In this article, the authors develop a new methodology for simulating wind-farm induced gravity waves in an LES framework at a considerably lower computational cost. This technique is based on the fact that gravity-wave induced pressure perturbations are dependent on the vertical displacement of the capping inversion. Therefore, the multi-scale coupled (MSC) model is adopted to compute this displacement, which is given as input to TOSCA, a LES solver. The latter only simulates the flow in the atmospheric boundary layer, as gravity-wave effects are taken into account by the MSC model. The results obtained with this technique, named AGW-modelled, are compared against LESs that resolve the full domain (i.e. including the free atmosphere), which typically necessitate lots of computational resources. The comparison shows that the AGW-modelled velocity and pressure-perturbation profiles are in good agreement with the ones predicted by the AGW-resolved simulations. Moreover, the extreme case of a rigid lid is also explored by fixing the vertical extent of the domain to the height of the capping inversion.

In my opinion, the manuscript is generally well written. Moreover, the ideas are clearly presented, which makes the reasonings easy to follow in most of the sections. I also believe that this article is of interest to the wind energy community as it shows a new methodology for investigating wind-farm operations in conventionally neutral boundary layers which considerably reduces the computational cost of LES while maintaining a good level of accuracy. In the following, I list comments on how the paper may be improved.

Scientific comments/questions

1. Abstract: the authors mention that this new technique demands less than 15% of the computational resources than traditional methods, which is an impressive achievement. However, this technique also reduces the accuracy of an LES to the one of the MSC model. I think this is quite an important limitation, which should be mentioned.

2. Line 34: the authors mention that AGWs have "extremely large spatial scales". I think the reader would benefit if a number in terms of km would be provided, as the word "extremely" can be subjective. I suggest trying to limit the use of this word throughout the manuscript (e.g. see lines 35 and 51).

3. Line 37: I would also refer to some earlier works about AGW excited by hills/mountains (e.g. Klemp and Lilly 1977, Teixeira 2014).

4. Line 54-67: the authors describe the Rayleigh damping layer and fringe region technique in this paragraph. However, the descriptions are mixed so that it is easy to confuse the role of the two techniques. I would propose to re-write this paragraph explaining first the Rayleigh damping technique, for instance, and afterwards the working principle and use of the fringe region technique.

5. Line 78: the authors define the rigid-lid approximation as a case with strong free atmosphere stability. However, this could also be seen as a case with a very strong (practically infinite)

capping-inversion strength. To avoid confusion, I would define the rigid-lid case as an approximation for a case with very strong thermal stratification above the atmospheric boundary layer.

6. Introduction: I suggest to include the general aim of the article in the first paragraph of the introduction. This will give the reader a hint about why AGW numerical models and boundary conditions are described in this section.

7. Line 107: I would mention this in the introduction, i.e. that the article only deals with CNBLs.

8. Section 2.2: the proposed technique uses the MSC model to predict the capping-inversion vertical displacement (eta) and only simulate the flow within the ABL with LES. Hence, it is implicitly assumed that H>H1, with H1=2*z_hub. However, H1 values are getting closer to heights at which the capping inversion is typically located (for instance, the Vestas V172-7.2MW turbine has a hub height of 199 m, a figure that may increase in the coming years) so that cases where the ratio H/H1 $\approx$ 1 are realistic. In the current work, the authors fix H/H1 to 2.77 and do not discuss this further. However, I believe that the performance of the new technique is sensitive to the H/H1 ratio. Would it be possible to add a few comments and/or simulations on what would happen when H/H1 $\approx$ 1 ? Also, how would this technique deal with cases where H/H1 $<$ 1? If this is a limitation, it should be reported in the text.

9. Line 184: the authors mention that “Then, the vertical displacement is linearly distributed to the underlying cells, deforming the mesh before starting the simulation”. This passage is not clear to me. Would it be possible to explain in more detail?

10. Line 193: in the LES framework, the vertical displacement of the capping inversion generates a cold anomaly, which in turn results in pressure perturbation. However, the authors mention that in CNBLs, the potential-temperature equation can be left out in the AGW-modelled simulations. Hence, are the pressure perturbations solely caused by the flow convergence/divergence caused by the irregular top edge of the main domain? How is the buoyancy term computed in the vertical momentum equation? Or is it neglected? I would appreciate a more detailed explanation.

11. Line 194: the authors mention that “This condition is only violated very close to the top boundary, where discrepancies in turbulent fluctuations produced by the absence of stability and by the physical boundary are deemed acceptable as they happen away from the wind farm”. I would note that this affirmation does not hold for low H/H1 ratios (as mentioned in comment 8).

