

Response to Referee Comment 1 (RC1)

We would like to thank Anonymous Referee #1 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 blue background, whereas our responses appear on an ordinary white background.

This manuscript couples the dynamic TenStream 3D radiative transfer solver to the PALM large-eddy simulation model and investigates how 3D versus 1D radiative transfer treatments affect shallow cumulus cloud field evolution over a full diurnal cycle. It also serves as an online validation of the dynamic TenStream solver against the original, more expensive TenStream implementation. The manuscript provides a number of novel and significant insights including:

- Demonstrating for the first time that dynamic surface heterogeneities from 3D radiative transfer (RT) with a realistic diurnal cycle can produce cloud streets
- Cloud cover with a 1D treatment of RT produces higher cloud coverage than simulations with a 3D RT treatment for high coverage (40-50%) shallow cumulus days
- The importance of 3D RT in the longwave for nighttime cloud field properties, with the clouds with 1D RT developing significantly higher average LWP and deeper fields, and converging toward full coverage sooner than clouds with 3D RT.
- Longwave 3D RT effects impact mean surface latent heat flux; reduced thermal emission from the surface in the 3D RT simulations (due to thermal emission from clear and not only cloudy regions in the 3D RT simulations) is offset mainly by increased latent heat flux, mediated by increases in surface skin temperature.

Separately from the physical insights provided by the manuscript, this work also includes validation of the dynamic TenStream solver against the original version:

- For a high cloud coverage shallow cumulus case in which the dynamic TenStream solver is expected to perform well, the dynamic TenStream shows similar daytime cloud cover, average LWP, and cloud depth as the original TenStream solver and modestly diverging cloud cover and LWP at nighttime.

Below are some suggestions to improve this already strong manuscript.

Manuscript organization

The manuscript tackles two objectives: (1) to validate the performance of the dynamic TenStream solver against the original, higher-fidelity TenStream solver in an online setting and (2) to discuss how 3D versus 1D treatments of radiation impact cloud field statistics of shallow cumulus. While there are some subtle lessons about the cloud field to be learned from the differences between the dynamic and original TenStream solvers (i.e. lingering surface enhancements in the dynamic version), for the most part these two objectives appear separate. It would benefit the reader for the authors to treat them as such, particularly within sections 3.2 and 3.3. Subsections might help here.

We agree that the two main objectives of the manuscript were not always separated clearly enough in the originally submitted version, particularly in Sects. 3.2 and 3.3. In the revised manuscript, we

therefore restructured the results section to make this distinction more explicit. In each relevant subsection or subsubsection, we now first discuss how treating radiation in 3D rather than 1D affects the cloud field, before separately assessing whether the dynamic TenStream solver is able to reproduce the corresponding original TenStream results.

We decided not to separate the two objectives into two completely independent sections, though, because the results section 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 objectives within each figure-based discussion while preserving the overall narrative flow. As part of this restructuring, Sect. 3.3 has also 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 the role of the dynamic TenStream validation clearer while maintaining the readability of the results section.

Emphasis of novel results

Several of the manuscript's most interesting results could be given greater prominence.

This work demonstrates for the first time that radiatively driven cloud streets can develop with 3D RT under *realistic diurnal conditions*, clearly distinguishing it from prior work by Jakub and Mayer (2017) in which solar zenith angle was kept fixed. The manuscript would benefit from better highlighting this result.

We thank the reviewer for pointing this out. We agree that this is an important distinction from the work of Jakub and Mayer (2017), and that the result deserved more emphasis in the originally submitted manuscript. We therefore revised the discussion in Sect. 3.1 to highlight that the formation of radiatively driven cloud streets is observed here under a realistic diurnal cycle in which the cloud streets follow the changing azimuthal position of the Sun. Specifically, we added the following sentence to the paragraph corresponding to lines 261–271 of the original manuscript:

“[...] This absence of organization in the simulation driven by 1D radiation together with organization perpendicular to the mean westerly flow in the simulations coupled to 3D radiative transfer indicates that the cloud streets in these simulations are not dynamically induced, but really driven by radiation. *Moreover, extending beyond the idealized setup of Jakub and Mayer (2017), this demonstrates, to our knowledge, for the first time that radiatively driven cloud streets can also form under a realistic diurnal cycle, with their orientation following the changing solar azimuth angle.* [...]”

