Title: *Comparison of dealiasing schemes in large-eddy simulation of neutrally-stratified atmospheric boundary-layer type flows*

Authors: *Fabien Margairaz, Marco G. Giometto, Marc B. Parlange, and Marc Calaf*

Submitted to: Geosci. Model Dev. (gmd-2017-272)

Reviewer: Anonymous Referee #1

Date: May 1, 2018

General response

Reviewer general comment: *This manuscript addresses the computational cost and the consequences of using different dealiasing schemes in the advective term of the momentum equations for pseudo-spectral discretizations, such as those frequently used in simulations of turbulence in the planetary boundary layer. In particular, two more economical approaches are contrasted with the exact* $3/2$ *rule. Overall the results presented may be relevant for the community interested in turbulence simulation in geophysical settings, such as the planetary boundary layer and the ocean mixed layer. The paper is well written, even though a more concise version would probably be more appealing to the reader. I am a bit surprised by the low cost of the Poisson solver, and the choice of a very low-cost subgrid-scale model certainly helps increasing the share of the convective term on the total cost (as briefly mentioned in the text). After reading the manuscript, I would probably still stick to the 3/2 rule, but that is just my opinion. Having a lower-cost alternative may be useful to other scientists working in the field. The topic may be a bit on the margins of interest of GMD, but I do recommend it for publication. Some important remarks follow below.*

Authors response: We thank the reviewer for his/her positive comments and the recommendation for publication. In this regard the reviewers' comments have made us realize that there were some aspects in the manuscript that required additional clarification and/or improvement. For this reason, down below we provide a clear response to each one of the issues and comments indicated by the reviewer, and an additional explanation on how these have now been addressed within the manuscript.

Based on the comments from both reviewers, we have added a new discussion section to improve the readability of the manuscript. We have also improved the results section by running additional simulations at the resolutions of $192, 192, 128$ and 192^3, as well as running the 256^3-case long enough to be able to report converged statistics. As a result Figures 7, 8 and 9 have been updated, and a new figure (Figure 10) has been included, which is discussed at the end of the results section.

Changes in manuscript: The new discussion section is reported here:

"In the development of this manuscript, focus has been directed to the study of the advantages and disadvantages of different dealiasing methods. In this regard, throughout the analysis we

have tried to keep the structure of the LES configuration as simple and canonical as possible, to remove the effect of other add-on complexities. Additional complications might arise when considering additional physics; here we discuss the potential effect that these different dealiasing methods could have on them. One of such elements of added complexity is for example the use of more sophisticated subgrid scale models based on dynamic approaches to determine the values of the Smagorinsky constant (Germano et al., 1991; Bou-Zeid et al., 2005). In most of these advanced subgrid models, information from the small-scale turbulent eddies is used to determine the evolution of the subgrid constant. However, in both the FT and FS method, the small turbulent scales are severely affected and hence use of dynamic subgrid models could be severely hampered unless these are accordingly modified and adjusted, maybe via filtering at larger scales than the usual grid scale. Another element of added complexity consists in using more realistic atmospheric forcing, considering for example the effect of the Coriolis force with flow rotation as a function of height and velocity magnitude. In this case, we hypothesize that the FT method could lead to stronger influences on the resultant flow field as this dealiasing technique not-only affects the distribution of energy in the small turbulent scales, but also in the large scales (as apparent from Fig. 3), being these later ones potentially more affected by the Coriolis force. This represents a strong non-linear effect, that is hard to quantify and hence further testing, including realistic forcing with a geostrophic wind and Coriolis force would be required to better quantify these effects. Also often in LES studies of atmospheric flows one is interested in including an accurate representation of scalar transport (passive/active). In this case the differential equations don't include a pressure term and hence most of the computational cost is linked to the evaluation of the convective term. As a result, the benefit of using alternative, cheaper dealiasing techniques (FT or FS) will be even more advantageous, yet the total gain is not trivial to evaluate *a priori*, and the effect on the scalar fields should also be further evaluated.

In general, we believe that it is not fair to advocate for one or other dealiasing method based on the results of this analysis. Note that the goal of this work is to provide an objective analysis of the advantages and limitations that the different methods provide, letting the readers the ultimate responsibility to choose the option that will adjust better to their application. For example, while having exact dealiasing (3/2-rule) might be better in studies focusing on turbulence and dispersion, one might be well-off using a simpler and faster dealiasing scheme to run the traditionally expensive warm-up runs, or to evaluate surface drag in flow over urban and vegetation canopies, where most of the surface force is due to pressure differences (Patton et al.,2016)."

