Response to Review Comment 4 (RC4)

Manuscript: egusphere-2023-2129
Title: A dynamic approach to three-dimensional radiative transfer in numerical weather prediction models: the dynamic TenStream solver v1.0
Authors: Richard Maier, Fabian Jakub, Claudia Emde, Mihail Manev, Aiko Voigt, and Bernhard Mayer

We thank Anonymous Referee #3 for his or her comments on our manuscript, which we will respond to below. To structure our response, the referee's comments are printed on a gray background color, while our answers are displayed on ordinary white background.

* General comments:

This is a welcome update of the TenStream 3D radiative transfer (RT) code that already fills a major gap in LES modeling capability, namely, to perform 3D RT broadband radiation budget estimation for Large-Eddy Simulation (LES) models. LES is now routinely used in cloud-scale process modeling to address some of predictive climate science's most stubborn issues, such as cloud feedbacks and aerosol-cloud interactions.

However, in spite of generating fully 3D (i.e., vertically-developed) clouds driven by convective dynamics, the RT parameterizations used in LES are still too often heritage codes from Global Climate Models (GCMs) where nothing less than ~50 to 100 km in scale is resolved, hence all clouds and many cloud systems. A typical aspect ratio for a GCM cloudy column is therefore on the order of 1-to-10, thus, some form of 1D RT that captures the internal variability of the clouds (e.g., McICA) is justified since little radiation will be leaked through the horizontal boundaries anyway. In sharp contrast, a cloudy column in an LES has the opposite aspect ratio: say, 5 km by 50 m, hence about 100-to-1. Even cloud-resolving models (CRMs), say, at 5 km by 0.5 km are 10-to-1. NWP models are heading into that kind of spatial resolution as well. So there is plenty of opportunity for net horizontal fluxes to develop across grid-cell facets, starting with direct shadowing of neighboring cells in the anti-solar direction. The TenStream model is purposefully designed to account for this 3D RT in terms of radiation energetics, hence fluxes, not radiances, as required for computing heating rates profiles and net fluxes through the top and bottom boundaries.

The new _dynamic_ TenStream model is designed to address the issue of computational efficiency that is in the way of the general acceptance of TenStream in the LES community for operational implementation. Specifically, it brings CPU time allocation down to ~3x the baseline cost of 1D RT, and does so by cutting a few corners, which could carry a cost in accuracy. Therefore, dynamic TenStream is benchmarked for accuracy against the original TenStream, as well as 1D RT (delta-Eddington) and full 3D RT (MYSTIC). Its accuracy is at par with the original TenStream, which is already a vast improvement in accuracy for radiation budget estimation using standard 1D RT.

The paper is well written and illustrated. It should be published by GMD after a minor revision that addresses the following issues.

* Specific comments:

(1) Careful attention is paid to the heating-rate profile and surface irradiance/flux. However, it seems to me that the outgoing TOA flux is also important. Maybe TenStream enforces radiant

energy conservation is such a way that the TOA flux is as accurate as the rest, but that isn't obvious to this reviewer. At a minimum, some kind of statement on TOA flux accuracy is in order.

You are right that we focused our evaluation on heating rates and net surface irradiances, as they are the main drivers of the weather. For the revised version, we extended Section 4.3 to also account for the performance of the new solver in determining net irradiances at top of atmosphere (TOA). The content of this extension is centered around two new plots. The first one shows the temporal evolution of the mean absolute error (MAE) in the net irradiance at TOA in an otherwise similar fashion as Fig. 8 in the preprint:

