



# Connecting the Dots

Analysis, development and applications  
of the SimpleX algorithm

Chael Kruip

**CONNECTING THE DOTS**  
**ANALYSIS, DEVELOPMENT AND APPLICATIONS OF**  
**THE SIMPLEX ALGORITHM**



# **ANALYSIS, DEVELOPMENT AND APPLICATIONS OF THE SIMPLEX ALGORITHM**

## **Proefschrift**

ter verkrijging van  
de graad van Doctor aan de Universiteit Leiden,  
op gezag van de Rector Magnificus prof. mr. P.F. van der Heijden,  
volgens besluit van het College voor Promoties  
te verdedigen op 20 Oktober 2011  
om 13:45 uur

door

Chael Kruij

geboren te Bergeijk op 8 maart 1981

## Promotiecommissie

Promotor: Prof. dr. V. Icke

Referent: Prof. dr. C. Dullemond (Heidelberg University)

Overige leden: Prof. dr. G. Mellema (Stockholm University)  
Prof. dr. L. Kaper (Universiteit van Amsterdam)  
Prof. dr. R. van de Weygaert (Rijksuniversiteit Groningen)  
Prof. dr. K. Kuijken  
Dr. A. van Genderen

*It is only by windows into a larger and less fretful cosmos  
that the more tragic parts of life become endurable.*

BERTRAND RUSSELL

**Cover: Large-scale structure of the universe (from a dark-matter simulation by Derek Groen) with, overlaid, a Voronoi tessellation and a Delaunay triangulation.**

---

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.0.1	Cosmological reionization . . . . .	1
1.0.2	The enigmatic object $\eta$ Carinae . . . . .	4
1.0.3	The radiative transfer equation . . . . .	7
1.1	Numerical radiative transfer methods . . . . .	8
1.1.1	Long characteristics . . . . .	9
1.1.2	Short- and hybrid characteristics . . . . .	9
1.1.3	Moment methods . . . . .	10
1.1.4	Monte Carlo methods . . . . .	10
1.2	This thesis . . . . .	11
1.2.1	Chapter 1 . . . . .	11
1.2.2	Chapter 2 . . . . .	11
1.2.3	Chapter 3 . . . . .	11
1.2.4	Chapter 4 . . . . .	11
1.2.5	Chapter 5 . . . . .	12
1.2.6	Chapter 6 . . . . .	12
1.2.7	Chapter 7 . . . . .	12
1.2.8	Chapter 8 . . . . .	12
	<b>Part I: Analysis and development of the SIMPLEX algorithm</b>	<b>13</b>
<b>2</b>	<b>The SIMPLEX algorithm</b>	<b>15</b>
2.1	Radiative transfer on unstructured grids . . . . .	16
2.1.1	A natural scale . . . . .	17
2.1.2	The grid . . . . .	17
2.1.3	Voronoi-Delaunay structures . . . . .	18
2.1.4	Three types of transport . . . . .	20
2.1.5	Combined transport . . . . .	25
2.2	Parallelization . . . . .	26
2.2.1	Scaling tests . . . . .	26
2.3	Summary . . . . .	28
2.A	Derivation of the escape fraction of diffuse recombination radiation . . . . .	29



<b>3</b>	<b>Systematic effects in the SimpleX algorithm</b>	<b>31</b>
3.1	Introduction	32
3.1.1	Outline	34
3.2	Anisotropy and its consequences	34
3.2.1	Error measure	35
3.3	Effects on diffusive transport	37
3.3.1	Physical slow down	37
3.3.2	Diffuse drift	37
3.3.3	Diffuse clustering	39
3.3.4	Correcting diffuse transport: weighting schemes	40
3.3.5	Numerical examples: diffuse drift	44
3.3.6	Effects on diffuse transport: clustering	45
3.4	Effects on ballistic transport	47
3.4.1	Ballistic decollimation	47
3.4.2	Ballistic deflection	48
3.4.3	Correcting ballistic transport: weighting schemes	50
3.4.4	DCT throughout the mesh	51
3.4.5	Numerical examples: ballistic decollimation	51
3.4.6	Numerical examples: ballistic deflection	54
3.5	Discussion	55
3.5.1	Prevention versus correction	55
3.5.2	Relative importance of systematic effects	55
3.6	Summary	56
<b>4</b>	<b>The SimpleX mesh</b>	<b>59</b>
4.1	Introduction	60
4.2	Sampling function	60
4.2.1	Mean free path sampling	61
4.2.2	Hybrid sampling	61
4.2.3	Constraints on $\alpha$ and $n_0$	61
4.2.4	Cosmological density field	64
4.3	Creating the SIMPLEX mesh: regularly gridded data	65
4.3.1	Case study 1: a simple regular grid	66
4.3.2	Case study 2: Undersampling in grid-based data	69
4.3.3	Case study 3: 3D hydro of Eta Car	70
4.3.4	Case study 4: Cosmological data	74
4.4	Creating the SIMPLEX mesh: particle based data	75
4.4.1	Case study 5: Undersampling in SPH data of Eta Carinae	76
4.4.2	Case study 6: Sampling cosmological data	78
4.5	Summary	81
<b>5</b>	<b>Heating, cooling and multiple frequencies</b>	<b>83</b>
5.1	Introduction	83
5.2	Ionisation-state of the gas	84
5.2.1	Ionisation	84

5.2.2	Recombination	86
5.2.3	Evolution of the ionization-state	87
5.2.4	Time stepping	88
5.3	Thermal state of the gas	88
5.4	Proper treatment of the spectrum	91
5.4.1	Two grey approximations	91
5.4.2	Spectral hardening	92
5.4.3	Choice of bins	92
5.4.4	Convergence of temperature and ionization structure	95
5.4.5	Testing	97
5.5	Implementation	99
5.6	Summary	100
5.A	Rates and cross sections	100
5.A.1	Cross sections	101
5.A.2	Recombination rates	102
5.A.3	Dielectric recombination	102
5.A.4	Collisional ionization rates	102
5.A.5	Cooling rates	103

## Part II: Applications

105

<b>6</b>	<b>A Markov chain description of SimpleX Radiation Transport</b>	<b>107</b>
6.1	Introduction	108
6.2	Markov processes	108
6.2.1	The transition matrix	109
6.2.2	The initial condition	110
6.2.3	Graphs and reducibility	111
6.2.4	Periodicity	111
6.2.5	The Perron-Frobenius theorem	112
6.2.6	Invariant distributions	113
6.3	Two solution methods	114
6.3.1	Method 1: explicit iteration	114
6.3.2	Method 2: diagonalisation	115
6.4	Correspondence with continuous transfer	115
6.5	Matrix construction	117
6.5.1	Diffuse transport	117
6.5.2	Ballistic transport	118
6.5.3	Direction conserving transport	118
6.5.4	Sources and sinks	118
6.6	Demonstration	119
6.6.1	Using the valence as weights	120
6.6.2	Single source in a homogeneous medium with diffuse transport	121
6.6.3	Single source in a heterogeneous medium with diffuse transport	122
6.6.4	Single source in an absorbing medium with diffuse transport	122

6.7	Convergence	124
6.8	Discussion	125
6.9	Summary	125
<b>7</b>	<b>Late ionization of filaments or early ionization of voids?</b>	<b>127</b>
7.1	Introduction	128
7.1.1	Outline	129
7.2	Method	129
7.3	Diagnostics	131
7.4	Resolution study	131
7.4.1	Convergence with resolution	132
7.4.2	Simulations using the total volume	133
7.4.3	Ionization maps	136
7.4.4	Bubble size distributions	137
7.4.5	Global ionized fractions	138
7.4.6	Age of reionization	139
7.5	Code comparison	140
7.5.1	Diffusivity and moment methods	141
7.5.2	Test 3	141
7.5.3	Test 4	143
7.5.4	Influence of box-size and source-prescription	144
7.6	Summary	145
7.A	Effects of taking a subset of SPH particles	145
7.A.1	Including clumping	146
7.B	Possible effects due to lack of resolution	147
<b>8</b>	<b>Mass and chemical composition of the Homunculus</b>	<b>151</b>
8.1	Introduction	152
8.2	The infrared spectrum of $\eta$ Carinae	153
8.2.1	Observational details	153
8.2.2	Features in the spectrum	153
8.3	Method	155
8.3.1	A three temperature model	155
8.3.2	The size and shape of the dust grains	156
8.3.3	Fitting procedure	157
8.3.4	The stellar spectrum	158
8.3.5	The unidentified sharp feature at $20.23\ \mu\text{m}$	158
8.3.6	A measure for elemental abundance	158
8.4	Results	159
8.4.1	Statistical significance of the fit results	161
8.4.2	The potential presence of MgS	162
8.4.3	Minimum dust mass model	163
8.4.4	Very cold dust beyond the ISO range?	164
8.5	Discussion and conclusions	165
8.5.1	Composition in the context of dust formation	166

---

8.5.2	The total mass in the Homunculus . . . . .	167
8.5.3	Energetics of the Great Eruption . . . . .	168
<b>9</b>	<b>Radiative transfer in the <math>\eta</math> Carinae system</b>	<b>171</b>
9.1	Introduction . . . . .	172
9.2	Hydrodynamic simulations . . . . .	173
9.2.1	SPH . . . . .	173
9.2.2	Parameters . . . . .	174
9.2.3	AMR hydrodynamics . . . . .	176
9.3	RT simulations . . . . .	177
9.3.1	Equilibrium chemistry . . . . .	178
9.3.2	Collisional ionizations . . . . .	179
9.4	Results . . . . .	180
9.4.1	One-dimensional explorations of the ionised region around the primary star . . . . .	180
9.4.2	Three-dimensional RT of the secondary star . . . . .	182
9.4.3	Evolution of the ionization structure with phase . . . . .	188
9.4.4	Preliminary study of RT on AMR data . . . . .	189
9.5	Summary . . . . .	190
	<b>Bibliography</b>	<b>198</b>
	<b>Nederlandse Samenvatting</b>	<b>199</b>
	<b>Curriculum Vitae</b>	<b>209</b>
	<b>Nawoord</b>	<b>211</b>



# CHAPTER 1

---

## Introduction

The emission and absorption of light are inseparably linked with astronomy. Stars emit light which travels through the interstellar medium before it passes through the earth's atmosphere and reaches our eyes and telescopes. As the light interacts with the gas and dust along the way, it gets absorbed and scattered, changing its directionality and spectral characteristics. Hence, the light we receive carries information of its origin as well as its journey and can be used to study both.

Besides its role as messenger, light is a key constituent of our universe dominating the total energy budget until the universe was about 70,000 years old. Knowledge of the interplay between radiation and matter is essential for our understanding of many phenomena studied by astronomers and astrophysicists. Radiation can influence matter (generally gas) dynamically, either directly by radiation pressure or indirectly through heating (cooling) leading to an increase (decrease) of the gas pressure. On a microscopic level, the ionization- and excitation state of gas can be changed by interaction with photons. Reversely, the gas itself can emit radiation which makes the interplay between matter and radiation a highly coupled process.

We proceed with an introduction into two fields of astrophysics that are relevant for this thesis and where the transport of radiation is especially important. These introductions are followed by a brief discussion of several methods designed to solve radiation transfer problems numerically. We postpone a description of our own radiative transfer algorithm, the `SIMPLEX` method, to the next chapter.

### 1.0.1 Cosmological reionization

#### **The birth of the universe**

According to the concordance cosmological model that tries to explain how structures like galaxies and clusters form, the universe came into existence about 13.7 billion years ago in an event that is referred to as the ‘Big Bang’ and evolved according to the so-called  $\Lambda$ -cold dark matter (shortly  $\Lambda$ -CDM) paradigm. With the expansion of the universe, the hot plasma that it contained cooled down to the temperature where atoms can form ( $T \approx 3000$  K). This phase transition is called *cosmological recombination* and happened roughly 380,000 years after the Big Bang, at a redshift of  $z \approx 1100$  (see [Sunyaev & Chluba 2008](#), for a recent review).

With the formation of the first atoms, the photons that had been ‘trapped’ by Thomson scattering off free electrons were free to stream through the now transparent universe. We observe this radiation today as the redshifted Cosmic Microwave Background (CMB) radiation which provides a direct probe of the physical conditions at the time of recombination. The temperature fluctuations (on a mean temperature of 2.725 K) in the CMB are of the order of  $0.2 \mu\text{K}$  which confirms the large degree of isotropy predicted by the theory of *cosmological inflation*.

### The constituents of the universe

Recent observations suggest that the expansion rate of the universe is increasing. This is attributed to a mysterious cosmological phenomenon which is known as *dark energy* or the *cosmological constant* for those who are more prosaic (see [Frieman et al. 2008](#), for a recent review). According to the  $\Lambda$ -CDM model, the present day universe consists of 74% dark energy, 26% matter and a trace amount ( $< 0.1\%$ ) of radiation. The matter-part, in turn, consists for 83% of an unknown material called *dark matter*. The remaining 17% of matter (and thus less than 5% of the total content of the universe) is the *baryonic* matter that we are familiar with and consist of ourselves.

### Structure formation

Under the influence of gravity, the small fluctuations in the density of matter (both dark and baryonic) start to contract. At first, when density contrasts are small, the evolution is linear and can be calculated analytically. Depending on the size of the fluctuation, the collapse enters a non-linear regime. This happens earlier for small-scale than for large-scale fluctuations; the so-called *hierarchical* build up of structure (see [Baugh 2006](#), for a recent review). In a geometrical sense, the collapse of matter follows a hierarchical sequence as well. From the three-dimensional clouds, the first structures to form are two-dimensional *walls*, which collapse into one dimensional *filaments* which in turn stream into zero-dimensional *knots*. The resulting structure is often referred to as the *cosmic web*, and follows from simple analytical considerations ([Icke & van de Weygaert 1987](#)).

### The first ionizing sources

The collapsed dark matter haloes at the nodes of the cosmic web provided the gravitational potential wells where baryonic matter could cool and settle to form the first proto-galaxies. The subsequent formation of the first sources of ionizing radiation marked the end of the cosmological *dark ages* when the universe was about 300 million years old ([Barkana & Loeb 2001](#); [Loeb & Barkana 2001](#); [Ciardi & Ferrara 2005](#)). Because of the primordial gas composition (roughly 76% hydrogen and 24% helium, with a very small fraction of trace elements), a portion of these first stars are thought to be much heavier than those that we see today ([Bromm et al. 2002](#); [Abel et al. 2002](#); [Bromm et al. 2009](#)). The properties of such *Population III* stars are not well understood and are the subject of active research (see [Bromm et al. 2009](#), for a recent review on the first stars and galaxies).

## The Epoch of Reionization

It is probably these sources, together with the subsequent generation of stars and accreting black holes (so-called *Quasi Stellar Sources* or quasars), that brought about the second major phase transition in the history of the universe: *cosmological reionization*. During the epoch of reionization (EoR), the neutral hydrogen gas around the sources was ionized to a high degree (of about one part in  $10^{-5}$  neutral fraction), which is the ionization state of the universe today.

As the radiation from the first sources heated their surroundings, gas clouds that had not yet collapsed may be prevented from doing so due to extra pressure. This extra pressure raises the *Jeans mass* below which clouds can collapse and form galaxies. This so-called *Jeans filtering* (Shapiro et al. 1994; Gnedin & Hui 1998) may thus have direct implications for galaxy formation and the mass function of galaxies in the present day universe.

If the emerging view that reionization is primarily due to stars in galaxies and quasars (and not exotic sources such as decaying dark matter or evaporating black holes) is true, the process is highly inhomogeneous. The morphology of reionization will be patchy, with the shape and size distributions of ionized regions determined by the spatial distribution and character of the sources and sinks of ionizing radiation. Most likely, the gas in the densest regions, close to the sources, ionizes first. It is not yet settled if ionization progresses to ever lower densities with the under-dense *voids* staying neutral longest (the so-called *inside-out* scenario e.g., Iliev et al. 2006b; McQuinn et al. 2007; Trac & Cen 2007; Zahn et al. 2007) or if the filaments, regions of intermediate density with lower source density, are ionized last (the *inside-out-middle* scenario e.g., Gnedin 2000; Ciardi et al. 2003; Finlator et al. 2009b; Petkova & Springel 2011). We will return to this question in Chapter 7 where we investigate if differences in morphology can be due to the methods of radiative transfer used.

## Observational constraints

Although the epoch of reionization is a key period in our universe's history, we know very little about it. Currently the most important observational constraints are from Gunn-Peterson troughs in absorption spectra towards distant ( $z \gtrsim 6$ ) quasars and the Thomson scattering optical depth for CMB photons. Analysis of the Ly $\alpha$  forest in absorption spectra of distant quasars revealed an increase in the optical depth for redshifts larger than  $z \approx 6$  (Fan et al. 2002, 2006). Although this result has been interpreted by many as marking the end of reionization, this interpretation is debated because the optical depth to  $z > 6$  is high ( $> 10^5$ ), making it very difficult to accurately constrain the neutral fraction of hydrogen (Mesinger 2010). The same spectra show that the evolution of the mean opacity of the Ly $\alpha$  forest, happens on a very short time-scale which is compelling evidence for the change of the ionization state of hydrogen (Fan et al. 2006). Such evolution may, however, also be possible without reionization coming to an end (Becker et al. 2007).

The second major constraint comes from observations of CMB photons that Thomson scattered off of the electrons freed by reionization. This scattering introduced the large-scale polarization anisotropies in the CMB which have been measured with the Wilkinson Microwave Anisotropy Probe (WMAP) satellite. A best fit model using the seventh year data results in a mean reionization redshift of  $10.5 \pm 1.2$  (Larson et al. 2011). The constraints on the WMAP results are such that only a time of reionization can be deduced under the assumption that reion-



ization was an instantaneous event. There are, however, strong indications that reionization was a gradual process, spanning hundreds of millions of years (e.g., [Bolton & Haehnelt 2007](#); [Bowman & Rogers 2010](#)). Together, these two observational results constrain the completion of reionization to between  $6 \gtrsim z \gtrsim 12$ , but reveal very little about the nature and morphology of the process.

## Outlook

The most promising upcoming observational results with respect to reionization will be those obtained by radio observatories such as LOFAR, SKA, MWA, GMRT and PAPERS. These telescopes will be able to detect the redshifted 21 cm signal due to the *hyperfine transition* in atomic hydrogen as a function of redshift. Probably the first observational handle, due to its relative simplicity, will be the power spectrum of the spatial fluctuations in neutral hydrogen. When the sensitivity and background-removal techniques improve, it will be possible to construct spatially extended H I maps, directly probing the morphology of reionization. For a recent discussion of observational probes of reionization we refer the reader to [McQuinn \(2010\)](#).

Numerical simulations are essential to give proper interpretation to the existing and upcoming observations. With the aid of simulations, we can identify the relationships between the observed characteristics and values of the physical parameters involved, identify possible degeneracies between parameters and explore the sensitivity of those parameters (see [Trac & Gnedin 2009](#), for a recent review of simulations of the EoR). The computational bottleneck in simulations of the EoR is the inclusion of sufficiently realistic radiative transfer. In the next chapter, we argue that the SIMPLEX algorithm that will be used in this thesis is well-suited for this type of computations.

### 1.0.2 The enigmatic object $\eta$ Carinae

On the scale of stars, many interesting problems involving radiative transfer can be found. For example, a whole field of research is dedicated to the detailed modeling of stellar atmospheres and the spectra that emerge from them. Most of the work has been done with one-dimensional codes that trace a wealth of spectral lines and subtle atomic physics such as charge exchange. Closely related with these atmospheric processes are stellar winds and outflows, the driving and shaping of which rely on momentum transfer between the stellar photons and the outer layers of the star.

In this thesis we will focus on a particularly interesting star from the perspective of radiative transfer: the putative binary system  $\eta$  Carinae. This enigmatic object has been classified as a *Luminous Blue Variable* (LBV) or S Dor variable, and is one of the most extreme stars in the nearby universe. It resides in the Carinae nebula roughly 2.3 kpc away ([Allen & Hillier 1993](#); [Davidson & Humphreys 1997](#); [Meaburn 1999](#); [Davidson et al. 2001](#)), and is one of the most luminous infrared (IR) sources in the sky ([Westphal & Neugebauer 1969](#); [Aitken et al. 1995](#)). Almost all of the  $5.0 \times 10^6 L_{\odot}$  is emitted at IR wavelengths, indicating that the light we observe is reprocessed by dust that is irradiated by the central source ([Cox et al. 1995](#)). The bulk of this dust is contained in the bipolar *Homunculus* nebula which resulted from a 20-year period of mass loss and extreme brightening that occurred in the 19<sup>th</sup> century and is known as the *Great*

### *Eruption.*

Despite a huge amount of research, many fundamental questions about this system remain unanswered. From the total light, a combined mass of the binary of at least  $100 M_{\odot}$  can be derived but it is unclear how this mass is divided between the two stellar components. The binary period has been determined by observations of the recurring *spectroscopic events* to be 5.538 yr (Damineli et al. 2008). From the position of the X-ray minimum in the lightcurve, the eccentricity of the orbit has been constrained to be higher than 0.8 with a best fit value of 0.9 (Corcoran et al. 2001). It is unclear if the orbits' extreme eccentricity is a result of the Great Eruption or has another origin such as a gravitational interaction with a third object.

### **The Great Eruption and the Homunculus nebula**


The Great Eruption is a puzzle in its own right; it is not clear what process is responsible for the sudden brightening in 1837. With a peak apparent magnitude of  $-1$ ,  $\eta$  Car was one of the brightest sources in the sky at the time. Moreover, the brightening coincided with the expulsion of a large amount of gas during a  $\sim 20$  year period centered around  $1841.2 \pm 0.8$  AD (Currie et al. 1996).

The resulting Homunculus (which translates to 'little man' and originates from observations at low resolution, Gaviola 1950) nebula consists of two lobes that appear to be nearly spherical, each with a diameter of about  $8.5'' \approx 0.1$  pc in 1950. The lobes are separated by an equatorial *skirt* of elongated *spikes* and debris normal to the symmetry axis (Morse et al. 1998; Smith & Gehrz 2000). The material in the lobes follows an approximate Hubble flow with an expansion rate of about 0.66% per year (Currie et al. 1996; Morse et al. 2001).

The bipolar shape has been interpreted as the result of high gas densities in the equatorial plane (Konigl 1982; Frank et al. 1995). There is no consensus about the origin of this material. Many authors suggest it to be induced by a close binary companion, while some authors claim that can be caused by a single, moderately rotating, star (e.g., Lamers & Pauldrach 1991; Bjorkman & Cassinelli 1993; Owocki et al. 1994). Another explanation hinging on *gravity darkening* (von Zeipel 1924) has also been put forward and models incorporating this effect reproduce the bipolar shape at least qualitatively (Owocki & Gayley 1996; Maeder & Desjacques 2001; Dwarkadas & Owocki 2002; González et al. 2004). These models need substantial enhancement of wind density and expansion velocity at the poles. This requires that the star rotates close to its brake-up velocity. Therefore, the shape of the homunculus may be a natural consequence of the combined effects of rapid rotation and the high luminosity of Eta Carinae. The ellipsoidal form of the directly observed stellar surface (aligned with the bipolar axis) supports this hypothesis (van Boekel et al. 2003).

### **Mass estimates of the Homunculus nebula**

The total luminous output during the Great Eruption amounts to about  $10^{49.5}$  ergs. Kinetic energy estimates based on measurements of the expansion velocities in the Homunculus and dust mass estimates (discussed below) are of the order of  $10^{49}$  ergs, only two orders of magnitude less than a supernova explosion! These extreme values are beyond what well known, radiation driven, mechanisms can produce, implying that the Great Eruption is the result of some unknown effect.



**Figure 1.1:** *Left:* Hubble Space Telescope image of Eta Car (Morse et al. 1998). *Right:* Eta Car in the IR (at  $24.5\ \mu\text{m}$ ; Smith et al. 2002).

Because of the many observational difficulties associated with the direct detection of highly obscured gas, mass estimates have almost exclusively been based on measurements of the IR radiation due to dust. The major uncertainties are the chemical composition and grain size distribution of the dust and the gas-to-dust ratio.

The first attempts to determine the mass of the Homunculus were based on a spectral energy distribution (SED) consisting of a handful of photometric flux measurements and yielded dust mass estimates of the order of  $M_d = 0.01\text{-}0.03\ M_\odot$  (Mitchell & Robinson 1978; Hackwell et al. 1986; Cox et al. 1995). Smith et al. (1998) inferred the spatial distribution of dust emission from thermal-IR observations. They found that the Homunculus should contain at least  $0.025\ M_\odot$  of dust. Three temperature components were derived; 200, 240 and 420 K for the bipolar lobes, the skirt, and the core, respectively. Data obtained with the Infrared Space Observatory (ISO; Kessler et al. 1996) led to substantially higher mass estimates of  $0.1\text{-}0.15\ M_\odot$  by Morris et al. (1999) because of the detection of a cool dust component of 110, in addition to a warm 190 K component. From MIRAC3 imaging and the ISO data from Morris et al. (1999), Smith et al. (2003b) found in excess of  $12\ M_\odot$  of mass of which about 90% is traced by a thin, cool ( $\sim 140$  K) shell of molecular Hydrogen (Smith 2006). From sub-millimetre SCUBA observations, Gomez (Née Morgan) et al. (2006) derives  $0.3\text{-}0.7\ M_\odot$  of dust around Eta Carinae, however, this may be situated in a volume larger than the Homunculus.

Assuming a canonical gas-to-dust ratio of 100, estimates of the Homunculus mass thus range from 1 to  $70\ M_\odot$ , where most of the uncertainty is due to insufficient knowledge of the dust composition and the gas-to-dust ratio. We go deeper into these issues in Chapter 8 where we derive possible chemical compositions for the Homunculus nebula and find a lower limit for its total mass.

## Orbital parameters

One well established observational result is the binary periodicity of 5.538 yr, (Damineli et al. 2008) which also has been found in near-IR photometry (Whitelock et al. 1994), and radio data (Duncan & White 2003). This periodicity could be related to the thermal timescale of the outer layers (Davidson et al. 2005; Martin et al. 2006; Davidson & Humphreys 1997) but the consensus view is that  $\eta$ Carinae consists of more than one star. On the basis of radial velocity variations, Damineli (1996) postulated that the system is a high eccentricity ( $e \gtrsim 0.8$ ) binary with an orbital period of about  $5.52 \pm 0.01$  yr. This has been confirmed by Davidson (1997), and the evidence has been traced back as far as 1948.

Changes in the IR emission of the central region at two different times in the 5.52 year period have been reported by Duncan et al. (1999) and Smith & Gehrz (2000). Observations in the X-Ray domain show a minimum with the right periodicity, confirming the binary hypothesis (Ishibashi et al. 1999; Corcoran et al. 2001).

Recently, significant progress has been made by fitting spatially resolved  $[\text{Fe}_{\text{III}}]$  emission using SPH simulations and Hubble Space Telescope (HST) data (Madura 2010). They were able to determine for the first time the orientation of the orbit in three dimensional space. Their results show that, within the uncertainty of the model of about  $10^\circ$ , the orbital axis of the  $\eta$ Carinae binary coincides with the symmetry axis of the Homunculus nebula. This strategy of combining high-quality spectrometric data with simulation seems to be the path forward in order to place further constraints on the binary parameters such as mass-ratio. To improve on the current results, detailed radiative transfer effects have to be incorporated in the models. We will return to this question in Chapter 9.

### 1.0.3 The radiative transfer equation

Apart from the aforementioned wealth of interaction mechanisms, inclusion of radiation in astrophysical simulations is further complicated by the high dimensionality of the equations of radiation transfer (e.g., Eq. (1.2)). A beam of radiation is described by three spatial, two angular one frequency and a time variable. Depending on the application, geometrical simplifications can be made or the explicit time- and frequency-dependence can be dropped. Despite such approximations, there are hardly any applications where an analytical treatment is possible. We therefore turn to numerical methods for solving problems including complex radiative transfer.

With the advent of the computer as an important catalyst, a myriad of efforts to solve the equations of radiative transport within a numerical framework have been developed. The resulting algorithms cover a wide range of applications where generally a method is tailored to a specific problem. In most cases, specialisation means either the choice of a specific physical scale (consider detailed models of stellar atmospheres or large-scale cosmological simulations) or emphasis on physical processes relevant to the problem at hand.

To describe the emission and absorption of radiation as a function of space and time, we consider the gain and loss of energy in a beam of radiation due to source and sink terms. To this end, we define the specific intensity  $I$ , by stating that the energy flowing through a surface area  $da$  located at position  $\mathbf{r}$  in a time interval  $dt$  in a solid angle  $d\Omega$  around direction vector  $\mathbf{n}$

in frequency interval  $d\nu$  is given by

$$dE = I_\nu(\mathbf{r}, \mathbf{n}, t) \cos \theta da d\Omega d\nu dt \quad (1.1)$$

where  $\theta$  is the angle between  $\mathbf{n}$  and the normal of the surface element  $da$ . The equation of radiative transfer for a medium whose properties can change both in space and time is then given by

$$\frac{1}{c} \frac{\partial I_\nu}{\partial t} + \mathbf{n} \cdot \nabla I_\nu = \eta_\nu - \chi_\nu I_\nu. \quad (1.2)$$

where  $\eta_\nu$  is the monochromatic emission and  $\chi_\nu$  is the total opacity coefficient. If  $\frac{1}{c} \frac{\partial I_\nu}{\partial t} \ll 1$ , in other words, if  $I_\nu$  is not explicitly time dependent *or* the time and space discretization is such that  $c$  can be considered infinite, this equation simplifies to

$$\mathbf{n} \cdot \nabla I_\nu = \eta_\nu - \chi_\nu I_\nu. \quad (1.3)$$

If we take the spatial derivative along the ray and divide by  $\chi_\nu$ , Eq. (1.3) can be rewritten as

$$\frac{\partial I_\nu}{\partial \tau} = S_\nu - I_\nu, \quad (1.4)$$

where we define the source function  $S_\nu \equiv \eta_\nu/\chi_\nu$  and optical depth  $d\tau \equiv \chi_\nu ds$  and  $s$  parametrises the distance along the ray. Eq. (1.4) can be solved numerically if  $\eta_\nu$  and  $\chi_\nu$  are known locally, the meaning of ‘local’ depending on the type of discretisation of the volume. The form of Eq. (1.4) suggests that the optical depth is the most natural variable to discretise. In general, this will produce cells of unequal volume, and therefore, an irregular computational mesh. In subsequent sections and Chapter 4 we explain in more detail how we generate a set of nuclei that reflects the underlying optical depth field and connect these nuclei with a Delaunay triangulation to create such a mesh.

## 1.1 NUMERICAL RADIATIVE TRANSFER METHODS

We now give a short overview of methods that have been developed to solve Eq. (1.3) in its various forms. As will be discussed in detail in Chapter 2, the `SIMPLEX` algorithm (the main subject of this thesis) is particularly well-suited for applications in cosmological radiative transfer. We therefore focus mainly on methods whose field of application is that of cosmological reionization and galaxy formation. However, the `SIMPLEX` algorithm is not limited to this area of application, as we will argue in the subsequent chapters. The application to interacting winds in a binary stellar system in Chapter 9 also bears witness to this statement.

There exists a huge body of work dedicated to solving the radiative transfer equation in one-dimensional geometries (e.g. stellar and planetary atmospheres) including a myriad of atomic and molecular lines. In this text we will not discuss these but rather focus on those methods whose main purpose is to transport ionizing radiation through complex geometries in three-dimensional space.

This class of ‘cosmological’ radiative transfer codes has a set of requirements that are typical for this area of research. First, they must be able to deal with a large range in length-scales

(more than 7 orders of magnitude, [Trac & Gnedin 2009](#)). This is because the relevant absorbing systems may be as small as a few kpc (Lyman limit systems), whereas the simulation volume should be of the order of a Gpc to have a fair representation of rare quasars and not be dominated by cosmic variance. To allow for this extreme dynamic range, the simulation method must be able to dynamically adapt its resolution. To execute simulations including many (often more than  $128^3$ ) resolution elements implies large memory requirements greatly surpassing those available on common desktop computers. Furthermore, the number of sources in these large volumes can increase well into the millions. Furthermore, if recombination radiation is taken into account, every resolution element might become a source implying billions of sources. For these reasons, an ideal cosmological radiative transfer code should naturally adapt its resolution, be able to deal with many sources, and run efficiently on distributed memory super-computers.

### 1.1.1 Long characteristics

Almost every radiative transfer method (with exception of so-called *moment* methods (see Sect. 1.1.3)) aims at solving the radiative transfer equation along a pencil beam, making it possible to drop the explicit angular dependence. The first attempts to do this go back to work of [Mihalas & Weibel Mihalas \(1984\)](#), who introduced the *long characteristics* (LC) method. In this approach, rays are cast from every source to every cell in the computational domain and Eq. (1.3) is solved by integrating the optical depth between the source and target cell. Although LC allows for very accurate results, the computational effort scales as  $N_s N_c$  where  $N_s$  is the number of sources and  $N_c$  the number of cells in the domain, making it computationally prohibitively expensive for large values of  $N_s$ .

It has been shown, however, that the computation time can be dramatically reduced if the time-averaged optical depth along the rays is used ([Mellema et al. 2006](#)). This allows for very long computational time steps in comparison to other methods.

### 1.1.2 Short- and hybrid characteristics

The LC method is inefficient in another sense as well; cells close to sources are traversed by many rays (because every cell is connected to every source), resulting in many redundant calculations. This redundancy can be alleviated by adding up optical depth when traversing cells in an upstream fashion. The resulting *short characteristics* (SC; e.g. [Kunasz & Auer 1988](#)) method is less accurate because of interpolation errors at cell boundaries, but this is outweighed for most applications by its significant speed-up. The necessity of traversing cells outward from the source complicates the parallelization of the SC methods, and approximations have to be introduced to achieve this goal.

One implementation that combines the speed of the SC with the accuracy of LC has been devised by [Rijkhorst et al. \(2006\)](#). This so-called *hybrid characteristics* method can work on adaptive mesh refinement (AMR) grids and has been coupled to the FLASH hydro-code ([Fryxell et al. 2000](#)).

Another method has been presented in [Abel & Wandelt \(2002\)](#). They use a ray-splitting criterion that guarantees a fixed number of ray-intersections per cell. Several approximations that reduce the linear scaling with  $N_s$  of this code have been introduced; grouping of adjacent

sources (Razoumov & Cardall 2005), limiting the splitting of rays (McQuinn et al. 2007) or merging near-parallel rays (Trac & Cen 2007).

With the exception of the hybrid characteristics method by Rijkhorst et al. (2006) and the method of Razoumov & Cardall (2005), these methods can be used exclusively on regular meshes severely limiting their scope of application. Motivated by this limitation and the popularity of smoothed particle hydrodynamics (SPH), several ray-tracing schemes that operate directly on the SPH particles have been developed (Alvarez et al. 2006; Susa 2006; Pawlik & Schaye 2008).

### 1.1.3 Moment methods

Another way to bring down the computation time for radiative transfer problems is to completely break the unfavorable scaling of computation time with the number of sources. One popular approach that achieves this goal solves the radiative transfer equation by taking its first two angular moments. These moments represent the photon number density and flux respectively, and extended with a closure term, constitute a system of differential equations that can be solved straightforwardly. The various implementations of this method differ mainly by their method to compute the Eddington tensor used as closure term (Gnedin & Abel 2001; Aubert & Teyssier 2008; Petkova & Springel 2009; Finlator et al. 2009a).

Moment methods have the advantage that the solving the RT equation does not scale with the number of sources. The computation of the Eddington tensor, however, does scale with the number of sources if done accurately. In its simplest form, this is done by time-independent ray-casting (i.e., Finlator et al. 2009a). This breaks the time-dependent nature of the solution, an approximation that is incorrect for rapidly varying emissivity. However, in most cases, the sources do not evolve on a very short time-scale and one can get away with updating the Eddington tensor only when the emissivity changes considerably.

More approximate calculations of the Eddington tensor such as the Optically Thin Variable Eddington Tensor (OTVET) approach (Gnedin & Abel 2001) or Flux Limited Diffusion (FLD) approximation (Aubert & Teyssier 2008) are less computationally intensive. These approximations are known to distort the shape of ionized bubbles around multiple sources.

Another advantage is that, assuming the Eddington tensor is computed, the solution is completely local. This simplifies domain-decomposition in parallel implementations because a CPU only needs to receive information from CPUs that host neighbouring domains. Moreover, the solution can be solved implicitly, which means that, in principle, a stable solution is attainable without having to satisfy the radiative Courant condition.

### 1.1.4 Monte Carlo methods

A large and versatile class of radiative transfer methods use probabilistic means of solving the radiative transfer equation; the so-called *Monte Carlo methods*. In this approach, discrete packets of radiation are sent in random directions and interact with the medium along the way (e.g Ciardi et al. 2001). Monte Carlo methods are popular because of their relatively straightforward implementation and the possibility to include physical processes in a probabilistic manner.

In a sense, Monte Carlo methods are similar to LC because the number of photon packages that every source needs to send scales with the number of resolution elements in the simulation. In practice, this prohibitive scaling can be improved upon using so-called *importance sampling* techniques, where more rays are cast into regions where higher resolution is needed. Another disadvantage of Monte Carlo methods is that the signal to noise ratio improves slowly (with the well-known inverse square-root of the number of rays cast). Monte Carlo methods for cosmological reionization simulations exist for both regular meshes (Maselli et al. 2003; Maselli et al. 2009), SPH particles (Semelin et al. 2007; Altay et al. 2008) and AMR grids (Cantalupo & Porciani 2011).

## 1.2 THIS THESIS

In this thesis we describe the analysis and improvements of the SIMPLEX algorithm for radiative transfer (Chapters 2, 3, 4 and 5). Several applications are presented in Chapters 6, 7, 8 and 9.

### 1.2.1 Chapter 1

The central ideas behind the SIMPLEX algorithm are laid out. As they are essential to the method, we give a concise introduction into Voronoi-Delaunay structures and show how SIMPLEX uses them as the basis for efficient radiative transfer calculations. Three modes of transport are described together with some examples of their application. We conclude with an overview of the parallelization strategy and scaling properties of the parallel implementation.

### 1.2.2 Chapter 2

Four distinct systematic effects of the SIMPLEX method that arise due to local anisotropy of the Voronoi-Delaunay grid are identified. Analytical descriptions of these effects are quantified and verified with numerical calculations. This analysis leads to suitable corrections to the method which are implemented and tested. Corrections include several weighting schemes and a grid construction procedure that minimizes unphysical effects.

### 1.2.3 Chapter 3

We delve deeper into the specifics of the construction of computational grids optimized for SIMPLEX radiative transfer. Examples of the most important systematic effects studied in Chapter 3 in realistic simulations are given. We start with a detailed description of our sampling technique and apply it to both grid- and particle-based data.

### 1.2.4 Chapter 4

We describe the implementation of various physical processes relevant for the treatment of gas of primordial composition. Specifically, the inclusion of heating and cooling prescriptions and



a multiple frequency treatment of the radiation field is described. We discuss our strategy for an optimal choice of frequency discretization.

### 1.2.5 Chapter 5

The SIMPLEX method can be re-formulated in the context of Markov chains on a Delaunay graph. We discuss this dual description and the construction of the associated transport matrix. For stationary solutions to radiative transfer problems in static media, it is possible to find the solution by direct diagonalization of the transport matrix. Preliminary tests suggest that this solution method is much more efficient for the solution of equilibrium scattering problems in inhomogeneous media than the iterative transport implemented in SIMPLEX.

### 1.2.6 Chapter 6

We apply the SIMPLEX algorithm on cosmological simulations from [Finlator et al. \(2009b\)](#) to verify their controversial results with an independent and fundamentally different method. With the SIMPLEX method, a consistently inside-out morphology is found whereas the method used in [Finlator et al. \(2009b\)](#) finds a distinct reversal to outside-in morphology. We conduct an elaborate resolution study and conclude that the results do not differ due to lack of spatial resolution in the original study. With a side-by-side comparison of the two radiative transfer methods used, we rule out several possible origins of the observed discrepancies.

### 1.2.7 Chapter 7

The composition and mass of the solid state matter in the Homunculus nebula of  $\eta$  Carinae is determined using fits of the spectral energy distribution observed with ISO. A comprehensive fitting method is used where the fits are constructed by including or rejecting dust species based on a  $\chi^2$  criterion. We derive a lower limit for the total mass of the nebula based on condensational arguments.

### 1.2.8 Chapter 8

We apply the SIMPLEX algorithm to AMR and SPH data of the wind-wind interaction region of the  $\eta$  Carinae system. This paves the way to improved analysis of forbidden line emission observed with the HST. In current studies, the size and shape of ionization regions in the system are estimated based on simple geometrical criteria combined with a density threshold. We conduct full RT post-processing of the hydrodynamical data, including the effects of both photo- and collisional ionization for hydrogen and helium.

## **Part I**

# **Analysis and development of the SIMPLEX algorithm**



## CHAPTER 2

---

# The SIMPLEX algorithm

C. J. H. Kruij, J.-P. Paardekooper & V. Icke

Part of this chapter consists of work that has been published in  
Astronomy & Astrophysics 515, A78, 2010 and  
Astronomy & Astrophysics 515, A79, 2010

**W**E present the latest incarnation of the SIMPLEX algorithm. This method uses Voronoi-Delaunay structures for its computational mesh. Because such meshes are scale-free, radiative transfer can be performed on complex opacity fields with highly adaptive spatial resolution. The local nature of SIMPLEX radiative transport makes the method computationally very efficient and easy to parallelize. The computational effort of this algorithm does not scale with the number of sources which makes SIMPLEX ideally suited for simulations of large-scale reionization, and allows for self-consistent treatment of diffuse recombination radiation.

In this chapter, we describe the basic workings of the `SIMPLEX` algorithm which is the topic of the next two chapters and the common theme of this thesis. Conceived by [Ritzerveld & Icke \(2006\)](#), and implemented by [Ritzerveld \(2007\)](#), the `SIMPLEX` algorithm solves the general equations of particle transport by expressing them as a walk on a graph. The method can consequently be considered to be a *Markov chain* on a closed graph, as will be discussed in Chapter 6. Transported quantities travel from node to node on the graph, where each transition has a given probability.

This approach has several advantages that we will now highlight. More specifically `SIMPLEX`

- does not increase its computational effort or memory use with the number of sources in a simulation and consequently treats, for instance, scattering by dust and diffuse recombination radiation without added computational effort;
- naturally adapts its resolution to capture the relevant physical scales, (expressed in photon mean free path lengths);
- works in parallel on distributed memory machines;
- is compatible with grid-based as well as particle-based hydrodynamics codes (where the latter is the more natural combination due to the particle-based nature of both `SIMPLEX` and SPH);
- is computationally cheap because of the local nature of the Delaunay transport

## 2.1 RADIATIVE TRANSFER ON UNSTRUCTURED GRIDS

In this section we introduce the fundamentals of the `SIMPLEX` algorithm and the Voronoi-Delaunay triangulation, which lies at the heart of our method and plays a central role in this text, specifically Chapters 3, 4 and 6.

Although the `SIMPLEX` algorithm is the first Voronoi-Delaunay-based radiative transfer method, the advantageous properties of Voronoi tessellations ([Dirichlet 1850](#); [Voronoi 1908](#)) and Delaunay triangulations ([Delone 1934](#)) have been recognized before in the context of (astro-)physical and geophysical applications.

The adaptive nature of the Voronoi-Delaunay structures allows for the accurate representation of systems with large dynamic range such as the large-scale cosmological structure. Ground-breaking work has been done by [Bernardeau & van de Weygaert \(1996\)](#) that laid the fundamentals for the successful *Delaunay Triangulation Field Estimator* (DTFE [Schaap & van de Weygaert 2000](#)) method. The DTFE has been shown to give density estimates far superior to typical SPH-kernel based methods ([Pelupessy et al. 2003](#)) and has seen applications and further development up to the present day ([Romano-Diaz & van de Weygaert 2007](#); [van de Weygaert & Schaap 2009](#); [Aragón-Calvo et al. 2010](#); [Cautun & van de Weygaert 2011](#)).

Another important field of application is that of hydrodynamics where either the Delaunay triangulation or the Voronoi tessellation is used as the hydro-mesh ([Whitehurst 1995](#)) and ([Springel 2010](#); [Duffell & MacFadyen 2011](#)) respectively. Such methods provide a natural implementation of Lagrangian hydrodynamics and alleviate some of the problems encountered in Eulerian and SPH codes (also see [Heß & Springel 2010](#)).

### 2.1.1 A natural scale

The fundamental idea behind the SIMPLEX algorithm is that there exists a natural scale for the description of radiative processes: the photon local *mean free path*

$$l_{\text{mfp}} \equiv \frac{1}{\sigma n}, \quad (2.1)$$

where  $\sigma$  is the total extinction cross-section and  $n$  the number density of the particles responsible for that extinction. Since photons typically travel one mean free path before interacting with the medium, information on a much smaller scale does not necessarily yield a deeper understanding of the physical problem at hand. Because the expectation value for the distance travelled after unit optical depth is the mean free path (by definition), this statement is equivalent to the observation that Eq. (1.4) suggests a discretisation of the radiative transfer equation based on optical depth.

Following this line of reasoning, the next step would be to choose a computational mesh that inherently carries this natural scale, in other words, use an irregular grid whose resolution adapts locally to the mean free path of the photons travelling over it. For clarity, we limit our current discussion to monochromatic radiation only, and postpone the description of our multi-frequency implementation to Chapter 5.

### 2.1.2 The grid

This computational mesh or transport graph is constructed by first defining a point process that represents the underlying (physical) problem and second by connecting these points according to a suitable prescription.

The point process is defined by prescribing the local point field density as a function of the scattering and absorption properties of the medium through which the particles propagate. To avoid confusion, we distinguish between the number density of extinguishing particles (e.g., atoms or dust particles),  $n(\mathbf{x})$ , and the number density of the points that constitute the basis for our computational mesh,  $n_p(\mathbf{x})$ .

To translate a given opacity field given by  $n$  to a point set, we make use of a modified Poisson process which is locally homogeneous in the mean but changes its number density according to  $n$ .

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

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

The only parameter in this process is the point intensity  $n_p$ , which is a global constant and sets the total number of points in the volume. Every region in the volume has the same probability that points are placed there, which in our case corresponds to a constant opacity.

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

$$P(N(A) = x) = \frac{n_p(A) |A| e^{-n_p(A) |A| x}}{x!}, \quad x = 0, 1, 2, \dots \quad (2.3)$$

where

$$n_p(A) = \int_A n_p(\mathbf{x}) d\mathbf{x}. \quad (2.4)$$

The point intensity function  $n_p(\mathbf{x})$  follows the opacity of the medium on a global scale, while locally it retains the properties of the homogeneous Poisson distribution. An alternative, and possibly more physically intuitive, recipe for constructing the non-homogeneous Poisson process can be written as

$$n_p(\mathbf{x}) = \Phi * f(n(\mathbf{x})), \quad (2.5)$$

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  $n(\mathbf{x})$ . We use this recipe to construct the `SIMPLEX` grid. Grid points are placed by randomly sampling the correlation function  $f(n(\mathbf{x}))$  that translates between  $n(\mathbf{x})$  and  $n_p(\mathbf{x})$ . We will discuss the possible choices for the correlation function shortly below, and in more depth in Chapter 4.

### 2.1.3 Voronoi-Delaunay structures

With the point-set obtained with Eq. (2.5) in hand we can construct the computational mesh which forms the basis for the `SIMPLEX` algorithm. We have chosen to use the Voronoi-Delaunay triangulation for this purpose.

The *Voronoi tessellation* of a set of nuclei  $x_i$  in a  $D$ -dimensional space is defined as  $V = \{C_i\}$ , where

$$C_i = \{y \in \mathbb{R}^d : \|x_i - y\| \leq \|x_j - y\| \ \forall x_j \neq x_i\}. \quad (2.6)$$


This means that every point inside a Voronoi cell is closer to the nucleus of that cell than to any other nucleus. The set of all points that have exactly two nearest nuclei  $n_1, n_2$  is the *Voronoi wall* between these nuclei. If  $D = 2$ , this wall is a line; if  $D = 3$ , it is a plane; and so on.

A Voronoi wall therefore separates two neighbouring Voronoi cells. The connection between these nuclei, called an *edge*, is the geometric dual of the wall. The set of all the edges is the *Delaunay triangulation* of the point set and connects every Voronoi nucleus to its neighbouring nuclei. An example of a Voronoi-Delaunay grid in the plane is shown in Fig. 2.1.

Both the Voronoi tessellation and the Delaunay triangulation can be used to partition space. Other than the Voronoi tessellation, whose cells have no fixed number of walls, the Delaunay triangulation exists of *simplices* which are a generalisation of triangles in  $\mathbb{R}^d$ . So a simplex is a triangle in  $\mathbb{R}^2$  and a tetrahedron in  $\mathbb{R}^3$ . The Delaunay triangulation can be defined as the unique triangulation of nuclei for which the interior of the circum-sphere of any Delaunay simplex contains no other nuclei.

Miles (1970, 1974) and Møller (1989) have derived several properties for Voronoi-Delaunay triangulations based on homogeneous Poisson processes. Additional properties and generalizations to non-Poissonian distributions have been derived by van de Weygaert (1991) and van de Weygaert (1994). Because Eq. (2.5) is approximately homogeneous locally (we will discuss the effects of non-homogeneity in Chapter 3), they are also applicable to our transport grid. Two results relevant 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

$$E_{2D}(E) = 6 \quad (2.7)$$



**Figure 2.1:** Voronoi tessellation of the plane (solid lines). Each cell contains all points that are closer to its nucleus (indicated by a dot) than to any other nucleus. The corresponding Delaunay triangulation is shown in dashed lines. Note: only visible nuclei are included in the triangulation.

and

$$E_{3D}(E) = \frac{48\pi^2}{35} + 2 \approx 15.54. \quad (2.8)$$

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

$$E_{2D}(L) = \frac{32}{9\pi} n_p^{-1/2} \approx 1.132 n_p^{-1/2} \quad (2.9)$$

and

$$E_{3D}(L) = \frac{1715}{2304} \left(\frac{3}{4}\right)^{1/3} \pi^{-1/3} n_p^{-1/3} \approx 1.237 n_p^{-1/3}. \quad (2.10)$$

For our application to transport theory, it is also important that the Delaunay triangulation has a minimax property, i.e. it is the triangulation with the largest smallest angle between adjacent triangle edges. Of all the possible triangulations of a given point set, the Delaunay triangulation is the one that maximises the expectation value of the smallest angle of its triangles. In more colloquial terms, the Delaunay triangulation has the least ‘sliver-like’ triangles, and the most ‘fat’ triangles.

### Construction of the grid

The construction of the grid itself is a task performed by dedicated software. Once the generating nuclei have been given, the Voronoi-Delaunay structure is unique. In creating the distribution of generating nuclei, we can manipulate the properties of our computational grid. The translation from a given density (opacity) field to a point-set suitable for SIMPLEX is thus a fundamental part of the algorithm, but the grid construction can be done by any capable routine.

For the results in this thesis we have used the QHull package<sup>1</sup>, based on the Quickhull algo-

<sup>1</sup>www.qhull.org



rithm (Barber et al. 1996). This package computes the Delaunay triangulation in any dimension up to 8, including simplex properties such as volume and surface. It performs at the theoretical limit of  $O(N \log N)$ , and is very stable against floating-point round-off errors in the event that points lie very close to each other.

### 2.1.4 Three types of transport

In `SIMPLEX`, transport is always between neighbouring Voronoi cells, i.e., those connected by the Delaunay triangulation. This choice is by no means dictated by the grid. We may as well transport photons through the Voronoi-Delaunay structures by means of long characteristics or randomly chosen directions as in Monte Carlo methods.

The main reason we use the Delaunay edges as paths for the radiation is that it allows us to directly use the information carried by the Voronoi-Delaunay structures for the transport process. Cell-to-cell distances (Delaunay edge-lengths), volumes (Voronoi volumes) and connectivity are all known. The locality of the Delaunay transport method (radiation travels only to neighbouring cells) makes the transport step itself very efficient. For every computational cycle, which we refer to as *sweep* in the remaining text, all vertices take one transport step. The global nature of the radiative transfer problem is retrieved after sufficient iterations of this local transport.

When photons travel through a cell, the optical path length,  $l$ , is taken to be the average length of the Delaunay edges of that cell. If the number density of atoms in the cell is given by  $n$ , the fraction of photons that are removed from the bundle,  $N_{\text{rem}}$ , is given by

$$N_{\text{rem}} = N_{\text{in}} e^{-n\sigma l}, \quad (2.11)$$


where  $\sigma$  is the total cross-section that may be the result of multiple extinction processes.

In general, extinction can be subdivided into absorption and scattering. Radiation that is removed by absorption processes will change both the temperature of the medium as well as its physical state, but does not need to be transported further at this point. Photons that are removed from the bundle by scattering should propagate to neighbouring cells, either isotropically or with a certain directionality. This can be accomplished using the *diffuse transport* method schematically depicted in the left panel of Fig. 2.2 and explained in more detail in Sect. 2.1.4. As we will see in Sect. 3.3.4, anisotropic scattering processes can be simulated straightforwardly by assigning weights to the outgoing edges so that more radiation is transported in the appropriate directions. The radiation that is not removed from the bundle, however, needs to travel straight onwards along the original incoming direction. For optically thick ( $\Delta\tau > 1$ ) cells, we simulate this using *ballistic transport* (see the central panel of Fig. 2.2 and Sect. 2.1.4).

In regions of the grid where the cells are optically thin ( $\Delta\tau < 1$ ), ballistic transport becomes too diffusive and we need to resort to *direction-conserving transport* or DCT (see the right panel of Fig. 2.2 and Sect. 2.1.4). We now proceed by describing these transport methods in more detail and show how they are combined in a general simulation.

#### Diffuse transport

We begin with the description of conceptually the most simple form of transport implemented in `SIMPLEX`. For every sweep, the content of each nucleus is distributed equally among its neigh-



**Figure 2.2:** Three principal means of transport used in the SIMPLEX method. *Left:* Diffuse transport, photons from the incoming edge (not shown) are distributed outward along all edges (including the incoming edge). *Centre:* Ballistic transport, photons are transported along the D edges directed most forward with respect to the incoming direction. *Right:* Direction-conserving transport, photons (indicated with the dotted arrows) are transported as in ballistic transport but their direction is stored indefinitely in a global set of solid angles.

bouring nuclei (see the left panel of Fig. 2.2.) We call this kind of transport *diffuse* because it has no memory of direction. For a homogeneous distribution of nuclei, the transported quantity will diffuse outwards, spreading spherically from the position of a source. This type of transport is appropriate for photons that are either scattered diffusely or absorbed and re-emitted in random directions.

Photons due to recombination of electrons and ions are an example of diffuse radiation relevant for simulations where the ionisation-state of the medium is important. In contrast to methods where the computational effort scales with the number of sources, SIMPLEX can treat recombination radiation self-consistently without significant added computational effort. However, because the recombination photons originating in a cell will also be absorbed by the gas in that cell, we use a sub-grid description of the escape fraction of diffuse radiation as a function of optical depth,  $\tau$ , of the cell

$$f_{\text{esc}}(\tau) = \frac{3}{8\tau^3}(e^{-2\tau}(1 + 2\tau) - (1 - 2\tau^2)), \quad (2.12)$$

which is derived in Appendix 2.A.

### Ballistic transport

We now consider a group of photons that is transported along a Delaunay edge to a certain nucleus. We assume that the nucleus represents a finite optical depth. A fraction of the photons will be removed from the group by the interaction and another fraction will fly straight onward. Diffusive transport is not suited to describing this behaviour, so we introduce *ballistic* transport.

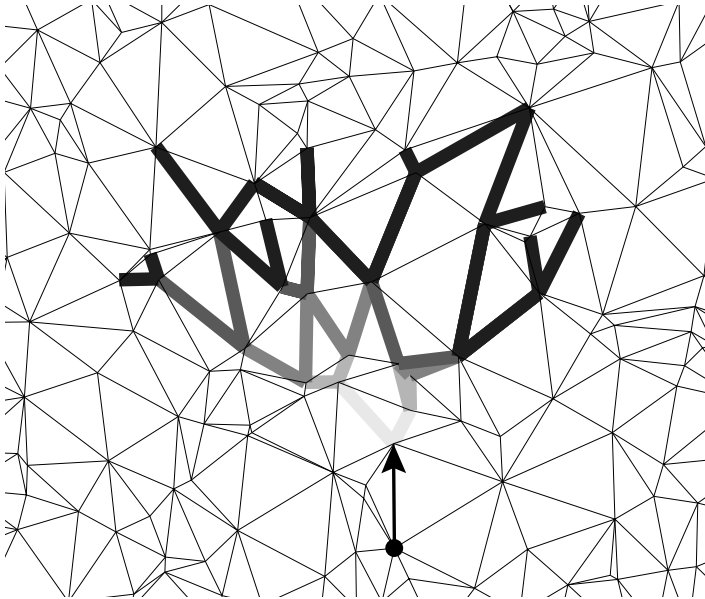
In the ballistic case, the incoming direction of the photons is used to decide the outgoing direction (introducing a memory of one step into the past). In the generic Delaunay triangulation, there is no outgoing edge parallel to the incoming one, so the outgoing photons are distributed

over the  $D$  most forward pointing edges, where  $D$  is the dimension of the propagation space (see centre panel of Fig. 2.2). As such, we ascertain that for an isotropically radiating source the complete ‘sky’ is filled with radiation because the opening angle associated with each edge on average corresponds to  $2\pi/\Lambda$  or  $4\pi/\Lambda$  in two and three dimensions, respectively<sup>2</sup>.

Because of the random nature of the directions in the Delaunay grid, we note that radiation will tend to lose track of its original direction after several steps, a property that we call *decollimation*, as it will steadily increase the opening angle of a beam of radiation as it travels along the grid (see Fig. 2.3). This property renders ballistic transport appropriate for highly to moderately optically thick cells only, where just a negligible amount of radiation has to be transported more than a few steps. If we take unity as a lower limit to the optical depth of a cell for which ballistic transport is used, at every intersection a fraction of  $(1 - 1/e)$  of the photons becomes absorbed and the cumulative average deflection (decollimation)  $\theta$  becomes

$$\theta = \theta_D \sqrt{\sum_{n=1}^{\infty} \frac{1}{e^n}} = \theta_D \sqrt{\frac{1}{e-1}} \approx 0.76 \theta_D, \quad (2.13)$$

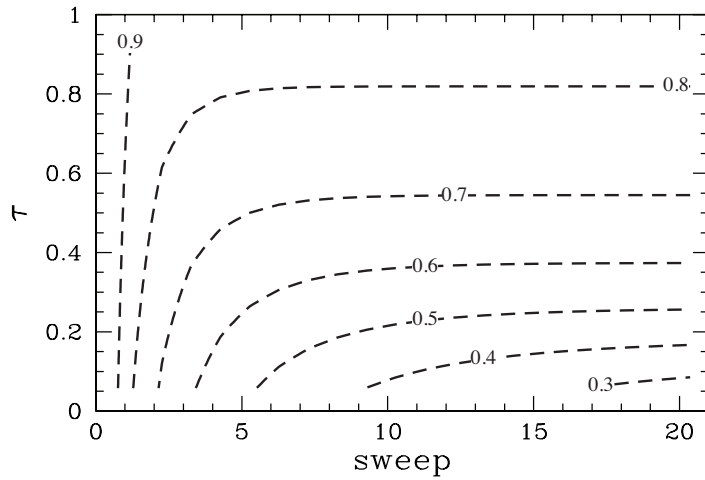
where  $\theta_D$  is the decollimation angle per ballistic step. In Sect. 3.4.1, we measure  $\theta_D$  and describe the consequences of decollimation in more detail.



**Figure 2.3:** Example of decollimation in the plane for five ballistic steps. The arrow indicates the initial influx of photons. According to the ballistic transport mechanism, photons are transported along the  $D$  most forward edges with respect to the incoming direction. The grey is as follows: with every step the photons acquire a shade that is darker. As the angle between adjacent edges is large in 2D (60 degrees on average), the bundle loses track of its original direction in only a few ballistic steps.


A visual extension of the statement given by Eq. (2.13) is shown in Fig. 2.4. From the figure, it is evident that the fraction of photons that is (ever) deflected more than  $45^\circ$  centred around the initial direction falls off sharply with the optical depth of a cell. Only for optically thin (say 0.2) cells, the fraction of photons whose deflection stay under  $45^\circ$  is lower than 0.5. In realistic cosmological simulations, the effect of decollimation results in diffusion that softens shadows behind opaque objects (e.g., filaments and halos). The diffuse radiation field will penetrate into

<sup>2</sup>Another way to look at this is that, for an isotropically emitting source, the number of nuclei that receive radiation must scale as the square (cube) of the travelled distance for two (three) dimensions.



