

Numerical non-LTE 3D radiative transfer using a multigrid method

Johan P. Bjørgen and Jorrit Leenaarts

Institute for Solar Physics, Department of Astronomy, Stockholm University, AlbaNova University Centre, 106 91 Stockholm, Sweden
e-mail: johan.bjorgen@astro.su.se

Received 13 December 2016 / Accepted 4 January 2017

ABSTRACT

Context. 3D non-LTE radiative transfer problems are computationally demanding, and this sets limits on the size of the problems that can be solved. So far, multilevel accelerated lambda iteration (MALI) has been the method of choice to perform high-resolution computations in multidimensional problems. The disadvantage of MALI is that its computing time scales as $O(n^2)$, with n the number of grid points. When the grid becomes finer, the computational cost increases quadratically.

Aims. We aim to develop a 3D non-LTE radiative transfer code that is more efficient than MALI.

Methods. We implement a non-linear multigrid, fast approximation storage scheme, into the existing Multi3D radiative transfer code. We verify our multigrid implementation by comparing with MALI computations. We show that multigrid can be employed in realistic problems with snapshots from 3D radiative magnetohydrodynamics (MHD) simulations as input atmospheres.

Results. With multigrid, we obtain a factor 3.3–4.5 speed-up compared to MALI. With full-multigrid, the speed-up increases to a factor 6. The speed-up is expected to increase for input atmospheres with more grid points and finer grid spacing.

Conclusions. Solving 3D non-LTE radiative transfer problems using non-linear multigrid methods can be applied to realistic atmospheres with a substantial increase in speed.

Key words. radiative transfer – Sun: chromosphere – methods: numerical

1. Introduction

Most of our understanding of astrophysical objects comes from analyzing their spectra. The comparison of models of those objects to observations therefore requires solving the radiative transfer (RT) problem with the model as input. In this paper, we investigate a non-linear multigrid method for solving the three-dimensional (3D) non-local thermodynamical equilibrium (LTE) RT problem in cool stellar atmospheres, and specifically, the solar atmosphere.

One-dimensional (1D) non-LTE radiative transfer has been possible for many decades. More recently, multidimensional radiative transfer has become possible. Now that 3D radiation magnetohydrodynamics (MHD) simulations of solar and stellar atmospheres are widely used (Asplund et al. 2003, 2004; Uitenbroek 2006; Trujillo Bueno & Shchukina 2007; Pereira et al. 2009; Holzreuter & Solanki 2013; de la Cruz Rodríguez et al. 2013; Pereira et al. 2015; Rathore & Carlsson 2015; Amarsi et al. 2016), it has become important to have efficient 3D non-LTE radiative transfer methods to compare these models with observations.

In the solar chromosphere, full 3D RT is important for strongly scattering lines (e.g. Leenaarts et al. 2009, 2012, 2013a; Štěpán et al. 2015; Sukhorukov & Leenaarts 2017). The next solar generation telescopes, such as the 4-m DKIST, will obtain a diffraction limit of $0.03''$ at 500 nm, corresponding to 32 km on the Sun. To be able to compare these high-resolution observations with MHD simulations, 3D radiative transfer is needed even for much weaker photospheric lines (Holzreuter & Solanki 2015; Judge et al. 2015). Some important chromospheric diagnostic lines are influenced by partial redistribution effects, such as Lyman- α , Mg II h&k, and Ca II H&K. These effects have been included in a 3D RT code (Sukhorukov & Leenaarts 2017), but

at the cost of an increase in computational time up to a factor 10. We therefore need a robust method that can solve the 3D non-LTE problem more efficiently than currently possible.

It is well known that Λ -iteration converges prohibitively slowly in optically thick non-LTE problems. Accelerated Λ -iteration (ALI, Cannon 1973) converges much faster when using an approximate Λ -operator. For 3D, one typically uses the diagonal of Λ -operator as the approximate operator (this choice is also known as Jacobi iteration), because it can easily be constructed without the need of computationally expensive matrix inversions. Other methods exist: Gauss-Seidel (GS) iteration (Trujillo Bueno & Fabiani Bendicho 1995; Trujillo Bueno & Manso Sainz 1999) converges four times as fast as the diagonal operator, but cannot be parallelized as efficiently as Jacobi iteration (Tsitsiklis et al. 1988), a serious disadvantage in parallel solvers that are used for 3D applications. Conjugate gradient methods are under development, but have not yet been implemented in a general 3D non-LTE RT code (Paletou & Anterrieu 2009).

Therefore, multilevel accelerated lambda iteration (MALI; Rybicki & Hummer 1991, 1992) with a diagonal approximate operator remains the current standard for 3D non-LTE radiative transfer problems for solar applications (Uitenbroek 2001; Leenaarts & Carlsson 2009; Štěpán & Trujillo Bueno 2013).

The disadvantage with MALI is that the convergence slows down as the grid spacing gets finer (Olson et al. 1986). A solution to this problem was proposed by Steiner (1991), who implemented a linear multigrid method for two-level atoms with coherent scattering in 1D and two-dimensional (2D) test atmospheres and found a speed-up of a factor 4 to 40. The idea behind multigrid methods is that Jacobi and Gauss-Seidel iteration quickly smooth out the high-spatial-frequency error in the

solution, but decrease the low-frequency error only slowly. After a few iterations, the error will therefore be smooth (i.e., contain only low frequencies). The problem can therefore be transferred to a coarser grid, where the low-spatial-frequency errors become higher-frequency errors. Iterations on the coarse grid quickly decrease these errors. A correction to the fine-grid solution is computed on the coarse grid, and interpolated up onto the fine grid, which now has a much smaller low-frequency error.

This method was expanded by Vath (1994), who applied it to a 3D problem with a smooth atmosphere. He showed that the method works, but does not provide a large speed-up on small domains in parallel setup. Fabiani Bendicho et al. (1997) combined the non-linear multigrid technique with a multilevel Gauss-Seidel iteration scheme (MUGA). The smoothing properties of MUGA do not deteriorate with increasing grid resolution, which makes it an excellent choice to use with multigrid. They found a speed-up of a factor 3 to 32 with a five-level Ca II atom and a 2D atmosphere that was isothermal in the vertical direction and contained a weak temperature inhomogeneity in the horizontal direction (see Auer et al. 1994, Eq. (18)). Léger et al. (2007) compared the performance of multigrid against Gauss-Seidel iteration with successive over-relaxation in smooth 2D prominence models. Štěpán & Trujillo Bueno (2013) proved that multigrid works in full 3D non-LTE with a MALI operator and parallel implementation using spatial-domain decomposition. They used a plane-parallel semi-empirical FAL model C atmosphere Fontenla et al. (1993) with a sinusoidal horizontal inhomogeneity in temperature with an amplitude of 500 K.

All previous studies used relatively smooth model atmospheres. So far, no one has applied non-linear multigrid to “real-life” atmospheres produced by radiation-MHD simulations, which are highly inhomogeneous in all atmospheric quantities. This paper presents a non-linear multigrid method that can handle such atmospheres.

In Sects. 2 and 3, we describe the method and our computational setup. In Sect. 4, we present numerical considerations for the 3D radiative transfer with multigrid. In Sect. 5, we present the speedup we obtain compared to regular MALI iteration. Section 6 contains our conclusions.

2. Method

2.1. The non-LTE radiative transfer problem

The time-independent non-LTE radiative transfer problem consists of solving the transport equation

$$\frac{dI_\nu}{d\tau_\nu(n_i)} = S_\nu(n_i) - I_\nu, \quad (1)$$

with I_ν the intensity, τ_ν the optical path along a ray, S_ν the source function, and n_i the atomic level populations, together with the equations of statistical equilibrium for the level populations:

$$n_i \sum P_{ij}(I_\nu) - \sum n_j P_{ji}(I_\nu) = 0, \quad (2)$$

with P_{ij} the sum of the radiative and collisional rates. We wrote the dependencies of τ_ν and S_ν on the level populations and P_{ij} on the intensity explicitly to emphasize the non-linearity of the problem. The radiation field couples various grid points together, making the problem also non-local. We assume, in this paper, that the collisional terms and the background opacity and emissivity are constant and fully determined by the input atmosphere.

To close the system, we replace one of the rate equations at each grid point with a particle conservation equation:

$$n_{\text{tot}} = \sum_{i=1}^{n_{\text{levels}}} n_i. \quad (3)$$

The statistical equilibrium equation must be solved at all grid points, and the transfer equation must be solved at all grid points for all frequencies and ray propagation directions. The combination of the equations can be written as a coupled system of non-linear equations:

$$\mathcal{W}(n) = f. \quad (4)$$

We note that these are $n_{\text{levels}}n_{\text{gridpoints}}$ equations that, in principle, depend on all $n_{\text{levels}}n_{\text{gridpoints}}$ level populations.

For a given grid point, if we replace the first rate equation with particle conservation, f looks like:

$$f = \begin{bmatrix} n_{\text{tot}} \\ 0 \\ \vdots \\ 0 \end{bmatrix}. \quad (5)$$

In the MALI scheme, these equations are preconditioned and solved by iteration (Rybicki & Hummer 1991, 1992).

2.2. The non-linear multigrid method

In this subsection, we present how the non-linear multigrid method is applied to the radiative transfer problem, that is, how it is used to solve Eq. (4).

The non-linear multigrid, fast approximation storage (FAS) scheme as formulated by Brandt (1977), preserves the mathematical structure of \mathcal{W} at the coarser grids. For non-LTE codes, this means that one can reuse the formal solver, computation of collisional rates, and so on, and that the iterative MALI method can be used on the coarser grids with only minor modifications.

We restate the rate equation:

$$\mathcal{W}^k(n^k) = f^k, \quad (6)$$

where k is an index that denotes the grid level: $k = \ell, \ell - 1, \dots, 1$ where ℓ is the number of grids, with $k = \ell$ the finest grid and $k = 1$ the coarsest grid. We provide an initial estimate for n and perform a number (ν_1) MALI iterations, which is called pre-smoothing. Because MALI iterations act as a smoothing operator, we have reduced the high-frequency error and obtained a smooth residual r^k :

$$r^k = f^k - \mathcal{W}_*^k(n^k). \quad (7)$$

The $*$, denotes that an extra formal solution has been performed to get the radiative rates consistent with the current population estimate.