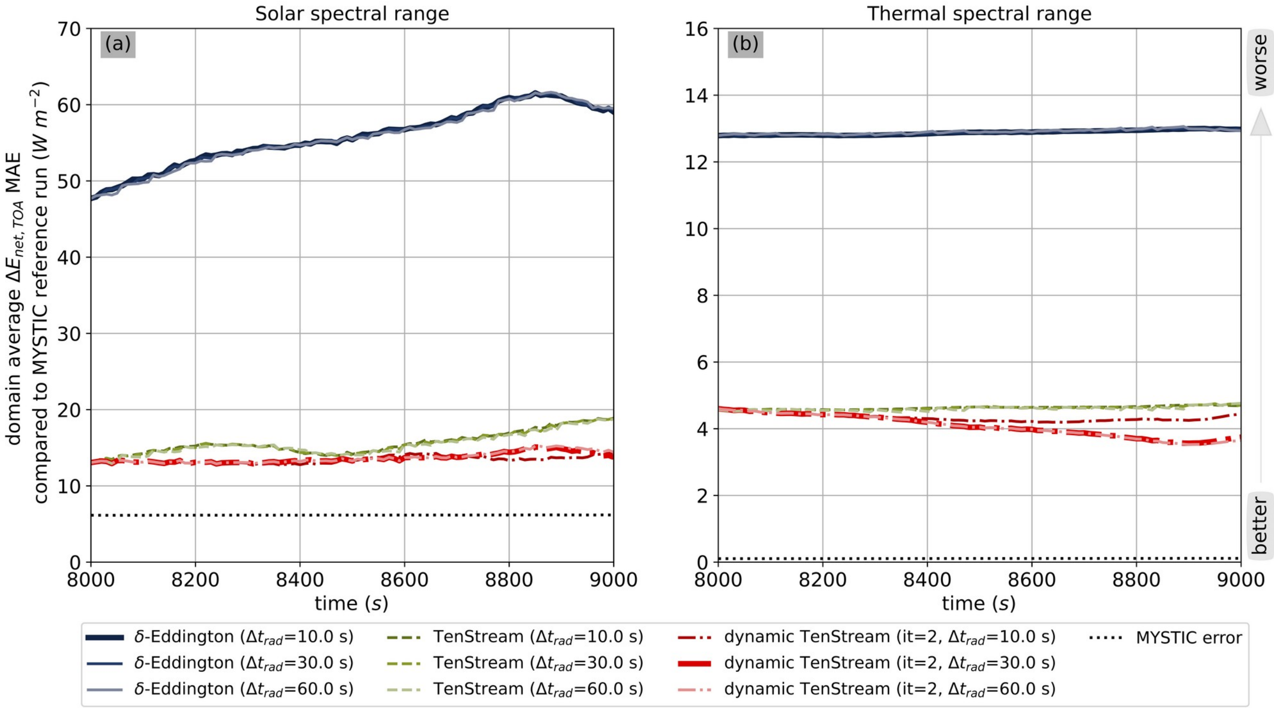

The plot shows that also at TOA, the 1D delta-Eddington solver performs worst (blue lines) in terms of the MAE, with the original TenStream solver (green lines) once more being a noticeable improvement, remaining significantly below the error of all 1D runs in both the solar and the thermal spectral range. Our newly developed dynamic TenStream solver (red lines) shows just slight deviations from the MAE of the full TenStream calculations, almost independent of the calling frequency used, and thus also stays significantly smaller than the MAE of any 1D delta-Eddington run throughout the entire time series – even at the lowest calling frequency of 60 s.

However, similar to the results obtained for the net surface irradiance, the performance of both the original, as well as our new dynamic TenStream solver is worse in terms of our other error measure, the mean bias error (MBE). The temporal evolution of this error measure in terms of the net irradiance at TOA is shown in the other new figure below. It shows that in the solar spectral range, the MBE for the new dynamic TenStream solver (red lines) clearly diverges from the MBE of the original TenStream solver (green lines). This spread from the original TenStream solver gets significantly larger at a calling frequency of 30 s compared to the dynamic TenStream run at a calling frequency of 10 s. Interestingly, however, the spread does not further increase when calling dynamic TenStream even more infrequently (bright red line). And in both cases, the MBE of the dynamic TenStream runs does not continuously increase, but stabilizes itself at some point in time.

And even for calling frequencies of 30 s and 60 s, the MBE peaks at values of around 8.5 W m^{-2}, which translates into a RMBE of about 1.2 % (not shown here).