Specific responses

Major comments:

1. **Reviewer comment:** *The authors perspective on the FT method is a bit different from mine, and to be honest, I am not sure which one is the most prevalent in the community. I have always considered the FT method (which is also known as the $2/3$ rule) as a slightly different implementation of the exact $3/2$ rule, but with a reduced number of effective grid points (actually $N_{eff} = 2N/3$ instead of N). In my view, performing a calculation with $3N/2$ points using the FT method should be identical (in simulation results but maybe not in computational cost) to a simulation with N points using the 3/2-rule, no? I also think the that simulations using the FT method should actually report grid resolution based on N_{eff} and not on N (i.e. the grid spacing should be $\Delta x_{eff} = 3\Delta x/2$ instead of Δx), but that is usually not done. Maybe the authors can comment on this?*

Authors response: This is a very important point in which we fully agree with the reviewer. Running a simulation with $3N/2$ points using the FT method will not provide the same fluid-flow results (physics) than a simulation with N points using the 3/2-rule. The FT method is equivalent of using a coarser grid for the convective term. Therefore, some spurious oscillations appear in the flow field as can be seen in figure 6 of the paper, and indicated therein. In addition, as mentioned also by the reviewer, there will be an important difference in computational cost. For the sake of discussion, we have run a comparison test case with a grid of 128^3 using the $3/2$-rule, and another case with a grid of $192, 192, 128$ (so $3N/2$ points in the horizontal direction). Results show that the latter is about 1.65 times more expensive than the former and the physical results are not the same or equivalent either. Specifically, figure 1 shows the instantaneous stream-wise velocity perturbation where some spurious oscillations appear in the flow field. In addition, figure 2 presents the velocity profile, velocity gradient, and variance profiles. The velocity profile obtained with the FT method shows an acceleration of the flow compare to the 3/2-rule, yielding a increase of MKE of 5%. The velocity gradient exhibits a large departure for the FT method at height between $0.2 < z/zi < 0.8$. Similarly, the FT methods yields larger variances throughout the domain.

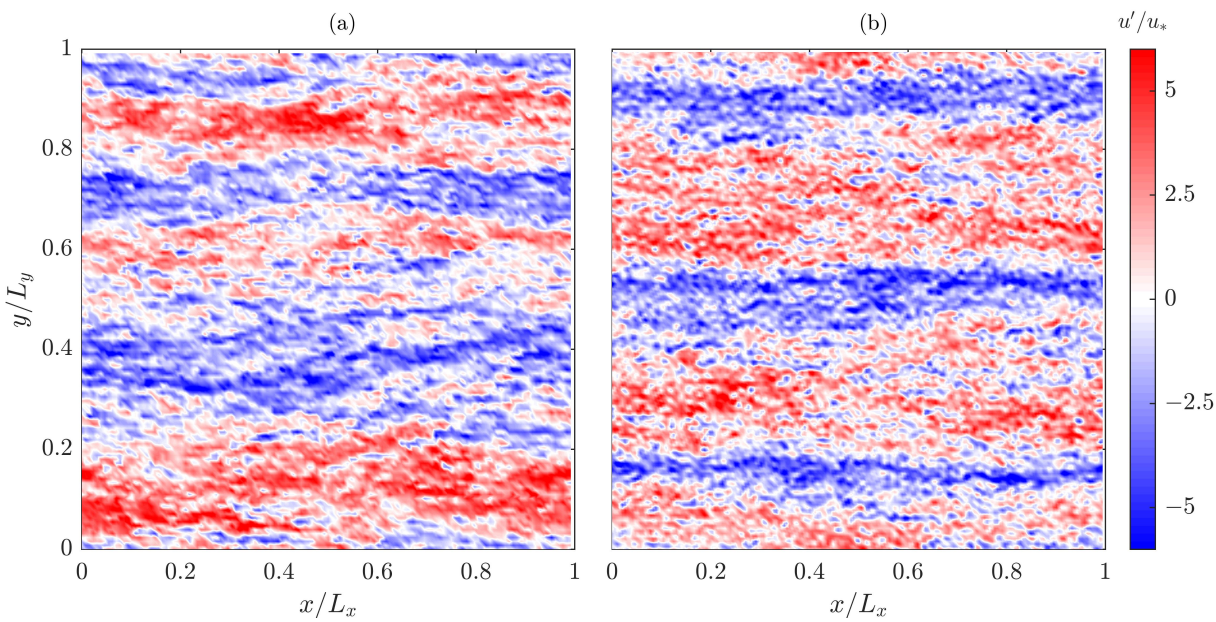

Figure 1: Instantaneous stream-wise velocity perturbation $u'(x, y, z, t) = u(x, y, z, t) - \bar{u}(x, y, z)$ at $z/z_i = 0.054$ for the 128^3-simulation with 3/2-rule (a), and the $192, 192, 128$-simulation with FT (b).

