Review of Geoscientific Model Development manuscript egusphere-2025-1135: "PortUrb: A Performance Portable, High-Order, Moist Atmospheric Large Eddy Simulation Model with Variable-Friction Immersed Boundaries" by Matthew Norman, Muralikrishnan Gopalakrishnan Meena, Kalyan Gottiparthi, Nicholson Koukpaizan, and Stephen Nichols.

This article introduces and demonstrates the portUrb model, an LES code designed with the intention of providing portability, performance, accuracy, simplicity, readability, robustness, extensibility, and ensemble capabilities. The code is written in C++ and advertised to work both on CPUs and different GPU architectures (Nvidia, AMD, and Intel). A number of canonical ABL cases are presented to demonstrate and verify the adequacy of the model implementation, some of which include comparisons to other existing LES results from the literature. I want to particularly commend the authors for the great work in this new LES model development, and for the clarity with which the article is mostly written and organized. That is not an easy task when needing to convey such a large amount of information and details. I am supportive of seeing this manuscript published and give portUrb visibility for the community. However, there are a number of aspects the authors need to address before the paper can be published in GMD.

Major comments:

1) Figure 4 and associated discussion. Spectral energy distribution from Fig. 4 does not support the claim of 4-8dx effective resolution. It appears to be closer to ~8dx, even starting to deviate from the theoretical -5/3 for slightly larger scales than 8dx.
2) Section 3.3.2, Figs. 10 and 11. It would be desirable to include digitized results from Morrison and Milbrandt (2011) here for reference. It is otherwise very difficult to assess the comparison, and it makes really difficult to follow the discussion.
3) Section 3.5. I am not sure what this test case is adding from a code verification standpoint. I would suggest removing it unless results are compared to previous studies. In addition, the very limited discussion does not provide any insight that contributes to the purpose of the paper.
4) Section 3.6.
    a. Lines 495-496. This vertical domain extent does not allow for a proper interaction between the urban features and the ABL. This needs to be mentioned, even if for the purpose of simply running an urban case this may be okay.
    b. Lines 501-502. This is another questionable setup choice. I am not sure this test case serves as much of a purpose as it could be. At least using a deeper domain, including Coriolis effects, and considering some sort of at least idealized vertical profiles for BCs would be a better way to assess the objective (1) in the first paragraph of Sect. 3.6. In addition, I would strongly recommend the authors select a field campaign experiment itself (e.g., OKC Joint Urban 2003 is the most common), to in addition have the opportunity to validate their results against some observations.
5) Grid spacing. It seems like portUrban is restricted to using uniform grid spacing in all directions (this is not specifically mentioned but appears to be the case based on all the examples presented in the paper). While this is okay in the horizontal directions, not being able to have a vertically stretched grid is a serious drawback for ABL simulations. Otherwise, it will be very computationally expensive to cover the required vertical extent

while maintaining sufficiently fine grid spacings near the surface. I strongly encourage the authors to work on this enhancement as their priority if they want to see portUrb being of scientific value for atmospheric simulations and used by the community. In the current state, the code is significantly hindered in its applicability to relevant problems.

6) Another aspect that needs to be clearly acknowledged is the apparent limitation of portUrb to not include terrain representation capabilities. That would make it really challenging to apply it to realistic atmospheric problems, since there are always terrain variations present, and that in some cases are dominant ABL drivers.
7) In addition to points #4 and #5, portUrb seems to only allow for the use of laterally periodic boundary conditions. Nothing but highly idealized ABLs can be reproduced with such setup, significantly limiting the applicability to relevant real-world scenarios.
8) While this requires some non-negligible development, I strongly suggest the authors consider at least the incorporation of the capabilities mentioned in #4 and #5 before releasing the code (and ideally #6, although that could come with a future release). These two are basic capabilities without which I am not sure there is enough value to bring the model to the open community, since the use would be highly restricted.
9) Code 'performance' is mentioned several times as one of the key targets aimed with portUrb. However, this aspect is not assessed to any extent on the paper. I recommend the authors find some examples of code performance experiments in some of the atmospheric LES codes they mention in the introduction and report in more detail about timings and scaling across processors.

Other comments:

- Line 67. Later in the section (Eqs. 5 and 6) there is viscous terms, so this should be Navier-Stokes equations then.
- Eq. 1. Do the horizontal directions use total pressure and the vertical is perturbation pressure? If not the d/dx and d/dy terms need to have p' instead of p.
- Lines 157-158. Looks like this sentence requires some grammatical fixes.
- Lines 215-216. I wonder how sensitive are the results to the choice of the roughness at the immersed boundary surfaces, especially given the general lack to prescribe their values? It could be beneficial to investigate this sensitivity and include some brief summary in an appendix.
- Lines 271-272. I am not sure this adds anything since a uniform value of SGS TKE will not trigger any SGS TKE production by itself. Initial potential temperature perturbations will be more effective for that matter.