Monte-Carlo methods for NLTE spectral synthesis of Supernovae.

M. Ergon[1], C. Fransson[1], A. Jerkstrand[2], C. Kozma[1], M. Kromer[3, 4], and K. Spricer[5]

[1] The Oskar Klein Centre, Department of Astronomy, AlbaNova, Stockholm University, 106 91 Stockholm, Sweden
[2] Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Str. 1, D-85748 Garching, Germany
[3] Heidelberger Institut für Theoretische Studien, Schloss-Wolfsbrunnenweg 35, D-69118 Heidelberg, Germany
[4] Institut für Theoretische Astrophysik am Zentrum für Astronomie der Universität Heidelberg, Philosophenweg 12, D-69120 Heidelberg, Germany
[5] Department of Mathematics, Stockholm University, 106 91 Stockholm, Sweden

To be submitted to Astronomy and Astrophysics.

ABSTRACT

We present JEKYLL, a new code for modelling of supernova (SN) spectra and lightcurves based on Monte-Carlo (MC) techniques for the radiative transfer. The code assumes spherical symmetry, homologous expansion and steady state for the matter, but is otherwise capable of solving the time-dependent radiative transfer problem in non-local-thermodynamic-equilibrium (NLTE). The method used was introduced in a series of papers by Lucy, but the full time-dependent NLTE capabilities of it have never been tested. Here, we have extended the method to include non-thermal excitation and ionization as well as charge-transfer and two-photon processes. Based on earlier work, we calculate the non-thermal rates by solving the Spencer-Fano equation. Using a method previously developed for the SUMO code, macroscopic mixing of the material is taken into account in a statistical sense. To save computational power we use a diffusion solver in the inner region, where the radiation field may be assumed to be thermalized. In addition, we use a statistical Markov-chain model to sample the emission frequency more efficiently, and introduce a method to control the sampling of the radiation field, which is used to reduce the noise in the radiation field estimators. Except for a description of JEKYLL, we provide comparisons with the ARTIS, SUMO and CMFGEN codes, which show good agreement in the calculated spectra as well as the state of the gas. In particular, the comparison with CMFGEN, which is similar in terms of physics but uses a different technique, shows that the Lucy method does indeed converge in the time-dependent NLTE case. Finally, as an example of the time-dependent NLTE capabilities of JEKYLL, we present a model of a Type IIb SN, taken from a set of models presented and discussed in detail in an accompanying paper. Based on this model we investigate the effects of NLTE, in particular those arising from non-thermal excitation and ionization, and find strong effects even on the bolometric lightcurve. This highlights the need for full NLTE calculations when simulating the spectra and lightcurves of SNe.

Key words. supernovae: general, radiative transfer

1. Introduction

Modelling the spectral evolution and lightcurves of supernovae (SNe) is crucial for our understanding of these phenomena, and much effort has been put into this during the last 50 years. To achieve realistic results local thermodynamic equilibrium (LTE) can generally not be assumed, and the full frequency-dependent, non-LTE (NLTE) problem has to be solved. Several paths exist and here we follow the one outlined in a series of papers by Lucy (2002, 2003, 2005, hereafter L02, L03, L05), in turn based on earlier work by Mazzali & Lucy (1993) and Lucy (1999). Using this method, hereafter referred to as the Lucy method, the radiative transfer is solved by a Monte-Carlo (MC) calculation, which is alternated with a NLTE solution for the matter until convergence is achieved (Λ-iteration). Basic tests were performed in the original papers, and a simplified version of the method, assuming LTE for the population of excited states, has been implemented in the code ARTIS (Sim 2007; Kromer & Sim 2009, hereafter K09 and S07). Several other codes, as e.g. TARDIS (Kerzendorf & Sim 2014), SEDONA (Kasen et al. 2006), SAMURAI (Tanaka et al. 2007) and the one by Mazzali (2000) are also based on the method, (or the early version of it), but are all restricted in one way or another. The Lucy method, or parts thereof, has also been used for modelling of other phenomena than SNe, e.g. in the code by Carciofi & Bjorkman (2006,

2008) and in the code by Long & Knigge (2002), upgraded by Sim et al. (2010), both for radiative transfer in different kinds of circum-stellar disks.

Here, we present JEKYLL, a C++ based code which implements the full NLTE-version of the method, extended to include also non-thermal excitation and ionization as well as charge-transfer and two-photon processes. These extensions are particularly important for modelling in the nebular phase, and for the calculation of the non-thermal rates we use the method developed by Kozma & Fransson (1992, hereafter KF92). Contrary to ARTIS, the initial version of JEKYLL is restricted to a spherical symmetric geometry. However, as the MC radiative transfer is performed in 3-D, this is not a fundamental restriction. The most fundamental limitation in JEKYLL, shared with most spectral codes, is the absence of hydrodynamics. As is often (but not always) justified, we instead assume the ejecta to be in homologous expansion.

Another code based on Λ-iteration and MC radiative transfer is the steady-state NLTE code SUMO (Jerkstrand et al. 2011, 2012, hereafter J11 and J12) aimed for the nebular phase. However, contrary to JEKYLL and other codes based on the Lucy method, conservation of MC packet energy is not enforced, so the MC technique used by SUMO is different. SUMO uses a novel statistical approach to represent the macroscopic mixing

of the ejecta occurring in the explosion, which we have also adopted in JEKYLL.

In addition to the MC based codes, there is a group of codes that use finite difference techniques to solve the radiative transfer equation in a more traditional way. Examples of such codes are PHOENIX (Hauschildt & Baron 1999) and the general purpose NLTE code CMFGEN (Hillier & Miller 1998, hereafter H98), which is similar to JEKYLL in terms of physics. In this paper, we compare JEKYLL with CMFGEN as well as with ARTIS and SUMO, which are also similar to JEKYLL in one way or another. These comparisons provide a thorough and critical test of the JEKYLL code. In a broader context, the comparison with CMFGEN also provides a test of the full time-dependent NLTE capabilities of the Lucy method.

In an accompanying paper (Ergon et al. in prep, hereafter Paper 2) we present an application of JEKYLL to Type IIb SNe, by modelling the early (before 150 days) evolution of a set of models previously evolved through the nebular phase with SUMO (Jerkstrand et al. 2015, hereafter J15). One of those is also presented in this paper as an example of a time-dependent NLTE calculation using a realistic ejecta model. However, for comparisons to observations and a deeper discussion of Type IIb SNe we refer to Paper 2.

The paper is organized as follows. In Sect. 2 we discuss the underlying physical problem, and in Sect. 3 we describe the method used to solve this problem and the design of the code. In Sect. 4 we provide the comparisons of JEKYLL with the ARTIS, SUMO and CMFGEN codes, and in Sect. 5 we provide the (example) application to Type IIb SNe, as well as some further tests based on it. Finally, in Sect. 6 we conclude and summarize the paper.

2. Physics

The general physical problem addressed is the time-evolution of the radiation field and the state of the matter, given the dynamical constraint of homologous expansion, and might be referred to as a radiation-thermodynamical problem. If the radiation field and the matter are in local thermodynamic equilibrium (LTE) this is simplified to a one-parameter (i.e the temperature) problem, and may be easily solved. Otherwise, we are in the non-LTE (NLTE) regime, and the number of parameters, as well as the complexity of the problem, increase drastically.

As is often done, we solve for the radiation field and state of the matter separately, and the problem is split into a radiative transfer and a thermodynamical part. The coupling, provided by radiation-matter interactions, is enforced through Λ-iterations, where the state of the matter and the radiation field are alternately and iteratively determined from each other. The Λ-iteration concept is at the heart of the method, and in Sect. 2.1 we provide some background and discuss the somewhat different meaning it has in traditional and MC based methods.

The state of the matter can be separated into a dynamical and thermodynamical part, where the former is trivially given by $\rho = \rho_0 (t/t_0)^{-3}$ and $v = r/t$ through the constraint of homologous expansion. The thermodynamical part is given by the temperature, and the populations of ionized and excited states, which are solved for using the thermal energy equation and the NLTE rate equations, respectively. To simplify we assume steady state, which is motivated if the thermodynamical time-scale is much smaller than the dynamical time-scale.

The radiation field is given by the specific intensity, which is solved for using an extended version of the MC based Lucy method, which is discussed in Sect. 3.3. In a traditional code like

CMFGEN, the specific intensity is solved for using the radiative transfer equation, whereas in a MC based code like JEKYLL, the radiative transfer is treated explicitly by propagating radiation packets which interact with the matter through absorption, emission and scattering. The different radiation-matter interactions supported are discussed in Sect. 2.4.

In addition, in SN ejecta radioactive decays emit high-energy photons or leptons, which give rise to a non-thermal electron distribution. Through collisions, these electrons contribute to the heating of the electron gas and the excitation and ionization of the ions. The problem may be broken up into two parts; deposition of the radioactive decay energy, and the partitioning of this energy into non-thermal heating, ionization and excitation.

2.1. Λ-iterations and convergence

In terms of the Λ-operator the radiative transfer equation may be written as $I = \Lambda[S]$[1], where I is the intensity and S the source-function. If the source function depends on the intensity, as in the case of scattering, solving the problem requires inverting the Λ-operator. This is typically a costly operation, and we may instead try an iterative procedure called Λ-iteration (see e.g. Hubeny & Mihalas 2014). In its original form an improved estimate of the intensity is then determined using the previous estimate of the source-function, i.e. $I_{i+1} = \Lambda[S_i]$. However, this method may converge extremely slowly if the source function is dominated by scattering, as the non-local coupling introduced only propagates one mean-free path per iteration. This may be solved by splitting the Λ-operator in two parts, one acting on the current iteration and one acting on the previous iteration, i.e. $I_{i+1} = \Lambda^*[S_{i+1}] + (\Lambda - \Lambda^*)[S_i]$. With an appropriate choice of Λ^*, e.g the local part of Λ, which is trivial to invert and still close to Λ, convergence could be accelerated, and the procedure is therefore known as accelerated Λ-iteration (see e.g. Cannon 1973b,a, Scharmer 1984, Werner & Husfeld 1985 and Olson et al. 1986).

It is important to realize that the explicit dependence of the scattering emissivity on the intensity does not cause slow convergence in the MC case. The reason for this is that the MC scattering emissivity depends directly on the current iteration of the MC radiation field. Actually, a MC Λ-iteration is similar to an accelerated Λ-iteration in the sense that the current iteration depends partly on itself. Note, however, that the implicit dependence of the MC emissivities on the intensity (via the matter quantities) may still cause slow convergence, and this is the problem addressed by the Lucy method. Enforcing the constraints of thermal and statistical equilibrium on the MC calculation, introduces a direct (but approximate) dependence of all MC emissivities on the current iteration of the MC radiation field. Although not formally proved, this is likely to accelerate the convergence in the general case, and as demonstrated in L03, Λ-iterations based on this method have excellent convergence properties.

It is worth noting that most MC based methods use the Sobolev approximation, which also helps to accelerate the convergence as line self absorption is already solved for. Note also that in a steady-state calculation, all locations are causally connected to each-other, whereas in a time-dependent calculation the causally connected region grows with time, which makes convergence less demanding in each individual step of a time-dependent calculation than in a steady-state calculation.

---

[1] **Note, that historically the Λ-operator was defined in terms of the mean intensity, but this distinction does not matter for the discussion.**

2.2. Statistical equilibrium

To determine the populations of ionized and excited states, the NLTE rate equations need to be solved. Assuming steady state, these equations simplify to the equations of statistical equilibrium, where the rates of transitions in and out of each state are in equilibrium. The equation of statistical equilibrium for level i of ion I may be written

$$\sum_{J=I\pm1} r^{\mathrm{BF}}_{J,j\to I,i}\, n_{J,j} + \sum_{j\neq i} r^{\mathrm{BB}}_{I,j\to i}\, n_{I,j} = \left(\sum_{J=I\pm1} r^{\mathrm{BF}}_{I,i\to J,j} + \sum_{j\neq i} r^{\mathrm{BB}}_{I,i\to j} \right) n_{I,i}$$
(1)

where r is the rate (per particle) for bound-free (superscript BF) and bound-bound (superscript BB) transitions, and n is the number density. Note that the system of equations is non-linear as (some) transition rates (per particle) depend on the number densities. Transitions may be caused by absorption/emission of photons (Sect. 2.4), or by collisions involving ions and thermal (Sect. 2.5) or non-thermal (Sect 2.6.1) electrons.

2.3. Thermal equilibrium

To determine the thermal state of the gas the thermal energy equation needs to be solved. Assuming steady state, this equation simplifies to the equation of thermal equilibrium, where the heating and cooling of the gas are in equilibrium. The equation of thermal equilibrium may be written

$$\sum_{J=I\pm1} g^{\mathrm{BF}}_{I,i\to J,j}(T)\, n_{I,i} + \sum_{i\neq j} g^{\mathrm{BB}}_{I,i\to j}(T)\, n_{I,i} + \sum g^{\mathrm{FF}}_{I}(T)\, n_{I,i} = H^{\mathrm{NT}}$$
(2)

where g is the net heating rate (per particle) for bound-free (superscript BF), bound-bound (superscript BB) and free-free (superscript FF) transitions, and H^{NT} is the heating rate by non-thermal collisions. Heating/cooling may arise through absorption/emission of photons (Sect. 2.4), or through collisions involving ions and thermal (Sect. 2.5) or non-thermal (Sect 2.6.1) electrons.

2.4. Radiation-matter interactions

In radiation-matter interactions, the radiation field and the matter (electrons and ions) exchange energy through absorption and emission of photons. Except for electron scattering, which is assumed to be coherent and isotropic in the co-moving frame (of the ejecta), and given by the Thomson cross-section, JEKYLL supports the following interactions.

Bound-bound Through detailed balance, the excitation and de-excitation rates are related and determined by a single quantity, e.g. the spontaneous emission coefficient. We assume that the Sobolev approximation (Sobolev 1957) applies, which is appropriate when expansion broadening dominates the thermal broadening. Expressions for the Sobolev optical depth as well as the transition rates are given in L02. In addition, we also support de-excitation through two-photon emission for bound-bound transitions otherwise radiatively forbidden.

Bound-free Through detailed balance, the ionization and recombination rates are related and determined by a single quantity, e.g. the photo-ionization cross-section. In bound-free transitions, the energy absorbed/emitted goes partly into ionization/recombination of the ion, and partly into heating/cooling of the electron gas. Expressions for the opacity, emissivity, transition rates and heating/cooling rates are given in L03.

Free-free (i.e. bremsstrahlung) Assuming thermal matter, the opacity and emissivity are related through Kirchoffs law. In free-free interactions, the energy of the photons absorbed/emitted goes solely into heating/cooling of the electron gas. Expressions for the opacity, emissivity, and heating/cooling rates are given in L03.

2.5. Matter-matter interactions

In matter-matter interactions, electrons and ions exchange energy through collisions. The collisions heat/cool the electron gas and result in bound-bound or bound-free transitions of the ions. Except for non-thermal collisions, which are discussed in Sect. 2.6.2, JEKYLL supports the following interactions.

Bound-bound and bound-free Through detailed balance, the collisional excitation and de-excitation rates are related and determined by a single quantity, e.g. the collisional strength. The same is true for the collisional ionization and recombination rates, and expressions for the transition rates and heating/cooling rates are given in L03.

Charge-transfer In collisions involving two ions, electrons may be transferred from one ion to another. This process is called charge-transfer and may be viewed as a recombination followed by a ionization. The charge transfer rates may be expressed in terms of a charge-transfer coefficient (α) that depends only on the temperature as

$$R_{\bar{I},\bar{J}\to\bar{U},\bar{L}} = \alpha_{\bar{I},\bar{J}\to\bar{U},\bar{L}}(T)\, n_{\bar{I}}\, n_{\bar{J}}$$

$$R_{\bar{U},\bar{L}\to\bar{I},\bar{J}} = \frac{\alpha_{\bar{I},\bar{J}\to\bar{U},\bar{L}}(T)}{\phi_{\bar{I},\bar{J},\bar{U},\bar{L}}(T)}\, n_{\bar{U}}\, n_{\bar{L}}$$
(3)

where $\phi_{\bar{I},\bar{J},\bar{U},\bar{L}}(T) = (n^*_{\bar{I}}\, n^*_{\bar{J}})/(n^*_{\bar{U}}\, n^*_{\bar{L}})$, the asterisk indicates the LTE value and $\bar{I} = (I,i)$ is an index vector specifying level i of ion I. The energy difference between the initial and final state of the process gives rise to heating or cooling of the electron gas with a rate given by $R_{\bar{I},\bar{J}\to\bar{U},\bar{L}}\ \left|\chi_{\bar{I},\bar{U}} - \chi_{\bar{L},\bar{J}}\right|$, where χ is the ionization energy.

