Modelling supernova light curves and spectra with Monte Carlo methods

May 14-16 block in "Radiative Transfer Across Disciplines" course, 2024.

A. Jerkstrand

May 16, 2024

Contents

1 Radiative transfer in a moving medium

When the medium is moving, Doppler effects impact the radiative transfer. Particularly important are Doppler-induced frequency changes in line transfer problems, because lines have such narrow absorption and emission profiles, set by the thermal velocity scale which is of order a few km/s, much smaller than Doppler velocities in stellar winds and explosions, where kinematic velocities vary on order 100 or 1000 km/s.

Moving media may be divided into **steady** and **non-steady** flows. A steady flow is for example a stellar wind - there is an important, and often differential, velocity field - but the domain looks the same at all times. A non-steady flow is for example a supernova explosion - a nebula is born that expands and the density field is different at different times.

1.1 Reference Frames

The laws of physics are the same in all inertial frame, and by the use of proper fictitious forces the laws can also be formulated in non-inertial frames. We can use this freedom to work in the frame which gives the most tractable formulation for the problem at hand. We will later see that, in fact, for moving media we often use different frames for different spatial points. In Monte Carlo simulations we follow photon packets through the medium - jumping between different frames that are "best for the moment".

There are three main reference frames which are of fundamental importance in radiative transfer:

1. **Atom frame ("AF")**. The frame in which the individual atom (or other particle, e.g. an electron) undergoing an interaction (emission, absorption, scattering) has zero velocity.
2. **Lab frame ("LF", or "observer-frame (OF)")**. The frame in which the observer, or sometimes the star (or corresponding system) as a whole, is at rest (meaning the momentum vector of the system as a whole is zero).
3. **Comoving frame ("CMF", or "Lagrangian frame (LF)")**. Frame at ($\mathbf{x}$,t) in which the integral of velocity vectors in a small volume around $\mathbf{x}$ is zero. Since the matter seen as "bulk" is at rest in this frame, the opacity, and in many cases also the emissivity, are isotropic. This is the key property that makes the CMF attractive to work in.

Note that the CMF in a strict sense is not an inertial frame, as the matter bundle, and therefore the frame itself, may accelerate. This acceleration effect can for almost all applications be ignored, however. The CMF of a given bundle can therefore be considered as a time-sequence of inertial frames, each moving with the instantaneous velocity of the fluid element under consideration.

The particularities of the CMF instead arises from another direction. If we write down any non-local physical law, such as the RT equation, in "the CMF", we are in fact writing a relation that describes something happening as we change frames from one CMF (e.g. starting point of a ds path segment) to another (ending point of the ds path segment). Terms will emerge due to this connection between two (quasi)-inertial frames.

Lorentz transformations give two important effects for RT, namely the **Doppler shift**:

$$\nu' = \gamma\nu\left(1 - \mathbf{n}\cdot\mathbf{v}/c\right), \tag{1}$$

$$\nu = \gamma\nu'\left(1 + \mathbf{n}'\cdot\mathbf{v}/c\right), \tag{2}$$

and the **aberration**

$$\mathbf{n}' = (\nu/\nu')\left(\mathbf{n} - \gamma\mathbf{v}/c\right)\left[1 - \frac{\gamma\mathbf{n}\cdot\mathbf{v}/c}{\gamma+1}\right], \tag{3}$$

$$\mathbf{n} = (\nu'/\nu)\left(\mathbf{n}' + \gamma\mathbf{v}/c\right)\left[1 + \frac{\gamma\mathbf{n}'\cdot\mathbf{v}/c}{\gamma+1}\right]. \tag{4}$$

Here, we adopt the notation that CMF frame quantities are primed. In 1D, the aberration formulae

simplify to ($\mu = \cos\theta$, where θ is the angle to the radial direction):

$$\mu' = \frac{\mu - \frac{v}{c}}{1 - \frac{v}{c}\mu}, \tag{5}$$

$$\mu = \frac{\mu' + \frac{v}{c}}{1 + \frac{v}{c}\mu'}. \tag{6}$$

Thomas [1930] derived the transformation laws for the specific intensity

$$I(\nu, \mu) = \left(\frac{\nu}{\nu'}\right)^3 I'(\nu', \mu'), \tag{7}$$

the emissivity

$$\eta(\nu, \mu) = \left(\frac{\nu}{\nu'}\right)^2 \eta'(\nu), \tag{8}$$

and the absorption coefficient

$$\chi(\nu, \mu) = \left(\frac{\nu}{\nu'}\right)^{-1} \chi'(\nu'). \tag{9}$$

Due to the direction-dependence of ν/ν' (Eq. 2), Eqs. 8 and 9 tell us that in the observer frame, emission and absorption coefficients are always angle-dependent.

1.2 The RT equation in the comoving frame

When the fluid is accelerating (as e.g. during a SN explosion) or there is velocity gradient (as in the coasting phase of a supernova), it is difficult to solve the RT equation in the observer frame, because the motions cause η_ν and χ_ν to become anisotropic as shown above. One approach is to simplify the RT equation with first-order expansions for η_ν and χ_ν. This method is described in section 93 in Mihalas and Mihalas [1984]. The limitation of this method is when lines are important; the first-order expansions are then insufficient. One then normally works in the CMF, as described in section 95 in Mihalas and Mihalas [1984]. Here χ_ν, and for many applications also η_ν, is isotropic and the matter-radiation interaction calculations become easier. The drawback is a more complex RT equation, induced by the change of frames.

The fully relativistic CMF RT equation is, even in spherical symmetry, a very long expression [Eq. 95.9 in Mihalas and Mihalas, 1984]. The moment equations are also lengthy (Eqs. 95.11 and 95.12). If we retain only terms to order v/c, and also ignore the fluid acceleration terms (involving dv/dt), we get in spherical symmetry (Eq. 95.17 MM84):

$$\begin{aligned} \frac{1}{c}\frac{D I'}{D t} + \frac{\mu}{r^2}\frac{\partial}{\partial r}\left[r^2 I'\right] + \frac{\partial}{\partial \mu'}\left((1-\mu'^2)\left[\frac{1}{r} + \frac{\mu'}{c}\left(\frac{v}{r} - \frac{\partial v}{\partial r}\right)\right] I'\right) \\ - \frac{\partial}{\partial \nu'}\left(\nu'\left[(1-\mu'^2)\frac{v}{cr} + \frac{\mu'^2}{c}\frac{\partial v}{\partial r}\right] I'\right) + \left[(3-\mu'^2)\frac{v}{cr} + \frac{(1+\mu'^2)}{c}\frac{\partial v}{\partial r}\right] I' \\ = \eta' - \chi' I', \end{aligned} \tag{10}$$

where $D/Dt = \partial/\partial t + v\partial/\partial r$. The many terms have their origin from that along a Lagrangian path, there are now changes in angle and frequency, in addition to space and time. Compare e.g. to the equation solved by the `CMFGEN` supernova code [Hillier and Dessart, 2012], which operates under the assumptions 1) $v/c \ll 1$ and 2) homology ($dv/dr = v/r$):

$$\frac{1}{c}\frac{\partial I'}{\partial t} + \frac{\mu' c + v}{c}\frac{\partial I'}{\partial r} + \frac{(1-\mu'^2)}{r}\frac{\partial I'}{\partial \mu'} - \frac{v\nu'}{rc}\frac{\partial I'}{\partial \nu'} + \frac{3v}{rc} I' = \eta' - \chi' I'. \tag{11}$$

Exercise 1: Derive Eq. 11 from Eq. 10 (by inserting the homology relation).

