Mem. S.A.It. Vol. 88, 22

Memorie della

Radiation transport in CO5BOLD

A short-characteristics module for local box models

M. Steffen

Leibniz-Institut für Astrophysik Potsdam, An der Sternwarte 16, 14482 Potsdam, Germany
e-mail: msteffen@aip.de

Abstract. We give a description of a new radiative transfer package based on an energy-conserving short-characteristics (SC) approach, developed for both CO5BOLD 3D radiation hydrodynamics local box simulations and NLTE3D line formation post processing. First test results for a realistic 3D solar atmosphere model indicate that the thermal structure obtained with the new radiation transport scheme closely matches the one found with the standard long-characteristics Feautrier (LCF) method that has served in CO5BOLD for many years. This is a very reassuring result, corroborating the validity of the CO5BOLD models computed previously with the LCF radiation transport.

Key words. Sun: atmosphere – Stars: atmospheres – Radiative transfer – Hydrodynamics

1. Introduction

One of the essential building blocks of any radiation hydrodynamics code is the radiative transfer (RT) module. Its main purpose is to compute the radiative heating rate ($Q_{\rm rad} = -\nabla \cdot \boldsymbol{F}_{\rm rad}$) for each grid cell in the model volume, given the 3D thermal structure. Among the challenges the numerical method has to meet, the most important requirement is that it must provide an accurate representation of the non-local radiation field in any optical depth regime, preferably without the need for an artificial selector switch between optically thick and thin conditions. In addition, particular care must be exercised to ensure that the numerical scheme is energy conserving, such that the volume integral of the radiative heating rate is consistent with the radiative surface fluxes.

Since the beginning, the radiative transfer scheme for CO5BOLD (Freytag et al. 2012) local box models was based on long rays. Later on, the same RT module was integrated in the 3D non-LTE line formation code NLTE3D (e.g., Steffen et al. 2012; Prakapavičius et al. 2013). In this context, the role of the RT calculations is to provide the angle averaged mean intensity $\overline{J}_\nu$ at the center of each grid cell, as well as the emergent fluxes. Here the main issue is to make sure that the resulting $\overline{J}_\nu$ are always non-negative, while energy conservation is not of primary importance.

In general, the long-characteristics radiative transfer method (MSrad) implemented in CO5BOLD was found to work well. One of the advantages is that, independently along each long ray, the non-local radiation field can be efficiently computed with the Feautrier method, a technique that is based on the differential form of the transport equation (for references see Sect. 2.1).

However, under extreme circumstances, the present implementation of the LCF scheme reveals its limitations. For example, in the presence of strong cell-to-cell variations of the thermodynamical conditions, the local radiation field is not always guaranteed to have the physically correct effect, such that hot pixels may be radiatively heated even further, or cool pixels may be cooled to even lower temperature, leading to a catastrophic runaway situation. This problem is particularly severe in the very dynamic atmospheres of cool giants, and puts a lower limit to range of surface gravity for which CO5BOLD local box models can be reliably computed. Similar effects might also be responsible for the development of excessively cool low-density regions inside magnetic flux concentrations found in some CO5BOLD MHD simulations of solar-type stellar atmospheres.

These and other reasons led us to the idea of developing an alternative and completely independent radiative transfer module for CO5BOLD *local box* models, with the main goal to overcome some of the deficiencies of the present LCF method. A short-characteristics scheme was considered the method of choice. In this approach, the transport equation is integrated along short rays which pass directly through the center of a grid cell. This immediate correspondence between ray system and hydrodynamical (HD) grid is not warranted in the case of the LC geometry.

It is worth mentioning that, in its star-in-a-box setup, CO5BOLD has been used successfully to compute *global models* of supergiants since many years (e.g., Freytag et al. 2002). In this mode, the radiation transport is computed with the SHORTradmodule, which is based on a short-characteristics scheme particularly designed for the geometry and boundary conditions of global convection simulations (Freytag et al. 2012). These global models demonstrate that a properly designed SC scheme is more robust in the presence of strong temperature and density fluctuations. It may thus be hoped that a similar approach will also allow local CO5BOLD 3D model atmospheres to be advanced towards lower surface gravities than previously possible with the LC RT module.

In the context of NLTE3D, the more direct correspondence between radiation field on the ray and in the cell centers is expected to lead to a better convergence of the Λ-iteration, which is essential for computing coherent continuum scattering and solving the statistical equilibrium equations. Last but not least, the availability of two independent RT options allows us to study the impact of the numerical implementation of the RT scheme on the resulting thermal structure of the 3D atmospheres, on the related synthetic spectra, and on the derived 3D spectroscopic abundances, which would otherwise be very hard to assess.

2. Basic description of the LC and SC radiative transfer schemes

In the present implementation of both RT modules, the computation of the 3D radiation field is decomposed into 2D sub steps where the radiative transfer is treated in 2D vertical slices, assuming periodic lateral boundary conditions. The grid is assumed to be equidistant in horizontal direction, but may be arbitrary in the vertical direction. To simplify matters, we consider only two sets of orthogonal 2D slices along the x and y direction of the hydrodynamical grid, i.e. we consider four angles for azimuthal integration of the radiation field ($\phi = 0, \pi/2, \pi, 3\pi/2$). The polar angles with respect to the vertical ($0 \leq \theta \leq \pi/2$, $\mu = \cos\theta$) may be arbitrary, in principle. In practice, they are given be the adopted quadrature formula for the μ-integration. By default, we use Lobatto's quadrature (e.g. Davis & Polonsky 1972) with typically between three and five μ-angles, including the vertical direction.

In the lower, optically very thick layers of the model, the problem can be simplified by computing the radiation field from the diffusion approximation, which is much faster and independent of the ray system: $Q_{\rm rad} = 4\pi/3\,\overline{\kappa\rho}\,(\partial^2 S/\partial\tau_x^2 + \partial^2 S/\partial\tau_y^2 + \partial^2 S/\partial\tau_z^2) = 4\pi\,\overline{\kappa\rho}\,(J - S)$, with τ_x, τ_y, τ_z denoting optical depth in x, y, z direction, respectively.

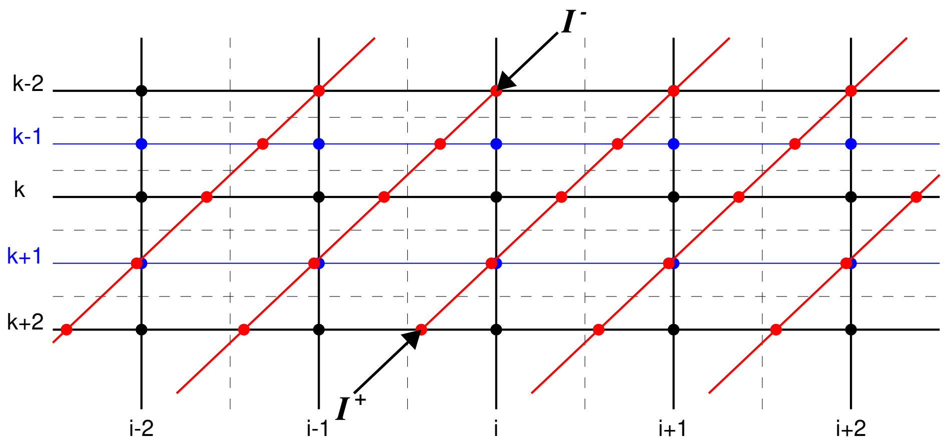

Fig. 1. Sketch of the different grids used with the long-characteristics radiative transfer method in a 2D vertical slice $(x - z)$. The hydrodynamics solver operates on a Cartesian grid (HD grid, black dots representing cell centers), while the radiative transfer equation is solved on a system of inclined rays (red dots representing the mesh points of the Feautrier scheme). For given direction, each HD grid point in the top layer coincides with the starting point of a long ray (number of rays equals number of HD cells in horizontal direction), and each ray extends from top to bottom, assuming periodic lateral boundary conditions. The HD grid can be refined in vertical direction by additional z-planes ($k-1$, $k+1$, blue) to provide sufficient sampling on strongly inclined rays.

