



Universiteit
Leiden
The Netherlands

Simulating Cosmic Reionisation

Pawlik, A.H.

Citation

Pawlik, A. H. (2009, September 30). *Simulating Cosmic Reionisation*. Retrieved from <https://hdl.handle.net/1887/14025>

Version: Corrected Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/14025>

Note: To cite this publication please use the final published version (if applicable).

Dies ist ein Seidel vorher, aber du kannst auch Spucke dazu sagen, dann müssen vielleicht alle Spucke dazu sagen, getrunken wird aber doch daraus.

Döblin, Berlin Alexanderplatz

CHAPTER 4

TRAPHIC - radiative transfer for smoothed particle hydrodynamics simulations

Andreas H. Pawlik & Joop Schaye

This chapter contains material that has been published together with the material presented in the next chapter, Chapter 5, in MNRAS 389 (2008), 651-677. It provides an updated version of Sections 1-4 in that publication.

WE present TRAPHIC, a novel radiative transfer scheme for Smoothed Particle Hydrodynamics (SPH) simulations. TRAPHIC (TRANsport of PHotons In Cones) is designed for use in simulations exhibiting a wide dynamic range and containing a large number of light sources. It is adaptive both in space and in angle and can be employed for application on distributed memory machines. The (time-dependent) radiative transfer equation is solved by tracing individual photon packets in an explicitly photon-conserving manner directly on the unstructured grid traced out by the set of SPH particles. To accomplish directed transport of radiation despite the irregular spatial distribution of the SPH particles, photons are guided inside cones. The expensive scaling of the computation time with the number of light sources that is encountered in conventional radiative transfer schemes is avoided by introducing a source merging procedure. In this chapter we limit ourself to the description of the radiative transfer scheme. Its implementation for use with state-of-the-art SPH simulations will be described and tested in Chapter 5.

4.1 INTRODUCTION

Radiation is one of the fundamental constituents of our Universe. Its interaction with baryons may lead to an energy exchange that can both heat and cool the matter, initiating pressure forces that may strongly influence the subsequent hydrodynamical evolution. Radiation may also exert a direct pressure force upon the matter through the exchange of momentum. Radiative interactions are furthermore often the dominating process in governing the excitation and ionization state of atoms and molecules. The inclusion of the transport of radiation into hydrodynamical simulations may therefore provide the key for interpreting the outcome of physical experiments and observational campaigns.

To perform hydrodynamical simulations, the Lagrangian technique Smoothed Particle Hydrodynamics (SPH; Gingold & Monaghan 1977; Lucy 1977) is often employed. In SPH, the continuum fluid is discretized using a finite set of point particles, each carrying its own collection of variables. The deformation of the fluid by internal and external processes manifests itself in a steady redistribution of the point particles in space. All it takes to determine the values of physical field variables at a given point in the simulation box is to perform a weighted average over the values these variables take on the particles in its surrounding. This elegant simplicity of the SPH technique is one of the many reasons for its success.

Although the first radiative transfer calculation was already included in SPH at the very birth of this numerical technique more than thirty years ago (Lucy 1977), the detailed treatment of radiation transport in SPH is still an enormous challenge. One of the main reasons for this is certainly the high dimensionality of the problem. In fact, the radiative transfer equation depends on no less than seven variables (three space coordinates, two angles, frequency and time). Moreover, existing numerical schemes to solve the radiative transfer equation have most often only been formulated for use with uniform grids. Despite these difficulties, there has been encouraging progress over the last years, with several interesting and rather different approaches (see Sec. 4.3). In this chapter we present a novel method to solve the radiative transfer equation in SPH. We specifically designed our method for use in hydrodynamical simulations exhibiting a wide dynamic range in physical length scales and containing a large number of light sources.

A prominent example for such simulations are cosmological simulations of large-scale structure formation. Performing cosmological simulations is an exceedingly demanding computational task. Difficulties in describing the growth of the initially tiny matter density perturbations produced before and during the event of recombination and their metamorphosis into the rich structure observable today do not only arise because of our ignorance of how to properly model the governing physical processes. Often, it is simply the lack of computational power which prevents us from faithfully representing the basic actions involved: Cosmological simulations are both time consuming and memory exhaustive. To overcome these computational challenges, one can make use of advanced techniques and resort to clever approximations, reducing the computational effort.

Consider the wide range of scales encountered in cosmological hydrodynamical simulations. According to the hierarchical model of structure formation, the first structures and building blocks of the evolving universe are expected to form at small scales. The non-linear evolution at these scales shapes the distribution of the matter at all times, thus necessitating simulations of high spatial resolution. On the other hand, we also require sufficiently large simulation boxes in order to properly account for the modulation of the small-scale nonlinear evolution by the large-scale structure formation processes (e.g. Barkana & Loeb 2004) and not to be deceived by the cosmic scatter.

To accommodate these two antithetic demands while keeping the number of particles representing the matter low enough to be computationally manageable, *spatially adaptive* SPH simulations have been invoked. It is then immediately clear that when solving the transport of radiation along with the gravito-hydrodynamics of the matter, one requires the radiative transfer scheme to be adaptive, too. Even spatially adaptive cosmological SPH simulations, however, make use of hundreds of millions of particles and are therefore still extremely memory-consuming. It is then indispensable to distribute the computational load over a large number of machines. For this reason, we require a radiative transfer scheme that is *parallel on distributed memory* machines.

When performing radiative transfer simulations, sources of light are assigned a special importance. As a result, the computation time of most of the available radiative transfer schemes scales linearly with the number of sources in the simulation box. However, in cosmological simulations, even at times as early as 1 billion years after the Big Bang, i.e. at a redshift of $z \sim 6$, a non-negligible amount ($\gtrsim 1$ per cent) of baryonic matter has undergone star formation. In addition to these stellar sources of light, the intergalactic gas emits photons too, producing a radiation component often referred to as diffuse radiation. Hence, without *breaking the linear scaling* with the number of light sources, radiative transfer simulations at the resolution of state-of-the-art cosmological hydrodynamic simulations need to be dispatched to the realm of the future.

In this paper we present a radiative transfer scheme for use in SPH simulations that is adaptive, parallel on distributed memory and that avoids the linear scaling of the computation time with the number of sources. Hence, it is ideally suited for application in large simulations covering a wide range of length scales and containing many sources. In our scheme we follow the propagation of individual photon packets. It is *explicitly photon-conserving* (Abel, Norman, & Madau 1999) and can be applied to solve the *time-dependent* radiative transfer equation. Because the photon packets are traced in cones, we refer to our scheme as TRAPHIC - TRANsport of PHotons In Cones. The introduction of cones is required in order to perform the transport of photon packets directly on the unstructured grid defined by the SPH particles.

Although we have designed our radiative transfer scheme to be readily coupled to the hydrodynamic evolution of the matter in SPH simulations, here we limit ourselves to the description of the radiative transfer scheme itself. We will present and tests its implementation in a state-of-the-art SPH code in Chapter 5. We will report on fully coupled radiation-hydrodynamical SPH simulations employing our implementation of TRAPHIC in future work.

The outline for the rest of this chapter is as follows. We start with a brief review of the principles of SPH that we consider crucial for the understanding of the present work. In Sec. 4.3 we then recall the radiative transfer equation and place our radiative transfer method in context by reviewing the approaches that have been used to solve the radiative transfer problem in SPH so far. In Section 4.4 we give a detailed description of the ideas behind our method. This is done in two steps: a brief overview over the method is followed by a more detailed discussion. Throughout we will emphasise how we satisfy the requirements set by our primary aim of performing radiative transfer in simulations exhibiting a wide range in length scales and a large number of light sources. Finally, in Sec. 4.5 we summarise our approach and conclude with an outlook.

4.2 SMOOTHED PARTICLE HYDRODYNAMICS

Excellent reviews of SPH exist (see e.g. Monaghan 2005; Monaghan 1992; Price 2005). Here, we just briefly outline the basic principles we consider critical for the understanding of our radiative transfer method.

At the basis of SPH lies the representation of a field $A(\mathbf{r})$ by its integral interpolant $A_I(\mathbf{r})$,

$$A_I(\mathbf{r}) = \int d^3r' A(\mathbf{r}')W(\mathbf{r} - \mathbf{r}', h), \quad (4.1)$$

where the smoothing length h determines the spatial resolution. The interpolation kernel W satisfies

$$\int d^3r' W(\mathbf{r} - \mathbf{r}', h) = 1 \quad (4.2)$$

$$\lim_{h \rightarrow 0} W(\mathbf{r} - \mathbf{r}', h) = \delta_D(\mathbf{r} - \mathbf{r}'), \quad (4.3)$$

where δ_D is the Dirac delta function, such that $A_I(\mathbf{r})$ coincides with $A(\mathbf{r})$ in the limit $h \rightarrow 0$. The last two conditions do not fix the functional form for W . An often adopted choice is the spherically symmetric compact spline

$$W(\mathbf{r} - \mathbf{r}', h) = \frac{8}{\pi h^3} \begin{cases} 1 - 6 \left(\frac{r}{h}\right)^2 + 6 \left(\frac{r}{h}\right)^3, & 0 \leq \frac{r}{h} \leq \frac{1}{2} \\ 2 \left(1 - \frac{r}{h}\right)^3, & \frac{1}{2} < \frac{r}{h} \leq 1 \\ 0, & \frac{r}{h} > 1, \end{cases} \quad (4.4)$$

where $r = |\mathbf{r} - \mathbf{r}'|$. From here on, when referring to the interpolation kernel W , we assume that it is of this form. The discrete representation of a fluid by SPH particles is achieved by approximating the integral interpolant (Eq. 4.1) by the summation interpolant,

$$A_I(\mathbf{r}) \approx A_S(\mathbf{r}) = \sum_j \frac{m_j}{\rho_j} A(\mathbf{r}_j)W(\mathbf{r} - \mathbf{r}_j, h), \quad (4.5)$$

where the summation is over the particles of mass m and mass density ρ . For a self-contained numerical treatment, any field needs to be discretized only at the positions of the particles i representing the fluid,

$$A_i \equiv A_S(\mathbf{r}_i) = \sum_j \frac{m_j}{\rho_j} A(\mathbf{r}_j)W(\mathbf{r}_i - \mathbf{r}_j, h). \quad (4.6)$$

Since the kernel W is compact, only local information needs to be accessed for the evaluation of the last sum, which can be carried out in a computationally efficient manner in parallel on distributed memory machines.

For the interpretation of Eq. 4.6, two main approaches can be taken, depending on the choice for h (Hernquist & Katz 1989). In the scattering approach each particle j is considered as a cloud of radius h_j and contributes to the field at the position of particle i with the weight $W(\mathbf{r}_i - \mathbf{r}_j, h_j)$. In the gathering approach, each particle i searches the sphere of radius h_i (in the following referred to as the sphere of influence, or *neighbourhood*) for particles (its *neighbours*) and obtains an estimate of the field value at its position by summing the field values at the positions of the neighbours j , each weighted by $W(\mathbf{r}_i - \mathbf{r}_j, h_i)$.

The two interpretations are identical only if the spatial resolution is fixed throughout the fluid, i.e. if $h_i = h_j = h$. However, the SPH formalism only unfolds its true strength by allowing the resolution to vary in space, according to the local density field. This is usually achieved by either directly fixing the number of neighbours N_{ngb} for all particles, or by requiring that the kernel volume contains a constant mass (Springel & Hernquist 2002).