The complexity of the CMF RT equation even in 1D (terms 4 and 5 on the LHS do not exist in the lab frame, and term 2 is here also more complex) means solving it will be yet more challenging - frequencies and angles are now coupled even if they are not through the emissivity function (see week 1 lecture notes). This hints at attractiveness of a Monte Carlo approach for moving media,

instead of solving the PDE by discretization and algebraic equation system solving. In MC, we follow packets along their trajectories and update the energy, frequency, and propagation angle(s) along the way; this is not really more difficult in a moving medium compared to a stationary one - there are just more transformations to be done. And it's not really more difficult in 3D compared to 1D. The RT equation, on the other hand, "explodes" in its dimensionality and number of variable couplings, as it by construction keeps track of "everything, everywhere, at once".

1.3 Line absorption in supernovae

We will now look at how transfer through a line occurs in a moving medium. A photon can be absorbed and cause a photoexcitation of a bound electron from a lower state l to an upper state u. This is also called a bound-bound transition. The cross section can be written

$$\sigma_\nu^{bb} = \Phi_0 \phi(\nu - \nu_0) \tag{12}$$

where ϕ is the line profile (unit s, per Hz), which is normalized. The strength of the line is given by the parameter

$$\Phi_0 = \frac{h\nu_0}{4\pi} B, \tag{13}$$

where B is the Einstein absorption coefficient which is of order 10^{10} for an optical resonance line. Thermal broadening is of order $v_{th} = \sqrt{3kT/m_p} \sim 10$ km s^{-1}, so the width of the line profile function is $\Delta\nu_{th} = \nu_0 v_{th}/c \sim 10^{13}$ Hz.

The excited atom typically relaxes back by spontaneous emission (rate A s^{-1}, where A is the Einstein emission coefficient), either in the same transition ("resonance scattering") or in a set of branching ones ("fluorescence"). It can also be collisionally deexcited by a free electron ("thermalization"), but the chance of this is typically quite small.

1.3.1 Line blocking

Material from an explosion, like the galaxies in the Universe (from Big Bang) or the layers in a supernova, move at different velocities described (approximately) by homology $v \propto r$. One may show that in such homologous flows, each point sees the other points moving radially away from it with velocity proportional to the distance (think of Hubble's law for galaxies). If a photon is emitted at a comoving wavelength λ_1 at point 1, when it arrives at point 2, moving away from point 1 with relative velocity Δv, it will have a redshifted comoving wavelength (ignoring the γ factor)

$$\lambda_2 \approx \lambda_1 \left(1 + \frac{\Delta v}{c}\right) \tag{14}$$

Thus, as a photon travels from point to point, it is continously redshifted in each local, comoving frame.

Consider what this means for transfer through lines. If emitted in a frame with comoving wavelength λ_{start}, the photon will come into resonance with a line at $\lambda_0 (> \lambda_{start})$ only after travelling a distance corresponding to velocity difference $\Delta v \approx c \times (\lambda_0/\lambda_{start} - 1)$. Giving the line profile a finite width $\Delta\lambda_{th}$ (given by thermal motions), the interaction region is a small length interval corresponding to where the comoving wavelength is in the range $\lambda_0 - \Delta\lambda_{th}/2$ to $\lambda_0 + \Delta\lambda_{th}/2$. This length is called the Sobolev length, L_{Sob}.

In supernovae $v/c \sim 0.01 - 0.1$. This is much larger than the thermal line widths: $v_{th}/c \lesssim 10^{-4}$. This means that

1. The Sobolev length is much smaller than the radial scale of the SN and a given photon interacts in a small local region with any given line.

2. A photon can sequentially come into resonance with many lines, at sequential points along its path. If there is a distance v_{edge} to the edge of the nebula (as perceived from the emitting frame), the photon will be exposed to absorption by all lines with rest wavelengths lying in the range from λ_{start} to $\lambda_{start} \times (1 + v_{edge}/c)$.

Property 1 means a *simplification* of the line transfer process as each interaction can treated as a local event and whether absorption happens or not depends just on local gas conditions, and

there is no dependency on detailed line profile shapes. Property 2 means a *complication* compared to static nebulae, where any given photon only interacts with a single line, here we need to consider the complex problem of line interlocking.

Local line interaction (The "Sobolev limit"). The photon will 'traverse' the line profile over a length $L_{Sob} = v_{th}/(dv/dr)$, where dv/dr is the velocity gradient. As $v = r/t$ in homology, $dv/dr = 1/t$ and $L_{Sob} = v_{th}t$. As $v_{th} \ll v$ this region is small compared to the size of the SN, vt. For a top-hat line profile ($\phi = 1/\Delta\nu_{th}$):

$$\boxed{\tau_{Sob} = \sigma^{bb} n_l L_{Sob} = \frac{h\nu_0}{4\pi} B \frac{1}{\Delta\nu_{th}} \times n_l \times \frac{\Delta\nu_{th}}{\nu_0} ct = \frac{hc}{4\pi} B n_l t} \tag{15}$$

where n_l is the number density of the lower (absorption) level of transition. The result holds, in fact, irrespective of the shape of ϕ, as seen from a more rigourous integration:

$$\tau_{Sob} = \int \sigma^{bb}(\nu'(s)) n_l(s) ds = \int \Phi_0 \phi(\nu') n_l(s(\nu')) \frac{ds}{dv}\frac{dv}{d\nu'} d\nu' = \frac{h\nu_0 B}{4\pi} n_l \int \phi(\nu') \times t \times \frac{c}{\nu_0} \times d\nu' = \frac{h\nu_0 B n_l ct}{4\pi\nu_0} = \frac{c^2}{2h\nu_0^3} A \frac{hct}{4\pi} = \frac{1}{8\pi} A \lambda_0^3 n_l t, \tag{16}$$

where in the last step we have assumed the resonance region to be small enough that $n_l(s)$ is constant over it. The optical depth depends only on the local number density (n_l) at the point of resonance. Ignoring changes in ionization/excitation, $n_l \propto t^{-3}$ in homology, so $\boxed{\tau_{Sob} \propto t^{-2}}$.

Expansion opacity. If there is a typical velocity separation Δv_{sep} between optically thick lines around wavelength λ, the mean-free path is $\lambda_{mfp} = \Delta v_{sep} t$ (note we use now λ for both mean-free-path and wavelength, the former has a $_{mfp}$ subscript). Write $\Delta v_{sep} = c\Delta\lambda_{sep}/\lambda$. Then, since $\kappa = 1/(\lambda_{mfp}\rho)$,

$$\boxed{\kappa_\lambda^{line,exp} \approx \frac{\lambda}{ct\rho\Delta\lambda_{sep}}} \tag{17}$$

A more refined formula can be obtained from knowing the probability that line i interacts is $1 - \exp\left(-\tau_{Sob}^i\right)$. Then refine the formula as

$$\boxed{\kappa_\lambda^{line,exp} \approx \frac{1}{ct\rho}\frac{1}{N}\sum_{i=1}^{N} \frac{\lambda_i}{\Delta\lambda_i}\left(1 - e^{-\tau_{Sob}^i}\right)} \tag{18}$$

where the set of N lines have to be chosen over some wavelength range centred on λ.

This is called an expansion opacity. Figures 1 shows an illustration. Note that this is not an exact local opacity like the usual ones - this is because line interaction is not a continuous process but occurs at discrete points in the Sobolev limit. As opacity can only be defined as some integral/summation over lines, there is always some degree of arbitrariness how this is done - N is basically a free parameter in Eq. 18. Line opacity is together with Thomson opacity typically the most important in SNe.

Finally, for line absorption fluorescence is often a more frequent deexcitation process that scattering back into the same transition (Fig. 2). Complex fluorescence means the emergent spectra of SNe are, in detail, complex to model. In particular at late times several years after explosion, models have shown that the whole optical spectrum of SN can come from fluorescence of UV emission.

