UNIVERSITEIT
GENT

Faculty of Sciences

Department of Physics and Astronomy

Monte Carlo radiative transfer simulations of galaxies using numerically generated models

Peter Camps

Supervisors:

Prof. Dr. Maarten Baes
Department of Physics and Astronomy, UGent

Prof. Dr. Tom Theuns
Institute for Computational Cosmology
Department of Physics, University of Durham, UK

Other members of the examination committee:

Prof. Dr. Jan Ryckebusch (chair)
Department of Physics and Astronomy, UGent

Prof. Dr. Sven De Rijcke
Department of Physics and Astronomy, UGent

Prof. Dr. Rony Keppens
Department of Mathematics, KU Leuven

Prof. Dr. Mika Juvela
Department of Physics, University of Helsinki

Prof. Dr. Herwig Dejonghe
Department of Physics and Astronomy, UGent

Dr. Alexandr Mosenkov
Department of Physics and Astronomy, UGent

Alex

The force will be with you, always

Contents

Acknowledgements

Not many doctoral theses are written by a grandparent, and I feel privileged to be one of the happy few. My own grandparents made a lasting impression on me. After loosing everything in the First World War, they went on to build a good life for their families, often rowing against the tides, and actively continuing after retirement. Their lifestyle projected the message that has become famous since Yoda delivered it in *The Empire Strikes Back*: "Do. Or do not. There is no try." I can only dream of walking in their footsteps.

While in a philosophical mood, let me explore the question: other than family, what would have been the defining moments in my life? I wrote my first computer program in the era of punched cards and paper tape under the impulse of my math teacher, Paul Joel, who sadly passed away since then. The program ran on a room-sized IBM 370 mainframe with less computing power and less storage capacity than today's cheapest mobile phones. Ever since, the love for software development has been a major driver in my every move.

Some years later, my best friend, Dirk Buffel, encouraged me to read *The World of Null-A*, a 1948 classic science-fiction book by A.E. van Vogt. I opened the book for a quick look just before going to bed, and ended up devouring it cover to cover the same night – causing me to skip classes the next morning. Many more books followed, including all science-fiction works by Isaac Asimov, one of the "Big Three" grandmasters of the genre, along with Robert A. Heinlein and Arthur C. Clarke. Although it is hard to pin down, I believe my interest in science, and astronomy in particular, may originate from the playful exploration of the universe offered in these stories. And then George Lucas created a galaxy far, far away...

When I returned to scientific research after a long stretch in commercial software development, it was only natural for many of my ex-colleagues and friends to ask about the "return of investment" in my new line of work. Clearly, a better understanding of the physical underpinnings of our world may some day lead to practical applications, as it has in the past, although we don't always know when that might happen. The true answer, however, lies elsewhere, and it can be easily found by observing the eagerness of a young child to learn, to explore new horizons, to ask questions, and to find the answers. This is the fundamental drive defining humankind.

It is now time to come to the point. First and foremost, I thank my family. Specifically, my parents for providing an open and nurturing environment; my wife, Chantal, for unconditionally supporting me through my various endeavors over the years, and

for so much more; our children, Steven and Karen, for putting up with a dad who was often absent and for giving back a lot more than I could ever give; our son-in-law Bruno, for being the best son-in-law imaginable; and our grandson Alex, for his smile, and for his enthusiastic and relentless exploration of the world.

My (ex-)colleagues in the commercial world deserve to be mentioned as well, although there are too many to list by name. Thank you for your help in creating the successes without which I would not have switched careers and embarked on this research project.

My new colleagues at the university have been simply wonderful. I've always felt welcome and accepted in spite of the differences in age and background. It has been truly motivating to be part of such a team. Thank you, as well, for your help with various aspects of my research work. A special mention goes to Gert De Geyter, Sébastien Viaene, Bert Vandenbroucke, Sam Verstocken, and Christian Peest, who shared an office with me during part of my stay, and participated in many interesting conversations.

Maarten Baes, my thesis supervisor here in Ghent, provided mountains of support over the years. Maarten, you are by far the best coach I have ever known. Your infinite supply of positive energy serves as an inspiration for everyone around you. Thank you for inviting me into your team, and for everything you've done for me since then. By the same token, I thank Tom Theuns, my thesis supervisor in Durham, for introducing me to the world of large-scale cosmological simulations and for the help and encouragement along the way of this project.

In astronomy, we study the world mostly through the photons or "packets of light" received from distant astrophysical objects. Photons emitted from stars may scatter off interstellar dust grains or molecules before they reach us, performing what resembles a random walk. The SKIRT software program described in this work simulates these effects so that we can better understand our observations. May the random walk of the photons reaching your eyes inspire you to enjoy the wonders of the universe.

Peter Camps, Thursday 26[th] May, 2016

Introduction

1

1.1	Galaxies and interstellar dust

In modern astrophysics, our aim is to comprehend the physical processes underlying the astronomical objects under study. This requires an understanding of the three-dimensional (3D) structure of these objects, and yet, by the very nature of the research subject, the observations are limited to two-dimensional (2D) projections on the plane of the sky. For each wavelength, all radiation arriving at the instrument along a particular line of sight is integrated into a single data point, whether it originated in the object under study or in the foreground. Furthermore, part of the original radiation may be absorbed or scattered by gas molecules and dust grains, and the absorbed energy is re-emitted as thermal radiation at a different wavelength. The presence of even a small fraction of dust can have a substantial impact. In a typical spiral galaxy viewed edge-on, for example, the central dust lane blocks most of the starlight in the UV and optical wavelength range and re-emits the absorbed energy in the infrared and sub-millimeter regime (De Looze et al. 2012a; Verstappen et al. 2013; De Geyter et al. 2014, 2015). While this process adds a layer of complexity, it also provides additional information about the medium interacting with the radiation.

On a philosophical note, when Albert Einstein published his famous theory of gravitation (Einstein 1916), the universe was still widely believed to consist of just one galaxy, the Milky Way. Only in the early 1920's, observations made by Edwin Hubble proved conclusively that some "nebulae" were much too distant to be part of the Milky Way and were, in fact, entire galaxies outside our own (Hubble 1925). Extrapolating from the extreme deep field recently observed by the space telescope named after Hubble (Illingworth et al. 2013), we estimate that there are over 225 billion galaxies in the observable universe (see Fig. 1.1). Given that civilizations have practiced astronomy for several millennia, the study of galaxies is a remarkably young science, and in view of their amazing abundance, these objects are certainly worth our attention.

---

[1] The short quotes at the start of each chapter are taken from the "Notebooks of Lazarus Long" in *Time Enough for Love* by Robert A. Heinlein (1973).

Figure 1.1: Color image of the deepest region of the extreme deep field observed by the Hubble space telescope. This image spans an area of ~ 4.7 arcmin^2 and shows more than 7100 distinct sources. Figure taken from Illingworth et al. (2013).

1.2   The radiative transfer problem

We first take a closer look at the radiative transfer (RT) problem for a dusty system, assuming that radiation traverses the system under study much faster than the time scale on which the system evolves. The problem is then described by the static, panchromatic RT equation, taking into account the physical processes of absorption, anisotropic scattering, and thermal emission:

$$\frac{\mathrm{d}I_\lambda}{\mathrm{d}s}(\mathbf{r}, \mathbf{k}) = j_\lambda^*(\mathbf{r}) \quad - \quad \rho(\mathbf{r})\, \kappa_\lambda^{\mathrm{ext}} I_\lambda(\mathbf{r}, \mathbf{k}) \tag{1.1}$$

$$+ \quad \rho(\mathbf{r})\, \kappa_\lambda^{\mathrm{sca}} \int_{4\pi} I_\lambda(\mathbf{r}, \mathbf{k}')\, \Phi_\lambda(\mathbf{k}, \mathbf{k}')\, \mathrm{d}\Omega' \tag{1.2}$$

$$+ \quad j_\lambda^{\mathrm{d}}(\mathbf{r}), \tag{1.3}$$

where s is the physical distance covered along a particular path; $I_\lambda(\mathbf{r}, \mathbf{k})$ is the radiation intensity at the wavelength λ, the position $\mathbf{r}$, and the direction $\mathbf{k}$; $j_\lambda^*(\mathbf{r})$ is the stellar emission; $\rho(\mathbf{r})$ is the density of the dust; $\kappa_\lambda^{\text{ext}}$, $\kappa_\lambda^{\text{sca}}$, and $\kappa_\lambda^{\text{abs}}$ are the extinction, scattering and absorption coefficients of the dust with $\kappa_\lambda^{\text{ext}} = \kappa_\lambda^{\text{sca}} + \kappa_\lambda^{\text{abs}}$; $\Phi_\lambda(\mathbf{k}, \mathbf{k}')$ is the scattering phase function of the dust; and $j_\lambda^{\text{d}}(\mathbf{r})$ represents the re-emission of the energy absorbed by the dust.

A first complexity in solving this equation arises from the anisotropic scattering term (1.2) which couples all lines of sight in a nontrivial manner. Furthermore, a solution that properly predicts the observations for realistic astrophysical models requires a full 3D treatment (Disney et al. 1989; Witt et al. 1992; Baes and Dejonghe 2001). This is one of the main reasons why most modern RT codes use the Monte Carlo technique; see Sect. 1.3.

The dust emissivity term (1.3) causes major headaches as well. Consider a population of dust grains with uniform size and composition, and assume that it reaches local thermal equilibrium in the radiation field. The dust would then emit as a modified blackbody,

$$j_\lambda^{\text{d}}(\mathbf{r}) = \rho(\mathbf{r})\, \kappa_\lambda^{\text{abs}}\, B_\lambda(T(\mathbf{r})), \tag{1.4}$$

where $B_\lambda(T)$ is the Planck function and $T(\mathbf{r})$ is the local equilibrium temperature, which is determined by the condition that the total amount of absorbed energy equals the total amount of emitted energy:

$$\int \mathrm{d}\Omega \int_0^\infty \kappa_\lambda^{\text{abs}}\, I_\lambda(\mathbf{r}, \mathbf{k})\, \mathrm{d}\lambda = 4\pi \int_0^\infty \kappa_\lambda^{\text{abs}}\, B_\lambda(T(\mathbf{r}))\, \mathrm{d}\lambda, \tag{1.5}$$

where $\int \mathrm{d}\Omega$ integrates over all directions $\mathbf{k}$ and we have eliminated $\rho(\mathbf{r})$ from both sides of the equation. This energy balance condition obviously couples all wavelengths, even under our very basic assumptions.

In many physically realistic scenarios the above simplifying assumptions are not valid. Rather than a uniform population, there is a mixture of dust grains with varying size and composition, and thus varying optical properties (e.g., Weingartner and Draine 2001a; Zubko et al. 2004; Draine and Li 2007). Small dust grains and polycyclic aromatic hydrocarbon (PAH) molecules may be transiently heated by the absorption of individual photons to much higher temperatures, so that there is no thermal equilibrium (e.g., Boulanger and Perault 1988; Draine and Li 2001). The dust emissivity $j_\lambda^{\text{d}}(\mathbf{r})$ then becomes a complex, non-linear function of the radiation intensity.

And finally, most astrophysical objects include components with complex geometries such as clumpy dust distributions, which need to be modeled (see Chapters 2 and 5) and discretized in a realistic RT simulation. Therefore, a crucial component of any modern RT simulation is the dust grid, which divides the dusty medium into a large number of tiny grid cells. Each of these cells is usually characterized by constant dust properties, which serve to determine the optical depth along a path traversing the cell and to calculate the radiation field caused by thermal emission in the cell. The

cell size sets the effective resolution of the simulation. However, since the memory and run-time requirements of the simulation scale with the number of cells, there is a practical limit to the total number of cells allowed in the grid. To optimize simulation resolution in the face of this limitation, the dust cells should be small where the dust density is high or where the radiation field shows a large gradient, and they can be bigger where the dust density is low and the radiation field does not change significantly. This is an active area of investigation in the RT research community (e.g., Brinch and Hogerheijde 2010; Lunttila and Juvela 2012; Saftly et al. 2013, 2014).

1.3   The SKIRT Monte Carlo radiative transfer code

Because of the complexities described in the previous section, most dust RT codes use the Monte Carlo technique to tackle the problem; see, e.g., the reviews by Whitney (2011) and Steinacker et al. (2013). The radiation field is represented as a stream of discrete photon packages. A simulation follows the individual path of each photon package through the dusty medium. The trajectory is governed by various events determined stochastically by drawing random numbers from the appropriate probability distribution. Typically, a photon package is emitted, undergoes a number of scattering events, and is finally either absorbed or leaves the system. A Monte Carlo simulation repeats this loop for each of the emitted photon packages and analyzes the results afterwards. The Monte Carlo technique is conceptually simple and allows efficient RT calculations for complex problems. However, due to the randomization process, the results inherently contain a certain level of Poisson noise.

SKIRT is a state-of-the-art Monte Carlo continuum radiative transfer code for simulating the effect of dust on radiation in static astrophysical systems, i.e. it is assumed that the radiation traverses the system much faster than the time scale on which the system evolves. SKIRT is developed by the research team headed by Prof. Maarten Baes, who also supervised this work. The code is registered in the Astrophysics Source Code Library with identifier *ascl:1109.003*. Earlier versions were described in Baes et al. (2003) and in Baes et al. (2011).

SKIRT handles multiple dust mixtures and arbitrary 3D geometries for radiation sources and dust populations, properly treats absorption and multiple anisotropic scattering by the dust, and offers a variety of simulated instruments for measuring the radiation field from any angle. The code implements the common RT optimization techniques, such as peel-off at emission and scattering events (Yusef-Zadeh et al. 1984), continuous absorption (Lucy 1999; Niccolini et al. 2003), and forced scattering (Cashwell and Everett 1959), and includes novel techniques such as the library mechanism described in Baes et al. (2011).

The current version of SKIRT is fully documented and is publicly available from a GitHub code repository.[2] As part of this thesis work, the SKIRT code was substantially

---

[2] SKIRT documentation: http://www.skirt.ugent.be
SKIRT code repository: https://github.com/skirt/skirt

extended and revised to meet the challenges related to performing RT simulations on numerically generated models, as described in Sect. 1.5. The most important additions and updates to the code are presented as part of this thesis; see Sect. 1.6 for an overview.

1.4	Radiative transfer simulations of analytical models

Many astrophysical objects have a complex geometry on various scales, ranging from arm structures in spiral galaxies (Misiriotis et al. 2000; Fritz et al. 2012), over filaments and clumps in star-forming regions (Goldsmith et al. 2008; Paron et al. 2013; Fallscheer et al. 2013), down to bow-shocks around evolved stars (Decin et al. 2012; Wang et al. 2013). Essentially the only way to properly probe the structure of these systems is to build a computer model of the object under study, have the computer trace the radiation from the sources through the interacting medium in the model towards the mock instruments, and confront the simulation results with actual observations. Recent applications include models of young stellar objects (Wolf et al. 1998), protostellar to protoplanetary disks (Indebetouw et al. 2006; Niccolini and Alcolea 2006), reflection nebulae (Witt and Gordon 1996), molecular clouds (Pelkonen et al. 2009; Steinacker et al. 2005), spiral galaxies (Bianchi 2008; Schechtman-Rook et al. 2012; De Looze et al. 2012a), interacting and starburst galaxies (Chakrabarti et al. 2007; Hayward et al. 2011), and active galactic nuclei (Schartmann et al. 2008; Stalevski et al. 2012a).

1.4.1   The dust energy balance in galaxies

SKIRT is often used to solve such an *inverse* RT problem, where the goal is to recover the actual 3D distribution of radiation sources and dust by fitting the results of RT simulations to observational data. This is a nontrivial task, since the underlying model typically has a large number of free parameters.

For example, De Looze et al. (2012b) used SKIRT to perform a detailed study of the dust energy balance in the Sombrero galaxy M104. The constructed models can reproduce images at optical and near-infrared wavelengths, the observed spectral energy distribution (SED), and the minor axis extinction profiles in the V and R_C bands. A basic model containing only an old stellar population to heat the dust, underestimates the observations of dust emission at infrared wavelengths by a factor of ~ 3 (see Fig. 1.2, left panel). Supplementing this basic model with a young stellar component of low star-formation activity resolves the discrepancy in the dust energy budget of the Sombrero galaxy at wavelengths shortward of 100 μm (see Fig. 1.2, right panel). To account for the increased far-infrared and submm emission beyond 100 μm, an additional dust component distributed in quiescent clumps was included (not shown). This last model with a clumpy dust structure predicts three-quarters of the total dust content to reside in compact dust clouds with no associated embedded sources.
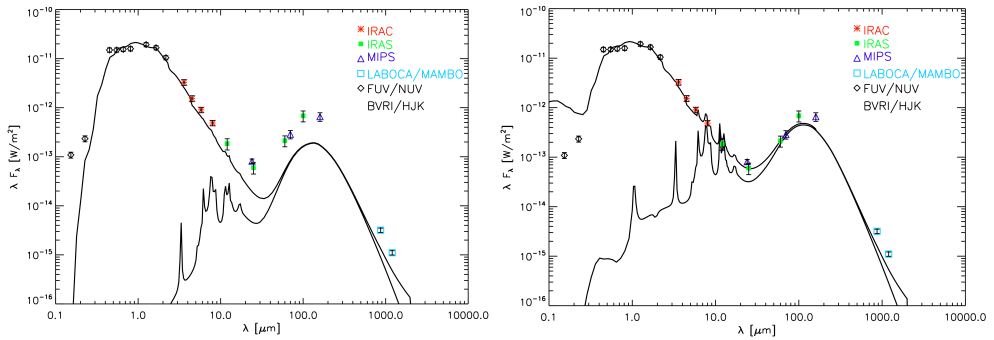
Figure 1.2: The observed fluxes for the Sombrero galaxy M104 (data points) confronted with the modeled SED (upper solid line) and the contribution of the dust (lower solid line) for the model with only an old stellar population (left panel) and for the model including a young stellar component (right panel). Figure taken from De Looze et al. (2012b).

SKIRT has also been used to study the dust energy balance in the edge-on spiral galaxies UGC 4754 (Baes et al. 2010) and NGC 4565 (De Looze et al. 2012a). These studies found a similar inconsistency in the dust energy budget, suggesting that a sizable fraction of the total dust reservoir consists of a clumpy distribution with no associated young stellar sources.

De Looze et al. (2014) construct a high-resolution SKIRT model that accounts for the absorption, scattering, and emission of dust in the face-on galaxy M 51 through a novel technique for deriving the 3D distribution of stars and dust from the observed 2D morphology in various wavelength bands. Extrapolating from this model, they present prescriptions for estimating the contribution of young stars to the global dust heating based on a tight correlation between the dust heating fraction and the specific star-formation rate.

1.4.2   Automated parameter fitting

Rather than using a manual trial and error procedure, which is quite tedious and potentially subjective, it seems preferable to automatically fit the 3D models to observational data. This is a challenging undertaking. As indicated in Sect. 1.4.1, the parameter space is quite large, easily going up to a dozen or more free parameters. Moreover the complexity of the stellar and dust geometries and the anisotropic scattering results in a non-linear, non-differentiable search space with multiple local optima. To make matters worse, noisy models must be fit to noisy data, since both the simulated images (produced by the inherently randomized Monte Carlo technique) and the observed images (created from reduced CCD data) contain a certain level of Poisson noise.

Stochastic search methods are better suited for this complex, high-dimensional optimization problem than classical procedures such as the downhill-simplex method. Genetic algorithms, which work on a set of solutions rather than a single point in parameter space, seem especially attractive because of their ability to handle noisy
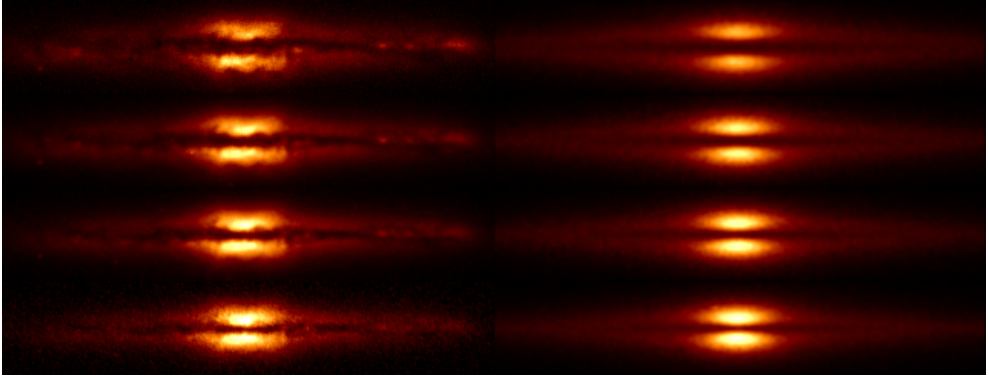
Figure 1.3: Radiative transfer model fits to images of NGC 3987 in four bands (g, r, i, and z from top to bottom). The left column shows the observed images; the right column shows the best fitting FitSKIRT images. Figure taken from De Geyter et al. (2014).

objective functions.

De Geyter et al. (2013) and De Geyter et al. (2014) developed FitSKIRT, a code to automatically fit a 3D model to images of a dusty galaxy in one or more bands. FitSKIRT uses the GAlib genetic algorithm library (Wall 1996) to optimize the output of the SKIRT code. For each new individual s_i created by the genetic procedure, FitSKIRT performs a SKIRT simulation using the parameters $\mathbf{p}_i$ defined by the genes of the individual. The resulting simulated image is convolved with the point spread function (PSF) of the observed image, and a pixel-by-pixel χ^2 value is calculated to compare simulated and observed images. This value is fed back into the genetic algorithm as the quality measure for individual s_i.

De Geyter et al. (2013) use FitSKIRT to determine the intrinsic distribution of stars and dust in the edge-on galaxy NGC 4013 by fitting to a single V-band image. They conclude that the resulting 11-parameter model compares favorably to similar models previously obtained by others. De Geyter et al. (2014) apply FitSKIRT on a sample of 12 edge-on galaxies. This time, a 19-parameter model was simultaneously fitted to four observed images (in bands g, r, i, and z) for each galaxy, as illustrated in Fig. 1.3 for NGC 3987, one of the galaxies in the sample. They show that such multi-band or *oligochromatic* RT fitting can disentangle some of the degeneracies which monochromatic fitting procedures cannot resolve.

1.4.3   Active galactic nuclei

Stalevski et al. (2012a) used SKIRT to investigate the emission of active galactic nuclei (AGN) accretion disks in the infrared domain, modeling the dusty torus as a two-phase medium with high-density clumps and a low-density medium filling the space between the clumps. They found that dust distribution, optical depth, clump size and the spatial arrangement of the clumps in the innermost region all have an impact on the shape of the near- and mid-infrared SED, and argued that

a two-phase dust medium might offer a natural solution to the lack of emission in near-infrared, compared to observed data. The resulting SED library has been made public as described in Stalevski (2012), and the models have been applied in several studies (Popović et al. 2012; Stalevski et al. 2012b, 2016).

<div style="border:1px solid">

1.5   Post-processing of hydrodynamical simulations

</div>

For the applications described in the previous section, the 3D model underlying the RT simulation was constructed using analytical components. Over the years, these computer models have become more sophisticated. For example, smooth analytical geometries have been augmented with stochastic artifacts to represent small-scale structure or clumpiness (Witt and Gordon 1996, 2000; Bianchi et al. 2000; Doty et al. 2005; Stalevski et al. 2012a; Schechtman-Rook et al. 2012). More recently, however, these rather artificial models are being replaced by more realistic, numerical models produced by computer programs that simulate the hydrodynamical evolution of an astrophysical system over time (Juvela and Padoan 2003; Bethell et al. 2004; Stamatellos and Whitworth 2005; Jonsson et al. 2010; Acreman et al. 2010; Hayward et al. 2011; Robitaille 2011; Lunttila and Juvela 2012; Juvela et al. 2012).

Performing RT on a model produced by numerical simulations poses several new challenges. Obviously, there is a need for transferring and interpreting the model data, depending on the simulation formalism used by the hydrodynamical code, and taking into account a variety of file formats. Also, the RT simulation will have to use dust grid structures that can automatically adapt to the unpredictable shapes and large dynamic ranges of the typical density distributions in hydrodynamical simulation results. We will come back to this in Sect. 1.6 and in subsequent chapters of this thesis. We now briefly review some applications in which the SKIRT code, updated as described in this thesis, is used to postprocess the results of hydrodynamical simulations.

Hendrix et al. (2015) model the formation of Kelvin-Helmholtz instabilities in the Orion nebula with a 3D hydrodynamical simulation that includes the dynamics of gas as well as dust. The output of these simulations is then used as input for the SKIRT code to obtain infrared images at several stages of the evolution, which can be compared to the observations. They confirm that a 3D Kelvin-Helmholtz instability is able to develop, and that the formation of the instability is not inhibited by the addition of dust. To get agreement with the observed Kelvin-Helmholtz ripples, the assumed geometry between the background radiation, the billows and the observer is seen to be of critical importance.

Deschamps et al. (2015) study systemic mass loss in Algol binary star systems through a "hotspot" ejection mechanism, where some of the material that is initially transferred from the companion star via an accretion stream is expelled from the system due to the radiative energy released on the gainer's surface by the impacting material. The calculated circumbinary dust mass distribution was fed into the SKIRT code for
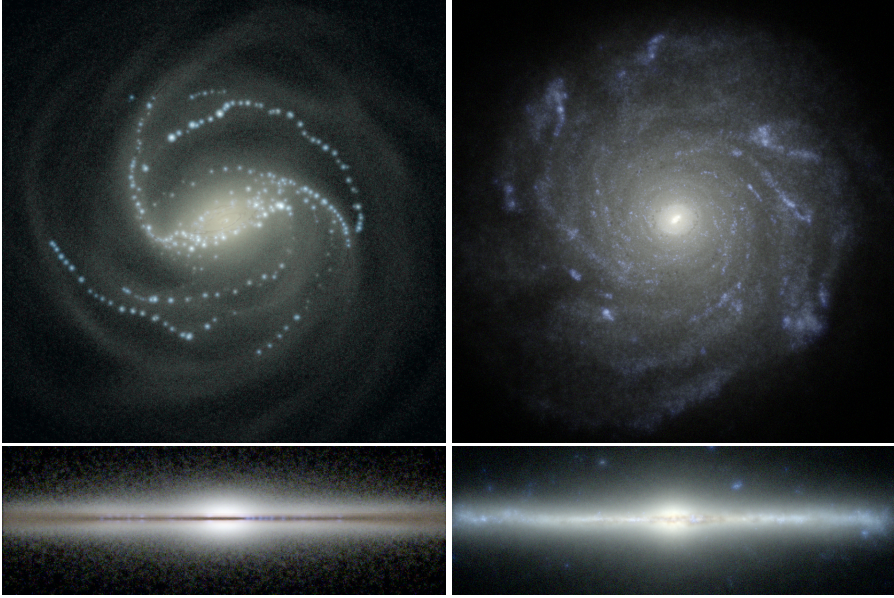
Figure 1.4: Mock face-on (top row) and edge-on (bottom row) views for the R13 galaxy (left column) and for the Eris galaxy (right column). The color images are based on the *r*, *g* and *u* band images produced by the SKIRT code. Figure taken from Saftly et al. (2015); snapshot data provided by Renaud et al. (2013) for R13 and Guedes et al. (2011) for Eris.

obtaining mock infrared observations. They find that systemic mass loss leads to clear observational imprints. Although these signatures are not expected to be found in genuine Algols, they should be present for other, closely related types of binary stars.

Saftly et al. (2015) study the energy balance problem in two artificial Milky Way-like galaxies produced by high-resolution hydrodynamical simulations. They create mock optical edge-on views of these simulated galaxies (using SKIRT; see Fig. 1.4), and then fit the parameters of a basic analytical spiral galaxy model to these images (using FitSKIRT). They find that the dust mass recovered by the fitted models is about three times smaller than the known dust mass of the hydrodynamical input models. This factor is in agreement with previous energy balance studies of real edge-on spiral galaxies (see Sect. 1.4.1). On the other hand, fitting the same basic model to less complex input models (e.g. a smooth exponential disc with a spiral perturbation or with random clumps), does recover the dust mass of the input model almost perfectly. Thus it seems that the complex asymmetries and the inhomogeneous structure of real and hydrodynamically simulated galaxies are a lot more efficient at hiding dust than the rather contrived geometries in typical analytical models. This effect may help explain the discrepancy between the dust emission predicted by RT models and the observed emission in energy balance studies for edge-on spiral galaxies.

1.6   Goals of this thesis

Cosmological simulations form a key ingredient in the study of galaxy formation and evolution. A recent contribution in this area is the EAGLE project (Schaye et al. 2015), which includes a suite of smoothed particle hydrodynamics (SPH) simulations that follow the formation of galaxies and large-scale structure in cosmologically representative volumes of a standard Λ cold dark matter universe. By studying the correspondences and differences between simulation results and observations we hope to further our understanding of the underlying physical processes. Because properties of real galaxies are derived from observed quantities (i.e. fluxes), they may be subject to systematic effects. Making mock observations of simulated galaxies enables direct comparison to observational data.

The key objective of this thesis work has been to use SKIRT for constructing and studying mock observations of the EAGLE galaxies, taking into account the effects of interstellar dust to properly handle a broad wavelength range including ultraviolet, optical, infrared, and submm radiation. While working towards this goal, we actually established a general framework for post-processing the results of hydrodynamical astrophysical simulations of any kind.

At the time when this thesis work started, SKIRT already was one of the leading 3D dust continuum RT codes (Sect. 1.3). It included powerful analytical models and had been successfully applied to various astrophysical problems (Sect. 1.4). However several challenges still needed to be overcome, as summarized in the following paragraphs.

User interface and software design. SKIRT already offered a large number of run-time options, for example to configure the built-in analytical models, and we needed to add many more to support various hydrodynamical simulation types and data formats. We decided to streamline this complexity by reorganizing the code. At the same time, we made sure that the design of the code allowed parallel execution on both shared memory and distributed memory systems, as will be reported by Verstocken et al. (in prep). Chapter 2 describes the SKIRT code in more detail, focusing on the software architecture and design principles introduced as part of this work. We argue that many scientific codes, like SKIRT, can benefit from careful object-oriented design and from a user-friendly interface.

Dust grids. As noted in Sect. 1.2, a proper discretization of the spatial domain is a key element of a RT simulation. Many analytical models are sufficiently smooth to allow a dust grid with a regular, preconfigured distribution of cell boundaries. However, to accommodate the arbitrary and unpredictable shapes and the large dynamic ranges of the typical density distributions in hydrodynamical simulation results, the size and placement of grid cells must automatically adapt to the requirements of each individual model. To this end, SKIRT now offers advanced adaptive dust grids, including the octree and k-d tree based grids implemented and studied by Saftly et al. (2013) and Saftly et al. (2014). These hierarchical grids partition the spatial domain by

recursively subdividing it into cuboidal subcells. The number of subdivisions in each area is automatically adjusted based on the local dust density. As part of this thesis work, we implemented an unstructured dust grid based on Voronoi tesselations of 3D space. Chapter 3 describes this dust grid, and presents a method for computing a straight path between two arbitrary points through a 3D Voronoi grid in the context of a RT simulation. We note that these unstructured grids have certain advantages over hierarchical grids in the context of post-processing hydrodynamical simulations, even if shooting photon packages through them is slower.

Stochastically heated dust grains. As part of the RT simulation, SKIRT self-consistently calculates the dust temperature distribution and the associated thermal re-emission in each dust cell. To properly reproduce the infrared emission of dusty galaxies, a full treatment of stochastically heated dust grains is needed, as noted in Sect. 1.2. The module handling this calculation in SKIRT supported only a single dust model, was very slow, and was extremely hard to maintain. Especially the first two limitations were unacceptable for our purposes, which included simulating many galaxies and testing multiple dust models. Chapter 4 reports on our reimplementation of this module, now fully integrated within the SKIRT code. We provide a self-contained guide for other implementors of such functionality, and we define a benchmark problem for testing the results. We perform this benchmark with six existing RT codes, including SKIRT, and offer insight into the effects of the various approximations and heuristics implemented by the participating codes to accelerate the calculations.

Importing snapshots. Chapter 5 describes the mechanisms added to SKIRT in the context of this work for importing snapshots produced by hydrodynamical simulation codes. These snapshots define the distribution of radiation sources (stars) and/or the interstellar medium (gas and dust) over a spatial grid or as smoothed particles. In case the hydrodynamical simulation does not trace the dust component itself, SKIRT offers some heuristics to derive the a dust density from the overall distribution of the interstellar medium.

Mock observations. The extensions mentioned in the previous paragraphs allow SKIRT to properly calculate mock observations for high-resolution numerical models. Trayford et al. (in prep) study the optical properties of thousands of EAGLE galaxies at redshift $z = 0.1$. They find that optical colors of EAGLE galaxies obtained using SKIRT show good agreement with observations, and that SKIRT reproduces face-on versus edge-on effects on dust attenuation that cannot be properly described by simpler models used in previous work. Chapter 6 reports on our study of the far-infrared and dust-related properties of the EAGLE galaxies. Specifically, we compare the EAGLE galaxies to the galaxies in the *Herschel* Reference Survey, a volume-limited sample of about 300 normal galaxies in the local universe, and we investigate the effect of adjusting certain parameters in our post-processing procedure.

Finally, in Chapter 7, we offer overall conclusions and we propose some topics for future research.

SKIRT code design[4]

<div style="text-align: right; font-size: 3em;">2</div>

ALWAYS LISTEN TO EXPERTS. THEY'LL TELL YOU
WHAT CAN'T BE DONE, AND WHY. THEN DO IT.

In this chapter, we discuss the software architecture and design principles that underpin the latest version of SKIRT, the 3D continuum dust radiation transfer code developed in our research group and introduced in Sect. 1.3. The configuration for a particular simulation is defined at run-time through a user-friendly interface suitable for both occasional and power users. These capabilities are enabled by careful C++ code design. The programming interfaces between components are well defined and narrow. Adding a new feature is usually as simple as adding another class; the user interface automatically adjusts to allow configuring the new options. We argue that many scientific codes, like SKIRT, can benefit from careful object-oriented design and from a friendly user interface, even if it is not a graphical user interface.

2.1 Introduction

In Chapter 1 we introduced the SKIRT code and the radiative transfer (RT) problem it solves, and we briefly described some research applications. We also referred to previous work that covers the Monte Carlo RT techniques implemented in the code. In this chapter, we focus on the software design choices involved with setting up the simulation model in SKIRT. Many scientific codes require a user to hard-code the model makeup for each distinct problem. In contrast, our strategy with SKIRT in recent years has been to continuously add new features without removing existing capabilities. Consequently, SKIRT now offers a wealth of configurable components that are ready to use without any programming at all, especially in the areas where the code has been most often applied. An ad-hoc approach to including all of this functionality would have lead to source code that is hard to understand, maintain, and use. Instead we developed a modular, generic software architecture that can support the wide range of built-in components and options in SKIRT in a developer- and user-friendly way.

In Sect. 2.2 we first provide an overview of SKIRT's features, including the user interface for configuring a particular simulation. In Sect. 2.3 we then discuss the

---

[4] Published as Camps and Baes (2015). My contribution includes the design and implementation of the software architecture and user interface described in this chapter, in addition to various optimizations to the existing core radiative transfer code and to the simulation items listed in Tables 2.1 through 2.4.
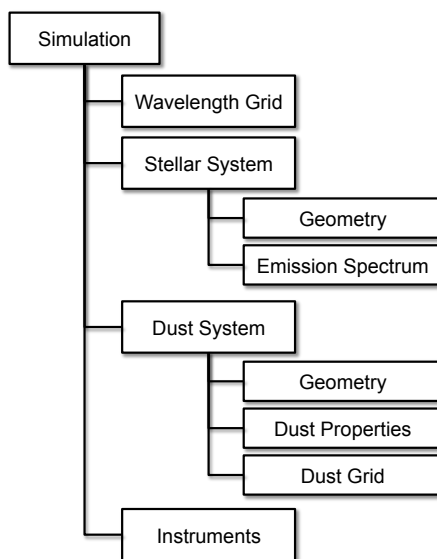
Figure 2.1: A schematic representation of the items to be configured for a particular SKIRT simulation.

design goals for the latest revision of the code, we describe the overall architecture, and we zoom in on a few key aspects of the design, such as the mechanism that automatically adjusts the user interface to accommodate new features. In Sect. 2.4 we finally argue that many scientific codes, like SKIRT, can benefit from careful object-oriented design and from a friendly user interface.

2.2	Features

SKIRT is a console application and it is completely written in C++. It can easily be deployed on any Unix system, including for example Ubuntu and Mac OS X. The code has no compile-time options; all built-in components and capabilities are configured at run-time. The first-time user would use the interactive query and answer mechanism (in a terminal window) to configure a particular simulation. The smart mechanism guides the user through all possible options, narrowing down the possibilities based on earlier choices. The complete configuration for the simulation is then saved as a SKIRT parameter file in XML (eXtensible Markup Language) format, which can be easily viewed and adjusted in a regular text editor, even by an occasional user.

2.2.1 Configuring a simulation

Figure 2.1 illustrates the structure of a SKIRT simulation. Each of the building blocks offers several alternatives and options that can be configured in the parameter file. At the top level, for example, SKIRT supports two simulation types: oligochromatic and

panchromatic. An *oligochromatic* simulation operates at just one or a small number of distinct wavelengths. It handles absorption and scattering by dust grains, but it doesn't support thermal dust emission. There is no way to compute the dust temperature without integrating the absorbed radiation energy over an appropriate wavelength range. This basic simulation mode is appropriate for studying optical wavelengths, since the dust emission is negligible there. A *panchromatic* simulation operates over a broad range of wavelengths. These simulations can handle thermal dust emission as well as absorption and scattering, and thus many more options need to be configured.

The wavelength grid for an oligochromatic simulation is simply a short list of distinct wavelengths. A panchromatic simulation employs a grid over a range that typically extends from UV to millimeter wavelengths. SKIRT offers a plain logarithmic wavelength grid and a nested logarithmic wavelength grid, providing a higher resolution in some subset of the range. The user can specify the wavelength range and the number of grid points. Alternatively SKIRT can read a custom wavelength grid from a text data file that lists the grid points.

The spatial distribution of radiation sources and dust is obviously an important part of the simulation model. For this purpose SKIRT offers a number of predefined geometries; the most important ones are listed in Table 2.1. Each geometry defines a spatial density distribution, which can be used for radiation sources as well as dust components. Choices include a point-like source and various theoretical models for distributed densities with spherical, cylindrical, or no symmetries. Decorator geometries adjust another geometry by shifting its center to an arbitrary location, deforming a spherical geometry into a spheroidal or triaxial distribution, or adding clumps in random locations. Other geometries can import a density distribution from a data file. Anisotropic radiation sources are supported as well. Multiple geometries can be combined in arbitrary ways, enabling the construction of complex models.

The stellar system describes the radiation sources in the simulation model. For each geometry, the configuration defines the emission spectrum and the luminosity. Table 2.2 lists the built-in spectral energy distributions (SEDs), including several well-known parameterized SED families, and the option to import an SED from file. The amount of radiation can be specified through the bolometric luminosity or through the spectral luminosity at the center of a standard wavelength band. SKIRT also includes specialized stellar system components to import a snapshot from a hydrodynamic simulation using smoothed particles (SPH) or an adaptive mesh (AMR). In this case, both the spatial distribution and the emission spectrum in each location are extracted from the input data, for example using a Bruzual-Charlot model based on stellar age and metallicity.

Similarly, the dust system describes the spatial distribution and the properties of the dust in the model. A dust system can have multiple components, each with its own geometry and dust characterization. The amount of dust in each component can be defined simply as a total mass, or by specifying the optical depth along a

Table 2.1: A selection of built-in components for defining the spatial distribution of radiation sources and dust components.

(#) Note: for compactness, "GeometryDecorator" is replaced by "GeoDec" in the decorator names.

Geometries	
Spherically symmetric	
PointGeometry	single point
PlummerGeometry	classical Plummer sphere (Plummer 1911; Dejonghe 1987)
SersicGeometry	spherical model with a Sérsic surface brightness profile (Sérsic 1963; Ciotti and Bertin 1999)
EinastoGeometry	spherical model with an Einasto density profile (Einasto 1965; Retana-Montenegro et al. 2012)
GammaGeometry	spherical model with a gamma density profile (Dehnen 1993; Tremaine et al. 1994)
ShellGeometry	spherical shell where the density behaves as a power law between an inner and an outer radius
Axisymmetric	
ExpDiskGeometry	optionally truncated exponential profile in both radial and vertical directions (van der Kruit 1986)
RingGeometry	ring with gaussian profile in the radial direction and exponential fall-off in the vertical direction
TorusGeometry	torus with radial power-law profile within opening angle (Stalevski et al. 2012a; Granato and Danese 1994)
GaussianGeometry	model with gaussian distribution in the radial and the vertical direction
MGEGeometry	geometry defined by a Multi-Gaussian Expansion (Emsellem et al. 1994; Cappellari 2002)
No symmetries	
SPHGeometry	density distribution defined by a set of smoothed particles, imported from a data file
AdaptiveMeshGeometry	density distribution defined over an adaptive mesh grid, imported from a data file
VoronoiGeometry	density distribution defined over a Voronoi tessellation, imported from a data file
Decorators	
OffsetGeoDec[#]	applies an arbitrary offset to any other geometry
RotateGeoDec	applies an arbitrary rotation to any other geometry
ClumpyGeoDec	replaces a portion of the mass in any geometry by randomly placed clumps
SphericalCavityGeoDec	carves out a central spherical cavity from any other geometry
SpheroidalGeoDec	transforms any spherically symmetric geometry to a spheroidal counterpart
TriaxialGeoDec	transforms any spherically symmetric geometry to a triaxial counterpart
SpiralStructureGeoDec	overlays a spiral arm perturbation onto any axisymmetric geometry (Misiriotis et al. 2000)
Anisotropic	
LaserGeometry	point source that emits in a single direction only
NetzerGeometry	point source with the anisotropic radiation profile of an accretion disk (Netzer 1987)
CubBackgroundGeometry	the surface of a cube from which radiation escapes in the inward direction only, such that all points inside the cube see the same intensity

Table 2.2: A selection of built-in components for defining the spectral energy distribution of radiation sources.

SEDs	
Simple	
BlackBodySED	classical black body spectrum for a given temperature
PegaseSED	SED templates for elliptical, lenticular and spiral galaxies (Fioc and Rocca-Volmerange 1997)
QuasarSED	SED template for a quasar (Schartmann et al. 2005)
StarburstSED	SED templates for a starbursting stellar population with given metallicity (Leitherer et al. 1999)
SunSED	the solar spectrum
FileSED	arbitrary spectrum imported from a data file
Families	
BruzualCharlotSED	stellar population SEDs parameterized on metallicity and age (Bruzual and Charlot 2003)
MarastonSED	stellar population SEDs parameterized on metallicity and age (Maraston 1998)
KuruczSED	stellar SEDs parameterized on metallicity, effective temperature and surface gravity (Kurucz 1993)
MappingsSED	starbursting region SEDs parameterized on metallicity, compactness, pressure and covering factor (Groves et al. 2008)

Table 2.3: A selection of built-in components for defining the properties of the dust mixture.

Dust Mixes	
Turn-key dust mixes	
DraineLiDustMix	mixture of graphite, silicate and PAH grains (Draine and Li 2007)
MRNDustMix	mixture of graphite and silicate grains (Mathis et al. 1977; Weingartner and Draine 2001a)
WeingartnerDustMix	mixture of graphite, silicate and PAH grains (Weingartner and Draine 2001a)
ZubkoDustMix	mixture of graphite, silicate and PAH grains (Zubko et al. 2004)
Custom dust mixes	
ConfigurableDustMix	custom-configured dust mix given a list of grain compositions and grain size distributions
Grain compositions	
DraineGrainComp	optical and calorimetric properties for graphite, silicate and PAH grains (Draine and Li 2007)
DustEmGrainComp	any of the dust grain properties provided with the DustEM code (Compiègne et al. 2011)
ForsteriteGrainComp	Forsterite crystalline silicate grain properties (Fabian et al. 2001; Min et al. 2005; Suto et al. 2006)
EnstatiteGrainComp	Enstatite crystalline silicate grain properties (Jaeger et al. 1998; Min et al. 2005)
Grain size distributions	
PowerLawGrainSize	modified power-law grain size distribution with a form inspired by Compiègne et al. (2011)
LogNormalGrainSize	modified log-normal grain size distribution with a form inspired by Compiègne et al. (2011)

Table 2.4: A selection of built-in components for defining the spatial discretization of the dust medium.

Dust Grids	
Regular grids	
Sphere1DDustGrid	spherically symmetric grid (each cell is a spherical shell)
Cylinder2DDustGrid	axisymmetric grid (each cell is a cylindrical shell)
CartesianDustGrid	regular cartesian grid (each cell is a cuboid)
Meshes for regular grids	
LinMesh	mesh with equidistant cell boundaries
PowMesh	mesh with cell boundaries spaced according to a power-law
LogMesh	mesh with logarithmically spaced cell boundaries
Hierarchical grids	
OctTreeDustGrid	octree grid that recursively subdivides cuboidal nodes into eight sub-nodes (Saftly et al. 2013)
BinTreeDustGrid	*k*-d tree grid that recursively subdivides cuboidal nodes into two sub-nodes (Saftly et al. 2014)
Unstructured grids	
VoronoiDustGrid	unstructured dust grid based on a Voronoi tesselation of 3D space (Camps et al. 2013b)

particular axis. The optical and chemical properties of the dust in each component can be configured in great detail, as listed in Table 2.3 and further described in Sect. 2.2.2. Again, there are specialized dust system components to import a snapshot from a hydrodynamic simulation (SPH or AMR). The spatial distribution of the dust is now calculated from the gas density in the input data, assuming that the amount of dust is proportional to the metal fraction in the gas, except in areas where the gas is too hot to form dust.

The dust system also configures the dust grid, i.e. the computational structure that is used to discretize the spatial domain under study. The grid partitions the spatial domain in individual dust cells, and all physical variables (dust density, optical properties, radiation field, dust temperature) are considered to be constant in each dust cell. During the RT simulation, photon packages propagate through the grid and interact with particular cells according to randomly generated events. Since memory requirements and computation time rapidly increase with the number of dust cells, a good grid has smaller cells in areas that require a higher resolution, and larger cells elsewhere. Table 2.4 lists the dust grids built into SKIRT. The regular spherical, cylindrical, and cartesian grids are perfect for simple models with the corresponding symmetries. The grid points are defined by configuring a "mesh" for each spatial dimension (respectively 1, 2 or 3). Linear meshes have equidistant grid points; logarithmic and power-law meshes place (much) smaller bins in the central areas of the model. Most state-of-the-art simulations use complex 3D models, however, and SKIRT offers a choice of smart structured and unstructured grids to help optimize accuracy and performance. This is an active field of study in our research group, as described in Sect. 2.2.3.
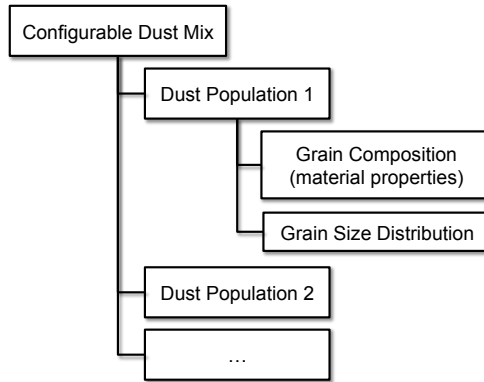
Figure 2.2: A schematic representation of a dust mix containing multiple dust populations. Each population describes a particular type of grain.

Finally, the configuration is completed with a number of synthetic instruments, which collect and write down information about the simulated radiation received at some specified viewpoint. The *SED* instrument outputs the spectral energy distribution of the received flux as a text file that can easily be plotted. The *frame* instrument collects a complete 3D data cube (a rectangular frame of flux samples at each simulated wavelength) and outputs the result as a FITS file (Flexible Image Transport System), enabling the use of the standard visualization and data manipulation tools. The instruments can treat the radiation differently depending on its source; for example direct radiation, scattered radiation and dust emission can be recorded separately. The default instruments assume that the distance to the model is very large, so that they can use parallel projection. The perspective instrument, however, can be placed anywhere, even inside the model. It is mostly used to create animations by specifying an instrument per movie frame, with slightly varying position and/or angles.

2.2.2   Dust properties

The dust system in a SKIRT simulation can hold multiple dust components, each with their own spatial distribution and their specific dust characterization. The dust properties applying to a particular dust component are bundled in a building block called a *dust mix*. SKIRT offers several options to configure a dust mix, ranging from very simple to quite involved. Table 2.3 lists some of the choices. Each of the turn-key dust mixes implements a particular dust model described in the literature, usually including some specific combination of silicate grains, graphite grains, and polycyclic aromatic hydrocarbon (PAH) molecules, with properties listed in data files and/or approximated by formulae. These dust mixes can be configured simply by supplying their name.

Alternatively, the user can configure a custom dust mix from basic building blocks, as illustrated in Fig. 2.2. A configurable dust mix holds a distinct dust population for each type of grain material in the mix. For each type of grain material, the
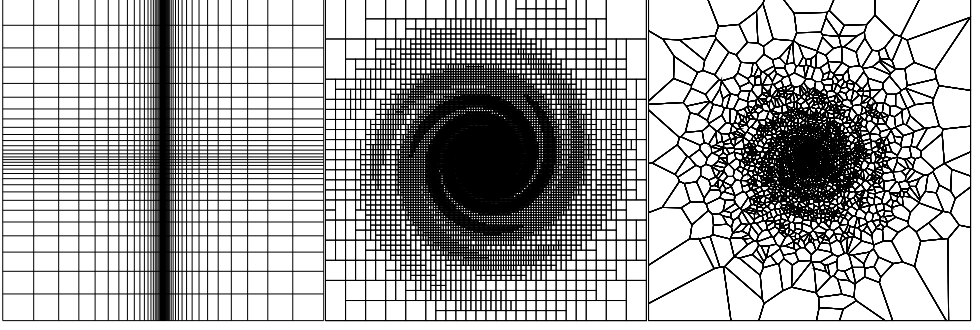
Figure 2.3: A planar cut through three of SKIRT's dust grids. *Left*: a regular grid with cell sizes distributed logarithmically on the horizontal axis and according to a power-law on the vertical axis. *Middle*: a cuboidal *k*-d tree grid with cell sizes that are adjusted to the dust density distribution in a simple spiral galaxy model. *Right*: an unstructured Voronoi grid where the generating sites are placed randomly following the dust density distribution of the same spiral galaxy model. Note that a planar cut through a 3D Voronoi tessellation is usually *not* a 2D Voronoi tessellation.

dust population specifies the optical and calorimetric material properties and a grain size distribution function. Optical dust properties include the scattering and absorption coefficients $\kappa^{\text{sca}}(\lambda, a)$ and $\kappa^{\text{abs}}(\lambda, a)$, and the asymmetry parameter $g(\lambda, a)$ determining the scattering phase function $\Phi_{\lambda, a}(\mathbf{k}, \mathbf{k}')$, for a range of wavelengths λ and a range of grain sizes a. Calorimetric properties include the heat capacity $C(T)$ or equivalently the internal energy $U(T)$ of the dust grain material at a range of temperatures T, and the bulk mass density ρ_{bulk} of the material. Several sets of standard material properties and often-used size distributions are built-in to SKIRT, as illustrated in Table 2.3, and new choices can be easily added.

Configurable dust mixes allow users to experiment with new dust models, or to specify a different spatial distribution for a particular dust population (by splitting it off into a separate dust mix corresponding to a dust component with its own geometry).