Furthermore, we now also emphasize this aspect in the summary:

“[...] To analyze the differences between the three simulations, daytime and nighttime conditions were considered separately. During the day, the clouds in the dynamic and full TenStream simulations were shown to organize into cloud streets aligned perpendicular to the angle of solar incidence, *extending the findings of Jakub and Mayer (2017) to a realistic diurnal cycle in which the cloud streets follow the changing azimuthal position of the Sun.* With 1D radiation, on the other hand, the

clouds remained more or less randomly positioned. [...]”

These changes give greater prominence to this novel aspect of the simulations while keeping the emphasis proportional to the scope of the analysis, since the cloud streets are discussed only qualitatively here rather than analyzed as a separate quantitative focus of the study.

This reader is fascinated by the result that for higher cloud cover days, cloud coverage for simulations with a 1D treatment of radiation is higher than in simulations with a 3D treatment of radiation (supported by the results of Tijhuis et al. 2024). Why might this be? Further discussion would be greatly appreciated.

We agree that the lower cloud cover in the 3D simulations under high-cloud-cover conditions is an interesting result, especially since it appears to be consistent with Fig. 3 of Tijhuis et al. (2024), although this aspect has not been emphasized much so far. At present, however, we cannot really identify the mechanism responsible for this behavior from the simulations presented here. Isolating this mechanism would likely require additional experiments specifically designed for this purpose, which is beyond the scope of the present study. We therefore clarified this point in the revised manuscript following the sentence corresponding to line 306 of the original manuscript:

“[...] This differs from findings in earlier studies, which usually found similar cloud cover in simulations driven by 1D and 3D radiative transfer (e.g., Jakub and Mayer, 2017; Veerman et al., 2020; Tijhuis et al., 2024). However, the overall cloud cover in these studies was usually much lower. And for values above approximately 30%, Tijhuis et al. (2024) actually also found mostly lower cloud cover in simulations coupled to 3D radiation, in good agreement with the results shown here (see their Fig. 3). *The mechanism behind this behavior, however, remains unclear and requires further research.*”

The manuscript shows that longwave 3D RT impacts the surface energy budget via reduced thermal emission, propagating into increased surface latent heat fluxes. This higher moisture flux into the atmosphere from the surface is hypothesized by the authors to partially explain the higher LWP in the 3D RT simulations. To this reviewer’s knowledge, this mechanism is novel in the literature and deserves greater focus. Do the authors believe it is equally significant to the mechanism whereby clouds in the 3D simulations are positioned over areas of enhanced net surface irradiance, strengthening rather than suppressing the associated updrafts? Can the two mechanisms be isolated to assess their relative importance? If computational resources allow, simulations with 3D RT for shortwave only and with 3D RT for longwave only would be of interest. Work from Veerman et al. 2022 and Tijhuis et al. 2024 suggests that the shortwave mechanism may be dominant as they see substantial LWP increases with 3D RT without longwave effects.

We thank the reviewer for pointing this out. We agree that the longwave-driven effect on the surface energy budget, and the resulting increase in latent heat flux, is a particularly interesting result of this study. To our knowledge, this mechanism has not been described in the literature before. In the revised manuscript, we therefore give this result greater prominence by discussing it in its own subsubsection

within Sect. 3.3 and by explicitly highlighting its novelty. The discussion in Sect. 3.3.2 now ends with:

“[...] Overall, the increased latent heat flux in the 3D simulations provides another potential explanation for why clouds grow thicker and larger during the day than their 1D-driven counterparts, as more water vapor is released into the atmosphere with 3D radiation. In contrast to the previously discussed mechanism related to the positioning of clouds relative to their shadows and sunlit regions, this longwave-driven effect has, to our knowledge, not been described in the literature before.”