2.6. Radioactive decays

2.6.1. Energy deposition

The energy released in the radioactive decays is carried by high-energy photons and leptons which deposit their energy in the ejecta mainly through Compton scattering on free and bound electrons. Although a detailed calculation is preferred, we use effective grey opacities determined through such calculations. We support the decay chains $^{56}\mathrm{Ni} \to {}^{56}\mathrm{Co} \to {}^{56}\mathrm{Fe}$, $^{57}\mathrm{Ni} \to {}^{57}\mathrm{Co} \to {}^{57}\mathrm{Fe}$ and $^{44}\mathrm{Ti} \to {}^{44}\mathrm{Sc} \to {}^{44}\mathrm{Ca}$, which are the most important for core-collapse SNe. For these decays we adopt the life-times and energies from Kozma & Fransson (1998) and the effective grey γ-ray opacities from J11, and assume that the positrons emitted are locally absorbed.

2.6.2. Energy partition

Through a cascade of collisions the deposited energy gives rise to a high-energy tail on the otherwise Maxwellian electron distribution. The shape of the non-thermal electron distribution and the fractions of the energy going into heating, excitation and ionization through non-thermal collisions can be calculated by solving the Spencer-Fano equation (Boltzman equation for electrons). This problem was solved by KF92 and for a further discussion we refer to this paper.

3. Method and design

Given the physical problem, we now describe the methods used to solve it, and provide an outline of how the code is designed. Except for the non-thermal solver, the code is written in C++, and the description therefore tends to reflect the object oriented structure of the code. The code is parallelized on a hybrid process (MPI) and thread (openMP) level, and we discuss this issue as well as the computational resources required in Sect. 3.7.

The SN ejecta are represented by a spatial grid of cells holding the local state of the matter and the radiation field. Although mostly geometry independent, the current version only supports spherically symmetric cells. To determine the state of the matter, JEKYLL provides several solvers with different levels of approximation (e.g. LTE and NLTE), and to determine the radiation field it provides a MC solver based on the Lucy method. As discussed, through Λ-iterations the matter and the radiation field are alternately determined from each other, a procedure which in JEKYLL is terminated after a fixed but configurable number of iterations. JEKYLL also provides a diffusion solver, intended for use at high optical depths where the matter and radiation field may be assumed to be in LTE. JEKYLL may be configured to run in steady-state or time-dependent mode, although the latter only applies to the radiative transfer. Steady-state breaks down if the diffusion time is large, and is therefore best suited for modelling in the nebular (optically thin) phase, or of the SN atmosphere in the photospheric (optically thick) phase.

3.1. Grid

The grid represents the SN ejecta and is spatially divided into a number of cells, which in the current version of the code are spherically symmetric. As mentioned, the code is mostly geometry independent, so cells with other geometries may easily be added in future versions. If macroscopic mixing is used, the cells may be further divided into compositional zones, geometrically realized as virtual cells. The grid provides functions to load the ejecta model, to load and save the state, as well as to export a broad range of derived quantities (e.g. opacities).

3.1.1. Cells

The cells hold the local state of the matter and the radiation field, and provide functions for the solvers to calculate derived quantities like opacities/emissivities and transition rates based on the local state and the atomic data. The local state of the matter is represented by the density, the temperature, and the number fractions of ionized and excited states. The local state of the radiation field is represented by the specific intensity, which is updated by the MC radiative transfer solver based on packet statistics following the method outlined by L03. In addition, JEKYLL supports simplified radiation field models based on pure or di-

luted blackbody radiation, given by $B_\nu(T_J)$ and $WB_\nu(T_R)$, respectively (see K09 for details).

JEKYLL also allows the radiation-field to be approximated by the source-function as $I = S(\mathbf{n}, T)$, which depends only on the local state of the matter. As discussed by Avrett & Loeser (1988), this approximation is essentially a generalized version of the classical on-the-spot approximation. The generalized on-the-spot approximation is intended for use bluewards ground-state ionization edges, which typically have high optical depths and dominate the source-function. When using this approximation, we only include bound-free opacities and emissivities in the source function, which is likely a good approximation.

3.1.2. Virtual cells

JEKYLL implements the concept of virtual cells, introduced by J11 to account for macroscopic mixing on a grid otherwise spherically symmetric. Each cell may be divided into zones occupying some fraction (filling factor) of the cell volume, and otherwise geometrically unspecified. These zones may have different densities and compositions, and the state is solved for separately by the matter-state solver. With respect to the MC-solver the zones are represented by virtual cells differing only in a geometrical and statistical sense. The virtual cells are spherical, have a size corresponding to some number of clumps, and their location is randomly drawn during the MC radiative transfer based on their size and the zone filling factor.

3.2. Atomic data

Once converted to the JEKYLL format, any set of atomic data may be loaded from file. The data is organized in a hierarchical structure of atoms, their isotopes and ions, and the bound states of the ions. Each ion holds a list of bound-bound transitions, and each atom holds a list of bound-free transitions. The atomic data also contains an (optional) list of charge-transfer reactions, which are mapped on two bound-free transitions, one recombination and one ionization. The specific atomic data used for the comparisons in Sects. 4.1-4.3 are discussed in Appendix A.1-A.3. The default choice, used for the application in Sect. 5, is inherited from SUMO (see J11 and J12), and has been extended as described in Appendix A.4.

3.3. MC radiative transfer solver

The MC radiative transfer solver determines the radiation field, and is based on the Lucy method. The radiation field is discretized as packets (Sect. 3.3.1), which are propagated on the grid (Sect. 3.3.2) and interact with the matter (Sect. 3.3.3). Note, that the packets are propagated in 3-D, so the constraint of spherical symmetry only applies to the grid they are propagated on. In the calculation, the constraints of statistical and thermal equilibrium are enforced, which accelerates the convergence of the Λ-iterations (see Sect. 2.1). The original method has been extended to include non-thermal ionizations and excitations, as well as charge-transfer and two-photon processes. In addition, we introduce an alternative, more efficient way to draw the emission frequency (Sect. 3.3.4), and a method to control the sampling of the radiation field (Sect. 3.3.5). Although we explain the basics, we refer to L02-L05 for the details of the original method.

3.3.1. Packets

The radiation field is discretized as packets, defined by their energy, frequency, position and direction. Following L03 and K09, we classify these as r-, i-, k- and γ-packets. The packets are indivisible and indestructible (but see Sect. 3.3.5 for a modified requirement), which enforce the constraint of thermal equilibrium on the MC calculation. Freely propagating photons are represented by r-packets, and upon absorption they are converted into i- and k-packets, representing ionization/excitation and thermal energy, respectively. The γ-packets are similar to the r-packets, but represent the γ-rays (or leptons) emitted in the radioactive decays, which are treated separately. Eventually, the i- and k-packets are converted into r-packets and re-emitted.

New r-packets are injected into the MC-calculation by sampling of the flux at the inner border (if any), and new γ-packets by sampling of the γ-ray emissivity. In addition, r-packets may be sampled from the initial intensity in each cell, as well as from the intensity in new cells taken over from the diffusion solver when the inner border is moved inwards.

3.3.2. Propagation

When the r- and γ-packets are propagated they undergo physical (radiation-matter interactions) and geometrical (border crossings) events. Whereas propagation is carried out in the rest frame, the physical events take place in the co-moving frame, and the packets are transformed back and forth to $O(v/c)$. After each event, a random optical depth for the next physical event is drawn as $\tau = - \ln z$, and the packet is propagated until the accumulated optical depth exceeds this value or a geometrical event occurs. Note that line-absorption may only occur at the resonance distance, and the (Sobolev) line-opacity may be regarded as a delta-function. In the case of a physical event, the packet is processed as described in Sect. 3.3.3, and in either case propagation continues as described above. Note, that in the case of γ-packets we use effective grey opacities (Sect 2.6.2), which differs from the more detailed procedure by L05. The r- and γ-packets leave the MC calculation by escaping through the outer border, where the r-packets are binned and summed to build the observed spectrum. When doing this we take light-travel time into account by defining the observers time as $t_O = t - (R/c)\,\mu$, where R is the radius of the grid and μ the cosine of the angle between the packet direction and the radius vector.

If the packet enters a cell with macroscopic mixing of the ejecta (Sect. 3.1.2), a randomly oriented virtual cell is drawn based on the filling factors for the compositional zones (see J11 for details). As long as the packet remains in the cell, the distance to the next geometrical event is given by the size and orientation of the virtual cell, and at each (virtual) border crossing the procedure is repeated.

3.3.3. Interactions

Once the packet has been absorbed, an interaction process is drawn in proportion to the opacities. As mentioned, the interactions take place in the co-moving frame, and in the case of (coherent) electron scattering, the frequency does not change. Otherwise, an emission frequency is drawn using the method described by L02 and L03, which enforces the constraints of statistical and thermal equilibrium on the MC calculation. Below we provide a summary of the original method and describe the extensions made for non-thermal, charge-transfer and two-photon
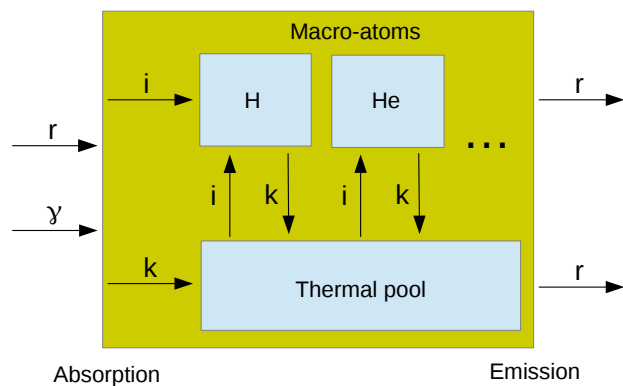

Fig. 1. Schematic figure of the MC state machine, showing the macro-atoms and the thermal pool, as well as the flows of r-, γ-, i- and k-packets.

processes. Before re-emission of the packet a new direction is drawn from an isotropic distribution.

Original method To enforce the aforementioned constraints on the MC calculation, L02 and L03 introduce the concepts of macro-atoms and the thermal pool, which are the MC analogues of the equations of statistical and thermal equilibrium. In computational terminology, a macro-atom is a (finite) state-machine, consisting of a set of states and the rules that govern internal transitions (between the states) and de-activation (of the machine). The states of a macro-atom corresponds to the energy levels of an atomic specie, and it is activated by ionizations and excitations and de-activated by recombinations and de-excitations. The rules that govern internal transitions and de-activation of a macro-atom are based on the equations of statistical equilibrium, and we give the details below. Together, the macro-atoms and the thermal pool constitute a single hierarchical state-machine, which we will refer to as the MC state-machine. In this context, the macro-atoms are sub-machines, and may be regarded as states with respect to a base-machine. In the base-machine, collisional activation and de-activation of the macro-atom sub-machines correspond to internal transitions between the thermal pool and the macro-atom states, and the rules that govern transitions in and out of the thermal pool are based on the equation of thermal equilibrium. The MC state-machine is activated by an absorption of an r- or γ-packet, and de-activated by the emission of an r-packet, and is illustrated in Fig. 1, where we also indicate the flows of i-, and k-packets within the machine.

Following the packets through the MC state-machine, the absorbed r- and γ-packets are converted into k- or i-packets in proportion to the energy going into heating and ionization/excitation[2]. In the former case, the k-packets are transferred to the thermal pool, and in the latter case, the i-packets activate the macro-atoms through radiative ionizations and excitations, drawn in proportion to their opacities. Eventually, the macro-atoms de-activates, and in de-activations through radiative transitions, the i-packets are converted to r-packets and re-emitted, whereas in de-activations through collisional transitions, the i-packets are converted to k-packets and transferred to the thermal pool. The k-packets enter the thermal pool through radiative and

---

[2] Note, that only the heating-channel is allowed for γ-packets in the original method.

collisional heating and leave through radiative and collisional cooling, in which case they are converted into r- or i-packets in proportion to the cooling rates. In the former case, the emission processes are drawn in proportion to their cooling rates and the r-packets are re-emitted, and in the latter case the i-packets activate the macro-atoms through collisional ionizations and excitations, drawn in proportion to their cooling rates. Note, that before the r-packets are re-emitted, their frequencies are drawn from the (normalized) emissivities of the de-activating processes.

Although the method is conceptually simple, it is a bit involved in the details, in particular with respect to the macro-atoms. As described in L02, the rules for the macro-atom state-machine are derived through a rewrite of the equations of statistical equilibrium in terms of energy. This leads to a number of terms that can be identified as the energy rates for activations, internal transitions and de-activations, and from this the probabilities can be calculated. A macro-atom is activated in level i by an physical upward transition (e.g. excitation) to this this level. Once in level i, each physical transition with number rate $R_{i \to j}$ corresponds to an internal state-machine transition with probability $P_{i \to j}^{I} \propto R_{i \to j} E_l$ (L02: Eq. 9), where E_l is the energy of level $l = min(i, j)$. In addition, each physical downward transition (e.g. de-excitation) may de-activate the macro-atom with probability $P_{i \to l}^{D} \propto R_{i \to l} (E_i - E_l)$ (L02: Eq. 7). If an internal transition is drawn, the state-machine proceeds to level j and the procedure is repeated. Note, that internal and de-activating transitions may result from several physical processes, and if not otherwise stated, $P_{i \to j}^{I}$ and $P_{i \to l}^{D}$ refer to the total probabilities for such transitions.

Non-thermal processes Upon absorption, γ-packets are converted into k- and i-packets in proportion to the energy going into heating and ionization/excitation. In the former case, the k-pakets are transferred to the thermal pool, and in the latter case, the i-packets activate the macro-atom state-machines by non-thermal transitions drawn in proportion to their energy rates. In the original method, only the heating-channel was allowed, and the addition of the ionization and excitation channels is one of our most important extensions to the Lucy method. The macro-atom state-machines are modified by adding non-thermal transitions, where the probabilities are calculated from their number rates as explained above. Non-thermal transitions are upward, and therefore correspond to internal transitions.

Charge-transfer processes As mentioned, charge-transfer is a collisional process that may be viewed as a recombination followed by an ionization, where the (small) energy difference results in either heating or cooling. The macro-atom state-machines are therefore modified by adding the corresponding ionizations and recombinations, where the probabilities are calculated from their number rates as explained above. Charge-transfer ionizations correspond to internal transitions, whereas charge-transfer recombinations correspond to internal and de-activating transitions.

De-activation of a macro-atom state-machine through a charge-transfer recombination results in either activation of another macro-atom state-machine through the corresponding ionization or in the conversion of the i-packet into a k-packet. The latter corresponds to the conversion of ionization energy into thermal energy, which may only happen if the reaction is exothermic, and is drawn in proportion to the energy going into heating. Correspondingly, if the reaction is endo-thermic, k-

packets may be converted into i-packets, in which case a macro-atom state-machine is activated by the corresponding ionization. This corresponds to the conversion of thermal energy into ionization energy, and is drawn in proportion to the cooling rate as described above.

Two-photon processes The macro-atom state-machines are modified by adding two-photon transitions, where the probabilities are calculated from their number rates as explained above. Two-photon transitions are downward, and might therefore be either internal or de-activating, and in the latter case the emission frequency is drawn from the (normalized) two-photon emissivity.

3.3.4. Markov-chain solution to the MC state-machine

A problem with the original method is that the number of transitions in the MC state-machine may become very large. This is particularly true when the collisional rates are high, causing the state-machine to bounce back and forth between macro-atoms and the thermal pool. To avoid this we use a statistical Markov-chain model to calculate the probability that the state-machine de-activates from a given state. This allows us to proceed to the state from which the state-machine de-activates in a single draw. A Markov-chain model can be constructed for the complete MC state-machine, as well as for its base- and sub-machine parts. Markov chain statistics have been used for MC radiative transfer before, e.g. for scattering in planetary atmospheres by Esposito & House (1978), but the application of it to the Lucy method described here is novel.

In the case of a macro-atom sub-machine, the state-machine consists of N states corresponding to the energy levels of an atomic specie, where the probabilities for internal and de-activating transitions, $P_{i \to j}^{I}$ and $P_{i \to j}^{D}$, are calculated as described in Sect. 3.3.3. In Markov-chain terminology, states are called transient if they are visited a finite number of times, recurrent if they are visited an infinite number of times, and absorbing if they can not be left. As the macro-atom will eventually de-activate, each state is only visited a finite number of times, and corresponds to a transient state in our Markov-chain model. In addition, we associate each transient state with an absorbing state, where a transition from a transient state into its associated absorbing state represents de-activation of the macro-atom (through any of the de-activating transitions).