**Figure 2.4:** The severity of decollimation as a function of two variables: the number of sweeps and the optical depth. Fractions of photons with a deflection angle within  $45^\circ$  of the initial direction are indicated by the contours. From the figure, it is evident that decollimation of photons is only potentially problematic when the cells are very optically thin

the opaque objects and ionise the high density gas inside. This results in too early ionisation of dense structures and the stalling of the ionisation front farther from the source. The problem of decollimation thus clearly necessitates the introduction of a means of transporting photons in the optically thin regime.



**Figure 2.5:** Shadow behind a dense filament irradiated by ionising radiation (indicated as a black dot). The hydrogen number density is plotted logarithmically in greyscale and ranges between  $2.8 \times 10^{-5}$  to  $0.12 \text{ cm}^{-3}$ . The contours indicate where the neutral fraction of hydrogen is 0.3. The dashed and solid white lines are for ballistic and direction-conserving transport (DCT) respectively. The black line is obtained with the  $C^2$ -ray code (Mellema et al. 2006).

## Direction-conserving transport

If the cells are optically thin, the loss of directionality introduced by ballistic transport over many steps becomes prohibitively large and we switch to *direction-conserving* transport (DCT for short). Here the radiation is confined in solid angles corresponding to global directions in space: if a photon has been emitted in a certain direction associated with a solid angle, it will remember this direction while travelling along the grid. This effectively decouples the directionality of the radiation field from the directions present in the grid (see right panel of Fig. 2.2.) The actual transport of photons still occurs along the three most forward directed Delaunay edges of the grid, where ‘forward’ is now with respect to the global directions of the solid angles<sup>3</sup>. These edges may lead to nuclei that lie outside the solid angle associated with the direction of the photons.

In this sense, we now have two types of angular resolution, in our method. The first is related to the size of the solid angle in which radiation is confined spatially, which is set by the number of Voronoi neighbours of a typical nucleus (15.54... in 3D). We call this *ballistic resolution* and emphasize that it depends solely on the nature of the Delaunay triangulation and as such, is not adjustable. The second type of angular resolution is set by the global division of the sky into arbitrarily many directions (not necessarily constant along the grid), and we call this the *directional resolution*. For the tests presented here, we use 40 directions (a directional resolution of 40), implying a solid angle of  $\pi/10$  sr for each unit vector.

A technical consequence of the approach sketched above is the need to divide space into equal portions of solid angle whose normal vectors are isotropic. There are many ways to divide the unit sphere into equal patches, but the requirement of isotropy in general cannot be met exactly. Another drawback of the introduction of global directions in the simulation is that they will give rise to artifacts much like those observed in hydrodynamical simulations on regular grids. This must in turn be counteracted by the randomisation of these global directions at appropriate time intervals, which makes DCT computationally relatively expensive.

For an example of the improvement of DCT over ballistic transport in a realistic case see Fig. 2.5. The contour for DCT shows a sharp ‘shadow’ (in good agreement with the result obtained with C<sup>2</sup>-ray (Mellema et al. 2006)) where the dense filament is left neutral whereas ballistic transport results in a more diffuse, softer, shadow. Moreover, the ballistic ionisation front stalls with respect to the DCT result. This is another symptom of spurious diffusion: the positive radial component of the diffuse radiation is smaller than it should be, resulting in more ionisations close to the source and less flux into the ionisation front.


As it is described above, DCT makes use of the most straightforward directions as prescribed by ballistic transport. The most straightforward directions are computed using the directions of the Delaunay edges along which the radiation streamed into the cell. More specifically, the procedure is as follows: Whenever radiation travels from one vertex using DCT, one first looks at the direction bin in which the radiation is stored. This direction bin has been associated with the Delaunay edge closest to it in an angular sense. The radiation is now divided over the three most straightforward neighbours of that closest edge.

This procedure does not maximally benefit from the accurate information concerning the directionality of the radiation field that has been stored in the direction bins. The largest loss of

<sup>3</sup>DCT can thus be viewed as ballistic transport with complete memory of direction.

information occurs at the moment when the high-resolution direction bin is associated to one of the (on average  $15.54 \dots$ ) Delaunay edges. It would be more sensible to compute the three most straightforward neighbours directly from the direction bin.

We have implemented this slightly costlier possibility and performed a test which shows that it leads to a substantial sharpening of shadows compared to default DCT. The test consists of a box of 8.2 kpc filled with a homogeneous hydrogen gas with number density  $n_{\text{H}} = 10^{-3} \text{ cm}^{-3}$ , which is initially fully neutral with a temperature of 100 K. A homogeneous, dense (200 times denser than the background) spherical slab with radius 0.56 kpc sits at a distance of 0.8 kpc in the x-direction from the source. The source has a  $10^5$  K black body spectrum and emits  $5 \times 10^{48}$  ionising photons per second. We have checked that the results are converged for a time step of 0.05 Myr. Fig. 2.6 shows contours of the ionisation front at 10, 30, 100, 200 and 500 Myr for simulations where the straightest directions are calculated from the Delaunay edges (grey contours) and from the direction bins (black contours). Both simulations use 42 direction bins. The shadows are much sharper when the straightest directions are calculated from the direction bins.



**Figure 2.6:** Slice through the domain at  $z = 4.1$  kpc for the shadowing test described in the text. The dense cloud is denoted in black. Grey contours show the position of the ionisation front at 10, 30, 100, 200 and 500 Myr for simulations where the straightest directions are calculated from the Delaunay edges. Black contours show the ionisation front at the same times but for straightest directions calculated from the direction bins. In this test, 42 direction bins are used. This image is taken from [Paardekooper \(2010\)](#), where a more complete set of tests (with different numbers of direction bins) can be found.

### 2.1.5 Combined transport

The three means of transport described above are applied simultaneously. If the total optical depth  $\tau_{\text{tot}}$  for a given cell is caused by multiple extinction processes

$$\tau_{\text{tot}} = \sum_i \tau_i, \quad (2.14)$$

the fraction,  $f_i$ , of the incoming ray of photons that will be removed by process  $i$  is given by

$$f_i \equiv \frac{\tau_i}{\tau_{\text{tot}}}. \quad (2.15)$$

Depending on the physical nature of this process, the photons are either removed from the bundle (to heat up the medium) or redistributed isotropically (or with some directionality) using diffuse transport. The remaining photons (those that are not removed from the ray) are transported with either ballistic or direction-conserving transport, depending on the total optical depth of the cell (ballistic if  $\tau_{\text{tot}} \geq 1$  and DCT if  $\tau_{\text{tot}} < 1$ ). The fraction of the incoming photons that is treated with one of the three transport methods is thus fully determined by the grid.

## 2.2 PARALLELIZATION

The SIMPLEX code is parallelized for distributed memory machines using MPI. Due to the local nature of the SIMPLEX transport, communications between processors are limited to the boundaries. The grid-construction is parallelized by dividing the domain into connected regions containing roughly equal number of vertices using space-filling Hilbert curve, which is also employed in other methods without a regular grid (Shirokov & Bertschinger 2005; Springel 2005, 2010).

Every processor constructs only that portion of the grid which belongs to its local points, plus a layer of border-points from neighbouring processors. This layer connects the subdomains that exist on separate processors. To determine if this layer is thick enough to guarantee a unique and unambiguous grid, the *empty circumsphere* property (also *Delaunay criterium*) is used. This principle states that the circumsphere that passes through each vertex of a Delaunay simplex, is devoid of vertices.

Communication of the radiation that crosses processor boundaries, happens at the end of every radiative time step. The resulting chemistry is done locally, so no physical quantities other than photons need to be communicated between processors.


### 2.2.1 Scaling tests

We have conducted three scaling tests that show how our implementation performs when the load per processor is increased (Fig. 2.7). The number of processors is increased for the same problem (left panel of Fig. 2.8, left panel), or the number of processors is increased at a constant load per processor (Fig. 2.8, right panel).

The physical set-up is a simple homogeneous medium illuminated by a single source (inclusion of more sources would not change the results) in the centre of the domain. All the simulations were conducted using AMD Opteron 246 64Bit CPUs of 2.6 Ghz with 4 GB of memory per node. Although we only had 8 nodes available for these tests, they give a general idea of the scaling.

From Fig. 2.7 we see that computation time increases linearly with the number of grid points  $N$  for most components of the simulation. The triangulation algorithm (which scales as  $\mathcal{O}(N \log N)$ ) and the combined transport scheme are exceptions. The latter will always have a computation time between that of ballistic (lower limit) and direction conserving transport (upper limit). Where it lies depends on the fraction of the cells that are optically thin (and thus treated with DCT).

The total computation time is dominated by that of the radiation transport, which takes



**Figure 2.7:** 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).

roughly one order of magnitude more time. This emphasizes that using the triangulation as the basis for radiation transport accelerates the whole computation because the triangulation provides many geometric quantities that are otherwise computationally intensive to calculate.


The left panel of Fig. 2.8 shows the strong scaling properties of the SIMPLEX algorithm. Here we simulate the same physical problem as before, but this time the number of grid points is held constant at  $128^3$ . Ideally, doubling the number of processors would halve the total execution time. If we were then to plot the product of the total computation time and the number of processors, the result would be a curve with zero slope. Deviations from this ideal case indicate that the parallelization itself and the resulting communications use up computational resources.

From the figure we see that the triangulation is the only component that deviates considerably from the ideal case. This can be understood directly when we realize that dividing the domain into more volumes leads to extra boundaries, and this increases the total number of points that need to be triangulated. This is not a serious issue because the computation time of the triangulation remains an order of magnitude smaller than the radiative transfer components, which is reflected in the slope of the accumulated time (dot-long dashed line).

The weak scaling properties of the SIMPLEX algorithm are shown in the right panel of Fig. 2.8. In this test we keep the number of vertices per processor constant while increasing the number of processors. In the ideal case, the total computation time stays constant and deviations therefrom can be directly related to extra work due to the parallelization of the problem.

The radiation transport components remains constant (within a few percent), showing marginal increase due to extra communications. The main difference between this test and the previous one, is that now the effects of a decreasing number of vertices per processor is eliminated from the equation. One effect remains on the triangulation time, however. The parallelization results in a change of the ratio of processors that share three, four, five or six walls with other processors






**Figure 2.8:** *Left:* Simulation time as a function of the number of processors for a constant number of grid points. The number of grid points is  $128^3$ . Most components of the simulation scale very close to linear (resulting in a constant curve) as the number of processors increases. An exception is the triangulation algorithm, due to the fact that every processor needs to triangulate extra points in the boundary between processors. *Right:* Simulation time as a function of the number of processors for a constant load per processor. The number of grid points at each processor is  $64^3$ . The computation time of the radiation transport components shows a marginal increase as the problem size gets bigger due an increase in the number of communications involved. The computation time of the triangulation increases more for reasons similar to those explained for the strong scaling (see text for details).

(and thus need to compute and communicate these boundaries). This effect is smaller than in the case of strong scaling because the number of boundary vertices is small compared to those in the domain.

## 2.3 SUMMARY

We have introduced our radiative transfer algorithm `SIMPLEX`. The `SIMPLEX` method naturally adapts its resolution, does not scale computationally with the number of sources, and is parallelized for distributed memory machines. For the number of processors used in our scaling test, the parallel scaling properties are close to ideal. It is therefore well-suited for applications in cosmological problems (for a realistic application see Chapter 7). Its scope of application is not limited to large scales, however. As we will see in Chapter 9, the virtues of `SIMPLEX` are also useful in the context of ionisation calculations in the complicated geometries that arise in wind-wind interaction regions in massive binaries.



**Figure 2.9:** Geometry of computational cell with radius  $r$ , as used in the derivation of Eq. (2.12). The path-length through the cell  $l$  depends on the impact parameter  $b$ .

## 2.A DERIVATION OF THE ESCAPE FRACTION OF DIFFUSE RECOMBINATION RADIATION

Because the recombination photons originating in a cell will also be absorbed by the gas in that cell, we use a sub-grid description for the escape fraction of such radiation. Assuming a constant emissivity of recombination photons,  $\eta$ , we can derive escape fractions as a function of optical depth of the cell. We have chosen the cell to be spherical with radius  $r$ , an assumption that is appropriate for SPH particles and a fair approximation of a typical Voronoi cell. Fig. 2.9 shows a schematic depiction the geometric variables used in the following derivation. With these assumptions, the total intensity of emitted recombination photons in the cell,  $I_e$ , is given by

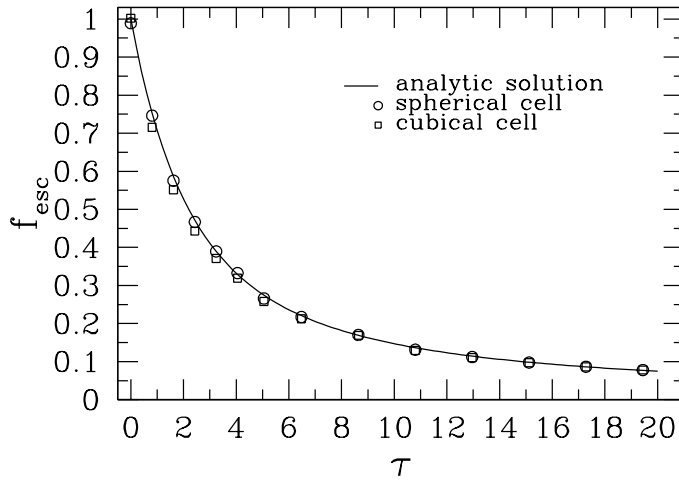
$$I_e = \frac{4\pi r^3 \eta}{3}. \quad (2.16)$$

The emerging intensity along a ray  $I_{\text{obs}}(l)$  through a medium with source function  $S \equiv \eta/\chi$  is given by

$$I_{\text{obs}}(l) = S(1 - e^{-\tau}), \quad (2.17)$$

where  $\tau = \chi l$  is the optical depth along the ray,  $\chi$  is the opacity and  $l$  is the path length through the sphere. We can find the emerging intensity by adding contributions of all lines of length  $l$  perpendicular to the impact parameter  $b$  (see Fig. 2.9). Each intensity  $I_{\text{obs}}(l)$  contributes  $2\pi b I_{\text{obs}}(l)$  to the total emerging intensity. This can be seen as the cylinder traced by rotating  $l$  over  $2\pi$  around the centre of the sphere. To obtain the total emerging intensity we integrate over impact parameter  $b$ :

$$\begin{aligned} I_{\text{obs,tot}} &= 2\pi \int_0^r I_{\text{obs}}(l(b)) b db \\ &= \frac{2\pi r^2 \eta}{\chi} \int_0^{\pi/2} (1 - e^{-2\chi r \sin \theta}) \cos \theta \sin \theta d\theta, \end{aligned}$$



**Figure 2.10:** The fraction of ionising recombination radiation that is able to escape from a computational cell for different optical depths. The solid line represents the analytical solution of Eq. (2.19), while the symbols represent simulations of cells with different geometries.

where we have used that  $l = 2r \sin \theta$  and  $b = r \cos \theta$ . This expression can be integrated analytically, yielding

$$I_{\text{obs,tot}} = \frac{\pi\eta}{2\chi^3} (e^{-2r\chi}(1 + 2r\chi) - (1 - 2r^2\chi^2)). \quad (2.18)$$

Normalizing Eq. (2.18) by Eq. (2.16) results in an escape fraction of

$$f_{\text{esc}}(\tau) = \frac{3}{8\tau^3} (e^{-2\tau}(1 + 2\tau) - (1 - 2\tau^2)), \quad (2.19)$$

where  $\tau = \chi r$  is the optical depth in the cell. This function is plotted as the solid line in Fig. 2.10. We have performed explicit numerical experiments using `SIMPLEX` to test to what extent the escape fraction depends on the geometry of the cell. To this end we have calculated escape fractions of radiation from a computational domain shaped as a single cubic cell with sides of length  $L$  (appropriate for grid-based methods) and a spherical cell (appropriate for particle-based methods). This has been done at several fixed optical depths where we have used  $\tau \equiv \chi L$  for the cubic domain. The results of these tests are shown in Fig. 2.10 together with the analytical result of Eq. (2.12). The excellent agreement between the simulation with spherical geometry and the analytical solution gives confidence in the correctness of the `SIMPLEX` radiative transfer method to be able to compute the transport of photons in a wide range of optical depths. The simulation of a cubic cell shows that the geometry of the cell does not have a significant influence, escape fractions differ no more than 4% between the simulations. This difference is largely due to the choice for the optical depth through the cubic cell which does not take directionality into account.

These results show that the majority of recombination radiation escapes from a cell if the optical depth is lower than  $\simeq 2$ . Recombination radiation will therefore play a role only in cells where a significant number of recombinations takes place, that is, in cells that are highly ionised and, consequently, have an optical depth much lower than 2. We conclude that the majority of ionising recombination radiation is *not* absorbed in the computational cell where it was produced, and that explicit treatment of such radiation may be important in realistic applications.

## CHAPTER 3

---

# Systematic effects in the SimpleX algorithm

C. J. H. Kruip, J.-P. Paardekooper, B. J. F. Clauwens & V. Icke

Part of this chapter consists of work that has been published in  
*Astronomy & Astrophysics* 515, A78, 2010

**W**E verify whether the SimpleX radiative transfer algorithm conforms to mathematical expectations and develop both an error analysis and improvements to earlier versions of the code. Our analysis leads us to conclude that it is possible to transport particles such as photons in a physically correct manner with the SimpleX algorithm. This requires the use of weighting schemes or the modification of the point process underlying the transport mesh. We explore and apply several possibilities.

### 3.1 INTRODUCTION

In the previous chapter, we introduced the fundamentals of the SIMPLEX algorithm. Three transport methods (diffuse- ballistic- and direction conserving transport) and their range of applicability have been discussed. All three transport methods are intimately connected to the transport mesh which consists of Voronoi-Delaunay structures. We have discussed several advantages of using Delaunay triangulations as the computational mesh for radiative transport. SIMPLEX transport is very fast due to its local nature (cells are only connected to neighbouring cells) and the resolution is highly adaptive. We have, however, not yet assessed the possible disadvantages of our approach. In this chapter we take a closer look at the geometrical properties of our transport algorithm to see if it conforms to our expectations. In other words, we want to obtain upper bounds to the error properties of our method.

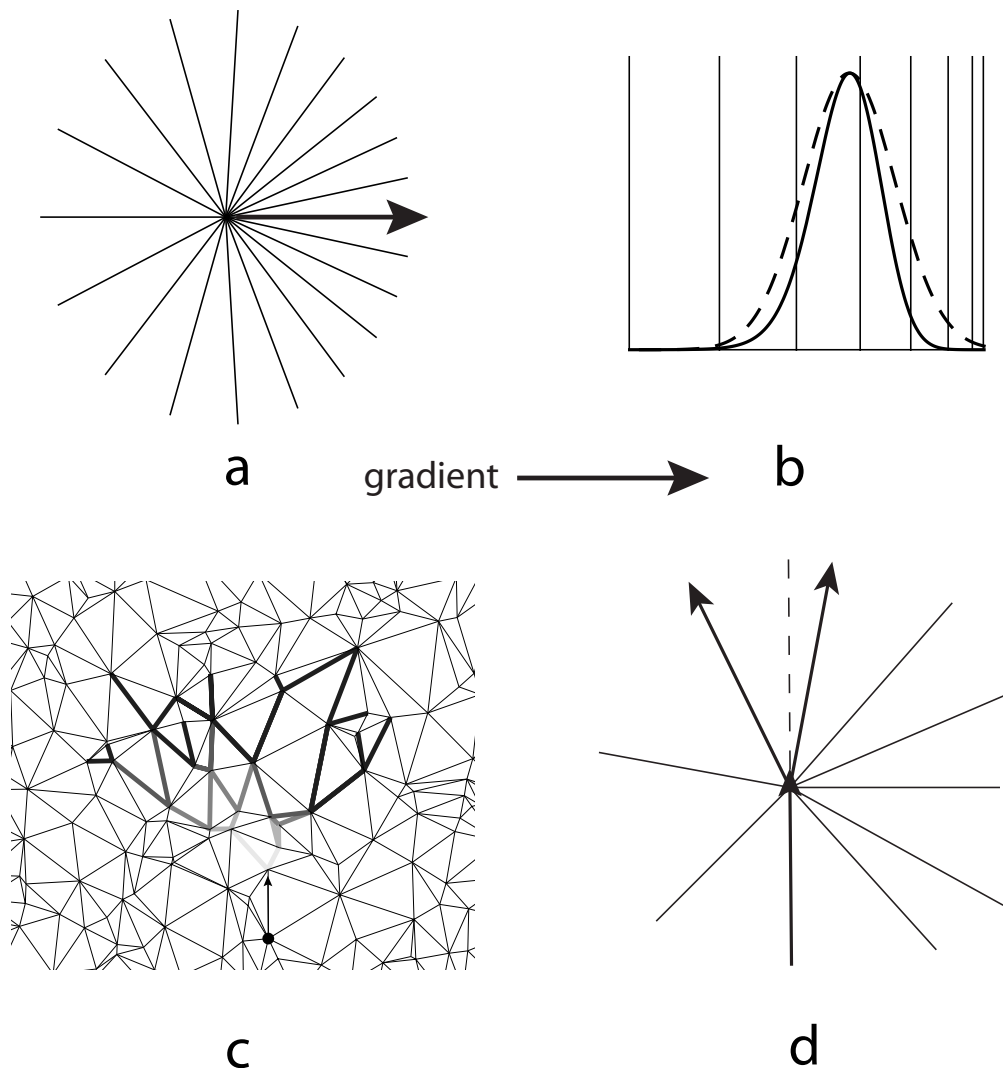
For the vast majority of numerical methods, errors are measured by comparing with a fiducial run of the code. This is usually a simulation wherein many more time steps, and/or a higher spatial resolution are used than would normally be feasible. Convergence of the result is studied which is accepted as the correct solution, at least within the limitations of the method. For SIMPLEX we cannot perform a similar convergence test<sup>1</sup>. For a spatial resolution that is significantly higher than the resolution dictated by the local mean-free-path length of the photons, a number of effects (described in Sect. 3.2) tend to make the radiation field more diffuse, *decreasing* instead of *increasing* accuracy. This property of SIMPLEX is not a weakness but a direct consequence of the fact that the method uses a *physically* motivated mesh, wherein deviations from its natural resolution can cause a deterioration in the solution. Because of its mathematically transparent nature, the SIMPLEX algorithm has the advantage that one can assess its error properties analytically. The resulting prescriptions are quite general, and can be applied to different regions of parameter space, an advantage over the ‘numerical converge approach’ usually applied in the error analysis of radiative transfer methods.

In this chapter we analyse the SIMPLEX algorithm in several idealized test problems using both analytical and numerical techniques. In all the tests, we introduce gradients in the number density of vertices which result in local anisotropies of the Delaunay edges. We study the effects of these local anisotropies on diffuse-, ballistic- and direction conserving transport. We will demonstrate that local anisotropy in the Delaunay triangulation introduces systematic errors that manifest themselves as four distinct effects: diffusive drift, diffusive clustering, ballistic decollimation, and ballistic deflection. For reasons of transparency we now briefly introduce the four systematic effects, postponing their rigorous treatment to later sections. Figure 3.1 shows a cartoon representation of the four effects. In the figure, the gradient in the number density of vertices points to the right.

- Diffuse drift: the systematic streaming of photons into a region with higher number density of vertices. This happens because locally more edges point toward that region. Because of diffuse drift, a cloud of photons will drift into the direction of the gradient.
- Diffuse clustering: the clustering of photons in regions with fluctuations in the number density of vertices. Clustering occurs because photons experience diffuse drift toward

---

<sup>1</sup>This is possible, however, if the direction-conserving transport (see Sect. 2.1.4) is used. Alternatively, several instances of the same simulation with a different random seed for the mesh construction can be averaged to obtain error estimates as well.



**Figure 3.1:** Cartoon description of the four systematic effects introduced by local anisotropy in the Delaunay mesh. Diffuse drift (a) is the systematic streaming of photons into a region with higher number density of vertices. This happens because more edges point toward that region. Diffuse clustering (b) happens because photons spread less easy in regions with fluctuations in the number density of vertices. The solid line shows the distribution of photons in a gradient whereas the dashed line shows the same distribution for a homogeneous vertex number density. The vertical lines indicate the vertex-to-vertex distance which decreases from left to right in the case of a gradient. Ballistic decollimation (c) is the loss of initial direction due to several ballistic steps. This also happens in a homogeneous Delaunay grid (see Fig. 2.3). Ballistic deflection (d) happens when photons move perpendicular to a gradient in the vertex number density using ballistic transport. The photons curve into the region of lower number density of vertices (to the left in this cartoon).

higher density regions and have trouble escaping them.

- **Ballistic Decollimation:** the widening of a beam of photons due to deflections from the original direction in ballistic transport. This effect has already been treated in Sect. 2.1.4.
- **Ballistic Deflection:** the deflection of photons moving perpendicular to a gradient. This effect occurs because the most straightforward neighbours in ballistic transport are affected by the gradient. The most straightforward neighbours to the left are (on average) fewer and make larger angles with the original direction whereas the most straightforward neighbours to the right are more numerous and make smaller angles with the original direction. The net effect is a slight deflection antiparallel to the gradient (to the left in Fig. 3.1).

### 3.1.1 Outline

We start with a general introduction into anisotropy in the context of inhomogeneous Delaunay triangulations (Sect. 3.2) and derive a measure for the *excess of edges* in the direction of the gradient (Sect. 3.2.1).

To treat the effects mentioned in the introduction, we adhere to the following structure: introduction and analytical treatment followed by solution and numerical examples. We first apply this structure to the effects on diffuse transport and then on ballistic transport. This seems the most logical choice because, as we will see below, the same solution can be applied to both diffuse drift and diffuse clustering. Remedies for ballistic decollimation and ballistic deflection turn out to be similar as well and are also described in conjunction.

After the treatment of the four systematic effects, we continue with a general discussion and summary. In Chapter 4, we demonstrate how the derived measures can be used to constrain the representation of a physical problem as a transport graph in such a way as to avoid or minimise errors.

## 3.2 ANISOTROPY AND ITS CONSEQUENCES

Notwithstanding the ‘ideal’ properties of the Delaunay triangulation, the probability density distribution of the angle between adjacent Delaunay edges is quite broad (Icke & van de Weygaert 1987; Okabe 2000, Sect.5.5.4). This ensures that even though the average triangle is the ‘fattest’ possible, many ‘thin’ triangles will occur. This situation may be changed by iteratively adapting the underlying point process in such a way that each nucleus comes to coincide with the centre-of-mass of its Voronoi region (see Lloyd 1982, for such an algorithm), resulting in so-called Centroidal Voronoi Tessellations (see e.g. Du et al. 1999). This procedure produces ‘most spherical’ Voronoi regions (hexagons if  $D = 2$ ), but is generally far too costly for practical computations in which the triangulation must be re-computed frequently. Moreover, the shifting of the nuclei implies that the connection with the physical properties of the underlying medium is no longer entirely faithful. This could be alleviated by resampling the original density field at the new (now Centroidally placed) nuclei.

Another reason to refrain from the use of Centroidal Voronoi Tessellations is the introduction of regularity and hence symmetry in the mesh. In two dimensions, for instance, the hexag-

onal Voronoi cells tend to align, introducing globally preferential directions in the transport of radiation. This will inevitably give rise to artifacts in the radiation field transported on such a mesh.

On the other hand, if only a few (typically two) Lloyd iterations are used, the most severe local irregularities disappear without regularizing the mesh too much. This approach can be used to suppress noise and artifacts in the outcome of a SIMPLEX simulation (work along these lines is also being done by Jakob van Bethlehem).

Although the Voronoi-Delaunay construction is the optimal choice for the tessellation of a random (Poisson process) point set, this may not be true when the point set is inhomogeneous or anisotropic. Inhomogeneity is, of course, the property we encounter in all practical cases. The probability distribution of the directions of Delaunay edges pointing to a Poisson nucleus is isotropic, but this is no longer the case when the distribution of the nuclei is inhomogeneous.

We must therefore face the consequences of the anisotropy bias on inhomogeneous point processes. We note in passing that anisotropy of Delaunay edges will exist *locally* even in the case of a homogeneous Poisson process. The reason is that because, due to shot noise, every instance of a random point process is locally anisotropic. The difference is that the latter anisotropy vanishes when many instances of the point distribution are averaged (or, alternatively, when several Lloyd iterations are applied to the point distribution) while the former anisotropy cannot vanish because it is inherent to the *macroscopic* distribution of nuclei.

### 3.2.1 Error measure

For a homogeneous Poisson distribution of nuclei, the expectation value for the number of Delaunay neighbours (and thus edges),  $\Lambda$ , is 6 in two- and  $15.54 \dots$  in three-dimensional space. These Delaunay edges have no preferential orientation and their statistical properties are well known (e.g., [van de Weygaert 1991, 1994](#); [Okabe 2000](#)).

We now consider an inhomogeneous distribution of nuclei. Spatial gradients then appear in the density of nuclei,  $n(\mathbf{x})$ , and the Delaunay edges connecting these nuclei are no longer distributed evenly over all possible orientations. For a given nucleus, there will be, on average, more edges pointing towards high-density regions than away from them.

This can be quantified as follows. Without loss of generality, we may assume a number density of nuclei that has a gradient in some fixed direction  $x$ , provided that the characteristic length scale of the gradient is much larger than the mean distance between nuclei (a provision we assume to be fulfilled from now on).

We take a cross-section perpendicular to the direction of the gradient through the box at an arbitrary position  $x_0$ . The resulting plane of surface  $S$  is pierced by Delaunay edges connecting nuclei on either side of the plane. The number of edges piercing the plane can be estimated as follows. The local density of edges is the product of the number density of nuclei and  $\Lambda$ . An edge is able to pierce the plane when two requirements concerning its orientation are fulfilled:

1. Its projected length must be larger than the distance between its originating nucleus and the plane.
2. It must point in the correct direction.



The expectation value of the Delaunay edge length,  $\mathcal{D}$ , is given by

$$\mathcal{D} = \xi n(x)^{-1/3}, \quad (3.1)$$

in three dimensions, where  $\xi = 1.237 \dots$  (Ritzerveld 2007, Eq. (3.22)). An edge emanating from a nucleus at position  $x$  will therefore only pierce the plane if its projected length exceeds  $\Delta x \equiv |x - x_0|$ . Consequently, the first requirement is fulfilled when

$$\Delta x \leq \mathcal{D} \cos \theta, \quad (3.2)$$

where  $\theta$  is the angle between the edge and the direction of the gradient. The orientation of these edges is random<sup>2</sup>, so we must average the cosine over a half sphere, which yields a factor of one half. As a result, the effective length of the edges is  $\mathcal{D}/2$ , which implies that only nuclei inside a slab of thickness  $\mathcal{D}$ , (centred at  $x_0$ ) contribute to the density of piercing edges.

The second requirement effectively excludes half (up to first order) of the edges because they point away from the plane. This statement is equivalent to noting that every piercing edge connects exactly two nuclei at opposite sides of the plane. By including both factors, we find that the number of piercing lines,  $N_p$ , is given by

$$N_p(x) = \frac{n(x)\Lambda\mathcal{D}S}{4}. \quad (3.3)$$

The surface density of lines piercing the slab,  $\sigma(x)$ , is now simply defined by

$$\sigma(x) = \frac{N_p}{S} = \frac{n(x)\Lambda\mathcal{D}}{4}. \quad (3.4)$$

The sought-after fractional excess,  $E(x)$ , of parallel (with respect to the gradient) over anti-parallel Delaunay edges is thus given by differencing  $\sigma(x)$  over a sufficiently small<sup>3</sup> interval  $\Delta x$  and division by  $\Lambda n(x)$

$$E(x) \equiv \frac{1}{\Lambda n(x)} \frac{\Delta\sigma(x)}{\Delta x} = \frac{\xi}{4n(x)^{4/3}} \frac{\Delta n(x)}{\Delta x}, \quad (3.5)$$

where we used Eq. (3.1) to eliminate  $\mathcal{D}$ . The excess of edges pointing towards the higher density regions may have a significant effect on quantities transported along these edges. In the next three sections, we describe and quantify these effects. We note in passing that in many practical applications (see also Chapter 4) the point density  $n(x)$  of nuclei is taken to be proportional to a power  $\alpha$  of the mass density  $\rho(x)$

$$n(x) \propto \rho(x)^\alpha. \quad (3.6)$$

In that case, Eq. (3.5) becomes

$$E(x) = \frac{\xi}{4n(x)^{1/3}} \frac{\Delta \log n(x)}{\Delta x} = \frac{\alpha\mathcal{D}}{4} \frac{\Delta \log \rho(x)}{\Delta x}. \quad (3.7)$$

If  $\alpha = 3$ , the length  $\mathcal{D}$  is proportional to the mean free path of the photons (Ritzerveld & Icke 2006). In many respects, this is the ‘ideal’ case, because it is in fact ‘transport-homogeneous’ as experienced by the photon, every step being of equal optical depth. However, Eq. (3.7) shows that this ‘ideal’ case has a gradient asymmetry that is three times stronger than is the case when  $n(x) \propto \rho(x)$ .

<sup>2</sup>The orientation may be correlated with the direction of the gradient, but we neglect this at the moment because it would only influence our results to second order.

<sup>3</sup>We emphasize the discrete nature of this derivative by noting that  $\Delta n = \frac{dn}{dx} \Delta x = \mathcal{D} \frac{dn}{dx} = \xi n(x)^{-1/3} \frac{dn}{dx}$  etc.

### 3.3 EFFECTS ON DIFFUSIVE TRANSPORT

In Sect. 2.1.4, we stated that for diffuse transport in a homogeneous distribution of nuclei, the radiation propagates spherically away from a source. On a graph corresponding to an inhomogeneous distribution of nuclei, however, this is not the case. The spreading of the transported quantity will no longer be spherical anymore for two reasons:

1. The Delaunay edges are shorter when the nuclei are spaced more closely together.
2. The orientation of the Delaunay edges is no longer isotropic: more edges point towards the overdense regions.

#### 3.3.1 Physical slow down

The first reason reduces the transport velocity and can be interpreted as a physical phenomenon. If we were to identify the length of a Delaunay edge with the local mean free path of the transported quantity (e.g. photons), the shorter edges would simply express that we have entered a region of increased optical depth where it takes a greater number of mean free path lengths to traverse a given physical distance. It has been shown (Ritzerveld & Icke 2006) that identifying the average Delaunay edge length with the local mean free path of the relevant processes is a natural choice when constructing the triangulation, and the observed behaviour is therefore both expected and physical.

#### 3.3.2 Diffuse drift

The second reason, quantified by Eq. (3.5), is an artifact of the Delaunay triangulation itself and causes unphysical behaviour. When too many edges are pointing into the overdense regions, the transported particles are deflected into those regions and the direction of propagation tends to align with the gradient (see Fig. 3.2 for an example in the plane). We call this effect *drift*, and now proceed to quantify its consequences for diffusive transport. A photon scattering at a nucleus has the following probabilities of moving in the dense (subscript  $d$ ) or underdense (subscript  $u$ ) direction:

$$\begin{aligned} p_d &= \frac{1}{2} + \frac{E(\mathbf{x})}{2} \\ p_u &= \frac{1}{2} - \frac{E(\mathbf{x})}{2}. \end{aligned} \quad (3.8)$$

The expectation value  $d_D$  of the drift per scattering event is therefore proportional to the  $E(\mathbf{x})$  part of an outgoing edge in the direction of the dense region. The multiplication factor, which can be found by integrating over a half-sphere, is  $1/2$  since not all edges point exactly to the right. This gives

$$d_D = \frac{\mathcal{D}}{2} E(\mathbf{x}), \quad (3.9)$$

in which  $\mathcal{D}$  is again the expectation value of the Delaunay edge length (see Eq. (3.1)).

We consider a scattering experiment where a number of photons are placed at one position of a triangulation with a density gradient. We expect the photons to diffuse outward with an

ever decreasing radial velocity (the distance travelled,  $L$ , scales with the root of the number of steps) while simultaneously drifting towards the dense region with a constant velocity. There comes a time (or distance) at which the drift is equal in magnitude to the diffusion radius. We define the drift length,  $L_{\text{drift}}$ , as the scale on which the diffusion distance is equal to the distance travelled through drift. Roughly speaking, diffusion dominates for  $L < L_{\text{drift}}$  and drift dominates for  $L > L_{\text{drift}}$ . Setting the diffusion distance equal to the distance travelled through drift,

$$\mathcal{D} \sqrt{N_E} = \frac{\mathcal{D}}{2} E(\mathbf{x}) N_E, \quad (3.10)$$

we find that the number of steps at equality is given by

$$N_E = 4/E^2(\mathbf{x}). \quad (3.11)$$

Using Eq. (3.1) and (3.5) gives an equality length of


$$L_{\text{drift}}(\mathbf{x}) = \frac{8n(\mathbf{x})}{|\nabla n(\mathbf{x})|}, \quad (3.12)$$

independent of  $\mathcal{D}$ . Therefore we have found a *local* expression for the relative importance of the unphysical drift.

For a given density distribution  $n(x)$ , the length  $L_{\text{drift}}(\mathbf{x})$  can be evaluated everywhere<sup>4</sup>. This parameter can be interpreted as follows. We define the minimum  $L_{\text{drift}}(\mathbf{x})$  for all  $\mathbf{x}$  to be  $N$  times the box side length, such that the drift can be at most  $1/N$  of a box side while the radiation scatters throughout the box. Regardless, we must ensure that  $L_{\text{drift}}(\mathbf{x})$  is much longer than the box side length. Therefore, even if the density contrast is very small, it must not fluctuate too severely.

We also note that it does not help to increase the number of nuclei, since both  $n(\mathbf{x})$  and  $\nabla n(\mathbf{x})$  in Eq. (3.12) scale with that number. Using more nuclei reduces the anisotropy per nucleus,

<sup>4</sup>It may surprise the reader that values for  $E(\mathbf{x})$  and  $\mathcal{D}$  are taken to be local, while the argument seems to involve a domain in which these values could change. What happens is that we adopt the *local* values for the *whole* domain in Eq. (3.10).



**Figure 3.2:** Schematic example of a nucleus and its edges subject to a gradient in the number density of nuclei in the positive  $x$ -direction. More edges point toward the over-dense region (along the gradient). If every edge were to transport an equal number of photons to neighbouring nuclei, the anisotropy of outgoing edges would produce an unphysical net flow along the gradient (indicated by the arrow, which is the vector sum of the edges scaled down by roughly a factor of three).

Eq. (3.5), but because the number of steps to be taken increases accordingly, the effect simply adds up to the same macroscopic behaviour. Only the shape of  $n(\mathbf{x})$  determines the magnitude of  $L_{\text{drift}}(\mathbf{x})$ .

To determine the resulting constraint on the mesh, we consider the case in which  $L_{\text{drift}}(\mathbf{x})$  does not depend on position. Setting  $L_{\text{drift}}(\mathbf{x}) = \text{constant}$  in Eq. (3.12) defines an exponential density distribution, where the anisotropy is smeared out maximally over the domain. In this case, if we wish the drift to be less than  $1/8$  of the box length, the density contrast must be less than a factor  $e \approx 2.71 \dots$  implying that the restrictions set by our isotropy demands are rather stringent.

### 3.3.3 Diffuse clustering

We have seen that the spurious drift for diffuse scattering places restrictions on the density contrasts that can be simulated by the plain implementation of SIMPLEX introduced in Sect.(3.3.2). We have assumed that the scale of the density fluctuations is comparable to the box size. This is not a restriction: if we are interested in a case where the density fluctuations are of a much smaller scale than the box size, we can just place an imaginary box around each density fluctuation and use all the quantitative results from above. In doing so, we see that for any given ‘snapshot’ too many photons will be present in the local overdense regions, but if there are no overall density contrasts on large scales, there will be no significant macroscopic drift. The effect of the local drift on small scales can, however, still influence results on a large scale. Not only does the drift influence the average position of the photons, it also influences the standard deviation around this average. Isotropic scattering maximizes the spreading of photons, but in an extreme case where, for example, at every nucleus 90% of the photons move in the same general direction, they will stick together for a longer period and cause the size of a ‘light cloud’ to grow more slowly. This effect occurs if we have a highly fluctuating density field on small scales, a case in point being the simulation of the filaments of large-scale cosmic structures.

We consider a density distribution that is homogeneous in the  $y$ - and  $z$ -direction but highly fluctuating in the  $x$ -direction. For the probabilities to travel into the dense or underdense regions, we again use Eq. (3.8) but with the difference that now  $d$  and  $u$  no longer denote global directions. We describe the transportation process by a binomial distribution. If there are no large scale-density contrasts, the drift is zero on average, but for the standard deviation we find

$$\begin{aligned} \sigma(N) &= \langle \sqrt{N p(1-p)} \rangle \\ &= \left\langle \sqrt{N \left( \frac{1}{2} + \frac{E(x)}{2} \right) \left( \frac{1}{2} - \frac{E(x)}{2} \right)} \right\rangle \\ &= \frac{\sqrt{N}}{2} \langle \sqrt{1 - E^2(x)} \rangle, \end{aligned} \quad (3.13)$$

where  $N$  is the number of sweeps. For small values of  $E$ , this reduces to


$$\sigma(N) \approx \frac{\sqrt{N}}{2} \left( 1 - \frac{\xi^2}{32} \langle |n'(x)n(x)^{-4/3}|^2 \rangle \right). \quad (3.14)$$

This factor is always smaller than unity, indicating a reduction in the spreading of the photons. This effect can be significant for small-scale fluctuations with either a very high amplitude or a

very short length. The effect slowly becomes smaller if the number of nuclei is increased. The value of  $E(x)$  comes close to unity or exceeds unity only if the characteristic density fluctuation length is smaller than a Delaunay length, which we had pointedly excluded.

### 3.3.4 Correcting diffuse transport: weighting schemes

A straightforward solution for the problems addressed in Sects. 3.3.2 and 3.3.3 (diffuse-drift and clustering, respectively) is to assign weights  $w_i$  to the edges emanating from a given nucleus in such a way that the anisotropy vanishes. This means that the fractions of the quantity transported to the neighbours are no longer equal to  $1/N$  but directly proportional to the solid angle that the corresponding Voronoi face spans. We refer to Fig. 3.3 for an example in two dimensions, the three-dimensional case being analogous. We explore this possibility and im-



**Figure 3.3:** Solid angles for two edges in the case of the Voronoi weighting scheme (cell 3) and the icosahedron weighting scheme (cell 1).

plement three different weighting schemes in our method, each with their specific virtues and drawbacks:


1. Voronoi weights: based on the natural properties of the triangulation. Its advantages are that it is automatically and ‘naturally’ adapted to the physics of the transport problem. Its disadvantage is the statistical noise inherent to the procedure.
2. icosahedron weights: based on a division of the unit sphere using the icosahedron. Its advantages are that it is flexible (the procedure does not depend on the type of triangulation) and that it can be easily refined. Its main disadvantage is that it is computationally more expensive.
3. Distance weights: based on the distance to a neighbour squared. Its main advantage is that it is very fast. Unfortunately we have only empirical evidence of its correctness.

### Centre of gravity weighting

In all three weighting methods, the sum of the weights equal unity, guaranteeing conservation of photons. In addition, the *vector* sum of the weighted Delaunay edges emanating from a nucleus should be zero to conserve momentum of the radiation field. We obtain this by adjusting the weights of the (two) edges most parallel and anti-parallel to the ‘centre of gravity’,  $\Psi$ , of the nucleus defined by

$$\Psi = \sum_i w_i \hat{r}_i, \quad (3.15)$$

such that in that direction the magnitude of  $\Psi$  vanishes. Here  $\hat{r}_i$  denotes the direction of the  $i$ -th Delaunay edge (see Fig. 3.4). The vector sum of all (weighted) Delaunay edges (grey arrows) is non-zero (short-dashed arrow). The Delaunay edges  $j$  and  $i$  are most parallel and anti-parallel to this vector sum respectively. Their weights are adjusted such that the vector sum vanishes in its current direction. To this end, the vector sum and its reflected counterpart (long dashed arrow) are projected on  $j$  and  $i$ , indicated with the smaller black arrows. Half of this projection is added to edge  $i$  and half is subtracted from  $j$ . The resulting vector (dot-dashed arrow) exactly cancels the vector sum. The adjustment of weights can still result in a non-zero (although smaller) vector sum in some other direction. The above procedure is repeated iteratively until the norm of the vector sum is smaller than a predefined tolerance. In our experience, this procedure reaches a (relative) tolerance of  $10^{-7}$  within twenty iterations for the Voronoi method and about ten for the icosahedron and distance weighting schemes.



**Figure 3.4:** Geometry of the COG weighting procedure. See text for details.

### Voronoi weighting

The Voronoi weighting scheme uses the faces of the Voronoi cells to calculate the  $w_i$  for each nucleus by estimating the solid angles subtended by the walls of the Voronoi regions, normalised to unity. The weights are then

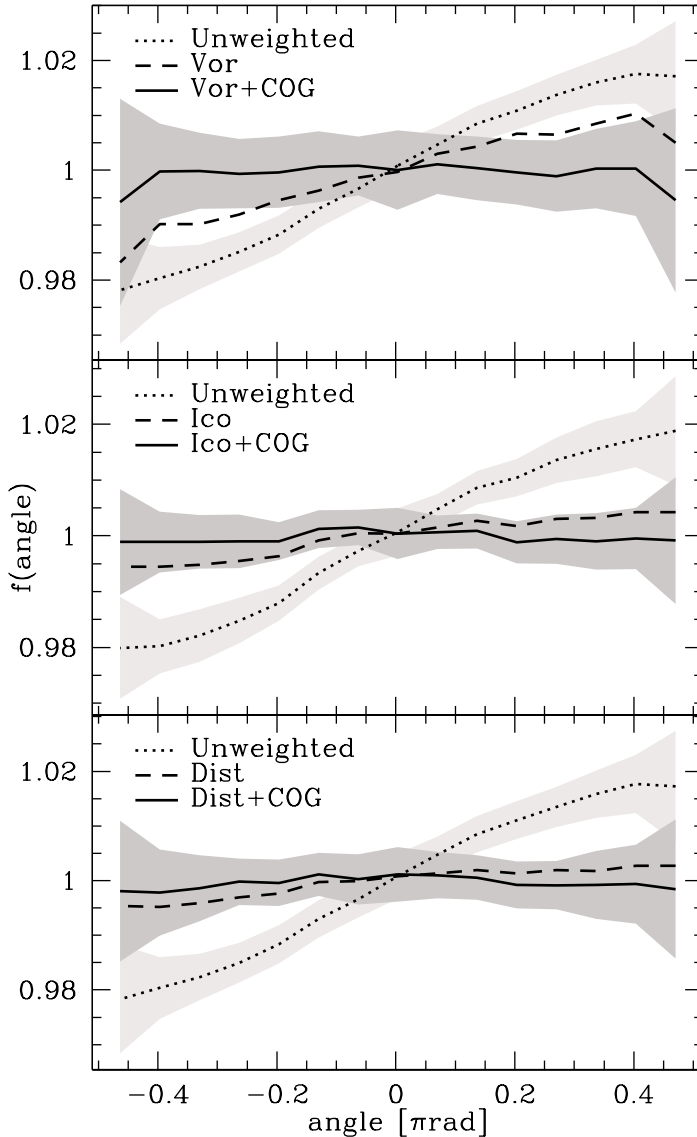
$$w_i = \frac{S_i d_i^{-2}}{\sum_j S_j d_j^{-2}}, \quad (3.16)$$

where  $S_i$  denotes the surface of the Voronoi face perpendicular to the Delaunay edge connecting the current nucleus and its  $i$ -th neighbour, and  $d_i$  is the length of that edge (the Voronoi face cuts the Delaunay edge through the middle,  $0.5d_i$  being the distance from the nucleus to the face). We note that Eq. (3.16) is an approximate expression neglecting projection effects for large angles. With some computational effort, this estimate could be refined. The solid angle subtended by a Voronoi face depends on the size of the cell but also on the precise position of the nuclei. Referring to Fig. 3.3, we see that cell ‘3’ has a relatively small surface because its nucleus is close to that of its neighbour ‘4’ and its nucleus is slightly further from the central nucleus than that of ‘4’. If that last statement were reversed, neighbour ‘4’ would have the smaller surface and would thus get the smaller weight. This property may seem harmful at first but it is of a stochastic nature, meaning that some noise will be introduced but no systematic error.

To demonstrate our method, we construct a Delaunay triangulation of  $10^5$  nuclei in three-space with a linear gradient in  $n(\mathbf{x})$  along the horizontal direction which runs from 0.005 on the left to 0.995 on the right. As a measure of the anisotropy in orientations of Delaunay edges, we take the angle between an edge and the direction of the gradient and plot the number of edges in an angular bin (see Fig. 3.5). Because we are interested in the relative deviation from a horizontal line (which would correspond to the isotropic situation), the results are given as a fraction of the average,  $f(\theta)$ . The data displayed in the topmost panel corresponds to the weighting schemes described above, which can be thought of as a basic correction plus a refinement thereof (the ‘centre of gravity’ correction). We show the effects on the angular distribution of the edges as we apply these corrections cumulatively.

The initial anisotropy apparent from the inclination of the line labeled ‘unweighted’ is reduced significantly (to about half its original value) after applying the Voronoi weighting alone (dashed line). After the application of the ‘centre of gravity’ correction, the scatter around the isotropic value of unity is below 0.25%, except at the outer edges where the normalisation dramatically increases the errors.

Apart from the anisotropy caused by the gradient, the triangulation itself shows some noise of order 0.5% (the light grey areas around the uncorrected lines in the three panels having a width of one standard deviation). After all weighting has been applied, some noise does remain (a unit-standard deviation area is shaded in dark grey). This noise is related to features of the triangulation itself. If the triangulation itself has more edges in a certain direction, this cannot be corrected with a *local* weighting scheme because it is a *global* property subject to chance. This noise can be reduced by either placing more points or by constructing several instances of the same triangulation and averaging the results. Because all the information needed for the Voronoi weighting scheme is intrinsic to the tessellation and its triangulation, the associated computational overhead is potentially small. Unfortunately, in most tessellation software, the areas of the Voronoi walls are not computed with the other properties of the tessellation. Calcu-



**Figure 3.5:** Fraction of edges,  $f(\theta)$ , with angle  $\pi - \theta$  with respect to the direction of the linear gradient as a function of  $\theta$  for different cumulatively applied corrections with the Voronoi (top panel), icosahedron (middle panel) and distance weighting scheme (bottom panel). The results are averaged over 100 different realisations of the mesh to suppress shot noise. Unit standard deviation regions around the ‘unweighted’ curve and the final result are shaded in light and dark grey, respectively. As a result of the anisotropy of the triangulation, more edges point towards the over-dense region (to the right) in the uncorrected case. Note that due to normalization, the results at the extreme ends are subject to noise.

lating these areas is computationally costly because in three dimensions the walls are generally irregular polygons with  $M$  vertices where  $M \geq 3$ .

### Icosahedron weighting

The second method is based on an ‘independent’ division of the unit-sphere into  $M$  (approximately) equal parts. We take the  $M$  vectors (originating in the nucleus under consideration) that point to these parts and assign weights  $w_i$  to the outgoing Delaunay edges as follows: we take a vector and calculate the  $N$  dot products with the Delaunay edges. The Delaunay edge that has the smallest dot product has a fraction  $1/M$  added to its weight. In the icosahedron scheme, the weight is thus proportional to the solid angle a Delaunay line occupies considering the angular vicinity of its neighbouring edges (in 2D the solid angle is consequently bound by the two bisectors shown as dotted lines in Fig. 3.3).



We chose the icosahedron as the basis of our weighting scheme. We use vectors to the middle of its 30 edges and its 12 vertices, yielding 42 reference vectors.

One significant advantage of the icosahedron method is that it is independent of the nature of the triangulation and is therefore very flexible. It becomes more accurate as the division of space is refined ( $M$  is increased) by taking other tessellations of the unit-sphere. The computational cost, however, scales linearly with  $M$ , forcing us to trade off accuracy for speed. Unfortunately this ‘independence’ also has a drawback: the reference vectors are oriented statically in space, introducing a systematic bias in those directions. To reduce this effect, one is forced to add some random noise to the procedure (for example, by applying random rotations of the whole icosahedron, another computationally costly operation).

After the correction with icosahedron weights (see middle panel of Fig. 3.5, dashed line), the anisotropy decreases below the 0.5% level. This immediately shows the strength of this method over the more noisy Voronoi scheme (compare also the unit standard deviation regions in dark grey).

### Distance weighting

Empirically we found that in 3D the square of the distance between nuclei is also a robust estimator of  $w_i$ . This quantity is readily calculated and provides by far the most rapid solution to the problem at hand. After the initial correction (see bottom panel of Fig. 3.5, dashed line), the anisotropy of edges diminishes to values lower than 0.5% even slightly better than in the icosahedron scheme. The deviations from the ideal isotropic case after the COG correction are of the same order as in the Voronoi and the icosahedron case.

We could not find a valid explanation for the success of this method. Intuitively, one would think that the opening angle,  $\Omega$ , of a Voronoi wall with respect to its nucleus would scale as  $r^{-2}$  rather than  $r^2$ . To find a mathematical reason for the proportionality between  $\Omega$  and  $r^2$ , one would have to delve more deeply into the field of computational geometry which is beyond the scope of this text.

### 3.3.5 Numerical examples: diffuse drift

We now proceed by describing the numerical experiment designed to show the effects of diffuse drift. The simulation domain of unity volume in three dimensions is filled with  $5 \times 10^5$  nuclei subject to a gradient in the point density of the form  $n(x) = x$ . We choose this linear form because it locally approximates every other type of gradient. The number of nuclei results in roughly 32 steps across the box along the gradient direction, and 38 in the directions perpendicular to the gradient. These numbers allow for most of the photons to travel through the box for more than 100 steps before being captured in the absorbing boundaries.

At the site closest to the centre of the domain, a number of photons are placed. Neither the outcome nor the speed of our simulation depends on this number as we use floating point numbers to represent photons.

Photons are transported over this mesh using diffusive transport without absorption. With every sweep, the photon cloud is expected to grow in size. In the case of a homogeneous point density, the photons will be distributed normally, as must be expected for pure diffusion. The

gradient in the point density will distort the form of the distribution function for two reasons. First, the mean free path on the mesh,  $\mathcal{D}$ , scales with the point density according to Eq. (3.1), allowing photons to diffuse faster into the underdense regions where the step sizes are larger. Second, the drift phenomenon described in Sect. 3.3.2 will counteract this physical diffusion and move the cloud into the overdense region.

To separate these two effects, we performed a one-dimensional Monte Carlo simulation of  $2 \times 10^6$  ‘random walkers’ that take steps with a size given by the recipe of Eq. (3.1). The experiment described above is emulated in one dimension but without the Delaunay mesh as an underlying structure. The random walkers are thus expected to experience the *physical* diffusion into the underdense region only. The unphysical drift is caused exclusively by the Delaunay mesh and will not be present in our results.

In Fig. 3.6, the intensity-weighted position of the photon cloud (along the direction of the gradient) versus the number of sweeps in the simulation is shown for SIMPLEX and the Monte Carlo experiments both with and without weights<sup>5</sup>. We can see that the behaviour conforms to our expectation. The weighting corrected SIMPLEX result coincides with the Monte Carlo as shown in the bottom panel.

Furthermore, we took our drift description Eq. (3.5) and applied it to the Monte Carlo experiment thus introducing a drift toward the overdense region similar to that experienced by a photon in SIMPLEX. The effect counteracts the physical diffusion into the underdense region resulting in a positive slope for the position of the photon cloud as a function of sweeps (see top panel of Fig.3.6). This procedure therefore provides a direct quantitative check of the correctness of Eq. (3.5).

The slopes of the Monte Carlo experiment agree to within a thousandth of a degree with the SIMPLEX results for both the uncorrected and weighted results. This implies that the weighting schemes discussed in Sect. 3.3.4 do not only correct the orientation of the edges in a statistical sense (as shown in Fig. 3.5) but also allow for correct transport over these edges. This may seem a trivial statement, but despite everything being fine in a global sense, local anomalies may prevail. The transport of photons, however, depends strongly on the local correctness of the weighting scheme and is thus a more stringent test. Furthermore, the recovery of the SIMPLEX results with the Monte Carlo experiments suggests that Eq. (3.5) describes the drift phenomenon accurately in a quantitative sense.

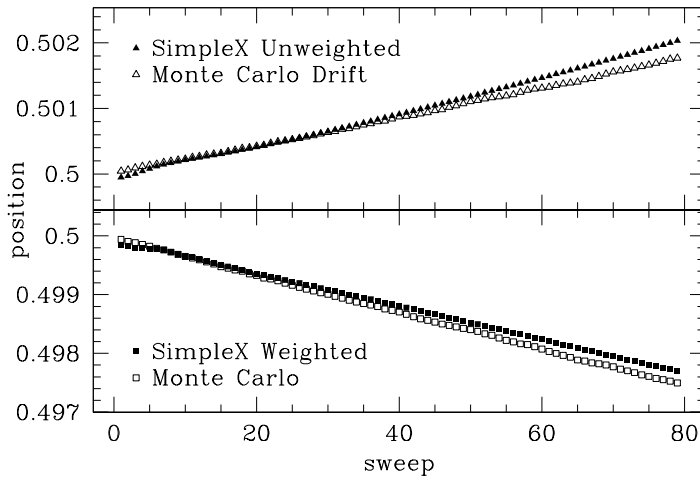
### 3.3.6 Effects on diffuse transport: clustering

As seen in Sect. 3.3.3 the expansion of a photon cloud is stalled by small (relative to the simulation domain) scale gradients in the mesh. To illustrate this effect performed the following experiment. The simulation domain is again given by the unit cube with a large number of photons at the centre. The  $5 \times 10^5$  nuclei are distributed according to the probability distribution

$$n(\mathbf{x}) = 0.2 + 0.8 \sin^2(\omega|\mathbf{x}|), \quad (3.17)$$

where  $\omega = 20$ . The density of nuclei thus inhibits concentric variations with a amplitude of 0.8 on a homogeneous background of 0.2.

<sup>5</sup>We used icosahedron weights in this case but this choice does not influence the results.

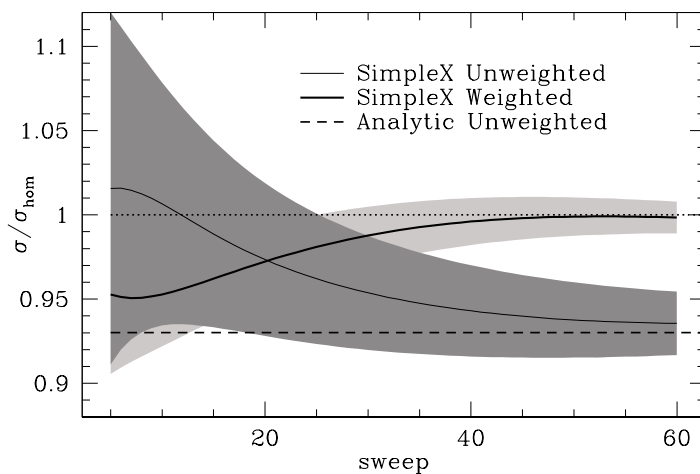


**Figure 3.6:** Intensity-weighted position of a photon cloud released in the centre of a simulation domain with a linear gradient in the number density of nuclei as a function of the number of sweeps. Results obtained with SIMPLEX are shown with filled symbols and results obtained by the Monte Carlo experiments are indicated with open symbols. We use triangles and squares in the corrected and uncorrected case, respectively.

We now prepare a similar simulation with a homogeneous distribution of nuclei where the box-size measured in units of  $\mathcal{D}$  is the same. A comparison of the spread of the photon cloud in the homogenous setup with that given by Eq. (3.17) provides a direct measure of the effect described by Eq. (3.14).

Application of a weighting scheme for diffusive transport described in Sect. 3.3.4 should remove the difference between the spread of the cloud in the two cases described above. As a check, we also compare the results to the analytical result given by Eq. (3.13) and the expected spread of the photon cloud in the homogenous case (which is simply  $\mathcal{D} \sqrt{N}/2$ ).

In Fig. 3.7, the spread of the photon cloud in terms of the normalised standard deviation in the intensity-weighted positions is shown for the inhomogeneous point density given by Eq. (3.17) (solid lines) and the analytical prediction (dashed line). The data are obtained from twenty runs with different realisations of the mesh.