The exact solution can be described by adding a correction to the approximation: $n_{\text{exact}}^k = n^k + e^k$. The exact fine grid equation is then:

$$\mathcal{W}^k(n^k + e^k) = f^k. \quad (8)$$

Using the residual equation (Eq. (7)), the fine grid equation (Eq. (8)) may be written as

$$\mathcal{W}^k(n^k + e^k) = r^k + \mathcal{W}_*^k(n^k). \quad (9)$$

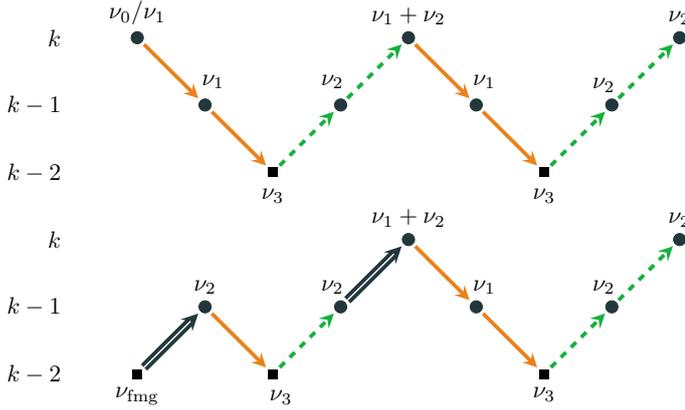


Fig. 1. Schematic representation of V-cycle and full multigrid with three grid levels. \Rightarrow means interpolating the full approximation, \rightarrow means restricting the approximation and the residual, and \dashrightarrow means interpolating the correction. The iteration parameters ν_{FMG} and $\nu_{1,2,3}$ are defined in Sect. 2, the parameter ν_0 in Sect. 4.12.

All these quantities can be mapped to a coarser grid $k-1$ using a restriction operator \mathcal{R}_k^{k-1} . We point out that we assume that only the residual is smooth and not the populations themselves. The populations from the fine grid are used as an initial approximation at the coarse grid. The coarse grid operator \mathcal{W}^{k-1} is obtained by performing a formal solution at the coarse grid. This results in the coarse-grid equation:

$$\mathcal{W}^{k-1} (n^{k-1} + e^{k-1}) = \mathcal{R}_k^{k-1}(r^k) + \mathcal{W}_*^{k-1} (\mathcal{R}_k^{k-1}(n^k)). \quad (10)$$

The right-hand side tries to force the solution of the system to maintain the fine-grid accuracy. The coarse-grid equation does not need to be solved exactly, since it is limited by the accuracy of the right-hand side. Instead, we compute an approximate solution by applying a number ν_3 MALI iterations, called coarse-grid iterations. The relative correction (see Sect. 4.6) is then computed by:

$$e^{k-1} = \frac{n_{\text{after}}^{k-1} - n_{\text{before}}^{k-1}}{n_{\text{before}}^{k-1}}, \quad (11)$$

where the subscripts before and after denote the populations before and after the ν_3 MALI iterations.

The correction is mapped on the operator to the finer grid with an interpolation operator \mathcal{I}_{k-1}^k :

$$n_{\text{new}}^k = n_{\text{old}}^k (1 + \mathcal{I}_{k-1}^k(e^{k-1})). \quad (12)$$

The interpolation introduces high-frequency errors on the fine grid. By applying ν_2 MALI iterations on the fine grid (called post-smoothing), these high-frequency errors are removed, and we have obtained a better approximation to the true solution on the finest grid. To obtain a converged solution, multiple cycles have to be performed.

The method we just described is called two-grid FAS. The same procedure can be applied to Eq. (10) because it has the same structure as Eq. (4). Doing this leads to the general multigrid procedure. The successive coarsening starting from the finest grid and the interpolation steps from the coarsest to the finest grid together are called a V-cycle.

2.3. Full multigrid

Full multigrid (FMG) can be seen as a method for obtaining a better initial approximation at the finest level. FMG is considered

to be the optimal multigrid solver (Trottenberg et al. 2000). We explain FMG using a three-grid ($\ell = 3$) example, as illustrated in Fig. 1. In FMG, Eq. (4) is solved at the coarsest grid:

$$\mathcal{W}^{k=1} (n^{k=1}) = f^{k=1}. \quad (13)$$

The initial population and the right-hand side can be obtained by restricting it from the finest grid or direct initialization at the $k = 1$ grid. After applying ν_{FMG} iterations, an initial approximation is acquired at the coarsest level. The full approximation is then interpolated up to the next finer grid $k = 2$:

$$n^{k=2} = \mathcal{I}_{k=1}^{k=2} (n^{k=1}). \quad (14)$$

The new approximation is used as the initial population for solving Eq. (4) at the $k = 2$ grid. Now, a V-cycle is applied at this level, instead of interpolating up to the finest grid $k = 3$. The reason for this is that the approximation at $k = 2$ will have both low-frequency and high-frequency errors due to the interpolation and discretization errors. These errors are efficiently reduced by a V-cycle. When the V-cycle is completed, we interpolate the population at $k = 2$ to the finest grid $k = 3$. Now we have an initial approximation at $k = 3$ that is closer to the true solution than a straight initialization at grid $k = 3$.

We reiterate that we are not aiming to get the full solution for the radiative transfer problem, but rather a better initial approximation at the finest grid. After reaching the finest level, the FMG is followed by regular FAS V-cycles to obtain the solution.

In Fig. 1, we show a schematic representation of the FAS and full-multigrid algorithms. Figure 2 shows an illustration of the behavior of the relative error in a smooth semi-empirical atmosphere and a non-smooth column from a radiation-MHD simulation, but with otherwise identical setup. Comparing the two cases shows that the radiation-MHD simulation column shows much more high-frequency error after the coarse grid correction than in FALC. These high-frequency errors stem from the inevitable inability to accurately represent a non-smooth atmosphere on a much coarser grid. We show in this paper that multigrid can still handle such atmospheres, but at the cost of an increase in computational time and thus a smaller increase in computational speed compared to MALI than reported in earlier works for smooth atmospheres.

3. Computational setup

We implemented the FAS multigrid method including FMG into the existing radiative transfer code Multi3D (Leenaarts & Carlsson 2009). This code solves the statistical equilibrium equations for one atomic species with a background continuum using MALI with a diagonal approximate operator (Rybicki & Hummer 1991, 1992). The code used 3D atmospheres on a Cartesian grid as input. It is parallelized with Message Passing Interface (MPI) using spatial domain decomposition as well as decomposition over frequency. The atomic bound-bound transitions can be treated in complete or partial redistribution (Sukhorukov & Leenaarts 2017). The transport equation is solved using short-characteristic (Kunasz & Auer 1988) with a linear or Hermite-spline (Auer 2003; Ibgui et al. 2013) interpolation scheme. We use the 24-angle quadrature (A4 set) from Carlson et al. (1963). All computations presented in this paper are done assuming complete redistribution and without frequency parallelization.

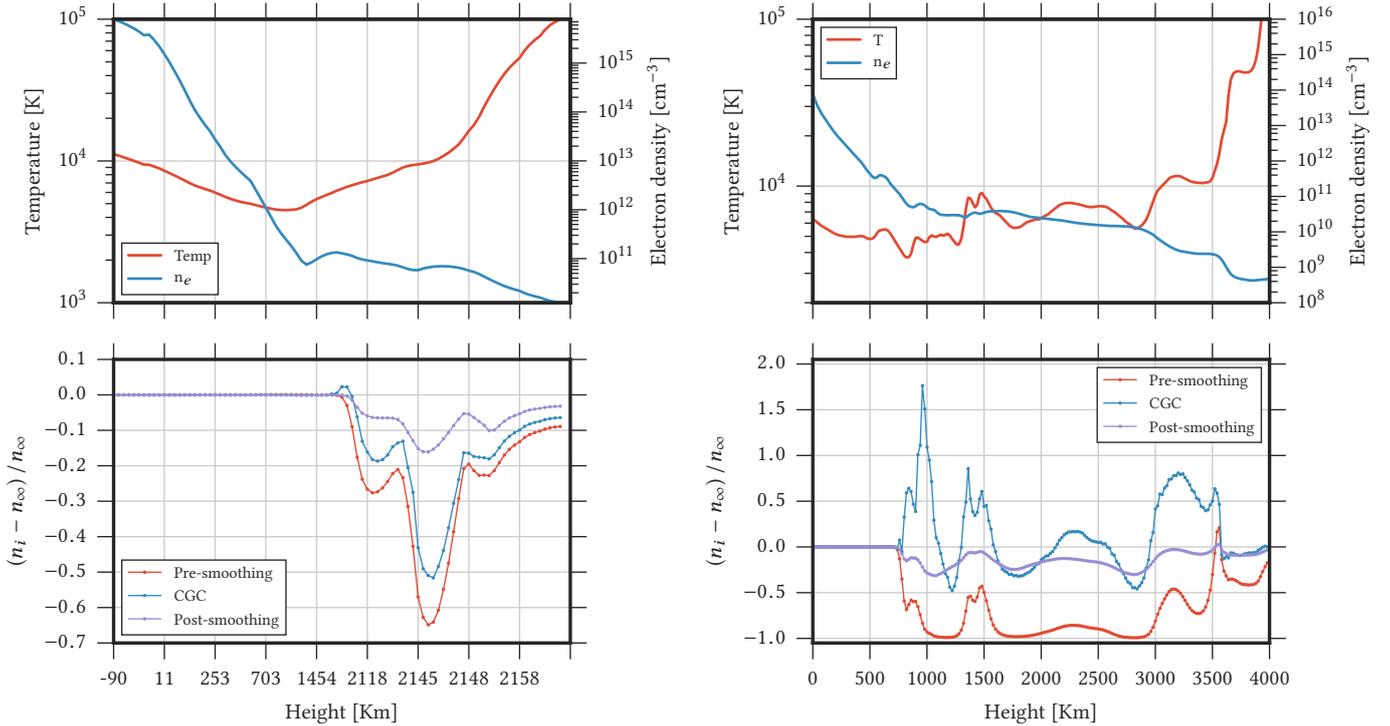


Fig. 2. Behavior of the relative error in population at the finest grid during various steps in a V-cycle with three grids, for the FAL-C atmosphere (*left-hand panels*) and a column extracted from a 3D radiation-MHD atmosphere (atmosphere Model 1, see Sect. 3). Both computations are done for a H I atom using 10 pre-smoothings, 32 coarse-grid iterations and 25 post-smoothings. *Upper panel:* temperature and electron density in the atmosphere. *Lower panels:* error for the H I $n = 2$ level after the pre-smoothing (red), the error after application of the coarse-grid correction (blue), and the error after application of the post-smoothing (purple). For both atmospheres the coarse grid correction has decreased the low spatial-frequency error, but increased the high-frequency error. The latter is largely removed by the post-smoothing.