2. **Reviewer comment:** *I am surprised by the high share of the cost carried by the convective term compared to a very low share carried by the Poisson solver. Is this something that is specific to the pencil decomposition parallelization? It may be the case that padding with the 3/2 rule is not as simple in a pencil decomposition approach given that one usually needs to pad the entire 2D "horizontal" wave-number space? Maybe some details of the padding in the context of pencil decomposition would be useful. In addition, is it possible that the Poisson solver is faster when pencil decomposition is used?*

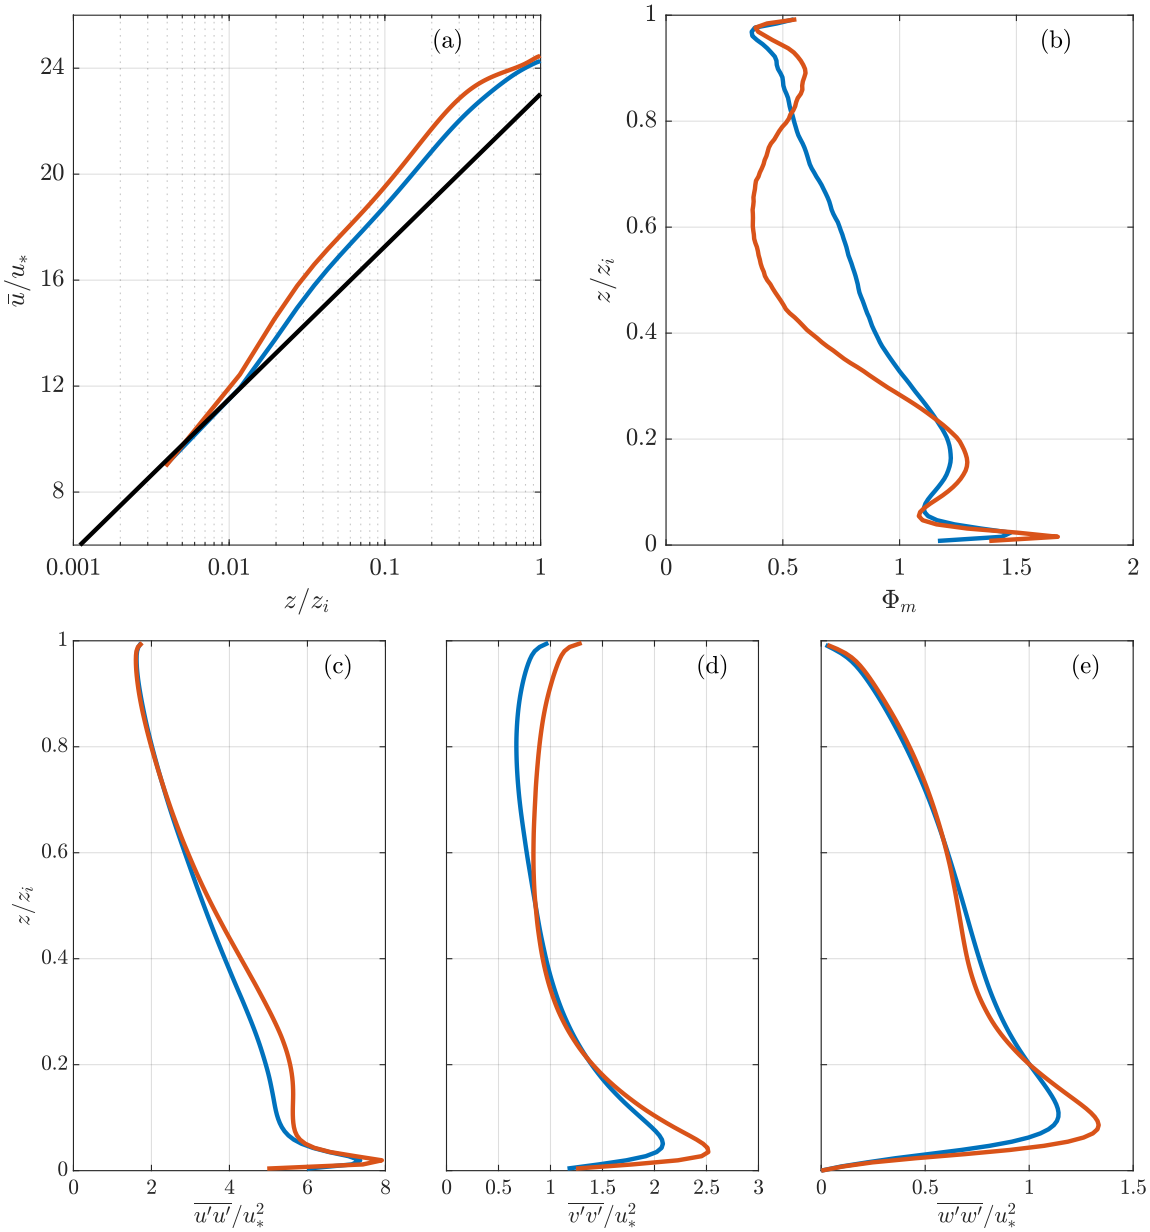

Figure 2: Plots of the non-dimensional mean stream-wise velocity profile (a), the mean stream-wise velocity gradient (b), and the non-dimensional variances (c/d/e). The lines represent the 128^3-simulation with 3/2-rule in blue line and, the $192, 192, 128$-simulation with FT in orange line.

Authors response: Indeed. The cost breakdown for the resolution of the convective term and the Poisson solver is also influenced by the pencil decomposition. In treating the convective term with the pencil decomposition the communication cost increases with respect to the traditional slice parallelization. In this case a total of nine transpositions are needed to compute the convective term, significantly increasing the computational cost.

On the other hand, the Poisson solver becomes faster when using the pencil decomposition in comparison to the slice parallelization. Note that in the pseudo-spectral method the hor-

izontal directions (x and y) are treated in Fourier space and only the vertical direction (z) remains in physical space, therefore each mode in k_x and k_y become independent of each other. In this case the system of equations, originally of size $n_x \times n_y \times n_z$ becomes $n_x \times n_y$ systems of n_z equations, making each vertical line in the domain independent. The pencil decomposition can take full advantage of this fact, making the resolution of the Poisson equation faster. Specifically, once the domain is transposed in the Z-pencil (square pipe aligned with the z-coordinate), the process of solving each of the $n_x \times n_y$ systems does not require any communication, making it very efficient, and limiting its cost to the transposition between the different pencils.