**Figure 3.7:** Normalised standard deviation of the intensity-weighted positions of a photon cloud expanding in a distribution of nuclei with small-scale gradients with and without weights (thick and thin solid lines, respectively) and the analytical expectations according to Eq. (3.13) (dashed line). A dotted line at  $\sigma/\sigma_{\text{hom}} = 1$  is included to guide the eye. open arrows is longer than the length of the dotted line, which is corrected for by a global factor.

As expected, the photon cloud expands (almost 6%) more slowly when no weights are used in the case of small-scale gradients (thin solid line). Application of the weighting scheme results in a cloud that (after a slow start) conforms to the expected size (dotted line at  $\sigma/\sigma_{\text{hom}} = 1$ ). The analytical model predicts this behaviour to within 3% after 30 sweeps, and increasingly poorly for fewer sweeps (although the deviation remains below 10%). The larger deviations for fewer sweeps are caused by the cloud size (expressed in standard deviations) being small and the results are divided by this small number. We note that the unit standard deviation region for the weighted case is narrower in the weighted case (light gray) than the unweighted case (dark gray). This is expected because the weighting scheme corrects for local anisotropies and thus reduces the noise in the diffusive transport.

In the homogeneous case we have verified that applying a weighting scheme does not alter the expansion speed of a photon cloud significantly, as expected. We can thus conclude that, although clustering can impose a substantial effect on the expansion of diffuse radiation, our weighting scheme corrects for it appropriately.

### 3.4 EFFECTS ON BALLISTIC TRANSPORT

For ballistic transport, problems similar to those of diffusive transport arise but the picture is further complicated due to locally anisotropy of the transport. We first describe some properties of this kind of transport on a homogeneous mesh to appreciate deviations from the ‘normal’ case later on.

We identify two distinct phenomena that photons travelling ballistically may experience: deflection and decollimation. Deflection is here understood as ‘loss of direction’, where the direction is given by the vector sum of the three most forward pointing directions. Decollimation is defined to be the effect of the increase in the opening angle of a beam of photons as they are transported ballistically. The first phenomenon depends on the evolution of the vector sum of the three most forward directions, whereas the second phenomenon is related to the angular separation between these directions individually. In the case of a homogenous distribution of nuclei, the net effect of deflection will vanish because there is no preferential direction in the mesh.

#### 3.4.1 Ballistic decollimation

Ballistic decollimation has already been described in Sect. 2.1.4 as the motivation for direction conserving transport. The discussion here merely gives additional information; primarily on the distribution of most forward directions. As previously stated in Sect. 2.1.4, the angular resolution, and therefore the minimal opening angle, is a property of the Voronoi cell and is thus fixed for the chosen triangulation. To exemplify this, we consider a typical Voronoi nucleus in 3D, connected to  $\Lambda$  neighbouring nuclei. The expectation value for the solid angle,  $\Omega$ , subtended by each edge is thus

$$\Omega = \frac{4\pi}{\Lambda} \approx \frac{\pi}{4}. \quad (3.18)$$

In Fig. 3.8, the distribution of angles is given for a mesh of  $10^5$  homogeneously placed nuclei in three dimensions. For the vector sum, the average departure from the incoming direction is

about  $15^\circ$  (solid line). This implies that a photon loses all knowledge of its original direction in typically  $(90/15)^2 = 36$  steps.


For the distribution of the three separate most forward edges, we find a standard deviation of about  $39^\circ$ , which means that after  $(90/39)^2 \approx 5$  steps the photon has lost all memory of its original direction. We note that the width of the vector sum of the weighted edges is smaller than that of the most forward directed edge (indicated by ‘1’ in the figure) alone.

In general, the gradual loss of direction is not a major concern since most sources emit isotropically anyway, but the effect becomes important when many edges are traversed, as it causes the photons to behave diffusively. Simulations in which the mean free path for scattering or absorption is much smaller than 5 edges are fine in this respect, as the ballistic photons never enter this random walk regime (see also the discussion in Sect. 3.5).


### 3.4.2 Ballistic deflection

We are now ready to quantify the effect of the anisotropy of the triangulation on the ballistic transport of photons, which travel along the three edges closest (in angular sense) to the incoming direction. For the sake of simplicity, we estimate the deflection for radiation travelling along the most forward pointing edge and discuss the applicability to three edges afterward.

We consider photons streaming perpendicular to the gradient direction (see Fig. 3.9 for the geometry of this situation). The standard deviation in the deflection of the outgoing edge with respect to the incoming direction is typically  $15^\circ$  (see Fig. 3.8) for a homogeneous mesh, but in general depends on the local value for the gradient. Referring to Eq. (3.18), the expected opening angle depends on the effective number of outgoing edges in that direction. In other words, as the anisotropy increases, the number of edges pointing toward the overdense region and the angular resolution increases accordingly. This motivates one to define the direction-



**Figure 3.8:** Normalised distribution of angles between the incoming direction and the vector sum of the (three) most forward pointing edges (solid line) and the separate most forward edges (dotted, dashed and dot-dashed lines) for a homogeneous distribution of nuclei. The distribution of the most forward edges added is shown as the long dot-dashed line with label ‘Added’ and is related to the decollimation effect. The standard deviation ( $\text{FWHM}/\sqrt{\ln 256}$ ) is about  $39^\circ$  for the added edges and  $15^\circ$  for the vector sum of the edges.



**Figure 3.9:** Geometry of radiation travelling perpendicular to the gradient direction. The deflection angle toward the underdense region,  $\theta_u$ , scales with the radius of the area  $S_u$  and similarly for the deflection angle towards the overdense region,  $\theta_d$ .

dependent number of outgoing edges,  $\Lambda_{\text{eff}}$ , to be

$$\Lambda_{\text{eff}}(\phi) = \Lambda[1 + E(x) \cos \phi], \quad (3.19)$$

where  $\phi$  is the angle with the gradient direction. If we interpret the solid angle as a projected circular area on the unit sphere, we can estimate the maximal deflection angle,  $\theta_d$ , for ballistic transport as

$$\theta_d = \arcsin \sqrt{\frac{4}{\Lambda_{\text{eff}}(\phi)}}. \quad (3.20)$$

For a homogeneous distribution of nuclei, this angle equals approximately  $30^\circ$  with a typical value of  $15^\circ$  being expected from the analysis shown in Fig. 3.8. If we include this empirical factor of one half in Eq. (3.20) and approximate the arcsine by its argument (correct to within 1% for angles smaller than  $\pi/12$ ), we obtain

$$\theta_d = (\Lambda[1 + E(x) \cos \phi])^{-1/2}, \quad (3.21)$$

where we have used Eq. (3.19) to substitute for  $\Lambda_{\text{eff}}$ . This result can be approximated to first order by

$$\theta_d \simeq \Lambda^{-1/2} \left[ 1 - \frac{E(x) \cos \phi}{2} \right]. \quad (3.22)$$

To obtain the deflection per step over the mesh, we average  $\theta_d$  over azimuthal angle while projecting along the direction of the gradient

$$\theta_{\text{eff}} = \frac{1}{2\pi} \int_0^{2\pi} \theta_d \cos \phi d\phi \quad (3.23)$$

$$= -\frac{E(x)}{4\sqrt{\Lambda}}. \quad (3.24)$$

The sign of  $\theta_{\text{eff}}$  is negative, which means that the resulting deflection can be found in the region of lower density. This result may at first surprise the reader. One could naively predict the photons to deflect into the direction of higher density (similar to the diffuse drift), but the situation is exactly the reverse for ballistic deflection.

We note that this effective deflection angle acts as an upper limit for the deflection encountered in a simulation for the following reasons. The effect is maximal for radiation travelling perpendicular to the direction of the gradient, and this is the situation we have used as a starting point for the above derivation of  $\theta_{\text{eff}}$ . Secondly, the fact that  $\theta_d < \theta_u$  increases the likelihood of selecting edges in the overdense region, effectively diminishing the deflection, an effect that we have neglected in the derivation above. To include this effect, one would have to know the distribution of outgoing edges as a function of angle with the gradient, and assign a probability to the selection of an edge accordingly. As we see in Sect. 3.4.6, the omission of this effect does not seem to be of much importance to our predictions.

Sending radiation along three edges rather than one decreases the effect of deflection. This is expected because the deflection of one edge is larger than that of the vector sum of three outgoing edges (as we saw in the case of a homogeneous mesh).

The growth of the deflection when traversing the mesh can be found by dividing Eq. (3.24) by the typical step-size,  $\mathcal{D}$

$$\frac{\Delta\theta_{\text{eff}}}{\Delta l} = \frac{1}{16\sqrt{\Lambda}} \frac{\nabla n(x)}{n(x)}. \quad (3.25)$$

As for the drift length of Eq. (3.12), we can define a deflection length,  $L_{\text{def}}$ , at which the cumulative deflection is equal to say,  $\pi/4$ ,

$$L_{\text{def}} = \frac{4\pi\sqrt{\Lambda}n(x)}{\nabla n(x)} \approx 50 \frac{n(x)}{\nabla n(x)}. \quad (3.26)$$

Comparing with Eq. (3.12), we see that both the diffuse drift and the ballistic deflection place approximately the same restrictions on the box size, the only difference being a factor of 50 instead of 8. The interpretation, again, is simple. If  $L_{\text{def}}(\mathbf{x})$  is strictly larger than say five times the box size, the deflection will be less than one-fifth of  $\pi/4$  when travelling across the box in a ‘straight’ line. We recall that this effect is added to the expected decollimation from Sect. 3.4.1.

### 3.4.3 Correcting ballistic transport: weighting schemes

Assigning weights to the outgoing edges as described above removes the drift and clustering problems for diffusive transport described in Sect. 3.2.

The effects of anisotropy in the ballistic case are intrinsically more challenging to correct because it is not a priori clear how the anisotropy of all outgoing edges of a nucleus should be used to choose three weights for the outgoing edges. As described in Sects. 3.4.1 and 3.4.2, in the ballistic case we must distinguish between decollimation and deflection, the first of which dominates the overall ‘loss of direction’. To diminish the decollimation of a beam, we can assign weights to the most forward pointing edges. If we make these weights somehow proportional to the inner product of the edge and the incoming direction, the most forward pointing edges transport most of the photons.

There are many possible ways to assign weights to the  $D$  most forward directed edges, each of which optimise different aspects of the transportation process. If the exact direction is important, for instance, the weights should be chosen in such a way that the vector sum of the resulting edges points straight ahead. On the other hand, minimising the decollimation will generally yield a different set of weights, namely those that maximise the length of the vector sum.

From this myriad of possibilities, we choose a very simple and computationally cheap approach. To edge  $j$  we assign the weight  $w_j$  given by

$$w_j = \frac{\cos(\theta_j)/\sin(\theta_j)}{\sum_j \cos(\theta_j)/\sin(\theta_j)}, \quad (3.27)$$

where  $\theta_j$  is the angle between the incoming direction and edge  $j$ . The division by the sine of the angle places extra emphasis on the edges with the smallest inner product.

Applying this weighting scheme to the point distribution used to produce Fig. 3.8 yields the distribution of edges shown in Fig. 3.10. As specified by Eq. (3.27), the ballistic weights have the desired effect of decreasing the width of the distributions of the edges used in ballistic transport.

In Fig. 3.11, evolution of the step-size of ballistic transport in the weighted case is shown with crosses. The decollimation angle per step is substantially smaller than in the case without weighting ( $26^\circ$  compared to  $39^\circ$ ). This decrease in decollimation angle consequently relaxes the requirement that the number of ballistic steps must remain below 5. Including ballistic weights typically allows for up to 12 ballistic steps before the original direction is lost. Finally, there are cases where the anisotropy of edges emanating from a cell cannot be solved by any weighting scheme. If there is no edge in a given direction, the radiation cannot go there, regardless of the weight values. One should begin with a reasonably isotropic triangulation to apply a weighting scheme in a useful way. In the case of Delaunay triangulations in three-space, this is almost always the case.

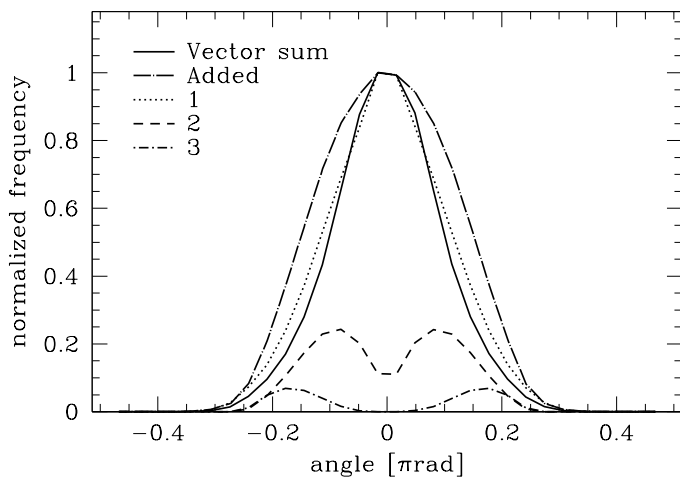
### 3.4.4 DCT throughout the mesh

A conceptually different solution to the problems described in Sect. 3.2 is to apply the direction-conserving transport from Sect. 2.1.4 throughout the simulation domain. As the angular direction of the radiation is effectively decoupled from the mesh in this approach, all drift, clustering, decollimation, and deflection are resolved at once. The cone containing the nuclei that receive radiation will, however, still be dependent on the mesh itself as already pointed out in Sect. 2.1.4.

### 3.4.5 Numerical examples: ballistic decollimation

To observe the loss of direction of photons transported over many edges with ballistic transport, we construct a triangulation of a random point distribution containing  $10^5$  homogeneously placed nuclei in a cube of unit dimension. In the centre, we define a small spherical volume of radius 0.05 and place a number of photons in each nucleus contained in this volume. The





**Figure 3.10:** Normalised distribution of angles between the incoming direction and the vector sum of the  $D$  most forward pointing edges (solid line), and the separate most forward edges (dotted, dashed and dot-dashed lines) for a homogeneous distribution of nuclei with ballistic weights. The distribution of the sum of the most forward edges is shown as the long dot-dashed line and labeled ‘Added’. The standard deviation ( $\text{FWHM}/\sqrt{\ln 256}$ ) is about  $26^\circ$  for the added edges, and  $16^\circ$  for the vector sum of the edges. Note that the width of the distribution of the single most forward pointing edge poses a definite lower limit to the width of the ‘Added’ lines.

photons are assigned a direction by sending them along the edge whose direction is maximally parallel to one of the coordinate axes (we select the positive x-axis). When applying our ballistic transport method, the photons begin to move in the positive x-direction with a step-size equal (on average) to  $\mathcal{D} \cos \theta_D$ , where  $\theta_D$  is the mean decollimation angle. Every step means a multiplication by a factor  $\cos \theta_D$ , until after infinitely many steps the effective step size in the x-direction is zero. In Fig. 3.11, the step-size as a function of sweeps is shown together with a fitting function of the form


$$f(N_s) = A \cos(\theta_D)^{N_s}, \quad (3.28)$$

where  $N_s$  is the number of sweeps. We find that without ballistic weights (filled squares; see Sect. 3.4.3 for the description of ballistic weights)  $\theta_D = 0.67 \pm 1\%$  ( $39^\circ \pm 1\%$ ), in accordance with the value obtained from the mesh statistics presented in Fig. 3.8.


### Sharp gradients

An extreme case of decollimation is encountered if strong gradients in the point density occur. In this case, a large cell may find itself surrounded by much smaller cells (see Fig. 3.12 for an example in 2D). Because of the sharp contrast in size, two out of the three most straightforward edges of the smaller cells in the direction of the large cell are actually directed *around* the large cell. If no ballistic weights are used, this scenario results in too few photons flowing into the large cell.

It is precisely this effect that caused the anomalous result of the first incarnation of the SIMPLEX algorithm in test 4 of the first cosmological radiative transfer comparison project (Iliev



**Figure 3.11:** A least squares fit of the projected step-size in the  $x$ -direction as a function of sweeps for the unweighted ballistic transport (solid line) and with ballistic weights as described in Sect. 3.4.3 (dashed line). A region indicating unit standard deviation is shaded in grey around the markers. Error bars correspond to  $\sigma/\sqrt{N-1}$ , where  $\sigma$  is the standard deviation and  $N$  the number of experiments executed. For the uncorrected case, the fitted value for the deflection angle  $\theta_D$  in Eq. (3.28) is  $39^\circ \pm 1\%$ . For the weighted case, we find  $\theta_D = 26^\circ \pm 3\%$ .

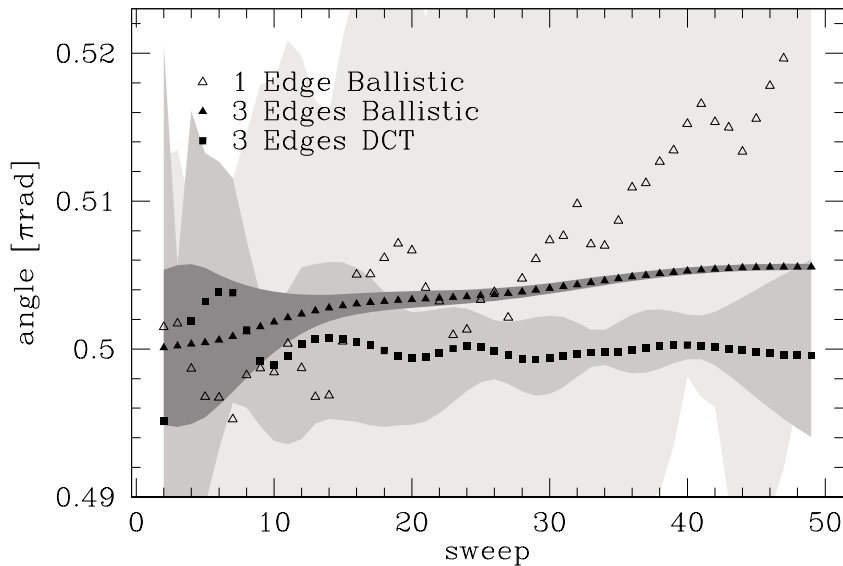


**Figure 3.12:** Voronoi-Delaunay mesh around a sharp gradient. The solid grey arrow indicates the incoming direction into the large cell. The three most straightforward directions with respect to the incoming direction are shown as dashed grey arrows. Note that we have used three most straightforward directions for this illustrative example whereas we normally would use only two in 2D.

et al. 2006a). Observe for instance the chunky artifact in the upper left corner of Fig. 33 in that work. Here the voids are under-resolved in the SIMPLEX mesh resulting in cells that are orders of magnitude larger in volume than the surrounding ones. We will return to this issue in Sect. 4.3.2.

### 3.4.6 Numerical examples: ballistic deflection

To study whether anisotropy has any effect on the direction of a beam of photons, we considered a triangulation of a random point distribution containing  $5 \times 10^5$  nuclei placed in a cube of unit dimensions with a linear gradient along the  $x$ -direction. In the centre, we included a spherical volume of radius 0.05 and place a number of photons in each nucleus contained in this volume. The photons were assigned a direction by sending them along the edge whose direction is maximally perpendicular to the  $x$ -axis. The simulation was executed with ballistic transport using both one and three outgoing edges and in addition with the direction-conserving implementation of SIMPLEX (as described in Sect. 2.1.4). Every run was repeated 10 times with different instances of the mesh to suppress shot noise and obtain error estimates.



**Figure 3.13:** Mean angular directionality of a photon cloud as a function of the number of sweeps. In this simulation, a linear gradient in the density of nuclei is present and the angle is measured relative to the direction of the gradient, so smaller values of the angle point toward the denser region. As expected, the direction-conserving implementation of SIMPLEX (as described in Sect. 2.1.4, indicated by filled squares) has no long-term loss of directionality, although it shows some oscillatory behaviour that dampens with time. Ballistic transport with one edge (open triangles) shows a deflection towards larger angles for number of sweeps of order  $5 \times 10^{-4} \pi$  rad/sweep in accordance with predictions. Using three outgoing edges (filled triangles) diminishes the effective deflection by a factor of five. Unit standard deviation regions are shaded in grey for the three simulations.

When applying our ballistic transport method, the photons undergo small angular deflections into the underdense region as predicted in Sect. 3.4.2. Using Eq. (3.24), the expected effective deflection for this setup is  $\theta_{\text{eff}} = 5 \times 10^{-4} \pi$  rad/sweep, which yields a cumulative result of approximately  $0.025 \pi$  rad after 50 sweeps. This estimate is consistent with the result of the simulation with one forward edge denoted by the open triangles in Fig. 3.13. As expected, sending radiation along three edges decreases the deflection and fluctuation of the angular di-

rection in general because radiation is distributed over many ( $3^{N_s}$ ) edges, which quickly reduces shot noise with time. When using the direction-conserving implementation of SIMPLE<sub>X</sub> (filled squares), the radiation does not suffer from ballistic deflection.

## 3.5 DISCUSSION

After identifying and correcting the various effects caused by anisotropy in the Delaunay mesh, we discuss their implications in a broader context.

### 3.5.1 Prevention versus correction

It is impossible in general to define the perfect mesh. As already pointed out, a large dynamic range inevitably leads to gradients that, in turn, produce the unphysical effects described in this text. Although we have shown that these effects can be corrected by weighting schemes, these corrections are not fail-safe if the gradients are too strong. In particular, when the density of nuclei changes appreciatively on length scales smaller than the Delaunay edge length  $\mathcal{D}$ , the lack of edges towards the underdense region may lead to extreme decollimation as there are simply hardly any edges that point that way. In this case, a weighting scheme is not feasible as there are no edges to assign weights to in the underdense direction. We therefore need to begin with a fairly well behaved mesh in the first place.

One could argue that if the mesh is constructed such that the unphysical effects are below a predefined tolerance, the deployment of weighting schemes can be circumvented. On the other hand, this places rather tight constraints on the mesh and it may be preferable to retain a greater dynamic range (and thus structure) in the mesh at the expense of the need for weighting.

An intermediate solution will often be the most viable option, where we retain the dynamic range of the data as much as possible while weighting schemes and DCT are employed at *certain locations* in the simulation domain only.

### 3.5.2 Relative importance of systematic effects

Although we have analysed the various systematic effects in isolation, in real-life simulations they will appear simultaneously. It is therefore relevant to know which effect will dominate the others in what situation. Both diffuse drift and ballistic deflection allow for a description where we can derive a typical length at which they start dominating the transport. The appropriate expressions are given by Eq. (3.12) and Eq. (3.26) where the former is a factor  $50/8 \approx 6$  more severe. So, diffuse drift is a bigger problem than ballistic deflection.

To put this statement into proper perspective, however, we must look at several aspects of our real-life simulation. First, we must ask ourselves if our computational mesh is comparable to the idealized linear gradient for which the above equations have been derived. In most cases, we will have gradients that are local (spanning maybe less than 10% of the box size) and not global. Therefore, the effects of drift and deflection may be only visible locally. Secondly, we must take into account the *amount* of radiation that we are sending with either diffuse or ballistic transport. In most cases, diffuse transport will only be used to transport diffuse recombination

radiation from ionized gas. The majority of photons will travel ballistically. If diffuse drift is a big problem compared to ballistic deflection, it will be a big problem for a small number of photons. Moreover, recombination radiation will often be absorbed nearby its position of emission and will thus not have time to experience much diffuse drift. We therefore conclude that diffuse drift is not likely to affect the outcome of realistic simulations unless they include many diffuse photons that can travel many edge lengths.

What about diffuse clustering? From Fig. 3.7, we can directly see that deviations in the spread of photons is maximally of the order of 5% for our test setup (see Sect. 3.3.3). We have tweaked this test to maximise the effect of diffuse clustering and are therefore confident that 5% is a robust upper limit for most applications. Moreover, diffuse clustering can be easily aided with a simple weighting scheme.

The systematic effect that we need to worry about in almost all practical applications is ballistic decollimation. The loss of direction of the ballistic photons results in softening of shadows and over-prediction of ionisation rates inside ionised regions. Fortunately, ballistic decollimation can be aided in several ways. First we must try to use this transport mode only in regions of high optical depth. This requirement can be met in two ways, either by reducing resolution (so cells become larger and, therefore, optically thick) or by using ballistic weights or DCT to conserve direction if the cells become optically thin. We have implemented this last idea as a switch between ballistic- and direction conserving transport in the `SIMPLEX` code. The switch is based on the simple requirement that DCT replaces ballistic transport as soon as the optical depth of a cell becomes smaller than unity.

## 3.6 SUMMARY

We now summarize the main insights and results from the above analysis.

- The mathematically transparent nature of the `SIMPLEX` algorithm allows a rigorous and general assessment of its systematic effects to be performed. Although only some of the described effects (diffuse drift and ballistic decollimation) need to be considered in typical applications (see [Paardekooper et al. 2010](#)), we have investigated all four in detail for completeness.
- The use of a random Delaunay triangulation in the radiative transfer method `SIMPLEX` introduces global errors when the point-density is inhomogeneous. We have identified and quantified four distinct effects: ‘drift’ and ‘clustering’ of photons in diffusive transport, and ‘decollimation’ and ‘deflection’ in ballistic transport.
- We have shown how diffuse drift and clustering can be adequately corrected for by adopting a weighting scheme. Three weighting schemes for diffusive transport have been discussed and compared.
- We have shown that decollimation become troublesome only when the number of traversed edges becomes larger than roughly 5 steps. This implies that one either preclude this regime or, when this is undesirable or impossible, correct for the unphysical behaviour.
- The use of direction-conserving-transport (introduced in the previous chapter) can prevent loss of angular direction in optically thin regions (i.e., ballistic decollimation). DCT

also provides a means of correcting for deflection, drift, and clustering. The computational cost and need for additional randomisation of the solid angle directions make the employment of DCT computationally more expensive than weighting schemes. We therefore adopt it only in situations where it is necessary.

- A weighting method for ballistic transport has been shown to decrease the effect of declination significantly by pushing the limit for correct transport from 5 to 12 steps. This is a computationally efficient alternative for direction-conserving-transport.

## ACKNOWLEDGMENTS

Part of this work was supported by the German *Deutsche Forschungsgemeinschaft*, DFG project number Ts 17/2–1. CJHK would like to extend his gratitude to Rien van de Weijgaert, Jakob van Bethlehem, Sven de Man en Marcel Kruip for proofreading the manuscript and valuable discussion. We also thank the anonymous referee for useful comments on the manuscript.



## CHAPTER 4

---

# The SimpleX mesh

C. J. H. Kruij, J.-P. Paardekooper & V. Icke

Part of this chapter consists of work that has been published in  
Astronomy & Astrophysics 515, A78, 2010 and  
Astronomy & Astrophysics 515, A79, 2010

**W**E present the construction procedure for the SIMPLEX radiative transfer mesh. The parameters of this function can be constrained by the diagnostics developed in Chapter 3. We present a total of six case studies that exemplify several issues related to the construction of a mesh suitable for radiative transfer with the SIMPLEX algorithm. The studies focus on the accuracy of density representations and the prevention of harmful effects due to geometrical properties of the mesh.



## 4.1 INTRODUCTION

The Voronoi-Delaunay mesh lies at the heart of the SimpleX method. Speed, accuracy and stability of the method all depend vitally on its properties. The local resolution, for example, depends on the *local* number density of mesh-points. But every mesh-point can be placed only once. Hence, high resolution in one place immediately implies lower resolution in another. The speed of the method depends heavily on the ionization and recombination time scales in the computational cells. These time scales in turn depend on the local radiation field and the opacity of the cells. Stability issues are rarely encountered in SimpleX but they may occur in the mesh construction phase where computational domains are stitched together allowing for distributed-memory parallel computations. Here it is vital to ensure that the boundaries of the various sub-domains are convex. If this is not the case, the topology of the triangulation of a region will differ between adjacent boxes prohibiting their connection. The central question is thus how to place the available points in the computational domain without sacrificing either speed, accuracy or jeopardizing the stability of the computation. In this chapter we will discuss various methods and strategies designed to obtain the point distribution that satisfies the aforementioned criteria.

To represent the given medium (generally a density-field) by a discrete point-distribution, we introduce in Sect. 4.2 a sampling function with adjustable parameters. We will describe its use and give some examples of the influence of its parameters and how they can be constrained meaningfully. In almost all cases of interest, the computational mesh represents a physical medium. This physical medium can be obtained from analytical models, mesh- or particle-based hydrodynamics. We should therefore take good care that the relevant properties of the medium are captured by our mesh. To this end we employ several diagnostic methods that allow us to compare the original medium with our mesh.

When we talk about the SIMPLEX mesh, we mean both the *geometrical* properties of the mesh (encoded in the Delaunay triangulation) and the *density field* that is to be stored on the nodes or vertices of that triangulation. In order to construct a computational mesh that is both suitable for accurate RT and captures the density field of the original data accurately we need to keep these two properties in mind. We will give working examples of how this can be achieved in Sect. 4.3 and Sect. 4.4.

We introduce our main sampling procedure in Sect. 4.2 and give some examples of its use. In Sect. 4.3 and Sect. 4.4 we present a total of six case studies that focus on issues such as sample power (Sect. 4.3.1 and Sect. 4.3.2) and mass inclusion methods (Sect. 4.3.3 and Sect. 4.4.2). A study of the effects of resolution and adaptivity of the mesh in the context of density representation (Sect. 4.3.4) and of geometry (Sect. 4.4.1) is also included.

## 4.2 SAMPLING FUNCTION

As already mentioned in Chapter 2, the connection between a density field given by  $n(\mathbf{x})$  and a point density field  $n_p(\mathbf{x})$  is determined by the currently un-specified function  $f(n(\mathbf{x}))$  used in Eq. (2.5). The accuracy and efficiency of the SIMPLEX algorithm depend vitally on the choice of this function. As we have seen in Chapter 3, strong gradients in the point-density  $n_p$  give rise to un-desirable systematic effects in the solution of the radiative transfer problem (e.g. ballistic

decollimation and ballistic deflection). On the other hand, gradients in the point density are a direct consequence of the adaptive nature of the SIMPLEX mesh and are therefore a desirable feature. We thus face a trade-off between maximizing dynamic range and minimizing the strong gradients this inevitably introduces. In other words, we must investigate the effects that different sampling functions have on the triangulation and its transport properties.

#### 4.2.1 Mean free path sampling

As mentioned in Sect. 3.3.2, the average Delaunay edge length can be associated with the mean free path of the photons in a very natural way. This is achieved when  $n(\mathbf{x})$  scales with the opacity (or number density) to the  $D$ -th power (Ritzerveld & Icke 2006, Eq. (14), cf. Eq. (2.5))

$$n_p(\mathbf{x}) = \Phi * n^D(\mathbf{x}), \quad (4.1)$$

where  $\Phi$  denotes a homogeneous Poisson process,  $D$  is the dimension of the propagation space, and  $*$  denotes the convolution. This particular choice for the sampling function has one especially elegant property. Because the mean free path for photons is now a global constant of the mesh, every cell thus has the same optical depth. This greatly simplifies the calculation of the absorption fraction during radiative transfer as it is always the same factor  $1/e$ .

#### 4.2.2 Hybrid sampling

One major drawback of using Eq. (4.1) is that the resulting point process, given a substantial density contrast, places many nuclei in dense regions and very few in underdense regions. For many applications, simulating the true mean free path may thus require a number of nuclei that is prohibitively high (see also the middle panel of Fig. 4.2). We therefore adopt a more flexible sampling function, based on the harmonic mean of two sampling functions, that behaves differently for the extremal densities of the mesh. According to this function, the number density of nuclei that generate the Voronoi mesh is given by

$$n_p(\mathbf{x}) = 2\Phi * \left(y^{-\alpha} + y^{-D}\right)^{-1}, \quad (4.2)$$

where  $y \equiv n/n_0$  and  $n_0$  is a reference density that marks the transition between the two regimes of sampling power. In our applications,  $\alpha$  will always be smaller than  $D$  (and even smaller than unity) to prohibit the over-emphasis on high density regions. Using the sampling function of Eq. 4.2, the overdense regions of the medium are sampled with the power  $\alpha$ , whereas the underdense regions are sampled with power  $D$ . The parameter  $n_0$  determines where the transition between over- and underdense is located.

#### 4.2.3 Constraints on $\alpha$ and $n_0$

According to Eq. (3.12) and (3.26), the quantity that links the properties of the mesh (number density of nuclei and gradients therein) to the systematic effects described in the previous sections is  $Q_n \equiv n_p/|\nabla n_p|$  (note that we have dropped the subscript  $p$  in  $Q_n$  for simplicity although

it describes the *point* number density). We are now in the position to define, in a way analogous to  $Q_n$ , the quantity

$$Q_y \equiv y/|\nabla y|, \quad (4.3)$$

which can be measured over the physical density field. As the sampling function of Eq. (4.2) maps the physical field to the number density of nuclei, the measured value of  $Q_y$  and the upper limit of  $Q_n$  (posed by the maximally acceptable value for either  $L_{\text{drift}}$  or  $L_{\text{def}}$ ) constrain the sampling parameters  $\alpha$  and  $n_0$ .

From Eq. (4.2), we can derive an expression that links the measured value of  $Q_y$  to  $Q_n$  and thus to the sample parameters, given by

$$Q_n = Q_y \frac{y^{-\alpha} + y^{-D}}{y(\alpha y^{-\alpha-1} + D y^{-D-1})}. \quad (4.4)$$

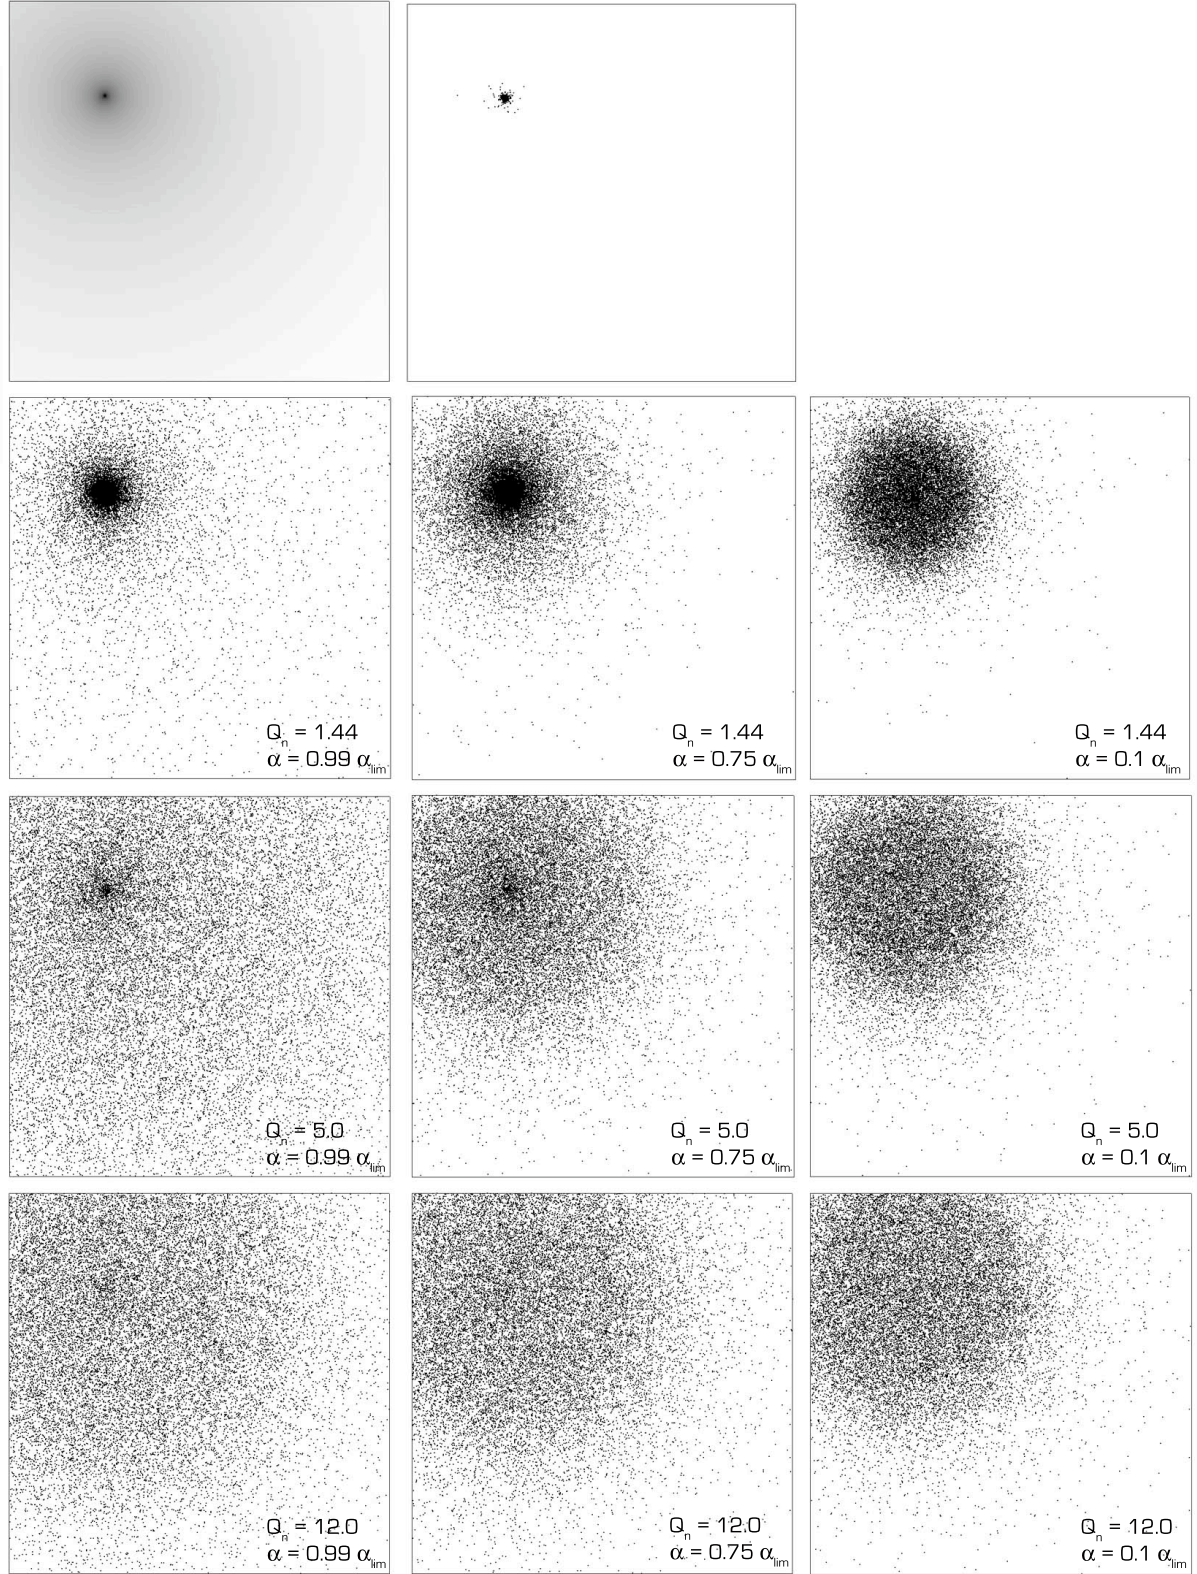
Solutions for  $Q_n = Q_y$  exist only for certain pairs of  $\alpha$  and  $n_0$  as can be seen from the limit of Eq. (4.4),

$$Q_n^{\text{lim}} = \lim_{y \rightarrow \infty} Q_n = \frac{Q_y}{\alpha}. \quad (4.5)$$

If we were to choose  $\alpha$  at its maximal value of  $\alpha_{\text{lim}} \equiv Q_y/Q_n^{\text{lim}}$ , we would therefore effectively enforce that  $n_0 \rightarrow 0$ , resulting in the limit where all density regimes are sampled with the same power  $\alpha$ . Choosing  $\alpha$  to be smaller but close to its extremal value maximises the dynamic range in the high density regions and allows  $n_0$  to attain values where a substantial volume of the density field (namely the voids) is sampled with the advantageous power  $D$  (keeping the optical depth per cell constant in those regions). In Fig. 4.1 we show the effect of different values for  $Q_n$  and  $\alpha$  (as a fraction of  $\alpha_{\text{lim}}$  on the sampling of a  $r^{-2}$  density distribution. The value of  $n_0$  is fixed by the choice of these two parameters. For comparison, we also show the result of the ‘mean free path’ sampling function of Eq. (4.1). Clearly, this sampling strategy has a tendency to under-resolve low density regions.

The columns in Fig. 4.1 show the influence of increasing  $Q_n$  (from top to bottom). Higher values of  $Q_n$  result in a more extended distribution of points. Looking at the rows, we increase  $\alpha$  from left to right at fixed  $Q_n$ . A lower value of  $\alpha$  implies a higher value for  $n_0$  resulting in a less pronounced density peak. This shift of emphasis has two reasons. First, lower values of  $\alpha$  place relatively fewer points in the high density peak. Second, lower values of  $\alpha$  imply higher values of  $n_0$  which means that a larger range of densities is sampled with the  $D$ -th power resulting in less emphasis on the lowest values in that range. In other words, the strong density dependence of Eq. (4.1) results in emphasis of the density values close to  $n_0$ .

In summary, the hybrid sampling function makes it possible to control the slope of gradients in the point density. Higher values of  $Q_n$  result in shallower gradients and thus higher effective resolution in low-density regions. Given the fact that the total number of points in a simulation is generally fixed by memory or speed requirements, this increased resolution at lower densities goes at the expense of the resolution in high-density regions. One should therefore choose the lowest  $Q_n$  value for which numerical errors due to undersampling in low density regions are within a predefined tolerance set by the requirements of the simulation (e.g., Eq. (3.12) and (3.26)). Furthermore, for fixed  $Q_n$ , the value of  $\alpha$  regulates at which  $n_0$  the transition between the two sampling powers occurs. This implies that it is crucial to choose an  $\alpha$  value that ensures that no strong density gradients exists at  $n < n_0$ .



**Figure 4.1:** Examples of application of our hybrid sampling function (Eq. (4.2)) for different values of  $\alpha$  and  $Q_n$ . *First row:* Density field with a single density peak and a  $r^{-2}$  profile with  $Q_y = 1.44$  and the sampling according to the  $d$ -th power of the density (mean-free-path sampling). In this example,  $d = 2$ . *Second row:* sampling with  $Q_n = 1.44$  and  $\alpha_{lim} = 1.0$ . *Third row:* sampling with  $Q_n = 5.0$  and  $\alpha_{lim} = 0.29$ . *Fourth row:* sampling with  $Q_n = 12.0$  and  $\alpha_{lim} = 0.12$ .

#### 4.2.4 Cosmological density field

We now apply this idea to a more realistic example and see what this all means in practice. A typical application wherein the strength of the `SIMPLEX` algorithm can be fully deployed is the epoch of reionization (see Chapter 7 as well as [Paardekooper et al. \(2010\)](#) for an account of relevant test cases). Both the wide range of densities (four to five orders of magnitude) and the large number of sources (hundreds to thousands in realistic cosmological volumes) can be treated naturally by *SimpleX* due to the adaptive nature of the Delaunay mesh and because the computational effort does not increase with increasing number of sources.

We consider a (0.5/h Mpc comoving) cosmological volume (as used in test 4 of [Iliev et al. \(2006a\)](#)) and construct a representative point set using constraints from analysis of the original data. In this specific case, the cosmological density field was presented as a regular grid of  $128^3$  equal size cells. Given the regular grid, the quantity  $Q_y$  can be determined for every pair of adjacent cells. The mean value of  $Q_y$  obtained is 1.6<sup>1</sup>.

This value implies that if we would use a linear correspondence between the mass density of the original data and the number density of the nuclei of the Delaunay mesh, the equality length (see Eq. (3.12)) would become  $8 \times 1.6 \simeq 13$  and the deflection length of Eq. (3.26) becomes  $50 \times 1.6 \simeq 81$ . This means that radiation can travel 13 box-lengths before diffuse drift starts to dominate the transport of radiation, and 81 box-lengths before ballistic deflection reaches a cumulative magnitude of  $\pi/4$ .


This linear sampling could be performed using  $\alpha = 1$  and  $\rho_0 \rightarrow 0$  in Eq. (4.2). Linear sampling is hardly ever desirable in a cosmological setting, however, as it tends to place the bulk of the nuclei in high density regions resulting in oversampling (many nuclei per cell of the original grid) of filaments and clumps and under-sampling of voids.

We proceed by finding optimal values of our sample parameters, given the measured value of  $Q_y$  and a chosen value for the maximum allowed equality length. An upper limit to the equality length, by means of Eq. (3.12), immediately implies a maximal  $Q_n$ .

For example, if we set  $L_{\text{drift}} = 40$  (and consequently  $L_{\text{def}} = 250$ ), we see that we need to choose  $\alpha$  and  $n_0$  such that  $Q_n = 5$ .

Solving Eq. (4.5) for  $Q_y = 1.6$ , we find  $\alpha = 0.32$  to be the maximal value. Taking the somewhat smaller value of 0.3, the corresponding value for  $n_0 = 3.7 \times 10^{-5} \text{ cm}^{-3}$  which is significantly higher than the lowest density in the cosmological field, which is about  $2.4 \times 10^{-5}$ . With these parameters, the point density in the lowest density regions is equivalent to a resolution of approximately  $77^3$  for  $128^3$  grid points. In the right panel of Fig. 4.2, a cut at  $z = 0.5$  through the sampling defined by these parameters is shown. The middle panel shows the point distribution one obtains when sampling with  $\alpha = D = 3$  throughout the grid. This distribution does not accurately reproduce the low-density regions, placing almost all points in the dense filaments and their intersections. In contrast, the ‘hybrid’ sampling method provides sufficient resolution in both the low- and high-density regions. For the  $\alpha = D$  sampling, the cell size in the filaments and dense clumps is many times smaller than that of the original regular grid. No information is carried, however, by this extra resolution, effectively wasting computational resources.

<sup>1</sup>If the mean is not representative for the whole grid, we can use the maximal value of  $Q_y$  but we have found that using the mean gives excellent results in all tested cases.



**Figure 4.2:** Cut through a cosmological density field (left panel; grey scale indicates number density in logarithmic scale) sampled with  $128^3$  points using the sampling function of Eq. (4.1; middle panel) and Eq. (4.2; right panel) with the parameters obtained from analysis of the underlying density field.

### 4.3 CREATING THE SIMPLEX MESH: REGULARLY GRIDDED DATA

Hydrodynamical data is either stored as discrete particles (in the case of SPH) or on a mesh, possibly with an AMR structure. In this section we discuss some of the issues that naturally occur in the process of translating hydrodynamical data to the SIMPLEX mesh used for RT. We focus on two distinct but related properties of the mesh that are important to keep in mind: the *geometry* and the *density representation*.

With geometry we mean everything related to the placement of the vertices that form the basis of the mesh. Their placement determines the local resolution but also the anisotropies in the connecting edges. As we have seen in Chapter 3, the central issue with the creation of point sets suitable for SIMPLEX transport is that of strong gradients and we have studied the consequences of these gradients in isolated and idealized problems.

Even if the geometrical properties of the point set are such that the RT is not affected by systematic effects, the outcomes of a simulation can be wrong if the original density field is not represented faithfully. There are different ways to quantify if a field is represented well. In zeroth order, we can check if the mass of the original mesh is conserved. The first order requirement could be that this mass is in the right places. We don't want that our translation 'smears out' sharp features in the density for example. On the other hand, we don't want anomalies in the tessellation to result in high density spikes not present in the original data.


In the following sections we give examples of these two issues and show how they can produce some of the unphysical systematic effects of Chapter 3. In this respect, this section is the practical counterpart of the more idealized descriptions in Chapter 3. We begin with the description of issues related to gridded data and continue with particle based data afterwards in Sect. 4.4.

### 4.3.1 Case study 1: a simple regular grid

To get from grid-based data to the `SIMPLEX` mesh, we simply extend our sampling prescription Eq. (2.5) to the case where  $n(\mathbf{x})$  is discrete. Vertices are placed randomly inside the cells of the original grid and attain the value of that cell. Therefore, the density is represented perfectly at the resolution of the original data.

If all the cells of the regular grid are sampled with one or more vertices, the mass of the original data is conserved almost perfectly. The conservation is not perfect because the (irregular) volumes of the Voronoi cells are in general not equal to the cells of the original data. Fortunately, these fluctuations tend to cancel each other out and no systematic bias of the total mass is expected.


With regard to the *geometry* of the mesh we still need to tread carefully, however. The Poissonian nature of the `SIMPLEX` mesh is retained in this procedure unless the original data has large density steps between adjacent cells and the sampling power in that region is a strong function of density. In such cases, the distribution of vertices exhibits discontinuities that may lead to the problems discussed in Chapter 3. A simple example of this issue is shown in Fig. 4.3 where Eq. (4.1) is used to sample a  $3 \times 3$  grid with a density contrast of 3.6 (the density increases by this factor across a cell boundary). The discontinuities in the point set are strong (almost a factor 4) when Eq. (4.1) is used. When linear sampling ( $n_p \propto n$ ) is deployed, these discontinuous jumps are much smaller (density jumps are lower than a factor 1.6 for all cell-cell interfaces).



**Figure 4.3:** Example of the problem that a strong dependence of the sampling function on density of the original data can lead to a discontinuous point set. *Left:* The density grid with grey-scales indicating a density contrast of 3.6. *Central:* Use of Eq. 4.1 on this density field leads to differences in the point density of a factor of 3.7 maximally between adjacent cells. *Right:* Same as central panel but with a sampling power,  $\alpha$ , of unity. The discontinuities do now stay under a factor 1.6.

Such sharp discontinuities are expected to yield artifacts in the RT solution because of the local anisotropy of Delaunay edges. We focus here on the problems related to ballistic transport because this is the primary transport method employed in realistic simulations. The two systematical effects that we may expect are *ballistic decollimation* and *ballistic deflection* as described in Sect. 3.4.1.

We have performed the following test to verify whether the expected systematical effects play a role in a mesh with sharp discontinuities in the vertex density. A single source of ionizing radiation is placed in the centre of the domain and is allowed to create a spherical ionized region that stabilizes at a diameter of roughly half the box-size. The ionization front thus travels from a moderate number density of vertices towards lower values (to the lower-right) and higher values (to the upper-left). The number density of atoms is kept constant throughout the mesh so we expect the resulting ionised region to be spherical. We have performed the test with the mesh obtained with a sampling power of 3. A mesh with  $5 \times 10^4$  and one with  $5 \times 10^5$  vertices has been used for simulations employing ballistic transport and DCT in separate runs. In Fig. 4.4, these runs are shown with contours indicating a neutral fraction of 0.5 for both ballistic transport (black dashed curves) and DCT (grey curves).




**Figure 4.4:** Example of the RT artifacts due to a discontinuous point set for the mesh obtained with a sampling power equal to 3. Contours show the equilibrium ionisation front. The solid contour indicates the DCT result and the dashed curves are for the ballistic transport result. *Left:* Grid containing  $5 \times 10^4$  vertices. *Right:* Grid containing  $5 \times 10^5$  vertices.

In the low resolution result (left panel of Fig. 4.4), the ionisation front is highly aspherical with a protrusion into the low resolution section of the mesh. The artifact is visible in both ballistic and DCT result suggesting that it has to do with a local anomaly in the mesh that influences both transport methods. This is not surprising because in both cases, the radiation is transported along the Delaunay edges. If these edges are very anisotropic, the transport in both cases will be hampered in one direction. It is difficult to say, even in this simple situation, how much of this a-sphericity is due to ballistic decollimation and how much to ballistic deflection but it is clear that we are dealing with problems due to local anisotropy of the mesh.

This becomes more obvious if we look at the inner ionisation structure of the Strömgren ‘spheres’ shown in Fig. 4.5. The ballistic result shows several artifacts due to local anisotropies in the mesh that are not present in the DCT result which is to be expected as DCT is less






**Figure 4.5:** Example of the RT artifacts due to a discontinuous point set for the mesh obtained with a sampling power equal to 3. *Left:* Ballistic result. *Right:* DCT result.

susceptible to such anisotropies.

One straightforward solution to this type of geometrical problem is to use a sampling function that depends less sensitively on the density and thus yields a smoother distribution of vertices. Another solution is to simply use more vertices (keeping the sampling function the same). The contours of the same test using 10 times as many points is shown in the right panel of Fig. 4.4. Increasing the number of vertices yields a Strömgren region with a more spherical shape. This is somewhat surprising because the discontinuous jumps in the number density are of the same magnitude. If the artifacts shown in Fig. 4.4 and Fig. 4.5 are indeed systematic and result from the jumps, we would see their effect in the high resolution result as well. We are thus led to conclude that the observed artifacts are due to chance ‘glitches’ in the mesh, the influence of which decreases as the number of vertices increases, providing more paths for the light to travel along.

The test presented above shows an ionisation front that travels along the boundaries of several cells of different vertex number density. For completeness we have performed a similar test but ‘zoomed in’ on one such cell boundary. In this case the contrast in number density is chosen to be twice as big as in the previous case, a factor of 8 between the low and the high number density regions. Again the number density of gas is kept constant throughout the mesh. We have placed two sources in the domain, one in the high and one in the low resolution region. The radiation travels with DCT only for this test. Any negative effects due to the sharp jump in resolution would distort the Strömgren spheres. Moreover, we can directly assess if there is a difference between the direction with which an ionisation front crosses such discontinuities.

In Fig. 4.6, the neutral fraction of the stable Strömgren spheres is shown for the setup described above. The employed direction conserving transport can cope surprisingly well with this challenging setup as the deviations from sphericity are only of the order of the local cell size. Although somewhat academic, this result has some relevance for the (post-)processing of AMR data. In AMR meshes, a single refinement step is usually achieved by dividing each dimension in two, which yields cells that are eight times smaller. If we were to generate a RT



**Figure 4.6:** Strömgren sphere test for a grid with a resolution discontinuity of a factor 8. The sources are shown as black dots and the expected ionisation-front position is shown as a dashed white circle. The source is placed in the high (left) and low (right) density half of the grid respectively. The grey scale indicates the neutral fraction of the material.


mesh that follows the refinement of the original mesh, steps of 8 in resolution occur naturally. The test shown here suggests that such a simple approach would not be unfeasible because of problems in the RT.

### 4.3.2 Case study 2: Undersampling in grid-based data

As already mentioned in Sect. 3.4.5, sharp gradients in the point number density can result in large cells receiving too few photons. The reason for this is simply that the  $D$  most straightforward pointing outgoing edges are directed almost perpendicularly to the direction into the large cell (see Fig. 3.12). This can be regarded as a problem of ballistic decollimation *in one step*. Such extreme gradients can occur if insufficient points are used to resolve the regions of lowest density.

In Fig. 4.7 a result of the first incarnation of the SIMPLEX algorithm in Test 4 of the first cosmological radiative transfer comparison project (Iliev et al. 2006a) is shown alongside the improved result using the same version of the method and ballistic transport without weights. The only difference is the sampling strategy used to obtain the point distribution. The result published in Iliev et al. (2006a) (their Fig. 33) is obtained with  $128^3$  points and a sampling function  $f(n(\mathbf{x})) = \sqrt{n(\mathbf{x})}$ . The new result uses Eq. (4.2) with the parameters derived in this section resulting in the point set depicted in the right panel of Fig. 4.2.

Evidently, there exist large cells in the ‘old’ mesh that stay neutral due to the lack of inflowing photons. These artifacts are completely absent in the ‘new’ result obtained with hybrid sampling. Although grid-based data is used in this example, the problem of undersampling low



**Figure 4.7:** Comparison between the result with a uniform sampling power as was performed for the Radiative Transfer Comparison Project (Iliev et al. 2006a, left) and the hybrid sampling function of Eq. (4.2) (right). The number of mesh points is in both cases  $128^3$ . To facilitate their comparison, both meshes have been interpolated to a regular grid of  $128^3$  cells. Note that the interpolation to the structured grid introduces some artefacts close to the boundary. Shown is a slice through the  $z = z_{\text{box}}/2$  coordinate of the computational domain at  $t = 0.2$  Myr, as the influence of the incorrect sampling is most pronounced at that time.

density regions can occur in SPH simulations as well. Especially when a subset of the total number of particles is used for the RT.

### 4.3.3 Case study 3: 3D hydro of Eta Car

In practice it is hardly ever desirable to place one or more vertices in every cell of the regular (AMR) mesh (as in the previous section) because of the memory requirements of RT are in general more stringent than for the hydrodynamics. Moreover, the full flexibility of the `SIMPLEX` mesh can only be exploited if we allow the sampling algorithm to represent less interesting regions with Voronoi cells that are (much) larger than those of the original grid.

If cells of the original data are not sampled by a `SIMPLEX` vertex, their mass must be accounted for somehow to ensure mass conservation. We could use either a zero-th or higher order interpolation method for this purpose. Below we explore two possibilities, one zero-th order and one first-order interpolation method which we will call the *Voronoi* and *Delaunay* method respectively. We incorporate the density estimate without mass-conservation in our discussion as a reference and we will call it the *density conserving* method because it simply copies the densities of the original grid to the vertices.

In all cases, the situation is initially as follows. A subset of the original cells has been sampled by placing vertices inside them. These vertices make up the RT mesh which, at this point, does not include all the mass of the original grid. The cells that have not yet been included in the RT mesh are flagged. Using the flag, the un-sampled cells are included in the RT mesh

except for the case of the density conserving method.

The Voronoi method consists of finding the un-sampled grid cells within each Voronoi cell of the RT mesh (that consists of all vertices that have been used to sample the data) and assigns the mass of these cells to the corresponding vertex. The density of the cell is found by dividing the mass of the cell by its Voronoi volume

$$\rho = \frac{m}{V_{Vor}}. \quad (4.6)$$

Because of the definition of a Voronoi cell, this is equivalent to a loop over all un-sampled cells and finding the closest vertex for each of them. We have implemented this using an octree which optimizes the search in 3D space. We note that we do not account for the geometrical overlap between the regular cells and the Voronoi cells. If the centre of a regular cell lies within the Voronoi cell, its mass goes completely to the nucleus of that Voronoi cell.