12. Section 3.2: if I understand correctly, the inflow data used for the AGW resolved and modelled simulations are computed with two different precursor techniques. Hence, differences in the wind-farm simulation results between the two techniques could be also attributed to this difference in inflow conditions. Why the authors do not drive the AGW-modelled simulations with inflow sections taken from the precursor simulations used for the AGW-resolved simulations?

13. Line 264: the authors mention that “we used the velocity inflow data of the subcritical case to prescribe an inlet for the rigid lid”. Is the velocity profile in the precursor domain for sub- and supercritical cases equal? Showing some vertical profiles of the precursor simulation would be beneficial.

14. Figure 2: I suggest to center the colorbar around the zero value (so that the background color is white).

15. Line 274: the authors mention that “the AGW-modeling technique requires a domain that is more than 85% smaller compared to the AGW-resolving approach”. In which terms? Number of cells? Does this number also consider the precursor domain (for instance, the AGW-modelled simulations use a precursor domain 7 times bigger in the y-direction)?

16. Figure 3: How would you explain the differences between AGW-modelled simulations and the MSC model? Are those due to the simplifications made in the MSC model (for instance, linearity, absence of resolved turbulence, etc..)?

17. Figure 3: I suggest extending the caption of this figure, mentioning for instance that the profiles shown are averaged in height between H1 and H and along the wind-farm width. This comment extends to the whole manuscript since I feel that the combination of figure and caption is often not self-explanatory.

18. Line 304: the domain in Figure 3 is too small to appreciate this behaviour.

19. Line 313: at which height are the profiles taken? I suggest mentioning it in the manuscript.

20. Figure 4: the match in terms of velocity is really good. This makes it hard to spot differences between AGW-modelled and resolved simulations in Figure 4, which obviously is a good sign. Could be an idea to also plot the relative error for both cases? This will allow the reader to easily understand where the two methods differ the most.

21. Table 3: Would it be possible to include the non-local, wake and farm efficiency values in this table (as defined in Lanzilao and Meyers (2024), for instance)? The total farm power (and farm efficiency) in the supercritical case is almost identical in the AGW-M and AGW-R cases. However, I expect some differences in non-local and wake efficiencies.

22. Section 4.2: I’m assuming that the rigid lid is located at H=500 m. Is this correct? I suggest to explicitly mention the vertical extent of the domain in the text.

23. Section 4.2: the pressure build-up in the rigid-lid case is solely attributed to flow confinement. However, even in neutral conditions, the cumulative turbine induction generates a pressure build-up and consequently, a flow slow down (typically much lower in values than the one observed in the presence of thermal stratification). Therefore, I would rephrase this sentence

and/or find an alternative method to evaluate the flow blockage solely induced by the flow confinement.

24. Line 350: Which is the MSC setup used to generate Figure 8? In the text, only the capping-inversion strength and free lapse rate are mentioned. In general, I would appreciate more details on the simulation setup, so that it would get easier and more intuitive to reproduce the results.

25. Figure 8: it would be interesting to split this figure into three panels, reporting the sensitivity of the relative error based on non-local, wake and farm efficiency to the capping-inversion strength and free lapse rate. I suggest this because at times the total farm power of two simulations can be almost identical, even though the power trends are very different (two behaviours that cancel out).

26. Table 4: the relative error remains positive for all cases when the MSC model is used. However, in the LES case, the error becomes negative for the subcritical case. Any idea about why this occurs?

27. Line 372-374: the fringe region is adopted in pseudo-spectral (or fully spectral) flow solvers to impose the inflow conditions. The presence of gravity waves does not imply the use of a fringe region, as inflow-outflow boundary conditions can be adopted (although the implementation is not trivial).

28. Line 385: As a future work, I would also suggest further validation of this technique, as it has been only applied to two idealized cases. For instance, Lanzilao and Meyers (2024) performed 40 LESs in different atmospheric conditions, for which the displacement of the capping inversion is computed. This comparison could offer further insights into the performance of the proposed method.

29. Appendix A: this construction looks quite artificial to me. From my point of view, it would be easier to drive the main domain using the same precursor simulation for both the AGW-modelled and AGW-resolved simulations. Is there a particular reason why the authors decided to not pursue this option? It would eliminate the need for the process described in this appendix together with ensuring equal inflow conditions in both cases.

Technical comments

1. Line 113: replace “If one wishes to resolve AGW within LES” with “When simulating AGWs in an LES framework” or similar

2. line 205: used -> use

3. Line 368: conventionally neutral boundary layers -> CNBLs