2 Line formation in early ("photospheric") phases of a supernova

Because $v = rt$ in homology, all points on a sheet perpendicular to the line of sight (LOS) have the same LOS velocity and give the same Doppler shift relative to the observer (Fig. 3). If the sheet is emitting in a line with rest wavelength λ_0, the observer will see emission at wavelength $\lambda_0 \times (1 + v_{los}/c)$, where positive projected velocities are away from the observer (to the right in

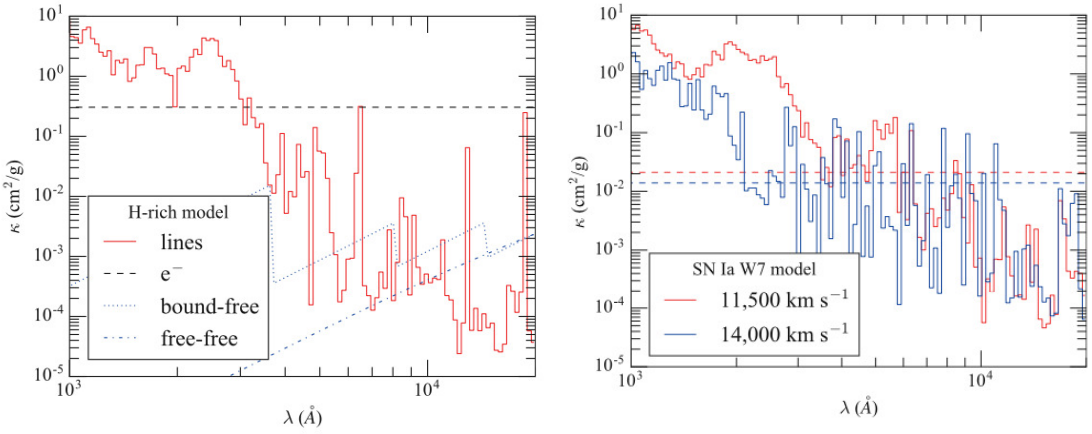

Figure 1: *Top*: Opacities in a Type IIP SN layer at 11,000 km/s at $t = 15$d. *Bottom*: Thomson opacity and line opacity in two layers in a Type Ia SN around maximum light. From Sim 2017.

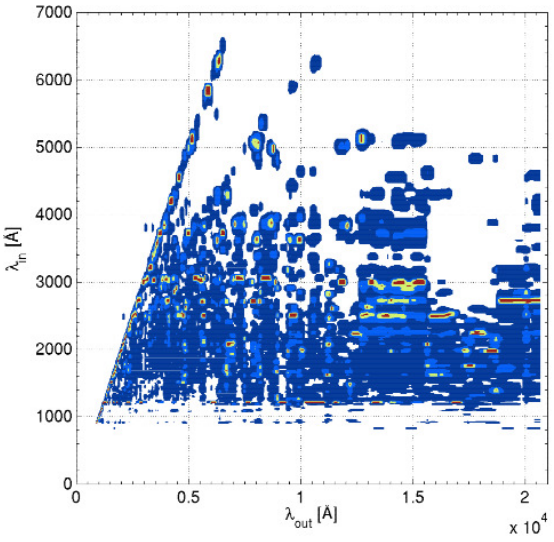

Figure 2: Fluorescence redistribution in a model of SN 1987A at an age of 8 years. From Jerkstrand et al 2011.

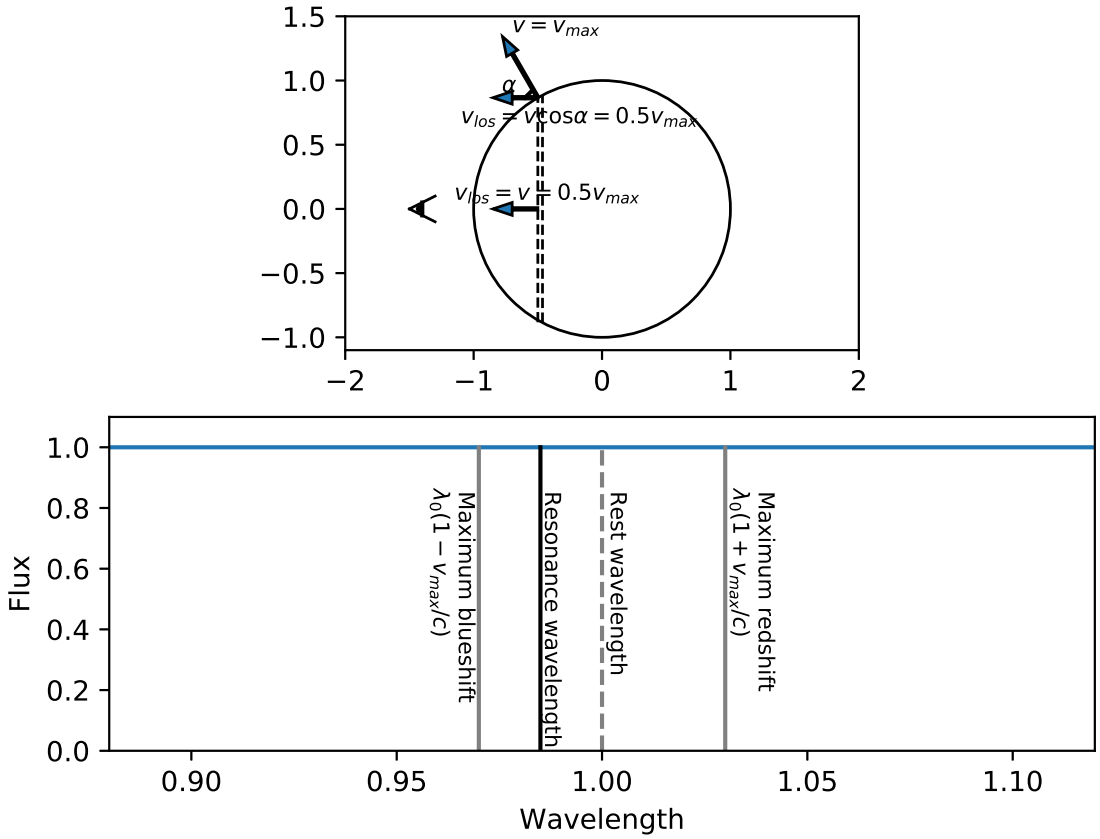

Figure 3: Illustration of how a sheet perpendicular to the LOS has constant LOS velocity, and therefore constant Doppler shift with respect to the observer. The bottom panel shows the corresponding wavelengths if each LOS point, in this example λ values are calculated for $v_{max}/c = 0.03$.

the figure). If the line absorbs, the sheet is capable of blocking photons at λ_0 in the CMF, equaling $\lambda_0 \times (1 + v_{los}/c)$ in the LF, travelling from the region it covers from the observer, giving an absorption line at the same wavelength.

The scattering atmosphere (also called the Schuster-Schwarzschild model) is a simplified model framework where blackbody radiation is emitted from an inner boundary - the photosphere - and then scatters (by electron scattering and line scattering) in the outer layers. It is a simplified concept because in a real supernova

- The location of the photosphere (formally defined as where optical depth integrating from the outside inwards is $\tau_\nu = 2/3$) varies with frequency.
- Even for a fixed frequency, the photosphere does not sharply divide into a scattering-only region on the outside and a thermalization-only region on the inside.

In fact, advanced radiative transfer models for SN spectra do not rely on the scattering atmosphere ansatz but model also the transition region and the deeper layers. Nevertheless, it is a framework that allows a good understanding of the basic processes forming the spectrum, and rough spectral models to be developed that by comparison with observations allow reasonably accurate estimates of the density profile and composition of the outermost layers. It is the foundation for the open-source supernova spectral modelling codes `SYNOW/SYN++` [Thomas et al., 2011] and `TARDIS` [Kerzendorf and Sim, 2014].

