Mem. S.A.It. Suppl. Vol. 24, 53

Memorie della
Supplementi

Opacities in CO^{5}BOLD

H.-G. Ludwig[1] and M. Steffen[2]

[1] Zentrum für Astronomie der Universität Heidelberg, Landessternwarte, Königstuhl 12, D-69117 Heidelberg, Germany; e-mail: HLudwig@lsw.uni-heidelberg.de
[2] Leibniz-Institut für Astrophysik Potsdam, An der Sternwarte 16, D-14482 Potsdam, Germany; e-mail: MSteffen@aip.de

Abstract. We describe the present treatment of the frequency-dependence of the radiative transfer in CO^{5}BOLD. This mostly refers to the way opacities are *binned* into groups. We discuss the basic ideas behind the Opacity Binning Method, give some details of the practical implementation in CO^{5}BOLD, and point to some issues where improvements are needed.

Key words. radiative transfer – methods: numerical – stars: atmospheres – stars: late-type

1. Introduction

One distinguishing feature of CO^{5}BOLD in comparison to other hydro codes is the detailed representation of the radiation field making CO^{5}BOLD a radiation-hydrodynamics code suitable for constructing stellar atmosphere models. The dependence of the radiation field on flow geometry and frequency (or equivalently wavelength) is taken into account. The treatment of the frequency-dependence in CO^{5}BOLD is the subject of this contribution. We shall discuss its present implementation, and point out a number of issues which need to be addressed in the near future. We start by describing the basics of how the frequency-dependence is handled in CO^{5}BOLD. We mostly follow the description given by Freytag et al. (2012).

Send offprint requests to: H.-G. Ludwig

2. Opacity Binning Method (OBM) – basic idea

The emission of radiation in a stellar atmosphere happens (largely) thermally, and the shape of the Kirchhoff-Planck function controls the spectral range in which a significant amount of radiation flows and energy is exchanged. Different from the smooth Kirchhoff-Planck function the radiative energy exchange is highly variable. The absorption coefficient strongly varies with frequency due to the presence of spectral lines, on top of the more gradual change of the continuous opacity. In cool stars like the Sun, spectral lines count in the millions so that an exact treatment of the frequency dependence in a complex multi-dimensional geometry is beyond present computer capacities, and one has to resort to an approximate treatment. An important simplification stems from the fact that one is not interested in the detailed frequency dependence of the heat exchange between stellar plasma and radiation field but only in its frequency-

integrated total effect $Q_{\rm rad}$. This motivates the main trick, namely to interchange the solution of the radiative transfer (which has potentially to be done for a large number of frequency points) with the integration over frequency. This amounts to operating on frequency integrals of the source function and suitable frequency-averages of the opacities to obtain frequency-integrated intensities directly, which can then be used to calculate the radiative heating (or cooling) rates. As we shall see this is only possible approximately. To this end, at present all multi-dimensional hydrodynamical stellar atmosphere codes employ the so-called *Opacity Binning Method (OBM)*. The method was first laid out by Nordlund (1982), and later refined in works by Ludwig (1992), Ludwig et al. (1994), and Vögler (2004).

The basic idea of opacity binning is the classification of frequency points by the similarity of their associated Λ-operator – the operator relating source function S_ν and mean intensity J_ν of the radiation field:

$$J_\nu = \Lambda_\nu [S_\nu] . \quad (1)$$

We added an index ν to the Λ-operator to emphasize that its form, written in geometrical coordinates, is different for different frequencies due to opacity variations. However, in cases where the operator happens to be similar, its linearity allows to operate on the sum of the source functions to obtain the integrated mean intensity, symbolically expressed as

$$\begin{aligned} J_{1+2} &\equiv J_1 + J_2 \\ &= \Lambda_1 [S_1] + \Lambda_2 [S_2] \approx \Lambda [S_1 + S_2] , \end{aligned} \quad (2)$$

where Λ is some suitable mean of Λ_1 and Λ_2. The problem now is to classify all frequencies into distinct sets Ω_i grouping together as similar as possible Λ-operators. The Λ-operator can be calculated from the monochromatic optical depth scale τ_ν so that the classification can be equivalently done by combining frequencies with a similar relation between geometrical and optical depth scales. This is in fact the way how one proceeds in practice.