In the Delaunay method, one searches for the Delaunay simplex in which every un-sampled cell is located. The mass of the cell is divided equally over the four vertices that define the simplex. Every vertex thus gets mass from un-sampled cells within its so-called *contiguous* Voronoi cell. The volume  $W_{Vor}$  of a vertex is now defined as the joint volume of all Delaunay simplices that share that vertex,  $W_{Vor} \equiv \sum_i^N V_{Del,i}$ , where  $N$  is the number of neighbours of the vertex. With this adjustment, Eq. (4.6) transforms to

$$\rho = \frac{m(D + 1)}{W_{Vor}} \quad (4.7)$$

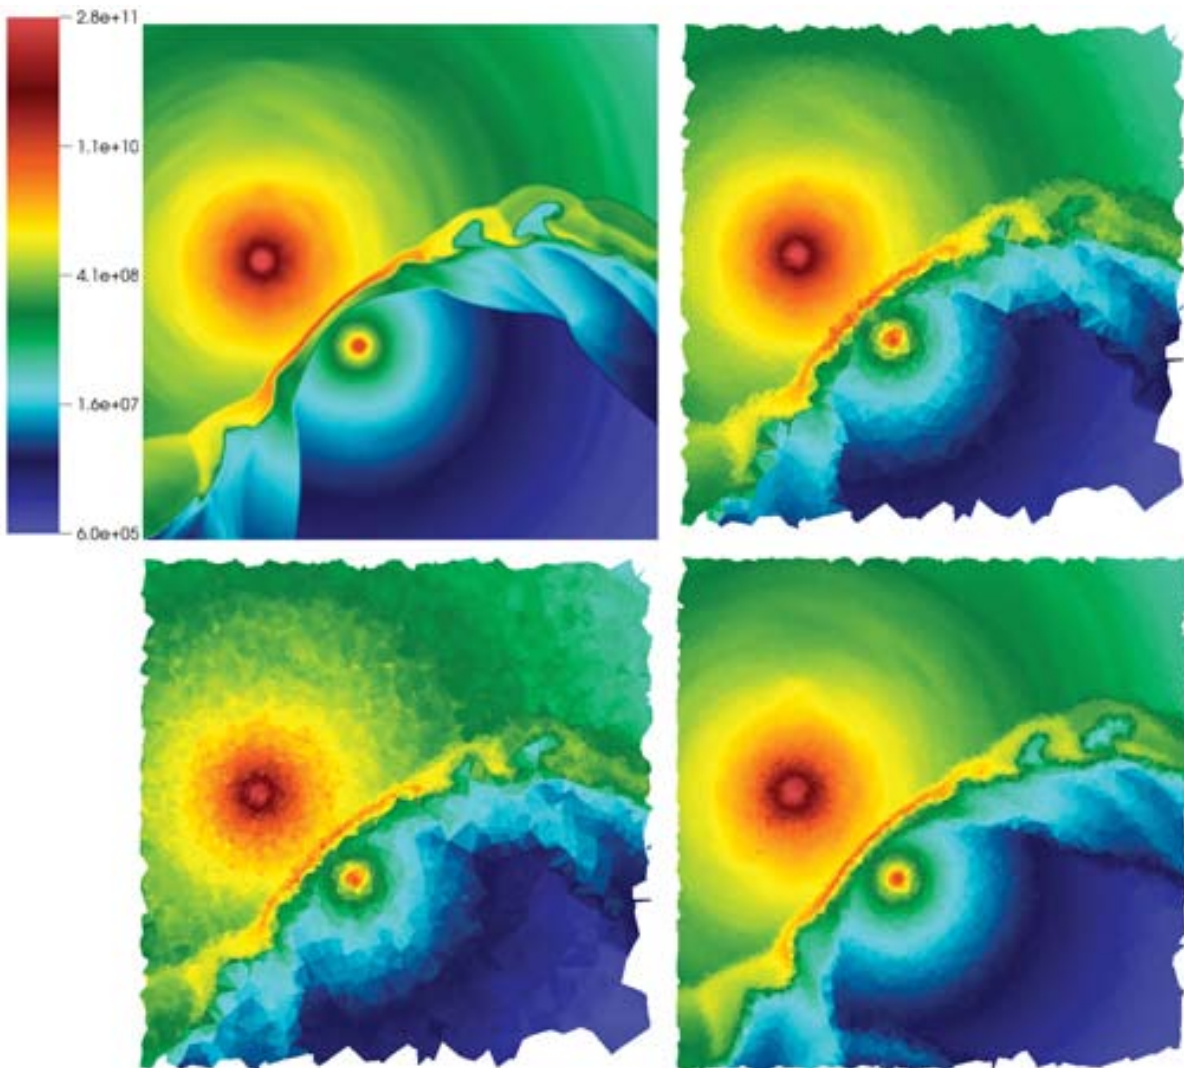
where the factor  $(D + 1)$  comes in because every Delaunay simplex has  $(D + 1)$  vertices where  $D$  is the dimension. We note that the manifestly mass-conserving *Delaunay Triangulation Field Estimator* (DTFE Schaap & van de Weygaert 2000) uses the same contiguous Voronoi cells combined with linear interpolation for accurate representation of point-based data. Given the fact that this Delaunay method uses a form of linear interpolation to assign ‘missed’ mass to the RT mesh, we expect the resulting density field to be smoother than the Voronoi equivalent. In Fig. 4.8, a visual comparison of the density fields obtained with the three methods described above is shown. The top left panel shows the original AMR data with an effective resolution of roughly  $200^3$ . A total of  $2 \times 10^5$  (approximately  $60^3$  and thus a factor 40 smaller than the original cells) vertices is used to construct the RT mesh. We have chosen a relatively low number of vertices in order to demonstrate more clearly the areas where issues do arise.

The vertices are placed according to a sampling function that has a shallow (square-root) dependency on the density in except for the high density region very close around the central source which was sampled with a constant number density. This density cutoff is a necessary evil because the density around the central source is a factor of 100 higher than that in the contact discontinuity and the majority of vertices would be placed there even with a low sampling power such as used here. Note that (with the obvious exception of the original data) all density fields shown in Fig. 4.8 use the same SIMPLEX mesh and differences are solely due to different density estimates. In this section we thus concern ourselves with an optimal representation of the original density field while keeping the geometry constant (in contrast to the previous section where the situation was reversed).

To obtain the field shown in the top right panel, every vertex was assigned the density of the cell it was placed in. Accordingly, the result is an almost literal copy of the original grid

but with a different resolution and no guarantee of mass conservation. In regions where the density changes gradually (in particular the subtle waves in the primary wind in the upper part of the simulation domain) are reproduced extremely faithfully as must be expected from this procedure. Problems arise, however, in the regions where density discontinuities occur, e.g., in the wind-wind interaction region. Because a vertex placed around a discontinuity can attain either the high or the low value but nothing in between, these regions look puffy and crudely represented.


The Voronoi method for including the un-sampled cells does a much better job in reproduc-



**Figure 4.8:** *Top left:* Density field of the original data in log-scale. *Top right:* Density conserving RT mesh (without mass conservation). *Bottom left:* Mass conserving RT mesh with octree-based mass inclusion (Voronoi method). *Bottom right:* Mass conserving RT with simplices-based mass inclusion (Delaunay method).

ing these regions of sharp density contrast. The density discontinuity is better resolved as the un-sampled cells can correct the initial density estimate of the sampling procedure. In regions where there are relatively few cells per vertex, however, the density of a Voronoi cell depends sensitively on the number of un-sampled cells that are included by it. This can be seen as the grainy structure in the un-shocked primary and secondary wind. The severity of this ‘noise’ depends on the problem at hand. If the number of un-sampled cells per Voronoi cell is large, the noise is small relative to the total density but if there are only few un-sampled cells per Voronoi cell, density spikes can emerge that were absent in the original data. If the recombination rate (which depends on the square of the density) is important, these density spikes may seriously affect the result of a RT simulation.

The Delaunay method seems to strike an optimal balance between the non-mass-conserving sampling and the noisy Voronoi method. Because of its first order interpolation (based on the contiguous Voronoi cell), noise is suppressed while detail on small scales is retained almost at the same level as in the Voronoi method. Apart from several small density spikes close to the primary star, the smooth un-shocked winds are represented faithfully without compromising the accuracy in the contact discontinuity.



**Figure 4.9:** *Left:* Volume weighted probability density function (equivalent to mass) as a function of number density of the hydro-data and corresponding SIMPLEX meshes using different mass-inclusion procedures. *Right:* Same as left panel but normalized by the PDF of the original data.

In Fig. 4.9, a more quantitative measure for the difference between the density representations discussed above is given. The figure shows the probability density function of the volume-weighted density. This is a graph of the fractional contribution from each density bin to the total mass. The density PDF shows us on a global scale if mass is ‘where it should be’ and which densities dominate the total mass.

The left panel of Fig. 4.9 shows the four PDFs plotted together. We can immediately see that the Delaunay method overlaps with the original data over the whole density domain. The other methods show some scatter at the high density part of the PDF, but this is inevitable and of no concern as the density cut-off was placed at  $5 \times 10^{10} \text{ cm}^{-3}$  and anything beyond that density is only sampled with a few vertices. The curve of the Voronoi mesh lies systematically above the

original data for densities higher than the  $\approx 5 \times 10^8 \text{ cm}^{-3}$  exactly where the spikes in the primary wind show up. This simply means that some mass is transferred (most probably from the mean) to this density range. Also, the Voronoi method does not reproduce most of the wiggles in the original data that are captured accurately by the Delaunay method and to a lesser extent by the density conserving method.

The right panel shows the same data as the left panel but normalized to the original result. The solid line is unity and shown for reference and the other line-types are as for the left panel. Deviations from the original data are largest for the density conserving method (up to almost a factor of ten) where the mass is over-estimated for the low and medium density and under-estimated for densities larger than  $10^9 \text{ cm}^{-3}$ . The Voronoi method performs better but fails to reproduce some of the more subtle wiggles in the PDF. The Delaunay method clearly shows superior performance with all fluctuations below the 5% level (except above the cut-off density where the result is not to be trusted anyway).

#### 4.3.4 Case study 4: Cosmological data


Although the density in the hydrodynamic data used in the previous section ranges about five orders of magnitude, strong density fluctuations are few except in the wind-wind interaction region. Such data is relatively easy to represent in stark contrast to the cosmological data used in this section. Cosmological density fields have a distinct web-like structure consisting of clumps connected by filaments which are in turn connected by walls. Density gradients increase on all scales with decreasing redshift due to the gravitational attraction of matter.

The data that we use in this section has  $64^3$  cells and serves as the basis for several runs in that chapter and is taken from [Finlator et al. \(2009b\)](#). It represents  $16h^{-1}$  comoving Mpc at redshift  $z = 6$  and contains densities that span more than two orders of magnitude. This is much less than the five orders of magnitude of the original data because of the relatively crude gridding.

In the left panel of Fig. 4.10 a slice of the SIMPLEX mesh consisting of  $64^3$  vertices is shown. The sampling power is set to zero in order to eliminate (down to the level of Poisson fluctuations) the possible influence of adaptivity of the mesh. The volume-weighted Probability Density Functions (PDFs) of the original grid are compared to those of three SIMPLEX meshes of different resolution in the right panel. If the regular grid is sampled with fewer points than the number of grid-cells (for example the ‘Vor 32’ PDF), content of high and low density cells is smeared out increasing the importance of average densities. Furthermore, the peak of the DPF shifts about 0.1 dex to higher densities. In this case, information of the original grid is not conserved, with possibly severe errors in the resulting radiation transfer solution. Recombinations, for instance will be systematically under-estimated by the ‘Vor 32’ mesh, which leads to incorrect speeds of ionisation fronts in simulations of reionization.

The solution to the problems described above is trivial, the number of points that sample the original data must be equal or larger (both the Vor 64 and 128 PDFs are almost perfectly converged with that of the original data) to capture all relevant details.

If we want to fully exploit the advantages of the adaptive resolution of the SIMPLEX mesh, we must abandon the constant sampling function ( $\alpha = 0$ ) and allow the point number density to adapt itself to fluctuation in the underlying medium. As an example, we have used a simple



**Figure 4.10:** *Left:* Slice through the homogeneous SIMPLEX mesh at  $z = L_{box}/2$ . The Delaunay simplices are shown. *Right:* Density probability density functions of the  $64^3$  grid used in [Finlator et al. \(2009b\)](#) (Fin 64) and the derived SIMPLEX meshes (Vor 32, 64 and 128) where the number indicates the cube-root of the number of vertices. A constant sampling is used ( $\alpha = 0$  in Eq. (4.2)). The PDFs are stored as 50 logarithmically spaced bins.

linear sampling function. The resulting mesh (again containing  $64^3$  vertices) is shown in the left panel of Fig. 4.11. In contrast to the mesh shown in Fig. 4.10, the number density of vertices is inhomogeneous. This inhomogeneity allows SIMPLEX to treat dense regions with higher precision, but it comes at a price. Because there are in this example just as many vertices as cells in the original data, every refinement implies an under-sampling of (lower density) regions. This under-sampling of low-density regions shows up clearly in the density PDF shown in the right panel of Fig. 4.11. The curve indicated by Vor 64 lies under the original curve for almost the entire underdense part of the abscissa.

#### 4.4 CREATING THE SIMPLEX MESH: PARTICLE BASED DATA


In the case that the density field is given by discrete SPH particles, we might obtain an estimate of the density at any position in the domain using a typical SPH kernel function  $W(r, h)$

$$\rho(r) = \sum_j m_j W(|r - r_j|, h_j), \quad (4.8)$$

where  $h$  is the smoothing length and  $m_j$  is the mass of particle  $j$ . We note that the summation is not necessarily over a fixed number of neighbours here, hence the range is omitted. Using this kernel function we can thus sample the data with our usual tools (e.g., Eq. (4.2)) and the statements of Sect. 4.3.2 apply directly to the resulting point set.

Unfortunately, there is a serious drawback to this approach. The majority of kernel functions in the literature and applications has a spherical symmetry. Such kernels have the tendency to





**Figure 4.11:** *Left:* Slice through the adaptive mesh at  $z = L_{box}/2$ . The Delaunay simplices are shown. *Right:* Density probability density functions of the  $64^3$  grid used in [Finlator et al. \(2009b\)](#) (Fin 64) and SIMPLEX meshes (Vor 32, 64 and 128) where the number indicates the cube-root of the number of vertices. The sampling function used looks like  $n_p \propto n$ . The PDFs are stored as 50 logarithmically spaced bins.

smooth out anisotropic features in the density field. A related disadvantage is that smoothing leads to poor resolution of physical density discontinuities. Differences with less dissipative density estimates such as the DTFE can become as large as several orders of magnitude ([Pelupessy et al. 2003](#)). Aside from these issues, the SPH kernel does not strictly conserve the total mass of a continuous density field. These features of SPH kernel based density estimates may directly influence the outcome of radiative transfer calculations which makes them less well-suited for our purposes.


Given these drawbacks and the fact that the original data is already particle-based, it might be more natural to use the particles themselves as the generating nuclei for the Voronoi-Delaunay mesh. In that case we do not need to use the intrinsically diffusive Eq. (4.8) to estimate the density but a more direct estimate can be provided by division of the particle mass by the Voronoi volume of its corresponding cell as described by Eq. (4.6) or Eq. (4.7). We come back to the differences in these density estimates in Sect. 4.2.4. Another advantage is that, due to pressure forces, the particles in an SPH simulation are in general positioned more regularly than for a pure Poisson process. Finally, we note that with future applications of radiation-hydrodynamics in mind, a coupling of SIMPLEX with an SPH method is most natural when the radiation transport is applied directly on the SPH particles so that no spurious interpolation is needed.

#### 4.4.1 Case study 5: Undersampling in SPH data of Eta Carinae

We have already seen in Sect. 3.4.5 and Sect. 4.3.2 how under-sampling can have a negative effect on the outcome of an RT simulation. The same problem holds in principle for particle based data.

To check if our results of Chapter 9 are prone to such issues, we use a typical snapshot including  $5 \times 10^5$  particles that show densities ranging roughly ten orders of magnitude. In the left panel of Fig. 4.12, a density cut through the SIMPLEX mesh of the original data at half the box size is shown. For the sake of clarity, densities above  $10^7 \text{ cm}^{-3}$  have not been included in the plot (they occur closely around the primary star and are not of interest to our study of the wind-wind interaction region). More important are the sharp discontinuities in both density and number density of SPH particles seen along the low-density ‘fingers’ that permeate the volume.


In this section we take a pragmatic approach and simply increase the resolution of the sparsely sampled regions in the data. If the RT result turns out to be sensitive to sharp gradients, we expect significant differences between the simulation with increased resolution and the original simulation.



**Figure 4.12:** Number density of hydrogen in the SPH particles in the WIND900 simulation (Madura 2010) at half the box-size. *Left:* Original data. *Right:* Data with increased resolution. Besides an increase in resolution, the number density becomes slightly higher due to the interpolation.

We augment the resolution of the sparsely sampled regions as follows. First, a triangulation of the SPH particles has been constructed. For every tetrahedron in this triangulation that is larger than a tolerance volume, an additional vertex is placed in the centre of the tetrahedron and  $1/5$  of the mass of the four vertices that constitute the tetrahedron is given to the new vertex. This procedure is thus manifestly mass conserving and regularizes the mesh in regions of low resolution.

In the right panel of Fig. 4.13, an example of mesh obtained with this procedure is shown. The lower limit of the resolution has been taken to be  $100^3$  in this case. So, tetrahedra larger than  $10^{-6}$  of the total volume (which is approximately unity) will be subdivided into four by the insertion of a vertex in their centre. The resulting (number) density field is compared to the original in Fig. 4.12. Although the difference is not large, the overall density in the aforementioned



**Figure 4.13:** *Left:* Cut through the Delaunay triangulation of the mesh shown in Fig. 4.12. *Right:* Same as left panel but with extra sample points in the low number density region East of the central source.


Eastern region is slightly enhanced by the interpolation.

To study the influence the procedure has on the radiative transfer we have performed a preliminary RT simulation including flux from the secondary star only. The result after  $0.1 \text{ yr}^2$  of physical time is shown in Fig. 4.14. Although the overall shape of the ionised regions does not change, the ionisation fraction in the region most affected by the added resolution is somewhat lower after the procedure. This is consistent with the notion that the interpolation increases the density somewhat in this region with a higher recombination rate as a consequence. These results suggest that (locally) increasing resolution does not change the overall shape and size of the ionised regions. It nevertheless affects the ionisation structure in the affected regions. At first glance, this change is related to the increased density in regions of lowest density resulting from interpolation. We are therefore confident that the RT results of Chapter 9 are not susceptible to systematic effects related to strong gradients in the particle number density.

#### 4.4.2 Case study 6: Sampling cosmological data

Because of the high computational cost of radiative transfer compared to, for example, gravity or hydrodynamics, it is often impossible to perform RT on all particles of the original data. We must therefore resample the density field to obtain a new point-set or select a subset of these particles. We focus here on the latter option because it retains many of the advantages of using the SPH particles directly for RT. The first most notable advantage is that each particle can be assigned a unique ID which is necessary when information must be carried from one snap-shot

<sup>2</sup>This ionisation structure does not evolve much after this time which corresponds roughly to the inter-snapshot time of the original SPH simulation.



**Figure 4.14:** Same as Fig. 4.13 but for the neutral fraction of hydrogen for the stable Strömgren region induced by the secondary star. The ionised part of the simulation volume is fairly robust to the increase in resolution. As expected, the Eastern region is less ionised after adding the extra vertices because the extra density results in increased recombinations.


to the next as is often the case in post-processing or radiation-hydrodynamics. The second is the previously mentioned regularity of SPH particles resulting from their hydrodynamic interactions.

When selecting a sub-set of SPH particles to obtain a suitable set, care must be taken that mass is conserved and, more subtly, mass is conserved ‘in the right places’. We now show an example of how we ensure mass-conservation in our mesh-creation procedure in a manner that minimizes the inevitable loss of information related to degrading resolution.

If the selected subset is representative of the original set, in other words: if the density PDF of the subset is close to the original, we could do the following. For every particle in the subset, take the density estimate of the original data (obtained with e.g., Eq. (4.8)) and assign this density to the Voronoi volume of that particle in the SIMPLEX mesh.

Unfortunately, this hardly ever gives the correct answer. With decreasing number of particles in the subset, ever larger structures present in the original data are replaced by a single particle. This single particle now takes the volume of, for example, a halo of 128 original particles. Assume now that the subset has been selected at random. This implies that the particle selected to represent the halo is probably one with a density above the average (because there are, per definition, more particles per unit volume in the high density regions). When its density is assigned naively to the according Voronoi volume, the mass of that region is always over-estimated.

The most direct way to deal with this issue is to abandon the original density estimate and work with the masses of the SPH particles instead and use Eq. (4.6) or Eq. (4.7) for the density estimate. To conserve mass, we now have to associate the mass of every un-sampled particle to its closest neighbour in the subset. We use the octree-based Voronoi method for this purpose.



**Figure 4.15:** Scatter plots of the quotient of volume estimates for the SPH data used in Chapter 7. Mean values with  $\pm 1\sigma$  ‘error’-bars are shown for 20 logarithmically spaced bins. *Top panel:* All particles are included (although only every 64-th particle is plotted to avoid clutter). *Centre panel:* As for the top panel but now a  $64^3$  particle subset of the data is used. All densities (and, accordingly, volumes) are scaled up linearly to make the total volume equal to unity. *Bottom panel:* As for the centre panel but now the masses of un-sampled particles are added to their nearest neighbour in the sampled set.

An example of the issue described above is shown in Fig. 4.15. Instead of the densities, we plot the volumes that are estimated by the SPH kernel and the Voronoi tessellation<sup>3</sup>. We show the quotient of SPH and Voronoi volume estimates for a typical cosmological density field as used in Chapter 7. We have taken a snapshot at redshift  $z = 6$  including  $512^3$  particles that show densities ranging roughly five orders of magnitude. The top panel of Fig. 4.15 shows the scatter plot for all data points. Grey dots represent values for the quotient of the Voronoi and SPH volume and mean values are indicated with black dots with  $\pm 1\sigma$  ‘error’-bars.

Even when all particles are included, the Voronoi and SPH volumes differ. The deviations from unity are largest for intermediate (over-)densities but these are small compared to the scatter. The Voronoi volumes are slightly (less than 0.15 dex) larger for intermediate densities. This is a consequence of the SPH kernel over-estimating the density in underdense regions close

<sup>3</sup>This is completely equivalent because in both cases, the density is defined as the quotient of mass and volume and mass is equal for the two methods.

to overdensities.

In the case that a subset of particles is chosen we have multiplied the original SPH volumes by a constant factor in order to make their sum equal to the total volume of the Voronoi cells for sake of easy comparison (middle panel). In the slightly overdense (roughly between 0 and 1) regions, we clearly see the effects of the sampling preference for particles with high densities (and thus small volumes), the Voronoi cells can be more than an order of magnitude larger. For the underdense regions the Voronoi volumes are generally smaller than the SPH volumes because they compensate for the lost volume in the high density regions when the SPH volumes are scaled up. Particles with the highest overdensities are surrounded by particles with similar densities. These particles have a high probability of being selected and in these regions the SPH and Voronoi estimates are in good agreement again.

The bottom panel of Fig. 4.15 shows the volume quotients for the subset with correct inclusion of the un-sampled mass. Although the scatter is somewhat larger than for the complete particle set, it stays under 0.2 dex and the mean is consistent with the original data. In Chapter 7 we will use this technique to reduce the number of SPH particles in cosmological density fields from  $512^3$  to  $128^3$  and even  $64^3$ .

## 4.5 SUMMARY

- We have introduced our main sampling procedure in Sect. 4.2 and given some examples of its use. The central idea is that often a density field exhibits two regimes in number density that need to be sampled with a different relation between density of vertices and number density for optimal result. Our sampling procedure smoothly links these two regimes based on two adjustable parameters.
- Unwanted geometrical effects due to (a too high) sample power have been investigated in Sect. 4.3.1. We showed how the shape of an ionised region can become deformed due to discontinuous jumps in the vertex number density. This can be prevented by increasing the number of vertices or by using DCT which is less sensitive to geometrical artifacts in the mesh.
- In Sect. 4.3.2 we have looked at the effects of incorrect sampling of cosmological data. The high dynamic range of the original data results in neighbouring cells with very different volumes. This can result in severe problems where radiation effectively by-passes large cells, leaving them too neutral.
- We have discussed several mass-inclusion methods (Sect. 4.3.3) and argue that the Delaunay method based on the contiguous Voronoi cell gives more satisfactory results than the non-mass-conserving ‘density conserving’ and noisier Voronoi method.
- A study of the effects of resolution and adaptivity of the mesh in the context of density representation (Sect. 4.3.4) and of geometry (Sect. 4.4.1) has also been included. These two sections are complementary to the studies of Chapters 7 and 9 and justify the mesh-construction criteria of those chapters.



## CHAPTER 5

---

# Heating, cooling and multiple frequencies

## 5.1 INTRODUCTION

In this chapter we present several extensions to and improvements of the physics module of SimpleX. Most notably the incorporation of helium and dust (in addition to hydrogen), heating and cooling functions appropriate for a primordial composition and a multiple frequency treatment of the radiation field.

The original version was limited to monochromatic ionization of hydrogen only without temperature evolution. Even with such approximate physics, state-of-the art cosmological simulations are possible because the refinements described in this chapter are not of vital importance for many applications. The large-scale distribution of ionized material as a function of redshift is fairly un-sensitive to these approximations for example.

There are, however, many problems in galaxy- and star formation where a detailed treatment of the temperature of the gas is essential. The collapse of gas clouds, subsequent formation of molecules and the resulting onset of nuclear processes that comprise the essential steps of star-formation are all highly sensitive to temperature. The same is true for feedback effects of stellar radiation on the gas such as the expansion of H II regions due to photo-heating and collisional destruction of molecules.

Also, the thermal state of the interstellar medium (ISM), intergalactic medium (IGM) and intracluster medium (ICM) is an important observable and its evolution can be studied in a cosmological context. The thermal evolution of the cosmic gas is, for instance, strongly determined by the details of reionization and, as such, can be used to distinguish between reionization scenarios (e.g. [Miralda-Escudé & Rees 1994](#); [Theuns et al. 2002](#); [Hui & Haiman 2003](#); [Tittley & Meiksin 2007](#)).

Because we live in the era where the first direct observational handles on reionization become within reach ([McQuinn 2010](#)), it is timely to concern ourselves with the modeling of the expected signal as to aid its interpretation. The first step to this goal is to include a proper treatment of relevant physical processes.

This chapter is structured as follows. We describe the processes that determine the ionization-state of the gas, such as photo- and collisional ionization and recombination, and how we evolve it in Sect. 5.2. In Sect. 5.3 we show how we follow the thermal state of the gas. How the ionization-state and temperature depend on the use of multiple frequencies is the topic of



Sect. 5.4. We end with a summary and give explicit expressions for the rates and cross sections used in the appendix.

## 5.2 IONISATION-STATE OF THE GAS

In this section we describe the physical processes that we include in our physics treatment and the specifics of their implementation.

### 5.2.1 Ionisation

In a gas with cross section for photo ionization  $\sigma(\mathbf{x}, \nu)$  at position  $\mathbf{x}$ , the local photoionization rate,  $\Gamma_{P,i}(\mathbf{x})$  (which gives the number of photoionizations per second per atom of species  $i$  in units  $[s^{-1}]$ ), is given by (e.g., [Osterbrock & Ferland 2006](#)):

$$\Gamma_{P,i}(\mathbf{x}) \equiv \int_0^\infty \frac{4\pi J_\nu(\mathbf{x})}{h\nu} \sigma_i(\mathbf{x}, \nu) d\nu, \quad (5.1)$$

where  $J_\nu(\mathbf{x})$  is the local mean intensity and the three ( $i = 1, 2, 3$ ) species capable of absorbing ionizing photons are H I , He I and He II . In addition to these species, a contribution to the opacity due to dust can be included analogously.

In addition to photoionization, we include collisional ionizations due to the interaction of free electrons and neutral atoms. As this is a kinetic process, the collisional ionization rate,  $\Gamma_C$  depends on the thermal state of the electrons and is given by

$$\Gamma_C = n_e \sum_i \Gamma_i(T) n_i, \quad (5.2)$$

where the collisional ionization rates  $\Gamma_i(T)$  are given in Tab. 5.2 and  $n_e$  is the electron number density. The total ionization rate is thus given by the sum of photo- and collisional ionization rates,  $\Gamma = \Gamma_P + \Gamma_C$ .

### Discretization of the ionization rate

Implementing ionization processes in a numerical code requires that the relevant equations can be expressed in a discretised form. In particular, we need to know the ionization rate in each cell of our computational grid. More precisely, we need the ionization rate *per species* and *per frequency*.

In SIMPLEX, ionizing radiation travels from cell to cell along the Delaunay edges. At the nucleus of each Voronoi cell, photons are taken away from the incoming radiation field and their energy is used to ionize the neutral atoms of that Voronoi cell and heat the gas. Given the number densities of these species,  $n_{\text{HI}}$ ,  $n_{\text{HeI}}$  and  $n_{\text{HeII}}$  and the path length through the cell  $l$ , the monochromatic optical depth of ionizing radiation  $\tau_\nu$  is given by

$$\tau_\nu \equiv (n_{\text{HI}}\sigma_{\text{HI}} + n_{\text{HeI}}\sigma_{\text{HeI}} + n_{\text{HeII}}\sigma_{\text{HeII}})l. \quad (5.3)$$

The *total* number of ionizations per unit time,  $\dot{N}_{\text{ion}}$  for a cell with optical depth  $\tau_\nu$  is then given by

$$\dot{N}_{\text{ion}} = \int_0^\infty \dot{N}_\gamma(\nu)(1 - \exp(-\tau_\nu))d\nu, \quad (5.4)$$

where  $\dot{N}_\gamma(\nu)$  is the number of ionizing photons per unit time streaming into the cell. To quantify how much of the resulting ionizations is due to a particular species, we can use the contribution to the total optical depth of that species. The number of ionizations of species  $i$  per unit time is given by

$$\dot{N}_{\text{ion},i} = \dot{N}_{\text{ion}} \int_0^\infty \frac{\tau_{\nu,i}}{\tau_\nu} d\nu \quad (5.5)$$

If we divide this by the number of neutral atoms of species  $i$  in the cell,  $N_i$ , we have the spatially discretized equivalent of Eq. (5.1)

$$\Gamma_{P,i} = \frac{\dot{N}_{\text{ion},i}}{N_i}. \quad (5.6)$$

Discretising this result in frequency can be done by simply evaluating Eq. (5.4) and (5.5) for a set of bins, each bounded by two values of  $\nu$ . The optical depth,  $\tau_{i,j}$ , thus carries two indices, one for the species and one for the frequency bin.

For a frequency bin  $j$ , the total number of ionizations per unit time is simply given by

$$\dot{N}_{\text{ion},j} = \frac{\tau_j}{\tau} \dot{N}_{\text{ion}}, \quad (5.7)$$

where  $\tau_j \equiv \sum_i \tau_{i,j}$  is the optical depth of the  $j$ -th frequency bin and  $\tau$  is the total optical depth (integrated over frequency and summed over species). To find the number of ionizations for a species  $i$ , we need to sum over all frequency bins and select the contribution of that species

$$\dot{N}_{\text{ion},i} = \sum_j \frac{\tau_{i,j}}{\tau_j} \dot{N}_{\text{ion},j} = \frac{\dot{N}_{\text{ion}}}{\tau} \sum_j \tau_{i,j}, \quad (5.8)$$


where we have substituted Eq. (5.7) in the last step. In other words; once we know  $\tau_{i,j}$ , we can quickly find the ionizations for a frequency bin and for a species.

## Dust extinction

The procedure described above can be extended trivially to include absorption due to dust particles. We use the empirical dust model described in [Gnedin et al. \(2008\)](#). The extra opacity due to dust particles effectively removes photons from the radiation field because we do not follow the possible re-emission of thermal radiation from the dust.

To avoid technicalities related to size distributions and shape of the grains, [Gnedin et al. \(2008\)](#) use measurements of extinction towards the Small and Large Magellanic Clouds to model dust extinction. The cross section of dust is taken relative to the either the neutral or total hydrogen abundance. The first choice is motivated by the observation that most dust will not form at temperatures typically seen in ionized gas.

Although we will not go into specifics here, we note that in most cosmological applications, dust will not have a large influence on the ionizing radiation field because of two reasons. First,



**Figure 5.1:** Cross section of hydrogen (solid), SMC dust (dotted) and LMC dust (dashed) as function of frequency normalised to the frequency of the Lyman limit.

the cross section of dust is orders of magnitude smaller than that of hydrogen at the ionization frequency (see Fig. 5.1). Second, the metallicity in the early universe is lower than in Milky Way. The difference in cross section between SMC and LMC dust is mainly due to the higher metallicity of the latter.

### 5.2.2 Recombination

The inverse process of ionization is recombination. This free-bound interaction of electrons and ions depends on temperature and number density of ions and electrons. The number of recombinations per unit time per hydrogen atom [ $s^{-1}$ ] is given by

$$R_i = n_e \alpha_i(T), \quad (5.9)$$

where  $\alpha_i(T)$  is the recombination coefficient of species  $i$ .

In many astrophysical applications, the recombination transition to the ground-state is excluded from the recombination coefficient because this transition produces a photon capable of ionizing hydrogen. This is called ‘case B’ in contrast to ‘case A’ where all recombinations are included. The assumption underlying use of the ‘case B’ recombination coefficient is that the radiation associated with this transition is absorbed nearby again resulting in a new ionization. This is called the ‘on-the-spot’ (OTS) approximation. In other words, the environment is assumed to be optically thick with respect to Ly-continuum radiation. The OTS approximation is popular because it circumvents the explicit treatment of recombination photons which is prohibitively difficult for most radiative transfer methods as every resolution element can become a source. In `SIMPLEX`, recombination photons can be included without additional computational

effort (because of the locality of cell-to-cell transport) and we have implemented both case A and case B recombinations rates.

The validity of the OTS approximation has been investigated analytically in [Ritzerveld \(2005\)](#), who concludes that recombination radiation can dominate source photons close to the ionization front, especially for sources with a spectrum that peaks around the Lyman limit. This statement is verified numerically in [Williams & Henney \(2009\)](#); [Cantalupo & Porciani \(2011\)](#) and [Paardekooper \(2010\)](#). In the latter study the `SIMPLEX` method, including the effects of heating on the gas, was used. They found only marginal influence of the OTS approximation on large-scale reionization simulations but concluded that effects on smaller scales, relevant for star-formation, can be important.

### 5.2.3 Evolution of the ionization-state

Together, ionizations and recombinations determine the ionization-state of the gas described by the following three coupled differential equations and three closure relations

$$\dot{n}_{\text{HI}} = n_{\text{HII}}R_{\text{HI}} - n_{\text{HI}}\Gamma_{\text{HI}} \quad (5.10)$$

$$\dot{n}_{\text{HeI}} = n_{\text{HeII}}R_{\text{HeI}} - n_{\text{HeI}}\Gamma_{\text{HeI}} \quad (5.11)$$

$$\dot{n}_{\text{HeIII}} = -n_{\text{HeIII}}R_{\text{HeII}} + n_{\text{HeII}}\Gamma_{\text{HeII}} \quad (5.12)$$

$$n_{\text{H}} = n_{\text{HI}} + n_{\text{HII}} \quad (5.13)$$

$$n_{\text{He}} = n_{\text{HeI}} + n_{\text{HeII}} + n_{\text{HeIII}} \quad (5.14)$$

$$n_e = n_{\text{HII}} + n_{\text{HeII}} + 2n_{\text{HeIII}}. \quad (5.15)$$

This set of equation does not have a general analytical solution and must thus be solved numerically. For this purpose we adopt a sub-cycling scheme described in [Pawlik & Schaye \(2008\)](#). In this scheme, ionizations and recombinations are evolved on a time-scale that is smaller than the ionization- or recombination times-scales  $t_{\text{ion}}$  and  $t_{\text{rec}}$ . During a radiative transfer time step, the ionizing flux is assumed to be constant making the procedure manifestly photon-conserving. This allows for radiative timesteps,  $\Delta t_{\text{rt}}$  that are much large than the dominant timescale governing the evolution of the ionization-state. The sub-cycling time step is harmonic mean of the ionization and recombination rates

$$\Delta t_{\text{sub}} \equiv \frac{t_{\text{ion}} + t_{\text{rec}}}{t_{\text{ion}}t_{\text{rec}}}. \quad (5.16)$$

Because the procedure is analogous for each species, we here give the explicit example for the integration step for hydrogen only. A time  $t_{\text{sub}} \in (t_{\text{rt}}, t_{\text{rt}} + \Delta t_{\text{rt}})$  the rate equation is given by

$$dn_{\text{HII}}^{(t_{\text{sub}})} = n_{\text{HI}}^{(t_{\text{sub}})}\Gamma_{\text{H}}^{(t_{\text{sub}})}\Delta t_{\text{sub}} - n_e^{(t_{\text{sub}})}n_{\text{HII}}^{(t_{\text{sub}})}\alpha_{\text{H}}(T)\Delta t_{\text{sub}}, \quad (5.17)$$

where the photoionization rate at  $t_{\text{sub}}$  is given by

$$\Gamma_{\text{H}}^{(t_{\text{sub}})} = \Gamma_{\text{H}} \left( \frac{1 - e^{-\tau^{(t_{\text{sub}})}}}{1 - e^{-\tau}} \right) \frac{n_{\text{HI}}}{n_{\text{HI}}^{(t_{\text{sub}})}}, \quad (5.18)$$

where  $\Gamma_{\text{H}}$  and  $\tau$  are the photoionization rate and optical depth at the beginning of the subcycling and  $\tau^{(t_{\text{sub}})} = \tau n_{\text{HI}}^{(t_{\text{sub}})}/n_{\text{HI}}$ . By defining the photionization rate in this way, the ionizing flux in

the cell is constant during the radiative transfer time step. This sub-cycling scheme becomes computationally expensive when  $\Delta t_{\text{sub}} \ll \Delta t_{\text{rt}}$ , but photoionization 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 Eq. (5.17) and (5.18) but the procedure can be trivially extended to multiple frequencies by treating every bin separately.

### 5.2.4 Time stepping

The current implementation of SIMPLEX solves the time-independent radiative transfer equation (Eq. 1.3). This implies that the speed of light is infinite. In practice this means that radiation must be able to travel through the simulation volume on a timescale that is much shorter than the fastest physical timescale that needs to be captured. For a simulation to correctly trace the propagation speed of ionization fronts, the ‘speed of light’ of that simulation must be much faster than that of those fronts (but could be much lower than the actual speed of light in vacuum,  $c$ , Abel et al. 1999). We must convince ourselves that the radiative transfer time step,  $\Delta t_{\text{rt}}$ , is sufficiently small to satisfy the time-independent transfer equation. For the physical tests presented in this thesis, we have checked this requirement explicitly.

The sub-cycling scheme described above allows for time steps that are much larger 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.

## 5.3 THERMAL STATE OF THE GAS

Energy exchange between the radiation field and the gas can occur through various interaction processes. We have implemented heating and cooling processes relevant for a primordial composition of gas. First collected by Black (1981), these are photo-heating, recombination cooling, dielectric recombination cooling, excitation cooling and collisional ionization cooling. We also include a prescription of cooling due to non-relativistic free-free interactions.

For a gas parcel of fixed volume, the change in internal energy per unit mass is

$$\frac{du}{dt} = \frac{n_{\text{H}}^2}{\rho} (\mathcal{H} + C), \quad (5.19)$$

where  $\mathcal{H}$  and  $C$  are the (normalised) radiative heating- and cooling function respectively. 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. Assuming an ideal gas, the internal energy and the temperature are related by

$$u = \frac{\frac{3}{2} k_{\text{B}} T}{\mu m_{\text{H}}}, \quad (5.20)$$

where  $k_B$  is Boltzmann's constant,  $m_H$  is the mass of the hydrogen atom and  $\mu \equiv \bar{m}/m_H$  is the mean molecular weight, with  $\bar{m}$  the average mass of a gas particle. The average mass per particle depends on the ionization-state of the gas through

$$\bar{m} \simeq \frac{m_H(n_{\text{HI}} + n_{\text{HII}} + 4n_{\text{HeI}} + 4n_{\text{HeII}} + 4n_{\text{HeIII}})}{n_{\text{HI}} + 2n_{\text{HII}} + n_{\text{HeI}} + 2n_{\text{HeII}} + 3n_{\text{HeIII}}}, \quad (5.21)$$

where we have used that  $m_{\text{He}} \simeq 4m_H$ .

### Heating processes

In this work we limit ourselves to photo-heating so this is the only contribution to  $\mathcal{H}$

$$\mathcal{H} = \sum_i \frac{n_i}{n_{\text{H}}^2} \epsilon_{\gamma,i} \quad \text{with } i = \text{H I}, \text{He I}, \text{He II}. \quad (5.22)$$

Here  $\epsilon_{\gamma,i}$  is the photoheating coefficient of species  $i$ , given by

$$\epsilon_{\gamma,i} = \int_{\nu_{0,i}}^{\infty} \frac{4\pi J_{\nu} \sigma_i(\nu)(h\nu - h\nu_{0,i})}{h\nu} d\nu = \Gamma_{P,i} \langle \mathcal{E}_i \rangle, \quad (5.23)$$

where we have used Eq. (5.1) to define the average excess energy of the ionizing photons

$$\langle \mathcal{E}_i \rangle \equiv \left( \int_{\nu_{0,i}}^{\infty} \frac{4\pi J_{\nu} \sigma_i(\nu)(h\nu - h\nu_{0,i})}{h\nu} d\nu \right) \left( \int_{\nu_{0,i}}^{\infty} \frac{4\pi J_{\nu}}{h\nu} \sigma_i(\nu) d\nu \right)^{-1}, \quad (5.24)$$

and  $\nu_{0,i}$  is the ionization energy of the atom. 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 ionization cooling, collisional excitation cooling and cooling from free-free emission:

$$C = \sum_i c_i, \quad (5.25)$$

where  $i$  runs over all contributions due to different processes and species. The different contributions and their origin are given in Tab. 5.3 in Sect. 5.A.

### Implementation


The evolution of the internal energy, Eq. (5.19), is explicitly integrated in step with the evolution of the ionization state of the medium. This yields accurate results in almost all cases because the heating and cooling time-scales are generally much larger than the ionization and recombination time-scales. In the rare case that this is not true, we solve Eq. (5.19) on a time-scale that guarantees a stable solution.

If the sub-cycling scheme described in Sect. 5.2.3 is used, the following situation may occur. The sub-cycling converges to ionization equilibrium before the end of the radiative transfer time step. In this case, we evolve the internal energy (on the thermal time scale) until the end of the time step, *assuming* a constant ionization state of the medium. Once the internal energy of the medium at that final time is known, we then set the ionization state of the medium to the equilibrium values appropriate for that internal energy. Because the radiative transfer time steps are not much larger than the thermal time scale in almost all cases, we only need to take several explicit steps to correctly solve for the internal energy.

### Comparison to CLOUDY

To test our implementation of helium cross sections and recombination rates, we compare the ionization fractions of hydrogen and helium with those obtained with the well-established code CLOUDY (Ferland et al. 1998). Because CLOUDY incorporates many additional processes, we can check whether our implementation captures the most important contributions.

The test set-up is as follows. A black-body source with  $5 \cdot 10^{54}$  photons capable of ionizing hydrogen per second placed within a homogeneous neutral medium. The medium consists of hydrogen with number density  $n_{\text{HI}} = 10^{-3} \text{ cm}^{-3}$  and helium with number density  $n_{\text{HeI}}/n_{\text{HI}} = 0.1$ . The CLOUDY results represent the equilibrium solution whereas we have ended the simulation at 2 Gyr when equilibrium is not fully established yet. We allow the temperature to evolve due to the cooling processes given in the Appendix and we include photo-heating. In Fig. 5.2 we show the results of this test.



**Figure 5.2:** Ionised and neutral fractions for our one-dimensional RT code (black curves) compared to CLOUDY (grey curves). Our results are calculated at 2 Gyr whereas CLOUDY results represent the equilibrium solution.

The positions of the ionization fronts of both helium and hydrogen are in good agreement between the two codes.

In the inner regions, deviations are all under several percent with the largest differences seen in the He I fractions. Outside of the ionization-front, the CLOUDY curves fall off less steep when compared to our results. This is because our result is not completely converged to the equilibrium value for this time.

## 5.4 PROPER TREATMENT OF THE SPECTRUM

As we discuss below in Sect. 5.4.1, a one-frequency-bin or ‘grey’ approach to the radiative transport of ionizing photons is a useful approximation in many applications. For problems where detailed heating or spectral hardening are important, the form and evolution of the radiation spectrum needs to be followed and we must use a multiple frequency description. We have a certain freedom in the discretization of the spectrum that we might turn to our advantage. In Sect. 5.4.3 we discuss possible choices and their merit.

### 5.4.1 Two grey approximations

In numerical simulations involving radiation it is often necessary to approximate the continuous spectrum of radiation with a finite number of discrete frequency bins due to memory requirements. The extreme (but often employed) limit of one single frequency bin is commonly referred to as the ‘grey approximation’. Although in the grey approximation all spectral information is lost, it is still possible to enforce the conservation of a quantity of importance such as the number of ionizations per unit time or the energy deposition into the medium per unit time.

The conservation of energy is accomplished by defining the effective (grey) cross section for species  $i$ ,  $\sigma_{E,i}$ , as

$$\sigma_{E,i} = 4\pi \int_0^\infty \sigma_i(\nu) J_\nu d\nu / J, \quad (5.26)$$

where the energy rate per surface area (which is the frequency integrated mean intensity),  $J$ , is defined by

$$J \equiv 4\pi \int_0^\infty J_\nu d\nu. \quad (5.27)$$

If we want to know the amount of energy injected into the medium per unit time due to ionizations of species  $i$ ,  $\dot{E}_i$ , we simply multiply the grey cross section with  $J$

$$\dot{E}_i = \sigma_{E,i} J. \quad (5.28)$$

Alternatively, the source function can be divided by the energy per photon,  $h\nu$ , to keep the number of ionizations per unit time constant

$$\sigma_{I,i} = 4\pi \int_0^\infty \frac{\sigma_i(\nu) J_\nu}{h\nu} d\nu / \dot{N} \quad (5.29)$$

where  $\dot{N}$  is the rate of ionizing photons per surface area defined by

$$\dot{N} \equiv 4\pi \int_0^\infty \frac{J_\nu}{h\nu} d\nu. \quad (5.30)$$

The phototization rate is thus given by

$$\Gamma_{P,i} = \sigma_{I,i} \dot{N}. \quad (5.31)$$

Both the energy- and the ionization-conserving definitions have their merit. Intuitively, the conservation of ionizations leads to a better solution if the problem is sensitive to the ionization



state of the medium. On the other hand, if the temperature of the medium is most important, the energy conserving choice would be more appropriate because in the vicinity of ionizing sources, photo-heating is the dominant source of thermal energy.

The correctness of the grey approximation depends on the spectrum and the cross sections that are used. It works best for spectra and cross sections that peak sharply around the same energy, and are thus well approximated with a single frequency bin covering that energy. For spectra that fall-off slowly (power-laws with small powers for instance) and cross sections with contributions separated by large ranges in frequency, the need for a multiple frequency description becomes more pressing. We furthermore note that the grey approximation does not allow for different types of sources to be used in the same simulation because the effective cross section is photon conserving for a single type of spectrum only.

### 5.4.2 Spectral hardening

In the vicinity of ionizing sources, photo-heating is the dominant source of thermal energy. For every ionization of a given atom, the liberated electron carries an energy of

$$E_{\text{thermal}} = h(\nu - \nu_{0,i}) \quad (5.32)$$

where  $\nu_{0,i}$  is the ionization energy of the atom. Reproducing Eq. (5.23) but now with frequency limits that represent multiple bins, the effective thermal energy added to the medium due to ionization of species  $i$ , for a photon in bin  $j$ , is

$$\epsilon_{\gamma,i,j} = \int_{\nu_j}^{\nu_{j+1}} \frac{4\pi J_\nu \sigma_i(\nu)(h\nu - h\nu_{0,i})}{h\nu} d\nu. \quad (5.33)$$

Replacing Eq. (5.23) with Eq. (5.33) and the related discretisation of the cross section, will have several related physical effects on the transport of ionizing radiation. Because the cross section is a strongly frequency-dependent function, the absorption of photons in different frequency bins will happen at different depths in the medium. Photons with energies much higher than  $\nu_{0,i}$  will be able to travel further than the ionization fronts, depositing their (significant) surplus of energy to the thermal bath of the medium. This effect is known as *spectral hardening* because the spectrum shifts to more energetic ‘harder’ photons when progressing through the medium. Spectral hardening results in significant heating beyond the ionization front and, consequently, a lower recombination rate there. The effect of this suppression of recombinations in turn is a more extended ionization front.

### 5.4.3 Choice of bins

In the previous section we have seen that the multiple frequency scheme in SIMPLEX compares excellently to other methods. The solutions for the ionization and temperature structure show significant change when we go to multiple frequency bins. But how many frequency bins should one use to obtain this solution (given some tolerance)? The answer is that the number of bins depends on the choice of intervals in frequency space, the cross section and the type of spectrum.

In general, we would like the frequency bins to capture the behaviour of the *product* of the spectrum and the cross section in sufficient detail. Bins can be wide when this product is small and/or changes slowly and must be narrow otherwise. For the moment we concentrate primarily on the amplitude of the product of  $\sigma$  and  $J$ , striving to keep either the number of *ionization events* (e.g., Eq. (5.31)) or *energy transfer* per unit time (e.g., Eq. (5.28)) constant across bins. We can choose the width of the frequency bins such that the energy-absorption rate is constant across bins

$$\Gamma_{E,i} = \Gamma_{E,j} \quad \forall \quad i \neq j \quad \text{where} \quad \Gamma_{E,i} = 4\pi \int_{\nu_i}^{\nu_{i+1}} \sigma(\nu) J_\nu d\nu, \quad (5.34)$$

where we have omitted reference to species to avoid cluttering. Equivalently, we can force the ionization rate is constant across bins

$$\Gamma_{I,i} = \Gamma_{I,j} \quad \forall \quad i \neq j \quad \text{where} \quad \Gamma_{I,i} = 4\pi \int_{\nu_i}^{\nu_{i+1}} \frac{\sigma(\nu) J_\nu}{h\nu} d\nu. \quad (5.35)$$


Alternatively, we can choose logarithmically spaced bins. This would make sense in the case that the product of cross section and spectrum behave according to a power-law. The width of the bins tend to increase in the same way as the strength of the spectrum diminishes. If the product is roughly constant, however, we could use bins of constant width. Finally, we take bins that just keep the radiated energy per bin constant, we call this possibility ‘spectrum-weighted’.

$$J_i = J_j \quad \forall \quad i \neq j \quad \text{where} \quad J_i = 4\pi \int_{\nu_i}^{\nu_{i+1}} J_\nu d\nu. \quad (5.36)$$


We show these five options for bin spacing in Fig. 5.3. The first three panels (from the top) are the same for hydrogen only (left figure) and hydrogen plus helium (right figure) because the spacing is independent of the spectrum. The bottom two panels of each figure show the bin-spacings given by Eq. (5.35) and Eq. (5.34) respectively. Because the extra factor of  $\nu$  in the denominator of Eq. (5.35), the ionization-weighted bins are more focussed on the peaks of the cross sections than the energy-weighted bins.

In Fig. 5.4 the distribution of frequency bins for a power-law  $\nu^{-1}$  (typical for quasar spectra e.g., Elvis et al. 1994) is shown. As mentioned previously, the spectrum-weighted bins are equivalent to the logarithmic case for a power-law spectrum. The ionization-weighted bins are even more focussed on low frequencies than in the blackbody case. The energy-weighted bin distribution is very similar to the blackbody case for pure hydrogen but is more extended to higher frequencies for simulations including helium as well. Because the maximum of the spectrum lies at the origin, the first part of the cross section (due to hydrogen) is sampled with three instead of two bins.

We note that regardless of our choice of bins, we are still at liberty to choose either the ionization or the energy conserving weighted cross sections. The only difference being that now, the weighting will be done over the width of one bin instead of the whole frequency interval. We use ionization-weighted cross sections for the tests described below but the results are not influenced significantly by this choice.



**Figure 5.3:** Spacing of frequency bins for the calculation of effective cross sections using  $J(\nu) = B_\nu(T = 50\text{K})$  for five different bin-size strategies. In contrast to the linear and logarithmic spacing, the other three strategies use information of the spectrum and the cross-sections to find an optimal spacing of the bins. In each case we use 10 bins ranging from 1 to  $6 \nu_{0,\text{HI}}$  for aesthetic reasons (we use an upper limit of  $15 \nu_{0,\text{HI}}$  in our runs). The width of the grey-shaded bars indicate the range of frequency of the corresponding bin. A scaled  $10^5$  K blackbody spectrum and the cross section are plotted for reference. *Left:* Pure hydrogen. *Right:* Hydrogen and helium.



**Figure 5.4:** Same as in Fig. 5.3 but now for  $J(\nu) = \nu^{-1}$ . *Left:* Pure hydrogen. *Right:* Hydrogen and helium.


#### 5.4.4 Convergence of temperature and ionization structure

We investigate how the convergence of two important observables, the ionization structure and the temperature in a hydrogen (and helium) gas, converge as a function of the number of frequency bins for the following five choices of bin-spacing:

1. Linearly spaced
2. Logarithmically spaced
3. Spectrum-weighted (the radiated energy per unit time is constant along bins)
4. Energy-weighted (the rate of energy transferred to the medium is constant)
5. Ionisation-weighted (the ionization rate is constant along bins)

We use the set-up of ‘test 2’ from the Cosmological Radiative Transfer Comparison Project published in [Iliev et al. \(2006a\)](#) and described in the previous section. Because the effects of different prescriptions are more pronounced for an ionized region that has not yet reached its equilibrium size (also because the box is too small to capture the temperature profile outside of the ionization front), we perform our tests after 30 Myr when the ionization front lies at roughly half the Strömngren radius. We perform this test both with the ‘default’  $10^5$  K blackbody curve used in the original test and a power-law of  $\nu^{-1}$ .

Figure 5.5 shows the temperature profiles for these two spectra without (left panel) and with (right panel) inclusion of helium. These show the numerical results with converged number of frequencies and time step. In the case of a pure hydrogen gas, the power-law spectrum results in



**Figure 5.5:** Temperature profile at 30 Myr for the two types of spectra used in the convergence tests in this chapter. *Left:* Pure hydrogen. *Right:* Hydrogen and helium.

a more extended temperature profile. This is expected because the spectrum falls off less steeply than the blackbody spectrum (which has the well-known  $\nu^{-3}$ -behaviour at larger wavelengths) and thus carries more energy in harder photons. This hard radiation, in turn, pierces further

beyond the ionization front and heats the gas further out. In the case helium is also taken into account, the temperature profile for the power-law spectrum shows an additional bump closer to the source.

We have performed simulations using a one-dimensional implementation of our chemistry solver to speed up the testing process and to eliminate possible noise due to radiative transfer effects. We assume spherical symmetry and use 100 equally spaced cells spanning 6.6 kpc. Radiation enters the first cell and is traced through the other cells until it is absorbed or leaves the domain. We do not include recombination radiation explicitly but adopt the on-the-spot approximation. We have checked convergence of our results for number of cells and time-step. To check for convergence, we perform a (fully converged) reference simulation (typically with 128 frequency bins) to compare against. To compare two results we can define an error measure that gives a total relative error

$$\sigma \equiv \frac{\sum_i |x_i - x_{\text{ref},i}|}{\sum_i 0.5(x_i + x_{\text{ref},i})}. \quad (5.37)$$

We will use Eq. (5.37) to check for convergence, allowing for maximally  $\sigma = 0.01$  for the ionization structure and  $\sigma = 0.05$  for the temperature. We allow this larger error in the temperature because it is more susceptible to numerical noise (especially close to the source) and we are more interested in the convergence behaviour than the absolute error at this point.

Also, there is freedom in choosing an upper limit for the frequency bins. The spectrum-, energy- and ionization-weighted results are hardly affected by this choice, as long as the upper limit is sufficiently high. For linear and logarithmic bins, the choice of maximal frequency is more critical. Many bins are needed when this upper limit is too high ( $100 \nu_{0,\text{HI}}$ , for instance, is prohibitively high because it requires over 100 linear bins). In the tests below we have used an upper limit of  $15 \nu_{0,\text{HI}}$ . At that frequency, the contribution of the product of the spectra and the cross section(s) has decreased to less than a millionth of its maximal value.

Binning	Temperature				Ionisation			
	blackbody		power-law		blackbody		power-law	
	H	H + He	H	H + He	H	H + He	H	H + He
Linear	16	14	5	4	10	20	17	17
Logarithmic	8	8	2	2	5	6	5	4
Spectrum weighted	4	4	2	2	3	4	5	4
Energy weighted	8	4	14	6	12	6	14	4
Ionisation weighted	44	8	35	14	15	4	35	5

**Table 5.1:** Number of frequency bins needed for convergence within 1% (for ionization fraction) and 5% (for temperature) depending on choice for bin-spacing (column 1). Results are for a blackbody spectrum (column 2,3,6 and 7) and a power-law spectrum (column 4,5,8 and 9).

In Tab. 5.1 we summarize our results. We see that in general, the ionization structure requires about the same number of bins to converge within the given tolerance as the temperature.

In many cases, more ionization-weighted bins are needed for convergence than logarithmic and energy-weighted. The number of bins is comparable to linearly spaced bins. For the ionization fractions this is somewhat surprising because they are chosen to keep the ionization-rate per bin constant. For both the blackbody and the power-law spectrum, ionization-weighting places (too) much emphasis on the frequencies close to  $\nu_{\text{HI},0}$  and makes it harder for the temperature structure to converge for pure hydrogen. This problem is alleviated when helium is added because there are contributions of the cross section at higher energies, which implies a better frequency resolution there.


Energy weighted bins are more efficient in all cases but do not out-perform logarithmic and spectrum-weighted bins. They also do not seem to be more efficient than the spectrum-weighted and logarithmic bins when it comes to temperature convergence, which is, again, surprising because they are designed to optimally resolve the ionization rate.

From Fig. 5.3 we see that logarithmic bins closely resemble the spectrum-weighted result in all cases (they are identical for the power-law spectrum of course). Therefore they are almost as efficient from the point of convergence as can be seen from Tab. 5.1. We caution against the naive use of logarithmic bins for general (i.e., non-power-law) spectra, however, because they tend to severely under-sample the low frequency part of the spectrum when the upper bounds are chosen large. Only linearly spaced bins perform worse in this respect.


Considering these results, we advocate the use of spectrum-weighted binning which respects the form of the spectrum. This type of binning does not suffer from the limitations of logarithmic and linear binning and is easily calculated. More involved methods that take the cross section into account do not increase the convergence of either the ionization state or the temperature distribution and are as such not recommended.

### 5.4.5 Testing

We have performed test 2 from the first Cosmological Radiative Transfer Comparison Project published in Iliev et al. (2006a). This test traces the ionization fraction and temperature for a single  $10^5$  K blackbody source in a homogeneous  $n_{\text{HI}} = 10^{-3} \text{ cm}^{-3}$  initially neutral hydrogen gas. The source emits  $10^{48}$  ionizing photons per second and is located in the centre of a 13.2 kpc box. In Fig. 5.6, we show the difference between the grey (1 frequency, dash-dotted line) and multiple frequency results obtained with SIMPLEX and using energy weighted bins (see Sect. 5.4.3). The results obtained with SIMPLEX show convergence with 5 frequency bins for both the 30 and 500 Myr result. All results (even the monochromatic ones) lie within the grey area that shows the range of results obtained by other radiative transfer codes that participated in the comparison project. This is due to the fact that some of these codes rely on approximate methods to treat spectral hardening, some of which deviate significantly from the true multi-frequency solution. For the temperature-related results, presented in Fig. 5.7, we see that only the completely monochromatic result lies outside the range of other results. The convergence is somewhat slower in the 30 Myr case than for the ionization structure. In contrast, the 500 Myr results seem converged for all but the monochromatic case. This is also a result of the problem setup, the box is too small to show the temperature structure beyond the ionization front (where the deviations are most notable) at 500 Myr.



**Figure 5.6:** Test 2. The effect of spectral hardening on the position of the ionization-front. Shown are the ionized and neutral fractions as a function of radial distance from the source. Results are shown at 30 Myr (*Left*) and 500 Myr (*Right*). The results of other radiative transfer codes that participated in the comparison project published in [Iliev et al. \(2006a\)](#) are shown as the shaded grey area. This image has been published earlier in [Paardekooper \(2010\)](#).



**Figure 5.7:** Test 2. The effect of spectral hardening on the temperature of the gas (otherwise identical to Fig. 5.6). This image has been published earlier in [Paardekooper \(2010\)](#).

## 5.5 IMPLEMENTATION

We now describe how the above ideas are put into practice in the chemistry module of the `SIMPLEX` code. As mentioned before, the natural computational unit for the `SIMPLEX` code is the *sweep*. The sweep represents a physical time-step  $\Delta t$  in the sense that the number of photons sent from a source with luminosity  $L$  in one sweep is given by  $L \times \Delta t$ . Furthermore, every sweep, the equations that describe the change of the chemical state of the gas are evolved over a period  $\Delta t$ . In one sweep, all nuclei of the grid are visited *once* in random order and the state of each nucleus is changed as follows.

First, all radiation that has arrived at the nucleus during the *previous* sweep is used to change the chemical state of the gas in the cell. We must be careful to exclude radiation that arrived in the current sweep because otherwise photons can take more than one step on the grid<sup>1</sup>.

To find the number of ionisations of species  $i$ , we evaluate Eq. (5.8) and multiply  $\dot{N}_{\text{ion},i}$  by  $\Delta t$ . The number of photons *for a given frequency bin* absorbed for these ionisations can be found by multiplying Eq. (5.7) with  $\Delta t$ .

The number of photons that needs to travel onward to neighbouring cells can be found by taking the original number of photons and subtracting the absorptions. If the number of photons in frequency-bin  $j$  before chemistry is given by  $N_{\gamma,j}$ , the number of outgoing photons after chemical interactions is given by

$$N_{\gamma,j}^{\text{out}} = N_{\gamma,j} - \dot{N}_{\text{ion},j} \Delta t. \quad (5.38)$$

This radiation is sent from each nucleus to its neighbours using diffuse-, ballistic- or direction conserving transport. We note that the chemical state of the gas in a cell is changed right before that cell sends its content to its neighbours. This allows us to use the updated optical depth of the cell for the radiative transport.

Because the time step  $\Delta t$  is often much larger than the typical chemical timescale (taken as the minimum of the ionisation-, recombination- and heating/cooling timescale) we apply a sub-cycling procedure as described in Sect. 5.2.3. The time step used for this sub-cycling is defined by Eq. (5.16), a harmonic mean of the ionisation- and recombination timescale. As mentioned in Sect. 5.3, the internal energy of the gas (Eq. (5.19)) is evolved on the sub-cycling timescale unless the thermal timescale is smaller than  $\Delta t_{\text{sub}}$  (which happens sporadically). In these rare cases, Eq. (5.19) is evolved on a fraction of the thermal timescale which is the minimum of the timescales of all heating and cooling processes.