Consider a scattering atmosphere where line absorption is treated in the Sobolev limit. We will study what line profile arises from a single optically thick line ($\tau_{Sob} \gg 1$), with the two parameters $h = v_{phot}/v_{max}$, where v_{max} denotes the maximum velocity at which the line is optically thick,

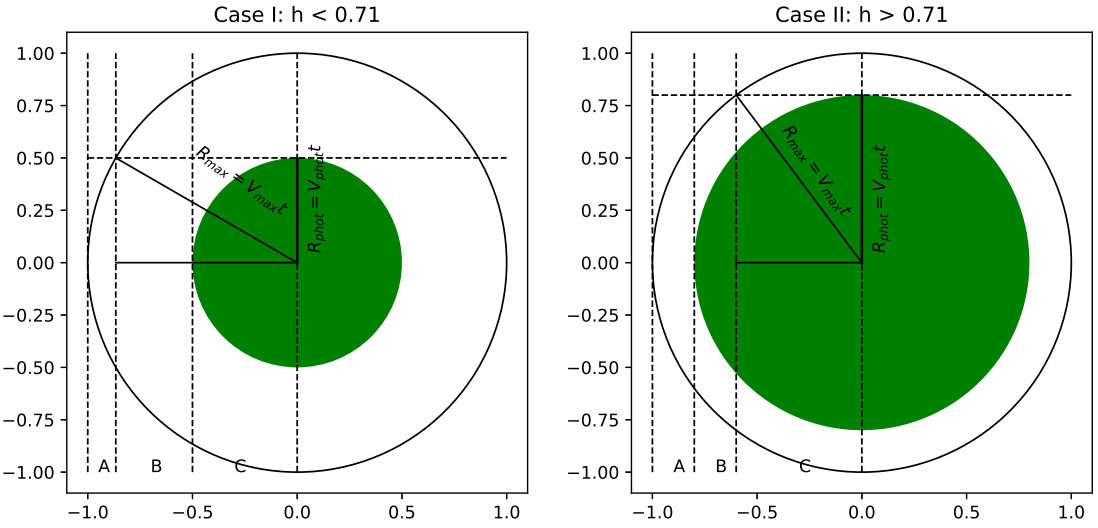

Figure 4: Geometry of P-Cygni line formation. The observer is towards the left.

and ϵ is the "destruction probability" (probability $1 - \epsilon$ for scattering).

Consider Fig 4. We delineate two cases.

- **Case I**. $h < 1/\sqrt{2} = 0.71$. The photosphere is fully blocked by certain sheets (in B region) $\rightarrow$ complete absorption at the Doppler velocity of those sheets.
- **Case II**. $h > 1/\sqrt{2}$. The photosphere is never fully blocked by a sheet $\rightarrow$ only partial absorption.

The ABC regions behave as following:

- **Region A**: The whole resonance sheet covers (part of) the photosphere. The area of the sheet grows going inwards (towards the rest wavelength), giving a deeper absorption as a larger fraction of the photosphere gets blocked.
- **Region B-Case I**: (Part of) the resonance sheet covers the whole photosphere. Complete absorption is produced throughout region B.
- **Region B-Case II**: The (whole) resonance sheet covers part of the photosphere. The sheet area is constant, giving a flat bottom in the absorption profile but at non-zero flux. The absorption depth increases with decreasing h.
- **Region C**. Part of the sheet covers part if the photosphere. A declining fraction of the photosphere is blocked moving towards the centre, giving declining degree of absorption moving towards the rest wavelength.

Exercise 2. P-Cygni lines arise when lines scatter photospheric radiation in a moving atmosphere. Use the code in mode "linescattering", to study P-Cygni line formation. Create a figure showing a simulated P-Cygni line, where you mark out how the parameters $(v_{phot}, v_{max}, \epsilon)$ can be determined.

By inferring v_{max}, one may be able to determine the density of the element at that velocity by equating the line optical depth to unity there.

Exercise 3. From a SN spectrum observed at 50d post-explosion, v_{max} for the Ca II triplet is inferred to 10,000 km/s. If $T = 6000$ K and you can assume all Ca is in the Ca II ionization stage, and excited levels are in LTE, what is the calcium density at this velocity?

You can find information on atomic energy levels here:
https://physics.nist.gov/PhysRefData/ASD/levels_form.html,
and on line transitions here:
https://physics.nist.gov/PhysRefData/ASD/lines_form.html.

One should note, however, than in more realistic models the interpretation of line profiles can be more complex (see e.g. Fig 6 in Sim 2017). The Sobolev line optical depth is a function of velocity (it's a local quantity) and does not give a sharp boundary as here. Also, there may be emission in the line in the atmosphere by additional processes than resonance scattering (e.g. recombination).

2.1 Thermal line broadening

We have so far assumed that scattering has been fully coherent in the comoving frame, so the photons just change direction but not frequency when they scatter. While this is an accurate description in the atom frame (for the type of interactions currently being considered), the atom has a thermal velocity in the comoving frame. This thermal velocity is $v_{therm} \approx \sqrt{\frac{3kT}{m}}$ (from $3/2kT = 1/2mv^2$), where m is the mass of the atom. This becomes $v_{therm}^{atom} = 11$ km s^{-1} $(T/5000\text{ K})^{1/2}$ $A^{-1/2}$, where A is the atomic mass. This is small compared to the Doppler broadening occurring due to the SN expansions which is typically $10^3 - 10^4$ km s^{-1}.

For an electron, however, the thermal broadening is $v_{therm}^{e-} = 480$ km s^{-1} $(T/5000\text{ K})^{1/2}$. Thus electron scattering, especially multiple ones, is capable of broadening (emission) lines with $\gtrsim 10^3$ km s^{-1}. This is still quite a minor effect in SNe emitting lines from regions moving with several thousand kilometers per second. But some SNe, in particular Type IIn, emit lines also from slower-moving material, ~ 100 km s^{-1}. If these layers are hot, the line profiles may become determined by thermal broadening rather than expansion broadening.

Exercise 4. Use the code in mode "electronscattering" to simulate emission profiles from a uniform sphere (take vlim = vmax = 1000 km/s) with different degrees of electron scattering opacity. Study what happens when the electron scattering optical depth is raised from 0 to 1 to 2. For a fixed optical depth (e.g. $\tau_{es} = 1$), how does the temperature impact the line profile (compare e.g. 0 K, 10,000 K, and 30,000 K)? Make two figures showing these comparisons.

3 Line formation in late ("nebular") phases of supernovae

At later times ($\gtrsim$ 100d), the photosphere disappears (it continuously recedes inwards in mass coordinate) and we enter the nebular phase. Now, the spectrum consists of emission lines. As powering comes from radioactivity, emission comes mainly from the inner regions where ^{56}Ni resides. Figure 5 shows an example of an observed nebular spectrum illustrating the quite different nature to photospheric spectra.

