



Capturing and explaining the effects of three-dimensional radiative transfer on cloud evolution with the dynamic TenStream solver

Richard Maier¹, Fabian Jakob¹, Fabian Hoffmann^{1,2}, and Bernhard Mayer¹

¹Chair of Experimental Meteorology, Faculty of Physics, Ludwig-Maximilians-Universität München, Munich, Germany

²Institute of Meteorology, Freie Universität Berlin, Berlin, Germany

Correspondence: Richard Maier (richard.maier@physik.uni-muenchen.de)

Abstract. Radiative transfer is an inherently three-dimensional (3D) process that, for computational reasons, is still approximated as one-dimensional (1D) in most atmospheric models. To address this limitation, Maier et al. (2024) introduced the dynamic TenStream solver, which reduces the cost of 3D radiative transfer calculations through incomplete solves. Here, we investigate how coupling dynamic TenStream to the large-eddy simulation model PALM affects cloud development compared to simulations using conventional 1D and full 3D radiation. Results show that during daytime, clouds driven by either of the 3D solvers organize into cloud streets oriented perpendicular to the solar incidence angle, whereas with 1D radiation they remain more or less randomly distributed. Moreover, daytime clouds grow larger, become thicker, and contain more liquid water with 3D radiative transfer. It is shown that these differences arise because, unlike in the 1D case, clouds coupled to 3D radiation are not positioned directly above their own shadows. Instead, they are located over areas of enhanced net surface irradiance, where values even exceed those in the clear-sky columns of the 1D simulation, strengthening rather than weakening the associated updrafts. Additionally, 3D radiation is shown to reduce the domain-averaged net thermal emission at the surface, which affects the surface energy budget and is primarily balanced by an increase in the domain-averaged latent heat flux, resulting in a greater release of water vapor into the atmosphere. Both effects are captured by dynamic TenStream, demonstrating its ability to represent 3D radiative effects on cloud development at a substantially lower computational cost.

1 Introduction

As numerical weather prediction (NWP) models move toward higher horizontal resolutions, inter-column three-dimensional (3D) radiative effects become increasingly important. At cloud-resolving scales, cloud shadows, for instance, are no longer confined to the vertical model columns of their respective clouds, but can extend into several neighboring ones. Additionally, radiation scattered from cloud sides can enhance the diffuse downward radiation in adjacent cloud-free columns (Hogan and Shonk, 2013), whereas radiation entering through gaps between clouds can become trapped between them and the surface, thereby increasing the diffuse downward radiation below clouds as well (Hogan et al., 2019). In the thermal spectral range, higher-resolution models feature more pronounced cloud-side cooling. Together, all of these 3D radiative effects influence the spatial distribution of sources and sinks of radiative energy in the atmosphere, which in turn drive the atmospheric circulation and the weather. Accurately capturing these effects is therefore crucial for predicting future atmospheric states.



25 However, the calculation of 3D radiative effects is computationally expensive, which has largely prevented their representation in operational forecasting. Instead, state-of-the-art NWP models still rely on one-dimensional (1D) independent column approximations, such as the Monte Carlo independent column approximation (McICA; Pincus et al., 2003) currently employed at both the Deutscher Wetterdienst (DWD) and the European Centre for Medium-Range Weather Forecasts (ECMWF) (DWD, 2021; Hogan and Bozzo, 2018). These models assume that radiative transport only takes place in the vertical and neglect any
30 horizontal energy exchange. While this approximation generally performs well in horizontally homogeneous situations such as clear skies or completely overcast conditions, it breaks down as soon as numerous individual clouds shape the sky. Broken cloud fields, however, are common. Measurements conducted in Cabauw, the Netherlands, for instance, show that conditions that are neither clear sky nor overcast accounted for roughly half of the daylight time over a 10-year period, indicating that situations where 3D radiative effects can be important arise quite frequently (Mol et al., 2023). In such broken cloud fields,
35 3D radiative transfer has been shown to strongly influence both the organization and evolution of clouds. Klinger et al. (2017) found that incorporating 3D radiative transfer in the thermal spectral range results in systematically larger cooling and stronger organizational effects compared to simulations performed with 1D radiative transfer approximations. Similarly, Jakub and Mayer (2017) demonstrated that inter-column 3D radiative transfer in the solar spectral range can foster the formation of cloud streets that are not observed to the same extent in 1D simulations. Furthermore, focusing on cloud characteristics, Veerman et al. (2020, 2022) and Tjihuis et al. (2024) showed that clouds in simulations coupled to 3D radiative transfer become thicker, grow larger in horizontal extent, and exhibit higher mean liquid water paths than their 1D-driven counterparts. Nonetheless, despite this demonstrated impact on cloud development, the implications of 3D radiative transfer on weather forecasts remain largely unexplored, primarily due to its high computational cost.

To address this issue, in recent years, considerable effort has been put into making 3D radiative transfer models computationally more feasible. Many such models simplify 3D radiative transfer calculations by considering only a discrete number
45 of angles (e.g., Lovejoy et al., 1990; Gabriel et al., 1990; Davis et al., 1990). More recently, the TenStream solver (Jakub and Mayer, 2015) built upon this idea. It is capable of calculating 3D radiative fluxes and heating rates in both the solar and thermal spectral ranges by extending the 1D two-stream formulation to ten streams, thereby allowing for horizontal transport of energy. In contrast to these angular discretization methods, the neighboring column approximation (NCA; Klinger and Mayer, 2016, 2020) offers a fast analytical method for computing 3D heating rates in the thermal spectral range. To do so, it estimates cloud-side effects by taking only the immediate neighbors of a specific grid box into account. Apart from these two approaches, significant progress has also been made in accelerating highly accurate 3D Monte Carlo solvers for the use in large-eddy simulation (LES) models, with Veerman et al. (2022), for example, speeding up the method by utilizing graphics processing units (GPUs). This allowed them to perform LES runs driven by a full Monte Carlo solver for the first time ever.

55 Despite these advances, 3D solvers remain too slow to be used operationally. For instance, the GPU-accelerated Monte Carlo solver of Veerman et al. (2022) is still at least 6.4 times slower than the two-stream model it was compared to — and that only when using just 32 photons per spectral band and model column. At this low photon count, results remain notably noisy, both in irradiances and heating rates. Achieving substantially more accurate results, with root-mean-square errors of 6.88 W m^{-2} in irradiances and 0.17 K d^{-1} in heating rates, requires increasing the photon count eightfold to 256 photons per spectral band



60 and model column, which raises the computational cost to 18.5 times that of the two-stream model. This high computational burden continues to prevent the use of 3D solvers in operational weather forecasting, especially given that radiation is already called far less frequently than the dynamical core in NWP models.

To overcome these limitations, Maier et al. (2024) took first steps toward a new "dynamic" 3D radiative transfer model: the dynamic TenStream solver. Currently designed for subkilometer-scale horizontal resolutions, where model grid boxes can be treated as homogeneous, it builds upon the original TenStream solver while introducing a novel approach that is speeding up inter-column radiative transfer by treating radiation more like model dynamics. Specifically, inspired by how NWP models integrate their primitive equations over time, the solver does not recalculate radiation from scratch each time it is called but instead updates the radiative field based on the result from the previous radiation time step. Starting from this updated field, it then performs only the first few steps of an iterative scheme toward the new solution, thereby limiting interactions to neighboring grid boxes. This approach not only enables a significant acceleration of radiative transfer calculations, but also facilitates the otherwise tedious parallelization of 3D radiative transfer methods, which typically require information from all parts of the domain. When coupled to a precalculated shallow cumulus cloud field, Maier et al. (2024) showed that this new solver is able to compute 3D heating rates and net surface irradiances at a noticeably faster speed than other 3D solvers, while also providing a significant improvement in terms of accuracy over currently employed 1D schemes.

75 Based on this work, the primary objective of this paper is to investigate how coupling the dynamic TenStream solver to model dynamics affects cloud development, and whether the new solver produces clouds that more closely resemble those driven by full 3D radiation than those coupled to a traditional 1D solver. In addition, this work aims to explore the mechanisms responsible for these differences. Because the current version of the dynamic TenStream solver assumes homogeneous grid boxes, we employed an LES model for this study. The setup of this model is described in Sect. 2 of this paper. Sect. 3 then discusses the differences between the simulations driven by 1D and 3D radiative transfer, whether the dynamic TenStream solver is able to capture them, and the mechanisms that may explain these differences. The paper concludes with a summary and outlook, given in Sect. 4.

2 Methodology

2.1 Simulation setup

85 All simulations discussed in this paper were performed with PALM (Raasch and Schröter, 2001; Maronga et al., 2015, 2020). Similar to the pre-calculated cloud fields used in Maier et al. (2024), the aim was to perform simulations in which shallow cumulus clouds develop over the course of the day. The purpose of this subsection is to provide an overview of how PALM was set up in order to generate these shallow cumulus cloud fields, as well as which simulations were ultimately performed.

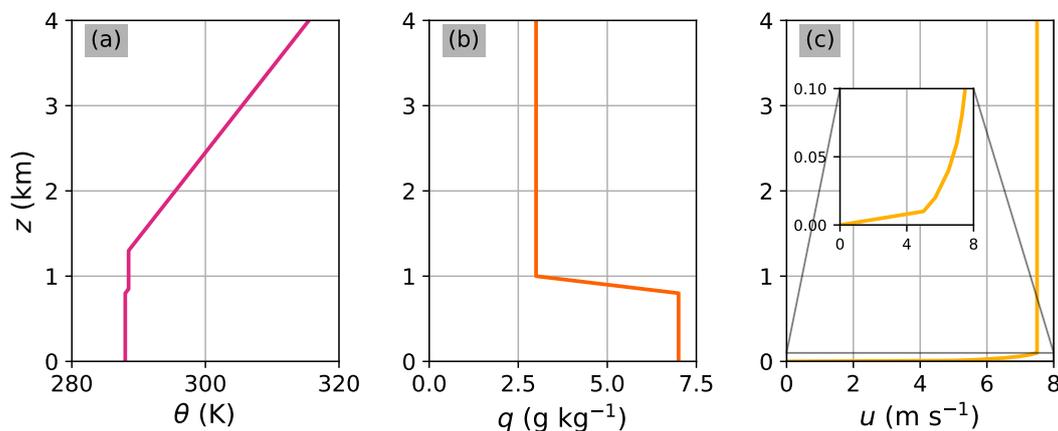


Figure 1. Initial profiles of the PALM simulations for potential temperature (a), total water mixing ratio (b), and zonal wind speed (c). The smaller graph in panel (c) shows a zoomed-in detail of the zonal wind speed profile between the surface and 100 m height.

2.1.1 Domain size

90 All simulations were performed with the same 100 m horizontal grid spacing that was used in Maier et al. (2024), but with
 256 × 256 grid boxes, resulting in a 25.6 × 25.6 km² large domain — that is sixteen times the size used in the aforementioned
 study, allowing for improved statistics when assessing cloud characteristics dependent on the radiative transfer model used.
 In the vertical, 80 grid boxes with a spacing of 50 m were used, resulting in a domain that extends up to 4 km height. This
 comparatively low vertical resolution was chosen in order to reduce the storage demand of the simulation results while still
 95 ensuring a z–x aspect ratio below 1, just like in a typical NWP model.

2.1.2 Model initialization

The simulations 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. They were initially started on 14 June 2023 at 22:00 UTC, which
 corresponds to exactly midnight local time (00:00 CEST on 15 June 2023).

100 In addition to location and time, PALM required initial vertical profiles for potential temperature, total water mixing ratio,
 and zonal and meridional wind speed, which are shown in Fig. 1. They were chosen to suppress cloud formation prior to sunrise
 while allowing shallow cumulus clouds to develop and reach a limited vertical extent after surface heating begins. A purely
 westerly wind profile was prescribed, and the Coriolis force was disabled to maintain a stationary mean flow.

The land–surface model of PALM was initialized with a flat, short grassland surface. Other than that, its setup was largely
 105 based on an example in the PALM documentation (PALM, 2026a), with only the soil temperature profile being modified to
 match the rest of the model initialization. Specifically, the soil temperature at the surface was set to 288 K, consistent with the
 potential temperature, and gradually decreased with depth to 280 K in the deepest soil layer. Additionally, the soil moisture



was set to $0.18 \text{ m}^3 \text{ m}^{-3}$ in all layers, which is somewhere between the wilting point ($0.13 \text{ m}^3 \text{ m}^{-3}$) and the saturation moisture ($0.43 \text{ m}^3 \text{ m}^{-3}$) of the medium-fine soil type used (PALM, 2026b).

110 For further details and to ensure that the simulations performed in this paper can be easily reproduced, you may refer to the parameter file that was used for all the model runs, which is provided in Maier et al. (2026).

2.1.3 Radiative transfer solvers

The main goal in terms of the setup was to perform different PALM simulations that vary only in the radiative transfer model used. To achieve this, use was made of the fact that the TenStream framework had been coupled to PALM in the past. Similar to
115 the libRadtran library employed in Maier et al. (2024), this framework allows for the application of different radiative transfer solvers using an otherwise identical environment, and all the main features of the dynamic TenStream solver presented in Maier et al. (2024) have been implemented into it. In addition to that, the TenStream framework is fully parallelized. Hence, using this framework, and similar to the evaluation in Maier et al. (2024), the following radiative transfer solvers were applied:

1. A 1D δ -Eddington approximation

120 This 1D approximation represents the type of radiative transfer scheme used in most models today and serves as a worst-case benchmark for evaluating the dynamic TenStream solver.