In recent years, the treatment of scattering within the OBM attracted some attention due to the significant influence of scattering on the temperature structure in metal-poor stars, here particularly in giants. Early on, Skartlien (2000) implemented a treatment of scattering within the OBM in the solar context for studying the effects on the chromospheric temperature structure. Targeting metal-poor atmospheres Hayek et al. (2010) implemented a treatment of scattering including efficient parallelization. We come back to the issue of scattering in Section 7.2.

3. OBM – constructing the bands

When trying to classify the frequencies, one is confronted with the problem that the optical depth scales depend on the atmospheric model under consideration, i.e. its geometry, the ensuing thermal conditions, and velocities. One has to choose a reference model for which the classification is performed. Naturally, this reference model is chosen to be close to the stellar atmosphere to be simulated, in the simplest case a one-dimensional model of the atmosphere in question. Other choices are possible, however, in any case the resulting classification is optimized for a particular set of atmospheric parameters and has to be repeated when numerical simulations in other parameter regimes are conducted. Since even for fixed atmospheric parameters a large variety of different thermodynamic conditions are met along various lines-of-sight in a 3D numerical model (with correspondingly different τ_ν), limits to the achievable accuracy by the opacity binning have to be expected. Thus only a reasonable similarity among τ_ν-scales within a class (often also referred to as opacity group) is aimed at in practice. Typically one is content if the τ_ν-scales of a group share the property to reach unity within a given range of depth – usually defined via the frequency-independent Rosseland optical depth. This emphasizes the emergent radiation intensity as the primary quantity to be captured correctly, obviously an important quantity linked to the overall flux properties of a stellar atmosphere.

Figure 1 illustrates schematically how the classification is performed when applying the $\tau_\nu = 1$ criterion only. It is possible to add further criteria. For instance, it turned out that it

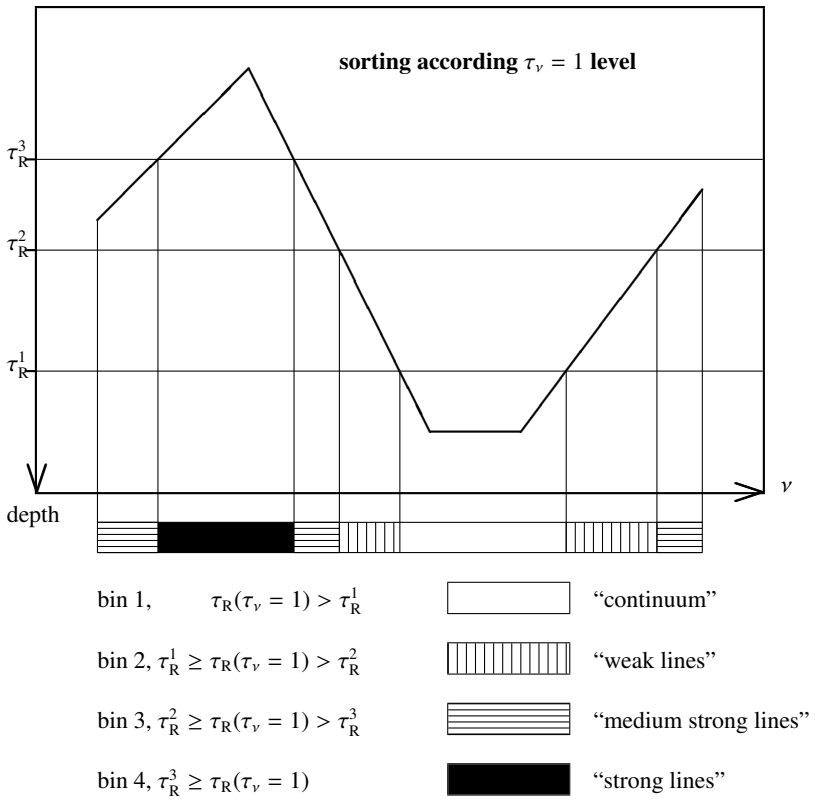

Fig. 1. Schematic illustration of the classification of frequency points in the OBM. In the case shown, 4 bins are constructed representing the continuum and increasingly stronger spectral lines. Note, that in general the bins do not form contiguous frequency bands (from Ludwig 1992).

is sometimes advantageous to split bins further according to frequency. In the 12-bin setup commonly used in CO5BOLD models for the Sun the three bins closest to the continuum are split into two sub-bins each.