3.1 Line profiles

Let us consider line profiles of emission lines arising for a few simple cases of geometry. Figure 6 shows line profiles in six cases, of which we comment on three:

- **Uniform sphere (upper left)**. Letting v_p denote the component of v_{max} perpendicular to the line-of-sight, each sheet contributes flux at wavelength $\lambda_0 v_{los}/c$ in proportion to its area

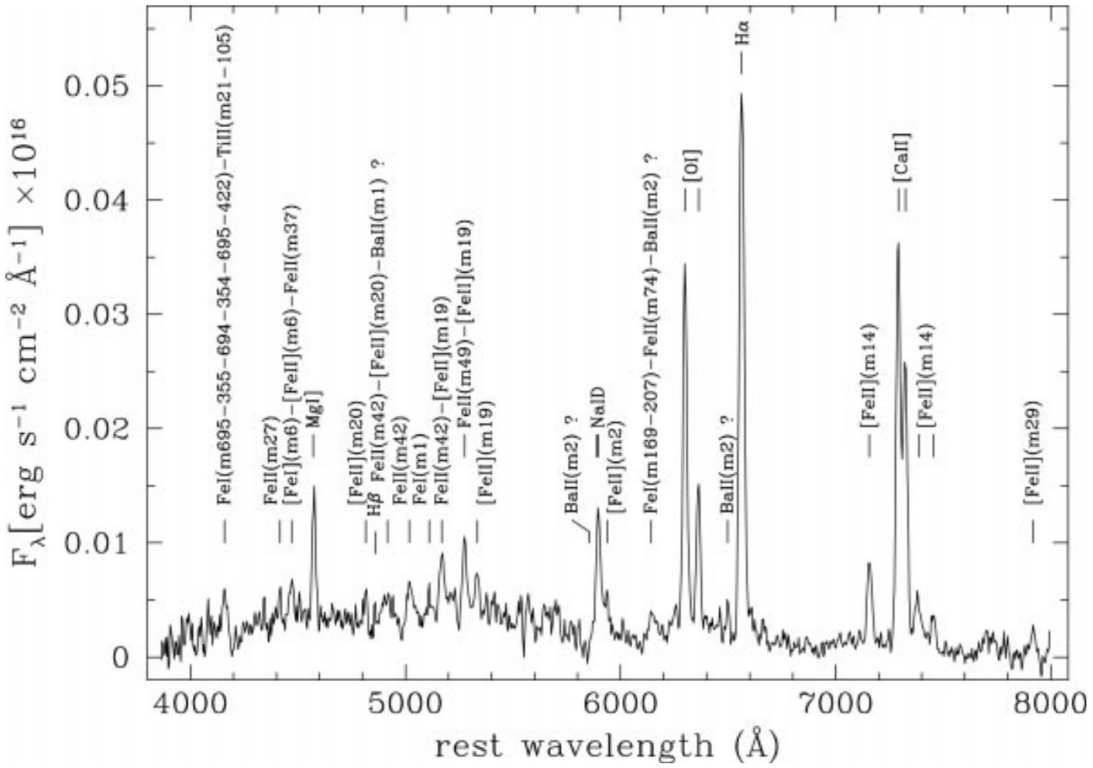

Figure 5: Example of an observed nebular spectrum from a low-velocity Type IIP SN, from Benetti et al. 2001. As the photosphere has disappeared there are no P-Cygni lines, instead the spectrum is made of emission lines, some of which are very prominent. Models show that what looks like "continuum" (e.g. between 6600-7000 Å) is in fact thousands of overlapping weaker emission lines.

which grows as $\pi\left(v_p t\right)^2 = \pi t^2\left(v_{max}^2 - v_{los}^2\right) = \pi t^2 v_{max}^2\left(1-(v_{los}/v_{max})^2\right)$. The line profile is therefore a parabola.

- **Thin shell (upper middle)**. A slice of a shell can be shown to give a constant area, so the line profile becomes a flat top.
- **Gaussian (bottom left)**. One can show that also the line profile becomes a Gaussian.

For further details of the derivations see Jerkstrand 2017, Handbook of SNe.

3.2 Luminosity

The profile-integrated emissivity of a line is

$$\tilde{j} = \frac{1}{4\pi} n_u h\nu A \beta_{Sob} \quad \text{erg s}^{-1}\text{cm}^{-3}\text{ster}^{-1} \tag{19}$$

where n_u is the number density of the upper level of the line (unit cm^{-3}), A is the Einstein A-coefficient (rate of spontaneous radiative deexcitation, unit s^{-1}) and $\beta_{Sob} = (1-e^{-\tau_{Sob}})/\tau_{Sob}$ is the Sobolev escape probability : the chance that an emitted photon will escape the local line resonance region rather than be reabsorbed. Consider two limits:

- **Optically thin line** ($\beta_{Sob} \rightarrow 1$). The total volume-integrated luminosity is then $L = 4\pi\tilde{j} \times V = N_u h\nu A = \frac{M_{ion}}{\mu_{ion}} \times f_u(T, n_e) h\nu A$, where the function $f_u(T, n_e)$ determines what fraction of the ions of this type are in state u. If this can be estimated, one may then be able to determine the mass of the ion.
- **Optically thick line** ($\beta_{Sob} \rightarrow 1/\tau_{Sob}$). The total luminosity is now $L = V \times \frac{n_u h\nu A}{hcBn_l t}$, where B is the Einstein B- coefficient. Because B can be related to A, this can be written $L = V \times t^{-1} \times g_{ul}(T, n_e)$ where the function g_{ul} determines the ratio of populations g_u/g_l. Thus, if this can be estimated one can determine the volume of the emitting region.

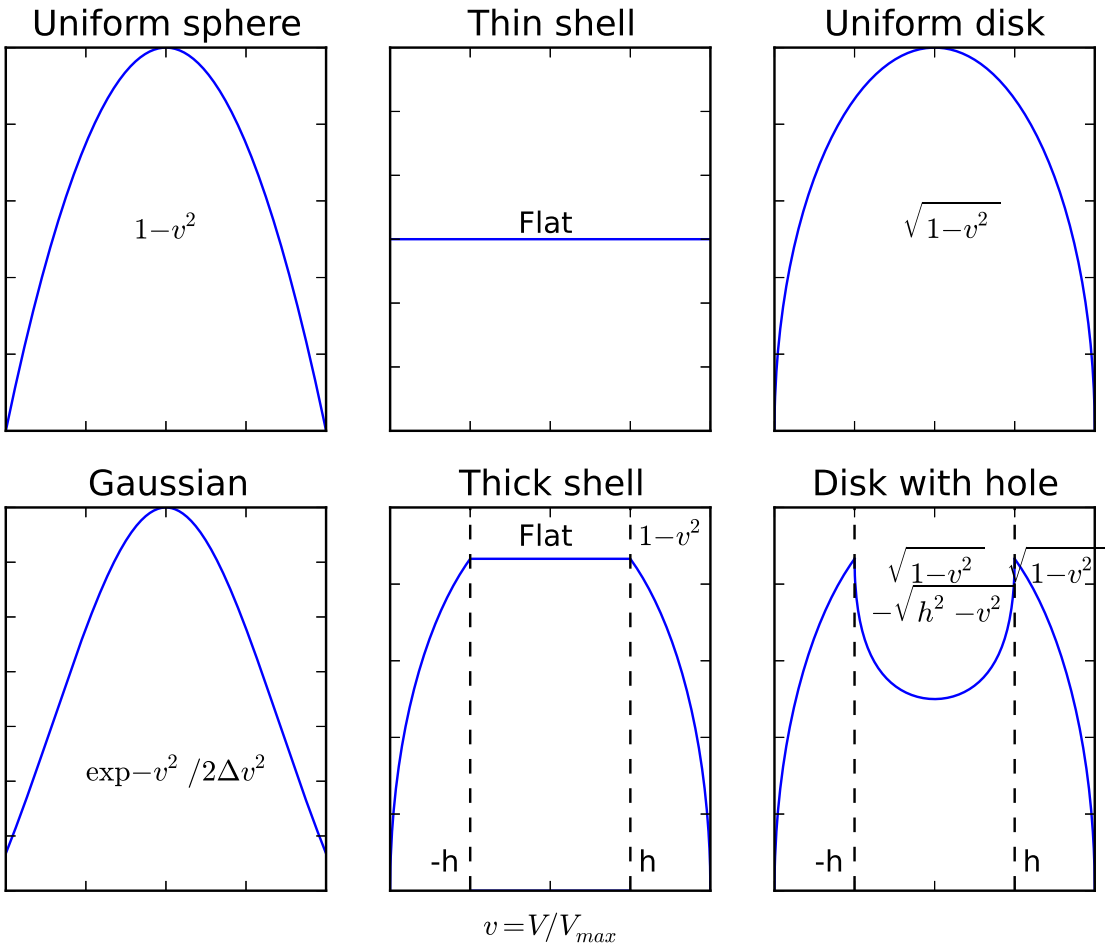

Figure 6: Nebular line profiles for 6 different cases. From Jerkstrand 2017 (Handbook of SNe).