The procedure described above requires several quantities and numbers to be known beforehand. For example, we must choose a number of frequency bins to represent the spectrum<sup>2</sup> of the sources. Once the source spectrum and bin spacing method have been chosen, the actual spacing is calculated using linear- or logarithmic spacing or frequency ranges that satisfy Eq. (5.36), Eq. (5.35) or Eq. (5.34). In the latter cases, we use numerical integration combined with a minimization routine to find the correct boundaries for the frequency bins.

<sup>1</sup>This is a consequence of the random treatment of the nuclei. If we were to relax this one-step-per-sweep requirement, the number of communications between processors in a parallel simulation would have to be increased from one to ‘every time a photon crosses a boundary between processors’ which can be a large number.

<sup>2</sup>In the current implementation, this spectrum can be either power-law or a black-body source. It is not difficult to use more general source-spectra, but in that case, the choice of frequency bin spacing may not be as simple as described above.



Once the bins are known, we can also calculate the average excess energy of an ionisation of an atom of species  $i$  from a photon of frequency  $j$  by generalizing Eq. (5.24) to more frequencies

$$\langle \mathcal{E}_{i,j} \rangle \equiv \left( \int_{\nu_j}^{\nu_{j+1}} \frac{4\pi J_\nu \sigma_i(\nu)(h\nu - h\nu_{0,i})}{h\nu} d\nu \right) \left( \int_{\nu_j}^{\nu_{j+1}} \frac{4\pi J_\nu}{h\nu} \sigma_i(\nu) d\nu \right)^{-1}. \quad (5.39)$$

We note that the spacing of the frequency bins, the excess energy per ionisation and the effective cross sections need to be computed only once, before the actual simulation. This initialization of the code takes several seconds in total. We have tested the convergence and stability of the frequency binning procedure up to 200 frequency bins.

## 5.6 SUMMARY

- We have presented the chemical module implemented in the SIMPLEX code including photo-heating and various cooling mechanisms for hydrogen and helium.
- The results of SIMPLEX compare well with those obtained by other radiative transfer methods.
- To account for spectral hardening and spectral evolution of the radiation field we have implemented a flexible multiple frequency binning. We have compared five binning methods on the subject of convergence to a reference solution. Our tests with two different source types (blackbody and power-law) lead us to conclude that a spectrum-weighted approach is favourable in general.

## 5.A RATES AND CROSS SECTIONS

Table. 5.2 summarizes the rates and cross sections used in this work. The various choices are based on comparison with other work focussing on range of validity and estimated error but reflect to a certain extent the personal preference of the author. We have found the elaborate discussion in Pawlik (2009) useful and have adopted their preferred rates in most cases. The selection of heating and cooling processes is not exhaustive but includes the most important ones (as first compiled by Black (1981)).

Photo ionization cross section	HI, HeI, HeII	Verner et al. (1996)
Dust cross sections		Gnedin et al. (2008)
Collisional ionization rate	HI, HeI, HeII	Theuns et al. (1998)
Spontaneous recombination	HII, HeIII	Hui & Gnedin (1997)
	HeII	Hummer & Storey (1998)
Dielectric recombination	HeII	Aldrovandi & Pequignot (1973)
Recombination cooling	HII, HeIII	Hui & Gnedin (1997)
	HeII	Hummer & Storey (1998)
Collisional ionization cooling	HI, HeI, HeII	Theuns et al. (1998)
Dielectric recombination cooling	HeII	Black (1981)
Line (collisional excitation) cooling	HI, HeII	Theuns et al. (1998)
Free-free cooling		Theuns et al. (1998)

**Table 5.2:** Summary of rates and cross sections for various physical effects.

### 5.A.1 Cross sections

#### Hydrogen and helium

We use the cross sections for hydrogen and helium from Verner et al. (1996). The photoionization cross sections for these species are well approximated by a fitting function

$$\sigma = \sigma_0 \left( \left( \frac{E}{E_0} - y_0 - 1 \right)^2 + y_w^2 \right) \frac{\left( \sqrt{\left( \frac{E}{E_0} - y_0 \right)^2 + y_1^2} \right)^{0.5p-5.5}}{\left( 1 + \sqrt{\frac{\sqrt{\left( \frac{E}{E_0} - y_0 \right)^2 + y_1^2}}{y_a}} \right)^p} \text{ cm}^2, \quad (5.40)$$

where the coefficients are given by

Species	$E_0$	$\sigma_0 [10^{-14}]$	$p$	$y_a$	$y_w$	$y_0$	$y_1$
H I	0.4298	5.475	2.963	32.88	0	0	0
He I	13.61	0.09492	3.188	1.469	2.039	0.4434	2.136
He II	1.72	1.369	2.963	32.88	0	0	0

#### Dust

The optical depth due to dust is given by

$$\tau_d = n_H \sigma_d l, \quad (5.41)$$

where  $l$  is the path length,  $n_H$  is the hydrogen number density and  $\sigma_d$  is the effective dust cross section per hydrogen atom. To account for a metallicity different from that in the SMC and LMC we simply scale the hydrogen number density with the metallicity fraction,  $Z/Z_0$ . Here  $Z_0$  is the reference metallicity of the SMC and LMC relative to solar metallicity. We use  $Z_{0,\text{SMC}} = 0.25$  and  $Z_{0,\text{LMC}} = 0.5$  (Welty et al. 1997, 1999).

## 5.A.2 Recombination rates

We use the recombination rates from [Hui & Gnedin \(1997\)](#) who present fits to data from [Ferland et al. \(1992\)](#) (for H I and He III) and our fit (linear in log-log) to the data of [Hummer & Storey \(1998\)](#) for He II. The reason for this is that [Hui & Gnedin \(1997\)](#) have used only three data-points (at 0,  $10^4$  and  $2 \cdot 10^4$  K) from [Burgess & Seaton \(1960\)](#) for their fit of the He II recombination rate (probably because at the time they were the only data available) whereas [Hummer & Storey \(1998\)](#) provide values in  $^{10}\log(T)$  from 1 to  $10^{4.4}$  K in 0.2 steps. The case A and case B recombination rates that we use are thus

$$\alpha_{\text{HeII}}^{\text{A}}(T) = 1.269 \cdot 10^{-13} \frac{(\lambda_{\text{HI}})^{1.503}}{\left[1.0 + (\lambda_{\text{HI}}/0.522)^{0.47}\right]^{1.923}} \quad (5.42)$$

$$\alpha_{\text{HeII}}^{\text{A}}(T) = 10^{-9.79782-0.638125 \log(T)} \quad (5.43)$$

$$\alpha_{\text{HeIII}}^{\text{A}}(T) = 2.538 \cdot 10^{-13} \frac{(\lambda_{\text{HeII}})^{1.503}}{\left[1.0 + (\lambda_{\text{HeII}}/0.522)^{0.47}\right]^{1.923}} \quad (5.44)$$

$$\alpha_{\text{HeII}}^{\text{B}}(T) = 2.753 \cdot 10^{-14} \frac{(\lambda_{\text{HI}})^{1.5}}{\left[1.0 + (\lambda_{\text{HI}}/2.74)^{0.407}\right]^{2.242}} \quad (5.45)$$

$$\alpha_{\text{HeII}}^{\text{B}}(T) = 10^{-9.79239-0.687189 \log(T)} \quad (5.46)$$

$$\alpha_{\text{HeIII}}^{\text{B}}(T) = 2.0 \cdot 2.753 \cdot 10^{-14} \frac{(\lambda_{\text{HeII}})^{1.5}}{\left[1.0 + (\lambda_{\text{HeII}}/2.74)^{0.407}\right]^{2.242}}, \quad (5.47)$$

where we follow the notation of [Hui & Gnedin \(1997\)](#) in using  $T_5 \equiv T/10^5$ ,  $\lambda_{\text{HI}} = 2T_{\text{HI}}/T$ ,  $\lambda_{\text{HeI}} = 2T_{\text{HeI}}/T$ ,  $\lambda_{\text{HeII}} = 2T_{\text{HeII}}/T$  and ionization threshold energies  $T_{\text{HI}} = 157807$  K,  $T_{\text{HeI}} = 285335$  K,  $T_{\text{HeII}} = 631515$  K.

## 5.A.3 Dielectric recombination

For the dielectric recombination rate (only for He II), we use the expression given by [Aldrovandi & Pequignot \(1973\)](#) who re-express results of [Brown \(1971\)](#).

## 5.A.4 Collisional ionization rates

Collisional ionizations have been taken into account using the coefficients (in units  $\text{cm}^3\text{s}^{-1}$ ) by [Theuns et al. \(1998\)](#) who revisited those published in [Cen \(1992\)](#).

$$\Gamma_{e\text{HI}} = 1.17 \cdot 10^{-10} \sqrt{T} \exp(-1.578091 \cdot 10^5/T) \left(1 + \sqrt{T/10^5}\right)^{-1} \quad (5.48)$$

$$\Gamma_{e\text{HeI}} = 4.76 \cdot 10^{-11} \sqrt{T} \exp(-2.853354 \cdot 10^5/T) \left(1 + \sqrt{T/10^5}\right)^{-1} \quad (5.49)$$

$$\Gamma_{e\text{HeII}} = 1.14 \cdot 10^{-10} \sqrt{T} \exp(-6.31515 \cdot 10^5/T) \left(1 + \sqrt{T/10^5}\right)^{-1} \quad (5.50)$$

### 5.A.5 Cooling rates


Table 5.3 summarizes the cooling rates used in this work.

$c_i$	Cooling rate [ $\text{erg cm}^3 \text{ s}^{-1}$ ]	Multiplication factor	Species
Collisional ionization cooling			
$c_1 =$	$2.54 \cdot 10^{-21} \sqrt{T} e^{-157809.1/T} (1 + \sqrt{T_5})^{-1}$	$n_e n_{\text{HI}}$	H I
$c_2 =$	$1.88 \cdot 10^{-21} \sqrt{T} e^{-285335.4/T} (1 + \sqrt{T_5})^{-1}$	$n_e n_{\text{He I}}$	He I
$c_3 =$	$9.90 \cdot 10^{-22} \sqrt{T} e^{-631515/T} (1 + \sqrt{T_5})^{-1}$	$n_e n_{\text{He II}}$	He II
Collisional excitation cooling			
$c_4 =$	$7.5 \cdot 10^{-19} e^{-118348/T} (1 + \sqrt{T_5})^{-1}$	$n_e n_{\text{HI}}$	H I
$c_5 =$	$5.54 \cdot 10^{-17} T^{-0.397} e^{-473638/T} (1 + \sqrt{T_5})^{-1}$	$n_e n_{\text{He II}}$	He II
Free-free cooling			
$c_6 =$	$1.42 \cdot 10^{-27} \sqrt{T} (1.1 + 0.34 e^{\{(-5.5 - \log_{10} T)^2 / 3\}})$	$n_e (n_{\text{H II}} + n_{\text{He II}} + 4n_{\text{He III}})$	H II, He II, He III
Dielectric recombination cooling			
$c_7 =$	$1.24 \cdot 10^{-13} T^{-1.5} e^{(-470000/T)} (1 + 0.3 e^{(-94000)})$	$n_e n_{\text{He II}}$	He II
Recombination cooling			
$c_8 =$	$1.778 \cdot 10^{-29} T \frac{\lambda_{\text{HI}}^{1.965}}{(1.0 + (\lambda_{\text{HI}}/0.541))^{0.502})^{2.697}}$	$n_e n_{\text{H II}}$	H II case A
$c_9 =$	$-25.6822 + 0.347135 \log(T)$	$n_e n_{\text{He II}}$	He II case A
$c_{10} =$	$8 \times 1.778 \cdot 10^{-29} T \frac{\lambda_{\text{He II}}^{1.965}}{(1.0 + (\lambda_{\text{He II}}/0.541))^{0.502})^{2.697}}$	$n_e n_{\text{He III}}$	He III case A
$c_8 =$	$3.435 \cdot 10^{-30} T \frac{\lambda_{\text{HI}}^{1.970}}{(1.0 + (\lambda_{\text{HI}}/2.250))^{0.376})^{3.720}}$	$n_e n_{\text{H II}}$	H II case B
$c_9 =$	$-25.9067 + 0.500511 \log(T) - 0.0414826 \log(T)^2$	$n_e n_{\text{He II}}$	He II case B
$c_{10} =$	$8 \times 3.435 \cdot 10^{-30} T \frac{\lambda_{\text{He II}}^{1.970}}{(1.0 + (\lambda_{\text{He II}}/2.250))^{0.376})^{3.720}}$	$n_e n_{\text{He III}}$	He III case B

**Table 5.3:** Implemented cooling processes. *First column:* identifier. *Second column:* coefficient as a function of temperature. *Third column:* multiplication factor that converts to units of [ $\text{s}^{-1} \text{cm}^{-3}$ ] or rate of energy loss per unit volume. *Fourth column:* species for the coefficient. We have used the notation from [Hui & Gnedin \(1997\)](#):  $T_5 \equiv T/10^5$ ,  $\lambda_{\text{HI}} = 2T_{\text{HI}}/T$ ,  $\lambda_{\text{He I}} = 2T_{\text{He I}}/T$ ,  $\lambda_{\text{He II}} = 2T_{\text{He II}}/T$  where  $T_{\text{HI}} = 157807$  K,  $T_{\text{He I}} = 285335$  K and  $T_{\text{He II}} = 631515$  K are the ionization energies of H I, He I and He II respectively.

#### Recombination cooling

Because the recombination cooling-rates depend on the recombination rates, the discussion about the He II recombination rate also applies here. For the reasons discussed earlier for the He II recombination rate, we choose to use a fit in log-log to the values of [Hummer & Storey \(1998\)](#) which is shown in Fig. 5.8. We use the dielectric recombination cooling formula from [Black \(1981\)](#), who refers back to the work by [Gould & Thakur \(1970\)](#) and [Burgess \(1964\)](#). Although the extrapolation above  $10^{4.4}$  K is of course a source of uncertainty, this does not pose a big problem because dielectric recombination cooling becomes dominant at temperatures  $T \gtrsim 5 \cdot 10^4$  K.



**Figure 5.8:** Fits to the He II recombination cooling data of [Hummer & Storey \(1998\)](#). We have used a linear fit in log-log for the case A and a quadratic fit in log-log for the case B recombination coefficient. Also plotted are the dielectric cooling terms from [Black \(1981\)](#) which dominate the cooling for  $T \gtrsim 5 \cdot 10^4$  K.

**Part II**

**Applications**



## CHAPTER 6

---

# A Markov chain description of SimpleX Radiation Transport

C. J. H. Kruip, R. Vooy's & V. Icke

WE formulate the basic transport methods for the SIMPLEX algorithm in the context of Markov chains. The Delaunay triangulation is interpreted as a graph from which a transport matrix can be distilled. We explicitly solve for the equilibrium solution of simple test problems with two distinct methods. The first method is equivalent to the explicit SIMPLEX transport in a static medium, whereas the second is based on the diagonalisation of the transport matrix. We argue that for scattering problems with optically thick regions, the second solution method is (much) more efficient. The possibility of using the eigenvalue spectrum of the transport matrix as a proxy for speed of convergence of the explicit solution method is briefly touched upon.



## 6.1 INTRODUCTION

In Chapter 3 we already exploited the advantage of mathematical clarity of the SimpleX method in the context of error analysis. Here we explore a mathematical description of the SimpleX algorithm in the context of graphs.

In a general sense, a graph is a system of connections between discrete points in space. In the Voronoi-Delaunay mesh, the set of all Voronoi nuclei are connected uniquely by the Delaunay edges. Photons travelling along these edges can thus be thought of as transported along the connecting lines of a graph. It is this duality, the equivalence of SimpleX transport and graph theory, that we touch upon in this chapter. To give the reader a feeling for the advantages of studying SimpleX transport in the context of graph theory and Markov processes, we outline two anticipated results. First, a dual description provides additional mathematical tools to study the original algorithm. Questions about the uniqueness and existence of stationary solutions to given transport problems are stated and answered more naturally in the language of graphs for instance. Second, for some transport problems it may be computationally (much) more efficient to use the solution methods presented below than to explicitly move the transported quantities to and fro over the computational mesh as in the SimpleX algorithm. Studying the duality between SimpleX transport and Markov chains on a graph can thus improve our understanding of the SimpleX algorithm and may lead to more efficient solution of certain transport problems.

The focus of this chapter is to introduce the concept of Markov chains relevant for the description of SimpleX transport and to illustrate the possible advantages of such a description rather than to give a thorough evaluation of the field. We limit our discussion in this text to equilibrium solutions of radiative transfer problems in static media. That is, the opacity of the medium does not change as a function of the local radiation field (as in ionisation chemistry).

Regardless of these limitations, the algorithm presented in this chapter can be applied to realistic problems of astrophysical interest. Possible applications include the determination of the temperatures in an inhomogeneous distribution of dust particles or the spatial photon density in a scattering-dominated stellar atmosphere. Because the properties of the mesh (volume, vertex-positions, connectivity) do not need to be stored in working memory, the requirements of the method are very modest. This allows for problems with many resolution elements.

In Sect. 6.2, we introduce Markov processes, which are most relevant for our study. We then take a step to translating the various means of RT within the SimpleX algorithm to a matrix-based representation in Sect. 6.4 and explicitly construct the matrices in Sect. 6.5. Subsequently we will demonstrate the equivalence of the Markov-matrix result with the direct RT solution for several simple toy-problems in Sect. 6.6. In Sect. 6.7 we shortly comment on the convergence of the RT and the matrix diagonalisation.

## 6.2 MARKOV PROCESSES

Markov processes (after the Russian mathematician Andrey Markov) are time varying phenomena for which the Markov property holds. Colloquially speaking, a process that satisfies the Markov property is memoryless, its future state depends only on its present state but never its past.

**Definition 1.** Let  $X_0, X_1, X_2, \dots$  be a sequence of random variables and let  $\mathbf{P}$  denote the probability. The sequence  $(X_0, X_1, \dots)$  is called a **Markov chain** if this random process possesses the **Markov property**:

$$\mathbf{P}(X_{n+1} = x | X_0 = x_0, X_1 = x_1, \dots, X_n = x_n) = \mathbf{P}(X_{n+1} = x | X_n = x_n) \quad (6.1)$$

Or in words: The conditional distribution of the state  $X_{n+1}$  given the states  $(X_0, X_1, \dots, X_n)$  depends only on  $X_n$ . Where we call  $X_0$  the initial (starting) state, and  $X_n$  the state after  $n$  steps.

Most games played with dice satisfy the Markov criterion in contrast to most card games where, for instance, the content of the stack reflects the games past. If a Markov process is defined on a discrete time basis, it is often referred to as a Markov *chain*. Note that we will introduce the concepts of Markov theory exclusively in the context of our transport problem on a graph. The definitions and theorems are much more general though and apply to any form of Markov chain imaginable.

## 6.2.1 The transition matrix

Whether it concerns weather types, outcomes of a dice-roll or positions in space, a Markov chain always describes a series of states that follow each other. In the context of our radiative transfer problem, we can imagine following a single photon that travels from one position to the next, tracing out its particular chain of states. Although this is a useful perspective, we might also look at the Markov chain from the states point of view, however. There might be many photons in our volume, each state hosting more than one photon. At each iteration of the Markov chain, these photons will change states with a given probability that depends on the process. All the relevant information about how to transition from one state to the next can be encoded in the *transition matrix*

$$M = \begin{pmatrix} M_{00} & M_{10} & \dots \\ M_{01} & M_{11} & \dots \\ \vdots & \vdots & \ddots \end{pmatrix}. \quad (6.2)$$

**Definition 2.** Let  $M$  be a  $k \times k$  matrix with elements  $\{M_{i,j} : i, j = 1, 2, \dots, k\}$ . A random process  $(X_0, X_1, \dots)$  with finite state space  $S = \{s_1, s_2, \dots, s_k\}$  is said to be a **homogeneous Markov chain with transition matrix**  $M$ , if for all  $n$ , for all  $i, j \in \{1, \dots, k\}$  and for all  $i_0, \dots, i_{n-1} \in \{1, \dots, k\}$  we have:

$$\begin{aligned} \mathbf{P}(X_{n+1} = s_j | X_0 = s_{i_0}, X_1 = s_{i_1}, \dots, X_{n-1} = s_{i_{n-1}}, X_n = s_i) &= \\ \mathbf{P}(X_{n+1} = s_j | X_n = s_i) &= M_{i,j} \end{aligned}$$

The matrix elements  $M_{i,j}$  are transition probabilities that express the probability of transitioning from state  $s_i$  to state  $s_j$ . In our context, a photon will have to do a Monte Carlo experiment and decide which of the available next states it will go into, each transition having its own probability. Alternatively, we might interpret the transition probabilities as fractions that decide which part of the light goes into a given state. Of course, this interpretation is equivalent

to the first if the number of photons is large and we will use it in our practical examples in this chapter.

Every transition matrix is nonnegative:  $M_{i,j} \geq 0$  for all  $i, j \in \{1, \dots, k\}$ , and each row of the matrix sums up to unity:  $\sum_{j=1}^k M_{i,j} = 1$  for all  $i \in \{1, \dots, k\}$ . Because of this property, the transition probabilities can also be interpreted as fractions of the content of vertex  $i$  to be sent to vertex  $j$  as discussed above. A matrix with these properties is also called a *stochastic matrix*.

The term *homogeneous* in the above definition means that the probability from one state to another does not change over time. An inhomogeneous Markov chain has transition matrices  $M^{(1)}, M^{(2)}, \dots$  for each time step. We will not be considering inhomogeneous Markov chains in this text.

The transition matrix  $M$  hence reflects the probability of transitioning between the possible states of the Markov chain. In our case, these states are the nodes of the graph which are in turn identified with the vertices that constitute the basis of the transport mesh. If we use the Voronoi-Delaunay triangulation as the connection between the vertices, the non-negative entries of the transition matrix are completely determined by this triangulation. Note that every edge has in general *two* numbers associated with it, one for each direction of the transport. In other words;  $M_{ij} \neq M_{ji}$  in general.

As we will see in Sect. 6.4, the probabilities of these connections are still a free parameter which can be associated with the concept of solid angle.

## 6.2.2 The initial condition

Now we give a definition of the starting point of the Markov process: the initial condition, and how we denote the distribution of states after several times.

**Definition 3.** *The start of the Markov chain is called the **initial distribution**, and is represented as a row vector given by:*

$$\mu^{(0)} = (\mu_1^{(0)}, \mu_2^{(0)}, \dots, \mu_k^{(0)}) \quad (6.3)$$

$$= (\mathbf{P}(X_0 = s_1), \mathbf{P}(X_0 = s_2), \dots, \mathbf{P}(X_0 = s_k)) \quad (6.4)$$

Similarly, the row vectors  $\mu^{(1)}, \mu^{(2)}, \dots$  denote the Markov chain at times 1, 2,  $\dots$ , so that for time  $n$  we have:

$$\mu^{(n)} = (\mu_1^{(n)}, \mu_2^{(n)}, \dots, \mu_k^{(n)}) \quad (6.5)$$

$$= (\mathbf{P}(X_n = s_1), \mathbf{P}(X_n = s_2), \dots, \mathbf{P}(X_n = s_k)) \quad (6.6)$$

Depending on the problem, the initial state might have one or all states occupied by photons. Sometimes, it is possible to guess an initial state that is similar to the final equilibrium state of the chain. This might improve the convergence of the solution and can save a substantial amount of time.

The next theorem will tell us how to compute future distributions, given an initial state.

**Theorem 1.** *For a Markov chain  $(X_0, X_1, \dots)$  with state space  $\{s_1, \dots, s_k\}$ , initial distribution  $\mu^{(0)}$  and transition matrix  $M$ , we have for any  $n$  that the distribution  $\mu^{(n)}$  at time  $n$  satisfies*

$$\mu^{(n)} = \mu^{(0)} M^n \quad (6.7)$$


*Proof.* The proof can be found in Haggström (2001). There induction by  $n$  is used to show that  $\mu^{(n+1)} = \mu^{(n)}M = \mu^{(0)}M^{n+1}$ . □

### 6.2.3 Graphs and reducibility

It is possible to picture a Markov chain by its so-called *transition graph*. The nodes of the graph represent the states and arrows between the nodes indicate the connectivity with the transition probabilities. Of course this also works the other way around, we can start with a graph, with transition probabilities between nodes, and list these probabilities in a matrix. We distinguish between two classifications for Markov chains (or their equivalent graphs): irreducible and reducible Markov chains.

**Definition 4.** A Markov chain  $(X_0, X_1, \dots)$  with state space  $S = \{s_1, s_2, \dots, s_k\}$  and transition matrix  $M$  is said to be **irreducible** if for all  $s_i, s_j \in S$ , there is a positive probability of ever reaching state  $s_j$  when we start from  $s_i$ . Or in other words: The chain is irreducible if for any  $s_i, s_j \in S$  we can find an  $n$  such that  $(M^n)_{i,j} > 0$ . Otherwise, the chain is called **reducible**.


The graph in Fig. 6.1 is irreducible, every state can be attained from every other state. In Fig. 6.2 we show the contrasting case of a reducible graph. Here state  $B$  cannot be reached from any other state.



**Figure 6.1:** Left panel: Irreducible graph. Right panel: Matrix representation of left panel.

### 6.2.4 Periodicity

Another important property of a Markov chain is its *periodicity*. This tells us whether there are preferential paths involved when we want to travel back to the state that we came from. It could be that there exist only even or odd paths from a state to itself for example. The formal definition uses the greatest common divisor of all possible paths as a measure for periodicity.



**Figure 6.2:** Left panel: Reducible graph. Right panel: Matrix representation of left panel. Note that the second row of the matrix contains only zeros because state  $B$  is cannot be reached from any other state.

**Definition 5.** A Markov chain is said to be **aperiodic** if all its states are aperiodic. A state  $s_i$  is aperiodic if its period  $d(s_i)$  is 1; where

$$d(s_i) = \gcd \{ n \geq 1 : (M^n)_{i,i} > 0 \}. \quad (6.8)$$

Otherwise, the chain is said to be **periodic**.

If we use the Voronoi-Delaunay triangulation as a transport graph, it can be shown straightforwardly that the resulting Markov chain is aperiodic. The argument goes as follows: the Delaunay triangulation can be thought of as a space-covering collection of simplices. These simplices allow for both a path consisting of three edges and a path of four edges linking every node to itself. Therefore, the greatest common divisor is unity and the graph is aperiodic. When we use a different prescription for the connectivity of the nodes, this argument no longer holds and we have to check the periodicity explicitly.

### 6.2.5 The Perron-Frobenius theorem

Now we state an important theorem on nonnegative, aperiodic, irreducible matrices: the Perron-Frobenius theorem. We adopt the version given in Meyer (2001) and Seneta (1981). The proofs can be found in Meyer (2001), and we omit them here.

**Theorem 2** (Perron-Frobenius theorem for nonnegative matrices). Let  $A = (a_{ij})$  be an  $n \times n$  nonnegative matrix, let it be irreducible with period 1. Then the following holds:

- There exists a positive real number  $r$  such that for all eigenvalues  $\lambda$ :  $|\lambda| \leq r$

- $r$  is a simple root of the characteristic polynomial of  $A$ , hence  $r$  is an eigenvalue of  $A$ , with algebraic multiplicity one.
- The left and right eigenvectors  $v$  and  $w$  corresponding to  $r$  have strictly positive entries.
- Any nonnegative left eigenvector of  $A$  is a scalar multiple of  $v$ , and any nonnegative right eigenvector of  $A$  is a scalar multiple of  $w$ .
- $r$  satisfies the inequalities  $\min_j \sum_i a_{ij} \leq r \leq \max_j \sum_i a_{ij}$

This is an important result. It tells us that, if we can formulate our RT problem as a conservative transport problem on a connected graph, there is only one, unique (invariant) solution to that problem.

### 6.2.6 Invariant distributions

If a Markov chain is irreducible, then the last assertion of the Perron-Frobenius theorem states that the maximal eigenvalue of the transition matrix equals 1, since the minimal and maximal rowsums are both unity. This immediately implies that there exists a (unique) positive eigenvector  $x$  such that  $Mx = x$ , which leads us to the following definition.

**Definition 6.** Let  $(X_0, X_1, \dots)$  be a Markov chain with state space  $S = \{s_1, s_2, \dots, s_k\}$  and transition matrix  $M$ . A row vector  $\pi = (\pi_1, \dots, \pi_k)$  is called an **invariant** (or: stationary, equilibrium) distribution for the Markov chain, if:

- $\pi_i \geq 0$  for  $i = 1, \dots, k$  and  $\sum_{i=1}^k \pi_i = 1$ , and
- $M\pi = \pi$ , meaning that  $\sum_{j=1}^k \pi_j M_{ij} = \pi_i$  for  $i = 1, \dots, k$

Now we state the most important theorem of this section. It tells us how to find the invariant distribution, if it exists, when the Markov chain is irreducible and aperiodic, independent of the choice of the initial distribution. Due to the Perron-Frobenius theorem, we already know for our case, such a stationary distribution indeed exists. The theorem is adapted from (Norris 1999, Sect 1.8, Theorem 1.8.3) to fit our notation.

**Theorem 3** (Convergence to invariant distribution). Let  $(X_0, X_1, \dots)$  be an irreducible and aperiodic Markov chain with state space  $S = \{s_1, \dots, s_k\}$ , any initial distribution  $\mu^{(0)}$  and transition matrix  $M$ . Suppose  $M$  has an invariant distribution  $\pi$ . Then

$$P(X_n = j) \rightarrow \pi_j \quad (6.9)$$

as  $n \rightarrow \infty$  for all  $j$ . In particular,

$$M_{i,j}^n \rightarrow \pi_i \quad (6.10)$$

as  $n \rightarrow \infty$  for all  $i, j$ .

*Remark:* The last assertion can be rephrased as when  $n \rightarrow \infty$ ,  $M^n \rightarrow \Pi$ , where  $\Pi$  is the limiting matrix whose entries are constant along the rows, i.e. the columns of  $\Pi$  are repetitions of the same vector  $\pi$ .

## 6.3 TWO SOLUTION METHODS

In this section we introduce two different methods to obtain the static solution to RT problems in Markov chain form. The first method makes use of the intuitive concept of repeated iterations of matrix multiplication. This method is easy to implement but can show slow convergence if the second eigenvalue lies close to unity. The second solution method is more involved but has a convergence rate that is independent of the eigenvalue spectrum.

### 6.3.1 Method 1: explicit iteration

The *power iteration method* explicitly implements the concepts of Theorem 3. There we computed powers of  $M$  and let the result act on the initial distribution vector, but this is equivalent with iteration. In fact, it is more efficient to form  $v_k$  by multiplying  $M$  with  $v_{k-1}$  rather than compute the next power of  $M$  at every step.

**Algorithm 1** (The power iteration method). 1. Choose an initial vector  $v_0$

2. For  $k = 1, 2, 3, \dots$  compute:

$$(a) \quad v_k = \frac{Mv_{k-1}}{\|Mv_{k-1}\|}$$

$$(b) \quad \lambda_k = \frac{v_k^T M v_k}{v_k^T v_k}$$

3. stop when  $\|Mv_k - v_k \lambda_k\| < \tau |\lambda_k|$

We can prove that whenever there exists a dominant maximal eigenvalue, we can obtain the equilibrium solution by the power iteration method. The proof holds for the case where the matrix is diagonalizable which always is the case in our applications.

**Theorem 4.** Application of the power iteration method will make  $\lambda_k$  converge to the dominant eigenvalue  $\lambda_1$  and  $v_k$  converge to the corresponding eigenvector of matrix  $M$ .

*Proof.* Let  $v_0$  be a random vector. Suppose the exact eigenvectors of  $M$  are  $\{x_1, x_2, \dots, x_n\}$ , with corresponding eigenvalues  $\{\lambda_1, \lambda_2, \dots, \lambda_n\}$ . Since the eigenvectors of  $M$  form a basis for  $\mathbb{R}^n$ ,  $v_0$  can be written as a linear combination of these eigenvectors:

$$v_0 = \sum_{k=1}^n c_k x_k \tag{6.11}$$

Multiplication of this vector with  $M$  gives:

$$Mv_0 = \sum_{k=1}^n c_k Mx_k = \sum_{k=1}^n c_k \lambda_k x_k \tag{6.12}$$

Repeating this  $m$  times yields :

$$M^m v_0 = \sum_{k=1}^n c_k M^m x_k = \sum_{k=1}^n c_k \lambda_k^m x_k = \lambda_1^m \left( c_1 x_1 + \sum_{k=2}^n c_k \left( \frac{\lambda_k}{\lambda_1} \right)^m x_k \right) \tag{6.13}$$

Since  $M$  has an dominant eigenvalue,  $\left(\frac{\lambda_k}{\lambda_1}\right)^m \rightarrow 0$  for  $k > 1$  as  $m \rightarrow \infty$ . The rate of convergence thus depends on the ratio  $\frac{\lambda_2}{\lambda_1}$ . From part 2(a) of the algorithm we then easily see:

$$v_k = \frac{Mv_{k-1}}{\|Mv_{k-1}\|} = \frac{M^k v_0}{\|M^k v_0\|} \rightarrow x_1 \quad (6.14)$$

as  $k \rightarrow \infty$ .

The second step, 2(b) follows from the definition of eigenvectors and eigenvalues:

$$Mv_k = \lambda_k v_k. \quad (6.15)$$

Multiplying both sides with the transpose of  $v_k$  and rearranging gives the assertion. And finally, since  $v_k \rightarrow x_1$  as  $k \rightarrow \infty$ ,  $\lambda_k \rightarrow \lambda_1$  as  $k \rightarrow \infty$ .  $\square$

We see from the proof that the convergence will depend on the second largest eigenvalue  $\lambda_2$ , for when  $\lambda_2$  is almost as large as  $\lambda_1$ , the ratio  $\frac{\lambda_2}{\lambda_1} \approx 1$  and it's powers will vanish slowly. The  $\tau$  parameter in the algorithm is the tolerance.

The sparsity of the matrix  $M$  is exploited in step (a) of the algorithm since the number of operations for multiplying the matrix with a vector is of the order of the number of nonzero entries in the sparse matrix.

We note that this iterative solution method is completely equivalent to SimpleX transport with static media. The stable solution is reached by repeated application of the transport prescription on the discretely changing distribution of photons.

### 6.3.2 Method 2: diagonalisation

The static solution of Eq. (6.14) corresponds to the unique non-negative eigenvector of eigenvalue 1 and can also be obtained by diagonalisation of the transport matrix. We use a *conjugate gradient* method (Press et al. 1992) to solve the eigenvalue problem. Because the system is very sparse (approximately 15 non-zero entries per row of the matrix), this can be done both fast and with little memory requirement.

## 6.4 CORRESPONDENCE WITH CONTINUOUS TRANSFER

The link between continuous transfer and the Markov chains under study is most natural when we realise that in nature, the seemingly continuous flow of light is actually a discrete process at hart. The SimpleX method models this discrete process on an interaction-by-interaction basis, in a sense, imitating nature. What we perceive as a continuous flow of light is in fact a macroscopic average over a discrete Markov process. In this section we will describe the correspondence of the discrete steps in the Markov chain and the familiar expressions used in the macroscopic radiative transfer equation (cf. Eq. (1.3)).

Recall that the specific intensity  $I = I(\mathbf{r}, \mathbf{n}, \Omega, \nu, t)$  describes the energy emitted through a surface  $d\mathbf{A}$  in direction  $\mathbf{n}$  into a solid angle  $\Omega$  located at  $\mathbf{r}$  per unit time  $t$  per unit frequency  $\nu$ . Hence

$$dE = I\mathbf{n} \cdot d\mathbf{A}d\Omega d\nu dt. \quad (6.16)$$



In the Markov description, the quantities featuring in Eq. (6.16) are defined as follows

- $dt$ : the time interval  $dt$  is associated with a single step of the Markov chain. This is equivalent to one matrix-multiplication of the transition matrix  $M$  with the state vector  $\mathbf{x}$ . This is also equivalent to one sweep in the SIMPLEX algorithm.
- $dv$ : we will limit ourselves in this chapter to the transport of a single frequency component. The generalisation to multiple frequencies is possible but not straightforward because the SimpleX mesh represents a physical scale associated with the local mean free path of the photons. The mean free path depends on frequency and the mesh should reflect this dependency. Strictly speaking, every frequency requires its own triangulation but we will see in Chapter. 5 that this requirement can be relaxed considerably. Transport in frequency space is in turn realised by interconnecting the different meshes. Because the Markov representation of the transport process is completely equivalent to the SimpleX method, the relaxation of this requirement is also applicable here.
- $\mathbf{r}$ : this quantity corresponds to the position of the vertex from which the radiation is transported away.
- $\mathbf{n}$ : the direction of the radiation is given by the unit vector,  $\mathbf{e}_{ij}$ , pointing from the vertex,  $i$ , from which the radiation is transported away to the receiving vertex,  $j$ .
- $d\mathbf{A}$ : in continuous radiative transfer theory this quantity describes an infinitesimal surface (area with an orientation). If we increase the surface area, more radiation will pass through it. If we orient its normal parallel to the direction of the photons, the flow of photons through the surface is maximised. In the SIMPLEX algorithm, however, radiation is discretised over a number of directions (the Delaunay edges). In this context, the surface associated with a connection is the Voronoi wall,  $\mathbf{A}_{ij}$ , that separates vertex  $i$  from its neighbour  $j$ . The normal of this wall is parallel to the unit vector  $\mathbf{e}_{ij}$  making the dot product  $\mathbf{n} \cdot d\mathbf{A} = \mathbf{e}_{ij} \cdot \mathbf{A}_{ij}$  equal to  $A_{ij}$ . This is not an infinitesimal but a fixed area with fixed orientation.
- $d\Omega_{ij}$ : There is some freedom in the definition of opening angle as we have already discussed in detail in Sect. 3.3.4. Again, in the context of Voronoi-Delaunay radiative transfer, this quantity is no longer infinitesimal and we adopt the simple notation  $\Omega_{ij}$ . The solid angle of a connection is approximately given by

$$\Omega_{ij} \simeq \frac{A_{ij}}{(0.5l_{ij})^2}, \quad (6.17)$$

where  $l_{ij}$  is the distance between the nuclei  $i$  and  $j$  (it is also the length of the Delaunay edge connecting these vertices). The solid angle  $\Omega_{ij}$  and  $A_{ij}$  are thus not independent in the SIMPLEX algorithm.

Because we are dealing with discrete photon numbers rather than energy, we will also express the fundamental quantities of radiative transfer in terms of photon numbers. This simply means that we divide Eq. (6.16) by  $h\nu$  where  $h$  Planck's constant and  $\nu$  the frequency of the radiation.

In this formalism, the specific intensity in direction  $\mathbf{e}_{ij}$ ,  $I$  is given by

$$\mathbf{I}_{ij} \equiv \frac{h\nu P_i M_{ij}}{A_{ij} \Omega_{ij} dt} \mathbf{e}_{ij}, \quad (6.18)$$

where  $P_i$  is the number of photons residing at vertex  $i$  and  $M_{ij}$  is the element of the transition matrix that holds the fraction of  $P_i$  that streams along the Delaunay edge connecting vertex  $i$  with vertex  $j$ . The mean intensity at vertex  $i$ ,  $J_i$  is consistently defined as the sum of  $I_{ij}$  over the complete sky

$$J_i = \frac{1}{4\pi} \sum_j \mathbf{I}_{ij} \Omega_{ij} \quad (6.19)$$

$$= \frac{h\nu P_i}{dt} \sum_j \frac{1}{A_{ij}} \quad (6.20)$$

where the last equality holds because the solid angle of the total sky is  $4\pi$  and the vector sum of the unit vectors  $\mathbf{e}_{ij}$  vanishes. The flux between vertex  $i$  and  $j$  is also easily obtained by considering *both* directions of the transport of the  $ij$  edge

$$\mathbf{F}_{ij} \equiv \frac{h\nu}{A_{ij} \Omega_{ij} dt} (M_{ij} P_i - M_{ji} P_j) \mathbf{e}_{ij}. \quad (6.21)$$

The flux in arbitrary direction  $\mathbf{n}$  can be found by summing all projected contributions of  $F_{ij}$  in direction  $\mathbf{n}$

$$F(\mathbf{n}) = \sum_j \mathbf{F}_{ij} \cdot \mathbf{n}. \quad (6.22)$$

Given the above definitions we are equipped to compare Markov chain calculations directly to results obtained with SimpleX, other RT codes or analytical solutions to known RT problems.

## 6.5 MATRIX CONSTRUCTION

In this section we will describe in more detail the construction procedure of the transport matrix  $M$  to the end of capturing within the Markov formalism radiation transport as practiced in SimpleX. Although the formalism is completely general, we will, for reasons of convenience, assume that photons are the transported quantity and we will use the Voronoi-Delaunay triangulation to connect the nodes of our graph. First we translate the various transport algorithms introduced in Chapter 2 to the Markov formalism. We then proceed by introducing sources and sinks to the mesh.

### 6.5.1 Diffuse transport

Diffuse transport (DT) is described straightforwardly in the Markov representation by assigning to every (outgoing) edge of a vertex  $i$  a fraction  $M_{ij}$ .

This means that the number of photons that is transported along edge  $ij$  per iteration (or multiplication by  $M$ ) is given by

$$P_{ij} = M_{ij} P_i. \quad (6.23)$$

## 6.5.2 Ballistic transport

In ballistic transport (BT), the  $D$  outgoing directions of radiation depend on the incoming direction. In this sense, BT does not satisfy the Markov criterion of memorylessness. We can, however, switch from a vertex-based description to an edge-based one and regain the Markov criterion. To this end we introduce the notion of an *arrow*. For every edge, we have two arrows, one for each direction of transport. So, the edge connecting vertex  $i$  to its neighbor  $j$  has two arrows:  $a_{ij}$  and  $a_{ji}$ .

For a given an arrow, the transport is again independent of history because the arrow is connected to three other arrows (the most straightforward directions from Sect. 2.1.4) and that is all we need to know to construct the transition matrix. The constants connecting an arrow to its unique other three arrows can be chosen to be  $1/3$  or something more involved as in Sect. 3.4.3.

Because there are on average  $15.54 \dots$  arrows per vertex, the matrix is going to be larger by the square of that factor, approximately 241. Because the matrix is very sparse and only non-zero entries have to be stored, the actual increase in memory needed to store  $M$  is only nine times larger than in the case of diffuse transport, however. In the case of diffuse transport, the number of entries per vertex is  $15.54 \dots$  which implies  $N \times 15.54 \dots$  entries in total. For BT this number is  $3 \times N \times 15.54 \dots$ , only three times larger.

Although we limit our explicit examples to diffuse transport in this preliminary text, there are no conceptual difficulties in application of ballistic transport in the Markov context.

## 6.5.3 Direction conserving transport

In DCT, the situation is only marginally more complicated than for BT.<sup>1</sup> This can be understood right away when we realise that DCT uses ballistic transport to move photons from vertex to vertex. The only difference is that the radiation is stored in  $N_b$  direction-bins that do not know of each others existence. A photon which is emitted in the  $i$ -th bin will never leave it.

This does not mean that the transport process is identical to ballistic transport. The choice for transport direction *does* depend on the direction bin because every direction bin is associated with a unique arrow at every site of the mesh.

Because the direction bins are independent, the problem can be described by  $N_b$  independent transport matrices. These matrices are even more sparse than the single matrix describing BT because every direction bin is associated with only one arrow. The matrices for DCT are thus a factor  $15.54 \dots$  more sparse.

## 6.5.4 Sources and sinks

Up to this point we have considered transport along the edges of graphs but we have not asked ourselves about the origin of the transported quantities. When we want to calculate more physically relevant examples, the graph needs to be extended with sources and sinks of radiation.

<sup>1</sup>In practice this is not entirely true. The random rotations needed in DCT in order to suppress imprinting of preferential directions in the mesh would necessitate the construction of many transport matrices rendering DCT very slow.

We define a source as a vertex of the graph that does not receive radiation from neighbours but sends its contents to all neighbours isotropically. A sink in contrast only receives radiation and sends nothing to its neighbours. In other words, sources have only arrows that are outgoing and sinks have only arrows that are incoming.

There is one extra requirement, however, the total number of photons must be a conserved quantity on the graph. To enforce this we connect the sinks directly to the sources with connections that are thus non-Delaunay (in contrast to all others). This is necessary to guarantee a stationary solution because otherwise all radiation would vanish into the sinks and we would end up with the trivial stationary solution of *no* radiation anywhere. These sink-to-source connections are uni-directional as would be expected for a pure sink.

Sources with different strength are readily implemented by connecting every sink to *all* sources with weights that express the relative strengths of the sources. Absorbing boundaries are hence accomplished by a boundary of sink-particles around the domain.

## Absorption

The sinks of radiation described above absorb *all* radiation. In nature, however, a photon passing through an optically thin parcel of gas has a probability smaller than unity to be absorbed. To model this, we can allow vertices in the domain to absorb a fraction of the radiation that has reached them. Conveniently, the absorption of radiation is proportional to the local intensity. It is thus trivial to associate the value of  $M_{ii}$  with the absorption coefficient of the vertex.

In the transport matrix we can model absorption by connecting the edge to itself by setting the element  $M_{ii}$  to a non-zero value. The vertex thus retains a non-zero fraction of its content. If we apply this methodology, all the radiation on our mesh will end up in the absorbing vertices and the stationary state will tell us where in the mesh the radiation has been absorbed.

Physically speaking this solution corresponds to tracing down the fate of a single pulse of radiation as it spreads through space and finds absorbing material on its path.

If we want to model a continuously emitting source surrounded by an absorbing medium, however, we must make sure the radiation is conserved on the mesh and we need to connect absorbing vertices to the source(s) instead in the same fashion as sink vertices. The only difference being that the absorbing vertices *also* have outgoing connections to their direct neighbours.

## 6.6 DEMONSTRATION

In this section we explicitly calculate the stationary state for the Markov chain of several simple test problems. The first test problem (Sect. 6.6.1) serves as a quirky example of how the nature of the graphs can leave its imprint on the solution. Consequently it stresses the importance of using a physical weighting scheme for BT as proposed in Sect 3.4.3.

The examples of Sect. 6.6.2 and 6.6.3 serve two purposes, to ascertain that the static solution for a single source in a non-absorbing medium behaves as expected (the number density of photons falling off with distance as  $r^{-2}$  and to show that solution method 1 converges less fast for inhomogeneous density distributions.

Finally, the test in Sect. 6.6.4 serves as an example of the inclusion of absorption.

### 6.6.1 Using the valence as weights

Before we turn to more physical examples, we shortly discuss a curious case where the equilibrium solution can be found directly and without any effort.

For diffuse transport, we can choose the matrix element  $M_{ij}$  of a vertex  $i$  to be equal to  $1/V_i$  where  $V_i$  is the number of edges of the vertex, also called *valence* in the context of Markov chains. With this choice,  $M$  is properly normalised (all rows sum to unity) and the solution is rather straightforward. The entries in  $\mathbf{x}_f$  will reflect the valency of the vertices in the sense that for every entry the following equality holds


$$\frac{\mathbf{x}_{f,i}}{\mathbf{x}_{f,j}} = \frac{V_i}{V_j}. \quad (6.24)$$

Why this must be the case is easily demonstrated by considering a single edge, connecting vertex  $i$  with vertex  $j$ . The state is invariant if there is no net transport along the edge. Because the transported fraction from vertex  $i$  to vertex  $j$  is given by  $1/V_i$ , and vice versa, we must have

$$\frac{P_i}{V_i} = \frac{P_j}{V_j}, \quad (6.25)$$

which is equivalent to Eq. (6.24).

An example is given in Fig. 6.3 where a Delaunay triangulation of  $10^3$  points in the plane is shown. The shading gives the relative occupation number of the final state where every triangle is coloured according to the mean of the valence of the (three) vertices at its corners.




**Figure 6.3:** *Left:* Delaunay triangulation of  $10^3$  points in the plane. Shading is lighter for higher valence. *Right:* histogram of the valence for this example. The number of neighbours for the two-dimensional Delaunay triangulation has expectation value 6 with considerable scatter.

However amusing, this property is not desirable in the context of radiative transport in SimpleX because the number of outgoing edges has no physical meaning. It proves very useful, however, to check if our methods for inverting the matrix equations to obtain  $\mathbf{x}_f$  work correctly.

From a physical point of view, the mesh is merely the canvas on which the medium is described. It should not matter for the outcome of a physical experiment what the valence of a particular vertex so we will use the Icosahedron weights (see Sect. 3.3.4) for the tests below.

### 6.6.2 Single source in a homogeneous medium with diffuse transport

In this section we take a modest step towards realistic application by considering the problem of a single source embedded in a homogeneous medium with absorbing boundary conditions. We place a single source vertex in the centre of a three dimensional mesh consisting of  $10^4$  vertices in the unit volume. All vertices further than 0.45 away from the centre are turned into sinks. We thus end up with a source in a spherical volume surrounded by a boundary layer of sinks (see left panel of Fig. 6.4). Next, the sinks are connected to the source and we do not allow any other connections towards the source.




**Figure 6.4:** Left panel: 2D projection of the homogeneous distribution of particles used in Sect. 6.6.2. The sinks that surround the domain have been omitted from the plot. The source particle has been plotted as a filled square. Right: the same as the left panel but for the particle distribution used in Sect. 6.6.3.

A physical scenario that fits this description would be a single source of radiation embedded in a homogeneous medium of perfectly scattering particles. All radiation arriving at a vertex is isotropically redistributed to its neighbouring particles without changing its frequency.

The initial state is a vector filled with every entry initialized to unity. The stationary solution does not depend on the initial state as long as the same amount of radiation is present on the graph. It *will* influence the convergence of the solution method, however as we will demonstrate in Sect. 6.7.

Figure 6.5 shows the flux distribution as a function of distance from the source for a simulation where the outgoing edges get equal weight ( $1/V$ ; open squares). We have checked that the result does not change if every edge is weighted with its associated solid angle (Icosahedron weights).

There are several things to note here. The flux-points show considerable scatter. This is a direct consequence of the stochastic nature of diffuse transport. Because the photons are continuously redistributed in all directions, local irregularities in the mesh can increase or decrease the



**Figure 6.5:** Radial flux as a function of radial distance from the central source for  $10^4$  nuclei in a spherical volume. The spread is a result of the stochastic nature of the transport. The flux shows the appropriate radial dependence as can be seen from the line with slope -2 shown for comparison.

local flux somewhat. Nevertheless, the combined effect of the photons is a net radial flux that shows the correct  $r^{-2}$ -dependence. This is direct evidence that our Markov chain formulation indeed solves the radiative transfer problem as we claimed before.

### 6.6.3 Single source in a heterogeneous medium with diffuse transport

In this section we repeat the experiment described above but we place a dense knot of vertices in the domain. We do this by distributing half of the total of  $10^3$  vertices randomly over the computational domain and the other half in a Gaussian distribution of points with a standard deviation of 0.02 centered on  $(x, y, z) = (0.75, 0.5, 0.25)$  (see right panel of Fig. 6.4).

From a physical point of view, this problem is expected to converge slower if we apply Eq. (6.3.1) because the photons can be trapped in the region of high vertex number density. This trapping inhibits the communication between all regions of the domain, slowing down the convergence to a stationary solution.

Thus, by introducing the clump, the average path-length between any two vertices in the graph (expressed in number of edges to be traversed) has increased dramatically with respect to the homogeneous case.

The converge of the conjugate gradient method does not depend on the connectivity of the transport matrix, however. It is certain to converge within  $N$  iterations, where  $N$  is the number of vertices in the mesh.

### 6.6.4 Single source in an absorbing medium with diffuse transport

In this section we add absorption to the medium, otherwise the setup is as in Sect. 6.6.2. We have assigned an optical depth of unity to every non-sink vertex of the graph. This amounts to adding a connection with weight  $1 - e^{-1} \simeq 0.63$  from every vertex to the source.

With this setup we expect the radiation-field to fall off with increasing distance from the source more quickly than expected from geometric dilution alone. Quantitatively speaking, the

flux must fall of as

$$F(r) = F_0 \frac{e^{-N(r)}}{r^2} \quad (6.26)$$


where  $F_0$  is the flux at the position of the source and  $N$  is the mean number of vertices between the source and position  $r$ . We note that  $N$  depends on the distance  $r$  as  $N \propto r^2$  because the diffuse transport is equivalent to a random walk along the mesh. Equation (6.26) can thus be rewritten to

$$F(r) = F_0 \frac{e^{-N_{lin}r^2}}{r^2} \quad (6.27)$$

where  $N_{lin}$  is the number of vertices passed if the path would *not* be a random walk. We have performed a fit to the flux (in logarithmic bins) as a function of  $\log(r)$  using the fitting function obtained by expressing Eq. (6.27) as a function of  $x \equiv \log(r)$ ,

$$\log(F(r)/F_0) = A \exp(Bx) + Cx. \quad (6.28)$$

Here  $A$  represents  $N_{lin}$ ,  $B$  is expected to be equal to 2 and  $C$  is the constant that describes the geometric component of the dilution and is expected to be equal to -2 (as the flux falls off with  $r^{-2}$ ). As can be seen from Fig. 6.6, the estimator for  $N_{lin}$  is of order ten, which is expected (although somewhat high) for  $10^3$  vertices. The factor  $B$  is indeed very close to 2 as expected from the random walk nature of this type of transport. For the geometric dilution term  $C$ , the value is surprisingly far from the expected -2. A plausible explanation is that geometrical dilution is completely dominated by the absorption in this case and  $C$  is just not constrained very well. For sake of comparison we also show the data for the same setup without absorption (Sect. 6.6.2) which conforms to expectation with a slope of about 2.



**Figure 6.6:** Flux as a function of distance with (open squares) and without (filled squares) absorption. A fit to the data is shown as well as unit standard deviation error-estimates.




## 6.7 CONVERGENCE

When we are interested in the stationary state, there is only one eigenvalue of importance, the one with eigenvalue 1. The rest of the eigenvalue spectrum, however, carries information on the convergence of the Markov chain. Roughly speaking, if many eigenvalues exist with values close to one, the solution of the matrix multiplication (e.g. the repeated application of Eq. (6.3.1)) is expected to converge slowly.

As an example, the eigenvalue spectrum of the simulations described in Sect. 6.6.2 and Sect. 6.6.3 are shown in the left panel of Fig 6.7. Eigenvalues with rank larger than 700 are all zero and omitted from the figure. For the homogeneous test, the one-but-highest eigenvalue (after unity) is 0.805, substantially further from unity than the typical distance between eigenvalues. This means that the stationary state is well defined and well separated from the other eigenstates. For the inhomogeneous test from Sect. 6.6.3 the second eigenvalue is 0.981, much closer to unity. For this test, there are many more non-zero eigenvalues (roughly twice as many, as can be seen from the left panel of Fig 6.7). Although we do not pursue this here, this observation strongly suggests that we can use the number of non-zero eigenvalues as a proxy for the speed of convergence of the associated problem.

The right panel of Fig 6.7 shows the (logarithm) of the absolute value of the difference of the current and the final state of the Markov chain as a function of the iteration of Eq. 6.3.1. The magnitude of this difference is a direct measure of the speed of convergence which falls in between quadratic and exponential for the Sect. 6.6.2 test but becomes sub-quadratic for a number of iterations larger than 11 for the Sect. 6.6.3 test.



**Figure 6.7:** *Left:* Norm of the eigenvalues for the ‘homogenous’ problem described in Sect. 6.6.2 and the ‘clumpy’ problem Sect. 6.6.3. *Right:* Log of the absolute value of the difference between the current and final state as a function of matrix multiplication. The convergence of the multiplications is faster than quadratic but somewhat slower than exponential for the homogenous simulation. The simulation with a clumpy medium shows exponential behaviour in the first 11 iterations but much slower decline (as  $10^{-0.01x}$ ) convergence at later times.

This behaviour is conform our intuition of the physical difference between the problems presented in Sect. 6.6.2 and Sect. 6.6.3. As noted at the end of Sect. 6.6.3, the conjugate gradient method converges within  $N$  iterations, where  $N$  is the dimension of the matrix. In contrast, we cannot set a limit for the convergence of the ‘iterative’ problem. The ‘clumpy’ simulation is

converging very slowly (less than quadratic) rendering brute force radiation transport computationally costly.

## 6.8 DISCUSSION

Given the results obtained in this preliminary study, it seems feasible to apply this method to more realistic RT problems. In its current form, the class of problems is limited to those where the opacity of the medium does not change due to interaction with the radiation field.

To overcome this limitation, we must allow the entries of the transport matrix to change during the solution process. Although we do not explicitly prove that this is possible in our framework, we are optimistic about such an extension because there are many parallels between our approach and the vast body of work that has been dedicated to the iterative solution of ‘changing matrix problems’ in the context of  $\Lambda$ -iteration (e.g., [Mihalas 1978](#)). Here, every resolution element in a transport problem is linked to every other, resulting in a large, non-sparse matrix.

It is worthwhile to note that it is possible to speed up this procedure considerably by considering only those ‘connections’ in the matrix that correspond to cells that are close to each other in terms of optical depth, the so-called *Accelerated*  $\Lambda$ -Iteration (ALI, [Ng 1974](#)) technique. This effectively makes the matrix significantly more sparse as only a few entries close to the diagonal remain non-zero. We note that this speedup is inherent to our approach because every cell is only connected to its Voronoi neighbours.

## 6.9 SUMMARY

- It is possible to re-formulate the explicit transfer of radiation on a Voronoi-Delaunay mesh as a Markov chain on a connected graph.
- For such a formulation, a stationary and unique solution is guaranteed by the Frobenius-Perron theorem.
- We have used two methods for solving simple scattering problems in this framework
  - Explicit iteration of the transport matrix (i.e. Eq. (6.3.1)). This is equivalent to SimpleX radiative transfer because the transported quantity is explicitly moved between connected vertices (with or without weighting factors).
  - Implicit solution using a *conjugate gradient* method. This approach is guaranteed to converge in a number of iterations that depends only on the size of the matrix.
- The simulations performed in this chapter are of a preliminary nature and at this stage a quantitative comparison of speed between the two solution methods is not in order. We can, however, predict that the implicit method will become (much) more efficient than the explicit method based on the convergence behaviour observed in Sect. 6.7.
- The results presented here suggest that for a scattering problems in (very) inhomogeneous media, solving the matrix-diagonalisation problem is computationally much more efficient than performing direct radiative transfer. The problems for which this is the case are characterized by one or more regions of the graph which are weakly connected to the

rest of the domain. For a photon present in such a region it takes many steps along the edges of the graph to communicate with other regions in the domain. In a physical situation these problems typically have optically thick clumps that trap radiation for many scatterings.

- We have also shown that absorption of radiation can be included trivially in this formalism.

## CHAPTER 7

---

# Late ionization of filaments or early ionization of voids?

C. J. H. Kruip<sup>1</sup>, K. Finlator<sup>2</sup>, B. D. Oppenheimer<sup>1</sup>,  
J.-P. Paardekooper<sup>3</sup> & V. Icke<sup>1</sup>

WE investigate whether the results obtained by [Finlator et al. \(2009\)](#) (FODO09) are reproducible with the `SIMPLEX` radiative transfer method. We focus on the morphology of the reionization scenario, which is the main theme of FODO09. These authors have used a ‘moment method’ (described in [Finlator et al. \(2009\)](#)), whereas `SIMPLEX` uses Delaunay radiative transfer as described in Chapter 2. We separate the effects due to resolution from those due to the use of a different radiative transfer method. We are not able to reproduce the inside-out-middle reionization scenario described in FODO09 and attribute this to the use of a different method for radiative transfer.

---

<sup>1</sup>Leiden University, Niels Bohrweg 2, 2333 CA Leiden, Leiden, the Netherlands

<sup>2</sup>University of California, Santa Barbara, Santa Barbara, CA 93106, USA

<sup>3</sup>Max-Planck-Institut für extraterrestrische Physik, Giessenbachstrae, 85748 Garching, Germany

## 7.1 INTRODUCTION

In the cosmological  $\Lambda$ -CDM model, the first sources of radiation capable of ionizing hydrogen are thought to form around  $z \approx 20 - 30$  (Barkana & Loeb 2001) and thus initiated the epoch of reionization (EoR). The nature of these first sources is uncertain but likely candidates are massive population III stars and, subsequently, population II stars and quasars.

According to linear theory, sources are expected to form in over-dense knots, the nodes of intersection of the cosmic web. It is in these environments that reionization begins. After the reionization of the densest structures, ionizing radiation streams into the regions of lowest density, the voids. Owing to their low number density of neutral hydrogen (and helium), voids ionize quickly and the resulting ionized bubbles overlap until the average neutral fraction is below  $10^{-4}$  (as observed in the present day universe, e.g. Fan et al. 2006). This is generally referred to as ‘inside out’ (IO) reionization (Iliev et al. 2006b; McQuinn et al. 2007; Trac & Cen 2007; Zahn et al. 2007).