2.1. LC scheme with Feautrier solver

The basic setup of the long-characteristics RT scheme implemented in CO5BOLD is shown in Fig. 1. In the following, we briefly summarize the main steps of the computation of Q_{rad}.

For each opacity bin n, the first step is to interpolate the source function S_n and opacity $\chi_n = \rho\kappa_n$ from the HD grid (black dots) to the mesh points of the ray system (red dots). The latter are defined as the intersection points of the inclined ray with the horizontal HD mesh lines. In this way, interpolation is restricted to the horizontal direction (x), where S_n and χ_n are assumed to be piecewise linear functions of x between two adjacent HD grid points. As indicated in Fig. 1 (blue dots), the original HD grid with vertical spacing Δz_k is refined in vertical direction by introducing additional sub-levels that globally divide the distance between all adjacent main z-levels into m equal intervals of size $\Delta z_k/m$. The refinement factor m is chosen to provide sufficient sampling of the horizontal fluctuations on strongly inclined rays, requiring the vertical resolution of the refined grid to satisfy $\Delta z_k^* = \Delta z_k/m \lesssim \mu\, \Delta z_k$ at all heights. Note that m is an integer which is independent of z but may be chosen to depend on μ. For the short-characteristics method, the refinement strategy is slightly different; here m is a function of Δz and μ, making the grid refinement more local (see Sect. 2.2.2).

The next step is the computation of the optical depth scale along each ray. For this purpose, χ_n is represented by a piecewise cubic function between adjacent mesh points along the ray. To avoid overshooting interpolation, and potentially negative optical depth increments, we restrict the slopes of the cubic function at the mesh points according to Steffen (1990) to enforce a monotonic run of the interpolating function, from which the optical depth increments $(\Delta\tau)$ along each ray are obtained by analytic integration.

Now the transport equation is solved along each ray independently, using a modified Feautrier method. Instead of solving for the

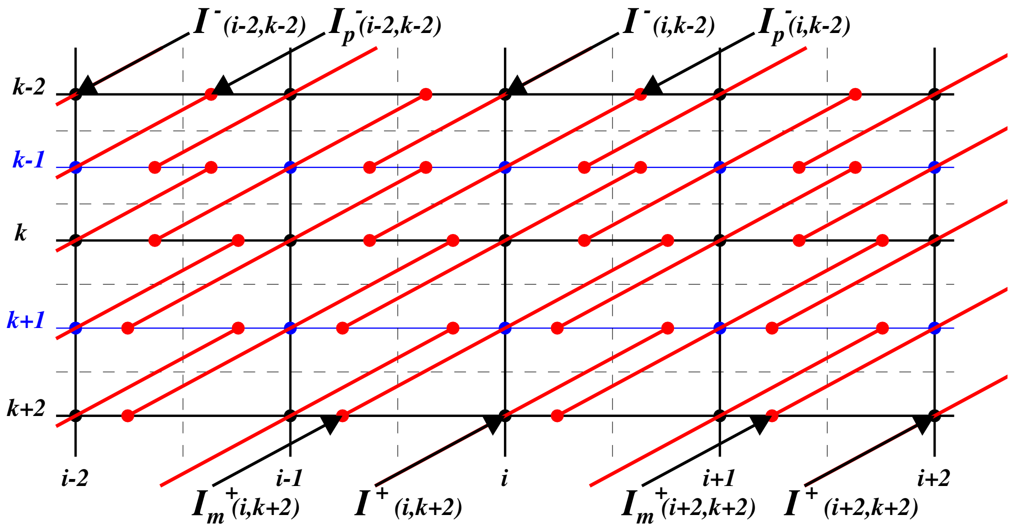

Fig. 2. Sketch of the 2D short-characteristics scheme for moderately inclined rays ($\mu = 0.67$). In this example, the HD grid (black dots representing the HD cell centers) needs to be refined in vertical direction by additional z-planes ($k-1$, $k+1$, blue) to ensure that the horizontal range of the inclined rays segments is less than the grid resolution in x direction. Each of the cell centers of the refined grid (black plus blue dots) is the mid point of a short ray that connects to the layers above and below. Red dots mark the start / end points of each short characteristics. On all ray segments, the radiative transfer is solved in both downward and upward direction. As in Fig. 1, dashed lines outline the cell boundaries of the refined grid. Note that the number of rays per unit surface area is effectively two times higher compared to the long-characteristics scheme.

unknown $u = (I^+ + I^-)/2$ (where I^+ and I^- denote the outgoing and incoming intensity, respectively), we solve the modified equation for $y = u - S$. By default (method LCF0), the second-order differential equation in y is solved on the optical depth scale (see, e.g., Mihalas 1978, p. 151 ff). However, we have recently added the alternative option (method LCF1) to solve the differential equation on the geometrical scale, as in fact the method was formulated in the original paper by Feautrier (1964). As pointed out by Hoffmann et al. (2014), the standard Feautrier solution on the optical depth scale fails to correctly describe the radiation field in the presence of large opacity variations from grid point to grid point. They demonstrate that this problem can be remedied by instead solving the transfer equation on the geometrical scale, since the spatial opacity variation can be correctly treated in this formulation only.

Both methods (LCF0 and LCF1) eventually provide the solution $y = u - S$ on all mesh points of the ray system, which, after multiplication with mean opacity $\overline{\chi}$, give the partial radiative heating rate on the ray, $q_{\rm ray} = \overline{\chi}\,(u - S)$. The definition of $\overline{\chi}$ is critical for ensuring energy conservation and depends on the method: for LCF0, $\overline{\chi}_k = (\tau_{k+1} - \tau_{k-1})/(s_{k+1} - s_{k-1})$ ($s =$ geometrical distance on the ray); for LCF1 $\overline{\chi}_k = \chi_k$.

Finally, the partial heating rates $q_{\rm ray}$ are interpolated back onto the HD grid in a conservative way, such that the horizontal averages of $q_{\rm HD}$ and $q_{\rm ray}$ are identical for all height levels k. The total radiative heating rate $Q_{\rm rad}$ is then built up by adding the individual contributions $q_{\rm HD}$ of the different ray directions (θ, ϕ) with