The probabilities for transitions between transient states (i.e. internal transitions in the Lucy terminology) form the matrix $\mathbf{P}^{T}$, where $P_{i,j}^{T} = P_{i \to j}^{I}$ is the probability for a transition from the transient state i into the transient state j. The probabilities for transitions from transient states into their associated absorbing states (i.e. de-activations in the Lucy terminology) forms the diagonal matrix $\mathbf{R}$, where $R_{i,i} = \sum P_{i \to j}^{D}$ is the probability for a transition from the transient state i into its associated absorbing state i. From Markov-chain theory (see e.g. Ross 2007) we obtain the probability to end up in the absorbing state j given that we start in the transient state i as $(\mathbf{SR})_{i,j}$, where the matrix $\mathbf{S}$ is given by

$$\mathbf{S} = (\mathbf{I} - \mathbf{P}^{T})^{-1} \tag{4}$$

where $\mathbf{I}$ is the identity matrix. Therefore, to find the probability that the macro-atom de-activates from state j given that it was activated in state i, we simply look up the j:th element in
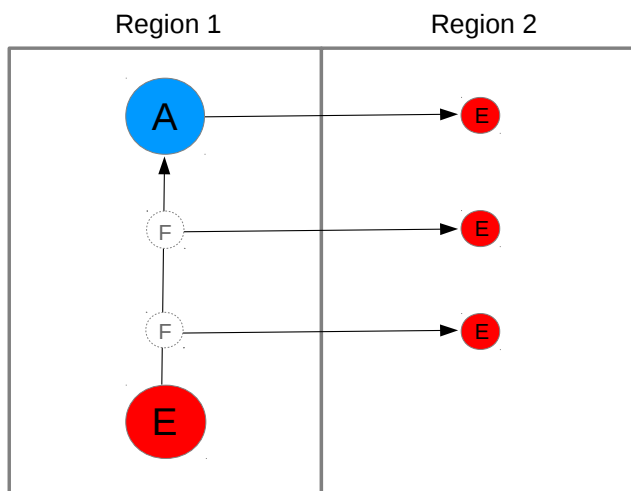
Region 1   Region 2

Fig. 2. Schematic figure of the re-sampling procedure for physical events. A packet emitted in region 1 with a large packet size (left), give rise to emission in region 2 with a small packet size (right) through fictitious absorption, and is eventually physically absorbed. Region 1 and 2 are assumed to be spatially coincident but separated in frequency. The circles indicates the size of the packets, emission is labelled with an "E", and physical and fictitious absorption is labelled with an "A" and an "F", respectively.

the i:th row of the **SR** matrix. Once the state from which the macro-atom de-activates has been drawn, we just draw the de-activating transition from their (normalized) probabilities. The implementation is based on two look-up tables for each state, one containing (a row of) **SR**, and one containing the (normalized) de-activating transition probabilities. Once a macro-atom is activated, the state from which it de-activates is drawn from the former table, the state-machine proceeds to this state, and the de-activating transition is drawn from the latter table.

The Markov-chain model for a macro-atom sub-machine described above is easily generalized to the complete MC state-machine, in which case the states are the energy levels of all macro-atoms plus the thermal pool. Note, however, that collisional activation and de-activation of the macro-atoms then correspond to internal transitions between the thermal pool and the macro-atoms. In case the method is applied to the complete MC state-machine, the **S** matrix has size $N \times N$, where N is the total number of energy levels for all macro-atoms plus one (the thermal pool), and the computational time to invert the matrix is a potential problem. This may be circumvented by splitting the MC state-machine into its base- and sub-machine parts, and calculate the corresponding **S** matrices separately. The procedure is similar to what is described above, but the computational time to invert the base- and sub-machine **S** matrices is much less than for the complete **S** matrix. We give the details on the split state-machine approach in Appendix B.

3.3.5. Packet sampling control

Another problem with the original method is that there is no (or limited) control of the number of packets as a function of frequency, space and time. This may result in too few packets, leading to noise in the radiation field estimators, or too many, leading to unnecessary computational effort. The number of packets can not be directly controlled, but this might instead be achieved by

adjusting the amount of energy they carry. Hereafter, we will refer to this as their size, and by conservation of energy the number of packets is inversely proportional to their size. We therefore introduce a method for continuous re-sampling of the radiation field through control of the packet size, which is allowed to vary as a function of frequency, space and time. This breaks the indivisibility and destructibility requirements introduced by L02, but conservation of energy, which is the essential physical property, is still maintained in an average sense.

A set of sampling regions (bounded in frequency, space and time) is defined, and each of these is assigned a packet size. When packets flow from one sampling region into another, their size is adjusted to that of the destination region. To maintain the rate of energy flowing into the destination region, the rate of packets flowing into it has to be adjusted with F, the ratio of the packet sizes in the source and destination regions. Consider now the series of events that occurs. In the original method, a leave event in the source region triggers an enter event in the destination region. However, as the number of leave and enter events in our method may differ, this procedure no longer apply. The simplest method to solve this, which we use for geometrical events, is to trigger F enter events for each leave event. More specifically, when crossing a border, the packets are split into F child packets[3] (if $F>1$) or terminated with probability $1-F$ (if $F<1$).

In the case of physical events, when the frequency changes, packets leave the source region through absorption and enter the destination region through emission, and we need to adjust the emission rates with F. A fundamental short-coming of the method used for geometrical events is that we are not adjusting the number of events in the source region, and this number may be small (or even worse, it may be zero). To overcome this, we introduce a class of fictitious absorption events, which occur F times as often as the physical absorption events, and which have the sole purpose to trigger the emission events. For an interaction process with physical opacity κ, we may then define a fictitious opacity κ_F corresponding to these fictitious absorption events. The interaction rate given by κ_F may be higher or lower than the one given by κ, and is the one required to produce the adjusted emission rates[4].

As we are adjusting the emission rate, whereas physical absorption needs to proceed at the original rate, the fictitious opacity to use is not κ_F but $max(\kappa_F, \kappa)$. Using this fictitious opacity, a packet is propagated as described in Sect. 3.3.2, and once an interaction is drawn, it is selected for absorption and emission with probabilities $max(\kappa/\kappa_F, 1)$, and $max(\kappa_F/\kappa, 1)$, respectively. If the packet is selected for absorption but not for emission it is terminated, and if the packet is selected for emission but not for absorption a child packet is created. Otherwise the interaction is handled as described in Sect. 3.3.3 (which also recovers the normal behaviour if $\kappa = \kappa_F$). As mentioned, the size of the emitted packets are always adjusted to that of the destination region. The re-sampling procedure for physical events outlined here is illustrated in Fig. 2, and as is possible to show, it gives the correct (average) energy flows in and out of each sampling region.

Although the basic idea is straightforward, the actual implementation is complicated by the way the emission frequency is drawn in the Lucy method (Sect. 3.3.3). In particular, we need to adjust the probabilities for all possible paths a packet may take through the MC state-machine. This is only possible if these are

---

[3] Ignoring here the fact that F is not an integer. In practice, the number of child packets is drawn to fulfil this condition in an average sense.
[4] Strictly speaking, this is only true as long as physical absorption proceeds at the original rate.

known, which is only the case if a Markov-chain solution to the MC state-machine (Sect. 3.3.4) has been obtained. We give the details of the implementation in Appendix C, where we explain how to adjust the MC state-machine probabilities, and how to calculate the fictitious opacity.

When using the method in JEKYLL, the number of packets is controlled by an adaptive algorithm, where the packet size in each sampling region is adjusted once per Λ-iteration. Assuming a Poisson distribution, the signal-to-noise ratio (SNR) in each sampling region is estimated based on the mean number of (physical and geometrical) packet events per frequency bin. Comparing this to a pre-configured target SNR, the packet size in each sampling region is adjusted based on the ratio of the target and estimated SNR. In addition, the packet size is bounded below and above by pre-configured minimum and maximum values.

3.4. Matter state solvers

To determine the state of the matter, JEKYLL provides the NLTE solver, as well as the more approximate LTE and Mazzali & Lucy (1993, hereafter ML93) solvers. It also provides an option to mix these solvers, e.g. by using the NLTE solver for the ionization and the LTE solver for the excitation, in a manner similar to what is done in ARTIS. In addition, JEKYLL provides a solver to determine the non-thermal electron distribution, used by the NLTE solver.

3.4.1. LTE solver

The LTE solver determines the state of the matter assuming that LTE applies. The populations of ionized and excited states are calculated using the Saha ionization and Boltzman excitation equations, respectively. The temperature used may be that associated with the pure or diluted blackbody radiation field models (T_J or T_R; see Sect. 3.1.1 and K09), or the matter temperature determined by some other method (e.g. thermal equilibrium).

3.4.2. ML93 solver

The ML93 solver determines the state of the matter assuming that the radiative rates dominate, and is based on the approximations for the populations of ionized and excited states derived by Mazzali & Lucy (1993) and Abbott & Lucy (1985). Following Mazzali & Lucy (1993), the temperature is assumed to be controlled by the radiation field and set to $0.9T_R$, where T_R is the temperature associated with the diluted blackbody radiation field model (see Sect. 3.1.1 and K09).

3.4.3. NLTE solver

The NLTE solver determines the state of the matter by solving the equations of statistical and thermal equilibrium for the level populations and the temperature, respectively. The solution is determined in two steps. First, thermal equilibrium is scanned for in a configurable temperature interval (centered on the solution from the previous Λ-iteration) using the bi-section method. In doing this, we solve for statistical equilibrium at each temperature step. Based on this estimate, thermal and statistical equilibrium are simultaneously iterated for until convergence is achieved, using a procedure similar to what is described by L03.

Statistical equilibrium The non-linear system of statistical equilibrium equations (Eq. 1) is solved by iteration on the level populations. In each step the system is linearised in terms of changes in the level populations, and the rates and their derivatives are calculated using the previous estimate of these. The linearised system is then solved for changes in the level populations using lower-upper (LU) decomposition and back-substitution. If all number derivatives (explicit and implicit) are included this is equivalent to a Newton-Raphson solver, but we leave this as a configurable choice, and in the simplest configuration only the explicit derivatives (i.e. rates per particle) are included.

The system of equations may be solved separately for the states of each atom, ignoring any coupling terms, or for all states at once. As the total number of states may be too large for a coupled solution, we provide the possibility to alternate a decoupled solution with a fully coupled solution for the ionization balance. Typically a decoupled solution works well, but charge-transfer reactions and the source-function radiation field model (see Sect. 3.1.1) may introduce strong coupling terms. Transition rates (Sects. 2.4 and 2.5) for bound-bound and bound-free radiative and collisional processes, as well as for non-thermal, charge-transfer and two-photon processes are all supported, but which ones to include is a configurable choice.

Thermal equilibrium The equation of thermal equilibrium (Eq. 2) is solved either using the bisection method (initial estimate) or Newton-Raphson's method (refined estimate), in which case an explicit temperature derivative is used. Heating and cooling rates (Sects. 2.4 and 2.5) for bound-bound and bound-free radiative and collisional processes, free-free processes, as well as non-thermal and charge-transfer processes are all supported, but which ones to include is a configurable choice.

3.4.4. Non-thermal solver

The non-thermal solver determines the non-thermal electron distribution resulting from the radioactive decays, and the fraction of the deposited energy going into heating, excitation and ionization. This is done by solving the Spencer-Fano equation (i.e. the Boltzman equation for electrons) as described in KF92. The fractions going into heating, excitation and ionization depend on the electron fraction, but as this dependence is rather weak a decoupled solution works well. The non-thermal solver is called by the NLTE solver at least twice each Λ-iteration (before each of its main steps), and otherwise whenever the electron fraction changes more than some pre-configured value.

3.5. Diffusion solver

The diffusion solver determines the temperature in each cell by solving the thermal energy equation assuming spherical symmetry, homologous expansion, LTE and the diffusion approximation for the radiative flux. This results in a non-linear system of equations for the temperature in each cell, which is solved by a Newton-Raphson like technique, similar to the one used by Falk & Arnett (1977). Two specific topics require some further discussion though; the Rosseland mean opacity used in the diffusion approximation, and the boundary where the diffusion solver is connected to the MC radiative transfer solver.

3.5.1. Opacity

The Rosseland mean opacity used in the diffusion approximation is calculated from the LTE state of the matter and the atomic data. This may sound straightforward, but the bound-bound opacity, and in particular the macroscopic mixing (see Sect 3.1.2) complicates things. In the latter case, if the clumps are all optically thin, the opacity may be calculated as an average over the compositional zones, but otherwise a geometrical aspect enters the problem. Therefore we calculate the Rosseland mean opacity using a Monte-Carlo method. In each cell a large number of packets are sampled based on the blackbody flux distribution and the filling factors for the compositional zones. These packets are then followed until they are absorbed, and their pathlength averaged to get the Rosseland mean free path. This gives the Rosseland mean opacity, including the bound-bound contribution, as well as the geometrical effects arising in a clumpy material.

3.5.2. Connecting boundary

If the diffusion solver is connected to the MC radiative transfer solver, appropriate boundary conditions must be specified for both solvers. As connecting boundary condition for the diffusion solver we use the temperature in the innermost cell handled by the MC radiative transfer solver. As connecting boundary condition for the MC radiative transfer solver we use the luminosity at this boundary determined with the diffusion solver. This is analogous to how the boundary between the diffusion and radiative transfer solvers is treated in Falk & Arnett (1977), except that in JEKYLL these calculations are not coupled and performed separately. To implement the connecting boundary condition for the MC radiative transfer solver we use an approximate method. During a time-step Δt, packets with total energy $L\Delta t$ are injected at the connecting boundary, whereas packets propagating inwards are simply reflected at this boundary. The frequency of the injected packets are sampled from a blackbody distribution at the temperature of the innermost cell.

3.6. Notes on the MC radiation field

As mentioned in Sect, 3.3.5, the sampling of the MC radiation field is a potential problem with the original method. In our experience this problem is most severe bluewards ~3000 Å and in the outer region of the ejecta. In principle, this could be solved by the method for packet sampling control, but in practice we use the generalized on-the-spot approximation (Sect. 3.1.1) bluewards ground-state ionization edges of abundant species (e.g. the Lyman break). The reason for this is twofold. First, to achieve a reasonable SNR in this region might require very large boost factors, which could potentially make the method for packet sampling control unstable. Second, in case a residual from two almost cancelling radiative rates (e.g. ionization and recombination in the Lyman continuum) is large enough to be important for the solution, even larger boost factors might be needed to achieve the required SNR. Due to this we use both packet sampling control and the generalized on-the-spot approximation in most of the simulations presented here and in Paper 2.

3.7. Computational resources and parallelization

JEKYLL is solving a complex problem with several thousands of independent quantities, so the computational resources (e.g. CPU time and memory) required to run a simulation are in-evitably large. To handle this, the code is parallelized and uses several methods to reduce the computational effort. The code is parallelized on a hybrid process (MPI) and thread (openMP) level, although the number of threads are limited by shared memory access. The CPU time required for the MC radiative transfer is (roughly) proportional to the number of MC packets. As these are independent, the MC radiative transfer is naturally parallelized on them, and the processing power scales nicely with the number of CPU cores up to some large number of MC packets.

The CPU time required for the NLTE solver is (roughly) proportional to the number of grid cells, and as these are independent, the NLTE solver is naturally parallelized on them. However, this limits the scaling of the processing power to the number of cells (which is much smaller than the number of MC packets). Therefore the code is further parallelized, which increases the scaling limit with a factor of at least a few. Note that the CPU time needed to solve the statistical equilibrium equations and to invert the Markov Chain S-matrix (see Sect. 3.3.4) is proportional to the third power of the number of energy levels, so the number of such levels is critical.

The CPU time needed for a typical simulation like the Type IIb model presented in Sect. 5 is a few thousand CPU hours, which using a few hundred CPU cores results in an execution time of about ten hours. The number of MC packets used (in each iteration) in such a typical simulation is a few millions, the number of grid cells between 50 and 100 and the number of atomic energy levels about ten thousand.

Due to the large number of atomic energy levels required for a realistic simulation, JEKYLL is memory intensive, in particular if the Markov-chain solution to the MC state-machine (Sect. 3.3.4) is used. While the MC packets as well as the processing of them are distributed over the available processes, only the processing of the grid cells is distributed. The reason for this is that during the radiative transfer, the MC packets potentially require access to all grid cells. This may be improved in future versions, but currently a typical model like the Type IIb model presented in Sect. 5 requires about 1 GB of memory per cell and node, which limits the total number of cells that can be used in a simulation.

JEKYLL uses several methods to reduce the computational effort. The most important are the Markov-Chain solution to the MC state-machine (Sect. 3.3.4) and the use of a diffusion solver in the inner region (Sect. 3.5). The speed-up achieved from the Markov-Chain solution to the MC state-machine is considerable and is discussed in Sect. 5.5.3. However, as mentioned above, this method comes with the caveat of an increased use of memory resources. The speed-up achieved from the use of a diffusion solver in the inner region is essential for the code, and potentially huge at early times when the optical depths are large.