2. The dynamic TenStream solver

The dynamic TenStream solver introduced in Maier et al. (2024) is the main focus of this evaluation. For the simulations presented here, it was configured to perform a full TenStream solve on its first use. After that, it was operated with a
125 minimum of just two Gauß–Seidel iterations for diffuse radiation and one for direct solar radiation each time it was called. Since this implementation of the dynamic TenStream solver is fully parallelized, the iterations are performed independently on each subdomain, with communication between the cores occurring only once at the end of each radiation scheme call. The motivation behind this setup mirrors that in Maier et al. (2024): by applying a full 3D solve in the beginning, one can assess whether subsequent incomplete solves lead to different results than simulations using
130 full TenStream solves throughout. Furthermore, as explained in Maier et al. (2024), using two iterations instead of one ensures that the iteration direction through the underlying system of linear equations is altered at least once per call.

3. The original TenStream solver

Simulations performed with the original TenStream solver, also referred to as *TenStream reference solver* from here on, serve as benchmark simulations. Since the dynamic TenStream solver is based on this original TenStream model,
135 reproducing its results despite using incomplete solves represents the best possible outcome.

The exact `tenstream.options` files used for the simulations with each of these solvers can be found in Maier et al. (2026). Regardless of the radiative transfer solver applied, the atmospheric trace gas concentrations were provided by the 1976 US standard atmosphere (Anderson et al., 1986). Furthermore, all PALM simulations were carried out with a radiation time step of 30 s, as this proved to be a good speed-accuracy trade-off for the radiative transfer calculations in Maier et al. (2024),
140 which were performed with the same horizontal grid spacing of 100 m.

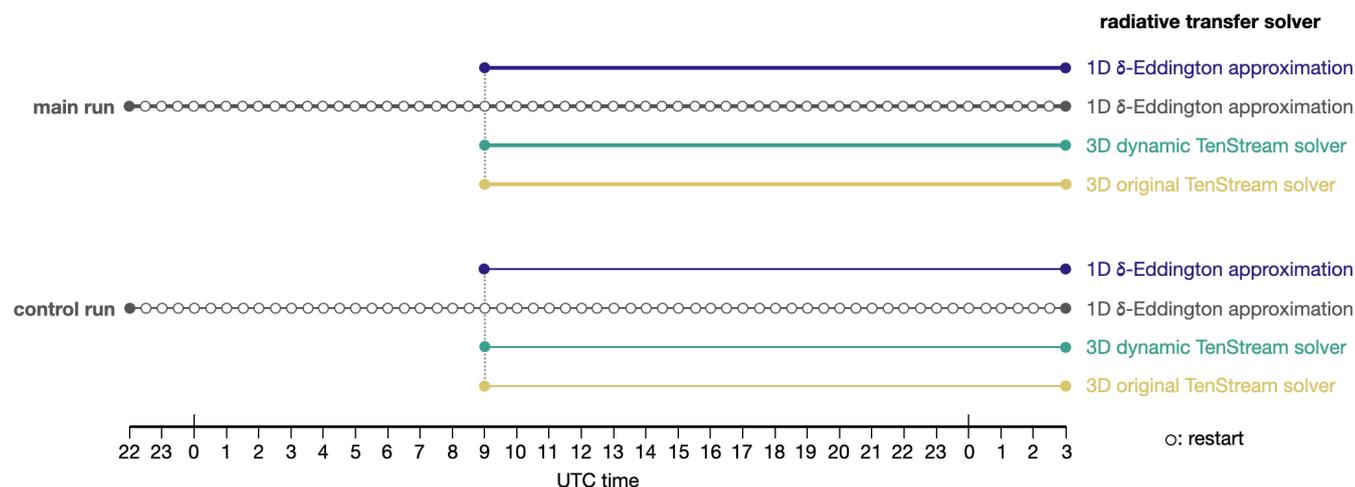


Figure 2. Schematic illustration of all PALM simulations discussed in this paper. On 14 June 2023 at 22:00 UTC, two simulations, a main run and a statistically independent but otherwise identical control run (both shown in gray here), were started. These two initial model runs were driven by 1D radiation and restarted every 30 minutes, as visualized by the open circles in their respective timelines. Starting at 09:00 UTC on 15 June 2023, three additional restart runs were launched from each of the two initial runs. All these in total six restart runs were calculated in one go. They only differed in the radiative transfer model applied: the purple runs were driven by a 1D δ -Eddington approximation, the turquoise ones by the dynamic TenStream solver, and the ochre ones by the original TenStream solver. Filled circles, in contrast to the open circles, visualize the start and end points of each PALM simulation.

2.1.4 Overview of the PALM simulation setup

Using these radiative transfer solvers, we can now discuss the different PALM simulations that were performed with them. A schematic overview of the setup is provided in Fig. 2. It shows that initially, two different simulations were conducted: a main run and a statistically independent but otherwise identical control run. Both of these simulations were started on 14 June 2023 at 22:00 UTC using the model initialization described earlier. The only difference between the two simulations is that, at certain times (every 150 s in our runs), different random seeds of perturbations were applied to the horizontal wind fields until a predefined perturbation kinetic energy limit was reached. This limit was set to the default value of $0.01 \text{ m}^2 \text{ s}^{-2}$. The resulting slightly different wind fields in the main and control run represent the uncertainties in the initial conditions of the setup and effectively create a two-member ensemble, enabling an estimation of whether certain features in the main run are robust or fall within the simulation's inherent uncertainty.

From these slightly differently perturbed initial states, both simulations then ran for 29 hours, ending on 16 June 2023 at 03:00 UTC. During this time, the simulations were first given enough time to spin up and adjust their initial state until the Sun rose on 15 June 2023 at 03:13 UTC. After sunrise, both model runs then encompassed almost a full diurnal cycle, including sunset at 19:15 UTC and ending shortly before the next sunrise on 16 June 2023 at 03:12 UTC. Throughout this entire time, both initial runs were driven by 1D radiation and restarted every 30 minutes. At each of these restart points, visualized by the



open circles in Fig. 2, PALM stopped the simulation and saved its current state, before restarting it from exactly this saved atmospheric state.

Now, remember that our objective was to conduct different PALM simulations that vary only in the radiative transfer model used. The restart mechanism allows such simulations to be initiated from any intermediate time step shown in Fig. 2 by
160 restarting the model run from the corresponding saved atmospheric state with modified runtime parameters, i.e., different radiative transfer solvers in this case. For this evaluation, the simulations were restarted from 09:00 UTC, i.e., 11:00 a.m. local time. At this point in time, shallow cumulus clouds have just started to form in the domain, making it a suitable moment for investigating how the subsequent development of these clouds and the surrounding atmosphere differs depending on the radiative transfer solver used. Consequently, starting from 09:00 UTC, three different restart runs were performed for both
165 the main and the control run, each of them coupled to one of the three different radiative transfer solvers introduced above: either the 1D δ -Eddington approximation (shown in purple in Fig. 2), the dynamic TenStream solver (turquoise), or the original TenStream solver (ochre). These restart runs form the main foundation for the evaluation in this paper. Performing them for both the main and the control run was important, as it will allow us to estimate whether differences between runs driven by different radiative transfer solvers are statistically significant or not. Apart from that, all restart runs were conducted in one go,
170 i.e., without further restarts. They ran exactly as long as the two initial runs, i.e., until 16 June 2023 at 03:00 UTC, that is 05:00 a.m. local time, covering approximately ten hours of daytime and eight hours of nighttime, thereby providing a robust dataset for analyzing both of these regimes.

In terms of the general runtime configuration, all simulations were performed with a model time step of 5 s. Moreover, they were all executed on 64 CPU cores, using an 8×8 grid of subdomains, each containing 32×32 vertical columns. This domain
175 decomposition is particularly important for the dynamic TenStream solver, since its incomplete solves are performed for each of the subdomains in parallel, with interaction between subdomains occurring only once at the end of each radiation scheme call. Information can thus only spread within individual subdomains during the radiative transfer calculations, and propagation to neighboring subdomains is delayed to subsequent calls of the scheme.

Regarding the cloud model, the built-in "morrison" scheme was applied, which uses two-moment cloud microphysics according to Seifert and Beheng (2005), Khairoutdinov and Kogan (2000), Khvorostyanov and Curry (2006) and Morrison and Grabowski (2007). Cloud water sedimentation was enabled, while aerosol concentration and cloud droplet number density
180 were prescribed using the default values of PALM. Lastly, the data output for all the model runs was written at a temporal resolution of 60 s. This is considerably coarser than the 10 s resolution used in Maier et al. (2024), but a necessary reduction in terms of the overall storage size of the simulations. Although this means that output is only available for every second radiation
185 time step, Sect. 3.2 will show that the averaged cloud characteristics primarily analyzed in this paper vary on much longer temporal scales, making this reduced output frequency also a scientifically acceptable compromise.

2.2 Evaluation methods

The entire evaluation in this paper is focused on the six restart runs that we just introduced, i.e., the simulations shown in color in Fig. 2, which, started from either the main or control run of the setup, are driven by the 1D δ -Eddington approximation, the



190 dynamic TenStream solver, or the original TenStream solver. For simplicity, they will be referred to as " δ -Eddington (control)",
"dynamic TenStream (control)" and "TenStream reference (control)" runs from here on. Using these simulations, our primary
objective is to investigate whether, starting from the same atmospheric state, the atmosphere and its clouds develop differently
depending on whether 1D or 3D radiation is applied. Additionally, and even more important for this work, our aim is to assess
whether the dynamic TenStream solver is able to reproduce these potential 3D-related differences.

195 2.2.1 Accuracy evaluation

One way to investigate the aforementioned differences is by comparing the temporal evolution of selected model quantities
across simulations driven by different radiative transfer solvers. Since the solvers fully interact with the model dynamics in
these simulations, this evaluation is not just limited to radiative quantities but can also include other variables, such as the cloud
water mixing ratio. For any such quantity ξ and at any point in time, the accuracy of a simulation relative to the TenStream
200 reference run can be quantified using the mean bias error (MBE), which is given by

$$\text{MBE} = \langle \xi \rangle - \langle \xi_{\text{ref}} \rangle = \langle \xi - \xi_{\text{ref}} \rangle, \quad (1)$$

where $\langle \dots \rangle$ denotes an arithmetic average over all model grid boxes. Point-based error measures, such as the mean absolute error
or root-mean-square error, however, are not utilized in this evaluation. This is because simulations coupled to different radiative
transfer solvers lead to diverging atmospheric states, particularly in terms of clouds. Hence, even if two simulations featured
205 clouds with nearly identical characteristics, slight positional differences could result in substantial errors when using point-
based error metrics. Consequently, since this evaluation focuses on differences in cloud properties rather than on differences
in the clouds' exact spatial placement, such error measures have not been applied. The MBE, by contrast, is based on domain-
averaged values and is therefore unaffected by these double-penalty errors.

At this point, it should also be noted that MBE values must be interpreted relative to the intrinsic variability of the simu-
210 lations, which is estimated from the differences between the respective main and control runs. A simulation is therefore only
considered to perform worse than the TenStream reference run if its MBE clearly exceeds the internal variability of the refer-
ence simulation. Likewise, differences between simulations driven by different radiative transfer solvers are only regarded as
meaningful if the differences between their MBEs exceed the corresponding main–control run differences.

2.2.2 Quantification of cloud characteristics

215 Apart from the temporal evolution of the MBE in certain model quantities, a particular interest lies in how the clouds in the
simulations develop depending on the radiative transfer solver used. Therefore, inspired by the work of Tijhuis et al. (2024),
three quantities are used to characterize the cloud fields at any point in time: cloud cover, average liquid water path (LWP) in
cloudy columns, and average cloud depth. A grid box is defined as cloudy if its cloud water mixing ratio q_c exceeds 0.01 g kg^{-1} .
The application of this threshold is necessary because many grid boxes in the simulations are subject to very small q_c values,
220 often as low as $10^{-35} \text{ g kg}^{-1}$. While physically negligible, the large number of these very small values strongly influences



quantities such as cloud cover, so they are excluded from the analysis by applying the aforementioned threshold. The three cloud characteristic measures can then be defined as follows:

1. Cloud cover

At any point in time, the cloud cover is given by the fraction of vertical columns in the domain that contain at least one grid box with $q_c > 0.01 \text{ g kg}^{-1}$.

2. Average liquid water path in cloudy columns

This quantity is defined as the mean LWP of all vertical columns that contain at least one grid box with $q_c > 0.01 \text{ g kg}^{-1}$.

3. Average cloud depth

For every cloudy column (i, j) , the index of the highest ($k_{\max, i, j}$) and lowest ($k_{\min, i, j}$) grid box with $q_c > 0.01 \text{ g kg}^{-1}$ is determined. The cloud depth in this column is then given by

$$\Delta z_{\text{cld}, i, j} = (k_{\max, i, j} + 1 - k_{\min, i, j}) \Delta z, \quad (2)$$

where Δz is the vertical grid spacing. The average cloud depth is subsequently calculated as the mean of $\Delta z_{\text{cld}, i, j}$ across all cloudy columns. Note that this method assumes only one cloudy layer per column, as cloud depths are calculated from the vertical extent between the highest and lowest cloudy model levels in each column, ignoring any cloud-free layers in between. This assumption, however, is valid for the shallow cumulus clouds used in this evaluation.

At this point, it is important to note that all these quantities describe general characteristics of cloud fields rather than properties of individual clouds. This is precisely the goal, as the objective of this paper is to identify systematic differences between cloud fields driven by 1D and 3D radiation. Moreover, unlike the work of Tjihuis et al. (2024), our evaluation also accounts for radiative transfer in the thermal spectral range, enabling the investigation of these differences during nighttime as well.

3 Discussion of the results

3.1 Simulation overview

Before getting into details, let us first take a look at the general evolution of the clouds in the different simulations. To this end, Fig. 3 shows five snapshots of the temporal evolution of the LWP in the simulations driven by the 1D δ -Eddington approximation, the dynamic TenStream solver, and the original TenStream model. For this overview, only simulations started from the main run of the setup are shown. Plots of the LWP are used because this quantity provides a good measure of both cloud position and thickness.

