



Universiteit
Leiden
The Netherlands

And there was light : Voronoi-Delaunay radiative transfer and cosmic reionisation

Paardekooper, J.P.

Citation

Paardekooper, J. P. (2010, December 16). *And there was light : Voronoi-Delaunay radiative transfer and cosmic reionisation*. Retrieved from <https://hdl.handle.net/1887/16247>

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/1887/16247>

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

CHAPTER 2

SIMPLEX2: Radiative transfer on an unstructured, dynamic grid

Jan-Pieter Paardekooper, Chael Kruip & Vincent Icke

Part of this chapter consists of work that has been published in
Astronomy & Astrophysics 515, A79, 2010

WE present an improved version of the SIMPLEX method for radiative transfer on an unstructured Delaunay grid. The grid samples the medium through which photons are transported in an optimal way for fast radiative transfer calculations. A scheme of dynamic grid updates is employed to ensure the grid stays optimal when the optical depth changes during a simulation. We have applied a direction conserving transport scheme that correctly transports photons in the optically thin regime, a regime where the original SIMPLEX algorithm lost its accuracy. For the application to large data sets, the method is parallelised for distributed memory machines using MPI.

2.1 INTRODUCTION

In many astrophysical applications radiative processes play an important, and occasionally dominant, role. The interaction between radiation and matter leads to possibly strong feedback effects on the medium caused by heating and cooling, radiation pressure and the change of the ionisation and excitation state of the medium. It is therefore of crucial importance to include these effects in the simulations of hydrodynamic flows. However, radiative transfer is a very complex process due to the high dimensionality of the problem. The specific intensity $I(\boldsymbol{\Omega}, \mathbf{x}, \mathbf{t}, \nu)$ depends on seven variables, and is impossible to solve for in an analytic way in general. With the exception of certain limiting cases where an analytical solution exists, one therefore has to rely on numerical methods to obtain the desired solution.

Several kinds of methods exist for this purpose, all of which have their specific advantages and shortcomings. A relatively straightforward way of solving the radiative transfer equation is the long characteristics method (e.g. Mihalas & Mihalas (1984)), where rays connect each source cell with all other cells and the transfer equation is solved along every ray. Although this method is relatively precise, it is computationally demanding, requiring N^2 interactions for N cells. A solution to this unfortunate scaling is implemented in the short characteristics method (e.g. Kunasz & Auer (1988)), where only neighbouring cells are connected by rays. Not only does this make the method computationally less expensive than the long characteristics method, it is also easier to parallelise. A drawback of this method is that it is known to cause numerical diffusion in the solution. In recent years, hybrid schemes combining the benefits of both approaches have been developed (Rijkhorst et al. 2006; Trac & Cen 2007). Instead of direct integration along the characteristics, Monte Carlo methods use a stochastic approach by sending photon packets along the rays. The properties of the photon packets and their interaction with the medium are determined by sampling the distribution functions of the relevant processes. Moment methods solve the moments of the radiative transfer equation, which allows for a computational speed-up in certain opacity regimes. In these methods there is a trade-off between computation time and numerical diffusion in the solution, depending on what method is used to obtain the closure relation.

What almost all of these methods have in common is that they use a predefined grid on which the radiative transfer calculation is done. Most often this grid is given by a hydrodynamics simulation, which is either a regular, adaptive mesh refinement (AMR) or smoothed particle hydrodynamics (SPH) grid. These grids are not optimised for radiative transfer but for hydrodynamics calculations, possibly resulting in unphysical behaviour of the numerical solution. Moreover, the computation time of almost all methods scales linearly with the number of sources, which severely limits the range of applications. Exceptions are moment methods that generally do not scale with the number of sources, but sacrifice precision by introducing numerical diffusion (e.g. Gnedin & Abel (2001); Cen (2002)) and the method by Pawlik & Schaye (2008), where a source merging procedure is used to avoid linear scaling with the number of sources. For many applications, the linear scaling of the computation time with the number of sources becomes prohibitive, for example when simulating scattering processes, where effectively every cell might become a source. In the case of simulations of the epoch of reionisation, which is the topic of the second half of this paper, it is necessary to include many sources. It is therefore essential to use a method for which the computation time is independent of the number of sources.

The approach to solve the radiative transfer equation taken in this thesis is radically different from the methods described earlier. Instead of using a predefined grid, the `SIMPLEX` method calculates the grid for the radiative transfer problem from the properties of the physical medium through which photons will be transported. This leads to a computationally fast method that does not scale with the number of sources, making it an ideal tool for simulations of the epoch of reionisation. A previous version of the method has been described in Ritzerveld & Icke (2006), where the general idea of transport on random lattices is laid down, with a small section on the application to cosmological reionisation. The first comprehensive set of tests of the method were performed for the Radiative Transfer Comparison Project (Iliev et al. 2006). In this project, the focus lay on comparing the performance of all participating codes in the test problems, making an in-depth analysis of the `SIMPLEX` results impossible.

The aim of this chapter is to describe the improvements of the `SIMPLEX` method since these previous two papers. Recently, we have performed a detailed study of the error properties of the method (Kruip et al. 2010) (KPCI10), which has led to some essential improvements to the method. The two main problems in ballistic transport that were addressed in KPCI10, decollimation and deflection, are minimised both by using an alternative transport scheme in the opacity regime where these problems occur and by adapting the grid to changes in the opacity during a simulation. In addition, the algorithm has been parallelised for distributed memory. In this chapter, we describe the working of the improved `SIMPLEX` method focussing on the new features.

The format of this chapter is as follows. In Sect. 2.2, we give an overview of the `SIMPLEX` method and a description of the new features of the method. We then describe the parallelisation strategy in Sect. 2.3 and present the computational scaling properties of the algorithm. Finally, in Sect. 2.4 we describe the application of the method to ionising radiation with a focus on cosmological radiative transfer.

2.2 THE `SIMPLEX` METHOD

In this section, we describe the basics of the `SIMPLEX` method and specifically the new features that were added to improve the performance in the lower opacity regimes. For the sake of clarity we repeat some essential information that was presented earlier in Ritzerveld & Icke (2006) and Ritzerveld (2007), which is necessary to appreciate the new features of the method. We start with a description of how the unstructured grid is created and how to optimise it for radiative transfer calculations with `SIMPLEX`. We then proceed by describing the different ways of transporting photons on this grid, governed by the physical properties of the problem at hand.

2.2.1 Grid calculation

At the basis of the `SIMPLEX` method lies the unstructured grid on which the photons are transported. The grid adapts to the physical properties of the medium through which the photons travel in such a way that more grid points are placed in regions with a higher opacity. The result is a higher resolution in places where it's needed most, there where the optical depth is highest.




Figure 2.1: Two-dimensional example of a point distribution based on a homogeneous Poisson process (*left*) and based on a non-homogeneous Poisson process (*right*).

Point process

The placement of the grid points is a stochastic process based on the Poisson process, which can be defined as follows. Suppose $N(A)$ is the number of points in a non-empty subset A of the volume $S \subset \mathbb{R}^d$, with d the dimension. Then the probability to that A contains x points is

$$\Phi = P(N(A) = x) = \frac{n_p |A| e^{-n_p |A| x}}{x!}, \quad x = 0, 1, 2, \dots \quad (2.1)$$

The only parameter in this process is the point intensity n_p , which is a global constant. Every region in the volume has the same probability that points are placed there, which in our case corresponds to a constant opacity. An example of a homogeneous Poisson process is shown on the left hand side of Fig. 2.1.