4. Code comparisons

In this section we compare JEKYLL to ARTIS (S07 and K09), SUMO (J11 and J12), and CMFGEN (H98), three codes which have similar, but not identical capabilities. ARTIS provides a good test of the time-dependent MC radiative transfer, which is very similar, but only supports partial NLTE. SUMO on the other hand, provides a good test of the full NLTE problem, but requires steady-state, so no test of the time-dependent MC radiative transfer is possible. However, CMFGEN, which is similar to JEKYLL in physical assumptions but different in technique, does provide a test of the full time-dependent NLTE problem. In particular, as it solves the coupled radiation-matter problem and
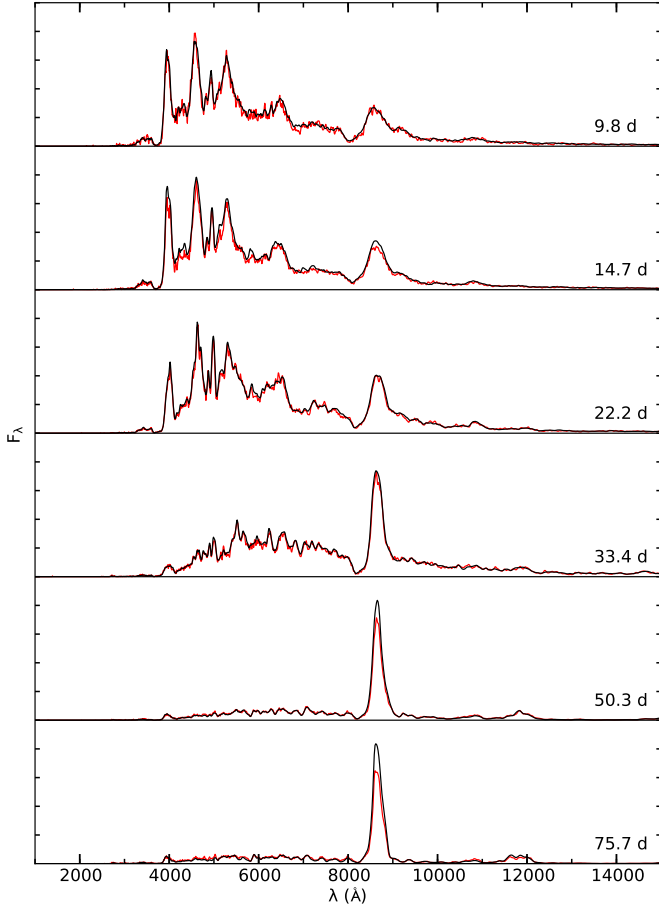
Fig. 3. Comparison of spectral evolution for model 12C as calculated with JEKYLL (black) and ARTIS (red). **For this comparison both codes use LTE estimates for the population and the ionization state of the gas.**

does not rely on Λ-iterations, it provides a way to show that the Lucy method actually converges to the correct solution. Here we present a comparison for a somewhat simplified test case, which still provides a good test of the full time-dependent NLTE problem. The comparisons to ARTIS, SUMO and CMFGEN are complementary, and taken together they provide a thorough test of the JEKYLL code.

4.1. Comparison with ARTIS

ARTIS is a spectral synthesis code aimed for the photospheric phase presented in S07 and K09. Both ARTIS and JEKYLL are based on the Lucy method, but ARTIS only supports a simplified NLTE treatment[5], where the excited states are populated according to LTE and the energy deposited by the radioactive decays goes solely into heating. On the other hand, the current version of JEKYLL assumes a spherical symmetric geometry, which is not a limitation in ARTIS. In addition, ARTIS calculates the deposition of the radioactive decay energy by Compton scattering, photo-electric absorption and pair production, whereas JEKYLL uses effective grey opacities (based on such calculations). There are also differences in the NLTE ionization treatment, in partic-

---

[5] A more general NLTE treatment and the inclusion of non-thermal processes in ARTIS are currently under development.
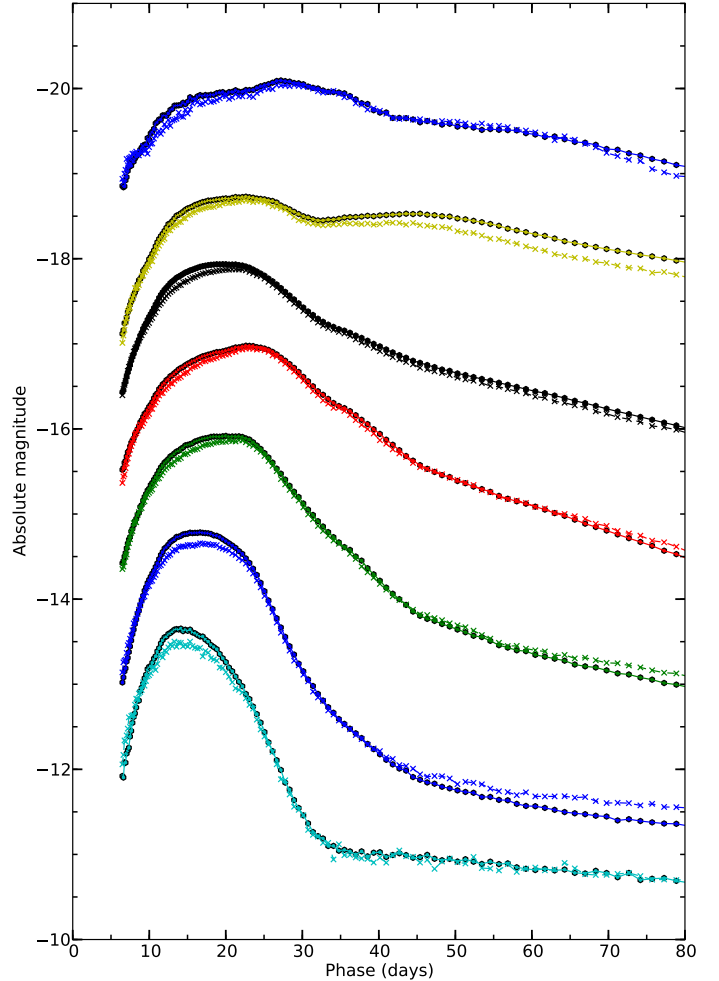
Fig. 4. Comparison of broadband and bolometric lightcurves for model 12C as calculated with JEKYLL (solid lines and circles) and ARTIS (dashed lines and crosses). From bottom to top we show the U (cyan), B (blue), V (green), R (red), bolometric (black), I (yellow) and J (blue) lightcurves, which for clarity have been shifted with 2.0, 2.0, 1.5, 0.5, -1.0, -1.0 and -3.0 mags, respectively

ular with respect to the calculation of photo-ionization rates, and due to this we decided to run ARTIS in its LTE mode. This still allows for a complete test of the time-dependent MC radiative transfer, which is the main purpose of the ARTIS comparison.

For the comparison we use the Type IIb model 12C from J15, which is also used for the application to Type IIb SNe in Sect. 5. The original model was converted to microscopically mixed form and re-sampled to a finer spatial grid as described in Sect. 5.1. To synchronize JEKYLL with ARTIS, it was configured to run in time-dependent (radiative transfer) mode using the LTE solver, and the ARTIS atomic data was automatically converted to the JEKYLL format. The details of the code configurations and the atomic data used are given in Appendix A.1, and we find the synchronization good enough for a meaningful comparison. Note that as non-thermal processes are crucial for the population of the excited He I states, the characteristic He I signature of Type IIb SNe is not reproduced.

In Figs. 3 and 4 we compare the spectral evolution and the lightcurves, respectively, whereas in Fig. 5 we compare the evolution of the temperature and the electron fraction. As the grey approximation used in ARTIS (see Appendix A.1) affects the
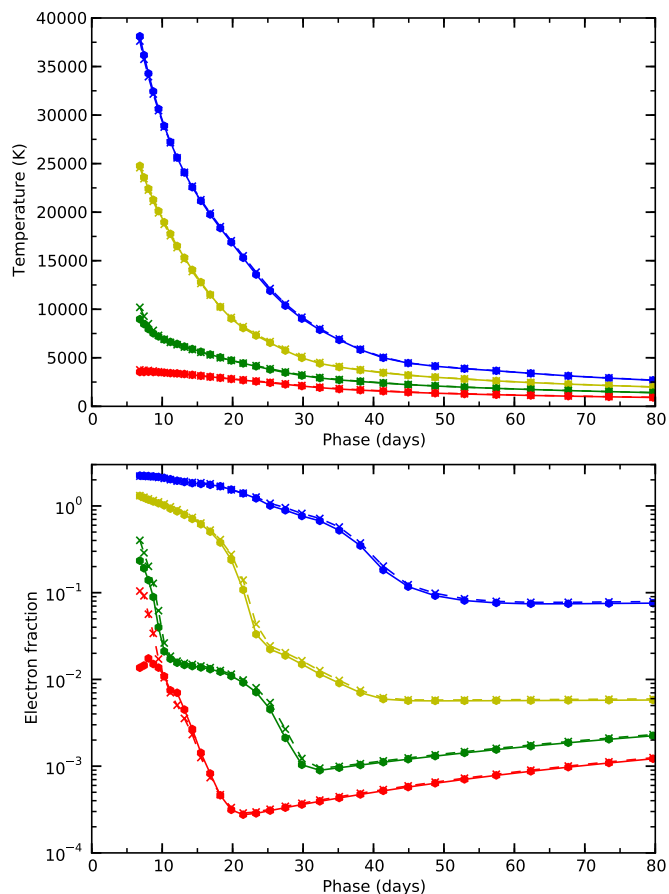
Fig. 5. Comparison of the evolution of the temperature (upper panel) and electron fraction (lower panel) in the oxygen core (blue), inner/outer (yellow/green) helium envelope and the hydrogen envelope (red) for model 12C as calculated with JEKYLL (circles and solid lines) and ARTIS (crosses and dashed lines).
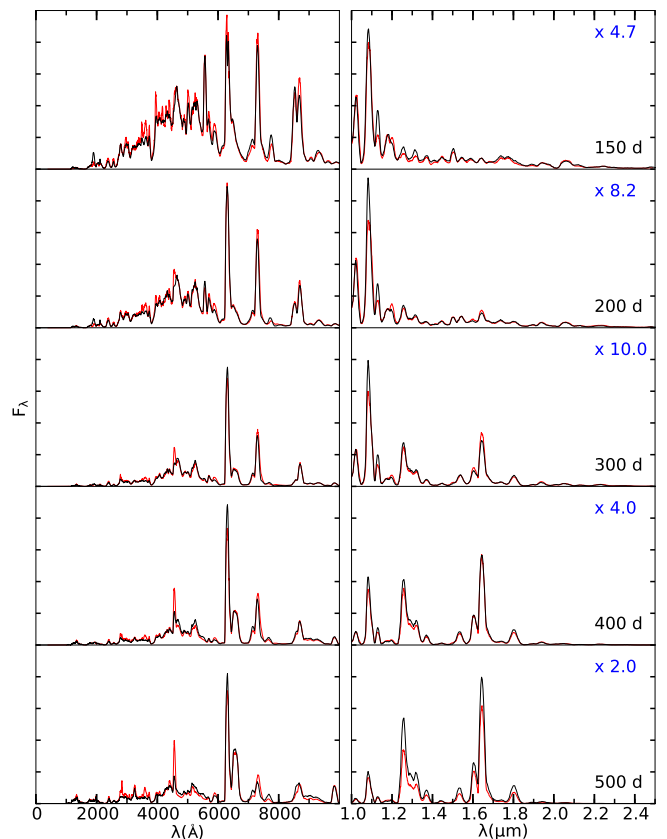

Fig. 6. Comparison of optical (left panel) and NIR (right panel) spectra for model 13G at 150, 200, 300, 400 and 500 days as calculated with JEKYLL (black) and SUMO (red). For clarity the NIR spectra have been scaled with respect to the optical spectra with the factor given in the upper right corner.

early evolution, we compare the models after 6 days, although the effect seems to last for a few more days in some quantities (e.g. the electron fraction). As can be seen, the general agreement is good in both the observed and the state quantities. The most conspicuous discrepancy appears in the Ca II 8498,8542,8662 Å line after ~40 days, and gives rise to a ~15 percent discrepancy in the *I*-band lightcurve. Another discrepancy appears after ~50 days in the *B*-band, growing towards ~15 percent at 80 days. There is also a small (<5 percent) but clear difference in the bolometric tail luminosity, reflecting a similar difference in the radioactive energy deposition. This is due to the the more approximate method for this used by JEKYLL, which may also explain the differences in the tail broad-band lightcurves. There are also minor differences in the diffusion peak lightcurves, most pronounced in the *U*- and *B*-bands, which could be related to the simplified treatment at high optical depths in ARTIS and JEKYLL (grey approximation and diffusion solver, respectively; see Appendix A.1). Summarizing, although there are some minor differences in the spectra and the lightcurves, we find the overall agreement to be good.

4.2. Comparison with SUMO

SUMO is a spectral synthesis code aimed for the nebular phase presented in J11 and J12. Similar to JEKYLL, it uses a Λ-

iteration scheme, where the radiative transfer is solved with a MC method and the state of the matter determined from statistical and thermal equilibrium. Except for the steady-state assumption (for the radiative transfer), which is required by SUMO and an option in JEKYLL, the main difference between SUMO and JEKYLL is the MC technique used. Whereas JEKYLL is based on the Lucy method, where conservation of packet energy is enforced, SUMO uses another approach. Except for electron scattering and excitations to high lying states, the packet energy absorbed in free-free, bound-free and bound-bound processes is not re-emitted. As long as these processes are included in the emissivity from which the packets are sampled, this gives the correct solution in the limit of convergence. However, it could be an issue for the rate of convergence, in particular at high absorption depths, and the method is probably not suited for the photospheric phase. There are also a few differences in the physical assumptions. Whereas JEKYLL correctly samples the frequency dependence of the bound-free emissivity, this is done in a simplified manner for all species but hydrogen by SUMO. On the other hand, JEKYLL does not take the escape probability from continua and other lines in the Sobolev resonance region into account. Yet another difference is that SUMO does not include the highest lying states in the NLTE solution (see Jerkstrand et al. 2012). However, in general the physical assumptions are similar.

For the comparison we use the Type IIb model 13G from J15, and run models with JEKYLL at 150, 200, 300, 400 and 500
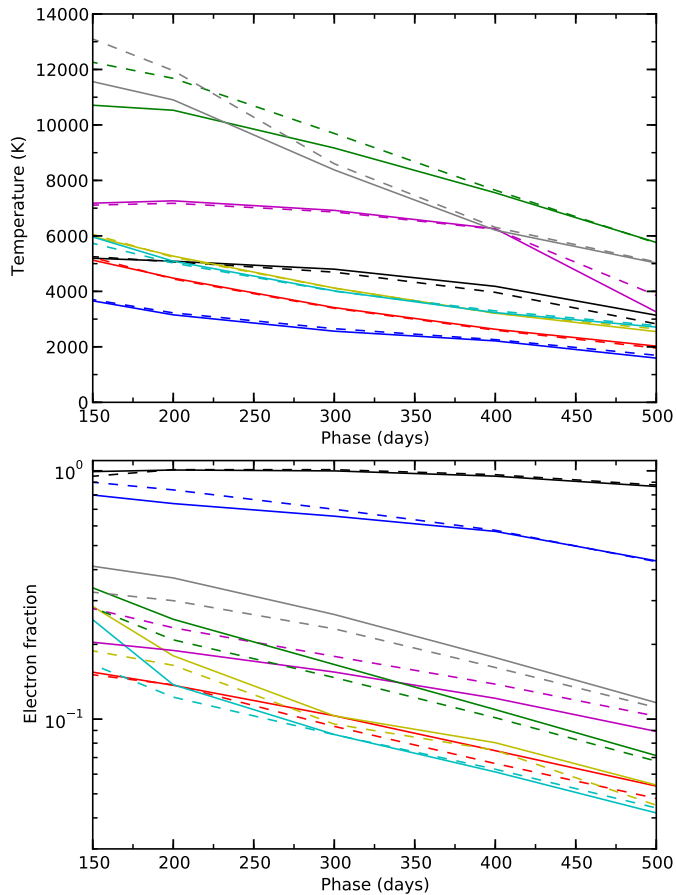
Fig. 7. Comparison of the evolution of the temperature (upper panel) and the electron fraction (lower panel) for model 13G in the Fe/Co/He (black), Si/S (blue), O/Si/S (red), O/Ne/Mg (yellow), O/C (cyan), He/C (magenta), He/N (green) and H (grey) zones as calculated with JEKYLL (solid lines) and SUMO (dashed lines).
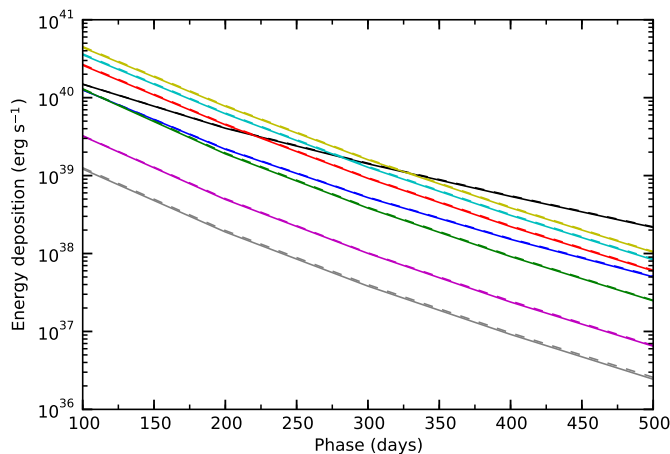

Fig. 8. Comparison of the evolution of the radioactive energy deposition for model 13G in the Fe/Co/He (black), Si/S (blue), O/Si/S (red), O/Ne/Mg (yellow), O/C (cyan), He/C (magenta), He/N (green) and H (grey) zones as calculated with JEKYLL (solid lines) and SUMO (dashed lines).

days. To synchronize JEKYLL with SUMO, it was configured to run in steady-state mode using the NLTE solver, and we have tried to synchronize the atomic data as much as possible. The

details of the code configurations and the atomic data used are given in Appendix A.2, and although not complete, we find the synchronization good enough for a meaningful comparison.