3.1. Model atmospheres

We use three different model atmospheres in this paper.

The first is a snapshot from the radiation-MHD simulation by Carlsson et al. (2016) computed with the Bifrost code (Gudiksen et al. 2011) taken at $t = 3850$ s.

This model extends from the upper convection zone to the corona in a box with $504 \times 504 \times 496$ grid points, with a horizontal grid spacing of 48 km, and a vertical grid spacing ranging from 19 to 100 km. The model contains strong gradients in temperature, velocity, and density and has been used before in several studies (e.g. de la Cruz Rodríguez et al. 2013; Leenaarts et al. 2013a,b; Štěpán et al. 2015; Pereira et al. 2015; Loukitcheva et al. 2015; Rathore & Carlsson 2015; Rathore et al. 2015). We used this model to see how multigrid behaves in a realistic-use case. To obtain an odd number of grid points in the vertical direction (see Sect. 4.3), we added one extra layer in the convection zone through linear extrapolation of all quantities, so we obtain $504 \times 504 \times 497$ grid points. This MHD snapshot will be called atmosphere Model 1.

We also use a Bifrost simulation snapshot with $768 \times 768 \times 768$ grid points (Carlsson & Hansteen, priv. comm.). In this snapshot, the x - and y -axis are equidistant with a grid spacing of 32 km. The vertical axis is non-equidistant, with a grid spacing of 13 km in the photosphere and the chromosphere, and increasing to 60 km in the convection zone and the corona. The physical extent is the same as for Model 1: $24 \text{ Mm} \times 24 \text{ Mm} \times 16.9 \text{ Mm}$. We only used the part of the atmosphere between $z_{\min} = -0.5 \text{ Mm}$ and $z_{\max} = 6 \text{ Mm}$ to cover the formation height of the Ca II lines. The grid size is therefore reduced to $768 \times 768 \times 473$. This simulation has a similar

magnetic field configuration as Model 1, but was computed with an LTE equation of state instead of the non-equilibrium equation of state of Model 1. This atmosphere will be called Model 2. Finally, we also use some computations using the plane-parallel semi-empirical FAL model C atmosphere (Fontenla et al. 1993), which we will refer to as FAL-C.

3.2. Model atoms

We use three different model atoms. Most test computations were done with a minimalistic three-level Ca II atom, containing the ground level, the upper level of the Ca II K line, and the Ca III ground state. We also use the six-level Ca II atom and the six-level hydrogen atom from Carlsson & Leenaarts (2012). All bound-bound transitions are treated in complete redistribution.

3.3. Convergence criteria and error estimates

To validate and compare our multigrid implementation with the standard one-grid MALI scheme, we compare with reference solutions computed using the one-grid MALI scheme. With the three-level Ca II atom, we converged the reference solution to $\max |\delta n/n| \leq 10^{-5}$. For the hydrogen atom, we converged to $\max |\delta n/n| \leq 10^{-4}$. The reason for the lower criterium for hydrogen was to limit computational expenses: to reach the limit for hydrogen, we needed approximately 300 000 CPU hours.

To characterize the absolute error after i iterations, we used the infinity norm:

$$E_{i,\infty} = \max \left| \frac{n_i - n_\infty}{n_\infty} \right|, \quad (15)$$

where the n_∞ populations are taken to be the reference solutions. Since one grid point can slow down or stall the convergence with multigrid, the infinity-norm is the best choice. Other norms, such as the 2-norm, can give a false picture of the convergence of the multigrid method.

The maximum relative correction in population during each iteration is given by

$$C_i = \max \left| \frac{n_i - n_{i-1}}{n_i} \right|. \quad (16)$$

For linear non-LTE radiative transfer problems, one can define the spectral radius as the largest eigenvalue of the amplification matrix of the iteration scheme. However, in the non-linear case, this is not possible. Therefore we define the spectral radius operationally as

$$\rho_{\text{sr}} = \frac{E_{i+1}}{E_i}. \quad (17)$$

It measures how much the error is reduced for each iteration.

For normal use cases, there is no reference solution, so the error cannot be computed directly. For MALI iterations, it was shown (Auer et al. 1994; Fabiani Bendicho et al. 1997) that one can use the maximum relative correction to estimate the error. When the relative correction reaches the asymptotic linear regime, the relation between the error, the spectral radius, and the relative correction is:

$$E_i \approx \frac{\rho_{\text{sr}}}{1 - \rho_{\text{sr}}} C_i. \quad (18)$$

Because there is no reference solution, the spectral radius is estimated using $\rho_{\text{sr}} \approx C_{i+1}/C_i$. We test whether this approximation is also valid for multigrid iteration in Sect. 4.13.

4. Numerical considerations

In the following subsections, we discuss a number of issues that are relevant for radiative transfer using multigrid methods.

4.1. Grid coarsening

We use a simple coarsening strategy for the grid coordinates, by removing every second grid point along the x , y , and z axes when coarsening from grid k to $k - 1$.

4.2. Grid sizes

In the vertical direction, we use fixed non-periodic boundaries. In the lower boundary, the incoming radiation is set to the Planck function at the local gas temperature. In the upper boundary, the incoming radiation is set to zero or to pre-specified values. We chose to keep the vertical boundaries at fixed heights. To obtain this, we require an odd number of grid points in the vertical direction at all grid levels. Therefore, the number of vertical grid points N_z at the finest level must fulfil the following relation:

$$N_z = A \cdot 2^\ell + 1, \quad (19)$$

where A is an integer and ℓ the number of grids.

In the horizontal plane, we use periodic boundary conditions. Because Multi3D requires constant grid spacing in the horizontal direction, we thus require that $N_{x,y}$, the number of grid points in the x or y direction, follows

$$N_{x,y} = B \cdot 2^\ell, \quad (20)$$

with B an integer.

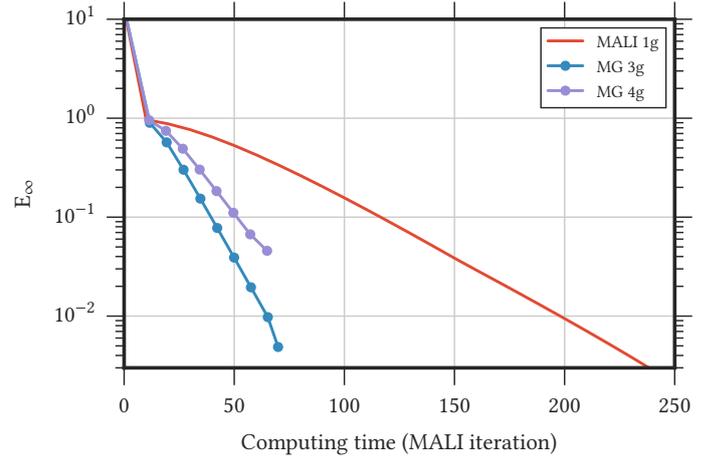


Fig. 3. Convergence behavior for MALI (red) and multigrid with three (blue) and four (purple) grids for the three-level Ca II atom within atmosphere Model 1. The computing time is expressed in units of the time it takes to perform one MALI iteration at the finest grid. The multigrid computations were performed with $\nu_1 = 2$, $\nu_2 = 2$, $\nu_3 = 32$, full-weighting restriction and trilinear interpolation.

4.3. Sub-domain sizes and number of grids

Multi3D divides the computational domain into sub-domains. The formal solver requires five ghost zones in the horizontal directions, and the number of internal points in the horizontal direction must be at least as large as the number of ghost zones. In the vertical direction, we have only one ghost zone. In single-grid applications, the sub-domains have typical sizes from 32^3 to 64^3 . In multigrid, we use the same domain decomposition as in the finest grid, and coarsen the grid local to each sub-domain. This limits the number of coarsenings that we can do. For a 32^3 sub-domain, we can achieve two coarsenings to 16^3 and 8^3 . A 64^3 sub-domain can be coarsened three times. Through various tests we found that three grids gives us the best and stable performance. Using more than three grids can lead to non-convergence or lower performance than when using three grids (see also Table 1 of Steiner 1991). In Fig. 3, we display the result of such a test for the three-level Ca II atom and atmosphere Model 1.

4.4. Restriction operator

We implemented three different restriction operators: injection, half-weighting, and full-weighting (e.g., Trottenberg et al. 2000). Injection is the most straightforward method, as it only removes grid points when moving a quantity from a fine to a coarse grid. Half-weighting is a smoothing operation, where a coarse grid point value is an average of the same point plus its six neighboring grid points located along the x , y , and z axes on the fine grid. Full-weighting is similar, but includes all 26 neighbor grid points. Full-weighting has the advantage that it conserves a quantity when it is integrated over the entire computational domain. Injection is computationally more efficient than the half- and full-weighting methods, but the extra computing time is negligible compared to the time spent on the formal solution.