Starting with the first row of panels, we can see that initially, all simulations were indeed started from the same cloud field, allowing differences in their subsequent development to be observed depending on the radiative transfer solver used. Results are shown for 09:01 UTC instead of 09:00 UTC for this first time step, as model output is not immediately available after the simulations start at 09:00 UTC. Three hours later, at 12:00 UTC, the clouds coupled to the different radiative transfer solvers

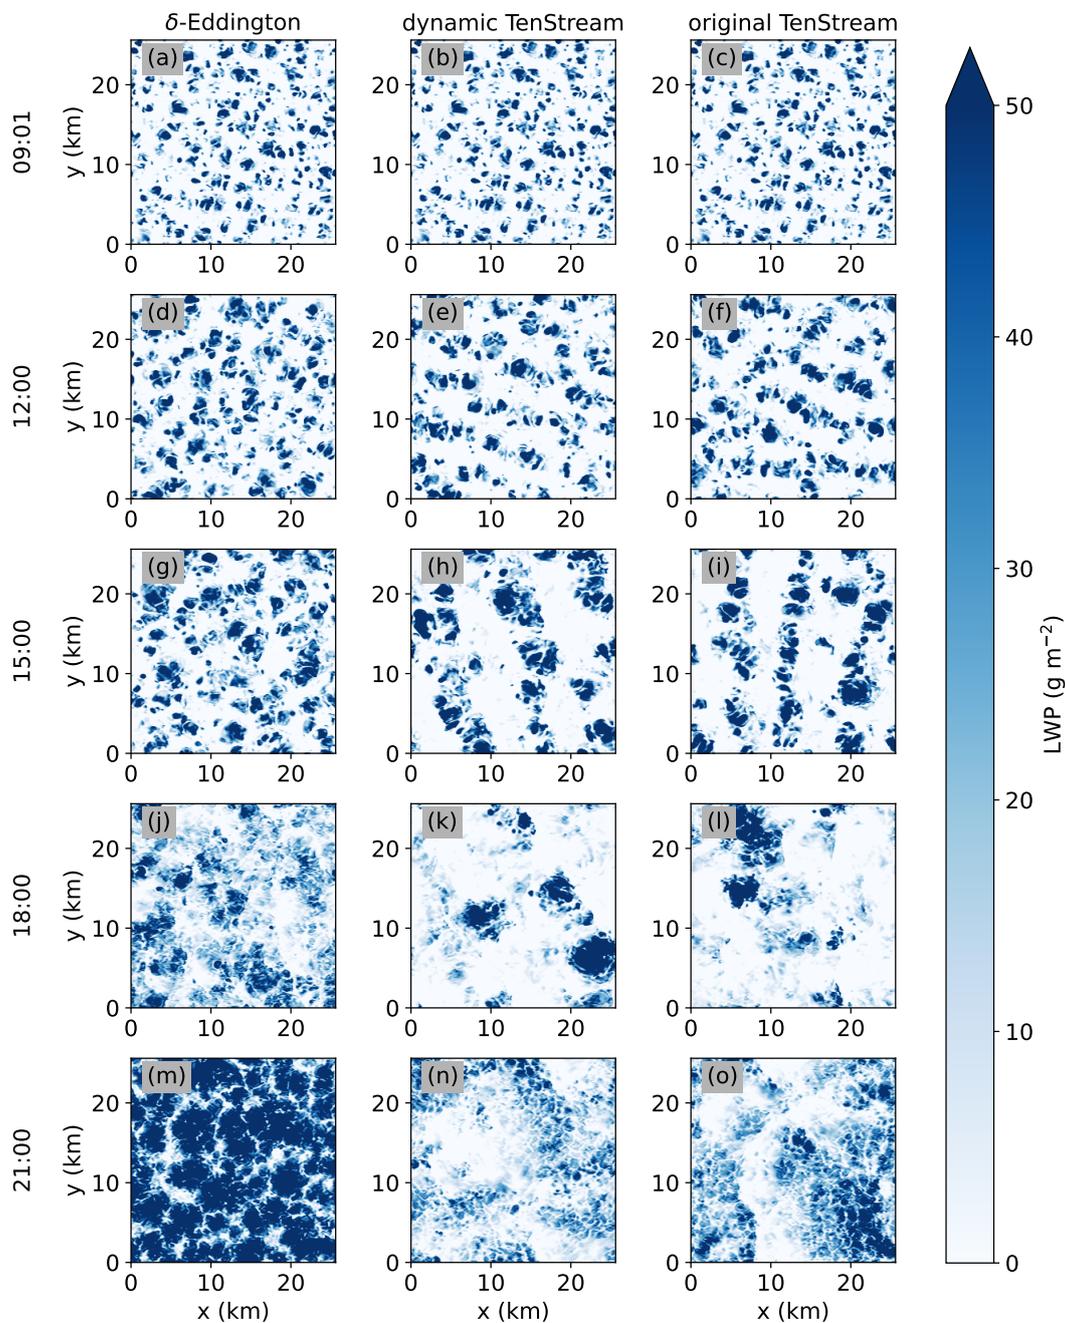


Figure 3. Temporal evolution of the liquid water path (LWP) in the PALM simulations driven by the 1D δ -Eddington approximation (left), the dynamic TenStream solver (middle), and the original TenStream model (right), shown for five time steps between 09:01 and 21:00 UTC. To enhance the contrast of the plots, the maximum LWP value for the colorbar was set to 50 g m^{-2} instead of the global maximum value of 363 g m^{-2} .



have already developed differently. While they all increased in size compared to the plots at 09:01 UTC, we can clearly see that their organization differs between the simulation driven by 1D radiation shown in panel (d) and the one driven by the original TenStream model shown in panel (f). The clouds in panel (d) are still pretty unorganized. In panel (f), however, we can observe the build-up of cloud streets, just as they were proposed in Jakub and Mayer (2017). These cloud streets are supposed to form perpendicular to the angle of solar incidence and parallel to the mean wind flow. The latter is a constant westerly flow, whereas the Sun is positioned at an azimuth angle of 205° at 12:00 UTC, i.e., in the south–south–west. And indeed, the clouds are oriented perpendicular to that direction, forming streets with an east–south–easterly orientation, which is particularly visible in the lower right part of panel (f). The cloud streets are also visible in the simulation coupled to the dynamic TenStream solver shown in panel (e), providing the first proof of a 3D-related effect that is captured by the dynamic treatment of radiation in that solver.

Moving on to panels (g)–(i), another three hours later, the cloud streets in the simulations coupled to 3D radiative transfer have become markedly more pronounced. Because the Sun moved to an azimuth angle of 262° in the meantime, which is almost exactly in the west, the cloud streets are oriented from north to south now, still perpendicular to the angle of solar incidence. In contrast to that, clouds in the simulation driven by 1D radiation are still more or less randomly positioned. 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. Apart from these organizational aspects, we can also see that individual clouds in both panels (h) and (i) are noticeably larger and thicker than in panel (g), although the increased thickness is not as apparent in the plots because the maximum value for the colorbar was set to a value well below the global LWP maximum to enhance contrast. Apart from that, however, this development of larger and thicker clouds during daytime in simulations driven by 3D radiative transfer is consistent with findings in other studies (Veerman et al., 2020; Tijhuis et al., 2024).

As the day proceeds, the organization of clouds in the simulations driven by 3D radiative transfer starts to break down and is no longer visible in plots (k) and (l) at 18:00 UTC. However, this is also only about an hour before sunset at 19:15 UTC, with the Sun already at a relatively low elevation of 10° . It is also around this point in time that the cloud cover starts to noticeably increase across all simulations. Experiments with the model initialization showed that this behavior is strongly influenced by the initial soil moisture content assigned to the simulations. When reducing its value from $0.18 \text{ m}^3 \text{ m}^{-3}$ to $0.16 \text{ m}^3 \text{ m}^{-3}$, which is still well above the wilting point of $0.13 \text{ m}^3 \text{ m}^{-3}$, clouds start to dissolve around sunset rather than fully covering the domain. While this dissolution in the evening may be more realistic for shallow cumulus clouds on a typical summer day, the presence of clouds at nighttime allows differences between simulations driven by 1D and 3D radiative transfer to be examined after sunset as well, when only radiative transfer in the thermal spectral range plays a role. And as we can see, there is a noticeable change in the characteristics of the clouds in that part of the day that could be observed across a variety of simulation setups experimented with. Namely, it is now the clouds in the simulation driven by 1D radiation that are becoming noticeably thicker than their 3D-driven counterparts, as can be seen in panels (m)–(o). Unlike during the day, only one time step is shown at nighttime, as the overall structure of the clouds remains largely unchanged thereafter: all simulations continue to be overcast with clouds and it is only their thickness that further increases throughout the night, with the simulation driven by 1D radiation



maintaining the thickest clouds. The extent to which this development originates from daytime cloud evolution, as opposed to nighttime radiative effects, is addressed in Sect. 3.2.

Apart from all of this, Fig. 3 demonstrates that, at least visually, the clouds in the dynamic TenStream run show very similar characteristics to the ones in the TenStream reference run. Unlike the simulation driven by the 1D δ -Eddington approximation, the dynamic TenStream solver is able to reproduce 3D-related features such as cloud streets and the development of larger and thicker clouds during daytime. However, if the solves produced by the two 3D solvers were identical, they would also lead to the exact same cloud fields in this setup. That said, the fact that the simulations, starting from the same cloud field, develop differently when compared point by point, shows that the small deviations between the solvers lead to differently positioned clouds when coupled to model dynamics. These slightly different positioned clouds, however, still represent the reference solution far better than the simulation coupled to the 1D δ -Eddington approximation does.

3.2 Effects of radiation on cloud characteristics

Next, let us quantify these observed differences in the cloud fields. To this end, Fig. 4 shows the temporal evolution of the cloud characteristics quantities introduced in Sect. 2.2.2 for the six PALM simulations that are considered in this evaluation. Let us discuss the results during daytime first. Panel (a) shows that in this time period, the cloud cover is at around 45% for the simulations driven by 3D radiative transfer (shown in turquoise and ochre), whereas it is noticeably higher, mostly above 50%, for the simulations driven by 1D radiation (shown in purple). The cloud cover also remains relatively constant over time in the 3D simulations, whereas it steadily increases in the 1D cases. 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; Tjihuis et al., 2024). However, the overall cloud cover in these studies was usually much lower. And for values above approximately 30%, Tjihuis 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).

Apart from that, panels (b) and (c) show that the clouds in simulations coupled to 3D radiative transfer also become thicker and feature a higher average liquid water path during daytime than their 1D-driven counterparts, similar to how it was shown in other studies (Veerman et al., 2020, 2022; Tjihuis et al., 2024). Both effects are not very large, but still bigger than the differences between the respective main and control runs, and thus significant. To put it in numbers, the average LWP in cloudy columns during daytime is 25.5 g m^{-2} in the δ -Eddington main run, whereas it is 28.5 g m^{-2} in the dynamic TenStream main run and 29.3 g m^{-2} in the TenStream reference main run. In terms of depth, these numbers are given by 169.6 m, 178.5 m and 180.8 m, respectively. Hence, the clouds in the simulations coupled to 3D radiation are about 5–7% thicker and feature an about 12–15% larger LWP during the day, with similar numbers for the control run. As we can see in Fig. 5, this increase in LWP is also not only caused by the increased depth of the clouds, but also by a higher overall cloud water mixing ratio, i.e., more liquid water per kilogram of air in the clouds coupled to 3D radiation.

Switching to nighttime, panel (a) of Fig. 4 shows that all simulations converge toward a completely overcast sky after sunset. However, the corresponding increase in cloud cover occurs significantly earlier in the 1D simulations. Alongside this, the cloud characteristics change substantially. Now, it is the clouds coupled to 1D radiation that become thicker and contain more

