it is demonstrated in Appendix D, TOSCA is parallel-efficient, i.e. its performance does not degrade – up to a certain limit, like any other solver – when the number of used processors or the problem size increase.

Regarding differences in solver speed, these can attributed to differences in the numerical procedures. SOWFA is based on OpenFOAM, which is an unstructured solver. This inherently requires to perform more operations to form the discretized system of equations (mainly to access mesh information through the connectivity data structure), and requires complex processor communication at processor interfaces. Hence, while SOWFA certainly is a fast and flexible solver in general, it is inherently less efficient when running at scale, i.e. with thousands of processors, than codes specifically deigned to run on HPC architectures. In curvilinear codes like TOSCA or rectilinear codes like AMR Wind, mesh elements and quantities of interest are directly accessed indexing 3D arrays, so there is no need for connectivity, nor for send/receive operations that involve more than a few processors. Additionally, if wind energy applications are considered, unstructured meshes are not the optimal choice as the majority of the domain is usually represented by the atmosphere. This means that a structured grid provides a faster and more accurate solution (hence the NREL choice to use AMR Wind as their atmospheric solver). Moreover, as explained in the paper, in order to write output fields, OpenFOAM/SOWFA generate a number of directories equal to the number of processors. When dealing with thousands of cores, the number of produced files increases rapidly, and may quickly saturate the maximum file counts allowed on HPC systems.

Regarding AMR Wind and NekRS, the reviewer correctly mentioned that on an absolute level, TOSCA may be slower. However, the reviewer should consider that TOSCA uses an implicit matrix free momentum solution. This yields a higher time per iteration than explicit methods, but allows larger time steps with CFL values that can exceed 2. Hence, it is still not clear how much and if TOSCA is slower than AMR Wind and NekRS. Specific comparisons on the same system solving the same case would be required in this sense. Another aspect to mention is that AMR Wind and NekRS are specifically designed for speed and parallel scalability and, regarding AMR Wind, it is still not clear how flexible this code is. Moreover, TOSCA has features that the two solvers do not have, namely the concurrent precursor method, which is a unique feature among existing finite volume codes, and the immersed boundary method to model complex terrains. Despite being already present in the open-source repository and working both in dynamic (for fully resolved turbine cases) and static (terrains or objects) modes, the latter has not been included in the present paper but, as mentioned, will be the focus of a follow up study.

The authors state that “Our choice to simulate a non-evolving flow (mainly in terms of BL height) is explained in the paper, and we believe that this is not more idealized than running a precursor for big

amount of time without temperature control, as in reality CNBLs are only observed for short lapses of time, especially in transition between night and day.” Perhaps this is a question of semantic, about what is meant by idealized. However, on the spectrum of idealized to realistic the simulations presented in the manuscript are leaning heavily to idealized. In the review of the original manuscript, I provided a number of references related to numerical simulations of wind turbines and wind plants under realistic atmospheric conditions. The issue is that the simulations of presented in the manuscript are idealized simulations and this should be clearly stated. I fully agree with the authors argument that “adopting the proposed controlling methodology” is beneficial for model comparison, however, this is still very different from simulating realistic atmospheric inflow conditions.

Fair enough. In this second revision, we modified the word “realistic” to “time-resolved”.

The authors state that “prior to our study, resolving gravity waves in an LES of a wind farm has only been achieved by the group of Meyers.” However, the question is what is meant by resolving gravity waves, are they resolved correctly – their wave length and magnitude, are the gravity waves real or a numerical artifact? The authors insistence in application of Boussinesq approximation to the domain that extends in vertical direction to more than 10 km is in conflict with assumptions approximation is based on. For example, Wood and Bushby (2016): “Specifically, the Boussinesq equations (e.g., Spiegel & Veronis 1960) are valid only for small perturbations to the thermodynamic variables in systems with small vertical length scales (in particular, the domain height must be much smaller than all of the thermodynamic scale heights).” Furthermore, Lilly (1996) states: “The principal defect of both systems is their inability to conveniently allow for arbitrary mean thermodynamic profiles. They both assume a reference state with neutral static stability. A non-neutral state can be imposed, but that represents a perturbation on the reference state which produces increasing error with layer depth.” The authors use as a reference from which density perturbations are computed, in Equation 4 “ θ_0 is a reference potential temperature, chosen as the ground temperature,” i.e., essentially a neutral reference state. Considering that this is the reference state, that the domain extends to 28 km in vertical, and that imposed stratification is 1 K/km this means that the perturbation through the troposphere and stratosphere the perturbation exceeds 10% of its reference value. Miesen et al. (1988, *Physica Scripta*) provides an analysis of the limits to application of the Boussinesq approximation for gravity waves in the atmosphere.

Again, the question is: can the magnitude and wavelength of gravity waves induced by a wind plant can be accurately captured with an incompressible solved based on Boussinesq approximation? In their response to my original comment about the use of Boussinesq approximation the authors provided a list of papers using this approximation. They state: “Regarding the results from Smith, we highlight that they employ linear gravity wave theory, which also features the Boussinesq approximation. Several books (Nappo 2012, Lin 2007) or articles (Allaerts and Meyers 2019, Smith 2002, Smith 2006, Smith 2010) use linear theory to model free atmosphere gravity waves, where the latter are produced by vertical disturbances in the flow generated by terrain features or wind farms.” Furthermore, the authors state that: “In particular, in order for the latter to hold [Boussinesq approximation], the following hypotheses should be verified (Lin 2007): (1) the vertical dimension of the fluid motion is much less than any scale height and (2) the motion-induced fluctuations in density and pressure do not exceed, in order of magnitude, the total static variations of these quantities. These hypotheses are verified, as wind farm gravity waves only slightly perturb the pre-existing equilibrium, inducing extremely small motions.” It is not clear where these hypotheses are verified, i.e. what is the ratio of vertical wave length to characteristic height, also since the perturbation in density is proportional to potential temperature perturbation and the reference temperature is ground temperature the “density” (i.e. temperature) perturbation is approximately 10% of the reference state. In the view of the analysis presented by Miesen et al. (1988) the question is if the Boussinesq approximation can be applied to this problem. The authors should show that the approximation is within the limits of the assumptions that the approximation is based on, or outside of its range of applicability. In particular in their simulations the ratio of density (i.e. temperature perturbation) to reference density (i.e. temperature) is 0.1 as indicated above and it is directly proportional to the relative amplitude of the gravity wave. At the same time the

ratio of vertical wave length to the characteristic height (in their reply the authors stat that “the scale height of the atmosphere [is], typically around 10 km) is likely larger than 0.1. However, the analysis by Miesen et al. (1988) shows that, for Boussinesq approximation to be applicable to analyses of atmospheric gravity waves, the relative magnitude of the gravity wave must be much smaller than the ratio of vertical wave length and characteristic height which in turns must be much less than one. In their reply the authors indicate that the vertical wavelength of gravity waves in their simulation is about 12 km “(around 12 in our study)” which means that the ratio of the vertical wave length to the characteristic length is larger than one.

Although this issue has already been addressed in our first answer, we apply the analysis proposed by Spiegel and Veronis (1960) specifically to our simulations to verify the appropriateness of the Boussinesq approximation by checking that both $\Delta\theta_0/\theta_m \ll 1$ and $\alpha \ll \Delta\theta_0/\theta_m$. The first inequality can be easily verified considering that the lapse rate is 1 K/km, and that the domain is 28 km height, yielding $\Delta\theta_0/\theta_m = 28/300 \approx 0.1 \ll 1$. The second inequality, considering that $\alpha = \theta'/\theta_m = 0.09/300$, where θ' are the maximum magnitude of the motion-induced fluctuations shown in Fig. 3, yields $3 \cdot 10^{-4} \ll \Delta\theta_0/\theta_m \approx 0.1$, which is also verified. Regarding the analysis proposed by Miesen et al. (1988), this requires $\lambda_z/H \ll 1$ and $\alpha \ll \lambda_z/H$, where λ_z is the gravity wave vertical wavelength (definition is given in the first answer), ≈ 12 km in our study, and H is the problem scale height. In our previous answer we used the atmosphere scale height, which is around 10 km, to show that the motion-induced fluctuations in boundary layer displacement η (on the order of 10 m near the capping inversion layer) are at least three orders of magnitude smaller, satisfying the condition that the motion happens in a layer whose thickness is much smaller than the atmospheric scale height. This is in fact another way of looking at the second hypothesis of Spiegel and Veronis (1960). Alternatively, in light of the analysis presented by Miesen et al. (1988), the problem scale height should be used, which is physically and numerically meaningful for the specific problem being simulated and input parameters chosen, i.e. $H = \theta_m/\gamma \approx 300$ km (this scale height is the one used for formal analysis by Miesen). Hence, $\lambda_z/H \approx 0.04 \ll 1$, and $\alpha = 3 \cdot 10^{-4} \ll \lambda_z/H \approx 0.04$ are both verified.