2.2.3   Dust grids

The construction of a proper dust grid is a key aspect of a RT simulation. Many astrophysical models feature small structures, such as dust clumps or star forming regions, which require a lot of cells to resolve properly. To minimize memory requirements and computation time, the grid should be adapted to the spatial structure of the model. Therefore, in addition to regular grids similar to the one shown in the leftmost panel of Fig. 2.3, SKIRT offers several types of adaptive grids, including the *k*-d tree and the Voronoi grid shown in middle and rightmost panels of the same figure.

Starting from a cuboidal root cell that spans the complete spatial domain, a typical adaptive mesh refinement (AMR) scheme recursively subdivides each cell into $a \times$

$b \times c$ cuboidal subcells until sufficient resolution has been reached in each region. In the special case where $a = b = c = 2$, each cell is subdivided into eight subcells and the data structure is called an octree. The octree implementation in SKIRT was optimized in the context of RT as reported in Saftly et al. (2013).

A k-d tree (k-dimensional tree) is a space-partitioning data structure where each cell is recursively split into just two subcells along a particular hyperplane. In 3D space, i.e. with $k = 3$, a k-d tree is similar to an octree. In fact, any octree of depth n has an equivalent k-d tree of depth $3n$. For each octree level, the k-d tree uses three consecutive levels with mutually orthogonal dividing planes. However, while an octree forces all eight subcells to be created at the same time, a k-d tree allows more fine-grained control over which of the two initial subcells are subdivided further. As reported in Saftly et al. (2014), this property gives k-d tree grids a relevant advantage over octree grids in the context of RT. The middle panel of Fig. 2.3 shows a cut through a k-d tree grid with cell sizes that are adjusted to the dust density distribution in a simple spiral galaxy model.

Adaptive grids with cuboidal cells have become popular mainly because of their relative ease of implementation. But there is no a priori reason to assume that the cuboidal cell form is optimal. To the contrary, the strict coordinate-plane alignment of cell boundaries makes it hard to represent steep gradients in arbitrary directions, raising the number of cells needed to properly resolve clumpy features. One could consider constructing a grid using polyhedra instead of cuboids, but in general this seems a daunting task. Fortunately George Voronoi (Voronoi 1908) provided a specific way of partitioning 3D space into convex polyhedra. The mathematical properties of a Voronoi tesselation greatly facilitate implementation of an unstructured Voronoi grid in the context of RT, as described in Camps et al. (2013b). Further research should determine whether Voronoi grids can indeed resolve astrophysical structures using less cells than cubodial adaptive grids. The rightmost panel of Fig. 2.3 shows a cut through a 3D Voronoi grid. Cells are placed randomly according to the dust density distribution of a simple spiral galaxy model.

2.2.4 User interface

The complete configuration for a particular SKIRT simulation is stored in a single parameter file, called a *ski file* (pronounced "skee file"). In view of the many features, options and interdependencies described in the previous sections, the contents of a ski file can become quite complex. To deal with this complexity, we opted for a file format based on XML (eXtensible Markup Language). This format has several advantages. XML elements can be nested to create hierarchies of features and options that reflect the natural makeup of the simulation's configuration. XML is stored as plain text, so it can be easily viewed and adjusted in a regular text editor, even by an occasional user; the human-readable XML tags make the format self-explanatory to a large degree. And finally, existing ski files remain compatible when new features are added (with appropriate defaults), or can be automatically upgraded when the structure changes in an incompatible way.
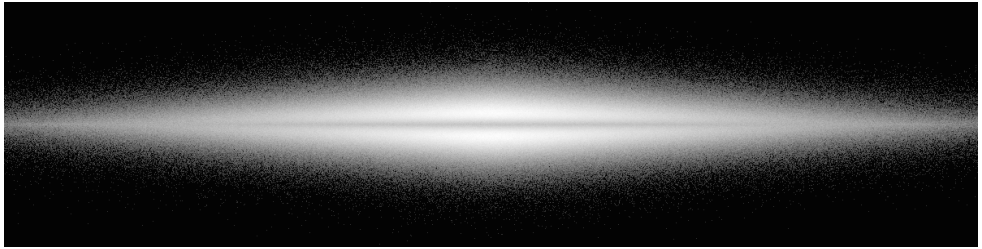
Figure 2.4: The output of a SKIRT simulation for a very simple spiral galaxy model with an instrument that registers the total flux of an edge-on view. The SKIRT parameter file for this simulation is shown in Fig. 2.6.

To perform a simulation, the user starts the code in a terminal window, supplying the name of the relevant ski file on the command line. The code runs fully unattended and all results are written to output files. A small number of command line options allow overriding some defaults in the run-time environment, such as the number of parallel threads or the location of input and output files. The makeup of the simulation itself is fully defined in the ski file. When SKIRT is started without any command line arguments, it enters an interactive query and answer mode that guides the user through the process of creating a new ski file.

To illustrate how this works, we configure a SKIRT simulation for a simple spiral galaxy model, with an instrument that produces the edge-on view shown in Fig. 2.4. The query and answer session in the terminal window is illustrated in Fig. 2.5. The resulting ski file is shown in Fig. 2.6, and a pretty-printed version is shown in Fig. 2.7.

When configuring a particular type of simulation for the first time, the query and answer mechanism guides the user through all possible options, narrowing down the possibilities based on earlier choices. This is similar to the concept of a *wizard* in graphical user interfaces. Subsequently, the user can easily adjust the constructed ski file in a text editor; a slightly more experienced user can copy and paste building blocks between different ski files. For each simulation performed, SKIRT produces a LATEX file describing the contents of the input ski file in a human-readable format that can be used for documentation purposes.

To further facilitate the configuration process, physical quantities such as distances, sizes or masses can be specified in units selected by the user. The default unit system for a simulation's input and output is specified early on in the ski file (e.g. extragalactic units on line 2 of Fig. 2.6), and individual parameter values can be specified with a units string that overrides the default. For example, a scale length of 6600 pc could be specified as `"6600 pc"`, `"6.6 kpc"`, or approximately `"2e20 m"`.

```
1  $ skirt
2  ? Enter the name of the ski file to be created: spiralgalaxy
3    Possible choices for the simulation:
4       1. An oligochromatic Monte Carlo simulation
5       2. A panchromatic Monte Carlo simulation
6  ? Enter one of these numbers [1,2] (1): 1
7    Possible choices for the units system:
8       1. SI units
9       2. Stellar units (length in AU, distance in pc)
10      3. Extragalactic units (length in pc, distance in Mpc)
11 ? Enter one of these numbers [1,3] (3):
12 ? Enter the number of photon packages per wavelength [0,2e13] (1e6): 1e7
13   Possible choices for the wavelength grid:
14      1. A list of one or more distinct wavelengths
15   Automatically selected the only choice: 1
16 ? Enter the wavelengths [0.0001 micron,1e6 micron]: 0.55
17   Possible choices for the stellar system:
18      ...
19   Possible choices for the geometry of the dust component:
20      1. A point source geometry
21      2. A Plummer geometry
22      ...
23      9. An exponential disk geometry
24      ...
25 ? Enter one of these numbers [1,42] (9): 9
26 ? Enter the radial scale length [0 pc,∞ pc]: 6600
27 ? Enter the axial scale height [0 pc,∞ pc]: 250
28      ...
29   Possible choices for the type of normalization for the dust component:
30      1. Normalization by defining the total dust mass
31      2. Normalization by defining the edge-on optical depth at some wavelength
32      ...
33 ? Enter one of these numbers [1,7] (1): 1
34 ? Enter the total dust mass of the dust component [0 Msun,∞ Msun]: 4e7
35   Possible choices for item #2 in the dust components list:
36      1. A dust component
37 ? Enter one of these numbers or zero to terminate the list [0,1] (1): 0
38   Possible choices for the dust grid:
39      ...
40      2. An axisymmetric dust grid in cylindrical coordinates
41      ...
42      8. A Voronoi dust grid
43 ? Enter one of these numbers [1,8] (5): 2
44 ? Enter the cylindrical radius of the grid [0 pc,∞ pc]: 25000
45      ...
46   Successfully created ski file: spiralgalaxy.ski
47 $
```

Figure 2.5: A partial transcript of the query and answer terminal session to configure a SKIRT simulation for a simple spiral galaxy model. The smart mechanism guides the user through all possible options, narrowing down the possibilities based on earlier choices. For example, on line 13 there is only one choice for the wavelength grid because on line 6 the user selected an oligochromatic simulation. Also, the dust grid choices on lines 39-42 are limited to 2D and 3D grids (omitting 1D grids) since the geometry selected on line 25 is axisymmetric. Furthermore, the options for the geometry in lines 26-27 and for the dust grid in lines 44 and beyond are tailored to the selected type of geometry/dust grid.

```
1   <OligoMonteCarloSimulation packages="1e7">
2      <units type="Units"> <ExtragalacticUnits/> </units>
3      <wavelengthGrid type="OligoWavelengthGrid">
4         <OligoWavelengthGrid wavelengths="0.55 micron"/>
5      </wavelengthGrid>
6      <stellarSystem type="StellarSystem">
7         <StellarSystem>
8            <components type="StellarComp">
9               <OligoStellarComp luminosities="1e11">
10                 <geometry type="Geometry">
11                    <ExpDiskGeometry radialScale="4400 pc" axialScale="500 pc"/>
12                 </geometry>
13              </OligoStellarComp>
14           </components>
15        </StellarSystem>
16     </stellarSystem>
17     <dustSystem type="OligoDustSystem">
18        <OligoDustSystem>
19           <dustDistribution type="DustDistribution">
20              <CompDustDistribution>
21                 <components type="DustComp">
22                    <DustComp>
23                       <geometry type="Geometry">
24                          <ExpDiskGeometry radialScale="6600 pc" axialScale="250 pc"/>
25                       </geometry>
26                       <mix type="DustMix"> <DraineLiDustMix/> </mix>
27                       <normalization type="DustCompNormalization">
28                          <DustMassDustCompNormalization dustMass="4e7 Msun"/>
29                       </normalization>
30                    </DustComp>
31                 </components>
32              </CompDustDistribution>
33           </dustDistribution>
34           <dustGrid type="DustGrid">
35              <Cylinder2DDustGrid maxR="25000 pc" minZ="-7000 pc" maxZ="7000 pc">
36                 <meshR type="Mesh">
37                    <LogMesh numBins="101" centralBinFraction="0.01"/>
38                 </meshR>
39                 <meshZ type="MoveableMesh">
40                    <SymPowMesh numBins="101" ratio="50"/>
41                 </meshZ>
42              </Cylinder2DDustGrid>
43           </dustGrid>
44        </OligoDustSystem>
45     </dustSystem>
46     <instrumentSystem type="InstrumentSystem">
47        <InstrumentSystem>
48           <instruments type="Instrument">
49              <FrameInstrument instrumentName="xz" distance="10 Mpc"
50                          inclination="90 deg" azimuth="-90 deg" positionAngle="0 deg"
51                          fieldOfViewX="56000 pc" pixelsX="1200" centerX="0 pc"
52                          fieldOfViewY="14000 pc" pixelsY="300" centerY="0 pc"/>
53           </instruments>
54        </InstrumentSystem>
55     </instrumentSystem>
56  </OligoMonteCarloSimulation>
```

Figure 2.6: The ski file (SKIRT parameter file) configured during the query and answer session shown in Fig. 2.5. While it would be hard for a human to create this file from scratch, it is surprisingly readable because of the self-explanatory tag names. For example, it is easy even for a casual user to adjust the scale height of the dust lane on line 24 or to add an extra instrument by copying lines 49-52 and modifying the inclination angle of the second instrument.

¹ **SKIRT parameter overview: spiralgalaxy**
² An oligochromatic Monte Carlo simulation
³ .   The random number generator: the default random generator
⁴ . .   The seed for the random generator: 4357
⁵ .   The units system: extragalactic units (length in pc, distance in Mpc)
⁶ .   The instrument system: an instrument system
⁷ . .   Item #1 in the instruments list: a basic instrument that outputs the flux in each pixel
⁸ . . .   The name for this instrument: xz
⁹ . . .   The distance to the system: 10 Mpc
¹⁰ . . .   The inclination angle θ of the detector: 90 °
¹¹ . . .   The azimuth angle φ of the detector: -90 °
¹² . . .   The position angle ω of the detector: 0 °
¹³ . . .
¹⁴ .   The number of photon packages per wavelength: 1×10^7
¹⁵ .   The wavelength grid: a list of one or more distinct wavelengths
¹⁶ . .   The wavelengths (λ): 0.55 μm
¹⁷ .   The stellar system: a stellar system composed of various stellar components
¹⁸ . . .
¹⁹ .   The dust system: a dust system for use with oligochromatic simulations
²⁰ . .   The dust distribution: a dust distribution composed of various dust components
²¹ . . .   Item #1 in the dust components list: a dust component
²² . . . .   The geometry of the dust component: an exponential disk geometry
²³ . . . . .   The radial scale length: 6600 pc
²⁴ . . . . .   The axial scale height: 250 pc
²⁵ . . . . .   The radial truncation length (zero means no truncation): 0 pc
²⁶ . . . . .   The axial truncation height (zero means no truncation): 0 pc
²⁷ . . . .   The dust mixture of the dust component: a Draine & Li (2007) dust mix
²⁸ . . . .   Output a data file with the optical properties of the dust mix: yes
²⁹ . . . .   Output a data file with the mean optical properties of the dust mix: yes
³⁰ . . . .   The type of normalization for the dust component: total dust mass
³¹ . . . . .   The total dust mass of the dust component: 4×10^7 M$_\odot$
³² . .   The dust grid: an axisymmetric dust grid in cylindrical coordinates
³³ . . .   Output data files for plotting the structure of the grid: yes
³⁴ . . .   The cylindrical radius of the grid: 25000 pc
³⁵ . . .   The start point of the cylinder in the Z direction: -7000 pc
³⁶ . . .   The end point of the cylinder in the Z direction: 7000 pc
³⁷ . . .   The bin distribution in the radial direction: a logarithmic mesh
³⁸ . . . .   The number of bins in the mesh: 101
³⁹ . . . .   The central bin width fraction: 0.01
⁴⁰ . . .   The bin distribution in the Z direction: a symmetric power-law mesh
⁴¹ . . . .   The number of bins in the mesh: 101
⁴² . . . .   The bin width ratio between the outermost and the innermost bins: 50
⁴³ . .   The number of random density samples for determining cell mass: 100
⁴⁴ . .   Output a data file with convergence checks on the dust system: yes
⁴⁵ . .   Output FITS files displaying the dust density distribution: yes
⁴⁶ . .   Output FITS file with a V-band optical depth map seen from the center: no
⁴⁷ . .   Calculate and output quality metrics for the dust grid: no
⁴⁸ . .   Output a data file with relevant properties for all dust cells: no
⁴⁹ . .   Output statistics on the number of cells crossed per path: no
⁵⁰ . .   Output FITS files displaying the mean intensity of the radiation field: no

Figure 2.7: A pretty-printed version of the ski file (SKIRT parameter file) shown in Fig. 2.6. SKIRT produces a LaTeX file describing the configuration in this way for each simulation performed. The description includes any default values that were omitted from the ski file; see for example lines 43-50.

2.3	Architecture

The latest version of SKIRT was re-architected with the following major design goals in mind:

- *Structured parameter file*: use a self-documenting ski file format that supports the complex configuration needs described above in a user-friendly manner.

- *Single point of definition*: define all information relating to a new feature only once and in the same place, including the code, the human-readable text strings used in the query and answer session, and the tags in the ski file.

- *Data-driven user interface*: conduct the query and answer session and handle the ski file based solely on these data definitions, so that the user interface adjusts automatically as new features are added.

- *Shared-memory parallelization*: make all code reentrant by eliminating the use of global variables; protect the remaining global resources or writable shared data with appropriate locking mechanisms.

- *Modularity*: minimize dependencies among different areas of the code by providing appropriate interfaces and data encapsulation.

In this section we describe the overall architecture of the code and we point out how it achieves these design goals. SKIRT is written in C++ using object-oriented design principles. Specifically we use several of the design patterns originally described by Gamma et al. (1994) in their classic work, including for example the *Composite*, *Builder*, *Visitor*, and *Decorator* patterns.

2.3.1   Simulation items

The core of the SKIRT code is obviously about performing RT simulations. A complete SKIRT simulation is represented at run-time as a hierarchy of objects called *simulation items*, similar to the structure illustrated in Fig. 2.1. This object hierarchy represents the configuration of the simulation (in its structural makeup and in some data members), offers the functionality to perform the simulation (through its member functions), and provides space for any intermediate and resulting data structures (in its data members). Multiple simulation object hierarchies can co-exist and are fully independent of each other.

The object hierarchy for a particular simulation closely mimics the structure of the corresponding ski file. For example, Fig. 2.8 shows the hierarchy that would be constructed for the ski file listed in Fig. 2.6. A solid rectangle represents a simulation item of the specified type; a dashed oval indicates the name and value of an property. Plain properties hold a single value (or a list of values); composite properties link other simulation items into the hierarchy. The simulation items and attributes in this hierarchy map directly to an XML element or attribute in the ski file with the same
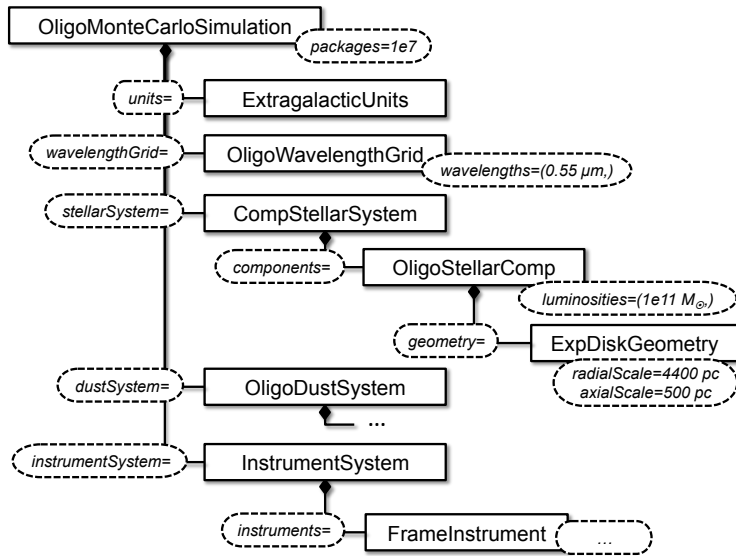
Figure 2.8: The run-time object hierarchy that would be constructed for the ski file shown in Fig. 2.6. A solid rectangle represents a simulation item of the specified type; a dashed oval indicates the name and value of a (plain or composite) property. Connections starting with a diamond indicate aggregation. Each simulation item instance and each property in this hierarchy maps directly to an XML element or attribute in the ski file with the same name.
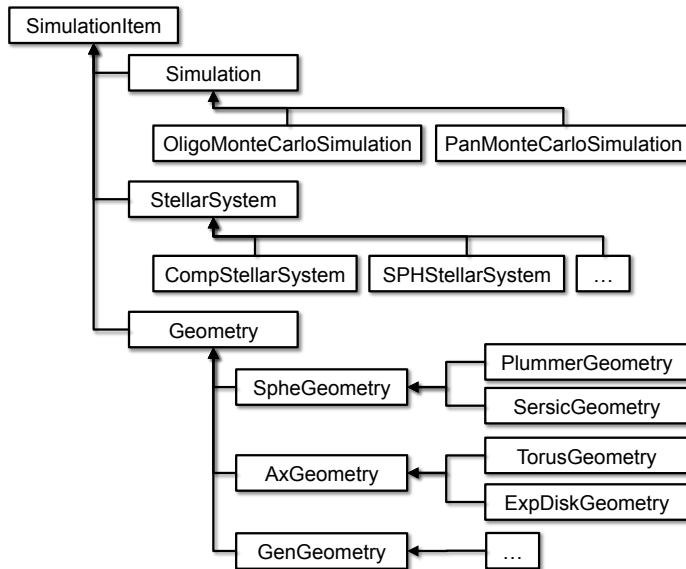
Figure 2.9: A small portion of the more than 150 simulation item classes in the compile-time inheritance hierarchy. A solid rectangle represents a simulation item class with the specified name. Connections starting with an inverted arrow indicate inheritance.
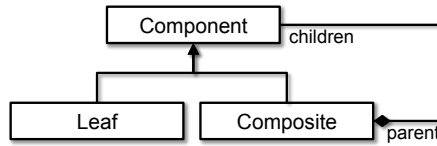
Figure 2.10: The *Composite* design pattern (Gamma et al. 1994). Connections starting with an inverted arrow indicate inheritance; connections starting with a diamond indicate aggregation. This pattern describes an aggregation of objects that all have the same base type.

```
1  void SimulationItem::setup()
2  {
3      setupSelfBefore();
4      for (SimulationItem* child : children())
5      {
6          child->setup();
7      }
8      setupSelfAfter();
9  }
```

Figure 2.11: The implementation of the `SimulationItem::setup()` function (ignoring some implementation details).

name. This correspondence plays an important role in the automation of the user interface, as we will discuss in Sect. 2.3.5.

Each simulation item is an instance of a class that directly or indirectly derives from the `SimulationItem` base class. The simulation item classes form a compile-time inheritance hierarchy, a small portion of which is shown in Fig. 2.9. The run-time object hierarchy representing a simulation is thus an aggregation of objects of the same type, reflecting the *Composite* design pattern (Gamma et al. 1994) illustrated in Fig. 2.10. The Component role is played by the `SimulationItem` class, and the Composite role is assumed by any `SimulationItem` subclass that has one or more composite properties.

The use of the *Composite* pattern is fundamental to the implementation of the user interface discussed in Sect. 2.3.5, and it substantially facilitates reducing dependencies between portions of the code, as described in Sect. 2.3.2 and Sect. 2.3.3.
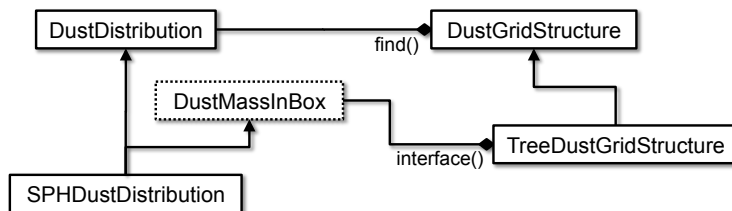
Figure 2.12: A specialty interface (dotted rectangle) connects two specific simulation items at run time, optimizing performance without creating undesirable dependencies in the respective base classes. Connections starting with an inverted arrow indicate inheritance; connections starting with a diamond indicate aggregation.
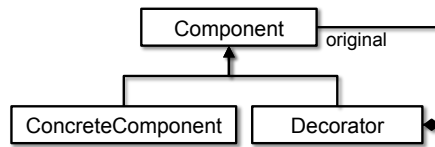
Figure 2.13: The *Decorator* design pattern (Gamma et al. 1994). Connections starting with an inverted arrow indicate inheritance; connections starting with a diamond indicate aggregation. This pattern describes a convenient way to adjust or *decorate* the behavior of another object of the same base type.

2.3.2 Simulation phases

A SKIRT simulation has three phases: *construction*, *setup* and *run*. In the *construction* phase, the code constructs the simulation item hierarchy corresponding to a particular ski file, initializing the values of all plain and composite properties, as described in Sect. 2.3.5. This process completes in a fraction of a second because it doesn't do much work. During the *setup* phase, each simulation item in the hierarchy gets a chance to perform further initialization, such as reading data from resource files or pre-computing frequently used information. This phase may require some processing power, for example, to set up a dust grid that is adapted to the specified dust distribution. Finally, the *run* phase performs the actual simulation and writes down the results. Usually this phase consumes the bulk of the computing resources, and it is fully parallelized.

The `SimulationItem` base class offers the `setup()` function; its implementation is shown in Fig. 2.11. The `children()` function used on line 4 returns a list of all simulation items held by any composite property of the current simulation item. The `setupSelfBefore()` and `setupSelfAfter()` functions are declared `virtual` in the `SimulationItem` base class and are overridden by subclasses that need to perform initialization during the *setup* phase. They are invoked respectively before and after any children of the simulation item have been set up.

The implementation of the `setup()` function follows the *Template Method* design pattern (Gamma et al. 1994) to delegate the actual initialization work to subclasses. To recursively invoke all simulation items in the hierarchy, it relies on the fact that all simulation item classes derive from the same class, which is ensured by the use of the *Composite* design pattern.

SKIRT requires that the root object of the simulation item hierarchy inherits from the `Simulation` class. This class offers the `run()` function, which, not surprisingly, executes the *run* phase. Thus, after constructing the run-time hierarchy, the code simply invokes the `setup()` and `run()` functions on the hierarchy's root object to complete all phases of the simulation.

Reducing dependencies

Most simulation item classes are organized in groups with a common purpose, e.g. wavelength grids, geometries, or dust mixes. All classes in a particular group inherit from the same base class, i.e. `WavelengthGrid`, `Geometry`, or `DustMix`. The base class offers the common interface for all classes in the group towards classes outside of the group. This design principle avoids undesirable dependencies between classes in different groups, enhancing modularity.

Some information about a simulation's configuration is accessed from many different places, and thus must be readily available. To facilitate access to other simulation items in the same hierarchy, the `SimulationItem` class offers the `T* find<T>()` template function, where `T` stands for the name of any class that derives from `SimulationItem`. This template function searches the object hierarchy in which the receiving simulation item resides for a simulation item of the specified type `T`, and returns a pointer to the first such object found after dynamically casting it to the requested type. If the hierarchy does not contain an object of the specified type, the function throws an exception. The implementation of the `find<>()` template function again relies on the fact that all simulation item classes derive from the same class.

For example, every run-time simulation hierarchy includes an instance of a particular `WavelengthGrid` subclass, such as `OligoWavelengthGrid` or `NestedLogWavelengthGrid`. Any simulation item in the hierarchy can call `find<WavelengthGrid>()` to retrieve a pointer to the common wavelength grid interface; the caller does not know the specific sub-type of the returned object. Also, the caller has no need to know where the returned object resides, so this mechanism replaces application-wide global data (which gets in the way of parallelization) by simulation-wide available data.

While modularity is important, the generic and narrow interfaces between different areas of the code sometimes hide information that can be relevant for optimal cooperation between components. As a first example, the dust mass in each grid cell is usually estimated by probing the dust density distribution in a number of random locations uniformly distributed over the spatial extent of the cell. This generic mechanism works for any cell shape and for any type of density distribution. However, for certain cell shapes combined with certain types of density distribution, it might be orders of magnitude faster to directly calculate the mass in the cell. As a second example, sometimes we would like to build a dust grid based on particular locations (such as SPH particles) defined as part of the input dust density distribution, rather than based on the density distribution itself. However, the generic interface does not offer particle information because the concept is meaningless for most density distributions.

These features can be accomplished without breaking modularity by using a specialty interface that is known only to the specific classes involved, and by providing a mechanism to dynamically connect the two players at run time. This is illustrated in Fig. 2.12 for the first example described above. The `DustMassInBox` interface declares pure virtual functions offering the relevant special capabilities. The
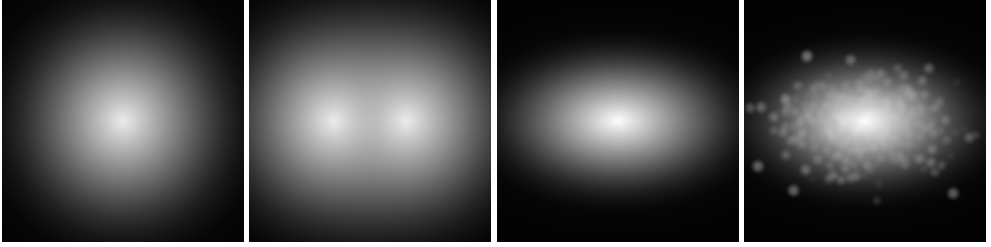
Figure 2.14: Planar cuts through four density distributions derived from the same underlying geometry using the *Decorator* design pattern. From left to right: a plain, spherical Einasto profile with index 1; two superposed Einasto profiles, each shifted aside using an offset decorator; the Einasto profile deformed into a spheroidal shape by a decorator; and a clumpy, spheroidal Einasto distribution derived from the original profile by applying a chain of two decorators.

`SPHDustDistribution` class inherits the interface and actually implements its functions. Finally, the `TreeDustGridStructure` simulation item recovers a pointer to the specialty interface as follows. First it uses the `find<>()` template function to retrieve the dust distribution object in the hierarchy; this could in fact happen in the `DustGridStructure` base class since the returned pointer is of the generic type `DustDistribution`. Then it invokes the `interface<>()` template function on the dust distribution object, which is in fact of type `SPHDustDistribution`, to return a pointer to the `DustMassInBox` interface implemented by that same object.

For this purpose, the `SimulationItem` base class provides the `T* interface<T>()`, where `T` stands for the name of the specialty interface to be recovered. In the example of Fig. 2.12, the function can be implemented with a simple dynamic cast. To support more complicated cases, the function also allows a simulation item to delegate the implementation of a specialty interface to a different object.

2.3.4 Reusing components

In a Monte Carlo RT code, the key function of a dust geometry (describing the distribution of the dusty medium) is to retrieve the dust density at a specified location in space. This function is called (quite often) during setup to build an appropriate dust grid and calculate the dust mass in each grid cell. On the other hand, the key function of a stellar geometry (describing the distribution of radiation sources) is to generate a random location in space, drawn from a probability distribution corresponding to the geometry's density distribution. This function is called repeatedly during the simulation to determine the point of emission for a new photon package.

It might seem that the functionalities of the two geometry types are thus rather disjunct, but this is not the case. For example, we want to build a Voronoi dust grid using generating sites placed according to the dust density distribution. In this case, we need the key stellar geometry functionality (generating random points) in the dust geometry. Therefore the recent SKIRT version has unified geometry classes that offer both functions, as listed in Table 2.1.

This change prompted us to invest in a number of geometry classes that modify other geometry's in interesting ways, following the *Decorator* design pattern (Gamma et al. 1994) shown in Fig. 2.13. An object in the Decorator role maintains a pointer to another object of the same base class, called the original object. The decorator implements the base class interface by calling corresponding functions in the original object, and returning the results after possible adjustment. In SKIRT, the `Geometry` class plays the Component role in this pattern, and any `Geometry` subclass can assume the ConcreteComponent role, i.e. the role of the original geometry being decorated. The `ClumpyGeometry` class is one of the classes that plays the Decorator role. It modifies the original geometry by replacing a fraction of its total mass allocation by randomly placed clumps. Other decorators relocate the original geometry's center, or deform a spherical geometry into a spheroidal or triaxial distribution. Multiple decorators can be chained to achieve the combined effects, as illustrated in Fig. 2.14.

2.3.5   Automating the user interface

The latest version of SKIRT is based on the Qt development framework[5], which includes a rich set of cross-platform C++ libraries and an integrated development environment (IDE) called Qt Creator. Although we don't need its graphical user interface (GUI) capabilities, the Qt environment offers substantial benefits in other areas as well[6]. Specifically, the Qt environment provides run-time introspection of classes and their member functions, assuming the appropriate declarations were added in the code. The Qt mechanism is a lot more advanced than the standard C++ run-time type information (RTTI) system. For example, the Qt library offers functions to retrieve compile-time information such as the class inheritance hierarchy and the type of function arguments or return values. It is also possible to invoke a function by specifying the function name at run-time as a string, or to construct a new class instance in a similar manner.

SKIRT relies on the Qt introspection features to automatically construct a user interface from the C++ class declarations in the code. To enable this process, all `SimulationItem` subclass declarations must be augmented with some extra information, as illustrated in Fig. 2.15. The keywords starting with Q_ are provided by the Qt development environment. The `Q_OBJECT` keyword on line 3 is required to enable the Qt introspection features for this class. The `Q_CLASSINFO` definitions on lines 4-25 associate an ordered list of key-value pairs with the compile-time class information; these strings can be retrieved at run-time through the Qt introspection system. The `Q_INVOKABLE` keyword on lines 28 and 35-48 enables Qt introspection for the constructor or member function declaration following the keyword.

In SKIRT, the `Q_CLASSINFO` key-value pairs are used to provide a human readable description for the class, and to define its configurable properties (i.e. the properties that can be specified in a ski file). The name of a property, e.g. `"clumpFraction"`,

---

[5] The Qt project: http://qt-project.org

[6] The new features in the recent C++11 language standard cover much of the functionality for which SKIRT uses Qt, with the exception of the introspection capabilities discussed in this section.

```
1   class ClumpyGeometry : public Geometry
2   {
3       Q_OBJECT
4       Q_CLASSINFO("Title", "a geometry that adds clumpiness to any geometry")
5
6       Q_CLASSINFO("Property", "geometry")
7       Q_CLASSINFO("Title", "the geometry to which clumpiness is added")
8
9       Q_CLASSINFO("Property", "clumpFraction")
10      Q_CLASSINFO("Title", "the fraction of the mass locked up in clumps")
11      Q_CLASSINFO("MinValue", "0")
12      Q_CLASSINFO("MaxValue", "1")
13
14      Q_CLASSINFO("Property", "clumpCount")
15      Q_CLASSINFO("Title", "the total number of clumps")
16      Q_CLASSINFO("MinValue", "1")
17
18      Q_CLASSINFO("Property", "clumpRadius")
19      Q_CLASSINFO("Title", "the scale radius of a single clump")
20      Q_CLASSINFO("Quantity", "length")
21      Q_CLASSINFO("MinValue", "0")
22
23      Q_CLASSINFO("Property", "cutoff")
24      Q_CLASSINFO("Title", "cut off clumps at the boundary of the underlying geometry")
25      Q_CLASSINFO("Default", "no")
26
27  public:
28      Q_INVOKABLE ClumpyGeometry();
29
30  protected:
31      void setupSelfBefore();
32      void setupSelfAfter();
33
34  public:
35      Q_INVOKABLE void setGeometry(Geometry* value);
36      Q_INVOKABLE Geometry* geometry() const;
37
38      Q_INVOKABLE void setClumpFraction(double value);
39      Q_INVOKABLE double clumpFraction() const;
40
41      Q_INVOKABLE void setClumpCount(int value);
42      Q_INVOKABLE int clumpCount() const;
43
44      Q_INVOKABLE void setClumpRadius(double value);
45      Q_INVOKABLE double clumpRadius() const;
46
47      Q_INVOKABLE void setCutoff(bool value);
48      Q_INVOKABLE bool cutoff() const;
49
50  public:
51      double density(Position bfr) const;
52      Position generatePosition() const;
53
54      ...
55  };
```

Figure 2.15: A typical simulation item class declaration. The keywords starting with `Q_` are provided by the Qt development environment, and serve to define the extra information needed to automatically build a user interface for the features offered by this class, as explained in Sect. 2.3.5. The setup functions declared on lines 31-32 are described in Sect. 2.3.2 and Fig. 2.11. The geometry-specific functions declared on lines 51-52 are described in Sect. 2.3.4.
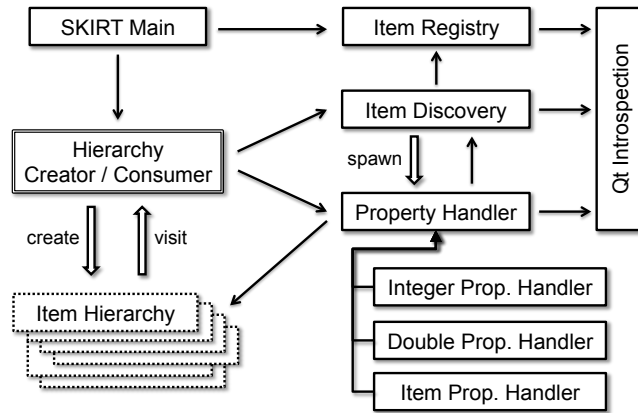
Figure 2.16: Schematic overview of the code that automatically builds the user interface from simulation item class declarations. In this diagram, the phrase *simulation item* has been shortened to *item*. A solid arrow indicates that the source module uses the target module. Connections starting with an inverted arrow indicate inheritance. The relationships are described in more detail in Sect. 2.3.5.

must match the name of a getter function `clumpFraction()` and a setter function `setClumpFraction()`, declared with the keyword `Q_INVOKABLE`. The type of the configurable property is derived from the getter's return value (`double` in this case). Additional key-value pairs can specify options such as a default value, or the type of physical quantity represented by this property, which determines the units used or accepted in the user interface.

Figure 2.16 depicts the overall organization of the code that automatically builds the user interface from the compile-time data. The architecture is inspired by – but doesn't correspond exactly to – the *Builder* and *Visitor* design patterns (Gamma et al. 1994). The *hierarchy creator* object in the figure plays the Builder role, and the *hierarchy consumer* object plays the Visitor role.

Just after program startup, the *item registry* is initialized with a list of all simulation item classes. This ensures that all classes are actually linked into the code, and it provides the starting point for the *item discovery* module to implement queries about the simulation item classes. Functions offered by this module include for example `title(itemType)`, `descendants(itemType)`, and `createPropertyHandlers(item)`. The latter function spawns a property handler of the appropriate type for each property of the specified simulation item.

A property handler combines a pointer to a particular simulation item object in the hierarchy, with knowledge about the compile-time attributes of one of the properties of the item. The handler can be used to get or set the property value directly from or into the target object, or to retrieve attributes such as its default value. The handler also knows how to convert a property value into a string for human consumption, and vice versa. There are handlers for various property types, including Boolean, integer, enumeration, floating point (with support for units), string, and pointer to

simulation item.

To create a simulation item hierarchy from a ski file, the command line handler in the *SKIRT main* module enlists an `XmlHierarchyCreator` object. This object uses the XML tags in the ski file, which correspond to simulation class names and property names, to recursively construct the corresponding simulation items and set their property values. The code heavily relies on the *item discovery* module and the property handlers spawned by it.

Similarly, a `ConsoleHierarchyCreator` object is used to create a hierarchy from scratch by conducting a query and answer session. The top of the hierarchy must be occupied by an instance of the `Simulation` class, so the algorithm obtains a list of concrete `Simulation` subclasses from the *item discovery* module, and asks the user to make a choice. The question is formulated using the titles provided in each class declaration; see lines 3-6 in Fig. 2.5. The algorithm constructs an instance of the selected subclass, and then loops over all of its configurable properties, asking the appropriate question(s) for each property depending on its type. Boolean, numeric and string properties only need a single question; see lines 12, 26, 27, 34, and 44 in Fig. 2.5. A property that points to another simulation item prompts a multiple choice question to select one of the available concrete subclasses that inherit the appropriate type, again obtained from the *item discovery* module; see, e.g., lines 7-11 and 19-25 in Fig. 2.5. A new simulation item of the selected type is created (and linked into the hierarchy), and the same mechanism is recursively applied to the new object.

By selecting the desired type of simulation item at each level in the recursion, the user's responses drive the nature of subsequent questions in the session. While this is sufficient for most purposes, the discovery process implements a few extra mechanisms to support specific needs. For example, the list of available dust grids depends on the (lack of) symmetries in the geometries selected earlier; e.g. the user can't select a 1D or 2D grid for a 3D geometry. Also, it is possible to skip questions that are deemed irrelevant based on the response to a previous question in the same class. All of these mechanisms are fully data-driven from the `Q_CLASSINFO` definitions in the simulation item class declarations.

Once a simulation item hierarchy is in place, the same underlying data can be used to reverse the process and write down the configuration in a human-readable form. In Fig. 2.16 the Creator object is now replaced by a Consumer object that recursively visits the items in the hierarchy to produce the corresponding output, using the information supplied by the *item discovery* module and the property handlers it spawns. Most importantly, SKIRT uses the `XmlHierarchyWriter` object to output a ski file (Fig. 2.6) after the user configured a simulation item hierarchy through a query and answer session (Fig. 2.5). A newly generated ski file is also stored with each set of simulation results, as a standard reference, explicitly listing the default values for properties that may have been omitted in the input ski file. Using the `LatexHierarchyWriter` object, SKIRT also writes a LaTeX source file that documents the configuration in an even more user-friendly format (Fig. 2.7).

Because of this automation, adding a new SKIRT feature is extremely straightforward. For example, to add a new geometry, a developer would copy one of the existing geometry classes, rename the class, adjust the implementation and documentation of the member functions, adjust the `Q_CLASSINFO` definitions in the class declaration, and add a single line in the `RegisterSimulationItems` class to register the new geometry to the discovery system. Except for this trivial registration requirement, all information about the new geometry is in a single place, and the user interface will be automatically adjusted to incorporate it.

2.4	Conclusions

We described the major features of SKIRT, a state-of-the-art Monte Carlo dust radiation transfer simulation code used to study spiral galaxies, accretion disks and other astrophysical systems. In addition to its core capability of tracing the radiation through the dust, SKIRT offers a large number of built-in options for configuring all aspects of the simulation model, including spatial and spectral distributions, dust grain characterizations, simulated detection systems, and discretization.

Providing a proper user interface to support these complex configuration requirements is a nontrivial undertaking. Most SKIRT simulations run for hours or days, often on remote servers, so there is no need for a fancy graphical user interface. Still, the typical user is an expert astrophysicist who interacts with the SKIRT code rather occasionally, and thus benefits greatly from a low-barrier interface. SKIRT addresses this challenge through the combination of a *wizard*-like query and answer session to guide a first-time user through the configuration process, and self-documenting XML-based parameter files that can be easily updated in a text editor.

We further described the overall architecture of the code. Inspired by standard software design principles and patterns, the latest version of SKIRT has a modular implementation that can be easily maintained and expanded. Programming interfaces between components are well defined and narrow. The user interface is automatically constructed from data provided in the C++ class declarations, allowing a single point of definition, and placing the user interface information right next to the code implementing the corresponding feature.

All too often, scientific codes are written without much concern for user interface or for modular software design. This is very unfortunate. Scientists may not need a *graphical* user interface, but, just like every one else, they do benefit from an interaction mechanism that hides the underlying complexity. As we have illustrated in this work, a well-designed non-graphical user interface may be a perfect fit, and can often be developed and maintained with limited resources. Similarly, adhering to proven software design principles pays off, even for small and mid-sized projects.

The SKIRT source code is publicly available, and it has already been applied to RT problems in various astrophysical domains. We welcome new applications, and we invite potential users and code contributors to join the SKIRT community.

Voronoi grids[8]

<div style="text-align: right; font-size: 3em;">3</div>

YIELD TO TEMPTATION; IT MAY NOT PASS YOUR WAY AGAIN.

Choosing an appropriate discretization of the spatial domain is crucial for any numerical simulation. Adaptive grids with cuboidal cells such as octrees have proven very popular; however, several recently introduced hydrodynamical and RT codes are based on a Voronoi tessellation of the spatial domain. An unstructured grid of this nature poses new challenges in laying down the rays (straight paths) needed in RT codes. In this chapter, we present a method for computing straight paths between two arbitrary points through a 3D Voronoi grid in the context of a RT code. We implement this method in SKIRT and we compare the results obtained through the Voronoi grid with those generated by an octree grid. We find that the presented algorithm produces correct results for our test models. Shooting photon packages through the geometrically much more complex 3D Voronoi grid is only about three times slower than the equivalent process in an octree grid with the same number of cells, while in fact the total number of Voronoi grid cells may be lower for an equally good representation of the density field. We argue that the benefits of using a Voronoi grid in RT simulation codes will often outweigh the somewhat slower performance.

3.1 Introduction

For the purpose of numerical computation, the domain under study must be discretized. Since memory requirements and computation time rapidly increase with the number of grid cells, modern 3D RT codes employ an adaptive grid, placing more and smaller cells in areas that require a higher resolution. Starting from a cuboidal root cell that spans the complete spatial domain, an adaptive mesh refinement (AMR) scheme recursively subdivides each cell into $k \times l \times m$ cuboidal subcells until the required resolution is reached. In the special case of an octree, $k = l = m = 2$ so that each cell is subdivided into eight subcells (hence the name of the data structure). Adaptive-mesh grids and especially octree grids are well established (Kurosawa and Hillier 2001; Steinacker et al. 2002; Wolf 2003b; Harries et al. 2004; Niccolini and Alcolea 2006; Jonsson 2006; Bianchi 2008; Laursen et al. 2009; Robitaille 2011; Lunttila and Juvela 2012; Heymann and Siebenmorgen 2012; Saftly et al. 2013) and several

---

[8] Published as Camps et al. (2013b). My contribution includes the invention of the straight path algorithm and the implementation and testing of the Voronoi grid in SKIRT. The co-authors promoted the use of Voronoi grids in radiative transfer to begin with, and proposed ideas for testing the implementation.
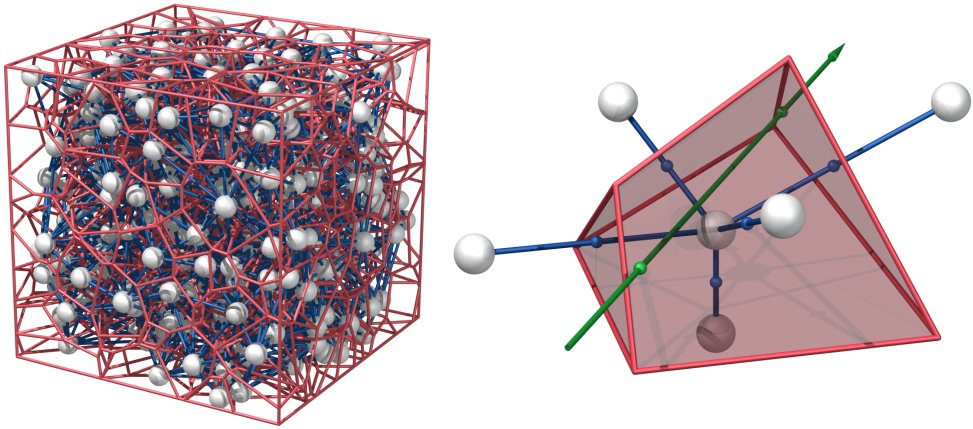
Figure 3.1: *Left*: A Voronoi tessellation for 400 random sites (in gray), bounded by a cube; Voronoi cell edges are shown in red, Delaunay edges in blue. *Right*: A single Voronoi cell (in red) with its neighboring sites (in gray) and corresponding Delaunay edges (in blue); a straight path through the cell is shown (in green) with its intersection points with the cell boundary at entry and exit.

methods have been investigated to make them as efficient as possible (Saftly et al. 2013).

Adaptive-mesh grids seem to be an obvious choice. It is straightforward to construct an appropriate grid for any density field, whether defined by an analytical model or by a collection of smoothed particles; and it is easy to calculate a straight path through the grid, since the boundaries of the cuboidal cells are lined up with the coordinate axes and each cell has a limited number of neighbors (Saftly et al. 2013). This second point is very important in the context of RT because ray tracing and Monte Carlo RT codes determine the radiation field in each grid cell by laying down random rays (i.e., straight paths) through the domain. The simulation run time is often dominated by the portion of the code that identifies the grid cells crossed by each path and calculates the lengths of the corresponding path segments.

Adaptive-mesh grids also have drawbacks. First of all, for a given density field and required resolution, an AMR grid may not be the kind of grid with the least number of cells. To illustrate this, consider a density field defined by a set of smoothed particles. An octree grid constructed such that each cell encloses at most one particle usually has over three times more cells than there are particles; i.e., two out of three cells are empty[9]. In contrast, an unstructured grid based on a Voronoi tessellation of the spatial domain (see Sect. 3.2.1 and Fig. 3.1), using the given particles as generating sites, has exactly the same number of cells as there are particles. While not an issue in many situations, minimizing the number of cells is sometimes crucial. For example, consider a panchromatic RT simulation involving small dust grains in non-LTE conditions. Because each cell stores radiation field data per wavelength bin, memory requirements

---

[9]   To verify this claim, we ran a few tests with particles distributed uniformly over the spatial domain, and with particles representing a galaxy generated by a hydrodynamical simulation.

are substantial. Moreover, the simulation run time is most likely dominated by the calculation of the non-LTE heating and re-emission of the dust grains in each cell. In this case, both memory usage and run time scale roughly linearly with the number of cells.

Furthermore, RT simulations frequently serve to predict the observable properties of artificial systems resulting from (magneto-)hydrodynamical (MHD) simulations (Juvela and Padoan 2003; Bethell et al. 2004; Stamatellos and Whitworth 2005; Jonsson et al. 2010; Acreman et al. 2010; Hayward et al. 2011; Robitaille 2011; Lunttila and Juvela 2012; Juvela et al. 2012). Hydrodynamical simulation codes historically employ one of two schemes: a Lagrangian formulation based on moving particles (smoothed particle hydrodynamics or SPH), for example Gadget (Springel 2005; Dolag and Stasyszyn 2009; Pakmor et al. 2012) and SEREN (Hubber et al. 2011); or a Eulerian approach based on a non-moving spatial grid, often an AMR grid, for example RAMSES (Fromang et al. 2006), Enzo (Collins et al. 2010; Bryan et al. 2014), and AMR-VAC (Keppens et al. 2012).

Recent codes including TESS (Duffell and MacFadyen 2011) and AREPO (Springel 2010) introduce a new scheme that employs a moving mesh based on a Voronoi tessellation of the spatial domain (see Sect. 3.2.1 and Fig. 3.1). This new scheme is claimed to combine the best features of SPH and the traditional Eulerian approach, and it is becoming increasingly popular. It has already been applied to various problems including the formation of stars, galaxies, and cosmological structures (Greif et al. 2011; Bauer and Springel 2012; Sijacki et al. 2012; Kereš et al. 2012; Vogelsberger et al. 2012; Torrey et al. 2012; Nelson et al. 2013; Marinacci et al. 2014). While the output from a moving mesh code can be re-gridded to an AMR grid to perform RT, the resampling process unavoidably introduces inaccuracies and represents additional overhead; it seems preferable to perform both aspects of the simulation (MHD and RT) on the same grid.

These considerations lead to the question of whether it is possible to perform accurate and efficient RT on unstructured Voronoi grids.

One approach is to approximate a straight path through the grid by a sequence of non-collinear segments connecting neighboring sites. For example in the SimpleX code (Paardekooper et al. 2010) and in the LIME code (Brinch and Hogerheijde 2010) radiation travels along the edges of the Delaunay triangulation corresponding to the Voronoi grid (see Sect. 3.2.1; the Delaunay edges are shown in blue in the leftmost panel of Fig. 3.1). While it facilitates calculating the paths, this approximation requires additional mechanisms to compensate for errors in path length (see Fig. 5 in Paardekooper et al. 2010) and direction (see Fig. 4 in Brinch and Hogerheijde 2010). The distance covered by the path inside a particular grid cell becomes a fuzzy concept, while this is an important quantity in many RT codes, e.g., for tracking the amount of energy absorbed in the cell. And finally the spread on direction makes it hard to produce high-resolution images of the simulated object.

In Sect. 3.2 we present instead an efficient method of calculating a straight path

between two arbitrary points through a 3D Voronoi grid, applicable in any RT code based on ray tracing or Monte Carlo techniques. The path segments inside each grid cell are calculated to high precision using a straightforward algorithm that relies on the mathematical properties of Voronoi tessellations. In Sect. 3.3 we introduce an implementation of the method in our dust RT code SKIRT (Baes et al. 2011). We demonstrate the method's reliability, accuracy, and efficiency by comparing results obtained through the Voronoi grid with those generated by existing well-tested grids. In Sect. 3.4 we summarize our conclusions.

3.2 Method

3.2.1   Voronoi tesselations of 3D space

Given a set of points $\{\vec{p}_1, \vec{p}_2, \ldots \vec{p}_n\}$ in 3D space, called *sites*, the corresponding Voronoi tessellation (Dirichlet 1850; Voronoi 1908) is a set of cells $\{C_i\}$ where each cell C_i consists of all the points $\vec{p}$ at least as close to $\vec{p}_i$ as to any other site. The corresponding Delaunay triangulation (Delaunay 1934) is a graph created by placing a straight edge between any two sites that share a cell boundary in the Voronoi tessellation. Thus every site is connected to its nearest neighbors.