To account for different opacity regimes inside the computational volume, we use the non-homogeneous Poisson process, defined as

$$P(N(A) = x) = \frac{n_p(A) |A| e^{-n_p(A) |A| x}}{x!}, \quad x = 0, 1, 2, \dots \quad (2.2)$$

where

$$n_p(A) = \int_A n_p(\mathbf{x}) d\mathbf{x}. \quad (2.3)$$

The point intensity function $n_p(\mathbf{x})$ follows the opacity of the medium on global scale, while on local scale the point distribution retains the properties of the homogeneous Poisson distribution. An example of a point distribution based on the non-homogeneous Poisson process is shown on the right hand side of Fig. 2.1. An alternative, possibly more physically intuitive, recipe for constructing the non-homogeneous Poisson process can be written as

$$n_p(\mathbf{x}) = \Phi * \mathbf{f}(n(\mathbf{x})), \quad (2.4)$$

that is, by defining the grid point distribution as a convolution of a homogeneous Poisson process and a function of the possibly inhomogeneous medium density distribution $n(\mathbf{x})$. It is this recipe for the non-homogeneous Poisson process that we use to construct the **SIMPLEX** grid. Grid points are placed by randomly sampling the correlation function $f(n(\mathbf{x}))$. We will discuss the exact shape of the correlation function in Sect. 2.2.1.




Figure 2.2: Two-dimensional example of a random point distribution, the Voronoi tessellation around the points, the corresponding Delaunay triangulation connecting the points and the combination of both.

The Delaunay triangulation

The grid points thus created form the nuclei of a Voronoi tessellation. Given a set of nuclei x_i , the Voronoi tessellation (Dirichlet 1850; Voronoi 1908) is defined as $V = \{C_i\}$, in which

$$C_i = \{y \in \mathbb{R}^d : \|x_i - y\| \leq \|x_j - y\| \forall x_j \neq x_i\}. \quad (2.5)$$

In other words, this means that every point inside a Voronoi cell is closer to the nucleus of that cell than to any other nucleus. By joining all nuclei that have a common facet (an edge in 2D, a wall in 3D), we create the Delaunay triangulation (Delaunay 1934). Thus, every nucleus is connected to its closest neighbours. A 2D example of a Voronoi tessellation and the corresponding Delaunay triangulation is shown in Fig. 2.2.

The Delaunay triangulation consists of simplices that fill the entire domain. A simplex is the generalisation of a triangle in \mathbb{R}^d , so a triangle in \mathbb{R}^2 and a tetrahedron in \mathbb{R}^3 . In a valid Delaunay triangulation, all simplices obey the empty circumsphere criterion. The circumsphere of a simplex is the unique sphere that passes through each of the vertices that make up the simplex. In a valid Delaunay triangulation, no vertex exists inside this circumsphere.

For Voronoi tessellations and Delaunay triangulations that are constructed from a point process based on a homogenous Poisson process, so-called Poisson Delaunay triangulations, it is possible to derive some general properties relevant for our radiative transfer method. These results were mainly derived by Miles (1970, 1974) and Møller (1989). Two important properties for our purposes are the average number of neighbours of a vertex and the average distance between two connected vertices. The expectation value for the number of neighbours of a typical vertex in \mathbb{R}^2 and \mathbb{R}^3 is

$$E_{2D}(E) = 6 \quad (2.6)$$

and

$$E_{3D}(E) = \frac{48\pi^2}{35} + 2 \approx 15.54. \quad (2.7)$$

The expectation value for the distance between two connected vertices in \mathbb{R}^2 and \mathbb{R}^3 is

$$E_{2D}(L) = \frac{32}{9\pi} n_p^{-1/2} \approx 1.132 n_p^{-1/2} \quad (2.8)$$

and

$$E_{3D}(L) = \frac{1715}{2304} \left(\frac{3}{4}\right)^{1/3} \pi^{-1/3} n_p^{-1/3} \approx 1.237 n_p^{-1/3}. \quad (2.9)$$

Note that these values are only exact for Delaunay triangulations constructed from a homogeneous Poisson process, while in `SIMPLEX` we use the non-homogeneous Poisson process to place the grid points. Except for regions in the domain with strong gradients in the point density, on local scale the point distribution resembles a homogeneous point distribution quite well. Therefore the properties of the Poisson Delaunay triangulation give a good qualitative idea of the properties of the grid on which we perform our radiative transfer calculations.

`SIMPLEX` is set up in such a way that once the point distribution is created according to Eq. (2.4), the Delaunay triangulation is calculated by an external software package. It is therefore possible to use any package that suits the application at hand. For all simulations presented in this paper, the Delaunay triangulation is calculated using the `QHull` package¹. This is a software package written in C that is able to calculate the Delaunay triangulation, the surfaces and the volumes of the simplices in up to 8 dimensions. `QHull` is based on the Quickhull algorithm (Barber et al. 1995), using the convex hull property of the Delaunay triangulation. `QHull` has the advantages that it computes the Delaunay triangulation in optimal time $\mathcal{O}(N \log N)$, it is very stable against floating point round off errors in case points lie very close to each other and it is easy to implement as modular plugin routine. One of the drawbacks of `QHull` is that it triangulates the entire point set in one call, so it's impossible to add or delete points after the triangulation has been computed. This results in extra computational overhead in the grid dynamics scheme presented in Sect. 2.2.1. However, for the post-processing simulations presented in this thesis the computation time of the triangulation is small compared to the computation time of the radiative transfer (see also Fig. 2.5), so in the present case the extra computational overhead is acceptable.

The correlation function

In the previous discussion we have not specified the exact shape of the correlation function $f(n(\mathbf{x}))$ with which we sample the density distribution of the medium. In order for the grid to adapt to the properties of the medium, the correlation function should be a monotonically increasing function in $n(\mathbf{x})$. Thus, the distance between two connected vertices will be smaller in regions with high density. From basic transfer theory, we know that the local mean free path in a medium relates to the local medium density in the following way:

$$\lambda(\mathbf{x}) = \frac{\mathbf{1}}{\mathbf{n}(\mathbf{x})\sigma}, \quad (2.10)$$

where σ is the total cross section, $\sigma = \sum_i \sigma_j$, consisting of different interaction cross sections σ_j . If we compare this to the expectation value of the Delaunay edge length in Eq. (2.8) and Eq. (2.9) it follows that if we choose the correlation function $f(n(\mathbf{x}))$ to sample the d -th power of the local medium density, e.g. $f(x) \propto x^d$, the local mean free path scales linearly with the expectation value of the Delaunay edge length via a constant c :

$$\langle L_D \rangle(\mathbf{x}) = c\lambda(\mathbf{x}). \quad (2.11)$$

¹www.qhull.org

Equation (2.11) is a global relation with a global constant c , given by

$$c = \xi(d, D, N) \sigma. \quad (2.12)$$

Here, D is the size of the computational domain and N the number of vertices.

This sampling recipe works very well in physical media where the density fluctuations are small. However, if the density fluctuations are significant, sampling the density to the d -th power will favour the high density regions, resulting in a possibly severe undersampling of the low density regions. This undersampling can have serious consequences for the radiative transfer calculation, for instance by causing preferential directions that lead radiation around low density regions. See Sect. 3.2.2 for an example of this effect. Moreover, KPCI10 showed that large gradients in the grid point distribution lead to systematic errors when transporting photons on this grid.

We therefore need a sampling recipe that retains the advantages of the adaptive grid by keeping the dynamic range as large as possible, while maximising the minimum resolution of the grid. This is achieved by defining a reference density $n_0(\mathbf{x})$ above which the density is no longer sampled to the d -th power but a different power α . Both $n_0(\mathbf{x})$ and α depend on the properties of the medium that needs to be sampled. The two sampling recipes are smoothly joined by taking the harmonic mean, resulting in the following sampling function:

$$f(n(\mathbf{x})) = \left(\left(\frac{n(\mathbf{x})}{n_0(\mathbf{x})} \right)^{-d} + \left(\frac{n(\mathbf{x})}{n_0(\mathbf{x})} \right)^{-\alpha} \right)^{-1}. \quad (2.13)$$

This sampling recipe favours low density regions by sampling those with a higher (d -th) power.

By choosing the sampling parameter Q_n and the reference density $n_0(\mathbf{x})$ we can create a grid where the dynamic range is maximal without causing numerical artefacts due to undersampling of the low density regions or large gradients in the point distribution. One has to be careful that by sampling different opacity regimes in a different way, the interaction coefficient of Eq. (2.12) is no longer a global constant, but we can take care of this by the way the interaction at each grid point is accounted for. This will be discussed in Sect. 2.2.2. We will take a closer look at how the grid is created from various (hydrodynamics) density grids and how the sampling parameters $n_0(\mathbf{x})$ and α affect the properties of the resulting point distribution in chapter 3.