4.3 RADIATIVE TRANSFER IN SPH - PREVIOUS WORK

Before describing TRAPHIC, our method to solve the radiative transfer equation in SPH, we briefly recall the radiative transfer equation and give a short overview of the main numerical methods that have been employed so far to obtain its solution in SPH simulations.

The classical equation of radiative transfer reads (see e.g. Mihalas & Weibel Mihalas 1984)

$$\frac{1}{c} \frac{\partial I_\nu}{\partial t} + \mathbf{n} \cdot \nabla I_\nu = \epsilon_\nu - \kappa_\nu \rho I_\nu. \quad (4.7)$$

In Eq. 4.7, $I_\nu \equiv I_\nu(\mathbf{r}, \mathbf{n}, t)$ is the monochromatic intensity (units $\text{ergs cm}^{-2} \text{s}^{-1} \text{Hz}^{-1} \text{sr}^{-1}$) of frequency ν at position \mathbf{r} , \mathbf{n} is a unit vector along the direction of light propagation and c is the speed of light. Sources and sinks of radiation are described by the emissivity ϵ_ν (units $\text{ergs cm}^{-3} \text{s}^{-1} \text{Hz}^{-1} \text{sr}^{-1}$) and the mass absorption coefficient κ_ν (units $\text{cm}^2 \text{g}^{-1}$), respectively. In general, both ϵ_ν and κ_ν are functions of \mathbf{r} , \mathbf{n} and time t .

If we define the photon number density ψ_ν such that $\psi_\nu d\Omega d\nu$ is the number of photons per unit volume with frequencies $(\nu, \nu + d\nu)$ travelling with velocity c into a solid angle $d\Omega$ around \mathbf{n} , the intensity is given by $I_\nu = ch_p \nu \psi_\nu$, where h_p is the Planck constant. Eq. 4.7 is a partial differential equation depending on seven variables, which in general is hard to solve. Different approximations have been employed to obtain its solution, giving rise to different numerical approaches. Below we discuss the main approaches that have been taken so far to accomplish the transport of radiation in SPH simulations.

In his study of optically thick proto-stars, Lucy (1977) modelled the transport of radiation as a diffusion process by including a corresponding term in the SPH formulation of the energy equation. Brookshaw (1994) pointed out drawbacks of the particular formulation used in Lucy (1977), and re-phrased the diffusion equation in SPH such as to reduce its sensitivity to particle disorder. This new formulation of the diffusion equation was employed by Viau, Bastien, & Cha (2006) to perform collapse simulations of centrally condensed clouds.

Treating the transport of radiation in the diffusion limit is, however, only a good approximation in optically thick media. In the optically thin regime, infinite signal propagation speeds result, since the diffusion equation is a partial differential equation of parabolic type. This defect can be remedied by supplementing the diffusion equation with a so-called flux-limiter. SPH simulations with flux-limited diffusion have been carried out by e.g. Herant et al. (1994), Whitehouse & Bate (2006), Fryer, Rockefeller, & Warren (2006) and Mayer et al. (2007), covering a wide range of physical applications, from supernovae explosions and proto-stellar collapse to the study of proto-planetary disks. The (flux-limited) diffusion approach to solve the radiative transfer equation fails in complex geometries. In particular, opaque obstacles illuminated by a point source do not cast sharp shadows, because the shadow region is filled in by the diffusion.

To solve the radiative transfer equation without being restricted to the optically thick regime or simple geometries, Monte Carlo techniques can be employed (Oxley & Woolfson 2003, Stamatellos & Whitworth 2005, Croft & Altay 2007, Semelin, Combes, & Baek 2007, Altay, Croft, & Pelupessy 2008). In Monte Carlo radiative transfer, individual photon packets emitted by each source are directly followed as they pass through the matter, thus simulating the physical process of radiation transport as encountered in nature. While Monte Carlo simulations allow realistic shadows to be cast behind opaque obstacles, they are computationally very demanding, since the Poisson noise inherent to the statistical description of the radiation field leads to a signal-to-noise ratio that grows only with the square-root of the number of photon packets emitted.

Ray-tracing schemes keep the advantages of Monte Carlo radiative transfer simulations, whilst not being affected by the statistical noise. In short, in ray-tracing schemes rays are cast from each source throughout the simulation box. The optical depth to each point in space is calculated along these rays, and the attenuation of the flux emitted by the sources is obtained. Variants of ray-tracing working with photons instead of fluxes blur the differences with Monte-Carlo schemes.

Ray-tracing schemes are most easily implemented on a regular grid superimposed on the SPH density field. However, this implies that all information about substructure on scales smaller than the size of the grid cells is ignored. As a cure, adaptive grids have been invoked (e.g. Razoumov & Sommer-Larsen 2006, Alvarez, Bromm, & Shapiro 2006). In SPH, *grid-less* ray-tracing has been introduced by Kessel-Deynet & Burkert (2000) to investigate the effects of ionizing UV radiation by massive stars on the surrounding interstellar medium. The Strömgren volume method (Dale, Ercolano, & Clarke 2007, compare also Johnson, Greif, & Bromm 2007) and the ionization front tracking technique employed by Yoshida et al. (2007) are related approaches. Another grid-less radiation tracing scheme that has been applied to SPH simulations is presented in Ritzerveld & Icke (2006). Inspired by the Kessel-Deynet & Burkert (2000) approach, Susa (2006) describes a radiation-hydrodynamics scheme for the transport of ionizing radiation in cosmological simulations of structure formation.

One major drawback most ray-tracing schemes share with the Monte-Carlo technique is that the computational effort to solve the radiative transfer equation scales linearly with the number of sources in the simulation box, because rays are traced for each source in turn (but see Sec.4.4.2 later on). For this reason, SPH radiation-hydrodynamical simulations employing the ray-tracing approach have mostly been restricted to the study of problems involving only a few sources.

Recently, so-called moment methods have also been implemented for use with SPH simulations (Petkova & Springel 2008 for an adaptive, grid-less implementation; Finlator, Özel, & Davé 2009 for an implementation on a super-imposed grid). These methods work by forming an infinite hierarchy of moments of the radiative transfer equation, which is then solved. The moment approach allows for a straight-forward merging of sources, which renders its computation time independent of the source number. Its numerical solution requires, however, a closure relation that cuts the moment hierarchy to obtain a finite number of moments. Practically feasible choices for this closure relation shift the hyperbolic character of the radiative transfer equation to become diffusion-like. Moment methods therefore typically do not produce shadows behind opaque obstacles.

In the following sections we will describe TRAPHIC, a novel method to solve the radiative transfer equation in SPH. TRAPHIC employs a photon tracing technique that works directly on the discrete set of irregularly distributed SPH particles. It is designed to overcome the challenges set by our goal of carrying out large radiation-hydrodynamics SPH simulations exhibiting a wide dynamic range in length scales and containing a large number of light sources.

4.4 TRAPHIC - TRANSPORT OF PHOTONS IN CONES

In this section we give a detailed description of TRAPHIC, our method to solve the radiative transfer equation in SPH. We start the presentation of our radiative transfer scheme by sketching the main idea, in order to give an overview and to introduce the reader to the underlying concepts, which will be illustrated and explained in more detail in the subsequent sections. TRAPHIC is designed to solve the radiative transfer equation in three-dimensional SPH simu-

lations. When illustrating concepts with schematic figures we will, however, restrict ourselves to two dimensions for the sake of clarity.

Although we ultimately aim at performing simulations in which the radiation transport is fully coupled to the hydrodynamics, here we assume that the SPH density field is static and that the radiation does not exert any thermodynamical or hydrodynamical feedback on the matter. We will report on the coupling of radiation and hydrodynamics in future work.

4.4.1 Overview

We solve the radiative transfer equation (Eq. 4.7) in finite steps $(t_r, t_r + \Delta t_r)$, where t_r is the current simulation time, by tracing photon packets radially away from their location of emission until a certain stopping criterion is fulfilled. We do not introduce a grid on which the photons are propagated. Instead, we propagate the photons directly on the discrete set of particles in the simulation. For the transport of photons we employ the same particle-to-neighbour communication scheme that is already used in the SPH simulation to solve the equations of hydrodynamics. That is, we accomplish the global transport of photons by propagating them only locally, between a particle and its \tilde{N}_{ngb} neighbours. We allow \tilde{N}_{ngb} to be different from N_{ngb} , the number of neighbours used during the SPH calculations. In SPH, N_{ngb} determines the adaptive spatial resolution of the hydrodynamical simulation. Similarly, in our scheme, \tilde{N}_{ngb} sets the adaptive spatial resolution of the radiation transport. It is the first of two numerical parameters in our method (see also the left-hand panel of Fig. 4.1).

Working directly on the set of SPH particles, our scheme exploits the full spatial resolution of the hydrodynamical simulation. A further advantage of our approach is that the radiation transport is automatically parallel on distributed memory, as long as the SPH simulation itself is parallel on distributed memory. This is in contrast to radiation-hydrodynamical schemes that employ individual communication schemes for the transport of radiation and the SPH computations. Our radiative transfer scheme can hence be coupled to the SPH simulation without introducing any additional computational structures.

The main challenge we face when performing radiation transport by propagating photons only from a particle to its neighbours is achieving a transport that is directed: Photons travel along straight rays, while the spatial distribution of the SPH particles is generally highly irregular. In the following we give an overview of the concepts employed in our scheme and describe how we overcome this and other challenges to accomplish the transport of radiation in SPH simulations.

We distinguish two types of particles present in SPH simulations that are relevant for the transport of photons: star particles and SPH (gas) particles. Only gas particles can interact with photons, via absorptions and scatterings. We assume, however, that both star and gas particles can be sources of photons and will therefore refer to them as source particles, or simply sources.

The transport of radiation at simulation time t_r over a single radiative transfer time step Δt_r starts with the emission of photons by the source particles. Subsequently, these photons are propagated downstream, radially away from the source, simultaneously for all sources. In our particle-to-neighbour transport scheme photons are propagated from gas particle to gas particle. On their way, photons may interact with the matter field discretized by the gas particles. Each gas particle can simultaneously receive photons, possibly emitted at different times, from multiple (in principle all) sources in the simulation. Their further propagation would require a loop over all propagation directions. We explicitly avoid the resulting linear scaling of the computation effort with the total number of sources by introducing a source

merging procedure.

The distribution of photons amongst the neighbours of the source particles must respect the angular dependence of the source emissivity. We achieve this by introducing for each source a set of randomly oriented, tessellating *emission cones*¹(Sec.4.4.2), as illustrated in the middle panel of Fig. 4.1. The fraction of the total number of photons emitted into each cone is proportional to its solid angle. The emission direction (which we will equivalently refer to as the propagation direction) associated with each emitted photon packet is determined by the central axis of the emission cone. The number of cones used in the emission cone tessellation, N_c , is the second numerical parameter in our method. It sets the angular resolution. The random orientation given to the cones ensures that photons are not restricted to propagate along a fixed number of directions and prevents artefacts introduced by the shape of the emission cones.

Some of the emission cones may not contain any neighbours, as is the case for the bottom-right cone of the tessellation shown in the middle panel of Fig. 4.1. To nevertheless emit photons into the corresponding directions, we create additional neighbours to which the photons can then be transferred (Fig. 4.1, right-hand panel). We refer to these neighbours as virtual particles (ViPs), since they do not take part in the SPH simulation. The properties of the ViPs are obtained from those of the neighbouring SPH particles using SPH interpolation. ViPs compensate for the lack of neighbours in certain solid angles around the emitting source. Hence, by employing ViPs we introduce the freedom of choosing emission directions independently of the geometry of the SPH particle distribution. As we will argue in Section 4.4.2 and investigate in more detail in Appendix 4.A, this freedom is a necessary requirement for achieving a desired angular dependence of the transport process, e.g. the isotropic emission of photons by star particles, in any particle-to-neighbour transport scheme.