Changes in manuscript: We realize that this is an important detail that should be also mentioned in the manuscript. For this reason we have included a couple of lines in section 3.3.

The text now reads as:"In addition, it is important to note that the low computational cost of the Poisson solver is related to the the use of the pencil decomposition, which takes full advantage the pseudo-spectral approach. Specifically, the Z-pencil combines with the horizontal treatment of the derivatives to make the implementation of the solver very efficient."

3. **Reviewer comment:** *The quantification of computation cost in section 4.1 is fine. Regarding the results in section 4.2, I wonder if the question could be posed the other way around. In my view, maybe the most relevant question would be: "given a set computational cost, do I get better results (a) using the $3/2$ rule or (b) running a finer grid using a less expensive dealiasing?". I know the experiment is not easy to design (because one needs to estimate the computational cost ahead of time), but that seems to be the more important question to be answered. If I save 30% on dealiasing and spend it on more resolution, do I get better results? From my comment above, I would expect that the FT method is almost equivalent to the $3/2$ rule (if smartly coded). How about the FS?*

   Authors response: This is a very interesting point that is being raised. It is true that *a-priory*, as the reviewer mentions, one could compute the associated computational gain linked to using the FT method and decide using a finer grid that would 'use' the time resources saved in the benefit of resolution. However this is a challenging endeavor for many different reasons. For example, if one was to consider an a-priory 20% gain when using the FT method on a given numerical grid of 128^3 points, and decided to use a more resolved grid to 're-invest' the saved computational resources, then one would have to use a grid of ~ 150 $(128 + 0.2 \times 128)$. Unfortunately, this number of grid points will not work well with the FFT given that it's not a power of 2 and hence will induce a slow down of the computations. Accounting for this factor is quite challenging, if not unfeasible, turning the reviewer's suggestion in a very hard challenge. In this regard, also note the response to the first point raised where it is clarified the fact that the $3/2$ rule does not provide the exact same results to the case using the $2/3$ rule with an equivalently larger grid.

