Response to Referee Comment 2 (RC2)

We would like to thank Anonymous Referee #2 for reviewing our manuscript and for their helpful feedback. Our responses to the individual comments are provided below. The reviewer's comments are shown on a light green background, whereas our responses appear on an ordinary white background.

The authors coupled the dynamic TenStream solver, a recently developed and computationally efficient 3D radiative transfer approximation, to the PALM large-eddy simulation model. The authors then study how 3D radiative transfers affects the development of cloud field properties and organization, and evaluate how well the dynamic TenStream solver perform compared to the original TenStream solver. The paper is interesting and generally well-written. Subject to the following comments, I consider it suitable for publication.

Main comments:

1. What are the initial profiles, surface values and other parameters based on? To me, the start of section 2.1.2 "simulations were set-up for the geographic location of Munich … started on 14 June 2023" suggests you are reproducing a realistic case study here, but it is not clear too what extent the settings and initial profiles are based on observations or whether it is somewhat idealized case study. If it is largely based on observations, are there any cloud (lidar, satellite) or surface irradiance observations available to compare against simulation results?

Thank you for pointing this out, as this could indeed create confusion in the originally submitted manuscript. The simulations are not intended to reproduce a specific observed case, and the initial profiles and surface parameters are not based on observations either. Instead, they represent an idealized case study. The geographic location and date were chosen to prescribe a realistic solar geometry and diurnal cycle, but not to simulate the actual weather conditions over Munich on that day. Accordingly, no comparison against cloud or surface irradiance observations was performed.

To make this clearer, we revised the first sentence of Sect. 2.1.2 as follows: "The simulations*, which represent an idealized case study,* were set up for the geographic location of Munich, Germany (48.1°N, 11.6°E), but with its altitude set to 0 m, i.e., to sea level with a surface pressure of 1013.25 hPa."

2. The results section alternates between discussing 1D vs 3D radiative effects and discussing the original vs dynamic Tenstream. In find that this alternation breaks the logical flow of the text and makes the discussion of the results sometimes a bit confusing, as it is not always clear whether "3D" refers to the dynamic or original TenStream solver. Since the 3D effects (compared to 1D) are sufficiently large compared to the differences between the two Tenstream solvers, I would separate the two parts more clearly: first discuss how and why switching from 1D to 3D (dynamic Tenstream) changes your cloud field, secondly discuss how any errors between the dynamic and original Tenstream solver may affect your results.

We agree that the alternation between the discussion of 1D–3D radiative effects and the comparison between the dynamic and original TenStream solvers made the logical flow of the original results section less clear than intended. In the revised manuscript, we therefore restructured the results section so that the two aspects are separated more explicitly. In each relevant subsection or subsubsection, we now first discuss the effects of switching from 1D to 3D radiative transfer on the cloud field and then separately assess whether the dynamic TenStream solver reproduces the corresponding original TenStream results.

We decided not to split the results into two fully independent sections, though, because the discussion is organized around the figures and the narrative they build. A complete separation would have required the reader to repeatedly return to figures discussed several pages earlier when reaching the dynamic TenStream validation. Instead, we now separate the two aspects within each figure-based discussion while preserving the overall narrative flow. In addition, Sect. 3.3 has been divided into three subsubsections, which distinguish between the two proposed mechanisms linking radiation and cloud characteristics and the final summary/limitations. We believe that this revised structure makes clearer when we discuss 1D–3D radiative effects and when we assess remaining differences between the dynamic and original TenStream solvers.

3. Throughout the manuscript, I would stress more explicitly how the understanding of 3D radiative effects from this work extend beyond the work of e.g. Jakub & Mayer (2017), Veerman et al. (2020,2022), Tijhuis et al. (2024). That is, better highlight the novel aspects regarding 3D effects. For example, the inclusion of 3D longwave radiative effects, cloud streets that change orientation following the sun, and, if considered realistic (see later comment), 3D effects that continue into the night.

Thank you for pointing this out. We agree that the originally submitted manuscript did not always highlight clearly enough where the present study extends beyond previous work on 3D radiative effects on shallow cumulus clouds. We have therefore revised several parts of the manuscript to emphasize the novel aspects more explicitly.

First, we now highlight more clearly that, to our knowledge, this study demonstrates for the first time that radiatively driven cloud streets can develop under a realistic diurnal cycle. This extends the idealized setup of Jakub and Mayer (2017), where the solar geometry was kept fixed.