The photon packets distributed amongst the neighbours of the sources are further transferred downstream until they experience an interaction. In contrast to the emission of photons by source particles, gas particles distribute photons only amongst the subset of neighbours located in the downstream direction (see Fig. 4.2). We distinguish this directed particle-to-neighbour transport from the emission process by referring to it as *transmission* (Sec.4.4.2). In more detail, each gas particle transmits photons downstream by distributing photon packets amongst the subset of neighbours located in certain regular cones only. These so-called transmission cones are centred on the emission (propagation) direction associated with each photon packet and have an opening angle determined by the angular resolution, i.e. by N_c . They confine the photon packets to the cone into which they were emitted by the corresponding source particle. As a result, the photon packets are propagated radially away from the sources, in a manner that is independent of the geometry of the SPH particle distribution. Like emission cones, transmission cones might not contain any neighbours. Again we employ ViPs to accomplish the photon transport into the corresponding directions.

By keeping the opening angle of the cones fixed, the solid angle subtended by a transmission cone as viewed from the original source decreases with the distance from the source. As a result, the photon transport becomes adaptive in angle, as in the ray-splitting technique employed in ray-tracing schemes (e.g. Abel & Wandelt 2002). We will demonstrate later on, when testing a numerical implementation of our scheme (see Chapter 5), that the use of implicitly adaptive transmission cones confining the photon propagation yields sharp shadows behind opaque obstacles.

Gas particles receive photon packets from other particles through the processes of emission

¹We use the word cone in a general sense. It does not necessarily describe, but includes in its meaning, a regular cone, which is defined as a pyramid with a circular cross section.

and transmission described above. A given gas particle can simultaneously receive multiple photon packets, which may each be associated with a different emission direction. In fact, our assumption that all star and gas particles in the simulation box can be source particles would imply that the number of directions from which photon packets can be simultaneously received by a given gas particle would scale linearly with the total number of particles in the simulation. The computational effort to accomplish the subsequent photon transmission along the associated directions would then also scale linearly with the total number of particles. Since this consideration applies to the transmission of photons by any SPH particle in the simulation, the total computational effort to solve the radiative transfer equation would scale no less than quadratically with the total number of SPH particles.

To avoid this computationally expensive scaling, we introduce a source *merging* procedure (cp. Fig. 4.4, which will be explained in more detail in Sec.4.4.2). We make use of the fact that source particles that are seen as close (in angle) to each other from a given gas particle can be approximated by, or merged into, a single point source. We average the emission directions associated with the photon packets received from sources in solid angle bins defined by a so-called reception cone tessellation which, except for a random rotation, is identical to the cone tessellation employed for the emission process. Consequently, photon packets need to be transmitted into at most N_c directions.

We distinguish two types of interactions with the matter field sampled by gas particles and ViPs that photons may experience. Absorption interactions change the number of photons contained in each packet. Scattering interactions change the propagation direction (and possibly the frequency) of the photons in a packet. All interactions are taken into account in an explicitly photon-conserving manner. We give a detailed description of interactions in Sec.4.4.3.

4.4.2 Transport of photons

We now expand on the description of our radiative transfer scheme given above. We comment in detail on the main concepts underlying TRAPHIC, i.e. the emission of photons by source particles, the transmission of photons by gas particles and the source merging procedure. The description of how these concepts are employed to advance the solution of the radiative transfer equation in time is deferred to Sec.4.4.4.

Photon emission by source particles

In this section we describe the emission of photons by source particles. Star particles emit photons according to their intrinsic luminosity, which can for instance be obtained from population synthesis models. An example of the emission of photons by gas particles is the emission of recombination radiation by a recombining ion.

Let us consider the emission of photons by source particle i , located at \mathbf{r}_i . We denote the number of photons emitted per unit time per unit frequency ν per unit solid angle Ω around the unit direction vector \mathbf{n} by $\dot{\mathcal{N}}_{\nu,i}(\mathbf{n}, \mathbf{r}, t)$, or simply $\dot{\mathcal{N}}_{\nu,\mathbf{n},i}$. With this notation, the number of photons $\dot{\mathcal{N}}_{\nu,i}$ emitted per unit time at frequency ν is $\dot{\mathcal{N}}_{\nu,i} \equiv \int d\Omega \dot{\mathcal{N}}_{\nu,\mathbf{n},i}$ and the total number of photons emitted per unit time is $\dot{\mathcal{N}}_{\gamma,i} \equiv \int d\nu \int d\Omega \dot{\mathcal{N}}_{\nu,\mathbf{n},i}$.

Source particle i emits $\int_t^{t+\Delta t_r} \dot{\mathcal{N}}_{\gamma,i}$ photons over the radiative transfer time step Δt_r . In agreement with our particle-to-neighbour based transport approach, the emission process is modelled by distributing these photons amongst the N_{ngb} nearest gas particles, residing in a sphere of radius \tilde{h}_i centred on \mathbf{r}_i . This sphere is schematically depicted in the left-hand panel of Fig. 4.1.

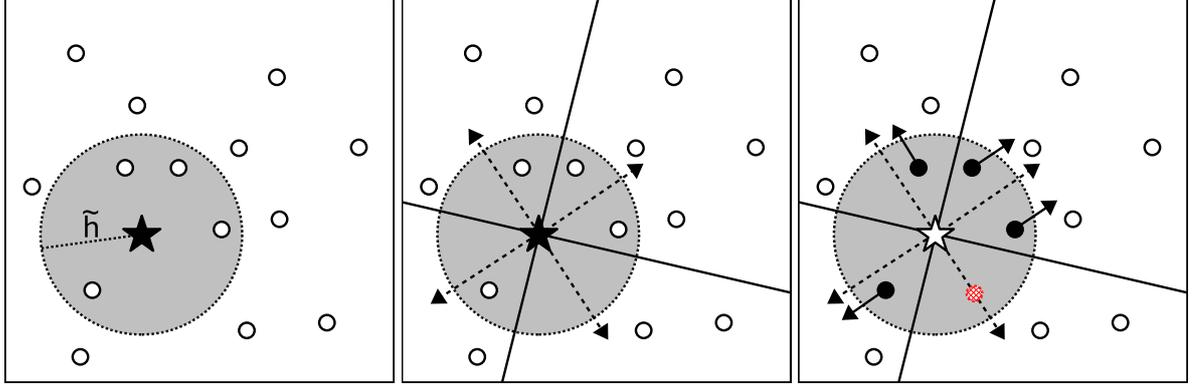


Figure 4.1: Emission of photon packets by source particles. *Left-hand panel:* A source particle (black star) and its neighbourhood (big grey disc). In this example the radius \tilde{h} of the neighbourhood is defined such that there are $\tilde{N}_{\text{ngb}} = 4$ neighbouring gas particles (white discs). *Middle panel:* The randomly oriented emission cone tessellation of the source particle (solid lines) is shown for the case $N_c = 4$. The dashed arrows indicate the central axes of the cones. Note that the bottom-right cone does not contain a neighbouring gas particle. *Right-hand panel:* The source particle has transferred its photon packets to its neighbours (black discs). The emission direction (small solid arrows) associated with each packet points in the direction of the axis of the cone in which the neighbour resides. For the empty cone, a virtual particle is created (red hatched disc), placed randomly along the central axis of the cone.

Source particle i distributes its photons amongst its neighbours with the help of a set of space-filling emission cones, defined as follows (middle panel of Fig. 4.1). We tessellate the simulation box into N_c cones of (generally not identical) solid angles Ω_i^k , $k = 1, 2, \dots, N_c$, with the apexes located at the position \mathbf{r}_i of particle i . The number of neighbours $\tilde{N}_{\text{ngb},i}^k$ in each cone k is determined, taking values in the range $0 \leq \tilde{N}_{\text{ngb},i}^k \leq \tilde{N}_{\text{ngb}}$.

For what follows we consider the emission of photons for each frequency ν separately. To each neighbour j in cone k a photon packet of characteristic frequency ν is emitted. The packet contains a fraction $w_i^{kj} > 0$ ($j = 1 \dots \tilde{N}_{\text{ngb},i}^k$) of the total number of photons $\dot{N}_{\nu,i}$ to be emitted per unit time at frequency ν , with

$$w_i^{kj} = \frac{\int_{\Omega_i^k} d\Omega \dot{N}_{\nu,\mathbf{n},i}}{\int d\Omega \dot{N}_{\nu,\mathbf{n},i}} \times \frac{w_i^j}{\sum_{l=1}^{\tilde{N}_{\text{ngb},i}^k} w_i^l}, \quad (4.8)$$

where w_i^j are weights attached to neighbour j in cone k . The first factor on the right-hand side of Eq. 4.8 determines which fraction of the total number of photons is emitted into cone k , whereas the second factor controls the fraction of photons that is transferred to neighbour j , that is, it controls the distribution of photons amongst the particles within cone k . Here we set $w_i^j = 1$, i.e. equal weights for all neighbours in a given cone. The number of cones used in the tessellation, the parameter N_c , determines the angular resolution of the radiation transport. A cone tessellation may consist of cones with different solid angles. We therefore define the angular resolution to be the size of the average solid angle, $\langle \Omega \rangle = 4\pi/N_c$.

For each emission cone k the central axis is defined, characterised by the unit vector \mathbf{n}_k pointing away from the source (Fig. 4.1, middle panel). We refer to this vector as the emission direction associated to photon packets emitted into cone k . When a photon packet is emitted to a neighbouring gas particle, the emission direction is transferred in addition to the number

of photons it contains (Fig. 4.1, right-hand panel). Since the orientation of the emission cone tessellation is randomised by applying a random rotation at each emission, there is no a priori limit on the values the emission directions can take. Note that while we transfer the emission direction, we do not transfer the position of the source. Photon packets are traced further downstream based on their emission direction only, as we will explain in Sec.4.4.2. Each photon packet has also an associated *clock* t^* . At emission the clock time is set to the time of the simulation, $t^* = t_r$. In Sec.4.4.4 we will see that the clock can be employed to propagate the photon packets at the speed of light.

Some cones may not contain any neighbouring gas particles. This is for instance the case for the bottom right cone in the middle panel of Fig. 4.1. For a fixed number of neighbours these empty cones will occur more frequently if the angular resolution is higher (i.e. if the solid angles of the cones are smaller). On the other hand, for a fixed angular resolution N_c , empty cones will occur more frequently if the number of neighbours \tilde{N}_{ngb} is smaller. Spatial clustering of the neighbours also increases the probability of the occurrence of empty emission cones. In the absence of a neighbouring gas particle photons cannot be transported along the emission direction of the empty cone. We therefore create a new neighbour, a so-called virtual particle (ViP), to which the photons are transferred. The ViP is placed along the cone axis at a random distance $< \tilde{h}_i$ from the source particle, such that the volume of the cone is uniformly sampled (see the right-hand panel of Fig. 4.1). The properties of ViPs (e.g. density) are determined from their neighbouring gas particles using SPH interpolation. ViPs are only employed for the transport of photons. We stress that ViPs are not used by the SPH simulation.