   Alternatively, the gain in computational time could be for example invested in running longer simulations, or reduce the time in 'warm-up' configurations.

   Changes in manuscript: We have included an additional comment related to the possible use of the FT or FS method as precursor simulations in the new discussion section (reported above).

4. **Reviewer comment:** *Regarding the interpretation of Figure 7. I am not convinced the log-law prediction is as good as described in the text (but perhaps I am missing something here?). First, it seems that there is one log-law on top of the solid line that may extend only for 2 or 3 vertical levels (ending around $z/z_i \leq 0.02$). Then there is a second log-law (only in $3/2$ and FT cases) that starts*

around $z/z_i = 0.04$ and goes beyond $z/z_i = 0.1$, This second log-law has the incorrect roughness (if one were to extrapolate it to $u = 0$, it would yield a lower value of z_0 than the one imposed on the simulation, I think). There is no clear second log-law in the FS method. I would see this as a concern for the FS method (which is being advocated here), except that not even the exact $3/2$ rule has a good log-law (as seen in Bou-Zeid et al, 2005). I am pretty sure this is due to the SGS model adopted here. In any case, if I had to choose between the FS and FT methods based on Figure 7, I would probably go with FT, since it does a reasonable job in the lower 20% of the domain (which is the region of interest in a simulation like this, I guess). Also, the FS method seems more over-dissipative near the wall (on panel b), which is opposite to what is described in the text?

Authors response: We agree with the reviewer that the log-law predictions are far from perfect, yet they match well with those presented in Bou-Zeid et al, (2005) when using the constant Smagorinsky coefficient (contrasted on a side work). This means that the deviations from the theoretical log-law are mostly due to the SGS model, as the reviewer suggests. These deviations are more prominent in the FS method whereas the FT shows excellent agreement with predictions from the 3/2-rule in the surface layer.

In regards to the second comment, we believe that there was a misunderstanding, given that the authors comments in the manuscript referred to the upper region of the BL, while the reviewer is referring to the surface layer region.

Changes in manuscript: We have now clarified both issues in the revised version of the manuscript. The interpretation of the figure 7 now reads:
"The horizontally- and temporally-averaged velocity profiles are characterized by an approximately logarithmic behavior within the surface-layer ($z \approx 0.15 z_i$, as apparent from Fig. 7, where results are illustrated for the three resolutions: 128^3,192^3, and 256^3). For the 128^3 case, the observed departure from the logarithmic profile for the 3/2-rule case is in excellent agreement with results from previous literature for this particular SGS model (Port-Agel et al., 2000; Bou-Zeid et al., 2005). When using the FT method the agreement of the averaged velocity profile with the corresponding 3/2-rule profiles improves with increasing resolution. While in the 128^3 case a good estimation of the logarithmic flow is obtained at the surface layer, there is a large acceleration of the flow further above. This overshoot does not occur for the higher resolution runs. When using the FS method, the mean velocity magnitude is consistently over-predicted throughout the domain, and the situation does not improve with increasing resolution (the overshoot is up to 7.5% for the 128^3, 8.5% for the 192^3 and 7% for the 256^3 run). Further comparing the results obtained by the FS and FT method with those obtained with the 3/2-rule, it is clear that while the FS method presents a generalized overestimation of the velocity with a an overall good logarithmic trend, the FT method presents a better adjustment in the surface layer with larger departures from the logarithmic regime on the upper domain region that get reduced with increasing numerical resolution. The mean kinetic energy of the system is overestimated by $\approx +2\%$ and $\approx +12\%$ by the FT and FS methods, when compared to that of runs using the 3/2-rule in the 256^3 case. Overall, the mean kinetic energy is larger for the FT and FS cases, when compared to the 3/2-rule case, even at the highest of the considered resolutions ($\approx +2\%$ and $\approx +12\%$ by the FT and FS methods for the 256^3 case). Such behavior can be related to the low-pass filtering operation that is performed in the near-wall regions, which tends to reduce resolved turbulent stresses in the near-wall region, resulting in a higher mass flux for the considered flow system. This is more apparent for the low resolution cases.