To get either a mass or volume then, we need to know what fraction of atoms for that ion is in the upper state u (optically thin case), or what the ratio of upper and lower state populations (n_u/n_l) are (optically thick case).

At high electron density, collisions dominate the populations and depopulations of states, and then the f_u and g_{ul} functions depend only on temperature ("Local Thermodynamic Equilibrium (LTE)" limit). Thus, in this limit only T needs to be estimated. However, one should be careful: the f_u and g_{ul} functions have exponential dependencies on the temperature. Thus, if $\Delta E \gg kT$, where ΔE is the excitation energy of the upper level E_u for optically thin lines and $E_u - E_l$ for optically thick ones, even small errors in the T estimate can give very different factors. The result will only be robust if $\Delta E \lesssim$ few $\times\, kT$.

4 Monte Carlo Radiative Transfer

Monte Carlo (MC) methods refer to a class of algorithms making use of random sampling to solve mathematical or physical problems. The method has its original applications based on physical experiments whose random outcomes would mimic the problem. For example, let's say you want to determine the area of an odd shape drawn on the floor in a room. One way you could do this would be to throw a bouncing ball say 100 or 1000 times into the room, and after bouncing around on the walls see what percentage of time the ball comes to rest within the domain.

With the rise of computers, and development of algorithms able to create (pseudo) random numbers, MC methods became increasingly popular in the 1950s. In the late 1960s, MC calculations

finally entered the astrophysics stage with studies such as Avery and House [1968], Auer [1968] and House and Avery [1968]. In the time since, MC methods have become established, successful, and reliable tools for the study of a variety of astrophysical radiative transfer phenomena.

What is random, exactly? A point is worth dwelling on for a moment is that the randomness in a MC method can refer to very different things. Quantum mechanics established that nature itself is random - a huge conceptual leap in physics. Therefore, MC methods can directly simulate this randomness e.g. by making random draws on which deexcitation paths an excited atom will take or how radioactive nuclei decay. The power of MC methods in this regime is that the MC packets become tracers automatically following the most important paths in the system. *MC noise will be smallest where it's most useful that it's small - in the dominant flow channels and output signals.* It allows a compute power optimization that qualitiatively does not exist when solving RT equations. But conversely, if one is interested in some rare or energetically unimportant signal, a MC approach may not be a good idea.

But these kind of problems - randomly sampling the true randomness in nature - are only one subgroup of MC methods. In other groups of methods, the random draws may instead be used for..

1. Selecting injection particles.
2. Choosing interactions outcomes which in principle are deterministic (e.g. scattering angle as depending on impact parameter), but are chaotic (e.g. very sensitive to initial conditions) and/or the incoming particles would fill a property space that one wants to avoid keeping track of the details for.
3. Random sampling as an approximate statistical description of a process or domain that is deterministic but the full details are not meaningful. For example, for transfer through a domain characterized by a large number of clumps, randomization of paths and clump encounters captures the physical effects without needing to specify and keep track of a complex specific realization [e.g. Jerkstrand et al., 2011].

The common denominator for the different MC method groups is really that one uses physical insight to replace complex, exact equation solving with statistical descriptions that are easier to code, and in addition gives the algorithm a built-in power to spend compute power where it most matters. The focus shifts from large equation system solving and numerics, to the microphysical interactions, particle-by-particle.

4.1 Monte Carlo basics

True randomness is difficult to achieve on a machine that is inherently deterministic (quantum computers are not yet here), but for essentially all practical purposes the "pseudo-randomness" that computers can produce are sufficient [see e.g. Kalos and Whitlock, 2008, §9]. Based on a starting value (referred to as a *seed*), these algorithms provide sequences of numbers ξ, typically uniformly distributed over the interval $[0, 1]$. Such sequences are referred to as "pseudo" random since they share statistical properties with true randomness but are still generated by relying on deterministic prescriptions.

The probability that a random variable X takes a value within $[x, x + \mathrm{d}x]$, is encoded in the so-called *probability density function (PDF)* $\rho_X(x)\mathrm{d}x$, which fulfills the normalization

$$\int_0^\infty \rho_X(x)\mathrm{d}x = 1. \tag{20}$$

The *cumulative distribution function (CDF)* is the probability that X takes a value between 0 and x:

$$f_X(x) = \int_0^x \rho(x')\mathrm{d}x'. \tag{21}$$

Unlike the PDF, which is always positive but not necessarily monotonous, the CDF per construction always is. Consequently, a 1-to-1 mapping between two cumulative probability distributions can be established via

$$f_X(x) = f_Y(y). \tag{22}$$

This is the fundamental relation needed for sampling a probability distribution $\rho_X(x)$ using draws from another one, $\rho_Y(y)$. Using the random numbers ξ, uniformly distributed between 0 and 1 and thus giving a CDF $f_\xi(\xi) = \xi$, this simplifies to

$$f_X(x) = \xi. \tag{23}$$

Inversion of this equation results in the central sampling rule

$$\boxed{x = f_X^{-1}(\xi).} \tag{24}$$

Example 1: Scattering angles. Consider the situation of isotropic scattering of a photon. In this case, no propagation direction after the interaction is preferred and the probability that the photon escapes into a specific solid angle element $\mathrm{d}\Omega = \mathrm{d}\phi\,\mathrm{d}\theta\sin\theta$ (with $\phi \in [0, 2\pi]$ and $\theta \in [0, \pi]$) is constant

$$\rho(\phi, \theta)\,\mathrm{d}\phi\,\mathrm{d}\theta\sin\theta = \mathrm{c}_1. \tag{25}$$

With the introduction of the μ parameter

$$\mu \equiv \cos\theta, \tag{26}$$

and writing $\rho(\phi, \theta) = c_2 \times \rho_1(\phi)\rho_2(\mu)$ (no correlation exists between the angles so the joint PDF is separable) this reduces to

$$\rho_1(\phi) = \mathrm{c}_3, \tag{27}$$
$$\rho_2(\mu) = \mathrm{c}_4. \tag{28}$$

Normalization gives

$$\int_0^{2\pi} c_3 d\phi = 1 \rightarrow c_3 = \frac{1}{2\pi}, \tag{29}$$

and

$$\int_{-1}^{+1} c_4 d\mu = 1 \rightarrow c_4 = \frac{1}{2}. \tag{30}$$

Then, $f_{\rho_1}(\phi) = \phi/(2\pi)$ and $f_{\rho_2}(\mu) = \frac{1}{2}(\mu + 1)$. Finally, equate the CDFs with the uniform random number ξ and solve for ϕ and μ, giving (careful to use two *different* random numbers ξ_1 and ξ_2)

$$\phi = 2\pi\xi_1, \tag{31}$$
$$\mu = 2\xi_2 - 1. \tag{32}$$

Example : Interaction point. Another important application of random sampling is the decision of when the packet will interact. The probability of interaction per path segment ds, *assuming the packet has survived to point* s, is given by the absorption coefficient χ (unit cm^{-1}). The probability to have survived to at least s is as function $g(s)$ that is simply the complement $g(s) = 1 - f(s)$ to the CDF function $f(s)$. The probability to interact in $[s, s + ds]$ is then given by $g(s) - g(s + ds) = g(s) \times \chi(s)ds$, or

$$-\frac{dg(s)}{ds} = g(s)\chi, \tag{33}$$

which, if χ is constant, has solution

$$g(s) = \exp(-\chi s). \tag{34}$$

Finally, obtain the random sampler for s by replacing $g(s)$ with $1 - f(s)$ and inverting, giving

$$s = -\frac{\ln(1 - \xi)}{\chi}. \tag{35}$$

By using an optical depth variable τ instead of a positional variable s, we can also allow for χ varying along the path, and the corresponding sampling becomes

$$\boxed{\tau = -\ln(1 - \xi).} \tag{36}$$

This is a powerful result because it means we can determine the "life span" τ of any packet at it's creation - and simply find the point in the domain where this τ will be reached. We can transport the packet directly to that point, and compute the interaction there, usually with another random sampling process.