Finally, in relation to capturing gravity waves the authors state: “We believe that WRF would face the same issues in a simulation where the background stratification is steady and the simulation is carried out for a sufficient amount of time. In fact, as pointed out by Klemp and Lilly 1978, wave reflectivity increases over time for a given steady background state. It may be that energy accumulation due to wave reflection might represent a smaller issue for those WRF simulations where the background flow is evolving. These issues are worth investigating, but we feel fall outside the scope of the present paper.” I agree that these issues fall outside the scope in terms of analysis, but they do not fall out of scope when the statements are made about the impacts that may not be realistic. Authors’ statement points potentially to the essence of the question related to how accurately gravity waves are captured in simulations with incompressible solver and Boussinesq approximation. As pointed out in the review of the original manuscript the simulations with stationary conditions that extends for 20 hours is an idealized simulation and in authors’ own words “It may be that energy accumulation due to wave reflection might represent a smaller issue for those WRF simulations where the background flow is evolving” while in their simulation it may result in unrealistic gravity waves.

While wind farm induced gravity waves are not a numerical artifact, as has been addressed in our first answer, it still remains to be seen what is their effect under evolving atmospheric conditions, mainly in terms of wind profile and potential temperature structure. We are currently leveraging TOSCA to simulate a full diurnal cycle using mean temperature and velocity forcing from ERA5 reanalysis data. The objective is to understand if and how results obtained from a temporal average of an evolving flow differ from results obtained by running a statistically steady state simulation forced by the same profiles averaged during the diurnal cycle. We re-emphasize that, although extremely important, such answer is not straightforward and requires dedicated analyses that fall outside of the present paper.

Regarding the WRF simulations mentioned by the authors which did not produce any gravity waves

(Sanchez Gomez et al., 2023) we note, as explained in our previous response, that these simulations are not properly setup to resolve gravity waves, as the domain height is less than $2\lambda_z$. The authors claim that they still observe global blockage without gravity waves. In fact, three forms of blockage are present around a wind farm:

1. Local blockage produced by individual turbine induction: this can be observed without modeling thermal stratification.
2. Global blockage produced by the flow confinement under the capping inversion layer, assuming that this is not perturbed (rigid lid approximation): this can be observed without modeling thermal stratification, i.e. placing the top boundary at the inversion height or with thermal stratification. In the latter case, if the domain is such that gravity waves are not resolved, the actual perturbation of the inversion layer will be incorrect, as it depends solely on gravity waves. Nevertheless, the large scale blockage might not be, in some cases, too far from the actual blockage observed when modeling gravity waves. Its perturbation though, i.e. beneficial/detrimental effects inside the wind farm, will be inaccurate.
3. Gravity wave blockage: this is produced by the vertical perturbation of the inversion layer w.r.t its equilibrium/freestream height. They are related to the physics of gravity waves aloft and interface waves within the inversion layer. The effect of this waves can be detrimental upstream and more or less beneficial inside the farm, hence it should be modeled.

If a simulation is conducted such that that it cannot capture gravity waves, only the first two contributions to the blockage will be observed, while the third becomes a result of the specific numerical setup. In fact, it may be affected by reflections from the side or top boundaries or, when using nested grids, gravity waves might be directly killed by relaxation regions outside of the innermost grid, where they are generated.

In their reply the authors state that: “On the same note, we would also like to highlight that the numerical implementation that is necessary to drive LES simulation with realistic mesoscale information is essentially the same that we propose, with the only difference that the average reference state changes in time (Allaerts at al. 2020, Allaerts at al. 2023).” They still ignore a significant recent body of work dealing with realistic mesoscale forcing through coupled mesoscale microscale simulations referenced in Haupt et al. 2023.

The referee correctly points out that, in addition to the "profile-assimilation" technique, there exists the "boundary-coupled" simulations in order to provide realistic mesoscale forcing. In the second revision, we have tried to improve the literature review such that it points out the importance of meso-micro scale coupling.

We are currently working on implementing profile-assimilation techniques within TOSCA. Here, instead of a uniform driving pressure gradient and a steady-state potential temperature profile, we would just have to render these quantities time-dependent, possibly including wind height dependency (temperature height dependency is already present). On the boundary-coupled side, TOSCA already features overset-mesh capabilities, so the structure to support nested domains already exists in TOSCA. Similarly to WRF, we use a relaxation region to drive the simulation results from the outer to inner grid. On the other hand, TOSCA does not currently feature a PBL scheme, required to run computations in the mesoscale domain with a grid spacing that falls above the *terra-incognita*. This is certainly a very interesting opportunity in the direction of extending TOSCA's capabilities, and we have to thank the reviewer for pointing it out.

There are a number of instances where the authors did not clarify the statements in the manuscript, but instead provided arguments that frequently are either incomplete or not accurate, for example, the authors claim that (line 231 in the original manuscript, line 823 in the revised manuscript) “If the mesh is uniform, the filtering operation on the face area vectors has no effect. Hence the equality holds exactly. If it is

stretched, this is an approximation as stated.” It is not clear how this statement can be supported, filtered product is different than product of filtered quantities. For example, if a filter is idempotent, i.e., yields the same result if it is applied once or multiple times, then the equality in line 822 can hold only if the difference between unfiltered and filtered quantity is exactly zero. Clearly, if equality really holds and it is not just an approximation than this needs an explanation.

We agree to the referee’s argument that filtered product is different than product of filtered quantities in general. However, in the case of uniform meshes, face area vectors are constant along each direction. For homogeneous LES filters (as in the case of TOSCA), the value of a filtered constant must be equal to the constant. For this reason, the filtering operation has no effect on the value of the filtered face area vector $\overline{S_j^k}$. Applying this to the product $S_j^k S_{ij}$, we can write

$$\overline{S_j^k S_{ij}} = \overline{S_j^k} \overline{S_{ij}} = S_j^k \overline{S_{ij}} \quad (6)$$

where $\overline{S_j^k}$ is identically equal to S_j^k since we cannot compute face area vectors at a resolution smaller than the mesh size. When the mesh is stretched, Eq. 6 does not hold exactly anymore, as face area vectors are now varying along curvilinear directions. Nevertheless, the tilde LES filter has a size of 3 mesh cells in each direction, hence $\overline{S_j^k}$ (face area vectors at the central cell) and $\overline{\overline{S_j^k}}$ (filtered face area vectors within the box) are almost identical provided that mesh grading is smooth enough. For instance, grid stretching over three cells is usually imperceptible, as rapid grid stretching introduces other sources of problems such as diffusion and numerical instabilities when using e.g. second-order centered schemes.

Another example is Marjanovic et al. (2017) reference that they consider of no relevance, however, they did not do a thorough literature search and missed Marjanovic et al. “Implementation of a generalized actuator line model for wind turbine parameterization in the Weather Research and Forecasting model” J. Renewable Sustainable Energy 9, 063308 (2017).

We still believe that the mentioned paper is irrelevant to establish whether or not gravity waves are observed in WRF. A single wind turbine does not produce a momentum deficit that is strong enough to be felt at the inversion height. Hence, as the capping layer is not displaced vertically, wind farm induced gravity waves will not be observed. Moreover, the following sentence mentioned in the paper is incorrect: "Forcing idealized LES with geostrophic wind speed and direction results in inertial oscillations which alter boundary layer characteristics over the first several hours of a simulation. For the TWICS cases (CBL and NBL), a spin-up LES is run for 15 h with the surface heat flux specified (20 and 0 W/m^2 for weakly convective and neutral, respectively) to allow the solution to come into balance with the geostrophic wind vector". How can the solution come into balance in 15 h within the free atmosphere if the inertial oscillation period – considering the latitude of Boulder – is roughly 37 h, and the geostrophic equations are the one of an un-damped oscillator? There is no physical term in the governing equations (apart from some additional damping term, like the one proposed in our paper) that can operate in the sense of damping these inertial oscillations. We clearly show in our paper that these oscillations are almost undamped and that they go on indefinitely for a classic precursor simulation setup.