A comparison of the spectral evolution is shown in Fig. 6, and in Figs. 7 and 8 we compare the evolution of the temperature, the electron fraction and the radioactive energy deposition in each of the different nuclear burning zones (see J15). As can be seen, the general agreement of the spectra is quite good, although the match is slightly worse at 500 days. The largest discrepancies are seen in the Mg I] 4571 Å line, the O I 11290,11300 Å line before 300 days, the He I 10830 Å line at 200-300 days, and a number of features originating from the Fe/Co/He zone at 500 days. That one of the largest discrepancies is seen in the Mg I] 4571 Å line is not surprising as magnesium is mainly ionized and the Mg I fraction is small (see J15). This makes the strength of the Mg I] 4571 Å line sensitive to this fraction, in turn sensitive to the network of charge transfer reactions.

The evolution of the temperature shows a good agreement and the differences are mainly below ~5 percent. An exception is the He/N and H zones at early times, and in particular at 150 days where the difference is ~15 percent. The evolution of the electron fraction shows a worse agreement, but the differences are mainly below ~10 percent. Again, the agreement is worst at early times, and in particular at 150 days when the electron fractions in the O/Ne/Mg and O/C zones differ by ~30 percent. This discrepancy is reflected in e.g. the O I 11290,11300 Å line discussed above, but in general the spectral agreement at 150 days is quite good.

The evolution of the radioactive energy deposition shows an excellent agreement. This shows that the radiative transfer of the γ-packets (Sect. 3.3.1 and 3.3.2), representing the γ-rays (and leptons) emitted in the radioactive decays, as well as the method for macroscopic mixing of the material (Sect. 3.1.2 and 3.3.2), works as intended. Summarizing, although there are some notable differences both in the spectra and the state variables, we find the overall agreement to be good, in particular as the data and the methods are not entirely synchronized.

4.3. Comparison with CMFGEN

CMFGEN is a general purpose spectral synthesis code presented in its steady-state version in H98, and extended with time-dependence in Dessart & Hillier (2008, 2010) and Hillier & Dessart (2012) and non-thermal processes in Dessart et al. (2012). It is similar to JEKYLL in the physical assumptions, but uses a different method to solve the NLTE problem, where the coupled system of differential equations for the matter and the radiation field is solved by a linearization technique. The potential difficulties with convergence in Λ-iteration based methods (Sect. 2.1) are therefore avoided, and the comparison provides a good test of the convergence properties of the Λ-iteration and MC based method used in JEKYLL (i.e. the one by Lucy). The main difference in the physical assumptions is that JEKYLL assumes steady state for the matter, whereas CMFGEN does not. In addition, CMFGEN does not rely on the Sobolev approximation, but this is likely of less importance at the high velocity gradients present in SN ejecta. Note, however, that in the Sobolev approximation, absorption in continua and other lines within the resonance region is ignored, which is a potential problem. **On the other hand, the choice of micro-turbulent velocity in CMFGEN might lead to an overestimate of the overlap between lines.**
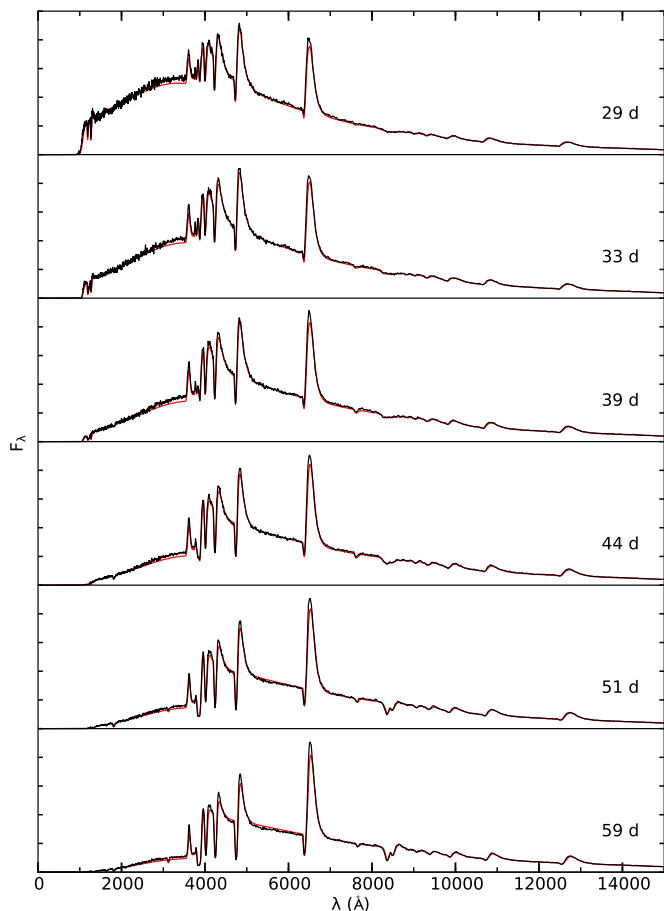
Fig. 9. Comparison of spectral evolution for the test model as calculated with JEKYLL (black) and CMFGEN (red).

Fig. 10. Comparison of the evolution of the temperature (upper panel) and electron fraction (lower panel) at the same epochs as in Fig. 9 for the test model as calculated with JEKYLL (black) and CMFGEN (red). The border between the regions handled by the diffusion solver and the MC radiative transfer solver have been marked with black circles.

To synchronize with JEKYLL, time-dependence for the matter needs to be switched off in CMFGEN. However, as this turned out to be difficult, we have instead added support for limited time-dependence in JEKYLL. This is done by adding an option to use the more general NLTE rate-equations, which is achieved by adding the time-derivative

$$\rho \frac{D(n_{I,j}/\rho)}{Dt} \qquad (5)$$

to the right-hand side of Eq. 1 (see Dessart & Hillier 2008). This accounts for the effect of time-dependence on the degree of ionization, which is the most important one, at least in the test model. Note that time-dependence is only added in this limited form to facilitate the comparison with CMFGEN, and is not explored further in the paper. A general upgrade of JEKYLL to full time-dependence will be presented in a forth-coming paper.

For the comparison we use a model of a red supergiant of 15 $M_\odot$ initial mass, evolved with MESA (Paxton et al. 2011, 2013) and exploded with an energy of 1 Bethe with the hydrodynamical code HYDE (E14). The JEKYLL and CMFGEN simulations begin at 25 days, and for the test we use a simplified composition consisting of hydrogen, helium, oxygen and calcium. To synchronize, the CMFGEN atomic data was automatically converted to the JEKYLL format, and the details of the code configurations and the atomic data used are given in Appendix A.3,

A comparison of the spectral evolution is shown in Fig. 9, and in Fig. 10 we compare the evolution of the temperature
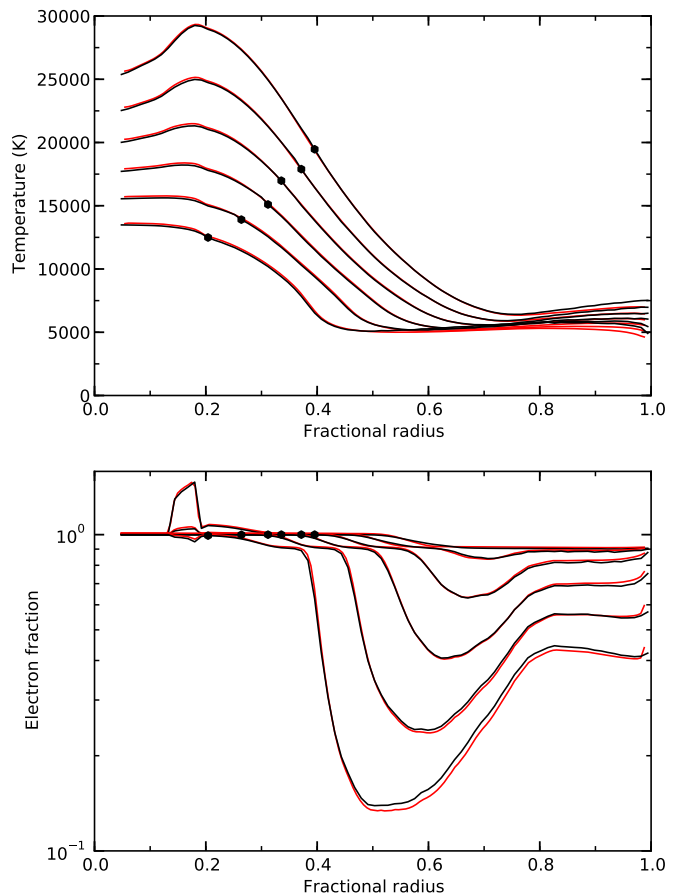
and the electron fraction. As can be seen, the overall agreement is good in both the spectra and the matter quantities. The largest differences in the spectra are a somewhat higher flux in the Balmer continuum and a bit stronger emission in the Balmer lines in the JEKYLL model. The electron fraction is in good agreement, but the temperature is slightly higher in the outer region in the JEKYLL model. Given that time-dependence is only partly implemented in JEKYLL, and is missing in the thermal energy equation, differences at this level are not surprising.

The good overall agreement found in both the spectra and the matter quantities shows that the Λ-iteration and MC based method used in JEKYLL (i.e. the one by Lucy) does indeed converge to a solution close to the correct one[6], at least in the specific case tested here. Departures from LTE are large (typically a factor of ten or larger) in the optically thin region, so although based on a model with simplified composition, the comparison provides a good test of the time-dependent NLTE capabilities.

---

[6] Note, that CMFGEN may not be free of bugs, and may also have other short-comings, so this has to be taken with a grain of salt. In addition, physics that might be required to correctly model SN ejecta (as e.g. 3-D) are still missing in both CMFGEN and JEKYLL.

5. Application and tests

In this section we provide an example of a time-dependent NLTE model based on a fully realistic ejecta model. The ejecta model (12C) is taken from the set of Type IIb models constructed by J15, and was found to give the best match to the observed nebular spectra (J15) and lightcurves (E15) of SN 2011dh. For a more detailed discussion of the model and a comparison to SN 2011dh we refer to Paper 2, where we also explore other models.

In addition to a brief discussion of the model and its evolution, we investigate the effect of NLTE on the spectra and the lightcurves, in particular with respect to non-thermal ionization and excitation. Based on the model, we also investigate the convergence of the Λ-iterations, and provide tests of our most important (technical) extensions to the original Lucy method, i.e. the use of a diffusion solver in the inner region, the packet sampling control and the Markov-chain solution to the MC state-machine.

5.1. Ejecta model

A full description of the Type IIb model 12C is given in J15, but we summarize the basic properties here. It is based on a model by Woosley & Heger (2007) with an initial mass of 12 M$_\odot$, from which we take the masses and abundances for the carbon-oxygen core and the helium envelope. We assume the carbon-oxygen core to be fully mixed and to have a constant (average) density, and the helium envelope to have the same density profile as the best-fit model for SN 2011dh by Bersten et al. (2012). In addition, a 0.1 M$_\odot$ hydrogen envelope based on models by Woosley et al. (1994) is attached. Note that the ejecta model explored here is a microscopically mixed version, in which the abundances in the nuclear burning zones (see J15) have been averaged. For the macroscopically mixed version we refer to Paper 2, where we also discuss the effect of this difference on the model evolution. To be suitable for modelling in the photospheric phase, the original ejecta model has also been re-sampled to a finer spatial grid.

5.2. Model evolution

JEKYLL was configured to run in time-dependent (radiative transfer) mode, using the NLTE solver based on an updated version of the J15 atomic data, and we give the details of the configuration and the atomic data in Appendix A.4. The model was evolved from 1 to 100 days, and the initial temperature profile was taken from the best-fit model for SN 2011dh from E15. Figure 11 shows the spectral evolution, whereas Fig. 12 shows the lightcurves and Fig. 13 the evolution of the temperature and the electron fraction. In Fig. 11 we also display the process giving rise to the emission, based on the last emission events for the MC packets (excluding electron scattering).

The main signature of a Type IIb SN is the transition from a hydrogen to a helium dominated spectrum, and this is well reproduced by the model. Initially, the hydrogen lines are strong and emission from the hydrogen envelope is dominating. Between 10 and 15 days the helium lines appear, grow stronger, and eventually dominate the spectrum at ~40 days. Hydrogen line emission disappears on a similar time-scale, completing the transition, although the Balmer lines remain considerably longer in absorption. After ~40 days the carbon-oxygen core gets increasingly transparent and the amount of realism in the microscopically mixed model starts to degrade, exemplified by the strong calcium NIR triplet at 100 days. As is demonstrated in Paper 2, the macroscopically mixed version of the model does a
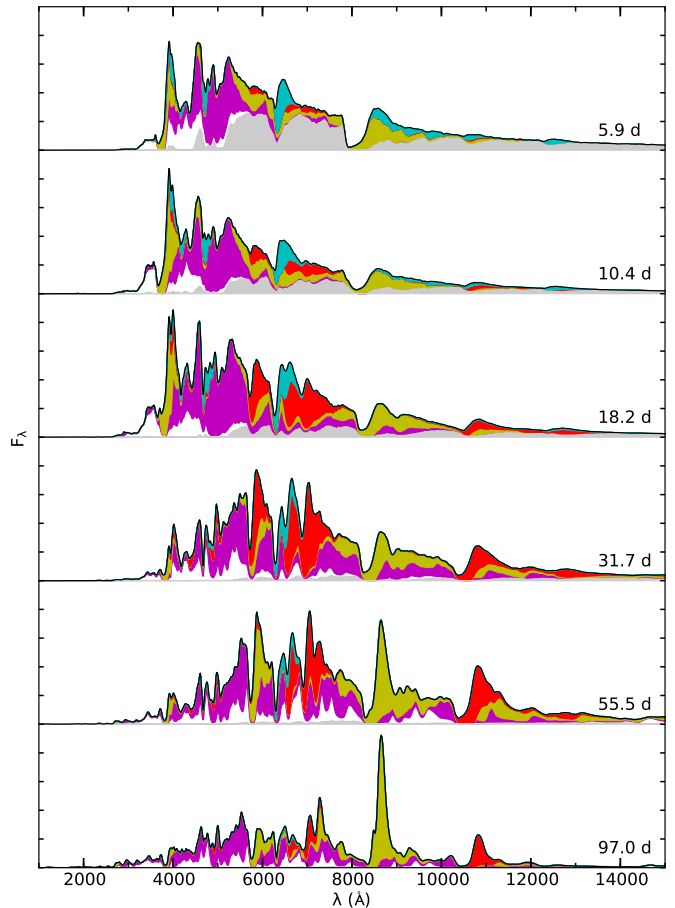
Fig. 11. Spectral evolution for model 12C as calculated with JEKYLL. In the spectra we show the contributions to the emission from bound-bound transitions of hydrogen (cyan), helium (red), carbon-calcium (yellow), scandium-manganese (white) and iron-nickel (magenta) as well as continuum processes (grey).

considerable better job in reproducing observations in the nebular phase.

The lightcurves show the characteristic bell shape of stripped-envelope (SE; Type IIb, Ib and Ic) SNe, and as discussed in Paper 2, their change in shape with effective wavelength (as e.g. a broader peak for redder bands) is in good agreement with observational studies. It is worth noting that the behaviour of model 12C is similar to that of the NLTE models of SE SNe presented by Dessart et al. (2015, 2016). Those models were evolved with CMFGEN, and in particular the Type IIb model 3p65Ax1 shares many properties with model 12C.

5.3. The effect of NLTE

Figures 14 and 15 show the bolometric lightcurve and the spectral evolution of model 12C calculated with JEKYLL with and without non-thermal ionization and excitation. Before 10 days both the bolometric lightcurve and the spectral evolution are very similar, after which they start to differ in several aspects. This turning point coincides with the time when the radioactive energy deposition becomes important outside the photosphere (see Paper 2). The most striking difference in the spectral evolution is the absence of (strong) helium lines in the model without non-thermal processes. This well-known effect was pointed out al-
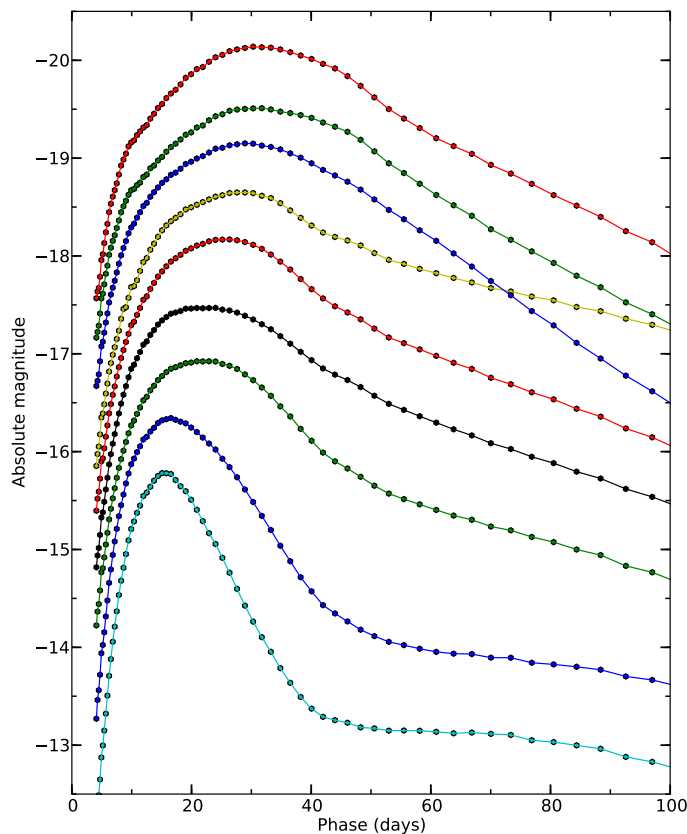
Fig. 12. Broadband and bolometric lightcurves for model 12C as calculated with JEKYLL . From bottom to top we show the U (cyan), B (blue), V (green), R (red), bolometric (black), I (yellow), J (blue), H (green) and K (red) lightcurves, which for clarity have been shifted with 0.4, 0.0, -0.9, 0.0, -0.9, -1.4, -1.7, -2.1 and -2.4 mags, respectively.