A dynamic grid

In the previous section we described how the SIMPLEX grid is created according to the properties of the medium. In this discussion, it was assumed that the medium is static and does not change during the radiative transfer calculation. However, in reality the properties of the medium change continuously under influence of, for example, gravity and radiation. For this reason, the SIMPLEX grid should be updated every time step in case of full radiation hydrodynamics simulations. In this paper we do not consider the application of SIMPLEX to radiation hydrodynamics but instead focus on post-processing static density fields. We will discuss the application of SIMPLEX to radiation hydrodynamics simulations in future work.

Even though the gas density is assumed to be static, the properties of the medium might still change during a radiative transfer calculation. For example, photo-ionisation lowers the optical depth for ionising radiation. Changes in the optical depth lead to a deviation from the

recipe for grid point distribution of Eq. (2.13). In other words, the grid is no longer an optimal representation of the physical properties relevant for the radiative transfer calculation. KPCI10 showed that this might have serious consequences for the transport of photons through regions where the optical depth between grid points is much lower than unity. In Sect. 2.2.2 we describe a new transport scheme that minimises errors in the optically thin regime. This section describes a different solution for transport through regions where the optical depth has severely changed, the removal of grid points.

One of the advantages of the adaptive grid that `SIMPLEX` uses is that resolution is put where it is needed most, in the regions with highest optical depth. If during a radiative transfer calculation the optical depth changes, we are effectively wasting computational resources in the regions with high resolution where the opacity has decreased. The high resolution is no longer necessary, since no interesting radiative transfer effects that need to be resolved at high resolution are taking place. The superfluous grid points only add to the computation time. We therefore remove unnecessary grid points from the regions where the optical depth has significantly decreased.

Another reason for the removal of superfluous grid points is that the transport of photons between grid points with an optical depth lower than unity is prone to numerical errors. KPCI10 showed that photons that travel through these regions are subject to decollimation and deflection. These effects are caused by the grid that is no longer an optimal representation of the properties of the medium. A straightforward solution for these problems is updating the grid in such way that the physical properties of the medium remain correctly accounted for. By ensuring that the optical depth between grid points is always close to unity, the grid remains optimally suited for the radiative transfer calculation.

In some cases the removal of grid points according to this scheme leads to regions devoid of grid points. An example is the photo-ionisation of a cloud of neutral hydrogen gas, which causes the optical depth for ionising photons to drop so dramatically that almost no grid points should be placed if the optical depth between grid points has to be of order unity. This extreme example leads to different errors in the solution than the ones previously described. For example, recombination rates will be incorrect as the density in the cloud is no longer resolved by any grid point. We circumvent this by imposing a minimum resolution below which no grid points are removed. This ensures that in every part of the simulation relevant structures are resolved and our requirements for accuracy are met. The effect of grid point removal and the optimal value for the minimum resolution in realistic applications will be further explored in Chapter 5.

The consequence of preventing undersampling errors is that there will remain grid points in the simulation domain between which the optical depth is lower than unity. Numerical errors in the transport of photons between these grid points are prevented by using the transport scheme described in Sect. 2.2.2. As we will show in that section, this transport scheme is more expensive both in computation time as in memory usage compared to the other transport schemes. For optimal computation time and memory usage it is therefore important to keep the number of superfluous grid points in regions with low opacity to a minimum.

2.2.2 Radiation Transport

We have shown how the unstructured grid is created on which the radiative transfer calculation will be performed. In this section we will show how we can employ the unique properties of this grid to efficiently and accurately solve the radiative transfer equation.

During a radiative transfer calculation photons are transported from grid point to grid point along the edges of the Delaunay triangulation. At every grid point an interaction takes place, with the interaction coefficient given by Eq. (2.12). According to the solution of the 1-dimensional radiative transfer equation the number of photons that interacts at this grid point is

$$N_{act} = N_{in}(1 - e^{-c}), \quad (2.14)$$

where N_{in} is the number of incoming photons. The number of photons that is not interacting at the grid point is

$$N_{out} = N_{in}e^{-c}, \quad (2.15)$$

these photons should continue along their original path.

In this photon transport scheme there is no difference between a grid point that is a source and a grid point that is not. A source is simply defined as a grid point that sends photons to all its neighbours, which has no influence on the number of computations involved. Thus, the SIMPLEX method does not scale with the number of sources.

After interaction at a grid point there are three ways of photon transport, depending on the opacity regime and the physical process at hand.

Scattering processes

In the case of an isotropic scattering process, or the absorption and re-emission of photons at a grid point, the outgoing photons have no memory of their original direction. In SIMPLEX these photons are isotropically redistributed to all neighbouring vertices, as depicted on the left hand side of Fig. 2.3. This type of transport does not increase the number of computations involved, SIMPLEX is therefore ideally suited to simulate scattering processes.

For the application to ionising radiation it is straightforward to include the diffuse radiation from recombining atoms in the simulation. Hydrogen ions that capture an electron directly to the $n = 1$ state emit photons capable of ionising other atoms. Most radiative transfer methods do not include this radiation but instead adopt the on-the-spot approximation (e.g. Osterbrock & Ferland (2006)), assuming these photons to be absorbed close to the emitting atom (see also Sect. 2.4.2). Even though the validity of this approach is not well established (Ritzerveld 2005) we will use the on-the-spot approximation for all the tests presented in this paper, in order to make a clean comparison between our results and the analytic solution or the results of other codes that all use the on-the-spot approximation.

Ballistic transport

The N_{out} photons in Eq. 2.15 that are not interacting at a grid point should continue travelling in their original direction. However, on the unstructured grid there is no outgoing Delaunay edge in the same direction as the incoming edge. This is solved by splitting the photon packets




Figure 2.3: Two-dimensional examples of the three modes of transport with which photons are transported from one grid point to another. Open arrows indicate incoming photons, solid arrows outgoing photons. Transport in case of a scattering process is shown on the left hand side. Incoming radiation loses all memory of the initial direction and is sent to all neighbours of the vertex. Ballistic transport is shown in the centre plot. The photons are redistributed to the 2 most straightforward neighbours of the vertex, with respect to the Delaunay edge of the incoming photons. On the right hand side direction conserving transport is shown. The photons are redistributed to the 2 most straightforward neighbours with respect to the Delaunay edge that is associated with the direction bin the incoming photons are in. The outgoing photons stay in the same direction bin and have thus a memory of their original direction.

into p equal parts and dividing these packets among the p most straightforward neighbours with respect to their incoming direction. The total number of photons is conserved in this process. Tests indicate that if we choose p equal to the dimension of the problem d , the solid angle that is represented by one Delaunay edge of the emitting vertex is best preserved. In other words, a source that sends out photons in all directions will fill the entire 'sky' with radiation. However, on the unstructured grid it might be the case that one of the most straightforward neighbours lies outside a cone of 90 degrees. To prevent photons from travelling backwards, we exclude those neighbours. In Fig. 2.3 the centre plot shows an example of ballistic transport.

The advantage of this transport scheme is that the most straightforward directions have to be calculated only once, at the start of the simulation. As long as the grid doesn't change these directions do not have to be recalculated. One important disadvantage of this transport scheme is that there is no memory of the original direction of the photons. At every interaction the outgoing direction was computed from the incoming direction in that step, so deflections from the original direction can add up, causing numerical diffusion to dominate after approximately 5 interactions (KPCI10). As long as the mean free path of photons is smaller than 5 Delaunay edges, this numerical diffusion has no influence on the results, since photons will be interacting with the medium before the diffusion becomes dominant. This means that during the grid calculation we have to be careful that the interaction coefficient c in Eq. (2.12) is close to one or larger. However, this may lead to too severe constraints on the number of grid points that can be placed in optically thin regions. Therefore, a different type of transport can be employed in optically thin media.




Figure 2.4: Example of the photon path deviating from a straight line due to the unstructured grid. Photons that are travelling in the direction of the Delaunay edge coming in from the left, should be travelling in a straight line along the dotted grey line. However, as this is impossible on the unstructured grid, photons travel along the Delaunay edges closest to their original direction, depicted by the open arrows. The total path depicted by the open arrows is longer than the length of the dotted line, which is corrected for by a global factor.