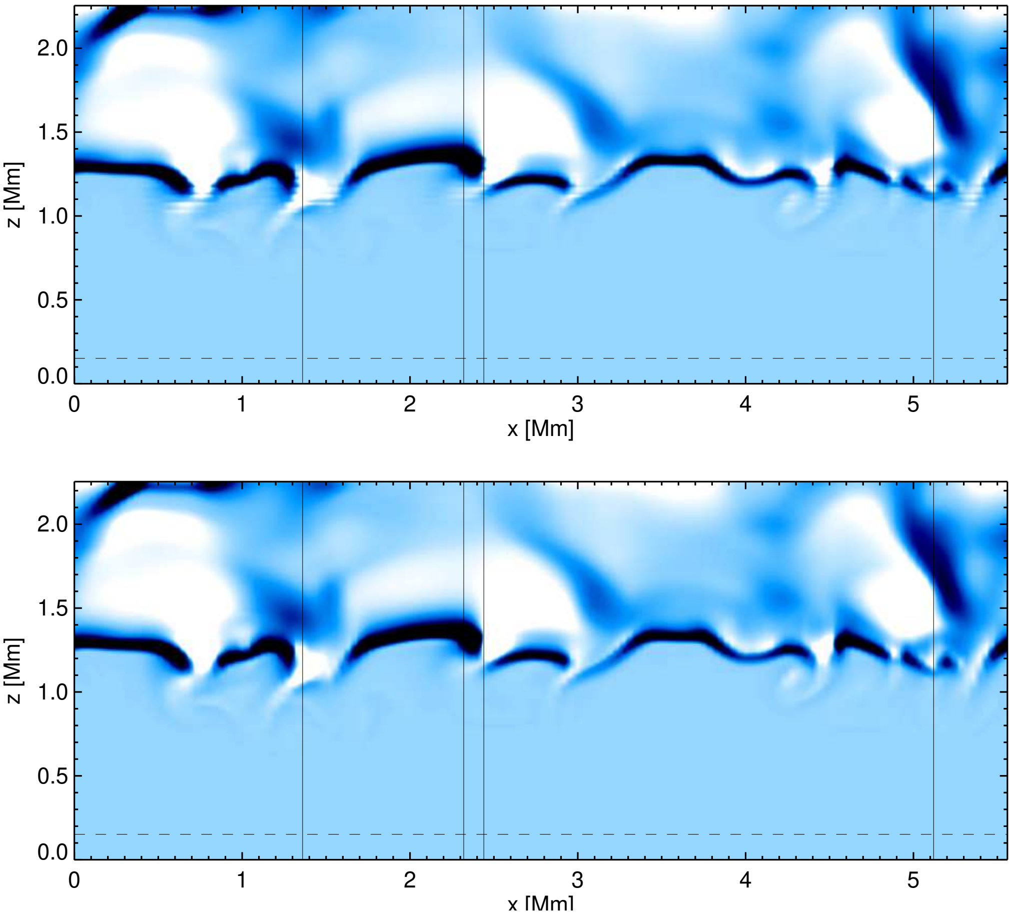

Fig. 3. Total radiative heating rate per unit mass, Q_{rad}/ρ, in a vertical slice of a 3D solar CO5BOLD model computed with 12 opacity bins and 4+1 ray inclinations μ. Dark/bright colors indicate radiative cooling/heating. This representation emphasizes the radiation field in the upper (photospheric) layers. Upper and lower panels show Q_{rad}/ρ recomputed with the LCF1 and the new SC scheme, respectively.

their appropriate integration weights, and summation over all opacity bins n. Further details may be found in Freytag et al. (2012).

In the context of NLTE3D, the quantity that is interpolated from the ray system to the HD grid is $\tilde{u}_{\mathrm{ray}} = u - \alpha\, S$ instead of q_{ray}. Note that α is constant for a given grid level k and depends on the mean optical depth $\overline{\tau}$ at the respective geometrical height, such that $\alpha \rightarrow 0$ for small $\overline{\tau}$ and $\alpha \rightarrow 1$ for large $\overline{\tau}$. Then we obtain the mean intensity at the HD grid points as $u_{\mathrm{HD}} = \tilde{u}_{\mathrm{HD}} + \alpha\, S_{\mathrm{HD}}$. The factor α was introduced to achieve a reasonable accuracy of the back interpolation for all optical depths.

2.2. SC scheme with piecewise linear source function

The general ideas for constructing the new short-characteristics radiative transfer method in 2D (SCrad) were developed independently of the existing SHORTrad concept. Different layouts of the ray system and interpolation pro-

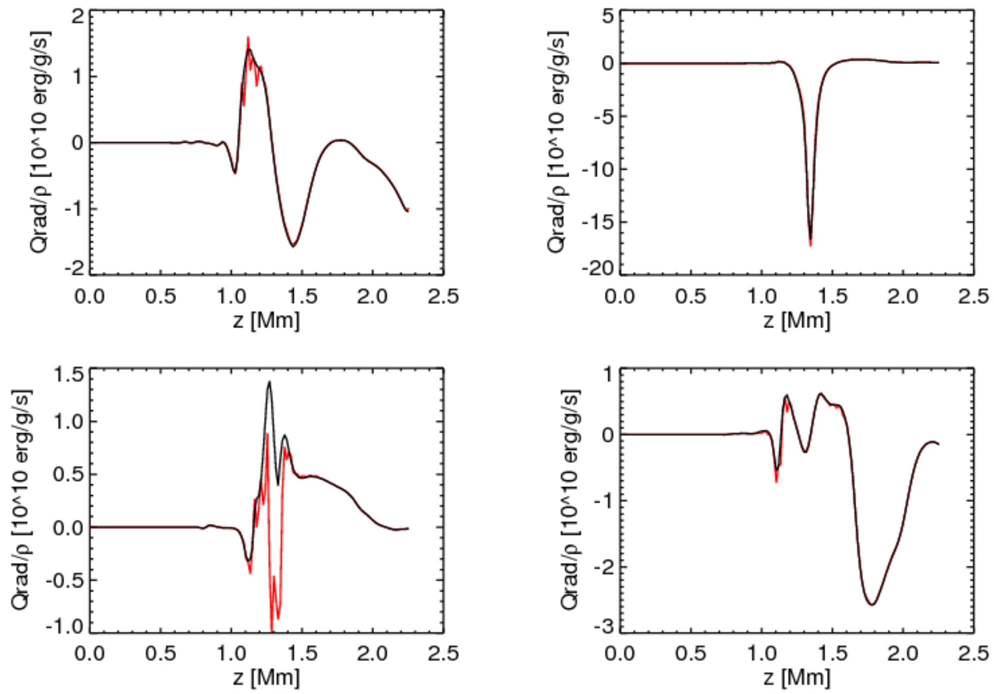

Fig. 4. Comparison of the depth dependence of $Q_{\rm rad}/\rho$ at the four horizontal positions indicated by vertical black lines in both panels of Fig. 3. Red and black lines refer to the LCF1 and SC method, respectively.

cedures have been considered. After several failures, the current design described below has passed several tests and its performance appears to be very promising.

As illustrated in Fig. 2, each HD grid point (i, k) is connected to the horizontal layer above $(k-1)$ and below $(k+1)$ by a short ray. In this way, two adjacent grid levels are connected by $2 \times n_x$ short rays, where n_x is the number of HD cells in the horizontal direction. Thus the number of rays per unit length (in x and y) is twice as large as in the LC scheme (Fig. 1).

The cell interfaces (dashed lines in Figs. 1 and 2) are defined at the middle between two mesh lines of the (refined) HD grid. In case of a non-equidistant grid (z-direction), the HD grid points (black dots) representing the cell centers of the hydrodynamics grid are not located exactly at the geometrical center of the cell. This is the usual geometry that was also adopted for the long-characteristics radiation transfer, and thus it is not necessary to change the setup of the hydro models or to re-interpolate them onto a new grid geometry when switching between the LC and SC method.