A more complicated picture, however, has arisen from semi-analytical work (Choudhury et al. 2009) and the numerical study by FODO09. The filaments, structures with intermediate density, were found to reionize last. This scenario is referred to as ‘inside-out-middle’ (IOM) reionization and is characterised by a reversal of the ionization front, progressing from voids into filaments in the later stages. Similar results have been observed in the simulations of Ciardi et al. (2003); Gnedin (2004) and more recently Petkova & Springel (2011).

The statistical properties of the reionization morphology has important observational consequences. Specifically, the redshift-dependent distribution of neutral hydrogen (observable from its redshifted 21-centimeter radiation) directly determines the brightness temperature power spectrum, likely one of the first major observational handles on the properties of reionization (e.g., McQuinn 2010). It is therefore important to identify which parameters are shaping the global reionization morphology and their relative importance.

The morphology of the EoR depends both on the type, abundance and distribution of the ionizing sources (e.g., McQuinn et al. 2007) and sinks of ionizing radiation (e.g., Zaldarriaga et al. 2004). A widely adopted approach to the problem of identifying sources in cosmological  $n$ -body simulations is to identify virialised structures (haloes) from the dark matter particle distribution and assign luminosities to these haloes according to some physically motivated prescription (e.g., McQuinn et al. 2007; Iliev et al. 2006b). FODO09 have taken a self-consistent approach to simulating reionization where cooling and metallicity-dependent star formation are treated in considerable detail. In their approach, the sources follow directly from the hydrodynamical simulation and need not be introduced a-posteriori. These hydrodynamical simulations have subsequently been post-processed using the RT method described in Finlator et al. (2009). They find an evolution of EoR morphology that agrees better with the IOM than the IO scenario.

Possible reasons for this result, which contrasts with the bulk of the numerical work so far, are the different source prescription used, the possibility that their RT grid does not capture all relevant physical scales (i.e. resolution) and the method for RT employed. As we will argue below, the source prescription is unlikely to be key factor in the observed difference in morphology and we, therefore, focus on the other two issues, resolution and RT method.

### 7.1.1 Outline

We first introduce the computational method used in Sect. 7.2. Here the emphasis lies on the construction of radiative transfer grids. In Sect. 7.3, several diagnostics that are used in this chapter to identify differences between the various simulations. For the sake of clarity, the resolution study and the method comparison are treated separately in Sect. 7.4 and Sect. 7.5.

## 7.2 METHOD

For the transportation of ionizing radiation we use the SimpleX method as described in [Paardekooper et al. \(2010\)](#). In this method, photons travel along the edges of an unstructured Delaunay triangulation. The transport of radiation along the connecting Delaunay edges of an unstructured Voronoi grid without memory of the original direction of emission (called Ballistic Transport) can lead to diffusion in optically thin regions of space. A direction conserving transport method (DCT for short) has been introduced ([Kruip et al. 2010](#); [Paardekooper et al. 2010](#)) to overcome this problem.


The Delaunay triangulation connects the nuclei of a Voronoi tessellation constructed from a discrete set of points in three dimensions. These points, in turn, represent the physical medium through which the radiation is traveling. SimpleX can work on continuous and grid-based density/opacity fields as well as fields described by discrete points (the SPH particles of the N-body simulation.) Both possibilities are exploited in this study.

Because this method can use the SPH particles directly as basis for the RT grid, the spatial resolution can be chosen to be equal to that of the original hydrodynamics. An example of this adaptive resolution is shown in Fig. 7.1 where we show a slice through the highest resolution grid used in FODO09 alongside a SIMPLEX grid using about equally many cells constructed directly from a subset of the SPH particles. FODO09 have used a RT grid of maximally  $96^3$  cells. Given that the size of the cosmological volume is  $16 \text{ h}^{-1} \text{ Mpc}$ , this results in a RT resolution of  $177 \text{ h}^{-1} \text{ kpc}$ . The dynamic range of the SimpleX grid spans roughly five orders in magnitude, equivalent to 16 levels of refinement in AMR jargon. The corresponding physical resolution is about  $160 \text{ h}^{-1} \text{ pc}$ , roughly a factor  $10^3$  higher than that of the original simulations. If resolution effects play a role, they are thus expected to manifest themselves in these simulations. We will study the effects of resolution on the morphology of reionization in Sect. 7.4. The comparison of the two RT methods, the moment method of [Finlator et al. \(2009\)](#) and SIMPLEX at equivalent resolution is performed in Sect. 7.5.

Several modifications to the method described in [Paardekooper et al. \(2010\)](#) have been made in this study. The most notable ones are that the evolution of distance and number density with the expansion of the universe is accounted for and a more precise treatment of the ionization of highly optically thick computation cells has been implemented.

### Correcting for cosmological expansion and optical thick cells

The snapshots describing the density field and the sources are taken at roughly 100 Myr intervals. The resulting discontinuous jumps in density and luminosity result in jumps of the ionization state as well (also see the inset of Fig. (6) of in FODO09). The differences in the



**Figure 7.1:** Comparing the RT grids at  $z = 6.49$  side-by-side. *Left:*  $96^3$  regular grid from FODO09. *Right:* SimpleX grid based on a random sampling of the original  $512^3$  SPH particles with  $10^6$  particles. The insets demonstrate the adaptive resolution of the Delaunay grid which typically captures a dynamic range of five orders of magnitude in length-scale.

density field are due to two mechanisms: dynamical (gravity and hydrodynamics) and cosmological (the expansion of the universe).

The first effect is non-linear for an increasing part of the simulation volume with decreasing redshift and can therefore not be corrected for other than taking more snapshots (but this is not possible as we are working with a given data-set). We do not interpolate density and luminosity between snapshots to allow for a fair comparison with the results in FODO09.

We can correct for the second effect by appropriate scaling of the relevant quantities: number density and distance. They depend on the scale factor,  $a$ , as  $a^{-3}$  and  $a$ , respectively.

Furthermore we have implemented a treatment for the ionization-state of a cell that distinguishes between optically thin and optically thick cells. If a cell is optically thin, the width of the ionization front is larger than the linear size of the cell. In this case it is valid to describe the ionization fraction as being constant throughout the cell. For cells with high optical depth, the thickness of the ionization front becomes much smaller than the size of the cell and this description is no longer valid. In this case, we interpret the ionization fraction as the fraction of the volume of the cell that is fully ionized. The remainder of the volume is considered to be neutral.

### 7.3 DIAGNOSTICS

The simulations, described in Sections 7.4 and 7.5 and summarized in Table 7.1, are compared on the basis of several diagnostics:

1. Ionization maps, Sect. 7.4.3
2. Bubble size distribution, Sect. 7.4.4
3. Mean mass- and volume-weighted ionization fraction (and their quotient), Sect. 7.4.5
4. Age of ionization, Sect. 7.4.6

We use ionization maps to gain insight in the general morphological properties of the different simulations. The resulting conclusions are necessarily of a qualitative nature. Comparing the size distributions of ionized regions is a much-applied diagnostic for the comparison of reionization simulations (Zahn et al. 2007; McQuinn et al. 2007; Mesinger & Furlanetto 2007). We have applied the widely-used bubble size estimator introduced by Zahn et al. (2007) for ease of comparison with other work (for a the detailed discussion of different method to obtain bubble sizes see Friedrich et al. 2011). The third diagnostic provides a quantitative handle on the global reionization morphology and is well-suited for the direct comparison of different simulations (Iliev et al. 2006b). We will use it to check the convergence of simulation results with resolution in the next section and it serves as our main tool to distinguish between the IO and IOM reionization scenarios. A useful diagnostic in the context of our study is the ‘age of ionization’ of different density bins introduced by FODO09. The shape of the resulting ‘age-curve’ is a very direct probe of reionization morphology and is less susceptible to noise than the quotient of the mass- and volume-weighted ionization fraction.

### 7.4 RESOLUTION STUDY

To study the effects of resolution on the morphology of reionization we perform a series of runs where we use the gas particles of the SPH/dark matter simulations described in Oppenheimer & Davé (2008) and FODO09 as the generating points of our Voronoi-Delaunay grid.

#### Number density in Voronoi cells

We assign number densities of hydrogen to the nuclei in our grid by dividing the gas-mass of gas particles through the volume of their Voronoi cell. Star forming gas has been omitted from this density field and their effect on the RT has been accounted for by adopting an escape fraction of 0.13, chosen to reach reionization at  $z=7$ . Because we use a subset of the particles, however, mass is no longer conserved and its is incorrect to assign densities by dividing the particles masses by their Voronoi volume. We have enforced mass conservation by assigning the mass of all un-sampled particles to their closest sampled particle. In order to account for the effect of density variations in the original simulation that are lost in the subset we assign local clumping factors to every particle in the subset. In Sect. 7.A we describe this procedure in more detail.



### Assigning luminosity to the sources

The simulations include, in addition to gas and dark matter particles, a third type of particle that represents gas converted into stars. In contrast to SPH particles, these *star particles* do not have a neutral hydrogen mass associated with them. In the original simulations of FODO09 they contributed a certain amount of emissivity to the cells of the regular RT grid. We have added the luminosity of each star particle to the closest gas particle, a procedure closest to the original approach of FODO09. This implies, however, that we do not aim at resolving the escaped fraction of ionizing radiation on ISM scales. The resolution of the simulations do not allow for a realistic description of the density distribution inside dark matter haloes anyway, so this approximation is the best we can do given the data. We have checked that explicitly including the star particles does not affect our results.

### Outlook

The main effect of the extra resolution will be to increase the importance of recombinations on the process of reionization. This might delay reionization and push the size distribution of ionized bubbles towards more disconnected smaller bubbles (see e.g., [McQuinn et al. 2007](#); [Choudhury et al. 2009](#); [Friedrich et al. 2011](#)).

A possible secondary effect of increased resolution is increased 'porosity' of the density field due to tunnels through the high density regions unresolved in the regular grid. This would result in radiation leaking earlier into the voids leading to a reionization scenario with a short initial inside-out stage followed by an outside-in scenario.

First we will check for the convergence of the global mass- and volume-ionized fractions for three simulations with different resolution in Sect. 7.4.1. For this purpose we ionize (parts of) a single snapshot at  $z = 6$  because at that time density contrasts are more pronounced than at higher redshift and the effect of recombinations (and thus resolution) is important in the densest structures.<sup>4</sup> Although FODO09 performed simulations with redshifts ranging from 14 to 5.5, the IOM morphology can easily be reproduced using only a single snapshot so this temporary simplification does not affect the validity of our result. We have explicitly verified this and discuss the results in Sect. 7.4.2. Subsequently we will investigate in Sect. 7.B the possible influence of (limited) resolution on the simulations presented in later sections of this Chapter.


#### 7.4.1 Convergence with resolution

To allow a resolution study to incorporate the highest possible resolution of the data we have taken a sub-volume of the data which spans 1/4 of the total box-size in every direction (so 1/64 of the total volume). The total number of particles in the volume is  $64^3$ , easily tractable on a modest desktop computer. We have taken sub-samples of the particles in this volume to get effective resolutions of 512 (original data), 256, 128 and 64 cubed. The lowest resolution box only has 512 vertices in this case.

Figure 7.2 shows the volume averaged ionization fraction and the quotient of mass- and volume averaged ionization fraction for the described sub-volume. The simulation starts with a

<sup>4</sup>Recombinations do not become more important overall because the expansion of the universe dilutes the gas on a global scale.

completely neutral box and the sources do not evolve. The simulation is run until the complete box is ionized. As mentioned in Sect. 7.4.5), the quotient of mass- and volume-ionized fraction describes the mean density of ionized regions. If the quotient is larger than unity, the morphology is IO and OI if the reverse is true. The various curves indicate different spatial resolutions and the labels refer to the equivalent resolution of the simulation *if the complete snapshot would have been used*. Except for the  $64^3$  result, all curves are consistent with each other within a few



**Figure 7.2:** *Bottom:* volume weighted ionized fractions as a function of time for three runs at  $z = 6$  at different resolution. *Top:* quotient of mass- and volume-weighted ionized fractions for the same runs. The different lines correspond to the effective resolution expressed as the cube root of the number of particles in the total simulation volume. The simulation converges for a resolution of 128-cubed particles.

percent. All curves except the  $64^3$  show a completely IO morphology. The fact that the  $64^3$  result yields an OI morphology after an initial short IO phase is some cause for worry because the regular grid used in FODO09 has exactly that resolution. This is no conclusive evidence that the morphology found in their study is due to insufficient resolution, however. We have used a small sub-volume of their data which can behave quite differently than the complete volume. We should therefore continue with a similar test using the complete volume at a resolution that shows convergence with the result including all particles. From Fig. 7.2 we can conclude (with some reserve because of the limited size of the volume) that a resolution of  $128^3$  is sufficient for this purpose.

## 7.4.2 Simulations using the total volume

From the previous section we have seen that for a sub-volume of the data, the morphology found with the SIMPLE algorithm is consistently IO, except for the lowest resolution of  $64^3$ . In this section we perform simulations to see if this result remains valid if the complete volume is used. We use a subset of the SPH particles of  $128^3$  effective resolution. We have seen in the previous section that at this resolution, the global ionization fraction is converged.

To assess whether the discretization procedure used in FODO09 is the origin of the discrepancy between our results and those of the original study, we use the regular grid used for the fiducial simulation of FODO09 as the basis of our RT grid. Because SIMPLE does not operate

on regular data, we need to employ the grid sampling techniques of Chapter 4 to translate this regular grid to a `SIMPLEX` grid.

We will compare our simulation to the results of FODO09 *and* to the simulation performed using the original SPH particles mentioned above.

### RT mesh construction from regular gridded data

To construct a RT grid from the regular data used in FODO09, we use the procedure described in Sect. 4.3.4 without adaptive resolution. The density field is stored on a regular grid with  $64^3$  cells (corresponding to  $250 \text{ comoving } h^{-1} \text{ kpc.}$ ) This regular grid is constructed by assigning the SPH particles of the original simulation to the grid cell that encompasses them (nearest grid-cell interpolation). For SPH particles that lie near cell boundaries the mass is divided between the neighbouring cells by summing incomplete gamma functions to their equivalent Plummer SPH smoothing kernels.

The `SIMPLEX` grid is constructed by randomly placing  $10^6$  points in the cells of the structured grid used by FODO09. We take more than one sample point per grid-cell of the original grid to minimize the effects of Poisson noise. We have checked that increasing the number of sample points does not change the result. The sample points obtain their densities directly from the cells (nearest grid-point interpolation). In this sense, the resolution of the `SIMPLEX` grid is maximally equal to the original regular grid. The density PDFs of the structured grid and the resulting `SimpleX` grid are shown in Fig. 7.12.


### One snapshot simulations

Although in the rest of this chapter we will consider simulations that span a range in redshifts and take account of the evolution of the ionizing sources, we first elaborate on the tests shown in Sect. 7.4.1. Instead of a resolution study with a sub-volume, we now ionize a single snapshot at  $z = 6$ . One simulation uses the regular grid sampled with  $100^3$  vertices and the other uses a  $128^3$  subset of the SPH particles.

To assess the converge of the morphology with box-size, we show the result for the total volume (solid black line) and the spread of results when subdividing the volume in 8 (light grey shaded region) and 64 sub-volumes (dark grey shaded region) in Fig. 7.3. From the light grey region, we immediately see that the `SIMPLEX` method gives consistently IO results on all 8 sub-volumes of the snapshot for both types of grids. This suggests that for this simulation volume, the morphology found with our method is converged with respect to the simulation volume. The dark grey region shows that for some of the 64 sub-volumes, the reionization morphology is decidedly consistent with the IOM scenario, however. For the regular grid, we also show the result obtained with the method used in FODO09 (dashed line). This clearly shows reversal to OI morphology after about 1075 Myr.

### Simulation overview

In order to keep track of the various simulations used in the following sections, we give each one a name where the number corresponds to whether the original data is a regular grid (series ‘1’) or SPH particle based (series ‘2’). Differences in red-shift range or RT method employed are indicated with letters.



**Figure 7.3:** Same as Fig. 7.2 but for the total volume of the simulation box. The box has been subdivided in 8 and 64 cubic sub-volumes to assess the sensitivity of the morphology to box-size. *Left:* simulation based on the regular grid used by FODO09. We have included the result obtained with the method used in FODO09 as the dashed line. *Right:* simulation based on a  $128^3$  subset of the SPH particles.

The RT simulations performed on the regular data (1A, 1B and 1C) span a range in redshift from  $z = 14$  to  $z = 5.5$ . For the original (particle based) data all snapshots before  $z=9$  have become corrupted unfortunately, so these range from  $z = 9$  to  $z = 4$ . To assess the impact of omitting the influence of sources before  $z = 9$ , we perform an additional simulation based on the regular gridded data but starting at  $z = 9$  instead of  $z = 14$ . This simulation, 1C, is in effect a consistency check between the two methods for grid construction. If 1C and 2A give consistent results, the difference in morphology cannot be attributed to the grid.

To test whether the diffusivity inherent to moment methods (see Sect. 7.5.1) is an important factor in the rate with which filaments re-ionize as observed in FODO09, we include run 1B. This simulation applies the more diffuse ballistic transport (BT) method throughout the simulation adding numerical diffusivity to the radiation field. This being inspired by the notion that moment methods are intrinsically more diffusive than ray-tracing methods. We note, however, that the *kind* of diffusivity introduced in BT may not be equal to the kind inherent to moment methods.

The corresponding SimpleX grid consists of  $10^6$  particles, randomly placed on the simulation domain. So there are on average 3.8 particles per regular cell. Using more particles does not change our results significantly because the resolution of the SimpleX grid is already higher than that of the underlying grid. Reducing the number of points below  $64^3$  would, however, necessitate interpolation from un-sampled grid-cells to the (now coarser) Voronoi grid resulting in loss of detail.

In order to study the effects of the absence of SPH snapshots prior to  $z = 9$  for the series ‘2’ runs, we have also performed run 1C which starts at this later time. Table 7.1 shows an

overview of the performed simulations and their parameters.

Identifier	Grid/SPH	res [# <sup>3</sup> ]	z-range	method
1A	grid	100	14-5.5	DCT
1B	grid	100	14-5.5	BT
1C	grid	100	9-5.5	DCT
2A	SPH	128	9-4	DCT

**Table 7.1:** Overview of simulation parameters. The second column indicates whether the SimpleX grid was constructed from a regular mesh or the SPH particles. The approximate number of nodes (nuclei of the Voronoi cells) in the grid is given in the third column. Column four indicates the range in redshift that the simulations span and the fifth column shows the RT method used where DCT and BT stand for direction consering- and ballistic transport respectively.

### 7.4.3 Ionization maps


We present ionization maps of the simulation volumes at  $x_V = 0.1, 0.3, 0.5, 0.7$  for simulations 2A and 1C in Fig. 7.4. This allows for a clean comparison of the morphology based on the distribution of matter and sources as described by the two types of grid.

The first observation is that, for the leftmost column ( $x_V = 0.1$ ), the large ionized regions are in the same positions although this is not true for the smaller bubbles. This discrepancy cannot be attributed to differences in the position of sources because they, by construction, cannot be displaced more than half a cell of the  $64^3$  regular grid of FODO09. We must therefore conclude that it is the small-scale distribution of matter resolved in run 2A that suppresses weak sources that are capable of forming ionization regions in run 1C where density is smoothed over the relatively coarse grid. It is expected that the high resolution run therefore will show more small ionization regions around single sources that have not yet merged.

The shape of the larger ionized regions in the  $x_V = 0.3$  snapshots are quite different as well, largely due to difference in the density field. Run 1C shows larger, more spherical regions owing to clustering of sources (again at the resolution of the grid) and a smoother density distribution. The bubbles in the 2A run contain more holes and tunnels and are, in general, more irregularly shaped.

As the ionized regions grow, the differences between the two runs decrease and the morphology becomes more similar. One of the major differences, hardly visible in the figure, is the presence of self-shielded clumps of high-density material in the 2A simulation that are absent in the 1C run. Again, this is a typical result of the difference in resolving power between the two grids.

Despite these differences, the overall IO progress of the ionization fronts is very similar between the runs. We do not see evidence for radiation avoiding filaments or ionization fronts reversing from voids toward regions of high-density (OI). A more quantitative description of these statements will be given in the following sections and in particular Sect. 7.4.5.




**Figure 7.4:** Contours of ionization fronts ( $x_v = 0.5$ ) of simulations 2A (turquoise line) and 1C (violet line) at volume-weighted ionization fractions  $x_v = 0.1, 0.3, 0.5, 0.7$ . The number density is plotted as background in logarithmic units.

#### 7.4.4 Bubble size distributions

We follow Zahn et al. (2007) in obtaining bubble size distributions using the method that was called the *spherical average* method by Friedrich et al. (2011). This procedure assigns a bubble-size to every cell of the grid by computing the largest sphere around that cell that is more than 90% ionized. The distribution of bubble-sizes for runs 1C and 2A are plotted in Fig. 7.5 for  $x_{\text{HII},v} = 0.1, 0.3, 0.5, 0.7$ . Although this measure is highly correlated (every ionized cell can contribute to many bubbles) it gives us a feel for the size of a typical HII-region. As expected, the high resolution 2A run shows smaller bubble-sizes at low global ionization fractions than the 1C run. When the bubbles become larger, this difference disappears completely. This confirms the qualitative observations drawn from inspection of the ionization maps in Sect. 7.4.3.

Given the relative modest size of our box, we cannot extend our analysis of bubble sizes beyond roughly  $10 \text{ Mpc h}^{-1}$  corresponding to  $x_{\text{HII},v} \sim 0.6$  (see also Fig. 7.4). The reader must therefore exercise some caution with the interpretation of the  $x_{\text{HII},v} = 0.7$  panel of Fig. 7.5. The



**Figure 7.5:** Bubble size distribution (PDF) for run 1C (solid lines) and 2A (dotted lines) at  $x_{HII,V} = 0.1, 0.3, 0.5, 0.7$ . The distribution of bubble sizes is very similar for higher ionization fraction but the 2A run shows more small bubbles at  $x_{HII,V} = 0.1, 0.3$  as a result of the more granular small scale structure of the grid.

characteristic peak of the PDF at high global ionization fractions has also been observed in other studies (e.g., [McQuinn et al. 2007](#)) and points to the dominance of one single ionized region dominating the simulation domain. Clearly, our simulations reach this point around  $x_{HII,V} \sim 0.7$  and we cannot identify sizes beyond this global ionization fraction.

#### 7.4.5 Global ionized fractions


As shown by [Iliev et al. \(2006b\)](#), the quotient of the mass- and volume-weighted ionized fraction  $x_M/x_V$  is equal to the mean density of ionized regions in units of the mean density of the universe. This quantity is a measure for the character of the reionization process. In the first stages, ionization progresses predominantly from inside high density regions outwards. The corresponding value of  $x_M/x_V$  is larger than unity because the dense clumps, where the sources reside, contribute more mass than volume when ionized. In several studies ([Gnedin 2000](#); [Ciardi et al. 2003](#); [Choudhury et al. 2009](#); [Finlator et al. 2009](#); [Petkova & Springel 2011](#)), this initial IO phase is followed by an inversion where the ionization proceeds inwards from the voids into the densest clumps that were not ionized in the initial IO part of the reionization process. This inversion results in values for  $x_M/x_V$  smaller than unity.

In [Fig. 7.6](#) we show the evolution of  $x_M/x_V$  (top panel) and  $x_V$  (bottom panel) as a function of cosmic time and redshift for a selection of the runs in [Table 7.1](#). Also included is the corresponding fiducial result from FODO09 also shown in their [Fig. \(6\)](#). Firstly, as noted in the beginning of this section, the series ‘1’ runs start at earlier times than the series ‘2’ runs. This results in the somewhat earlier ( $\Delta t \approx 100\text{Myr}$ ) onset of reionization as can be seen from the bottom panel of [Fig. 7.6](#). When the series ‘1’ runs would have started at the same (later) time as the series ‘2’ runs, the evolution of  $x_V$  at later times is hardly affected (run 1C). This suggests that the main effect that postpones reionization in the series ‘2’ runs is the higher (globally averaged) recombination rate of the adaptive simulations.

The initial values of  $x_M/x_V$  are much higher for the series ‘2’ runs. This seems logical as the higher (adaptive) resolution of the density field resolves high density gas around the sources that is smoothed in the structured grid used by FODO09. This is not the reason for the discrepancy, however, because in that case the value of  $x_M/x_V$  in the 1C run should not behave similarly to the series ‘2’ runs, as it does now. So, we conclude that the sudden onset of reionization at  $z = 9$  (because it is the first snapshot) is responsible for the observed effect. At this redshift, density peaks around sources are higher in both the series ‘1’ and ‘2’ grids and the difference in resolution apparently does not play a big role in this initial stage.

In FODO09, reionization is complete around  $z = 6$  whereas, we find that reionization extends to  $z \approx 5.5$  for the series ‘2’ runs.

Finally, we fail to see the significant dip below unity of  $x_M/x_V$  observed in Fig. (6) of FODO09 in any of our simulations. In other words, we only see evidence for IO reionization in all our runs.



**Figure 7.6:** Mass- and volume-weighted ionization fractions as indication of reionization morphology. *Bottom:* volume weighted ionized fractions as a function of time for several of our runs. *Top:* quotient of mass- and volume-weighted ionized fractions. The results from FODO09 are included as the dot-dashed lines.


#### 7.4.6 Age of reionization

Similar to Fig. (4) in FODO09 we have plotted in Fig. 7.7 the redshift of reionization as a function of over-density for a selection of the simulations as well as the fiducial run from FODO09. This provides us with a clear diagnostic for the IOM reionization scenario.

The IOM behaviour is clearly demonstrated by the result of [Finlator et al. \(2009\)](#), plotted as the dashed line in Fig. 7.7). Our runs 1A and 1B show a similar behaviour in the over-dense regions but under-dense regions tend to reionize much ( $\Delta z \approx 2$ ) later. This would suggest that in the study by FODO09 is characterized by early ionization of voids rather than late ionization of filaments.

For run 1A and 1B, the curve describing the reionization red-shift as a function of over-density has positive second derivative everywhere (except for regions where noise dominates),





**Figure 7.7:** Redshift of reionization simulations. *Left:* simulations 1A, 1B and FODO09. *Right:* simulations 2A and 1C. Note that, due to the low number of high density cells, the scatter in the high density region of the plot is due to noise rather than physical effects.

whereas the 2A and 1C runs show negative second derivatives over a large range of the plot. This is due to the fast reionization of moderately over-dense regions after the start of the simulation. Switching on the sources at a later time thus hardly changes the age of reionization of structures with overdensities lower than roughly 10. The instantaneous radiation-field hence dominates the historical evolution of the ionization field. Note that this might be different if realistic feedback is included for a coupled radiation-hydro simulation. Again, our results indicate a purely inside out evolution where the high density regions ionize first, at the onset of the simulation, and the voids last.

## 7.5 CODE COMPARISON

Now that we have established that SIMPLEX does not find an IOM reionization morphology for either the high-resolution SPH data as the gridded data used in FODO09 we are led to extend our investigation beyond the scope of a resolution study. In other words, we should check if the code presented in Finlator et al. (2009) and SIMPLEX give the same answers for simple test-problems that have, preferably, been solved with other methods as well.

A good starting point for the comparison of our code with that presented in FODO09 is the first RT comparison project (Iliev et al. 2006a) where initially 11 codes participated with several others following. We pick out the tests that are most relevant for our current goal: to get a handle on why the codes give different morphology in the EoR. The self-shielding clump test (Test 3) and the cosmological density field test (Test 4) are most likely to uncover relevant differences between the codes and we refer the reader to that paper for details.

### 7.5.1 Diffusivity and moment methods

But what should we look for in these tests? What are the characteristics of moment methods that could possibly be responsible for the observed discrepancies?


Moment methods store information of the local radiation field in 10 variables: the number density, three components of the flux and the three-by-three Eddington tensor (with only six independent entries out of nine). All directional information of the radiation field beyond the (local) flux-vector, is thus encoded in this tensor. Because of this limited capacity to store information, a radiation field dominated by high multipoles will not be captured precisely. Such a radiation field might arise from a medium with small-scale variations in opacity related to clumps and tunnels in the density field. Because high frequency components of the radiation field are not accounted for, the method is effectively blind for the small-scale structures responsible for it. These small-scale structures in turn reside predominantly around the highly clustered sources that form in regions of high density. The radiation field may thus break out of the over-dense regions more easily and stream into the voids, effectively bypassing the filaments.

### 7.5.2 Test 3


Test 3 is a cosmological shadowing test consisting of an incoming (planar) radiation field and a spherical clump of high density material. The test is constructed in such a way that the ionization front is trapped inside the clump and a neutral shadow behind the clump is formed. The test can be used as a diagnostic for both the correct ionization chemistry which traps the ionization front in the prescribed position and the ability of a RT method to produce a sharp shadow behind a self-shielding object. For a more detailed description of the test we refer the reader to (Iliev et al. 2006a).

Because both the moment method used in FODO09 and SIMPLEX do not naturally incorporate planar waves, we have altered the setup slightly by using a single point-source to emit the ionizing photons. We have chosen the strength of the source such that the flux at the dense obstacle equals that of the planar wave of the original setup.

Figure 7.8 shows a cut through the box at (from left to right) 1, 3 and 15 myr. The ionization front is trapped halfway in the clump as required by the test (see Iliev et al. 2006a) although, due to the irregular nature of the SIMPLEX grid, the result is not completely symmetrical. The shadow cast behind the dense clump is stable after about 7 myr and is of the size expected from the angular resolution used for this test (solid angles of  $4\pi/84$ , or equivalently, a linear angle of about 23 as indicated by the angle shown in the upper left corner of the leftmost panel). The method used in FODO09 shows a different behaviour in this test. The result of the shadowing test (which is the same as described in this section) which appeared in the method paper Finlator et al. (2009) (their Fig. 14) is reproduced in Fig. 7.9. We have performed the shadowing test again with the more diffusive ballistic transport method throughout the simulation volume (whereas it actually should be used exclusively in optically thick regions where radiation is absorbed in a few steps). As can be seen from Fig. 7.10, a more diffusive mode of transport results in total absence of a shadow behind the obstacle. Moreover, the obstacle is completely ionized long before the end of the simulation time of 15 Myr. Such diffusivity is not likely to produce a reversal to OI reionization, however, because it tends to ionize dense structures more easily




**Figure 7.8:** Neutral fraction of hydrogen for the shadowing test described in the text at (from left to right) 1, 3 and 15 Myr. In the left panel, the theoretical limit for the shadow and the source are indicated with white lines and a white dot respectively. The white disk with the black lines shows the angular resolution of the method for this test (opening angles of  $4\pi/84$ ).



**Figure 7.9:** Result for the shadowing test taken from [Finlator et al. \(2009\)](#). The initial shadow disappears almost completely to the end of the simulation time due to diffusion. Interestingly, the ionization-front hardly penetrates the dense obstacle in conflict with the results of ([Iliev et al. 2006a](#)) and those presented in this section.

and much sooner. Moreover, we have also used the diffuse BT mode of transport throughout the simulation volume for simulation run 1B which shows a purely IO morphology. We have checked that application of diffuse transport (which is maximally diffusive, see Sect. 2.1.4) instead of BT does not change this result.

There is an interesting difference between the results shown in Fig. 7.9 and those in Fig. 7.10. In the first, the ionization front hardly penetrates the dense obstacle whereas in the second, the



**Figure 7.10:** As in Fig. 7.8 but for a simulation with the more diffuse ballistic transport method applied throughout the grid.

obstacle is completely ionized. This behaviour may be the first hint to a solution of the issue of this chapter. If the moment method applied in FODO09 leaves over-dense obstacles intact while ionizing the under-dense regions around them, late ionization of filaments is easily explained.


### 7.5.3 Test 4

The second test that we subject both methods to is that of multiple sources in a cosmological density field, or Test 4 from (Iliev et al. 2006a). This test is closest to the cosmological reionization simulations presented in FODO09. The simulation volume is much smaller, however,  $0.5h^{-1}$  versus  $16h^{-1}$  comoving Mpc and contains only 16 sources (whereas the data used in this section contains of tens of thousands of sources for the later snapshots). The simulation time is 0.4 Myr, which is enough to ionize the box for 80% to 90% in volume.

In Fig. 7.11 we show the quotient of mass- and volume weighted ionized fraction,  $x_m/x_v$ , as a function of time. Again, if this quotient is smaller than unity, the mean over-density of ionized regions is smaller than unity and the reionization morphology is OI. As a reference, we also show curves for two well-established codes that participated in the comparison project of (Iliev et al. 2006a), C<sup>2</sup>-Ray (Mellema et al. 2006) and FTTE (Razoumov & Cardall 2005).

Clearly, the SimpleX result shows a significant dip below unity for the quotient  $x_m/x_v$ , in accordance with results from the other two codes albeit at a slightly later time. We thus conclude that SimpleX agrees excellently with the other codes shown here when it comes to discrimination between IO and OI morphology as diagnosed by  $x_m/x_v$ .

The method used in FODO09, indicated with the label ‘Finlator’, shows the same general behaviour but completely ionizes the box within the simulation time. The dip below unity of  $x_m/x_v$  happens at a slightly earlier time than for the other codes. There is no clear tendency for the method towards a reversal to OI at an earlier point than the other codes nor is the dip more pronounced. If anything, the opposite is true; the other codes show values of  $x_m/x_v$  as low as 0.96, whereas the ‘Finlator’ result stays always above 0.98. The method used in FODO09 does not show a more pronounced OI morphology than the other codes in this test. Test 4 therefore does not further our understanding of the observed discrepancy other than to show that the reversal to OI morphology is not a consistent property of the moment method.



**Figure 7.11:** *Bottom:* volume weighted ionized fractions as a function of time for three codes ( $C^2$ -Ray, FTTE and SimpleX) from (Iliev et al. 2006a). *Top:* quotient of mass- and volume-weighted ionized fractions for the same runs.

#### 7.5.4 Influence of box-size and source-prescription

On scales smaller than approximately 30 comoving Mpc/h, the effect of cosmic variance on the morphology can no longer be neglected. If box-sizes are smaller than this value, the simulation can become biased towards an OI morphology because high-mass sources are under-represented in small volumes allowing dense structures to become self-shielding. Also, if the volume is dominated by a single dense cluster, the reionization morphology will be biased towards OI after an short initial IO period. The reason for this is the high recombination rate of the over-dense gas in the cluster which allows it to absorb many more photons per baryon before reionizing compared to non-cluster gas (see e.g., Ciardi et al. 2003, who use a 10/h comoving Mpc box). Similar behaviour is seen in Gnedin (2004) where a simulation volume spanning 4/h comoving Mpc is used.

Although our simulation volume is significantly larger than used in these early studies, it is still a factor of two smaller than the 30/h comoving Mpc required to exclude all effects of cosmic variance (except for the influence of quasars which needs volumes larger than at least 100/h comoving Mpc). We *can* rule out that the IOM scenario seen by FODO09 is a result of cosmic variance, however, because we are not able to reproduce the OI part of the reionization epoch observed in FODO09, even when using boxes of 1/8 the total volume.

Another important variable in the characteristics of reionization morphology is the distribution and mass-luminosity dependence of sources (e.g., McQuinn et al. 2007). It is therefore not far-fetched to ascribe the differences in reionization scenario to the source prescription adopted by FODO09 which deviates from that practiced by the bulk of reionization studies. However interesting this subject is, we feel that it is irrelevant for this work because we fail to recover the IOM scenario *adopting the same sources as in FODO09*.

Based on the above discussion, we argue that both box-size and source prescription are not important for the results presented here. The underlying reason is that we have performed a

*relative* study where only differences with the results of FODO09 are important. With respect to the box-size, we argue that the IO result obtained with SIMPLEX is robust.

## 7.6 SUMMARY


- We have performed two types of RT simulations based on SPH simulations described in FODO09. The first type uses the regular grid as the basis for the unstructured Voronoi-Delaunay grid, the second uses the underlying SPH particles directly. These two types of simulations allow us to disentangle the effects of RT method and resolution.
- We have not observed the IOM morphology described in FODO09 and observed in [Petkova & Springel \(2011\)](#) and [Choudhury et al. \(2009\)](#). Contrarily, we have observed a purely inside-out evolution of reionization.
- The extra resolution in the density field of the series ‘2’ runs, leads to smaller ionized bubble sizes at low (0.1-0.3) global ionization fractions and a delayed time of reionization.
- Our results are robust to resolution and type of grid construction (grid or SPH based) pointing to the difference in RT methods as the source of the different morphology.
- We have compared the method used in FODO09 with SIMPLEX using several standard test-problems. This lead us to conclude that the diffusivity often attributed to moment methods is not likely to be the source of the observed discrepancy in morphology of reionization.
- The most plausible reason for the ‘late ionization of filaments’ observed in FODO09 is the behaviour seen in Sect. 7.5.2. For the results obtained with the moment method, the shadow behind the dense obstacle is eaten away but the obstacle itself remains almost neutral.
- Because we do not retrieve the IOM morphology, we argue that a comprehensive study of the effect of source type and distribution is beyond the scope of this work.

## 7.A EFFECTS OF TAKING A SUBSET OF SPH PARTICLES

Due to the large intrinsic computational cost of radiative transfer, it is seldom feasible to include all particles from an SPH/DM simulation in the RT simulation. Therefore, it is necessary to choose a subset of the particles as the basis for the RT grid. As discussed in Chapter 4 care must be taken of proper representation of the original density field and conservation of mass. Sect. 4.3.4 and Sect. 4.4.2 of that chapter discuss some of these issues applied to the data of this chapter. Here we present a small extension of these sections by showing the volume-weighted Probability Density Functions (PDFs) for the number density in the various grids used in this study in Fig. 7.12. The Voronoi grids constructed from a subset of the SPH particles are seen to under-represent the highest density regions because of the lower resolution. The mass that is not represented at the highest densities is shifted to intermediate densities shifting the peak of the function rightwards. This effect is obviously more pronounced for the  $64^3$  simulation than for the  $128^3$  simulation. In the latter, there exists some volume at densities lower than present in the original data. This is due to the fact that the volumes of Voronoi cells can be larger than

the corresponding SPH volume. This is even true if all SPH particles would be included in the triangulation.

The Voronoi grid consisting of  $10^6$  particles (indicated as Vor 1e6) can be seen to represent the structured density field almost perfectly. This is due to the modest dynamic range of the original grid and the fact that the resolution of the SimpleX grid is higher than that of the grid everywhere.



**Figure 7.12:** Probability density functions of the raw SPH data (All) and the Voronoi grid resulting from taking a random subset of  $64^3$  SPH particles from the raw data (Vor 64). Also shown are the PDFs for the  $64^3$  grid used in FODO09 (Fin 64) and the resulting Voronoi grid (Vor 1e6), both shifted down for clarity.

### 7.A.1 Including clumping

As far as ionizing radiation is concerned, (local) mass conservation alone suffices to guarantee the correct number of ionizations on the resolution of the subset. The recombination rate, however, depends on the square of the density and will be systematically under-estimated due to the 'lumping together' of gas particles into a cell with their mean density.

To correct for this lack of sub-cell variations in the density, we assign to every cell a clumping factor based on the distribution of un-sampled particles in that cell. Clumping factors have been widely adopted in recent simulations of cosmological reionization. Sometimes to add structure that was not present in the original simulation (e.g., [McQuinn et al. 2007](#)), or to compensate for the de-refinement of the original simulation to allow for radiative transfer calculations ([Kohler et al. 2007](#), e.g., and references therein). Most studies adopt a global clumping factor based on the neutral gas in the total simulation volume, but this generally yields over-estimates of the clumping for two main reasons: heating ([Pawlik et al. 2009](#)) and the non-locality of clumping ([Raičević & Theuns 2011](#)).

In this work, we compute local clumping factors based on the sets of un-sampled particles assigned to the subset particles implemented as a combination of masses and densities given by the last step in the following equation

$$f_c = \frac{\langle \rho^2 \rangle}{\langle \rho \rangle^2} = \frac{\frac{\sum \rho^2 V}{\sum V}}{\left(\frac{\sum \rho V}{\sum V}\right)^2} = \frac{\frac{\sum \rho m}{\sum m/\rho}}{\left(\frac{\sum m}{\sum m/\rho}\right)^2} = \frac{\sum \rho m \sum m/\rho}{M^2}. \quad (7.1)$$

Here we have used that the volume, mass and density of a particle are related by  $m = \rho V$  and we have defined  $M \equiv \sum m$ . For notational clarity the indices are omitted from the above equation.

## 7.B POSSIBLE EFFECTS DUE TO LACK OF RESOLUTION

As shown in Chapter 3 and Chapter 4, gradients in the number density of Voronoi nuclei can give rise to undesirable systematic effects in the transport of ionizing radiation in SimpleX. The four systematic effects induced by gradients in the point distribution as identified in Chapter 3 are drift and clustering of diffuse radiation, ballistic deflection and ballistic de-collimation. The first two effects have relevance for diffuse radiation only. In this work we do not include diffuse recombination radiation although the SimpleX method can do this without added computational effort. The effect coined ballistic deflection leads to radiation (on average) deflecting away from the original direction into regions of lower number density (of particles). The second effect, ballistic de-collimation, describes the general loss of direction of radiation. This leads to stalling of the ionization-front position and loss of shadows behind dense irradiated structures.

In this section we assess to what extent the simulation in this chapter are affected by either ballistic decollimation or deflection. To this end we study in detail a region where possible effects of point density gradients would be most prominent.

### A region with possible issues

From the upper right in-set of Fig. 7.1 it is clear that the approach of random selecting a subset of SPH particles results in a point-distribution with considerable gradients. We select a region of  $2.4 h^{-1}$  co-moving Mpc around the dense knot shown in the upper right inset of Fig. 7.1 and re-simulate the RT on this small sub-volume in two different ways. Once including all particles of the original simulation and once with the same subset as used in the 2A run. The point distribution based on all SPH particles is not expected to give rise to strong gradients because of the regularizing nature of the SPH method.

In Fig. 7.13 the density field is plotted logarithmically for the low- and high-resolution simulation. Although the low-res subset contains a factor 512 less particles than the original data, the most prominent structures are reproduced. If we were to see effects of either ballistic deflection or de-collimation, they would show up in this sub-volume of the grid.


### Results

It has been demonstrated that the introduction of direction conserving transport (DTC) cures both ballistic de-collimation and deflection in the idealized situations where the effects are well observable (Kruip et al. 2010; Paardekooper et al. 2010). In a more complex example as presented here, the effects are harder to analyse. But we can look for three distinct features when comparing positions of the ionization front of the two simulations (see Fig. 7.14). The lower resolution simulation may suffer from

- Stalling of the ionization front due to de-collimation
- Loss of shadowing due to de-collimation
- Bypassing of filaments due to deflection

When we inspect the ionization front positions in Fig. 7.14 we see that, apart from the difference in resolution, the results are surprisingly similar. The ionization front in the low resolution simulation may be somewhat advanced with respect to the high resolution result






**Figure 7.13:** Logarithmic plot of the density structure of the sub-volume described in the text. *Left:* low resolution subset. *Right:* field including all particles.

(never more than 20%) but there is no evidence of stalling. We have checked that the box is fully ionized after 45(50) Myr for the low(high) resolution simulation.

This cumulative retardation of reionization for the higher resolution is most certainly due to the fact that it resolves all density fluctuations and therefore has a global recombination factor which is higher than that of the low resolution simulation. Note that we incorporated a local clumping factor in the low-res simulation to compensate for this effect but our results suggest that a clumping factor is not a replacement for resolution in this case.

We see no evidence for the bypassing of filaments at any time but the ionization front has fewer features for the low resolution case. This is expected because the densest structures are largely unresolved here. If the observed loss of shadowing must be (partially) attributed to decollimation is not clear, however. We argue that the effect is rather moderate considering the enormous difference in resolution and that it will not likely influence our global results significantly. To check this statement, we have used the more diffusive ballistic transport throughout the whole simulation (run 1B, see Tab. 7.1) where we normally employ DCT when the cells become optically thin to prevent any spurious diffusion of the radiation field. Ballistic transport is susceptible to both decollimation and deflection and, if present, these effects should be visible in the aforementioned runs.



**Figure 7.14:** Ionization-front of hydrogen in a sub-volume of the simulation domain containing a dense clump that might give rise to artifacts in the RT. Cuts are made halfway the  $2.4h^{-1}$  Mpc box in the  $z$ -direction. From top to bottom, time runs from 5 to 20 Myr in 5 Myr intervals. Thick solid contours are for the subset of particles and the thin dashed contours are for the simulations including all particles.



## CHAPTER 8

---

# Mass and chemical composition of the Homunculus

C.J.H. Kruip<sup>1</sup>, P. W. Morris<sup>2</sup>, A. de Koter<sup>34</sup>, M. Min<sup>34</sup>,  
L.B.F.M. Waters<sup>35</sup> and S. Hony<sup>6</sup>

**T**HE Luminous Blue Variable star  $\eta$  Carinae experienced an outburst in the 19th century creating the Homunculus nebula seen today. We aim to derive the chemical composition and mass of the dust in the Homunculus nebula, and to constrain the kinetic energy liberated in the Great Eruption. We have fitted a model describing optically thin dust at three temperatures to the infrared spectrum from 2 to 170  $\mu\text{m}$ . Our fits describe several compositional scenarios, each with its own dust composition and mass. The total mass of the Homunculus corresponding to these scenarios ranges from 45 and 75  $M_{\odot}$ .

---

<sup>1</sup>Leiden University, Niels Bohrweg 2, 2333 CA Leiden, Leiden, the Netherlands

<sup>2</sup>California Institute of Technology, 1200 East California Boulevard, Pasadena California, US

<sup>3</sup>Astronomical Institute 'Anton Pannekoek', University of Amsterdam, Kruislaan 403, 1098 SJ, Amsterdam, the Netherlands

<sup>4</sup>Astronomical Institute, Utrecht University, Princetonplein 5, 3584 CC Utrecht, the Netherlands

<sup>5</sup>Instituut voor Sterrenkunde, Katholieke Universiteit Leuven, Celestijnenlaan 200B, B-3001 Leuven, Belgium

<sup>6</sup>Laboratoire AIM, CEA/DSM - CNRS - University Paris Diderot, DAPNIA/SAP, 91191 Gif-sur-Yvette, France

## 8.1 INTRODUCTION

The enigmatic object  $\eta$  Carinae is one of the most luminous sources in the sky in infrared (IR) (Westphal & Neugebauer 1969; Aitken et al. 1995). Almost all of the  $5.0 \times 10^6 L_{\odot}$  is emitted at IR wavelengths, indicating that the light we observe is reprocessed by dust that is irradiated by the central source (Cox et al. 1995). The bulk of this dust is contained in the bipolar ‘Homunculus’ nebula which resulted from a 20-year period of mass loss and extreme brightening that occurred in the 19<sup>th</sup> century and is known as the ‘Great Eruption’.

The mechanism underlying the Great Eruption is under debate for many decades (e.g. Davidson & Humphreys 1997; Owocki et al. 2004). An important quantity that would help identify the mechanism is the total kinetic energy of the Homunculus. In order to derive this quantity one has to know the mass of the ejected material as a function of velocity. Whereas the kinematics are well known (see e.g. Currie et al. 1996; Morse et al. 2001; Smith 2006), the absolute value for the total mass in gas and dust is uncertain. The amount of gas is difficult to measure directly because the Homunculus is mainly a reflection nebula at UV and optical wavelengths (see e.g. Meaburn et al. 1993).

At mid- and far-infrared wavelengths, the Homunculus is approximately optically thin (Davidson & Humphreys 1997) with the exception of the power law component between 2 and 8  $\mu\text{m}$  which is characteristic of optically thick material concentrated around the unresolved source. This property of the nebula makes it possible to assess the amount of dust in the whole nebula from the observed IR luminosity. The total mass can subsequently be derived from the measured dust mass using an appropriate value for the gas-to-dust ratio.

The total dust mass of the ejected material has been determined from ground-based near- and mid-infrared observations (e.g. Smith et al. 1998, 2003b); the spectral energy distribution (SED) obtained with *Infrared Space Observatory* (ISO) (Morris et al. 1999) and flux points at far-infrared up to centimeter wavelengths (e.g. Mitchell & Robinson 1978; Hackwell et al. 1986; Cox et al. 1995; Gomez (Née Morgan) et al. 2006). These attempts have yielded values for the dust mass ranging from 0.02  $M_{\odot}$  up to 0.7  $M_{\odot}$ , i.e. they show a range of a factor 35.

All dust studies listed above make simplifying assumptions on the chemical composition of the dust. This is immediately clear when the fits to the SED are scrutinized: they do not reproduce the spectral signatures. In this study we aim, for the first time, to constrain the composition of the Homunculus by fitting the detailed individual features of the spatially unresolved spectrum over a large range of infrared wavelengths (from 2 to 170  $\mu\text{m}$ ). We present solutions that reproduce the SED in great detail and analyze these solutions in terms of quality of fit and dust formation considerations. We discuss the gas-to-dust ratio and the total kinetic energy given to the Homunculus material during the Great Eruption.

In Sect. 8.2 the observations are discussed. Section 8.3 outlines the model and general strategy. Results are presented in Sect. 8.4 and discussed in Sect. 8.5.

## 8.2 THE INFRARED SPECTRUM OF $\eta$ CARINAE

### 8.2.1 Observational details

The infrared SED as observed with the short wavelength spectrometer (SWS) and the long wavelength spectrometer (LWS) on board ISO (Kessler et al. 1996) was published by Morris et al. (1999).

SWS covers the wavelength range 2.38–45.2  $\mu\text{m}$  at spectral resolution  $R \equiv \lambda/\Delta\lambda \approx 1000$ –2000; LWS covers wavelengths from 43.0 up to 197.5  $\mu\text{m}$  with  $R \approx 150$ –200. The aperture used for the SWS instrument consists of several windows which are shown in Fig. 8.1. In contrast to what one could assume from the image (which simply shows aperture overlays), the instrument response at each wavelength range is not spatially flat in the rectangular apertures. In fact in the cross-dispersed dimension, which is along each apertures semi-axis and coincidentally parallel to the  $\eta$  Carinae bipolar axis, the spatial response is decidedly non-uniform, variously peaked, and asymmetric (non-Gaussian). This holds to a lesser degree for the dispersion direction (slit major axis) as well but at least peak responses are better centered there.


The strange response curves are due to an optical design flaw in the instrument (Beintema et al. 2003). The aperture for the LWS is broader than that of the SWS, (20'' instead of 14'') and consequently, contains more of the cool lobes. This still holds if this simple geometric statement is combined with orientation and shapes of each of the response functions in both directions. If these portions would contribute significantly to the total flux, one would see a sharp jump in the SED. Such jumps were commonly observed in SWS spectra of extended sources such as PNe like NGC 7027, galactic center sight-lines, etc., but not the case for  $\eta$  Car.  $\eta$  Carinae on the other hand calibrates like a point source after removal of hysteresis and saturation effects, and matches surprisingly well with the LWS spectrum (to within each instruments photometric uncertainties of 25% for SWS and 30% for LWS at the 45–50  $\mu\text{m}$  frontier).

This does not mean, however, that there is not some contribution from the lobes. At the point where the SWS beam response functions are falling off, the mid-IR imaging clearly reveals the lobes so there is thermal and line emission through each passband. The SWS and LWS spectrographs are divided in a number of sub-bands (Clegg & The LWS Consortium. 1999). At the boundaries of these sub-bands the shape of the spectrum may be affected due to calibration issues. Flux jumps at sub-band edges are typically of the order of a few percent. Instrumental artifacts may also play a role, notably at the edges of band 3E (from 27.5 to 29  $\mu\text{m}$ ) in the LW section.


### 8.2.2 Features in the spectrum

The continuum subtracted spectrum shown in Fig. 8.2 shows several striking features. The near-IR part of the spectrum (from about 2.3 to 7  $\mu\text{m}$ ) contains spectral lines originating in the present-day stellar wind of  $\eta$  Carinae in addition to some line emission from the nebula, in particular Fe lines. For a discussion of the central source and its wind the reader is referred to Hillier et al. (2001); Smith et al. (2003a); van Boekel et al. (2003); Weigelt et al. (2007).

The 10  $\mu\text{m}$  feature is shown in the top panel of Fig. 8.2. The feature is peculiar in the sense that it shows a ‘shoulder’ at long wavelengths extending to  $\sim 15.5 \mu\text{m}$ . Several studies



**Figure 8.1:** The position of the aperture for the SWS and LWS parts of the ISO data obtained and reduced by [Morris et al. \(1999\)](#) and used in this chapter. The dimensions (in arcseconds projected on the sky) of the sub-apertures and the wavelength intervals for which they are used are indicated. The figure is simplified in the sense that the response of the instrument is not spatially flat in the depicted rectangular apertures (see text for details). Note that beyond  $27 \mu\text{m}$  a larger part of the Homunculus is included in the aperture (see text for discussion). The arrows in the image indicate the supposed massive torus as described in [Morris et al. \(1999\)](#).



**Figure 8.2:** Continuum subtracted ISO SED of Eta Car. The subtracted continuum is a combination of four blackbodies at 110, 190, 300 and 610 K which are multiplied with a power-law of the form  $\nu^{1.22}$ . The top panel shows the region with clear contributions from the star and its wind and the  $10 \mu\text{m}$  feature. The bottom panel shows the wavelength-range containing the  $20.23\text{-}\mu\text{m}$  feature and  $24\text{-}35\mu\text{m}$  complex. The stellar part of the spectrum shows strong blanketing because of the strong continuum due to hot dust.

report that the feature broadens with distance to the central star, which is a direct indication of varying dust properties throughout the Homunculus. This has been modelled by variations in dust compositions with location ([Mitchell & Robinson 1978](#)) and by increasing the grain

size with distance from the star (Mitchell & Robinson 1978; Robinson et al. 1987). The model used by these authors comprises two components, silicates and a featureless component such as metallic iron. More recently, Chesneau et al. (2005) (see also de Koter et al. 2005) fitted the feature convincingly with amorphous corundum ( $\text{Al}_2\text{O}_3$ ) and silicates. Assuming a continuous distribution of ellipsoidal (CDE) particles, a mixture of amorphous corundum and olivine was superimposed on a smooth continuum. This yielded excellent fits. The use of CDE particles removes the need for very large grains ( $\geq 2\mu\text{m}$ ) required by e.g. Robinson et al. to match the observed slope of the red side of the feature.

At  $20.23\mu\text{m}$  (see bottom panel Fig. 8.2), the SED shows a strong and very sharp symmetrical peak. A sharp feature peaking at  $20.1 \pm 0.1\mu\text{m}$  – dubbed the ‘ $21\mu\text{m}$  feature’ – is also found to arise from the ejecta around a small number of carbon-rich evolved stars (Kwok et al. 1989). A comparison of the feature found in  $\eta$  Carinae and in the C-rich ejecta of the post Asymptotic Giant Branch (AGB) object HD 56126 attributed to TiC reveal significant differences. It is much narrower (we measure a full width at half maximum of somewhat less than  $1.0\mu\text{m}$  vs  $3.0\mu\text{m}$  in the C-rich source) and symmetric in  $\eta$  Carinae while the  $21\mu\text{m}$  feature in HD 56126 (Hony et al. 2003) is distinctly asymmetric with a prominent red wing. Hony et al. (2003) argue that most likely TiC is not responsible for the feature seen in HD 56126. Likely the carrier of the feature in  $\eta$  Carinae is not the same as the carrier of the  $21\mu\text{m}$  feature in C-rich evolved stars, although we can of course not rule out that the carriers may be chemically related.

The spectrum shows a broad feature starting with a steep rise at about  $24\mu\text{m}$ , followed by a ‘plateau’ from  $\sim 25$  to  $31\mu\text{m}$ . A rich spectrum of substructure is seen on top of this plateau. The feature ends with a more or less linear decline from roughly  $31$  to  $35\mu\text{m}$ . A feature with a similar profile and width at roughly the same position is observed in carbon-rich AGB sources and has been accounted for with MgS by Hony et al. (2002).

## 8.3 METHOD

The spatial distribution of the dust in the homunculus of  $\eta$  Carinae is complicated and dynamic, possibly involving a (disrupted) torus (Morris et al. 1999; Smith et al. 2002). This complex geometry will present a challenge to any chemical analysis of the dust that includes a realistic treatment of dust optical depth at different wavelengths in a single integrated spectrum. As a first approach, however, it is natural to start with a model that assumes the dust to be optically thin.

Limb brightening effects, clearly present in images at  $\lambda = 4.8, 8.8, 10.3, 12.5, 18.0$  and  $24.5\mu\text{m}$  (see Smith et al. 2002, 2003b), suggest that at mid- and far-infrared wavelengths, this approximation is valid. The spatial distribution and solid state properties of the dust are then the prime constituents of the model. In this section we discuss our optically thin dust model.

### 8.3.1 A three temperature model

In an optically thin medium the temperature of a given dust species is determined by its distance from the central source. If the dust material is about equally distributed over distance, one would find a continuum of temperatures. However, in the case of  $\eta$  Carinae, we find that the



thermal infrared spectrum is well represented by only three characteristic temperatures, suggesting in turn (three) main locations of the dust. That this is indeed the case appears to be confirmed by imaging observations (e.g. [Smith et al. 2003b](#)). We have taken the three dimensional Homunculus model proposed by [Smith \(2006\)](#) and the temperature law given in Eq. 3 of [Smith et al. \(2003b\)](#) to find such a description. Inspection of the global morphology of the Homunculus as given in e.g. [Smith \(2006\)](#), suggests that it consists of a ‘hot’ inner core, a ‘cool’ region comprising the Homunculus ‘exterior’ and polar caps, and a ‘warm’ interior of the lobes.

We note that this subdivision in geometrically separated regions is in conflict with the findings of [Morris et al. \(1999\)](#) who conclude that most of the (dust) mass is located in a torus close to the central source. If much of the contribution to the IR flux would be due to the lobes, a jump in the SED should occur when switching between the SWS and LWS apertures (see Fig. 8.1). The absence of such a jump suggests thus that the lobes do not have a significant contribution to the total IR flux.

This puzzling contradiction between the findings of [Morris et al. \(1999\)](#) and, for instance, those of [Smith \(2006\)](#) remains unresolved. This is the main motivation to make our model independent of geometrical effects. It is unclear how to account for the extra geometrical constraints on the SED fitting in a way that does not jeopardize the objectivity of our results and we choose to leave this to future work.

We find that adding a continuum of temperatures does not improve the fit significantly. Two-component models, however, fail to account for the short wavelength part (2-7 $\mu$ m) of the SED. Moreover, models based on more than three temperature components increase the degeneracy of the fitting problem (allowing for several solutions with different parameters but very similar chi-squared values). These findings is consistent with the empirical approach taken by [Morris et al. \(1999\)](#), who use a three component fit based on an optically thin treatment of the dust.

### 8.3.2 The size and shape of the dust grains

To describe the shape of the dust grains, we use a Continuous Distribution of Ellipsoids (see e.g. [Min et al. 2003](#)). The actual shape distribution will not affect our findings much – as long as it breaks the perfect symmetry of homogeneous spheres ([Min et al. 2005](#)) – because it has little influence on the emission efficiency of silicates and carbonaceous particles (and therefore the derived mass). Furthermore, we assume the Rayleigh limit to be valid, i.e. the grain sizes are much smaller than the wavelength of the incident photons. To test this last assumption we have performed fits with grains of 2 micron in size that are represented by hollow spheres (see [Min et al. 2005](#)). Such large grains have been proposed to be present in the Homunculus in order to account for the observed grey extinction towards the central source ([Davidson & Humphreys 1997](#)). The derived dust mass is found to increase some tens of percents when such grains are used. We conclude that the assumptions on particle size and shape have only a modest influence on the derived mass.

### 8.3.3 Fitting procedure

The flux,  $F_\nu$ , due to an optically thin cloud of dust at temperatures  $T_i$  ( $i = 1, 2, 3$ ) is given by

$$F_\nu = \sum_{i=1}^3 \sum_{j=1}^N \frac{\kappa'_{\nu,j} B_\nu(T_i) M_{i,j}}{D^2}. \quad (8.1)$$

where  $\kappa'_{\nu,j}$  [ $\text{cm}^2\text{g}^{-1}$ ] is the mass absorption coefficient of the  $j^{\text{th}}$  material and  $B_\nu(T_i)$  is the Planck function at the  $i^{\text{th}}$  temperature.  $D$  is the distance to the source and  $M_{i,j}$  is the mass of the  $j^{\text{th}}$  material at the  $i^{\text{th}}$  temperature. The set of equations (one for every flux point) of this form is solved with the DLSEI routine from the SLATEC library (Vandevender & Haskell 1982) to obtain the weights  $M_{i,j}$ . Note that this approach leaves room for inhomogeneity of the medium, i.e. the temperature components can have different composition. The three temperatures are also treated as free parameters and are solved for with an independent minimisation routine.