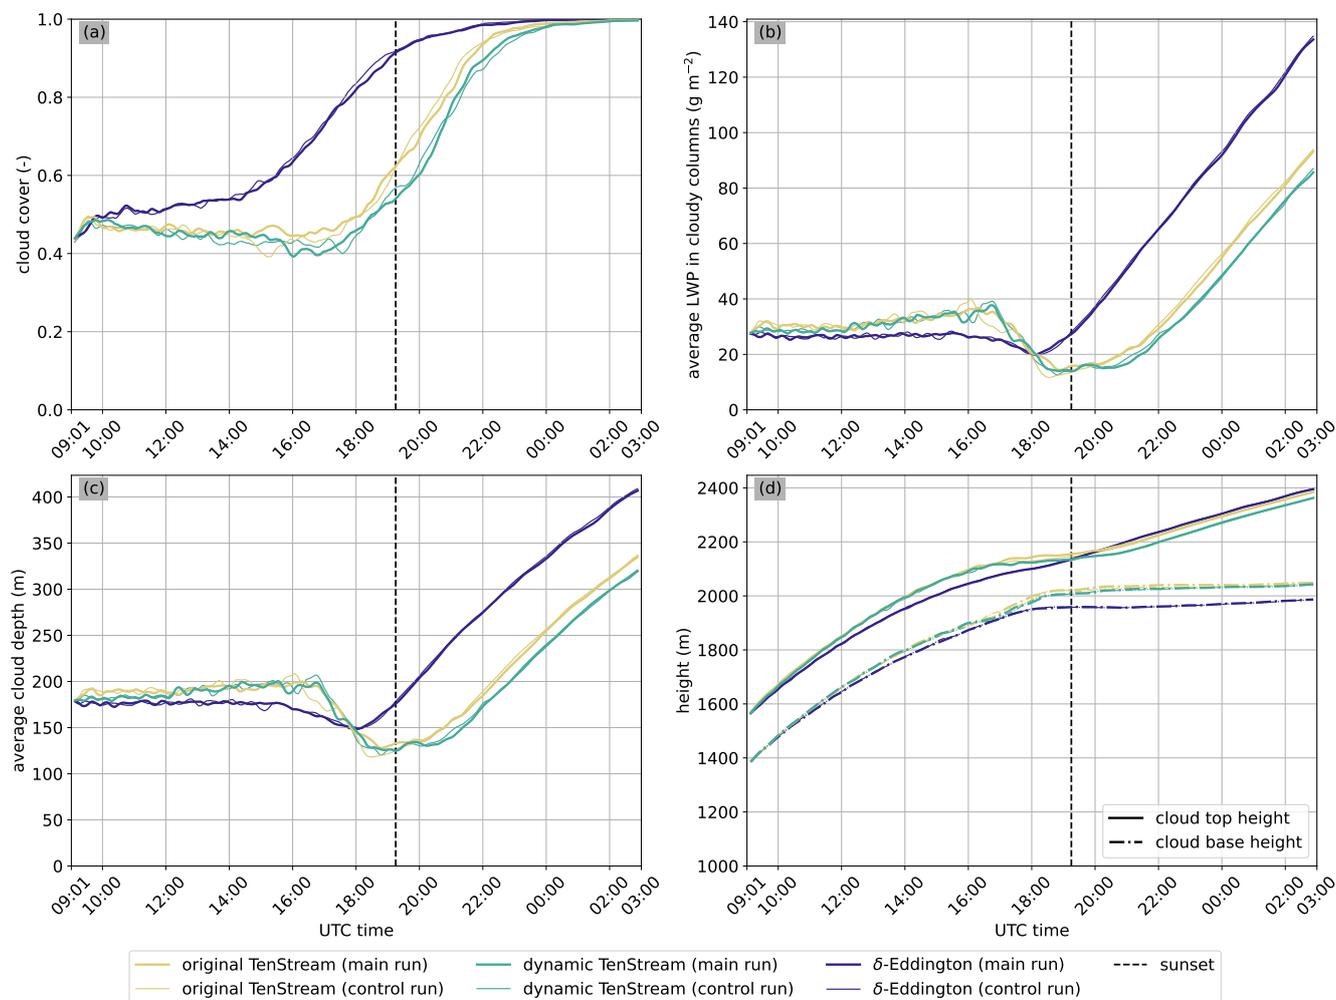


Figure 4. Temporal evolution of cloud cover (a), average liquid water path in cloudy columns (b), average cloud depth (c), and cloud top and base heights (d) for the six restart runs introduced in Sect. 2.1.4. The vertical line in all the plots indicates the sunset time. To improve readability, a 15-minute running mean was applied to the data.

320 liquid water. As shown in panel (d) of Fig. 4, this increased thickness is primarily due to the cloud base height stabilizing
 earlier in the 1D simulations, while the cloud top height develops similarly across all simulations. Together, this results in
 thicker clouds in the simulations driven by 1D radiation, although the development at night itself is not much different from the
 other simulations. Hence, the thicker clouds at night might actually be related to the higher cloud cover before sunset and the
 correspondingly reduced radiative energy at the surface, which could lead to an earlier suppression of convection. However,
 325 similar model runs with lower initial soil moisture, which did not feature a comparable increase in cloud cover near sunset, still
 lead to the development of the same nighttime cloud characteristics, i.e., the development of thicker clouds that contain more
 liquid water in simulations coupled to 1D radiative transfer, which suggests that these features are fairly robust overall.

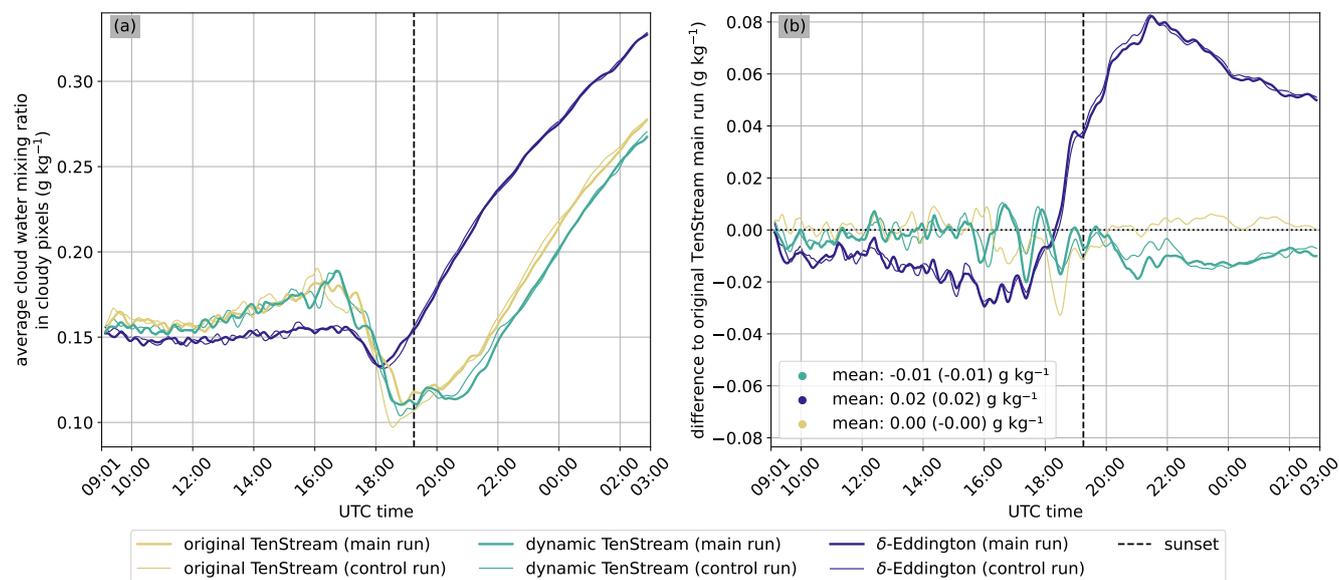


Figure 5. Temporal evolution of the average cloud water mixing ratio q_c in cloudy pixels (a) and differences relative to the TenStream reference run (b) for all six (five in panel (b)) PALM simulations considered in this evaluation. The legend in panel (b) provides temporal mean values, with the numbers for the corresponding control runs given in brackets. By isolating changes in in-cloud condensate from changes in cloud depth, this plot complements Fig. 4 (b). As in this preceding figure, the vertical line marks the time of sunset, and a 15-minute running mean was applied to the data.

At this point, the top row of Fig. 6 summarizes the results discussed so far using an alternative visualization motivated by the work of Tijhuis et al. (2024). As shown in panel (a) of this figure, we found that simulations coupled to the original
 330 TenStream solver feature a considerably lower cloud cover than those coupled to 1D radiation, although this may only apply to simulations with cloud covers above about 30%. We have also seen that in terms of liquid water path and cloud depth, we need to distinguish between day and night. During the day, simulations coupled to the original TenStream solver produce thicker clouds with more liquid water. Correspondingly, in panels (b) and (c), all data points before approximately 18:00 UTC lie to the left of the identity line. After sunset, however, these characteristics change, as clouds in the TenStream reference run become
 335 thinner and contain less liquid water than their 1D-driven counterparts. Consequently, all post-sunset data points in panels (b) and (c) shift to the right of the identity line, although this shift may not be driven by nighttime radiative effects alone but could also reflect changes in cloud characteristics shortly before sunset, as discussed above.

Up to this point, the discussion has primarily focused on differences between simulations driven by 1D and 3D radiative transfer, without addressing whether the incomplete solves in the dynamic TenStream solver can adequately reproduce these
 340 3D-related differences. To evaluate this, the bottom row of Fig. 6 compares the average cloud characteristics between the simulation driven by the dynamic TenStream solver and the original TenStream model. The plots show that the cloud characteristics agree pretty well between the two solvers, with small differences between daytime and nighttime. During the day, the

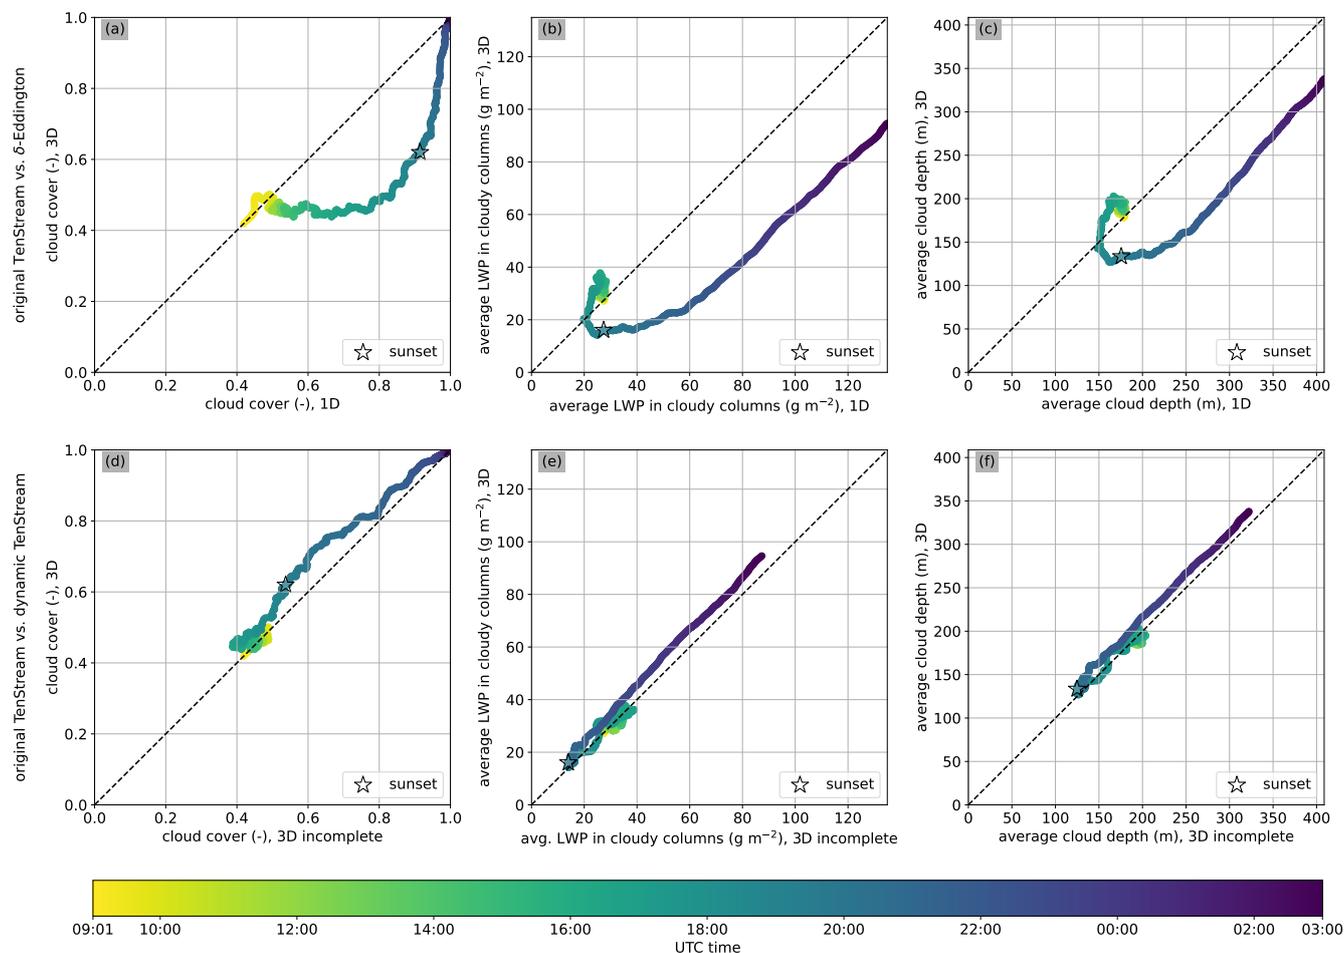


Figure 6. Comparison of average cloud characteristics between the δ -Eddington and the TenStream reference run (top row), as well as between the dynamic TenStream and the TenStream reference run (bottom row). The individual scatter plots show these comparisons for cloud cover (a, d), average LWP in cloudy columns (b, e) and average cloud depth (c, f). Every data point is color-coded with respect to time. In addition to that, in each panel, the data point at sunset is highlighted with a star. Only data from the simulations performed from the main run are shown.

clouds driven by the two solvers exhibit nearly identical characteristics, with differences falling within the variability of the simulations themselves, as one can see by going back to Fig. 4. At night, however, the incomplete solves in the dynamic TenStream solver slightly underestimate cloud cover, average liquid water path in cloudy columns, and cloud depth. Nonetheless, these discrepancies remain small compared to those observed in the δ -Eddington runs and suggest that the incomplete solves effectively capture all relevant 3D-related effects in terms of cloud characteristics.

It is important to note, however, that the performance of the dynamic TenStream solver is highly dependent on the simulation setup, and particularly on factors such as resolution, domain decomposition, and cloud cover. The relatively high cloud cover



350 in the simulations here, for example, is very beneficial for its performance, as the clouds are relatively close to each other. This proximity reduces the distance — measured in grid boxes — over which information crucial for 3D radiative transfer must be transported. When the overall cloud cover drops below 30% — which is common for shallow cumuli, with Neggers et al. (2003), for instance, reporting cloud covers between 10% and 20% on a prototypical shallow cumulus day — the performance of the solver weakens, as the increased spacing between clouds requires information to traverse more grid boxes, and potentially even multiple subdomains. At this point, performing just two Gauß–Seidel iterations per radiation scheme call may not be sufficient anymore. However, the dynamic TenStream solver was developed with NWP models in mind. At their kilometer-scale resolution, clouds are inherently closer together in terms of grid box distance, which should significantly mitigate this issue.

3.3 Differences in cloud–radiation interactions

360 So far, the analysis has primarily focused on the different development of clouds in simulations driven by 1D and 3D radiative transfer, showing that the dynamic TenStream solver successfully captures most of the 3D effects on clouds, despite relying on incomplete solves. However, the causes of these differences have not yet been investigated. Identifying them is challenging, as the clouds in the fully interactive PALM simulations are not only shaped by the corresponding radiative fields but also actively influence them, making it difficult to separate cause from effect. Nevertheless, this subsection still aims to identify at least some of the connections between the different radiative fields and the resulting cloud characteristics.