Direction conserving transport

If the interaction coefficient c in Eq. (2.12) is smaller than one, it is no longer sufficient to determine the direction of the photons from the direction in the previous step, but a memory of the initial direction of the photon is needed. If every photon would remember its initial direction and at every interaction point the next interaction point would have to be calculated from this direction, the computation time in optically thin regions would grow unacceptably. Instead, the original direction is preserved by confining photons to solid angles corresponding to global directions in space. Unless interacting with the grid, photons stay in the same solid angle as they travel along the grid. Even though the direction of the photons is now decoupled from the grid, the photons still travel along the edges of the triangulation in the same manner as during ballistic transport. Direction conserving transport is shown on the right hand side of Fig. 2.3.

Since photons still travel along the edges of the triangulation, the photon path deviates from an exact straight line in which the photons should be travelling, see Fig. 2.4. Even though the direction of the photons is preserved, their paths are longer than physically possible. In other words, photons travel slower than the light speed on the unstructured grid. We solve this problem by applying a global correction factor to the distance between grid points, thus ensuring photons travel with the correct speed.

Introducing these global directions on the unstructured grid gives rise to preferential directions, one of the problems the unstructured grid was meant to solve. By rotating the solid angles over random angles in between photon transport, the preferential directions disappear. The drawback of this procedure is that it makes direction conserving transport computationally more expensive than ballistic transport. For the latter, the most straightforward directions are calculated from the grid and thus have to be calculated only once, at the start of the simulation. For direction conserving transport it is necessary to recalculate the most straightforward directions every time the direction bins rotate. Another drawback of direction conserving transport is

that the photons now have to be stored in n direction bins instead of on average 16 neighbours. Typically, $n = 42$ gives converged results, but this depends on the number of optically thin grid points the photons traversed. Thus, the memory requirements for direction conserving transport are higher.

Combined transport

The three modes of transport described above are in general applied simultaneously during a simulation. Depending on the physical process at hand, photons are transported to all neighbours (diffuse transport), or to the d most straightforward neighbours (ballistic or direction conserving transport). In regions where the optical depth is higher than or close to one, ballistic transport is used, while in the optically thin regions direction conserving transport is applied.

Of the three modes of transport, direction conserving transport is computationally the most expensive (see Sect. 2.3.3 for a comparison between the computation time of ballistic and direction conserving transport). By applying this scheme only in the regions where it is necessary, the computation time is drastically reduced. As mentioned earlier, numerical diffusion starts to dominate in ballistic transport after approximately 5 steps. A first guess would therefore be to switch from ballistic to direction conserving transport when the optical depth after 5 interactions is less than one. That way, we are sure that the majority of photons does not take more than 5 steps in ballistic transport. The influence of the optical depth at which is switched in a realistic simulation is studied in more detail in Chapter 5. Another way to reduce the computation time is by applying the grid dynamics scheme from Sect. 2.2.1. Removing superfluous grid points in the low opacity regime limits the number of vertices at which direction conserving transport is performed.

In combined transport we need to convert from one transport scheme to another. This is straightforward because every Delaunay edge of an optically thin vertex is associated with a solid angle in a global direction. When this vertex sends photons to an optically thick vertex the photons are transported along the Delaunay edges, so the optically thick vertex stores the photons according to the Delaunay edge associated with the solid angle. In the opposite case, when an optically thick vertex sends photons to an optically thin vertex, the photons are converted to the solid angles associated with the Delaunay edge along which the photons were sent.

2.3 PARALLELISATION STRATEGY

Even though the radiative transfer scheme presented in the previous sections is computationally efficient, in order to do large simulations involving a very high number of grid points it is necessary that the algorithm can run in parallel on distributed memory machines. This will not only reduce the computation time involved, it also reduces the amount of memory needed at each processor to store the physical properties of the grid points. The transport algorithm described in Sect. 2.2.2 has the advantage that it is local: the only information needed to do a radiative transfer calculation is stored at the neighbours of the vertex. This makes the method relatively easy to parallelise. By choosing a smart domain decomposition we can minimise the number of communications involved.

2.3.1 Domain decomposition

The computation time of a SIMPLEX calculation is independent of the number of sources, it is therefore sufficient to have a domain decomposition that assigns every processor an approximately equal number of grid points. Dividing space into equal volumes and assigning each volume to a processor is not sufficient because the number of points in each volume may differ dramatically due to the adaptive grid. We therefore use a domain decomposition based on the space-filling Hilbert curve, which is also employed in other methods without a regular grid (Shirokov & Bertschinger 2005; Springel 2005, 2010). The Hilbert curve is a fractal that completely fills a cubic rectangular volume. A Hilbert curve is uniquely defined by its order m and its dimensionality d , filling every cell of a d -dimensional cube of length 2^m . The following properties of the Hilbert curve are beneficial when using it for domain decomposition:

- **Locality:** points that are close along the 1D Hilbert curve are in general also close in 3D space.
- **Compactness:** a set of cells defined by a continuous section of the Hilbert curve has a small surface to volume ratio.
- **Self-similarity:** the Hilbert curve can be extended to arbitrarily large size.

The first two properties minimise the number of communications between processors, while the third property ensures that we can use an arbitrarily large number of cells to determine the domain decomposition.

The first step in the domain decomposition is dividing the domain into 2^{md} equal, regular cells, where d is again the dimension and m the order of the Hilbert curve. We then step through the cells along the Hilbert curve, counting the number of grid points inside each cell until the number of grid points equals the total number of grid points divided by the number of processors. In this way, every processor approximately holds an equal number of grid points, thus dividing the work load evenly, while the necessary communications between processors are minimal due to the locality and compactness property of the Hilbert curve.

2.3.2 Parallel radiative transfer

The QHull algorithm that is applied to calculate the triangulation works only in serial. In the case of parallel execution of SIMPLEX, every processor calculates the triangulation of the vertices that belong to that processor and vertices in a border around it belonging to neighbouring processes. This border is used to connect the triangulation between different processes. We ensure that the border is large enough by using the empty circumsphere principle. For all simplices that contain at least one vertex inside the domain of the current process, we ensure that the circumsphere of the simplex lies entirely within the boundary around the domain. Thus, we are certain that no vertex exists on another process that lies inside the circumsphere of this simplex and the triangulation is valid. After the triangulation algorithm has been applied, every process keeps only the vertices assigned to the process and a local copy of vertices assigned to other processes that are neighbour to a vertex on this process. These local copies are strictly used to send photons from one process to another, no physical interaction is taking place.

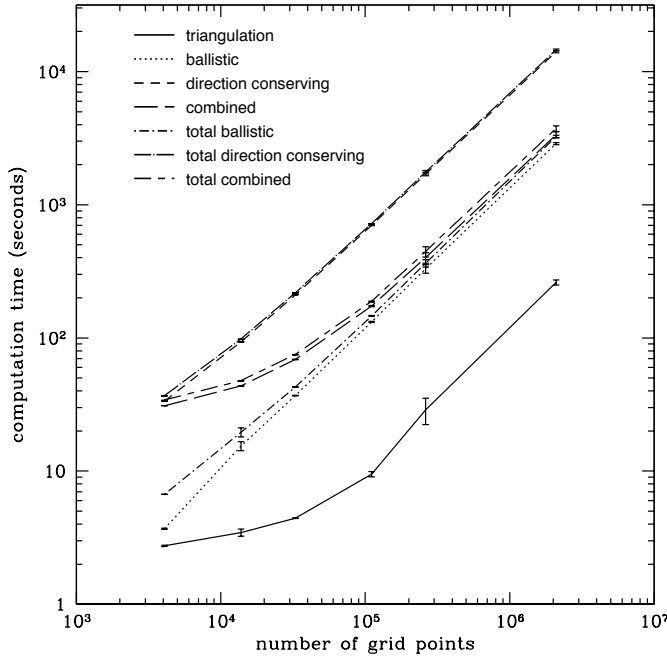


Figure 2.5: Simulation time as a function of the number of grid points. Shown are the computation time of the triangulation (solid curve), ballistic radiation transport (dotted), combined radiation transport (long dashed), direction conserving transport (short dashed), total simulation time with ballistic transport (dot-short dashed), total simulation time with combined transport (short dash-long dashed) and total simulation time with direction conserving transport (dot-long dashed).