An example Voronoi tessellation and a single Voronoi cell are illustrated in Fig. 3.1. A Voronoi cell is delimited by a convex polyhedron. A Delaunay edge, i.e., a line segment that connects two sites sharing a polygonal face, is perpendicularly bisected by the plane containing the face, although the bisection point may lie outside the face. For a set of sites chosen randomly from a uniform distribution, the number of nearest neighbors (or equivalently the number of cells sharing a face with any given cell) has an expectation value of 15.54 (van de Weygaert 1994).

To obtain an optimal grid in the context of a RT simulation, the Voronoi sites should be more densely packed (generating smaller cells) in regions where a higher resolution is desired. For example one could select the positions randomly from a probability distribution that follows the density of the RT medium, and perhaps place extra sites near sharp edges or large gradients. If the density field is defined by a set of smoothed particles, the particle locations form natural Voronoi sites. And of course if the density field is already defined on a Voronoi mesh, the original site locations can be used.

3.2.2   A straight path through a Voronoi grid

We consider a cuboidal spatial domain $\mathcal{D}$ defined by its corner points $(\vec{D}_{\min}, \vec{D}_{\max})$, and a set of sites $\{\vec{p}_m \in \mathcal{D}, m = 1 \ldots M\}$. All sites are inside the domain, and the corresponding Voronoi cells are clipped by the domain walls, as illustrated in the left-most panel of Fig. 3.1. Given a ray describing the path of a photon package, defined by a starting point $\vec{r}_0 \in \mathcal{D}$ and a direction $\vec{k}$, our aim is to calculate the ray's consecutive intersection points with the Voronoi cell walls – or equivalently, the distance travelled

in each cell – until the ray leaves the domain. This is illustrated in the rightmost panel of Fig. 3.1 for a single cell. The presented method can easily be adjusted for other domain geometries, or for rays that originate outside the domain.

During a setup phase, before any straight paths are calculated, the following data is prepared:

1. The domain boundaries $(\vec{D}_{\min}, \vec{D}_{\max})$.

2. The positions of the sites $\{\ \vec{p}_m,\ m = 1 \ldots M\ \}$.

3. For each site $\vec{p}_n$, the indices $\{\ m_{n,i},\ i = 1 \ldots I_n\ \}$ of all sites neighboring that site or, equivalently, of all cells neighboring the cell corresponding to that site. Domain walls are represented by special (negative) index values.

Data items (1) and (2) are externally specified as part of the problem definition. The neighbor lists (3) can easily be derived from a Voronoi tessellation or Delaunay triangulation for the specified set of sites, since nearest neighbor information is the defining characteristic of these concepts. *No information is needed on the vertices, edges or faces of the polyhedra delimiting the Voronoi cells.*

To begin calculating a straight path, the current point $\vec{r}$ is set to the starting point, and the current cell index m_r is set to the index of the cell containing the starting point. By definition of a Voronoi tessellation, finding the cell containing a given point $\vec{r}$ is equivalent to locating the site $\vec{p}_i$ nearest to $\vec{r}$. This is a straightforward operation that can easily be optimized as described later in Sect. 3.2.3. For the time being we assume that there is a function $\mathcal{C}(\vec{r})$ that returns the index m of the cell containing a given point.

Once initialized, the method loops over the algorithm that computes the exit point from the current cell, i.e., the intersection of the ray formed by the current point $\vec{r}$ and the direction $\vec{k}$ with the current cell's boundary. The algorithm also produces the index of the neighboring cell without extra cost. If an exit point is found, the loop adds a path segment to the output, updates the current point and the current cell index, and continues to the next iteration. If the exit is towards a domain wall, the loop is terminated. Because of computational inaccuracies it may occur that no exit point is found. In that case, no path segment is added to the output, the current point is advanced in the direction $\vec{k}$ over an infinitesimal distance $\epsilon \ll ||\vec{D}_{\max} - \vec{D}_{\min}||$, and the new current cell index is determined by calling the function $\mathcal{C}(\vec{r})$.

The algorithm computing the exit point from the current cell requires the following input data: the current point $\vec{r}$; the ray's direction $\vec{k}$ as a unit vector; the index m_r of the current cell; the indices $\{\ m_i,\ i = 1 \ldots I\ \}$ of all cells neighboring the current cell, with domain walls represented by special (negative) values; the positions of the sites $\{\ \vec{p}_m,\ m = 1 \ldots M\ \}$; and the domain boundaries $(\vec{D}_{\min}, \vec{D}_{\max})$.

The line containing the ray under consideration can be written as $\mathcal{L}(s) = \vec{r} + s\vec{k}$ with $s \in \mathbb{R}$. The exit point can similarly be written as $\vec{q} = \vec{r} + s_q \vec{k}$ with $s_q > 0$, and the distance covered within the cell is given by s_q. The index of the cell next to the exit

point is denoted m_q and is easily determined as follows.

1. Calculate the set of values $\{s_i\}$ for the intersection points between the line $\mathcal{L}(s)$ and the planes defined by the neighbors m_i (see below for details on this calculation).

2. Select the smallest nonnegative value $s_q = \min\{s_i|s_i > 0\}$ in the set to determine the exit point and the corresponding neighbor m_q.

3. If there is no nonnegative value in the set, no exit point has been found.

To calculate s_i in step (1) for a regular neighbor $m_i > 0$, intersect the line $\mathcal{L}(s) = \vec{r} + s\vec{k}$ with the plane bisecting the sites $\vec{p}(m_i)$ and $\vec{p}(m_r)$. An unnormalized vector perpendicular to this plane is given by

$$\vec{n} = \vec{p}(m_i) - \vec{p}(m_r) \tag{3.1}$$

and a point on the plane is given by

$$\vec{p} = \frac{\vec{p}(m_i) + \vec{p}(m_r)}{2}. \tag{3.2}$$

The equation of the plane can then be written as

$$\vec{n} \cdot (\vec{x} - \vec{p}) = 0. \tag{3.3}$$

Substituting $\vec{x} = \vec{r} + s_i\vec{k}$ and solving for s_i provides

$$s_i = \frac{\vec{n} \cdot (\vec{p} - \vec{r})}{\vec{n} \cdot \vec{k}}. \tag{3.4}$$

If $\vec{n} \cdot \vec{k} = 0$ the line and the plane are parallel so that there is no intersection, and the above equation produces $s_i = \pm\infty$. When using standard IEEE 754 floating point arithmetic there is no reason to test for this special case, since the infinite value will never be selected as the exit point in step (2).

To calculate s_i in step (1) for a domain wall $m_i < 0$, substitute the appropriate normal and position vectors for the wall plane in Eq. 3.4. For example, for the left wall one has $\vec{n} = (-1, 0, 0)$ and $\vec{p} = (D_{\min,x}, 0, 0)$ so that

$$s_i = \frac{D_{\min,x} - r_x}{k_x}. \tag{3.5}$$

In an actual implementation of this algorithm there is no need to accumulate the complete set of $\{s_i\}$ values; one can simply keep track of the smallest nonnegative value. As a further optimization, part of the intersection calculation can be avoided for about half of the planes by noting that the sign of $\vec{n} \cdot \vec{k}$ determines the sign of s_i in Eq. 3.4. Indeed, the site $\vec{p}(m_r)$ and the current point $\vec{r}$ are on the same side of the plane defined by Eq. 3.3 (unless $\vec{r}$ lies *in* the plane), so that the numerator of Eq. 3.4 is always positive (or zero if $\vec{r}$ lies in the plane).

3.2.3   Finding the cell containing a given point

We now return to the implementation of the function $\mathcal{C}(\vec{r})$ that identifies the Voronoi cell containing a given query point. By definition of a Voronoi tessellation, this operation is equivalent to finding the site closest to the query point. There are many sophisticated ways to accelerate this nearest neighbor search, for example by building a kd-tree (Friedman et al. 1977) or an R-tree (Guttman 1984) data structure. We chose to use a simple mechanism, since this function is usually invoked only once per path and thus its performance is not overly critical.

We assume the domain $\mathcal{D}$ is partitioned in a set of cuboidal *blocks* $\{\,\mathcal{B}_k,\,k=1\ldots K\,\}$ according to a regular linear grid. During the setup phase described in the beginning of Sect. 3.2.2, an additional data structure is constructed containing, for each block $\mathcal{B}_k$ in the partition of the domain $\mathcal{D}$, the indices $\{\,m_{k,j},\,j=1\ldots J_k\,\}$ of all Voronoi cells that possibly overlap that block. Determining these lists in principle requires an intersection test between each block and each cell. In practice it suffices to consider the cell's bounding box, which can be easily intersected with the blocks.

One might be tempted to derive a Voronoi cell's bounding box from the positions of the neighboring sites; however, the convex hull of a cell's neighboring sites does not necessarily fully enclose the cell. Because a Voronoi cell is convex, its bounding box can be easily calculated from the list of its vertices. This requires fully constructing the cell; however, this is needed anyway to calculate the cell volume for use in other areas of the RT simulation (e.g., determining the specific energy absorbed per unit mass by the medium in the cell). Regardless, the cell geometry is needed solely during setup and does not have to be retained thereafter.

The function $\mathcal{C}(\vec{r})$ receives the following input data: the query point $\vec{r}$; for each block $\mathcal{B}_k$, the indices $\{\,m_{k,j},\,j=1\ldots J_k\,\}$ of all Voronoi cells that possibly overlap that block; the positions of the sites $\{\,\vec{p}_m,\,m=1\ldots M\,\}$; and the domain boundaries $(\vec{D}_{\min}, \vec{D}_{\max})$.

For each query, the function $\mathcal{C}(\vec{r})$ performs these steps:

1. Verify that the query point is inside the domain; if not return a special (negative) index value.

2. Locate the block containing the query point; this is trivial since blocks are on a regular linear grid.

3. Retrieve the list of cells possibly overlapping that block, and thus possibly containing the query point.

4. Calculate the squared distance from the query point to the sites for each of these cells. By definition of a Voronoi tesselation, the closest site determines the Voronoi cell containing the query point.

3.3	Tests, results, and discussion

3.3.1   Implementation

We implemented a Voronoi dust grid in SKIRT according to the method presented in Sect. 3.2. This allowed us to use the built-in geometries for creating synthetic test models, and to compare the results with those produced by the existing and well-tested grids.

We employed the open source library Voro++ (Rycroft 2009) to set up the input data described in Sects. 3.2.2 and 3.2.3. The library and its data structures are used only during setup. All relevant information is extracted and stored in our own data structures for reference after setup.

3.3.2   Test models

We tested the Voronoi dust grid with two synthetic models of our own making, called *torus* and *spiral*, and we ran the RT benchmark described by Pascucci et al. (2004). We first present the results for our models, and in Sect. 3.3.5 we discuss the results for the Pascucci benchmark.

The *torus* model consists of a central light source surrounded by an axisymmetric dusty torus, as might be present in the center of active galactic nuclei. The dust geometry is described by a radial power-law density from a given inner to outer radius, with an opening angle of 50 degrees. A cut through the dust distribution is shown in the top row of Fig. 3.2. The sites for the Voronoi dust grid are selected randomly from a uniform distribution over the cuboidal domain enclosing the torus. Since the model is axisymmetric, we can compare the results of the Voronoi grid with those produced by a regular two-dimensional (2D) cylindrical grid, in addition to those produced by an adaptive (3D) octree grid. In Fig. 3.3 we illustrate the effect of the number of Voronoi grid cells for the torus model.

The *spiral* model represents an idealized spiral galaxy with three arms, similar to the spiral model presented in Saftly et al. (2013). The stellar distribution includes a flattened Sérsic bulge and a double-exponential disk with a spiral arm perturbation. The dust is distributed in a thinner, similarly perturbed double-exponential disk. Cuts through the dust distribution are shown in the top row of Figs. 3.4 and 3.5. In this case the sites for the Voronoi dust grid are selected randomly from the dust distribution, as opposed to a uniform distribution. Areas with a higher dust density are thus, on average, covered with smaller cells.

3.3.3   Test grids

For the torus model we ran simulations with three different dust grids: a regular 2D cylindrical grid with $250^2 = 62\,500$ cells; an adaptive octree grid with $\approx 950\,000$ cells; and a Voronoi grid with about the same number of uniformly distributed cells. The
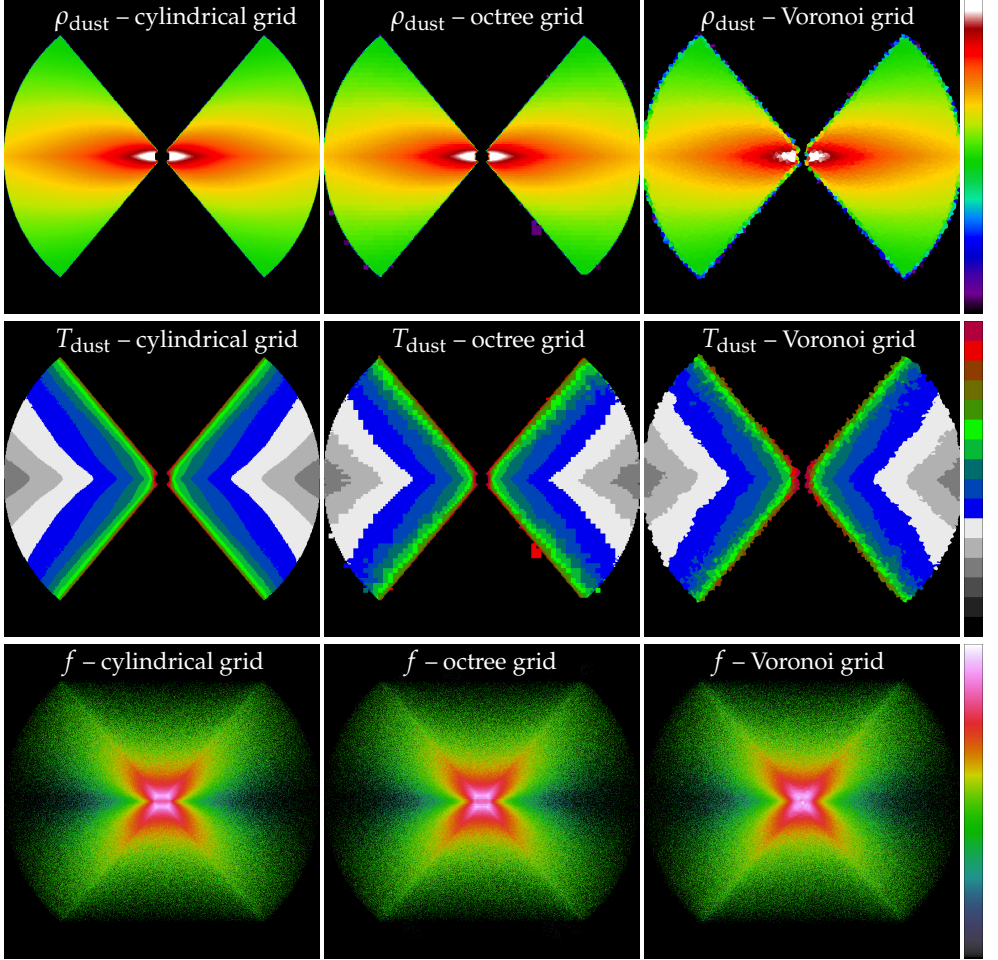
Figure 3.2: Illustration of the results for the *torus* model with three different dust grids. **Rows** – *top*: the dust density distribution (cut through the central edge-on plane); *middle*: the calculated dust temperature (cut through the central edge-on plane); *bottom*: the calculated flux density escaping from the model (edge-on view). **Columns** – *left*: regular 2D cylindrical grid with $250^2 = 62\,500$ cells; *middle*: adaptive octree grid with $\approx 950\,000$ cells; *right*: Voronoi grid with $\approx 950\,000$ uniformly distributed cells.
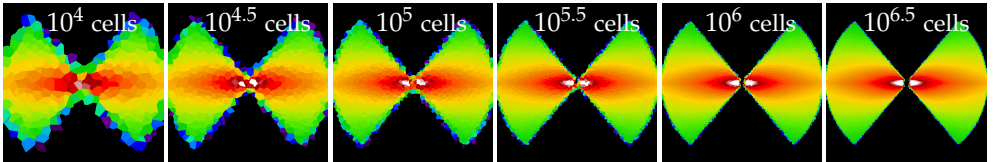
Figure 3.3: A cut through the dust density distribution of the torus model, discretized on Voronoi grids with a resolution varying from 10^4 cells (*left*) to $10^{6.5}$ cells (*right*). All grids were constructed from a set of uniformly distributed sites.
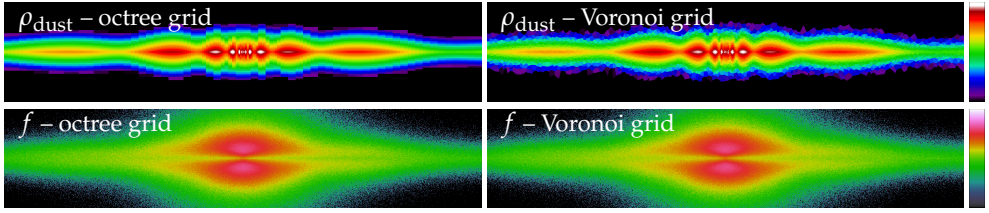
Figure 3.4: Illustration of the results for the *spiral* model, edge-on view. **Rows** – *top*: the dust density distribution (cut through the central edge-on plane); *bottom*: the calculated flux density escaping from the model (edge-on view). **Columns** – *left*: adaptive octree grid with ≈ 1 350 000 cells; *right*: Voronoi grid with ≈ 1 350 000 cells with a non-uniform, weighed distribution.
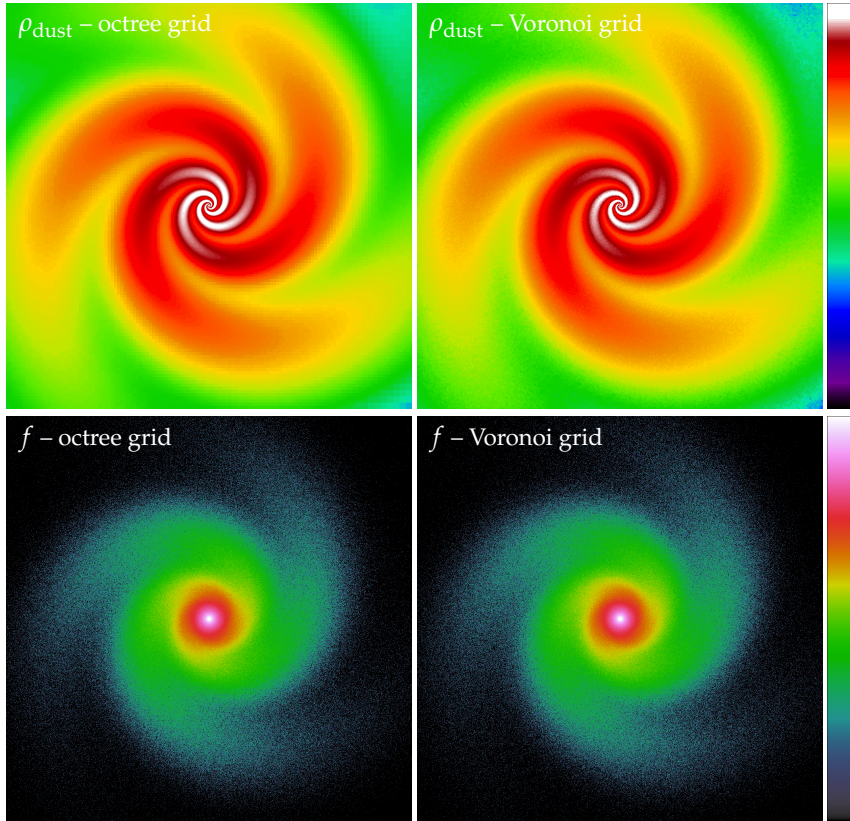
Figure 3.5: Illustration of the results for the *spiral* model, face-on view. **Rows** – *top*: the dust density distribution (cut through the central face-on plane); *bottom*: the calculated flux density escaping from the model (face-on view). **Columns** – *left*: adaptive octree grid with ≈ 1 350 000 cells; *right*: Voronoi grid with ≈ 1 350 000 cells with a non-uniform, weighed distribution.

Table 3.1: *Grid quality.* The difference between the theoretical and gridded dust density is sampled at a large number of random points, uniformly distributed over the domain. The standard deviation on this difference is used as a quality measure for the grid. In the table, the value for the octree grid is normalized to unity for each model. *Smaller numbers indicate better quality.*

Model	Cylindrical	Octree	Voronoi
Torus	0.82	1	1.75
Spiral	–	1	1.68

top row of Fig. 3.2 shows a cut through the gridded dust density distribution for each of these grids. The cylindrical grid captures the sharp edges of the model perfectly, because the cylindrical coordinate axes are lined up with the edges. The octree grid does a fine job as well because of its adaptive nature: smaller cells are automatically created along the sharp edges. The Voronoi grid does not do particularly well at the edges because of the random placement of its cells. This would not be an issue when importing a grid from a moving mesh code, because the cell sizes would already be properly adjusted to the underlying gradients.

For the spiral model we ran simulations with two different dust grids: an adaptive octree grid with $\approx 1\,350\,000$ cells; and a Voronoi grid with about the same number of cells, placed using a weighed distribution according to the dust density (smaller cells in higher density areas). The top rows of Figs. 3.4 and 3.5 show a cut through the gridded dust density distribution for each of these grids. The differences between the grids are most easily seen in the lower density areas.

Although this study does not focus on grid quality, we still need to ensure that our Voronoi grid implementation properly represents the theoretical dust densities defined by the synthetic models. To obtain an objective quality measure, we sample the theoretical dust density ρ_t and the gridded dust density ρ_g at a large number of random points uniformly distributed over the domain. We use the standard deviation of the difference $\rho_t - \rho_g$ as a measure for how well the grid reflects the theoretical density distribution. Table 3.1 lists the resulting numbers for the various grids and models. For each model the value for the octree grid is normalized to unity.

Taking into account our naive cell placement, the Voronoi grid compares well with the highly tuned adaptive octree grid, thus verifying this aspect of our implementation.

3.3.4 Results

Shooting photon packages through the grid is the most important test in the context of this study.

The middle row of Fig. 3.2 shows the dust temperature calculated by a panchromatic simulation for the torus model, using the three grids describe above. All quantities, including the radiation field and the amount of dust absorption, are discretized

Table 3.2: *Run time.* The elapsed time for the photon shooting phase of a simulation is divided by the number of grid cells crossed during that phase. The result is an indication of the time spent per cell crossing, including grid traversal calculations and some overhead for generating the random paths and for storing results. The tests were performed on a typical desktop computer using a single core. The last column lists the ratio between the run times for the Voronoi and octree grids. *Larger numbers indicate slower performance.*

| | | Time per cell crossing (ns) | | |
Model	Simulation type	Octree	Voronoi	Vor./Oct.
Torus	monochromatic	219	693	3.2
Torus	panchromatic	400	1006	2.5
Spiral	monochromatic	309	903	2.9
Spiral	panchromatic	442	1095	2.5

on the same grid as the dust density. In each simulation, the central light source emits 10^5 photon packages for each of 100 wavelength bins on a logarithmic grid. Scattering events cause additional photon packages to be created, which is particularly relevant for this model because of the high optical depth of the torus. In the end, each simulation traces about 700 million photon packages through the dust grid.

The bottom rows of Figs. 3.2, 3.4, and 3.5 show the flux density calculated by a monochromatic simulation for each model and grid combination. The Poisson noise is caused by the statistical nature of the Monte Carlo technique. In each simulation, the light sources emit 10 million photon packages at a fixed wavelength, and scattering events again cause additional photon packages to be created.

Other than the effects of grid resolution and the unavoidable noise, the calculated temperature and flux density maps are the same for the various grids. In particular, as noted in Sect. 3.3.3, the Voronoi grid does not resolve the central area of the dust distribution as well as the other grids, causing some deviation in the central area of the calculated flux density field. This effect is ultimately due to the naive placement of the Voronoi cells in our tests, and would not be present for a properly adjusted grid.

These results validate the accuracy of our straight path calculation method for Voronoi grids.

Table 3.2 provides an indication of the processing time spent per cell crossing for each simulation. To obtain these numbers, the elapsed time for the photon shooting phase of a simulation is divided by the number of grid cells crossed during that phase. The result thus includes some overhead for generating the random paths and for storing results, in addition to the grid traversal calculation itself. The tests were performed on a typical desktop computer using a single core. The last column lists the ratio between the cell crossing times for the Voronoi and octree grids. The Voronoi grid performs roughly three times slower than our highly optimized octree implementation (which maintains, for example, a neighbor list for each cell to accelerate the process of finding the next cell on a path). This seems surprisingly fast in view of the high geometric
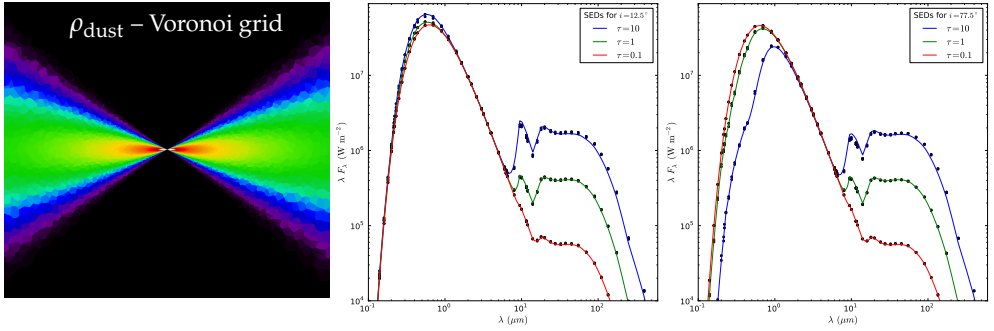
Figure 3.6: Illustration of the results for the *Pascucci* benchmark ([Pascucci et al. 2004](#)). The *left* panel shows a cut through the central edge-on plane of the dust density distribution discretized on a 3D Voronoi grid with one million cells randomly placed according to a $1/r$ distribution. The other panels show the simulated spectral energy distribution (SED) for disk inclinations equal to $12.5°$ (*center*) and $77.5°$ (*right*), for optical depths $\tau = 0.1, 1$, and 10. Dots indicate benchmark reference points; solid lines represent our simulation results using the 3D Voronoi grid shown in the left panel.

complexity of a Voronoi grid (illustrated in Fig. 3.1) compared to the cuboidal cells in an octree. Moreover, as noted in the introduction, an octree grid may need many more cells than the Voronoi grid to represent a particular density field, further balancing performance in favor of the Voronoi grid.

As discussed in Sect. 3.2.2, the cell crossing algorithm may occasionally fail to find an exit point because of computational inaccuracies. In our tests this occurred at most once per 50 million cell crossings, so this issue does not affect the algorithm's performance.

3.3.5   The Pascucci benchmark

The *Pascucci* benchmark ([Pascucci et al. 2004](#)) models a star embedded in a circumstellar disk with an inner cavity free of dust, prescribing an analytical 2D distribution and a set of optical depths and viewing angles. A cut through the central edge-on plane of the dust density distribution is shown in the left panel of Fig. 3.6.

We ran panchromatic simulations for this model with optical depths $\tau = 0.1, 1$, and 10 using a 3D Voronoi grid consisting of one million cells randomly placed according to a $1/r$ distribution. This distribution serves to properly resolve the intense radiation field in the center of the model. In each simulation, the central light source emits 10^5 photon packages for each of 150 wavelength bins on a logarithmic grid.

The two rightmost panels of Fig. 3.6 compare the spectral energy distribution (SED) produced by our SKIRT simulations with the corresponding benchmark results published in [Pascucci et al.](#) (2004). The center panel shows the SEDs for the various optical depths at a nearly face-on disk inclination of $12.5°$, the right panel at a nearly edge-on inclination of $77.5°$. Dots indicate benchmark reference points; solid lines represent our simulation results.

For higher optical depths our simulation results deviate slightly because the SKIRT code is not optimized for operation in this regime; running the benchmark with a 2D axially symmetric logarithmic grid results in the same deviation (not shown). These results further validate our method for calculating straight paths through a Voronoi grid.

3.3.6   Applicability

From Sect. 3.2 it follows that the presented method requires as input data solely the coordinates of the Voronoi sites, plus any relevant physical properties (such as mass densities) for the cell surrounding each site. In other words, the interface between the input model and the RT code is very thin, opening up a wide range of possibilities. An input model can be defined by SPH particles, serving as Voronoi sites; or by a Voronoi mesh produced by an MHD code; or by appropriately distributed random points generated from (semi)-analytical density or opacity fields, similar to the approach in Paardekooper et al. (2010), for example.

We also note in Sects. 3.2.2 and 3.2.3 that the path calculation algorithm itself requires no information on a Voronoi cell other than its bounding box and the locations of its own site and all neighboring sites. The required data structures can be easily built from the input data using a publicly available Voronoi library. The library code is invoked only during the initialization phase, minimizing its impact on performance and robustness, and allowing it to be easily replaced by another code if the need arises. For example, while we are happy with the Voro++ library's ease of use and with its performance during the tests, we may in the future consider using a parallelized method (Lo 2012; Springel 2010).

As a consequence, the presented method allows the RT code to support a Voronoi grid while remaining uncoupled from the code producing the input model. This decreases the complexity of the interface, and allows cooperation even when source code is not publicly available. In contrast to this approach, for example, the Sunrise RT code (Jonsson 2006) directly invokes parts of the non-public Arepo moving mesh code (Springel 2010) to implement the interface and to build its Voronoi grid.

3.4   Conclusions

The choice of an appropriate discretization is crucial in any numerical simulation code. Because of the large dynamic range of the physical quantities, in most problems the resolution of the grid must scale with the field densities or gradients. Adaptive grids with cuboidal cells, such as octrees or more generally AMR grids, have proven very popular in part because of their relative ease of implementation. However, several recent codes have adopted unstructured grids based on Voronoi tessellations, or equivalently, Delaunay triangulations. These grids tend to more closely reflect dynamic ranges in the model with fewer cells, presenting cell boundaries that are more adjusted to the underlying gradients. Since a Voronoi grid is defined solely by

its generating points, the cell size and distribution can be easily fine-tuned by placing these sites in the appropriate locations.

In a RT simulation the Voronoi grid can be a very flexible tool. Appropriate sites can be generated randomly, distributed according to the input model's density or opacity fields; if needed extra sites can be added in high-gradient areas. In the case of a particle-based input model, the particle locations themselves can serve as sites; and for an input model already based on a Voronoi mesh no re-gridding is required at all.

In this work we have shown that it is straightforward to implement accurate and efficient RT on Voronoi grids. In spite of the geometric complexity of the cell boundaries, calculating straight paths between two arbitrary points through a 3D Voronoi grid is only about three times slower than a highly optimized octree implementation with the same number of cells, while in practice the total number of Voronoi grid cells may be lower for an equally good representation of the density field. The presented method automatically yields the precise distance covered by the path inside each grid cell, and eliminates the need for approximate corrections or work-arounds required by alternate approaches where the radiation travels only along the Delaunay edges. The method requires only a thin interface with the input model and with the actual construction of the grid, allowing codes to remain largely uncoupled and enabling the use of a publicly available Voronoi library.

While we implemented and tested the method in our continuum RT code SKIRT, focusing on the effects of dust, it is widely applicable to all RT codes using ray tracing or Monte Carlo techniques.

We conclude that the benefits of using a Voronoi grid in RT simulation codes will often outweigh the somewhat slower performance.

Stochastic heating of dust grains[11]

4

THE TRUTH OF A PROPOSITION HAS NOTHING TO DO WITH ITS CREDIBILITY. AND VICE VERSA.

Thermal emission by stochastically heated dust grains (SHGs) plays an important role in the radiative transfer (RT) problem for a dusty medium. In this chapter, we define an appropriate problem for benchmarking dust emissivity calculations in the context of RT simulations, specifically including the emission from SHGs. We process this problem using six RT codes participating in the benchmark effort and compare the results to a reference solution computed with the publicly available dust emission code DustEM. Our aim is to provide a self-contained guide for implementors of such functionality, and to offer insight into the effects of the various approximations and heuristics implemented by the participating codes to accelerate the calculations. We conclude that the relevant modules in RT codes can and do produce fairly consistent results for the emissivity spectra of SHGs. This work will pave the way for a more extensive benchmark effort focusing on the RT aspects of the various codes.

4.1	Introduction

While larger dust grains can often be assumed to be in local thermal equilibrium (LTE) with the surrounding radiation field, this assumption does not typically hold for very small grains (VSGs) or for polycyclic aromatic hydrocarbon molecules (PAHs). The absorption of a single optical photon substantially boosts the internal energy of such a small collection of atoms, causing its emission spectrum to vary over time (see, e.g., Sellgren 1984; Dwek 1986; Boulanger and Perault 1988; Helou et al. 2000; Draine et al. 2007; Smith et al. 2007). We use the term *stochastically heated grains* (SHGs) to collectively indicate dust grains and PAH molecules that cannot be assumed to be in LTE with the radiation field.

Since SHGs spend a significant amount of time at much higher energy levels than if

---

[11] Published as Camps et al. (2015). I wrote the method summary (based on information available in the literature), implemented the method in SKIRT, defined the benchmark input fields, produced the reference solutions using DustEM, coordinated the benchmark effort, collected the various code results from the respective co-authors, produced the plots, and wrote most of the analysis in the results and discussion section, often based on suggestions from the co-authors. Karel Misselt contributed the dust model definition, including the related data files. Co-author(s) contributed the sections describing each participating code, and ran the benchmark problem with their own code, providing the results to me and adjusting their codes as needed in an iterative process.

they are in LTE, they emit at shorter wavelengths, which can have a strong effect on the observed spectrum. It is therefore essential to verify that RT codes properly calculate this emission before studying the effects of spatial distribution and other model parameters on the simulated observables. In this chapter we present a benchmark problem for this purpose, including a dust model, a series of input radiation fields, and a reference dust emission spectrum for each input field. The dust model described here has been designed for use in the TRUST benchmarks[12] (Transport of Radiation through a dUSTy medium), which test the actual RT aspect of various codes.

A typical 3D RT simulation calculates the dust emission spectra for millions of dust cells (or at least for many thousands of library items that are representative of the cells). When calculating the emission spectrum for the dust grain population in a particular cell, the first task is to determine the temperature probability distribution of the grains, given the grain sizes and chemical compositions, and given the radiation field in the cell. This is also the computationally most demanding part of the calculation. In the current epoch, performing a full treatment of vibrational quantum modes for the dust grains in each cell is computationally prohibitive. In practice, RT codes use the continuous cooling assumption, which was shown to provide a good approximation in areas relevant for RT through the interstellar medium (Draine and Li 2001).

The reference solutions for this benchmark were produced with the public version of the DustEM code (Compiègne et al. 2011). DustEM determines the grain temperature distribution by iteratively solving an integral equation (Desert et al. 1986). We compare the reference solutions with the dust emission spectra calculated by six distinct RT codes. These codes determine the grain temperature distribution by solving a set of linear equations (Guhathakurta and Draine 1989). While this method is inherently faster (Guhathakurta and Draine 1989), it still becomes very expensive when the grains are in LTE with the radiation field. To avoid this problem, care must be taken to properly transition the calculation from the stochastic to the equilibrium regime.

More generally, the need for fast dust emission calculations has prompted the authors of RT simulation codes to implement various acceleration techniques, discretization choices, approximations, and heuristics. We study the effects of these mechanisms on the calculated solutions and quantify the level of (dis)agreement between the participating codes, with the objective of helping to inform the interpretation of RT simulation results that include SHG dust emission calculations of the type presented here.

The information provided in this chapter and in the accompanying data files[13] is self-contained. Readers can implement the code to calculate the SHG emission, set up the benchmark tests, and verify the results, based solely on the information provided here, without referring to other sources.

In Sects. 4.2 and 4.3 we define the dust model and the input radiation fields used in

---

[12] TRUST benchmarks: http://ipag.osug.fr/RT13/RTTRUST/
[13] SHG benchmark data: http://www.shg.ugent.be

Table 4.1: Symbols used in this chapter for various physical quantities, with the corresponding SI units.

Symbol	Units	Description
λ	m	Wavelength
s	m	Distance along a path
V	m^3	Volume
M	kg	Dust mass
T	K	Temperature
τ	s	Interaction timescale
ρ	$kg\,m^{-3}$	Dust mass density
$\sigma^{abs,sca,ext}$	m^2	Cross section (absorption, scattering or extinction)
$\mathcal{N}_H$	H	Number of hydrogen atoms
$n_H = \mathcal{N}_H/V$	$H\,m^{-3}$	Hydrogen atom number density
$\mu = M/\mathcal{N}_H$	$kg\,H^{-1}$	Dust mass per hydrogen atom
$\varsigma = \sigma/\mathcal{N}_H$	$m^2\,H^{-1}$	Cross section per hydrogen atom
$\kappa = \sigma/M = \varsigma/\mu$	$m^2\,kg^{-1}$	Mass coefficient
a	m	Dust grain size
$\mathcal{N}_D$	1	Number of dust grains
$n_D = \mathcal{N}_D/V$	m^{-3}	Dust grain number density
$\Omega(a) = (\frac{dn_D}{da})/n_H$	$m^{-1}\,H^{-1}$	Dust grain size distribution per hydrogen atom
$Q^{abs,sca,ext}$	1	Efficiency
ρ^{bulk}	$kg\,m^{-3}$	Bulk mass density of grain material
c	$J\,K^{-1}\,kg^{-1}$	Specific heat capacity of grain material
h	$J\,kg^{-1}$	Specific enthalpy (internal energy per unit mass)
H	J	Enthalpy (internal energy) of a dust grain
J	$W\,m^{-3}\,sr^{-1}$	Mean spectral radiance (intensity of radiation field)
B	$W\,m^{-3}\,sr^{-1}$	Black-body spectral radiance (Planck's law)
U	1	Radiation field strength relative to solar neighborhood
j	$W\,m^{-1}\,sr^{-1}$	Spectral dust emission
$\varepsilon = j/\mathcal{N}_H$	$W\,m^{-1}\,sr^{-1}\,H^{-1}$	Spectral dust emission per hydrogen atom

the benchmark problem. Section 4.4 presents the linear equation method used by the RT codes represented in this work to calculate the SHG emission spectra, and Sect. 4.5 briefly introduces each of these codes. In Sect. 4.6 we compare the results produced by the RT codes with the reference solution, and we discuss the differences between the methods used by the various codes and how they influence the results. Finally, in Sect. 4.7 we summarize and conclude.

For ease of reference, Table 4.1 lists the symbols used in this chapter for various physical quantities, with the corresponding SI units.

4.2	Dust model

The exact choice of dust grain model (e.g., Weingartner and Draine 2001a; Zubko et al. 2004; Jones et al. 2013) is not critical for benchmark purposes, as long as all codes employ the same model. Specifically, our choices do not imply a preference for or an endorsement of a particular model. Still, to properly evaluate the results of our calculations and the effects of approximations in the context of real-world RT simulations of astrophysical systems, it is important to use grain compositions and size distributions that are in agreement with observational constraints, rather than defining arbitrary synthetic grain properties. With this in mind, we have elected to utilize the simple BARE-GR-S model of Zubko et al. (2004). This model uses a mixture of spherical, single composition (BARE) graphitic (GR) and silicate (S) dust grains. The graphitic grains include both graphite grains and astronomical PAH molecules. The relative populations of the three components and their size distribution has been optimized to reproduce observations characteristic of the diffuse Milky Way interstellar medium with $R_V = 3.1$. The observational constraints on the model include the shape of the wavelength-dependent extinction, the infrared emission, and the abundance of elements locked up in the solid phase.

Assuming that the scattering process is elastic, so that the internal energy of the interacting dust grain remains unaffected, scattering is irrelevant for the benchmark problem presented in this work. Thus the subsequent discussion does not focus on the scattering properties of the dust model.

4.2.1   Optical grain properties

For both the graphite and silicate components, optical properties were computed directly from the dielectric functions (refractive indices) of the bulk material using a Mie code originally provided by Viktor Zubko with some small modifications. The optical properties include absorption efficiencies $Q_k^{\mathrm{abs}}(a, \lambda)$ as a function of composition k, grain size a, and wavelength λ.

While the dielectric functions are in general functions of both temperature and size (Draine and Lee 1984), to minimize free parameters and in keeping with essentially all astronomical applications of grain properties, we adopt a single set of dielectric functions at a specific temperature and size for each component. The graphitic indices were taken from Draine (2003) for 0.1 μm grains at 20 K. Calculations were carried out for both parallel and perpendicular orientations relative to the basal plane and combined using the 1/3-2/3 approximation (Draine 1988). For the silicate grains, the dielectric functions for smoothed astronomical silicates (0.1 μm) (Laor and Draine 1993; Weingartner and Draine 2001a) were used in the Mie calculations. With these sets of dielectric functions, and assuming spherical grains, the efficiencies can be calculated for each component as a function of size and wavelength. For each component, the Mie calculation was carried out for 121 logarithmically spaced sizes between 0.00035 μm and 100 μm. For each size, optical properties were computed at

1201 wavelength points logarithmically spaced between $0.001\,\mu m$ and $10000\,\mu m$. The efficiencies thus derived can be cast as either cross sections or mass coefficients:

$$\sigma_k^{abs,sca}(a,\lambda) = \pi a^2 Q_k^{abs,sca}(a,\lambda) \tag{4.1}$$

$$\kappa_k^{abs,sca}(a,\lambda) = \frac{3 Q_k^{abs,sca}(a,\lambda)}{4 a \rho_k}. \tag{4.2}$$

Because there are no refractive indices for the class of materials commonly referred to as astronomical PAH, the optical properties for these materials are simply constructed so as to reproduce the available observations. Therefore, the PAH cross sections were computed following Draine and Li (2007). The utilization of the Draine and Li (2007) formulation for the cross sections represents a deviation from a pure Zubko et al. (2004) BARE-GR-S model in the same way as Zubko et al. (2004) utilized the Li and Draine (2001) PAH cross sections.

In Draine and Li (2007), the PAH cross sections were updated to reflect knowledge gained regarding the shape of the PAH emission spectrum from the wealth of Spitzer observations available. Also of note is that the PAH component in the model defined here consists of pure neutral PAHs. No attempt was made to account for the varying ionization fraction of a PAH molecule as a function of effective PAH size and radiation field. The PAH optical properties were computed on the same wavelength grid as the graphite and silicate components of the model. However, the size grid was of course altered with properties computed for 28 logarithmically spaced sizes between $0.00035\,\mu m$ and $0.006\,\mu m$.

4.2.2 Grain size distributions

The relative contribution of each material in a dust model is set by the grain size distribution for that particular material. The shape of the size distributions in the BARE-GR-S model of Zubko et al. (2004) matches observations of the interstellar medium in the solar neighborhood, including the amount of refractory material available and the wavelength dependence of the Milky Way's diffuse ($R_V = 3.1$) extinction and emission spectrum. Zubko et al. (2004) provide convenient analytic functional approximations to the size distributions for each component of the model, which take the form

$$\log O(a) = c_0 + b_0 \log\left(\frac{a}{\mu m}\right) - b_1 \left|\log\left(\frac{a}{a_1}\right)\right|^{m_1} - b_2 \left|\log\left(\frac{a}{a_2}\right)\right|^{m_2}$$
$$- b_3 \left|\frac{a - a_3}{\mu m}\right|^{m_3} - b_4 \left|\frac{a - a_4}{\mu m}\right|^{m_4} \tag{4.3}$$

where a is the grain size and a_i, b_i, c_i, m_i represent constant parameters. In the interest of clarity, the notation in Eq. 4.3 omits a subscript k indicating the type of material being referenced. For the BARE-GR-S model, we require three sets of parameters, one for each of the materials comprising the dust model. The appropriate values can be

Table 4.2: Parameters for the analytical approximation to the BARE-GR-S size distribution defined in Eqs. 4.3 through 4.5 and the bulk densities for the three components in our dust model. With these parameter values, the grain size a substituted in the equations must be expressed in μm and the resulting size distribution $\Omega_k(a)$ is expressed in number of dust grains per hydrogen atom and per μm.

Parameter	Units	PAH	Graphite	Silicate
$a^{\min}$	μm	0.00035	0.00035	0.00035
$a^{\max}$	μm	0.005	0.33	0.37
A	$\mu\text{m}^{-1}\,\text{H}^{-1}$	2.227433×10^{-7}	1.905816×10^{-7}	1.471288×10^{-7}
c_0	1	-8.02895	-9.86000	-8.47091
b_0	1	-3.45764	-5.02082	-3.68708
b_1	1	1.18396×10^3	5.81215×10^{-3}	2.37316×10^{-5}
a_1	μm	1	0.415861	7.64943×10^{-3}
m_1	1	-8.20551	4.63229	22.5489
b_2	1	0	0	0
a_2	μm	1	1	1
m_2	1	1	1	1
b_3	1	1.0×10^{24}	1.12502×10^3	2.96128×10^3
a_3	μm	-5.29496×10^{-3}	0.160344	0.480229
m_3	1	12.0146	3.69897	12.1717
b_4	1	0	1.12602×10^3	0
a_4	μm	0	0.160501	0
m_4	1	1	3.69967	1
ρ^{bulk}	$\text{kg}\,\text{m}^{-3}$	2240	2240	3500

found in Table 7 of Zubko et al. (2004) and are reproduced here in Table 4.2.

By construction, Eq. 4.3 has been normalized such that

$$\int_{a_k^{\min}}^{a_k^{\max}} O_k(a)\,\mathrm{d}a = 1 \tag{4.4}$$

where $a_k^{\min}$ and $a_k^{\max}$ are the minimum and maximum sizes over which the size distribution for component k is defined. The complete size distribution function is then given by

$$\Omega_k(a) = A_k\,O_k(a) \tag{4.5}$$

where A_k is the overall normalization of component k, as listed in Table 4.2.

With this definition of the grain size distribution, the total number of dust grains per hydrogen atom and the total dust mass per hydrogen atom are given by

$$\mathcal{N}_\mathrm{D}/\mathcal{N}_\mathrm{H} = \sum_k \int_{a_k^{\min}}^{a_k^{\max}} \Omega_k(a)\,\mathrm{d}a \tag{4.6}$$

$$\mu = M/\mathcal{N}_\mathrm{H} = \sum_k \int_{a_k^{\min}}^{a_k^{\max}} \frac{4}{3}\pi a^3 \rho_k^{\text{bulk}} \Omega_k(a)\,\mathrm{d}a. \tag{4.7}$$

We also provide the grain size distributions defined by Eqs. 4.3 and 4.5 and Table 4.2 in tabulated form (see Sect. 4.2.4), on a size grid with 24, 62, and 63 samples for the PAH, graphite, and silicate component, respectively.

4.2.3   Calorimetric grain properties

The internal energy of a dust grain of composition k and its temperature are related via the heat capacity of the grain material:

$$H_k(a, T) = \frac{4}{3}\pi a^3 \rho_k^{\text{bulk}} h_k(T) \tag{4.8}$$

$$h_k(T) = \int_0^T c_k(T')\,\mathrm{d}T' \tag{4.9}$$

where $H(a, T)$ is the internal energy (enthalpy) of a grain with size a at temperature T, $h(T)$ is the specific enthalpy (per unit mass), $c(T)$ is the specific heat capacity, and ρ^{bulk} is the bulk density of the grain material.

We elected to use the heat capacity functions proposed in Draine and Li (2001) and Li and Draine (2001). To avoid subtle differences between implementations, we provide this information in tabulated form (see Sect. 4.2.4) rather than expecting each code to implement the equations. One table describes the graphitic components (PAH molecules and graphite grains) and another one describes the silicate component. Each table lists the specific enthalpy and the specific heat capacity at 1000 logarithmically spaced temperature points ranging from 1 K to 2500 K.

Finally, the bulk densities for the three components in our dust model are listed in Table 4.2.

4.2.4   Data files

The data files defining the dust properties described above can be downloaded from the web site indicated in footnote 13. They are contained in the `DustModel` directory, which is organized in the following subdirectories:

- `GrainInputs`: optical and calorimetric properties for each of the three grain types k. Optical properties include the absorption efficiencies $Q_k^{\text{abs}}(a, \lambda)$ tabulated on a grid of 1201 wavelengths λ and 28 (for PAH) or 121 (for graphite and silicate) grain sizes a. Calorimetric properties include the specific enthalpy $h(T)$ and the specific heat capacity $c(T)$ of the grain material, tabulated on a grid of 1000 temperatures T. The calorimetry data for graphitic grains should be used for PAH molecules as well.

- `SizeInputs`: tabulated grain size distributions $\Omega_k(a)$ for each grain type k, on a size grid with 24, 62, and 63 samples for the PAH, graphite, and silicate component, respectively. Implementations may choose to compute the size distribution from the functional form defined in Sect. 4.2.2, or to load the tabulated data.

- `EffectiveGrain`: size-integrated values of the optical properties. This information is not needed for calculating dust emission.

- `ScatMatrix`: scattering matrix elements. This information is not needed for calculating dust emission.

4.3	Radiation fields

4.3.1   Basic definitions

The spectral radiation for a black body in thermal equilibrium at temperature T in function of wavelength λ is given by the Planck function

$$B(\lambda, T) = \frac{2hc^2}{\lambda^5} \frac{1}{\exp\left(\frac{hc}{\lambda kT}\right) - 1} \tag{4.10}$$

where h denotes the Planck constant, c the speed of light in vacuum, and k the Boltzmann constant.

We define the solar-neighborhood interstellar radiation field (ISRF) given in Table A3 of Mathis et al. (1983) through the following functional form inspired by (but not identical to[14]) Eq. 31 in Weingartner and Draine (2001b):

$$J^{\mathrm{Mat}}(\lambda) = \begin{cases} 0 \\ \qquad\qquad\qquad\qquad \lambda < 0.0912\,\mu\mathrm{m} \\ 3069\,\mathrm{W/m^3/sr} \times (\lambda/\mu\mathrm{m})^{3.4172} \\ \qquad\qquad 0.0912\,\mu\mathrm{m} \leqslant \lambda < 0.110\,\mu\mathrm{m} \\ 1.627\,\mathrm{W/m^3/sr} \\ \qquad\qquad 0.110\,\mu\mathrm{m} \leqslant \lambda < 0.134\,\mu\mathrm{m} \\ 0.0566\,\mathrm{W/m^3/sr} \times (\lambda/\mu\mathrm{m})^{-1.6678} \\ \qquad\qquad 0.134\,\mu\mathrm{m} \leqslant \lambda < 0.250\,\mu\mathrm{m} \\ 10^{-14}\,B(\lambda, 7500\,\mathrm{K}) + 10^{-13}\,B(\lambda, 4000\,\mathrm{K}) \\ \qquad\qquad + 4 \times 10^{-13}\,B(\lambda, 3000\,\mathrm{K}) \\ \qquad\qquad\qquad\qquad 0.250\,\mu\mathrm{m} \leqslant \lambda \end{cases} \tag{4.11}$$

The recipes in the referenced papers prescribe the total radiation field $4\pi J^{\mathrm{Mat}}$, whereas we prescribe the mean radiation field J^{Mat}. Based on this reference field, we define the strength U of an arbitrary radiation field $J(\lambda)$ as

$$U = \int_0^\infty J(\lambda)\,\mathrm{d}\lambda \bigg/ \int_0^\infty J^{\mathrm{Mat}}(\lambda)\,\mathrm{d}\lambda \tag{4.12}$$

---

[14] The Weingartner and Draine (2001b) equation is formulated in function of frequency rather than wavelength, and the dilution factor for the 4000 K black body listed in their Table 1 is not adjusted to the value specified in Sect. 2.1 of Mathis et al. (1983).