We emphasise that the introduction of cones is essential for accomplishing the emission process in our particle-to-neighbour transport scheme. To see this, consider the alternative of distributing the photons directly amongst the neighbours of gas particle i . This amounts to setting $\Omega_i^1 = 4\pi$ and $\Omega_i^j = 0$ for $j > 1$ in Eq. 4.8. Without any reference system, the emission directions associated to the photon packets would then be given by the unit vectors pointing from source i to neighbour j . In Appendix 4.A we show that in this case the net emission direction in general would correlate with the direction towards the centre of mass of the neighbours. As a result, the emission process would depend on the clustering of the neighbours in space as set by the geometry of the SPH simulation. The use of emission cones combined with ViPs gives us the freedom to choose emission directions independently of the spatial clustering of the neighbours. This allows us to model any angular dependence of the source emissivity, within the bounds of the chosen angular resolution.

In summary, source particles transfer photon packets to their neighbouring gas particles using a set of randomly oriented, tessellating cones. Each cone defines a direction of transport, i.e. the emission direction, associated with the packets. The random orientation applied to the cone tessellation ensures that photon packets are not transferred along a fixed set of directions only. Virtual particles (ViPs) are placed in emission cones not containing any neighbours, to which the photon packets are then transferred to. The combination of emission cones and ViPs makes the radiation transport independent of the spatial distribution of the neighbours.

Photon transmission by gas particles and ViPs

In the last section we described how a source particle distributes photon packets amongst the gas particles in its neighbourhood. For each packet we defined an emission direction, independently of the spatial distribution of the neighbouring gas particles by employing a randomly oriented set of emission cones. In this section we describe how the packets are propagated through the simulation box along their emission directions, by employing directed particle-to-

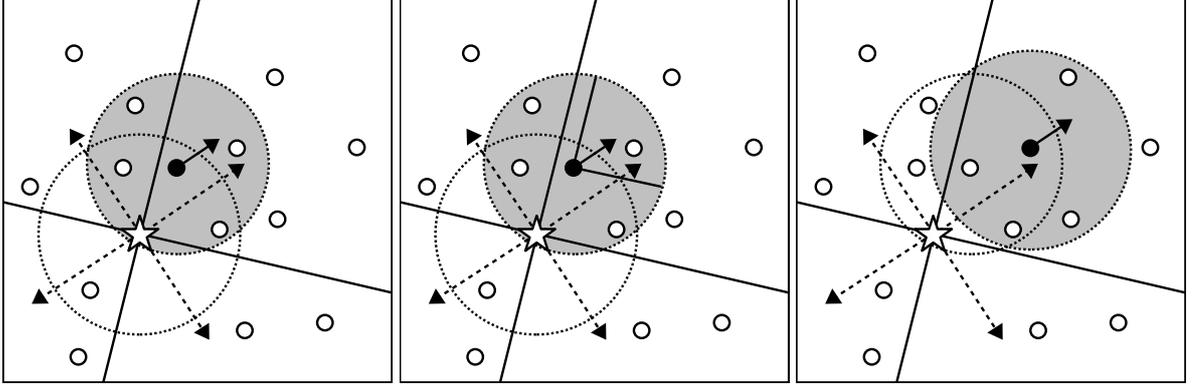


Figure 4.2: Transmission of photon packets by gas and virtual particles. The particle positions shown are the same as in Fig. 4.1, as are the neighbourhood (dashed circle) and emission cone tessellation of the source particle (white star). The transmission of photons is illustrated for one of the neighbours that received radiation from the star in Fig. 4.1. *Left-hand panel:* The neighbourhood (big grey disc) of the transmitting gas particle (black disc) is defined. As in Fig. 4.1, $\tilde{N}_{\text{ngb}} = 4$. The emission direction associated with the photon packet to be transmitted is shown as the short solid arrow. *Middle panel:* The transmission cone is shown, centred on the emission direction of the photon packet that is to be transmitted. In this cone one downstream neighbour is found. *Right-hand panel:* The photon packet is transmitted to the downstream neighbour, turning it into a transmitting particle. The cycle repeats with defining the neighbourhood for this particle. As a result, the photon packet is propagated downstream, radially away from the source particle.

neighbour transport, a process which we refer to as transmission.

Consider a gas particle i , which receives a photon packet. Recall that the neighbours of particle i are the \tilde{N}_{ngb} nearest gas particles located in the sphere with radius \tilde{h}_i , centred on particle i (Fig. 4.2, left-hand panel). Analogous to the case of emission, particle i re-distributes the photons contained in the packet amongst its neighbours. In contrast to emission, photons are only distributed amongst the subset of neighbours $\tilde{N}_{\text{t,ngb}} \leq \tilde{N}_{\text{ngb}}$ located *downstream*. These are the neighbours residing in a *regular*² cone centred on the direction of propagation of the packet (see the middle panel of Fig. 4.2). We refer to this cone as transmission cone.

The apex of the transmission cone is located at the position of gas particle i . The solid angle of the cone is set by the angular resolution³, $\Omega_{\text{t}} = 4\pi/N_c \equiv \langle \Omega \rangle$, and determines the apex angle ω through the standard relation

$$\omega = 2 \arccos \left(1 - \frac{\Omega_{\text{t}}}{2\pi} \right) \times \frac{180}{\pi}. \quad (4.9)$$

We show this relation in Fig. 4.3. By definition, a neighbour j with position \mathbf{r}_j is interior to the transmission cone with apex at the position \mathbf{r}_i of the transmitting particle i if the inner angle between the transmission cone axis and the vector $\mathbf{r}_j - \mathbf{r}_i$ is less than $\omega/2$.

The received photon packet is split into $\tilde{N}_{\text{t,ngb}}$ packets, each of which is transferred to one of the downstream neighbours j of particle i . The number of photons contained in each packet is set by the weights w_i^j , such that each neighbour j receives a photon fraction $w_i^j / \sum_k w_i^k$ of

²A regular cone is a pyramid with a circular cross-section.

³In principle it is possible to choose the solid angle of the transmission cones independently of the angular resolution $4\pi/N_c$ implied by the reception cone tessellation. However, here we do not explore this option.

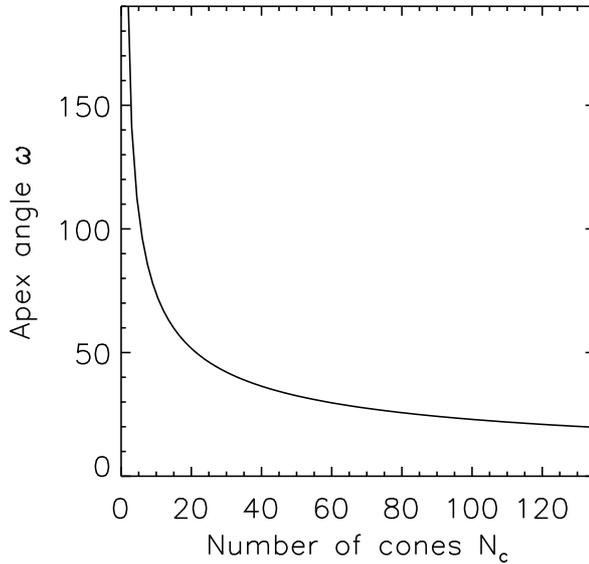


Figure 4.3: The apex angle ω of the transmission cones as a function of the angular resolution N_c (see Eq. 4.9).

the parent photon packet, where the sum is over all downstream neighbours $k = 1, \dots, \tilde{N}_{t,\text{ngb}}$. Here we assume $w_i^j = 1$, giving equal weight to each neighbour. If there are no downstream neighbours, i.e. if the transmission cone is empty, we simply create a neighbour as in the case of empty emission cones described in the previous section. That is, we place a virtual particle (ViP) along the emission direction of the photon packet, a random (in volume) radial distance $< \tilde{h}_i$ away from the transmitting gas particle. The properties of the ViP are determined by SPH interpolation from its neighbours, and the photon packet is transferred.

Each photon packet inherits the emission direction from its parent packet (Fig. 4.2, right-hand panel). After all packets have been transferred to the downstream neighbours, the transmission is finished. Subsequently, the transmission procedure can be applied again. The transmission process thus generally splits each photon packet into multiple packets, propagating with the same emission direction. With each subsequent transmission individual photon packets are confined to a smaller fraction of the solid angle traced out by the emission cone they were emitted into (see Fig. 4.2, middle panel). This is similar to the technique of ray splitting employed in ray-tracing codes. Hence, in TRAPHIC the photon transport takes place adaptively in angle.

The transmission procedure specified above for gas particles is also applied for ViPs. Again, the transmission cone of the ViP can either contain gas neighbours or be empty. If it is empty, the ViP creates another ViP, as in the case of transmission by gas particles. After the ViP has performed the transmission, it is deleted. ViPs are therefore temporary constructs. For $N_c \gg \tilde{N}_{\text{ngb}}$ the total number of ViPs in the simulation is proportional to N_c , whereas for $\tilde{N}_{\text{ngb}} \gg N_c$ the total number of ViPs approaches zero.

The main purpose of the transmission cones is to confine the downstream propagation of photons to the solid angles into which they were emitted by the source particles, which is the main challenge for schemes using particle-to-neighbour transport. Just as emission cones were introduced to make the emission of photons by source particles independent of the geometry of the SPH particle distribution, transmission cones are introduced to further propagate the photons downstream, independently of the geometry of the SPH particle distribution. As a

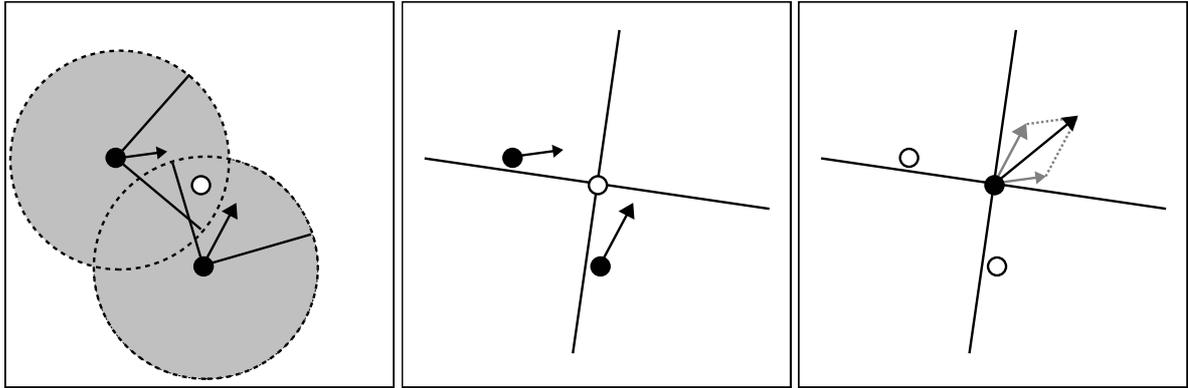


Figure 4.4: Merging of photon packets by gas particles. *Left-hand panel:* A gas particle (white disc) is receiving photon packets simultaneously from two transmitting gas particles (black discs). The emission direction associated with each photon packet is indicated by an arrow, whose length is given by the number of photons contained in the packet. The big grey discs indicate the neighbourhoods of the transmitting gas particles, the solid lines are the transmission cones (we assume $N_c = 4$). For clarity, we do not show particles and photon packets that are not directly involved in the merging event. *Middle panel:* A randomly oriented reception cone tessellation is defined for the receiving gas particle. We omitted the transmission cones and the neighbourhoods of the transmitting particles. *Right-hand panel:* The photon packets have been transferred to the receiving gas particle, turning it into a transmitter. Their emission vectors (grey arrows), after translation to the position of the receiver, are both found to fall into the upper-right reception cone. Note that the particles itself fall into different reception cones. Their positions are irrelevant for the merging, since transmitted photon packets rather than transmitting particles are merged. The emission vectors are merged through a weighted average as described in Section 4.4.2, resulting in a single photon packet (black arrow).

consequence, shadows will be thrown behind opaque obstacles, and their sharpness will be in agreement with the chosen angular resolution. In fact, as we will see in Chapter 5, the shadows are much sharper than implied by the formal angular resolution.