We performed various tests using both the FAL-C and Model 1 atmospheres. We found that injection performs well for the FAL-C atmosphere, but not for Model 1. Figure 4 shows a comparison between single-grid MALI, three-grid injection, and three-grid full-weighting using the three-level Ca II atom in Model 1. The convergence of multigrid using injection stalls

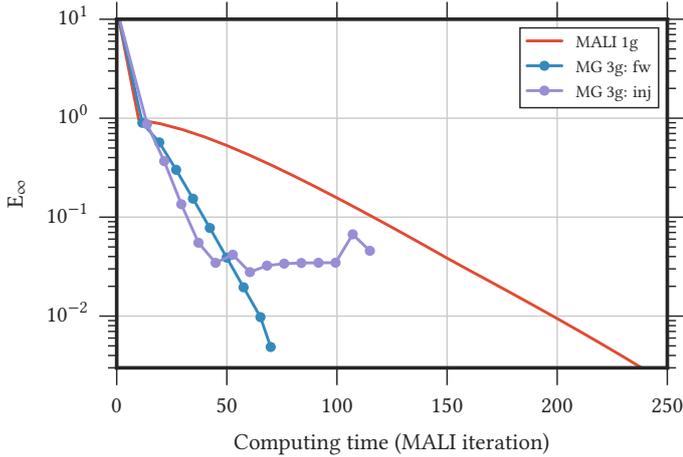


Fig. 4. Convergence behavior for MALI (red) and multigrid with injection (purple) and full-weighting (blue) as restriction operators, using the three-level Ca II atom and atmosphere Model 1. The multigrid computation was performed with $\nu_1 = 2$, $\nu_2 = 2$, $\nu_3 = 32$, three grids and trilinear interpolation.

around $E_\infty = 0.03$, while the full-weighting computation converges to the reference solution. We therefore chose to use full-weighting for the atomic level populations and residuals for all further computations.

4.5. Interpolation operator

We implemented two interpolation operators: linear and cubic Hermite (Auer 2003; Ibgui et al. 2013). The coefficients for linear interpolation depend only on the grid spacing and can be precomputed. A full 3D interpolation uses information from the eight surrounding grid points. Cubic Hermite interpolation needs a $4 \times 4 \times 4$ stencil for a 3D interpolation and requires the computation of the derivatives of the quantity to be interpolated on the inner eight grid points. This makes cubic Hermite interpolation somewhat more involved to implement and more computationally expensive. We performed a number of 1D test calculations using columns from Model 1 and found that, on average, both interpolation methods perform equally well. We therefore chose to use the simpler and faster linear interpolation for our 3D computations.

4.6. Correction

The radiative transfer problem is highly non-linear, and a level population can, in principle, change by orders of magnitude from one grid point to the next. This poses a problem for the interpolation of the correction from a coarser to a finer grid (Eqs. (11) and (12)) because the interpolation operation can introduce a large amount of interpolation noise and even lead to the unphysical prediction of negative populations at the fine grid. Both cases destroy the good convergence behavior of multigrid. We found that interpolating the relative correction (Eq. (11)) decreases the interpolation noise and largely eliminates the occurrence of negative populations as a consequence of interpolation.

4.7. The atmosphere at various grid levels

Multigrid requires solving the formal solution at all grid levels. We thus need a method to generate coarse-grid model atmospheres from the fine-grid atmosphere. We tested both injection

and full weighting to do so. Our tests indicate that injection can lead to the best convergence rate, especially in smooth atmospheres like FAL-C, but it might also cause stalls in convergence or even divergence. Full-weighting has a lower best-case convergence, but does not lead to stalls or divergence. The reason for this difference is that injection more accurately represents the fine-grid problem, and so has better best-case performance. But it also enhances gradients on the coarser grid, leading to stronger non-linearities and larger probability of stalling. Full-weighting on the other hand leads to coarse-grid atmospheres that are more different from the fine-grid atmosphere, leading to lower performance, but also has smaller gradients leading to increased stability. Because stability is usually the greatest concern when using radiation-MHD atmospheres, we use full-weighting for the atmosphere initialization in our computations.

4.8. Initializing the radiation field at various grid levels

Multi3D treats background processes using LTE opacities and a source function with a thermal part and a coherent-scattering part. At each grid level, we thus need to perform a small number of Λ -iterations to ensure the background processes are initialized correctly. Because a Λ -iteration takes a similar amount of time as a MALI iteration, we want to minimize the number of Λ -iterations. Empirically, we found that one Λ -iteration is sufficient at the finest grid. At the coarser grids, we obtain an initial estimate for the radiation field by restricting it from the next finer grid and performing one or two Λ -iterations.

4.9. Computational cost

In this paper, we are mainly concerned with the theoretical performance of multigrid compared to MALI iteration. We therefore express computation time in units of one MALI iteration at the finest grid in our figures. The computing time of a formal solution at coarse grid $k-1$ is therefore one eighth of the computing time at grid k .

However, Multi3D does not follow this theoretical perfect scaling. The reason is that the short-characteristic solver in Multi3D requires five ghost zones in the horizontal direction and one ghost zone in the vertical direction. As the number of grid points per sub-domain gets smaller with increasing coarsening, the relative amount of time spent on updating the ghost zones increases. In Fig. 5, we compare the theoretical maximum speed-up with the actual speed-up measured on a HP Cluster Platform 3000 with Intel Xeon E-5 2600 2.2 GHz cores using three-grids. At grid $k=2$, the performance penalty is a factor 3 for the $32 \times 32 \times 31$ sub-domain, but interestingly the $96 \times 96 \times 48$ sub-domain shows a speed-up better than the theoretical value. At the $k=1$ grid, the performance penalty is visible for all sub-domain sizes that we tested, but with the expected behavior that the larger the sub-domain, the smaller the performance penalty.

4.10. Occurrence of negative populations

Standard MALI iteration keeps atomic level populations positive as long as the iteration is started with positive populations. This is, however, not true for multigrid iterations. The reason is that coarse grid iterations have a modified right-hand side in the rate equations:

$$\mathcal{W}^k(n^k) = \begin{cases} f^k, & \text{if } k = \ell \\ \mathcal{R}_{k+1}^k (f^{k+1} - \mathcal{W}^{k+1}(n^{k+1})) + \mathcal{W}^k(\mathcal{R}_{k+1}^k(n^{k+1})), & \text{if } k < \ell. \end{cases} \quad (21)$$

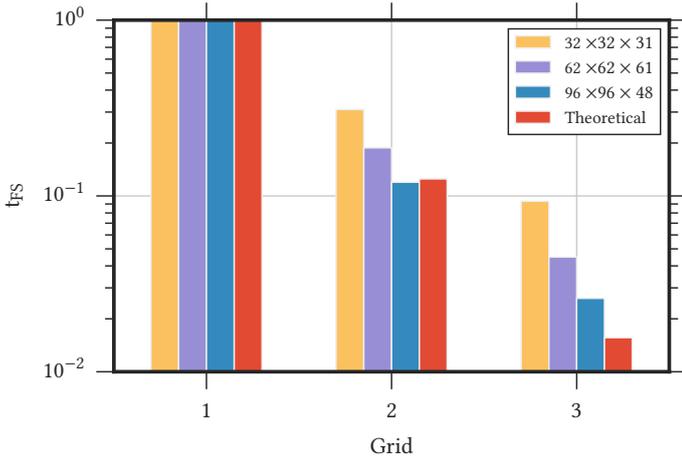


Fig. 5. Computational cost for one formal solution at each grid in multigrid with different sub-domain sizes. We set the cost of a formal solution at the finest grid to 1. The theoretical value is calculated as $(1/8)^{\ell-k}$, with k being the grid level.

The equation for $k < \ell$ can, in principle, have a negative population as a solution, or produce negative estimates of the solution during iteration. Negative populations on the coarse grids will propagate to the finest grid and destroy the convergence.

By experimentation, we found that negative populations at the coarse grid typically occur when the error $(n_i - n_\infty)/n_\infty$ is not smooth enough, but we did not manage to get an unambiguous definition of what is “smooth enough”. The error is sometimes not smooth at the location of strong gradients in the atmospheric parameters, but sometimes large error fluctuations also occur in smooth parts of the atmosphere models. The error is typically less smooth in our H I atom than in the Ca II atoms. This is caused by the larger spectral radius for hydrogen (see Sect. 4.13), which lowers the smoothing capability of Jacobi iterations.

Figure 6 illustrates this in an $\ell = 3$ test computation using a non-smooth 1D atmosphere and the H I atom. The left-hand columns show the error for the first two energy levels of the atom after two V-cycles but with a different number of post-smoothing iterations. The error for the $n = 2$ level exhibits a strong spike at a height of 1000 km for $\nu_2 = 2$, but this spike is much smaller for $\nu_2 = 10$ and $\nu_2 = 20$. In contrast, the error in the $n = 1$ population is relatively smooth, and at 1000 km in height, it is very small.

The black dot in the left-hand panels indicates the grid point that we investigate in more detail in the right-hand panels. The right-hand panels show the relative change in the $n = 1$ and $n = 2$ populations as a function of iteration at the coarsest grid ($k = 1$). The corrections are small for $n = 1$, and the $n = 1$ population remains positive for all coarse-grid iterations, as expected from the small error seen in the left-hand panels.

The $n = 2$ population behaves remarkably differently. For $\nu_2 = 2$, the first coarse-grid iteration predicts a relative change of -1.7 and thus produces a negative population. Successive corrections are smaller, but the population remains negative. When $\nu_2 = 10$ (second row), the population is positive for the first 3 iterations, but then turns negative until iteration 22, after which it is positive again. For $\nu_2 = 20$, the behavior is similar, but the number of iterations where the populations are negative is shorter.

For comparison, in the fourth row, we show the relative change at the coarsest grid with $\nu_2 = 2$ but where we perform

standard MALI iterations (Eq. (4)) instead of coarse-grid iterations (Eq. (10)). Here the populations are always positive, demonstrating again that MALI iterations keeps a positive solution positive.

If negative populations occur after application of the correction (Eq. (11)), we simply replace it with 10% of the uncorrected population and continue as usual with post-smoothing iterations. If this happens at isolated points in the atmosphere, then the convergence of the multigrid iteration is not affected, because a single outlying grid-point represents high-frequency noise, which is damped away efficiently. However, if there are too many negative populations in the same area of the atmosphere then convergence will be slower or the iteration might even diverge.