3.1. Testing opacity tables

The performance of the opacity tables is commonly checked by comparing the radiative flux and radiative heating rates with the result of the calculation when using the full frequency resolution. In CO5BOLD, we take the 1D structure employed during classification also as reference structure for which the testing is done. The Figs. 2 and 3 show the situation for an OBM scheme assuming five bins. In the shown case, the OBM leads to a dramatic improvement of the radiative flux and radiative heating rate relative to a simple grey treatment. However, the case is particularly benign, and often one needs more bins, or has to be content with a larger deviation from the result at full frequency resolution. Moreover, it is sometimes difficult to obtain an improvement of the situation by merely increasing the number of bins. This is related to the fact that the OBM does not converge to the exact solution in the limit of an infinite number of bins. This points to the need of an improved scheme if one wants

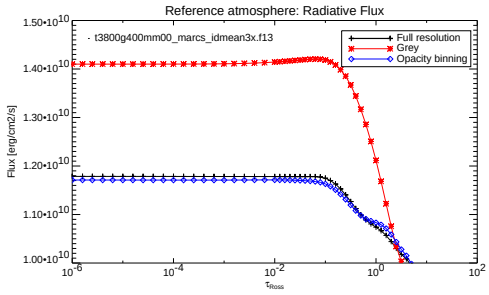

Fig. 2. Total radiative flux as a function of optical depth employing high frequency resolution (black), opacities based on a 5-bin OBM (blue), and a grey Rosseland opacity (red). The five bins used in the OBM lead to a significant improvement of the flux profile relative to the grey case.

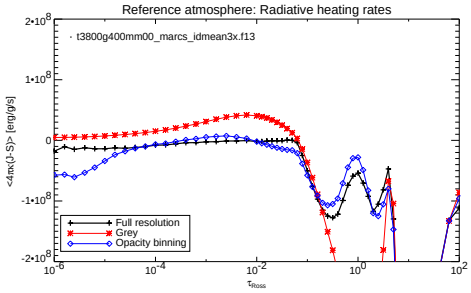

Fig. 3. As Fig. 2, but for the radiative heating rates per unit mass.

to reach significantly higher accuracies than achieved at present.

When trying to interpret Figs. 2 and 3 one problem becomes apparent: it is not immediately clear how significant deviations in radiative flux or heating rate actually are. One would be most interested in the associated temperature errors. In the simplest case, one might ask by how much the radiative equilibrium temperature would change given a certain deviation in flux between OBM and a calculation at full frequency resolution. Unfortunately, corresponding diagnostics have not been implemented into the binning procedure yet. What is possible is to compare equilibrium models based on the OBM with other models calculated with a detailed treatment of the frequency-dependence (e.g., comparing against ATLAS9 models). However, such an approach has the disadvantage that differences in the numerical treatment and microphysics enter the comparison.

4. Contents of the opacity tables

The binned opacities are obtained from a suitable average of the opacities grouped into a particular bin and stored in look-up tables as a function of thermodynamic variables – in CO5BOLD as a function of gas pressure and temperature. In addition, the Kirchhoff-Planck function (as source function), integrated over the frequencies of the bin, are stored as a function of temperature. This approach only works if the opacities and the source function can be calculated from the thermodynamic conditions alone, i.e. are thermodynamic equilibrium quantities. While this is often fulfilled to good approximation, there are exceptions. For instance, the formation of dust clouds in cool stellar atmospheres is a non-equilibrium process, and actual particle properties are only known after solving the governing kinetic equations, taking into account the history of the evolution of a particular mass element in the flow. In CO5BOLD, we proceed by separating the equilibrium part (gas opacities) from the non-equilibrium part (dust opacities). The gas opacities are binned into groups in the usual way, and the dust opacities are calculated during the simulation "on-the-fly" and added to the gas opacities. Obviously, this increases the computational demands.