To capture the spatial variation of opacity and source function in a simple and robust way, the optical depth increments of the (inclined) short rays between levels k and $k+1$ are computed as $\Delta\tau = \Delta s\,(\chi_k + \chi_{k+1})/2$, where $\Delta s = (z_k - z_{k+1})/\mu$, and χ_k and χ_{k+1} are the opacities at the end points of the short ray. The optical distance to the interface level $(k{+}1/2)$ is assumed to be $\Delta\tau/2$. The source function is treated as a piecewise linear function of optical depth (τ) along the rays (as in SHORTrad). The source function gradient $\mathrm{d}S/\mathrm{d}\tau = (S_{k+1} - S_k)/\Delta\tau$ is constant and identical on each of the two connected short ray segments ($k \rightarrow k{+}1/2$) and ($k{+}1/2 \rightarrow k{+}1$),

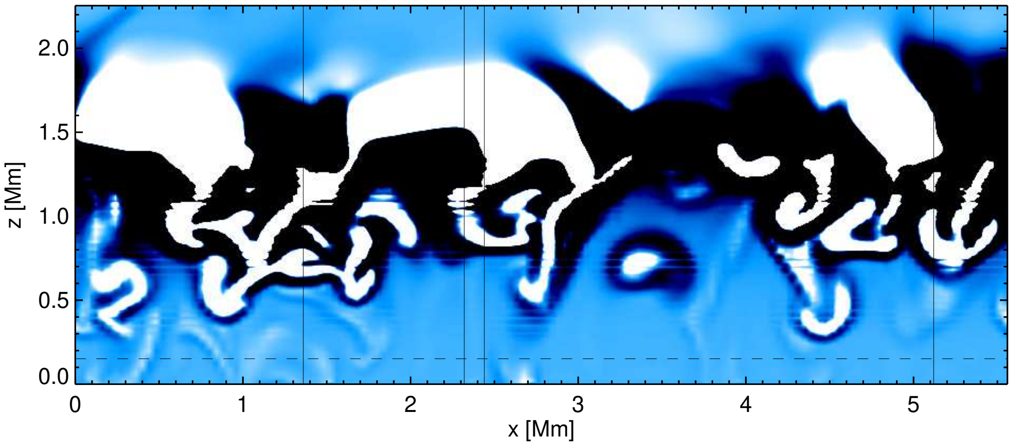

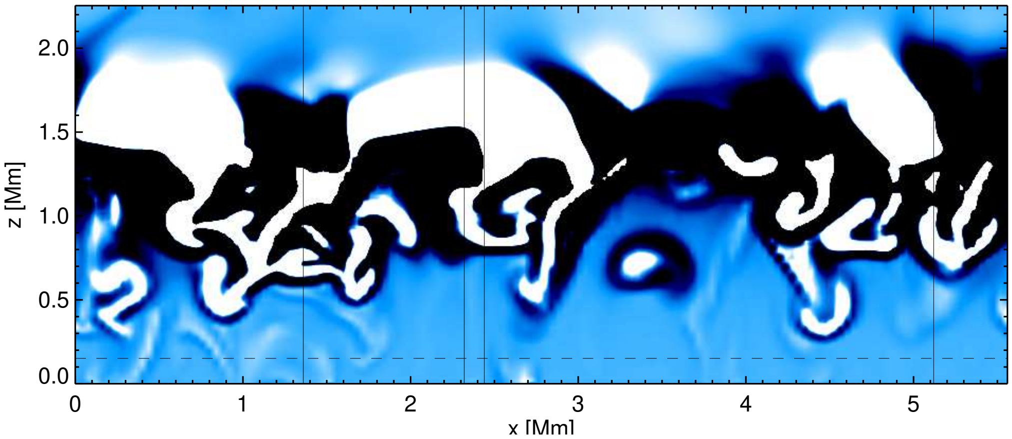

Fig. 5. Total radiative heating rate per unit volume, $Q_{\rm rad}$, in the same vertical slice as shown in Fig. 3. This representation emphasizes $Q_{\rm rad}$ in the deeper (sub-photospheric) layers (dark/bright indicate cooling/heating) Below the dashed horizontal line, $Q_{\rm rad}$ is computed from the diffusion approximation, while in the bulk of the model, it is computed from the LCF1 (top) and SC (bottom) ray systems.

respectively; $\mathrm{d}S/\mathrm{d}\tau$ is thus continuous across the cell interfaces.

On each short characteristics centered at level k, the transport equation is solved in four steps in each direction ($k{-}1 \rightarrow k{-}1/2 \rightarrow k \rightarrow k{+}1/2 \rightarrow k{+}1$ and reverse). Assuming a constant source function gradient $\mathrm{d}S/\mathrm{d}\tau$ on each ray segment, the transfer equation can be solved analytically. The formulation in terms of the quantity $y = I - S$ turns out to be most convenient, giving more compact expressions.

In a first sweep, the transfer equation is integrated from the top layer ($k = 1$) to the bottom layer ($k = k_{\rm max}$). Starting with the downward intensities at the top layer, $I^-(i, 1)$, known from the upper boundary condition (Sect. 2.2.1), and $I^-_{\rm p}(i, 1)$, obtained by interpolation, the transfer equation determines the intensities on the next lower level, $I^-_{\rm m}(i, 2)$ and $I^-(i, 2)$ (here $I^-_{\rm m}(i, k)$ and $I^-_{\rm p}(i, k)$ denote the downward intensity localized on level k at the x-position of the ray passing through HD grid point $(i, k{-}1)$ and $(i, k{+}1)$, respectively). The energy absorbed or

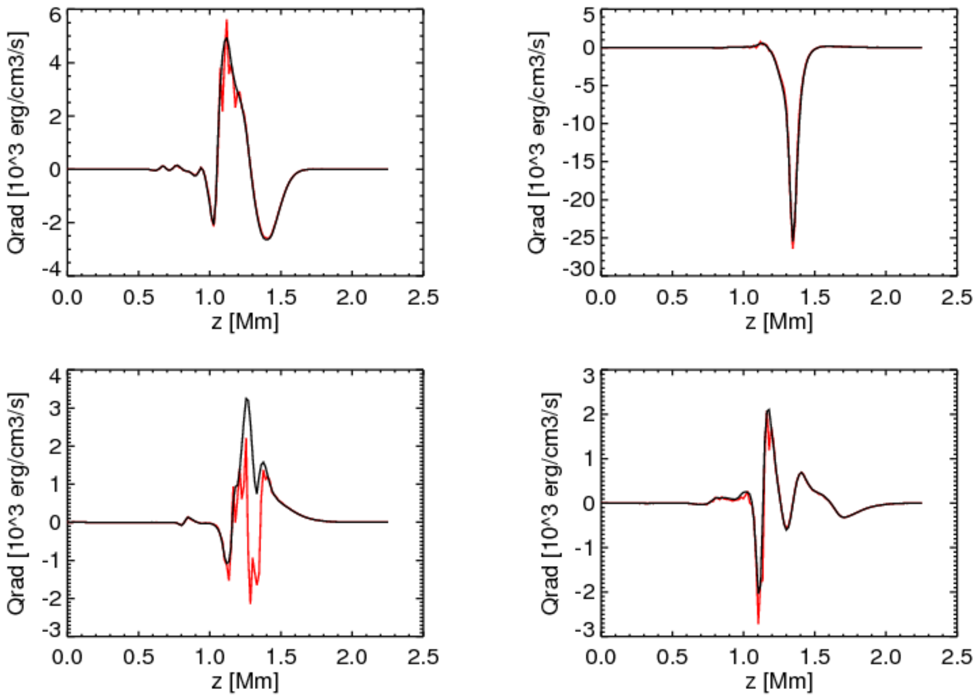

Fig. 6. Comparison of the depth dependence of $Q_{\rm rad}$ at the same four horizontal positions as in Fig. 4 (cf. vertical black lines in Fig. 5). Red and black curves refer to LCF1 and SC RT, respectively.

emitted per unit length on each short ray segment, $q^-(\theta,\phi)$, is computed on the fly and stored for later combination with the corresponding contribution from the outgoing intensity obtained in the second sweep.