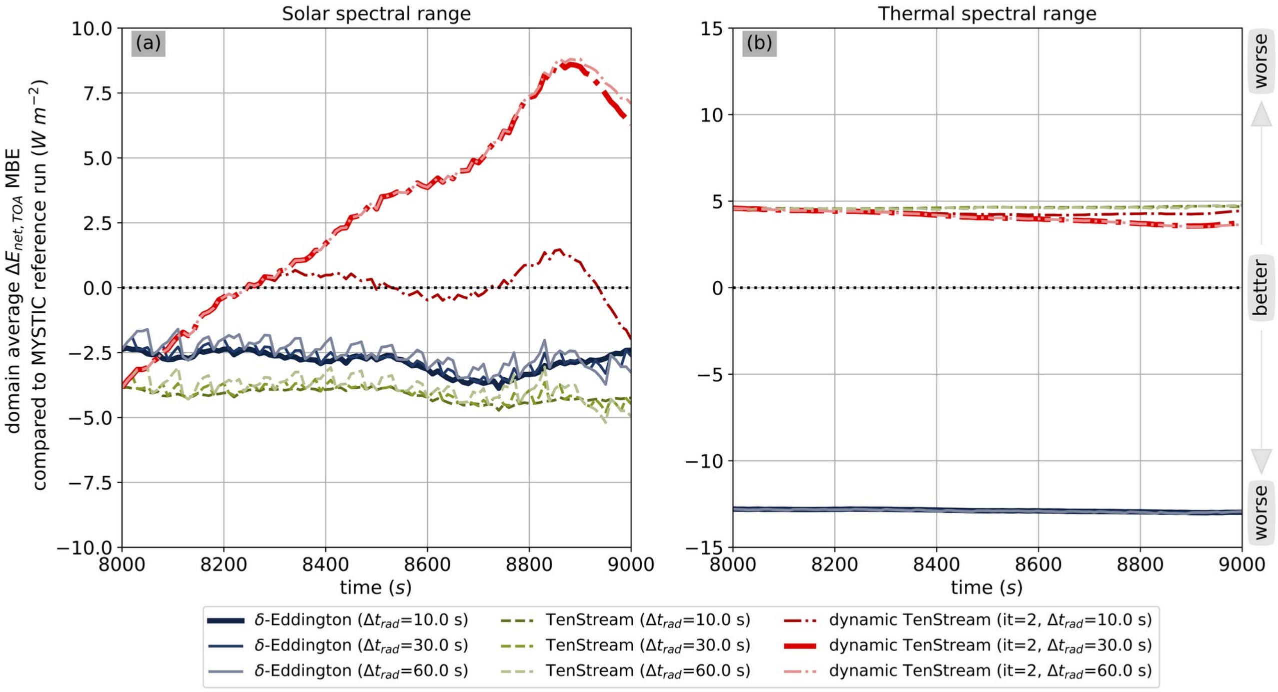

In the thermal spectral range, on the other hand, the MBE in the net TOA irradiance for both the original, as well as the new dynamic TenStream solver stays significantly below the error of the 1D delta-Eddington runs throughout the entire time series, peaking at values of around 5 W m^{-2} (-2 %) for the 3D solvers compared to -13 W m^{-2} (5 %) for the delta-Eddington solver.

(2) The temporal down-sampling and the incomplete solves naturally cause the new model to drift away from the original counterpart. Would it not be beneficial to occasionally "reset" this drift to zero by calling the original TenStream? Of course there is a whole study to perform about when to do this operationally, without the benchmark information at hand.

This is a good idea, and one that we definitely had in mind when thinking about future couplings of our new solver to LES or NWP models. The implications of such resets would be relatively straightforward, as the error metrics would simply reduce to those of the original TenStream solver whenever such a reset was performed. For this paper, however, we wanted to focus on how our new solver performs when applied with the lowest computational cost possible – that is, with a low number of two Gauß-Seidel iterations per call and no intermediate resets of the new model.

In the future, however, it would certainly be interesting to investigate the trade-off between increased accuracy due to occasional model resets on one side, and the additional computational cost that these resets introduce on the other side. We have included this thought in the outlook of the revised version of the paper: “In this context, it would also be interesting to investigate whether occasional full solves are a computationally feasible means of ensuring that the results of our new solver do not deviate too much from those of the original TenStream solver.”

(3) Although it should have been done when documenting the original TenStream model, it would be good to look into the past to find models with similar mathematical structure in terms of radical angular simplification compared to standard 3D RT solvers, more precisely with improved efficiency in mind. Can I suggest a few?

- an original "6-flux" model, applied to homogeneous plane-parallel media (but with potential for heterogeneous media):

Chu, C.M. and Churchill, S.W., 1955. Numerical solution of problems in multiple scattering of electromagnetic radiation. The Journal of Physical Chemistry, 59(9), pp.855-863.

- a discrete-angle RT formalism predicated on regular tessellations of 2D and 3D spaces, seeking the minimal number of directions to capture 3D RT effects:

Lovejoy, S., Davis, A., Gabriel, P., Schertzer, D. and Austin, G.L., 1990. Discrete angle radiative transfer: 1. Scaling and similarity, universality and diffusion. Journal of Geophysical Research: Atmospheres, 95(D8), pp.11699-11715.

- a 2D (4-stream) RT model in a deterministic fractal medium, emphasizing numerical implementation (successive over-relaxation scheme):

Davis, A., Gabriel, P., Lovejoy, S., Schertzer, D. and Austin, G.L., 1990. Discrete angle radiative transfer: 3. Numerical results and meteorological applications. Journal of Geophysical Research: Atmospheres, 95(D8), pp.11729-11742.