Non-analytic CDFs. In these examples so far, the underlying cumulative distribution function could be inverted analytically to give direct formulas for the random sampling. Naturally, this is not always feasible, and in such cases one has to rely on numerical integration and bin finding, which raises the cost of the random sampling significantly. Significant literature exists on doing this as efficiently as possible for commonly used non-integratable PDFs, such as the Planck function [Fleck and Cummings, 1971, Carter and Cashwell, 1975, Bjorkman and Wood, 2001].

Exercise 5: Derive the CDF for Thomson scattering, which has a PDF $\rho_2(\theta) = 3/8\left(1+\cos^2\theta\right)\sin\theta$. Is this CDF invertable to give an equivalent expression to Eq. 32? Plot the Thomson CDF against the isotropic one and discuss with your neighbor whether you think it's worth implement a numeric sampling algorithm.

Packet propagation and interaction. Figure 7 shows an illustration of how the next interaction point, and process, is selected. While line interactions are local and "stand-alone" (Sobolev limit), if interaction occurs by continuum absorption, a random draw determines which of the several contributing opacities should be selected. This would in the simplest case be done by assigning selection probabilities in proportion to the contribution to the total opacity:

$$p_{bf} = \frac{\chi^{\rm bf}}{(\chi^{\rm bf}+\chi^{\rm ff}+\chi^{\rm es})}, \tag{37}$$

$$p_{ff} = \frac{\chi^{\rm ff}}{(\chi^{\rm bf}+\chi^{\rm ff}+\chi^{\rm es})}, \tag{38}$$

$$p_{es} = \frac{\chi^{\rm es}}{(\chi^{\rm bf}+\chi^{\rm ff}+\chi^{\rm es})}. \tag{39}$$

Whenever a packet experiences an interaction, its properties are updated in accordance with the physical laws underlying the interaction as outlined in the next section. Following these principles, each packet is moved through the domain until a termination condition is reached. Depending on the problem, this may be an absorption interaction, or the packet reaching a domain surface or simply that a pre-defined amount of physical time has elapsed. The propagation process is complete after all MC packets have been processed in this manner. During the propagation process, various events may be recorded or the change of certain packet properties tracked. These may then be used to reconstruct physical properties of the radiation field as described further below.

4.2 Interaction Processes

In moving media it is common to adopt a so called **mixed frame approach** for the calculation of the packet trajectories and interactions. All tasks regarding radiation-matter interactions are handled in the CMF as motivated earlier. The energy and wavelength of the packet are maintained in the current CMF, continuously being transformed to new CMF as the packet reaches new points. The spatial packet propagation is carried out in the lab frame since the discretization of the computational domain is done in this frame.

If $v/c \ll 1$, as typically the case for both SNe $v/c \lesssim 0.03$) and KNe ($v/c \lesssim 0.2$), we really don't need to worry about the effects on relativistic beaming or movement of the grid over the photon flight time, these lead to changes of the emergent intensities of order v/c. The only important effect induced by the motion is the frequency Doppler shift as $v_{thermal} \ll v$, v/c shifts easily shift photons into and out of resonance with lines also when $v/c \ll 1$. Thus, this limit can follow the scheme

1. From a starting (E'_0, ν'_0, μ'_0) in the CMF at radius r_0, take the lab frame angle as $\mu_0 = \mu'_0$.

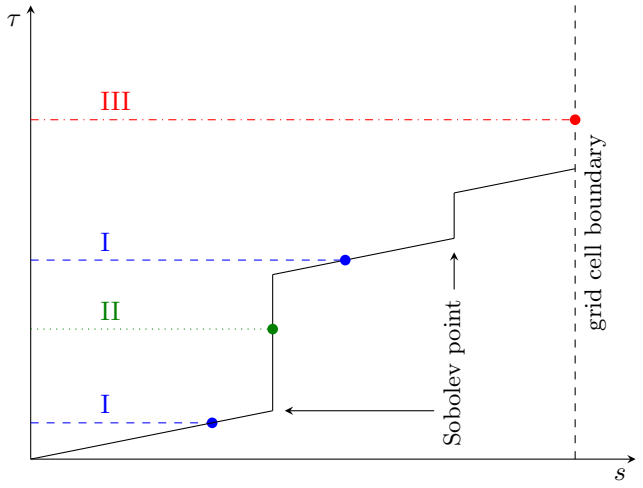

Figure 7: Schematic diagram of the accumulated optical depth τ along a packet trajectory [taken from Noebauer and Sim [2019], in turn adapted from Mazzali and Lucy [1993]]. The optical depth τ exhibits a linear increase with the traveled distance s due to continuum opacity from bound-free, free-free and Thomson scattering processes. Discontinuous jumps occur at the Sobolev points, where the packet comes into resonance with a line. At the start of a propagation step, a random optical depth τ_n is assigned to the packet. The next event is then identified by finding the position at which the accumulated optical depth of the packet $\tau(s) = \tau_{\rm n}$. This is illustrated for four exemplary random optical depths, resulting in (I) a continuum absorption,(II) a line absorption and (III) escaping the cell boundary.

2. Compute the LF distance d to the next interaction point (ignoring movement of the nebula during the photon flight[1])

3. Move packet to the point of next interaction (depends on d and μ_0), i.e. update the values for r and μ. This may involve stopping at cell boundaries to update opacity values.

4. Transform the pre-interaction packet properties from the first CMF ("0") to the new CMF ("1"), $(E'_0, \nu'_0, \mu'_0) \rightarrow \left(E'_{1,\rm in}, \nu'_{1,\rm in}, \mu'_{1,\rm in}\right)$.

5. Impose energy conservation in the new CMF; $E'_{1,\rm in} = E'_{1,\rm out}$ to fix the post-interaction energy $E'_{1,\rm out}$. For coherent scattering in the CMF the post-interaction frequency is given by $\nu'_{1,\rm in} = \nu'_{1,\rm out}$. If, for example, the interaction is a line absorption with radiative rate A to resonance scatter and $A/2$ to fluoresce, giving two photons of frequency $1/3\nu_0$ and $2/3\nu_0$ in the down-cascade, we can keep the condition $E'_{1,\rm out} = E'_{1,\rm in}$ by assigning probabilities

$$p_{res} = \frac{A h\nu_0}{A h\nu_0 + \frac{1}{2} A h\nu_0 \frac{1}{3} + \frac{1}{2} A h\nu_0 \frac{2}{3}} = \frac{4}{9}, \tag{40}$$

$$P_{f1} = \frac{\frac{1}{2} A h\nu_0 \frac{1}{3}}{A h\nu_0 + \frac{1}{2} A h\nu_0 \frac{1}{3} + \frac{1}{2} A h\nu_0 \frac{2}{3}} = \frac{1}{9}, \tag{41}$$

$$P_{f2} = \frac{\frac{1}{2} A h\nu_0 \frac{2}{3}}{A h\nu_0 + \frac{1}{2} A h\nu_0 \frac{1}{3} + \frac{1}{2} A h\nu_0 \frac{2}{3}} = \frac{2}{9}, \tag{42}$$

   By weighting the transition probabilities with the transition energies like this, one can avoid needing to split the MC packets, which makes codes easier to construct and the runs are more controllable. By enforcing $E_{out} = E_{in}$ in each interaction, one also enforces the radiative equilibrium condition, which acts as a Lambda accelerator for MC algorithms.