4.3.2   Benchmark input fields

In the benchmark described in this work, the dust grains are exposed to two sets of distinct radiation fields. The first set consists of eleven scaled versions of the Mathis ISRF, ranging from *weak* to *strong*, defined as

$$J^{\mathrm{SHG},i}(\lambda) = U_i \times J^{\mathrm{Mat}}(\lambda) \text{ with } U_i = 10^{-4}, 10^{-3}, ..., 10^5, 10^6 \qquad (4.13)$$

The second set consists of the following six diluted black body fields with varying temperatures, ranging from *soft* to *hard*:

$$J^{\mathrm{SHG},j}(\lambda) = \begin{cases} 8.28 \times 10^{-12} \, B(\lambda, 3000\,\mathrm{K}) \\ 2.23 \times 10^{-13} \, B(\lambda, 6000\,\mathrm{K}) \\ 2.99 \times 10^{-14} \, B(\lambda, 9000\,\mathrm{K}) \\ 7.23 \times 10^{-15} \, B(\lambda, 12000\,\mathrm{K}) \\ 2.36 \times 10^{-15} \, B(\lambda, 15000\,\mathrm{K}) \\ 9.42 \times 10^{-16} \, B(\lambda, 18000\,\mathrm{K}) \end{cases} \qquad (4.14)$$

The dilution factors were chosen so that the far-infrared peak of the dust emissivity is at the same level for all fields in this set (for ease of visualization), and so that all fields in the set have a strength of $1 \lessapprox U < 10$.

4.3.3   Calculation and wavelength grid

Using the dust model described in section 4.2, the codes participating in this benchmark calculate the spectral dust emissivity $\varepsilon(\lambda)$ for each of the input radiation fields specified in section 4.3, taking stochastic heating of small grains into account. The radiation emitted by the dust itself is ignored with respect to the input field; i.e. it is not our intention to calculate a self-consistent radiation field. The calculations were performed and the results written down using the wavelength grid on which the optical properties have been tabulated. This is a logarithmic grid with 1201 points in the range $0.001\,\mu\mathrm{m} \leqslant \lambda \leqslant 10000\,\mu\mathrm{m}$.

4.4	Dust emission

4.4.1   Emission from a dust mixture

The thermal emission of a dust grain depends nonlinearly on the grain size a, even in LTE conditions. It is therefore impossible to calculate the emission for a dust mixture with varying grain sizes from effective grain properties that would somehow represent the whole mixture (Steinacker et al. 2013; Wolf 2003a). Instead, we define a grid of grain size bins b for each dust model component k, and we choose an average, representative grain for each bin. We then proceed to calculate the emission as if each bin would contain only representative grains. For a sufficiently large number of bins,

this procedure converges to the proper result.

A simple approach is to represent each bin by a grain size at the arithmetic or geometric center of the bin. In a somewhat more sophisticated approach, the absorption cross section per hydrogen atom representative for a particular bin can be calculated by an integration over the size distribution:

$$\varsigma_{k,b}^{\mathrm{abs}}(\lambda) = \int_{a_{k,b}^{\mathrm{min}}}^{a_{k,b}^{\mathrm{max}}} \pi a^2 \, Q_k^{\mathrm{abs}}(a, \lambda) \, \Omega_k(a) \, \mathrm{d}a \tag{4.15}$$

where $[a_{k,b}^{\mathrm{min}}, a_{k,b}^{\mathrm{max}}]$ specifies the size range of bin b for dust model component k.

The representative mass of a dust grain in a particular bin can similarly be obtained from

$$M_{k,b} = \int_{a_{k,b}^{\mathrm{min}}}^{a_{k,b}^{\mathrm{max}}} \rho_k^{\mathrm{bulk}} \frac{4\pi}{3} a^3 \, \Omega_k(a) \, \mathrm{d}a \, \bigg/ \, \int_{a_{k,b}^{\mathrm{min}}}^{a_{k,b}^{\mathrm{max}}} \Omega_k(a) \, \mathrm{d}a \tag{4.16}$$

so that the enthalpy of a representative dust grain at temperature T is given by

$$H_{k,b}(T) = M_{k,b} \, h_k(T) \tag{4.17}$$

where $h_k(T)$ is the specific enthalpy of the grain material of dust component k at temperature T.

The emissivity per hydrogen atom from a dust mixture with grain type components k and grain size bins b exposed to a radiation field $J(\lambda)$, called the *input field*, can be expressed as a function of the representative grain properties as

$$\varepsilon(\lambda) = \sum_{k,b} \varsigma_{k,b}^{\mathrm{abs}}(\lambda) \int_0^\infty P_{k,b,J}(T) \, B(\lambda, T) \, \mathrm{d}T \tag{4.18}$$

where $B(\lambda, T)$ is the Planck function defined in Eq. 4.10, and $P_{k,b,J}(T)$ is the probability of finding the representative grain of bin k, b at temperature T.

The emission originating in a dust mixture with specified total mass M can then be written as

$$j(\lambda) = \frac{M}{\mu} \, \varepsilon(\lambda) \tag{4.19}$$

with μ given by Eq. 4.7. When combining the emission from various dust mixes, it is useful to recall that it is physically meaningful to add cross sections and masses, while mass coefficients, in general, cannot be added meaningfully:

$$\kappa_1 + \kappa_2 = \frac{\varsigma_1}{\mu_1} + \frac{\varsigma_2}{\mu_2} \neq \frac{\varsigma_1 + \varsigma_2}{\mu_1 + \mu_2}. \tag{4.20}$$

The challenge is thus to compute the probability distribution of grain temperatures, $P_{k,b,J}(T)$, which depends on the input radiation field in addition to the grain properties. See, for example, Fig. 4 of Draine and Li (2007) for an illustration of various temperature distribution curves. In this discussion, we characterize P as a function

of grain temperature. The temperature of a grain and its internal energy are related through Eq. 4.8, so we could equivalently characterize P as a function of internal grain energy.

4.4.2   Equilibrium heating dust emission

When the representative grain in bin k, b is in LTE with the surrounding radiation field $J(\lambda)$, the temperature probability distribution $P_{k,b,J}(T)$ becomes a delta function at the grain equilibrium temperature

$$P_{k,b,J}(T) = \delta(T - T_{k,b,J}^{\text{eq}}) \tag{4.21}$$

and the equilibrium temperature can be determined via the energy balance equation

$$\int_0^\infty \varsigma_{k,b}^{\text{abs}}(\lambda)\, J(\lambda)\, \mathrm{d}\lambda = \int_0^\infty \varsigma_{k,b}^{\text{abs}}(\lambda)\, B(\lambda, T_{k,b,J}^{\text{eq}})\, \mathrm{d}\lambda. \tag{4.22}$$

4.4.3   Stochastic heating dust emission

When a single photon absorption can significantly change the internal energy of a representative grain, the grain is not in LTE with the surrounding radiation field $J(\lambda)$. The grain is *stochastically heated*, and its state can no longer be characterized by a single temperature. In that case, we need to solve for the temperature probability distribution $P_{k,b,J}(T)$ to calculate the grain emission. The six RT codes benchmarked in this work employ the method described in Guhathakurta and Draine (1989), Siebenmorgen et al. (1992), Manske and Henning (1998), and Draine and Li (2001). For ease of reference, this section summarizes the method using the quantities and notation introduced in the previous sections of this work. We focus on a single grain size bin and a specific radiation field, dropping the indices k, b and J from the notation.

We select an appropriate temperature grid with N bins T_i, $i = 0, \ldots, N - 1$ (see Sect. 4.4.1). Our goal is to determine the probabilities $P_i \equiv P(T_i)$ that a grain resides in temperature bin i. To this end we construct a transition matrix $A_{f,i}$ that describes the probability per unit time for a grain to transfer from initial temperature bin i to final temperature bin f. Because the temperature of a grain and its internal energy are related through Eq. 4.8, we can define an equivalent frequency grid $h\nu_i \equiv H(T_i)$, $i = 0, \ldots, N - 1$. Similarly, the input radiation field J can be formulated as a function of frequency rather than wavelength.

The matrix element for a heating transition $(f > i)$ is then given by width of the final frequency bin, $\Delta\nu_f$, multiplied by the number of photons absorbed from the input field that carry the appropriate energy for the transition, $h\nu_{fi} \equiv h(\nu_f - \nu_i)$, or

$$A_{f,i} = \Delta\nu_f \times \frac{\varsigma^{\text{abs}}(\nu_{fi})\, 4\pi J(\nu_{fi})}{h\nu_{fi}}. \tag{4.23}$$

Using $h\nu = H$, $\lambda\nu = c$, and $\nu J(\nu) = \lambda J(\lambda)$, this can be easily transformed into

$$A_{f,i} = 4\pi \, \varsigma^{\text{abs}}(\lambda_{fi}) \, J(\lambda_{fi}) \, \frac{hc \, \Delta H_f}{\left[H(T_f) - H(T_i) \right]^3}, \tag{4.24}$$

where $H(T_f)$ and $H(T_i)$ are the enthalpies of a dust grain in the final and initial temperature bins, $\Delta H_f = H(T_f^{\text{max}}) - H(T_f^{\text{min}})$ is the enthalpy width of the final temperature bin, and λ_{fi} is the transition wavelength given by

$$\lambda_{fi} = \frac{hc}{H(T_f) - H(T_i)}. \tag{4.25}$$

We assume that a dust grain cools by radiating photons with an energy that is very small compared to the internal energy of the grain. Under this *continuous cooling* assumption, cooling transitions occur only to the next lower level, so that $A_{f,i} = 0$ for $f < i - 1$. We further assume that all energy is radiated at the same temperature, so that the transition matrix element for $f = i - 1$ can be approximated as

$$A_{i-1,i} = \frac{1}{h\nu_i - h\nu_{i-1}} \int_0^\infty \varsigma^{\text{abs}}(\nu) \, 4\pi B(\nu, T_i) \, d\nu. \tag{4.26}$$

Using the same substitutions as above this easily leads to

$$A_{i-1,i} = \frac{4\pi}{H(T_i) - H(T_{i-1})} \int_0^\infty \varsigma^{\text{abs}}(\lambda) \, B(\lambda, T_i) \, d\lambda. \tag{4.27}$$

The diagonal matrix elements are defined as $A_{i,i} = -\sum_{f \neq i} A_{f,i}$. However there is no need to explicitly calculate these values because they are not used in the final procedure.

Assuming a steady state situation, the probabilities P_i can be obtained from the transition matrix by solving the set of N linear equations

$$\sum_{i=0}^{N-1} A_{f,i} \, P_i = 0 \qquad f = 0, ..., N - 1 \tag{4.28}$$

along with the normalization condition

$$\sum_{i=0}^{N-1} P_i = 1 \tag{4.29}$$

where N is the number of temperature bins. Because the matrix values for $f < i - 1$ are zero, these equations can be solved by a recursive procedure of computational order $\mathcal{O}(N^2)$. To avoid the numerical instabilities caused by the negative diagonal elements, the procedure employs a well-chosen linear combination of the original equations. This leads to the following recursion relations for the adjusted matrix

elements $B_{f,i}$

$$B_{N-1,i} = A_{N-1,i} \qquad i = 0, \ldots, N-2 \qquad (4.30)$$

$$B_{f,i} = B_{f+1,i} + A_{f,i}, \quad f = N-2, \ldots, 1; i = 0, \ldots, f-1, \qquad (4.31)$$

for the unnormalized probability distribution X_i

$$X_0 = 1 \qquad (4.32)$$

$$X_i = \frac{\sum_{j=0}^{i-1} B_{i,j} X_j}{A_{i-1,i}} \qquad i = 1, \ldots, N-1, \qquad (4.33)$$

and finally for the normalized probabilities P_i

$$P_i = \frac{X_i}{\sum_{j=0}^{N-1} X_j} \qquad i = 0, \ldots, N-1. \qquad (4.34)$$

While this method seems rather straightforward, specific algorithmic approaches differ between codes. One important characteristic that tends to differ between implementations is the grid that discretizes the grain temperatures (or equivalently internal energy states) during the construction of $P(T)$. With a fixed grid, a range of probable temperatures is defined, e.g. $2.7\,\mathrm{K} < T < T_{\max}$, where $T_{\max}$ is chosen to exceed the sublimation temperature of the bulk material being considered and the interval is divided into N temperatures (e.g., Bianchi 2008). The distribution function is evaluated at that set of fixed internal energy states for all grains considered to be in the stochastic heating regime. The advantage of this approach is that many of the quantities used to generate $P(T)$ can be precomputed, reducing the computational requirements of the solution method. A disadvantage of this approach is that calculations are done over the full defined temperature range, including regions where $P(T)$ is negligible. Not only is this computationally inefficient, but it also results in poor resolution of the form of $P(T)$, especially for grains of intermediate size where $P(T)$ will be relatively narrowly distributed with T. One alternate approach is to dynamically define the temperature grid (e.g., Manske and Henning 1999; Misselt et al. 2001). In this iterative approach, a coarse and broad temperature grid is defined and $P(T)$ computed. The grid is refined based on $P(T)$; temperatures with low $P(T)$ are removed from the grid and $P(T)$ is recomputed on the new, smaller, more densely sampled grid. The grid refinement is continued until energy balance is achieved or the number of temperature points exceeds a predefined threshold. The advantage of this approach is that $P(T)$ is properly sampled for all grain sizes. The disadvantage is of course an increase in the computational load. This is amplified by the fact that the algorithm will naturally increase the number of temperature samples as the grain approaches the equilibrium regime, because $P(T)$ is increasingly peaked as the grain size increases.

A second important characteristic that tends to differ between implementations is the mechanism to transition from stochastic to equilibrium heating regime. The simplest

Table 4.3: Overview of the discretization parameters and heuristics used by participating codes. Refer to Sects. 4.5.2 through 4.5.7 for a more extensive description. Also, Fig. 4.9 further supplements the information in the last column of this table.

Code	Grain size bins (Sil/Gra/PAH)	Temperature bins	Heuristic to select or determine temperature grid	Heuristic to transition from stochastic to equilibrium regime
SKIRT	15/15/15	20/625/1250	one of 3 grids based on width of $P(T)$	based on width of $P(T)$
DIRTY	121/121/28	50-1000	iterative range & resolution adjustment	based on $\tau_{\mathrm{abs}}/\tau_{\mathrm{rad}}$
TRADING	20/20/8	80	fixed predefined grid	fixed at $a_{\mathrm{trans}} = 0.05\,\mu\mathrm{m}$
CRT	15/15/15	128	one of 6 grids based on $\mathcal{P}_{\mathrm{abs}}$	based on $\mathcal{P}_{\mathrm{abs}}$
MCFOST	63/62/24	300	fixed predefined grid	based on $\tau_{\mathrm{abs}}/\tau_{\mathrm{rad}}$
DART-Ray	63/62/24	200	iterative range adjustment	based on σ_T of Gaussian approximation

approach is to fix the grain size of the transition so that all grains with $a < a_{\mathrm{trans}}$ are considered to be stochastically heated and all those with $a \geq a_{\mathrm{trans}}$ are considered to be in the equilibrium regime, regardless of the true state of the grain. Since the appropriate transition point is a function of the radiation environment in addition to composition, this approach leads to errors in the treatment of the emission. However, with judicious selection of a_{trans} , e.g. $a_{\mathrm{trans}} \sim 0.01\,\mu\mathrm{m}$ (Bianchi 2008), the results can be acceptable at least for non-extreme field strengths. Alternatively, the characteristics of $P(T)$ can be used to terminate the stochastic heating treatment and transition to equilibrium heating. For example, in the case of a dynamic temperature grid as described above, the size at which the heating algorithm fails can be defined as a_{trans} (Misselt et al. 2001). A third method for determining a_{trans} is to compute the absorption and radiative timescales for each grain size in the considered radiation field (Draine and Li 2001). These timescales are a natural physical metric since stochastic heating occurs when the mean time between photon absorptions is long compared to the time the grain takes to radiatively cool. The ensemble of grains will then be found at a wide range of temperatures, resulting in a broad probability distribution. With this approach, if the absorption timescale is significantly shorter than the radiative timescale at a given grain size, the stochastic heating regime is terminated for that and all larger sizes.

A third characteristic that may differ between implementations is the discretization of the grain size distribution, as discussed in Sect. 4.4.1.

4.5   Reference code and participating codes

4.5.1   DustEM

The DustEM code is described in Compiègne et al. (2011). DustEM is a stand-alone code (i.e., it is not a RT code) that calculates the emission and extinction of dust grains given their size distribution and their optical and thermal properties. It determines the grain temperature distribution $P(T)$ using the formalism of Desert et al. (1986), and it then computes the dust SED and associated extinction for given dust types and size distributions. To correctly describe the dust emission at long wavelengths the original algorithm has been adapted to better cover the low temperature region. Using an adaptive temperature grid, DustEM iteratively solves the integral equation Eq. 25 from Desert et al. (1986) in the approximation where the grain cooling is fully continuous. The temperature distribution calculation is performed for all grain populations and sizes including those for which the thermal equilibrium approximation would apply.

We produced the reference solutions in this work with the public version of DustEM (v3.8, dated Spring 2010). To this end, we converted the dust properties defined in Sect. 4.2 and the input radiation fields defined in Sect. 4.3 into the data format expected by DustEM. We adjusted the values of the physical constants in the DustEM code, raised the maximum number of grain size and temperature bins to accommodate our input data, and fixed a minor problem in the routine that imports the grain size distribution. More important, to obtain accurate reference solutions, we substantially increased the number of temperature bins and the number of numerical iterations (see Sect. 4.6.2). Other than this, the DustEM code was used without modifications.

Our use of DustEM for producing reference solutions should not be understood to imply that it necessarily produces the physically correct results. The DustEM implementation relies on the continuous cooling assumption just like the RT codes participating in this benchmark. However, since it is not a RT code by itself, DustEM's focus is solely on calculating dust properties and emission, and it has a neutral status in the context of this benchmark.

The following sections describe each of the participating codes with a focus on the specific heuristics employed for calculating the results presented in this work. Table 4.3 offers a (very) concise overview of this information.

4.5.2   SKIRT

SKIRT closely follows the method presented in Sect. 4.4 to calculate dust emission. The computation time for the emissivity of a stochastically heated dust mixture exposed to a particular radiation field scales roughly with BN^2, where B is the total number of size bins in the dust mixture (for all grain types combined), and N is the number of temperature bins used in the calculation.

For the results shown in this work, we use 15 size bins b for each grain type k, distributed logarithmically over the complete size range, so that $B = 45$. The absorption efficiencies loaded from the tables described in Sect. 4.2.4 are interpolated logarithmically as needed to perform the integrations over grain size presented in Sect. 4.4.1 on a logarithmic grain size grid with 201 points within each bin.

Because of the N^2 dependency in the computation time, the choice of the temperature grid is fairly crucial. By varying N in a number of experiments, it can be easily shown that, for most input fields, the temperature probability distribution $P(T)$ for very small grains can be calculated accurately on a rather coarse grid. Larger grains require a finer grid because $P(T)$ is narrower and has steep flanks. As discussed in Sect. 4.4.2, for sufficiently large grains $P(T)$ approaches a delta function, so that the procedure described in Sect. 4.4.3 requires an exceedingly refined grid with $N > 5000$ to produce accurate results, which becomes computationally prohibitive. Thus, in the interest of both speed and accuracy, we need to switch from transient to equilibrium calculations for grains that can be considered to be in LTE.

A RT simulation typically calculates the emissivity of a certain (fixed) dust mixture for a large number of radiation fields. This computation can be accelerated substantially by precalculating and storing the elements of the matrix $A_{f,i}$ defined in Sect. 4.4.3, insofar as they don't depend on the radiation field. The memory requirements scale with BN^2, just like the computation time. More importantly, precalculating these values requires a predefined temperature grid that remains fixed for all emissivity calculations. However, performing all calculations on the finest grid would be very inefficient.

The SKIRT implementation handles these conflicting requirements as follows. We predefine three separate temperature grids (A, B, and C) that can be used for any of the emissivity calculations, and we precalculate and store all radiation-field-independent values on each of these three grids for each size bin. The temperature range is the same for all grids; it is usually set to $[2\,\mathrm{K}, 3000\,\mathrm{K}]$, but if needed the upper limit is decreased to the highest temperature for which enthalpy data is available in the dust properties.

Grid A has only 20 bins. The widths are distributed according to a power law, providing a lot more resolution at low temperatures where most of the action is. The ratio of the largest bin (at high temperatures) over the smallest bin (at low temperatures) is set to 500. This grid is used to find a quick estimate of the range in which the temperature probability distribution is nonzero (or rather, larger than a very small fraction of the maximum probability).

All bins in grid B are 4 K wide. This medium-resolution grid is used to calculate the temperature probability distribution for dust grains with a very wide temperature range (i.e., very small grains and essentially all PAH grains).

Grid C has an average bin width of 2 K, with a power-law ratio of 3 between the largest and smallest bins. This provides a fine resolution of less than 1 K at low temperatures, while still offering decent resolution at high temperatures. This high-resolution grid

is used to calculate the temperature probability distribution for dust grains with a rather narrow temperature range (but not so narrow that they would be considered to be in equilibrium).

SKIRT implements the following heuristic to select the appropriate calculation for each representative dust grain:

1. calculate the equilibrium temperature T_{eq} for this grain;

2. use grid A to estimate the temperature range $\Delta T = T_{max} - T_{min}$ in which $P(T)/P_{max} > 10^{-20}$;

3. if $\Delta T < 10\,\text{K}$ or if $T_{max} < T_{eq}$ then calculate the emissivity assuming equilibrium at temperature T_{eq} and exit;

4. calculate $P(T)$ using grid B (if $\Delta T > 200\,\text{K}$) or grid C (if $\Delta T < 200\,\text{K}$);

5. update the temperature range $\Delta T = T_{max} - T_{min}$ in which $P(T)/P_{max} > 10^{-20}$ based on the new calculation;

6. if $\Delta T < 10\,\text{K}$ or if $T_{max} < T_{eq}$ then calculate the emissivity assuming equilibrium at temperature T_{eq} and exit;

7. calculate the emissivity using $P(T)$ from step 4 over the range $[T_{min}, T_{max}]$ determined in step 5.

The conditions in steps 3 and 6 are designed to avoid numerical instabilities when the temperature probability distribution approaches a delta function ($\Delta T < 10\,\text{K}$), and to capture situations where the result is clearly inaccurate since the equilibrium temperature lies outside the calculated temperature range ($T_{max} < T_{eq}$).

Further experiments with the SKIRT implementation show that for $\Delta T \gtrsim 25\,\text{K}$, the result is highly sensitive to the exact value of ΔT; for lower values the result converges to a stable solution. For values down to $\Delta T \approx 10\,\text{K}$, the result is numerically stable in the sense that performing the calculation on higher-resolution grids essentially produces the same solution.

4.5.3 DIRTY

DIRTY (Gordon et al. 2001; Misselt et al. 2001) is a Monte Carlo RT code designed to study dust and its effect on radiation in arbitrary astrophysical systems. DIRTY is a fully 3D code allowing for the specification of arbitrary density distributions of both dust and radiation sources. It implements an adaptive mesh allowing for the efficient allocation of computing resources amongst regions in the model space depending on the physical characteristics of the system. Dust absorption, temperature distribution, and emission are handled self-consistently and multiple, anisotropic scattering is implemented. The dust heating implementation supports both equilibrium and stochastic processes based on the local radiation field and dust properties at each grid in the model space.

Like other codes presented here, DIRTY follows the approach presented in Sect. 4.4. Internal to DIRTY, the dust grain size distribution is not further discretized beyond the input discretization of the model. For the benchmark dust model, we computed the heating and emission for all sizes in the input mesh (28 for PAH, 121 for graphitic and silicate components; see Sect. 4.2.4).

The heuristics employed by DIRTY in calculating the dust emission from each grain size of each component exposed to the local radiation field at a point in the model space are as follows:

1. The equilibrium temperature, T_{eq}, cooling timescale, τ_{rad}, and heating time scale, τ_{abs} are computed according to section 7 of Draine and Li (2001);

2. If the time scales computed in step 1 satisfy the inequality $\tau_{abs} < \tau_{rad}$, the grain is considered to be in equilibrium with the local field; it is assigned a temperature of T_{eq} and its emission is calculated following Eq. 4.18 with $P = \delta(T - T_{eq})$. All grains of the same composition that are larger than the size for which the inequality is first satisfied are by default treated as being in equilibrium with the local field.

3. For those grains found to be in stochastic regime in step 2, $P(T)$ is computed on an initial coarse temperature grid. The coarse grid is defined on 50 points linearly spaced in the interval $[0.3, 3.0]\, T_{eq}$.

4. Depending on the results of step 3, the algorithm proceeds in one of two directions:

   (a) If $P(T)$ is not below a specified tolerance at the endpoints of the temperature grid ($P(T_{min}), P(T_{max}) < 10^{-15}$), the temperature limits on the grid are expanded by 50% and we return to step 3 with the new, expanded temperature grid.

   (b) If $P(T)$ is below the tolerance at the endpoints of the temperature grid, compute the energy emitted by the stochastically heated grain using Eq. 4.18. If the emitted energy is within 1% of the energy absorbed from the radiation field by the grain, consider the calculation converged and return the emitted energy spectrum. Otherwise, proceed to step 5.

5. The grid is now refined by a series of moves that refine the temperature limits and increase the number of temperature samples if necessary.

   (a) Remove all points on the ends of the temperature range that have low probabilities ($P(T) < 10^{-15}$).

   (b) Recompute the probabilities on the smaller grid with the same number of samples.

   (c) Compute the emission and emitted energy. If the emitted energy matches the absorbed energy to 1%, consider the calculation converged and return the emitted energy spectrum. Otherwise;

   (d) Keeping the same temperature limits, increase the number of samples by 50% and recompute the probabilities. Trim temperature endpoints with $P(T) <$

10^{-15} and return to step 5b. If this step would result in the number of bins exceeding a predefined maximum (1000), record the failure of the stochastic heating algorithm and proceed to the next grid in the model space.

In practice, the failure described in step 5d occurs in model bins for which the local field has not been well defined, generally because there are very few photons interacting in that cell, either through a poor definition of the adaptive mesh or from insufficient photons being run in the Monte Carlo simulation of the radiation. These cells are generally unimportant in the overall energy budget of the model and can be masked in post-processing. Such cells are rare in most model runs; their number and distribution in the model space can be used as a metric of the overall quality of the simulation.

The approach in step 2 results in all PAH sizes being treated in the stochastic regime in the radiation fields explored in this benchmark. Silicate and graphitic grains generally achieve equilibrium between sizes of 0.006 and 0.020 μm.

4.5.4 TRADING

TRADING (Bianchi 2008) is a 3D Monte Carlo dust continuum RT code with characteristics similar to those of SKIRT and DIRTY. Originally designed to study the effects of clumping in the disks of spiral galaxies, it uses a binary-tree adaptive grid (octree) for the dust distribution.

TRADING computes stochastic heating following the method described in Sect. 4.4. A single temperature grid is used to precompute the field-independent terms of the matrix elements (Eq. 4.27 and most factors in Eq. 4.24). Since the RT models of clumpy galactic disks in Bianchi (2008) need a few million dust cells, and each cell requires the calculation of dust heating for the full grain size distribution, Bianchi (2008) uses a limited temperature grid of 80 logarithmically spaced bins between 2.7 and 2000 K. As mentioned in Sect. 4.4.3, faster thermal equilibrium calculations were performed for grains larger than $a_{\text{trans}} = 0.01\,\mu$m. The original setup resulted in SEDs that were estimated to be within 10% of full solutions for wavelengths up to 1000 μm.

For this benchmark we left the number of temperature bins at 80, but we extended the temperature range to 3000 K. While the previous choice of a_{trans} produces accurate results for the typical interstellar radiation fields encountered in spiral galaxies (with $U \gtrsim 0.1$, see Aniano et al. (2012); Hunt et al. (2015)), we adopted here $a_{\text{trans}} = 0.05\,\mu$m in order to improve the solution for fields with lower intensities.

The dust grain size distribution was discretized using a grid with 20 size bins for graphite, 20 for silicates, and 8 for PAHs, logarithmically spaced over their size range. The choice allows for similar bin widths for all materials. Optical properties were derived by interpolating logarithmically over those of the full size table. The grain size grid used for this work is similar to the grid adopted in Bianchi (2008) though for a different dust model.

For $a_{\text{trans}} = 0.01\,\mu$m, only about half of the size bins pass through the stochastic

heating calculation. For $a_{\text{trans}} = 0.05\,\mu\text{m}$, this goes up to 75%: all of the PAH bins and 14 of the 20 graphite and silicate bins. If a grain with $a < a_{\text{trans}}$ attains the condition of thermal equilibrium, the adopted temperature grid might fail to compute the resulting narrow probability function. Thus, in addition to the grain size cut-off, we chose to assume thermal equilibrium whenever the range of the computed temperature distribution defined by $P(T) > 10^{-15}$ is smaller than the equilibrium temperature.

4.5.5   CRT

CRT (Juvela and Padoan 2003; Juvela 2005; Lunttila and Juvela 2012) is a Monte Carlo dust continuum RT program. It solves the RT equation self-consistently with a full treatment of scattering, absorption and emission of radiation in 1D, 2D, and 3D geometries. The program allows using an external component, for instance DustEM, for calculating the dust emission spectrum for a given input radiation field. This feature has been used in Ysard et al. (2011), for example, to study the microwave emission from spinning dust grains. However, CRT itself also has optimized routines for fast dust emission calculations, including the treatment of stochastically heated grains. Although CRT can calculate dust emission with fully discrete cooling, the results presented in this work use the continuous cooling approximation, and the following discussion focuses on the algorithms for continuous cooling computations.

The basic algorithms employed by CRT follow the outline presented in Sect. 4.4. To allow the use of precomputation to speed up the construction of the transition matrix A, the temperature grid is not defined fully dynamically according to the input field. Instead, each dust type and size uses one of several predefined grids, for which the precomputations are done at the beginning of calculation. The predefined grids are linear in T, and their upper and lower limits are chosen to allow a good representation of the grain temperature distribution in the types of radiation fields that are found in the model. In particular, the grids should be built for reference fields that approximately span the range of radiation energy densities expected to be found in the model. If the hardness of the radiation field varies significantly, it may also be useful to include grids for different spectra with the same energy density. For the calculations presented in this work we used six temperature grids that were built for scaled Mathis ISRF with $U = 10^{-4}, 10^{-1}, 1, 10, 10^3$, and 10^6. For large grains and strong radiation fields, the predefined grid has a special entry that triggers equilibrium calculation, otherwise full stochastic calculations are used. The number of temperature bins for stochastic calculation can be set by the user, and in this work we use 128 bins.

To select the temperature grid that is used for calculating the emission for a grain size and type in a given input field, we calculate the absorption time scale τ_{abs} and the mean energy of the absorbed photon $\langle E_{\text{abs}} \rangle$. We choose the predefined grid that has been built for the reference field most like the current input field. In the calculations presented in this work the selected grid is simply the one for whose reference field the mean absorbed power per grain $\mathcal{P}_{\text{abs}} = \langle E_{\text{abs}} \rangle / \tau_{\text{abs}}$ is closest to

the values calculated for the input field. Although the selected temperature grid is not necessarily optimal, it is good enough to allow accurate results using a modest number of energy bins.

The computation of emission from stochastically heated grains in CRT differs slightly from the description given in Sect. 4.4.3. Instead of using Eq. 4.24 for calculating the upward transition rates, we apply Eqs. 15–25 and 28 from Draine and Li (2001), which include corrections for the finite size of an enthalpy bin. Similarly, instead of Eq. 4.27, we use Eq. 41 from Draine and Li (2001) for calculating the cooling part of the transition matrix. Including these finite bin size corrections allows for using a lower resolution grid, which substantially benefits the computation time.

The cooling matrix elements ($f < i$) are independent of the radiation field and they are precomputed. The heating elements ($f > i$) depend on the radiation field and must be calculated separately for each input field. Moreover, when using the equations from Draine and Li (2001), instead of using the radiation field strength at a single wavelength as in Eq. 4.24, we must integrate numerically over the wavelength grid. We use precomputed integration weights $w_{f,i}$ corresponding to each grid point of the wavelength grid. The integral in Eq. 15 in Draine and Li (2001) can then be evaluated as $A_{f,i} = \sum_{k=1}^{N_\lambda} w_{f,i}(k) J(\lambda_k)$. If there are a lot of points in the wavelength grid, calculating the full sum is slow. However, for given energy bins i and f, radiation within only a narrow range of wavelengths can induce transitions $i \to f$. Therefore, $w_{f,i}(k)$ is non-zero only for a few k and only non-zero integration weights are stored and used in the summation.

Discretization of the particle size distribution can be defined by the user. In this work we employ the same discretization as SKIRT: 15 logarithmic bins for each of the three dust components. Dust properties for each size bin are calculated according to Eqs. 4.15-4.17 using numerical integration with 256 grid points.

4.5.6 MCFOST

MCFOST (Pinte et al. 2006, 2009) is a 3D continuum and line RT code. It relies on the Monte Carlo method to compute the local specific intensities and related quantities (e.g. temperature, molecular levels) and computes observables via a ray-tracing method. The emerging fluxes are calculated by formally integrating the RT method along rays using the specific intensities and source functions computed during the Monte Carlo run.

MCFOST computes the stochastic heating of small dust grains following the method presented in section 4.4, with a few refinements to ensure numerical stability and speed.

We first compute the time between two successive absorptions of a photon and compare it to the cooling time of the grain, following the method described in Draine and Li (2001). For dust grains where the time between two absorptions is shorted than the cooling time, only the equilibrium temperature is calculated.

For those grains that have a shorter cooling time, we compute the full temperature probability distribution using a fixed temperature grid with 300 points logarithmically distributed between 1 and 3000 K. The cooling terms of the transition matrix (Eq. 4.27), which are independent of the radiation field, are precomputed, as are most of the heating terms factors (Eq. 4.24). We estimate the specific intensity $J(\lambda_{fi})$ for each term in Eq. 4.24 by interpolating the radiation field computed by the Monte Carlo run. Since the interpolation coefficients are identical for every cell in the model, they are also precomputed to speed the calculations up.

For dust in radiative equilibrium, MCFOST solves the problem of self-consistent dust heating and re-emission using the immediate re-emission concept of Bjorkman and Wood (2001). This method eliminates the need for iteration and ensures a perfect conservation of energy. For non-equilibrium dust grains, however, this procedure is prohibitive because it requires a temperature calculation at each absorption or re-emission. Instead, we use the classical iterative scheme where we store the energy absorbed by the dust, compute the temperature probability distribution in all cells once all the packets have been propagated, and re-emit the absorbed energy via new packets according to the new temperature probability. The procedure is iterated until a desired convergence on the temperature or energy is reached. Because only the fraction of radiation that is absorbed by the non-equilibrium grains needs to be re-emitted in an iterative way, convergence is usually reached after only a few iterations. In practice, when a packet is absorbed inside a cell, its energy is split: the fraction absorbed by dust grains in equilibrium is immediately re-emitted, while the fraction absorbed by non-equilibrium grains is stored to be re-emitted during the next iteration. For the benchmark presented in this work, the input radiation field is fixed so no iteration is required.

In the presented calculations, the grain size distributions were discretized using 63 logarithmically spaced grain sizes for silicates, 62 for graphite and 24 for PAHs.

4.5.7   DART-Ray

DART-Ray is a ray-tracing 3D dust RT code that implements the RT algorithm described in Natale et al. (2014). It can be used to derive radiation-field energy-density distributions and outgoing radiation surface brightness maps for arbitrary 3D distributions of dust mass and stellar emission. It includes treatment of both absorption and anisotropic scattering. For the dust emission calculations, DART-Ray uses the prescription initially incorporated into the 2D RT model of Popescu et al. (2000) and later updated in Popescu et al. (2011). However, unlike the 2D models where the stochastically heated dust emission could be explicitly computed for each individual position, the calculations for the 3D models, often containing millions of cells, can be accelerated by using an adaptive SED library approach (see Natale et al. 2014b, in press).

For a given radiation field intensity spectrum found for a particular model cell, the stochastically heated dust emission is derived following Voit (1991). The method used

to determine the probability distribution $P(T)$ combines the numerical integration of Guhathakurta and Draine (1989) with a step-wise analytical solution. The algorithm provides accurate and swift results on a relatively coarse grid. This is particularly useful for larger grain sizes where the probability distribution $P(T)$ converges to a narrow distribution around the equilibrium temperature. In the case of the first-order integration of Guhathakurta and Draine (1989), the width of the energy bins needs to be considerably smaller than the mean deposited energy to preserve energy balance. An increasing number of energy bins would be required to avoid energy losses in the heating process, which would make the calculation inefficient (see Fischera 2000, and Sect. 4.4.3).

The absorption of CMB photons is assumed to provide a continuous heating source. It is taken into consideration by subtracting the heating rate related to the CMB from the cooling rate (Eq. 4.27), which limits the temperatures to values not lower than the CMB temperature. As further discussed in Sects. 4.6.3 and 4.6.6, the cooling assumed in the dust model is only valid as long as the emitted photon energy of the modified black body is low relative to the enthalpy of the grain. This assumption no longer applies at very low temperatures of small dust grains or PAH molecules (Fischera 2000). While negligible for most cases, we find that the CMB becomes a considerable heating source for the very low radiation field strength. For $U = 10^{-4}$, the heating of silicate grains by CMB photons is for all sizes larger than 10%. For graphites the contribution is considerably lower with only a few percentage points and still negligible for PAH molecules.

The temperature distribution $P(T)$ for each grain size of a given composition is obtained consecutively by starting with the largest grain size and utilizing the basic characteristic that the distributions broaden with decreasing grain size. The probability distribution is only considered for all grains below a certain size where the stochastic heating leads to a considerable distribution of the dust temperatures. Above this critical grain size the grains are assumed to radiate at the equilibrium temperature T_{eq}. To estimate the transition from equilibrium to non-equilibrium, we apply the Gaussian approximation in the limit of large grains as derived by Voit (1991). We consider the grains to be stochastically heated if $2\sigma_T/T_{\mathrm{eq}} > 0.05$. For $0.05 < 2\sigma_T/T_{\mathrm{eq}} < 0.1$, we apply the Gaussian approximation, and the full $P(T)$ distributions are derived for $2\sigma_T/T_{\mathrm{eq}} > 0.1$.

The temperature distributions of stochastically heated dust grains are derived for dynamically determined temperature intervals $[T_{\mathrm{min}}, T_{\mathrm{max}}]$. For the results presented in this work, we subdivided the temperature interval into 200 temperature bins equally spaced in $\log T$. The interval boundaries were determined iteratively by increasing T_{min} or decreasing T_{max} by 30% until the probabilities $P(T)$ for the lowest and highest temperature bin were lower than 10^{-20}, ensuring that the emission at higher and lower temperatures can be neglected. To accelerate the iterative process, we used the derived interval of the previous larger dust grain as the initial estimate of the temperature interval for each dust grain. For the first grain for which the temperature distribution was derived, we used a width based on the Gaussian approximation in

the limit of large dust grains as initial estimate of the temperature interval.

For the results presented in this work, we derived the dust emission for each grain species at each grain size of the tabulated size distribution described in Sect. 4.2.4. To calculate the total emission, we then integrated over the size distribution and summed the contributions from different grain species. By comparing the total dust emission with the total absorbed energy, we ascertained that the energy balance for every grain species is fulfilled with an accuracy better than a few percentage points. The largest discrepancies (3-4%) are found for the most extremely scaled Mathis fields considered in the benchmark ($U = 10^{-4}$ and $U = 10^6$).

4.6	Results and discussion

4.6.1   Data files

The data files representing the benchmark results can be downloaded from the web site indicated in footnote 13. For each participating code, and for each input radiation field, the calculated solution is stored in a separate text file with columns specifying the wavelength λ (in μm); the mean intensity J_λ of the input field (in $W\,m^{-3}\,sr^{-1}$); and the silicate, graphite and PAH emissivities $\lambda\,\varepsilon_\lambda$ (in $W\,sr^{-1}\,H^{-1}$), in that order. The file naming scheme and the precise file format are described on the web site.

4.6.2   Reference solutions

As mentioned in Sect. 4.5.1, we compare the results below from the codes participating in this benchmark against reference solutions generated with DustEM. Figure 4.1 shows these solutions for a selection of the input fields defined in Sect. 4.3.

To ensure proper accuracy, we increased the number of temperature bins and the number of iterations in the DustEM integral equation solver until the calculated emission converged to a stable solution; see Fig. 4.2. Specifically, the number of temperature bins was raised from 200 to 3500, and the number of iterations from 80 to 250. These changes dramatically increased the computation time, however this is acceptable for calculating a reference solution.

As an extra sanity check, we verified that the emissivities calculated by DustEM (using the same number of temperature bins and iterations as for the reference solutions) indeed converge to the corresponding equilibrium emissivities. Figure 4.3 shows this comparison for 0.05 μm grains exposed to radiation fields ranging from extremely weak (left) to strong (right). For a strong field, where we expect the grains to be in equilibrium, the DustEM solutions indeed match the equilibrium emissivities. For a weak field, the solutions differ since the grains are no longer completely in equilibrium. This shows the importance of performing the full stochastic calculation in the presence of extremely weak fields, even for grain sizes up to 0.05 μm. Codes transitioning from stochastic to equilibrium calculation at a fixed grain size should
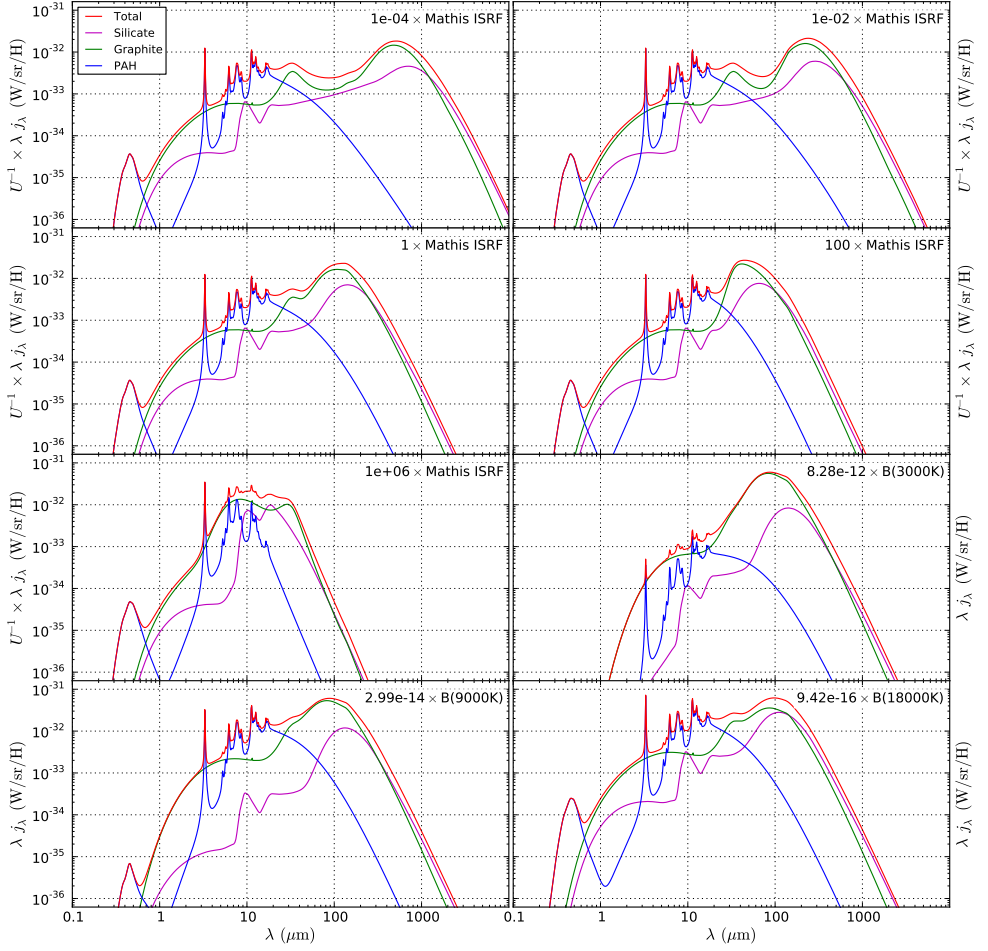
Figure 4.1: Reference solutions generated with the public version of DustEM (see Sect. 4.5.1) using 3500 temperature bins and 250 iterations in the integral equation solver. The panels show the calculated dust emissivity for a selection of the input fields defined in Sect. 4.3. In each panel, the red curve represents the total emissivity and the other curves represent the portion of the emissivity for each grain type, silicate (magenta), graphite (green) and PAHs (blue). For the scaled Mathis input fields, the emissivity is divided by the input field strength U to allow identical axis ranges for all plots.

thus set a sufficiently high value of a_{trans} (see Sects. 4.4.3 and 4.5.4).

4.6.3 Benchmark solutions

Figure 4.4 compares the total emissivities calculated by each of the codes participating in this benchmark to the corresponding reference solutions for a selection of the input fields defined in Sect. 4.3. Subsequent figures zoom in on the emissivities for each dust component separately: silicate (Fig. 4.5), graphite (Fig. 4.6), and PAH (Fig. 4.7) grains.
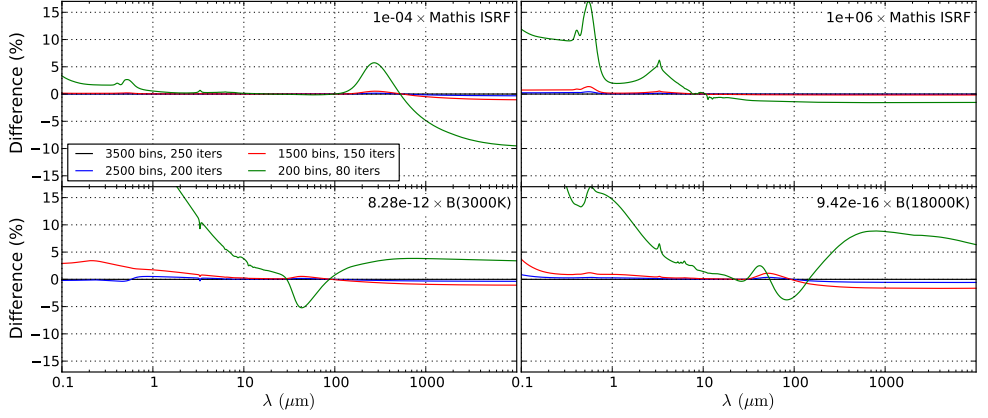
Figure 4.2: Comparison of DustEM solutions for the most extreme input fields defined in Sect. 4.3, calculated with a varying number of temperature bins and iterations in the DustEM integral equation solver. The solutions employed as a reference for our benchmark are calculated with 3500 temperature bins and 250 iterations; these solutions are represented in this figure by the zero lines. The solutions calculated with the standard DustEM values of 200 temperature bins and 80 iterations are represented by the green curve. For these extreme fields, the standard solution deviates by up to 20% (and even more for wavelengths shorter than 1 μm). The solutions using 2500 temperature bins and 200 iterations (the blue curve) differ by less than 1% from the reference solution, indicating numerical convergence at these parameter values. The contribution of each grain type separately has a similar convergence behavior (not shown).
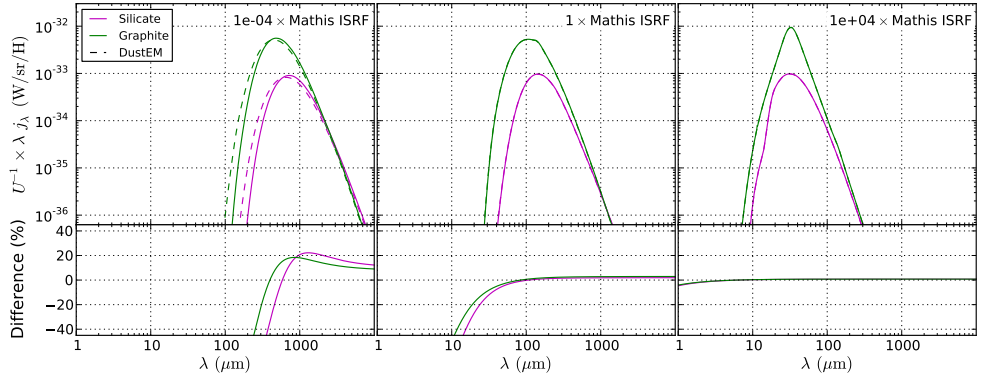
Figure 4.3: Comparison of the emissivities calculated by DustEM (using 3500 temperature bins and 250 iterations) for single-size, near-LTE grain populations to the corresponding equilibrium emissivities. The panels show the comparison for input fields ranging from extremely weak (left) to strong (right). The emissivity is divided by the input field strength U to allow identical axis ranges for all plots. We used a dust mixture consisting of 0.05 μm silicate (magenta) or graphite (green) grains with a total dust mass per hydrogen atom of 10^{-30} kg/H for each grain type. The solid curves represent the emissivities calculated by one of our codes (SKIRT) under the assumption of LTE. The dashed curves represent the solutions calculated by DustEM, without any LTE assumptions. The lower panels show the deviation of the equilibrium solutions from the corresponding full solutions. In a strong field, where we expect the grains to be in equilibrium, the solutions are indeed virtually identical. In a weaker field, the solutions differ since the grains are no longer completely in equilibrium.
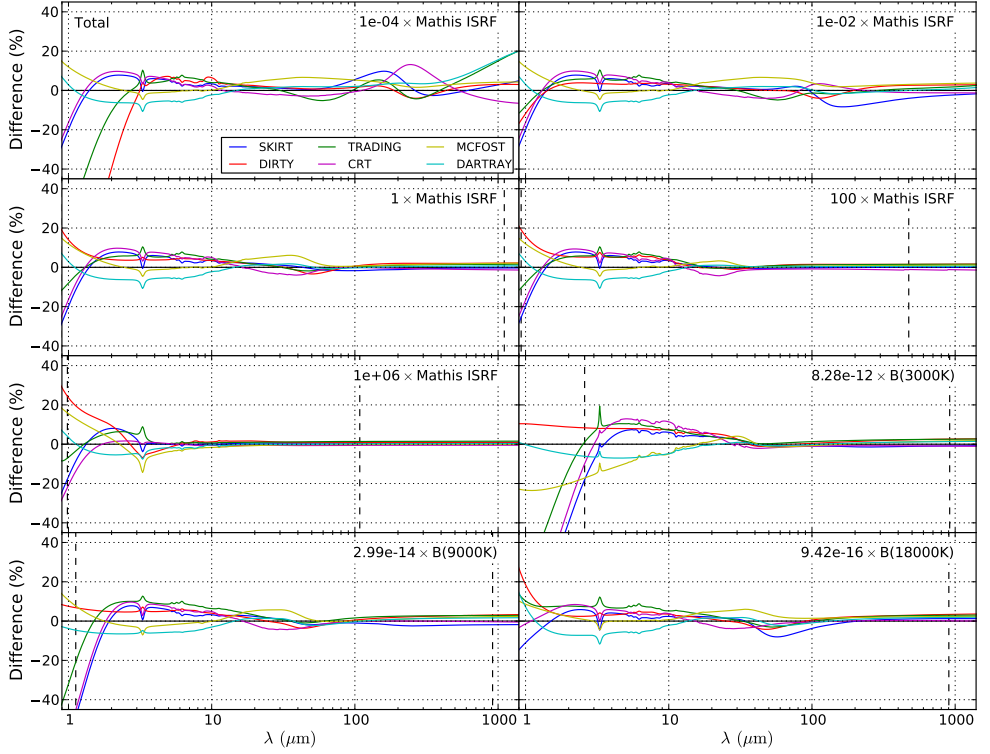
Figure 4.4: Relative differences between the total emissivities calculated by each of the codes participating in this benchmark and the corresponding reference solutions. The panels show the results for a selection of the input fields defined in Sect. 4.3. In each panel, the reference solution is represented by the zero line. Positive percentages indicate results above the reference solution. The vertical dashed lines indicate where the reference solution becomes three orders of magnitude smaller than its peak value.