Fig. 13. Evolution of the temperature (upper panel) and electron fraction (lower panel) in the oxygen core (blue), inner/outer (yellow/green) helium envelope and the hydrogen envelope (red) for model 12C.

ready by Lucy (1991), and was later confirmed using CMFGEN by Dessart et al. (2012). Non-thermal excitation and ionization are essential to populate the excited levels of He I, which is in turn required to produce the lines observed. As discussed by Lucy (1991), the population process is subtle, as ionization of He I is amplified by photo-ionization from the excited levels, which proceeds at a rate far exceeding the non-thermal one.

Less known is the quite strong effect on the bolometric lightcurve, where the diffusion peak of the model with non-thermal processes is considerably broader. The reason for this is the increased degree of ionization, and therefore the increased electron scattering opacity. This is illustrated by Fig. 16, which shows the evolution of the electron fraction in the carbon-oxygen core and the helium and hydrogen envelopes. In the model with non-thermal processes, the electron fraction in the helium envelope drops much slower than in the model without. Due to the lower ionization potential, the effect is much less pronounced in the carbon-oxygen core and the hydrogen envelope. In Fig. 14 we also show the bolometric lightcurve for the LTE version of model 12C used for the comparison with ARTIS in Sect. 4.1. This model shows an even narrower bolometric lightcurve, which is related to an even lower degree of ionization. Note, that the atomic data for this model differs somewhat from that used for the NLTE models in this section, but this difference does not significantly affect the bolometric lightcurve.
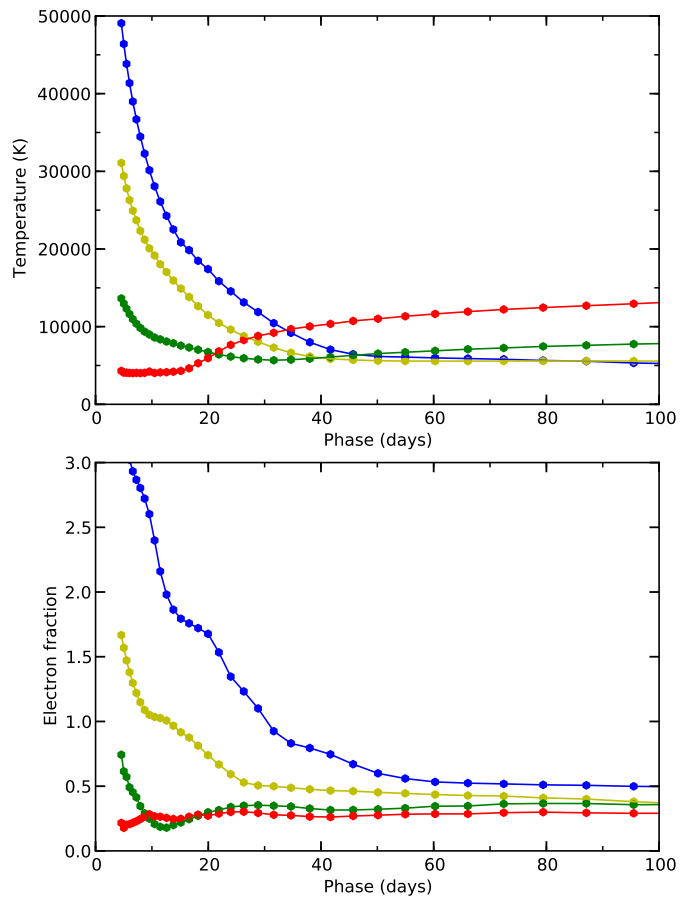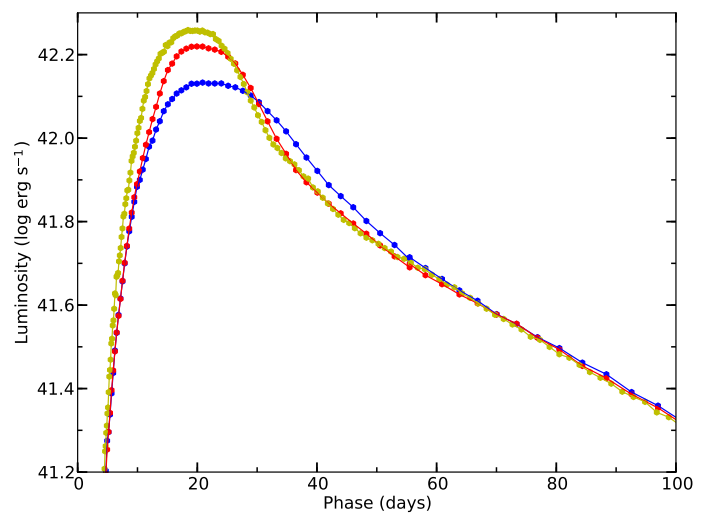

Fig. 14. Bolometric lightcurve for model 12C calculated with (blue) and without (red) non-thermal ionization and excitation. We also show the bolometric lightcurve for the LTE version of model 12C (yellow) presented in Sect. 4.1.

5.4. Convergence of the Λ-iterations

As mentioned in Sect. 3, JEKYLL uses a fixed (but configurable) number of Λ-iterations. In time-dependent (radiative transfer)
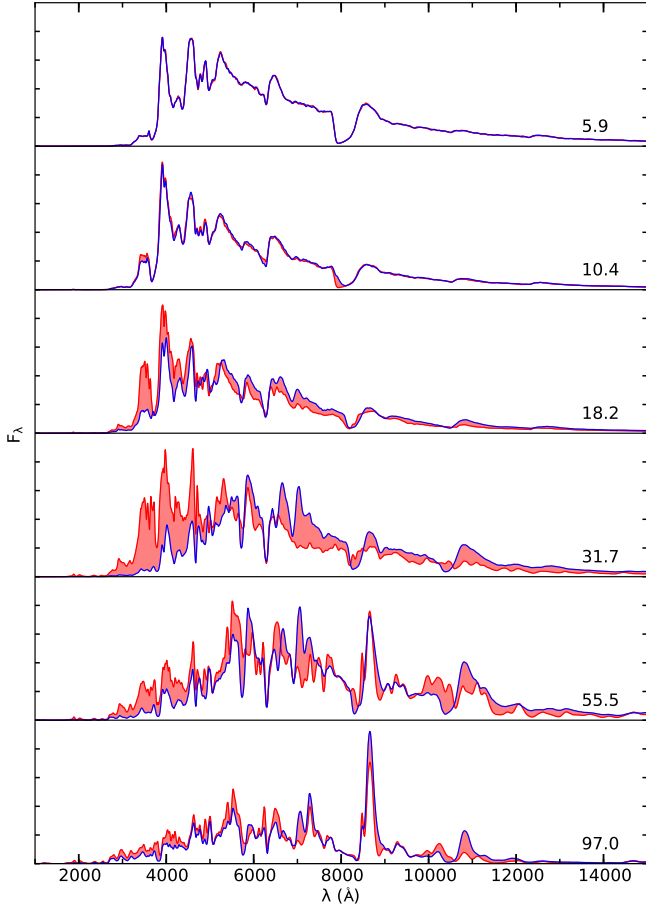
Fig. 15. Spectral evolution for model 12C calculated with (blue) and without (red) non-thermal ionization and excitation, where the difference has been highlighted in shaded red.
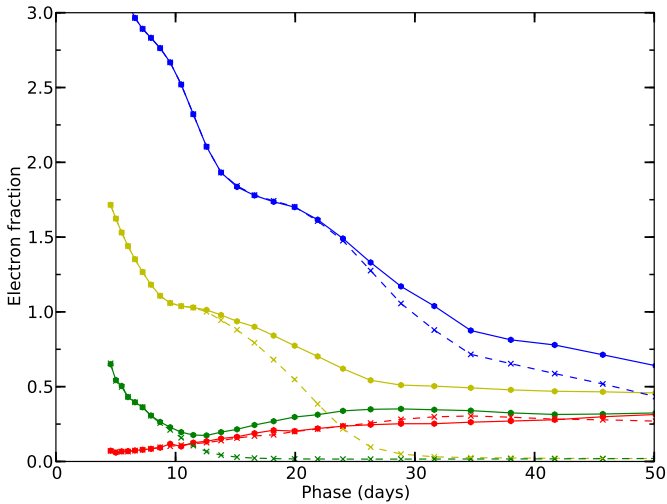

Fig. 16. Evolution of the electron fraction in the oxygen core (blue), inner/outer (yellow/green) helium envelope and the hydrogen envelope (red) for model 12C calculated with (circles and solid lines) and without (crosses and dashed lines) non-thermal ionization and excitation.

mode, this is the number of Λ-iterations per time-step, and corresponds to some (unknown) number of effective Λ-iterations depending on the length of the time-step and the rate at which the
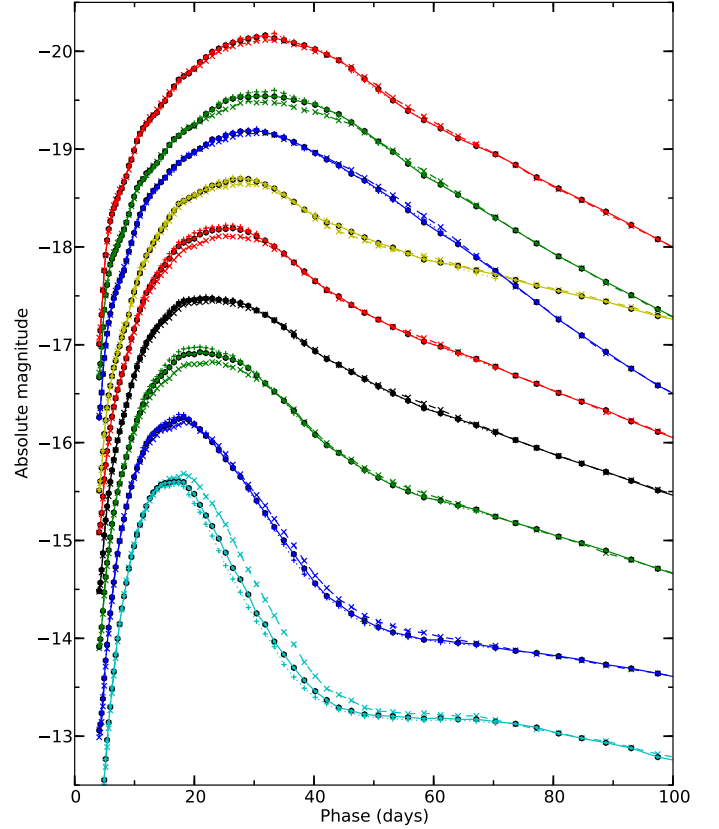

Fig. 17. Broad-band and bolometric lightcurve for model 12C as calculated with JEKYLL using two (dashed lines and crosses), four (solid lines and circles) and eight (dotted lines and pluses) Λ-iterations per time-step. Otherwise as described in Fig. 12.

state is changing. The time-dependent NLTE run in this section uses a logarithmic time-step of 5 percent and four Λ-iterations per time-step. In Figs. 17 and 18 we show the lightcurves and the temperature and electron fraction, respectively, for three such runs using two, four and eight Λ-iterations per time-step. Note, that to speed up the calculations we used coarser spatial sampling, slightly simplified atomic data and fewer packets than in the original model. As can be seen, convergence is fast, and more than four iterations per time-step does not make a significant difference. Even two iterations are good enough for most purposes, although there is a ∼25 percent difference in the U-band during the drop from the peak onto the tail.

To further investigate the convergence, we have plotted the correction per Λ-iteration for the lightcurves in Fig. 19. The corrections have been averaged between 15 and 40 days, where we have included the sign in the average to reduce the MC noise. As exemplified by the U and V-band, which are the bands with the largest deviations, convergence appears to be geometric with a ratio of successive corrections close to 1/2. Given this ratio, the remaining error equals the last correction, which is less than one percent in all bands after eight Λ-iterations.

The comparisons to ARTIS in Sect. 4.1 behave in a similar way, but the shorter time-step of 1 percent, and possibly a faster convergence in the LTE case, make even a single-iteration run well converged, showing less discrepancy than the two-iterations run in Figure 17.

The comparisons to CMFGEN in Sect. 4.3 also behave similarly. Using more than four Λ-iterations per time-step does not
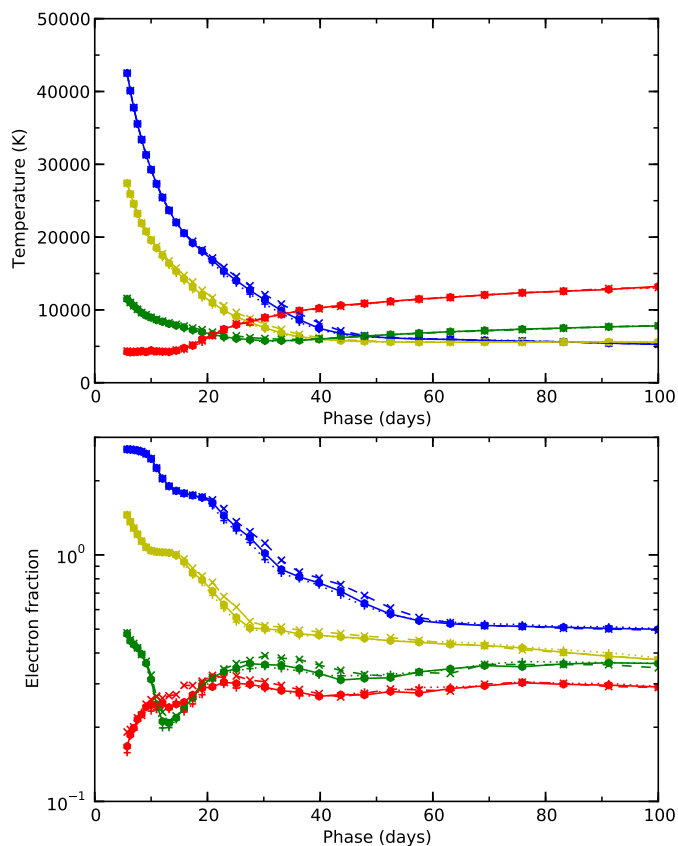
Fig. 18. Temperature (upper panel) and electron fraction (lower panel) for model 12C as calculated with JEKYLL using two (dashed lines and crosses), four (solid lines and circles) and eight (dotted lines and pluses) Λ-iterations per time-step. Otherwise as described in Fig. 13.

make a significant difference, and using four instead of two Λ-iterations only marginally changes the spectra. Furthermore, in this case we know that the Λ-iteration is converging to a solution that is (most likely) close to the correct one. Together with the nice convergence properties for model 12C, this is assuring, although further comparisons to CMFGEN would be interesting.

5.5. Testing extensions of the Lucy method

5.5.1. The use of a diffusion solver

The use of the diffusion solver in the inner region speeds up calculations in the early phase, and it is used in all simulations except the comparison with SUMO. It is therefore of interest to investigate how the diffusion solver, and the depth at which it is coupled to the MC radiative transfer solver, influence the solution. To achieve this, we have run model 12C with the diffusion solver coupled at a (Rosseland mean) continuum optical depth of 50, 100 and 200, using the same simplified set-up as described in Sect. 5.4. At these coupling depths, the diffusion solver is only used until 21, 15 and 12 days, respectively, so in the last case most of the diffusion peak is actually calculated using the MC radiative transfer solver alone.