Regarding the differences in the radiative fields, Maier et al. (2024) showed that sources and sinks of radiative energy in the atmosphere — i.e., heating rates — are distributed very differently in simulations driven by 1D and 3D radiative transfer. This also affects the radiation at the ground, which we will focus on now. To this end, Fig. 7 shows five time steps of the temporal evolution of the solar net surface irradiance for the three simulations coupled to the three different radiative transfer solvers considered in this evaluation. Similar to the other results so far, we can directly see distinct differences between the fields produced by the 1D solver (left panels) and those created by the 3D solvers (middle and right panels).

To discuss these differences, let us first look at panels (d)–(f), which show the net surface irradiance fields at 11:00 UTC. Starting with panel (d), remember that in 1D radiative transfer calculations, the atmosphere is divided into independent vertical columns. If such a column contains a cloud, it casts a shadow directly beneath it. These are the purple pixels in panel (d), which perfectly match with the corresponding cloud field. The remaining columns are all subject to very similar clear-sky conditions, resulting in the pretty uniform background color in that panel. In contrast to that, the surface irradiance field of the original TenStream solver, shown in panel (f), is far more complex. Cloud shadows in this simulation are displaced according to the angle of solar incidence. At 11:00 UTC, the Sun is positioned at an azimuth angle of 173° , or roughly in the south. The cloud shadows are thus shifted slightly to the north of the clouds. Additionally, the net surface irradiance values in the shadows are lower than in panel (d), which is due to the thicker clouds that contain more liquid water in this simulation, as well as generally darker cloud shadows in simulations coupled to 3D radiation (e.g., Gristey et al., 2020). Even more striking are the bright areas at the southern edges of the shadows, where the net surface irradiance exceeds the clear-sky values in panel (d). These so-called cloud enhancements (Tijhuis et al., 2022) occur because radiation coming from the cloud sides or entrapped

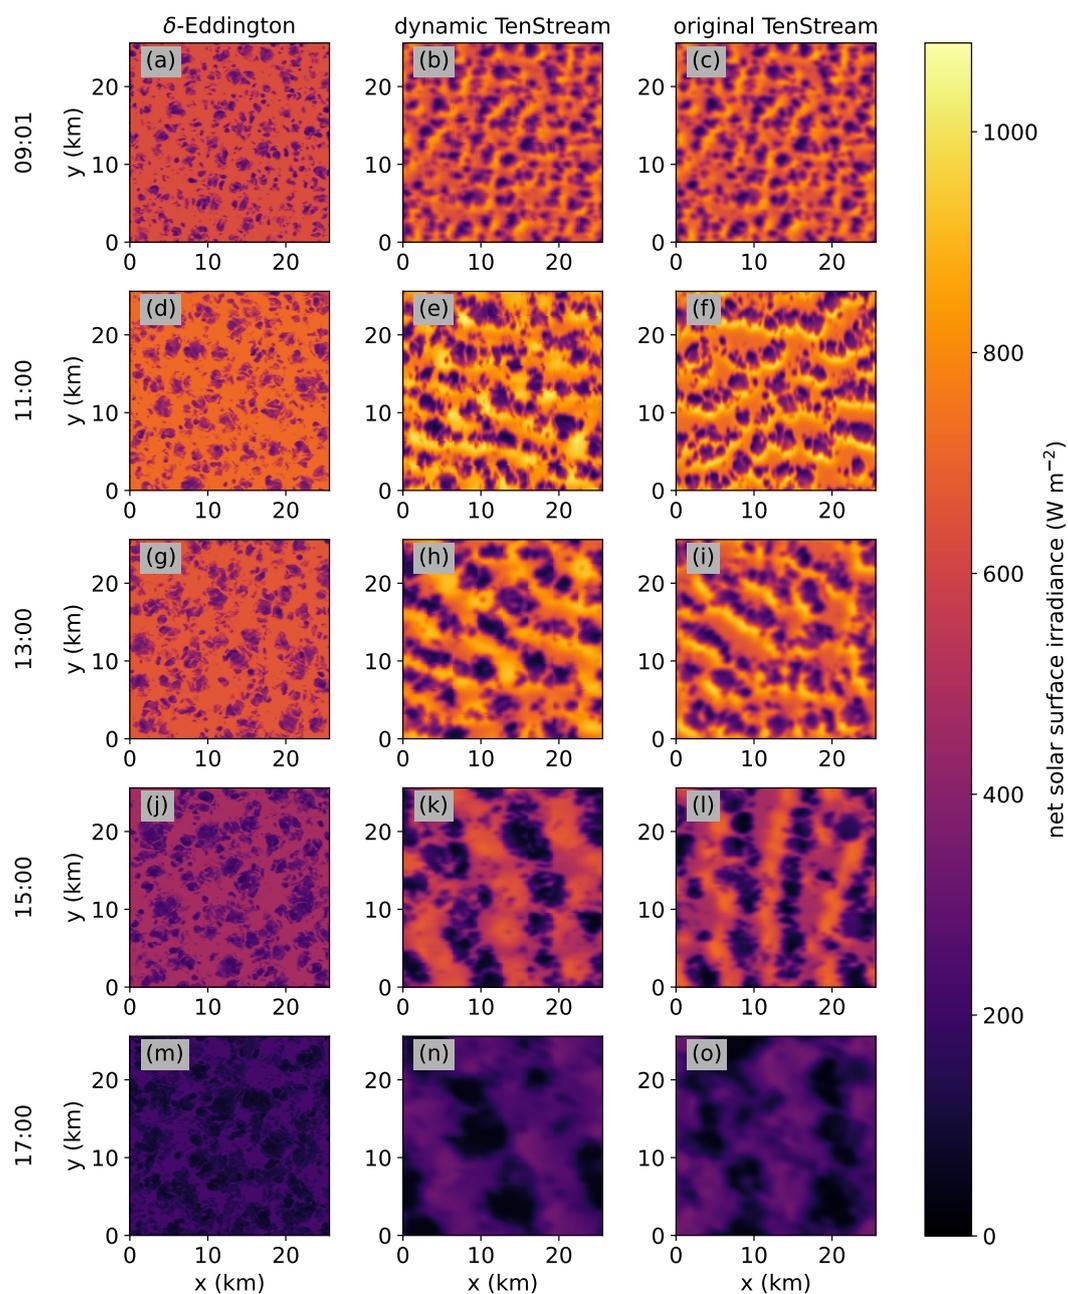


Figure 7. Temporal evolution of the net surface irradiance in the solar spectral range for the simulations driven by the δ -Eddington approximation (left), the dynamic TenStream solver (middle), and the original TenStream solver (right), shown for five time steps between 09:01 UTC and 17:00 UTC. Only simulations performed from the main run of the setup are shown.

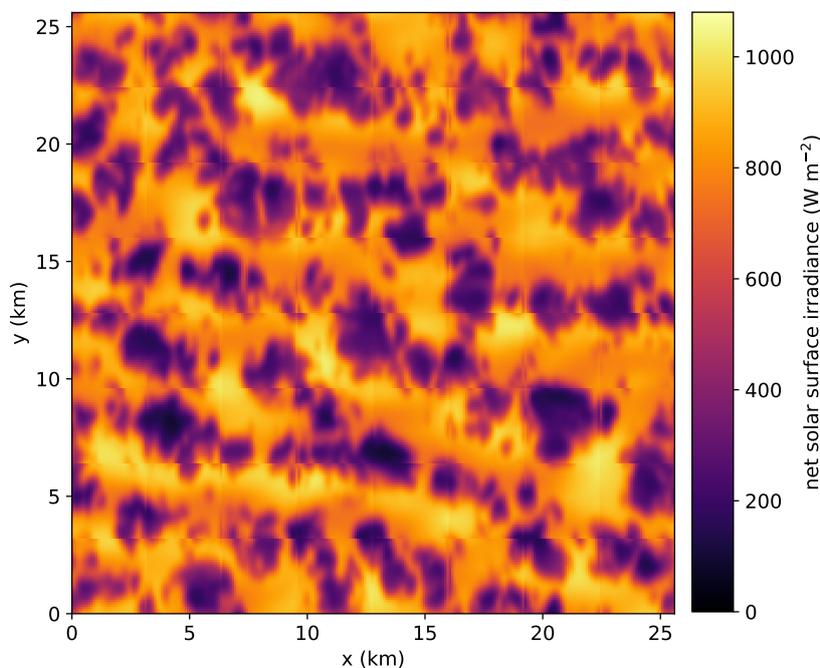


Figure 8. Magnified version of panel (e) in Fig. 7, highlighting the discontinuities between subdomains in the dynamic TenStream solver.

between the surface and the cloud bases enhances the diffuse downward radiation (Hogan and Shonk, 2013; Hogan et al.,
 385 2019). Examining the plot at different times confirms that these cloud enhancements consistently appear at the sunward edges
 of the cloud shadows. For instance, in panel (c), the Sun is positioned at an azimuth angle of 121° (i.e., in the east–south–east),
 whereas in panel (l), it is at an angle of 262° (i.e., in the west), causing the enhancements to appear to the east of the shadows
 in panel (c) and to the west of them in panel (l).

The dynamic TenStream solver in the middle panels of Fig. 7 successfully reproduces all these features, although it slightly
 390 overestimates the size of the cloud enhancements. This overestimation can be at least partially attributed to the delayed interac-
 tion between radiative fluxes at the clouds and the surface when using incomplete solves, especially when this interaction spans
 multiple subdomains. In such cases, radiative enhancements caused by a cloud may persist at a specific location on the ground
 even after the cloud has moved on with the mean wind flow, enlarging the associated features on the surface. Following this
 argumentation, these artifacts are expected to be most pronounced when the solar zenith angle is low, resulting in interactions
 395 that involve multiple subdomains, and when the clouds move away from the direction of solar incidence and thus away from
 their radiatively enhanced areas. And indeed, the overestimation is relatively small in panel (b), where the clouds are moving
 toward the Sun, but become much more pronounced in panel (k), where the Sun is at a zenith angle of 51° and the clouds move
 away from the direction of solar incidence, allowing the cloud enhancement areas to increase in size.

By providing a close-up of panel (e), Fig. 8 highlights another artifact of the incomplete solves. In this magnification of
 400 the net surface irradiance field, we can clearly identify the underlying domain decomposition into 8×8 subdomains through

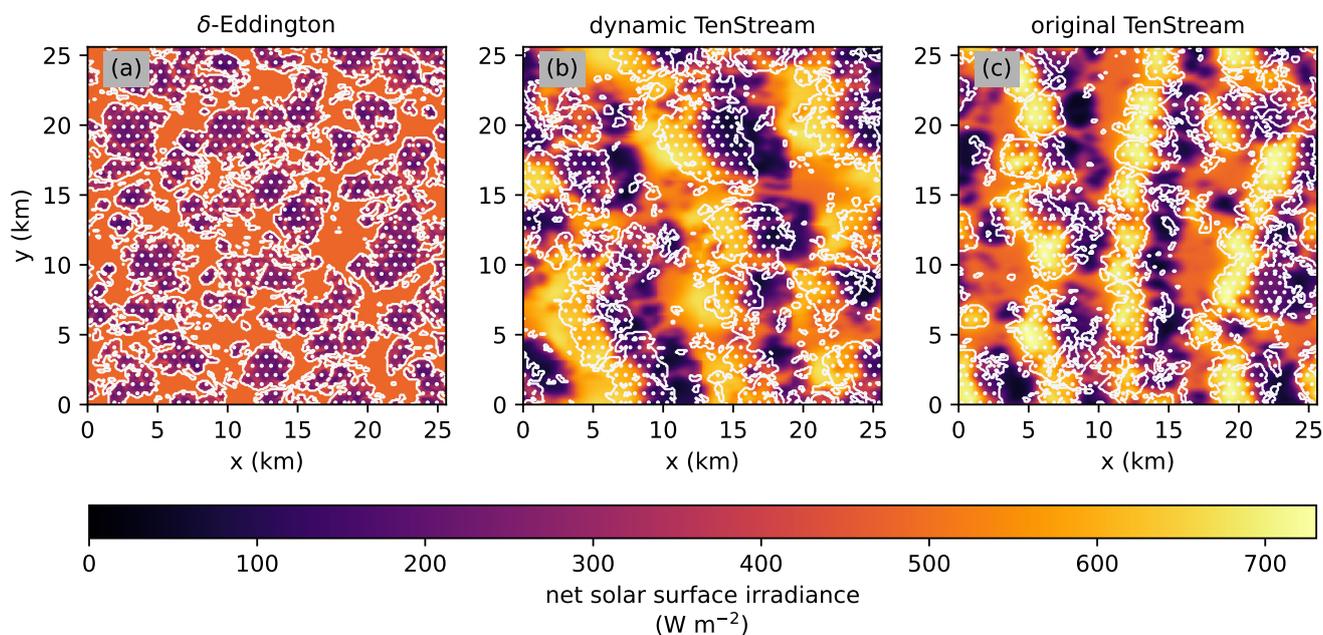


Figure 9. Positioning of the clouds relative to the net solar surface irradiance fields at 15:00 UTC in the simulations coupled to the δ -Eddington (a), dynamic TenStream (b) and original TenStream (c) solvers. To this end, the white contour lines in each plot enclose areas where the vertical columns above each grid point contain at least one grid box with $q_c > 0.01 \text{ g kg}^{-1}$, classifying them as cloudy. To improve the overall contrast of the plots, the maximum value of the colorbar was set to the highest net surface irradiance across the three panels, rather than the global maximum from all simulations, as it was done in Fig. 7.

discontinuities at the respective domain boundaries. These discontinuities occur because the incomplete solves in the dynamic TenStream solver are performed independently for all the subdomains, with interactions between them taking place just once at the end of each radiation scheme call. As a result, the individual subdomains have not yet updated their radiative fluxes based on the new boundary conditions at the end of each call, leading to the observed artifacts. However, these discontinuities are relatively small and, for this setup, do not significantly impact the representation of any of the 3D effects discussed previously.