The OBM relies on "raw" opacities provided as extensive tables of monochromatic opacities as a function of gas pressure and temperature. Currently, three different sources are used. For K-type and hotter, non-degenerate stars mostly opacities from the MARCS package of the Montpellier-Uppsala group (Gustafsson et al. 2008) are used. For cool M-type and later stars opacities from the PHOENIX stellar atmosphere code are provided by the Lyon group (Allard et al. 2012). The opacities for the CO5BOLD models of DA

white dwarfs come from the Montreal atmosphere code (Tremblay & Bergeron 2009).

4.1. Opacity table format

The binned opacities are stored in ASCII format. Since the number of opacity bins is not large the data files are compact, not exceeding 2 Mb in size. The file structure follows a scheme where meta-data describing the contents are tagging the numerical data blocks. Comments are also added, in particular a header summarizing the file contents which is usually copied into the CO5BOLD "out" files. For all N_Ω opacity bins, the band averaged opacities $\log\chi_i(\log P, \log T)$, $i \in \{1, 2, \ldots, N_\Omega\}$ are stored as a function of the logarithmic pressure and temperature. In addition, the globally averaged, grey Rosseland average $\log\chi_0(\log P, \log T)$ is stored. These Rosseland opacities are used in grey CO5BOLD models, or for the calculation of average model structures on τ-levels.

The band integrated source function S_i is stored as a function of $\log T$, normalized to the frequency-integrated Kirchhoff-Planck function as $w_i = S_i/\frac{\sigma}{\pi}T^4(\log T)$, $i \in \{0, 1, \ldots, N_\Omega\}$. Case number 0 corresponds to the global integral and is added as sanity check. At high temperatures, it can happen that part of the Kirchhoff-Planck function is shifted out of the frequency-range used in the integration. In this case, the w_i with $i > 0$ are renormalized during read-in to ensure that they sum-up to unity.

The table files contain further auxiliary quantities for the interpolation in the pressure-temperature-grid. Moreover, the correspondence between frequency index in the raw opacity data and OBM bin is stored. This allows to trace back which monochromatic frequencies contributed to a particular bin.

There are three main routines operating on the opacity table files. For each of them there are versions in Fortran77, Fortran90, and IDL available. These are: `dfopta(⟨filename⟩)` for reading an OBM table; `xkaros(P,T,ibin)` for interpolating the opacity at given pressure, temperature, and bin number; and `wtbbd(T,ibin)` for interpolating the weight of the source function.

Table 1. Presently available chemical compositions of raw opacity data. Empty positions indicate zero, i.e. scaled solar values.

[M/H]	[α/Fe]	[C/Fe]	[N/Fe]	[O/Fe]
0.5				
0.0				
-0.5	+0.2			+0.2
-1.0	*+0.4*			*+0.4*
-1.5	+0.4			+0.4
-2.0	*+0.4*			*+0.4*
-2.5	+0.4			+0.4
-3.0	*+0.4*			*+0.4*
-4.0	+0.4			+0.4
-4.0	+0.4	+1.0		+0.4
-4.0	+0.4	+2.0		+0.4
-4.0	+0.4	+3.0	+3.0	+0.4
-4.0	+0.4	+3.0	+3.0	+3.0

5. Chemical compositions

Table 1 summarizes which chemical abundance mixtures are presently available. This refers to the most comprehensive set of data from the MARCS package. Opacities from PHOENIX are only available for solar metallicity, the opacities for white dwarf models are restricted to a pure hydrogen composition. The entries in Tab. 1 set in italics refer to the base opacities which have been primarily used when constructing the CIFIST-grid. Notably, there is now a dataset of super-solar metallicity available, and carbon-enhanced mixtures of different carbon-to-oxygen ratios at [M/H] = −4. Effects on the atmospheric structure of 3D models need to be investigated.

6. Interlude: opacities and spectral synthesis

In this section we briefly digress, and want to comment on opacities in the spectral synthesis-related codes Linfor3D and NLTE3D which use CO5BOLD model snapshots as input. Linfor3D is adapted to synthesize individual line profiles or small spectral ranges including Doppler shifts and take into account the 3D geometry. NLTE3D solves the NLTE problem for a given

model atom also taking into account the 3D geometry.