To proceed from level k to level $(k{+}1)$, we need to compute intensities $I_{\rm p}^-(i,k)$. This is done by interpolation of the known intensities $I^-(i,k)$ and $I_{\rm m}^-(i,k)$ $(i{=}1\ldots n_x)$ in x-direction. In this crucial interpolation step, care must be taken to (i) account for flux conservation, (ii) avoid negative intensities, and (iii) prevent excessive diffusion in horizontal direction. Once the $I_{\rm p}^-(i,k)$ are defined, it is straightforward to integrate all short rays to level $(k{+}1)$ below.

In a second sweep, the transfer equation is integrated from bottom $(k{=}k_{\rm max})$ to top $(k{=}1)$, starting with the upward intensities $I^+(i,k_{\rm max})$ known from the lower boundary condition (Sect. 2.2.1), and the interpolated intensities $I_{\rm m}^+(i,k_{\rm max})$. In analogy with the downward sweep, the transport step gives us the intensities on next higher level, $I_{\rm p}^+(i,k_{\rm max}{-}1)$ and $I^+(i,k_{\rm max}{-}1)$. The energy absorbed or emitted on each short ray segment, $q^+(\theta,\phi)$, is combined with the corresponding contribution from the incoming intensity obtained and stored in the first sweep, $q^-(\theta,\phi)$, to evaluate the net heating rate on each individual short ray segment as $q(\theta,\phi) = q^-(\theta,\phi) + q^+(\theta,\phi)$.

In addition, the same steps are performed for the mirrored ray system $(\phi+\pi)$. After completion of the downward and upward sweep, we know the net radiative heating rates $q(\theta,\phi)$ and $q(\theta,\phi+\pi)$ on each segment of the two complementary sets of rays. The mean radiative heating rate of a HD grid cell, $Q(i,k,\theta,\phi)$, is determined by adding up the heating rates $q(\theta,\phi)$, of all ray segments crossing the respective grid cell, weighted by their length inside the cell. Some ray segments may contribute to two different HD cells. In addition, we have all information to compute the mean intensity $u(i,k,\theta,\phi) = [I^+(i,k,\theta,\phi) + I^-(i,k,\theta,\phi)]/2$ at all HD grid points.

Finally, the total radiative heating rate $Q_{\rm rad}$ is built up by adding the individual contributions $Q(\theta,\phi)$ of the different ray directions (θ,ϕ) with their appropriate integration weights, and summation over all frequency bins n. In the same way, the angle-averaged intensity J is assembled from the individual $u(\theta,\phi)$.

2.2.1. Boundary conditions

The physical assumptions behind the upper and lower boundary conditions for the short-characteristics radiation transport scheme are essentially the same as for the long-characteristics method (see also Freytag 2017, this volume).

At the *upper boundary*, we first estimate the optical depth from the local opacity, χ_1, as $\tau_1 = \chi_1 H_\tau$, where H_τ is the effective scale height of the opacity $\chi = \rho\,\kappa$ at the top layer of the model, assumed to be identical for all opacity bins. Besides H_τ (or the scaling factor with respect to the local pressure scale height, $f = H_\tau/H_p$), we also specify a boundary temperature $T_0 = T(\tau = 0)$, from which we can compute the source function $S_{n,0} = S_n(\tau = 0)$ for each opacity bin n. We may further specify an external irradiation with intensity $I^-_{n,0}(\mu) = I^-_n(\tau = 0, \mu)$ or $y^-_{n,0}(\mu) = I^-_{n,0}(\mu) - S_{n,0}$. Assuming a linear dependence of the source function on optical depth between $\tau = 0$ and τ_1, we obtain the incident radiation at the upper boundary $y^-_{n,1}(\mu) = I^-_n(\tau_1,\mu) - S_n(\tau_1)$ as

$$y^-_{n,1}(\mu) = y^-_{n,0}(\mu)\, e^{-\tau_1/\mu} - \\ - \frac{S_{n,1} - S_{n,0}}{\tau_1/\mu} \left(1 - e^{-\tau_1/\mu}\right).$$

The *lower boundary* ($z = z_{\rm b}$) is assumed to be located at large optical depth such that the diffusion approximation is valid. There are different options.

(1) In case of an *open* lower boundary, we choose the 'neutral' radiative boundary condition $\nabla \cdot \boldsymbol{F}_{\rm rad} = 0$. Here the radiative flux through the boundary is determined by the local temperature gradient. The boundary condition is realized by setting $y^+_n(z_{\rm b},\mu) = -y^-_n(z_{\rm b},\mu)$ for all inclinations μ and opacity bins n, which is equivalent to $[I^+_n(z_{\rm b},\mu) + I^-_n(z_{\rm b},\mu)]/2 = S_n(z_{\rm b})$ or $J_n = S_n$ for all grid cells of the bottom layer.

(2) For a *closed* bottom, we can (a) enforce a given vertical radiative flux, F^*_z, through the lower boundary by setting $y^+_n(z_{\rm b},\mu) = y^-_n(z_{\rm b},\mu) + (2\pi/3)^{-1} F^*_z\, w_n\, \mu$, where the weight w_n depends on the opacity bin and is given by $w_n = (\mathrm{d}S_n/\mathrm{d}T)/(\mathrm{d}S/\mathrm{d}T)\,(\bar\kappa_{\rm Ross}/\bar\kappa_{{\rm Ross},n})$, satisfying $\sum_n w_n = 1$. Alternatively, we can (b) specify the radiative flux to be transported through the cell interface below the lower boundary. This is achieved by setting $y^+_n(z_{\rm b},\mu) = (4\pi/3)^{-1} F^*_z\, w_n\, \mu$.

(3) If the radiation field in lower part of the computational domain is computed in the diffusion approximation, boundary conditions 1 and 2b are currently available. At the lower boundary of the ray system, at $z = z_{\rm d}$, the required incident intensities are obtained from the diffusive fluxes $F^{\rm D}_{n,x}(z_{\rm d})$ and $F^{\rm D}_{n,z}(z_{\rm d} - \Delta z/2)$ as $y^+_n(z_{\rm d},\mu) = (4\pi/3)^{-1} \left(F^{\rm D}_{n,x}\, \sqrt{1-\mu^2} + F^{\rm D}_{n,z}\, \mu\right)$ such that the energy transport across the interface between diffusion regime and ray system is smooth and conservative.

2.2.2. Grid refinement

The scheme described above is restricted to inclinations θ satisfying the condition

$$\tan\theta \le \frac{\Delta x}{z(k) - z(k+1)}\,, \;\; k = 1 \ldots k_{\rm max} - 1\,. \quad (1)$$

For larger θ, the short rays would extend horizontally by more than $\pm\Delta x$, which would lead to under-sampling problems and would undermine the close coupling between short rays and individual HD grid cells.