The authors also state that: “Random perturbations do not satisfy continuity and would be killed by the Poisson iteration. We follow the same approach used in SOWFA, where the introduced oscillations are sinusoidal and divergence free.” However, random solenoidal field can be generated.

That is correct, but we do not understand the relevance of the comment. Whatever type of initial perturbations can be used provided that they are solenoidal. In fact, these are just used to trigger turbulence in the precursor and their presence is lost after a few recycling time. In particular, the final ABL state does not –

and should not – depend on the type of used turbulence-triggering method.

Finally, based on Table 3 it is still not clear how cases G and S2 differ. This needs to be clearly stated. Based on Table 3 these two cases are identical, this can be confusing for a reader.

Thanks for pointing this out, the caption has been modified. The two cases have in fact the same set-up, but case S2 is from Allaerts and Meyers (2017). Specifically, the difference between these two cases consists in the code used to run them. Moreover, when comparing case S2 with the other cases (Figure 6 in the paper), it should be noted that data from case S2 have been averaged at an earlier stage of the simulation (we did not possess the time history, so we used their published data). For this reason, differences are not only due to the code, but also to the fact that the ABL has evolved less. This is exactly one of the messages that our paper tries to convey: when running idealized CNBLs used to feed wind farm successor simulations aiming at reaching a statistically steady state, results depend on the simulated time if both temperature and velocity are not controlled in the mean. Such drift is not physical and, since simulations are already idealized as the reviewer states, it should be avoided.

References

- Allaerts, D. and Meyers, J.: Boundary-layer development and gravity waves in conventionally neutral wind farms, *Journal of Fluid Mechanics*, 814, 95–130, <https://doi.org/10.1017/jfm.2017.11>, 2017.
- Boussinesq, J.: *Théorie analytique de la chaleur.*, (Paris: Gathier-Villars), 2, 1903.
- Miesen, R. H. M., Kamp, L. P. J., and Sluijter, F. W.: On the application of the Boussinesq approximation for nonlinear gravity waves in the atmosphere, *Physica Scripta*, 38, 857, <https://doi.org/10.1088/0031-8949/38/6/018>, 1988.
- Nappo, C. J., ed.: Copyright, vol. 102 of *International Geophysics*, Academic Press, <https://doi.org/10.1016/B978-0-12-385223-6.00014-8>, 2012.
- Sanchez Gomez, M., Lundquist, J. K., Mirocha, J. D., and Arthur, R. S.: Investigating the physical mechanisms that modify wind plant blockage in stable boundary layers, *Wind Energy Science*, 8, 1049–1069, <https://doi.org/10.5194/wes-8-1049-2023>, 2023.
- Spiegel, E. A. and Veronis, G.: On the Boussinesq Approximation for a Compressible Fluid., *The Astrophysical Journal*, 131, 442, URL <https://api.semanticscholar.org/CorpusID:121803237>, 1960.

TOSCA - An Open-Source Finite-Volume LES Environment for Wind Farm Flows

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

¹University of British Columbia, Okanagan Campus, CA

²Delft University of Technology, NL

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

Abstract. The growing number and 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 on 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 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 ~~correct~~ eliminate the dependency of the potential temperature profile evolution on the code architecture observed in previous studies, we ~~propose a method~~ introduce a method that allows 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 ~~overcomes these issues~~ allows 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 are also presented to assess TOSCA's ability to model large-scale wind farm flows accurately and with high parallel efficiency.

20 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/farm induction (Bleeg et al., 2018), is defined as the wind

25 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 gained importance only recently (Lundquist et al., 2019; Ahsbahs 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
30 important aspect to account for and model (Nygaard et al., 2020).

Turbine-level induction has been researched for many years (Troldborg 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
35 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.

40 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
45 negative impact on all companies' financial estimates.

Reduced-order models need to be thoroughly validated, but comprehensive observation datasets are difficult to obtain. Nu-
merical analyses, in particular large eddy simulations (LES), 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
50 of scales involved, spanning from resolved turbulence eddies of a few meters, to gravity waves characterized by wavelengths
of several kilometers. ~~In addition, many numerical aspects have to be carefully treated, such as wave reflections produced by
the domain boundaries, and realistic inflow turbulence generation.~~

~~In order to efficiently and accurately simulate the flow around a finite-size wind farm under thermal stratification, an LES
solver. To successfully tackle these problems, LES solvers must possess good parallel efficiency -and an optimized code input-
55 output (I/O), and a method for defining initial and boundary conditions that accurately reflect the spatio-temporal ABL state.
This includes realistic inflow turbulence modeling, and a system to avoid gravity wave reflections at the physical boundaries
-These 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)~~

60 The last two tasks can be achieved ~~at once 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 waves 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
65 solutions matching at the fringe region exit. More details on precursor techniques are given in Sec. 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 PALM model (Maronga et al., 2015), developed
70 by the Institute of Meteorology and Climatology at Leibniz Universität of Hannover (Germany), or SOWFA (the Simula-
tor 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 waves effects at the same time avoiding inlet/out-
let 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 ~~not sufficiently parallel efficient when running based on~~
75 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 ~~cores, as processors~~, the number of produced files
~~increases drastically with processor count. While some of these shortcomings have been addressed and solved in the NREL~~
~~can easily saturate, within a few checkpoint iterations, the maximum file count on many HPC architectures. To address these~~
80 shortcomings, the NREL has started the Exawind project (Min et al., 2022), but the latter is not yet at a ~~production stage.~~
~~Moreover, all SOWFA, PALM and Exawind platforms do final production stage, it does~~ not feature the concurrent precursor
technique ~~, making TOSCA the only finite-volume open-source code to possess such capability nor a method to model complex~~
terrains.

85 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 OpenMPI (Gabriel et al., 2004), PETSc (Balay et al., 2022), HYPRE (Fal-
90 gout and Yang, 2002) and HDF5 (The HDF Group, 2000-2010) for the parallel solution of partial differential equations and
handling of intense I/O operations. TOSCA is specifically designed to enable LES simulations of large finite wind farms ~~under~~
~~realistic turbulence inflow and thermal stratification~~. 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 method-

95 ology 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 extensively used in pseudo-spectral methods (Wu and Porté-Agel, 2017; Allaerts and Meyers, 2017, 2018). For this reason, gravity waves 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
100 a uniform grid resolution, leading to high cell counts. Conversely, the finite-volume method allows for grid stretching, enabling 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 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 Sec. 2, we describe the developed LES framework. Next, Sec. 3 presents
105 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 Sec. 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 Sec. 5, we present the simulated flow field around a reference
110 100-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 Sec. 6.

2 Methodology

TOSCA is a finite-volume code, formulated in generalized curvilinear coordinates, allowing it to take as input also non-Cartesian structured meshes. The present section is organized as follows. ~~We first report the~~ The governing equations in cartesian coordinates Cartesian coordinates are reported in Sec. 2.1. ~~Wind turbines are represented using actuators models, while actuator models used to represent wind turbines in the domain are~~ described in Sec. 2.2. ~~but can also be fully resolved, together with complex terrain, through a sharp-interface immersed boundary method (IBM) based on Haji Mohammadi et al. (2019), to be detailed in a follow-up paper.~~
115 ;

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 App. A, while the LES turbulence
120 model in the curvilinear frame is detailed in App. B.

An overview of TOSCA’s parallel efficiency is given in Appendix D, where we analyze the time per iteration with increasing number of nodes and mesh elements on the Niagara (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 cluster, demonstrating its capability to
125 handle massively-parallel computations.