4.11. Initial population

A good initial approximation for the level populations is critical for good convergence. There are two popular methods to initialize the populations: LTE and the zero-radiation-field approximation. The latter means solving the statistical equilibrium equations with the radiation field set to zero everywhere in the atmosphere. We found that LTE-populations as an initial approximation typically lead to negative populations at the coarse grid.

We suspect that the reason is that the residual and the error are less smooth using LTE initialization than using the zero-radiation approximation. To obtain a smooth residual and error, many more iterations are required. We demonstrate our conjecture by comparing the smoothing number (Hackbush 1985; Fabiani Bendicho et al. 1997) as a function of MALI iterations at the fine grid for both initializations. The smoothing number is defined as

$$\rho_L(i) = \max |\mathcal{D}_2 [n_i - n_\infty]|, \quad (22)$$

where the operator \mathcal{D}_2 is the second-order derivative and i is the number of MALI iterations. In Fig. 7, we show that the zero-radiation initialization leads to a much smoother error than LTE initialization. A similar result is obtained for the residual. We therefore use zero-radiation initialization for all our computations.

4.12. Selection of pre- and post-smoothing parameters

The convergence and stability of the multigrid method depends on the selection of the multigrid parameters ν_1 , ν_2 and ν_3 (see Sect. 2.2). To ensure a good initial smoothing before the first V-cycle, we define an additional parameter, ν_0 , which is the number of MALI iterations performed after initialization. This means that, at the finest grid level, we perform ν_0 iterations during the first V-cycle, but use ν_1 during all later V-cycles.

In order to analyze the multigrid behavior as a function of the ν parameters, we use two metrics. The first is the spectral radius. The other metric is the speed-up S_i , which we define as the ratio

$$S_i = t_{\text{MALI}}^i / t_{\text{MG}}^i, \quad (23)$$

where t_{MG}^i is the computing time required to perform i V-cycles (with an associated error $E_{i,\infty}$), and t_{MALI}^i the computing time to reach the same error. A small spectral radius means one needs fewer V-cycles to reach convergence, but this does not necessarily mean that S_i also increases because one might need many iterations (as defined by the value of ν_1, ν_2 and ν_3) within a V-cycle to obtain the small spectral radius.

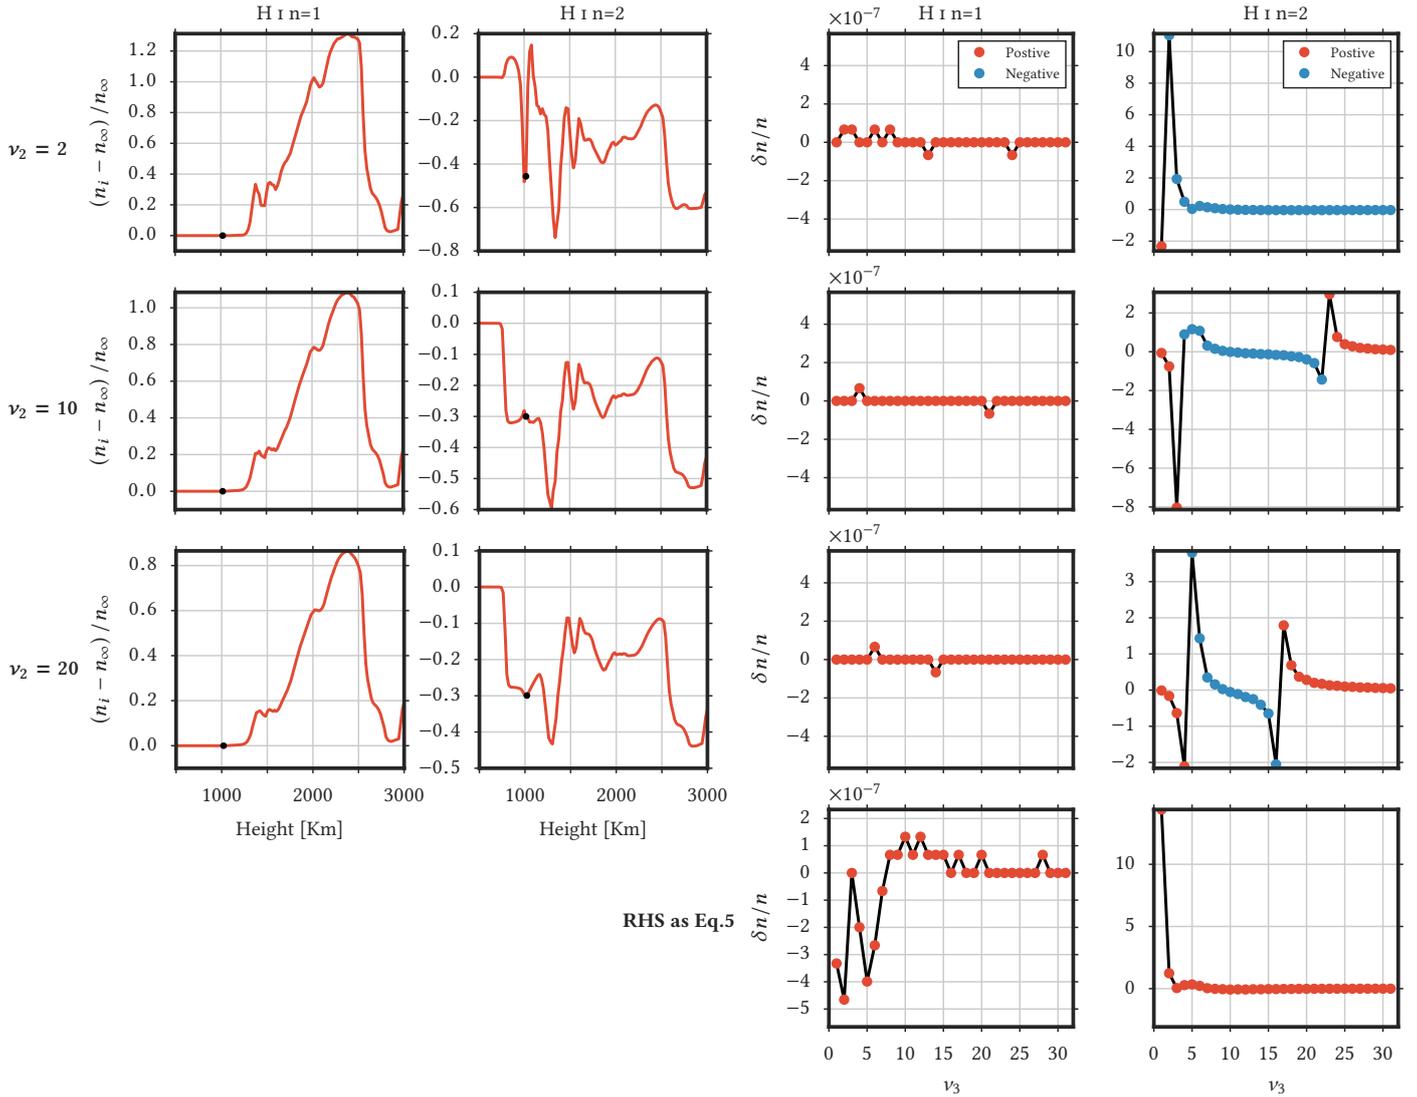


Fig. 6. Dependence on the number of post-smoothing iterations of the occurrence of negative populations at the coarsest grid. The *left-hand columns* show the relative error at the finest grid after two full V-cycles. The *right-hand columns* show the relative change of the population at the coarsest grid in a single grid point in the atmosphere. Red dots indicate positive populations; blue dots negative populations. The chosen grid point is indicated by the black dot in the *left-hand columns*. For each row of panels, a different number of post-smoothing iterations is applied. First row: $\nu_2 = 2$; second row: $\nu_2 = 10$; third row: $\nu_2 = 20$. In the fourth row, we show the relative change of the population at the coarsest grid with $\nu_2 = 2$, but we solve Eq. (4) instead of Eq. (10), that is, we perform standard MALI iterations instead of coarse-grid iterations. The computation was performed in a 1D plane-parallel atmosphere constructed from column $x = 0$, $y = 244$ of Model 1, with the H I atom. The other multigrid parameters are $\nu_1 = 2$, $\nu_3 = 32$, full-weighting restriction and linear interpolation.

We found that every atmosphere and model atom behaves differently, and it is impossible to give a set of parameters that always provides the fastest stable multigrid iteration. Empirically, we found that two to four pre-smoothings ($\nu_1 = 2-4$) worked well in all our test cases, but for Models 1 and 2, with the hydrogen atom we required extra initial smoothing for the first cycle ($\nu_0 \geq 15$). At the coarsest grid, we found that $\nu_3 \approx 32$ is a good number, independent of the problem. Iterating further did not significantly influence the convergence speed or stability.

The most sensitive parameter is the post-smoothing number. Our tests with the Ca II atoms in Model 1 and FAL-C showed stable and fast convergence with $\nu_2 = 2$, but for Model 2, we used $\nu_2 = 8$ to avoid excessive occurrence of negative populations. In Fig. 8, we show the speedup and spectral radius in a 1D calculation where we used a column from Model 1 as a plane-parallel 1D atmosphere. The spectral radius decreases with increasing ν_2 ,

but the speedup gets smaller owing to the increase in computational cost per V-cycle.