In order to find the dust species for our fits we take a comprehensive approach. We apply a library of optical constants of about 100 individual measurements of solid state materials. It comprises the common species of both carbon and oxygen dominated mineralogy. The most important oxides, silicates and sulfides are included along with carbonaceous species and metallic iron. A table including all dust types together with references to the original laboratory measurements is provided in the on-line material.

These species are fitted one-by-one to the SED and those that yield the best fit are selected. To assure an equal weighing we resampled the observed ISO spectrum using  $\Delta \log \lambda$  [nm] = 0.23. As a measure for the quality of fit we use  $q \equiv \sqrt{\chi^2}$ . This quantity is not normalized; the error in each wavelength bin is assumed to follow a Poisson distribution. A species that results in a quality of fit  $q$  is selected if the following criterion is satisfied:

$$\frac{q - q_{\min}}{\bar{q}} < 0.05 \quad (8.2)$$

where  $q_{\min}$  is the smallest  $q$  that occurs and  $\bar{q}$  is the mean over the whole library. The adopted relative difference of 5 percent is chosen empirically. It is taken sufficiently large to avoid a plethora of selected species, but small enough to assure a detailed fit to the spectrum. The materials selected form the basis for a new fit. Every material in the library is added one-by-one to this basis and selected if Eq. (8.2) holds. For several dust species the library includes measurements at different conditions (e.g. density and temperature) yielding a ‘family’ with very similar optical properties. If such a species satisfies Eq. (8.2) the whole family in general does. We only include the best candidate in this case.

From all fits computed we also identify the one with the lowest total dust mass. We discuss this fit in Sect. 8.4.3.

#### Degeneracy of the fitting procedure

The relatively smooth nature of the SED increases the chance of ‘over-fitting’ resulting in a set of solutions that are statistically indistinguishable and thus lack any discriminative power. We check for degeneracy of our fits by applying a ‘bootstrapping procedure’ (Efron & Tibshiani 1986). Every fit is repeated a hundred times with a different random selection of the data-points.

From the resulting distribution of  $\chi^2$  results we can, under the assumption of independence of the individual data-points, infer the statistical significance of each fit. We discuss the degeneracy of our fit results in Sect. 8.4.1.

### 8.3.4 The stellar spectrum

In the near-IR, the ISO spectrum shows spectral lines originating in the present day stellar wind of  $\eta$  Carinae itself (e.g. Hillier et al. 2001). Also, in the sub-millimeter to centimeter regime the thermal free-free emission of this wind is observed. We therefore include the SED of the stellar component of  $\eta$  Carinae, as computed by Hillier et al. (2001), in our model. They account for non-LTE line blanketing which is especially important given the dense forest of Fe lines in  $\eta$  Carinae's wind. The Hillier et al. (2001) model therefore provides a more accurate description of the stellar source than, for instance, an LTE Kurucz model at corresponding  $T_{\text{eff}}$  and gravity.<sup>7</sup> Apart from these advantages, these authors explain that the near-IR flux of their model (having a luminosity equal to the IR luminosity of the Homunculus) is too large compared to observations. In order to match the near-IR flux they adopt a fictitious distance of  $\sim 6$  kpc. We adopt the canonical distance to  $\eta$  Carinae of 2.3 kpc, implying that we need to resolve the discrepancy in an alternative way. We do so by introducing - in addition to the reddening by the dust in our model with a Cardelli et al. (1989) reddening law - a grey extinction component. The nature of this component remains unknown but may point, as mentioned before, to very large (several  $\mu\text{m}$ ) grain sizes in the Homunculus.

### 8.3.5 The unidentified sharp feature at $20.23 \mu\text{m}$


In the ISO spectrum a strong and sharp feature occurs at approximately  $20.23 \mu\text{m}$ . This feature is too sharp and symmetrical to be attributed to any amorphous or crystalline solid state material that is observed in astronomical environments. We therefore model it with an array of identical, damped, harmonic oscillators, i.e. a Lorentz model for a crystalline material (see e.g. Bohren & Huffman 1984). We do not assign a continuum to this component which contributes less than a percent to the total mass in all our fits.

### 8.3.6 A measure for elemental abundance

We calculate the elemental abundance of the solid state material relative to solar abundances (adopting Allen 1973), by dividing the total mass of element  $i$  in a fit,  $m_i$ , by the sum of the masses of all the elements in that fit. This procedure is also applied to the same element in the sun and the fractions are divided to give the relative number. The relative mass abundance for the  $i^{\text{th}}$  element,  $\Phi_i$ , is thus given by

$$\Phi_i \equiv \frac{m_i / \sum_i m_i}{m_{\odot,i} / \sum_i m_{\odot,i}}. \quad (8.3)$$

<sup>7</sup>The latter would better match the near-IR continuum in an overall fit of the SED, but would over-predict the ionizing UV field.




**Figure 8.3:** Best fit composition to the dust in the Homunculus of  $\eta$  Carinae presented in a ‘tree diagram’. The ‘a’ and ‘c’ indicate amorphous and crystalline material respectively. a-ISM-silicate  $\equiv$  a-Mg<sub>1.447</sub>Fe<sub>0.02</sub>Na<sub>0.036</sub>Al<sub>0.036</sub>SiO<sub>3.538</sub>.

Expressing the abundances in this way ‘weighs’ the mass of a specific element relative to the total mass. Note that  $\Phi_i$  uses the mass of an element as it is found in the solid state material, which might be very different from the mass of that element in the gas phase. Deviations from solar abundance are therefore expected in this formulation.

## 8.4 RESULTS

The procedure described above leads to a ‘composition tree’ of best fitting dust compositions, shown in Fig. 8.3, in which every branch corresponds to a fit. Fits I and II have at the top of the branch a featureless component of metallic iron which is replaced by amorphous carbon for fits III and IV. All fits include amorphous olivine (MgFeSiO<sub>4</sub>) and have pyroxenes at the end of their branch. As there are several pyroxenes that satisfy the criterion of Eq. (8.2), we only show the ones that improve the fit the most. We note that in branch II also crystalline magnesium sulfide (c-MgS) is found to satisfy the criterion of Eq. (8.2) (though only just) and that the same is true for crystalline carbon in branch III. We note that branch III is in closest agreement with the three-component fit of (Morris et al. 1999).

Crystalline corundum (c-Al<sub>2</sub>O<sub>3</sub>; featuring in branch I and IV) and silicon carbide (SiC; in branch II and III) play an important role in the fitting of the red shoulder of the 10  $\mu$ m feature which extends up to  $\sim$  15  $\mu$ m. If we were to continue adding branches to the tree, adding a fifth species, fit I would be supplemented with c-MgSiO<sub>3</sub>, the crystalline equivalent of the amorphous pyroxene already present. This crystalline material mainly improves fitting of the blue side of the 10  $\mu$ m feature. Branch II and III however would both include crystalline corundum which suggests that although very similar to SiC, corundum is better suited to account for the red



**Figure 8.4:** Fit with the composition of branch I (our best fit). Both axes are logarithmic. Species that are listed in the legend but do not appear in the figure have a peak contribution less than 300 Jy. The bottom panel shows the relative deviations plotted on a linear scale.

shoulder of the  $10\mu\text{m}$  feature. Finally, branch IV would include amorphous corundum (in addition to the crystalline equivalent). These additions seem to represent mere refinements of the composition and are hence not included in Fig. 8.3.


Except for the featureless components the composition of the four fits is very similar. All four fits share olivine and pyroxenes and either SiC or corundum which, although important for compositional considerations, contribute negligibly to the mass. Basic parameters of the fits in the four branches of the composition tree of Fig. 8.3 are given in Tab. 8.1, a fifth fit (denoted by ‘min.’ and discussed below) is also included.

In Fig. 8.4, the components that make up the fit of branch I are shown together with the SED of  $\eta$  Carinae. The coding is as follows: materials are grouped by line type (and color) and temperature is indicated by line width where thinner is hotter (and shade where lighter is hotter). The first part of the SED (from say 2 to  $5\mu\text{m}$ ) is due to light from the star imposed on a continuum of hot (502 K) and warm (254 K) metallic iron. The cool (90 K) component which completely dominates the dust mass ( $> 99\%$ ) contains predominantly pyroxene, olivine and iron and has a mass of  $0.22 \pm 0.07 M_{\odot}$ . The error-estimate is the combined effect of a small measurement error (estimated to be  $< 3\%$  by bootstrapping) and an error due to the uncertainty of the absolute flux calibration that can reach 30%. The  $10\mu\text{m}$  feature is due to olivine and corundum added to a smooth continuum of iron, in agreement with the results of Chesneau


et al. (2005). In the bottom panel, the relative difference between the fit and the data (residuals) are plotted in linear scale. Deviations are typically only a few percent, with peaks reaching ten percent. These deviations may signal two types of flaws in the fit. First, one or several specific materials are left out that ought to be included to account for (part of) these particular spectral signatures. Second, the optical properties of the dust can sensitively depend on the conditions of the laboratory measurements and the details of the chemical composition (i.e. purity issues) and shape distribution of the grains (for a discussion see e.g. Kimura et al. 2005).

### 8.4.1 Statistical significance of the fit results

As mentioned in Sect. 8.3.3, we have performed a bootstrapping procedure to evaluate the statistical significance of our fit results. The resulting distributions of the quality of fit (QOF) are shown in Fig. 8.7. The initial assumption of uncorrelated data points (and thus a Poisson distribution) seems validated by the shape of the distributions. From the figure we can immediately see that the fit with lowest QOF (Branch I) is significantly better (more than  $6\sigma$ ) than the runner up (Branch II). Branch IV is not significantly different from either Branch II or Branch III. These results indicated that it would not increase our knowledge of the problem if we were to augment our models with more variables because this would bring the various results closer together in QOF space, making them statistically indistinguishable.




**Figure 8.5:** Close-up view of the fit (red dotted line) of branch I. Both the near-IR part (top panel) and the  $10\mu\text{m}$  feature (middle panel) are fitted convincingly. The region from  $22$  to  $35\mu\text{m}$  (bottom panel) has some apparent shortcomings. In the green dashed line  $\text{Mg}_{0.9}\text{Fe}_{0.1}\text{S}$  is added to the dust composition. This improves the fit only slightly in the region from  $\sim 22$  to  $\sim 35\mu\text{m}$ .




**Figure 8.6:** Fit with a cool component that is dominated by metallic iron. As metallic iron is a very efficient emitter, the total mass associated with this fit is about 7 times lower than that of our branch I. The abundance pattern implied by this fit is representative for non-equilibrium chemistry predictions by [Gail et al. \(1999\)](#). Species that are listed in the legend but do appear in the figure have a peak contribution less than 300 Jy. The bottom panel shows the relative deviations plotted on a linear scale.

#### 8.4.2 The potential presence of MgS

In Fig. 8.5, a detailed view of the SED fit of Fig. 8.4 (representing branch I) is shown. In the near-IR (top panel) the fit (red dotted line) including the stellar model agrees well with the observations. The  $10\ \mu\text{m}$  feature (middle panel) is also fitted convincingly but the region between approximately 22 to  $35\ \mu\text{m}$  shows some obvious deviations. As discussed in Sect. 8.2, the broad feature from  $\sim 24$  to about  $30\ \mu\text{m}$  is similar to that seen in carbon-rich AGB stars and has been ascribed to MgS. To investigate the possibility of MgS as the carrier for the feature observed in  $\eta$  Carinae, we fitted the SED between 16 and  $40\ \mu\text{m}$  using a continuum of metallic iron (and the Lorentz model) and adding every material of the library separately to this continuum. We found that  $\text{Mg}_{0.9}\text{Fe}_{0.1}\text{S}$  best reproduces the general form of the flat topped feature (see Fig. 8.8). If we extend the fitting range to shorter and longer wavelengths, silicates tend to supersede the magnesium silicate because they perform better at shorter wavelengths (though, as mentioned, in branch II MgS also fulfills the criterion of Eq. (8.2) for the fourth species). In fits in which silicates are not a major component (such as the minimum dust mass model, see below) MgS appears naturally. Adding MgS does not alter the total dust mass significantly.



**Figure 8.7:** Distributions of quality-of-fit (QOF) for our four fit results obtained with bootstrapping of the original data. Differences are expressed in standard deviations  $\sigma$  as obtained from the distribution of Branch I.



**Figure 8.8:** Fit of the region between 16 to 40  $\mu\text{m}$  with metallic iron and  $\text{Mg}_{0.9}\text{Fe}_{0.1}\text{S}$  and a Lorentzian.  $\text{Mg}_{0.9}\text{Fe}_{0.1}\text{S}$  seems to do justice to the steep rise at  $\sim 26 \mu\text{m}$  and the overall flatness of the feature. The sharp signatures of the feature are not accounted for by the relatively smooth fit.

### 8.4.3 Minimum dust mass model

Of the four branches in Fig. 8.3, branch I provides a lower limit to the total dust mass. It is, however, possible to make a fit with much less mass than  $0.22 M_{\odot}$  (see Fig. 8.6). The freedom to do this exists because any material that is featureless at long wavelengths – and most are – can be used to represent the cool component, i.e. the component that is dominating the dust mass. Because metallic iron is a very efficient emitter, the total dust mass taking this species to represent the cold dust is merely  $\sim 0.03 \pm 0.01 M_{\odot}$ , roughly a factor 7 lower than the mass of fit I. More than 95% of the mass is due to iron while silicates contribute less than 1% to the mass (see Tab. 8.2). Interestingly, such a peculiar dust pattern in the Homunculus of  $\eta$  Carinae



## FIT PARAMETERS: TEMPERATURE AND MASS

Fit	Quality	T <sub>C</sub>	T <sub>W</sub>	T <sub>H</sub>	M <sub>tot</sub>	$\sigma$	M <sub>C</sub>	M <sub>W</sub>	M <sub>H</sub>
I	328.5	89.8	253.6	501.7	0.22	0.07	0.219	0.00154	0.000224
II	358.7	81.3	205.1	468.3	0.3	0.1	0.298	0.00143	0.000210
III	386.3	82.7	173.9	373.1	0.5	0.2	0.483	0.0336	0.000576
IV	369.4	90.9	295.4	590.8	0.7	0.2	0.652	0.00105	0.000175
min.	417.7	122.2	345.0	729.5	0.03	0.01	0.0337	0.000580	0.0000666

**Table 8.1:** Summary of properties of the fits presented in Sect. 8.4. Values in the first column refer to the compositional branch in Fig. 8.3 except for the last line where ‘min.’ indicates the fit with predominantly iron in its cool component. The second column gives the quality of the fit as the square root of  $\chi^2$  where lower is better. T<sub>C</sub>, T<sub>W</sub> and T<sub>H</sub> are the temperature of the cool, warm and hot component respectively in Kelvin. M<sub>X</sub> refers to the dust mass in the X component where X can be C (cool), W (warm), H (hot) and tot (total) in solar masses. The error on the derived dust mass (column 7) is estimated to be 30 percent and is dominated by flux calibration uncertainties.

is predicted by Gail et al. (2005, see Sect. 8.5.1 for a discussion). These authors predict, in conjunction with a large quantity of metallic iron, FeSi to condense out in the ejecta of the Great Eruption. Although it does not occur ‘naturally’ in our fitting procedure the long wavelength behaviour of (cold) FeSi is very similar to that of metallic iron. Cold FeSi can therefore not be distinguished from iron spectrally and could very well be present. As the emission efficiency is also similar to iron, the mass of the fit would not be altered significantly if FeSi were to replace iron.

#### 8.4.4 Very cold dust beyond the ISO range?

We have looked for contributions from very cold dust at wavelengths beyond 170  $\mu\text{m}$  that are not probed by the ISO data by extending our observations up to 3.3 cm. Figure 8.9 shows the ISO SED supplemented by sub-millimeter points from Gomez (Née Morgan) et al. (2006) at 450 and 850  $\mu\text{m}$ , millimeter observations at 1.3 and 2.9 mm from Cox et al. (1995) and a 3.3 cm measurement by White et al. (1994). Overplotted is our fit I which does *not* use the stellar spectrum as a constraint. The sub-millimeter points nicely match the Rayleigh-Jeans tail of the dust emission. At wavelengths beyond about a millimeter the free-free emission of the present-day stellar wind dominates the emission. We find that in this regime the *intrinsic* stellar wind model from (Hillier et al. 2001) – so without applying a grey extinction – adequately accounts for the observed flux.

This match of data and model is consistent with the ISO spectrum being exclusively due to dust emission from the Homunculus.

## CHEMICAL COMPOSITION AND ELEMENTAL ABUNDANCE OF THE DUST

Fit	$\Phi_{\text{O}}$	$\Phi_{\text{C}}$	$\Phi_{\text{Fe}}$	$\Phi_{\text{Mg}}$	$\Phi_{\text{Al}}$	$\Phi_{\text{Si}}$	Material and mass (percentage)
I	0.7	0.0	1.7	5.9	5.7	4.0	c-Fe (10), a-MgFeSiO <sub>4</sub> (32) a-MgSiO <sub>3</sub> (58), a-Al <sub>2</sub> O <sub>3</sub> (< 1)
II	0.6	0.005	3.0	4.0	0.004	3.2	c-Fe (4), a-MgFeSiO <sub>4</sub> (94), a-SiC (< 1) c-Mg <sub>0.5</sub> Fe <sub>0.43</sub> Ca <sub>0.03</sub> Al <sub>0.04</sub> SiO <sub>3</sub> (< 1)
III	0.5	2.4	0.0	2.8	0.0	2.3	a-C (42), a-MgFeSiO <sub>4</sub> (44), a-SiC (< 1) a-Mg <sub>0.5</sub> Fe <sub>0.43</sub> Ca <sub>0.03</sub> Al <sub>0.04</sub> SiO <sub>3</sub> (13)
IV	0.2	3.8	0.4	0.9	5.7	0.6	a-C (79), a-Al <sub>2</sub> O <sub>3</sub> (4), a-MgFeSiO <sub>4</sub> (13), a-ISM silicate (4)
min	0.01	0.0	8.3	0.2	0.8	0.03	c-Fe (97), a-ISM silicate (< 1), a-Al <sub>2</sub> O <sub>3</sub> (< 1), c-Fe <sub>0.1</sub> Mg <sub>0.9</sub> S (1)

**Table 8.2:** Summary of properties of the fits presented in Sect. 8.4.  $\Phi_e$  is the abundance relative to solar of element  $e$  according to Eq. (8.3). As an example, consider the contribution from iron to the minimum dust model (bottom row, fourth column). The value of 8.3 means that in the solid state material of the Homunculus, iron is 8.3 times more abundant than in the solar gas. This implies that, because the abundances are calculated relative to the total mass, the other materials have been suppressed relative to Fe. This is reflected in the fractions of the other species in the fit which are all smaller than one. This overabundance of a species can signal enhanced efficiency of its condensation or suppressed condensation of the other species, or both. As C and O are depleted in CNO-processed material, the relative overabundance of Fe is not surprising. Note that, although Si is found to be relatively underabundant in our fit its value is uncertain as FeSi can replace Fe as the dominant species in the cool component. The materials in the fit and their (approximate) contribution to the dust mass (in percentage) are listed in the last column. The Lorentzian component which was added to all models is not mentioned explicitly since its carrier is unidentified. Its contribution to the total dustmass never exceeds a percent.


## 8.5 DISCUSSION AND CONCLUSIONS

So, what is the dust composition in the Homunculus of  $\eta$  Carinae, what is the dust mass, and what is the total mass in gas and dust that can be associated with the Great Eruption? Even though the different branches differ substantially in dust composition, the resulting models perform virtually equally well in reproducing the observed SED (See Table 8.1, column 2). If we add to this the significant uncertainty on the flux calibration and the dependence of the dust optical properties on the laboratory conditions and purity of the samples, it appears presumptuous to rely only on the quality of fit to deduce which branch best represents the dust in the nebula of  $\eta$  Carinae. In this section we discuss the plausibility of the four branches (I through IV) and the minimum dust mass model in terms of dust formation in a CNO processed environment.

### 8.5.1 Composition in the context of dust formation

The central source in  $\eta$  Carinae shows CNO-processed material at its surface (Hillier et al. 2001) and several ejecta associated with the system and probably ejected *prior* to the Great Eruption (Walborn et al. 1978; Walborn & Alvarez 1988) also show severe depletion in carbon and oxygen in favour of nitrogen (see e.g. Davidson et al. 1982; Smith & Morse 2004). This suggests that a mineralogy with significant contributions from amorphous carbon (as in branches III and IV) is unlikely. In this respect it is also important that no emission from polycyclic aromatic hydrocarbons is detected in spite of the presence of a strong UV radiation field, in stark contrast with carbon-rich planetary nebulae, i.e. amorphous carbon based ejecta that are illuminated by a nearby UV source.

Corundum (branch I) or silicon carbide (branch II) are primarily required to fit the extended red wing of the  $10\ \mu\text{m}$  feature, with silicates producing the remainder of the profile. Dust emit-



**Figure 8.9:** ISO SED and several data points beyond the ISO range. The flux points from Gomez (Née Morgan) et al. (2006) are well matched by a fit without stellar model (thick dotted line) and the flux points from Cox et al. (1995) and White et al. (1994) are convincingly accounted for by the stellar model.

ting at such short wavelengths is part of the hot component, accounting for only a minor fraction (less than one percent) of the total dust mass. Dust in the hot component – seen relatively close to the central star – may not represent dust that is formed in the Great Eruption. Though SiC is predicted to form in Giant Eruption like conditions (see below), it appears plausible that in a relatively quiescent situation dust formation might proceed more like that seen in outflows from cool stars, where corundum is formed at the start of the condensation sequence (and SiC only in reducing conditions, i.e.  $C/O > 1$ ; e.g. Lodders & Fegley 1999). Because  $\eta$  Carinae is thought to be a high-eccentricity binary (Damineli 1996) a situation could occur where the *present-day* winds of the central source and its putative companion collide at peri-astron passage. The short duration of these passages may prohibit the condensation sequence to progress beyond the formation of  $\alpha$ - $Al_2O_3$  leaving it as an important species close to the star. Chesneau et al. (2005), for instance, use corundum to fit the red wing of the  $10\ \mu m$  feature in the central  $3''$ . On these grounds, dust formation in the present-day (colliding) outflow favors branch I.

In Sect. 8.4.2 we discuss the potential presence of MgS. This species is observed in the outflows of normal carbon stars on the AGB (Hony et al. 2002) and recently in S-type stars (Smolders et al., in preparation). S-type stars are transitional objects between oxygen dominated and carbon dominated stars. They have  $C/O \sim 1$  and are known to produce silicates and oxides.

Perhaps MgS may also form in the oxygen depleted but still oxidizing ( $C/O < 1$ )  $\eta$  Carinae environment, where the small surplus of oxygen is bound in mostly simple oxides (such as corundum) and remaining metals drive a chemistry forming MgS.

Gail et al. (2005) present predictions of dust formation in an  $\eta$  Carinae like outburst. They point out that due to the lack of sufficient quantities of oxygen Mg-Fe-silicates form only in very modest quantities. No soot is formed. Abundant dust species are found to be FeSi and metallic Fe – consistent with our minimum dust mass model – with a small amount of forsterite ( $c$ - $Mg_2SiO_4$ ). These authors predict that SiC will be the first abundant condensate, but that, in chemical equilibrium, it rapidly disappears again in favor of FeSi. Under non-equilibrium conditions, that may occur in a rapidly expanding and cooling gas, it may nonetheless survive because SiC and FeSi start to form at similar temperatures and dust chemistry may freeze-in before equilibrium is reached.

We conclude that the CNO processed envelope and environs of  $\eta$  Carinae likely favor a carbon-less dust formation. The hot dust contains amorphous silicates, small amounts of pyroxenes and corundum and/or silicon carbide. A distinct 'flat topped' feature at  $24$ - $30\ \mu m$  is best explained by warm MgS dust. The composition of cold dust can not be constrained spectroscopically because the spectrum at long wavelengths is featureless. Silicates may be responsible. Calculations of dust condensation in the  $\eta$  Carinae outburst, however, predict dust that is dominated by Fe and FeSi. We can also achieve a good fit using such a composition for the cold component, requiring relatively limited amounts of dust ( $0.03 M_\odot$ ), i.e. the minimum model in Table 8.1.

## 8.5.2 The total mass in the Homunculus

Knowledge of the gas-to-dust ratio  $f_{gd}$  is required in order to convert the total dust mass of the Homunculus to a total mass. It is however not possible to empirically determine the value of  $f_{gd}$  as the nebula is mainly a reflection nebula at ultraviolet and optical wavelengths making it

difficult to measure the gas content (see e.g. [Meaburn et al. 1993](#)). Sensitive as this quantity is to the abundance pattern and local dust forming conditions, the gas-to-dust ratio in the nebula of an evolved star such as  $\eta$  Carinae is likely to deviate appreciatively from the canonical value of 100. The material in the Homunculus is known to be CNO-processed and therefore depleted in carbon and oxygen in favour of nitrogen. As O (and C in the case of carbon rich chemistry) is important in dust formation while N is essentially inert,  $f_{\text{gd}}$  is likely to be larger than 100. Evolutionary tracks of an initially  $120 M_{\odot}$  star ([Schaller et al. 1992](#)) show a carbon and oxygen depletion of about a factor 20 and 10 respectively at the time the mass fractions of hydrogen and helium are about equal. The latter is typical for the Luminous Blue Variable phase  $\eta$  Carinae is thought to be in. In normal conditions oxygen accounts for about half the mass in dust forming elements (O, C, Mg, Si, Fe, S, Al, Ca and Ti). These depletions thus suggest  $f_{\text{gd}} \sim 200$ .

If indeed large amounts of metallic Fe and FeSi are produced in the Giant Eruption (see above), the minimum dust mass model shows that the total dust mass may be an order of magnitude less compared to a situation in which the dust is dominated by silicates. In case the dust is dominated by Fe and Si, the gas to dust ratio  $f_{\text{gd}}$  by complete condensation of these two elements is  $\sim 350$ . This would correspond to a total mass of  $10 M_{\odot}$ . However, [Gail et al. \(2005\)](#) report degrees of condensation of only 10 to 15 percent, implying an  $f_{\text{gd}}$  of at least  $\sim 2500$ . Such a gas-to-dust ratio corresponds to a total mass of  $\sim 75 M_{\odot}$ .

Due to the approximate black-body behaviour of the overall SED ([Smith et al. 2003b](#); [Hillier et al. 2006](#); [Smith & Ferland 2007](#); [Smith 2010](#)), the total mass of our model fits with mainly contribution from silicates (Fit I and II) is in good agreement with earlier results obtained by less rigorous methods than employed here (e.g. [Morris et al. 1999](#); [Smith et al. 2003b](#)).

We conclude that the likely total mass is at least some  $\sim 45 M_{\odot}$  if the grains are dominated by silicates (branch I). This assumes  $f_{\text{gd}} \sim 200$ . If Fe and FeSi dominate the grain content the dust mass is as low as  $0.03 M_{\odot}$ . Assuming full condensation of Fe and FeSi results in a total mass of the Homunculus of  $10 M_{\odot}$ , that can be identified as the minimum mass of the nebula. Predictions of Fe and FeSi condensation show that only part of the nebular gas is converted into dust, and imply a total mass that may even be larger than  $45 M_{\odot}$ .

### 8.5.3 Energetics of the Great Eruption

[Smith \(2006\)](#) proposes a model to account for the spatial distribution of velocity and mass in the Homunculus. In this model, most mass is situated at angles  $\lesssim 45^{\circ}$  from the bipolar axis and the velocity follows an approximate Hubble law (i.e. the velocity is proportional to distance from the star), reaching  $\sim 650 \text{ km sec}^{-1}$  in the polar caps. We derive for the kinetic energy of the ejecta

$$E_{\text{kin}} \simeq 5.8 \times 10^{49} \left( \frac{f_{\text{gd}}}{100} \right) \left( \frac{M_{\text{d}}}{0.1 M_{\odot}} \right) \text{ ergs} \quad (8.4)$$

where  $M_{\text{d}}$  is the total dust mass in solar masses. For  $M_{\text{d}} = 0.22 M_{\odot}$  and  $f_{\text{gd}} = 200$  this yields  $E_{\text{kin}} \simeq 10^{50.4}$  ergs. Such an energy almost rivals that of the radiative output of typical type-II supernovae ( $\sim 10^{51}$  ergs) and exceeds the luminous energy presumed to be radiated by  $\eta$  Carinae during the 20 yr duration of the Great Eruption, which was  $E_{\text{rad}} \simeq 10^{49.5}$  ergs ([Davidson & Humphreys 1997](#)). Equating  $E_{\text{kin}}$  to  $E_{\text{rad}}$  would yield a Homunculus mass of  $5.5 M_{\odot}$ , assuming  $f_{\text{gd}} = 200$ . It has been proposed by [Owocki et al. \(2004\)](#) that a power-law porosity-moderated

---

continuum driven wind may lead to the very large mass-loss rate implied in the Great Eruption. Such models lead to mass-loss rates that approach the photon-tiring limit. A by-product of such models is that there is little remaining radiative energy to be observed as emergent luminosity, explaining why the dust masses found here imply  $E_{\text{kin}} > E_{\text{rad}}$ .



## CHAPTER 9

---

# Radiative transfer in the $\eta$ Carinae system

C. J. H. Kruip<sup>1</sup>, T. I. Madura<sup>2</sup>, J. H. Groh<sup>2</sup>, N. Clementel<sup>1</sup>,  
S. Verdolini<sup>1</sup>, T. R. Gull<sup>3</sup> & V. Icke<sup>1</sup>

**W**E perform radiative transfer post-processing of three-dimensional hydrodynamical simulations of the interacting winds in the  $\eta$  Car system. This results in ionization-maps of hydrogen and helium that constrain the regions where forbidden lines can form. We explore the difficulties that arise when paving the way for the next generation of synthetic observations of these forbidden lines.

---


<sup>1</sup>Leiden University, Niels Bohrweg 2, 2333 CA Leiden, Leiden, the Netherlands

<sup>2</sup>Max-Planck-Institut fuer Radioastronomie, Auf dem Huegel 69, Bonn, Germany

<sup>3</sup>Laboratory for Extrasolar Planets and Stellar Astrophysics, Exploration of the Universe Division, Code 667, Goddard Space Flight Center, Greenbelt, MD 20771, USA



## 9.1 INTRODUCTION



**Figure 9.1:** *Left image:* Eta Carinae as seen in visible light with the Hubble Space Telescope (Morse et al. 1998). *Right inset:* Cut through SPH simulation of the interacting wind region in the plane of the orbit (Madura 2010) at apastron. The brightness indicates the gas density in logarithmic scale where light indicates high and dark indicates low density. As a result of the violent gravito-hydrodynamical interaction at periastron, large under-dense regions open up toward the observer.

As argued in Chapter 8, the enigmatic object  $\eta$  Carinae is one of the most extreme stars in the nearby universe. Despite a huge amount of research, many fundamental questions about the system remain unanswered. From a dynamical point of view, perhaps the most fundamental parameters are the orientation, eccentricity and period of the binary orbit.

The binary period has been determined by observations of the recurring *spectroscopic events* to be 5.52 years (Damineli 1996). From timing of the spectroscopic event, the eccentricity of the orbit has been constrained to be higher than 0.8 (Corcoran et al. 2001). It is unclear if the orbits' extreme eccentricity is a result of the Great Eruption or if it has another origin such as a gravitational interaction with a third object.

Several authors have modeled the X-ray light-curve assuming the radiation emanates from the shocked wind-wind collision interface (Pittard & Corcoran 2002; Okazaki et al. 2008; Parkin et al. 2011, and references therein). Since X-rays are spatially unresolved, they cannot be used to constrain the absolute orientation of the orbital axis (the axis perpendicular to the orbital plane). Most studies assume that the orbital axis is aligned with the symmetry axis of the Homunculus (Gull et al. 2009).

Recently, high-resolution HST spectroscopy of [FeIII] has been used to constrain the orientation of the binary orbit in three dimensions (Madura et al. 2010; Madura et al. 2011). A combination of SPH simulations (see right panel of Fig. 9.1) with optically thin radiative transfer has been deployed to generate synthetic-observations which were, in turn, confronted with

the data. This approach breaks the degeneracy and for the first time successfully constrained the orientation of the orbit in three dimensional space. One important result is the alignment of the orbital direction with the bipolar Homunculus nebula, in strengthening the predicted link between the binary orbit and the Great Eruption.

Although very successful, the approach of Madura et al. (2010); Madura et al. (2011) can be improved upon considerably. Particularly, the origin of the forbidden line emission has been based on simple geometrical criteria combined with a density threshold although, in reality, the population of forbidden states depends on the local ionization state of the medium. No radiative transfer of ionising radiation has been included in these simulations.

In this chapter we strive to improve on this simplified approach by post-processing the SPH simulations of (Madura et al. 2010; Madura et al. 2011) with full three-dimensional radiative transfer of radiation capable of ionising hydrogen and helium. We use `SIMPLEX` to obtain detailed ionization fractions of both species at the resolution of the original SPH simulations. This allows us to predict much more precisely where, and to what extent forbidden iron and nitrogen lines are formed. This will ultimately lead to improved synthetic observations that yield more accurate measurements of the orbital parameters and perhaps even unconstrained parameters like the mass fraction of the binary stars.<sup>4</sup>

## 9.2 HYDRODYNAMIC SIMULATIONS

There are two main approaches to the hydrodynamics of WWIRs that can be used complementary to each other: SPH and AMR based methods. The major advantage of SPH is its ability to naturally cover orders of magnitude in resolution within the same simulation which makes it possible to study the WWIR at distances far removed from the binary system itself. This is very useful, because observationally, this region is far simpler to handle and understand than the core region. Moreover, the spatial scale of the available HST data requires a large region around the stars to be simulated.

The inevitable drawback of large simulation volumes is the relatively poor representation of shocks and contact-discontinuities in the inner part of the simulation. Grid based hydrodynamics, on the other hand, is ideally suited to study these inner parts where accurate shock capturing is important.

In this chapter we use both approaches for the reasons discussed above. We have to note, however, that the AMR hydrodynamics is of a very preliminary nature and will play a much more significant role in future work on this topic.

### 9.2.1 SPH

We use the SPH simulations from Madura et al. (2010); Madura et al. (2011) as the basis for our radiative transfer calculations. The code is a modified version of the one used in Okazaki et al. (2008). We use simulations with two different box-sizes, 210 and 21 semi major axes.

---

<sup>4</sup>From the total light, a combined mass of the binary of at least  $100 M_{\odot}$  can be derived but it is unclear how this mass is divided between the two stellar components.

With a semi major axis of about 15.4 AU, this results in simulation domains of 3234 and 323.4 AU respectively.

Snapshots of the large-box simulation are available at intervals of roughly one week or, equivalently, 0.00318 in phase. Every snapshot contains approximately  $5 \times 10^5$  particles. This number is not fixed as the particles are allowed to stream out of the simulation domain, and the stars emit new particles as they orbit each other. The stars themselves are modeled by single SPH particles with the (approximate) mass of the stellar sources.

## 9.2.2 Parameters

The major hydrodynamical parameters are the wind outflow speed and mass-loss rate of the two stars. For the radiative transfer, an ionising luminosity must be assigned to both stars. The parameters used in our simulations are summarized in Tab. 9.1.


Parameter	Value	reference
Primary $\dot{M}$	$1 \times 10^{-3} M_{\odot}/\text{yr}$	( <a href="#">Hillier et al. 2001, 2006</a> )
Secondary $\dot{M}$	$10^{-5} M_{\odot}/\text{yr}$	( <a href="#">Pittard &amp; Corcoran 2002</a> )
Primary ionising flux	$1 \times 10^{51}/\text{s}$	
Secondary ionising flux	$3.02 \times 10^{49}/\text{s}$	( <a href="#">Martins et al. 2002</a> )

**Table 9.1:** Overview of simulation parameters.


### Grid construction

To convert the SPH simulations to a `SIMPLEX` mesh, we use every SPH particle as the nucleus of a Voronoi cell whose density is defined as the mass of the particle divided by the volume of the cell. This procedure yields density estimates that are less smooth than those obtained with typical kernel functions of the type of Eq. (4.8) but guarantee mass conservation and represent small scale structures in the density field more accurately. The various issues related to density estimates and mesh construction for the data used in this chapter are discussed in Chapter 4 and specifically Sect. 4.4.1.

In Fig 9.2, two cuts through the density fields of the large box simulation are shown, one at apastron (left panel) and one at periastron (right panel). The density fields clearly show the spherical shells that are deposited during every orbit and slowly coast outwards until they leave the computational domain. The shells are disrupted towards the right due to the secondary wind which blows the shell apart in that direction. The momentum of the secondary wind to the left is balanced by that of the primary wind which is much denser and slower. A close up view of the inner tenth of the large computation volume is shown in Fig. 9.3. It is now possible to see that the momentum balance between the wind of the primary star (in the left panel at apastron, the primary is the larger star that sits almost in the centre) and that of the secondary star is such that the WWIR lies closer to the secondary star (at about 75% of the separation distance). The periastron image (right panel) shows the characteristic conical shape of the WWIR, with its tip bent around the primary star.




**Figure 9.2:** Density in logarithmic scale for different phases: apastron (left) and periastron (right) for the large simulation volume. The box spans about 3234 AU (or 210 semi-major axis of the binary). The cut is made through the orbital plane.



**Figure 9.3:** Same as Fig. 9.2 but for the small size box. The box spans approximately 304.5 AU (or 21 semi-major axis of the binary).

As described in Chapter 4, all the SPH particles (approximately  $5 \times 10^5$ ) are converted to vertices of the RT mesh. Due to the modest number of particles in these simulations it is not necessary to take subsamples of the data. An example of the resulting `SIMPLEX` mesh (using the large box simulation at apastron) is shown in Fig. 9.4. The sharp variations in number density of SPH particles (and thus vertices) that occur as we pass from a dense shell into the rarified

region between shells can be problematic considering the analysis presented in Chapter 3. In Sect. 4.4.1 we have assessed to what extent these sharp gradients influence the outcome of our RT simulations. Because it turned out that their effect is negligible, we will proceed to use the SPH directly as the generating vertices of our mesh throughout this chapter.



**Figure 9.4:** Slice through the RT mesh for the same data as shown in the left panel of Fig. 9.2. The resolution of the `SIMPLEX` mesh follows the resolution of the original data everywhere.

### 9.2.3 AMR hydrodynamics


We have performed preliminary AMR hydro simulations of the  $\eta$  Carinae system using the Flash code (Fryxell et al. 2000). For these tests, a small computation volume has been used with a maximum of 5 levels of refinement. The computational domain spans approximately 36 AU in every direction and contains a total of roughly  $8 \times 10^6$  cells depending on the phase of the orbit.

The driving mechanism for the stellar winds is different from that employed in the SPH (where the particles are ejected from the stellar surface with the terminal velocity) by that is *pressure driven*. In practice this means that the winds pick up speed over a non-zero distance set by the hydrodynamical conditions within and around the star. In these simulations, the stars are represented by a *stencil* that has an approximately spherical shape. At every hydro-step this stencil is used to initialize the cells within the stellar radius to their (constant) values.

Also, the mass-loss-rate of the primary star is different from that used in the SPH data, namely  $2.5 \times 10^{-4} M_{\odot}/\text{yr}$  instead of  $1 \times 10^{-3} M_{\odot}/\text{yr}$ . This is expected to make the WWIR less dense and thus easier to ionise.

In Sect. 4.3.3 we have compared various methods to include the mass of un-sampled AMR cells. In Fig. 9.5 one of the resulting `SIMPLEX` meshes is shown alongside the original AMR data and the resulting Delaunay triangulation (which is the geometrical basis for the RT mesh). We have used a sampling function of the form Eq. (4.2) with  $\alpha = 0.5$  and  $n_0 = 10^6 \text{ cm}^{-3}$  and  $2 \times 10^5$  vertices. To prevent the majority of vertices being placed very close to the density peak around

the central star, we have implemented a density cut-off at  $5 \times 10^{10}$ . These choices result in a moderate increase of resolution with density, accurately resolving the contact discontinuity and features therein without depriving the less dense regions of their resolution. In Sect. 9.4.4 we use this snapshot to gain initial insight into the differences between RT applied on (relatively) high resolution AMR data and the SPH data used in the rest of this chapter. To give an idea




**Figure 9.5:** Grid construction from AMR data. *Left:* Density field of the original data. *Centre:* SIMPLEX mesh using the Delaunay mass inclusion method (see Sect. 4.3.3 for details) *Right:* Delaunay triangulation of the SIMPLEX mesh where the resolution is increased in the WWIR in order to capture the relevant details.

of the resolution of this simulation, we compare it side-by-side to the small (high resolution) SPH simulation box in Fig. 9.6. Although the orientation and mass-loss-rate are different (especially the mass-loss-rate), this gives a good indication of the difference in detail in the WWIR and the relative smoothness of the mesh-based hydro at this scale. The graininess of the SPH representation of the primary wind is particularly conspicuous. This is a direct consequence of the particle based nature of the method. The contact discontinuity, although not nearly as detailed as in the mesh-based hydro still shows the basic characteristics of a density jump. Given the significant differences in size of the simulation volume (the smaller SPH simulation is a factor  $323.4/36 \simeq 9$  larger in diameter) and resolution (the same simulation uses a factor  $5 \times 10^5/8 \times 10^6 = 1/16$  fewer resolution elements), the SPH simulation captures the essence of the system remarkably well.

### 9.3 RT SIMULATIONS

We have performed the radiative transfer calculations in post-processing using the SIMPLEX code as described in Chapter 2 with one important alteration. For a fraction of cells close to the primary source, the ionization- and/or recombination time-scale are exceedingly short (in the order of seconds). Because the non-equilibrium chemistry module described in Chapter 5 essentially integrates the chemical network on the shortest chemical timescale, these cells need an unfeasibly large number of sub-cycle steps to ensure a correct solution. To overcome this



**Figure 9.6:** Comparing the SPH to the AMR data. *Left:* Density field of the AMR hydro simulation. *Right:* Delaunay triangulation of the SPH data of the small box simulation.

problem, we have implemented a equilibrium chemistry solver that deals with those cells which cannot be treated with the non-equilibrium module.

### 9.3.1 Equilibrium chemistry

The equilibrium solution effectively puts the ionization fractions of hydrogen and helium to their equilibrium values under the assumption that the incoming flux of ionising photons is constant.

As already mentioned in Chapter 5, the ionization-state of the gas is described by the following three coupled differential equations and three closure relations

$$\dot{n}_{\text{HI}} = n_{\text{HII}}R_{\text{HI}} - n_{\text{HI}}\Gamma_{\text{HI}} \quad (9.1)$$

$$\dot{n}_{\text{HeI}} = n_{\text{HeII}}R_{\text{HeI}} - n_{\text{HeI}}\Gamma_{\text{HeI}} \quad (9.2)$$

$$\dot{n}_{\text{HeIII}} = -n_{\text{HeIII}}R_{\text{HeII}} + n_{\text{HeII}}\Gamma_{\text{HeII}} \quad (9.3)$$

$$n_{\text{H}} = n_{\text{HI}} + n_{\text{HII}} \quad (9.4)$$

$$n_{\text{He}} = n_{\text{HeI}} + n_{\text{HeII}} + n_{\text{HeIII}} \quad (9.5)$$

$$n_e = n_{\text{HII}} + n_{\text{HeII}} + 2n_{\text{HeIII}}. \quad (9.6)$$

As before, the ionization rates are the sum of both photo- and collisional ionization rates. We can derive the equilibrium equations by putting  $\dot{n}_{\text{HI}} = \dot{n}_{\text{HeI}} = \dot{n}_{\text{HeII}} = 0$ . After some algebra this

yields

$$x_{\text{HI}} = (1 + \Gamma_{\text{HI}}/R_{\text{HI}})^{-1} \quad (9.7)$$

$$x_{\text{HII}} = 1 - x_{\text{HI}} \quad (9.8)$$

$$x_{\text{HeI}} = [1 + \Gamma_{\text{HeI}}/R_{\text{HeI}} \times (1 + \Gamma_{\text{HeII}}/R_{\text{HeII}})]^{-1} \quad (9.9)$$

$$x_{\text{HeII}} = x_{\text{HeI}} \Gamma_{\text{HeI}}/R_{\text{HeI}} \quad (9.10)$$

$$x_{\text{HeIII}} = x_{\text{HeII}} \Gamma_{\text{HeII}}/R_{\text{HeII}}, \quad (9.11)$$

where  $x_i$  is the fraction of species  $i$  and we have used  $n_i = x_i n_j$  where  $j \in (\text{H}, \text{He})$ . These equations are coupled by the free electron density (given by the last equation in Eqs. (9.6)).

Unfortunately, the set of equations presented in Eqs. (9.11) cannot be solved analytically. This is due to the non-linear dependence on ionization fractions of the photo-ionization rate through the optical depth. More specifically, the photo-ionization rate in a cell is given by Eq. (5.6), where the monochromatic analog of Eq. (5.4) is given by

$$\dot{N}_{ion} = \dot{N}_{\gamma}(1 - \exp(-\tau)), \quad (9.12)$$

with  $\tau = (x_{\text{HI}} n_{\text{H}} \sigma_{\text{HI}} + x_{\text{HeI}} n_{\text{He}} \sigma_{\text{HeI}} + x_{\text{HeII}} n_{\text{He}} \sigma_{\text{HeII}}) l$ .

Because of this non-linear dependence, the equilibrium fractions must be iteratively solved for. If the iterative procedure converges (the neutral fractions of the different species differ less than a given tolerance between iterations), the neutral fractions are assigned to the cell under treatment and the flux is diminished by the number of absorptions during that time-step  $\Delta t$  ( $n_{\text{HI}} \Gamma_{\text{HI}} \Delta t$ ).

We apply the equilibrium chemistry to cells for which the non-equilibrium chemistry solver would take more than  $10^5$  sub-cycling steps. In a sense, this number is arbitrary and simply reflects how much time the user is willing to spend on finding the non-equilibrium solution. In our application, the cells for which the equilibrium solution is employed are primarily confined to a small region around the primary star. This region is expected to be highly ionised regardless and application of the equilibrium chemistry will therefore likely be a decent approximation.

### 9.3.2 Collisional ionizations

We incorporate collisional ionizations in addition to photo-ionizations to achieve a more complete census of the ionizations in the system. The collisional ionization fraction depends only on the temperature of the medium, which is in principle a function of the hydrodynamical motion of the gas (i.e., adiabatic terms), photo-heating and multiple cooling terms.

In this preliminary study, we do not use a coupled radiation-hydrodynamics code and are therefore unable to treat heating and cooling in a self-consistent fashion. As a first approximation, however, we can use the temperature calculated by the SPH code based on adiabatic expansion and compression only to estimate the importance of collisional ionizations. In future study, we will treat the temperature evolution of the system self-consistently and refine this effort. To assess which of the two is the dominant effect we perform both simulations with and without collisional-/photo-ionizations.



## 9.4 RESULTS

We investigate the influence of the primary and secondary sources on the ionization-structure of the wind-wind-interaction region (WWIR). Although extremely luminous, the primary source has a spectrum representative of a much cooler star than the secondary source (35000-40000 K). The effective temperature of the primary is lower but uncertain. This is mainly because measurements of the effective temperature for LBV stars strongly depend on the depth of the measured region (Hillier et al. 2001). Because the primary star is enshrouded in a dense wind, its ionising flux is diminished substantially before reaching the WWIR.

The secondary star on the other hand has a much harder spectrum which is hardly attenuated by its relatively tenuous wind. We therefore expect the secondary star to primarily shape the ionization structure of the WWIR which is the region of interest when predicting the forbidden line emission from the system.

We start with an exploration of the possibility to include ionising flux of the primary star in our simulations in Sect. 9.4.1 using one-dimensional RT simulations with hydrogen only. In Sect. 9.4.2, we investigate the influence of resolution, collisional ionizations and the inclusion of helium on the 3D data. In Sect. 9.4.3, we look at the differences in morphology at both apastron and periastron for the small and large SPH volumes CHECK. Finally, in Sect. 9.4.4 we present preliminary results of RT on the AMR mesh data and discuss several differences with the SPH data.


### 9.4.1 One-dimensional explorations of the ionised region around the primary star

From the detailed fitting of the visual and UV spectrum by Hillier et al. (2001, 2006), the region of ionised hydrogen around the primary star has been found to extend radially about 120 AU from the stellar source. The density of hydrogen is expected to fall off approximately as  $r^{-2}$  due to the conservation of mass and because the primary wind is more or less constant.

To explore the dependence of the position of the ionization front on the ionising luminosity of the primary star, we have performed one-dimensional experiments using the equilibrium chemistry described above. The density profile is an analytical fit (see Fig. 9.7) to the spherically averaged density profile of the WIND900 SPH simulation box at apastron. To deal with the pole at the origin, a central core of constant density and a radius of 0.1 AU is used. This does not affect our findings, however, as the core is negligibly small compared to the total size of the box. The analytic density profile is thus characterized by


$${}^{10}\log n(r) = \begin{cases} 13 \text{ cm}^{-3} & \text{if } r < 0.1 \text{ AU} \\ 13 - 2({}^{10}\log r + 4) \text{ cm}^{-3} & \text{otherwise} \end{cases} \quad (9.13)$$

The 1D code simulates radiation traveling through spherically symmetric shells with a maximal radius of half the box-size of the WIND900 simulation  $\approx 1522$  AU. The radiation is injected in the first shell and then travels outward until it is either absorbed or exits the last shell. For the results shown below we have used  $3 \cdot 10^4$  shells. We note that time-stepping is arbitrary because of the equilibrium chemistry. The only variable is therefore the luminosity of the source.



**Figure 9.7:** Fit to the WIND900 radial density profile in log-log scale as described by Eq. 9.13. The fit is used in the 1D tests described in the text.

The neutral (solid line) and ionised (dashed line) fraction as a function of the radius are plotted for several qualitatively different configurations in Fig. 9.8. The density profile is plotted with a dotted line and quantified by the right axis. For luminosities below  $\approx 4.8 \cdot 10^{51} \text{ s}^{-1}$ , the box is neutral and the Strömberg radius is confined to the central core. With this setup, the ionization front is located somewhere between the centre and the outside of the box for a very small range in values centred around  $4.7636 \cdot 10^{51}$ . The centre panel gives an example of this configuration. The slightest increase in the luminosity results in a completely ionised box (right panel) while further increase results only in a lower neutral fraction throughout the simulation volume. This behaviour is completely expected, however, for ionization fronts in power-law



**Figure 9.8:** Neutral and ionised fractions for the 1D RT simulations described in the text. *Left:* Luminosity =  $10^{51} \text{ s}^{-1}$ . *Centre:* Luminosity =  $4.7635 \cdot 10^{51} \text{ s}^{-1}$ . *Right:* Luminosity =  $4.7636 \cdot 10^{51} \text{ s}^{-1}$ . This behaviour signifies that the ionization front is highly unstable and thus changes *qualitatively* with small perturbations in either the density or the luminosity.

density profiles with powers smaller than  $-2/3$  (Franco et al. 1990; Shapiro et al. 2006). For such profiles, the circumstellar medium simply cannot support stable ionization fronts.

However simplified, we can still derive several useful conclusions from this 1D approach. We know for instance that, because ionised hydrogen shows up in UV spectra, the ionising flux

of  $\eta$  Carinae A in our models should be above the magical value of  $\approx 4.8 \cdot 10^{51} \text{ s}^{-1}$ . Furthermore, we should abandon hope of constraining the ionization front for hydrogen to the value derived by Hillier et al. (2001), as it is made practically impossible by its intrinsically unstable nature.

So, what could be the origin of the paradox posed by the stable front at roughly 120 AU derived by Hillier et al. (2001) and the unstable front predicted theoretically and obtained using a 1D numerical study? We must realize that our result is over-simplified; the instability is real in a *pure* hydrogen gas, but will disappear with the introduction of the myriad spectral lines (mostly by Fe) that have a so-called *line blanketing* effect on stellar spectra.

Given these difficulties, perhaps the most sensible choice for a initial effort to model the ionised WWIR is to omit the primary star altogether. This may seem an oversimplification at first glance, but there are several arguments for this approximation. First, the primary source will sustain an ionised region (in hydrogen) that spans roughly 250 AU in diameter and less than 0.05% of the volume of the large box (see Fig. 9.9). This volume is too close to the central source to *directly* affect the ionization fraction at the regions where the forbidden line emission is to be studied. It may, however, influence the ionization structure further away indirectly by reducing the opacity for photons from the secondary source. We expect that this would primarily result in UV flux from the primary penetrating the WWIR more easily, effectively enlarging the the ionised fraction on the far side of the primary source (to the left). For an observer on earth, this region is, at apastron, located behind the two sources and therefore obscured by the dense primary wind anyway. Second, one may argue that it is likely that, because the ionised regions encompasses both stars and the densest part of the WWIR, photons from the primary also reach the right part of the simulation volume. This argument relies, however, on the assumption that the ionised region is indeed spherical and therefore penetrates the WWIR toward the secondary star. This assumption is likely incorrect given the high optical depth of the WWIR. In other words, we would be applying a model based on spherical symmetry to a region that clearly has a very asymmetrical geometry.


### 9.4.2 Three-dimensional RT of the secondary star

In this section we explore the ionization structure resulting from the secondary star's UV flux. Motivated by the difficulties discussed in the previous section, we neglect the influence of the primary star for the moment, a simplification the validity of which will be further assessed in future work.


#### Convergence with resolution

We apply RT on both the large and the small simulation volume at apastron to verify if the large volume has sufficient resolution in the central region to describe the RT accurately around the ionising source. These simulations are performed with photo-ionization and hydrogen only. These simplifications are not likely to influence the results of this section, however.

Figure 9.10 shows the small box and the corresponding part of the large box side-by side. The size of Delaunay simplices gives a good idea of the resolution difference between the two meshes. One possible effect of insufficient resolution in the WWIR is that features which may be critical to the RT are smoothed away. The WWIR will therefore be less compressed with a less




**Figure 9.9:** Position of the ionization front of hydrogen from the 1D model of Hillier et al. (2001). *Left:* Large box. *Right:* Small box.



**Figure 9.10:** Comparison of the resolution of the large and the small box (inset). In the small box, the secondary star is easily visible whereas in the large box simulation it is almost lost in the WWIR.

resolved contact discontinuity. For the optical depth through the WWIR this is not likely to be very important as it essentially reflects the amount of matter present along a line of sight and is therefore not sensitive to small scale features. For the recombinations, which play an important role in this system, the precise density distribution is important, because the recombination rate scales with the square of the density. If the difference in resolution is important we expect flux from the secondary source to penetrate the WWIR to a lesser degree in the small box.

Figure 9.11 shows the same region for the two simulation boxes after the ionization region of hydrogen has stabilized. Apart from the obvious differences in resolution, the shape and size of the ionised regions are remarkably similar. We therefore conclude that, at least at apastron, the large simulation box captures all relevant detail of the inner region to guarantee plausible results at larger distances from the sources. Now that we have established this important fact,




**Figure 9.11:** Comparison of the RT solution for the central region of the large box (left) and the complete small box (right).

we turn to the physical interpretation of the ionization region. The first interesting observation is that the ionization front is not trapped within the WWIR but rather extends into the primary wind. This results in an approximate hour-glass configuration of moderately ionized (neutral fractions larger than  $10^{-3}$ ) hydrogen.

### The influence of collisional ionizations


The shocks induced by the violent wind-wind interaction are able to heat the gas to temperatures well above those where collisional ionizations become important. Although the SPH simulations do not include anything beyond adiabatic cooling, they give a fair idea of the location and temperatures of the shock-heated gas. In Fig. 9.12 we show a slice through the box with density and temperature in logarithmic scale. There is a strong anti-correlation between density and temperature. The low density regions that are heated by shocks appear dark in density and

bright in temperature. The highest temperatures are seen in the region close to the stars where the gas has not had time to cool by adiabatic expansion.



**Figure 9.12:** *Left:* Number density of hydrogen for the WIND900 box. *Right:* Temperature of the same cut.

As a first approach, we apply collisional ionization equilibrium as an initial condition to the RT and investigate its influence on the final result. In Fig. 9.13 we show three cuts through the WIND900 box: the first with only collisional ionizations, the second with only photo-ionizations and the third a combination of both. All three simulations are hydrogen only.




**Figure 9.13:** Neutral fractions of hydrogen for the WIND900 box. *Left:* Collisional ionizations only. *Centre:* Photo-ionizations only. *Right:* Collisional- and photo-ionizations combined.

Adding photo-ionizations results in a higher ionization fraction in the region right of the stars where the conical WWIR opens up towards the observer. The contribution from the collisional ionizations is primarily in the shocked but tenuous ‘fingers’ that extend into the left part of the simulation volume.

### The influence of helium

Until now, we have treated the gas as if it consisted of pure hydrogen. In reality the helium abundance (in number) is 0.2 relative to hydrogen (Hillier et al. 2001). In this section we include helium and study its influence on the shape and size of the ionization region. One expected result is that, because the recombination rate for doubly ionized helium is roughly an order of magnitude larger than that of hydrogen, the ionized volume will be smaller and dense regions will be able to sustain a higher neutral fraction.


First we will look at the result of photo-ionization only. Figure 9.14 shows the hydrogen ionization fronts superimposed on a cut through the simulation box at apastron, including hydrogen only (black contours) and both hydrogen and helium (white contours). The differences



**Figure 9.14:** Ionisation front position for simulations including hydrogen (black contours) and both hydrogen and helium (white contours) *Left:* Large box. *Right:* Small box.

between the results with and without the addition of helium are dramatic. Including helium reduces the ionised region considerably. Moreover, the ionised region is now confined to the low density secondary wind, implying that the WWIR is capable of shielding the ionising flux from the secondary star.

In Fig. 9.15 we show the fractions of (from left to right) HI, HII, HeI, HeII and HeIII for the large box simulation at apastron (top row) and the small box at apastron (bottom row). The regions of doubly ionised helium overlap almost completely with those of ionised hydrogen. In other words, HIII and HeIII are strongly correlated. As expected, HI and HeI share a similar dependence. The structure of HeII is more involved. It is seen in the WWIR as a barrier between neutral and doubly ionised helium.




**Figure 9.15:** Fractions of (from top to bottom) HI, HeI, HeII and HeIII for the large box simulation at apastron (left) and the small box at apastron (right).



### 9.4.3 Evolution of the ionization structure with phase


In this section we study the shape and size of the ionised regions at two distinctly different parts of the orbit, apastron and periastron. In Fig. 9.2 we can already see that at periastron, the secondary star is embedded inside a cocoon of dense material as it ‘plunges’ through the wind of the primary. At apastron the situation is very different in that the line of sight from an observer on earth towards the secondary star passes through the tenuous cone filled with the un-shocked secondary wind. From these observations, one would expect the ionization structure to be more confined during the periastron passage, and more extended around apastron.

In Figure 9.16 we show the ionized regions used in Madura et al. (2010) and Madura et al. (2011) at apastron (left) and periastron (right) for the large box simulation. These regions have been based on simple geometrical criteria combined with a density threshold. In Figure 9.17



**Figure 9.16:** Ionisation regions at apastron (left) and periastron (right) as used in Madura et al. (2010) and Madura et al. (2011). Spatial dimensions are indicated in semi-major axes (horizontal scale) and degrees on the sky (vertical scale). Although simplified, this approach conforms to our basic intuition that the ionized regions are small at periastron, large at apastron and are confined to the right side of the domain due to shielding of the WWIR.

we show the hydrogen-only ionization-front of our simulations at the same phase as Fig. 9.16. The ionized region at periastron is significantly smaller than at apastron. This is expected as the secondary star is now embedded in the primary wind at a distance less than 2 AU from the primary star. At this location, the wind is much denser (about a factor of 200) than at apastron. However, it is significantly larger than the virtually non-existent ionization region at periastron shown in Fig. 9.16 .




**Figure 9.17:** Ionisation front position for simulations including hydrogen only at apastron (left) and periastron (right).


#### 9.4.4 Preliminary study of RT on AMR data

We have performed a hydrogen-only RT run on the FLASH snapshot shown in Fig. 9.5. This snapshot is located somewhere between apastron and periastron, which is expected to give a result that lies between these two extremes. In Fig. 9.18 we show the stable ionization structure of hydrogen for this snapshot next to the density cut from Fig. 9.5. We can directly see that the only completely neutral part of the simulation is the region of high density around the primary source and a sharp shadow behind that structure. The rest of the volume is ionised to a high degree (with neutral fractions lower than  $10^{-5}$  everywhere except in the completely neutral zone).