In order to run ABL simulations, we developed a novel methodology, described in Sec. 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 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, Sec. 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

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

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

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

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 gas specific constant, c_p is the specific heat at constant pressure and p_0 is the reference pressure), g_i is the gravitational acceleration vector, 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 App. 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 velocity controller, and it is explained in Sec. 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 (AL), actuator disk (AD), and uniform actuator disk (UAD) models. Following (Sørensen and Shen, 2002; Sørensen et al., 2015;

Porté-Agel et al., 2010; Jimenez et al., 2007), the first two models require detailed blade information (i.e. airfoils, twist and chord), while the UAD only requires 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 AL and AD models 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 two 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 AL model, each rotor blade is represented by a line of points, which are physically rotated at each iteration, making it an unsteady model. In AD and UAD models, 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 AL and AD models, the point force is calculated exploiting the blade element 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, 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 radius because the airfoil type usually changes 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, then use the momentum theory to infer the corresponding freestream velocity (Meyers and Meneveau, 2010). In our framework, since AD and AL models 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

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

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 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 to prescribe both magnitude and direction at a specified height h_{ref} . In particular, given a desired velocity $u_{ref,i}$, 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 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

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

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

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

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 to remove these inertial oscillations, enabling the use of the pressure controller also in those cases where a steady state geostrophic wind is preferred. First, we note that the system of equations 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 (Eq. 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 one-tenth 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} inside the BL. 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

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

255 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, depend on the mixing history experienced inside the boundary layer. This is in turn affected by the specific LES
 260 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-spectra codes, such as NCAR-LES (Pedersen et al., 2014) or Wire-LES (?), which use for example a second-order central scheme in the vertical direction (this can be appreciated in Fig. 3). Besides,
 265 in finite-volume codes, like TOSCA or SOWFA, an upwind-biased advection scheme is usually preferred, as it stabilizes the numerical method, but it does not allow 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,
 270 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
 275 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 LES simulations **with realistic**
 280 **mesoscale information from mesoscale** [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 set-up 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 fringe region, must ensure that turbine wakes are ~~removed to avoid~~ prevented from being re-introduced into the domain by periodic boundaries, and that ~~a realistic~~ 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 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 to avoid wave reflections, while prescribing a ~~realistic~~ 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—both in the spanwise and vertical directions, it is usually oversized from the turbulence generation point of view. This leads to a considerable amount of computational resources being consumed when starting up the unperturbed ABL. In fact, domains of much smaller size are used in literature when the sole ABL is of interest, or when the inflow data is 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 both in the streamwise and 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.

315 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 ten highest cells, and slowly merge the instantaneous data to this
 320 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 ten highest cells.

After the concurrent precursor has run for one flow-through time, the turbulent inflow has reached the outlet, and stream-
 325 wise boundary conditions are switched to periodic. At this point, the simulation is self-sustained and we run 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 Sec. 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 105000 s to 120000 s.

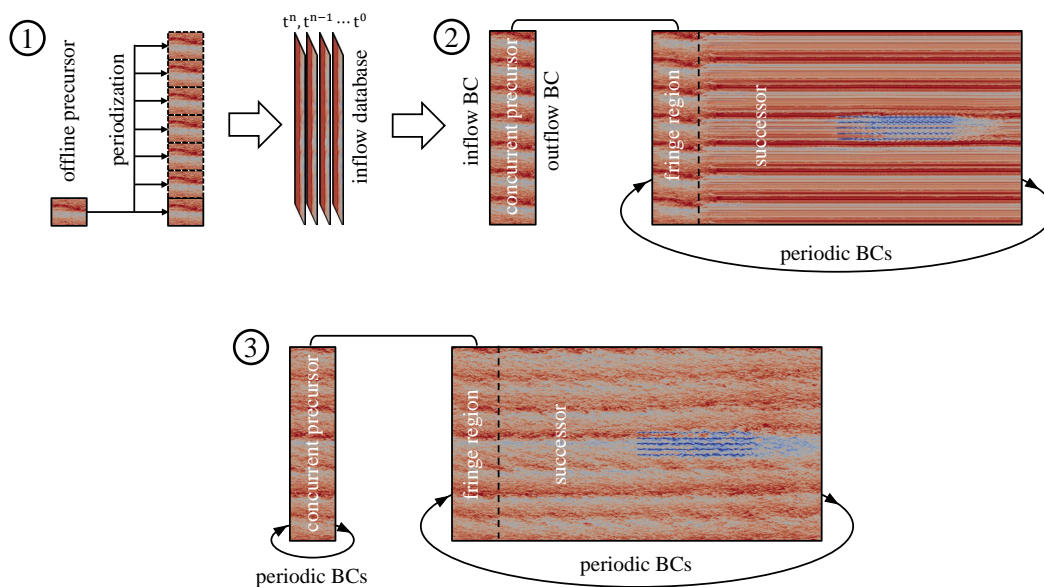


Figure 1. Sketch of the hybrid off-line/concurrent precursor method.

330 The drawback of such method is that a spanwise periodicity is introduced in the concurrent-precursor and successor 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
335 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 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
340 later simulation stages.

3 Validation

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

3.1 Isolated Rotor in Uniform Inflow

In this validation case, we perform two simulations of the NREL 5MW reference wind turbine (Jonkman et al., 2009) using
350 the ADM and the ALM techniques, with a uniform inflow velocity of 8 m/s. 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
355 extends beyond the edge of the rotor for 1 diameter.




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

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 in the radial and azimuthal direction. 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. Fig. 2 shows normalized wind speed deficit at 1 and 4 downstream rotor diameters for ALM and ADM simulations performed with both 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 was 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 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 Tab. 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).

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

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.

The ADM matches well with results from Martínez-Tossas et al. (2015), while TOSCA's ALM predicts a slightly higher power. ~~We attribute such differences to how the velocity is sampled at the actuator points. In TOSCA, nearest neighbor interpolation from the background mesh is adopted, which can result in small differences in the sampled wind, as the actuator point does not coincide with the closest cell center in general. Moreover, as pointed out by Churchfield et al. (2017), failing to sample the velocity at the geometric center of the projection function would determine self-induction at the sampling location. Using for example a linear interpolation at the actuator point would probably solve such a mismatch. The reason for such difference is presently unknown to the authors.~~

380 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 K/km and 10 K/km. In both validation cases, the geostrophic wind is $G = 10$ m/s, the Coriolis parameter is set to $f_c = 10^{-4} s^{-1}$, the surface roughness 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 (?), and NCAR-LES (Pedersen et al., 2014), which all use a pseudo-spectral horizontal discretization.

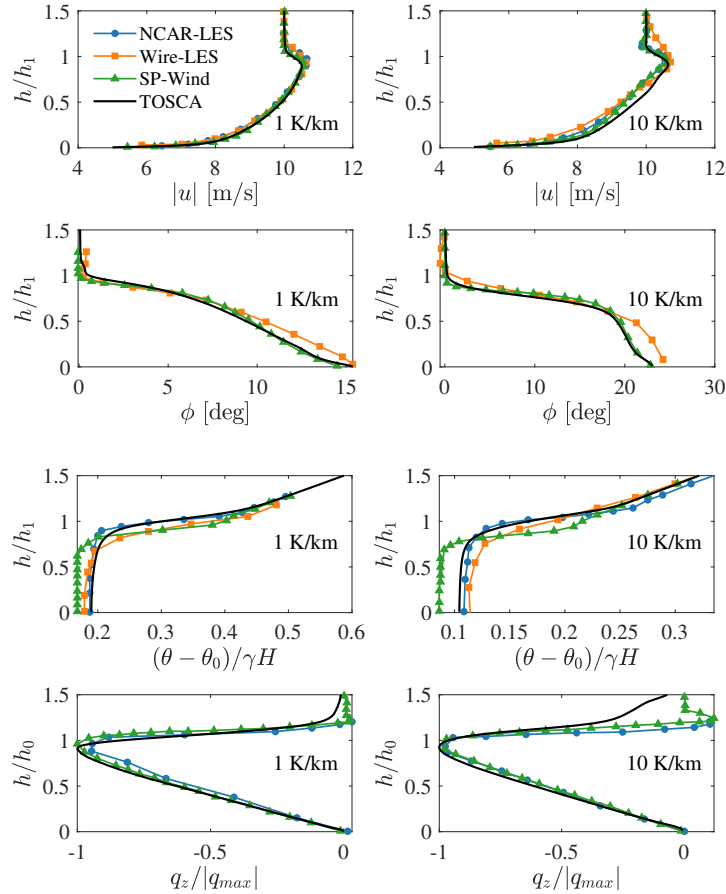


Figure 3. Vertical profiles averaged over the last simulation hour of, from top to bottom, velocity magnitude, wind angle, potential temperature and heat flux. The initial lapse rate is (left) 1 K/km and (right) 10 K/km.