For hydrogen, we found that one needs many more post-smoothing iterations ($\nu_2 = 25$). If a smaller number is chosen, then the iteration at the coarsest grid tends to produce large areas with negative populations. We found that this is related to the smoothness of the error at the finest grid. We illustrate this in Fig. 6. The two left-hand columns show the error for the first two levels of our hydrogen atom at the start of the third V-cycle for computations with a different number of post-smoothings ($\nu_2 = 2, 10$ and 20). The error becomes smoother with increasing ν_2 . The right-hand columns show the sign of the level populations and the maximum correction as function of the iteration number at the coarsest grid during the third V-cycle. For $\nu_2 = 2$, a negative population occurs already during the first coarse-grid iteration and this remains so for all other iterations. For $\nu_2 = 10$, negative populations occur between iterations 3 and 23, after

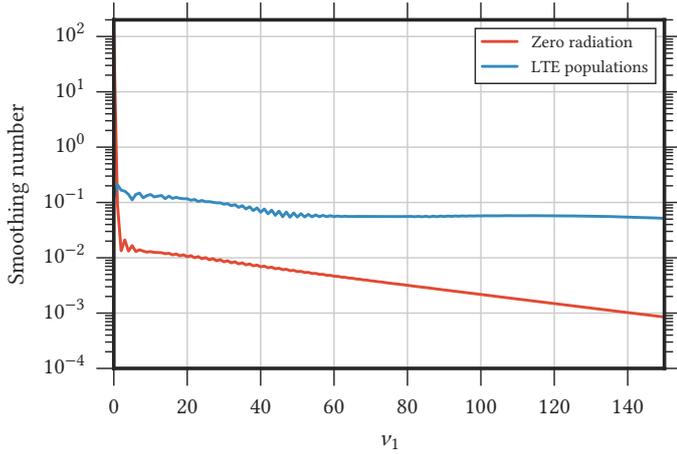


Fig. 7. Dependence of the smoothing number $\rho_L(v_1)$ (see Eq. (22)), on the initial condition of the populations. The results were computed using column $x = 0, y = 271$ from Model 1 as a 1D atmosphere and the hydrogen model atom.

which all populations become positive. For $\nu_2 = 20$, negative populations only occur during iterations 4 through 17. For comparison, in the fourth row, we show iterations at the coarsest grid for $\nu_2 = 2$, where we solve Eq. (4) instead of Eq. (10), that is, we perform standard MALI iterations instead of coarse-grid iterations. Here, no negative populations occur, demonstrating that negative populations occur owing to the right-hand-side of Eq. (10) and not owing to the coarse grid spacing.

4.13. Spectral radius and convergence criteria

The spectral radius for the three-level Ca II with MALI one-grid iteration in Model 1 is $\rho_{sr} = 0.971$, for and HI atoms, it is $\rho_{sr} = 0.9923$. An optimal multigrid method should achieve a spectral radius of 0.1–0.2 (Trottenberg et al. 2000). We obtained a spectral radius in the range of 0.3–0.4 with the pre- and post-smoothing parameters as discussed in Sect. 4.12 using Eq. (17).

The spectral radius of one-grid MALI iterations increases with decreasing grid spacing (Olson et al. 1986). Combined with Eq. (18), this means that one needs to iterate longer and longer to reach the same level of accuracy as the grid spacing gets finer. In contrast, the spectral radius of multigrid remains constant irrespective of the grid spacing (Steiner 1991; Fabiani Bendicho et al. 1997; Trottenberg et al. 2000), so that the speed-up of multigrid is expected to increase with decreasing grid spacing.

The relation between the maximum relative correction and the error (Eq. (18)) is valid for one-grid MALI iteration (Auer et al. 1994). Fabiani Bendicho et al. (1997) showed that it holds for multigrid in a smooth 2D atmosphere too. The 3D multigrid computations presented in Štěpán & Trujillo Bueno (2013) did not test the validity of the relation in 3D. Therefore we tested the validity in atmosphere Model 1. The results are shown in Fig. 9, where we compare the true error, the error estimate from Eq. (18), and the maximum relative correction. For Ca II we find that the formula still holds. We also find that the maximum relative correction is equal to the error. This is accidental, because the spectral radius is close to 0.5. For different pre- and post-smoothing parameters, a different result would be obtained. We also note that the maximum relative correction in Eq. (18) must be for the whole V-cycle, and not for the individual MALI iterations in the post-smoothing at the finest grid. In the hydrogen case, the approximation underestimates the true error by a

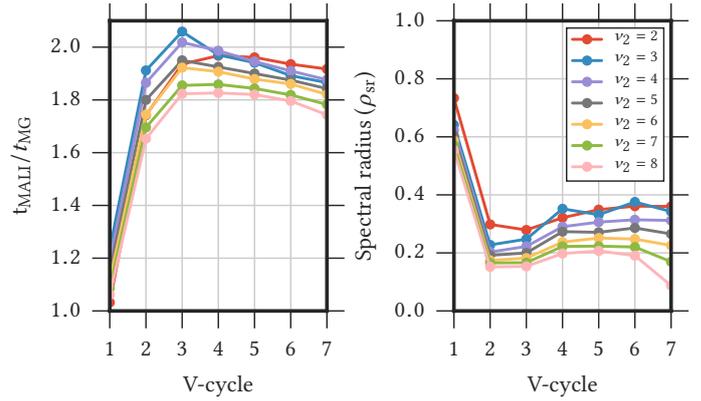


Fig. 8. Speed-up (left-hand panel) and spectral radius (right-hand panel) of multigrid as a function of the number of V-cycles. The computations were performed using column $x = 0, y = 271$ from Model 1 as a 1D atmosphere and the three-level Ca II model atom. The different colors indicate a different number of post-smoothing iterations ν_2 . The other multigrid parameters are: three grids, $\nu_1 = 2, \nu_3 = 32$, linear interpolation, and full-weighting restriction.

factor 2. Again, the maximum relative correction is only similar to the error because of the spectral radius being close to 0.5. The Ca II computation reaches the linear regime after five V-cycles. For HI it requires only one cycle. This is owing to the much higher number of post-smoothing iterations for hydrogen: $\nu_{2, \text{HI}} = 25$ and $\nu_{2, \text{CaII}} = 2$.

5. Results

In Table 1, we show a summary of our 3D non-LTE multigrid computations. We obtained a speed-up of 3.3 for Ca II and 4.5 for HI with atmosphere Model 1 and standard multigrid. Using full multigrid, we obtain a speed-up of 6 for Ca II, an improvement of a factor 1.8 compared to ordinary multigrid. A similar result for a smooth atmosphere was found by Fabiani Bendicho et al. (1997). In the hydrogen case, we needed 25 post-smoothing iterations to avoid too many occurrences of negative populations. This considerably reduced the speed-up. The convergence behavior of these three runs is shown in Fig. 11.

We also performed multigrid computation in atmosphere Model 2 to test whether multigrid can handle a finer grid spacing ($\Delta_{x,y} = 32$ km, $\Delta_z = 13$ km) and a different atmosphere. Obtaining a reference solution for this model was too computationally expensive. Therefore, we estimate the convergence based on Eq. (18), which we showed to be accurate for Ca II in Model 1 (see Fig. 9). The computation with Model 2 proves that 3D non-LTE radiative transfer using multigrid can handle different atmospheres and grid spacings.

Figure 10 shows the brightness temperature in the Ca II K line core for Models 1 and 2. The image computed from Model 2 shows substantially more fine-structure. This highlights the need for higher-resolution radiation-MHD simulations and the capacity to efficiently compute synthetic diagnostics using 3D non-LTE radiative transfer from them.

6. Summary and conclusions

In this article, we investigated whether or not a non-linear multigrid method can be applied to 3D non-LTE radiative transfer problems using snapshots from radiation-MHD simulations as