6. Sample a new propagation direction $\mu'_{1,\rm out}$, e.g. via $\mu'_{1,\rm out} = 2\xi - 1$ if isotropic scattering.

[1]The errors in r and μ for the new CMF, ignoring the nebula motion, is of order v/c, which is $\lesssim 0.1$ for SNe and KNe. Consideration of this effect requires, in general, a time-dependent MC simulation where packets are moved in small time steps, the spatial grid is updated, etc.

7. Loop over.

When the packet reaches the outer edge of the domain, a final transformation to the LF (=observer frame) gives the recorded properties of the packet for the observer.

4.3 Estimation of Radiation Field Quantities

The MC packet trajectories can be used to reconstruct radiation field quantities. In general, we need these to compute radiative ionization, excitation and heating rates, which all depend only on various moments of the radiation field, such as J_ν.

We can compute the **angle-averaged mono-chromatic photon number density** in cell i as [Lucy, 1999]:

$$4\pi\psi_{i,\nu} = \frac{\sum_{p=1}^{pass,i} \frac{E_p}{h\nu_p}\frac{\delta t_p}{\Delta t}}{V_i}, \tag{43}$$

where V_i is the cell volume and summation occurs over all packet path segments in the cell, with lengths $l_p = c\delta t_p$. Δt is the time duration of the MC simulation (usually 1 s). If we don't divide out the photon energies $h\nu_p$ we get an estimator for the **monochromatic radiation energy density**

$$E_{i,\nu} = \frac{\sum_{p=1}^{pass,i} E_p \frac{l_p/c}{\Delta t}}{V_i}, \tag{44}$$

Making use of the relation $J_\nu = \frac{c}{4\pi}E_\nu$ (see introduction week lecture notes, section 3.1) we get

$$J_{i,\nu} = \frac{1}{4\pi\Delta V_i \Delta t}\sum_{p=1}^{pass,i} E_p l_p. \tag{45}$$

In analogy, higher moments of the specific intensity may be reconstructed. In 1D:

$$H_{i,\nu} = \frac{c}{4\pi V_i \Delta t}\sum_{p=1}^{pass,i} E_p \mu_p l_p, \tag{46}$$

$$K_{i,\nu} = \frac{1}{4\pi V_i \Delta t}\sum_{p=1}^{pass,i} E_p \mu_p^2 l_p. \tag{47}$$

One can simplify even further. If the main goal of the simulation is to estimate for example a photoionization rate, we can directly accumulate this estimator from

$$\gamma_i = \int_0^\infty \frac{J_\nu}{h\nu}\sigma_\nu d\nu = \frac{1}{4\pi V_i \Delta t}\sum_{p=1}^{pass,i} \frac{E_p l_p}{h\nu_p}\sigma(\nu_p) \tag{48}$$

So the radiation field J_ν does not even need to be reconstructed and stored in each cell - not an issue for 1D modelling but perhaps for 3D.

Note that with these volume-based estimators, we (1) do not rely on any boundary crossings, and (2) packets contribute to estimators also when they travel between interaction points.

Exercise 6. Modify the code to emit a blackbody spectrum from the photosphere, and then compute the mean intensity as function of velocity coordinate. This means

- Develop a random sampler for $T = 8000$ K blackbody emission from the photosphere. Set vlim(= vphot) = 5000 and vmax = 10000 km/s. Note that the wlmin and wlmax parameters have to be changed to now allow binning of a full blackbody range of photon energies.
- Choose an atmosphere - meaning define some sort of obstacles to the radiation. Easiest will be to add a bunch of lines (adding line and electron scattering opacity together will require some code reworking), with ϵ values of your choice.
- Insert an accumulator to compute $J_{i,\nu}$ in each shell i, following Eq. 45.
- Run the simulation, save the $J_{i,\nu}$ vectors.
- Post-process these to compute the radiative heating rate in erg/s/cm^3 (derive the formula first) due to ground state ionization of H, for some assumed density of H, and that all H is in the ground state (this process does not need to be part of the actual MC simulation).

Bonus (if you have time) only): Replace the single-shell format of the code with multiple shells (take e.g. 10 equally spaced), and compute the heating rate as function of velocity coordinate.

References

L. H. Auer. Transfer of Lyman Alpha in Diffuse Nebulae. *ApJ*, 153:783, September 1968. doi: 10.1086/149705.

L. W. Avery and L. L. House. An Investigation of Resonance-Line Scattering by the Monte Carlo Technique. *ApJ*, 152:493–+, May 1968. doi: 10.1086/149566.

J. E. Bjorkman and K. Wood. Radiative Equilibrium and Temperature Correction in Monte Carlo Radiation Transfer. *ApJ*, 554:615–623, June 2001. doi: 10.1086/321336.

L. L. Carter and E. Cashwell. Particle-transport simulation with the Monte Carlo method. Technical report, 1975. URL http://www.iaea.org/inis/collection/NCLCollectionStore/_Public/07/227/7227109.pdf.

J. A. Fleck and J. D. Cummings. An implicit monte carlo scheme for calculating time and frequency dependent nonlinear radiation transport. *Journal of Computational Physics*, 8(3):313 – 342, 1971. ISSN 0021-9991. doi: 10.1016/0021-9991(71)90015-5. URL http://www.sciencedirect.com/science/article/pii/0021999171900155.

D. J. Hillier and L. Dessart. Time-dependent radiative transfer calculations for supernovae. *MNRAS*, 424:252–271, July 2012. doi: 10.1111/j.1365-2966.2012.21192.x.

L. L. House and L. W. Avery. The Monte Carlo Technique Applied to Radiative Transfer. In *Resonance Lines in Astrophysics*, page 133, December 1968.

A. Jerkstrand, C. Fransson, and C. Kozma. The ^{44}Ti-powered spectrum of SN 1987A. *A&A*, 530: A45, June 2011. doi: 10.1051/0004-6361/201015937.

M. H. Kalos and P. A. Whitlock. *Monte Carlo Methods: Second Revised and Enlarged Edition*. Wiley-VCH Verlag, 2008.

W. E. Kerzendorf and S. A. Sim. A spectral synthesis code for rapid modelling of supernovae. *MNRAS*, 440:387–404, May 2014. doi: 10.1093/mnras/stu055.

L. B. Lucy. Computing radiative equilibria with Monte Carlo techniques. *A&A*, 344:282–288, April 1999.

P. A. Mazzali and L. B. Lucy. The application of Monte Carlo methods to the synthesis of early-time supernovae spectra. *A&A*, 279:447–456, November 1993.

D. Mihalas and B. W. Mihalas. *Foundations of Radiation Hydrodynamics*. New York: Oxford University Press, 1984.

Ulrich M. Noebauer and Stuart A. Sim. Monte Carlo radiative transfer. *Living Reviews in Computational Astrophysics*, 5(1):1, June 2019. doi: 10.1007/s41115-019-0004-9.

L. H. Thomas. The radiation field in a fluid in motion. *The Quarterly Journal of Mathematics*, os-1(1):239–251, 1930. doi: 10.1093/qmath/os-1.1.239. URL http://qjmath.oxfordjournals.org/content/os-1/1/239.short.

R. C. Thomas, P. E. Nugent, and J. C. Meza. SYNAPPS: Data-Driven Analysis for Supernova Spectroscopy. *PASP*, 123:237–248, February 2011. doi: 10.1086/658673.