Now, let us examine how these effects might relate to the differences in cloud characteristics observed in the previous subsection. To this end, we compare panels (j)–(l) in Fig. 7 with panels (g)–(i) in Fig. 3. Figure 9 facilitates this comparison by projecting the outlines of the corresponding clouds onto the underlying net surface irradiance fields. Starting with panel (a), we can see that in the simulation coupled to 1D radiation, cloud shadows are unsurprisingly placed directly below the corresponding clouds. This placement of shadows, however, suppresses the updrafts responsible for cloud formation, thereby reducing both the size and lifetime of the clouds (e.g., Schumann et al., 2002; Horn et al., 2015). In contrast, the original TenStream solver in panel (c) shows a displacement of the shadows to the east of the corresponding clouds. Beneath them, we instead observe regions of increased net surface irradiance that even exceed the clear-sky values in panel (a). Panel (b), which illustrates the results for the dynamic TenStream solver, shows a similar positioning of the clouds above areas with enhanced net



415 surface irradiance. However, unlike in panel (c), the clouds in panel (b) clearly extend into the shadowed areas as well. This behavior can again 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. The enhanced surface heating in these areas subsequently likely strengthens, rather than weakens, the corresponding updrafts, potentially explaining why the clouds in both 3D
420 simulations are larger and thicker compared to those in the 1D simulation during daytime.

Ultimately, this hypothesis is difficult to prove, as the clouds are not only shaped by the corresponding radiative fields but also actively influence them through their evolving characteristics. However, at least some evidence that the observed changes in cloud characteristics are primarily driven by the radiative field can be found in panels (a)–(c) of Fig. 7, where the clouds have
425 not yet evolved but are already subject to the discussed effects. Of course, this it is not enough to prove our hypothesis, though. Exactly these challenges associated with the coupling of radiation and model dynamics were also highlighted by Tjihuis et al. (2024), who investigated a similar hypothesis. They came to comparable conclusions, although their more statistical approach found the wind–sun angle to be an even more important factor for the development of larger and thicker clouds during daytime than just the shadow displacement. The mechanism behind this dependence on the wind–sun angle is that clouds driven by
430 3D radiation cannot significantly outgrow their 1D counterparts if they move into their own shadows, as this suppresses their updrafts, similar to what occurs in 1D simulations. Conversely, if the clouds do not move into their own shadows, the updrafts can persist or even intensify, allowing the clouds to live longer and grow larger (Tjihuis et al., 2024). 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
435 moving into their own shadows, starting with the panels at 13:00 UTC in Fig. 7. However, this does not necessarily negate the existence of this mechanism. Instead, it suggests that in the simulations here, the surface response to cloud movements may be rapid enough to maintain the updrafts in sunlit regions. 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. The representation of these cloud–enhanced areas relies on unique 3D effects, such as cloud–side escape (Hogan and Shonk, 2013) and entrapment (Hogan et al., 2019), which can only be accounted for with full 3D radiative transfer, underscoring the
440 importance of 3D radiative transfer for the correct representation of cloud characteristics in numerical models.

Up to this point, our analysis has primarily focused on differences in the spatial structure of the net surface irradiance fields and their potential relation to the differences in cloud characteristics observed in Sect. 3.2. What remains to be investigated, however, is whether the simulations coupled to the δ -Eddington approximation or the dynamic TenStream solver exhibit inherent biases when compared to the TenStream reference solution. To investigate this aspect as well, we turn to Fig. 10, where
445 panel (a) illustrates the temporal evolution of the domain-averaged net surface irradiances presented in Fig. 7. Despite the substantial differences in the spatial structure of the corresponding fields, we can see that the domain averages evolve quite similarly across all simulations. Panel (c) quantifies the deviations to the TenStream reference run, showing that for most of the day, both the δ -Eddington and the dynamic TenStream runs differ by less than 10 W m^{-2} from the reference solution. It is only after about 14:00 UTC that the δ -Eddington run becomes subject to larger deviations, peaking at around 25 W m^{-2} .

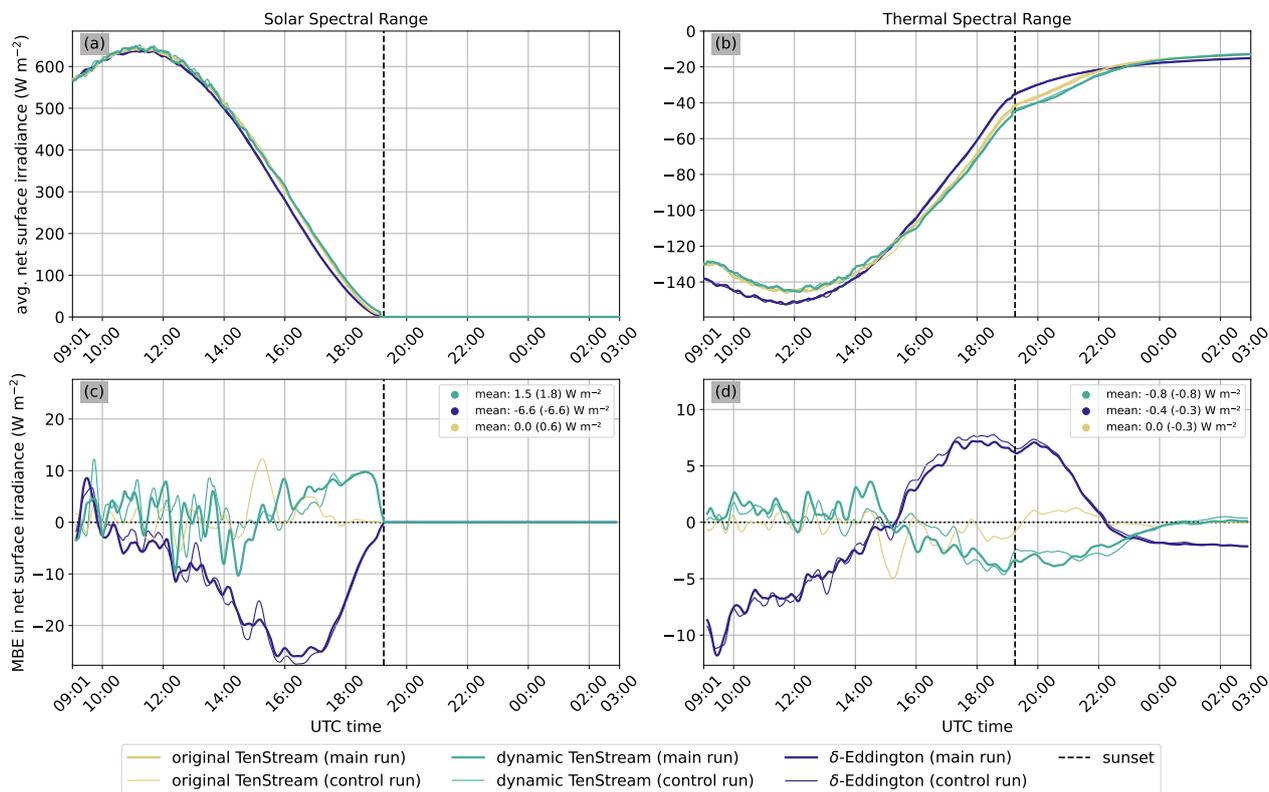


Figure 10. Temporal evolution of the domain-averaged net surface irradiance in the solar (a) and thermal (b) spectral ranges for the six simulations considered in this evaluation. Panels (c) and (d) show the corresponding differences to the TenStream reference run, with temporal mean differences for the simulations coupled to the different solvers listed in the legends (values for the control runs are given in brackets). To improve readability, a 15-minute running mean was applied to the data in (c) and (d). The vertical line in each plot indicates the time of sunset.

450 However, going back to Fig. 4, we can see that this divergence coincides with a noticeable increase in cloud cover in the 1D simulations after 14:00 UTC, which naturally reduces the net surface irradiance compared to the 3D simulations, where the cloud cover remains lower. Apart from this specific difference, we can conclude that neither the δ -Eddington approximation nor the dynamic TenStream solver introduces a substantial bias into the net surface irradiance in the solar spectral range. In particular, no significant accumulation of bias over time is observed when using the dynamic TenStream solver, in contrast to

455 the findings of Maier et al. (2024). It should be noted, however, that in their Fig. 9 this bias accumulation mostly remained below 10 W m⁻², and hence within the variability of the bias seen in panel (c) of Fig. 10.

Before comparing these results to those of other studies, let us first recall which differences in terms of the domain-averaged solar net surface irradiance we generally expect between 1D and 3D radiative transfer solvers. As illustrated in Fig. 10 of Maier et al. (2024), 1D radiative transfer tends to underestimate the net surface irradiance at small solar zenith angles, as it cannot



460 account for cloud-side escape and entrapment, which both enhance the diffuse downward radiation. Conversely, at large solar
zenith angles, 1D solvers typically overestimate the net surface irradiance, because they cannot cast shadows at slant angles,
thereby underestimating their size. In the fully interactive PALM simulations, however, especially the underestimation at small
solar zenith angles is counteracted by the development of larger and thicker clouds in the 3D simulations, which increase
absorption and backscattering to space, as explained in Tjihuis et al. (2024). As a result, and consistent with their findings, no
465 substantial differences in the domain-averaged net surface irradiance are observed between simulations coupled to 1D and 3D
radiation. This underscores that, at least in the solar spectral range, 3D radiative transfer is primarily important to accurately
model the spatial structure of the radiative fields. However, it does not substantially affect the domain-averaged properties of
these fields.

This is slightly different in the thermal spectral range. Panel (b) of Fig. 10 shows the temporal evolution of the domain-
470 averaged net surface irradiance for the different simulations in this spectral range. We can see that initially, the simulations
coupled to 1D radiation experience systematically stronger cooling, similar to what was observed in Fig. 9 of Maier et al.
(2024). This is most likely because the increased thermal emission of clouds affects not only the cloudy columns themselves
when using 3D radiative transfer, but also neighboring ones due to the horizontal transport of energy in these solvers. Especially
early on, when the surface temperature — which primarily determines the upward thermal irradiance at the surface — is still
475 similar across all simulations, the correspondingly increased downward irradiance at the surface is likely the main factor
explaining the systematically weaker net cooling observed with 3D radiation. This pattern changes, however, at around 15:00
UTC, when the cloud cover in the 1D simulations noticeably increases. Along with this shift in cloud cover, the net thermal
emission in the 1D simulations decreases and becomes lower than in the 3D simulations. As shown in Fig. 16 of Maier et al.
(2024), net thermal emission is significantly reduced below clouds when using 1D radiative transfer, explaining this noticeable
480 decrease in net thermal emission as the cloud cover starts to rise. In contrast, the simulations coupled to the dynamic TenStream
solver show much smaller deviations from the reference run, with differences staying below 5 W m^{-2} throughout the entire
time, as one can see in panel (d) of Fig. 10. This once more indicates that the dynamic TenStream solver, despite its use of
incomplete solves, is still able to capture all the key features of the full 3D solutions.

Since the temporal evolution of the domain-averaged net solar surface irradiance showed only minor differences between the
485 solvers, the observed differences in the domain-averaged net thermal surface irradiance must be balanced by other components
of the surface energy budget. In general, the net solar irradiance at the ground is balanced by net thermal emission, the ground
heat flux, and the release of sensible and latent heat into the atmosphere. Hence, the observed differences in the net thermal
irradiance can result in changes in any of the latter three components. To investigate which of them compensates for the
observed differences, Fig. 11 illustrates the temporal evolution of the sensible and latent heat fluxes across all simulations
490 considered in this evaluation. Looking at panels (a) and (b), first note that the Bowen ratio — i.e., the ratio of the sensible to
the latent heat flux (Bowen, 1926) — is larger than one in all simulations. This is rather unusual for vegetated surfaces such
as the flat grassland used here, and more characteristic of urban environments or semi-arid regions (e.g., Stull, 2006; Kotthaus
and Grimmond, 2014), suggesting that the soil in the model is relatively dry. Apart from this general remark, panel (a) of
Fig. 11 shows that the sensible heat flux evolves quite similarly across all simulations until about 15:00 UTC. At that point,