The cones making up the emission tessellation will generally be irregular - in three dimensions there exists no tessellation of space with regular cones. The use of regular cones for the downstream propagation of photon packets emitted into a certain irregular emission cone is therefore an approximation. In principle, each photon packet could be transmitted into a cone having the shape of the emission cone it originates from. However, tracing photons within irregular cones is computationally more demanding. Furthermore, in the next section we will see the necessity for combining multiple photon packets propagating along different directions into a single photon packet that propagates along a correspondingly averaged direction. Since there is no unique way of defining the shape of a transmission cone associated with this average direction, we use a regular shape.

Photon merging by gas particles

In the previous section we described the transmission of a *single* photon packet arriving at a gas particle. In principle, each gas⁴ particle can simultaneously receive *multiple* photon packets, possibly emitted by different sources at different times, propagating along different emission directions, a situation which is depicted schematically in the left-hand panel of Fig. 4.4. Ac-

⁴We do not consider ViPs here, since by construction each ViP can receive only a single photon packet.

completing the transmission would then require executing a loop over all received packets for each transmitting gas particle. Clearly, for each gas particle this would imply a linear scaling of the computation effort with the number of packets that need to be transmitted, or in other words, with the total number of source particles present in the simulation.

In cosmological simulations of structure formation, for instance, where a significant number of star particles can usually be found already at high redshift, the simultaneous reception of multiple photon packets will be the rule rather than the exception. Even without taking into account the diffuse radiation of the intergalactic gas resulting from radiative de-excitations, recombinations and scatterings, a linear scaling with the number of sources would quickly turn the radiation transport into a computational task that would be too expensive to be carried out even with the most advanced supercomputer available today.

A few radiation tracing schemes have attempted to tackle this scaling problem by replacing sources that are close to each other with a single point source (e.g. Cen 2002; Razoumov & Sommer-Larsen 2006; Gnedin & Abel 2001). Here we introduce a photon packet merging procedure that strictly limits the number of packets that need to be transmitted per particle to at most N_c , without restricting the number of directions along which each individual packet can propagate.

For each gas particle i receiving a photon packet we define a so-called reception cone tessellation, as shown in the middle panel of Fig. 4.4. A reception cone tessellation is a tessellating set of N_c cones attached to a gas particle. It is identical to the set of N_c cones employed for the emission of photons by source particles, but generally has a different orientation. As their name indicates, and in contrast to the emission cones, reception cones are used to collect photons. Whenever a photon packet p is transferred to gas particle i , the packet is binned into one of the reception cones by examining into which reception cone its emission direction \mathbf{n}_p falls, after translation to the location of the gas particle. Photon packets whose emission directions fall in one and the same cone are merged, leaving only a single packet for transmission (Fig. 4.4, right-hand panel). The reception cone tessellation is given a random orientation to prevent artefacts.

A merged photon packet created in cone k is assigned the new, merged emission direction $\mathbf{n}_{m,k}$. This new emission direction is defined as the weighted sum⁵ $\mathbf{n}_{m,k} = \sum w_p \mathbf{n}_p / |\sum w_p \mathbf{n}_p|$, where the sum is over all photon packets received in emission cone k ($k = 1, 2, \dots, N_c$), and the weights w_p are the number of photons per packet p that are to be transmitted into direction \mathbf{n}_p to the downstream neighbours of particle i (Fig. 4.4, right-hand panel). Similarly, the merged photon packet is associated a merged clock $t_{m,k}^*$ by averaging over the individual clock times t_p^* in reception cone k , i.e. $t_{m,k}^* = \sum w_p t_p^* / \sum w_p$.

As a result of the photon packet (resp. source) merging at most N_c packets have to be transmitted per gas particle, fully consistent with the angular resolution of our transport scheme. The computation effort required to perform radiative transfer simulations with TRAPHIC can therefore be readily controlled: In the most demanding situation, i.e. in the optically thin limit, when the whole box is filled with radiation, it merely scales with the product of spatial and angular resolution, $N \times \tilde{N}_{\text{ngb}} \times N_c$, where N is the total number of particles (SPH + stars). The merged photon packets are not associated with any existing source particle in the simulation. They should be thought of as emerging from an imaginary source, whose properties are implicitly defined by our merging prescription. The merged photon packets are traced further downstream, radially away from this imaginary source, into the merged emission direction,

⁵We note that this expression contains a typo in the original publication, i.e. in Pawlik & Schaye (2008). The sum used in that publication does not result in a unit vector, which is inconsistent with the employed notation.

exactly as was described for non-merged emission directions in the previous section.

4.4.3 Photon interactions with gas particles

Photons are propagated radially away from each source until they interact with the matter represented by the SPH particles. Here we distinguish between absorptions and scattering interactions and describe how they are accounted for in TRAPHIC. We remind the reader that in this chapter we do not discuss the effect of interactions on the thermodynamical and hydrodynamical evolution of the SPH particles.

Absorption

Absorptions remove photons from the beam emitted by the source. This process is described in terms of the mass absorption coefficient κ_ν . From the formal solution of the radiative transfer equation it can easily be seen (e.g. Mihalas & Weibel Mihalas 1984) that the fraction of photons of frequency ν that is absorbed while travelling along the straight line connecting the positions \mathbf{r}_1 and \mathbf{r}_2 is given by $1 - \exp(-\tau)$, where

$$\tau = \int_{\mathbf{r}_1}^{\mathbf{r}_2} dr \kappa_\nu(r) \rho(r), \quad (4.10)$$

is the optical depth over the distance $d = |\mathbf{r}_1 - \mathbf{r}_2|$.

In TRAPHIC, absorptions are accounted for by removing the photon fraction $1 - \exp(-\tau_{ij})$ from photon packets propagating between particle i and its neighbouring particles j . For the calculation of the optical depth τ_{ij} between these two particles we need to numerically evaluate the integral Eq. 4.10. This is computationally expensive, since it involves the calculation of the density ρ (and the absorption coefficient κ_ν) at a large number of points along the photon path using the SPH formulation. Similar to the approach of Kessel-Deynet & Burkert (2000), we therefore approximate the density field near a gas or virtual particle i by the SPH density ρ_i evaluated at its position.

Since photon packets are propagated between particle i , located at \mathbf{r}_i , and particle j , located at \mathbf{r}_j , along their emission direction, i.e. in the direction of the unit vector \mathbf{n} , the distance d they cover is generally not equal to the distance between the particles. Instead, we set $d = |\mathbf{n} \cdot (\mathbf{r}_i - \mathbf{r}_j)|$, i.e., we employ the projection of the particle distance along the emission direction. This is the correct limit far away from the source from which the photon packets originate, but is not applicable close to it. Therefore, we only employ the projected distance when propagating photons between a transmitting particle and its neighbours. On the other hand, when propagating photons from a source particle to its neighbours, we use the particle distance, i.e. we set $d = |\mathbf{r}_j - \mathbf{r}_i|$, corresponding to radially outward propagation.

The distance d is used to obtain the optical depth τ_{ij} between particle i and particle j . We set $\tau_{ij} = \kappa_\nu \rho_j d$. This approximation neglects the existence of any substructure between particle i and its neighbouring particle j . However, this does not result in a loss of information if the radius \tilde{h}_i of the neighbourhood of particle i is chosen such that $\tilde{h}_i \lesssim h_i$ (resp. $\tilde{N}_{\text{ngb}} \lesssim N_{\text{ngb}}$).

The distance d is furthermore used to advance the clock t^* of each photon packet according to the rule

$$t^* \rightarrow t^* + d/c. \quad (4.11)$$

By synchronising the clock t^* with the simulation time t_r , the clock can be employed to advance photon packets at the speed of light, as we discuss in Sec. 4.4.4. We note in passing that the clock t^* can also be used to implement the cosmological redshifting of photons.

The number of absorbed and transmitted photons is now calculated as follows. From the photon packet emitted or transmitted by particle i to particle j , a fraction $1 - \exp(-\tau_{ij})$ is absorbed by particle j . The photon packet is subsequently transmitted by particle j containing a fraction $\exp(-\tau_{ij})$ of the original number of photons. By construction, this procedure explicitly conserves photons, since

$$1 - e^{-\tau_{ij}} + e^{-\tau_{ij}} = 1. \quad (4.12)$$

Photons absorbed by ViPs are distributed amongst all neighbours of the ViP using photon-conserving SPH interpolation. This is necessary, because ViPs are a temporary construct employed for the transport of photons; permanent information is only stored at the particles in the SPH simulation. For further reference, we denote the total number of photons impinging on and the total number of photons absorbed by particle i over a single time step Δt_r with $\Delta \mathcal{N}_{\text{in},i}$ and $\Delta \mathcal{N}_{\text{abs},i}$, resp.

Scattering

In contrast to absorptions, scatterings change the direction and possibly the frequency of the interacting photons. A scattering event can be thought of as an absorption event followed by an immediate re-emission. Scatterings can hence be described by Eq. 4.7 after adapting the absorption coefficient κ_ν and the emissivity ϵ_ν . The re-emission re-distributes the photons in angle (and possibly frequency). For this reason the scattered photons are sometimes referred to as diffuse. Adapting the emissivity for scatterings requires evaluating an integral over the angular dependence of the intensity. Therefore, scatterings turn the radiative transfer equation (Eq. 4.7) into an integro-differential equation.

In TRAPHIC, scatterings, that is the transport of diffuse photons, are modelled by a combination of absorption and re-emission events. Photons to be scattered travelling between two particles are first removed from the photon packet as described in the last section. For a true absorption event, the photon energy would be lost to the thermal bath provided by the matter. In contrast, in case of scattering, a gas particle that absorbed photons re-emits them, closely following the description for source particles in Sec.4.4.2.

In particular, we employ randomly oriented emission cones in combination with Eq. 4.8 to re-emit the absorbed photons. Since modelling the scattering process means following photons emitted earlier by a source, the clock of the photon packets that are to be re-emitted is not set to the simulation time t_r , as was the case for the intrinsic emission of photons by source particles. Instead, it is determined by the clock of the photon packets that are scattered and the details of the scattering interactions (e.g. there could be a time delay between absorption and re-emission, or photons could have travelled a distance greater than the (projected) particle distance d when inferring the properties of the scattering process from a sub-resolution model). If the scattering details are taken into account by correspondingly adjusting d , Eq. 4.11 can be employed to find the new clock time. Note that it is crucial for the modelling of scatterings to employ a radiative transfer scheme which does not scale with the number of sources, since every gas particle will soon re-emit photons.