The transport scheme described in Sect. 2.2.2 is local, because photons are only transported from one grid point to another. We can therefore do a full radiative transfer time step without any communication between processes. During a time step, photons might be sent to a neighbour of a vertex that does not exist on the current process. After each time step, these photons are then communicated to the appropriate processes and the cycle starts anew. The local copies of vertices are only used for the transport of photons, all the physical interactions are taking place on the process that the vertex is assigned to. Hence, we are certain that physical interactions take place exactly once for every vertex during a radiative transfer time step.

2.3.3 Scaling tests

To check how well the described parallelisation strategy works, we have conducted several scaling tests. The test consists of a simulation of a single source in a homogeneous medium, similar to the set-up of the first test in Chapter 5 with the difference that the simulation is stopped at 10 Myr. No grid points were removed during these tests. All the simulations were conducted using AMD Opteron 246 64Bit CPUs of 2.6 Ghz with 4 GB of memory per node. Unfortunately, we had only 8 nodes available for these tests, but it will give a general idea of the scaling. Note that even though we have used only one source for these scaling tests, the inclusion of more sources would not have influenced the timings presented here.

As a first test we used only one node with an increasing number of grid points, shown in Fig. 2.5. This figure shows that an increase of the number of grid points N increases the computation time linearly for most components of the simulation. Exceptions are the triangulation algorithm, which scales $O(N \log N)$ and the combined transport scheme. The computation time of the combined scheme will always be between the computation time of ballistic transport and that




Figure 2.6: Simulation time as a function of the number of processors for a constant number of grid points. The number of grid points is $128^3 = 2097152$. Shown are the computation time of the triangulation (dashed curve), ballistic radiation transport (dotted), combined radiation transport (long dashed), direction conserving transport (short dashed), total simulation time with ballistic transport (dot-short dashed), total simulation time with combined transport (short dash-long dashed) and total simulation time with direction conserving transport (dot-long dashed).

of direction conserving transport. It depends highly on the number of optically thin grid points. For the low resolution simulations, the number of optically thin grid points is relatively high and therefore the computation time is comparable to the computation time of direction conserving transport. Increasing the number of grid points decreases the relative number of optically thin grid points and therefore the computation time of combined transport comes closer to that of ballistic transport in the case of more grid points. Note that this is a feature of the set-up that we chose for the scaling test. If a larger region of the computational volume would be ionised, the computation time would be longer. On the other hand, if the ionised region were smaller, the computation would be shorter. Fig. 2.5 also shows that the computation time of the entire simulation is dominated by the radiation transport and not by the construction of the triangulation, since the computation time of the triangulation is one order of magnitude shorter than that of the radiation transport.

Fig. 2.6 shows the strong scaling properties of the SIMPLEX algorithm. We simulate the same physical problem as the previous test, but this time the number of grid points is held constant at $128^3 = 2097152$. By increasing the number of processors we can analyse the extra work that needs to be done when more processors are employed. In the ideal case no extra work would have to be done at all, this would result in a horizontal line in the figure. The only component of the simulation that does not follow the linear scaling very well is the triangulation, which can be easily understood by the way the triangulation is constructed on multiple processors. Every processor has to calculate the triangulation of the grid points assigned to the processor and an additional number of grid points in a boundary around this domain. Increasing the number of processors thus means effectively increasing the number of grid points that needs to be triangulated and therefore the scaling is not as favourable as one might hope. However, the computation time of the triangulation remains an order of magnitude smaller than the radiative




Figure 2.7: Simulation time as a function of the number of processors for a constant number of grid points per processor. The number of grid points at each processor is $64^3 = 262144$. Shown are the computation time of the triangulation (solid curve), ballistic radiation transport (dotted), combined radiation transport (long dashed), direction conserving transport (short dashed), total simulation time with ballistic transport (dot-short dashed), total simulation time with combined transport (short dash-long dashed) and total simulation time with direction conserving transport (dot-long dashed).

transfer components that do scale almost linearly, so this presents no serious issue.

Finally, Fig. 2.7 shows the weak scaling of the algorithm. With the same problem set-up as before, we now increase the number of processors while keeping the number of grid points per processor constant. Ideally, the amount of work per processor would stay the same and the computation time would remain constant. Within a few percent the computation time for the radiation transport components remains constant, it only increases marginally due to an increase in the amount of communication necessary. As with the strong scaling, the triangulation algorithm needs to triangulate more boundary points as the number of processors increases and therefore the computation time increases with number of processors. However, the number of boundary points is small compared to the total number of grid points that needs to be triangulated, therefore the curve of the computation time flattens with an increase of the number of processors.

2.4 RADIATIVE TRANSFER OF IONISING RADIATION

Having laid down the basics of the transport mechanism of SIMPLEX, in this section we will describe the application to ionising radiation. Generally, the input for the radiative transfer calculation will be given by a hydrodynamics simulation that can be in the form of a regular, AMR or SPH grid, from which the SIMPLEX grid is calculated according to the recipe described in Sect. 2.2.1. Photons are then transported on this grid as described in Sect. 2.2.2. Our aim is to apply SIMPLEX to cosmological applications, we therefore start this section with the cosmological radiative transfer equation. However, the method can be used for all applications that require the transport of ionising photons.

2.4.1 Cosmological radiative transfer equation

The cosmological radiative transfer equation reads (e.g. Gnedin & Ostriker (1997); Norman et al. (1998))

$$\frac{1}{c} \frac{\partial I_\nu}{\partial t} + \frac{\hat{\mathbf{n}} \cdot \nabla \mathbf{I}_\nu}{\bar{a}} - \frac{H(t)}{c} \left(\nu \frac{dI_\nu}{d\nu} - 3I_\nu \right) = j_\nu - \alpha_\nu I_\nu, \quad (2.16)$$

where $I_\nu = I(\hat{\mathbf{n}}, \mathbf{x}, \mathbf{t}, \nu)$ is the specific intensity with frequency ν along the unit propagation direction vector $\hat{\mathbf{n}}$ at position \mathbf{x} and time t , $j_\nu(\mathbf{x}, \hat{\mathbf{n}})$ is the emission coefficient, and $\alpha_\nu(\mathbf{x}, \hat{\mathbf{n}})$ is the extinction coefficient. The latter includes the scattering coefficient $\alpha_\nu^{scat}(\mathbf{x}, \hat{\mathbf{n}})$ and the pure absorption coefficient $\alpha_\nu^{abs}(\mathbf{x}, \hat{\mathbf{n}})$. Furthermore, $H(t) \equiv \dot{a}/a$ is the time-dependent Hubble parameter, and $\bar{a} = \frac{1+z_{em}}{1+z}$ is the ratio of cosmic scale factors between photon emission and present time t . If the mean free path of photons is much smaller than the horizon, the classical transfer equation is a valid approximation:

$$\frac{1}{c} \frac{\partial I_\nu}{\partial t} + \hat{\mathbf{n}} \cdot \nabla \mathbf{I}_\nu = \mathbf{j}_\nu(\mathbf{x}, \hat{\mathbf{n}}) - \alpha_\nu(\mathbf{x}, \hat{\mathbf{n}}) \mathbf{I}_\nu. \quad (2.17)$$

This approximation holds fairly well during the beginning of reionisation. However, care must be taken when ionised bubbles start to overlap and the photon mean free path increases dramatically. In this case the expansion of the Universe does have to be taken into account.