The natural procedure to deal with strongly inclined ray systems would be to rotate the

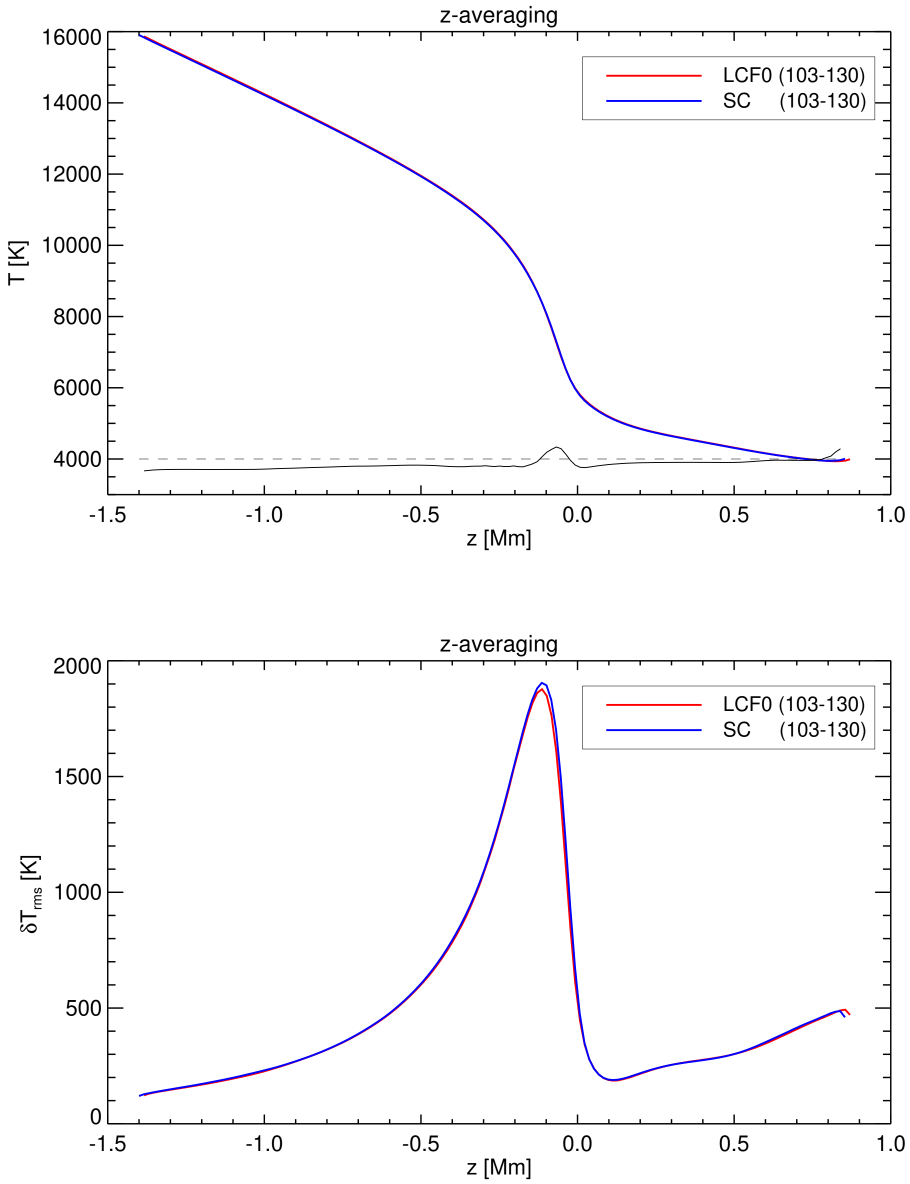

Fig. 7. Comparison of the mean temperature stratification resulting from the LCF0 (red) and SC (blue) based simulations, respectively, as a function of geometrical height z. Top: $\overline{T}(z) = \langle T(x,y,z)\rangle_{x,y,t}$, where $\langle\cdot\rangle_{x,y,t}$ denotes horizontal averaging over planes of constant geometrical height and over 140 snapshots. The lower black curve shows the tenfold temperature difference $10 \times (\overline{T}(\mathrm{SC}) - \overline{T}(\mathrm{LCF0})) + 4000\,\mathrm{K}$. Bottom: Amplitude of horizontal temperature fluctuations $\delta T_{\mathrm{rms}}(z) = \left[\langle\left(T(x,y,z) - \overline{T}(z)\right)^2\rangle_{x,y,t}\right]^{1/2}$.

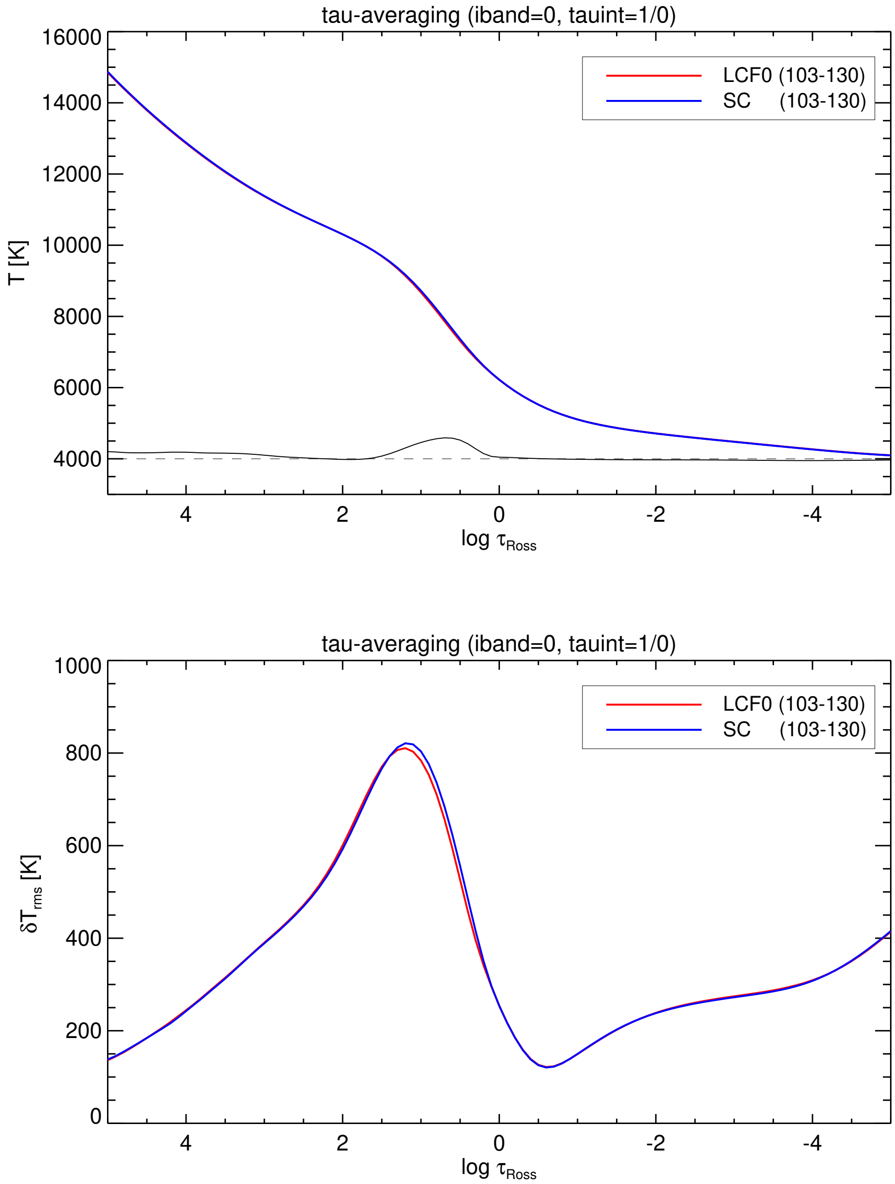

Fig. 8. Comparison of the mean temperature stratification resulting from the LCF0 (red) and SC (blue) based simulations, respectively, as a function of Rosseland optical depth $\tau_{\rm Ross}$. Top: $\overline{T}(\tau_{\rm Ross}) = \langle T(x,y,\tau_{\rm Ross})\rangle_{x,y,t}$, where $\langle\cdot\rangle_{x,y,t}$ denotes horizontal averaging over surfaces of constant $\tau_{\rm Ross}$ and over 140 snapshots. The lower black curve shows $10\times(\overline{T}(\mathrm{SC})-\overline{T}(\mathrm{LCF0}))+4000\,\mathrm{K}$. Bottom: Amplitude of horizontal temperature fluctuations on iso-$\tau_{\rm Ross}$ surfaces, $\delta T_{\rm rms}(\tau_{\rm Ross}) = \left[\langle\left(T(x,y,\tau_{\rm Ross})-\overline{T}(\tau_{\rm Ross})\right)^2\rangle_{x,y,t}\right]^{1/2}$.