In these figures we limited the displayed wavelength range to $1\,\mu m \leq \lambda \leq 1000\,\mu m$. Outside of this range, other sources or processes usually dominate the radiation emanating from astrophysical objects, so it is not relevant to evaluate the results of a dust emissivity calculation in that spectral range. Moreover, some of the assumptions underlying the computations are no longer valid, rendering the results physically meaningless.

First considering the shorter wavelength range, a black body with peak emission at $\lambda = 1\,\mu m$ has a temperature of $T \approx 2900\,K$. The sublimation temperature of a dust grain is estimated at $1200\,K$ for silicates and at $2100\,K$ for graphites (Kobayashi et al. 2009). Evaporation rates rise roughly exponentially with increasing grain temperature (Guhathakurta and Draine 1989; Kobayashi et al. 2009), i.e. with decreasing emission wavelength. It is clear that a relevant portion of the dust that would emit at $\lambda \lesssim 1\,\mu m$ is destroyed by evaporation, so that the grain size distribution in our dust model is no longer valid under these conditions. Consequently, the calculation would substantially overestimate the resulting dust emission.

Figure 4.5: Relative differences between the emissivities of the silicate component calculated by each of the codes participating in this benchmark and the corresponding reference solutions. The panels show the results for a selection of the input fields defined in Sect. 4.3. In each panel, the reference solution is represented by the zero line. Positive percentages indicate results above the reference solution. The vertical dashed lines indicate where the reference solution becomes three orders of magnitude smaller than its peak value.

We now consider the longer wavelength range. For a sufficiently strong input field, say $U \gtrsim 1$, we can expect the calculated emissivity results to be correct because most of the grains emitting at wavelengths longer than 1000 μm are in LTE. However, the emissivity peaks at much shorter wavelengths, so that the level at 1000 μm is already several orders of magnitude below the peak level (see all panels in Fig. 4.1 except for the first two). The situation is different for extremely weak input fields approaching the level of the cosmic microwave background (CMB). A black body with peak emission at $\lambda = 1000\,\mu$m has a temperature of $T \approx 2.9\,$K, just above the temperature of the CMB. For small dust grains at such low energies, the continuous cooling assumption no longer holds,[15] meaning that the method presented in Sect. 4.4.3 does not necessarily yield correct results. In conclusion, the dust emission at wavelengths $\lambda \gtrsim 1000\,\mu$m is either calculated assuming equilibrium conditions (which renders comparison uninteresting) or calculated improperly (which renders comparison meaningless).

---

[15] With the properties of our dust model, the internal energy of a 0.007 μm graphite grain at 2.9 K is insufficient to emit a 1000 μm photon.
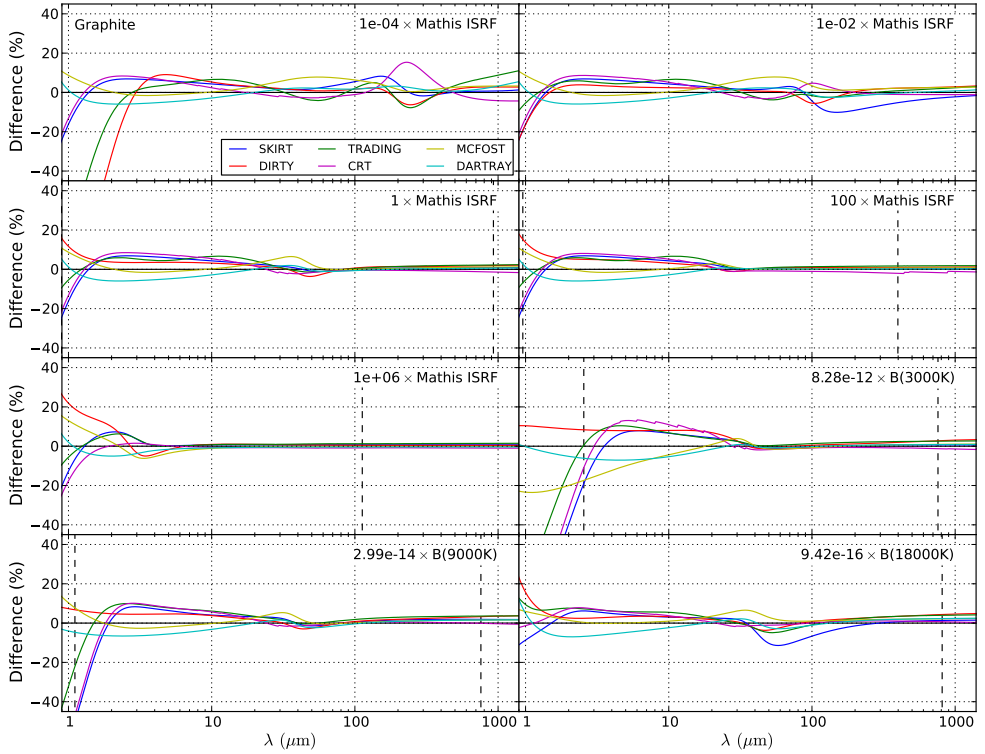
Figure 4.6: Relative differences between the emissivities of the graphite component calculated by each of the codes participating in this benchmark and the corresponding reference solutions. The panels show the results for a selection of the input fields defined in Sect. 4.3. In each panel, the reference solution is represented by the zero line. Positive percentages indicate results above the reference solution. The vertical dashed lines indicate where the reference solution becomes three orders of magnitude smaller than its peak value.

4.6.4   Evaluation of benchmark results

In the wavelength range $3\,\mu\text{m} \leq \lambda \leq 1000\,\mu\text{m}$, all participating codes reproduce the total dust emissivity within 20% of the reference solution for all input fields used in this benchmark (see Fig. 4.4). Excluding the weakest ($U \lesssim 10^{-4}$) and the softest ($T \lesssim 3000\,\text{K}$) fields, the correspondence in the same wavelength range is within 10%. The larger relative deviations at wavelengths shorter than $3\,\mu\text{m}$ are caused in part by the much lower absolute emissivity values in that range (two to three orders of magnitude below peak values; see Figs. 4.1 and 4.4).

The emissivities calculated for the silicate and graphite components (Figs. 4.5 and 4.6) show a similar pattern, although the deviations in the individual components are sometimes slightly larger. The emissivities calculated for the PAH component (Fig. 4.7) show the largest deviations. If we restrict the analysis to the wavelength range in which the emissivity of the reference solution is within three orders of magnitude of its peak value, the correspondence is still within 40% for all codes and for all input fields, and often a lot better.
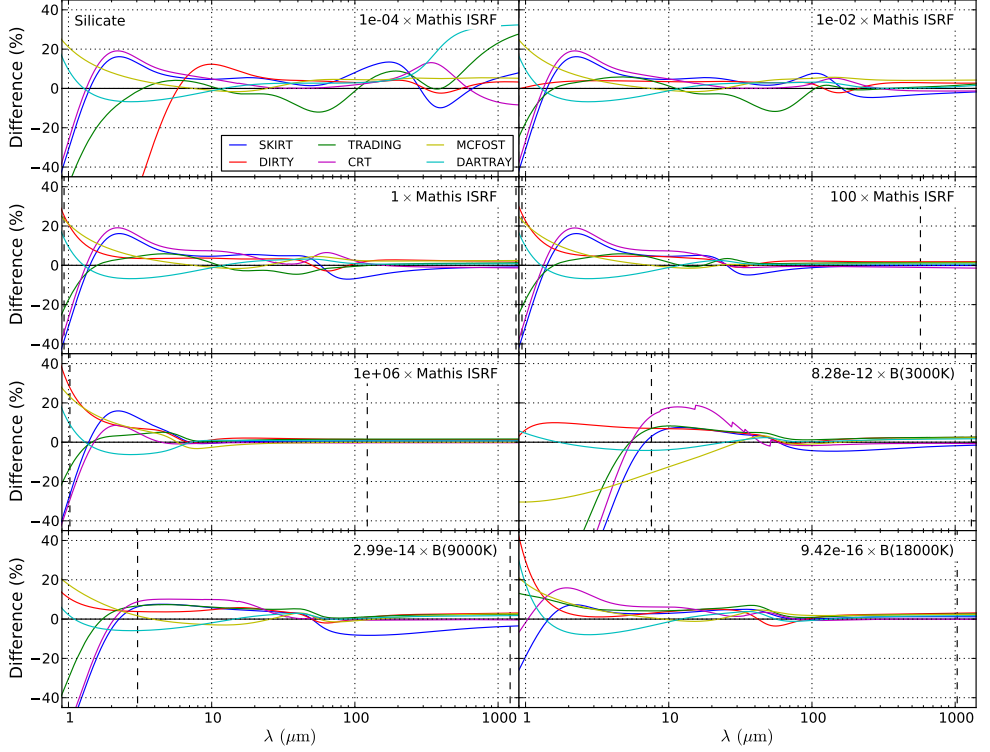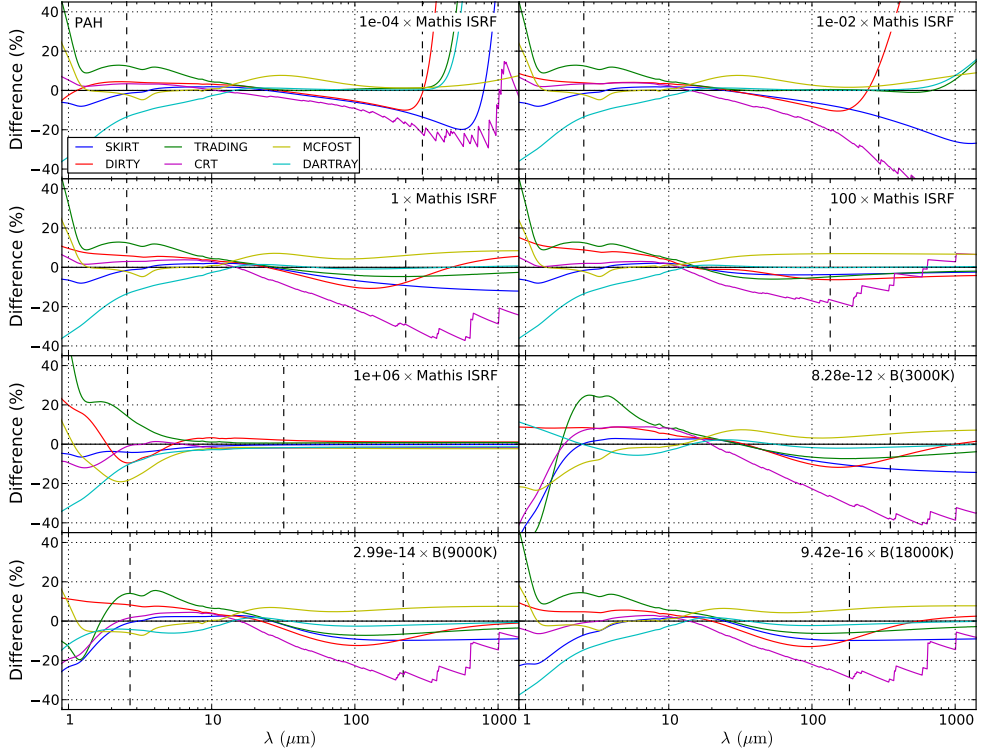
Figure 4.7: Relative differences between the emissivities of the PAH component calculated by each of the codes participating in this benchmark and the corresponding reference solutions. The panels show the results for a selection of the input fields defined in Sect. 4.3. In each panel, the reference solution is represented by the zero line. Positive percentages indicate results above the reference solution. The vertical dashed lines indicate where the reference solution becomes three orders of magnitude smaller than its peak value.

The larger discrepancies between the various codes for the PAH component can be traced to the PAH molecules generally being substantially smaller than silicate and graphite grains: see the $a^{\max}$ values in Table 4.2. As noted in Sect. 4.4.3 and further discussed in Sect. 4.6.5, smaller grains are more likely to remain in the stochastic regime, requiring complex calculations that are more sensitive to differences in discretization (choice of grids) and concrete implementation (even if the same overall method is employed), as compared to equilibrium calculations.

Specifically, the PAH emissivities calculated by CRT (magenta curves in Fig. 4.7) deviate from the other codes because, as described in Sect. 4.5.5, CRT implements the Draine and Li (2001) equations rather than Eqs. 4.24 and 4.27 for calculating the heating and cooling transition rates. This method does not allow the emission of photons with an energy higher than the enthalpy content of the emitting grain (see Eq. 56 in Draine and Li 2001). The upward jumps in the emission spectrum appear when, going toward longer wavelengths (lower photon energy), a new enthalpy bin enters the emission calculation (see also Figs. 14 and 15 in Draine and Li 2001). These
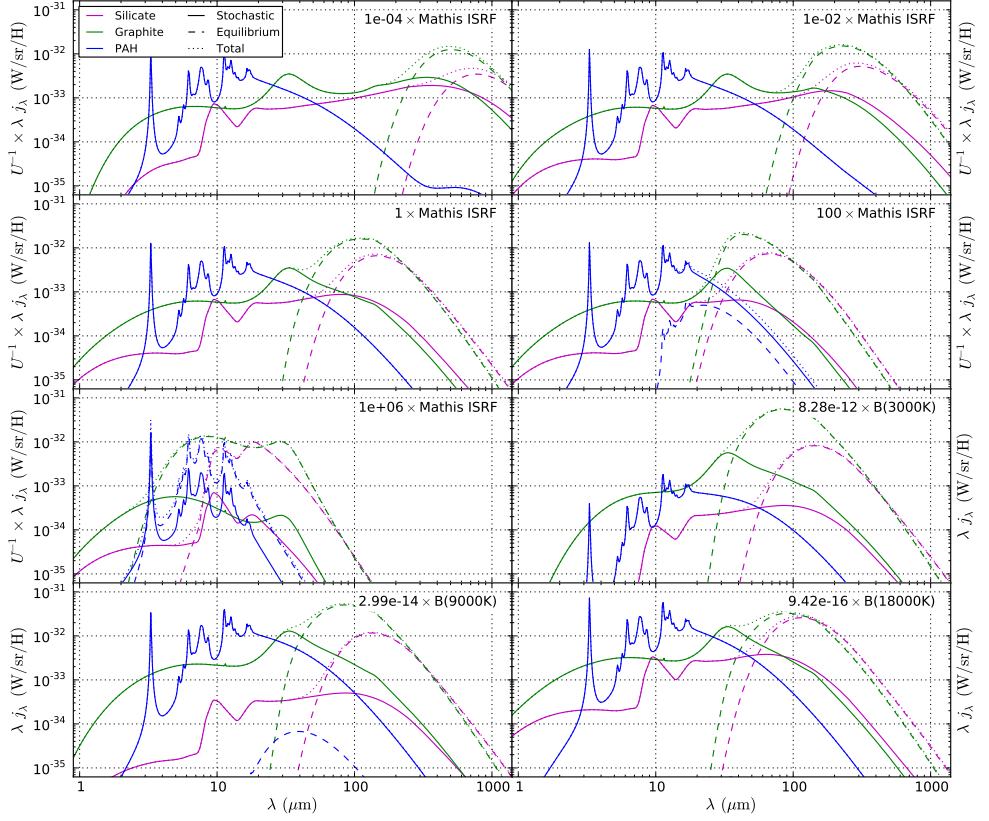
Figure 4.8: Contributions calculated in the stochastic (solid) and equilibrium (dashed) regimes to the total emissivity (dotted) for each of the grain types silicate (magenta), graphite (green) and PAHs (blue), by one of our codes (DIRTY). The panels show the results for a selection of the input fields defined in Sect. 4.3. For the scaled Mathis input fields, the emissivity is divided by the input field strength U to allow identical axis ranges for all plots. The relative contributions for each regime may vary between codes because of the differences in the schemes to transition from one regime to the other.
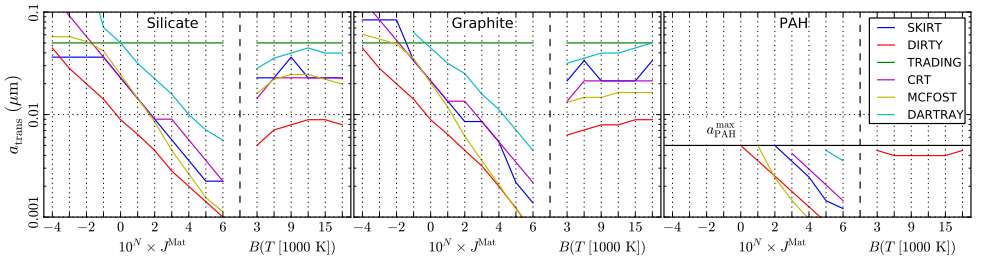
Figure 4.9: Grain size a_{trans} for which the participating codes transition from the stochastic to the equilibrium calculation regime. The panels show the smallest grain size that was considered to be in equilibrium for silicate (left), graphite (right) and PAH (right) grains, for each of the codes, across all of the input fields defined in Sect. 4.3 (within each panel, the scaled Mathis ISRF fields to the left, and the diluted black body fields to the right).

discontinuities appear only in the wavelength range where emission is largely from grains with very low enthalpy, for which the continuous cooling approximation is not valid, as described in Sect. 4.6.3.

Figure 4.8 shows the contributions to the total emissivity calculated by one of our codes (DIRTY) in the stochastic and equilibrium regimes for each of the three grain types, and for a number of input fields.[16] While the silicate and graphite components show a significant equilibrium contribution for all fields, the PAHs remain in the stochastic regime for all but the strongest fields.

With respect to our conclusion at the end of Sect. 4.6.3, it is worth noting in Fig. 4.8 that, at $\lambda = 1000\,\mu$m, the equilibrium contributions of the silicate and graphite components dominate the total stochastic contribution for all input fields. Only for the weakest fields ($U \approx 10^{-4}$) does the stochastic contribution at that wavelength become a noticeable fraction of the total emission.

Considering the contributions of the different grain types to the total spectrum (Fig. 4.1), it turns out that, for our dust model, the graphite emission dominates in most of the wavelength range for most input fields. The PAHs dominate in a very small region (their peak) and the silicates only at the longest wavelengths. Since the agreement between the different codes is significantly better for graphite and silicates than for the PAHs, our choice of dust mixture (unintentionally) benefits the global agreement between the different codes.

4.6.5 Transition to equilibrium

As introduced in Sect. 4.4.3, and further elaborated upon in the code descriptions in Sect. 4.5, each code handles the transition from the stochastic to the equilibrium calculation regime in its own way. Figure 4.9 shows the grain size a_{trans} for which the participating codes transition from the stochastic to the equilibrium calculation regime, for each grain type, and for each of the input fields defined in Sect. 4.3. While the details differ between codes, as a general trend, small grains (e.g. $a < 0.01\,\mu$m) are considered to be in equilibrium only when exposed to the strongest fields ($U > 10^2$).

Because most of the PAHs remain in the stochastic regime (Fig. 4.8 and right panel of Fig. 4.9), the transition differences are most easily seen in the results for the silicate and graphite grains (Figs. 4.5 and 4.6). As a general trend, the results converge for longer wavelengths, because the equilibrium emission dominating in that range is calculated in the same straightforward manner across all codes. The larger discrepancies are found at shorter wavelengths, where the stochastic regime dominates. Depending on the input field, the transition point shifts on the wavelength scale. Interestingly, for some fields, the discrepancies show extra "wiggles" near the transition points, most likely caused by differences in handling the transition. This is particularly evident in

---

[16] The precise form of these respective contributions varies between codes because of differences in transitioning from one regime to the other, but the general trend is similar.

the silicate emissivities for the scaled Mathis fields with strengths $U = 10^{-2}$ to 10^2 (see the top half of Fig. 4.5).

The total emissivities are influenced by these transition differences mostly in the wavelength range just shortward of the large submm emission peak, which is dominated by LTE emission from silicate and graphite grains (see Fig. 4.4). The position of this peak is determined by the temperatures of the dust grains, and thus depends on the input field. For the strongest field in our benchmark ($U = 10^6$), the broad peak even overlaps the PAH features in the 3 to 30 μm wavelength range (see the third panel in the leftmost column of Fig. 4.4).

4.6.6   Weak fields

As discussed in Sect. 4.4.3, the transition size a_{trans} above which a grain can be considered to be in LTE depends on the radiation field to which the grain is exposed. For a given grain composition, a_{trans} tends to be higher for weaker fields because photon interactions are less frequent, which keeps larger grains in the stochastic regime. Consequently, the dust emission calculations are more complex for weaker fields, especially for small grains in weaker fields. In addition to the computation time, this complexity affects the accuracy of the results, which explains why the largest discrepancies between the various solutions occur for the weakest field in the benchmark (see the top left panels in Figs. 4.4 through 4.7).

In fact, the weakest field in our benchmark ($U = 10^{-4}$) may be unrealistically weak, since its peak intensity is below the peak of the CMB, albeit in a different wavelength regime (UV versus mm wavelengths). To evaluate the effect of neglecting the CMB, we added the CMB to the $10^{-4} \times J^{\text{Mat}}$ field and used DustEM to recalculate the emissivity of our dust model exposed to this new input field. For wavelengths $\lambda \leq 100 \, \mu$m, the results are essentially identical to those shown in the top lefthand panel of Fig. 4.1. The submm peak is a notch higher and slightly shifted to longer wavelengths, causing an emissivity increase of about 35% at $\lambda = 1000 \, \mu$m. While this effect may not be negligible, it does not invalidate the benchmark test.

However, as argued in Sect. 4.6.3, our computations may no longer be physically founded for these weak fields, especially for small grains with internal energies comparable to those of the CMB photons. In conclusion, the $10^{-4} \times J^{\text{Mat}}$ input field is benchmarking the various codes properly, but the calculations may be collectively incorrect because the continuous cooling approximation is inappropriate in this regime.

4.6.7   Temperature discretization

The participating codes implement various ways to discretize the grain temperature (or equivalently, the grain enthalpy), as described in Sect. 4.5. The different schemes are mostly driven by the aim to increase performance while preserving accuracy. Here we discuss the impact of the minimum and maximum temperature values allowed

on the grid.

DustEM, which was used to calculate our reference solutions, does not impose a lower temperature limit other than the zero point. Indeed, under the continuous cooling assumption, a small dust grain does not have to be in equilibrium with the CMB, and thus there seems to be no reason the grain should not have, at any given moment in time, a temperature below 2.73 K. In other words, we need to calculate the temperature probability distribution as usual. As argued in Sects. 4.6.3 and 4.6.6, this line of reasoning breaks down for small grains at very low energies, since the continuous cooling assumption no longer holds.

This is why most codes participating in this benchmark impose a lower temperature limit of 2.73 K, or even, rather arbitrarily, 1, 2, or 3 K. A limit of 2.73 K, for example, causes a bump in the PAH emission peaking at $\lambda \approx 1060 \, \mu$m (the peak CMB wavelength) because all the probabilities for lower temperatures are bunched together in the 2.73 K temperature bin. This effect is to some extent responsible for the discrepancies between the codes and the reference solution seen in the top lefthand panel of Fig. 4.7, in the wavelength range to the right of the dashed line, where the absolute value of the emissivity has become small anyway. The effect is negligible for all but the weakest fields.

At the other end of the scale, all codes in this benchmark use the complete temperature range for which the dust properties are defined, i.e. up to 2500 K for our dust model. As described in Sect. 4.6.3, this is well above the sublimation temperature of the dust material, although a fraction of the grains may survive at these temperatures for some time. Because the method used for this benchmark ignores dust grain destruction, we expect it to overestimate the emissivity for shorter wavelengths. To evaluate this effect, we reran the benchmark calculations with one of the codes (SKIRT) using a maximum grid temperature of 2250 K instead of 2500 K. As expected, the emissivities in the range $\lambda > 3 \, \mu$m are essentially unaffected by this change. At wavelengths shorter than 3 μm, the total emissivity for the hardest black body input fields ($T \gtrsim 15000$ K) decreases noticeably (by about 30% at 1 μm), while there is no perceptible change down to 1 μm for the softer fields or for the scaled Mathis fields. The silicate and graphite components behave similarly. Interestingly, the emissivity of the PAH component shows a noticeable decrease for all but the softest black body fields ($T \lesssim 6000$ K) and for all scaled Mathis fields. This is can be understood by recalling that the PAH particles are, on average, a lot smaller than those in the other components, so that they are more easily propelled to higher temperatures.

4.6.8   Wavelength discretization

The method described in Sect. 4.4 for calculating dust emission involves wavelength discretization in several distinct areas. The optical dust properties are tabulated on some predefined wavelength grid. We also need to configure the wavelength range and sampling resolution for the input fields and for the output emissivity. And finally, the calculation of the cooling coefficients defined in Eq. 4.27 requires an integration

of the optical properties over wavelength, which implies a grid as well. As long as proper interpolation procedures are in place, these wavelength grids do not need to be identical.

To keep matters simple for the benchmarks presented in this work, Sect. 4.3.3 specifies the same wavelength grid for the optical properties as for the calculated emissivities, i.e. a logarithmic grid with 1201 points in the range $0.001 \, \mu m \leqslant \lambda \leqslant 10000 \, \mu m$. In this section we discuss the impact of the resolution and of the lower and upper limits of the wavelength grid on the emissivity calculations. We reran the benchmark calculations with one of the codes (SKIRT) using some wavelength grid variations as reported below. SKIRT employs a single (configurable) wavelength grid for all aspects of the calculation. The optical dust properties are interpolated to this grid, and the grid is subsequently used for the input/output fields and for the integration to obtain the cooling coefficients.

In practice, the lower wavelength limit is determined by the need to properly capture the input field in the calculation of the heating coefficients (Eq. 4.24) and the equilibrium temperature (Eq. 4.22). For a scaled Mathis field, the lower limit can be increased to $0.09 \, \mu m$ (see Eq. 4.11). For the hardest black body fields ($T \gtrsim 12000 \, \mathrm{K}$), the limit should be lower. With a lower limit of $0.01 \, \mu m$, there is no noticeable difference in the calculated emissivities even for the hardest field in this benchmark ($T = 18000 \, \mathrm{K}$)

The upper wavelength limit is mostly determined by the need to calculate the emissivity up to mm wavelengths. The upper limit also affects the calculation of the heating coefficients and the equilibrium temperature, but this effect is smaller because the absorption coefficients for the grain material are much lower at longer wavelengths. With an upper limit of $2000 \, \mu m$ instead of $10000 \, \mu m$, there is no noticeable difference in the calculated emissivities for any of the input fields. With an upper limit of $1000 \, \mu m$, the submm emissivity peak is overestimated by 20% for the weakest scaled Mathis field ($U = 10^{-4}$). The emissivity peak for the $U = 10^{-3}$ field shows a similar but much smaller effect. For all other input fields there is no noticeable difference.

Finally, we reran the benchmark calculations for logarithmically distributed wavelength grids with successively lower resolution, always using a range of $0.01 \, \mu m \leqslant \lambda \leqslant 2000 \, \mu m$. First of all, lowering the wavelength resolution affects the shape of the sharp PAH-dominated emissivity peaks in the 3 to $30 \, \mu m$ wavelength range; if there is no wavelength point at the center of a peak, the peak cannot be resolved. However, this does not affect the accuracy of the emissivity at other wavelength points unless the resolution becomes too low, as described in what follows. Other than this peak resolution effect, using 601 wavelength points instead of 1201 does not noticeably influence the results for any of the input fields. Lowering the resolution to 301 wavelength points causes minor deviations in the calculated PAH emissivities for wavelengths $\lambda < 3 \, \mu m$, which however do not noticeably affect the total emissivities down to $\lambda \geq 1 \, \mu m$. With only 151 wavelength points the deviation in the total emissivities is still limited (a few percent at $1 \, \mu m$), but the PAH features

are now clearly under-resolved and fairly smoothed out. This can be improved by concentrating more grid points in the wavelength range of the PAH features. For example, a specialty grid with a total of 151 points, 61 of which are concentrated in the 3 to 30 μm wavelength range, seems to provide an acceptable compromise.

4.6.9   Grain size discretization

The total calculation time for the emissivity of a dust population is roughly proportional to the number of grain size bins used to represent the population (see Sect. 4.4.1). Therefore, some of the participating codes recompute the optical grain properties on an internally defined grain size grid rather than using the size bins tabulated in the dust model data. We used two of the codes (SKIRT and CRT) to investigate the effect of the number of grain size bins on the calculated emission. As mentioned in Sects. 4.5.2 and 4.5.5, these codes were configured with 15 size bins per grain type to calculate the benchmark results presented in Figs. 4.4 through 4.7.

The predominant effect of changing the number of grain size bins appears for wavelengths $\lambda < 10 \, \mu$m and increases for shorter wavelengths. This is to be expected, because with a coarse grid, the effective size of grains in the smallest bin is relatively large, and it is difficult to heat the grains to the high temperatures that are needed for emission at shorter wavelengths. Specifically, when the number of size bins per grain type is reduced from 15 to 10, the calculated silicate and graphite emissivities show substantial deviations from the reference solutions. The PAH emissivities are virtually unaffected, which is easily understood because their size range is much smaller. When the number of size bins per grain type is increased to 30, the calculated silicate and graphite emissivities in the wavelength range $1 \, \mu$m $< \lambda < 10 \, \mu$m approach the reference results. Again the PAH emissivities are much less affected because the dust model data only has 28 PAH size bins anyway.

A secondary effect occurs for longer wavelengths because the rebinning influences the heuristic for transitioning between the stochastic and equilibrium calculation regimes (Sect. 4.6.5). This effect seems to be somewhat random in nature, causing deviations that remain within the accuracy limits described in Sect. 4.6.4.

4.6.10   Calculation time

A typical 3D RT simulation calculates the dust emission spectra for a large number of dust cells. In cases where dust self-absorption is a relevant factor, this calculation is repeated for each iteration of the loop that self-consistently determines dust heating and re-emission (see e.g. Sect. 4.5.6). The time spent on calculating dust emission might thus become a significant or even dominant fraction of the total RT simulation time. The aim of reducing the dust emission calculation time has guided many of the choices in the implementations of the RT codes participating in this benchmark. Most fundamentally, all codes adopt the continuous cooling approximation. In addition, all codes select specific discretization schemes, most codes employ heuristics to transition between stochastic and equilibrium regimes, and some precompute field-

independent data. Often these choices affect not only the calculation time, but also the accuracy of the results. In principle at least, the results can be made to match perfectly by increasing grid sizes and removing the heuristics, at the expense of calculation time.

It thus seems appropriate to consider calculation time when evaluating benchmark results. Unfortunately, it is not meaningful to compare the dust emission calculation times between the codes outside of the context of a RT simulation. For example, moving the relevant data for each dust cell from memory into the processor cache and back may represent a significant portion of the total calculation time, depending on the memory layout chosen by the RT code and depending on the architecture of the computer system. Consequently, a performance comparison would be more appropriately conducted as part of a RT benchmark.

Just to provide an order of magnitude, with the prescriptions provided in this work, a code can calculate a few hundred dust emission spectra per second on a modern desktop computer. This means that the calculation for five million dust cells can be completed in a matter of hours rather than days.

4.7   Conclusions

We defined an appropriate problem for benchmarking dust emissivity calculations in the context of RT simulations, specifically including the emission from stochastically heated dust grains (SHGs). The problem's definition includes the optical and calorimetric material properties, and the grain size distributions, for a typical astronomical dust mixture with silicate, graphite, and PAH components. It also includes a series of analytically defined radiation fields to which the dust population is to be exposed and instructions for the desired output.

We summarized a popular method for calculating the emission from SHGs with the intention of providing a self-contained guide for implementors of such functionality. The method is frequently used in RT codes because of its good performance and relative ease of implementation, although it assumes continuous cooling of the dust grains, which may be inaccurate in extreme environmental conditions. We then described the six RT codes participating in this benchmark effort, focusing on how their implementation of the SHG calculation differs, presenting relevant heuristics for accelerating the calculation, and studying the effects on the accuracy of the solutions. We also presented some practical hints with regards to the discretization of temperature and wavelength in the calculations. Most importantly, we processed the benchmark problem with each of the participating codes, and presented the results.

We reported in detail on the similarities and differences between the results from the participating codes and a reference solution. In the important wavelength range $3\,\mu\mathrm{m} \leq \lambda \leq 1000\,\mu\mathrm{m}$, all participating codes reproduce the total dust emissivity within 20% of the reference solution for all input fields used in this benchmark. Excluding the weakest and the softest input fields, the agreement in the same wavelength

range is within 10%.

Our discussion offered hints to how RT codes could be set up to properly calculate dust emission for a wider wavelength range. For example, when investigating systems with a lot of hot dust, such as circumstellar disks or accretion disks, it may be relevant to properly calculate dust emission for wavelengths shorter than 1 μm. To accomplish this, RT codes will need to model environment-dependent destruction of dust grains and to adjust the grain size distribution used in the dust emission calculation accordingly.

In conclusion, this benchmark effort shows that the relevant modules in RT codes can and do produce fairly consistent results for the emissivity spectra of SHGs, which have a significant impact on the final result of a multiwavelength RT simulation. We offer concrete, quantitative information on the level of (dis)agreement between RT codes, which will help inform the interpretation of RT simulation results that include SHG dust emission calculations of the type presented here. Specifically, this work paves the way for a more extensive benchmark effort focusing on the RT aspects of the various codes. And finally, we intend this work to serve as a reference for implementors of existing and new dust RT codes.

Importing snapshots from hydrodynamical simulations[18]

5

A GENERATION WHICH IGNORES HISTORY
HAS NO PAST − AND NO FUTURE.

*Radiative transfer (RT) simulations are exceedingly performed on input models gener-
ated by computer programs that simulate the hydrodynamical evolution of an astrophys-
ical system over time. In this so called "post-processing" paradigm, the two simulation
codes run independently, and the only link is the set of data defining a "snapshot" of the
simulated system at a particular time in its evolution. In this chapter, we describe the
mechanisms offered by SKIRT for importing a snapshot produced by a hydrodynamical
simulation, and for configuring the desired RT simulation. We conclude that SKIRT's
capabilities are very flexible and enable post-processing snapshots from essentially any
hydrodynamical code with minimal data conversion and little or no programming
requirements.*

5.1	Introduction

As noted in Chapter 1, (quasi-)analytical models often fall short in describing the geo-
metric complexity of the real world. More realistic numerical models can be produced
by computer programs that simulate the (magneto-)hydrodynamical evolution of an
astrophysical system over time. Such codes usually output a series of *snapshots*. The
data provided in each snapshot define the state of the simulated system at a particular
time in its evolution, including the spatial distribution and other relevant properties of
the dynamical components, such as for example the stars and the interstellar medium
in a galaxy.

Because of the nonlocal and nonlinear behavior of the radiative transfer (RT) problem,
also noted in Chapter 1, predicting the observable properties of these hydrodynamical
models requires a full 3D RT simulation in all but the most trivial cases. Assuming
that the radiation crossing times are much shorter than the dynamical time scales,
and that radiation pressure is negligible, the RT and hydrodynamical simulations can
be completely uncoupled. In other words, the RT simulation can then be performed
as a *post-processing* step based solely on the information contained in a particular
snapshot. The RT code extracts the spatial distribution of the radiation sources (e.g.,

---

[18] I implemented all features described in this chapter in SKIRT. All design and implementation choices are
mine. The snapshot data formats generated by the hydrodynamical codes under consideration were of
course designed and documented by their original authors.

stars) and the obscuring medium (e.g., dust), plus the corresponding properties, from the hydrodynamical snapshot data, and then performs a RT simulation on the resulting setup.

The radiation sources and the obscuring medium are obviously handled differently during a Monte Carlo RT simulation, implying different needs for the corresponding data structures. The most frequent demand on the data structure representing the radiation sources (Sect. 5.2) is to generate a photon package launch position that is randomly sampled from the luminosity distribution at a given wavelength. This operation usually occurs millions or even billions of times during a single simulation, and thus should execute fairly quickly. On the other hand, the key operation for the dusty medium (Sect. 5.3) is to determine the dust density at a given location, or the dust mass in a given volume. These functions are needed while constructing the dust grid, i.e. for deciding where to subdivide cells in adaptive grids, and for calculating the dust mass in each cell of the final dust grid. Although this is an intensive process, it is limited to the setup phase of the simulation, because, once constructed, the dust grid includes all information required for tracing photon packages through the spatial domain.

Hydrodynamical simulation codes historically employ one of two schemes: a Lagrangian formulation based on moving particles (smoothed particle hydrodynamics or SPH), for example Gadget (Springel 2005) and Gasoline Wadsley et al. (2004); or a Eulerian approach based on a non-moving spatial grid, often an adaptive grid with multiple refinement levels (AMR), for example RAMSES (Teyssier 2002; Fromang et al. 2006), Enzo (Collins et al. 2010; Bryan et al. 2014), and AMR-VAC (Keppens et al. 2012). Recent codes including AREPO (Springel 2010), TESS (Duffell and MacFadyen 2011), and Shadowfax (Vandenbroucke and De Rijcke submitted) introduce a new scheme that employs a moving mesh based on a Voronoi tessellation of the spatial domain (see Chapter 3). This new scheme is claimed to combine the best features of SPH and the traditional Eulerian approach, and it is becoming increasingly popular. It has already been applied to various problems including the formation of stars, galaxies, and cosmological structures (Greif et al. 2011; Bauer and Springel 2012; Sijacki et al. 2012; Kereš et al. 2012; Vogelsberger et al. 2012; Torrey et al. 2012; Nelson et al. 2013; Marinacci et al. 2014).

In this chapter we discuss SKIRT's mechanisms for importing distributions of radiation sources (Sect. 5.2) and dust (Sect. 5.3) from snapshots produced by each of these three schemes. The final section (Sect. 5.4) offers some conclusions.

5.2	Radiation sources

A SKIRT simulation always includes one or more *stellar components* describing the radiation sources in the model; see Fig. 2.1 in Chapter 2. Each stellar component defines (a) the spatial density distribution of the sources and (b) the spectral energy distribution (SED) at each location. In the context of post-processing hydrodynamical

simulations, there are three distinct ways of accomplishing this:

- The imported snapshot defines the spatial density distribution and provides additional properties for each smoothed particle or grid cell from which an appropriate SED can be determined. For this purpose, SKIRT includes two parameterized SED families; see Table 2.2 in Chapter 2. The Bruzual-Charlot family represents young and evolved stellar populations (Bruzual and Charlot 2003). The Mappings III SED family represents star-forming regions (Groves et al. 2008). More families can be added with minimal programming effort, provided the data for the SED templates is available in analytical or tabulated form.

- The imported snapshot defines just the spatial density distribution, and a particular SED (constant throughout the spatial domain) is configured from the SED components built into SKIRT and listed in Table 2.2 in Chapter 2.

- The radiation sources are completely defined within SKIRT by configuring a built-in geometry and SED; see respectively Tables 2.1 and 2.2 in Chapter 2. In case a hydrodynamical simulation does not include any sources, this option can be used to complete the RT model by adding, for example, a central star to a stellar dust disk, or background radiation to the dust in a molecular cloud.

One can even combine these options by using two different stellar components, for example, for adding an active galactic nucleus to a spiral galaxy snapshot.

5.2.1   Smoothed particles

In spite of claims that the technique suffers from fundamental problems (Agertz et al. 2007; Bauer and Springel 2012), SPH is still a popular hydrodynamics technique, especially for cosmological simulations of galaxy formation (e.g., Guedes et al. 2011; Feldmann and Mayer 2015; Schaye et al. 2015). An SPH snapshot consists of a set of "particles" (or rather anchor points in a co-moving grid), each characterized by a large suite of physical quantities, of which only a selected few are relevant for SKIRT. Rather than attempting to support a wide variety of file formats, SKIRT reads the properties for a list of SPH particles from a text file in a simple format that can be easily generated by a conversion script, for example in a programming language such as Python. As an added benefit, this allows additional processing such as unit conversions to occur along the way. In Chapter 6 we describe a much more involved example of such a conversion script.

The spatial luminosity distribution in an SPH snapshot is defined by a list of N smoothed particles and an assumed smoothing kernel $W(r, h)$, with each smoothed particle characterized by a position $\mathbf{r}_j$, a smoothing length h_j, and a luminosity contribution L_j. The total luminosity density at an arbitrary position $\mathbf{r}$ is then given by

$$\Sigma(\mathbf{r}) = \sum_{j=1}^{N} L_j \, W(|\mathbf{r} - \mathbf{r}_j|, h_j) \tag{5.1}$$
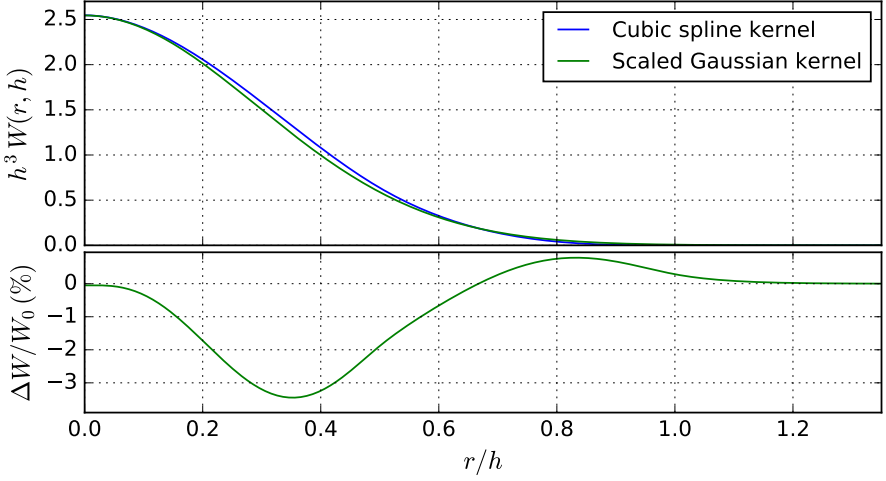
Figure 5.1: Comparison of the scaled Gaussian kernel (Eq. 5.2, green curve) used in SKIRT for generating random positions inside a smoothed particle and the standard cubic spline kernel (Eq. 5.6, blue curve) used in SKIRT for calculating the density of a smoothed particle at a given location. The bottom panel shows the difference between the two kernels as a percentage of the kernel value at $r = 0$.

In practice, the kernels used in SPH simulations almost always have a finite support (e.g., Monaghan and Lattanzio 1985; Desbrun and Gascuel 1996; Müller et al. 2003), so that only a relatively small number of terms in the sum have a non-zero contribution. SKIRT employs smoothing kernel implementations optimized for each specific task. For shooting photon packages, the code needs to generate random launch positions sampled from the radiation source's density distribution. This is rather straightforward in this case, thanks to the composition method (Baes and Camps 2015). The first step is the choice of a random smoothed particle, based on a discrete distribution where each particle is weighted by its relative luminosity contribution. The second step is generating a random position according to the luminosity distribution of the chosen particle. Although it is in principle possible to sample from any kernel using tabulated values, SKIRT samples a random position from a scaled Gaussian smoothing kernel,

$$W(r,h) = \frac{a^3}{\pi^{3/2}\,h^3}\,\exp\left(-\frac{a^2 r^2}{h^2}\right) \tag{5.2}$$

with the value of the scaling factor set to $a = 2.42 \approx 2\pi^{1/6}$ (Altay and Theuns 2013). Sampling from a Gaussian kernel can be implemented easily using the inversion method (Baes and Camps 2015). For consistency, the value of a has been chosen so that the scaled Gaussian kernel approximates (within a few per cent across the domain) the cubic spline kernel which is used in SKIRT for determining dust densities; see Eq. (5.6) in Sect. 5.3.1 and Fig. 5.1.

5.2.1.1   Spatial distribution and SED from snapshot

We first consider the case where both the spatial distribution and the SED at each location are determined from the SPH snapshot. The `SPHStellarComp` class in SKIRT expects a text file with a single line for each SPH particle. The first four columns specify the x, y and z coordinates and the smoothing length h for the particle (all in parsec). The number and interpretation of the subsequent columns depends on the SED family configured for this `SPHStellarComp` instance.

For the Bruzual-Charlot SED family, the remaining three columns provide the properties of the stellar population represented by the particle. The first extra column specifies the initial mass of the stellar population, i.e. the particle's mass M_{init} at the time it was born (in $M_\odot$). The second extra column specifies the metallicity Z of the stellar population (as a dimensionless fraction), and the third extra column represents the age t of the stellar population (in years).

For the Mappings III SED family, the remaining five columns provide the properties of the star-bursting region represented by the particle, in the following order: the star formation rate $\dot{M}$, assumed to be constant over the past 10 Myr (in $M_\odot\,yr^{-1}$); the metallicity Z (as a dimensionless fraction); the logarithm of the compactness $\log C$ (dimensionless); the pressure of the surrounding medium p (in Pa); and the photo-dissociation-region covering factor f_{PDR} (dimensionless).

Using these properties as parameters for the appropriate set of built-in SED templates, SKIRT calculates the luminosity contribution $(L_j)_\ell$ of each particle (used in Eq. 5.1) for each wavelength bin ℓ in the simulation's wavelength grid. This information is stored in memory during setup, and is used during the Monte Carlo cycle to determine the (wavelength-dependent) probability that a given particle will be selected as the launch site for a new photon package.

5.2.1.2   Velocity and Doppler shift

The `SPHStellarComp` class can be configured to take into account the Doppler shift caused by the velocity of the radiation sources relative to the center of mass of the system. If this option is enabled, the input text file must provide three additional columns (after the smoothing length and before the additional properties for the SED family), specifying the v_x, v_y and v_z components of the particle's velocity (in km/s). Given this velocity $\mathbf{v}$, the redshift z experienced by a photon package launched from the particle in direction $\mathbf{k}$ (given as a unit vector) is determined by

$$z = -\frac{\mathbf{v} \cdot \mathbf{k}}{c} \tag{5.3}$$

where c is the speed of light in vacuum. In function of the rest wavelength λ_0, the redshifted wavelength λ_z is given by $\lambda_z = (1 + z)\,\lambda_0$. For $z \ll 1$, the observed luminosity at each wavelength in the simulation's wavelength grid can then be written as $L_z[\lambda_\ell] = L_0[(1 - z)\,\lambda_\ell]$.

Since the redshift value depends on the angle between the particle and photon package trajectories, the emission is no longer isotropic. When launching a photon package from a given particle, the Monte Carlo code must randomly pick a direction from a non-uniform distribution. To this end, we express the photon package's direction in spherical coordinates (θ, ϕ) relative to a polar axis along the particle's trajectory. While the azimuth ϕ is still distributed uniformly, the probability distribution of the polar angle θ for wavelength bin ℓ is determined by the form of the SED assigned to the particle in the wavelength range $[(1 - v/c)\,\lambda_\ell, (1 + v/c)\,\lambda_\ell]$, where $v = ||\mathbf{v}||$ is the magnitude of the particle's velocity. The need to store this distribution at an acceptable resolution of, say, 1 degree in θ, for each particle and for each wavelength bin substantially boosts memory requirements. Otherwise the implementation is fairly straightforward.

5.2.1.3   Spatial distribution from snapshot with built-in SED

If the hydrodynamical simulation does not track the relevant properties for defining the local emission spectrum, we can still import the spatial luminosity distribution from the snapshot and assign a built-in SED that is constant across the spatial domain. In this case, we use the `GeometricStellarComp` class, and we configure it with an instance of the `SPHGeometry` class to import the snapshot. Similar to the `SPHStellarComp` class described in Sect. 5.2.1.1, the `SPHGeometry` class expects a text file with a single line for each SPH particle. The first four columns specify the x, y and z coordinates and the smoothing length h for the particle (all in parsec). The product of the values in the fifth and sixth columns specifies the density contribution of this particle, in arbitrary units since the `SPHGeometry` class normalizes the total weight of the density distribution to unity. (The definition uses two columns for compatibility with dust density distributions; see Sect. 5.3.1.1).

Because the `SPHGeometry` class normalizes the spatial distribution to unity, we need to separately supply the actual luminosity of the source through one of the standard normalization options offered in SKIRT. Because the SED is constant across the spatial domain, it is sufficient to set, for example, the luminosity at one particular wavelength.

5.2.2   Hierarchical grids

Apart from smoothed particle hydrodynamics, the main other technique that is used to perform hydrodynamics simulations is Eulerian mesh-based hydrodynamics (Stone and Norman 1992; Fryxell et al. 2000). A fundamental ingredient of this technique is the use of adaptive mesh refinement (AMR), where the hierarchical grid is adjusted depending on the resolution requirements in various regions of the model (Kurosawa and Hillier 2001; Niccolini and Alcolea 2006; Keppens et al. 2012). While the resolution requirements and thus the AMR grid may change as the system evolves, each particular snapshot corresponds to a unique and well-defined hierarchical grid. The SKIRT class names and documentation do not properly make

this distinction and use the terms "AMR", "adaptive mesh" and "hierarchical grid" interchangeably.

A hierarchical grid can be defined using spherical or cylindrical coordinates (e.g., Chen 1990; Park et al. 2010), perhaps to benefit from certain quasi-symmetries in the model, however, we will limit the discussion here to grids using Cartesian coordinates. In a Cartesian hierarchical grid, the cuboidal spatial domain is recursively subdivided into smaller cuboids according to some scheme until the desired resolution in each area is reached. Thus, an AMR snapshot includes some description of the structure of the hierarchical grid, implicitly or explicitly defining the spatial extent of each grid cell, and for each cell the values of a suite of physical quantities of which only a selected few are relevant for SKIRT.

5.2.2.1   Snapshot data format

SKIRT can directly read the binary data format generated by the mesh-based MPI-AMRVAC code developed by Keppens et al. (2012). This capability was used, for example, to facilitate the studies presented by Hendrix et al. (2015) and by Keppens et al. (in press). The specifics of this import mechanism are described in the SKIRT documentation. In this section we discuss a more general import mechanism based on a plain text file format specifically designed for this purpose. For most mesh-based snapshots (including the MPI-AMRVAC snapshots), writing a conversion script to generate this format should be rather straightforward.

The text file format we designed describes the structure of the hierarchical grid and lists the relevant physical properties in each cell. In other words, a file in this format represents one or more scalar fields over a given cubodial spatial domain (assuming Cartesian coordinates) which is recursively subdivided into cuboidal cells. The size of the spatial domain, the meaning of the fields, and the units in which the values are expressed, must be defined elsewhere (i.e. this information is not part of the data format itself).

The hierarchical grid structure is organized into a tree. Each tree node represents a cuboidal portion of the domain, called its *extent*. The *root* node's extent is the complete domain. A *nonleaf* node distributes its extent over its child nodes using a regular linear grid. The number of subdivisions is defined separately for each node and can differ for each spatial dimension. An octree, for example, would subdivide each nonleaf node into $2 \times 2 \times 2$ child nodes. Typical AMR schemes have much larger subdivision counts that sometimes vary within the grid; for example, the root node may be subdivided differently. A *leaf* node represents a cell that is not subdivided any further and that holds a data value for each field; the field values are considered to be constant over the leaf node's extent. Collectively the leaf nodes form a partition of the domain, i.e. their extents cover the complete domain without overlapping one another.