One more simplification can be made if j_ν and α_ν change on time scales larger than the light crossing time in the simulation volume. In this case the time-dependence can be dropped:

$$\boldsymbol{\Omega} \cdot \nabla \mathbf{I}_\nu = \mathbf{j}_\nu(\mathbf{x}, \hat{\mathbf{n}}) - \alpha_\nu(\mathbf{x}, \hat{\mathbf{n}}) \mathbf{I}_\nu. \quad (2.18)$$

It is this equation that SIMPLEX solves. In most astrophysical fluid flows this approximation holds fairly well as long as time scales over which the radiative transfer equation is solved are sufficiently short. However, one must take care that this equation implicitly assumes an infinite speed of light, which might result in unphysical behaviour. For example, as was pointed out by Abel et al. (1999) close to ionising sources an ionisation front might travel faster than the speed of light. In the tests presented in this paper, we have checked the validity of this approximation.

2.4.2 Ionisation processes

Photons with frequencies above the Lyman limit ν_0 can be absorbed by neutral hydrogen atoms. The number of photoionisations per unit time per hydrogen atom is (Osterbrock & Ferland 2006):

$$\Gamma_{\text{H}} = \int_{\nu_0}^{\infty} \frac{4\pi J_\nu}{h\nu} \sigma_{\text{H}}(\nu) d\nu, \quad (2.19)$$

where J_ν is the mean intensity and $\sigma_{\text{H}}(\nu)$ is the ionisation cross section for H by photons with energy $h\nu > h\nu_0$. This cross section can be approximated by

$$\sigma_{\text{H}}(\nu) = A_0 \left(\frac{\nu_0}{\nu} \right)^3 \quad \forall \nu \geq \nu_0, \quad (2.20)$$

with reference value $A_0 = \sigma_{\text{H}}(\nu_0) = 6.3 \cdot 10^{-18} \text{ cm}^2$. The total ionisation rate is the photoionisation rate Γ_{H} plus the collisional ionisation rate Γ_{C} , given by

$$\Gamma_{\text{C}} = n_{\text{e}}C_{\text{H}}(T), \quad (2.21)$$

where n_{e} is the electron density and $C_{\text{H}}(T)$ the collisional ionisation coefficient. The inverse process is the recombination of electrons with the ions. The number of recombinations per hydrogen atom per unit time is:

$$R_{\text{H}} = n_{\text{e}}\alpha_{\text{H}}(T), \quad (2.22)$$

where $\alpha_{\text{H}}(T)$ is the recombination coefficient of hydrogen atoms. The evolution of the number density of ionised atoms at a certain point in space can then be written as

$$\frac{dn_{\text{HII}}}{dt} = n_{\text{HI}}\Gamma - n_{\text{HII}}R_{\text{H}}, \quad (2.23)$$

where n_{HII} , n_{HI} and n_{H} are respectively the number density of ionised hydrogen atoms, neutral hydrogen atoms and the total number of atoms and $\Gamma = \Gamma_{\text{H}} + \Gamma_{\text{C}}$ is the total ionisation rate. Relevant time scales for these processes are the photoionisation time scale $t_{\text{ion}} \equiv 1/\Gamma$ and the recombination time scale $t_{\text{rec}} \equiv 1/R$.

When a hydrogen ion captures an electron directly to the ground level, radiation is emitted with a frequency above the Lyman limit. In almost all radiative transfer codes, it is assumed that this radiation is absorbed close to the emitting atom, the so-called on-the-spot approximation. One can therefore ignore recombinations to the ground level, as they are cancelled out by the emitted radiation, and use the corresponding recombination coefficient $\alpha_{\text{B}}(T)$. However, Ritzerveld (2005) showed that, depending on the density distribution, if the source produces radiation just above the Lyman limit, this approximation is not valid. Most radiative transfer codes are incapable of including this diffuse recombination radiation since effectively every grid cell that is ionised becomes a source. In the SIMPLEX algorithm it is straightforward to include the diffuse recombination radiation self-consistently. However, the analytic solutions and the results from other codes with which we will be comparing the SIMPLEX results in Chapter 5 all rely on the on-the-spot approximation. Therefore, we will use the on-the-spot approximation for all tests presented there. We will study the validity of this assumption in Chapter 7.

2.4.3 Assigning sources

On the SIMPLEX grid, a source is defined as a grid point that sends photons to all of its neighbours. If we assume the source to emit monochromatic radiation we can use a single frequency bin to transport the photons in. In some applications in Chapters 5 and 7 we will use this approximation. However, for more realistic modelling we need to take the spectrum of the sources into account. In SIMPLEX this can be done in two ways, using a single frequency bin or using multiple frequency bins. In all cases is the number of photons that a source sends out at every step is determined by the source luminosity and the radiative transfer time step Δt_{rt} .

Single frequency bin

If a single frequency bin is used the luminosity of the source is obtained by integrating the source spectrum S_{ν} divided by the energy of each photon over frequencies higher than the Lyman limit

ν_0 :

$$L_{ion} = \int_{\nu_0}^{\infty} \frac{S_{\nu}}{h\nu} d\nu. \quad (2.24)$$

By integrating the source spectrum over frequency, we neglect the influence of the spectrum on the interaction of the photons with the medium. We compensate for this by using a mean opacity representation (see Mihalas & Mihalas (1984)):

$$\bar{\sigma} = L_{ion}^{-1} \int_0^{\infty} d\nu \frac{S_{\nu}}{h\nu} \sigma_{\nu}. \quad (2.25)$$

This mean opacity gives a reasonable approximation to the total number of absorptions in the gas, which is the focus of reionisation simulations. How well this mean opacity approximates the number of absorptions in a cell depends on how well the spectrum of the radiation that travels through that cell is represented by the source spectrum. In order to fully account for the spectrum of the source more frequency bins are necessary.

The drawback of the mean opacity of Eq. (2.25) is that the influence of the source spectrum on the energy of the photons is not well approximated. This may result in incorrect heating rates when one frequency bin is used, leading to inaccuracies in the temperature estimation. It is therefore crucial to use multiple frequency bins in calculations of the heating of the gas.

Multiple frequency bins

To get a better representation of the source spectrum the radiation should be sent in separate frequency bins. The width of these bins is a matter of choice and depends on the problem that needs to be solved. Choices include bins of equal width, logarithmically spaced bins, bins in which the ionisation rate is approximately constant and bins in which the energy of the radiation field is divided equally among bins. Tests indicate that the latter choice gives the fastest convergence in problems with a black body source for different numbers of frequency bins. In all the multi-frequency simulations described in this thesis we therefore use the energy-weighted frequency bins. A more detailed description of the multi-frequency implementation of SIMPLEX will be provided in a forthcoming paper (Kruip et al., in preparation). Fig. 2.8 shows the width of the frequency bins for different numbers of bins. Similar to the case of a single frequency bin we use a mean opacity representation to get the correct number of absorptions inside every frequency bin. In this case the integral in Eq. (2.25) is over the frequency range of the bin.

In Sect. 2.2.1 we described that the grid points are placed in such a way that the distance between grid points scales with the optical depth. In case of multiple frequency bins the optical depth between grid points is different in every bin. Formally one would therefore need a separate grid for every frequency bin, which would dramatically increase the memory and computational costs of the method. In practise it turns out that separate grids are needed only for frequency bins in which the optical depth differs by orders of magnitude, for example if very hard photons are present. Furthermore, the direction conserving transport makes sure that photons are transported correctly even when the grid is not optimally adapted to this frequency range. The implementation of the multi-grid SIMPLEX will be described in future work (Kruip et al., in preparation).




Figure 2.8: The width of energy-weighted frequency bins for different numbers of frequency bins. The sources are assumed to be black body sources of 10^5 K, this spectrum is depicted by the solid line. The dotted line shows the hydrogen cross section as function of frequency. Top panel shows 2 frequency bins, middle panel 5 bins and the bottom panel 10 bins.