whole grid by 90°, and to perform the same steps outlined above on the rotated grid, as it is in fact done in SHORTrad, in accordance with the original ideas of the SC method (see, e.g., Kunasz & Auer 1988; Kunasz & Olson 1988, and references therein). While this procedure works fine in isotropic global models, there are several problems with this approach when applied to the local box models where the boundary conditions in vertical and horizontal direction are fundamentally different. This is the main reason why different radiation transport modules exist in CO5BOLD for global and local geometry.

In the local models, the HD grid is allowed to have a variable step size in the vertical direction, so the rotated grid would have a variable step size in the horizontal direction. This would lead to additional complications for which the scheme described above is not prepared. At the same time, the original upper and lower boundaries would become the lateral boundaries of the rotated grid, violating the basic assumption of periodic lateral boundaries. Even worse, the lateral boundaries of the original grid play the role of the upper and lower boundaries of the rotated grid. This gives rise to the problem that the values of the incident intensities are not explicitly known at the periodic upper and lower boundaries but would need to be determined iteratively (e.g., Hayek et al. 2010) or by some recursive procedure (Steiner 1990). This is a serious disadvantage.

To circumvent such problems, we adopted a different strategy to deal with strongly inclined ray systems: as in the case of the LC transport scheme, we introduce a refinement of the grid in the vertical direction, such that the aspect ratio of the cells of the refined grid are large enough to satisfy the inclination condition (Eq. 1). Additional horizontal planes are introduced only locally where necessary (see extra z-levels shown in blue in Fig. 2). After interpolation of the necessary thermodynamic quantities to the extra z-levels by monotonic cubic polynomials, our SC scheme can be applied to the refined grid without any modification, providing the radiation field at the cell centers of the refined grid. To obtain the radiation field on the original HD grid, the result on the refined grid needs to be remapped (in vertical direction only) on the original grid in a conservative way. This is achieved by adding up the contributions of all refined grid cells overlapping with the original HD cell, weighted by their fractional geometrical volume.

2.2.3. Exponential terms

The solution of the transfer equation requires the evaluation of terms involving the exponential function. In our present implementation of the SC method, the direct use of the exponential function is avoided. Instead, the exponential term $f_{\rm exp} = 1 - \exp(-x)$ is approximated by the expression

$$f_{\rm exp}^{(5)} = \frac{x + x^2/2 + x^3/6 + x^4/24 + x^5/120}{1 + x + x^2/2 + x^3/6 + x^4/24 + x^5/120}$$

over the full range of $0 \le x \le \infty$. This 5th order approximation has a maximum relative error of 0.52%. Exactly the same error is valid for the term $g_{\rm exp} = f_{\rm exp}/x$:

$$g_{\rm exp}^{(5)} = \frac{1 + x/2 + x^2/6 + x^3/24 + x^4/120}{1 + x + x^2/2 + x^3/6 + x^4/24 + x^5/120}\,.$$

Similar approximations are implicitly inherent in the Feautrier method. The advantage of these expressions is that they are accurate for arbitrarily small values of x (optical depth increments). In practice, however, the direct evaluation of $f_{\rm exp}$ and $g_{\rm exp}$ using the exponential function (in double precision) might work as well, and would be preferable if computationally faster.

3. First test results

After developing the SC radiation transport module in an IDL environment, the routines were ported to Fortran90, debugged, optimized, and implemented in the CO5BOLD code. This allows us to compare the results of the three different RT schemes, LCF0, LCF1, and SC, under realistic conditions.

Table 1. Emergent flux and intensity contrasts computed on the same single snapshot with methods LCF0, LCF1, and SC.

	LCF0	LCF1	SC
$F_{\rm rad}/F_{\odot}$	1.05623	1.05475	1.04453
$\delta I_{\rm b}$ [1)]	0.15441	0.15355	0.14618
$\delta I_{\rm c}$ [2)]	0.22792	0.22698	0.21587

[1)] bolometric intensity contrast; [2)] intensity contrast in opacity bin#3 ("blue continuum").

3.1. Comparison of $Q_{\rm rad}$, $F_{\rm rad}$, and $\delta I_{\rm c}$ for a given snapshot

We compared the radiative heating rates, $Q_{\rm rad}$, emergent fluxes, $F_{\rm rad}$, and two intensity contrasts ($\delta I_{\rm b}$ and $\delta I_{\rm c}$, defined below) obtained from the different RT methods for an arbitrary snapshot taken from a CO5BOLD 3D simulation of the solar surface layers. The simulation was run with method LCF1 with five μ-angles (including the vertical) in 12 opacity bins. For comparison, the radiation transport was recomputed on the selected snapshot with the standard method LCF0 and with the new SC method, using the same μ-angles and opacity bins.

Figures 3 – 6 show the comparison of $Q_{\rm rad}$ between methods LCF1 and SC in a 2D vertical slice of the 3D snapshot. It is obvious that the two methods produce very similar results at all optical depths, which must not be taken for granted since the two methods are based on completely different approaches. The only obvious difference is that the SC method produces a slightly smoother distribution of $Q_{\rm rad}$ and shows significantly less artifacts. It is also remarkable that the SC method gives the correct asymptotic behavior at large optical depths, nicely reproducing the diffusion approximation (see features crossing the dashed horizontal line in Fig. 5). The $Q_{\rm rad}$ differences between methods LCF0 and LCF1 are found to be negligible at this level of comparison.

Table 1 shows the comparison of the emergent bolometric fluxes (normalized to the solar flux $F_{\odot} = \sigma T_{\rm eff}^4$, $T_{\rm eff} = 5770$ K), and intensity contrasts, $\delta I_{\rm b}$ and $\delta I_{\rm c}$, for methods LCF0, LCF1, and SC. Here $\delta I_{\rm b}$ refers to the contrast computed from the emergent bolometric intensities ($I_{\rm b} = \sum I_n$ over all opacity bins n) of the vertical rays as

$$\delta I_{\rm b} = \left[\left\langle \left(\frac{(I_{\rm b}(x,y)}{\langle I_{\rm b}(x,y)\rangle} - 1 \right)^2 \right\rangle \right]^{1/2} , \qquad (2)$$

where $\langle\cdot\rangle$ denotes horizontal averaging. The contrast in the "blue continuum" is obtained by replacing $I_{\rm b}$ by I_3, the emergent intensity in opacity bin#3. While differences between LCF0 and LCF1 are minor, method SC yields slightly but significantly lower values of radiative surface flux and intensity contrast by roughly 1% with respect to LCF.