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
390 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
395 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. Fig. 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); ? 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.

400 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 for the 10 K/km case. We do not have a clear explanation for such behavior.

In Tab. 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

405 of the capping inversion layer. Both vertical profiles of Fig. 3 and quantitative ABL parameters reported in Tab. 2 demonstrate that TOSCA is well aligned with results from Pedersen et al. (2014); ?, thus capable of conducting CNBL simulations.

	γ [K/km]	θ_0 [K]	$\Delta\theta$ [K]	Δh [m]	H [m]	u^* [m/s]	$q_{min}/10^{-4}$ [Km/s]
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

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

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-

410 conservative advection schemes.

3.3 Infinite Wind Farm in Neutral Conditions

In this section, we run the same infinite wind farm simulation that has been conducted in Stevens et al. (2018), corresponding to the wind tunnel experiments performed by ?. 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,

415 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 . 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

420 value of $C_{t,max}$ from their reported BEM calculations. Therefore, since it wouldn't 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 to the p th actuator disk point, $\hat{\mathbf{e}}_t$ is a vector normal to the rotor disk, 425 pointing in the upstream direction, and U_∞ is evaluated from the average disk velocity U_{disk} exploiting the momentum theory. For instance, $U_\infty = U_{disk}/(1-a)$ 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.

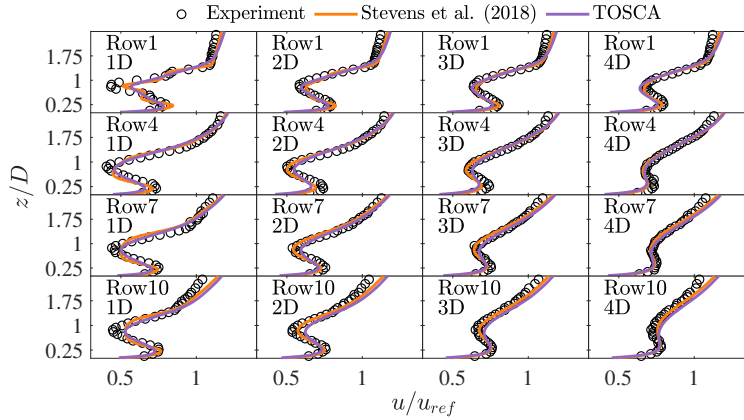


Figure 4. Velocity profiles made non-dimensional with the hub-height velocity for rows 1, 4, 7, and 10, averaged on the three columns. For each row, the wake evolution is reported at 1,2,3 and 4 diameters downstream.

430 The wind tunnel model used in the experiment by ? is the GWS/EP-6030 turbine, 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. 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 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 435 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 Sec. 2.4. The precursor domain is 1.8 m \times 1.8 m \times 0.675 m, with $129 \times 129 \times 145$ cells in each direction in order to match their cell size. The flow is driven by the

440 pressure controller described in Sec. 2.3.1, with a desired flow velocity u_{ref} of 3 m/s 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 of Stevens et al. (2018), with an equivalent roughness height of 0.03 mm.

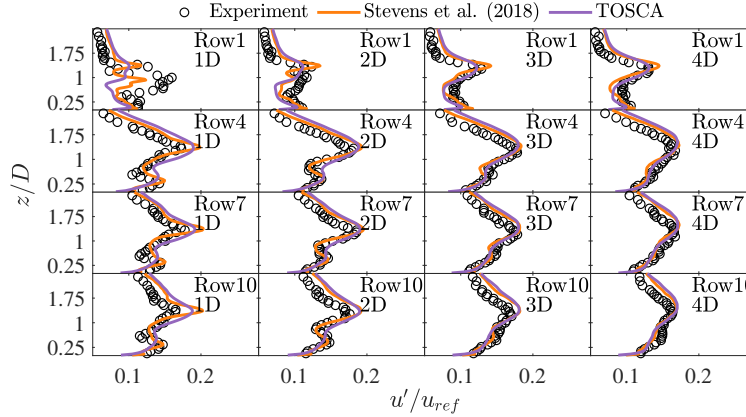


Figure 5. Profiles of streamwise fluctuations non-dimensionalized with hub-height velocity for rows 1, 4, 7, and 10, averaged on the three columns. For each row, the wake evolution is reported at 1,2,3 and 4 diameters downstream.

The precursor is run for 100 s (corresponding to ≈ 160 flow through times), after which we saved the inflow field at each time
 445 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 on 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 $588 \times$
 129×145 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
 450 time (≈ 3 s).

In Fig. 4 we show the velocity profiles, averaged among the wind farm columns, for rows 1, 4, 7, and 10, together with
 experimental data from ? 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
 455 area blockage in the numerical domain, which causes the flow to accelerate close to the upper boundary in order to conserve
 mass. In the 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 of 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
 460 et al. (2018), both predicting higher fluctuations than experiments in the top-most 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, both in the mean and in the fluctuations, making it suitable for the simulation of wind turbines immersed in a turbulent boundary layer.

4 CNBL Simulations with Different Controllers

465 In this section, we present CNBL results obtained using the different velocity and temperature controllers described in Sec. 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 Tab. 3. The simulations employ periodic boundary conditions in the horizontal directions and a slip boundary condition at the upper boundary. Classic
470 Monin-Obukhov similarity theory is enforced at the ground.

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

Following Allaerts and Meyers (2017), the domain size is $9.6 \text{ km} \times 4.8 \text{ km} \times 1.5 \text{ km}$ in the streamwise, spanwise and vertical directions respectively, discretized using $320 \times 320 \times 300$ cells in each direction. Case G is forced with a geostrophic wind speed of 12 m/s, matching the setup used by Allaerts and Meyers (2017) in case S2. Conversely, in P and PT cases the pressure controller aims to maintain a wind speed of 10.871 m/s at $h_{\text{ref}} = 100 \text{ m}$, thus matching the hub height wind speed
475 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 Tab. 4.

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

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).

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

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

480 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.




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 92800 s to 100000 s, while in Allaerts and Meyers (2017) (case S2) data are averaged from 54000 s to 72000 s.

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 can't 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 54000 s to 72000 s (over the last five hours of simulated time). Geostrophic damping in cases P and PT starts at $T_D = 6 \cdot 10^4$ s (this value is close to the oscillation period of the geostrophic wind), and it will be later shown the wind angle controller of case G stabilizes the wind angle at around $6.5 \cdot 10^4$ s. Hence, we average flow statistics from cases G, P and PT from 92800 s to 100000 s, while in Allaerts and Meyers (2017) results are averaged from 54000 s to 72000 s. In Fig. 6, we report vertical profiles of velocity magnitude, direction, shear stress and potential temperature obtained from the four different case.

As shown before in Sec. 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 obtained profile of potential temperature. 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 between case G and S2. For cases P and PT such parameter is an input, and it has been set according to results from case G. In Tab. 5, output quantities extracted from the four different simulations are reported, averaging flow statistics in the above-mentioned 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.

	u_{ref} [m/s]	G [m/s]	θ_0 [K]	$\Delta\theta$ [K]	H [m]	Δh [m]	u^* [m/s]	$q_{min}/10^{-4}$ [Km/s]	ϕ_G [deg]
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	-	-

Table 5. ABL parameters obtained by fitting the Rampanelli and Zardi (2004) model for the CNBL cases presented in this section, together with resulting friction velocity, minimum heat flux and geostrophic wind angle. Data from cases G, P, and PT are averaged from 92800 s to 100000 s, while in Allaerts and Meyers (2017) (case S2) data are averaged from 54000 s to 72000 s.

Fig. 6, together with quantitative data reported in Tab. 5, demonstrate 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 $G = 12$ m/s. Fig. 6 also highlights how sensitive the 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 time scale of the diurnal cycle instead. As a consequence, since such temperature drift is physical but arises from an idealization, we believe that ~~a better—and more reproducible—way of studying wind farms under specific atmospheric conditions would be to fix-fixing~~ the average potential temperature profile ~~. On top of that 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 ~~geostrophic damping the proposed geostrophic damping method.~~ This would ~~lead to a better agreement on wind farm power predictions and on the actual~~ allow precursor simulations to reach a truly statistically steady state, leading to better agreement between

different codes on inflow conditions used in successor simulations, allowing to compare precursor simulations that did not run for the same amount of time for successor simulations.