The leaf nodes or cells in this three-dimensional data structure can be arranged in
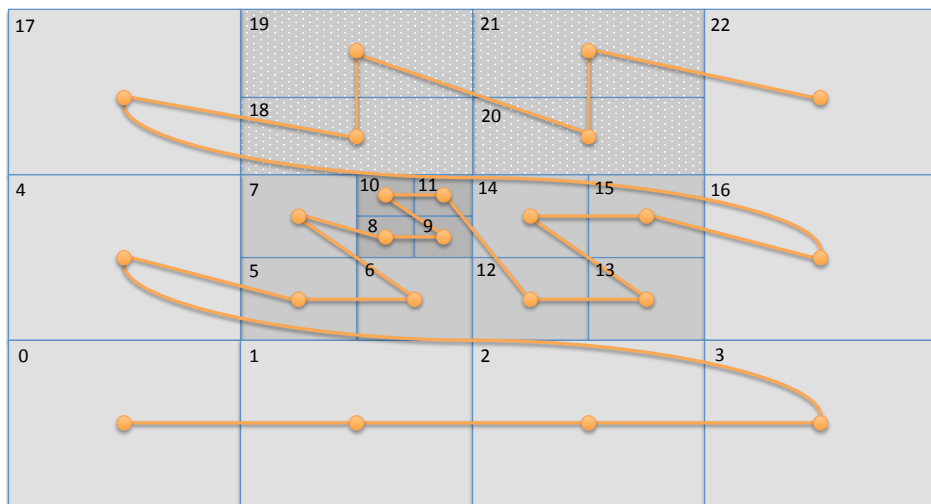
Figure 5.2: Illustration of the Morton order for the cells in a 2D hierarchical grid. The root node is subdivided into 4×3 subnodes; some of the nodes near the center are subdivided into 2×2 subnodes, and two of the cells in the upper row are subdivided into 1×2 subnodes. The solid line connects the leaf nodes in Morton order, starting at the lower left node with Morton index 0 and performing a depth-first traversal with each nonleaf node visiting its children in the order first x (horizontal) and then y (vertical).

```
1    # Snapshot text data file        16    11 0.8
2    ! 4 3 1                          17    ! 2 2 1
3    0 0.4                            18    12 0.6
4    1 0.4                            19    13 0.6
5    2 0.4                            20    14 0.6
6    3 0.4                            21    15 0.6
7    4 0.4                            22    16 0.4
8    ! 2 2 1                          23    17 0.4
9    5 0.6                            24    ! 1 2 1
10   6 0.6                            25    18 0.5
11   7 0.6                            26    19 0.5
12   ! 2 2 1                          27    ! 1 2 1
13   8 0.8                            28    20 0.5
14   9 0.8                            29    21 0.5
15   10 0.8                           30    22 0.4
```

Figure 5.3: Representation of the 2D grid shown in Fig. 5.2 in the snapshot text file format read by SKIRT. Because SKIRT deals with 3D grids, we've added a third dimension that is never subdivided. The lines starting with an exclamation mark (in red) indicate a subdivided node, e.g., the root node is subdivided in $4 \times 3 \times 1$ cells. The other lines (in black) specify the field values for a particular cell. In this example, the first field value specifies the serial number or *Morton index* of the cell, and the second value specifies the cell's "density" reflected by the gray level in Fig. 5.2. In an actual snapshot, these illustrative values would be replaced by the relevant physical quantities.

a linear sequence using so-called *Morton ordering* (Morton 1966). This ordering is obtained by performing a depth-first traversal of the tree, where each nonleaf node visits its children in the order x-first, then y, then z. The process is illustrated in Fig. 5.2 for a two-dimensional structure. The overall makeup of the snapshot text file reflects this structure and ordering. Each line describes a particular tree node (nonleaf or leaf), and the lines are given in Morton order. Specifically, each line in the file can be of one of the following types:

- *Comment*: lines having a number sign or hash (#) as the first non-whitespace character, lines containing only whitespace and empty lines are ignored (and do not count in the Morton order).

- *Nonleaf*: a nonleaf line has an exclamation mark (!) as the first non-whitespace character, followed by optional whitespace and then three whitespace-separated positive integer numbers N_x, N_y, N_z. These three numbers specify the number of child nodes carried by this node in each spatial direction. The child nodes are on a regular linear grid as described above.

- *Leaf*: a leaf node contains one or more whitespace-separated floating point numbers, specifying the values of the fields in the cell represented by this leaf node. All leaf node lines in the file must contain the same number of field values.

Note that there is no need to provide the cell positions because the geometry follows from the Morton order (and the total domain size, which is passed to SKIRT independently of the snapshot file). As an example, Fig. 5.3 shows a representation of the 2D grid from Fig. 5.2 in this format.

5.2.2.2   Spatial distribution and SED from snapshot

As before, we first consider the case where both the spatial distribution and the SED at each location are determined from the snapshot. The `AdaptiveMeshStellarComp` class in SKIRT expects a snapshot data file in the format described above, supplemented with the size of the spatial domain in each direction, because that information is not stored in the file. In addition, the user must configure which of the scalar fields in the snapshot file carry the relevant physical quantities for determining the SED assigned to each cell.

The current implementation supports just the Bruzual-Charlot SED family, so the `AdaptiveMeshStellarComp` class expects three values, respectively specifying the initial stellar density ρ_{init} (in $M_\odot \, pc^{-3}$), the metallicity Z of the stellar population (dimensionless fraction), and the age t of the stellar population (in years). Using these properties as parameters for the Bruzual-Charlot SED templates, SKIRT calculates the luminosity contribution of each cell for each wavelength bin. This information is used during the Monte Carlo cycle to determine the (wavelength-dependent) probability that a given cell will be selected as the launch site for a new photon package. Generating a random launch position from a uniform distribution within the selected cuboidal cell is of course trivial.

When the need arises, we can add support for other SED families, similar to what was described for the `SPHStellarComp` class dealing with SPH snapshots in Sect. 5.2.1.1.

5.2.2.3   Spatial distribution from snapshot with built-in SED

Similar to the procedure described for SPH snapshots in Sect. 5.2.1.3, we can import the spatial luminosity distribution from a mesh-based snapshot and assign a built-in SED that is constant across the spatial domain. In this case, we configure an instance of the `AdaptiveMeshGeometry` class to import the snapshot, separately specifying the extent of the spatial domain as before. The density contribution of each cell is given by the product of two field values with configurable indices, for compatibility with dust density distributions; see Sect. 5.3.2.1. Because the `AdaptiveMeshGeometry` class normalizes the spatial distribution to unity, we need to separately supply the actual luminosity of the source through one of the standard normalization options offered in SKIRT.

5.2.3   Voronoi grids

Recently, a new Lagrangian technique that solves the hydrodynamics equations on a moving, unstructured Voronoi grid is gaining popularity. It is claimed to avoid some of the difficulties of smoothed particle hydrodynamics on the one hand and Eulerian grid-based hydrodynamics on the other hand. This technique has been used for many years in the computational fluid dynamics community (Mavriplis 1997), and a number of novel codes based on this principle have recently been developed in the astrophysics community (Springel 2010; Duffell and MacFadyen 2011; Vandenbroucke and De Rijcke submitted). Moving mesh hydrodynamics is mainly applied to simulations of galaxy formation and evolution (e.g., Marinacci et al. 2014; Vogelsberger et al. 2014).

Due to the nature of a Voronoi tesselation (see Chapter 3), the geometry of the grid is completely defined by the positions of the generating sites. It is hence not necessary for a snapshot to store all the vertices and edges of each of the cells. Similar to our approach for other snapshot types, SKIRT reads the properties for a list of Voronoi cells from a text file in a simple format that can be easily generated by a conversion script. SKIRT constructs the Voronoi grid from the positions of the generating sites, and assigns the corresponding physical values to each cell, assuming that the values are constant across the cell's extent.

5.2.3.1   Launching photon packages

The generation of random positions for launching photon packages is similar to the procedure for hierarchical grids. In the first step, we pick a random cell using a discrete distribution where each cell is weighted by its relative (possibly wavelength-dependent) luminosity contribution. The second step, generating a random position

from within the chosen cell, is significantly more complex than in the case of a cuboidal cell. To the best of our knowledge, there are no dedicated techniques to generate a random point from a Voronoi cell. There are two possible options.

The first option is to partition the cell into a set of tetrahedra, subsequently select a random tetrahedron from a discrete distribution where every tetrahedron is weighted by its relative volume, and finally generate a random position from the selected tetrahedron. Specific algorithms are available for both the tetrahedrization of convex polyhedra (Edelsbrunner et al. 1990; Max 2001) and the generation of random positions from a tetrahedron (Rocchini and Cignoni 2000).

The second option, which is more straightforward and which we adopted in SKIRT, is to use the rejection technique. As the reference distribution we use a uniform density in a cuboidal volume, defined as the 3D bounding box of the cell. During construction of the Voronoi grid, for each cell, we store the cell bounding box (easily obtained from the cell vertices because Voronoi cells are convex polyhedra) and a list of all neighboring cells; also see Sects. 3.2.2 and 3.2.3. After generating a candidate random position uniformly from the cell bounding box, we decide whether the position is actually inside the Voronoi cell by checking that it is closer to the cell's generating site than to the generating sites of all its neighbors.

Our tests have shown that, depending on the distribution of the generating Voronoi sites, the average ratio of the volume of the bounding box of a Voronoi cell over the actual cell volume is about 3 to 4. This ratio represents the average rejection rate for the random position generation.

5.2.3.2   Loading snapshots

The procedures for loading a spatial and/or spectral distribution from a Voronoi-tesselation-based snapshot into SKIRT are very similar to what was described for the other snapshot types in Sects. 5.2.1 and 5.2.2.

The `VoronoiStellarComp` class expects a snapshot data file in a straightforward text column format, containing one line per cell. The first three columns provide the x, y, and z coordinates of the generating site for the cell, and the subsequent columns specify the relevant physical quantities for determining the SED assigned to each cell. The current implementation supports just the Bruzual-Charlot SED family, so the class expects three values specifying the initial stellar density, the metallicity of the stellar population, and the age of the stellar population. The order in which these are specified can be configured by the user.

The `VoronoiGeometry` class imports just the spatial luminosity distribution from a Voronoi-based snapshot, allowing to separately assign a built-in SED that is constant across the spatial domain. The file format is similar as described above, however the density contribution of each Voronoi cell is now given by the product of two field values with configurable indices. Because the class normalizes the spatial distribution

to unity, we need to separately supply the actual luminosity of the source through one of the standard normalization options offered in SKIRT.

| 5.3 | Dust properties |

Apart from radiation sources, any nontrivial SKIRT simulation also includes a *dust system* with one or more *dust components* describing the dusty medium in the model; see Fig. 2.1 in Chapter 2. Each dust component defines (a) the spatial density distribution of the dust and (b) the relevant material properties of the dust. In addition, a dust system also configures a *dust grid* that will be used to discretize the dusty medium for the purposes of the RT simulation. We will discuss each of these aspects in turn.

SKIRT offers two distinct ways to define the dust density distribution:

- The imported snapshot fully defines the dust density distribution, either directly or as a simple formula in function of the properties of the gas density as explained in Sect. 5.3.1. With this option, the dust system is configured with an instance of one of the `SPH-`, `AdaptiveMesh-`, or `VoronoiDustDistribution` classes. The latter two classes allow configuring multiple dust components (each with their own distribution and dust properties) to be loaded from the same snapshot.

- The dust density distribution is defined by configuring one or more built-in geometries (see Table 2.1 in Chapter 2) through the `CompDustDistribution` class, including one or more instances of the `SPH-`, `AdaptiveMesh-`, or `VoronoiGeometry` classes, and assigning separate dust properties to each of those components. With this option, it is possible to combine quasi-analytic dust geometries with densities loaded from a snapshot.

Each of the dust components configured as described above must be assigned a specific *dust mixture*, which defines the optical and calorimetric properties of the dust grains; see Table 2.3 and Sect. 2.2.2 in Chapter 2. The dust properties are considered to be constant across the spatial domain of the dust component, but evidently different dust components can be assigned different dust mixtures.

Note that the stellar system and the dust system each have their own distinct import mechanism. It is perfectly possible to import, say, radiation sources from SPH particles and dust densities from an AMR grid. Such a combination is not uncommon. The RAMSES code (Teyssier 2002), for example, can use an N-body solver for the pressure-less dark matter and stellar particles while employing a mesh-based hydrodynamical solver for the gas medium in the same simulation (e.g., Renaud et al. 2013).

Finally, a dust grid must be properly configured to provide adequate spatial resolution in areas with high dust densities or large density gradients. When importing a hydrodynamical simulation snapshot, the best option is to select an adaptive grid, such as the hierarchical octree and *k*-d tree grids described in Saftly et al. (2013) and Saftly et al. (2014), or the Voronoi grid discussed in Chapter 3. SKIRT offers several
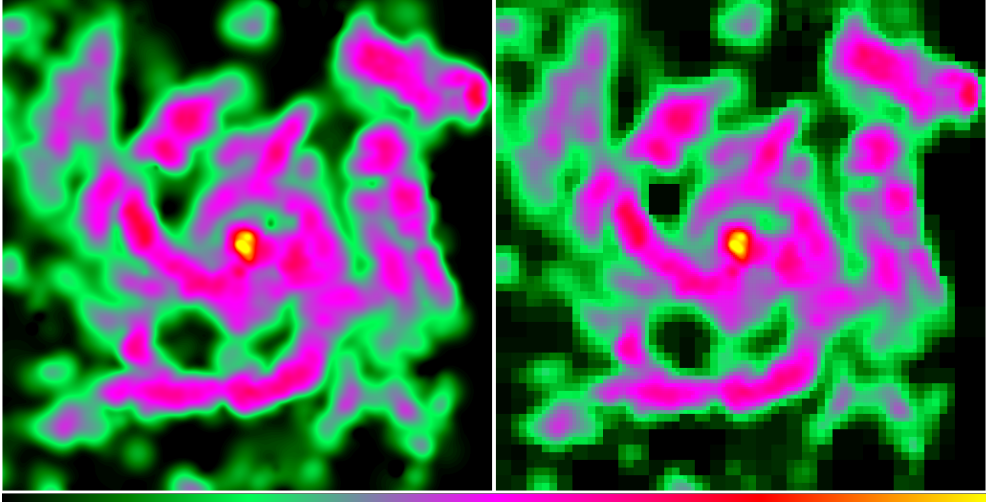
Figure 5.4: Cut through the dust density distribution in one of the EAGLE galaxies (see Chapter 6) as seen by SKIRT. The left panel shows the "theoretical" distribution as it is interpolated from the SPH particles in the EAGLE snapshot. The right panel shows the "gridded" distribution after discretization with the octree grid built into SKIRT (for illustration purposes, the grid is constructed with larger cells than what would be used in an actual simulation). The colors indicate dust density on an arbitrary scale (color bar at the bottom).

options to help evaluate and configure the quality of the dust grid; see, e.g., Saftly et al. (2014). It is also instructive to review the dust density cuts along the coordinate planes generated by SKIRT upon request; see Fig. 5.4. The "theoretical" density maps shows the dust density as it is defined through one of the `DustDistribution` classes, while the "gridded" density shows the dust density after discretization in the dust grid.

5.3.1   Smoothed particles

The mechanisms for loading the dust density from an SPH snapshot are similar to those discussed in Sect. 5.2.1 for radiation sources. There are, however, some differences and additional considerations. The total dust density at an arbitrary position is obtained through the equivalent of Eq. 5.1,

$$\rho(\mathbf{r}) = \sum_{j=1}^{N} M_j \, W(|\mathbf{r} - \mathbf{r}_j|, h_j) \tag{5.4}$$

where M_j now indicates the dust mass contribution for each particle. This value may simply be provided in the snapshot as a particle property, but just as likely, the hydrodynamical simulation traces the evolution of the gas medium as a whole rather than the dust per se. To facilitate those cases, SKIRT can calculate the dust mass from the gas mass in the input data, assuming that the amount of dust is proportional to the metal fraction in the gas, except in areas where the gas is too hot to form dust. In

other words, dropping the particle index, we have

$$M_{\text{dust}} = \begin{cases} f_{\text{dust}} \, Z \, M_{\text{gas}} & \text{if } T < T_{\text{max}} \\ 0 & \text{otherwise} \end{cases} \tag{5.5}$$

where M_{gas}, Z, and T are the gas mass, metallicity, and temperature given by the gas particle's properties in the snapshot, and f_{dust} and T_{max} are constant parameters specified when configuring SKIRT. A fixed dust-to-metal fraction f_{dust} is widely assumed to be an appropriate approximation for a variety of environments, although observed values range from $f_{\text{dust}} = 0.2$ to 0.4 (Dwek 1998; James et al. 2002; Brinchmann et al. 2013; Zafar and Watson 2013). If the cutoff temperature T_{max} is set to zero, the temperature criterion is ignored and the gas temperature does not need to be specified in the snapshot data.

As already mentioned in Sect. 5.2.1, SKIRT employs implementations for the smoothing kernel $W(r, h)$ that are optimized for each specific task. To calculate the dust density at a given location according to equation (5.4), SKIRT uses the cubic spline kernel,

$$W(r, h) = \frac{8}{\pi \, h^3} \times \begin{cases} 1 - 6 \, u^2 \, (1 - u) & \text{for } 0 < u < \frac{1}{2} \\ 2 \, (1 - u)^3 & \text{for } \frac{1}{2} < u < 1 \\ 0 & \text{otherwise} \end{cases} \tag{5.6}$$

with $u = r/h$. This kernel has a finite support so that the operation can be limited to particles that potentially overlap the location of interest. To facilitate this process, the setup phase of the simulation places a rough grid over the spatial domain and constructs a list of overlapping particles for each grid cell.

As described in Sect. 2.3.3 in Chapter 2, a further optimization is provided to calculate the mass within a given box (a cuboid lined up with the coordinate axes), as an alternative to sampling the density in various locations across the volume of the box. In this case, the calculation uses the analytical properties of the scaled and clipped Gaussian kernel defined by Eq. 5.2 in Sect. 5.2.1, which was designed to approximate the cubic spline kernel, to directly determine the mass in the box. The advantage of this kernel is that the integration over a box can be written in terms of the error function,

$$\int_{x_1}^{x_2} \int_{y_1}^{y_2} \int_{z_1}^{z_2} W(\sqrt{x^2 + y^2 + z^2}, h) \, dx \, dy \, dz = \frac{1}{8} \left(\text{erf}(\frac{a \, x_2}{h}) - (\text{erf}(\frac{a \, x_1}{h})) \right)$$
$$\times \left(\text{erf}(\frac{a \, y_2}{h}) - (\text{erf}(\frac{a \, y_1}{h})) \right) \times \left(\text{erf}(\frac{a \, z_2}{h}) - (\text{erf}(\frac{a \, z_1}{h})) \right) \tag{5.7}$$

Assuming the error function can be calculated quickly or, more likely, is tabulated at sufficient resolution, this optimization accelerates the density calculation for typical Cartesian grids by an order of magnitude.

5.3.1.1   Loading snapshots

The `SPHDustDistribution` class expects a text file in a format similar to what was described in Sect. 5.2.1.1 for radiation sources. Again, the first four columns specify the x, y and z coordinates and the smoothing length h for each particle (all in parsec). The subsequent columns now specify the gas mass M_{gas} (in $M_{\odot}$), metallicity Z (dimensionless fraction) and temperature T (in K). The last column is optional; if the temperature value is missing, the particle is considered to contain dust regardless of the temperature cutoff value T_{max}.

We can also configure the `CompDustDistribution` class with an an instance of the `SPHGeometry` class. The text file expected by this class has the same format as the one expected by the `SPHDustDistribution` class. However, the `SPHGeometry` class normalizes the total weight of the density distribution to unity, which means that we need to supply the total dust mass through one of the standard normalization mechanisms offered by SKIRT. Note that the `SPHGeometry` class can also be configured as part of a stellar system (see Sect. 5.2.1.3). In that case, the values in the fifth and sixth columns no longer have the meaning of mass and metallicity, and their product simply defines the relative luminosity contribution.

5.3.2   Hierarchical grids and Voronoi grids

The SKIRT classes for importing grid-based snapshots currently do not implement the temperature cutoff criterion described by Eq. 5.5 in Sect. 5.3.1. For each cell in the grid, the dust density is thus simply calculated according to

$$\rho_{\mathrm{dust}} = f_{\mathrm{dust}} \, Z \, \rho_{\mathrm{gas}} \tag{5.8}$$

where ρ_{gas} and Z are the gas density and metallicity given by the cell properties in the snapshot, and f_{dust} is a constant parameter specified when configuring SKIRT. In case the hydrodynamical simulation directly traces a dust population, we can configure SKIRT to use $f_{\mathrm{dust}} = 1$ and to ignore metallicity (effectively setting $Z = 1$ as well).

Determining the dust density at an arbitrary position comes down to identifying the cell that contains that position and returning the density associated with this cell. For a hierarchical grid this entails a simple search that starts at the root node and recursively descends into the child node that happens to contain the given position until a leaf node has been reached. With the cuboidal cells in a Cartesian grid this is rather straightforward. In the case of a Voronoi grid, however, the cell identification is not as simple. Due to the nature of Voronoi grids, locating the appropriate cell is essentially a nearest neighbor search. Rather than looping over all possible cells (or equivalently, all generating sites), SKIRT implements an approach using cuboidal blocks, as explained in Sect. 3.2.3 in Chapter 3.

5.3.2.1 Loading snapshots

The file formats expected by the `AdaptiveMesh-` and `VoronoiDustDistribution`
classes are very similar to those described for radiation sources respectively in
Sects. 5.2.2.1 and 5.2.3.2. The properties determining the SED for radiation sources
are in this case replaced by the gas density ρ_{gas} and metallicity Z required for
Eq. 5.8.

Alternatively, we can configure an instance of the `AdaptiveMesh-` or `VoronoiGeometry`
classes, which normalize the total weight of the density distribution to unity, so that
we need to supply the total dust mass through one of the standard normalization
mechanisms offered by SKIRT.

5.3.2.2 Configuring the RT dust grid

For grid-based hydrodynamical snapshots, SKIRT offers the option to perform the RT
simulation using the imported grid. In this case, rather than constructing a RT dust
grid using some configured scheme based on sampling the density in the imported
distribution, SKIRT directly adopts the grid on which the snapshot has been defined.
However, while this seems a logical thing to do, it is often not the best choice for
several reasons.

Most importantly, the resolution requirements for the RT treatment usually differ
from those for the hydrodynamical simulation. We might not need the same overall
resolution, and fine-grained cells may be needed in different areas of the domain.
Also, sometimes the SKIRT dust grid can be made smaller than the full snapshot
domain, cutting off the outer areas that do not contain significant amounts of dust
anyway. And finally, the octree and k-d tree dust grids in SKIRT are highly optimized
for photon package shooting, so the grid construction time is usually easily regained
during the actual RT simulation phase.

5.4	Conclusions

We described the various mechanisms, implemented in SKIRT as part of this thesis,
for importing snapshots generated by hydrodynamical simulation codes. We can
conclude that SKIRT's capabilities are very flexible and enable RT post-processing of
snapshots from essentially any hydrodynamical code with minimal data conversion
and little or no programming requirements.

At the same time, it is fair to point out some limitations. As it stands, several capabili-
ties are implemented only for SPH snapshots, e.g., support for a configurable SED
family (Sect. 5.2.1.1) or for velocity Doppler shift (Sect. 5.2.1.2), and other capabilities
are implemented only for grid-based snapshots, e.g., support for binary data formats
(Sect. 5.2.2.1). Adding these capabilities for the other snapshot types would be fairly

straightforward, but doing so within the current class structure would cause substantial code duplication and maintenance headaches. Similarly, the smoothed particle kernels employed in SKIRT are hard-coded and optimized for each use case, e.g., sampling random locations (Sect. 5.2.1) or calculating the density at a given location (Sect. 5.3.1). Instead, the user should be able to select a kernel that matches the kernel used in the originating SPH simulation, even if this would have some performance implications.

To address these issues, it is our intention to design and implement a cleaner structure of SKIRT's import modules in future work, which will benefit the code, the developer, and more importantly, the user.

Dust in EAGLE galaxies[20]

6

A TOUCHSTONE TO DETERMINE THE ACTUAL WORTH OF AN
"INTELLECTUAL" – FIND OUT HOW HE FEELS ABOUT ASTROLOGY.

The EAGLE cosmological simulations reproduce the observed galaxy stellar mass func-tion and many galaxy properties. In this work, we study the dust-related properties of present-day EAGLE galaxies through mock observations in the far-infrared and submm wavelength ranges obtained with the 3D dust radiative transfer code SKIRT. To prepare an EAGLE galaxy for radiative transfer processing, we derive a diffuse dust distribution from the gas particles and we re-sample the star-forming gas particles and the youngest star particles into star-forming regions that are assigned dedicated emission templates. We select a set of redshift-zero EAGLE galaxies that matches the K-band luminosity distribution of the galaxies in the Herschel *Reference Survey (HRS), a volume-limited sample of about 300 normal galaxies in the Local Universe. We find overall agreement of the EAGLE dust scaling relations with those observed in the HRS, such as the dust-to-stellar mass ratio versus stellar mass and versus* $NUV - r$ *color relations. A discrepancy in the* f_{250}/f_{350} *versus* f_{350}/f_{500} *submm color-color relation implies that part of the simulated dust is insufficiently heated, likely because of limitations in our sub-grid model for star-forming regions. We also investigate the effect of adjusting the metal-to-dust ratio and the covering factor of the photodissociation regions surrounding the star-forming cores. We are able to constrain the important dust-related parameters in our method, informing the calculation of dust attenuation for EAGLE galaxies in the UV and optical domain.*

6.1   Introduction

Cosmological simulations are a valuable tool in the study of how galaxies form and evolve. Recently, hydrodynamical simulations of the formation of galaxies in cosmologically representative volumes have succeeded in reproducing many – but not all – observed properties of galaxies and of the intergalactic medium to unprecedented levels of agreement (e.g., Le Brun et al. 2014; Vogelsberger et al. 2014; Schaye et al. 2015). The mass resolution for baryonic matter in these simulations is

---

[20] Submitted to MNRAS as Camps et al. (submitted). Other authors (as referenced in the text) performed and documented the EAGLE simulation suite. James Trayford contributed the procedure preparing EAGLE galaxies for postprocessing with SKIRT in the optical wavelength range. I extended the procedure for use in the wider range from UV to submm and performed all other work reported in this chapter. All co-authors contributed suggestions and text fragments to the manuscript.

on the order of 10^6 solar masses. Physical processes on unresolved scales (including star formation and stellar feedback) are handled through sub-grid prescriptions. Zoom-in simulations (e.g., Hopkins et al. 2014; Wang et al. 2015; McKinnon et al. 2016; Sawala et al. 2016) offer a better resolution, however, they still use similar sub-grid prescriptions. Inevitably these limitations lead to uncertainties in some of the simulation predictions.

By comparing simulation results and observations we hope to examine the empirical scaling laws, deduce improved sub-grid prescriptions, and eventually, to further our understanding of the underlying physical processes. Because properties of real galaxies are derived from observed quantities (i.e. fluxes), they may be subject to unknown systematic biases. Making mock observations of simulated galaxies enables direct comparison to observational data, and helps to characterize the systematics involved in the transformation between intrinsic and observed quantities (see, e.g., Hayward and Smith 2015; Guidi et al. 2015).

Extinction by dust grains residing in the interstellar medium (ISM) can substantially influence the flux detected from a galaxy in the UV and optical wavelength ranges. It is very hard to estimate the dust mass in a galaxy based solely on the information at these wavelengths, and thus it is difficult to account accurately for the dust obscuration effect (e.g., Disney et al. 1989; Byun et al. 1994). To alleviate this limitation, one can turn to the far-infrared (FIR) to submm wavelength range. In this window, the continuum spectra of star-forming galaxies are dominated by thermal emission from dust grains that reprocess the UV/optical radiation, providing an independent and more direct measurement of the amount of dust in a galaxy. This additional information is especially useful for constraining the dust modeling of numerically simulated galaxies that have no explicit dust component. On the other hand, accurately predicting dust emission from a simulated galaxy requires solving a nontrivial 3D radiative transfer problem (see, e.g., Whitney 2011; Steinacker et al. 2013).

In this work we concentrate on the FIR and dust-related properties of the present-day galaxies produced by the EAGLE simulations (Schaye et al. 2015; Crain et al. 2015). EAGLE is a suite of hydrodynamical simulations of the formation of galaxies in cosmologically representative volumes, with sub-grid models for radiative cooling, star formation, stellar mass loss, and feedback from stars and accreting black holes. The sub-grid physics recipes are calibrated to reproduce the present-day galaxy stellar mass function and galaxy sizes, and show good agreement with many observables not considered in the calibration, including present-day specific star-formation rates, passive fractions, the Tully-Fisher relation (Schaye et al. 2015), and the neutral gas content (Bahé et al. 2016). The simulations also track the observed evolution of the galaxy stellar mass function out to redshift $z = 7$ (Furlong et al. 2015) and reproduce the observed optical colors for galaxies in the Local Universe (Trayford et al. 2015; Trayford et al. in prep).

We use the *Herschel* Reference Survey (Boselli et al. 2010) (HRS), a volume-limited sample of about 300 'normal' galaxies in the Local Universe, as a reference for observed dust properties. We select a set of redshift-zero EAGLE galaxies that matches

Table 6.1: We use the redshift-zero snapshots of the three EAGLE simulations listed in this table. We refer to them through the labels in the first column. The second column shows the corresponding full EAGLE simulation name as defined in Tables 2 and 3 in Schaye et al. (2015). The remaining columns list the simulation's co-moving box size, the initial number of baryonic particles, the initial baryonic particle mass, and the maximum proper gravitational softening length (i.e. at redshift zero).

Label	EAGLE name	L (cMpc)	N	m_g ($M_\odot$)	ϵ_{prop} (kpc)
Ref100	Ref-L100N1504	100	1504^3	1.81×10^6	0.70
Recal25	Recal-L025N0752	25	752^3	2.26×10^5	0.35
Ref25	Ref-L025N0752	25	752^3	2.26×10^5	0.35

the *K*-band luminosity distribution of the HRS galaxies, and we use the 3D dust radiative transfer code SKIRT (Baes et al. 2011; Camps and Baes 2015) to calculate observable properties for these galaxies from UV to submm wavelengths. We compare the stellar mass, dust mass, and star-formation rate derived from our mock observations through standard tracers with the intrinsic EAGLE values, and we compare the EAGLE dust scaling relations with those observed for HRS galaxies presented by Boselli et al. (2012) and Cortese et al. (2012). Finally, we investigate the effect of varying dust-related parameters in our post-processing procedure. This allows us to constrain these parameters, thus informing the calculation of dust attenuation for EAGLE galaxies in the UV and optical domain by Trayford et al. (in prep).

In Sect. 6.2 we provide some background on the EAGLE simulations and we describe how the EAGLE results were exported to and post-processed by SKIRT. In Sect. 6.3 we present and discuss the results of our analysis, and in Sect. 6.4 we summarize and conclude.

6.2   Methods

6.2.1   The EAGLE simulations

The Evolution and Assembly of GaLaxies and their Environments (EAGLE) project (Schaye et al. 2015; Crain et al. 2015) is comprised of a suite of smoothed particle hydrodynamics (SPH) simulations that follow the formation of galaxies and large-scale structure in cosmologically representative volumes of a standard Λ cold dark matter universe. EAGLE uses the hydrodynamics code GADGET (first described by Springel 2005), but employs an improved hydrodynamics scheme, referred to as ANARCHY, described by Schaye et al. (2015) and Schaller et al. (2015). The sub-grid models used in EAGLE are based on those developed for OWLS (Schaye et al. 2010). They are described in detail in Schaye et al. (2015) and summarized very briefly below.

Hydrogen reionization is modelled by turning on the time-dependent, spatially uniform UV/X-ray background from Haardt and Madau (2001) at redshift $z = 11.5$.

Radiative cooling and photo-heating are implemented element by element following Wiersma et al. (2009a), including all 11 elements that they found to dominate the radiative rates. Star formation follows Schaye and Dalla Vecchia (2008), but with the metallicity-dependent density threshold of Schaye (2004). Stellar mass-loss and chemical enrichment is based on Wiersma et al. (2009b) and tracks the elements H, He, C, N, O, Ne, Mg, Si, and Fe individually, while fixed abundance ratios relative to Si are assumed for Ca and S. Energetic feedback from star formation uses a stochastic thermal feedback scheme following Dalla Vecchia and Schaye (2012), with a variable efficiency depending on local gas density and metallicity. A super-massive black hole seed is placed at the center of every halo above a threshold mass (Springel et al. 2005) and is allowed to grow through gas accretion and mergers (Rosas-Guevara et al. 2015; Schaye et al. 2015). Feedback from these accreting black holes quenches star formation in massive galaxies, shapes the gas profiles in the inner parts of their host halos, and regulates the growth of the black holes themselves.

A drawback for the purpose of this work is that the EAGLE simulations do not model the cold gas phase in the ISM (see Sect. 4.3 of Schaye et al. 2015). To limit the pressure of star-forming gas particles, the EAGLE simulations impose a temperature floor, $T_{\rm eos}(\rho)$, as a function of the local gas density, ρ, corresponding to the polytropic equation of state $\rho\, T_{\rm eos} \propto P_{\rm eos} \propto \rho^{4/3}$ (Schaye and Dalla Vecchia 2008). As a consequence, there are no resolved molecular clouds. Instead, the simulated ISM consists of smoothly distributed, warm gas. We address this issue to some extent by employing a separate sub-grid model for star-forming regions in our post-processing procedure (see Sect. 6.2.3.4), and by assigning dust to star-forming gas particles regardless of their imposed, unphysical temperature (see Sect. 6.2.3.3). It remains important, however, to keep this limitation in mind when interpreting our results.

To enable numerical convergence studies, the EAGLE suite includes simulations with varying spatial resolution and simulation volume. In this work, we use the redshift-zero snapshots of the three EAGLE simulations listed in Table 6.1. The sub-grid prescriptions in the EAGLE reference simulation ('Ref100' in Table 6.1) are calibrated to reproduce the present-day galaxy stellar mass function. One of the higher-resolution simulations ('Ref25' in Table 6.1) employs the same sub-grid parameter values, i.e. calibrated for the resolution of the Ref100 simulation. For the other simulation ('Recal25' in Table 6.1), the sub-grid prescriptions have been re-calibrated to compensate for the effects of the increased numerical resolution. This approach allows investigating the 'weak' and 'strong' convergence properties of the simulations, as explained in Schaye et al. (2015).

The public database presented by McAlpine et al. (2016) lists a wide range of properties for the galaxies in the EAGLE simulations, including intrinsic quantities obtained by integrating over particle properties, luminosities in various optical and near-infrared bands (ignoring extinction by dust), and mock optical thumbnail images. When referring to a specific galaxy in this work, we specify the unique identifier ('GalaxyID') associated with that galaxy in the public EAGLE database.
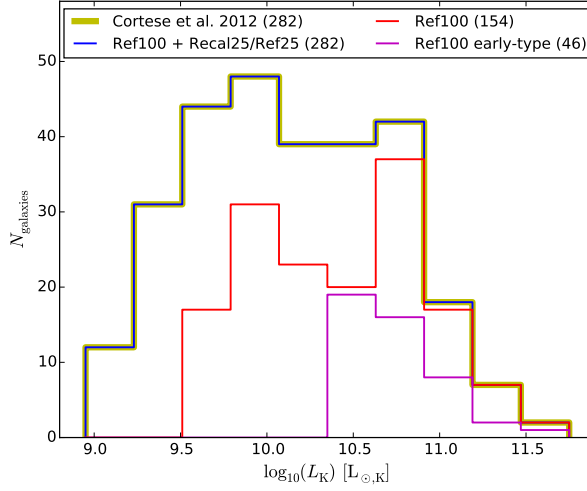
Figure 6.1: The *K*-band luminosity distribution of the galaxies in the Cortese et al. (2012) HRS sub-sample (dark yellow) and of the EAGLE galaxies selected for this work (blue) to match that sample. The curves are identical for both of the sets $\mathcal{C}$ and $\mathcal{F}$ listed in Table 6.2. The distribution of the galaxies taken from the standard-resolution snapshot Ref100 is shown in red; the remainder of the galaxies are taken from one of the higher-resolution snapshots, i.e. either Recal25 or Ref25. The distribution of the early-type galaxies, which are all taken from the Ref100 snapshot, is shown in magenta.

6.2.2   Galaxy selections

6.2.2.1   The HRS galaxies

In Sect. 6.3 we compare the dust-related properties of EAGLE galaxies with the observed properties of the galaxies in the *Herschel* Reference Survey (HRS, Boselli et al. 2010), and more specifically the subset presented by Cortese et al. (2012).

The HRS consists of a volume-limited sample ($15 \leqslant D \leqslant 25$ Mpc) including late-type galaxies with 2MASS (Skrutskie et al. 2006) *K*-band magnitude $K_S \leqslant 12$ mag and early-type galaxies with $K_S \leqslant 8.7$ mag. The total sample consists of 322 galaxies (260 late- and 62 early-type galaxies). As argued by Boselli et al. (2010), this sample is representative of the Local Universe and it spans different density regimes from isolated galaxies to the center of the Virgo cluster.

The HRS sub-sample analyzed by Cortese et al. (2012) includes only those galaxies for which *Herschel* as well as HI, NUV and SDSS observations are available, i.e. a total of 282 galaxies (234 late- and 48 early-type galaxies). As argued by Cortese et al. (2012), the sub-sample is representative of the full HRS sample, and it is thus representative of the local galaxy population as well.

According to Hughes et al. (2013) and Viaene et al. (2016), only 5 to 8 per cent of the HRS galaxies potentially host an active galactic nucleus (AGN), depending on the criteria used. Furthermore, Viaene et al. (2016) argue that the dust attenuation properties of the potential AGN hosts (and thus their FIR emission) do not differ

Table 6.2: Characteristics of the two sets of EAGLE galaxies for which we present results in this work. The first two columns show a symbol to identify the set and a mnemonic for the origin of this symbol. Subsequent columns list the total number of galaxies in the set and the number of galaxies extracted from each of the EAGLE snapshots used in this work (see Table 6.1). The final column shows the number of early-type galaxies in each set.

Set	Mmenonic	Total	Ref100	Recal25	Ref25	Early-type
$\mathcal{C}$	Re**C**al25	282	154	128	–	46
$\mathcal{F}$	Re**F**25	282	154	–	128	46

fundamentally from those of the other galaxies in the sample. Consequently, we do not exclude or single out these galaxies.

6.2.2.2   Selecting EAGLE galaxies

To enable a proper comparison between our mock observations and the HRS data, we construct a random sample of 282 present-day EAGLE galaxies mimicking the selection criteria described for HRS in Sect. 6.2.2.1, based on the intrinsic galaxy properties provided in the public EAGLE database (McAlpine et al. 2016). We first restrict our sample to galaxies with a minimum stellar mass of $10^{9.4}\,M_\odot$ for galaxies drawn from the Ref100 snapshot, and $10^{8.5}\,M_\odot$ for galaxies drawn from the Recal25 and Ref25 snapshots. These mass cutoffs ensure a minimum numerical resolution of roughly 2000 stellar particles per galaxy, as can be seen from the initial particle masses in Table 6.1, taking into account the mass transfer due to feedback processes over the lifetime of the stellar populations represented by the particles. As a consequence, our selection favors the high-resolution snapshots Recal25 or Ref25 for galaxies at the lower end of the mass range.

We then use the galaxy-type-dependent *K*-band selection criteria described in Sect. 6.2.2.1, assuming that all EAGLE galaxies are placed at a distance of 20 Mpc (the median distance of the HRS sample). We employ the specific star-formation rate $\dot{M}_* / M_*$ (sSFR) as a simple proxy for galaxy type, considering galaxies with sSFR $< 10^{-11}\,\mathrm{yr}^{-1}$ to be early-type (see, e.g., Fig. 8 in the review by Kennicutt and Evans 2012). Finally, we randomly reject galaxies until the sample matches the *K*-band luminosity distribution of the HRS sub-sample studied by Cortese et al. (2012), as shown in Fig. 6.1.

In fact, we construct two sets of EAGLE galaxies, named $\mathcal{C}$ and $\mathcal{F}$, that each match these criteria. Table 6.2 and Fig. 6.1 illustrate the make-up of these sets. Both sets contain *the same* collection of 154 galaxies drawn from the Ref100 snapshot, including 46 early-type galaxies. In addition, set $\mathcal{C}$ includes 128 galaxies drawn from the Recal25 snapshot, and set $\mathcal{F}$ likewise includes 128 galaxies drawn from the Ref25 snapshot. These two additional subsets have an identical *K*-band luminosity distribution, and contain no early-type galaxies.

Our analysis in Sect. 6.3 is mostly based on set $\mathcal{C}$. However, we evaluate the effects of the recalibration and numerical resolution of the EAGLE simulations by also
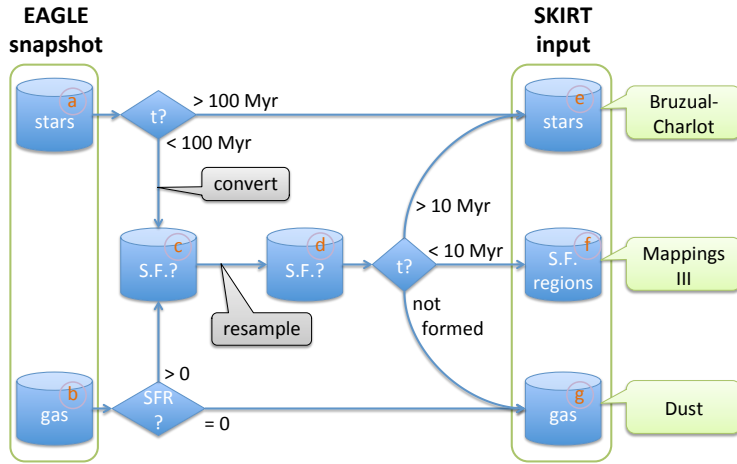
Figure 6.2: Schematic overview of the procedure used for preparing EAGLE galaxies for SKIRT. See the text in Sect. 6.2.3 for more details.

investigating some of the key results for set $\mathcal{F}$.

6.2.3 Preparing EAGLE galaxies for SKIRT

6.2.3.1 Extracting galaxies from an EAGLE snapshot

Trayford et al. (in prep) present a procedure for modeling EAGLE galaxies from optical to near-infrared wavelengths (0.28 − 2.5 μm) using SKIRT, generating spectra, broad-band photometry, line indices, and multi-band images for a large population of galaxies at redshift $z = 0.1$. We follow the same procedure, paying attention to dust emission and producing spectra and photometry over a much broader wavelength range (0.02 − 2000 μm).

Figure 6.2 illustrates the overall process of extracting the data for a galaxy from the EAGLE snapshot and preparing them for post-processing in SKIRT. For our purposes, a galaxy in an EAGLE snapshot is defined as a gravitationally bound substructure in a halo of dark and baryonic matter represented by particles. These structures are identified by the friends-of-friends and SUBFIND (Springel et al. 2001; Dolag et al. 2009) algorithms, which are run on the output of the EAGLE simulations. To study a particular galaxy, we extract the corresponding sets of star particles and gas particles (items *a* and *b* in Fig. 6.2). Following the convention used by Schaye et al. (2015), any particles outside a spherical aperture with radius of 30 kpc are ignored. The origin of the coordinate system is positioned at the galaxy's stellar center of mass. Unless noted otherwise, we retain the galaxy's original orientation, resulting in a 'random' viewing angle. In those few cases where we study the results for specific viewing angles, the face-on view looks down from the positive net stellar angular momentum vector of the galaxy, and the edge-on view observes from an arbitrary direction perpendicular to this vector.

6.2.3.2   Re-sampling star-forming regions

Star formation in EAGLE occurs stochastically: at each time step, a gas particle has a certain probability of being wholly converted to a star particle. Because individual particle masses are rather high (of the order of 10^6 M$_\odot$ for the reference simulation, see Table 6.1), a typical EAGLE galaxy contains only a small number of young star particles, unrealistically clumping all of the galaxy's young stars in a few point-like regions. This introduces sampling issues, which we alleviate by reprocessing the star-forming gas particles and the youngest star particles before feeding them into the SKIRT radiative transfer code, as illustrated in Fig. 6.2.

As a first step, we build a set of star-forming region candidates (item *c* in Fig. 6.2), including all star particles younger than 100 Myr, and all gas particles with a nonzero star-formation rate (SFR). All other particles, i.e. older star particles and non-star-forming gas particles, are transferred directly to the corresponding SKIRT input sets (items *e* and *g* in Fig. 6.2). The young star particles are converted back to star-forming gas particles. The SFR at the time of birth of these particles is calculated using the relation between pressure and SFR described in Sect. 4.3 of Schaye et al. (2015) and originally in Schaye and Dalla Vecchia (2008), which is based on the empirical Kennicutt-Schmidt law (Kennicutt 1998).

In the second step, the star-forming region candidates are re-sampled into a number of sub-particles (item *d* in Fig. 6.2) with lower masses drawn randomly from the power-law mass distribution function,

$$\frac{\mathrm{d}N}{\mathrm{d}M} \propto M^{-1.8} \quad \text{with } M \in [700, 10^6]\,\mathrm{M}_\odot. \tag{6.1}$$

This distribution of masses is inspired by observations of molecular clouds in the Milky Way reported by Heyer et al. (2001) and reviewed in Sect. 2.5 of Kennicutt and Evans (2012). Once a sufficient number of sub-particles have been generated to approximately represent the parent particle's mass, the sub-particle masses are proportionally adjusted to ensure exact mass conservation. The resulting sub-particles are assigned a formation time sampled randomly to represent their parent's SFR and mass, assuming a constant SFR over a 100 Myr lifetime. The sub-particles that formed more than 10 Myr ago are recast as star particles (item *e*); those that have not yet formed are recast as gas particles (item *g*); and those that formed less than 10 Myr ago are placed into a new SKIRT input set defining star-forming regions (item *f*).

Finally, the smoothing lengths and positions of the star-forming sub-particles are adjusted to match our post-processing assumptions as explained in Sect. 6.2.3.4.

6.2.3.3   Deriving the diffuse dust distribution

Table 6.3 offers an overview of the parameters defining the SKIRT radiative transfer model for each type of input particles, as discussed in the current and the following section. We derive a dust mass, M_{dust}, for each particle in SKIRT's 'gas' input set

Table 6.3: Input parameters of the SKIRT radiative transfer model for each type of EAGLE particle, in addition to the particle position. The procedure for deriving a dust distribution from the gas particles (item g in Fig. 6.2) is discussed in Sect. 6.2.3.3. The procedures for the particles representing stellar populations and star-forming regions (items e and f in Fig. 6.2) are discussed in Sect. 6.2.3.4.

Param	Description	Origin
	Dust distribution	
h	Smoothing length	Particle
M	Current gas mass	Particle
Z	Gas metallicity	Particle
T	Temperature of the gas	Particle
SFR	Star-formation rate of the gas	Particle
T_{max}	Highest temperature at which gas contains dust	Preset value
f_{dust}	Fraction of the metallic gas locked up in dust	Free param
	Young and evolved stars	
h	Smoothing length	Particle
M_{init}	Birth mass of the stellar population	Particle
Z	Metallicity of the stellar population	Particle
t	Age of the stellar population	Particle
	Star-forming regions	
h	Smoothing length	Calculated
M	Mass of the HII region	Sampled
SFR	Star-formation rate of the HII region	Calculated
Z	Metallicity of the HII region	Parent particle
ρ	Gas density at the HII region's position	Parent particle
P	Pressure of the ambient ISM	Calculated
C	Compactness of the HII region	Calculated
f_{PDR}	Dust covering fraction of the PDR region	Free param

(item g in Fig. 6.2) according to

$$M_{dust} = \begin{cases} f_{dust}\, Z\, M & \text{if } T < T_{max} \text{ or SFR} > 0 \\ 0 & \text{otherwise,} \end{cases} \qquad (6.2)$$

where Z, M, T, and SFR are the metallicity (metal mass fraction)[21], current mass, temperature, and star-formation rate given by the gas particle's properties in the EAGLE snapshot, and f_{dust} and T_{max} are free parameters. The characterization of gas particles based on the conditions of Eq. (6.2) is illustrated in Fig. 6.3 for an EAGLE disc galaxy. The star-forming (blue) and cold (cyan) gas particles trace the spiral arms in the galactic disk, while the hot gas (red) is located in the outskirts, as expected.

In summary, Eq. (6.2) assumes that a constant fraction f_{dust} of the metallic gas is locked up in dust, as long as the gas is forming stars or the gas is colder than the

---

[21] We use the SPH smoothed metallicity rather than the particle metallicity; see Wiersma et al. (2009b) and Schaye et al. (2015) for more information.
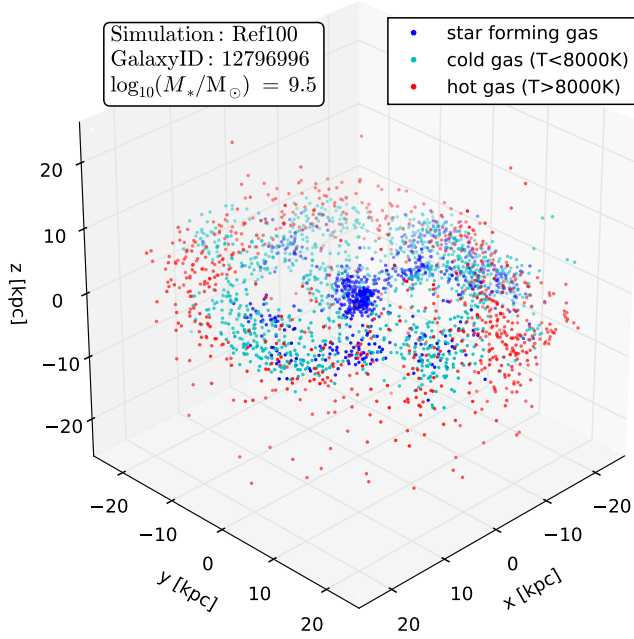
Figure 6.3: A projection of the gas particle positions in an EAGLE disc galaxy hand-picked for illustrative purposes. Our procedure allocates dust for star-forming gas (blue) and for cold gas (cyan). Hot gas (red) is deemed not to contain any dust.

cutoff temperature T_{max}. The assumption of a fixed dust-to-metal fraction f_{dust} is observed to be an appropriate approximation for a variety of environments (Dwek 1998; James et al. 2002; Brinchmann et al. 2013; Zafar and Watson 2013). We will vary this parameter as part of our analysis in Sect. 6.3.

The condition SFR > 0 captures the re-sampled gas particles that are eligible for star formation but were not actually converted into a star-forming region (the 'not formed' arrow between items d and g in Fig. 6.2). We need this condition because the EAGLE simulations assign a non-physical temperature to star-forming gas particles (see Sect. 6.2.1). However, by definition, the star-forming gas can be assumed to be sufficiently cold to form dust.

The temperature cutoff $T < T_{max}$ for the non-star-forming gas particles accounts for the fact that dust cannot form, or is rapidly destroyed, in hot gas (e.g., Guhathakurta and Draine 1989). We need to determine an appropriate temperature cutoff value. Unfortunately, since the EAGLE simulations do not model the cold gas phase in the ISM (see Sect. 6.2.1), we cannot properly constrain T_{max} using a physically motivated procedure. For our analysis in Sect. 6.3, we select a value of $T_{max} = 8000$ K, corresponding to the value of T_{eos} at $n_H = 0.1$ cm^{-1} used in the EAGLE simulations (Sect. 4.3 of Schaye et al. 2015).

6.2.3.4 Assigning SEDs to particles

Each particle in SKIRT's 'stars' input set (item e in Fig. 6.2) is assigned a stellar population SED from the Bruzual and Charlot (2003) family, using the birth mass, metallicity[21] and age given by the particle properties in the EAGLE snapshot (see Table 6.3 for an overview). We use the low resolution version of the Padova1994/Chabrier model, which is one of the two models recommended by Bruzual and Charlot (2003).

For the particles in the 'star-forming regions' input set (item f in Fig. 6.2) we follow the procedure described by Jonsson et al. (2010). Each particle is assigned an appropriate starburst SED from the MAPPINGS III family (Groves et al. 2008). These templates model both the HII region and the photodissociation region (PDR) surrounding the star-forming core, including the dust contained in those regions. We attempt to compensate for the additional dust mass assumed by the MAPPING III model by removing the equivalent amount of dust from the diffuse dust component. We now discuss this process in more detail.