3.2. Comparison of full 3D solar atmosphere simulations

As the next important test, we performed full 3D CO5BOLD simulations of the solar atmosphere with the three different RT methods (LCF0, LCF1, SC). A preliminary run with the new SC RT module showed that the inflow entropy had to be increased slightly with respect to the standard LCF setup to ensure that the resulting model has the same effective temperature (see Table 2).

After a short relaxation phase, all three runs were continued for 7 hours of simulated time, saving full 3D snapshots every 180 s. Table 2 summarizes the time-averaged radiative output properties of the three models and lists the adopted inflow entropy. By design, all model produce almost identical mean radiative surface fluxes (effective temperatures). In agreement with the previous test on a frozen snapshot, we find again only very minor differences in the intensity contrast between LCF0 and LCF1, but a slightly lower value in the SC simulation. Bolometric (sum over all 12 opacity bins) and monochromatic (bin#3 = "blue continuum") intensity contrasts show the same weak dependence on the employed RT method.

In Fig. 7 (top), we compare the mean temperature stratification, $\overline{T}(z)$, resulting from the LCF0 and SC based simulations, respectively, on the geometrical scale. $\overline{T}(z)$, is ob-

Table 2. Mean emergent radiative flux, mean intensity contrasts, and inflow entropy, for three full 3D CO5BOLD simulations of the solar atmosphere, employing RT methods LCF0, LCF1, and SC, respectively. The radiative quantities are averaged over a simulation time interval of 7 hours.

	LCF0	LCF1	SC
$\overline{F}_{\rm rad}/F_\odot$	1.0007	0.9993	0.9978
$\overline{\delta I}_{\rm b}$ [1)	0.1452	0.1451	0.1395
$\overline{\delta I}_{\rm c}$ [2)	0.2224	0.2227	0.2155
$s_{\rm inflow}$ [3)	1.7774	1.7774	1.7810

[1)] bolometric intensity contrast; [2)] intensity contrast in opacity bin#3 ("blue continuum"); [3)] in units of [10^9 erg/g/K].

tained by averaging temperature on horizontal planes and over time (140 snapshots covering 7 hours). The temperature differences are very small, nowhere exceeding 50 K. Likewise, the amplitude of the horizontal temperature fluctuations (Fig. 7, bottom) is almost identical for the two simulations at all heights.

Slightly larger differences are found when comparing the SC and LCF0 models on the Rosseland mean optical depth scale (Fig. 8). Here $\overline{T}(\tau_{\rm Ross})$ is obtained by horizontal averaging of temperature on iso-$\tau_{\rm Ross}$ surfaces and over time. To be consistent with the τ-integration used in the respective simulations, $\tau_{\rm Ross}$ was computed from cubic (LCF0) and linear (SC) $\chi(z)$ interpolation, respectively,

It can be seen that the sub-photospheric temperature gradient is slightly steeper in the SC based model, leading to a maximum temperature difference of $\approx$ 70 K near $\log\tau_{\rm Ross}$=0.75. In the photosphere, the amplitude of the horizontal temperature fluctuations on iso-$\tau_{\rm Ross}$ surfaces is essentially in perfect agreement, while in the sub-photospheric layers, the SC models shows slightly higher fluctuation amplitudes. One might therefore expect the SC model to exhibit a higher intensity contrast than the LCF0 model, but we actually find the opposite (see Table 2). This is because, for a given temperature field, SC gives a lower contrast than LCF (see Table 1), and this effect dominates over the opposite effect expected from the slightly enhanced fluctuation amplitude in the continuum forming layers of the model computed with the SC RT method.

Finally, we note that the differences between the mean LCF0 and LCF1 solar models and their fluctuation amplitudes may be considered as negligible for all practical purposes.

4. Conclusions

We have developed and implemented in the CO5BOLD code a new radiative transfer module based on a short-characteristics approach with piecewise linear source function. One of the essential features of the new SC RT scheme is its built-in strict energy conservation, ensuring consistency between volume-integrated heating rates and radiative surface fluxes. While this is also true for our LC Feautrier implementation, achieving strict energy conservation proved to be more challenging with the SC approach.

First tests on a frozen snapshot from a realistic 3D solar atmosphere simulation show that the new SC method produces remarkably similar radiation fields, at all optical depths, as the standard long-characteristics Feautrier solver used so far, and correctly reproduces the diffusion limit at large optical depths.

To study the effects of the numerical details of the RT treatment on the thermal structure of the resulting 3D model atmosphere, we have run full 3D CO5BOLD solar atmosphere simulations with the three RT methods currently available for the local box setup (LCF0, LCF1, and SC). Starting from the same initial model, all three (12 bin, low viscosity) simulations are continued for 7 hours after an initial short relaxation phase. Comparison of the mean thermal structures, obtained by horizontal and temporal averaging, indicates that the SC radiation transport scheme closely reproduces the mean structure and fluctuation amplitudes produced by the standard LCF method, both on the geometrical and on the $\tau_{\rm ROSS}$ optical depth scale. This is a very reassuring result, corroborating the validity of the CO5BOLD models computed previously with the LCF radiation transport module. In terms of computing time, the new SC module is slightly slower than LCF1, but somewhat faster that LCF0. In view of the fact that the number of rays is effectively twice as large in SC compared to LCF, this implies a clear improvement in performance. Also, it seems that the OpenMP scaling properties of the new RT code are significantly better (to be confirmed with systematic studies using more than 8 cores).

While the first results of solar atmosphere simulations reported here look very promising, further tests are still necessary before the SC RT module can be released for general usage by the CO5BOLD community. One of the next steps is to examine the performance of the SC radiative transfer solver in the context of NLTE3D. The ultimate test, however, is the application of the new SC method to more challenging conditions, encountered, for example, in simulations of near-surface magnetoconvection in solar-type stars or the violent convection in the atmospheres of red giants.

Acknowledgements. I am grateful to B. Freytag, H.-G. Ludwig, and O. Steiner for many helpful discussions on the topic of radiative transfer in general and in CO5BOLD in particular, as well as for useful suggestions for improving the present manuscript.

References

Davis P.J. & Polonsky I. 1972, in Handbook of Mathematical Functions, Abramowitz, M., & Stegun, I. A. eds. (Dover Publ., New York), 888
Feautrier, P. 1964, Comptes Rendus Academie des Sciences Paris, 258, 3189
Freytag, B., Steffen, M., & Dorch, B. 2002, Astron. Nachr., 323, 213
Freytag, B., Steffen, M., Ludwig, H.-G., et al. 2012, J. Comp. Phys., 231, 919
Hayek, W., Asplund, M., Carlsson, M., et al. 2010, A&A, 517, A49
Hoffmann, T. L., Sauer, D. N., Pauldrach, A. W. A., & Hultzsch, P. J. N. 2014, A&A, 569, A62
Kunasz, P., & Auer, L. H. 1988, J. Quant. Spec. Radiat. Transf., 39, 67
Kunasz, P. B., & Olson, G. L. 1988, J. Quant. Spec. Radiat. Transf., 39, 1
Mihalas, D. 1978, Stellar Atmospheres (W. H. Freeman and Co., San Francisco)
Prakapavičius, D., Steffen, M., Kučinskas, A., et al. 2013, MSAIS, 24, 111
Steffen, M. 1990, A&A, 239, 443
Steffen, M., Cayrel, R., Caffau, E., et al. 2012, MSAIS, 22, 152
Steiner, O. 1990, A&A, 231, 278