Mean velocity gradient profiles ($\Phi_m = \kappa \frac{z}{u_*} \partial_z \langle U \rangle_{xy}(z)$) are also featured in Figure 7 (d, e,

f). Profiles at each of the considered resolutions present a large overshoot near the surface, which is a well known problem in LES of wall-bounded flows and has been extensively discussed in the literature (Bou-Zeid et al., 2005; Brasseur and Wei, 2010; Lu and Port-Agel, 2013). In comparing the results between the FS and FT method with the 3/2-rule, it can be observed that there are stronger gradients in the mean velocity profile within the surface layer when using the FS method. This leads to the observed shift in the mean velocity profile. Conversely, when using the FT method, departures are of oscillatory nature, leading to less pronounced variations in the mean velocity profile when compared to the reference ones (the 3/2-rule cases). This behavior is consistently found across the considered resolutions, but the situation ameliorates as resolution is increased (*i.e.* weaker departures)."

5. **Reviewer comment:** *Regarding Figure 9. I do not agree that the filtering only affects the small-scale end of the spectra. For the FT case, it seems clear to me that there is significant damping of the large scales as well. This is related to the underestimation in the variance of the streamwise velocity seen in Figure 8. This is worrisome, and probably related to the fact that the true resolution here ($N_{eff} = 85$ points) is too coarse to model the ABL. I wonder if this situation would persist in the 256^3 simulations?*

   Authors response: Thank you very much, this is a very interesting and important point. The underestimation of the variance of the streamwise velocity did not persist in the 256^3 simulations. In addition, to clarify the effect of the FT and FS methods on the spectra, we have developed an additional analysis using the spectra presented in the paper. Although the effect of the FT and FS methods on the small scales can be clearly observed on the spectra, their effect on the large scales cannot be directly assessed from the figure, as pointed out by the reviewer. To compute a direct comparison scale by scale, the following ratio was used for the 128^3, 192^3, and 256^3-simulations,

$$\rho^{XX}(k) = \frac{E_{u,k}^{XX} - E_{u,k}^{3/2}}{E_{u,k}^{3/2}} \tag{1}$$

   where $E_{u,k}$ denotes the power spectral density of the u velocity component at wavenumber k, XX stands for the dealiasing method FT or FS. Hence, if $\rho(k) < 0$ energy is removed at that scale, and if $\rho(k) > 0$ energy is added at that scale. Figure 3 presents the ratio $\rho(k)$ for both methods where it can be observed that the effect of FT methods is very large at all scales. The large scales ($0 \leq k/k_{max} \leq 0.2$) are affected with a reduction of energy of ∼25%. The mid-range scales ($0.2 \leq k/k_{max} \leq 0.6$), corresponding to the inertial sub-range, exhibit an overestimation of their energy of about ∼50% on average. Therefore, this method redistributes the energy of the small scales into the inertial sub-range scales. On the contrary, in the FS method, the energy from the filtered small-scales is redistributed more or less uniformly throughout with an averaged overall variation of less the 13%.

   Changes in manuscript: We have added figure 3 and its interpretation in the manuscript. This now reads as: "Although the effect of the FT and FS methods on the small scale can be clearly observed in figure 9, their effect on the large scales also needs to be quantified. To compute a direct comparison scale by scale, the following ratio was used (equation 1) for the 128^3, 192^3, and 256^3-simulations,