- the same 2D (4-stream) RT model but in a random multifractal medium, emphasizing numerical implementation (Monte Carlo scheme):

Davis, A.B., Lovejoy, S. and Schertzer, D., 1991, November. Discrete-angle radiative transfer in a multifractal medium. In Wave Propagation and Scattering in Varied Media II (Vol. 1558, pp. 37-59). SPIE.

- vastly faster solution of the 4-stream model using sparse matrix inversion:

Lovejoy, S., Watson, B.P., Grosdidier, Y. and Schertzer, D., 2009. Scattering in thick multifractal clouds, Part II: Multiple scattering. Physica A: Statistical Mechanics and its Applications, 388(18), pp.3711-3727.

Thank you for these suggestions. We also think that these papers should have been primarily mentioned in the documentation of the original TenStream solver. Nonetheless, we included some of these papers into the introduction of the revised version of our paper.

* Technical corrections:

Title: The application to NWP models is both inspirational and aspirational. Here, however, the authors only get as far as LES, or CRM (100 m grid spacing). A more accurate title is in order.

You are certainly right with that. We will change the title to “A dynamic approach to three-dimensional radiative transfer in subkilometer-scale numerical weather prediction models: the dynamic TenStream solver v1.0” for the revised version.

l. 99: i.e., e.g., (need commas, I think)

We have changed the sentence containing this expression to clarify the next comment and added commas behind “i.e.” and “e.g.” elsewhere in the document.

l. 100: "n1" --> what is the "1" for?

We have clarified the meaning of the “1” by adding more explanation to the corresponding example: “For example, the emissivity $e_{0,i,j,k}$ of grid box (i,j,k) in upward direction is equal to the fraction of the downward facing radiative flux $\Phi_{1,i,j,k+1}$ that is absorbed on the way through that grid box, which in turn is one minus the sum of all fractions $a_{n1,i,j,k}$ of $\Phi_{1,i,j,k+1}$ exiting grid box (i,j,k), i.e. $e_{0,i,j,k} = 1 - \Sigma^{9}_{n=0}\, a_{n1,i,j,k}$, where $a_{n1,i,j,k}$ refers to the corresponding entries in the second column of matrix $T_{i,j,k}$.”

l. 104: first "out" --> not italics

We have changed that as suggested.

Fig. 3: For SZA near 45 deg, one could use a diagonal sweep through the grid? Same for ~45 deg in azimuth? Admittedly more tricky to code, but it would follow more closely the propagation of direct sunlight. No? [...]

```
 _____ _____ _____ _____ ___ etc.
|     |     |     |     |
|  1  |  2  |  6  |  7  |  15
|     |     |     |     |
 _____ _____ _____ _____ ___
|     |     |     |     |
|  3  |  5  |  8  |  14 |
|     |     |     |     |
 _____ _____ _____ _____ ___
|     |     |     |     |
|  4  |  9  |  13 |     |
|     |     |     |     |
etc.
```

Thank you for this suggestion. Indeed, one could think about more sophisticated patterns of propagating through the model grid boxes in order to follow the propagation of direct radiation even more closely. However, it would likely not improve convergence, since direct radiation is only represented by three independent streams pointing in x, y and z direction in the dynamic TenStream solver. By properly sorting the resulting three loops due to solar incidence angle, one already ensures that the ingoing direct fluxes of any grid box are always updated before calculating the corresponding outgoing fluxes – even at 45° zenith and azimuth angle. To illustrate that, let us look at a simplified version of Fig. 3 showing only direct streams:

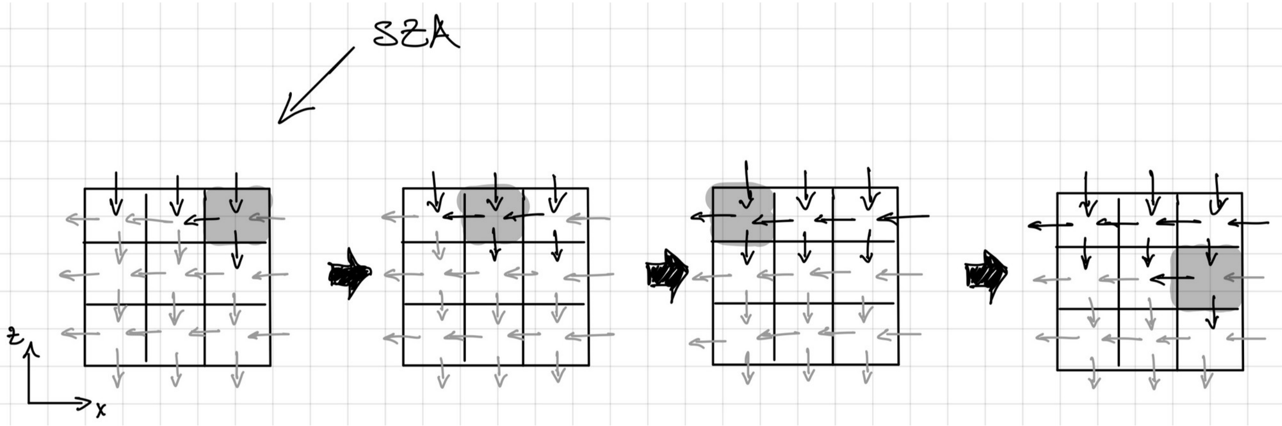

Similar to Fig. 3, this sketch shows the first four steps of one Gauß-Seidel iteration in two dimensions only. In every step, ingoing fluxes are used to update the outgoing fluxes of the corresponding grid box (highlighted in grey). Grey arrows in contrast to black arrows indicate fluxes that have not yet been updated in this Gauß-Seidel iteration. We consider a solar zenith angle of 45° with the Sun shining from the top-right. We can clearly see that even with our not as sophisticated way of iterating through the domain, we always use already updated ingoing fluxes to update the corresponding outgoing fluxes – except for fluxes at the borders of the domain, that are subject to boundary conditions. Due to the definition of the direct streams in the solver, a diagonal sweep through the grid boxes would actually even slow down convergence, as we would not always use already updated ingoing fluxes following such a pattern, although these diagonal sweeps seem to follow the propagation of solar incidence more closely at first.

l. 190: "this direction" -->? horizontal scan

We actually reverse the iteration direction in every other Gauß-Seidel iteration in all three dimensions. To clarify that, we have adjusted the corresponding sentence: “Thus, we reverse the direction of iteration in every other Gauß-Seidel iteration in all three dimensions to not favor propagation of information in one direction.”

S. 3.1 (beginning): specify domain size in cells _and_ km

We have changed that as suggested: “The data set originally features both a high temporal resolution of 10 s and 256 × 256 grid boxes with a high spatial resolution of 25 m in the horizontal.”

l. 262: specify domain height (in km too)

As we pointed out in the paper, the total domain is constructed using two different sources: For the first 220 layers, we use the high vertical resolution of 25 m provided by the LES runs. We clarified the domain height of that part referred to in l. 262: “In the vertical, the modified cloud data set consists of 220 layers with a constant height of 25 m, thus reaching up to a height of 5.5 km.”

Above this vertically highly resolved grid, we use atmospheric levels provided by the 1976 US standard atmosphere, as pointed out further down. To clarify the total domain height, we have thus also extended that part in l. 271: “Apart from the cloud field, the 1976 US standard atmosphere (Anderson et al., 1986) interpolated onto the vertical layers given by the cloud data grid serves as background atmosphere. Above the cloud data grid, the native US standard atmosphere levels as they are provided by libRadtran are used, so that the full grid features 264 vertical layers up to a height of 120 km.”.

Eqs. (6)-(7): why not look at TOA fluxes as well?

We have changed that as suggested.

l. 546: My first encounter with the notion of thermal "shadows". Is there a reference in the literture?

For now, we have not found a reference to these thermal “shadows” in the literature, as investigations are often solely focused on cloudy regions. In addition to that, these thermal shadows are also very small in magnitude, as one has to keep in mind that we used a logarithmic color scale in order to visualize them, as we explicitly stated in the paper.

l. 575: Clarify "feedback effect". Are the LES dynamics driven by a 3D RT model? Or is this a purely (instantaneous) 3D RT effects? BTW, what radiation scheme was used in the LES runs? Should be specified in Section 3.1 (I'm assuming a standard 1D RT model, but may be wrong).

The term “feedback effects” is now explained in more detail: “The reason for these artifacts are the incomplete solves, which can delay lower-order 3D effects, such as feedback effects from other clouds or the surface. The term "feedback effects" thereby refers to the fact that the 3D radiative effects of a cloud can theoretically alter the conditions determining the 3D radiative effects of any other cloud in the domain. Because these feedback effects require multiple back and forth transports of information, they cannot be fully accounted for when solving radiation incompletely.”

We also clarified how the dynamics were driven in the original LES data set by adding the following sentence to the beginning of Section 3.1: “Dynamics in this LES simulation were not driven by radiation, but by a constant net surface flux as described in the namelist input files in Jakub and Gregor (2022).”