4.4.4 Solving the radiative transfer equation

TRAPHIC solves the radiative transfer equation (Eq. 4.7) in discrete time steps Δt_r , the size of which depends on the details of the problem under study and has to be decided on a case by case basis. Hence, we defer its discussion to chapter 5, where we present an implementation of

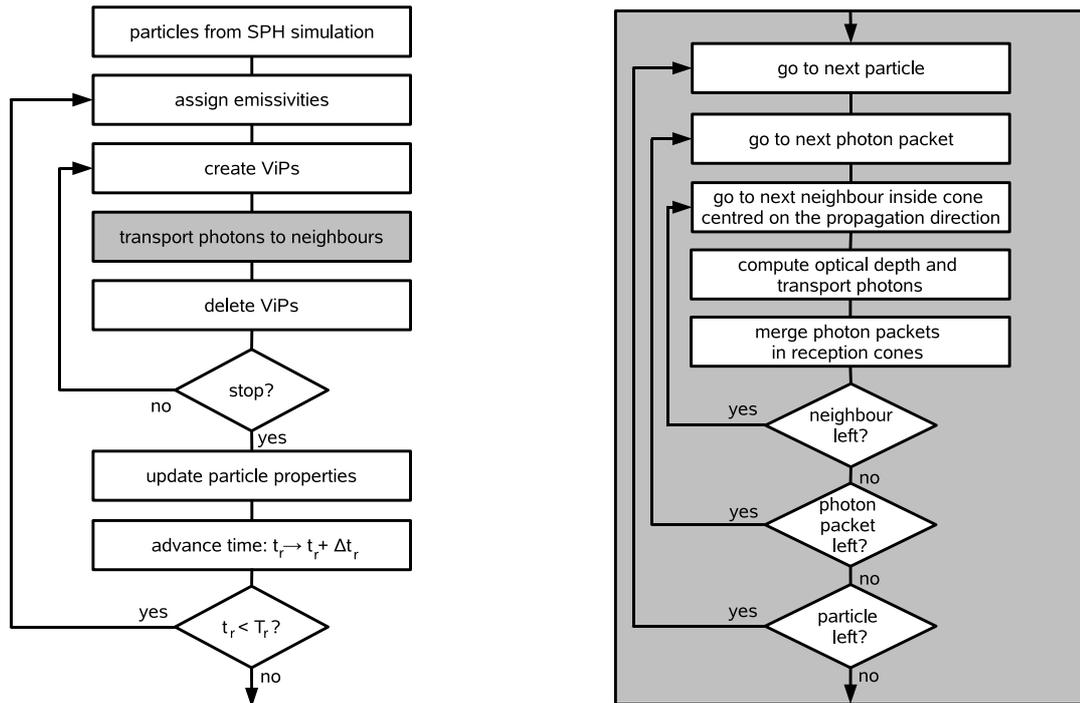


Figure 4.5: Flow charts. *Left-hand chart:* Overview of the radiation transport with TRAPHIC. The radiative transfer simulation starts with the assignment of emissivities to source particles. Thereafter, photon packets are transported from all particles to their neighbours (see right-hand chart for details). ViPs, if needed, are created in advance of the transport and are deleted immediately afterwards. The transport cycle continues until the user-defined stopping criterion is satisfied. Then, the properties of the particles are updated according to the photon-matter interactions that occurred. Finally, the radiative transfer time step is advanced. The radiative transfer simulation ends at time T_r . *Right-hand chart:* Details of the transport of photons to neighbours (grey box in the left-hand panel). Each particle holds at most N_c photon packets. The photons in these packets are distributed amongst the \tilde{N}_{ngb} neighbouring particles inside (tessellating) emission or (regular) transmission cones centred on the corresponding propagation directions. The optical depth to each neighbour is computed and the fraction of transmitted and absorbed/scattered photons is determined. Multiple photon packets received by individual neighbours are merged in (tessellating) reception cones, which limits the number of photon packets held by each particle to at most N_c .

our method to solve specific problems. In this section we briefly review the concepts described in the previous sections used to obtain the solution to the radiative transfer equation over a single time step before discussing how this solution is advanced in time. A schematic summary is provided by the flow chart in Fig. 4.5.

Starting at simulation time t_r , the size of the radiative transfer time step Δt_r is determined. Source particles emit photon packets to their \tilde{N}_{ngb} neighbouring gas particles employing N_c cones. The total number of photons contained in the packets is determined by the number of photons $\dot{N}_{\nu,i}$ to be emitted per unit time in frequency bin ν and by the size of the time step Δt_r . For each neighbouring gas particle the optical depth to the source particle is evaluated and the number of photons interacting with the matter on their way from the source particle is inferred. Some photons may be absorbed, some may be scattered; the remaining photons are left unimpeded for transmission.

For the subsequent transmission of photons by the gas particles, photons associated to sources that are seen as close (in angle) to each other are merged, limiting the number of directions into which photons have to be transmitted to at most N_c . Next, for each transmitting and scattering gas particle the \tilde{N}_{ngb} neighbours are found. Photons to be transmitted are distributed amongst the subset $\tilde{N}_{\text{t,ngb}}$ of neighbours that are located downstream, as determined by the corresponding (merged) emission directions. Photons to be scattered are distributed amongst all \tilde{N}_{ngb} neighbours following the prescription in Sec.4.4.3. Virtual particles (ViPs), which are placed in cones that do not contain any neighbours, are deleted after having transmitted or scattered their photons. Photons that were absorbed by a ViP are distributed amongst the neighbours of the ViP using photon-conserving SPH interpolation.

The cycle described in the previous paragraph repeats until a stopping criterion is satisfied. The form of the stopping criterion depends on whether one solves the time-*independent* or the time-*dependent* radiative transfer equation. When solving the time-independent radiative transfer equation, i.e. when neglecting the first term on the left hand side of Eq. 4.7, one formally assumes $c \rightarrow \infty$. Accordingly, the stopping criterion becomes independent of the speed of light. The cycle can then for example be repeated until all photons have been absorbed or have left the simulation box. In contrast, when solving the time-dependent radiative transfer equation, the stopping criterion is directly tied to the speed of light c : the cycle is repeated until all photons have propagated a distance $c\Delta t_r$.

To decide whether or not a photon packet has travelled over the distance $c\Delta t_r$, we employ its clock t^* . For further illustration it is convenient to explicitly follow the clock of a photon packet as it is propagated through the simulation box. Upon emission, the clock of the photon packet is initialised with the simulation time t_r , as already mentioned in Sec. 4.4.2. After the emission, when the photon packet has been transferred from the source to one of the neighbouring SPH particles (or to a ViP), its clock is advanced according to Eq. 4.11. It is then checked whether the clock has reached or crossed the threshold value

$$t_{\text{thr}}^* = t_r + \Delta t_r. \quad (4.13)$$

If not, the photon packet is transmitted further downstream. As before, its clock is advanced according to Eq. 4.11. The photon packet is repeatedly transmitted and its clock is advanced until it has reached or crossed the threshold value defined in Eq. 4.13. At this point, the value displayed by its clock can in full generality be expressed as $t^* = t_{\text{thr}}^* + \epsilon_0$, with $\epsilon_0 \geq 0$.

The propagation error, ϵ_0 , is typically larger than zero because the set of particles on which the photon packets are propagated is discrete, with the particle-to-neighbour distances generally not related to $c\Delta t_r$. Employing Eq. 4.13 to stop photon packets therefore results in the pho-

ton packets typically being propagated too far. Due to the Lagrangian nature of the SPH simulation, ϵ_0 is individual for each photon packet. If d is the particle-to-neighbour distance corresponding to the emission or transmission event after which the photon packet was stopped, then $\epsilon_0 < d/c$. The propagation error ϵ_0 is therefore strictly limited, $\epsilon_0 < \max_i \tilde{h}_i/c$, where the maximum is over all emitting and transmitting particles. Note that ϵ_0 becomes smaller with higher spatial resolution.

Once stopped, a photon packet is held (and hence its clock stands still) until the simulation time progresses to a value $t_r + \Delta t_r > t^*$, at which point the packet is transmitted and its clock is advanced again. The clocks of photon packets are thus kept in synchronisation with the current simulation time, which effectively matches their propagation speed to become the speed of light⁶. Photon packets may be held over several radiative transfer time steps, as ϵ_0 may be (several times) larger than Δt_r . While the photon packet is held, its propagation error with respect to the current simulation time steadily decreases with each passing radiative transfer time step. Since t^* is fixed while t_r is increasing in steps of Δt_r , the current propagation error decreases according to $\epsilon_n = \epsilon_{n-1} - \Delta t_r$, where n indicates the number of radiative transfer time steps that have passed since the photon packet was stopped. Immediately before the photon packet is propagated again, its propagation error is strictly limited by the size of the time step, $-\Delta t_r \leq \epsilon_n < 0$. The photon packet is then propagated such as to again bring the clock into synchronisation with the simulation time, as described above.

When multiple photon packets are merged into a single packet, the clock of the merged packet is determined by the averaging procedure described in Sec.4.4.2. As a result, photons may be propagated over distances that differ somewhat from the case without merging. We have seen above that, due to the synchronisation of the photon packet propagation with the current simulation time t_r , clocks display only values in the interval $t_r \leq t^* \leq t_r + \Delta t_r + \epsilon_0$. The difference in the clocks of different photon packets and hence the photon propagation error introduced by the merging procedure can therefore be controlled to become arbitrarily small by increasing the temporal (i.e. decreasing Δt_r) and spatial (leading to smaller ϵ_0) resolution.

After the stopping criterion (time-independent or time-dependent) has been fulfilled for all photon packets, the properties of gas particles are updated according to the radiative interactions that occurred. We emphasise that the properties of the gas particles (e.g. the ionization state) are therefore only updated at the end of each time step. Within each time step the order in which photon packets from different sources arrive at gas particles is therefore irrelevant. After updating the particle properties, the simulation time is advanced, $t_r \rightarrow t_r + \Delta t_r$. Finally, the size of the next time step is determined and the radiation is transported as described above.

In summary, we have presented a radiative transfer scheme for use in SPH simulations that works directly on the unstructured grid formed by the discrete set of irregularly distributed SPH particles. TRAPHIC thus employs the full spatial resolution of the SPH simulation. It achieves directed transport of radiation by adaptively tracing photon packets in cones. TRAPHIC can be used to solve both the time-independent and the time-dependent radiative transfer equation in an explicitly photon-conserving way. Our scheme is by construction parallel on distributed memory machines if the SPH simulation itself is parallel on distributed memory machines. Furthermore, the computation time necessary to accomplish the radiation transport does not scale linearly with the number of sources. Instead, it merely scales with the product of spatial and angular resolution, making our scheme suitable for simulations contain-

⁶We note that the clocks can also be employed to propagate photon packets at speeds \tilde{c} different from the speed of light (cp. the reduced speed of light approximation suggested in Gnedin & Abel 2001), by simply making the replacement $c \rightarrow \tilde{c}$.

ing a large number of sources as well as for taking into account a diffuse radiation component.

4.4.5 Reduction of particle noise

The radiative transfer equation describes the propagation of photons within a continuum. In our scheme, however, photons are propagated on the discrete set of SPH particles, localised at irregular positions dictated by the SPH simulation. Consequently, numerical noise may arise from the discreteness and irregularity of the spatial distribution of SPH particles and this could influence the numerical solution of the radiative transfer equation obtained with TRAPHIC.

To see how this *particle noise* can be reduced, it is helpful to employ a formal analogy (e.g. Monaghan 2005) between estimating the density in SPH simulations and estimating a probability density from sample points: We can consider the positions of the SPH particles as a random⁷ sample drawn from a probability density function proportional to the mass density. We can therefore Monte Carlo resample⁸ the density field without sacrificing its information content. Periodically assigning new positions to the SPH particles during the radiation transport by resampling the density field should lead to a better representation of the continuum physics by the discrete set of SPH particles and hence lead to a reduction of particle noise.