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 hub-height wind angle, hub-height velocity magnitude, and capping inversion center from cases P, G (no potential temperature control), and PT (with potential temperature controller described in Sec. 2.3.2). Time is non-dimensionalized with the start time of the geostrophic damping action T_D .

520 Fig. 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 Sec. 2.3.1 for definition). Moreover, in Fig. 7b we report wind angle and velocity magnitude horizontally averaged at the reference height, together with

525 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 exactly maintains the wind at the desired speed and 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

530 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 5MW wind turbines, aligned in 20 rows and 5 columns, with streamwise and spanwise spacing of 5 and 4.76 rotor diameters respectively. We

535 include thermal stratification to assess the effects of gravity waves blockage for a lapse rate of 1 K/km, 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 $40 \times 21 \times 28$ km in the streamwise, spanwise and vertical direction respectively. All directions are graded to reach a mesh resolution of $30 \times 12.5 \times 10$ m around the wind turbines. The hybrid off-line/concurrent precursor technique described in Sec. 2.4 has been used to spin-up turbulence in the precursor, providing a ~~realistic~~a time-resolved CNBL inflow for the successor
540 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 fringe, Rayleigh, and advection damping region coefficients are given in Appendix C.

Fig. 8 shows hub-height instantaneous velocity and pressure contours around the wind farm. The gravity wave footprint
545 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 8. Contours of instantaneous velocity (left) and pressure (right) at the wind turbine hub-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
550 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 ~~do not~~are not expected to alter the simulation results from the wind farm ~~performance point of view. Nevertheless, as statistics convergence can become extremely slow, this issue can be alleviated by using the so-called shifted periodic boundary conditions in the concurrent precursor simulation, where a spanwise offset is applied in the streamwise periodicity to artificially break the locking in position of such structures~~
555 ~~(Munters et al., 2016).~~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. Fig. 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

560 force, the direction of propagation of interfacial waves in the inversion layer is not aligned with the wind farm streamwise symmetry axis.




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




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

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

Fig. 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.

~~These~~ Another aspect that is noticeable from the average fields of Fig. 9 and Fig. 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 conducting finite-size wind farm simulations, capturing gravity waves 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, 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 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 ~~only differ slightly in the geostrophic wind between different codes, but which~~ 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 ~~of specific to~~ the adopted code, while the final temperature profile depends on both numerical dissipation and precursor simulated time. ~~Using the proposed methodology instead would ultimately enable better agreement on wind farm power estimates using LES.~~ 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 concurrent method is adopted for the turbine simulation. In fact, we found that the concurrent precursor, combined with a fringe region, are crucial elements to simultaneously avoid spurious gravity wave reflections ~~while providing a realistic turbulent inflow at the same time~~ and provide a time-resolved turbulent inflow. The off-line precursor data is 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 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

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

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

$$\mathbf{S}^\xi = \frac{1}{J} \left(\frac{\partial \xi}{\partial x} \hat{\mathbf{x}} + \frac{\partial \xi}{\partial y} \hat{\mathbf{y}} + \frac{\partial \xi}{\partial z} \hat{\mathbf{z}} \right), \quad \mathbf{S}^\eta = \frac{1}{J} \left(\frac{\partial \eta}{\partial x} \hat{\mathbf{x}} + \frac{\partial \eta}{\partial y} \hat{\mathbf{y}} + \frac{\partial \eta}{\partial z} \hat{\mathbf{z}} \right), \quad \mathbf{S}^\zeta = \frac{1}{J} \left(\frac{\partial \zeta}{\partial x} \hat{\mathbf{x}} + \frac{\partial \zeta}{\partial y} \hat{\mathbf{y}} + \frac{\partial \zeta}{\partial z} \hat{\mathbf{z}} \right) \quad (\text{A3})$$

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} , 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 \mathbf{S}^i), and the relation is again evaluated at cell faces, contravariant fluxes are obtained as $V^i = \mathbf{u} \cdot \mathbf{S}^i$. 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 that it can be partially written in terms of contravariant fluxes as

$$\frac{\partial V^q}{\partial l_q} = 0, \quad (\text{A4})$$

$$\begin{aligned} 640 \quad \frac{\partial V^q}{\partial t} + \frac{\partial l_q}{\partial x_i} \frac{\partial}{\partial l_r} (V^r u_i) &= -\frac{1}{\rho_0} \frac{\partial p}{\partial l_r} g^{rq} + \frac{\partial l_q}{\partial x_i} \frac{\partial}{\partial l_r} \left[\frac{\nu_{eff}}{J} \left(\frac{\partial u_i}{\partial l_k} g^{rk} + \frac{\partial u_j}{\partial l_k} \frac{\partial l_r}{\partial x_j} \frac{\partial l_k}{\partial x_i} \right) \right] - \frac{1}{\rho_0} \frac{\partial p_\infty}{\partial l_r} g^{rq} + \\ &+ \frac{1}{J} \frac{\partial l_q}{\partial x_i} \frac{\rho_k}{\rho_0} g_i - \frac{2}{J} \frac{\partial l_q}{\partial x_i} \epsilon_{ijk} \Omega_j u_k + \frac{1}{J} \frac{\partial l_q}{\partial x_i} (f_i + s_i^v + s_i^h), \end{aligned} \quad (\text{A5})$$

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

Eq. 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 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, 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, as 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 poses 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 hybrid staggered/non-staggered formulation facilitates the application of boundary conditions, which are prescribed on the Cartesian velocity using ghost cells.

665 Pressure-velocity coupling is provided using a second-order fractional step method similar to 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

670 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
 675 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 term

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

where $\overline{\cdot}$ is the filtering operation defined as

$$680 \quad \overline{\cdot} = \int \overline{\cdot}(\underline{x} - \underline{r}, t) G(|\underline{r}|) d\underline{r}, \quad (\text{B2})$$

and τ_{ij}^D is the deviatoric part of the sub-grid stresses, as the isotropic part is absorbed in the pressure variable. In curvilinear coordinates, Eq. B1 reads

$$\frac{\partial}{\partial l_k} (\overline{V^k u_j}) = \frac{\partial}{\partial l_k} (\overline{V^k} \overline{u_j}) + \frac{\partial}{\partial l_k} (\sigma_j^k), \quad (\text{B3})$$

where

$$685 \quad \sigma_j^k = \overline{V^k u_j} - \overline{V^k} \overline{u_j} \quad (\text{B4})$$

and assuming a linear eddy viscosity model,

$$\sigma_i^k = -2\nu_t S_j^k S_{ij} \quad (\text{B5})$$

$$\nu_t = C_s \Delta^2 \overline{S_{ij} S_{ij}} \sqrt{2 \overline{S_{ij} S_{ij}}}, \quad (\text{B6})$$

where $S_j^k = 1/J \partial l_k / \partial x_j$ are 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
 690 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

$$T_j^k = \widetilde{\overline{V^k u_j}} - \widetilde{\overline{V^k}} \widetilde{\overline{u_j}} \quad (\text{B7})$$

that 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
 695 tilde-filtered Eq. B3 from Eq. B7

$$G_j^k = T_j^k - \widetilde{\sigma_j^k} = \widetilde{\overline{V^k u_j}} - \widetilde{\overline{V^k}} \widetilde{\overline{u_j}}. \quad (\text{B8})$$

Using Eq. B5 and B6 to express $\widetilde{\sigma}_j^k$ and T_j^k reads

$$\widetilde{\sigma}_j^k = -2C_s \overline{\Delta}^2 |\widetilde{S}| \widetilde{S}_j^k S_{ij} \quad (\text{B9})$$

$$T_j^k = -2C_s \overline{\Delta}^2 |\widetilde{S}| \widetilde{S}_j^k \widetilde{S}_{ij}, \quad (\text{B10})$$

700 where in Eq. B10 the approximation $\widetilde{S}_j^k S_{ij} = \widetilde{S}_j^k \widetilde{S}_{ij}$ has been used (~~good enough for smooth spatial face area vector variation, exact for uniform meshes~~). In fact, as the LES filter has a size of 3 mesh cells in each direction, \widetilde{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 an homogeneous filter and a uniform mesh are considered. Inserting Eq. B9 and Eq. B10 into Eq. B8 leads to

$$705 \quad G_j^k = -2C_s \overline{\Delta}^2 |\widetilde{S}| \widetilde{S}_j^k \widetilde{S}_{ij} + 2C_s \overline{\Delta}^2 |\widetilde{S}| \widetilde{S}_j^k S_{ij} = C_s M_j^k \quad (\text{B11})$$

where

$$M_j^k = 2 \left[\overline{\Delta}^2 |\widetilde{S}| \widetilde{S}_j^k S_{ij} - \overline{\Delta}^2 |\widetilde{S}| \widetilde{S}_j^k \widetilde{S}_{ij} \right]. \quad (\text{B12})$$

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

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

710 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

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

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

715 along streamlines as

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

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

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

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

720 The integrals of Eq. B14 and Eq. 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 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

$$Pr = \frac{1}{1 + 2l/\Delta}, \tag{B19}$$

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

$$s = g_i \frac{\partial \theta}{\partial x_i}.$$

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 of the mixing length scale under stable conditions (Schumann, 1991).