In Linfor3D continuous opacities are needed, to which opacities of individual spectral lines are added in the line synthesis. The continuous opacities come from a different source than those applied in CO5BOLD. They are calculated with the help of two (Fortran) routines called `iondis` and `opalam`. `iondis` calculates ionization and dissociation equilibria to produce level populations assuming LTE. `opalam` uses these as input to compute monochromatic continuum opacities (absorption and scattering coefficients). While in cool stars the continua are dominated by H^--opacity which is well understood for some time, one should keep in mind that there are residual inconsistencies with the opacities applied in CO5BOLD due to differences in the implementation.

In a similar vein, NLTE3D is using `iondis` and `opalam` to compute LTE number densities and continuous opacities. The latter are combined with ATLAS9 opacity distribution functions representing the contribution of the line opacity in order to compute the 3D radiation field for estimating photoionization rates. Again, strictly speaking the employed opacities are not fully consistent with the opacities used in CO5BOLD runs as such. However, in view of the presently achievable precision we do not consider this a vital shortcoming since other uncertainties (like the OBM approximation proper) are likely to dominate the overall error budget.

7. OBM problems & improvements

When constructing the binned opacities several choices have to be made where the best option to be taken is non-obvious. Among them the choice of the opacity average in the optically thin regime. The OBM as implemented in CO5BOLD adopts Planck-means while Nordlund (1982) originally took intensity-weighted averages. Another choice is the switching depth separating optically thick and thin conditions.

7.1. Failure when introducing several opacity moments per bin?

For eliminating the switching depth in the OBM the authors experimented with an idea which was communicated to them by P. Woitke (2005, priv. comm.): instead of using a single mean opacity per bin one is using two with the property of representing the Planck- as well as Rosseland-mean simultaneously, here called κ_1 and κ_2. After calculating the bin-averaged Planck and Rosseland opacities, κ_1 and κ_2 are defined via the properties

$$\frac{1}{\kappa_{\rm R}} = \frac{1}{2}\left(\frac{1}{\kappa_1} + \frac{1}{\kappa_2}\right), \quad \text{and} \quad \kappa_{\rm P} = \frac{1}{2}\left(\kappa_1 + \kappa_2\right).$$

Note, that in the relations above a weighting is assumed which assigns the same weight to the two opacities, i.e., they each represent the same wavelength interval in the particular bin of interest. While this appears rather arbitrary here one can show that this choice leads to the most extreme ratio between opacities κ_1 and κ_2. Intuitively, one might expect that this maximizes the additional "information content" that a solution for both opacities might have. The two relations above can be solved for κ_1 and κ_2 giving

$$\kappa_{1,2} = \kappa_{\rm P} \pm \sqrt{\kappa_{\rm P}(\kappa_{\rm P} - \kappa_{\rm R})}.$$

Since generally, $\kappa_{\rm P} > \kappa_{\rm R}$ one obtains two well defined solutions for κ_1 and κ_2. Moreover, one solution is greater than $\kappa_{\rm P}$, the other smaller than $\kappa_{\rm R}$. The radiative transfer equation now has to be solved twice for each bin in which one chooses to use the two opacities κ_1 and κ_2, meaning that the computational demand has also doubled. As a benefit, the largely arbitrary switching depth is eliminated, and one may hope that the overall accuracy of the radiative energy exchange may improve. We implemented the procedure but, unfortunately, tests have shown that the overall accuracy was not improved. However, the actual reason for this failure is unclear, and one might speculate that it is related to the somewhat erratic convergence properties of the OBM. Nevertheless, we consider the idea of using several different moments still an interesting possibility to potentially improve the OBM.

7.2. Metal-poor giants and scattering

It is well known (e.g., Mihalas 1978) that Rayleigh scattering by neutral hydrogen atoms has an important impact on the atmospheric temperature structure of metal-poor stars, particularly in red giants. Scattering makes the solution of the radiative transfer equation computationally expensive since the independence between directions (or rays) is lost. Hayek et al. (2010) implemented coherent isotropic scattering in the solution of the radiative transfer equation exactly in their 3D code. Subsequently, Collet et al. (2011) tested the exact solution against the treatment of scattering opacities as true absorption or leaving out scattering opacities all together in the optically thin region of the atmosphere. In their metal-poor giant models Collet et al. found a large difference between the exact treatment and treating scattering as true absorption. Leaving out the scattering opacities gave results fairly close to the exact treatment. In the following, we will informally refer to this approximation as "Hayek's approximation".