$$\rho^{XX}(k) = \frac{E_{u,k}^{XX} - E_{u,k}^{3/2}}{E_{u,k}^{3/2}} \tag{2}$$

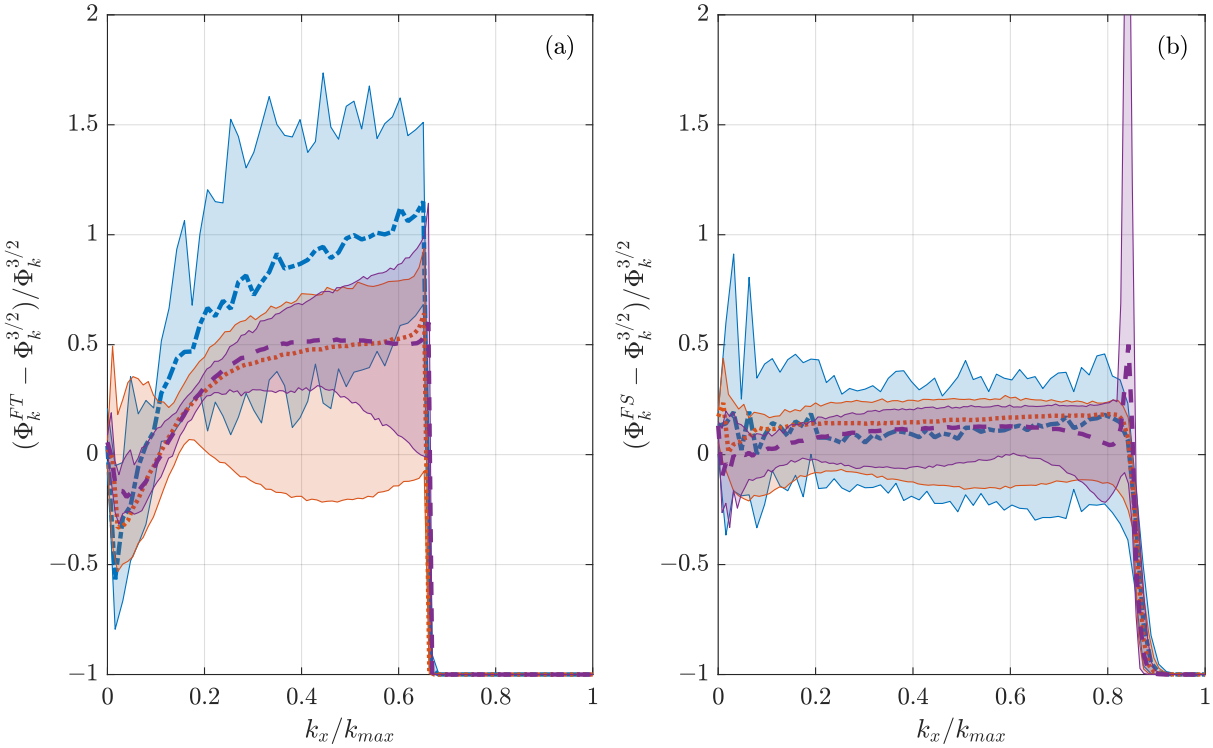

Figure 3: Effect of the FT (a), and the FS (b) methods of the stream-wise spectra of the stream-wise velocity compare to the 3/2-rule. The solid line represent the average value and the shaded area represent the extreme values. The resolutions are: 128^3 in blue dot-dashed line, 192^3 in red dotted line, and 256^3 in purple dashed line.

where $E_{u,k}$ denotes the power spectral density of the u velocity component at wavenumber k, XX stands for the dealiasing method FT or FS. If $\rho(k) < 0$ the energy density at that given wavenumber (k) is less than the corresponding one for the run using the $3/2$ rule, viceversa if $\rho(k) > 0$. Figure 3 presents the ratio $\rho(k)$ for both methods.

When using the FT method, energy at the low wavenumbers is underpredicted, whereas energy at the large wavenumbers is overpredicted. Departures are in general larger with decreasing resolution, with an excess of up to 100% for the 128^3-simulations in the wavenumber range close to the cutoff wavenumber. On the contrary, when using the FS method, the energy from the filtered (dealiased) small-scales is redistributed quasi-uniformly throughout the spectra with an averaged overall variation of less than 13%."

6. **Reviewer comment:** *I see two facts that could benefit from more discussion in the manuscript (maybe in the conclusions):*
(a) The importance of reducing computational cost in the advection term for simulations that use more sophisticated SGS models (maybe this can be brought up again in the conclusions?)
(b) The consequence for including one or many additional scalar fields (temperature, water vapor, etc.). These also require dealising and do not require additional pressure solvers and/or very expensive SGS models. I would probably anticipate that in simulations with several scalars, the savings would be significantly larger. Is that correct?

Authors response: Thank you, (a) is a great point. In this regard, we have now run some ad-

ditional simulations using the dynamic Smagorinsky model. Results are illustrated in Figure 4. In this figure, it can be observed that the dynamic model fails to compute the Smagorinsky constant when using either the FS or the FT methods. In addition, the FT methods strongly suppresses the turbulence, resulting in the lamieraizition of the flow. Alternatively, the consequences of using the FS method are less dramatic, although the flow also exhibits a large acceleration at the top of the domain. As mentioned by the reviewer, these results are not surprising given that the dynamic models are using a relation between the small scales to compute the Smagorinsky constant. Therefore, the FT and FS methods cannot be used