To employ the resampling, we think of particle i at position \mathbf{r}_i in the simulation box as being de-localised within its sphere of influence of radius h_i centred on \mathbf{r}_i and assume that the probability of finding it in the volume d^3r around a particular point \mathbf{r} in that sphere is given by the value of the interpolation kernel at that point (cp. the scattering approach of Sec.4.2). For the radiative transfer on a static set of particles with positions \mathbf{r}_i we then periodically Monte Carlo re-distribute the particles within their sphere of influence according to this probability. We note that while we change the positions of the SPH particles, we keep all their other properties (e.g. density) fixed in order to avoid introducing scatter in the physical variables. In the tests performed in Chapter 5 we will demonstrate that the application of this recipe leads to a significant reduction of particle noise. There we will, however, also see that for most applications the noise is small. We therefore do not employ the resampling technique in our simulations, unless explicitly stated.

Since the resampling randomly offsets the SPH particles from their positions provided by the hydrodynamical simulation, the apexes of the transmission cones attached to them are randomly offset, too. The resampling procedure may therefore lead to a slight diffusion of photons out of the emission cone they were emitted into, effectively decreasing the angular resolution. Note that a similar diffusion of photons could occur when solving the radiative transfer equation coupled to the hydrodynamics (on which we will report in a future publication), because of the physical particle motion. We will study this diffusion in our numerical implementation of TRAPHIC that we present in Chapter 5, but note already that it does not lead to any noticeable degradation of the effective angular resolution and hence leaves TRAPHIC's shadowing properties unaffected.

⁷We stress that SPH itself is not a Monte Carlo method, as first noted in Gingold & Monaghan (1978). We only appeal to the Monte Carlo picture for use with the radiative transfer, not for the hydrodynamic evolution of the particles in the SPH simulation.

⁸We explicitly note that the resampling does not affect the hydrodynamical simulation, for which the original positions are used.

4.5 SUMMARY

In this chapter we have presented TRAPHIC (TRANsport of PHotons In Cones), a novel radiative transfer scheme for SPH simulations. TRAPHIC has been designed for use in radiation-hydrodynamical simulations exhibiting a wide dynamic range and containing a large number of light sources, as for instance encountered in cosmological simulations. The requirements for this are high: Due to the limits on the computational power available today, radiative transfer simulations at the resolution of state-of-the-art hydrodynamical simulations need to be adaptive both in space and in angle and parallel on distributed memory machines.

TRAPHIC meets these requirements. The radiative transfer equation is solved by transporting photons in an explicitly photon-conserving manner directly on the discrete set of SPH particles. Photons are transported globally by propagating them locally, between a particle and its neighbours. For a fixed number of SPH particles, the number of neighbours employed for the transport, the parameter \tilde{N}_{ngb} , sets the adaptive spatial resolution. The conceptual similarity between the photon transport and the calculation of particle properties in SPH simulations allows a straightforward numerical implementation of TRAPHIC in SPH simulations for application on distributed memory machines.

In TRAPHIC photon packets are transported inside cones. Because source particles emit into a set of cones that tessellates the simulation box, the angular dependence of the emission can be modelled independently of the spatial distribution of the neighbouring SPH particles. The number of cones, N_c , is the second parameter employed in TRAPHIC and determines the formal angular resolution of the radiative transfer simulation.

After their emission, the photon packets are transported downstream from SPH particle to SPH particle. They are confined to the solid angle they were originally emitted into by regular cones of solid angle $4\pi/N_c$. Mimicking the ray-splitting technique used in ray-tracing schemes, the photon transport is adaptive in angle. The effective angular resolution is therefore much higher than the formal one. As we will demonstrate in the next chapter, Chapter 5, radiative transfer simulations employing TRAPHIC can thus be performed using only a relatively small number of cones.

The propagation of photons inside cones that do not contain a neighbour requires the introduction of so-called virtual particles (ViPs). It is the concept of cones combined with ViPs that enables the directed transport of photons on the unstructured grid defined by the discrete set of irregularly distributed SPH particles.

A practical problem often encountered when solving the radiative transfer equation is the linear scaling of the computational effort with the number of sources. Such a scaling limits the application of most of the existing radiative transfer schemes to simulations containing only a few sources. In TRAPHIC this problem is circumvented by introducing a source merging procedure. Photon packets emitted by different sources are merged in accordance with the angular resolution employed, such that at any point in space at most N_c photon packets need to be propagated. This renders the photon transport effectively independent of the number of sources in the simulation and makes it feasible to include a diffuse radiation component emitted by the SPH particles. In the most extreme case of radiation from an arbitrary number of sources completely filling the simulation box, the computational effort required by TRAPHIC merely scales as the product of spatial and angular resolution, $N \times \tilde{N}_{\text{ngb}} \times N_c$, where N is the total number of particles (SPH + stars).

Applications of TRAPHIC include the solution of both the time-independent and time-dependent radiative transfer equation in large-scale simulations of cosmological reionisation. In the next chapter, Chapter 5, we will show that TRAPHIC is able to accurately reproduce the

expected growth of the ionised sphere around a single ionising point source, both in a homogeneous and in a centrally peaked medium, and to cast sharp shadows behind opaque obstacles. Furthermore, we will test our scheme in a physical setting of complex geometry, simulating the growth of ionised regions around multiple point sources in a cosmological density field. We will find that the results obtained with our implementation of TRAPHIC are in excellent agreement with results obtained with other radiative transfer codes for the same test problem.

ACKNOWLEDGEMENTS

We thank Claudio Dalla Vecchia for stimulating discussions and Huub Röttgering for his valuable advice. We are grateful to Craig Booth and Benedetta Ciardi for a thorough reading of the draft version of Pawlik & Schaye (2008). This work was supported by Marie Curie Excellence Grant MEXT-CT-2004-014112.

REFERENCES

- Abel T., Norman M. L., Madau P., 1999, *ApJ*, 523, 66
Abel T., Wandelt B. D., 2002, *MNRAS*, 330, L53
Altay G., Croft R. A. C., Pelupessy I., 2008, *MNRAS*, 386, 1931
Alvarez M. A., Bromm V., Shapiro P. R., 2006, *ApJ*, 639, 621
Barkana R., Loeb A., 2004, *ApJ*, 609, 474
Brookshaw L., 1994, *MmSAI*, 65, 1033
Cen R., 2002, *ApJS*, 141, 211
Ciardi B., Ferrara A., Marri S., Raimondo G., 2001, *MNRAS*, 324, 381
Croft R. A. C., Altay G., 2007, *arXiv*, 709, [arXiv:0709.2362](https://arxiv.org/abs/0709.2362)
Dale J. E., Ercolano B., Clarke C. J., 2007, *MNRAS*, 1056
Dale J. E., Bonnell I. A., Whitworth A. P., 2007, *MNRAS*, 375, 1291
Dopita M. A., Sutherland R. S., 2003, *Astrophysics of the diffuse universe*, Springer
Finlator K., Özel F., Davé R., 2009, *MNRAS*, 393, 1090
Fryer C. L., Rockefeller G., Warren M. S., 2006, *ApJ*, 643, 292
Gingold R. A., Monaghan J. J., 1977, *MNRAS*, 181, 375
Gingold R. A., Monaghan J. J., 1978, *MNRAS*, 184, 481
Gingold R. A., Monaghan J. J., 1982, *J. Comput. Phys.*, 46, 429
Gnedin N. Y., Abel T., 2001, *NewA*, 6, 437
Gritschneder M., Naab T., Heitsch F., Burkert A., 2007, *IAUS*, 237, 246
Herant M., Benz W., Hix W. R., Fryer C. L., Colgate S. A., 1994, *ApJ*, 435, 339
Hernquist L., Katz N., 1989, *ApJS*, 70, 419
Hockney R. W., Eastwood J. W., 1988, *Computer Simulations Using Particles*, Taylor & Francis
Iliev I. T., Mellema G., Pen U.-L., Merz H., Shapiro P. R., Alvarez M. A., 2006, *MNRAS*, 369, 1625
Iliev I. T., et al., 2006, *MNRAS*, 371, 1057
Johnson J. L., Greif T. H., Bromm V., 2007, *ApJ*, 665, 85
Kessel-Deynet O., Burkert A., 2000, *MNRAS*, 315, 713
Kohler K., Gnedin N. Y., Hamilton A. J. S., 2007, *ApJ*, 657, 15
Lucy L. B., 1977, *AJ*, 82, 1013

- Maselli A., Ferrara A., Ciardi B., 2003, *MNRAS*, 345, 379
- Mayer L., Lufkin G., Quinn T., Wadsley J., 2007, *ApJ*, 661, L77
- Shampine L. F., Gear C. W., 1979, *SIAM Review*, Vol. 21, No.1, pp. 1-17
- Mellema G., Iliiev I. T., Alvarez M. A., Shapiro P. R., 2006, *NewA*, 11, 374
- Mihalas D., Weibel Mihalas B., 1984, *Foundations of radiation hydrodynamics*, Oxford University Press, New York
- Monaghan J. J., 1992, *ARA&A*, 30, 543
- Monaghan J. J., *Rep. Prog. Phys.*, 68, 1703 (2005)
- Okabe A., Boots B., Sugihara K., Chiu S. N., 2000, *Spatial Tessellations*, Second edition, Wiley
- Osterbrock D. E., 1989, *Astrophysics of gaseous nebulae and active galactic nuclei*, Palgrave Macmillan
- Oxley S., Woolfson M. M., 2003, *MNRAS*, 343, 900
- Paschos P., Norman M. L., Bordner J. O., Harkness R., 2007, *arXiv*, 711, arXiv:0711.1904
- Pawlik A. H., Schaye J., 2008, *MNRAS*, 389, 651
- Petkova M., Springel V., submitted
- Price D., 2005, *astro*, arXiv:astro-ph/0507472
- Razoumov A. O., Cardall C. Y., 2005, *MNRAS*, 362, 1413
- Razoumov A. O., Sommer-Larsen J., 2006, *ApJ*, 651, L89
- Ritzerveld J., Icke V., 2006, *PhRvE*, 74, 026704
- Semelin B., Combes F., Baek S., 2007, *arXiv*, 707, arXiv:0707.2483
- Soneira R. M., Peebles P. J. E., 1978, *AJ*, 83, 845
- Springel V., Hernquist L., 2002, *MNRAS*, 333, 649
- Springel V., 2005, *MNRAS*, 364, 1105
- Stamatellos D., Whitworth A. P., 2005, *A&A*, 439, 153
- Susa H., 2006, *PASJ*, 58, 445
- Susa H., Umemura M., 2006, *ApJ*, 645, L93
- Trac H., Cen R., 2006, *astro*, arXiv:astro-ph/0612406
- Viau S., Bastien P., Cha S.-H., 2006, *ApJ*, 639, 559
- Whalen D., Norman M. L., 2007, *arXiv*, 708, arXiv:0708.2444
- Wise J. H., Abel T., 2007, *arXiv*, 710, arXiv:0710.3160
- White S. D. M., 1996, *Cosmology and large scale structure*, Proceedings of the "Les Houches Ecole d'Ete de Physique Theorique" (Les Houches Summer School), p. 349, Elsevier Scientific Publishing Company, Amsterdam
- Whitehouse S. C., Bate M. R., 2006, *MNRAS*, 367, 32
- Yoshida N., Oh S. P., Kitayama T., Hernquist L., 2007, *ApJ*, 663, 687

4.A THE ANISOTROPY OF PARTICLE-TO-NEIGHBOR TRANSPORT

In Sec.4.4.2 we introduced an emission cone tessellation to accomplish the emission of photons by a source particle to its neighboring gas particles. Then we also mentioned that the cones are required to emit in agreement with the angular dependence of the emissivity of the source, independently of the spatial distribution of the neighbors. In this appendix we explain why this is the case. In particular, assuming that particles have equal mass, we show that when transporting photons (or more generally, any extensive quantity) in particle simulations from a particle to its neighbors, the net transport direction is generally correlated with the direction

towards the centre of mass of the neighbouring particles. As a result, the transport is partly governed by the geometry of the simulation, in addition to the intrinsic emissivity of the source.