Second, we now emphasize more strongly the longwave-driven surface-energy-budget mechanism identified in this study. In particular, the simulations suggest that 3D radiative transfer reduces the domain-averaged net thermal emission at the surface, which affects the surface energy budget and is primarily balanced by an increased latent heat flux into the atmosphere. To our knowledge, this pathway by which 3D radiative transfer can affect cloud development has not been described in the literature before. We therefore discuss it in its own subsubsection in Sect. 3.3 and also highlight its novelty in the summary.

Third, we clarified that the present setup also allows differences between 1D and 3D radiative transfer to be examined after sunset, when only thermal radiative transfer is active. However, we discuss these nighttime results more cautiously, because the observed nighttime cloud characteristics may still partly reflect the imprint of daytime cloud evolution. We therefore state that additional experiments, such as simulations restarted at sunset, would be required to separate genuinely nocturnal radiative effects from the memory of daytime development.

Together, these revisions should make clearer which aspects of the present study confirm previous work and which aspects extend it.

additional comments

1. 325 It may be worth further looking into this nighttime cloud layer, whether we can consider it realistic and why the cloud do not disappear. Besides looking at soil moisture, did you consider the initial profiles. Could you maybe quantify the convection in the sub-cloud layer, perhaps time series of the vertical velocity variance may indicate when the convections dies out.

We agree that the persistence and realism of the nighttime cloud layer is an interesting question that deserves further investigation. In the present study, we mainly used the nighttime period to assess whether differences between simulations driven by 1D and 3D radiative transfer also occur when only thermal radiative transfer is active. We certainly agree that a more detailed analysis of the nighttime cloud layer would be valuable, but consider it beyond the scope of the present study, since such an analysis should probably also directly address the causes of the observed nighttime cloud characteristics. This, however, would require additional experiments, such as restarting the simulations at sunset, to separate the imprint of daytime evolution from genuinely nocturnal radiative effects.

2. Consider omitting section 2.2.1, to me a separate section detailing the MBE is not necessary.

We agree that, in the originally submitted manuscript, a separate subsection describing the mean bias error may have seemed somewhat excessive. However, following a comment by Reviewer #1 regarding the statistical rigor of the dynamic TenStream validation, we expanded the original main/control-run setup to a five-member ensemble. In the revised manuscript, the biases are thus no longer calculated only with respect to a single TenStream reference run. Instead, they are calculated pairwise for each ensemble member relative to the corresponding original TenStream ensemble member initialized from the same restart state, and the resulting paired bias time series are then averaged. Because this procedure is less straightforward than the original calculation, we decided to retain Sect. 2.2.1, but revised it to describe the new ensemble-based bias calculation more clearly.

3. Runtime is briefly discussed in the summary, referring to Maier et al. 2024. Please elaborate on computational performance and limitations (run time, memory) of the simulations with 1D and with both 3D solvers earlier on in the manuscript (methods or results section).

Thank you for this suggestion. We agree that the computational performance and limitations of the different radiative transfer solvers should be mentioned earlier in the manuscript. We therefore added a short discussion of the computational cost of the dynamic TenStream solver to the solver description in Sect. 2.1.3. Specifically, we now state that the dynamic TenStream solver reduces the cost of 3D radiative transfer calculations through incomplete solves and that, in the offline benchmarks of Maier et al. (2024), it was shown to be about three times slower than a traditional 1D δ-Eddington approximation, but substantially faster than other currently available 3D solvers.

We did not add a detailed wall-clock-time comparison based on the PALM simulations performed in this study, however, because these simulations were conducted on a highly heterogeneous computing cluster. As a result, wall-clock times varied strongly depending on the assigned nodes and therefore do not provide a reliable basis for such an assessment.

4. Figure 6: it is not clear to me what the added value is of this figure over figure 4

We agree that the former Fig. 6 was partly repetitive and provided only limited additional information beyond Fig. 4. We therefore removed the figure, together with the corresponding discussion in lines 328–337 of the original manuscript, as this part mainly summarized results that had already been discussed beforehand. The subsequent discussion of the differences between the dynamic and original TenStream solvers was revised accordingly and now refers directly to Fig. 4(a)–(c), rather than to the lower row of the former Fig. 6.

5. Since you have 3D longwave and that turns out to be important, could you add spatial plots of net longwave irradiance (similar to fig 7 or 9) to aid the reader.