2.4.4 Interaction

Photons that were sent out by a source travel in one radiative transfer time step from grid point to grid point, where an interaction with the medium takes place. The photons travel a distance ΔL between grid points, and thus encounter an optical depth

$$\Delta\tau(\nu) = n_{\text{HI}}\bar{\sigma}(\nu)\Delta L = (1 - x)n_{\text{H}}\bar{\sigma}(\nu)\Delta L, \quad (2.26)$$

in which x is the ionised fraction inside the Voronoi cell through which the photon travels. Except for the ionised fraction, all these quantities have been calculated during the creation of the grid. Eq. (2.26) is equivalent to the interaction coefficient in Eq. (2.12). Thus, if the incoming number of ionising photons is N_{in} , the number of ionising photons that is absorbed at this grid point is

$$N_{abs} = N_{in}(1 - e^{-\Delta\tau(\nu)}), \quad (2.27)$$

and the number of ionising photons that is propagating onwards is

$$N_{out} = N_{in}e^{-\Delta\tau(\nu)}. \quad (2.28)$$

In case of multi-frequency transport, this is done for every frequency bin separately. The photons that are absorbed ionise the medium, thereby changing the local ionised fraction. As the medium gets ionised, the optical depth at this grid point changes, which means we should either use the direction preserving scheme at this grid point, or remove grid points to preserve the relation between optical depth and Delaunay line length, as described in Sect. 2.2.2 and Sect. 2.2.1.

2.4.5 Solving the photoionisation rate equation

Having established the number of photons that is available for ionising the cell, it is straightforward to convert this to a photoionisation rate and solve Eq. (2.23). However, care must be taken since by doing this we implicitly assume that the neutral density stays constant during a radiative transfer time step. This is only true for $\Delta t_{\text{rt}} \ll t_{\text{ion}}$ and $\Delta t_{\text{rt}} \ll t_{\text{rec}}$. We therefore adopt the scheme described in Pawlik & Schaye (2008) and subcycle the rate equation on time steps $\Delta t_{\text{chem}} \leq \Delta t_{\text{rt}}$ assuming that the total ionising flux is constant during a radiative transfer time step. This ensures photon conservation even if the radiative transfer time steps are large. The rate equation at time $t_{\text{chem}} \in (t_{\text{rt}}, t_{\text{rt}} + \Delta t_{\text{rt}})$ is then

$$dn_{\text{HII}}^{(t_{\text{chem}})} = n_{\text{HI}}^{(t_{\text{chem}})} \Gamma_{\text{H}}^{(t_{\text{chem}})} \Delta t_{\text{chem}} - n_{\text{e}}^{(t_{\text{chem}})} n_{\text{HII}}^{(t_{\text{chem}})} \alpha_{\text{H}}(T) \Delta t_{\text{chem}}, \quad (2.29)$$

where the photoionisation rate at t_{chem} is given by

$$\Gamma_{\text{H}}^{(t_{\text{chem}})} = \Gamma_{\text{H}} \left(\frac{1 - e^{-\tau^{(t_{\text{chem}})}}}{1 - e^{-\tau}} \right) \frac{n_{\text{HI}}}{n_{\text{HI}}^{(t_{\text{chem}})}}, \quad (2.30)$$

where Γ_{H} and τ are the photoionisation rate and optical depth at the beginning of the subcycling and $\tau^{(t_{\text{chem}})} = \tau n_{\text{HI}}^{(t_{\text{chem}})} / n_{\text{HI}}$. By defining the photoionisation rate in this way, the ionising flux in the cell is constant during the radiative transfer time step. This subcycling scheme becomes computationally expensive when $\Delta t_{\text{chem}} \ll \Delta t_{\text{rt}}$, but photoionisation equilibrium is generally reached after a few subcycles. It is then no longer necessary to explicitly integrate the rate equation, but instead use the values of the preceding subcycle step. This way of subcycling ensures photon conservation even for large radiative transfer time steps. For clarity we have omitted the frequency dependence in eqs. (2.29) and (2.30), the procedure can be trivially extended to multiple frequencies by treating every bin separately.

2.4.6 Time stepping

As we discussed in Sect. 2.4.1 we are primarily interested in solving the time-independent radiative transfer equation, which means that the speed of light is assumed to be infinite. In all the tests presented in this paper, photons are moved only one Delaunay edge during a radiative transfer time step, thus interacting only at one grid point in that time step. We therefore have to be careful that the radiative transfer time step Δt_{rt} is sufficiently small to satisfy the time-independent transfer equation. In the limit that Δt_{rt} goes to zero, the time it takes for photons to leave the simulation box goes to zero as well, satisfying the condition of infinite speed of light. For all tests presented in this paper, we have checked that the radiative time step is sufficiently small to be in agreement with this limit.

The subcycling scheme that is used to calculate the evolution of the neutral fraction at a grid point during a radiative transfer time step allows for much larger time steps than needed to satisfy the time-independent transfer equation. This is very useful in simulations where the photons are allowed to travel more than one Delaunay edge per time step, for example in case one needs to solve the time-dependent transfer equation. However, this was not done for the work presented in this thesis.

2.4.7 Solving for the temperature state of the gas

During a radiative transfer simulation the internal energy of the gas changes by heating and cooling processes. Especially when the radiative transfer is directly coupled to the gas dynamics it is crucial to determine the exact temperature state of the gas, to properly account for the energy budget. The change in internal energy per unit mass of a gas parcel is (in case the volume in which the gas resides does not change, which is always the case during a radiative transfer time step)

$$\frac{du}{dt} = \frac{n_{\text{H}}^2}{\rho} (\mathcal{H} + C), \quad (2.31)$$

where \mathcal{H} is the normalised radiative heating function and C is the normalised radiative cooling function. These functions are defined such that $n_{\text{H}}^2 \mathcal{H}$ gives the rate of energy gain per unit volume and $n_{\text{H}}^2 C$ gives the rate of energy loss per unit volume. For an ideal gas the internal energy is related to the temperature by

$$u = \frac{\frac{3}{2} k_{\text{B}} T}{\mu m_{\text{H}}}, \quad (2.32)$$

where k_{B} is Boltzmann's constant, m_{H} is the mass of the hydrogen atom and $\mu \equiv \bar{m}/m_{\text{H}}$ is the mean molecular weight, with \bar{m} the average mass of a gas particle. For a pure hydrogen gas μ depends only on the ionised fraction and can be written as $\mu = 1/(1+x)$. In the following we will discuss the various physical processes that contribute to the normalised heating and cooling functions. We will discuss the exact rates that we have used in chapter 5.

Heating processes

The normalised heating function \mathcal{H} is determined by the sum over the heating rates of individual processes that contribute to the heating. In this work we only consider photo-ionisation heating, which for a pure hydrogen gas is:

$$\mathcal{H} = h_{\gamma, \text{H}} = \frac{n_{\text{HI}}}{n_{\text{H}}^2} \epsilon_{\gamma, \text{H}}, \quad (2.33)$$

where $\epsilon_{\gamma, \text{H}}$ is the photoheating coefficient, given by

$$\epsilon_{\gamma, \text{H}} = \int_{\nu_0}^{\infty} \frac{4\pi J_{\nu} \sigma_{\text{H}}(\nu)(h\nu - h\nu_0)}{h\nu} d\nu = \Gamma_{\text{H}} \langle \mathcal{E}_{\text{H}} \rangle. \quad (2.34)$$

Here we have used Eq. (2.19) to define the average excess energy of the ionising photons

$$\langle \mathcal{E}_{\text{H}} \rangle = \left(\int_{\nu_0}^{\infty} \frac{4\pi J_{\nu} \sigma_{\text{H}}(\nu)(h\nu - h\nu_0)}{h\nu} d\nu \right) \left(\int_{\nu_0}^{\infty} \frac{4\pi J_{\nu}}{h\nu} \sigma_{\text{H}}(\nu) d\nu \right)^{-1}. \quad (2.35)$$

In case of multiple frequency bins the excess energy is calculated in every bin separately.

Cooling processes

The normalised cooling function C is determined by the sum over the cooling rates of individual processes that contribute to the cooling. Here we consider recombination cooling, collisional ionisation cooling, collisional excitation cooling and cooling from free-free emission:

$$C = c_{\text{rec,H}} + c_{\text{ci,H}} + c_{\text{ce,H}} + c_{\text{ff}}. \quad (2.36)$$

The processes of recombination and collisional ionisation have been described earlier in this section. In these processes energy is released. The energy released per unit volume per unit time due to hydrogen recombination is

$$c_{\text{rec,H}} = \zeta_{\text{rec,H}}(T)n_e n_{\text{HII}}, \quad (2.37)$$

where $\zeta_{\text{rec,H}}(T)$ is the recombination cooling coefficient. Energy released in collisional ionisations is

$$c_{\text{ci,H}} = \zeta_{\text{ci,H}}(T)n_e n_{\text{HI}}, \quad (2.38)$$

where $\zeta_{\text{ci,H}}(T)$ is the collisional ionisation cooling coefficient.