Consider a particle located at the origin O of a 3-dimensional coordinate system (the emitter frame). The particle emits photons to its \tilde{N}_{ngb} nearest neighbours, residing in the sphere of radius \tilde{h} centred on the emitting particle. Throughout this section we assume that all particles have equal mass, $m = 1$. The emission is performed by transferring the fraction $w_i / \sum_{j=1}^{\tilde{N}_{\text{ngb}}} w_j$ of the total number of photons to neighbour i ($i = 1, 2, \dots, \tilde{N}_{\text{ngb}}$), where the $w_i \geq 0$ are weights to be specified. The emission properties will depend both on the weights and on the spatial distribution of the neighbours.

To study the angular dependence of the emission process, we introduce the vector sum $\mathbf{s} = \sum_i w_i \mathbf{u}_i / \sum_i w_i$ over all unit vectors $\mathbf{u}_i \equiv \mathbf{r}_i / |\mathbf{r}_i|$, where \mathbf{r}_i is the position of neighbour i . This vector sum can be interpreted as the net direction of the emission. Let us denote the angle enclosed by \mathbf{s} and the vector to the centre of mass \mathbf{r}_{cm} of the neighbours, $\mathbf{r}_{\text{cm}} = \sum_i \mathbf{r}_i / \tilde{N}_{\text{ngb}}$, with α . We ask for the probability density function (pdf) $p(\cos \alpha)$ of the cosine of α , where

$$\cos \alpha = \frac{\mathbf{s} \cdot \mathbf{r}_{\text{cm}}}{|\mathbf{s}| |\mathbf{r}_{\text{cm}}|} = \frac{\sum_i w_i \mathbf{u}_i \cdot \sum_j \mathbf{r}_j}{|\sum_i w_i \mathbf{u}_i| |\sum_j \mathbf{r}_j|}. \quad (4.14)$$

The meaning of the pdf can be understood by looking at its extremes. If $p(\cos \alpha)$ is strongly peaked around $\cos \alpha \approx 1$, the emission has a net direction biased towards the centre of mass direction, and photons will be preferentially transported into the high (particle) density regions. On the other hand, if $p(\cos \alpha)$ is strongly peaked around $\cos \alpha \approx -1$, photons will be preferentially transported into the direction opposite to the centre of mass. Finally, a flat pdf $p(\cos \alpha) = 1/2$ indicates that the net emission direction and the direction towards the centre of mass are uncorrelated.

As an illustration, assume that all neighbours have identical weights $w_i = 1$ and that their positions \mathbf{r}_i are drawn from a probability distribution that uniformly samples the surface of the sphere of radius \tilde{h} surrounding the emitter at O . To obtain $p(\cos \alpha)$, we note that our assumptions imply that the centre of mass vector \mathbf{r}_{cm} and the vector sum \mathbf{s} point in the same direction, so that $p(\cos \alpha) = 2\delta_{\text{D}}(\cos \alpha - 1)$, where $\delta_{\text{D}}(x)$ is the Dirac δ -function. Hence, there is a perfect correlation between \mathbf{s} and \mathbf{r}_{cm} .

Another example is given by assuming that the particles of unit mass are distributed randomly within the sphere of radius \tilde{h} around the emitter. The resulting pdf $p(\cos \alpha)$ as obtained through a Monte Carlo simulation is shown in its cumulative form $F(\cos \alpha) = \int_{-1}^{\cos \alpha} dx p(x)$ in Fig. 4.6 for two choices of the number of neighbours, $\tilde{N}_{\text{ngb}} = 16$ (thin dotted curve) and $\tilde{N}_{\text{ngb}} = 64$ (thick dotted curve; falling nearly on top of the thin dotted curve). Clearly, there is a very strong correlation between the net direction of emission and the direction towards the centre of mass of the neighbour distribution.

For the important case of isotropic sources of photons the net emission direction should not be correlated with the direction towards the centre of mass, i.e. $p(\cos \alpha) = 1/2$, independent of $\cos \alpha$, and thus $F(\cos \alpha) = (1 + \cos \alpha)/2$. From here on, we will refer to this case as isotropic emission of radiation. We emphasize the statistical character of this statement. An individual particle may still emit anisotropically, which could be revealed by studying the amplitude $|\mathbf{s}|$. For emission that is isotropic for individual particles, $|\mathbf{s}| = 0$. Consider the neighbour distribution of the last example, for which we showed the emission pattern to be far from isotropic. Can the emission be isotropized by tuning the weights?

For instance, consider solid angle weighting, defined as follows. We project all neighbours radially onto the unit sphere. To each neighbour we attach a weight proportional to its area of

influence, which we define by the collection of all points on the unit sphere which are closer to the projected neighbour for which we calculate the weight than to all other projected neighbours. Hereby, the distance of two points on the sphere is given by the arclength of the connecting geodesic (great circle). This definition of the area of influence results in a specific tessellation of the surface of the sphere, the so-called Voronoi tessellation (e.g. Okabe et al. 2000), and provides us with well-defined non-overlapping solid angles.

The resulting cumulative distribution $F(\cos \alpha)$ is shown in Fig. 4.6 (dashed curves). Although the weighted emission has weakened the correlation between \mathbf{s} and \mathbf{r}_{cm} as compared to the case of equal weights, the net emission direction is still markedly peaked towards the direction to the centre of mass. We note that weighting by solid angle is the most natural choice and its inability to isotropize the transport might be surprising. Alternatively, one could introduce *ad hoc* weights $\hat{w}_i = w_i^r$, where $r > 1$ is some exponent and the w_i are solid angle weights, to re-shape the distribution $p(\cos \alpha)$ to become less and less peaked around $\cos \alpha \sim 1$.

However, the reason why the solid angle weights fail to isotropize the transport is not that the weights are inappropriate, but that the directions to the neighbours are fixed and thus cannot be freely chosen along with the weights, as mentioned in Sec.4.4.2. This strongly limits one's ability to influence the shape of $p(\cos \alpha)$. For illustration, consider an emitter having all its neighbours located in one hemisphere. The net emission direction \mathbf{s} will then necessarily lie in that hemisphere, too, *regardless* of the chosen weights. For a statistical ensemble of emitters, \mathbf{s} will then necessarily be correlated with the direction towards the centre of mass.

This also explains the dependence of $p(\cos \alpha)$ on the number of neighbours. The larger \tilde{N}_{ngb} , the greater the probability of finding a neighbour in any chosen solid angle. Through this increase in directional sampling, $p(\cos \alpha)$ can be expected to approach the uncorrelated case. This behaviour can indeed be observed in Fig. 4.6.

We also compute the statistic $p(\cos \alpha)$ for the case of neighbours clustered in space, which is more typical for cosmological simulations of structure formation. We did this both using a toy clustering model (Soneira & Peebles 1978) and for a cosmological density field obtained from a Λ CDM hydrodynamical simulation (the same field that is used in Sec.5.4.5, Chapter 5). For the weighting schemes we investigated the dependence of the shape of $p(\cos \alpha)$ on the chosen weights is qualitatively very similar to that in the case of the random distribution of neighbours discussed earlier.

We emphasize that our discussion is not limited to the transport of photons. Indeed, particle-to-neighbour transport is for example routinely employed in SPH simulations of galaxy formation, where a star particle distributes its metals and energy amongst its neighbours. Different recipes for distributing metals and energy can be found in the literature. Most of these apply a weighting related to the SPH formalism. As an example, to each neighbour i we associate the weight

$$w_i = \frac{\frac{m_i}{\rho_i} W(r_i, \tilde{h})}{\sum_j \frac{m_j}{\rho_j} W(r_j, \tilde{h})}. \quad (4.15)$$

Here, \tilde{h} is the radius of the neighbourhood of the emitting particle (at O), $r_i = |\mathbf{r}_i|$ and W is the spline kernel Eq. 4.4. Performing the transport on the same random neighbour distribution as before, we obtain the distribution $p(\cos \alpha)$. As can be seen in Fig. 4.6 (solid curves), the direction of net transport and the direction to the centre of mass are found to be almost uncorrelated. We note that this result, the reason of which at the moment is not clear to us, is only a statistical statement. We calculated $|\mathbf{s}|$ and found that the emission of individual particles is still not close to isotropic. We arrive at qualitatively similar results for different definitions of

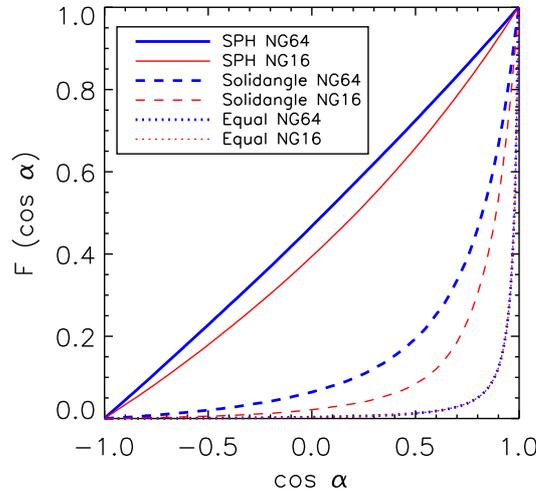


Figure 4.6: Monte Carlo simulation of particle-to-neighbour transport. We show the cumulative probability distribution of $\cos \alpha$, where α is the angle between the net emission direction \mathbf{s} and the direction towards the centre of mass \mathbf{r}_{cm} of the neighbours. The curves correspond to the following weighting schemes, with the number of neighbours indicated in brackets (from top to bottom): SPH (64), SPH (16), Solidangle (64), Solidangle (16), Equal (64), Equal (16). The last two curves are nearly on top of each other.

the SPH weights, both for random and clustered neighbour distributions. A more quantitative statement, however, must depend on both the specific properties of the spatial distribution of the neighbours, and the precise form of the adopted SPH weights.

In summary, in this appendix we have shown that when transporting a quantity from a particle to its neighbours, the net transport direction is generally governed by the spatial distribution of the neighbours, in addition to the intrinsic properties of the emitting particle. We emphasize that our observations apply to the particle-to-neighbour transport of arbitrary quantities in *any* particle simulation. We would like to remind the reader that in Sec.4.4 we presented a specifically designed particle-based transport scheme, TRAPHIC, that overcomes the problems of particle-to-neighbour based transport discussed in this appendix, not just statistically, but also on a particle by particle basis. We achieved this by employing cones to confine photons to the solid angle into which they are emitted, making the transport independent of the geometry of the SPH simulation, on the scale of the chosen angular resolution.