We think, however, that the relative importance of this longwave-driven surface-energy-budget mechanism compared to the shortwave mechanism related to cloud positioning within the solar net surface irradiance field cannot be isolated from the present simulations. We therefore added a corresponding limitation to Sect. 3.3.3:

“The relative importance of these two effects, i.e., the extent to which each of them contributes to the observed changes in cloud characteristics, however, has not yet been investigated. Quantifying their individual contributions would require additional sensitivity experiments, for example simulations in which 3D radiative transfer is applied only in the solar or only in the thermal spectral range. Such simulations would make it possible to assess how much the longwave-driven surface-energy-budget effect contributes relative to the positioning of clouds within the solar net surface irradiance fields. [...]”

Such sensitivity experiments would indeed be highly valuable. At present, however, the model setup does not allow different radiative transfer solvers to be assigned independently to the solar and thermal spectral ranges, so these experiments would require additional implementation work as well as substantial additional simulations. We therefore consider this beyond the scope of the present study, but now explicitly mention it as an important direction for future work.

The differences in 3D versus 1D shortwave surface flux fields, as described at the start of section 3.3, are well documented in the literature. A briefer introduction of these concepts might allow the authors to place stronger emphasis on (a) the connection between surface flux heterogeneities and cloud field evolution (b) what the model artifacts of the dynamic TenStream solver and resulting cloud field differences can teach us about this mechanism.

We agree that the introduction of the 1D–3D differences in the shortwave surface flux fields at the beginning of Sect. 3.3 was overly detailed, especially since these effects have already been described in the literature. In the revised manuscript, we therefore shortened the corresponding discussion in lines 366–388 of the preprint and now only briefly introduce the concepts needed for the subsequent interpretation. The revised text thus places stronger emphasis on how the spatial heterogeneity of the surface irradiance fields may affect the cloud field evolution, in particular through the positioning of clouds relative to their own shadows and to regions of enhanced net surface irradiance. We also revised the subsequent discussion of the dynamic TenStream artifacts to more clearly connect them to the delayed cloud–surface radiative coupling in the dynamic solver.

Statistical rigor of dynamic TenStream validation

Dynamic TenStream and original TenStream solver differences should not be evaluated solely by comparing their differences with those between 1D and original TenStream; instead the authors should leverage their (modest) ensemble, asking if the observed differences are greater or less than the variability observed for original TenStream. Using additional ensemble members, if computationally feasible, would allow for a more statistically rigorous analysis of the signal to noise ratio of these differences.

We agree that the originally submitted manuscript relied on a rather small ensemble to assess whether differences between the dynamic and original TenStream solvers are robust compared to the intrinsic variability of the simulations. In the original version, this variability was estimated from a main/control-run comparison. While this comparison already provided a first estimate of internal variability, we agree that a larger ensemble provides a more transparent and statistically more meaningful basis for the evaluation.

We therefore expanded the ensemble from two to five members for each radiative transfer solver, which was the largest ensemble size that we could justify given the substantial storage requirements of the 3D model output. The simulation setup described in Sect. 2.1.4 was revised accordingly. To reflect that change, in the revised manuscript, Figs. 4, 5, 9 (former Fig. 10), and 10 (former Fig. 11) now show ensemble means and one standard deviation across the corresponding five-member ensemble for each radiative transfer solver, rather than showing the main and control runs separately.

We also revised the bias evaluation. Instead of calculating MBEs only with respect to the TenStream reference main run, the revised analysis calculates the bias pairwise for each ensemble member relative to the corresponding original TenStream run initialized from the same restart state. The resulting five paired bias time series are then averaged, and their standard deviation is shown as a measure of ensemble spread. This revised procedure is described in Sect. 2.2.1.

The text and figure captions throughout the manuscript were updated to reflect this ensemble-based evaluation. Overall, the main conclusions are unchanged, but the revised analysis now provides a clearer estimate of the signal-to-noise ratio of the solver differences and better supports the conclusion that the dynamic TenStream solver reproduces the relevant 3D-related effects within the ensemble variability.

Is the top row of Figure 6 a useful and fair comparison for the performance of the dynamic TenStream model? Does it provide additional information beyond Figure 4?