The MAPPINGS III templates are parametrized by the SFR and the metallicity of the star-forming region, the pressure of the ambient ISM, the HII region compactness, and the covering fraction of the associated PDR (see Table 6.3 for an overview). The SFR is determined from the mass assigned to the star-forming particle (as discussed in Sect. 6.2.3.3), assuming a constant SFR during the HII region's lifetime of 10 Myr (following Groves et al. 2008). The metallicity, Z, is taken directly from the particle properties[21] in the EAGLE snapshot. The ambient pressure of the ISM, P, is calculated from the particle's density, ρ, using the polytropic equation of state imposed on star-forming particles (see Sect. 6.2.1). The HII region compactness, C, is designed to reflect the dust temperature distribution in the HII region (time-averaged over its lifetime), so that it predominantly controls the form of the FIR continuum dust emission. In our procedure the value of this parameter is derived from the ambient pressure, P, and our assigned particle mass, M (see Eq. 6.1), using Eq. (13) of Groves et al. (2008), i.e.,

$$\log_{10} C = \frac{3}{5} \log_{10} \left(\frac{M}{M_\odot} \right) + \frac{2}{5} \log_{10} \left(\frac{P/k_B}{cm^{-3} \, K} \right), \tag{6.3}$$

where k_B is the Boltzmann constant. Finally, the parameter f_{PDR} is defined as the time-averaged dust covering fraction of the photodissociation region (PDR) surrounding the star-forming core over the HII region's lifetime. Starbursts in which the PDR's dust entirely envelops the HII region have $f_{PDR} = 1$, while uncovered HII region complexes have $f_{PDR} = 0$. The covering fraction is treated as a free parameter, which we will vary as part of our analysis in Sect. 6.3.

Following Jonsson et al. (2010), we consider the region represented by the MAPPINGS III templates (including PDR and HII region) to be ten times as massive as the star-forming core represented by the particle. To determine the spatial extent of the region's emission, we assume that the region's center has the same density as the local

ambient ISM. For the cubic spline kernel employed in SKIRT, this leads to the easily inverted relation, $10M = (\pi/8)\rho h^3$, between the HII region mass, M, the ambient density, ρ, and the particle smoothing length, h. We also randomly shift the positions of the star-forming sub-particles within the smoothing sphere of the parent particle (see Sect. 6.2.3.2) to avoid overlap between the modelled regions.

As indicated above, the MAPPINGS III templates model the dust residing in the PDR region in addition to the core HII region itself. To avoid double counting, we subtract this PDR dust from the diffuse dust distribution derived as discussed in Sect. 6.2.3.3. We insert a 'ghost' gas particle with negative mass in the SKIRT gas input set (see Sect. 6.2.3.3) for each star-forming particle. The ghost particle receives the (negative) mass of the corresponding PDR region, i.e. ten times the mass of the star-forming particle. When sampling the gas (or dust) density field, SKIRT combines the negative ghost densities with the positive densities defined by the other particles, clipping the total density to zero if needed. To lower the probability of this occurring, we artificially increase the smoothing length of the ghost particle by a factor of three. According to our tests, this sufficiently alleviates the issue without otherwise affecting the results.

6.2.4   Radiative transfer on EAGLE galaxies

This section describes the SKIRT configuration used to perform the radiative transfer simulations on the EAGLE galaxies.

6.2.4.1   Dust grid

The SKIRT radiative transfer procedure requires the dust density distribution of the system under study to be discretized over a dust grid. Within each grid cell, the dust density and all other physical quantities, such as the radiation field, are assumed to be constant. SKIRT implements a performance-optimized mechanism to calculate the dust mass in each grid cell from the smoothed particles defining a galaxy. The particles are interpolated using a scaled and truncated Gaussian kernel designed to approximate a finite-support cubic spline kernel (Altay and Theuns 2013; Baes and Camps 2015).

Here, we use an adaptive, hierarchical Cartesian grid that encloses the 30 kpc aperture considered for each galaxy (see Sect. 6.2.3.1). Specifically, we use an octree grid (Saftly et al. 2013) that automatically subdivides cells until each cell contains less than a fraction $\delta_{\mathrm{max}} = 3 \times 10^{-6}$ of the total dust mass in the model, with a maximum of 10 subdivision levels (see Fig. 6.4). The smallest possible cell is thus about 60 pc on a side, which offers 5-10 times better resolution than the typical gravitational softening length in the EAGLE simulations (see Table 6.1).

Figure 6.5 provides some relevant statistics on the discretization of the diffuse dust density for the EAGLE galaxies analyzed in this work. The leftmost panel shows that about half of the galaxies in our selection have more than 3000 gas particles that
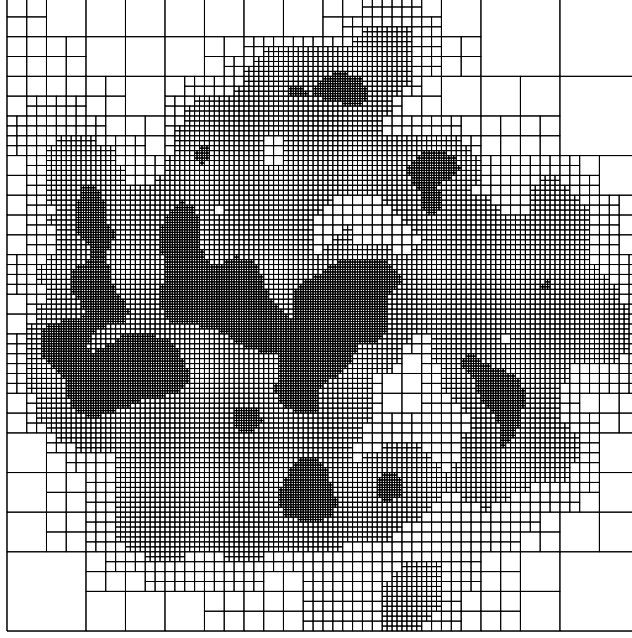
Figure 6.4: A cut along the galactic plane through an octree dust grid constructed for the EAGLE galaxy shown in Fig. 6.3. The darker areas trace regions of higher dust density (the grid has smaller dust cells and thus more cell boundaries). For presentation purposes, the illustrated grid uses fewer refinement levels and covers a smaller aperture than the grid actually used by SKIRT in this work.
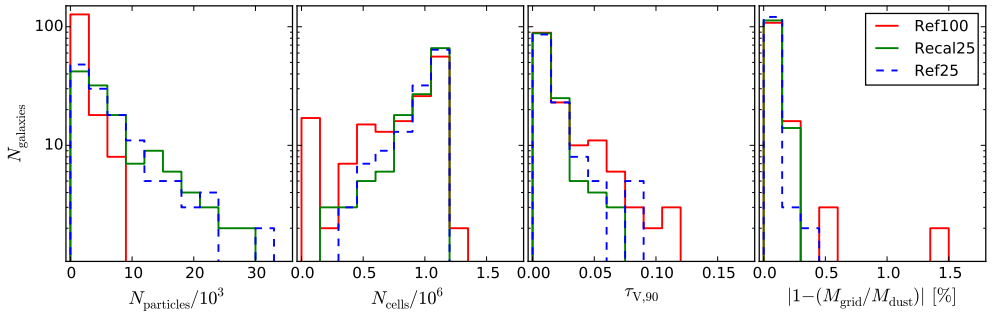
Figure 6.5: Distribution of SKIRT dust discretization properties for the EAGLE galaxies analyzed in this work (see Table 6.2), processed with $f_{dust} = 0.3$. From left to right: the number of gas particles that include dust (i.e. cold or star-forming gas particles); the number of cells in the dust grid constructed by SKIRT; the 90% percentile V-band optical depth of the cells in the dust grid; the discretization error on the total dust mass (i.e. difference between the dust mass in the grid and in the incoming particles).

include dust, i.e. particles representing cold or star-forming gas (see Sect. 6.2.3.3), which is sufficient to spatially resolve the diffuse dust distribution. Further analysis (not shown) indicates that about 100 of our galaxies have less than 100 'dusty' gas particles, however this includes most of the early-type galaxies, which do not contain much dust anyway.

The two middle panels of Figure 6.5 show properties of the dust grids constructed by SKIRT to perform radiative transfer on our EAGLE galaxies. Most dust grids have more than 250 000 cells, which is more than sufficient to resolve the imported smoothed particles. Also, over 90% of the dust cells in each grid have a V-band optical depth of less than 0.12 (and most have much lower optical depth), indicating that the grid properly resolves even the densest regions in the dust mass.

The rightmost panel of Figure 6.5 shows the difference between the dust mass obtained by summing over all cells in the dust grid, and the dust mass obtained by summing over the incoming particles. For most galaxies, this dust discretization error is limited to less than a third of a per cent, with some outliers of up to 1.5 per cent. While part of this error is caused by grid resolution limitations, further analysis (not shown) indicates that the larger discrepancies in the outliers are caused by the negative dust masses which are introduced to compensate for the dust modeled by star-forming regions (see Sect. 6.2.3.4). Specifically, the imported dust density becomes negative in some areas, and is then clipped to zero when building the dust grid.

6.2.4.2   Dust model

To represent the dust in the EAGLE galaxies, we use a dust mixture of non-composite graphite and silicate grains and neutral and ionized polycyclic aromatic hydrocarbon (PAH) molecules, designed so that the global dust properties accurately reproduce the extinction, emission and abundance constraints of the Milky Way. The optical properties are taken from Bruce Draine's website[22] (Draine and Lee 1984; Laor and Draine 1993; Li and Draine 2001). The calorimetric properties follow the prescription of Draine and Li (2001). The grain size distributions for each population are taken from Zubko et al. (2004).

The dust emission spectrum is calculated for each dust cell based on the stellar radiation absorbed by the dust in that cell. The calculation includes the effects of stochastically heated grains, i.e. dust grains and PAH molecules that are not in local thermal equilibrium with the radiation field, using the scheme described by Camps et al. (2015). To facilitate this calculation, SKIRT discretizes the size range of the dust grains into several size bins, for each type of grain material separately. For this work, following the recommendations of Camps et al. (2015), SKIRT uses 15 size bins for each of the graphite and silicate components, and 10 size bins for each of the neutral and ionized PAHs.

---

[22] http://www.astro.princeton.edu/~draine/dust/dust.diel.html
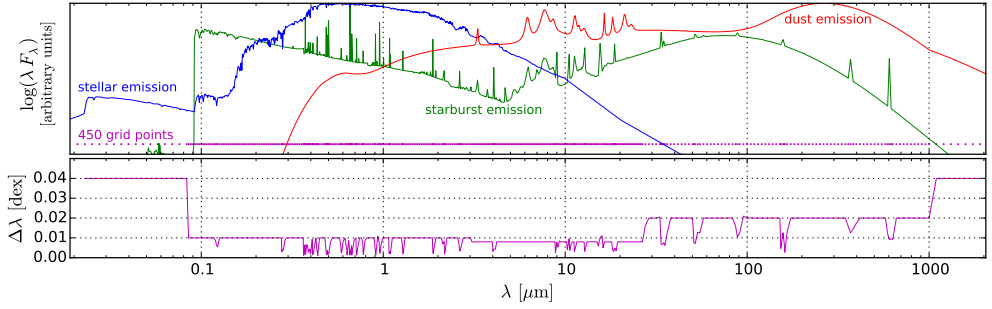
Figure 6.6: Characteristics of the wavelength grid used in all SKIRT calculations for this work. The three curves in the top panel illustrate typical SEDs for an evolved stellar population (blue), a star-forming region including young stellar objects and dust (green), and stochastically heated diffuse dust (red), plotted on an arbitrary logarithmic scale. The dots (magenta) represent the wavelength grid points. The curve in the bottom panel (magenta) indicates the distance between successive wavelength points on a logarithmic scale.

6.2.4.3 Wavelength discretization

The SKIRT code employs a single wavelength grid for all calculations. The input SEDs and dust properties are sampled on this grid, photon packages are given wavelengths corresponding to the grid points, dust absorption and re-emission are calculated for the wavelength bins defined by the grid, and the output fluxes are recorded on the same grid.

The wavelength grid used in all SKIRT calculations for this work is illustrated in Fig. 6.6. It resolves the relevant features in the input SEDs (see Sect. 6.2.3.4) and in the emission spectrum of the dust population (see Sect. 6.2.4.2). The grid has 450 wavelength points from 0.02 to 2000 μm laid out on a logarithmic scale. The bin widths are 0.04 dex in the outer wavelength ranges where fluxes are low, 0.02 dex in the dust emission continuum, 0.01 dex in the optical range, and under 0.01 dex in the PAH emission range and for specific emission or absorption features in the employed input spectra.

To further inspect this discretization, we compare band-integrated fluxes (see Appendix 6.5) calculated on our default 450-point wavelength grid with those calculated on a high-resolution grid with 20 000 points. For this purpose, we select a typical SED for a stellar population, one for a star-forming region, and one for stochastically heated dust (see Fig. 6.6). We calculate the fluxes for these SEDs in a set of bands essentially covering the complete wavelength range, using Eqs. (6.9) or (6.11). The results calculated on our 450-point wavelength grid are accurate to within 0.1 mag for all bands, and often much better. The results for the bands used in this work are listed in Table 6.4.

Table 6.4: Evaluation of the wavelength grid and numerical convergence for the SKIRT simulations in this work. The first two columns list the name of the instrument for which mock broadband fluxes are calculated and the corresponding pivot wavelength according to Eqs. (6.10) or (6.12). The next three columns show the differences between the magnitude calculated on a high-resolution wavelength grid and on our default wavelength grid, for the three SEDs shown in Fig. 6.6. The last column shows the maximum magnitude differences for SEDs calculated from SKIRT simulations with different dust grid resolutions and numbers of photons. The dashed line separates photon counters (top) and bolometers (bottom).

| Band | λ_{pivot} (μm) | ——— Wavelength grid ——— | | | Dust grid & photons (Δ mag) |
		SF region (Δ mag)	Stellar (Δ mag)	Dust (Δ mag)	
GALEX FUV	0.1535	0.002	0.041	–	0.004
GALEX NUV	0.2301	0.003	0.007	–	0.003
SDSS u	0.3557	0.047	0.027	0.001	0.002
SDSS g	0.4702	0.087	0.005	0.001	0.001
SDSS r	0.6176	0.054	0.003	0.001	0.002
SDSS i	0.7490	0.011	0.002	0.002	0.002
SDSS z	0.8947	0.016	0.005	0.001	0.001
2MASS J	1.239	0.016	0.014	0.011	0.001
2MASS H	1.649	0.011	0.012	0.003	0.002
2MASS K_S	2.164	0.003	0.013	0.002	0.001
WISE W1	3.390	0.021	0.003	0.018	0.001
WISE W2	4.641	0.005	0.005	0.001	0.001
WISE W3	12.57	0.001	0.001	0.001	0.003
WISE W4	22.31	0.001	0.001	0.009	0.003
MIPS 24	23.59	0.001	0.001	0.008	0.003
MIPS 70	70.89	0.001	0.001	0.001	0.003
MIPS 160	155.4	0.001	0.001	0.001	0.004
PACS 70	70.77	0.001	0.001	0.001	0.002
PACS 100	100.8	0.001	0.001	0.001	0.003
PACS 160	161.9	0.001	0.008	0.001	0.004
SPIRE 250 ext	252.5	0.001	–	0.001	0.034
SPIRE 350 ext	354.3	0.001	–	0.001	0.034
SPIRE 500 ext	515.4	0.026	–	0.001	0.035

6.2.4.4   Photon packages

The SKIRT radiative transfer simulation proceeds in two phases. In the first phase, SKIRT launches photon packages randomly originating at the stars and the star-forming regions, and traces these packages through the dusty medium. The simulation loop accounts for the effects of scattering off dust grains, and keeps track of the radiation absorbed in each dust cell. After this phase completes, the code calculates the emission spectrum of the dust population in each dust cell based on the established radiation field, taking into account the probabilistic thermal emission of small grains and PAH molecules (Camps et al. 2015). In the second phase, SKIRT launches photon packages originating from the dust distribution, corresponding to the calculated emission spectra, and traces these packages through the dusty medium as well.

Table 6.5: Properties of the *Herschel* SPIRE 250/350/500 instruments used in our mock flux derivation. The beam FWHM and beam area are taken from Ciesla et al. (2012). For the flux limit, we use the confusion noise level from Nguyen et al. (2010).

	Units	$250\,\mu m$	$350\,\mu m$	$500\,\mu m$
Beam FWHM	arcsec	18.2	24.5	36.0
Beam area	$arcsec^2$	423	751	1587
Flux limit	mJy/beam	5.8	6.3	6.8

For this work, we instruct SKIRT to ignore dust heating by photon packages emitted from the dust, substantially reducing the calculation time. This is justified because the body of dust in a normal galaxy is essentially transparent to infrared radiation. We verified this assumption for our EAGLE sample by comparing the simulation results with and without dust self-heating for the highest dust-mass galaxies. Finally, we configure SKIRT to launch 5×10^5 photon packages for each of the 450 points in the wavelength grid during each of the two phases. Thus the SKIRT simulation for each EAGLE galaxy traces 4.5×10^8 photon packages. In Sect. 6.2.4.6 we confirm that this choice is appropriate.

6.2.4.5   Mock fluxes

Mock detectors are placed along two of the coordinate axes at a fixed distance of 20 Mpc from the model, using parallel projection. If the model has been properly rotated (see Sect. 6.2.3.1), this results in a face-on and an edge-on view of the galaxy. The selected detector distance matches the median distance of the HRS galaxies; see Sect. 6.2.2.2. Each detector records an integral field data cube (a 400×400 pixel frame at each of the wavelength grid points) in addition to the spatially integrated fluxes at each wavelength grid point. From these results, we produce band-integrated fluxes and absolute magnitudes corresponding to the following filters (see also Table 6.4): GALEX FUV/NUV (Morrissey et al. 2007); SDSS *ugriz* (Doi et al. 2010); 2MASS *JHK* (Cohen et al. 2003); WISE W1/W2/W3/W4 (Wright et al. 2010); *Spitzer* MIPS 24/70/160 (Rieke et al. 2004); *Herschel* PACS 70/100/160 (Poglitsch et al. 2010); and *Herschel* SPIRE 250/350/500 for extended sources (Griffin et al. 2010).

To obtain the integrated fluxes, we convolve the simulated SED with the instrument's response curve. The precise procedure depends on whether the instrument counts photons or measures energy (bolometer); the formulae are summarized in Appendix 6.5. The GALEX, SDSS, 2MASS and WISE instruments are photon counters; the *Spitzer* MIPS and the *Herschel* PACS and SPIRE instruments are bolometers.

Because our analysis in Sect. 6.3 relies heavily on the *Herschel* SPIRE 250/350/500 fluxes, and because actual observations in these submm bands suffer from fairly severe observational limitations, we perform an additional procedure for these fluxes. Table 6.5 lists the relevant instrument properties, taken from Ciesla et al. (2012) and Nguyen et al. (2010). We first perform a convolution with the corresponding instrument response function for each of the 400×400 pixels in the recorded frames.

We then bin the pixels in the resulting frame to match the beam area of the instrument, and we convolve the rebinned frame (spatially) with a Gaussian filter scaled to the full width at half maximum (FWHM) of the instrument's beam. From this frame, we eliminate all pixels with a flux value below the sensitivity level of the instrument, and we finally sum over the remaining pixels to obtain the total flux.

6.2.4.6   Numerical convergence

A numerical convergence test can help ascertain that our discretization settings are appropriate. To this end, we perform the SKIRT simulations for galaxy set $\mathcal{C}$ (Table 6.2) using a higher-resolution dust grid and shooting more photon packages than for our default setup. Specifically, we set the maximum mass fraction per cell to $\delta_{\mathrm{max}} = 2 \times 10^{-6}$ rather than 3×10^{-6} (see Sect. 6.2.4.1), and we increase the number of photon packages launched per wavelength grid point to 10^6 from 5×10^5 (see Sect. 6.2.4.4). We then calculate the fluxes in the various bands used for this work according to the procedure described in Sect. 6.2.4.5, and we compare the results from the high-resolution simulation with those from the default setup. The rightmost column in Table 6.4 shows the absolute value of the resulting magnitude differences. The fluxes are accurate to within 0.05 mag for all bands, and even to within 0.005 mag for all but the four longest-wavelength bands. The somewhat larger errors for the *Herschel* SPIRE 250/350/500 bands are caused by our implementation of the observational limits in these bands (see Sect. 6.2.4.5), which heavily depends on the precise 2D distribution of the fluxes in the simulated images.

Overall, we conclude that the quality of the dust grid and the number of photons in our default setup are sufficient for our purposes.

6.3	Results and discussion

In the following subsections we present results for EAGLE galaxies that were post-processed according to Eq. (6.2) with $f_{\mathrm{dust}} = 0.3$ and $f_{\mathrm{PDR}} = 0.1$. We will further justify these parameter values in Sect. 6.3.5.

6.3.1   Intrinsic properties

Although our aim in this work is to evaluate mock observations of the EAGLE galaxies, it is instructive to briefly review some intrinsic properties, even if only to confirm that these values fall in the appropriate range. To this end, Figure 6.7 shows selected intrinsic properties of the EAGLE galaxies analyzed in this work, i.e. properties that can be calculated from the particles extracted from the snapshot without radiative transfer processing. Consistent with our selection criteria (Fig. 6.1), most high-mass and all early-type EAGLE galaxies are extracted from the Ref100 snapshot (red points). The remaining galaxies are extracted from the Recal25 snapshot (green points) or from
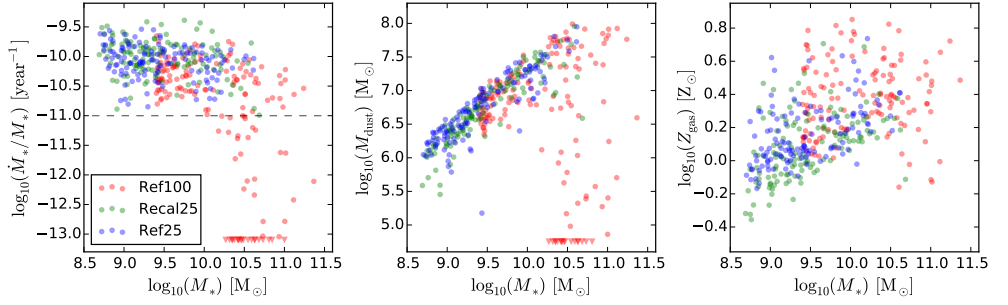
Figure 6.7: Selected intrinsic properties of the EAGLE galaxies analyzed in this work (see Table 6.2) plotted versus intrinsic stellar mass. Left panel: specific star-formation rate (sSFR); galaxies below the dashed line are deemed to be early-types; galaxies with sSFR below $10^{-13.1}$ are plotted as upper bounds at that value. Middle panel: dust mass, assuming a dust-to-metal fraction $f_{\mathrm{dust}} = 0.3$; galaxies with dust mass below $10^{4.75}\,\mathrm{M}_\odot$ are plotted as upper bounds at that value. Right panel: overall metallicity of the gas that includes dust, in units of $Z_\odot = 0.0127$; galaxies with zero dust mass are omitted.

the Ref25 snapshot (blue points) depending on the galaxy set under consideration (Table 6.2).

The leftmost panel of Fig. 6.7 plots specific star-formation rate (sSFR) versus stellar mass. Comparing this diagram to, e.g., Fig. 8 of Kennicutt and Evans (2012), we conclude that both sSFR and stellar mass values are in the expected range, and we can clearly recognize a blue cloud of star-forming galaxies above the horizontal dashed line. The red sequence of quiescent galaxies below the dashed line is less prominent because our selection disfavors these galaxy types to reflect the HRS sample (see Sects. 6.2.2.1 and 6.2.2.2).

The middle panel of Fig. 6.7 plots dust mass versus stellar mass. The dust mass is calculated by summing the result of Eq. 6.2 over all gas particles, using a dust-to-metal fraction $f_{\mathrm{dust}} = 0.3$. Comparing this figure to, e.g., Fig. 16 of Bourne et al. (2012), we conclude that these dust masses are within the expected range. The low dust masses for some of the high-stellar-mass (early-type) galaxies are also consistent with observations (di Serego Alighieri et al. 2013).

The rightmost panel of the same figure plots the metallicity of the gas that contains dust versus stellar mass. The galaxies in our sample have a fairly high metallicity compared to observations (Tremonti et al. 2004; Hughes et al. 2013; Zahid et al. 2014). For example, the metallicities of the HRS galaxies shown in Fig. 4 of Hughes et al. (2013) do not exceed $\log_{10}(Z/Z_\odot) = 0.2$, assuming $12 + \log_{10}(\mathrm{O/H})_\odot = 8.69$ (Allende Prieto et al. 2001). The high metallicities in our sample are, however, consistent with the mass-metallicity relation of the EAGLE galaxies reported in Fig. 13 of Schaye et al. (2015). It is noted there that the Ref100 EAGLE simulation systematically over-predicts metallicity in the stellar mass range $M_* < 10^{9.5}\,\mathrm{M}_\odot$. It is also evident from the right panel of Fig. 6.7 that galaxies in Ref25 tend to have higher metallicities than galaxies in Recal25, again consistent with the findings of Schaye et al. (2015). Because we use a constant dust-to-metal fraction, see Eq. (6.2), this leads to a slightly higher
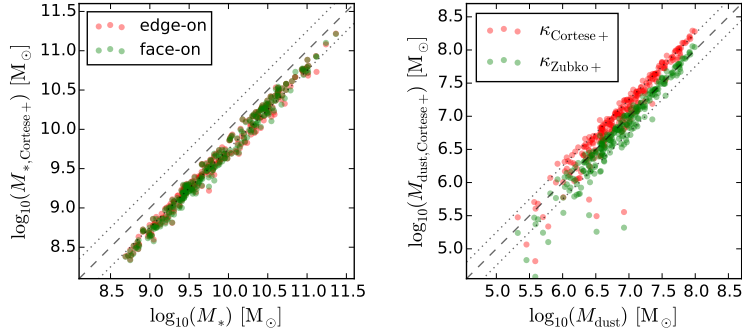
Figure 6.8: Comparison of stellar and dust mass derived from mock observations of the EAGLE galaxies in set $\mathcal{C}$ (Table 6.2) with the corresponding intrinsic properties. The EAGLE galaxies were post-processed using $f_{dust} = 0.3$ and $f_{PDR} = 0.1$. The dashed diagonal in each panel indicates the one-to-one relation; the dotted lines indicate ± 0.25 dex offsets. Left panel: stellar mass estimated through Eq. (6.4) following Cortese et al. (2012) and Zibetti et al. (2009) using edge-on (red) and face-on fluxes (green). Right panel: dust mass estimated through Eq. (6.5) following Cortese et al. (2012) using $\beta = 2$ and the two values of κ_{350} defined in the text; galaxies with dust mass below $10^{4.75}\,\mathrm{M_\odot}$ are omitted.

dust content for most Ref25 galaxies (middle panel of Fig. 6.7).

6.3.2 Inferred stellar and dust masses

We compare the stellar mass derived from mock observations of our EAGLE galaxies to the intrinsic stellar mass calculated by summing over all stellar particles. We mimic the procedure employed by Cortese et al. (2012), determining the 'mock' stellar mass M_* from the i-band luminosity L_i and the $g - i$ color through

$$\log_{10} \frac{M_*}{\mathrm{M_\odot}} = \log_{10} \frac{L_i}{L_{i,\odot}} + a + b \times (g - i), \tag{6.4}$$

with coefficients $a = -0.963$ and $b = 1.032$ taken from Table B1 of Zibetti et al. (2009). The result is shown in the left panel of Fig. 6.8. The mock observations of our EAGLE galaxies underestimate the stellar mass by about 0.25 dex, for both edge-on and face-on fluxes. The Zibetti et al. (2009) recipe assumes the Chabrier (2003) initial mass function (IMF), as do the EAGLE simulations (see Sect. 4.3 of Schaye et al. 2015) and the SED templates we assign to stellar particles (see our Sect. 6.2.3.4 and Sect. 2.3 of Bruzual and Charlot 2003), so there is no need to compensate for offsets between different IMFs in the model. However, the Zibetti et al. (2009) calibration of the stellar mass-to-light ratio relation was derived for resolved parts of galaxies. Several authors have proposed different values for the coefficients a and b (e.g., Gallazzi and Bell 2009; Taylor et al. 2010, 2011; Baldry et al. 2012), resulting in a systematic shift of up to 0.3 dex in the relation. In the following sections, we use the Zibetti et al. (2009) calibration because we will be confronting the mock observations with the results presented by Cortese et al. (2012).

We now compare the dust mass derived from mock observations with the dust mass calculated by summing over all gas particles according to Eq. (6.2). Following Cortese et al. (2012) and many other authors, the flux f_ν emitted by a modified black body at the frequency ν can be written as

$$f_\nu = \frac{M_{\mathrm{dust}}}{d^2}\,\kappa_\nu\,B_\nu(T_{\mathrm{dust}}) \qquad \mathrm{with} \quad \kappa_\nu = \kappa_{350}\left(\frac{\nu}{\nu_{350}}\right)^\beta, \qquad (6.5)$$

where M_{dust} is the dust mass, d is the distance, $B_\nu(T)$ is the Planck function, T_{dust} is the dust temperature, κ_ν is the dust mass absorption coefficient, assumed to depend on frequency through a power law with index β, and κ_{350} is the dust mass absorption coefficient at a wavelength of 350 μm. Cortese et al. (2012) use the values $\beta = 2$ and $\kappa_{350} = \kappa_{\mathrm{Cortese+}} = 0.192\ \mathrm{m^2\,kg^{-1}}$. The Zubko et al. (2004) dust model used in this work (see Sect. 6.2.4.2) has the same power-law index, $\beta = 2$. However, the absorption coefficient $\kappa_{350} = \kappa_{\mathrm{Zubko+}} = 0.330\ \mathrm{m^2\,kg^{-1}}$ differs substantially, causing a shift of 0.24 dex in the inferred dust mass.

Cortese et al. (2012) use the three *Herschel* SPIRE 250/350/500 fluxes to estimate the dust mass, employing a recipe presented in their Appendix B. The right panel of Fig. 6.8 plots the dust mass estimates calculated from SKIRT fluxes for our EAGLE galaxies according to this recipe, using $\beta = 2$ and the two values of κ_{350} defined in the previous paragraph. When using $\kappa_{350} = \kappa_{\mathrm{Cortese+}}$, the Cortese et al. (2012) recipe overestimates the dust mass. With the $\kappa_{350} = \kappa_{\mathrm{Zubko+}}$ appropriate for our dust model, however, the estimates are fairly accurate, although there is significant scatter in the low mass range. Because we will be confronting our mock observations with the results presented by Cortese et al. (2012), we use their dust mass recipe with $\kappa_{350} = \kappa_{\mathrm{Cortese+}}$ in the following sections.

6.3.3   Inferred star-formation rates

We compare in Fig. 6.9 three of the star-formation-rate (SFR) indicators listed in Table 1 of Kennicutt and Evans (2012), calculated for mock observations of our EAGLE galaxies, to the intrinsic SFR provided in the public EAGLE database (McAlpine et al. 2016). The leftmost panel of Fig. 6.9 shows the SFR based on the GALEX NUV flux (Hao et al. 2011; Murphy et al. 2011) using edge-on (red) and face-on fluxes (green). At these short wavelengths, the edge-on fluxes suffer significantly more from dust extinction than the face-on fluxes, especially in more active galaxies, and thus yield a correspondingly lower SFR. However, even the indicator based on face-on fluxes slightly underestimates the SFR for most galaxies. For a small number of outliers, mostly in the lower SFR regime, the indicator substantially overestimates the SFR. These outliers are passive galaxies with a low dust content (the edge-on and face-on fluxes are equal so there is little extinction), where the NUV radiation emitted by the evolved star population is interpreted as a sign of star formation by the indicator. This so-called UV-upturn is also found in observations (e.g., Brown et al. 1997, 2003).

The middle panel of Fig. 6.9 shows the SFR based on the integrated total infrared flux

Figure 6.9: Comparison of three of the star-formation-rate (SFR) indicators summarized in Table 1 of Kennicutt and Evans (2012)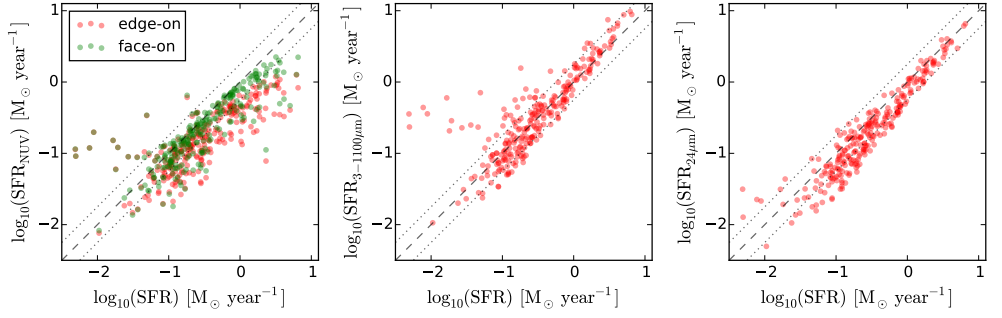, calculated for mock observations of the EAGLE galaxies in set $\mathcal{C}$ (Table 6.2), to the intrinsic SFR provided in the public EAGLE database (McAlpine et al. 2016). The EAGLE galaxies were post-processed using $f_{\text{dust}} = 0.3$ and $f_{\text{PDR}} = 0.1$. The dashed diagonal in each panel indicates the one-to-one relation; the dotted lines indicate ± 0.25 dex offsets. Left panel: SFR based on GALEX NUV flux (Hao et al. 2011; Murphy et al. 2011) using edge-on (red) and face-on fluxes (green). Middle panel: SFR based on integrated 3-1100 μm flux (Hao et al. 2011; Murphy et al. 2011). Right panel: SFR based on *Spitzer* MIPS 24 μm flux (Rieke et al. 2009). In all panels, galaxies with intrinsic or inferred SFR below $10^{-2.5}$ $M_\odot$ year^{-1} are omitted.

(3-1100 μm; Hao et al. 2011; Murphy et al. 2011). Because dust is mostly transparent to radiation at these wavelengths, the emission is isotropic and there is no need to compare edge-on and face-on fluxes. This indicator is fairly accurate, except for a number of outliers mostly in the lower SFR regime. In these cases, the emission from diffuse dust residing in the outskirts of those galaxies is interpreted by the indicator as a sign of star formation, while the dust is in fact being heated by an evolved star population. This phenomenon is also found in observations (Bendo et al. 2015).

The rightmost panel of Fig. 6.9 shows the SFR based on the *Spitzer* MIPS 24 μm flux (Rieke et al. 2009). Except at the lowest SFRs, this indicator consistently underestimates the intrinsic SFR of our galaxies, which can be understood as follows. The EAGLE simulations do not model the cold ISM phase, and the adjustments made by our post-processing procedure have limitations as well. For example, our model contains isotropically emitting star-forming regions that may not represent the strong variations in the radiation field near star-forming regions sufficiently or accurately. As a consequence, at least some fraction of the diffuse dust in the simulated galaxies is heated insufficiently, resulting in a 24 μm flux lower than observed. In addition, the flux in this wavelength range is very sensitive to the precise properties of the dust, and thus, the dust model used for the simulations (Fanciullo et al. 2015). For example, if the dust grain population contains a larger fraction of small grains, more grains will be stochastically heated to higher energy levels, shifting some of the dust emission to shorter wavelengths and into the 24 μm band, while the total infrared flux remains unchanged.
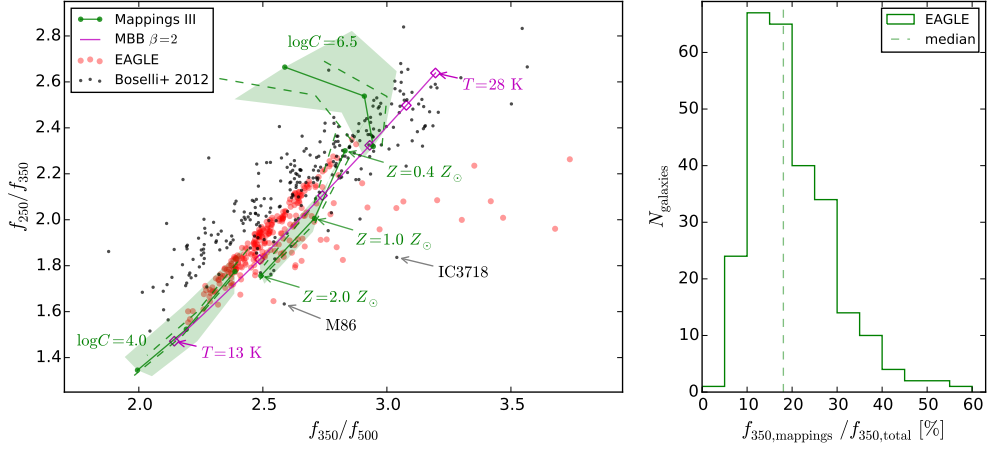
Figure 6.10: LEFT PANEL: *Herschel* SPIRE color-color relation f_{250}/f_{350} versus f_{350}/f_{500} for EAGLE galaxies (red points) from set $\mathcal{C}$ (Table 6.2) compared to HRS observations (black points) taken from Boselli et al. (2012) and Ciesla et al. (2012). The EAGLE galaxies were post-processed using $f_{dust} = 0.3$ and $f_{PDR} = 0.1$. Galaxies for which one or more SPIRE fluxes are below the detection limit are omitted. The magenta curve traces a modified black body (MBB) with $\beta = 2$ for temperatures ranging from 13 K to 28 K; the diamonds are spaced by 3 K. The green annotations indicate the flux ratios of the MAPPINGS III templates used to model star-forming regions for a range of input parameter values (see Table 6.3). The solid lines show ratios for fixed $f_{PDR} = 0.1$ and for the extreme compactness values supported by the templates ($\log_{10} C = 4.0, 6.5$) plus an intermediate value ($\log_{10} C = 5.25$). The dots on the solid lines show the ratios for varying metallicity ($Z = 0.4, 1, 2 \times Z_\odot$) with the given compactness. The dashed lines indicate the variation resulting from adjusting f_{PDR} to values of 0.05 (higher temperature) or 0.15 (lower temperature). The shaded areas indicate the variation resulting from adjusting the ambient pressure to the extreme values supported by the templates. RIGHT PANEL: histogram of the MAPPINGS III contribution to the total flux at 350 μm for the EAGLE galaxies shown in the left panel. The vertical dashed line indicates the median.

6.3.4 Dust scaling relations

In Fig. 6.10, left panel, we compare mock observations of our EAGLE galaxies (red points) to HRS observations taken from Boselli et al. (2012) and Ciesla et al. (2012) (black points) for a submm color-color relation involving the SPIRE fluxes f_{250}, f_{350}, and f_{500}. These fluxes characterize the downwards slope of the dust continuum emission, and thus are sensitive to the cold dust contents. Smaller flux ratios f_{250}/f_{350} and f_{350}/f_{500} indicate a flatter slope of the dust emission curve and thus a larger contribution from colder dust. This is illustrated in the figure by the magenta curve, which traces the flux ratio relation for the emission of a modified black body (MBB, see Eq. 6.5) with $\beta = 2$ (the value assumed by the dust model used in this work) for temperatures ranging from 13 K to 28 K. Data points away from this curve indicate contributions from dust at various temperatures, resulting in a superposition of MBB curves and thus a broader dust continuum emission peak.

The EAGLE and HRS scaling relations show similar slopes. However, the EAGLE flux ratios are generally smaller than the corresponding HRS flux ratios, i.e. the simulated
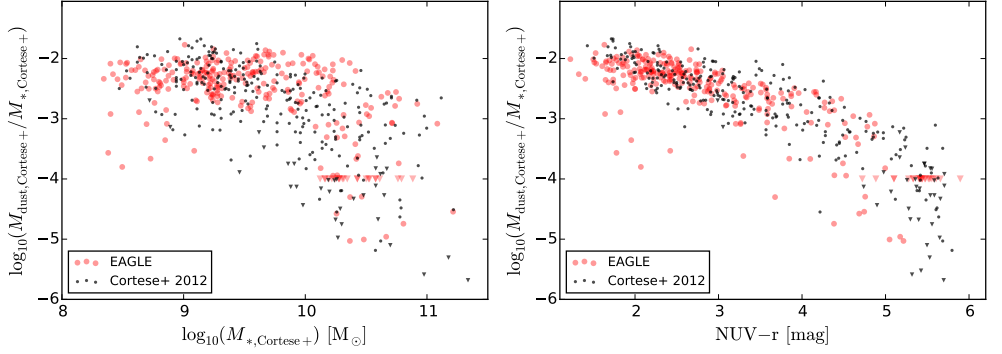
Figure 6.11: Dust scaling relations for EAGLE galaxies (red points) from set $\mathcal{C}$ (Table 6.2) compared to HRS observations (black points) taken from Cortese et al. (2012). The EAGLE galaxies were post-processed using $f_{\mathrm{dust}} = 0.3$ and $f_{\mathrm{PDR}} = 0.1$, and all EAGLE data points are obtained through mock observations (see text). The left panel shows the dust-to-stellar mass ratio versus stellar mass. The right panel shows the dust-to-stellar mass ratio versus $NUV - r$ color. In both panels, galaxies for which one or more SPIRE fluxes are below the detection limit are plotted as filled triangles rather than circles. For the simulated galaxies, these non-detections are plotted at an arbitrary value of $M_{\mathrm{dust}}/M_* = 10^{-4}$.

points are shifted to the area of the plot indicating lower dust temperatures. The EAGLE galaxies thus show a larger contribution from very cold dust ($T \lesssim 18$ K), confirming our finding in Sect. 6.3.3 that part of the simulated dust is insufficiently heated, plausibly because of limitations in the sub-grid models of the simulations. It appears that, even with the sub-grid techniques in our procedures, our model does not fully capture the clumpiness of the dust distribution in galaxies, so that an insufficient amount of dust is irradiated by the strong radiation fields present within and near star-forming regions.

The EAGLE outliers to the right of the MBB curve are caused by the simulated observational limitations built into our procedure. This can be understood by noting that the observational limitations are more severe for longer wavelengths, and furthermore depend on the absolute flux level and its spatial distribution across the sky, so that the effect on the observed flux ratios is strongly nonlinear. In fact, in a version of the plot (not shown) using convolved but unlimited submm fluxes, i.e. skipping the operations described in the last paragraph of Sect. 6.2.4.5, the EAGLE outliers move to the left of the MBB curve. We may thus surmise that the outlying HRS data points to the right of the MBB curve are similarly caused by observational limitations. The two labeled outliers, IC3718 and M86, are close to the SPIRE detection limit (Boselli et al. 2010; Gomez et al. 2010), supporting this assumption.

To help elucidate the contribution of the star-forming regions to the dust emission from the EAGLE galaxies, the green annotations in the left panel of Fig. 6.10 indicate the submm flux ratios of the MAPPINGS III templates used to model these regions for a range of input parameter values (see Table 6.3). The effective dust temperature of an isolated MAPPINGS III star-forming region is mostly determined by the specified metallicity, Z, and compactness, C. The other parameters, including the

covering fraction, f_{PDR}, and the ambient pressure, P, have a much smaller impact. The higher-than-observed EAGLE metallicities may thus contribute to the lower dust temperatures. At the same time, the temperature increases rapidly with the HII region's compactness. In fact, the temperature range $T \gtrsim 22$ K is reached only for the highest C values supported by the MAPPINGS III templates. The compactness in turn increases with the HII region mass (see Eq. 6.3). We sample these masses for our star-forming regions from a power-law distribution in the range $M \in [10^{2.8}, 10^6]\,\mathrm{M}_\odot$ (see Eq. 6.1), while Groves et al. (2008) quote a range of $M \in [10^{3.5}, 10^7]\,\mathrm{M}_\odot$, almost an order of magnitude higher. Our lower masses, and thus lower C values, might contribute to the apparent lack of warmer dust in our simulated results. These effects are limited, however, because for most galaxies the contribution of the MAPPINGS III templates to the total flux in the submm range is less than 25 per cent, as shown in the right panel of Fig. 6.10.

In Fig. 6.11 we compare the dust scaling relations for mock observations of our EAGLE galaxies (red points) to HRS observations taken from Cortese et al. (2012) (black points). In order to make the comparison as unbiased as possible, rather than employing intrinsic properties obtained by summing over the smoothed particles, all EAGLE data points in Fig. 6.11 are derived from mock observations based on the SEDs generated by SKIRT. Specifically, the stellar masses and dust masses are obtained following the recipes employed by Cortese et al. (2012) and summarized in Sect. 6.3.2. Galaxies for which one or more SPIRE fluxes are below the detection limit are plotted as filled triangles rather than circles. For the simulated galaxies, these non-detections are plotted at a fixed value of $M_{\mathrm{dust}}/M_* = 10^{-4}$.

The dust-to-stellar mass ratio versus stellar mass scaling relation (left panel of Fig. 6.11) is reproduced well, including the turn-off above $M_* = 10^{10}\,\mathrm{M}_\odot$, although there is a slight positive offset in the dust-to-stellar mass ratio. The relation of dust-to-stellar mass ratio versus $NUV - r$ color (right panel of Fig. 6.11) is also reproduced well, including the turn-off. As noted by Cortese et al. (2012), the dust-to-stellar mass ratio anti-correlates strongly with stellar mass and with $NUV - r$ color, a proxy for sSFR (e.g., Schiminovich et al. 2007). The relations are remarkably similar to the scaling relation involving the HI-to-stellar mass ratio (Catinella et al. 2010; Cortese et al. 2011; Fabello et al. 2011), suggesting that the dust and atomic hydrogen content of galaxies might be directly linked.

These results show that our EAGLE galaxies can indeed reproduce observed scaling relations based on FIR and submm fluxes. In the next section, we will present a quantitative measure for the correspondence between the simulations and data, and we will investigate the effect of changing the parameter values in our post-processing procedure.

6.3.5 Parameter study

To quantify the correspondence between two sets of data points, such as the mock observations and the HRS data for one of the scaling relations shown in Figs. 6.10
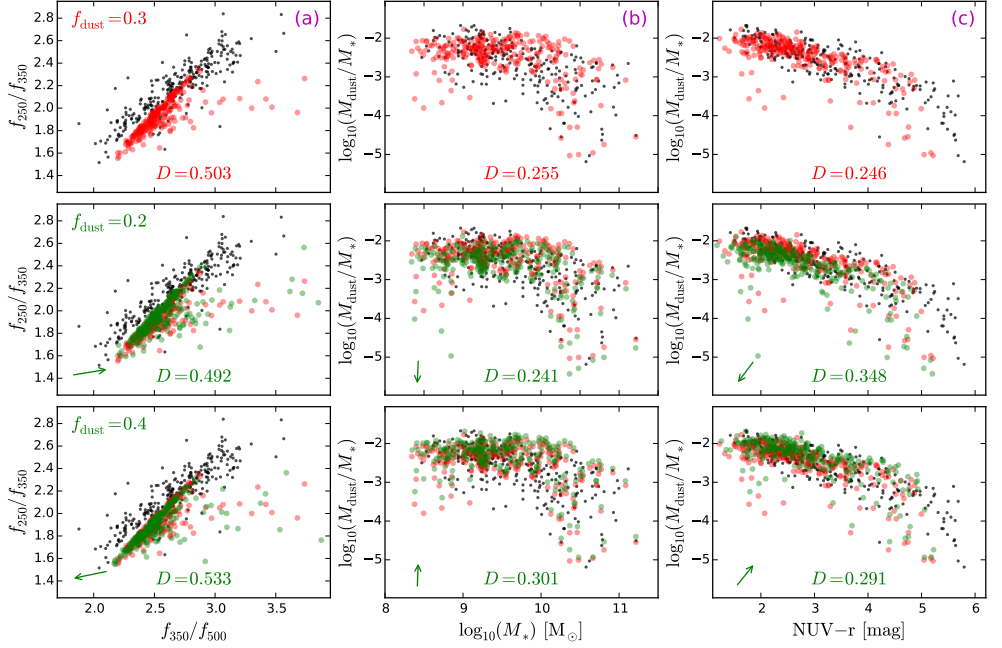
Figure 6.12: The scaling relations of Fig. 6.10, shown in column (a), and Fig. 6.11, shown in columns (b) and (c), for three values of the post-processing parameter f_{dust} (shown across the rows) with a fixed $f_{PDR} = 0.1$. The mock observations (red and green points) are overplotted on the HRS observations (black points) as before. The top row corresponds to the parameter settings of Figs. 6.10 and 6.11, i.e. $f_{dust} = 0.3$. These red points are repeated in the other rows for reference; the overplotted green points represent the mock observations with $f_{dust} = 0.2$ (middle row) and $f_{dust} = 0.4$ (bottom row). The green arrow in the lower left corner of these panels points in the direction of the average shift from red to green points (the length of the arrow is fixed). The D value is a measure for the correspondence between the mock and HRS observations (smaller is better). Galaxies for which one or more SPIRE fluxes are below the detection limit are omitted from this figure and from the calculation of D.

and 6.11, we use a generalization of the well-known Kolmogorov-Smirnov test (K-S test, Kolmogorov 1933; Smirnov 1948) to two-dimensional distributions. This generalization is due to Peacock (1983) and Fasano and Franceschini (1987), and its detailed implementation is described by Press et al. (2007). The 2D K-S test computes a metric D which can be interpreted as a measure of the 'distance' between two sets of two-dimensional data points, with smaller D values indicating better correspondence.

Figures 6.12 and 6.13 illustrate the effect of varying one of the post-processing parameters (respectively f_{dust} and f_{PDR}) on the scaling relations of Figs. 6.10 and 6.11, i.e. relations between submm flux colors in column (a), dust-to-stellar mass ratio versus stellar mass in column (b), and dust-to-stellar mass ratio versus $NUV - r$ color in column (c). Each panel in these figures shows the K-S test D value quantifying the discrepancy between the mock observations plotted in that panel and the underlying HRS data points. Because each column shows a particular scaling relation, D values
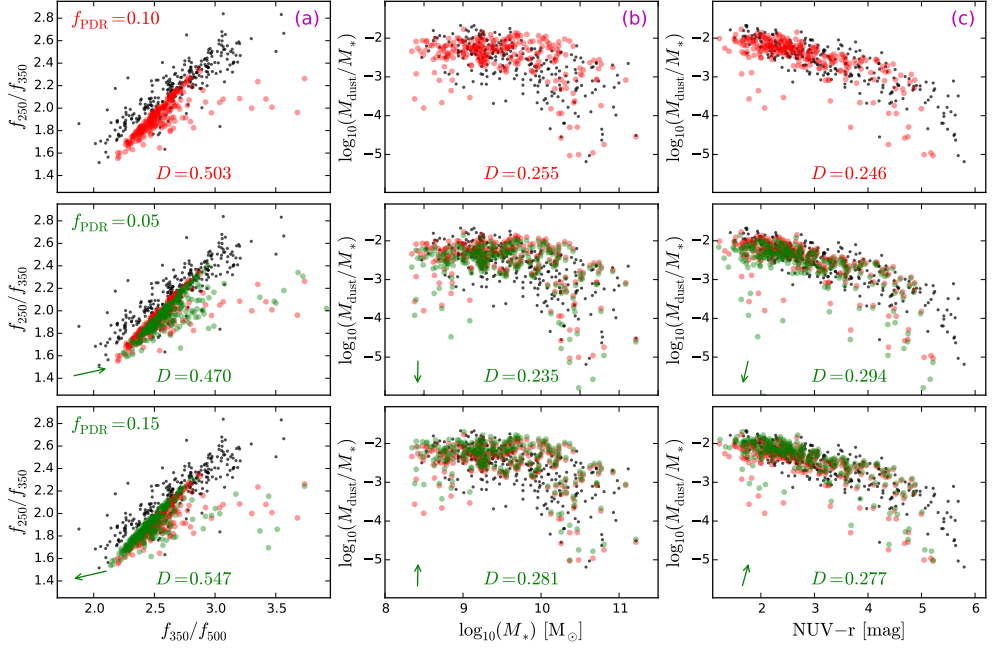
Figure 6.13: As Fig. 6.12, but now varying the post-processing parameter f_{PDR} with a fixed $f_{dust} = 0.3$. The top row again corresponds to the parameter settings of Figs. 6.10 and 6.11, i.e. $f_{PDR} = 0.1$. The overplotted green points in the other rows represent the mock observations with $f_{PDR} = 0.05$ (middle row) and $f_{PDR} = 0.15$ (bottom row).

are comparable across the rows within each column, but not between columns.