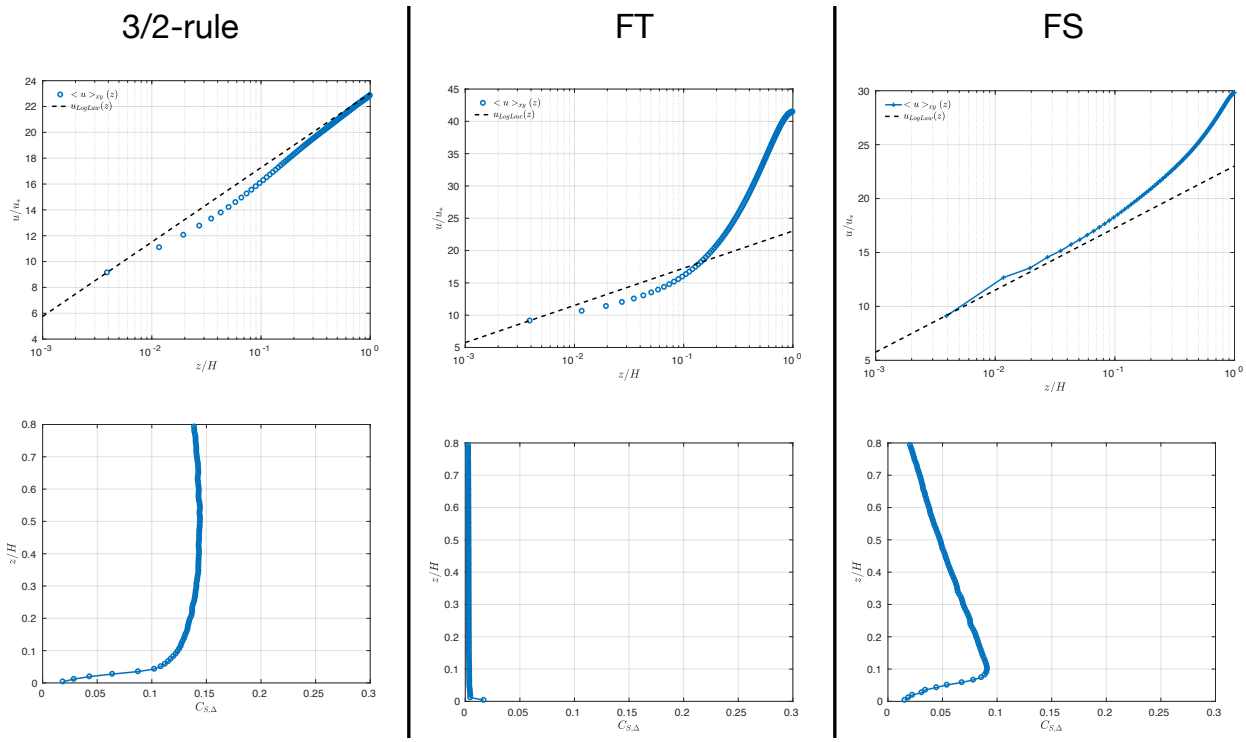

Figure 4: Profiles of the horizontal velocity and the Smagorinsky constant for the three deasliasing methods at a resolution of 128^3.

with dynamic SGS models unless the dynamic models are properly adjusted with filtering at scales larger than those affected by the truncation or smoothing. Therefore, it is impossible to *re-invest* the computation gain due to the FT or FS methods into a more sophisticate SGS model (at lease not a dynamic Smagorinsky-type model from Germano (1991)).

In relation to the second comment, adding extra scalars will decrease the cost of the momentum solver with respect to the total cost of the simulation. Each scalar field is advanced in time using an advection-diffusion equation that also requires dealiasing. However, the cost of the dealiasing of the latter equation is less expensive than the NS equations (especially in rotational form). To summarize, each dealiasing operation becomes less expensive with the FT or FS method, and the total gain will be more important as more operations are required. However, the total gain is not trivial to evaluate *a priori*.

Changes in manuscript: We agree that both issues are relevant, and hence we have included some additional comment in the new discussion section (reported above).

Minor comments:

1. **Reviewer comment:** *Page 4, Line 33 - please check that the N^3 term in the cost is correct here.*

   Authors response: Thank you for pointing out this mistake. This cost should not have a N^3 but only N.

2. **Reviewer comment:** *Page 5, Line 10 - delta implicit does not really correspond to a top-hat filter. The properties of an implicit filter are tied to the discretization scheme. As an example, for a true 3D spectral code the implicit filter is a spectral cutoff filter.*

   Authors response: Thank you this is a good point. This has been corrected in the text.