We agree that the former Fig. 6 was partly repetitive and provided only limited additional information beyond Fig. 4. In particular, the top row mainly compared the average cloud-characteristic quantities between the TenStream reference solver and the 1D δ -Eddington approximation, thereby summarizing results that were already discussed based on Fig. 4. We therefore removed the former Fig. 6, together with the corresponding discussion in lines 328–337 of the original manuscript. 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).

General suggestions and questions

The authors should take care to distinguish between surface field lag artifacts of the dynamic TenStream solver and the physical displacement of cloud shadow from cloud root that occurs at moderate solar zenith angles in the original TenStream results (lines 415–420). The former is a solver artifact arising from incomplete solves; the latter is a real 3D radiative effect. The current text could more clearly distinguish between the two.

Thank you for pointing this out. We agree that this paragraph was not phrased precisely enough. Both the dynamic and the original TenStream solver show a physically correct displacement of the cloud shadows to the east when the Sun is shining from the west at a moderate solar zenith angle, as in this figure. The main point here, however, was that the regions of enhanced net surface irradiance are less concentrated underneath the clouds in the dynamic TenStream simulation than in the original TenStream simulation. This difference is caused by the already discussed delayed cloud-surface radiative coupling when using incomplete solves. To make this distinction clearer, we have revised this part of the manuscript as follows:

“Since all the main features of the original TenStream fields are thus captured by the dynamic TenStream solver, we can now turn to Fig. 7 to investigate whether, similar to the original TenStream run, the clouds in the dynamic TenStream simulation are also maintained over regions of enhanced net surface irradiance. And indeed, we can observe a similar positioning of the clouds relative to their shadows, which are displaced to the east as a physical 3D radiative-transfer effect also in the dynamic TenStream simulation. In contrast to the original TenStream simulation, however, the regions of enhanced net surface irradiance are less concentrated underneath the clouds and appear spatially broader in the dynamic TenStream simulation. This broadening of the radiatively enhanced areas has already been discussed and can be attributed to the incomplete solves, which delay the interaction between radiative fluxes at the clouds and the surface. Specifically, as the clouds move eastward, both the response of the surface to the clouds and the subsequent radiative feedback from the surface back to the clouds can lag behind. Despite this lag, however, substantial portions of the clouds remain positioned over regions of increased net surface irradiance. Thus, also in the dynamic TenStream simulations, the surface irradiance pattern likely strengthens rather than weakens the corresponding updrafts, thereby promoting the observed differences in cloud characteristics.”

What is the surface heat capacity C_0 used in the PALM land surface model for this setup? If it is small, would the authors not expect a near-instantaneous response of the surface energy budget to changes in the radiative field (lines 425–435)?

Thank you for pointing this out. We checked the setup and found that the PALM surface skin-layer heat capacity used in our simulations is indeed $C_0 = 0 \text{ J m}^{-2} \text{ K}^{-1}$. We agree that this implies an effectively instantaneous adjustment of the surface energy budget to changes in the radiative field. We have therefore revised lines 422–440 of the originally submitted manuscript to explicitly mention this

point and to clarify that the correspondingly rapid surface response may explain why the updrafts can be maintained in sunlit regions even when the clouds move into their own shadows:

“[...] In the qualitative analysis presented here, no such strong dependency on the wind–sun angle is observed, since the clouds continue to grow even when moving into their own shadows, as can be seen starting with the panels at 13:00 UTC in Fig. 6. However, this does not necessarily negate the existence of this mechanism. *Instead, it suggests that in our simulations, the surface response to cloud movements is sufficiently rapid to maintain the updrafts in sunlit regions. This is consistent with the surface skin-layer heat capacity of $C_0 = 0 \text{ J m}^{-2} \text{ K}^{-1}$ used in these simulations, which causes an instantaneous adjustment of the surface energy budget to changes in the radiative field.* Otherwise, the results align with those of Tjihuis et al. (2024), although the connection between the clouds and areas of increased net surface irradiance has been emphasized much more strongly here. [...]”

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.