Appendix C: Finite Wind Farm Set-up

In this section, we describe in detail the setup of the finite wind farm case presented in Sec. 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.

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

(a) Vertical 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

(b) Streamwise (left) and spanwise (right) discretization parameters.

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

The size of the successor domain is $40 \text{ km} \times 21 \text{ km} \times 28 \text{ km}$ in the streamwise, spanwise and vertical direction 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 m 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 deg/s when flow misalignment exceeds 1 deg. 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, which 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 km \times 1 km domain, characterized by a resolution of 15 m \times 15 m in the streamwise and spanwise directions. The vertical
760 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 Sec. 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 15000 s.

765 The off-line precursor simulation uses the pressure and temperature controllers described in Sec. 2.3, while in both the
concurrent precursor and successor simulations, velocity is controlled using a constant source term, obtained by averaging
the off-line precursor source from 100000 s to 120000 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 Tab. C2. They are calculated based on the sensitivity
770 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 ABL, respectively. The
temperature profile is initialized according to Rampanelli and Zardi (2004).

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

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

775 In Fig. C1 we show vertical profiles of 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 (Tab. C2), while the resulting friction velocity corresponds to
780 0.432 m/s.




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 $10 \cdot h_{ref}$. All profiles are uniform above, except from the potential temperature profile which exhibits a linear increase equal to the lapse rate γ .

Regarding the successor case, we followed the approach of Lanzilao and Meyers (2022a) to chose 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 waves reflectivity. The dominant vertical wavelength of the gravity waves is estimated as $\lambda_z = 2\pi G/N \approx 11.8$ km (Allaerts and Meyers, 2017), where N is the Brunt-Väisälä frequency ($N = 5.72 \cdot 10^{-3} s^{-1}$ based on parameters listed in Tab. C2), and $G = 10.815$ m/s. 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 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 from 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 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.

x_s [km]	x_e [km]	Δ_s [km]	Δ_e [km]
-20	-15	1	1

(a) Fringe region parameters.

x_s [km]	x_e [km]	Δ_s [km]	Δ_e [km]
-18	-11	1	1

(b) Advection damping region parameters.

Table C3. Fringe and advection damping region information.

We used the same damping functions as Lanzilao and Meyers (2022a), and in Tab. C3 their parameters are reported for our 800 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 805 from $300 \times 300 \times 100$ cells. As a consequence, they consist of 9M, 72M, and 576M elements in total. Tab. D1 reports the number of nodes for each run, which only consisted of two hours of wall-clock time.

	CNBL 9M	CNBL 72M	CNBL 576M
Number of Niagara Nodes	5	10	100
	2	10	200
	4	20	400
	8	40	800
	16	80	-

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

Tests have been performed on Compute Canada's Niagara 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. Fig. D1 shows the time per iteration as the node count 810 increases for each of the CNBL meshes. TOSCA's strong scaling performance remains close to linear until roughly 25k cells per core are reached, which we identify as a reasonable trade-off between efficiency 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 815 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 one (we used 0.9 for these analyses), while explicit methods are usually limited to a value close to 0.5.




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.

Regarding actuator models, their solution and I/O operations are also parallelized in TOSCA. Specifically, we define a sphere
820 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, 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 Sec. 5 (100 wind turbines), the turbine update time was less than 0.1 s.
825 Individual turbine update time depends on cell size and on 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.

Code availability. TOSCA is available at <https://osf.io/q4vaf/>, DOI 10.17605/OSF.IO/Q4VAF

Author contributions. Conceptualization, S.S, D.A., J.B.; methodology, S.S.; software, S.S, A.A.; validation, S.S.; formal analysis, S.S.;
830 investigation, S.S.; computational resources, J.B.; data curation, S.S.; writing–original draft preparation, S.S.; writing–review and editing,
J.B., D.A.; visualization, S.S.; supervision, J.B., D.A.; project administration, J.B.; funding acquisition, J.B.. All authors have read and agreed
to the published version of the manuscript.

Competing interests. No competing interests are present.

Acknowledgements. The present study is supported by UL Renewables and the Natural Science and Engineering Research Council of
835 Canada (NSERC) through Alliance grant no. 556326. Computational resources provided by the Digital Research Alliance of Canada (www.alliancecan.ca) and Advanced Research Computing at the University of British Columbia (www.arc.ubc.ca) are gratefully acknowledged.