The results show no significant differences in the spectra and only small differences in the matter quantities, which justifies the use of the diffusion solver, at least at an optical depth of 50 or more. It is also interesting to return to the comparison with CMFGEN in Fig. 10, where we have marked the border between
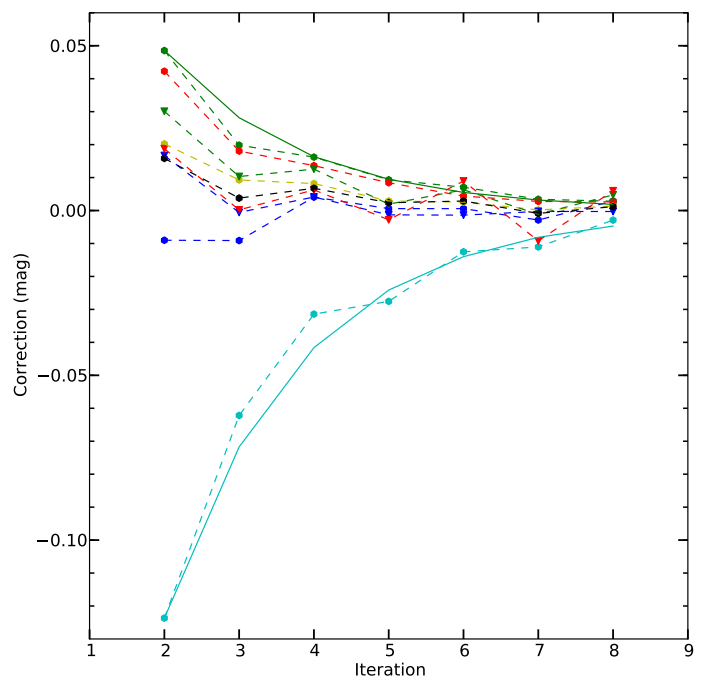

Fig. 19. Correction per Λ-iteration averaged between 15 and 40 days for the broad-band and bolometric lightcurves of model 12C. The colours used for the bands are the same as in Fig. 12, and the optical and NIR bands have been marked with circles and triangles, respectively. For the *U*- and *V*-bands we also show the curves of geometric convergence for a ratio of successive corrections of 0.58.
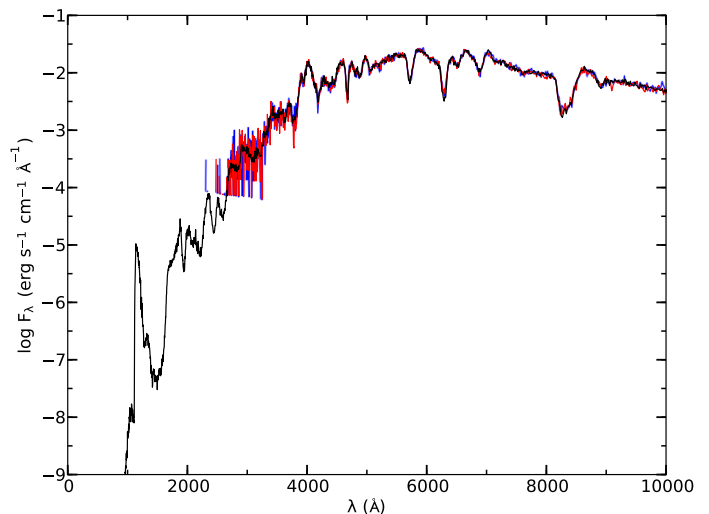
Fig. 20. Spectra for model 12C at 34.7 days with (black) and without (red) packet sampling control activated, as well as with (red) and without (blue) the Markov-Chain solution activated.

the diffusion solver and the MC radiative transfer solver. As we can here compare to as solution that is (most likely) close to the correct one, it is evident that the use of a coupled diffusion and MC radiative transfer solver works well.

5.5.2. The packet sampling control

The method for packet sampling control was introduced in Sect. 3.3.5, and is used to decrease the noise in the radiation field estimators. As a test, we have re-run the MC radiative transfer

for model 12C with and without packet sampling control activated. In doing this we have loaded the matter state from the original model 12C run and kept it fixed. Figure 20 shows the spectrum for model 12C at 34.7 days with and without packet sampling control activated. The packet sampling control was configured to maintain a SNR of 3 percent between the Lyman break and 25000 Å with a minimum boost factor of 1 and and a maximum boost factor of 10^{12}. As seen in Fig. 20, the two spectra agree well as long as the SNR is good in both, which shows that the method reproduces the correct radiation field. In addition, in the region bluewards ∼3000 Å, a SNR of ∼3 percent is maintained all the way to the Lyman break in the model with packet sampling control, but in the model without there are no MC packets at all below ∼2500 Å.

In the model with packet sampling control, the average boost of the number of packets is ∼4, whereas the boost factors in the blue region approach ∼10^9 near the Lyman break. Without packet sampling control the same SNR in the blue region could only have been achieved by increasing the total number of packets by this huge factor. Finally, we repeated the test, but decreased the number of packets injected into the calculation with a factor of ten. The resulting spectrum is indistinguishable from that from the original test, the only difference being the ten times higher boost factors needed to achieve the required SNR.

5.5.3. The Markov-Chain solution

The Markov-Chain solution to the MC state-machine was introduced in Sect. 3.3.4 as a way to sample the emission frequency more efficiently. As a test, we have re-run the MC radiative transfer for model 12C with and without the Markov-Chain solution activated. As in Sect. 5.5.2, we have loaded the matter state from the original model 12C run and kept it fixed. As for most simulations in this paper, we have used the split state-machine approach (see Sect. 3.3.4 and Appendix B) when calculating the Markov-chain solution, as this speeds up the calculation by a large factor.

Figure 20 shows the spectrum for model 12C at 34.7 days with and without the Markov-Chain solution activated. The spectra agree well, which shows that the Markov-Chain solution produces the same radiation field as the original state-machine. The speed-up of the MC radiative transfer when using the Markov-Chain solution is a factor of ∼50 at 1 day, increases to ∼150 at 4 days and then declines and levels out at ∼10. The overhead from calculating the Markov-chain solution (using the split state-machine approach) is small at early times, but increase with time, and at 40 days it decrease the effective speed-up with ∼30 percent. The number of packets as well as the depth at which the diffusion solver is coupled affects the measured speed-up, but the case investigated here should be fairly representative. Clearly the method is essential at early times, and otherwise most helpful in reducing the computational effort.

6. Conclusions

We present and describe JEKYLL, a new code for modelling of SN spectra and lightcurves. The code assumes homologous expansion, spherical symmetry and steady state for the matter, but is otherwise capable of solving for the time-evolution of the matter and the radiation field in full NLTE. The method used is an extension of the MC based Lucy method, here tested for

the first time in its time-dependent NLTE version. In particular, it includes a detailed treatment of non-thermal excitation and ionization, where the rates are calculated as described in KF92. Another important feature is a method to account for the macroscopic mixing that occurs in the explosion, which was previously introduced in J11. We also describe how to speed up the calculation by using a diffusion solver in the inner region, and by using Markov-chains to sample the packet frequency more efficiently. In addition, we introduce a novel method to control the sampling of the radiation field, which is used to reduce the noise in the radiation field estimators.

We also present comparisons with the ARTIS, SUMO and CMFGEN codes. The ARTIS and SUMO codes are similar to JEKYLL in some, but not all, aspects, and the comparisons provide tests of the time-dependent MC radiative transfer and the steady-state NLTE capabilities, respectively. The CMFGEN code is similar in terms of physics, but uses a different method, where the coupled system of differential equations for the matter and the radiation field is solved by a linearization technique. This comparison, which is done with a somewhat simplified ejecta model, provides a test of the time-dependent NLTE capabilities of JEKYLL. All comparisons show a good agreement in the observed quantities, as well as the state variables, and together they provide a thorough test of the JEKYLL code. In particular, the comparison with CMFGEN shows that the MC based Lucy method, where the de-coupled radiation-matter problem is solved through Λ-iteration, does indeed converge in the time-dependent NLTE case. This has previously only been shown for the steady-state NLTE case in a hydrogen SN atmosphere (L03).

Finally, we present an example of the time-dependent NLTE capabilities of JEKYLL using a realistic ejecta model for a Type IIb SN. This model belongs to the set of Type IIb models presented by J15, which are explored in more detail and compared to observations in Paper 2. Based on the model we investigate the effect of NLTE, and find strong effects even on the bolometric lightcurve. This casts some doubts on LTE-based modelling of the bolometric lightcurve commonly used in the literature. For example non-thermal ionization turns out to have a strong effect on the ionization level in the helium envelope, which introduces a coupling between the mixing of the radioactive ^{56}Ni and the diffusion time not accounted for in LTE-based models.

We also use the model to test our most important (technical) extensions of the Lucy method. Among others, these tests show that the Markov-chain solution to the MC state-machine may speed up the calculations in the early phase by large factors, and that the method for packet sampling control may improve the SNR in noisy wavelength regions by huge factors at a modest computational cost.

7. Acknowledgements

Appendix A: Configuration and atomic data

Appendix A.1: Comparison to ARTIS

To synchronize JEKYLL with ARTIS, both codes were configured to use a LTE solution for the matter based on T_J, the temperature associated with the pure blackbody radiation field model (see Sect. 3.1.1 and K09), and the MC radiative transfer solver

used by JEKYLL was configured to include radiative and collisional bound-bound and bound-free processes as well as free-free processes. In addition, ARTIS was configured to use its grey approximation (see K09) before ~6 days and below an optical depth of 100, and JEKYLL was configured to use the diffusion solver below an optical depth of 50.

The atomic data used by ARTIS is described in K09, but was restricted to the first four ionization stages, using the fourth as closure. As all of the atomic data is stored in data-files in a well-defined format, it was quite straight-forward to automatically convert it to the JEKYLL atomic data format, and it should be fully synchronized. ARTIS and JEKYLL were both configured to use a logarithmic time-step of 1 percent and single Λ-iteration per time-step, which is the standard procedure in ARTIS. As discussed in Sect 5.4, due to the short time-step, these runs are still well converged.

Appendix A.2: Comparison to SUMO

As much as possible, we have synchronized the configuration and the atomic data used by JEKYLL with that used for the modelling in J15. To achieve this, JEKYLL was configured to run in steady-state mode, and to use a full NLTE solution including the following; radiative bound-bound, bound-free and free-free processes, collisional bound-bound processes, non-thermal excitation, ionization and heating, as well as charge-transfer and two-photon processes. JEKYLL was also configured to use a recombination correction in a manner similar to SUMO (see J11), in which case detailed balance was not enforced.

The atomic data used for the modelling in J15 is described in J11 and Jerkstrand et al. (2012). In the case it was stored in data-files in a well-defined format, as for e.g. energy levels and spontaneous emission rates, it was automatically converted to the JEKYLL atomic data format, and otherwise it was added manually to the JEKYLL atomic data files based on the descriptions in J11 and Jerkstrand et al. (2012). Although not complete, the synchronization of the atomic data and the methods should be good enough for a meaningful comparison.

Appendix A.3: Comparison to CMFGEN

To synchronize JEKYLL with CMFGEN, JEKYLL was configured run in time-dependent (radiative transfer) mode, and to use a full NLTE solution including the following; radiative bound-bound, bound-free and free-free processes, as well as collisional bound-bound and bound-free processes. JEKYLL was also configured to use the time-dependent NLTE rate equations and to use the diffusion solver below an optical depth of 100. In addition, to assure good sampling of the radiation field, packet control (see Sect. 3.3.5) was turned on and the generalized on-the-spot approximation (see Sect. 3.1.1) was used bluewards the Lyman break.

The atomic data for the simplified composition of hydrogen, helium, oxygen and calcium were automatically converted from the well-defined format of CMFGEN to that of JEKYLL, and should therefore be fully synchronized. CMFGEN and JEKYLL were both configured to use a logarithmic time-step of 2.5 percent, and JEKYLL was configured to use 4 Λ-iterations per time-step. As discussed in Sect. 5.4, this gives a well converged solution.

Appendix A.4: Application to Type IIb SNe

JEKYLL was configured to run in time-dependent (radiative transfer) mode, and to use a full NLTE solution including the following; radiative bound-bound, bound-free and free-free processes, collisional bound-bound and bound-free processes, non-thermal excitation, ionization and heating, as well as two-photon processes. JEKYLL was also configured to use the diffusion solver below an optical depth of 50, and to use a recombination correction while still enforcing detailed balance. In addition, to assure good sampling of the radiation field, packet control (see Sect. 3.3.5) was turned on and the generalized on-the-spot approximation (see Sect. 3.1.1) was used bluewards the Lyman break. The logarithmic time-step was set to 5 percent and the number of Λ-iterations per time-step was set to 4. As discussed in Sect. 5.4, this gives a well converged solution.

The atomic data used is the same as for the comparison with SUMO (Sect. A.2), but with the following modifications. The highest ionization stage was increased to VI for all species, and the stage III ions were updated to include at least 50 levels for elements lighter than Sc, and at least 200 levels for heavier elements, using online data provided by NIST[7] and R. Kurucz[8]. Total recombination rates for the stage III ions were adopted from the online table provided by S. Nahar[9] whenever available, and otherwise from Shull & van Steenberg (1982). For ionization stages IV to VI we only included the ground-state multiplets, adopted the photo-ionization cross-section by Verner & Yakovlev (1995) and Verner et al. (1996) and assumed the populations to be in LTE with respect to stage IV.

Appendix B: Splitting the MC state-machine

To split the MC state-machine and the corresponding Markov-chain model into its base- and sub-machine parts (Sect. 3.3.3), we proceed as follows. First, consider the base-machine, where the states correspond to the macro-atoms and the thermal pool. The probabilities for internal transitions and de-activation from the thermal pool are given by the cooling rates. More specifically, the probability for an internal transition from the thermal pool (labelled T) to macro-atom state I is given by

$$P_{T \to I}^{I} = \sum_i P_{I,i}^{A,C} \tag{B.1}$$

where $P_{I,i}^{A,C}$ is the probability for collisional activation of macro-atom sub-machine I in state i, in turn given by the probability for collisional cooling through an upward transition to state i of macro-atom sub-machine I. The probability for de-activation from the thermal pool, P_T^D, is just the probability for radiative cooling. In the case that the base-machine is activated in its thermal pool state (which is the case we are interested in), the probabilities for internal transitions and de-activations from the macro-atom states are given by the cooling rates and the **SR** matrices (Sect. 3.3.4) for the macro-atom sub-machines. More specifically, the probability for an internal transition from macro-atom state I to the thermal pool is

$$P_{I \to T}^{I} = \sum_{i,j} P_{I,i}^{A,C} (\mathbf{SR})_{I,i,j}^{C} \tag{B.2}$$

where $\mathbf{SR}_I^C$ is the **SR** matrix for macro-atom sub-machine I with respect to *collisional* de-activation. Similarly, the probability for

---

[7] www.nist.gov
[8] http://www.cfa.harvard.edu/amp/ampdata/kurucz23/sekur.html
[9] http://www.astronomy.ohio-state.edu/~nahar/_naharradiativeatomicdata/

de-activation from macro-atom state I is given by

$$P_I^{\mathrm{D}} = \sum_{i,j} P_{I,i}^{\mathrm{A,C}} (\mathbf{SR})_{I,i,j}^{\mathrm{R}} \tag{B.3}$$

where $\mathbf{SR}_I^{\mathrm{R}}$ is the $\mathbf{SR}$ matrix for macro-atom sub-machine I with respect to *radiative* de-activation. From the probabilities for internal transitions and de-activations we can calculate the base-machine $\mathbf{SR}$ matrix as described in Sect. 3.3.4. Note, that this is done for the case when the base-machine is activated in its thermal pool state, which is all we need.

The procedure is now as follows. If the base-machine is activated in a macro-atom state, a macro-atom sub-machine is activated by a radiative transition. The de-activating transition is then drawn based on the $\mathbf{SR}$ matrix and the de-activation probabilities for this macro-atom sub-machine, as described in Sect. 3.3.4. Radiative de-activation corresponds to de-activation of the base-machine, whereas collisional de-activation corresponds to an internal base-machine transition to the thermal pool. In the latter case, or if the base-machine was activated in its thermal pool state, the state from which the base-machine de-activates is drawn based on its $\mathbf{SR}$ matrix, as described in Sect. 3.3.4. If it de-activates from the thermal pool, an emission process is drawn in proportion to the radiative cooling rates. If it de-activates from a macro-atom state, the de-activating transition is drawn based on the $\mathbf{SR}^{\mathrm{R}}$ matrix and the *radiative* de-activation probabilities for the macro-atom sub-machine, as described in Sect. 3.3.4. The split state-machine approach is more involved and comes at the expense of (two) more draws, but at a greatly reduced cost to calculate and store the $\mathbf{SR}$ matrices, and has been used for most of the simulations in this paper.

Appendix C: Adjusting the MC emission rates

As explained in Sect. 3.3.5, control of the number of packets is achieved by adjusting their size (i.e. the amount of energy they carry, see Sect. 3.3.5) as a function of frequency[10]. This may be expressed through a boost factor $B(\nu)$, in terms of which the packet size is given by $E_{\mathrm{P}}(\nu) = E_{\mathrm{P},0}/B(\nu)$, where $E_{\mathrm{P},0}$ is some reference size. If $B(\nu) > 1$, this corresponds to a boost of the number of packets, and if $B(\nu) < 1$ it corresponds to a reduction. As discussed in Sect. 3.3.5, the values for the boost factors[11] used in JEKYLL are determined by an adaptive algorithm, in which they are adjusted once per Λ-iteration to achieve a pre-configured target SNR.