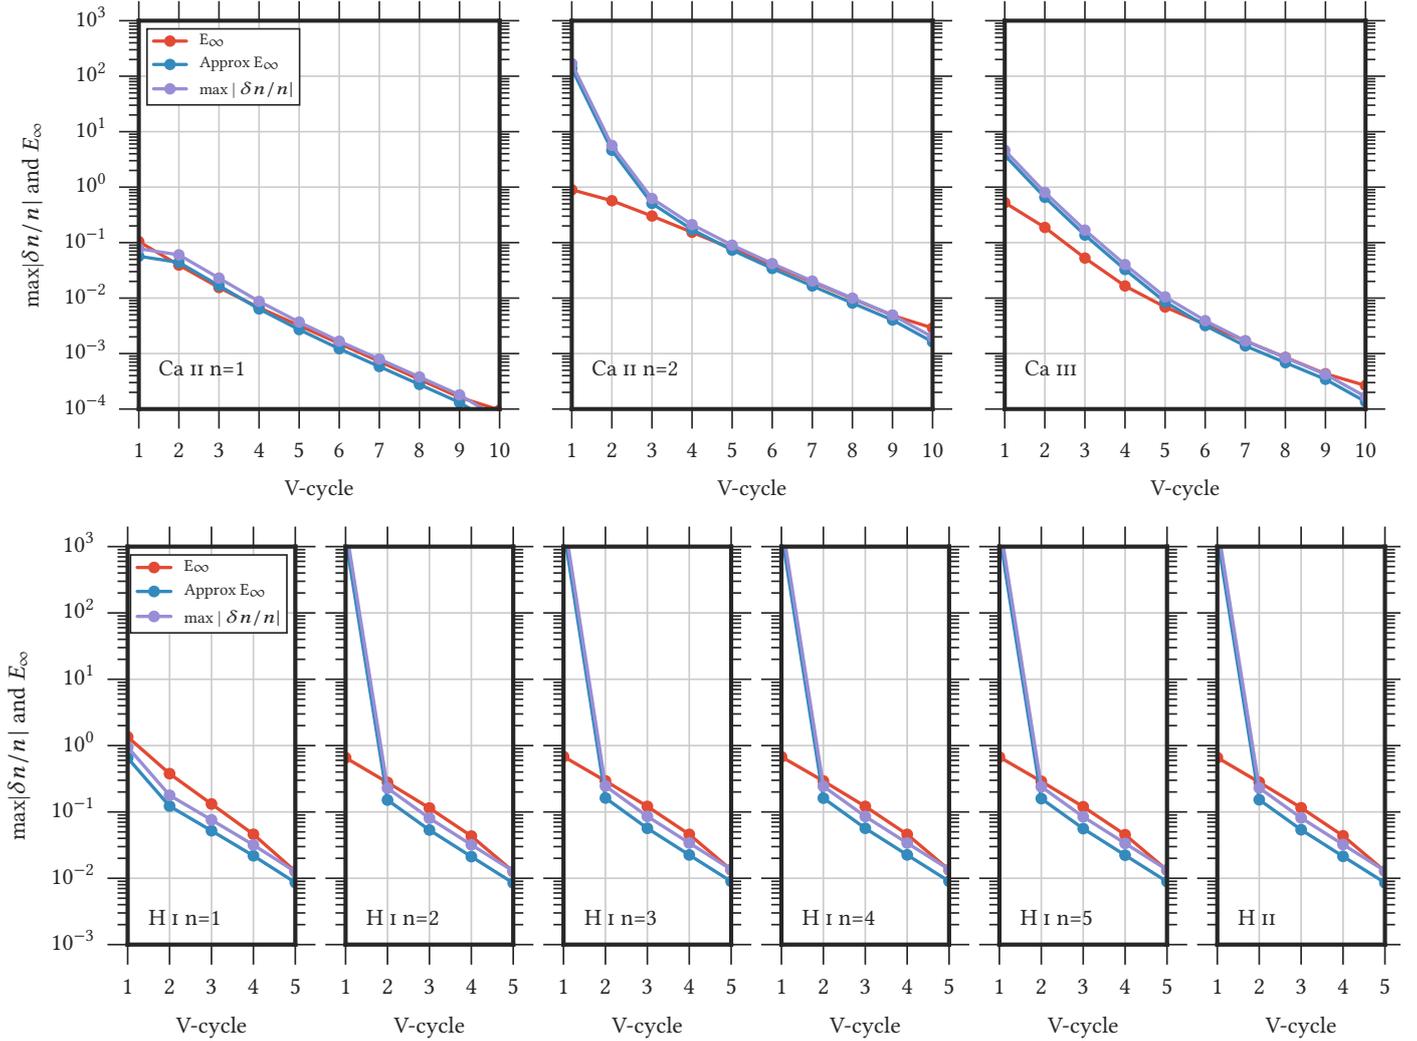


Fig. 9. Convergence behavior of multigrid, showing the maximum relative error and the maximum relative error estimated by Eq. (18) for each V-cycle, as well as the maximum relative change in population per V-cycle. *Upper panels:* three-level Ca II atom; *lower panels:* H I, both for atmosphere Model 1. For three-level Ca II we used $\nu_1 = 2$, $\nu_2 = 2$, $\nu_3 = 32$ and for H I, we used $\nu_0 = 15$, $\nu_1 = 2$, $\nu_2 = 25$, $\nu_3 = 32$.

Table 1. Summary of 3D non-LTE multigrid runs.

Atmosphere	Atom	Method	$\nu_{\text{FMG}}, \nu_0, \nu_1, \nu_2, \nu_3$	Speed-up
Model 1	3-level Ca II	MALI		1
		3-grid MG	–, 0, 2, 2, 32	3.3
		3-grid full-MG	16, 0, 2, 2, 32	6
	6-level H I	MALI		1
		3-grid MG	–, 15, 2, 25, 32	4.5
Model 2	6-level Ca II	3-grid MG	–, 20, 4, 8, 32	unknown

input. We showed that this is indeed the case: we obtain a speed-up factor of between 4.5 and 6 for a $504 \times 504 \times 496$ atmosphere (48 km horizontal and 19 km vertical resolution) computed with the radiation-MHD code Bifrost for Ca II and H I atoms. We also solve the 3D non-LTE radiative transfer problem using multigrid for a $768 \times 768 \times 768$ (32 km horizontal and 13 km vertical resolution) Bifrost snapshot. We conclude that the multigrid method can handle model atmospheres with very strong gradients in atmospheric parameters as well as strongly scattering lines such as hydrogen Lyman- α and Ca II H&K.

The implementation of multigrid requires choosing suitable numerical methods for the sub-tasks. We investigated several of

those choices using 1D and 3D test computations. Our findings are summarized as follows:

- we suggest using full-weighting as restriction operator;
- linear interpolation works well. Third-order Hermite interpolations appear to give higher convergence in 1D problems and should be investigated in 3D problems;
- one should interpolate the relative correction to the populations, and not the absolute correction;
- we suggest using full-multigrid if possible;
- initializing with the zero-radiation approximation is substantially better than initializing with LTE;

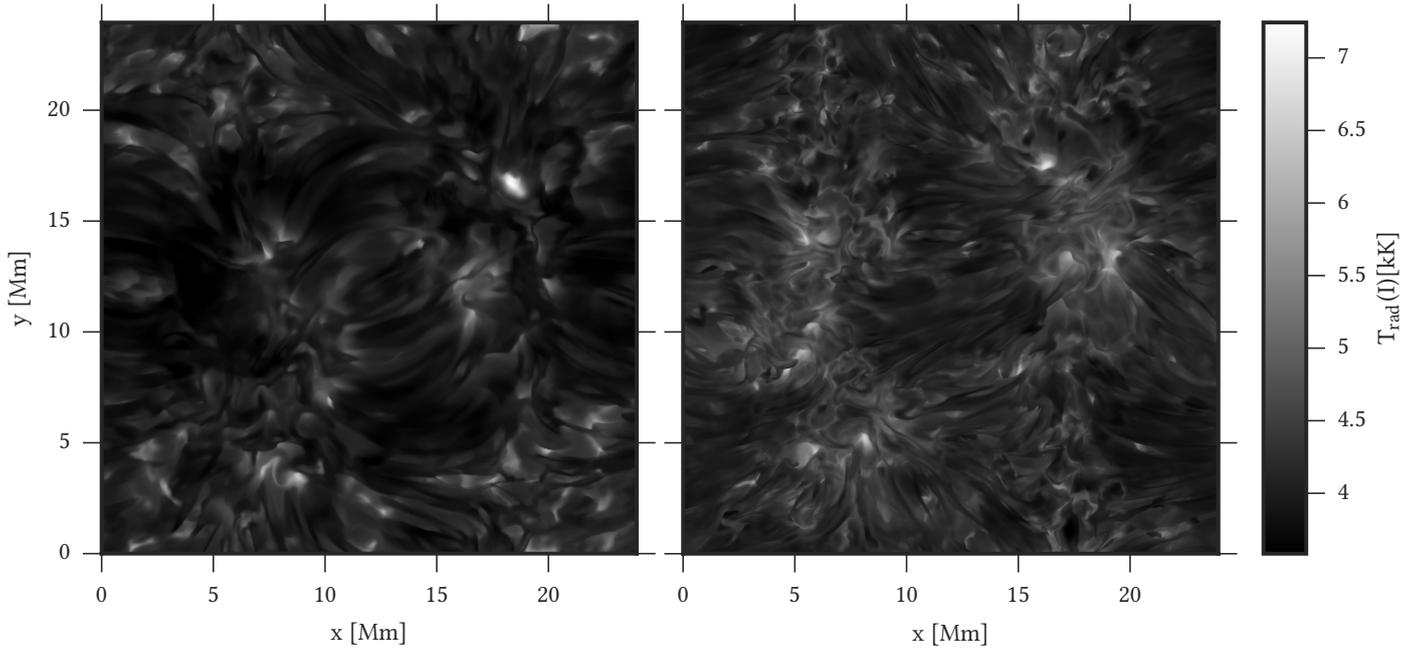


Fig. 10. Intensity of the Ca II K line core computed from atmosphere Model 1 (*left*) and Model 2 (*right*) at $\mu_z = 1$. The intensity is shown as the brightness temperature T_{rad} computed from $B_\nu(T_{\text{rad}}) = I_\nu$.

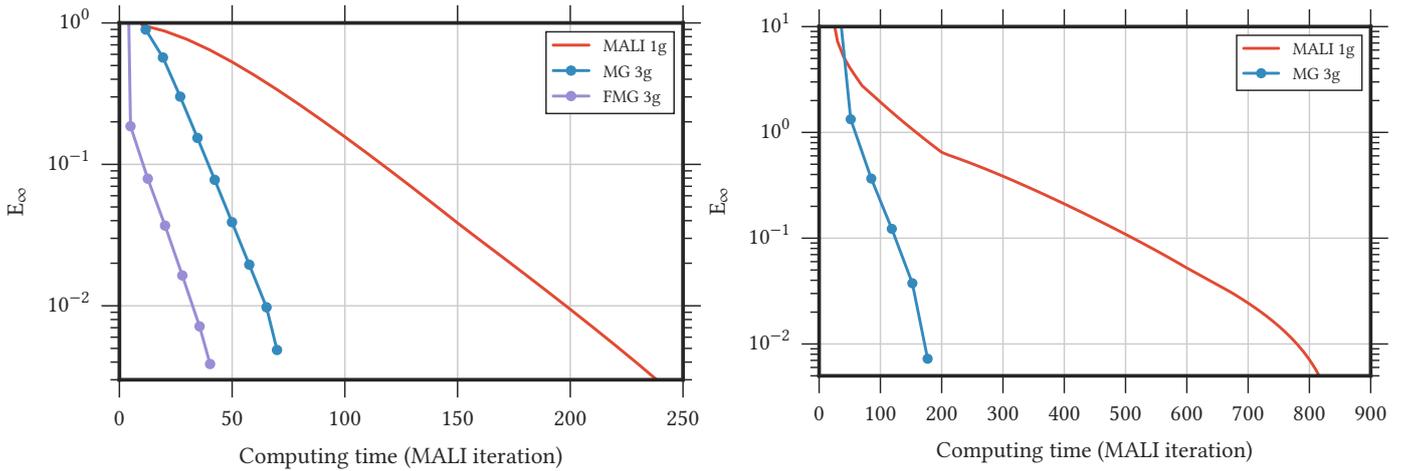


Fig. 11. Convergence behavior for MALI and multigrid for three-level Ca II (*left-hand panel*) and hydrogen atom (*right-hand panel*) for atmosphere Model 1. The computation was performed with three grids, full-weighting restriction, and trilinear interpolation. For three-level Ca II we used $\nu_1 = 2$, $\nu_2 = 2$, $\nu_3 = 32$ and for HI we used $\nu_0 = 15$, $\nu_1 = 2$, $\nu_2 = 25$, $\nu_3 = 32$.