Thank you for this suggestion. We agree that a spatial plot of the net thermal surface irradiance is interesting and can help illustrate the differences between 1D and 3D radiative transfer in the thermal spectral range. We therefore generated such a plot and include it in this response document for reference. Since thermal radiative transfer remains relevant after sunset, the plot covers a slightly longer time period than the corresponding shortwave figure in the manuscript.

However, we decided not to include this figure in the manuscript itself. The main longwave-related result discussed in the paper concerns the domain-averaged reduction in net thermal emission and its effect on the surface energy budget. This is already shown quantitatively in Fig. 10 and connected to the sensible and latent heat fluxes in Fig. 11 of the originally submitted manuscript. The additional spatial plot is consistent with this interpretation, but does not add information that is essential for the main narrative. Since the revised manuscript already contains several figures and has also been restructured to improve focus, we hence decided to keep the results section more streamlined and not add this figure to the manuscript.

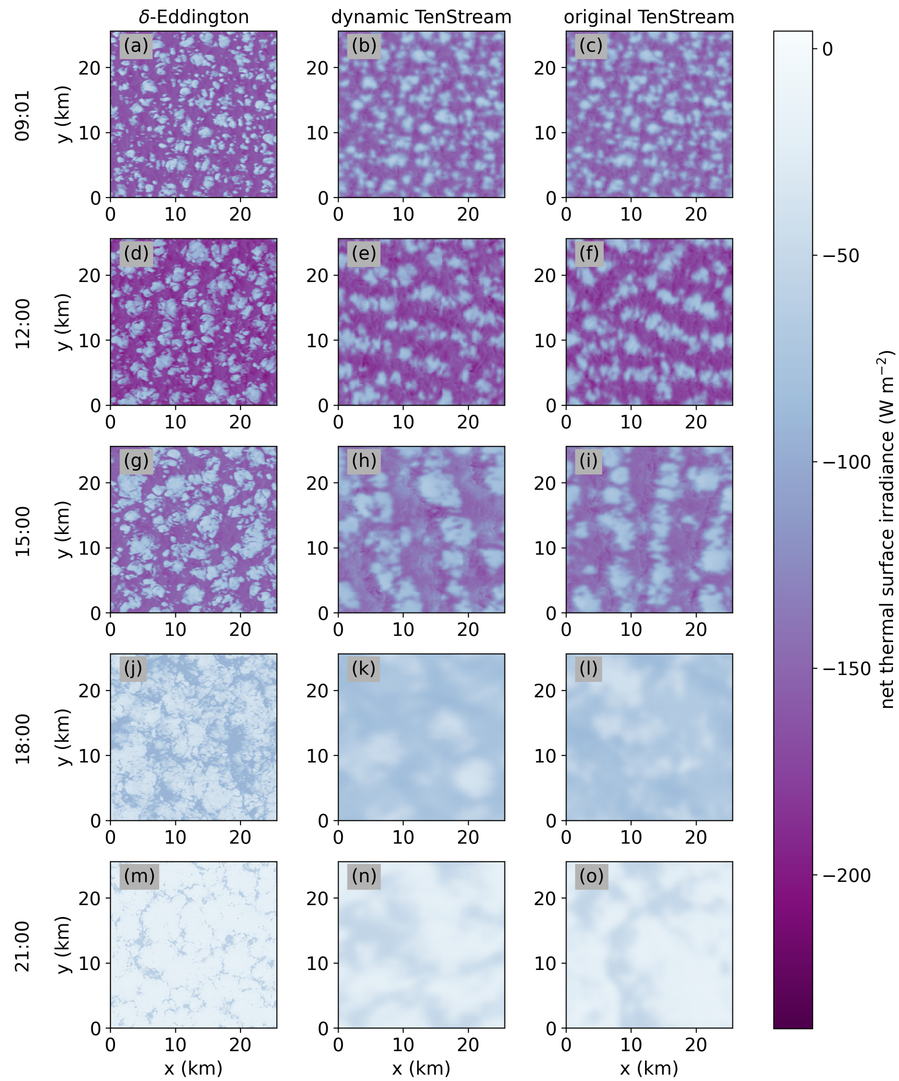

Figure: Temporal evolution of the net surface irradiance in the thermal spectral range for the simulations driven by the δ-Eddington approximation (left), the dynamic TenStream solver (middle), and the original TenStream solver (right), shown for five time steps between 09:01 and 21:00 UTC. Only simulations performed from the main run of the setup are shown. In contrast to the corresponding shortwave figure, the time period extends beyond sunset because thermal radiative transfer remains relevant at night.