Additional sources of cooling we consider are collisional excitation cooling and cooling from free-free emission. Hydrogen atoms can be excited by collisions with electrons. In subsequent de-excitation this energy can be radiated away, thereby providing a means of cooling. This process is referred to as collisional excitation cooling. The energy released is

$$c_{\text{ce,H}} = \zeta_{\text{ce,H}}(T)n_e n_{\text{HI}}, \quad (2.39)$$

where $\zeta_{\text{ce,H}}(T)$ is the collisional excitation cooling coefficient. Free-free emission provides another way of cooling. In case of non-relativistic electrons the energy released is given by

$$c_{\text{ff}} = \zeta_{\text{ff,H}}n_e n_{\text{HII}}, \quad (2.40)$$

where $\zeta_{\text{ff,H}}$ is the free-free emission cooling coefficient.

Implementation

During a radiative transfer calculation the internal energy in Eq. (2.31) is solved in step with the photo-ionisation rate equation, also in an explicit manner. Care must be taken that the time step at which the internal energy is updated is smaller than the characteristic heating time. In general this presents no problem since this time scale is much larger than the ionisation time scale at which the rate equation is solved. In the rare event that the heating time scale is smaller than the ionisation time scale we solve Eq. (2.31) on time steps sufficiently small to ensure a stable solution. A potential speed-up would be to use a dedicated integration scheme instead of the explicit scheme in these cases. However, we find that this rarely occurs and that therefore the speed-up will only be marginal.

In case the photon-conservation scheme described in Sect. 2.4.5 is applied the photoionisation rate equation is subcycled on time steps smaller than the ionisation time scale until photoionisation equilibrium is reached, after which the solution is advanced immediately to the end of the time step. However, when heating and cooling are included the temperature might

still change during the rest of the time step. We assume that once photoionisation equilibrium has been reached in the cell, the cell remains in equilibrium when the temperature changes. We can then proceed by solely advancing Eq. (2.31) at the appropriate time scale, which is in most cases much larger than the ionisation time scale at which was subcycled earlier. After each such subcycle step we adjust the species fractions to the equilibrium values at the new temperature. Again we could speed up the procedure by using a dedicated integration scheme to solve for the internal energy in photoionisation equilibrium. However, the heating time scale is generally larger than the remaining time step, so we only need to integrate over a few steps.

2.5 SUMMARY

In this chapter we have presented a new version of the `SIMPLEX` method. We have described two different ways in which numerical diffusion that was present in the original method can be overcome, by dynamically changing the grid and by using a different transport mode in the optically thin regime. Dynamically updating the grid during the simulation ensures that the scaling of the distance between grid points and the optical depth in the grid cells remains intact. This way photons will not travel farther than a few grid points and cannot deviate far from their original direction. Drawback of this approach is that the resolution in the low optical depth regimes will be insufficient to capture the relevant physical processes. We have therefore introduced a new way of transporting photons that ensures that photons keep their original direction even when they traverse many grid points. This direction conserving transport has the drawback that it is computationally more expensive than the original ballistic transport scheme. Optimal performance is reached with a combination of grid updates and direction conserving transport in the optically thin regime.

Another new feature of the method is the parallelisation for distributed memory machines. The local nature of the transport in `SIMPLEX` makes it relatively easy to parallelise, contrary to ray-tracing methods. When a simulation is performed on multiple processors the domain is split in such a way that approximately the same number of grid points is present on each processor, which ensures optimal scaling properties. The triangulation is performed on each processor separately and domains are glued together with help of ghost vertices. Tests show excellent weak and strong scaling properties of the algorithm.

Finally, we have added a subcycling routine that ensures photon conservation even for large radiative transfer time steps and included a module that solves for the heating and the cooling of the gas. The latter is essential in case of full radiation hydrodynamics simulations.

Extensive tests of the new features of the method in test problems tailored towards cosmological reionisation will be presented in Chapter 5. We will first in Chapter 3 have a closer look on how to create the `SIMPLEX` grid from grids used in hydrodynamics codes, focussing on both structured (regular and Adaptive Mesh Refinement) grids and Smoothed Particle Hydrodynamics grids.

ACKNOWLEDGMENTS

The authors would like to thank Milan Raicevic and Tom Theuns for useful discussions and Jakob van Bethlehem and Rien van de Weygaert for carefully reading the manuscript. We would also like to thank the referee for very useful comments. JPP thanks ICC Durham for their hospitality. JPP acknowledges support from the European Science Foundation (ESF) for the activity entitled 'Computational Astrophysics and Cosmology'.

REFERENCES

- Abel, T., Norman, M. L., & Madau, P. 1999, *The Astrophysical Journal*, 523, 66
- Barber, C. B., Dobkin, D. P., & Huhdanpaa, H. 1995, *ACM Transactions on Mathematical Software*, 22, 469
- Cen, R. 2002, *The Astrophysical Journal Supplement Series*, 141, 211
- Delaunay, B. 1934, *Classe des Sciences Mathematiques et Naturelles*, 7, 793
- Dirichlet, G. L. 1850, *Journal für die reine und angewandte Mathematik*, 40, 209
- Gnedin, N. Y. & Abel, T. 2001, *New Astronomy*, 6, 437
- Gnedin, N. Y. & Ostriker, J. P. 1997, *Astrophysical Journal* v.486, 486, 581
- Iliev, I. T., Ciardi, B., Alvarez, M. A., et al. 2006, *Monthly Notices RAS*, 371, 1057
- Kruip, C. J. H., Paardekooper, J.-P., Clauwens, B. J. F., & Icke, V. 2010, *A&A*, 515, 78
- Kunasz, P. & Auer, L. H. 1988, *J. Quant. Spectrosc. Radiat. Transfer*, 39, 67
- Maselli, A., Ferrara, A., & Ciardi, B. 2003, *Monthly Notices RAS*, 345, 379
- Mihalas, D. & Mihalas, B. W. 1984, *Foundations of Radiation Hydrodynamics*, Oxford University Press, New York
- Miles, R. E. 1970, *Izvestilia Akademii Nauk Armlianskoi SSR Matematika*, 5, 263
- Miles, R. E. 1974, A synopsis of 'Poisson flats in Euclidian space', *Stochastic Geometry* (New York: John Wiley), 202-227
- Møller, J. 1989, *Advances in Applied Probability*, 21, 37
- Norman, M. L., Paschos, P., & Abel, T. 1998, *Memorie della Societa Astronomia Italiana*, 69, 455
- Okabe, A., Boots, B., Sigihara, K., & Chiu, S. N. 2000, *Spatial tessellations : concepts and applications of voronoi diagrams*. 2nd ed. Chichester, John Wiley & Sons Ltd
- Osterbrock, D. E., & Ferland, G. J. 2006, *Astrophysics of Gaseous Nebulae and Active Galactic Nuclei*, 2nd edn. Univ. Science Books, Sausalito, CA
- Pawlik, A. H. & Schaye, J. 2008, *Monthly Notices RAS*, 389, 651
- Rijkhorst, E.-J., Plewa, T., Dubey, A., & Mellema, G. 2006, *A&A*, 452, 907
- Ritzerveld, J. 2005, *A&A*, 439, L23
- Ritzerveld, J. & Icke, V. 2006, *Physical Review E*, 74, 26704
- Ritzerveld, N. G. H. 2007, PhD thesis, Leiden Observatory, Leiden University, P.O. Box 9513, 2300 RA Leiden, The Netherlands

Shirokov, A. & Bertschinger, E. 2005, arXiv, astro-ph:0505087

Springel, V. 2005, Monthly Notices RAS, 364, 1105

Springel, V. 2010, Monthly Notices RAS, 401, 791

Trac, H. & Cen, R. 2007, The Astrophysical Journal, 671, 1

Voronoi, G. 1908, Journal für die reine und angewandte Mathematik, 134, 198