This ionization structure is in stark contrast to that obtained with the SPH-based simulations shown above. There the ionization front was effectively halted by the dense WWIR in the apastron snapshots. There are two reasons that could account for this inconsistency. First, the mass loss rate for this density field is a factor of four smaller than the one used in the SPH simulations. Second, the stars are not at their maximum separation and therefore the incident flux on the WWIR is somewhat higher. To see if these two effects are the reason behind the observed behaviour, we performed a simulation with the same mass-loss-rate as used in the SPH simulations, and post-processed a snapshot (almost) at apastron. The resulting ionization structure is shown in Fig. 9.19. The region to the left of the WWIR is now much less ionized than in Fig. 9.5. Nevertheless, ionizing radiation from the secondary star still penetrates the WWIR, which is in accordance with the results of Sect. 9.4.2. A next step will be to incorporate helium in these results as well and see if the ionization front is trapped in the WWIR as observed for the SPH simulations.



**Figure 9.18:** Neutral fractions of hydrogen for the AMR FLASH simulation of Fig 9.5. *Left:* Total hydrogen number density. *Right:* Neutral fraction of hydrogen.



**Figure 9.19:** Same as Fig 9.18, but at apastron and with a mass-loss-rate of  $1 \times 10^3 M_{\odot}/\text{yr}$  for the primary wind.

## 9.5 SUMMARY

- We have performed RT post-processing of both SPH and AMR data, although the latter is very preliminary.
- The SPH results converge between simulations with different scales and we are, therefore, confident that the large box simulations capture the necessary detail close to the sources in order to correctly predict the ionization structure further away from the centre.
- The explorations of Sect. 9.4.2 suggest that the inclusion of collisional ionizations changes

the ionization structure of hydrogen most notably in the under-dense ‘fingers’ that form between the dense shells around the binary system. These regions are effectively shielded from the ionising flux of the secondary star and therefore do not appear if only photo-ionizations are used.

- The influence of the primary star on the overall ionization structure is not expected to be large but this has not been investigated extensively. The results of Sect. 9.4.1 suggest, however, that inclusion of the primary source is a non-trivial endeavor which is beyond the scope of this explorative study. We plan to come back to this issue in future work.
- The inclusion of helium has a profound impact on the size and shape of the ionised region. If the gas is treated to consist only of hydrogen, the WWIR is completely ionised, whereas it is able to trap the ionization front of both hydrogen and helium if the latter element is added in the abundance  $n_{\text{He}}/n_{\text{H}}=0.2$ .
- In Section 9.4.3, we compare the simplified approach used in Madura et al. (2010) and Madura et al. (2011) to full radiative transfer simulations. Although the general trends are similar and conform to expectation, the differences are substantial. The extent of the ionization region at apastron is more confined in the result based on RT. Also, the ionised region obtained with the SIMPLEX method does not disappear at periastron in contrast to the simpler model.
- The preliminary results of Sect. 9.4.4 suggest that the (hydrogen only) ionization front has the tendency to penetrate significantly deeper into the primary wind for the AMR hydro. We have explicitly verified that this is not due to a more tenuous wind.

Given the results presented in this chapter we are confident that the hydrodynamical data available is suitable for post-processing with the SIMPLEX method and that a further study is feasible. From the ionization structures presented here, improved models for the formation of forbidden line emission can be constructed with little effort. The resulting synthetic observations will refine those presented in Madura et al. (2010); Madura et al. (2011) in two important ways. First, they will provide a more detailed and self-consistent spatial constraint for the forbidden line emission. Second, the addition of helium will immediately provide an extra observable in the form of blue-shifted line emission from helium.

## REFERENCES

- Abel, T., Bryan, G. L., & Norman, M. L. 2002, *Science*, 295, 93
- Abel, T., Norman, M. L., & Madau, P. 1999, *ApJ*, 523, 66, (c) 1999: The American Astronomical Society
- Abel, T. & Wandelt, B. D. 2002, *MNRAS*, 330, L53
- Aitken, D. K., Smith, C. H., Moore, T. J. T., & Roche, P. F. 1995, *MNRAS*, 273, 359
- Aldrovandi, S. M. V. & Pequignot, D. 1973, *A&A*, 25, 137
- Allen, C. W. 1973, *Astrophysical quantities* (London: University of London, Athlone Press, —c1973, 3rd ed.)
- Allen, D. A. & Hillier, J. D. 1993, *Proc. of the Astron. Soc. of Australia*, 10, 338
- Altay, G., Croft, R. A. C., & Pelupessy, I. 2008, *MNRAS*, 386, 1931, (c) Journal compilation © 2008 RAS
- Alvarez, M. A., Bromm, V., & Shapiro, P. R. 2006, *ApJ*, 639, 621
- Aragón-Calvo, M. A., van de Weygaert, R., & Jones, B. J. T. 2010, *MNRAS*, 408, 2163
- Aubert, D. & Teyssier, R. 2008, *MNRAS*, 387, 295
- Barber, C. B., Dobkin, D. P., & Huhdanpaa, H. 1996, *ACM Transactions on Mathematical Software*, 22, 469
- Barkana, R. & Loeb, A. 2001, *Phys. Rep.*, 349, 125
- Baugh, C. M. 2006, *Reports on Progress in Physics*, 69, 3101
- Becker, G. D., Rauch, M., & Sargent, W. L. W. 2007, *ApJ*, 662, 72
- Beintema, D. A., Salama, A., & Lorente, R. 2003, in *ESA Special Publication, Vol. 481, The Calibration Legacy of the ISO Mission*, ed. L. Metcalfe, A. Salama, S. B. Peschke, & M. F. Kessler, 215–+
- Bernardeau, F. & van de Weygaert, R. 1996, *MNRAS*, 279, 693
- Bjorkman, J. E. & Cassinelli, J. P. 1993, *ApJ*, 409, 429
- Black, J. H. 1981, *Royal Astronomical Society*, 197, 553, a&AA ID. AAA030.161.007
- Bohren, C. F. & Huffman, D. R. 1984, *Nature*, 307, 575
- Bolton, J. S. & Haehnelt, M. G. 2007, *MNRAS*, 382, 325
- Bowman, J. D. & Rogers, A. E. E. 2010, *Nature*, 468, 796
- Bromm, V., Coppi, P. S., & Larson, R. B. 2002, *ApJ*, 564, 23
- Bromm, V., Yoshida, N., Hernquist, L., & McKee, C. F. 2009, *Nature*, 459, 49
- Brown, R. T. 1971, *ApJ*, 170, 387
- Burgess, A. 1964, *ApJ*, 139, 776
- Burgess, A. & Seaton, M. J. 1960, *MNRAS*, 121, 471
- Cantalupo, S. & Porciani, C. 2011, *MNRAS*, 411, 1678
- Cardelli, J. A., Clayton, G. C., & Mathis, J. S. 1989, *ApJ*, 345, 245
- Cautun, M. C. & van de Weygaert, R. 2011, *ArXiv e-prints*
- Cen, R. 1992, *Astrophys. Journal Supp. Ser. (ISSN 0067-0049)*, 78, 341
- Chesneau, O., Min, M., Herbst, T., et al. 2005, *A&A*, 435, 1043

- Choudhury, T. R., Haehnelt, M. G., & Regan, J. 2009, *MNRAS*, 394, 960
- Ciardi, B. & Ferrara, A. 2005, *Space Sci. Rev.*, 116, 625
- Ciardi, B., Ferrara, A., Marri, S., & Raimondo, G. 2001, *MNRAS*, 324, 381
- Ciardi, B., Stoehr, F., & White, S. D. M. 2003, *MNRAS*, 343, 1101
- Clegg, P. E. & The LWS Consortium. 1999, in *ESA Special Publication*, Vol. 427, *The Universe as Seen by ISO*, ed. P. Cox & M. Kessler, 39–+
- Corcoran, M. F., Ishibashi, K., Swank, J. H., & Petre, R. 2001, *ApJ*, 547, 1034
- Cox, A. N., Mezger, P. G., Sievers, A., et al. 1995, *A&A*, 297, 168
- Currie, D. G., Dowling, D. M., Shaya, E. J., et al. 1996, *AJ*, 112, 1115
- Damineli, A. 1996, *ApJ*, 460, L49
- Damineli, A., Hillier, D. J., Corcoran, M. F., et al. 2008, *MNRAS*, 384, 1649
- Davidson, K. 1997, *New Astronomy*, 2, 387
- Davidson, K. & Humphreys, R. M. 1997, *ARA&A*, 35, 1
- Davidson, K., Martins, J. C., Humphreys, R. M., et al. 2005, *AJ*, 129, 900
- Davidson, K., Smith, N., Gull, T. R., Ishibashi, K., & Hillier, J. D. 2001, *AJ*, 121, 1569
- Davidson, K., Walborn, N. R., & Gull, T. R. 1982, *ApJ*, 254, L47
- de Koter, A., Min, M., van Boekel, R., & Chesneau, O. 2005, in *Astronomical Society of the Pacific Conference Series*, Vol. 332, *The Fate of the Most Massive Stars*, ed. R. Humphreys & K. Stanek, 313–+
- Delone, B. N. 1934, *Bull. Acad. Sci. USSR: Classe Sci. Mat.*, 793, 7
- Dirichlet, G. L. 1850, *Journal für die reine und angewandte Mathematik*, 40, 209
- Du, Q., Faber, V., & Gunzburger, M. 1999, *SIAM Review*, 41, 637
- Duffell, P. C. & MacFadyen, A. I. 2011, *ArXiv e-prints*
- Duncan, R. A. & White, S. M. 2003, *MNRAS*, 338, 425
- Duncan, R. A., White, S. M., Reynolds, J. E., & Lim, J. 1999, in *Astronomical Society of the Pacific Conference Series*, Vol. 179, *Eta Carinae at The Millennium*, ed. J. A. Morse, R. M. Humphreys, & A. Damineli, 54–+
- Dwarkadas, V. V. & Owocki, S. P. 2002, *ApJ*, 581, 1337
- Efron, B. & Tibshirani, R. 1986, *Bootstrap Methods for Standard Errors, Confidence Intervals, and Other Measures of Statistical Accuracy*, Vol. 1 (Institute of Mathematical Statistics), 54–77
- Elvis, M., Wilkes, B. J., McDowell, J. C., et al. 1994, *ApJS*, 95, 1
- Fan, X., Narayanan, V. K., Strauss, M. A., et al. 2002, *AJ*, 123, 1247
- Fan, X., Strauss, M. A., Becker, R. H., et al. 2006, *AJ*, 132, 117
- Ferland, G. J., Korista, K. T., Verner, D. A., et al. 1998, *PASP*, 110, 761
- Ferland, G. J., Peterson, B. M., Horne, K., Welsh, W. F., & Nahar, S. N. 1992, *ApJ*, 387, 95
- Finlator, K., Özel, F., & Davé, R. 2009, *MNRAS*, 393, 1090
- Finlator, K., Özel, F., & Davé, R. 2009a, *MNRAS*, 393, 1090

- Finlator, K., Özel, F., Davé, R., & Oppenheimer, B. D. 2009b, MNRAS, 400, 1049, (c) Journal compilation © 2009 RAS
- Finlator, K., Özel, F., Davé, R., & Oppenheimer, B. D. 2009, MNRAS, 400, 1049
- Franco, J., Tenorio-Tagle, G., & Bodenheimer, P. 1990, ApJ, 349, 126
- Frank, A., Balick, B., & Davidson, K. 1995, ApJ, 441, L77
- Friedrich, M. M., Mellema, G., Alvarez, M. A., Shapiro, P. R., & Iliev, I. T. 2011, MNRAS, 413, 1353
- Frieman, J. A., Turner, M. S., & Huterer, D. 2008, ARA&A, 46, 385
- Fryxell, B., Olson, K., Ricker, P., et al. 2000, ApJS, 131, 273
- Gail, H.-P., Duschl, W. J., Ferrarotti, A. S., & Weis, K. 2005, in Astronomical Society of the Pacific Conference Series, Vol. 332, The Fate of the Most Massive Stars, ed. R. Humphreys & K. Stanek, 317–+
- Gail, H.-P., Duschl, W. J., & Weis, K. 1999, in Astronomische Gesellschaft Abstract Series, Vol. 15, Astronomische Gesellschaft Abstract Series, ed. R. E. Schielicke, 103–+
- Gaviola, E. 1950, ApJ, 111, 408
- Gnedin, N. Y. 2000, ApJ, 535, 530
- . 2004, ApJ, 610, 9
- Gnedin, N. Y. & Abel, T. 2001, New Astronomy, 6, 437
- Gnedin, N. Y. & Hui, L. 1998, MNRAS, 296, 44
- Gnedin, N. Y., Kravtsov, A. V., & Chen, H.-W. 2008, ApJ, 672, 765
- Gomez (Née Morgan), H. L., Dunne, L., Eales, S. A., & Edmunds, M. G. 2006, MNRAS, 372, 1133
- González, R. F., de Gouveia Dal Pino, E. M., Raga, A. C., & Velazquez, P. F. 2004, ApJ, 600, L59
- Gould, R. J. & Thakur, R. K. 1970, Annals of Physics, 61, 351
- Gull, T. R., Nielsen, K. E., Corcoran, M. F., et al. 2009, MNRAS, 396, 1308
- Hackwell, J. A., Gehrz, R. D., & Grasdalen, G. L. 1986, ApJ, 311, 380
- Hägström, O. 2001, in in London Mathematical Society Student Texts (Cambridge University Press)
- Heß, S. & Springel, V. 2010, MNRAS, 406, 2289
- Hillier, J. D., Davidson, K., Ishibashi, K., & Gull, T. R. 2001, ApJ, 553, 837
- Hillier, J. D., Gull, T. R., Nielsen, K., et al. 2006, ApJ, 642, 1098
- Hony, S., Tielens, A. G. G. M., Waters, L. B. F. M., & de Koter, A. 2003, A&A, 402, 211
- Hony, S., Waters, L. B. F. M., & Tielens, A. G. G. M. 2002, A&A, 390, 533
- Hui, L. & Gnedin, N. Y. 1997, Roy. Astron. Soc., 292, 27
- Hui, L. & Haiman, Z. 2003, ApJ, 596, 9
- Hummer, D. G. & Storey, P. J. 1998, MNRAS, 297, 1073
- Icke, V. & van de Weygaert, R. 1987, A&A, 184, 16
- Iliev, I. T., Ciardi, B., Alvarez, M. A., et al. 2006a, MNRAS, 371, 1057
- Iliev, I. T., Mellema, G., Pen, U.-L., et al. 2006b, MNRAS, 369, 1625
- Ishibashi, K., Corcoran, M. F., Davidson, K., et al. 1999, ApJ, 524, 983

- Kessler, M. F., Steinz, J. A., Anderegg, M. E., et al. 1996, *A&A*, 315, L27
- Kimura, Y., Kurumada, M., Tamura, K., et al. 2005, *A&A*, 442, 507
- Kohler, K., Gnedin, N. Y., & Hamilton, A. J. S. 2007, *ApJ*, 657, 15
- Konigl, A. 1982, *ApJ*, 261, 115
- Kruip, C. J. H., Paardekooper, J.-P., Clauwens, B. J. F., & Icke, V. 2010, *A&A*, 515, A78+
- Kunasz, P. & Auer, L. H. 1988, *J. Quant. Spec. Radiat. Transf.*, 39, 67
- Kwok, S., Volk, K. M., & Hrivnak, B. J. 1989, *ApJ*, 345, L51
- Lamers, H. J. G. L. M. & Pauldrach, A. W. A. 1991, *A&A*, 244, L5
- Larson, D., Dunkley, J., Hinshaw, G., et al. 2011, *ApJS*, 192, 16
- Lloyd, S. P. 1982, *IEEE Transactions on Information Theory*, IT-28, no. 2, 129
- Lodders, K. & Fegley, Jr., B. 1999, in *IAU Symposium*, Vol. 191, *Asymptotic Giant Branch Stars*, ed. T. Le Bertre, A. Lebre, & C. Waelkens, 279–+
- Loeb, A. & Barkana, R. 2001, *ARA&A*, 39, 19
- Madura, T., Gull, T. R., Owocki, S., Okazaki, A. T., & Russell, C. M. P. 2010, in *Bulletin of the American Astronomical Society*, Vol. 42, *American Astronomical Society Meeting Abstracts #215*, 426.06–+
- Madura, T. I. 2010, PhD thesis, University of Delaware
- Madura, T. I., Gull, T. R., Owocki, S. P., Okazaki, A. T., & Russell, C. M. P. 2011, *Société Royale des Sciences de Liège*, 80, 694
- Maeder, A. & Desjacques, V. 2001, *A&A*, 372, L9
- Martin, J. C., Davidson, K., Humphreys, R. M., Hillier, D. J., & Ishibashi, K. 2006, *ApJ*, 640, 474
- Martins, F., Schaerer, D., & Hillier, D. J. 2002, *A&A*, 382, 999
- Maselli, A., Ciardi, B., & Kanekar, A. 2009, *MNRAS*, 393, 171
- Maselli, A., Ferrara, A., & Ciardi, B. 2003, *MNRAS*, 345, 379
- McQuinn, M. 2010, in *Astronomical Society of the Pacific Conference Series*, Vol. 432, *New Horizons in Astronomy: Frank N. Bash Symposium 2009*, ed. L. M. Stanford, J. D. Green, L. Hao, & Y. Mao, 65–+
- McQuinn, M., Lidz, A., Zahn, O., et al. 2007, *MNRAS*, 377, 1043
- Meaburn, J. 1999, in *Astronomical Society of the Pacific Conference Series*, Vol. 179, *Eta Carinae at The Millennium*, ed. J. A. Morse, R. M. Humphreys, & A. Damineli, 89–+
- Meaburn, J., Walsh, J. R., & Wolstencroft, R. D. 1993, *A&A*, 268, 283
- Mellema, G., Iliev, I. T., Alvarez, M. A., & Shapiro, P. R. 2006, *New Astronomy*, 11, 374, elsevier B.V.
- Mesinger, A. 2010, *MNRAS*, 407, 1328
- Mesinger, A. & Furlanetto, S. 2007, *ApJ*, 669, 663
- Meyer, C. D. 2001, *Matrix Analysis and Applied Linear Algebra* (SIAM)
- Mihalas, D. 1978, *Stellar atmospheres /2nd edition/* (San Francisco, W. H. Freeman and Co., 1978. 650 p.)
- Mihalas, D. & Weibel Mihalas, B. 1984, *Foundations of radiation hydrodynamics* (New York: Oxford University Press)



- Miles, R. E. 1970, *Izvestilia Akademii Nauk Armlianskoi SSR Matematika*, 5, 263
- . 1974, A synopsis of 'Poisson flats in Euclidian space', *Stochastic Geometry* (New York: John Wiley), 202
- Min, M., Hovenier, J. W., & de Koter, A. 2003, *A&A*, 404, 35
- . 2005, *A&A*, 432, 909
- Miralda-Escudé, J. & Rees, M. J. 1994, *MNRAS*, 266, 343
- Mitchell, R. M. & Robinson, G. 1978, *ApJ*, 220, 841
- Møller, J. 1989, *Advances in Applied Probability*, 21, 37
- Morris, P. W., Waters, L. B. F. M., Barlow, M. J., et al. 1999, *Nature*, 402, 502
- Morse, J. A., Davidson, K., & Ebbets, D. C. 1998, *AJ*, 116, 2443
- Morse, J. A., Kellogg, J. R., Davidson, K., & Ebbets, D. C. 2001, *ApJ*, 548, L207
- Ng, K. 1974, *J. Chem. Phys.*, 61, 2680
- Norris, J. R. 1999, *Markov Chains*, 1st edn. (Cambridge u.a.: Cambridge University Press), 237 S.
- Okabe, A., ed. 2000, *Spatial tessellations : concepts and applications of voronoi diagrams*
- Okazaki, A. T., Owocki, S. P., Russell, C. M. P., & Corcoran, M. F. 2008, *Massive Stars as Cosmic Engines*, 250, 133
- Oppenheimer, B. D. & Davé, R. 2008, *MNRAS*, 387, 577
- Osterbrock, D. E. & Ferland, G. J. 2006, *Astrophysics of gaseous nebulae and active galactic nuclei* (University Science Books)
- Owocki, S. & Gayley, K. 1996, in *Bulletin of the American Astronomical Society*, Vol. 28, *American Astronomical Society Meeting Abstracts*, 1337–+
- Owocki, S. P., Cranmer, S. R., & Blondin, J. M. 1994, *ApJ*, 424, 887
- Owocki, S. P., Gayley, K. G., & Shaviv, N. J. 2004, *ApJ*, 616, 525
- Paardekooper, J.-P. 2010, PhD thesis, Leiden Observatory, Leiden University, P.O. Box 9513, 2300 RA Leiden, The Netherlands
- Paardekooper, J.-P., Kruip, C. J. H., & Icke, V. 2010, *A&A*, 515, 79
- Parkin, E. R., Pittard, J. M., Corcoran, M. F., & Hamaguchi, K. 2011, *ApJ*, 726, 105
- Pawlik, A. H. 2009, PhD thesis, Leiden Observatory, Leiden University, P.O. Box 9513, 2300 RA Leiden, The Netherlands
- Pawlik, A. H. & Schaye, J. 2008, *MNRAS*, 389, 651, (c) Journal compilation © 2008 RAS
- Pawlik, A. H., Schaye, J., & van Scherpenzeel, E. 2009, *MNRAS*, 394, 1812
- Pelupessy, I., Schaap, W. E., & van de Weygaert, R. 2003, *Astronomy and Astrophysics*, 403, 389
- Petkova, M. & Springel, V. 2009, *MNRAS*, 396, 1383
- . 2011, *MNRAS*, 412, 935
- Pittard, J. M. & Corcoran, M. F. 2002, *A&A*, 383, 636
- Press, W., Flannery, B., Teukolsky, S., & Vetterling, W. 1992, *Numerical Recipes in C: The Art of Scientific Computing* (Cambridge University Press)
- Raičević, M. & Theuns, T. 2011, *MNRAS*, 412, L16

- Razoumov, A. O. & Cardall, C. Y. 2005, MNRAS, 362, 1413
- Rijkhorst, E.-J., Plewa, T., Dubey, A., & Mellema, G. 2006, Astronomy and Astrophysics, 452, 907
- Ritzerveld, J. 2005, A&A, 439, L23
- . 2007, PhD thesis, Leiden Observatory, Leiden University, P.O. Box 9513, 2300 RA Leiden, The Netherlands
- Ritzerveld, J. & Icke, V. 2006, Phys. Rev. E, 74, 026704
- Robinson, G., Mitchell, R. M., Aitken, D. K., Briggs, G. P., & Roche, P. F. 1987, MNRAS, 227, 535
- Romano-Diaz, E. & van de Weygaert, R. 2007, ArXiv e-prints
- Schaap, W. E. & van de Weygaert, R. 2000, Astronomy and Astrophysics, 363, L29
- Schaller, G., Schaerer, D., Meynet, G., & Maeder, A. 1992, A&AS, 96, 269
- Semelin, B., Combes, F., & Baek, S. 2007, A&A, 474, 365
- Seneta, E. 1981, Non-negative matrices and Markov chains (Springer-Verlag)
- Shapiro, P. R., Giroux, M. L., & Babul, A. 1994, ApJ, 427, 25
- Shapiro, P. R., Iliev, I. T., Alvarez, M. A., & Scannapieco, E. 2006, ApJ, 648, 922
- Shirokov, A. & Bertschinger, E. 2005, arXiv, astro-ph
- Smith, N. 2006, ApJ, 644, 1151
- Smith, N. 2010, MNRAS, 402, 145
- Smith, N., Davidson, K., Gull, T. R., Ishibashi, K., & Hillier, J. D. 2003a, ApJ, 586, 432
- Smith, N. & Ferland, G. J. 2007, ApJ, 655, 911
- Smith, N. & Gehrz, R. D. 2000, ApJ, 529, L99
- Smith, N., Gehrz, R. D., Hinz, P. M., et al. 2003b, AJ, 125, 1458
- . 2002, ApJ, 567, L77
- Smith, N., Gehrz, R. D., & Krautter, J. 1998, AJ, 116, 1332
- Smith, N. & Morse, J. A. 2004, ApJ, 605, 854
- Springel, V. 2005, MNRAS, 364, 1105
- . 2010, MNRAS, 401, 791
- Sunyaev, R. A. & Chluba, J. 2008, in Astronomical Society of the Pacific Conference Series, Vol. 395, Frontiers of Astrophysics: A Celebration of NRAO's 50th Anniversary, ed. A. H. Bridle, J. J. Condon, & G. C. Hunt, 35--+
- Susa, H. 2006, PASJ, 58, 445
- Theuns, T., Leonard, A., & Efstathiou, G. 1998, MNRAS, 297, L49, (c) 1998 The Royal Astronomical Society
- Theuns, T., Schaye, J., Zaroubi, S., et al. 2002, ApJ, 567, L103
- Tittley, E. R. & Meiksin, A. 2007, MNRAS, 380, 1369
- Trac, H. & Cen, R. 2007, ApJ, 671, 1
- Trac, H. & Gnedin, N. Y. 2009, arXiv, astro-ph.CO
- van Boekel, R., Kervella, P., Schöller, M., et al. 2003, A&A, 410, L37

- van de Weygaert, M. A. M. 1991, PhD thesis, Ph. D. thesis, University of Leiden (1991)
- van de Weygaert, R. 1994, *A&A*, 283, 361
- van de Weygaert, R. & Schaap, W. 2009, in *Lecture Notes in Physics*, Berlin Springer Verlag, Vol. 665, *Data Analysis in Cosmology*, ed. V. J. Martínez, E. Saar, E. Martínez-González, & M.-J. Pons-Bordería, 291–413
- Vandevender, W. H. & Haskell, K. H. 1982, *SIGNUM Newsl.*, 17, 16
- Verner, D. A., Ferland, G. J., Korista, K. T., & Yakovlev, D. G. 1996, *ApJv.*465, 465, 487
- von Zeipel, H. 1924, *MNRAS*, 84, 665
- Voronoi, G. 1908, *J. Reine Angew. Math.*, 134, 198
- Walborn, N. R. & Alvarez, M. A. 1988, *PASP*, 100, 797
- Walborn, N. R., Alvarez, M. A., & Thackeray, A. D. 1978, *ApJ*, 219, 498
- Weigelt, G., Kraus, S., Driebe, T., et al. 2007, *A&A*, 464, 87
- Welty, D. E., Frisch, P. C., Sonneborn, G., & York, D. G. 1999, *ApJ*, 512, 636
- Welty, D. E., Lauroesch, J. T., Blades, J. C., Hobbs, L. M., & York, D. G. 1997, *ApJ*, 489, 672
- Westphal, J. A. & Neugebauer, G. 1969, *ApJ*, 156, L45+
- White, S. M., Duncan, R. A., Lim, J., et al. 1994, *ApJ*, 429, 380
- Whitehurst, R. 1995, *MNRAS*, 277, 655
- Whitelock, P. A., Ebbets, D. C., Koen, C., Roberts, G., & Carter, B. S. 1994, *MNRAS*, 270, 364
- Williams, R. J. R. & Henney, W. J. 2009, *MNRAS*, 400, 263
- Zahn, O., Lidz, A., McQuinn, M., et al. 2007, *ApJ*, 654, 12
- Zaldarriaga, M., Furlanetto, S. R., & Hernquist, L. 2004, *ApJ*, 608, 622

---


# Het eerste licht

HET ontstaan van het heelal heeft de mensheid al sinds de vroegste beschavingen bezig gehouden. Toch heeft het tot de vorige eeuw geduurd voor een coherent model ontwikkeld was voor het ontstaan en de evolutie van het heelal, de zogenaamde *oerknaltheorie*. Deze theorie staat aan de basis van het model dat de vorming van de structuren in het heelal, zoals planeten, sterren en melkwegstelsels beschrijft. Een belangrijk ingrediënt dat nog ontbreekt in de theorie van de vorming van sterrenstelsels is de invloed van de straling van de eerstgevormde sterrenstelsels op het gas in het heelal en hoe dat de vorming van latere sterrenstelsels beïnvloed heeft.

## DE VORMING VAN STRUCTUUR IN HET HEELAL

De oerknaltheorie beschrijft het ontstaan van het heelal vanuit een beginpunt, ongeveer 13,7 miljard jaar geleden, waarin tijd, ruimte, materie en energie hun oorsprong vonden. Tijdens de hierop volgende uitdijning van het heelal ontstonden in de loop van miljarden jaren sterrenstelsels, sterren en planeten, waar de aarde er één van is. De oerknaltheorie ontleent haar fundament aan de algemene relativiteitstheorie die Albert Einstein heeft ontwikkeld aan het begin van de vorige eeuw. Uit deze theorie volgt dat het heelal niet noodzakelijkerwijs statisch is, maar kan uitdijen of inkrimpen. Het heeft lang geduurd voordat een uitdijend heelal een brede wetenschappelijk steun kreeg, maar door de vele waarnemingen die de theorie ondersteunen vormt de oerknaltheorie nu het kader waarbinnen sterrenkundigen werken.

Het eerste observationele bewijs voor de oerknaltheorie werd geleverd door Edwin Hubble in 1929. Hij zag dat sterrenstelsels van ons af bewegen met een snelheid die groter wordt naarmate ze verder van ons af staan. Hieruit volgt dat het heelal niet statisch is maar uitdijt. Recente waarnemingen hebben de uitdijingsnelheid van het heelal met toenemende nauwkeurigheid bepaald. Een ander belangrijk bewijsstuk van de oerknaltheorie wordt gevormd door de *kosmische achtergrondstraling*. Deze werd voor het eerst waargenomen in 1965 door Penzias en Wilson, die voor deze ontdekking de Nobelprijs ontvingen. De achtergrondstraling is ontstaan toen het heelal ongeveer 380 000 jaar oud was. Het extreem hete plasma van vrije protonen en elektronen waaruit het heelal tot dan toe had bestaan was op dat moment ver genoeg afgekoeld om hen te laten recombineren tot waterstofatomen. Dit resulteerde erin dat de hoeveelheid vrije elektronen, dus elektronen die niet in een atoom zitten, drastisch afnam. Vrije elektronen hebben een veel grotere kans dan atomen om een interactie met straling aan te gaan en het daarmee



**Figuur 1:** De vorming van het kosmische web sinds de oerknal, ongeveer 13,7 miljard jaar geleden. Met dank aan Ben Oppenheimer.

van richting te veranderen. Voordat de waterstofatomen zich konden vormen was het dus onmogelijk voor straling om vrij door het heelal te reizen. Omdat de meeste vrije elektronen gebruikt werden voor de vorming van atomen, kon straling uit dit tijdperk ontsnappen in een plotseling doorzichtig geworden heelal. Het licht uit dit markante moment in de geschiedenis van ons heelal nemen we vandaag de dag waar als de kosmische achtergrondstraling.

Een uitdijend heelal zorgt ervoor dat de totale dichtheid van de materie in het heelal steeds kleiner wordt. Immers, de hoeveelheid materie in het heelal blijft dezelfde, maar het heelal zelf wordt groter. Het gas in het heelal zal op deze manier steeds ijler worden en daardoor afkoelen, zonder dat sterren gevormd worden. Hoe kan het dan dat er toch sterrenstelsels en planeten als de aarde zijn? Om die vraag te beantwoorden moeten we eerst bekijken uit welke materie het heelal precies bestaat. Het grootste gedeelte van alle materie in het heelal, ongeveer 83 %, bestaat uit materie waarvan we de aard niet kennen. In tegenstelling tot de ‘gewone’ materie, die we kunnen waarnemen doordat het licht uitstraalt, kunnen we deze materie niet zien. Het wordt daarom ook wel *donkere materie* genoemd. De enige manier waarop het bestaan van donkere materie afgeleid kan worden is door de werking van de zwaartekracht. De precieze aard van donkere materie is voorlopig nog een raadsel. Het vormt echter een belangrijk onderdeel van de gravitationele instabiliteitstheorie, welke het ontstaan van sterrenstelsels beschrijft.

De gravitationele instabiliteitstheorie is gebaseerd op uiterst precieze waarnemingen van de kosmische achtergrondstraling. Deze laten zien dat er in het heel vroege heelal minieme dichtheidsfluctuaties in de materie waren. Omdat de zwaartekrachtswerking van gebieden met hogere dichtheid groter is, wordt materie uit minder dichte gebieden hiernaartoe getrokken. Door dit zichzelf versterkende proces ontstaat er langzaam een spinnenweb-achtige structuur van materie in het heelal, die het kosmische web genoemd wordt. Figuur 1 toont hier een voorbeeld van.

In tegenstelling tot ‘gewone’ materie zoals gas, heeft donkere materie alleen maar interactie

door middel van zwaartekracht. Hierdoor zal het proces van gravitationele samentrekking doorgaan tot een gravitationeel gebonden object gevormd is. Deze objecten worden halo's genoemd. De halo's vormen de plek waar sterren en sterrenstelsels ontstaan. Omdat de dichtheidsfluctuaties waaruit het kosmische web ontstaat op kleine schaal een grotere zwaartekrachtswerking hebben dan op grote schaal, zullen de eerste gravitationeel gebonden objecten klein zijn. Dit proces wordt ook wel *hiërarchische structuurvorming* genoemd. Omdat eerst kleine en daarna grote structuren ontstaan, zijn de grote structuren in het heelal opgebouwd uit kleinere structuren. Zo zijn sterren de eerste objecten die in het heelal gevormd zijn. Later zijn grotere objecten als sterrenstelsels ontstaan, die uit individuele sterren bestaan.

Voor de vorming van objecten van 'gewone' materie zijn naast de zwaartekracht ook andere processen belangrijk, zoals de temperatuur en druk van het gas. Dit maakt het veel moeilijker te beschrijven dan donkere materie. Uiteindelijk zal het gas in de halo's van donkere materie onder invloed van de zwaartekracht steeds verder samentrekken tot uiteindelijk de ineenstorting leidt tot de vorming van sterren en daarmee de geboorte van de eerste lichtbronnen in het heelal.

## HET TIJDPERK VAN REÏONISATIE

De eerste lichtbronnen hebben een grote invloed gehad op de vorming van latere structuren, zoals sterren en sterrenstelsels. De details van dit proces zijn nog grotendeels onbekend. De straling van de eerste bronnen zorgde voor verhitting en ionisatie van het inmiddels afgekoelde, neutrale gas. Dat betekent dat de elektronen en protonen, die tijdens het tijdperk van recombinaatie atomen hadden gevormd, opnieuw gescheiden werden, een proces dat reïonisatie genoemd wordt. Deze fase in de leeftijd van het heelal wordt daarom het *tijdperk van reïonisatie* genoemd. Gedurende dit tijdperk kreeg het heelal zijn huidige aanblik van hooggeïoniseerd gas met hier en daar dichte wolken neutraal gas waarin sterren gevormd worden.

Het tijdperk van reïonisatie is met veel vragen omgeven. Het is bijvoorbeeld onbekend welke lichtbronnen verantwoordelijk waren voor dit proces of wanneer het allemaal precies plaatsvond. Het enige wat we echt zeker weten is dat reïonisatie heeft plaatsgevonden, en dat het ongeveer 12,8 miljard jaar geleden geëindigd is. Omdat reïonisatie zo lang geleden plaatsgevonden heeft, zijn directe waarnemingen met telescopen op zijn zachtst gezegd uitdagend. Er zijn echter wel indirecte observaties die ons begrip van reïonisatie hebben vergroot. Waarnemingen van heldere kernen van sterrenstelsels (zogenaamde quasars) die zeer ver weg staan laten zien dat er rond 12,8 miljard jaar geleden een overgang van neutraal naar geïoniseerd gas plaatsvond. Dit wijst erop dat reïonisatie rond die tijd geëindigd is. Daarnaast zijn er sporen van reïonisatie te zien in waarnemingen van de kosmische achtergrondstraling. Hieruit volgt dat reïonisatie ongeveer 13,3 miljard jaar geleden gebeurde. Waarschijnlijk is reïonisatie dus niet plotseling gebeurd maar verspreid over langere tijd, waarbij eerst het gas rond lichtbronnen geïoniseerd werd. Wellicht zijn we in de nabije toekomst in staat hiervan de eerste waarnemingen te doen. Een voorbeeld van een veelbelovende bron van observationele kennis over reïonisatie is *Low Frequency Array* (LOFAR), een radiotelescoop die grotendeels in Nederland staat. Er worden momenteel pogingen gedaan om met dergelijke radiotelescopen de verdeling van neutraal en geïoniseerd gas tijdens reïonisatie waar te nemen.

Omdat het zo moeilijk is het tijdperk van reïonisatie direct waar te nemen met telescopen, komt de meeste van onze kennis hierover van computersimulaties. Het simuleren van reïonisatie

is echter een grote uitdaging, zelfs met de snelste supercomputers. De reden hiervoor is niet alleen dat er veel ingewikkelde fysische processen gesimuleerd moeten worden, maar ook de enorme schaal waarop het plaatsvindt. Het is noodzakelijk om voor de simulaties een volume te gebruiken dat groot genoeg is om representatief te zijn voor het hele heelal. Alleen op deze manier kun je garanderen dat de melkwegstelsels in de simulatie de diversiteit aan stelsels in het echte heelal benaderen. Dit is noodzakelijk omdat alle soorten en maten sterrenstelsels bijdragen aan het reïonisatieproces. Het is dus ook zaak om sterrenstelsels met een lage massa in de simulatie op te nemen. Door hun grote aantal zijn deze verantwoordelijke voor een behoorlijke bijdrage, al is hun lichtkracht per stuk niet groot. Dat betekent dat in de simulatie naast hele grote ook relatief kleine structuren moeten voorkomen. Hierdoor wordt er met deze simulaties hoge eisen gesteld aan zowel de supercomputer waarop de simulaties uitgevoerd worden als de computercode die alle fysische processen beschrijft. Pas de laatste jaren zijn simulaties in de buurt gekomen van deze eisen.

Er zijn drie belangrijke ingrediënten voor een succesvolle reïonisatiesimulatie. De eerste is donkere materie, die ongeveer 83% van alle aanwezige materie vormt. Omdat donkere materie alleen maar interactie heeft door middel van zwaartekracht is dit relatief eenvoudig te simuleren. Desondanks is het pas recentelijk gelukt om in een groot genoeg volume de vorming van kleine halo's correct te representeren. Het tweede ingrediënt is gas. Zonder gas vormen zich geen sterren en zal er dus ook geen straling zijn om reïonisatie te doen plaatsvinden. Gas is echter een stuk moeilijker te simuleren dan donkere materie omdat er meer fysische processen een rol spelen dan bij donkere materie. Het is bijvoorbeeld belangrijk de invloed van druk en temperatuur van het gas te berekenen. Hierdoor is het tot nu toe onmogelijk geweest de evolutie van gas mee te nemen in reïonisatiesimulaties op grote schaal. In plaats daarvan wordt vaak aangenomen dat het gas de evolutie van de donkere materie volgt. Op grote schaal is dit waarschijnlijk een goede aanname, omdat van alle krachten die werken op het gas, de zwaartekracht domineert. Op kleinere schaal gaat deze aanname echter niet op, met name in de gebieden waar sterren en sterrenstelsels gevormd worden. In plaats van de evolutie van het gas exact te berekenen gebruikt men hier de kennis over de vorming van sterrenstelsels om een uitspraak te kunnen doen over de stelsels die ontstaan in de gravitationeel gebonden halo's van donkere materie. Deze sterrenstelsels vormen de bronnen die het heelal geïoniseerd hebben.

Het derde belangrijke ingrediënt van reïonisatiesimulaties is stralingstransport. Stralingstransport beschrijft hoe licht door het gas reist en wat voor invloed het daarbij uitoefent op het gas. Straling zorgt er niet alleen voor ionisatie maar ook voor verhitting van het gas. Van al deze ingrediënten is stralingstransport het lastigste om te simuleren. Straling reist namelijk met de lichtsnelheid en kan dus heel makkelijk het totale simulatievolume doorkruisen. Dat betekent dat op elk punt in het gesimuleerde heelal de invloed van straling van elk sterrenstelsel meegevoerd moet worden. Gezien de grote hoeveelheid sterrenstelsels is dit een moeilijk probleem voor simulatiecodes. Tot nog toe heeft de complexiteit van stralingstransport ervoor gezorgd dat er maar weinig simulaties gedaan zijn van reïonisatie in een volume dat representatief is voor het heelal. Dit zal waarschijnlijk in de nabije toekomst veranderen.

---

# Het raadselachtige object $\eta$ Carinae

Ook op de schaal van sterren zijn er veel fenomenen waarin stralingstransport een belangrijke rol speelt. In dit proefschrift bekijken we een uitzonderlijke *dubbelster*  $\eta$  Carinae waarvan nog belangrijke vragen onbeantwoord zijn.

$\eta$  Carinae bevindt zich ongeveer 2.3 kpc van ons vandaan in de Carinae nevel en is een van de helderste bronnen van infrarode straling aan de hemel. Dit duidt erop dat bijna de totale lichtkracht van (equivalent aan grofweg  $5.0 \times 10^6$  zonnen) eerst geabsorbeerd en vervolgens weer uitgestraald is door stofdeeltjes. Het overgrote deel van dit stof bevindt zich in de bipolaire *Homunculus* nevel (zie Figuur 2), het resultaat van een 20-jarige periode van zwaar massaverlies en toename van lichtkracht in de 19<sup>e</sup> eeuw, de *Great Eruption* (grote uitbarsting).

## De ‘Great Eruption’ en de Homunculus nevel


In 1837 werd  $\eta$  Carinae plotsklaps een van de helderste sterren aan de hemel. Tegelijkertijd werd er een grote hoeveelheid gas uitgestoten door de dubbelster. De resulterende Homunculus (wat zich laat vertalen als ‘mannetje’, geïnspireerd op lage resolutie waarnemingen) nevel bestaat uit twee lobben die vrijwel sferisch lijken en in 1950 elk een diameter hadden van ongeveer 0.1 parsec (ofwel  $3 \times 10^{12}$  km, grofweg 400 keer groter dan ons zonnestelsel). Door zijn uitdijings-snelheid van ongeveer 550 km/s wordt de nevel elk jaar 0.66% groter. De ‘Great Eruption’ is omgeven door raadselen. Het is onduidelijk welk mechanisme verantwoordelijk is voor de waargenomen verschijnselen en dit blijft tot op de dag vandaag een actief onderzoeksgebied.

## Massa-schattingen van de Homunculus nevel

De totale, als licht uitgestraalde energie tijdens de ‘Great Eruption’ bedraagt zo’n  $10^{49.5}$  ergs ( $3.2 \times 10^{42}$  Joules). Schattingen van de kinetische energie gebaseerd op de expansie-snelheid en massa-schattingen zijn van de orde van  $10^{49}$  ergs, slechts twee ordegrotten kleiner dan een supernova explosie! Zulke extreme hoeveelheden energie zijn niet te verklaren binnen het kader van bekende mechanismen voor straling-gedreven massaverlies. Dit suggereert dat de ‘Great Eruption’ het resultaat is van een nog onbekend sterrenkundig fenomeen.

Omdat het gas in de Homunculus voor het overgrote deel verscholen ligt achter een laag stof, is het zeer lastig om de aard en massa ervan direct te bepalen. Schattingen van de totale massa zijn daarom bijna uitsluitend gebaseerd op waarnemingen van het infrarode licht dat uitgestraald wordt door de stofdeeltjes. Om de waargenomen infrarode straling te vertalen naar een massa





**Figuur 2:** *Links:* Hubble Space Telescope waarneming van  $\eta$  Carinae (Morse et al. 1998). *Rechts:*  $\eta$  Carinae in het IR (op  $24.5 \mu\text{m}$ ; Smith et al. 2002).

van het gas moeten er aannames worden gedaan met betrekking tot de chemische samenstelling van de stof-deeltjes, hun maat en vorm en, uiterst belangrijk, een realistische verhouding van stof en gas. Elk van deze aannames is tot op heden niet uitputtend onderzocht resulterend in een grote spreiding van massa-schattingen in de literatuur van 1 tot 70 zonsmassa's.

### Baanparameters

Een observationeel resultaat dat inmiddels wijde acceptatie geniet is de periode van de dubbelster van 5.538 jaar. Ook lijkt het er inmiddels sterk op dat de baan sterk *eccentrisch* is; de sterren volgen een baan om hun gemeenschappelijke zwaartepunt met de vorm van een sterk afgeplatte ellips.

Er is recentelijk substantiële vooruitgang geboekt in de bepaling van de baanparameters (die de vorm, duur en oriëntatie van de baan beschrijven) door het combineren van *Hubble Space Telescope* waarnemingen met gas-dynamische simulaties. Deze techniek maakte het mogelijk om de draai-as van de dubbelster in drie-dimensies vast te leggen. Uit deze studie blijkt de draai-as (binnen de onnauwkeurigheid van de methode) samen te vallen met de symmetrie-as van de Homunculus nevel, een resultaat dat een sterk verband suggereert tussen de baan van de dubbelster en de nevel. Om deze methode te verfijnen zal stralingstransport moeten worden toegepast op de gas-dynamische simulaties. We komen hier in Hoofdstuk 9 op terug.

## DIT PROEFSCHRIFT

In dit proefschrift behandelen we verschillende aspecten van de stralingstransportcode `SIMPLEX`. In het eerste deel analyseren we de werking van de methode en beschrijven recente uitbreidingen. In het tweede deel passen we `SIMPLEX` toe op verschillende astrofysische problemen.

### Inleiding

We beginnen in het inleidende hoofdstuk met een algemene introductie, waarbij we wat dieper op de stof ingaan dan in deze Nederlandse samenvatting.

### Hoofdstuk 1

In dit hoofdstuk beschrijven we de `SIMPLEX` simulatiecode voor stralingstransport welke de rode draad door dit proefschrift vormt. Om de reis van lichtdeeltjes oftewel *fotonen* door gas te beschrijven gebruikt `SIMPLEX` een speciaal rekenrooster zonder vaste structuur. `SIMPLEX` gebruikt, in plaats van de rechthoekige roosters die in veel traditionele stralingstransportmethodes toegepast worden, een rooster dat gebaseerd is op een verdeling van discrete punten. Deze punten representeren het fysische medium (bestaande uit gas en eventueel stofdeeltjes) en de concentratie van punten is gerelateerd aan de dichtheid van dat medium. Dit betekent dat meer punten komen te staan op plekken waar de gasdichtheid hoger is, wat ervoor zorgt dat de relevante fysische processen beter opgelost kunnen worden. De punten worden verbonden door middel van een triangulatieprocedure, de zogenaamde *Delaunay triangulatie*. Dit houdt in dat in een 3-dimensionale simulatie elk punt verbonden wordt met zijn (ongeveer) 16 dichtsbijzijnde punten. Straling reist over deze triangulatie van roosterpunt naar roosterpunt, waarbij op elk punt de invloed van de straling op het gas berekend wordt. Het heeft verschillende voordelen om straling over de triangulatie te laten reizen. Het eerste voordeel hebben we zojuist besproken: door de dichtheid van punten afhankelijk te maken van de dichtheid wordt de resolutie van de simulatie vergroot in gebieden waar een hogere nauwkeurigheid wenselijk is. Een ander belangrijk voordeel is dat, doordat straling van buur tot buur reist, de rekentijd niet afhankelijk is van het totale aantal lichtbronnen dat aanwezig is in de simulatie, een complicatie waar vrijwel alle traditionele methodes last van hebben. Omdat er in reïonisationsimulaties enorm veel bronnen aanwezig zijn, zal dit aspect kritiek zijn in de uitvoerbaarheid van de nieuwe generatie simulaties.

De nadruk van dit hoofdstuk ligt op de recente aanpassingen aan de code. Het is nu bijvoorbeeld mogelijk om `SIMPLEX` op meerdere computers tegelijk te laten rekenen. Dit is essentieel om simulaties te kunnen doen op supercomputers. Deze aanpak verkort de totale rekentijd en zorgt ervoor dat de grootte van simulaties niet langer beperkt wordt door het geheugen van een enkele computer. Ook beschrijven we een nieuwe manier om fotonen over het rooster te transporteren in gebieden waar er weinig interactie is tussen de straling en het gas, bijvoorbeeld als het gas hooggeïoniseerd is. Deze nieuwe transportmodus zorgt ervoor dat de straling in de correcte richting blijft reizen.

## Hoofdstuk 2

De SIMPLEX methode is een relatief nieuwe simulatie-code die op sommige vlakken sterk verschilt van meer traditionele methodes. Het is daarom nodig dat we vertrouwen hebben in de uitkomsten van SIMPLEX en uitspraken kunnen doen over de mogelijke onnauwkeurigheden van de code. In dit hoofdstuk doen we een uitgebreide analyse van de ongewenste effecten die kunnen optreden als gevolg van de eigenschappen van het rekenrooster dat SIMPLEX gebruikt. Hoewel dat rekenrooster vele voordelen kent, zijn er ook nadelen te noemen. Zo blijkt het dat straling voorkeursrichtingen kan krijgen omdat er variaties optreden in de resolutie van het rekenrooster. Verder blijkt het dat straling zijn originele richting kan vergeten als het te veel stappen over het rooster moet zetten. We introduceren oplossingen voor deze problemen en laten zien hoe ze te voorkomen zijn.

## Hoofdstuk 3

Een zeer belangrijk onderdeel van een SIMPLEX simulatie is het genereren van de roosterpunten waartussen de straling reist. De plaatsing van de punten is afhankelijk van de lokale eigenschappen van het gas. Deze eigenschappen worden meestal geëxtraheerd uit een gasdynamica-simulatie. Er zijn verschillende methoden die de evolutie van gas beschrijven en elk een ander rekenrooster gebruiken. Het is dus belangrijk dat de roosterpunten voor SIMPLEX gemaakt kunnen worden uit allerlei gasdynamica-simulaties. Dit hoofdstuk beschrijft hoe op een efficiënte wijze een rekenrooster gemaakt kan worden waarop stralingstransport met SIMPLEX gedaan kan worden.

## Hoofdstuk 4

De invloed van straling op het gas is tweeledig: niet alleen wordt het gas geïoniseerd, ook de temperatuur verandert onder invloed van licht. In dit hoofdstuk beschrijven we fysische processen die aan de SIMPLEX code zijn toegevoegd. De belangrijkste zijn verhittings- en afkoelingsprocessen waarbij we zowel waterstof als helium gas in acht nemen. Om de verhitte van het gas doormiddel van bestraling met licht correct te beschrijven moeten we meerdere frequenties (de algemenere term voor 'kleuren' die ook toepasbaar is op licht buiten het visuele spectrum) straling meenemen. We beschrijven in detail hoe we dat doen.

## Hoofdstuk 5

In dit hoofdstuk verkennen we een alternatieve formulering van stralingstransport over een Delaunay triangulatie. Het is mogelijk om de reis van fotonen (of andere deeltjes) te beschrijven als een reis over een zogenaamde *graaf*. Een graaf is een verzameling punten met verbindingen ertussen. De Delaunay triangulatie voldoet aan deze beschrijving en dus kan stralingstransport op deze triangulatie worden behandeld met de wiskundige gereedschappen van *grafentheorie*. We laten zien dat deze formulering voor bepaalde soorten simulaties zeer veel efficiënter is dan het 'gewone' stralingstransport dat we in rest van dit proefschrift gebruiken.

## Hoofdstuk 6

Zoals eerder in de samenvatting aan bod kwam, zijn computersimulaties tot op heden de grootste bron van kennis over het tijperk van reïonisatie. Het is daarom van vitaal belang dat verschillende studies, met verschillende methodes dezelfde resultaten produceren. Alleen dan kunnen we vertrouwen hebben dat deze resultaten iets met de werkelijkheid te maken hebben. Een van de onderwerpen waarover momenteel nog geen overeenstemming is bereikt heeft te maken met de volgorde waarop de verschillende componenten van het kosmische web werden geïoniseerd. De meeste studies beschrijven een situatie waarin straling eerst de dichte gebieden waar de bronnen huizen ioniseert, en vervolgens gebieden van steeds minder hoge dichtheid. Er zijn echter ook auteurs die een ommekeer waarnemen waarbij filamenten, structuren van middelmatige dichtheid, het allerlaatst geïoniseerd worden.

In dit hoofdstuk verifiëren we deze afwijkende resultaten door SIMPLEX toe te passen op de gas/donkere materie simulaties van de belangrijkste ‘afwijkende’ studie. Het blijkt dat we de resultaten van de originele studie niet kunnen reproduceren met SIMPLEX en gaan vervolgens dieper in op de mogelijke oorzaak van de waargenomen discrepantie.

## Hoofdstuk 7

De dramatische Homunculus nevel rond de dubbelster  $\eta$  Carinae is in een periode van minder dan 20 jaar tot stand gekomen in een zeer explosieve episode van massaverlies. Om dit fenomeen beter te begrijpen moeten we eerst weten hoeveel energie er voor nodig was om de Homunculus nevel te maken. We doen dit door te kijken naar de *kinetische energie* van de nevel te bepalen. Daarvoor moeten we twee dingen te weten zien te komen, de expansiesnelheid en de totale massa van de Homunculus. De snelheid is vrij nauwkeurig bekend maar de totale massa is nog onzeker. De massa van het gas is heel lastig direct te bepalen omdat het voor een groot deel afgeschermd wordt door een laag stof. De meest veelbelovende methode om er toch achter te komen hoeveel gas er in de Homunculus zit, is door de totale stof-massa te bepalen en daaruit de massa van het gas af te leiden. Het stof is namelijk veel makkelijker waar te nemen.

In dit hoofdstuk passen we de voorgestelde techniek toe op een manier die uitgebreider en preciezer is dan die van andere studies tot op heden. We laten de computer een grote hoeveelheid stofmodellen produceren en kijken systematisch welke het beste de waarnemingen beschrijven. Hieruit leiden we een ondergrens voor de totale massa in de Homunculus af tezamen met voorspellingen voor de chemische samenstelling van het stof. Onze bevindingen dat de Homunculus zwaarder is dan 20 zonsmassa's bevestigen het vermoeden dat de vorming van de nevel een onbegrepen fenomeen is. De energie die opgeslagen ligt in de Homunculus is namelijk vergelijkbaar met die van een *supernova-explosie*. In tegenstelling tot het lot van sterren die een supernova-explosie ondergaan, hebben de sterren in  $\eta$  Carinae de vorming van de Homunculus nevel overleefd. Er is vooralsnog geen mechanisme bekend dat zoveel energie kan vrijmaken in zo'n korte tijd en de ster laat voortbestaan.

## Hoofdstuk 8

Recentelijk zijn er gedetailleerde waarnemingen gedaan met de Hubble Space Telescope aan het binnengebied van  $\eta$  Carinae en met name de bewegingen van het gas aldaar. Door gasdynamische simulaties van de interactie van de *sterrenwinden* van de dubbelster te vergelijken

met deze waarnemingen zijn er belangrijke aspecten van  $\eta$  Carinae aan het licht gekomen. De gedane analyse zou echter verder verfijnd moeten worden om meer informatie uit de waarnemingen te kunnen halen. De belangrijkste tekortkoming van de huidige analyse is dat er geen straling meegenomen wordt. In dit hoofdstuk onderzoeken we of het mogelijk is om de gasdynamische simulaties te bewerken met SIMPLEX teneinde deze tekortkoming teniet te doen. We laten zien dat toepassing van SIMPLEX op de simulaties mogelijk is en bediscussieren verdere verfijningen. De belangrijkste is het toevoegen van helium en in achtname van botsingsionisaties (behalve foto-ionisaties). Vooral de toevoeging van helium heeft een grote invloed op de resultaten.

---

# Curriculum Vitae

Op 8 maart 1981 zag ik in het Brabantse dorp Bergeijk voor het eerst het licht. Ik heb mijn lagere en middelbare schooltijd doorgebracht aan de (toenmalige) Vrije School in Eindhoven en ben Natuur- en Sterrenkunde gaan studeren aan de Vrije Universiteit te Amsterdam. Na het afronden van mijn bachelor met een onderzoek naar fase-discontinuïteiten in Young's experiment bij Taco Visser ben ik de master Astronomy and Astrophysics gaan volgen aan de Universiteit van Amsterdam. Daar heb ik een master-onderzoek naar de chemische samenstelling van de Homunculus nevel rond Eta Carinae gedaan onder de bezielende leiding van Alex de Koter en Michiel Min.

Sinds mei 2006 ben ik verbonden aan de Universiteit Leiden als promovendus in de groep van Vincent Icke. Ik heb mijn werk gepresenteerd op conferenties in Nederland, het Verenigd Koninkrijk, Duitsland en India. Als werkcollegeassistent was ik betrokken bij de colleges Sterrenstelsels en Kosmologie eerst met Marijn Franx en later met Huub Röttgering als docent.

Na mijn promotietijd heb ik een half jaar gewerkt aan het incorporeren van SIMPLEX in het AMUSE ([www.amusecode.org](http://www.amusecode.org)) project. Verder heb ik documentatie geschreven en SIMPLEX met bijbehorende software in een publiekelijk toegankelijke vorm gegoten.



---

# Nawoord

DE laatste bladzijde van dit proefschrift wil ik, geheel volgens traditie, gebruiken om iedereen die bijgedragen heeft aan dit proefschrift te bedanken.

Allereerst wil ik mij richten tot de leden van de theoriegroep. Jullie waren van groot belang voor mijn (wetenschappelijk) welzijn. Jan-Pieter, gelukkig kon jij me altijd uitleggen waar ik mee bezig was, je inzicht en helderheid hebben dit proefschrift mede mogelijk gemaakt. Nicola, I'm convinced you will teach us a thing or two about Eta Car! Inti, het was een genoegen om les van je te krijgen in de wondere werelden van AMUSE, Python en hard-core scripting. Silvia, it has been a pleasure to work with you.

Op dit punt wil ik graag stilstaan bij de mensen die mij enthousiast hebben gemaakt voor de sterrenkunde. Nathalie: soms is het verstandig om het advies van een vrouw serieus te nemen. Dank je wel voor deze en andere wijze levenslessen! Veel dank ben ik verschuldigd aan de sterrenkundigen bij het Anton Pannekoek Instituut: Alex de Koter, Michiel Min, Rens Waters en de andere API's. Een speciale plek wil ik reserveren voor Ralph Weijers die aan de wieg stond van dit proefschrift. Ralph, zonder jou enthousiasme was ik misschien helemaal niet aan een promotie begonnen. Dank voor je vertrouwen en je colleges!

De Sterrewacht is een plek die bol staat van de wetenschap zonder dat de menselijke maat uit het oog wordt verloren. Ik heb veel gehad aan twee pijlers van de Sterrewacht: de computer-groep en het secretariaat. Erik, David, Tycho jullie hebben me vaak heel dom maar tegelijkertijd blij doen voelen. Kirsten, Jeanne, Liesbeth en Anita, het was een geruststelling dat jullie altijd wisten hoe de dingen echt in elkaar steken op de Sterrewacht. Xander, je bent erg belangrijk geweest in de laatste fases van dit proefschrift. Hartelijk dank voor je hulp en diplomatieke tovenarij. Zowel binnen als buiten de Sterrewacht heb ik veel plezier beleefd aan gesprekken en drankjes met Sterrewachters. Ben, Craig, Milan, Richard: in your company I experience a delightful mix of relaxation and uneasiness. Nathan, Arjen, Maryam, Ali, Olivera, Steven, Jeroen, Marcell, Freeke, Thibaut and many others, thanks for tea and chats. Rutger, Jet, Mamta, Anne-Marie, jullie waren fijne kamergenoten met een mooie balans tussen werk en vermaak. Rien en Kakob, de vruchtbare discussies met jullie hebben mij veel opgeleverd. Dank ook voor jullie gedetailleerde commentaar op het geschreven werk. Ook Sven en Sel wil ik bedanken voor de discussies en het lezen.

My colleagues abroad have helped me widen my view in more than one respect. Kristian, your sanity-checks were indispensable. Tom and Jose, thank you for the insight and support. I hope we will crack a problem or two in the future. Ted, it has been wonderful to learn from you why theorists should keep talking to observers!

Het niet wetenschappelijke deel van mijn leven werd de afgelopen jaren veraangenaamd door een groep mensen die ik voor het gemak onder de noemer vrienden en familie schaar.



Jullie zijn allemaal erg belangrijk voor me en ik hoop jullie weer veel te ontmoeten. Wel wil ik in het bijzonder mijn ouders voor hun steun en warmte en mijn zusje voor haar baksels en gedanste knuffels bedanken.

Tamar, dank je wel voor alles wat is geweest en alvast bedankt voor alles wat nog komen gaat.