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.
- 840 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, *Journal of Wind Engineering and Industrial Aerodynamics*, 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., 2016.
- 845 Allaerts, D. and Meyers, J.: Large eddy simulation of a large wind-turbine array in a conventionally neutral atmospheric boundary layer, *Physics of Fluids*, 27, 065 108, <https://doi.org/10.1063/1.4922339>, 2015.
- Allaerts, D. and Meyers, J.: Boundary-layer development and gravity waves in conventionally neutral wind farms, *Journal of Fluid Mechanics*, 814, 95–130, <https://doi.org/10.1017/jfm.2017.11>, 2017.
- 850 Allaerts, D. and Meyers, J.: Gravity Waves and Wind-Farm Efficiency in Neutral and Stable Conditions, *Boundary-Layer Meteorology*, 166, <https://doi.org/10.1007/s10546-017-0307-5>, 2018.
- Allaerts, D. and Meyers, J.: Sensitivity and feedback of wind-farm-induced gravity waves, *Journal of Fluid Mechanics*, 862, 990–1028, <https://doi.org/10.1017/jfm.2018.969>, 2019.
- Allaerts, D., Quon, E., Draxl, C., and Churchfield, M.: Development of a TimeHeight Profile Assimilation Technique for Large-Eddy Simulation, *Boundary-Layer Meteorology*, 176, <https://doi.org/10.1007/s10546-020-00538-5>, 2020.
- 855 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*, n/a, <https://doi.org/https://doi.org/10.1002/we.2811>, 2023.
- Armenio, V. and Piomelli, U.: A Lagrangian Mixed Subgrid-Scale Model in Generalized Coordinates, *Flow Turbulence and Combustion*, 65, 51–81, <https://doi.org/10.1023/A:1009998919233>, 2000.
- 860 Balay, S., Abhyankar, S., Adams, M. F., Benson, S., Brown, J., Brune, P., Buschelman, K., Constantinescu, E., Dalcin, L., Dener, A., Eijkhout, V., Faibussowitsch, 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, 2022.
- 865 Bastankhah, M. and Porté-Agel, F.: A new analytical model for wind-turbine wakes, *Renewable Energy*, 70, 116–123, <https://doi.org/https://doi.org/10.1016/j.renene.2014.01.002>, special issue on aerodynamics of offshore wind energy systems and wakes, 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, <https://doi.org/10.3390/en11061609>, 2018.
- Branlard, E. and Gaunaa, M.: Cylindrical vortex wake model: right cylinder, *Wind Energy*, 524, 2014.
- 870 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/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, *Journal of Physics: Conference Series*, 1618, 062 036, <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
875 simulations using large eddy simulation, *Philosophical Transactions of The Royal Society A Mathematical Physical and Engineering
Sciences*, 375, 20160 097, <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, *Physics of Fluids*,
22, <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 Block-
880 age engineering parametrizations with LES, *Journal of Physics: Conference Series*, 1934, 012 021, <https://doi.org/10.1088/1742-6596/1934/1/012021>, 2021.
- Chamorro, L. P. 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
885 Turbine Dynamics, *Journal of Turbulence - J TURBUL*, 13, <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>, 2012b.
- Churchfield, M., Schreck, S., Martínez Tossas, L., Meneveau, C., and Spalart, P.: An Advanced Actuator Line Method for Wind Energy
Applications and Beyond, <https://doi.org/10.2514/6.2017-1998>, 2017.
- 890 Etling, D.: Modelling the vertical ABL structure, pp. 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., pp. 632–641, Springer Berlin Heidelberg, Berlin, Heidelberg, 2002.
- Fleming, P. A., Gebraad, P. M., Lee, S., van Wingerden, J.-W., Johnson, K., Churchfield, M., Michalakes, J., Spalart, P., and Mori-
arty, P.: Evaluating techniques for redirecting turbine wakes using SOWFA, *Renewable Energy*, 70, 211–218, <https://doi.org/https://doi.org/10.1016/j.renene.2014.02.015>, special issue on aerodynamics of offshore wind energy systems and wakes, 2014.
- 895 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*, pp. 97–104, Budapest, Hungary, 2004.
- Ge, L. and Sotiropoulos, F.: A numerical method for solving the 3D unsteady incompressible NavierStokes equations in curvi-
900 linear domains with complex immersed boundaries, *Journal of Computational Physics*, 225, 1782–1809, <https://doi.org/https://doi.org/10.1016/j.jcp.2007.02.017>, 2007.
- Germano, M., Piomelli, U., Moin, P., and Cabot, W. H.: A dynamic subgridscale eddy viscosity model, *Physics of Fluids A: Fluid Dynamics*,
3, 1760–1765, <https://doi.org/10.1063/1.857955>, 1991.
- Glauert, H.: *Airplane Propellers*, pp. 169–360, Springer Berlin Heidelberg, Berlin, Heidelberg, https://doi.org/10.1007/978-3-642-91487-4_3, 1935.
- 905 Gribben, B. J. and Hawkes, G. S.: A potential flow model for wind turbine induction and wind farm blockage, *Systems and Engineering
Technology*, 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/https://doi.org/10.1002/fld.4711>, 2019.
- 910 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 Energy Science*, 8, 1251–1275, <https://doi.org/10.5194/wes-8-1251-2023>, 2023.
- 915 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, *Monthly Weather Review*, 142, 3418 – 3424, <https://doi.org/https://doi.org/10.1175/MWR-D-13-00400.1>, 2014.
- Jensen, N.: A note on wind generator interaction, 1983.
- Jimenez, A., Crespo, A., Migoya, E., and Garcia, J.: Advances in large-eddy simulation of a wind turbine wake, *Journal of Physics: Conference Series*, 75, 012 041, <https://doi.org/10.1088/1742-6596/75/1/012041>, 2007.
- 920 Jimenez, A., Crespo, A., Migoya, E., and Garcia, J.: Large-eddy simulation of spectral coherence in a wind turbine wake, *Environmental Research Letters*, 3, 015 004, <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/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, 925 National Renewable Energy Laboratory (NREL), <https://doi.org/10.2172/947422>, 2009.
- Knoll, D. and Keyes, D.: Jacobian-free NewtonKrylov methods: a survey of approaches and applications, *Journal of Computational Physics*, 193, 357–397, <https://doi.org/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, *Boundary-Layer Meteorology*, <https://doi.org/10.1007/s10546-022-00772-z>, 2022a.
- 930 Lanzilao, L. and Meyers, J.: Effects of self-induced gravity waves on finite wind-farm operations using a large-eddy simulation framework, vol. 2265, <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., 1988.
- Leonard, B.: A stable and accurate convective modelling procedure based on quadratic upstream interpolation, *Computer Methods in Applied Mechanics and Engineering*, 19, 59–98, [https://doi.org/https://doi.org/10.1016/0045-7825\(79\)90034-3](https://doi.org/https://doi.org/10.1016/0045-7825(79)90034-3), 1979.
- 935 Lilly, D. K.: A proposed modification of the Germano subgrid-scale closure method, *Physics of Fluids A: Fluid Dynamics*, 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, *Journal of Physics: Conference Series*, 256, 012 026, <https://doi.org/10.1088/1742-6596/256/1/012026>, 2010.
- 940 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, <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 co-ordinates, *International Journal for Numerical Methods in Fluids*, 4, 519–537, <https://doi.org/https://doi.org/10.1002/flid.1650040606>, 1984.
- Maronga, B., Gryschka, M., Heinze, R., Hoffmann, F., Kanani-Sühring, F., Keck, M., Ketelsen, K., Letzel, M., Sühring, M., and Raasch, 945 S.: The Parallelized Large-Eddy Simulation Model (PALM) version 4.0 for atmospheric and oceanic flows: Model formulation, recent developments, and future perspectives, *Geoscientific Model Development*, 8, 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/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, *Journal of Fluid Mechanics*, 319, 353385, <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, <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, <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, *Journal of Atmospheric Sciences*, 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 Geophiz. Inst.*, 151, 163–187, 1954.
- Munters, W., Meneveau, C., and Meyers, J.: Shifted periodic boundary conditions for simulations of wall-bounded turbulent flows, *Physics of Fluids*, 28, 025 112, <https://doi.org/10.1063/1.4941912>, 2016.
- Niagara: Niagara, <https://docs.alliancecan.ca/wiki/Niagara>.
- Niyafar, 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, *Journal of Physics: Conference Series*, 1618, 062 072, <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, vol. 2265, p. 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, 2018.
- Paulson, C. A.: The Mathematical Representation of Wind Speed and Temperature Profiles in the Unstable Atmospheric Surface Layer, *Journal of Applied Meteorology and Climatology*, 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, *Boundary-Layer Meteorology*, 153, <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, vol. 2265, p. 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>, 2019.
- Porté-Agel, F., Lu, H., and Wu, Y.-T.: A large-eddy simulation framework for wind energy applications, 2010.
- Rampanelli, G. and Zardi, D.: A Method to Determine the Capping Inversion of the Convective Boundary Layer, *Journal of Applied Meteorology*, 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, 1992.
- Saad, Y. and Schultz, M. H.: GMRES: A Generalized Minimal Residual Algorithm for Solving Nonsymmetric Linear Systems, *SIAM Journal on Scientific and Statistical Computing*, 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, *Theoretical and Computational Fluid Dynamics*, 2, 279–290, 1991.
- 990 Segalini, A.: An analytical model of wind-farm blockage, *Journal of Renewable and Sustainable Energy*, 13, 033 307, <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, *Quarterly Journal of the Royal Meteorological Society*, 140, <https://doi.org/10.1002/qj.2266>, 2014.
- 995 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 Energy Science*, 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/https://doi.org/10.1002/we.366>, 2010.
- Sockeye: UBC ARC: Sockeye - Detailed Technical Specifications, <https://arc.ubc.ca/sockeye-techspecs>.
- 1000 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, *Renewable Energy*, 116, 470–478, <https://doi.org/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, *Renewable Energy*, 170, 1342–1352, <https://doi.org/https://doi.org/10.1016/j.renene.2021.02.018>, 2021.
- 1005 Stull, R.: *Practical Meteorology: An Algebra-based Survey of Atmospheric Science*, BC Open Textbook Collection, AVP International, University of British Columbia, <https://books.google.nl/books?id=xP2sDAEACAAJ>, 2016.
- Sørensen, J. N. and Shen, W. Z.: Numerical Modeling of Wind Turbine Wakes , *Journal of Fluids Engineering*, 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, *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 373, 20140 071, <https://doi.org/10.1098/rsta.2014.0071>, 2015.
- 1010 The HDF Group: Hierarchical data format version 5, <http://www.hdfgroup.org/HDF5>, 2000-2010.
- Troldborg, N. and Meyer Forsting, A.: A simple model of the wind turbine induction zone derived from numerical simulations, *Wind Energy*, 20, <https://doi.org/10.1002/we.2137>, 2017.
- 1015 van der Laan, M. P., Storey, R., Sørensen, N., Norris, S., and Cater, J.: A CFD code comparison of wind turbine wakes, *Journal of Physics: Conference Series*, 524, 012 140, <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 Journal on Scientific and Statistical Computing*, 7, <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.
- 1020 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, 104 601, <https://doi.org/10.1103/PhysRevFluids.2.104601>, 2017.
- Ørsted: Ørsted presents update on its long-term financial targets, 2019.