We used Hayek's approximation for a number of opacity tables to test the impact of scattering in CO5BOLD models of metal-poor giants. Figure 4 shows one example for a metal-poor giant model at $T_{\rm eff} \approx 5000\,{\rm K}$, $\log g = 2.5$, and $[{\rm M/H}] = -3$ using six bins in the OBM. In comparison to the results of Collet et al. we find a significantly reduced sensitivity of the atmospheric structure to the treatment of scattering. While Collet et al. find a temperature change of $\approx 600\,{\rm K}$ at $\log \tau_{\rm Ross} = -4$, with CO5BOLD we find 120 K. The reason for this different sensitivity is not clear yet but may be rooted in the details of the actual sorting of frequency points into bins. A clarification of this point is likely to provide an explanation of the differences in 3D-1D abundance corrections which the two groups derive for metal-poor giant stars.

The opacity tables calculated at present usually come in two versions including scattering as true absorption, or applying Hayek's approximation. The user can choose between the options. On the long run, a more rigorous implementation of scattering is planned.

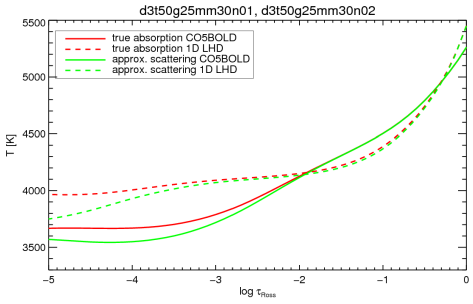

Fig. 4. Comparison of 1D radiative-convective equilibrium models (dashed lines) and 3D models (solid lines). Models depicted in red treat scattering opacities as true absorption, those in green ignore scattering opacities in the optically thin part of the atmosphere. (from Ludwig & Steffen 2012)

However, there are conceptual problems related to scattering in the OBM which need to be clarified first.

8. The future: OBM, ODF, or OS?

By now, the OBM has served its purpose for over three decades. The computing power increased tremendously during that period, and one might tackle the treatment of the frequency-dependence of the radiative transfer differently. As possible alternative techniques opacity distribution functions (ODFs), or opacity sampling (OS) come to mind which are now routinely applied in 1D stellar model atmospheres. Further refinements of the OBM might also be possible, as tried (unsuccessfully) by the procedure described in Sect. 7.1. Irrespective of the approach chosen, besides physical correctness the possibility of *automatization* should be one of the guidelines in the construction of the technique. Automatization shifts into focus since the construction of 3D model atmosphere *grids* is now undertaken, and the construction of opacity tables on a case by case basis becomes impracticable.

Acknowledgements. HGL acknowledges financial support by the Sonderforschungsbereich SFB 881 "The Milky Way System" (subproject A4) of the German Research Foundation (DFG).

References

Allard F., Homeier D., Freytag B., Sharp C. M., 2012, EAS, 57, 3
Collet, R., et al., 2011, A&A, 528, A32
Freytag, B., et al., 2012, J. Comp. Phys., 231, 919
Gustafsson, B., et al., , 2008, A&A, 486, 951
Hayek, W. et al., 2010, A&A, 517, A49
Ludwig, H.-G., 1992, PhD Thesis, Univ. of Kiel
Ludwig H.-G., Jordan S., Steffen M., 1994, A&A, 284, 105
Mihalas, D., 1978, *Stellar Atmospheres*, 2nd edition, W.H. Freeman
Ludwig H.-G., Steffen M., 2012, in: *Red Giants as Probes of the Structure and Evolution of the Milky Way*, p. 125, eds. A. Miglio, J. Montalbán & A. Noels
Nordlund, Å., 1982, A&A, 107, 1
Skartlien, R., 2000, ApJ, 536, 465
Tremblay P.-E., Bergeron P., 2009, ApJ, 696, 1755
Vögler A., 2004, A&A, 421, 755