As explained in Sect. 3.3.5, to conserve energy, the emissivity in the destination region (in terms of packets) has to be adjusted with F, the ratio of the packet sizes in the source and destination regions. In terms of the boost factors in these regions, F may also be written as a boost ratio, i.e. $F(\nu_A, \nu_E) = B(\nu_E)/B(\nu_A)$, where ν_A and ν_E are the absorption and emission frequencies of the packet. As is also explained in Sect. 3.3.5, each emission event in the destination region is triggered by a fictitious absorption event in the source region. For a given interaction process, these fictitious absorption events correspond to an fictitious opacity κ_{F}, which gives the interaction rate required to obtain the adjusted emissivities.

Beginning with absorption, a packet takes one of the possible paths through the MC state-machine, and is eventually emit-

---
[10] The variation of the packet size in space and time is ignored in this section.
[11] The discussion in Sect. 3.3.5 is in terms of packet size, but that makes no difference.

ted. To adjust the rates of packets flowing through the MC state-machine, we need to proceed in the reverse order. First, based on the boost factor (B), we calculate boost factors for the emissivities (B^{E}) and adjust these accordingly. Second, based on the boost factors for the emissivities, we calculate boost factors for the MC state-machine probabilities (B^{D}, B^{A} and B^{C}) and adjust these accordingly. Finally, based on the boost factors for the MC state-machine probabilities, we calculate boost factors for the opacities (B^{O}). For a specific packet with frequency ν, this gives the corresponding boost ratios $F^{\mathrm{O}}(\nu) = B^{\mathrm{O}}/B(\nu)$, which in turn give the fictitious opacities $\kappa_{\mathrm{F}}(\nu) = F^{\mathrm{O}}(\nu)\,\kappa(\nu)$. However, as absorption has to proceed at the original rate, the fictitious opacity to use is not κ_{F} but $max(\kappa_{\mathrm{F}}, \kappa) = max(F^{\mathrm{O}}, 1)\,\kappa$. Using this fictitious opacity, packets are then selected for emission, absorption or both according to the rules described in Sect. 3.3.5.

Note, that in this appendix we follow the convention to number energy levels with respect to the macro-atoms instead of the ions. Note also, that the boost factors and the adjustment of the macro-atom probabilities in Appendix C.2 apply to the complete MC state-machine, and have to be adjusted in case the split MC state-machine approach (see Appendix. B) is used.

Appendix C.1: Emissivities

The boost factor for emission as a function of frequency is just $B(\nu)$, and from this the boost factors for bound-bound, bound-free and free-free emission are calculated as

$$B_{I,i\to j}^{\mathrm{E,BB}} = B(\nu_{I,i\to j})$$
$$B_{I,i\to j}^{\mathrm{E,BF}} = \int B(\nu)\, f_{I,i\to j}^{\mathrm{BF}}(\nu)\, d\nu$$
$$B^{\mathrm{E,FF}} = \int B(\nu)\, f^{\mathrm{FF}}(\nu)\, d\nu \tag{C.1}$$

where $f_{I,i\to j}^{\mathrm{BF}}(\nu)$ and $f^{\mathrm{FF}}(\nu)$ are the distribution functions for bound-free and free-free emission, respectively. The distribution functions for bound-free and free-free emission are then adjusted as

$$f_{I,i\to j}'^{\mathrm{BF}}(\nu) = B(\nu)\, f_{I,i\to j}^{\mathrm{BF}}(\nu)\, / B_{I,i\to j}^{\mathrm{E,BF}}$$
$$f'^{\mathrm{FF}}(\nu) = B(\nu)\, f^{\mathrm{FF}}(\nu)\, / B^{\mathrm{E,FF}} \tag{C.2}$$

Appendix C.2: Macro-atom probabilities

The boost factor for de-activation from state i of macro-atom I is calculated as

$$B_{I,i}^{\mathrm{D}} = \sum_j B_{I,i\to j}^{\mathrm{E}} P_{I,i\to j}^{\mathrm{D}} \tag{C.3}$$

where $P_{I,i\to j}^{\mathrm{D}}$ is the probability for a de-activating (bound-bound or bound-free) transitions from state i to state j of macro-atom I. The probability for de-activating transitions is then adjusted as

$$P_{I,i\to j}'^{\mathrm{D}} = B_{I,i\to j}^{\mathrm{E}} P_{I,i\to j}^{\mathrm{D}} / B_{I,i}^{\mathrm{D}} \tag{C.4}$$

The boost factor for activation of macro-atom I in state j is calculated as

$$B_{I,i}^{\mathrm{A}} = \sum_{J,j} B_{J,j}^{\mathrm{D}}\,(\mathbf{SR})_{(I,i),(J,j)} + B^{\mathrm{C,R}}\,(\mathbf{SR})_{(I,i),T} \tag{C.5}$$

where the $\mathbf{SR}$ matrix is obtained from a Markov-chain solution to the MC state-machine (see Sect.3.3.4), and $B^{\mathrm{C,R}}$ is the boost

factor for radiative cooling (see Appendix C.3). The indices of **SR** are labelled (N, n) if they correspond to state n of macro-atom N or T if they correspond to the thermal pool. The **SR** matrix is then adjusted as

$$(\mathbf{SR})'_{(I,i),(J,j)} = B^{\mathrm{D}}_{J,j}\,(\mathbf{SR})_{(I,i),(J,j)}\,/B^{\mathrm{A}}_{I,i}$$
$$(\mathbf{SR})'_{(I,i),T} = B^{\mathrm{C,R}}\,(\mathbf{SR})_{(I,i),T}\,/B^{\mathrm{A}}_{I,i} \tag{C.6}$$

Appendix C.3: Thermal pool probabilities

The boost factor for collisional cooling is calculated as

$$B^{\mathrm{C,C}} = \sum_{I,i,j} B^{\mathrm{A}}_{I,j} P^{\mathrm{C,C}}_{I,i\to j} \tag{C.7}$$

where $P^{\mathrm{C,C}}_{I,i\to j}$ is the probability for collisional cooling through a transition from state i to state j of macro-atom I. This probability is then adjusted as

$$P'^{\mathrm{C,C}}_{I,i\to j} = B^{\mathrm{A}}_{I,j} P^{\mathrm{C,C}}_{I,i\to j}\,/B^{\mathrm{C,C}} \tag{C.8}$$

The boost factor for radiative cooling is calculated as

$$B^{\mathrm{C,R}} = \sum_{I,i,j} B^{\mathrm{E,BF}}_{I,i\to j} P^{\mathrm{C,BF}}_{I,i\to j} + B^{\mathrm{E,FF}} P^{\mathrm{C,FF}} \tag{C.9}$$

where $P^{\mathrm{C,BF}}_{I,i\to j}$ is the probability for radiative cooling through a bound-free transition from state i to state j of macro-atom I and $P^{\mathrm{C,FF}}$ is the probability for radiative cooling through free-free emission. These probabilities are then adjusted as

$$P'^{\mathrm{C,BF}}_{I,i\to j} = B^{\mathrm{E,BF}}_{I,i\to j} P^{\mathrm{C,BF}}_{I,i\to j}\,/B^{\mathrm{C,R}}$$
$$P'^{\mathrm{C,FF}} = B^{\mathrm{E,FF}} P^{\mathrm{C,FF}}\,/B^{\mathrm{C,R}} \tag{C.10}$$

The boost factor for (the total) cooling is calculated as

$$B^{\mathrm{C}} = B^{\mathrm{C,C}} P^{\mathrm{C,C}} + B^{\mathrm{C,R}} P^{\mathrm{C,R}} \tag{C.11}$$

where $P^{\mathrm{C,C}}$ and $P^{\mathrm{C,R}}$ are the probabilities for collisional and radiative cooling, respectively. The probabilities for collisional and radiative cooling are then adjusted as

$$P'^{\mathrm{C,C}} = B^{\mathrm{C,C}} P^{\mathrm{C,C}}\,/B^{\mathrm{C}}$$
$$P'^{\mathrm{C,R}} = B^{\mathrm{C,R}} P^{\mathrm{C,R}}\,/B^{\mathrm{C}} \tag{C.12}$$

Appendix C.4: R-packet opacities

The boost factor for absorption through a bound-bound transition from state i to state j of macro-atom I is calculated as

$$B^{\mathrm{O,BB}}_{I,i\to j} = B^{\mathrm{A}}_{I,j} \tag{C.13}$$

and the fictitious Sobolev optical depth for bound-bound transitions is then calculated as

$$\tau^{\mathrm{F}}_{I,i\to j} = max(F^{\mathrm{O,BB}}_{I,i\to j}, 1)\,\tau_{I,i\to j} \tag{C.14}$$

The boost factor for absorption through a bound-free transition from state i to state j of macro-atom I is calculated as

$$B^{\mathrm{O,BF}}_{I,i\to j} = B^{\mathrm{A}}_{I,j} P^{\mathrm{I}}_{I,i\to j} + B^{\mathrm{C}} P^{\mathrm{H}}_{I,i\to j} \tag{C.15}$$

where $P^{\mathrm{I}}_{I,i\to j}$ and $P^{\mathrm{H}}_{I,i\to j}$ are the probabilities that the bound-free transition results in ionization and heating, respectively. The probabilities for ionization and heating is then adjusted as

$$P'^{\mathrm{I}}_{I,i\to j} = B^{\mathrm{A}}_{I,j} P^{\mathrm{I}}_{I,i\to j}\,/B^{\mathrm{O,BF}}_{I,i\to j}$$
$$P'^{\mathrm{H}}_{I,i\to j} = B^{\mathrm{C}} P^{\mathrm{H}}_{I,i\to j}\,/B^{\mathrm{O,BF}}_{I,i\to j} \tag{C.16}$$

The boost factor for bound-free absorption is calculated as

$$B^{\mathrm{O,BF}}(\nu) = \sum_{I,i,j} B^{\mathrm{O,BF}}_{I,i\to j}\kappa_{I,i\to j}(\nu)/\sum_{I,i,j}\kappa_{I,i\to j}(\nu) \tag{C.17}$$

and the fictitious opacity for bound-free absorption is then calculated as

$$\kappa^{\mathrm{BF}}_{\mathrm{F}}(\nu) = max(F^{\mathrm{O,BF}}(\nu), 1)\,\kappa^{\mathrm{BF}}(\nu) \tag{C.18}$$

The boost factor for free-free absorption is calculated as

$$B^{\mathrm{O,FF}} = B^{\mathrm{C}} \tag{C.19}$$

and the fictitious opacity for free-free absorption is then calculated as

$$\kappa^{\mathrm{FF}}_{\mathrm{F}}(\nu) = max(F^{\mathrm{O,FF}}, 1)\,\kappa^{\mathrm{FF}}(\nu) \tag{C.20}$$

Appendix C.5: γ-packet opacities

The boost factor for absorption of a γ-packet is calculated as

$$B^{\mathrm{O,G}} = B^{\mathrm{C}} P^{\mathrm{NT,H}} + \sum_{I,i,j} B^{\mathrm{A,BB}}_{I,i\to j} P^{\mathrm{NT,E}}_{I,i\to j} + \sum_{I,i,j} B^{\mathrm{A,BF}}_{I,i\to j} P^{\mathrm{NT,I}}_{I,i\to j} \tag{C.21}$$

where $P^{\mathrm{NT,H}}$ and $P^{\mathrm{NT,E/I}}_{I,i\to j}$ are the probabilities for non-thermal heating and non-thermal excitation/ionization through a (bound-bound/bound-free) transition from state i to state j of macro-atom I. The probabilities for non-thermal heating, excitation and ionization are then adjusted as

$$P'^{\mathrm{NT,H}} = B^{\mathrm{C}} P^{\mathrm{NT,H}}\,/B^{\mathrm{O,G}}$$
$$P'^{\mathrm{NT,E}}_{I,i\to j} = B^{\mathrm{A,BB}}_{I,j} P^{\mathrm{NT,E}}_{I,i\to j}\,/B^{\mathrm{O,G}}$$
$$P'^{\mathrm{NT,I}}_{I,i\to j} = B^{\mathrm{A,BF}}_{I,j} P^{\mathrm{NT,I}}_{I,i\to j}\,/B^{\mathrm{O,G}} \tag{C.22}$$

and from the effective opacity used for the g-packet (which is different for different decays and decay products), the corresponding fictitious opacity is calculated as

$$\kappa^{\mathrm{G}}_{\mathrm{F}}(\nu) = max(F^{\mathrm{O,G}}, 1)\,\kappa^{\mathrm{G}}(\nu) \tag{C.23}$$

References

Abbott, D. C. & Lucy, L. B. 1985, ApJ, 288, 679
Avrett, E. H. & Loeser, R. 1988, ApJ, 331, 211
Bersten, M. C., Benvenuto, O. G., Nomoto, K., et al. 2012, ApJ, 757, 31
Cannon, C. J. 1973a, J. Quant. Spec. Radiat. Transf., 13, 627
Cannon, C. J. 1973b, ApJ, 185, 621
Carciofi, A. C. & Bjorkman, J. E. 2006, ApJ, 639, 1081
Carciofi, A. C. & Bjorkman, J. E. 2008, ApJ, 684, 1374
Dessart, L. & Hillier, D. J. 2008, MNRAS, 383, 57
Dessart, L. & Hillier, D. J. 2010, MNRAS, 405, 2141
Dessart, L., Hillier, D. J., Li, C., & Woosley, S. 2012, MNRAS, 424, 2139
Dessart, L., Hillier, D. J., Woosley, S., et al. 2015, MNRAS, 453, 2189
Dessart, L., Hillier, D. J., Woosley, S., et al. 2016, MNRAS, 458, 1618
Ergon, M., Jerkstrand, A., Sollerman, J., et al. 2015, A&A, 580, A142
Ergon, M., Sollerman, J., Fraser, M., et al. 2014, A&A, 562, A17
Esposito, L. W. & House, L. L. 1978, ApJ, 219, 1058
Falk, S. W. & Arnett, W. D. 1977, ApJS, 33, 515

Hauschildt, P. H. & Baron, E. 1999, Journal of Computational and Applied Mathematics, 109, 41

Hillier, D. J. & Dessart, L. 2012, MNRAS, 424, 252

Hillier, D. J. & Miller, D. L. 1998, ApJ, 496, 407

Hubeny, I. & Mihalas, D. 2014, Theory of Stellar Atmospheres

Jerkstrand, A., Ergon, M., Smartt, S. J., et al. 2015, A&A, 573, A12

Jerkstrand, A., Fransson, C., & Kozma, C. 2011, A&A, 530, A45

Jerkstrand, A., Fransson, C., Maguire, K., et al. 2012, A&A, 546, A28

Kasen, D., Thomas, R. C., & Nugent, P. 2006, ApJ, 651, 366

Kerzendorf, W. E. & Sim, S. A. 2014, MNRAS, 440, 387

Kozma, C. & Fransson, C. 1992, ApJ, 390, 602

Kozma, C. & Fransson, C. 1998, ApJ, 496, 946

Kromer, M. & Sim, S. A. 2009, MNRAS, 398, 1809

Long, K. S. & Knigge, C. 2002, ApJ, 579, 725

Lucy, L. B. 1991, ApJ, 383, 308

Lucy, L. B. 1999, A&A, 345, 211

Lucy, L. B. 2002, A&A, 384, 725

Lucy, L. B. 2003, A&A, 403, 261

Lucy, L. B. 2005, A&A, 429, 19

Mazzali, P. A. 2000, A&A, 363, 705

Mazzali, P. A. & Lucy, L. B. 1993, A&A, 279, 447

Olson, G. L., Auer, L. H., & Buchler, J. R. 1986, J. Quant. Spec. Radiat. Transf., 35, 431

Paxton, B., Bildsten, L., Dotter, A., et al. 2011, ApJS, 192, 3

Paxton, B., Cantiello, M., Arras, P., et al. 2013, ApJS, 208, 4

Ross, S. 2007, Introduction to Probability Models

Scharmer, G. B. 1984, Accurate solutions to non-LTE problems using approximate lambda operators, ed. W. Kalkofen, 173–210

Shull, J. M. & van Steenberg, M. 1982, ApJS, 48, 95

Sim, S. A. 2007, MNRAS, 375, 154

Sim, S. A., Miller, L., Long, K. S., Turner, T. J., & Reeves, J. N. 2010, MNRAS, 404, 1369

Sobolev, V. V. 1957, Soviet Ast., 1, 678

Tanaka, M., Maeda, K., Mazzali, P. A., & Nomoto, K. 2007, ApJ, 668, L19

Verner, D. A., Ferland, G. J., Korista, K. T., & Yakovlev, D. G. 1996, ApJ, 465, 487

Verner, D. A. & Yakovlev, D. G. 1995, A&AS, 109, 125

Werner, K. & Husfeld, D. 1985, A&A, 148, 417

Woosley, S. E., Eastman, R. G., Weaver, T. A., & Pinto, P. A. 1994, ApJ, 429, 300

Woosley, S. E. & Heger, A. 2007, Phys. Rep., 442, 269