- three-grid iteration converges faster than four-grid iteration;
- the coarse-grid iterations can converge to a negative solution. This is not a problem for isolated grid points, but if extended regions have negative populations, one should increase the number of post-smoothing iterations;
- each atom and atmosphere requires some testing to find the optimal number of pre-smoothings and post-smoothings. Our findings in Table 1 can be used as a starting point.

Since each problem is unique, other atmospheres and atoms could require different approaches. Therefore, the multigrid method should be implemented into radiative transfer codes in a modular way so that methods can be easily changed.

We did not obtain the high convergence rate (as measured in spectral radius of the multigrid iteration) as reported in Fabiani Bendicho et al. (1997). There are two reasons for this. First, these authors used a static smooth 2D with a weak horizontal

temperature inhomogeneity and no vertical temperature gradient, while we are using moving atmospheres with very large gradients in all atmospheric parameters. Second, they use Gauss-Seidel (GS) iterations, while we use Jacobi iteration in Multi3D. The smoothing properties and the convergence speed of GS iterations are superior to Jacobi iteration. Unfortunately, no MPI-parallelization scheme exists for GS iteration that scales well to thousands of computing cores, and we are forced to use Jacobi iteration. The lower convergence speed per iteration for Jacobi iteration can fortunately be offset by increasing the number of computing cores, but, ideally, one should develop an efficient parallel GS iteration scheme. A similar conclusion was reached by Štěpán & Trujillo Bueno (2013).

So far, we have only tested our multigrid method using complete frequency redistribution. Because partial frequency redistribution (PRD) can increase the computing time in non-LTE problems by more than an order of magnitude, the obvious next

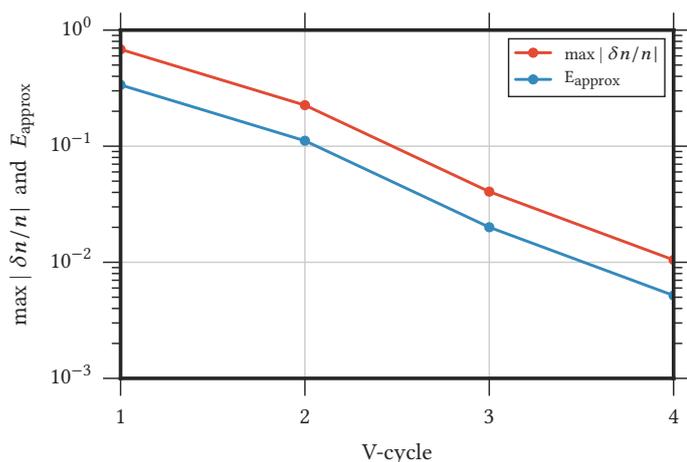


Fig. 12. Convergence properties for atmosphere Model 2 with six-level Ca II as a function of V-cycle. The E_{approx} is calculated based on the relation between the spectral radius and maximum relative population change (Eq. (18)). The computation was performed with three grids, full-weighting, and linear interpolation using $\nu_0 = 20$, $\nu_1 = 4$, $\nu_2 = 8$, $\nu_3 = 32$.

step will be to test the combination of 3D PRD (Sukhorukov & Leenaarts 2017) and multigrid in realistic-use cases. In this paper, we have used small model atoms with up to six levels. Multigrid method should also be tested with more complex atoms, such as Fe I, which are important for stellar photospheric abundance determinations (e.g., Nordlander et al. 2016).

The speed-up factor of multigrid increases with decreasing grid spacing. With the ever-increasing grid sizes and decreasing grid spacing of solar and stellar radiation-MHD models will make the use of multigrid methods in 3D non-LTE radiative transfer more and more desirable to keep computing costs manageable.

Acknowledgements. Computations were performed on resources provided by the Swedish National Infrastructure for Computing (SNIC) at the National Supercomputer Centre at Linköping University, at the High Performance Computing Center North at Umeå University, and at the PDC Centre for High Performance Computing (PDC-HPC) at the Royal Institute of Technology in Stockholm. J.P.B. and J.L. thank Mats Carlsson and Viggo Hansteen for providing a 768³ Bifrost snapshot. J.P.B. thanks Andrii Sukhorukov for answering questions regarding radiative transfer.

References

Amarsi, A. M., Asplund, M., Collet, R., & Leenaarts, J. 2016, *MNRAS*, **455**, 3735
 Asplund, M., Carlsson, M., & Botnen, A. V. 2003, *A&A*, **399**, L31
 Asplund, M., Grevesse, N., Sauval, A. J., Allende Prieto, C., & Kiselman, D. 2004, *A&A*, **417**, 751

Auer, L. 2003, in *Stellar Atmosphere Modeling*, eds. I. Hubeny, D. Mihalas, & K. Werner, ASP Conf. Ser., 288, 3
 Auer, L., Bendicho, P. F., & Trujillo Bueno, J. 1994, *A&A*, **292**, 599
 Brandt, A. 1977, *Math. Comput.*, **31**, 333
 Cannon, C. J. 1973, *ApJ*, **185**, 621
 Carlsson, M., & Leenaarts, J. 2012, *A&A*, **539**, A39
 Carlson, B., Adler, B., Fernbach, S., & Rotenberg, M. 1963, in *Statistical Physics*, eds. B. Alder, S. Fernbach, & M. Rotenberg (New York: Academic), 1
 Carlsson, M., Hansteen, V. H., Gudiksen, B. V., Leenaarts, J., & De Pontieu, B. 2016, *A&A*, **585**, A4
 de la Cruz Rodríguez, J., De Pontieu, B., Carlsson, M., & Rouppe van der Voort, L. H. M. 2013, *ApJ*, **764**, L11
 Fabiani Bendicho, P., Trujillo Bueno, J., & Auer, L. 1997, *A&A*, **324**, 161
 Fontenla, J. M., Avrett, E. H., & Loeser, R. 1993, *ApJ*, **406**, 319
 Gudiksen, B. V., Carlsson, M., Hansteen, V. H., et al. 2011, *A&A*, **531**, A154
 Hackbush, W. 1985, in *Multi-Grid Methods and Applications* (Berlin Heidelberg: Springer)
 Holzreuter, R., & Solanki, S. K. 2013, *A&A*, **558**, A20
 Holzreuter, R., & Solanki, S. K. 2015, *A&A*, **582**, A101
 Ibgui, L., Hubeny, I., Lanz, T., & Stehlé, C. 2013, *A&A*, **549**, A126
 Judge, P. G., Kleint, L., Uitenbroek, H., et al. 2015, *Sol. Phys.*, **290**, 979
 Kunasz, P., & Auer, L. H. 1988, *J. Quant. Spectr. Rad. Transf.*, **39**, 67
 Leenaarts, J., & Carlsson, M. 2009, in *The Second Hinode Science Meeting: Beyond Discovery-Toward Understanding*, eds. B. Lites, M. Cheung, T. Magara, J. Mariska, & K. Reeves, ASP Conf. Ser., 415, 87
 Leenaarts, J., Carlsson, M., Hansteen, V., & Rouppe van der Voort, L. 2009, *ApJ*, **694**, L128
 Leenaarts, J., Carlsson, M., & Rouppe van der Voort, L. 2012, *ApJ*, **749**, 136
 Leenaarts, J., Pereira, T. M. D., Carlsson, M., Uitenbroek, H., & De Pontieu, B. 2013a, *ApJ*, **772**, 89
 Leenaarts, J., Pereira, T. M. D., Carlsson, M., Uitenbroek, H., & De Pontieu, B. 2013b, *ApJ*, **772**, 90
 Léger, L., Chevallier, L., & Paletou, F. 2007, *A&A*, **470**, 1
 Loukitcheva, M., Solanki, S. K., Carlsson, M., & White, S. M. 2015, *A&A*, **575**, A15
 Nordlander, T., Amarsi, A. M., Lind, K., et al. 2016, *A&A*, **597**, A6
 Olson, G. L., Auer, L. H., & Buchler, J. R. 1986, *J. Quant. Spectr. Rad. Transf.*, **35**, 431
 Paletou, F., & Anterrieu, E. 2009, *A&A*, **507**, 1815
 Pereira, T. M. D., Asplund, M., & Kiselman, D. 2009, *A&A*, **508**, 1403
 Pereira, T. M. D., Carlsson, M., De Pontieu, B., & Hansteen, V. 2015, *ApJ*, **806**, 14
 Rathore, B., & Carlsson, M. 2015, *ApJ*, **811**, 80
 Rathore, B., Carlsson, M., Leenaarts, J., & De Pontieu, B. 2015, *ApJ*, **811**, 81
 Rybicki, G. B., & Hummer, D. G. 1991, *A&A*, **245**, 171
 Rybicki, G. B., & Hummer, D. G. 1992, *A&A*, **262**, 209
 Steiner, O. 1991, *A&A*, **242**, 290
 Sukhorukov, A. V., & Leenaarts, J. 2017, *A&A*, **597**, A46
 Trottenberg, U., Oosterlee, C. W., & Schuller, A. 2000, in *Multigrid* (Orlando: Academic Press)
 Trujillo Bueno, J., & Fabiani Bendicho, P. 1995, *ApJ*, **455**, 646
 Trujillo Bueno, J., & Manso Sainz, R. 1999, *ApJ*, **516**, 436
 Trujillo Bueno, J., & Shchukina, N. 2007, *ApJ*, **664**, L135
 Tsitsiklis, J. N. et al. 1988, *A comparison of Jacobi and Gauss-Seidel parallel iterations* (Massachusetts Institute of Technology, Laboratory for Information and Decision Systems)
 Uitenbroek, H. 2001, *ApJ*, **557**, 389
 Uitenbroek, H. 2006, *ApJ*, **639**, 516
 Štěpán, J., & Trujillo Bueno, J. 2013, *A&A*, **557**, A143
 Štěpán, J., Trujillo Bueno, J., Leenaarts, J., & Carlsson, M. 2015, *ApJ*, **803**, 65
 Vath, H. M. 1994, *A&A*, **284**, 319