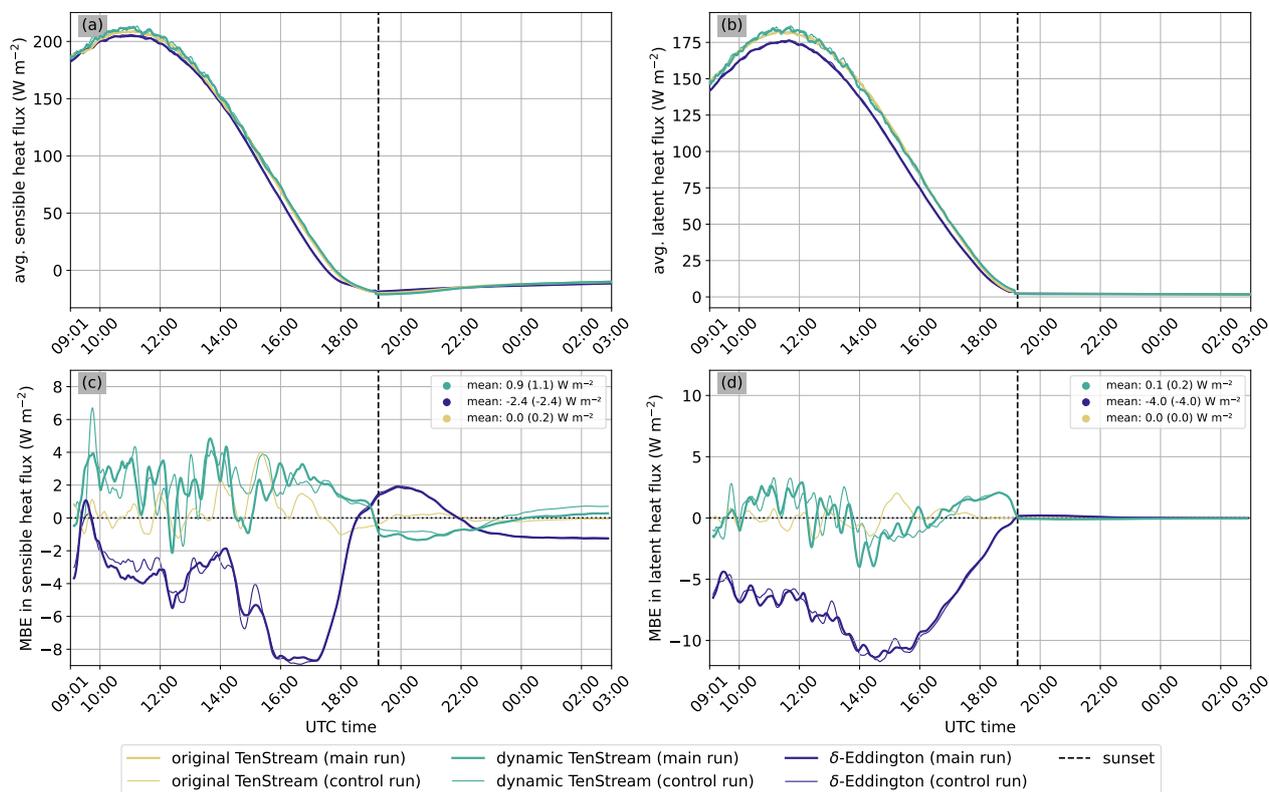


Figure 11. Temporal evolution of the domain-averaged sensible (a) and latent (b) heat fluxes for all six simulations considered in this evaluation. Panels (c) and (d) show the corresponding differences to the TenStream reference run, with temporal mean differences for the simulations coupled to the different solvers listed in the legends (values for the control runs are provided in brackets). To improve readability, a 15-minute running mean was applied to the data in panels (c) and (d). The vertical line in each plot marks the time of sunset.

495 the simulations coupled to 1D radiation begin to slightly diverge from those using 3D radiation. As discussed earlier, this divergence coincides with a substantial increase in cloud cover in the 1D simulations at that time, subsequently leading to a reduction of the net solar surface irradiance. Panel (a) indicates that this reduction in solar energy input also affects the release of sensible heat into the atmosphere, which is subsequently reduced by up to $8 W m^{-2}$ in the 1D simulations, as can be seen in panel (c).

500 A much more striking difference between the 1D and 3D simulations is revealed in panel (b), which shows that the domain-averaged latent heat flux is considerably higher in the simulations coupled to 3D radiation, accompanied by a slightly faster decrease in soil moisture (not shown here). This finding suggests that the reduced thermal emission in the 3D simulations observed in Fig. 10 — which could be balanced by a decrease in net solar irradiance, or by increases in the ground, sensible, or latent heat fluxes, or a combination thereof — is primarily offset by an increased release of water vapor into the atmosphere.



505 In PALM, this latent heat flux is parameterized as

$$LE = -\frac{\rho l_v}{r_a + r_s} \left[q_v \left(\frac{\Delta z}{2} \right) - q_{v,\text{sat}}(T_0) \right], \quad (3)$$

where ρ is the density of dry air, $l_v = 2.5 \times 10^6 \text{ J kg}^{-1}$ is the specific latent heat of vaporization, r_a and r_s are the aerodynamic and surface resistances (in units of s m^{-1}), $q_v \left(\frac{\Delta z}{2} \right)$ is the water vapor mixing ratio at height $z = \frac{\Delta z}{2}$, with Δz denoting the vertical grid spacing, and $q_{v,\text{sat}}(T_0)$ is the saturation vapor mixing ratio at the radiative temperature T_0 of the surface skin layer (Maronga et al., 2020). Since the domain-averaged $q_v \left(\frac{\Delta z}{2} \right)$ differs by less than 0.7% between the simulations (not shown here), the enhanced latent heat flux in the 3D simulations is unlikely to be driven by changes in near-surface humidity. Instead, it is most consistent with an increase in $q_{v,\text{sat}}(T_0)$, and thus in the radiative temperature T_0 of the surface skin layer. In PALM, such an increase in T_0 is governed by imbalances in the ground energy budget, following the relationship

$$C_0 \frac{dT_0}{dt} = R_n - H - LE - G, \quad (4)$$

515 where C_0 denotes the heat capacity of the surface skin layer (in units of $\text{J m}^{-2} \text{ K}^{-1}$), R_n is the net radiative intake at the surface, H is the sensible heat flux, and G is the ground heat flux (Maronga et al., 2020). In the 3D simulations, R_n is increased across both spectral ranges, while H , LE , and G initially remain unchanged. This is consistent with an increase in T_0 in these simulations, which subsequently leads to increases in H , LE , and G . Since $q_{v,\text{sat}}$ is directly proportional to the saturation vapor pressure e_w^* , which itself depends exponentially on T_0 , even small increases in T_0 can lead to substantial increases in $q_{v,\text{sat}}(T_0)$ and, consequently, in LE , explaining why the elevated skin layer temperature affects the latent heat flux more strongly than the sensible one. Moreover, this relationship supports the interpretation that the higher latent heat fluxes in the 3D simulations are primarily driven by increases in $q_{v,\text{sat}}(T_0)$, since changes in the aerodynamic resistance r_a — which could also affect LE — would tend to affect both H and LE more uniformly. 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. And looking at panel (d), we can see that the simulations coupled to the dynamic TenStream solver also capture this effect, with deviations from the TenStream reference solution remaining below 4 W m^{-2} compared to deviations of more than 10 W m^{-2} for the 1D simulations.

525 All in all, two potential links have now been identified between radiation and cloud characteristics that help explain why clouds in simulations coupled to 3D radiative transfer grow larger, become thicker, and contain more liquid water during the day than their 1D-driven counterparts:

1. First, clouds in simulations coupled to 3D radiative transfer turned out to be positioned above areas of increased net surface irradiance, rather than above their own shadows. This placement causes the associated updrafts to persist or even strengthen instead of weakening.
2. Additionally, 3D radiative transfer reduces the domain-averaged net thermal emission at the ground, most likely because the increased thermal emission of clouds enhances the downward longwave radiation not only in cloudy, but also in neighboring columns. This reduction in net thermal surface irradiance affects the ground energy budget and appears to



be primarily balanced by an increase in the domain-averaged latent heat flux, resulting in a greater release of water vapor into the atmosphere.

Both of these effects are captured not only by the full original TenStream model but also by the dynamic TenStream solver. This demonstrates that the incomplete solves in the latter provide a computationally efficient approach that not only captures the essential features of 3D radiative fields, but also reproduces most of the resulting cloud characteristics in fully interactive PALM simulations. These include the development of larger and thicker clouds containing more liquid water during the day, as well as thinner clouds with less liquid water at night. The causes of these nighttime cloud characteristics, however, are more difficult to disentangle and would require additional experiments, such as restarting the simulations at sunset, to separate the imprint of daytime evolution from genuinely nocturnal radiative effects, which is beyond the scope of this study.

Finally, it should also be noted that the original TenStream solver used as a reference in this paper is an approximation itself. The differences between this approximation and highly accurate 3D radiative transfer solvers such as the Monte Carlo model MYSTIC have been discussed in Maier et al. (2024) and are mostly much smaller than the 3D radiative effects and their subsequent influences on cloud development observed in this paper. Especially the results during daytime have furthermore been shown to be in good agreement with those of Tjihuis et al. (2024), which were obtained with a GPU-accelerated Monte Carlo ray tracer. Thus, while the exact magnitude of some of the effects discussed in this paper might slightly vary when using a highly accurate 3D radiative transfer model such as MYSTIC, the key features and underlying mechanisms are expected to remain unchanged.

4 Summary and outlook

The primary objective of this paper was to investigate how coupling the dynamic TenStream solver to model dynamics affects cloud development and how the resulting clouds compare to those driven by classical 1D and full 3D radiation. To answer these questions, PALM (Raasch and Schröter, 2001; Maronga et al., 2015, 2020) was used to perform a large-eddy simulation (LES) in which shallow cumulus clouds developed over the course of the day. Once the clouds had formed, three restart runs of the simulation were conducted, each coupled to a different radiative transfer model: either a classical 1D δ -Eddington approximation, the original TenStream model, or the new dynamic TenStream solver. Starting from the same cloud field, this setup then enabled a direct comparison of how the simulated clouds evolved depending on the radiative transfer model used. 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, consistent with the findings of Jakub and Mayer (2017), whereas in the 1D simulation, they remained more or less randomly positioned. Clouds driven by the two 3D solvers were also shown to grow larger, become thicker, and contain more liquid water than those coupled to 1D radiation, in agreement with Veerman et al. (2020, 2022) and Tjihuis et al. (2024). After sunset, however, this pattern reversed, with the clouds in the 1D simulation becoming thicker and containing more liquid water than their 3D-driven counterparts. Throughout both day and night, the dynamic TenStream solver successfully captured all these 3D effects, with discrepancies from the original TenStream model remaining small compared to



570 the differences observed for the 1D simulation. In contrast to these averaged cloud characteristics, the positioning of individual clouds in the dynamic and full TenStream simulations, however, indeed differed over time, illustrating that even the small differences in their corresponding radiative fields influenced individual cloud development over time, despite sharing very similar overall characteristics.

Building on these observations, this paper then sought to identify the mechanisms driving the described differences in cloud characteristics, with a particular focus on the daytime regime. A comparison of the corresponding net solar surface irradiance fields revealed first key differences. As expected, cloud shadows in the 3D simulations were displaced according to the angle of solar incidence, in contrast to the 1D simulation, where they were positioned directly beneath the corresponding clouds. At the sunward edges of the shadows, furthermore, areas of enhanced net surface irradiance were identified in the 3D simulations, with values even exceeding those in the clear-sky columns of the 1D simulation. Unlike their 1D-driven counterparts, the clouds in the 3D simulations turned out to be located above these areas of enhanced net surface irradiance. This positioning likely caused the associated updrafts to persist rather than weaken, promoting the observed differences in cloud characteristics. In addition, 3D radiative transfer was found to reduce the domain-averaged net thermal emission at the ground, most likely because the increased thermal emission of clouds enhances the downward longwave radiation not only in cloudy, but also in neighboring columns with this type of solver. This reduction in net thermal surface irradiance affected the ground energy budget in these simulations and appeared to be primarily balanced by an increase in the domain-averaged latent heat flux, resulting in a greater release of water vapor into the atmosphere. Together, both of these effects helped explain why clouds in the 3D simulations grew larger, became thicker, and contained more liquid water during the day than those in the 1D simulation. And even more importantly, both the original TenStream model and the dynamic TenStream solver captured them, further underscoring the latter's ability to represent 3D radiative effects on cloud development at a significantly lower computational cost.

590 It should be noted, however, that this study focused exclusively on shallow cumulus clouds with a relatively high cloud cover of more than 40%. In reality, though, clouds occur in a wide variety of forms. On the one hand, there are clear-sky or completely overcast situations, which are essentially one-dimensional, with little to no horizontal variability and thus negligible 3D radiative effects. Multi-layer cloud fields and deep convection, on the other hand, represent more complex scenarios than the shallow cumulus clouds discussed here, as their larger vertical extent introduces radiative interactions across greater distances. Such cloud fields may therefore pose a stronger challenge for the incomplete solves in the dynamic TenStream solver. However, since all vertical layers above a given grid point belong to the same model subdomain, these vertical interactions can still occur within a single call to the radiation scheme even with incomplete solves, alleviating these concerns to some extent. This is not necessarily the case for cloud fields with different cloud covers. As discussed in Sect. 3.2, the relatively high cloud cover in this study, for example, likely favored the performance of the dynamic TenStream solver, because the clouds were close together in terms of grid-box distance, reducing the distance over which information crucial for 3D radiative transfer needed to be exchanged. Lower cloud covers, by contrast, increase this spacing and may thus weaken the solver's performance. Future studies should therefore examine more diverse cloud configurations, including cases with more complex vertical structures and lower cloud cover, to evaluate the robustness of the solver under a wider range of conditions. These future studies could then



also adopt more realistic simulation setups, including real-world initial profiles and external large-scale forcing, to assess the
605 3D radiative effects observed in this paper under more realistic dynamical conditions as well.

At this point it is also worth mentioning that while this work focused on high-resolution LES output, the dynamic Ten-
Stream solver was developed with numerical weather prediction (NWP) models in mind. At their current kilometer-scale
resolution, clouds are inherently closer together in terms of grid-box distance, which should significantly mitigate the expected
performance degradation in low-cloud-cover scenarios. In this regime, however, model grid boxes can no longer be treated as
610 homogeneous, as it is currently assumed in the solver. Hence, the TenStream lookup tables used to determine the transport
coefficients in the underlying system of linear equations will need to be extended to account for cloud fraction as well. In
addition, although the dynamic TenStream solver is markedly faster than other 3D radiative transfer solvers, it remains about
three times slower than a classical 1D δ -Eddington solver (Maier et al., 2024). Since a new radiative transfer scheme should
ideally operate at a comparable cost to the 1D independent-column approximations employed in today's NWP models, further
615 optimization is essential. One promising direction in that way would be to accelerate the retrieval of the transport coefficients
from the TenStream lookup tables, which currently accounts for roughly one-third of the solver's total runtime. Achieving such
improvements would then finally enable the application of the dynamic TenStream solver at the scale it was originally intended
for, paving the way toward addressing the open question of how 3D radiative transfer affects weather forecasts.

Code and data availability. The simulations were performed with PALM version 25.04 (PALM, 2025). The model input and restart files
620 required to rerun these simulations, as well as the complete set of model output files, are available from Maier et al. (2026). Details on the
methodology used to analyze the output data and reproduce the figures in Sect. 3 are given in Sect. 2.2.

Author contributions. RM designed and conducted the study and carried out its analysis, with conceptual input from BM and FJ. FJ imple-
mented the dynamic TenStream solver into the TenStream framework, provided technical support regarding PALM, assisted in setting up the
simulations, and, together with BM and FH, contributed ideas for their evaluation. RM prepared the manuscript with contributions from all
625 co-authors.

Competing interests. The authors declare that they have no conflict of interest.

Disclaimer. This paper is based on material from the first author's dissertation (Maier, 2025), with revisions made to fit the journal-article
format and to reflect input from the co-authors.

Acknowledgements. The wording and phrasing in parts of this paper were improved with the assistance of OpenAI's ChatGPT models.



630 References

- Anderson, G., Clough, S., Kneizys, F., Chetwynd, J., and Shettle, E.: AFGL atmospheric constituent profiles (0-120 km), Tech. Rep. AFGL-TR-86-0110, Air Force Geophys. Lab., Hanscom Air Force Base, Bedford, Mass., <https://apps.dtic.mil/sti/pdfs/ADA175173.pdf>, last access: 20 February 2026, 1986.
- Bowen, I. S.: The Ratio of Heat Losses by Conduction and by Evaporation from any Water Surface, *Physical Review*, 27, 779–787, <https://doi.org/10.1103/physrev.27.779>, 1926.
- 635 Davis, A., Gabriel, P., Lovejoy, S., Schertzer, D., and Austin, G. L.: Discrete angle radiative transfer: 3. Numerical results and meteorological applications, *Journal of Geophysical Research: Atmospheres*, 95, 11 729–11 742, <https://doi.org/10.1029/jd095id08p11729>, 1990.
- DWD: Operational NWP System: Model configuration upgrade of ICON, https://www.dwd.de/DE/fachnutzner/forschung_lehre/numerische_wettervorhersage/nwv_aenderungen/_functions/DownloadBox_modellaenderungen/icon/pdf_2021/pdf_icon_14_04_2021.pdf?__blob=publicationFile&v=3, last access: 20 February 2026, 2021.
- 640 Gabriel, P., Lovejoy, S., Davis, A., Schertzer, D., and Austin, G. L.: Discrete angle radiative transfer: 2. Renormalization approach for homogeneous and fractal clouds, *Journal of Geophysical Research: Atmospheres*, 95, 11 717–11 728, <https://doi.org/10.1029/jd095id08p11717>, 1990.
- Gristey, J. J., Feingold, G., Glenn, I. B., Schmidt, K. S., and Chen, H.: Surface Solar Irradiance in Continental Shallow Cumulus Fields: Observations and Large-Eddy Simulation, *Journal of the Atmospheric Sciences*, 77, 1065–1080, <https://doi.org/10.1175/jas-d-19-0261.1>, 2020.
- 645 Hogan, R. J. and Bozzo, A.: A Flexible and Efficient Radiation Scheme for the ECMWF Model, *Journal of Advances in Modeling Earth Systems*, 10, 1990–2008, <https://doi.org/10.1029/2018ms001364>, 2018.
- Hogan, R. J. and Shonk, J. K. P.: Incorporating the Effects of 3D Radiative Transfer in the Presence of Clouds into Two-Stream Multilayer Radiation Schemes, *Journal of the Atmospheric Sciences*, 70, 708–724, <https://doi.org/10.1175/jas-d-12-041.1>, 2013.
- 650 Hogan, R. J., Fielding, M. D., Barker, H. W., Villefranque, N., and Schäfer, S. A. K.: Entrapment: An Important Mechanism to Explain the Shortwave 3D Radiative Effect of Clouds, *Journal of the Atmospheric Sciences*, 76, 2123–2141, <https://doi.org/10.1175/JAS-D-18-0366.1>, 2019.
- Horn, G. L., Ouwersloot, H. G., Vilà-Guerau de Arellano, J., and Sikma, M.: Cloud Shading Effects on Characteristic Boundary-Layer Length Scales, *Boundary-Layer Meteorology*, 157, 237–263, <https://doi.org/10.1007/s10546-015-0054-4>, 2015.
- 655 Jakub, F. and Mayer, B.: A three-dimensional parallel radiative transfer model for atmospheric heating rates for use in cloud resolving models—The TenStream solver, *Journal of Quantitative Spectroscopy and Radiative Transfer*, 163, 63–71, <https://doi.org/10.1016/j.jqsrt.2015.05.003>, 2015.
- Jakub, F. and Mayer, B.: The role of 1-D and 3-D radiative heating in the organization of shallow cumulus convection and the formation of cloud streets, *Atmospheric Chemistry and Physics*, 17, 13 317–13 327, <https://doi.org/10.5194/acp-17-13317-2017>, 2017.
- 660 Khairoutdinov, M. and Kogan, Y.: A New Cloud Physics Parameterization in a Large-Eddy Simulation Model of Marine Stratocumulus, *Monthly Weather Review*, 128, 229–243, [https://doi.org/10.1175/1520-0493\(2000\)128<0229:ancppi>2.0.co;2](https://doi.org/10.1175/1520-0493(2000)128<0229:ancppi>2.0.co;2), 2000.
- Khvorostyanov, V. I. and Curry, J. A.: Aerosol size spectra and CCN activity spectra: Reconciling the lognormal, algebraic, and power laws, *Journal of Geophysical Research: Atmospheres*, 111, D12 202, <https://doi.org/10.1029/2005jd006532>, 2006.



- 665 Klinger, C. and Mayer, B.: The Neighboring Column Approximation (NCA) – A fast approach for the calculation of 3D thermal heating rates in cloud resolving models, *Journal of Quantitative Spectroscopy and Radiative Transfer*, 168, 17–28, <https://doi.org/10.1016/j.jqsrt.2015.08.020>, 2016.
- Klinger, C. and Mayer, B.: Neighboring Column Approximation—An Improved 3D Thermal Radiative Transfer Approximation for Non-Rectangular Grids, *Journal of Advances in Modeling Earth Systems*, 12, e2019MS001 843, <https://doi.org/10.1029/2019ms001843>, 2020.
- 670 Klinger, C., Mayer, B., Jakub, F., Zinner, T., Park, S.-B., and Gentine, P.: Effects of 3-D thermal radiation on the development of a shallow cumulus cloud field, *Atmospheric Chemistry and Physics*, 17, 5477–5500, <https://doi.org/10.5194/acp-17-5477-2017>, 2017.
- Kotthaus, S. and Grimmond, C.: Energy exchange in a dense urban environment – Part I: Temporal variability of long-term observations in central London, *Urban Climate*, 10, 261–280, <https://doi.org/10.1016/j.uclim.2013.10.002>, 2014.
- Lovejoy, S., Davis, A., Gabriel, P., Schertzer, D., and Austin, G. L.: Discrete angle radiative transfer: 1. Scaling and similarity, universality and diffusion, *Journal of Geophysical Research: Atmospheres*, 95, 11 699–11 715, <https://doi.org/10.1029/jd095id08p11699>, 1990.
- 675 Maier, R.: Development and application of a dynamic approach to 3D radiative transfer in subkilometer-scale numerical weather prediction models, Ph.D. thesis, Ludwig-Maximilians-Universität München, <https://doi.org/10.5282/EDOC.35575>, 2025.
- Maier, R., Jakub, F., Emde, C., Manev, M., Voigt, A., and Mayer, B.: A dynamic approach to three-dimensional radiative transfer in subkilometer-scale numerical weather prediction models: the dynamic TenStream solver v1.0, *Geoscientific Model Development*, 17, 3357–3383, <https://doi.org/10.5194/gmd-17-3357-2024>, 2024.
- 680 Maier, R., Jakub, F., and Mayer, B.: PALM-LES dataset for evaluating the impact of radiative transfer solvers on shallow cumulus cloud development, <https://doi.org/10.57970/NP5R3-T5X93>, LMU Munich, Faculty of Physics [data set], 2026.
- Maronga, B., Gryscha, M., Heinze, R., Hoffmann, F., Kanani-Sühring, F., Keck, M., Ketelsen, K., Letzel, M. O., Sühring, M., and Raasch, S.: The Parallelized Large-Eddy Simulation Model (PALM) version 4.0 for atmospheric and oceanic flows: model formulation, recent developments, and future perspectives, *Geoscientific Model Development*, 8, 2515–2551, <https://doi.org/10.5194/gmd-8-2515-2015>, 2015.
- 685 Maronga, B., Banzhaf, S., Burmeister, C., Esch, T., Forkel, R., Fröhlich, D., Fuka, V., Gehrke, K. F., Geletič, J., Giersch, S., Gronemeier, T., Groß, G., Heldens, W., Hellsten, A., Hoffmann, F., Inagaki, A., Kadasch, E., Kanani-Sühring, F., Ketelsen, K., Khan, B. A., Knigge, C., Knoop, H., Krč, P., Kurppa, M., Maamari, H., Matzarakis, A., Mauder, M., Pallasch, M., Pavlik, D., Pfafferott, J., Resler, J., Rissmann, S., Russo, E., Salim, M., Schrempf, M., Schwenkel, J., Seckmeyer, G., Schubert, S., Sühring, M., von Tils, R., Vollmer, L., Ward, S., Witha, B., Wurps, H., Zeidler, J., and Raasch, S.: Overview of the PALM model system 6.0, *Geoscientific Model Development*, 13, 1335–1372, <https://doi.org/10.5194/gmd-13-1335-2020>, 2020.
- 690 Mol, W. B., Knap, W. H., and van Heerwaarden, C. C.: Ten years of 1 Hz solar irradiance observations at Cabauw, the Netherlands, with cloud observations, variability classifications, and statistics, *Earth System Science Data*, 15, 2139–2151, <https://doi.org/10.5194/essd-15-2139-2023>, 2023.
- 695 Morrison, H. and Grabowski, W. W.: Comparison of Bulk and Bin Warm-Rain Microphysics Models Using a Kinematic Framework, *Journal of the Atmospheric Sciences*, 64, 2839–2861, <https://doi.org/10.1175/jas3980>, 2007.
- Neggens, R. A. J., Duynkerke, P. G., and Rodts, S. M. A.: Shallow cumulus convection: A validation of large-eddy simulation against aircraft and Landsat observations, *Quarterly Journal of the Royal Meteorological Society*, 129, 2671–2696, <https://doi.org/10.1256/qj.02.93>, 2003.
- PALM: PALM model system, version 25.04, https://gitlab.palm-model.org/releases/palm_model_system/-/tree/v25.04, last access: 20 February 2026, 2025.
- 700 PALM: Land surface model, <https://palm.muk.uni-hannover.de/trac/wiki/doc/tec/lsm>, last access: 20 February 2026, 2026a.



- PALM: Land surface parameters, https://palm.muk.uni-hannover.de/trac/wiki/doc/app/land_surface_parameters, last access: 20 February 2026, 2026b.
- 705 Pincus, R., Barker, H. W., and Morcrette, J.-J.: A fast, flexible, approximate technique for computing radiative transfer in inhomogeneous cloud fields, *Journal of Geophysical Research: Atmospheres*, 108, 4376, <https://doi.org/10.1029/2002JD003322>, 2003.
- Raasch, S. and Schröter, M.: PALM - A large-eddy simulation model performing on massively parallel computers, *Meteorologische Zeitschrift*, 10, 363–372, <https://doi.org/10.1127/0941-2948/2001/0010-0363>, 2001.
- Schumann, U., Dörnbrack, A., and Mayer, B.: Cloud-shadow effects on the structure of the convective boundary layer, *Meteorologische Zeitschrift*, 11, 285–294, <https://doi.org/10.1127/0941-2948/2002/0011-0285>, 2002.
- 710 Seifert, A. and Beheng, K. D.: A two-moment cloud microphysics parameterization for mixed-phase clouds. Part 1: Model description, *Meteorology and Atmospheric Physics*, 92, 45–66, <https://doi.org/10.1007/s00703-005-0112-4>, 2005.
- Stull, R.: The Atmospheric Boundary Layer, in: *Atmospheric Science: An Introductory Survey*, chap. 9, pp. 375–417, Elsevier, second edn., ISBN 978-0-12-732951-2, <https://doi.org/10.1016/b978-0-12-732951-2.50014-4>, 2006.
- Tijhuis, M., van Stratum, B. J. H., Veerman, M. A., and van Heerwaarden, C. C.: An Efficient Parameterization for Surface Shortwave 3D Radiative Effects in Large-Eddy Simulations of Shallow Cumulus Clouds, *Journal of Advances in Modeling Earth Systems*, 15, e2022MS003 262, <https://doi.org/10.1029/2022ms003262>, 2022.
- 715 Tijhuis, M., van Stratum, B. J. H., and van Heerwaarden, C. C.: The impact of coupled 3D shortwave radiative transfer on surface radiation and cumulus clouds over land, *Atmospheric Chemistry and Physics*, 24, 10 567–10 582, <https://doi.org/10.5194/acp-24-10567-2024>, 2024.
- Veerman, M. A., Pedruzo-Bagazgoitia, X., Jakub, F., Vilà-Guerau de Arellano, J., and van Heerwaarden, C. C.: Three-Dimensional Radiative Effects By Shallow Cumulus Clouds on Dynamic Heterogeneities Over a Vegetated Surface, *Journal of Advances in Modeling Earth Systems*, 12, e2019MS001 990, <https://doi.org/10.1029/2019ms001990>, 2020.
- 720 Veerman, M. A., van Stratum, B. J. H., and van Heerwaarden, C. C.: A Case Study of Cumulus Convection Over Land in Cloud-Resolving Simulations With a Coupled Ray Tracer, *Geophysical Research Letters*, 49, e2022GL100 808, <https://doi.org/10.1029/2022gl100808>, 2022.