Figure 6.12 illustrates the effect of varying the dust-to-metal fraction f_{dust}. According to Eq. (6.2), the diffuse dust mass in a simulated galaxy scales linearly with f_{dust} (the star-forming regions model a separate body of dust that is independent of f_{dust}, see Sect. 6.2.3.4). In columns (b) and (c) we indeed see that, with increasing dust-to-metal fraction, the 'observed' dust-to-stellar mass ratio increases, and, for many galaxies, the $NUV - r$ color shifts red-ward because of the extinction caused by the extra dust. Also, as expected, there is very little effect on the 'observed' stellar mass. The effect on the submm colors shown in column (a) is not pronounced, although there seems to be a slight shift towards flatter slopes, i.e. colder dust temperatures, with increasing f_{dust}. This is because the larger body of diffuse dust is heated by the same stellar radiation to a lower average temperature.

Figure 6.13 illustrates the effect of varying the covering fraction f_{PDR} of the PDRs modelled in our simulations by the MAPPINGS III templates (see Sect. 6.2.3.4). In column (a) we see a shift towards flatter slopes, i.e. colder dust temperatures, with increasing f_{PDR}, caused by the more dispersed obscuration of the star-forming cores by the dust in the PDRs. In columns (b) and (c) we see an effect on the 'observed' dust mass similar to the effect of varying the dust-to-metal fraction (Figure 6.12). This effect is caused by the additional dust emission modelled by the MAPPINGS III SEDs for increasing covering fractions. For example, it is clear from the top panel of

Figure 6.14: The scaling relations of Fig. 6.10, shown in panel (a), and Fig. 6.11, shown in panels (b) and (c), for all galaxies in sets $\mathcal{C}$ and $\mathcal{F}$ (see Table 6.2). The points for our galaxies from snapshot Ref25 (blue) are plotted on top of those from Ref100 (red) and Recal25 (green) and the HRS observations (black). As in Figs. 6.10 and 6.11, the EAGLE galaxies were post-processed using $f_{\mathrm{dust}} = 0.3$ and $f_{\mathrm{PDR}} = 0.1$. Galaxies for which one or more SPIRE fluxes are below the detection limit are omitted from this plot.

Fig. 6 in Groves et al. (2008) that the continuum dust emission increases in addition to shifting to longer wavelengths. The apparent dust mass derived from a modified black body fitted to a MAPPINGS III SED increases by more than 80 per cent when f_{PDR} changes from 0.05 to 0.1, and by about another 50 per cent when f_{PDR} changes from 0.1 to 0.15, for fixed values of the other parameters (the precise percentages depend on the metallicity and other properties of the PDR). Varying the PDR covering fraction has only a small effect on the $\mathrm{NUV} - r$ color because the dust mass is added in compact regions and does not contribute much to the overall extinction. In fact, in our simulations, the dust in the PDR regions does not contribute at all to the extinction of radiation originating outside of the PDR region, because each region is handled individually using a MAPPINGS III model.

While adjusting f_{dust} or f_{PDR} produces clear overall trends in these scaling relations, the effect varies substantially between individual galaxies, as is especially noticeable for some of the outliers. These differences are caused by varying degrees of extinction by diffuse dust and different levels of contribution from star-forming regions. Also, our emulation of the observational limitations, including instrument resolution and sensitivity, cause a nonlinear response of the 'measured' flux to the actual submm emission of an EAGLE galaxy. This is especially true for galaxies with lower dust masses and thus less submm emission.

We used the D value of the Cortese et al. (2012) scaling relation in column (c) of Figs. 6.12 and 6.13 to determine our standard set of post-processing parameter values, even though the D value for the scaling relations in the other columns can be slightly more optimal for somewhat smaller amounts of dust (i.e. lower values of f_{dust} or f_{PDR}). This leads to adopted values of $f_{\mathrm{dust}} = 0.3$, fully compatible with the observed range of dust-to-metal fractions from 0.2 to 0.4 (Dwek 1998; Brinchmann et al. 2013; Zafar and Watson 2013), and $f_{\mathrm{PDR}} = 0.1$, which is lower than the fiducial value of 0.2 used by Jonsson et al. (2010). It is worth noting once more in this context that the dust masses derived from our mock observations scale inversely with the assumed value of κ_{350} (see Sect. 6.3.2). Using $\kappa_{\mathrm{Zubko+}}$ instead of $\kappa_{\mathrm{Cortese+}}$ would cause a 0.24 dex downward shift on the vertical axis in columns (b) and (c) of Figs. 6.12 and
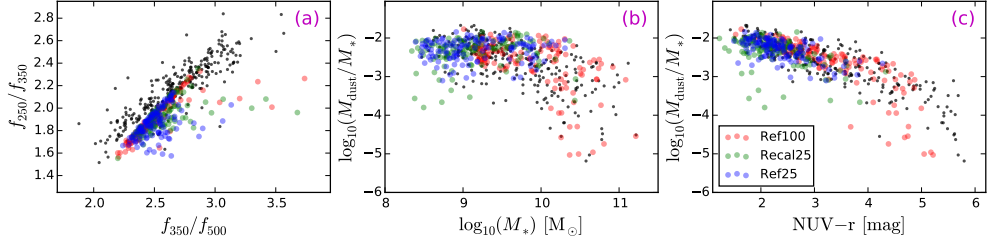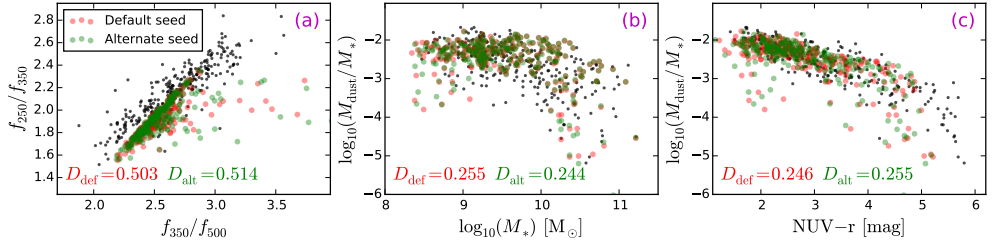
Figure 6.15: The scaling relations of Fig. 6.10, shown in panel (a), and Fig. 6.11, shown in panels (b) and (c), for the galaxies in set $\mathcal{C}$ using $f_{\text{dust}} = 0.3$ and $f_{\text{PDR}} = 0.1$. The star-forming regions are re-sampled with the default (red) and an alternate (green) pseudo-random sequence to evaluate the effect of these random variations on the SKIRT input model. The D value is a measure for the correspondence between the mock and HRS observations (smaller is better). Galaxies for which one or more SPIRE fluxes are below the detection limit are omitted from this plot.

6.13. Adjusting our post-processing parameters to compensate for this shift would evidently affect the absolute dust masses assigned to our galaxies, without, however, changing the actual extinction levels.

We now briefly review the mock observations shown in Fig. 6.14 for the galaxies extracted from the Ref25 snapshot (part of set $\mathcal{F}$, see Table 6.1). Based on the rightmost panel of Fig. 6.7, we already noted in Sect. 6.3.1 that the Ref25 galaxies on average have a higher metallicity than the Recal25 galaxies, resulting in a larger dust mass. This effect is evident in columns (b) and (c) of Fig. 6.14, where the Ref25 galaxies (blue points) are positioned slightly higher, on average, compared to the Recal25 galaxies (green points). Overall the convergence between the lower-resolution Ref100 snapshot and the higher-resolution Ref/Recal25 snapshots is very good.

For the analysis presented so far, the star-forming region re-sampling procedure described in Sect. 6.2.3.2 was performed only once for each galaxy. In other words, for a given galaxy, exactly the same particle input sets were presented to SKIRT for all radiative transfer simulations of that galaxy. This approach has enabled us to focus on the effects of varying the values of f_{dust} and f_{PDR} in the radiative transfer model without interference from random changes in the input model. To verify that no biases were introduced by our specific instantiation of the star-forming regions, we reran the re-sampling procedure with a different pseudo-random sequence (i.e. a different seed). Fig. 6.15 shows the calculated scaling relations, using our standard values of $f_{\text{dust}} = 0.3$ and $f_{\text{PDR}} = 0.1$, for the default and alternate input models. While individual galaxies can differ substantially, especially the outliers, the overall results and conclusions remain unchanged.

6.4   Conclusions

We calculated mock observations in the wavelength range from UV to submm for simulated galaxies extracted from the EAGLE suite of cosmological simulations

using the radiative transfer code SKIRT. To help overcome some of the resolution limitations of the simulations, we employed sub-grid models for the star-forming regions and for the diffuse dust distribution. We also took special care to mimic the effects of instrumental properties and observational limitations when calculating band-integrated fluxes, which are important especially in submm bands.

To validate our method, and at the same time confront the properties of the simulated galaxies with observed galaxies, we selected a set of present-day EAGLE galaxies that matches the *K*-band luminosity distribution and overall morphological type classification (using the sSFR as a proxy) of the galaxies in the *Herschel* Reference Survey (HRS), a volume-limited sample of about 300 normal galaxies in the Local Universe. We evaluated some intrinsic properties of our selected galaxies (Fig. 6.7), calculated by summing over the particles, and confirmed that the stellar masses, star-formation rates and dust masses fall in the expected range, while the average gas metallicities are above the metallicities observed in comparable galaxies.

We evaluated some relevant tracers by comparing the values derived from our mock observations to the corresponding intrinsic values (Figs. 6.8 and 6.9). We found that the Zibetti et al. (2009) calibration of stellar mass versus *g* and *i* band luminosity used by Cortese et al. (2012) underestimates stellar mass by about 0.25 dex, in line with the systematic uncertainty on the stellar mass-to-light ratio relation reported by other authors. Furthermore, the Cortese et al. (2012) recipe for deriving dust mass from the three SPIRE fluxes produces an offset that seems to be mostly caused by differences in the assumed dust absorption coefficient at 350 μm. The star-formation indicators based on respectively NUV, 24 μm, and integrated infrared fluxes perform fairly well, although the 24 μm estimates are consistently low, most likely because some of the dust in our model is too cold.

We then studied dust scaling relations, including the f_{250}/f_{350} versus f_{350}/f_{500} submm color-color relation (Fig. 6.10) and the dust-to-stellar mass ratio versus stellar mass and versus NUV $-$ r color relations (Fig. 6.11), comparing the properties derived from mock observations of the EAGLE galaxies with those observed for HRS galaxies. Using our 'standard' set of post-processing parameter values, we found good correspondence between the EAGLE and HRS scaling relations, with one important caveat. The submm colors indicate that the average temperature of the dust in our EAGLE galaxy models is lower than observed. We concluded that, even with the sub-grid techniques in our procedures, our model does not fully capture the clumpiness of the dust distribution in galaxies, so that an insufficient amount of dust is irradiated by the strong radiation fields present within and near star-forming regions.

We investigated the effects of varying the assumed dust-to-metal fraction, f_{dust} (Fig. 6.12), and PDR covering fraction, f_{PDR} (Fig. 6.13), in our post-processing procedure. The first parameter controls the diffuse dust, while the latter controls the dust near young stellar populations. We found that the effects on the scaling relations are consistent with expectations, although it is hard to determine unambiguously optimal values for both parameters because of the partial degeneracy of the effects. We settled on $f_{\mathrm{dust}} = 0.3$ and $f_{\mathrm{PDR}} = 0.1$ as our standard values, noting that these

values also depend on the properties of the dust in our model. These parameter values are compatible with observed values of f_{dust} (Dwek 1998; Brinchmann et al. 2013; Zafar and Watson 2013) and values of f_{PDR} suggested by other authors (Jonsson et al. 2010).

In conclusion, our analysis has shown that, in spite of some limitations, the EAGLE simulations can reproduce infrared and submm observations through a physically motivated post-processing procedure. Furthermore, we have used the mock submm observations to constrain the important dust-related parameters in our method, leading to a consistent calculation of dust attenuation in the UV and optical domain by Trayford et al. (in prep). While we studied present-day galaxies in this work, our post-processing method is equally applicable to galaxies at higher redshifts, and could be readily adapted to other hydrodynamical simulations.

The method presented here opens many opportunities for future work. We plan to add infrared and submm fluxes for a large subset of the EAGLE galaxies at various redshifts to the public EAGLE database (McAlpine et al. 2016) as a service to other researchers. We could study the morphology and structure of mock EAGLE galaxy images in various wavelength bands from UV to submm, comparing those properties with resolved observations in the same bands. It would also be instructive to post-process some of the more resolved galaxies in the zoom-in simulations (Sawala et al. 2016; Oppenheimer et al. 2016) based on the EAGLE code.

In a more distant future, the astrophysical community will undoubtedly develop more advanced simulation techniques. Future simulations of galaxy evolution and assembly may use sub-grid recipes on a smaller scale to model a cold phase in the ISM, which will allow more accurate modeling of the clumpy dust structure during post-processing. Radiative transfer codes may include sub-grid models of star-forming regions that are connected to the overall RT model in a self-consistent manner, as opposed to employing 'disconnected' SED templates. We hope that our current work will help inform such future efforts.

6.5   Appendix: Simulated broad-band photometry

This appendix summarizes the formulae used to convolve a simulated SED with an actual instrument's response curve to obtain a band-integrated flux. We essentially follow the treatments given on Ivan K. Baldry's web page[23] and in the appendix of Tokunaga and Vacca (2005). The procedure depends on whether the instrument counts photons or measures energy.

6.5.1   Photon counters

For a photon counting detector, the number of photons detected per unit time, t, and per unit area, S, from a source with an intrinsic spectral energy distribution $F_\lambda(\lambda)$,

---

[23] http://www.astro.ljmu.ac.uk/~ikb/research/mags-fluxes/

or equivalently $F_\nu(\nu)$, by a photon counter with total system response $R(\lambda)$, or $R(\nu)$, can be written as

$$\frac{dN_p}{dt\,dS} = \int \frac{\lambda}{hc} F_\lambda(\lambda) R(\lambda)\,d\lambda = \int \frac{1}{h\nu} F_\nu(\nu) R(\nu)\,d\nu. \tag{6.6}$$

We define the source's mean intrinsic flux $\langle F_\lambda \rangle$ and $\langle F_\nu \rangle$ through

$$\frac{dN_p}{dt\,dS} = \langle F_\lambda \rangle \int \frac{\lambda}{hc} R(\lambda)\,d\lambda = \langle F_\nu \rangle \int \frac{1}{h\nu} R(\nu)\,d\nu, \tag{6.7}$$

and we define the pivot wavelength λ_{pivot} and frequency ν_{pivot} connecting these two flux representations through

$$\langle F_\nu \rangle = \langle F_\lambda \rangle \frac{\lambda_{\text{pivot}}^2}{c} \quad \text{and} \quad \langle F_\lambda \rangle = \langle F_\nu \rangle \frac{\nu_{\text{pivot}}^2}{c}. \tag{6.8}$$

Combining Eqs. (6.6) and (6.7) yields

$$\langle F_\lambda \rangle = \frac{\int \lambda F_\lambda(\lambda) R(\lambda)\,d\lambda}{\int \lambda R(\lambda)\,d\lambda} \quad \text{and} \quad \langle F_\nu \rangle = \frac{\int F_\nu(\nu) R(\nu)\,d\nu/\nu}{\int R(\nu)\,d\nu/\nu}. \tag{6.9}$$

Substituting $d\nu/\nu = d\lambda/\lambda$ in Eq. (6.7) and combining with Eq. (6.8) yields

$$\lambda_{\text{pivot}} = \sqrt{\frac{\int \lambda R(\lambda)\,d\lambda}{\int R(\lambda)\,d\lambda/\lambda}} \quad \text{and} \quad \nu_{\text{pivot}} = \sqrt{\frac{\int R(\nu)\,d\nu/\nu}{\int R(\nu)\,d\nu/\nu^3}}. \tag{6.10}$$

6.5.2   Bolometers

For an energy measuring device, the quantities $\lambda R(\lambda)/hc$ and $R(\nu)/h\nu$ in the above analysis must be replaced by the total system transmission $T(\lambda)$ or $T(\nu)$. Eqs. (6.9) and (6.10) then become

$$\langle F_\lambda \rangle = \frac{\int F_\lambda(\lambda) T(\lambda)\,d\lambda}{\int T(\lambda)\,d\lambda} \quad \text{and} \quad \langle F_\nu \rangle = \frac{\int F_\nu(\nu) T(\nu)\,d\nu}{\int T(\nu)\,d\nu} \tag{6.11}$$

and

$$\lambda_{\text{pivot}} = \sqrt{\frac{\int T(\lambda)\,d\lambda}{\int T(\lambda)\,d\lambda/\lambda^2}} \quad \text{and} \quad \nu_{\text{pivot}} = \sqrt{\frac{\int T(\nu)\,d\nu}{\int T(\nu)\,d\nu/\nu^2}}. \tag{6.12}$$

Conclusions

7

7.1   Objectives

To improve our understanding of the physical processes underlying astronomical objects under study, astrophysical research often resorts to studying the correspondences and differences between numerical simulation results and astronomical observations. Constructing "mock observations" of simulated objects at multiple wavelengths enables direct comparison to observational data. There is thus a need for post-processing "snapshots" of sophisticated hydrodynamical evolution models, for example to mimic the effects of an obscuring medium on the radiation emitted by the sources in the model.

As part of this thesis work, we established a general framework for post-processing the results of hydrodynamical astrophysical simulations using the 3D dust continuum radiative transfer (RT) code SKIRT. We specifically applied our methods to construct mock observations of the galaxies in the EAGLE cosmological simulations, over a broad wavelength range including ultraviolet, optical, infrared, and submm radiation, taking into account the effects of interstellar dust on the stellar radiation. We finally compared these results to far-infrared and submm observations of a set of present-day galaxies. In Sect. 7.2 we provide a brief summary and formulate our main conclusions, and in Sect. 7.3 we offer some ideas for future research.

7.2   Summary and conclusions

User interface and software design. In Chapter 2 we described the design changes we implemented in SKIRT to cope with the large number of built-in options and the growing configuration requirements of this state-of-the-art RT simulation code. In addition to its core capability of tracing the radiation through the dust, SKIRT offers options for configuring all aspects of the simulation model, including spatial and spectral distributions, dust grain characterizations, simulated detection systems, and discretization. Inspired by standard software design principles and patterns, the

latest version of SKIRT has a modular implementation that can be easily maintained and expanded. Programming interfaces between components are well defined and narrow. The user interface is automatically constructed from data provided in the C++ class declarations, allowing a single point of definition, and placing the user interface information right next to the code implementing the corresponding feature.

All too often, scientific codes are written without much concern for user interface or for modular software design. This is very unfortunate. Scientists may not need a *graphical* user interface, but, just like every one else, they do benefit from an interaction mechanism that hides the underlying complexity. As we have illustrated in this work, a well-designed non-graphical user interface may be a perfect fit, and can often be developed and maintained with limited resources. Similarly, adhering to proven software design principles pays off, even for small and mid-sized projects.

Voronoi dust grids. A proper discretization of the spatial domain is a key element of a RT simulation. To accommodate the complex density distributions in hydrodynamical simulation results, the size and placement of grid cells must automatically adapt to the requirements of each individual model. In Chapter 3 we described our implementation in SKIRT of an unstructured grid based on Voronoi tesselations of 3D space, a grid type which has been successfully used in some recent hydrodynamical simulation codes. These grids tend to more closely reflect dynamic ranges in the model with fewer cells, presenting cell boundaries that are more adjusted to the underlying gradients.

In a RT simulation the Voronoi grid can be a very flexible tool. Since a Voronoi grid is defined solely by its generating points, the cell size and distribution can be easily fine-tuned by placing these sites in the appropriate locations. For example, sites can be randomly distributed according to the input model's density or opacity fields; if needed extra sites can be added in high-gradient areas. In the case of a particle-based input model, the particle locations themselves can serve as sites; and for an input model already based on a Voronoi mesh no re-gridding is required at all.

In this work we have shown that it is straightforward to implement accurate and efficient RT on Voronoi grids. In spite of the geometric complexity of the cell boundaries, calculating straight paths between two arbitrary points through a 3D Voronoi grid is only about three times slower than a highly optimized octree implementation with the same number of cells, while in practice the total number of Voronoi grid cells may be lower for an equally good representation of the density field. The presented method automatically yields the precise distance covered by the path inside each grid cell, and eliminates the need for corrections or work-arounds required by alternate, approximate approaches. We concluded that the benefits of using a Voronoi grid in RT simulation codes will often outweigh the somewhat slower performance.

Stochastically heated dust grains. To properly reproduce the infrared emission of dusty astrophysical objects in a RT simulation, it is necessary to calculate the

temperature probability distribution of the dust grains in each cell of the dust grid, including a full treatment of stochastically heated dust grains (SHGs). In Chapter 4 we reported on our implementation of this calculation in SKIRT. We provided a self-contained guide for implementors of such functionality through a method that is frequently used in RT codes because of its good performance and relative ease of implementation, although it assumes continuous cooling of the dust grains, which may be inaccurate in extreme environmental conditions.

We defined an appropriate problem for benchmarking this calculation in the context of RT simulations. The problem's definition includes the optical and calorimetric material properties, and the grain size distributions, for a typical astronomical dust mixture. It also includes a series of analytically defined radiation fields to which the dust population is to be exposed and instructions for the desired output. We then processed the benchmark problem with six RT codes participating in this benchmark effort, including SKIRT. We reported in detail on the similarities and differences between the results from the participating codes and a reference solution. In the important wavelength range $3\,\mu\mathrm{m} \leq \lambda \leq 1000\,\mu\mathrm{m}$, all participating codes reproduce the total dust emissivity within 20% of the reference solution for all input fields used in this benchmark. Excluding the weakest and the softest input fields, the agreement in the same wavelength range is within 10%.

In conclusion, this benchmark effort showed that the relevant modules in RT codes can and do produce fairly consistent results for the emissivity spectra of SHGs, which have a significant impact on the final result of a multiwavelength RT simulation. We offered concrete, quantitative information on the level of (dis)agreement between RT codes, which will help inform the interpretation of RT simulation results that include SHG dust emission calculations of the type presented here.

Importing hydrodynamical snapshots.   In Chapter 5 we described the mechanisms we added to SKIRT for importing snapshots produced by hydrodynamical simulation codes. We implemented configurable methods to read or derive the properties of the radiation sources and of the obscuring dust from snapshots of various simulation types and in various data formats. These properties included the spatial density distribution, in addition to the spectral energy distribution for the sources and the optical material properties for the dust. We concluded that SKIRT's capabilities are very flexible and enable RT post-processing of snapshots from essentially any hydrodynamical code with minimal data conversion and little or no programming requirements.

Dust in the EAGLE galaxies.   Using the SKIRT code updated as described above, we studied the far-infrared and dust-related properties of the present-day EAGLE galaxies, while, in related work, Trayford et al. (in prep) studied the optical properties. In Chapter 6 we presented our method to calculate mock observations in a wavelength range from UV to submm for galaxies extracted from the EAGLE suite of cosmological simulations. To help overcome some of the resolution limitations of the simulations,

we employed sub-grid models for the star-forming regions and for the diffuse dust distribution. We also took special care to mimic the effects of instrumental properties and observational limitations when calculating band-integrated fluxes, which are important especially in submm bands.

To validate our method, and at the same time confront the properties of the simulated galaxies with observed galaxies, we selected a set of present-day EAGLE galaxies that matches the *K*-band luminosity distribution and overall morphological type classification of the galaxies in the *Herschel* Reference Survey (HRS), a volume-limited sample of about 300 normal galaxies in the Local Universe. We evaluated some intrinsic properties of our selected galaxies, calculated by summing over the smoothed particles, and confirmed that the stellar masses, star-formation rates and dust masses fall in the expected range, while the average gas metallicities are above the metallicities observed in comparable galaxies.

We evaluated some relevant tracers by comparing the values derived from our mock observations to the corresponding intrinsic values. We found that the Zibetti et al. (2009) calibration of stellar mass versus *g* and *i* band luminosity used by Cortese et al. (2012) underestimates stellar mass by about 0.25 dex, in line with the systematic uncertainty on the stellar mass-to-light ratio relation reported by other authors. Furthermore, the Cortese et al. (2012) recipe for deriving dust mass from the three SPIRE fluxes produces an offset that seems to be mostly caused by differences in the assumed dust absorption coefficient at 350 μm. The star-formation tracers based on respectively NUV, 24 μm, and integrated infrared fluxes perform fairly well, although the 24 μm estimates are consistently low, most likely because some of the dust in our model is too cold.

We then studied dust scaling relations, including the f_{250}/f_{350} versus f_{350}/f_{500} submm color-color relation and the dust-to-stellar mass ratio versus stellar mass and versus NUV $- r$ color relations, comparing the properties derived from mock observations of the EAGLE galaxies with those observed for HRS galaxies (Boselli et al. 2012; Cortese et al. 2012). Using our 'standard' set of post-processing parameter values, we found good correspondence between the EAGLE and HRS scaling relations, with one important caveat. The submm colors indicate that the average temperature of the dust in our EAGLE galaxy models is lower than observed. We concluded that, even with the sub-grid techniques in our procedures, our model does not fully capture the clumpiness of the dust distribution in galaxies, so that an insufficient amount of dust is irradiated by the strong radiation fields present within and near star-forming regions.

We investigated the effects of varying the assumed dust-to-metal fraction, f_{dust}, and photodissociation-region covering fraction, f_{PDR}, in our post-processing procedure. The first parameter controls the diffuse dust, while the latter controls the dust near young stellar populations. We found that the effects on the scaling relations are consistent with expectations, although it is hard to determine unambiguously optimal values for both parameters because of the partial degeneracy of the effects. We settled on $f_{\text{dust}} = 0.3$ and $f_{\text{PDR}} = 0.1$ as our standard values, noting that these values also

depend on the properties of the dust in our model. These parameter values are compatible with values observed and suggested by other authors.

In conclusion, our analysis has shown that, in spite of some limitations, the EAGLE simulations can reproduce infrared and submm observations through a physically motivated post-processing procedure. Furthermore, we have used the mock submm observations to constrain the important dust-related parameters in our method, leading to a consistent calculation of dust attenuation in the UV and optical domain by Trayford et al. (in prep). While we studied present-day galaxies in this work, our post-processing method is equally applicable to galaxies at higher redshifts, and could be readily adapted to other hydrodynamical simulations.

7.3	Outlook

The methods and results presented in this thesis open many opportunities for future work. Here is a selection of projects that come to mind, and in which I would gladly be involved in the near future. The items are sorted roughly in the order of the chapters to which they relate.

- SKIRT has the potential to become the hydro post-processing code of choice within the astrophysical community (for dust RT at least). To fully realize this potential, we should implement a cleaner software architecture of the post-processing features, making the code more flexible and easier to configure for various astrophysical problem domains. Furthermore, we should continue to improve the overall design and the online presence of the code, further lowering the barrier to install, configure and use it.

- The special properties of Voronoi grids could reduce the number of grid cells required in a RT simulation to properly resolve the discretized physical quantities. However we need to develop specific heuristics for placing the generating sites in appropriate positions to achieve optimal grid quality, for example, by ensuring that cell borders are perpendicular to the gradients in the model's density distributions.

- Current RT codes, including SKIRT, construct the dust grid during setup of the simulation, based on the static model properties. An optimal dust grid should properly resolve the radiation field, which by definition is not yet available during setup. We should therefore look into an iterative or adaptive approach, where the grid is rebuilt or adjusted based on partial simulation results.

- The current method for calculating the emission of stochastically heated dust grains (in SKIRT and in many RT codes) assumes that dust grains cool "continuously", i.e. through a stream of photons each carrying off an energy that is small compared to the internal energy of the grain. Lifting this restriction requires solving a much more complex problem involving transitions between all energy states of the dust grain rather than just adjacent states, but is necessary when considering dust grains in more extreme radiation conditions.

- When investigating systems with a lot of hot dust, such as circumstellar disks or accretion disks, it may be relevant to properly calculate dust emission for wavelengths shorter than $1\,\mu$m. To accomplish this, RT codes will need to model environment-dependent destruction of dust grains and to adjust the grain size distribution used in the dust emission calculation accordingly.

- Now that we have established and verified our method, we should add infrared and submm fluxes for a large subset of the EAGLE galaxies at various redshifts to the public EAGLE database as a service to other researchers. Using this information we could, for example, study the evolution of the infrared and submm properties of relevant galaxy groups.

- Similarly, we could study the morphology and structure of mock EAGLE galaxy images in various wavelength bands from UV to submm, comparing those properties with resolved observations in the same bands. In that context, it might be instructive to post-process some of the more resolved galaxies in the zoom-in simulations based on the EAGLE code, or, for that matter, galaxies from other high-resolution simulations.

In a more distant future, the astrophysical community will undoubtedly develop more advanced simulation techniques. Future hydrodynamical simulations of astrophysical objects will use sub-grid recipes on a smaller scale, for example to properly model a cold phase in the interstellar medium, which will allow more accurate modeling of the clumpy dust structure during post-processing. Future radiative transfer codes will include much improved sub-grid models, for example star-forming regions that are connected to the overall RT model in a self-consistent manner, as opposed to employing 'disconnected' SED templates. We hope that our current work will help inform such future efforts.

Samenvatting 8

PROBEER NIET OM HET LAATSTE WOORD TE HEBBEN.
JE ZOU HET KUNNEN KRIJGEN.[24]

8.1 Inleiding

Toen Albert Einstein zijn bekende zwaartekrachtstheorie publiceerde (Einstein 1916), nu 100 jaar geleden, ging hij er zoals velen van uit dat het heelal uit één enkele galaxie bestond, onze Melkweg. Pas in de twintiger jaren toonde Edwin Hubble aan dat sommige van de waargenomen "nevels" in feite galaxieën zijn die ver buiten onze eigen Melkweg liggen (Hubble 1925). Nu weten we dat er miljarden galaxieën zijn (Illingworth et al. 2013), dus deze objecten verdienen zeker onze aandacht.

De moderne sterrenkunde heeft tot doel om de fysische processen beter te begrijpen die aan de grondslag liggen van de bestudeerde objecten. Om dit te kunnen doen, moeten we de driedimensionale (3D) structuur van die objecten in kaart brengen, terwijl we uiteraard slechts een tweedimensionale projectie van het object waarnemen. Bovendien wordt in veel galaxieën een belangrijk deel van de door sterren uitgezonden straling geabsorbeerd door stofdeeltjes in het interstellair medium, en vervolgens opnieuw uitgezonden onder de vorm van thermische straling met een veel langere golflengte (De Looze et al. 2012a; Verstappen et al. 2013; De Geyter et al. 2014, 2015). Deze complicatie zorgt tegelijkertijd ook voor extra informatie over het stof waarmee de straling interageert.

Voor realistische modellen van astrofysische systemen is het stralingsoverdracht-probleem onmogelijk analytisch op te lossen. De verstrooiing van fotonen aan stofdeeltjes zorgt voor een koppeling van alle mogelijke richtingen of zichthoeken in het systeem, en de absorptie van optische straling en heruitzending als thermische straling betekent ook een vermenging van de verschillende golflengten in de vergelijkingen. We kunnen er zelfs niet van uitgaan dat de stofdeeltjes in lokaal thermisch evenwicht zijn met de omgevingsstraling. Het blijkt namelijk dat kleine stofdeeltjes en polycyclische aromatische koolwaterstoffen soms tot ver boven de evenwichtstemperatuur kunnen opgewarmd worden door één enkel foton (Boulanger and Perault 1988; Draine and Li 2001), waardoor de thermische emissie van een stofpopulatie veel moeilijker te bepalen is.

---

[24] De korte citaten aan het begin van elk hoofdstuk zijn overgenomen uit de "Notebooks of Lazarus Long" in het boek *Time Enough for Love* door Robert A. Heinlein (1973).

Om al deze redenen gebruiken de meeste stralingsoverdracht-simulatiecodes een Monte Carlo techniek om de levenscyclus van een groot aantal fotonpaketten na te bootsen, waarbij elke gebeurtenis (zoals ontstaan, verstrooiing of absorptie) gestuurd wordt door een toevalsveranderlijke met de gepaste waarschijnlijkheidsverdeling (Whitney 2011; Steinacker et al. 2013). SKIRT (Baes et al. 2011) is zo'n code, ontwikkeld door de onderzoeksgroep van Prof. Baes, ook promotor van dit proefschrift. SKIRT kan willekeurige 3D modellen behandelen en voorziet de belangrijke optimalisatie-technieken zoals die in de literatuur beschreven zijn (Yusef-Zadeh et al. 1984; Lucy 1999; Niccolini et al. 2003; Cashwell and Everett 1959). Als deel van het werk voor dit proefschrift hebben we de SKIRT code substantieel uitgebreid en aangepast, zoals we verderop in deze samenvatting nog zullen aanduiden.

SKIRT werd reeds met succes gebruikt om bijvoorbeeld de energiebalans in galaxieën te onderzoeken. We verwachten uiteraard dat het stof dezelfde hoeveelheid energie uitzendt (voornamelijk in het verre infrarood) als het absorbeert (voornamelijk in het optische bereik). Deze balans kunnen we nagaan door een computermodel te bouwen voor een bepaalde galaxie, en de resultaten van een stralingsoverdracht-simulatie te vergelijken met de observaties. De studies door bijvoorbeeld Baes et al. (2010) en De Looze et al. (2012a,b) wijzen erop dat een belangrijk deel van het interstellair stof in een klonterige structuur vervat moet zitten om de energiebalans te doen kloppen.

Bij deze studies werden analytische modellen gebruikt om de ruimtelijke structuur voor te stellen, eventueel aangevuld met enkele numerieke technieken om bijvoorbeeld een aantal stofklonters toe te voegen op willekeurige plaatsen. Maar zelfs dan hebben deze modellen eerder kunstmatige vormen die de klonterige structuur onvoldoende benaderen. Daarom wordt er de laatste jaren meer en meer gebruik gemaakt van volledig numerieke modellen, geproduceerd door hydrodynamische simulatiecodes die de vorming en evolutie van één of meerdere galaxieën nabootsen. Zo hebben Saftly et al. (2015) bijvoorbeeld de energiebalans bestudeerd in twee numeriek gesimuleerde galaxiemodellen. Hieruit blijkt dat de klonterige structuur in deze modellen de realiteit voldoende dicht benadert om de energiebalans te corrigeren.

Het EAGLE project (Schaye et al. 2015) omvat een reeks van hydrodynamische simulaties die de vorming en evolutie van galaxieën volgen in een kosmologisch representatief volume. Door de overeenkomsten en verschillen te bestuderen tussen de simulatieresultaten en waarnemingen hopen we de fysische processen beter te begrijpen die geleid hebben tot het hedendaagse heelal. Maar om deze vergelijking op een correcte manier te kunnen doen, moeten we ook "waarnemingen" nabootsen van onze gesimuleerde galaxieën (Hayward and Smith 2015; Guidi et al. 2015).

Het hoofddoel van dit proefschrift is om, gebruik makend van SKIRT, dergelijke kunstmatige waarnemingen te kunnen uitvoeren voor de EAGLE galaxieën, rekening houdend met de effecten van interstellair stof in een breed golflengtebereik van ultraviolet, over optisch en infrarood, tot sub-millimeter. Op de weg om dit doel te bereiken hebben we in feite een algemeen bruikbaar raamwerk ontwikkeld om kunstmatige waarnemingen te doen voor de resultaten van om het even welke
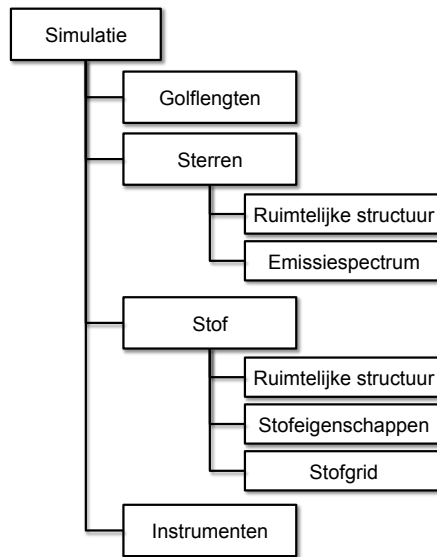
Figuur 8.1: Een overzicht van de configuratie voor een SKIRT simulatie.

hydrodynamische simulatie. Daartoe hebben we een aantal uitdagingen overwonnen die we in de volgende secties verder toelichten.

8.2 Software ontwerp

Figuur 8.1 geeft een algemeen overzicht van de onderdelen die moeten vastgelegd worden voor elke SKIRT simulatie. Zelfs vooraleer het werk aan dit proefschrift begon, bood de SKIRT stralingsoverdracht-simulatiecode al een hele waaier mogelijkheden aan voor elke onderdeel, bijvoorbeeld om de ruimtelijke structuur via één van de ingebouwde analytische modellen voor te stellen. Om ook numerieke modellen te kunnen behandelen van verschillende types en in allerlei bestandsformaten, moesten we nog een groot aantal bijkomende opties voorzien. Om het geheel beheersbaar te houden voor zowel de ontwikkelaar als de gebruiker vonden we het gepast om de C++ code grondig te reorganiseren. Hierbij hebben we de volgende doelstellingen gerealiseerd:

- De volledige configuratie voor een SKIRT-simulatie wordt vastgelegd in een gestructureerd XML bestand dat geschikt is voor automatische verwerking, maar dat ook toegankelijk is voor een expert-gebruiker via een gewone teksteditor.

- De informatie gerelateerd aan een bepaalde voorziening is op één plek verzameld, met inbegrip van de C++ code die de optie implementeert en de tekstsegmenten die dienen om de optie aan een gebruiker te omschrijven.

- De vraag-en-antwoord sessie waarmee de gebruiker een specifieke configuratie opbouwt, wordt volledig aangestuurd door de informatie bedoeld in het vorige
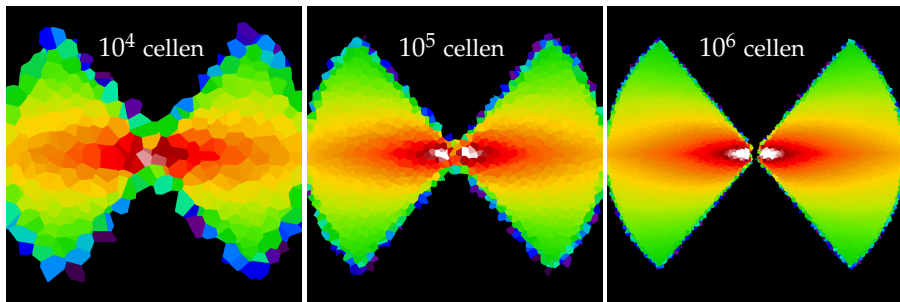
Figuur 8.2: Doorsnedes van de ruimtelijke stofverdeling in een torusmodel, gediscretiseerd via een Voronoi rooster met een resolutie gaande van 10^4 cellen (*links*) tot 10^6 cellen (*rechts*). In de drie gevallen werden de genererende sites willekeurig bemonsterd uit een uniforme waarschijnlijkheidsverdeling.

    item; met andere woorden, de gebruikersinterface past zich volautomatisch aan wanneer een nieuwe optie toegevoegd wordt.

- De code is op een modulaire manier georganiseerd, met minimale en goed-gedefinieerde afhankelijkheden tussen de verschillende onderdelen, onder meer door gebruik te maken van standaard ontwerppatronen (Gamma et al. 1994).

Door deze verbeteringen is de SKIRT code nu gemakkelijk uit te breiden en heel gebruikersvriendelijk. We durven te stellen dat andere wetenschappelijke codes, en de gebruikers ervan, ook zouden gebaat zijn met een dergelijke doordachte software architectuur en gebruikersinterface. De SKIRT code is publiek beschikbaar.[25]

8.3	Voronoi stofroosters

Een heel belangrijk onderdeel van elke stralingsoverdracht-simulatie is het stofrooster, met andere woorden de manier waarop het ruimtelijk domein in aparte cellen opgedeeld wordt. Zoals de meeste codes veronderstelt SKIRT dat alle fysische grootheden, zoals bijvoorbeeld de dichtheid en de optische eigenschappen van het stof, constant blijven binnen elke stofcel. Aangezien de simulatietijd en het benodigde geheugen stijgen met het aantal stofcellen, is het relevant om kleine cellen enkel te voorzien in gebieden waar dat nodig is om de fysische processen correct te vatten, en grotere cellen in de andere gebieden. Dit is een actief onderzoeksdomein in de stralingsoverdracht-gemeenschap (Brinch and Hogerheijde 2010; Lunttila and Juvela 2012; Saftly et al. 2013, 2014).

Bij de meest voor de hand liggende roosterstructuren hebben alle cellen de vorm van een kubus of een balk. Maar er bestaan ook "ongestructureerde" roosters, zoals bijvoorbeeld de ruimtelijke opdeling die bestudeerd werd door Voronoi (1908). Gegeven een lijst van genererende punten of sites s_i, wordt de Voronoi opdeling van de ruimte

---

[25] SKIRT documentatie: http://www.skirt.ugent.be
SKIRT code: https://github.com/skirt/skirt

met evenveel cellen als sites gedefinieerd door te stellen dat cel C_i bestaat uit alle punten in de ruimte die dichter bij site s_i liggen dan bij elke andere site. Zie Fig. 8.2 voor een voorbeeld. Voronoi roosters hebben een aantal eigenschappen die doen vermoeden dat, althans voor sommige stofverdelingen, eenzelfde rooster-kwaliteit kan bekomen worden met minder cellen dan wat nodig is met balkvormige cellen. Bovendien gebruiken sommige recente hydrodynamische simulatiecodes Voronoi roosters om de systeemevolutie te volgen (Duffell and MacFadyen 2011; Springel 2010; Vandenbroucke and De Rijcke submitted), en zou het zinvol zijn om hetzelfde rooster ook voor de stralingsoverdracht-simulatie te kunnen gebruiken. Anderzijds hebben Voronoi cellen een veel complexere vorm (elke cel is een convex veelvlak) waardoor het volgen van fotonpakketten door zo'n rooster behoorlijk ingewikkeld dreigt te worden.

Als onderdeel van dit proefschrift stelden we een eenvoudige methode voor om rechte paden te berekenen doorheen een willekeurig 3D Voronoi rooster. Deze methode gebruikt de specifieke eigenschappen van de Voronoi opdeling om complexe meet-kundige berekeningen te vermijden. De enige informatie vereist voor de berekening is een lijst van de dichtste naburen voor elke cel. Deze informatie kan zeer compact bijgehouden worden via de indices van de corresponderende genererende sites. De methode heeft geen nood aan het bijhouden van de hoekpunten of de grensvlakken van de cellen.

We hebben de voorgestelde methode geïmplementeerd in SKIRT en getest met een aantal modellen, met inbegrip van de Pascucci et al. (2004) benchmark. Daaruit blijkt dat het volgen van fotonpakketten doorheen een Voronoi rooster slechts ongeveer drie maal langer duurt, per overgestoken cel, dan doorheen een octree rooster (Saftly et al. 2013) met balkvormige cellen. Rekening houdend met de andere mogelijke voordelen van een Voronoi rooster is dit een veelbelovend gegeven.

8.4 Stochastische opwarming van stofdeeltjes

Zoals aangehaald in de inleiding, kunnen kleine stofdeeltjes en polycyclische aroma-tische koolwaterstoffen (PAH-moleculen) soms tot ver boven de evenwichtstempera-tuur opgewarmd worden door één enkel foton (Boulanger and Perault 1988; Draine and Li 2001). Daardoor wordt het een stuk moeilijker om de thermische emissie van een stofpopulatie te berekenen; we moeten namelijk de waarschijnlijkheidsverdeling van de stoftemperatuur (of beter gezegd, van het interne energieniveau) opstellen voor elk soort stofdeeltje in de populatie. Aangezien deze temperatuursverdeling af-hangt van het stralingsveld waarin het stof zich bevindt, moeten we deze berekening bovendien herhalen voor elke cel in het stofrooster van de simulatie.

De vroegere SKIRT module die instond voor deze berekeningen was traag en had een aantal beperkingen die onaanvaardbaar waren voor onze doelstellingen. Daarom hebben we de berekening van de thermische stofemissie in SKIRT opnieuw geïmplementeerd, deze keer volledig geïntegreerd in de rest van de code. We kozen
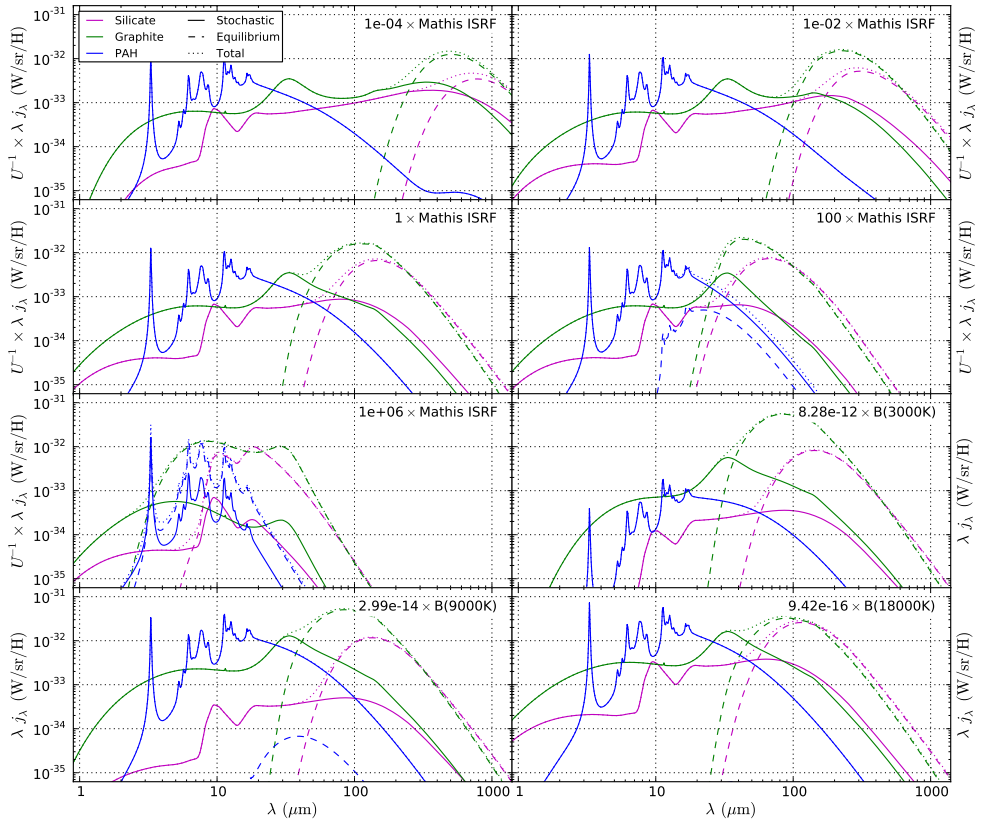
Figuur 8.3: Bijdragen tot de totale emissie (puntjeslijn) berekend in het stochastische regime (volle lijn) en in het evenwichtsregime (stippellijn) voor elk van de stofsoorten in ons model (silicaten in magenta, grafiet in groen en PAH moleculen in blauw). Elk paneel toont de resultaten voor één van de input-stralingsvelden in onze benchmark, t.t.z. een aantal schalingen van het lokaal interstellaire veld volgens Mathis et al. (1983) (bovenaan) en enkele zwarte stralers met verschillende temperaturen (onderaan).

voor een veel-gebruikte methode (Guhathakurta and Draine 1989) die veronderstelt dat de stofdeeltjes op een continue manier afkoelen, met andere woorden dat elk uitgestraald foton slechts een kleine fractie van de interne energie van een deeltje meedraagt. Met deze veronderstelling kan de berekening gevoelig vereenvoudigd en vooral versneld worden, wat voor een stralingsoverdracht-simulatie uitermate belangrijk is, maar ze is enkel gejustifieerd in bepaalde, gelukkig veel voorkomende omstandigheden (Draine and Li 2001).

Ironisch genoeg wordt de berekening van de temperatuursverdeling heel traag of zelfs onmogelijk als het bestudeerde stofdeeltje bijna in lokaal thermisch evenwicht is, omdat in dat geval de temperatuursverdeling ontaardt in een uiterst smalle piek, wat voor numerieke onstabiliteit zorgt. De methode moet dus in de gepaste omstandigheden overschakelen van het stochastische regime, en dus de berekening van de temperatuursverdeling, naar het evenwichtsregime, waar het stofdeeltje een

temperatuur heeft die heel eenvoudig kan bepaald worden uit de energiebalans. Dit wordt geïllustreerd in Fig. 8.3, waaruit blijkt dat de optimale overgang sterk afhangt van het stralingsveld waarin het stof zich bevindt.

We publiceerden een benchmark probleem voor deze berekening, met duidelijke definities voor alle relevante inputs en met resultaten berekend door zes stralingsoverdacht-codes waaronder SKIRT. Deze benchmark kan anderen helpen om de berekening te implementeren. Nog belangrijker is dat de vergelijking tussen de resultaten van de verschillende codes ons heeft geleerd bij welke stralingsvelden en in welke golflengtegebieden we de berekeningen kunnen vertrouwen in het kader van toekomstige stralingsoverdacht-simulaties.

8.5   Importeren van de resultaten van hydrodynamische simulaties

Tijdens een Monte Carlo stralingsoverdacht-simulatie, zoals in SKIRT, zijn de noden voor de datastructuren die de stellaire en stofcomponenten voorstellen uiteraard verschillend. Voor de stralingsbronnen is de belangrijkste opdracht om een willekeurige startpositie te bepalen voor het volgende fotonpakket, bemonsterd uit de ruimtelijke lichtsterktedichtheid van het model (voor een gegeven golflengte). Dit gebeurt miljoenen of zelfs miljarden malen tijdens een simulatie, dus snelheid is van belang. Voor de stofcomponent, daarentegen, is de belangrijkste opdracht om snel de dichtheid op een gegeven locatie te bepalen, of de massa in een gegeven volume. Deze operaties worden heel intens gebruikt tijdens het opbouwen van het stofrooster, maar zijn niet meer van belang tijdens de eigenlijke simulatie aangezien het stofrooster alle nodige informatie bevat.

Hydrodynamische simulatie-codes gebruiken verschillende technieken om de ruimtelijke verdeling van de bestudeerde systeemcomponenten voor te stellen. Een eerste veel-gebruikte techniek steunt op deeltjes met een uitgesmeerde dichtheid (smoothed particles hydrodynamics of SPH, bvb., Wadsley et al. 2004; Springel 2005), en een tweede op hiërarchische roosters waarvan het verfijningsniveau automatisch aangepast wordt aan de vereisten van het model (adaptive mesh refinement of AMR, bvb., Teyssier 2002; Collins et al. 2010; Keppens et al. 2012). Een meer recente techniek gebruikt bewegende Voronoi cellen (Duffell and MacFadyen 2011; Springel 2010; Vandenbroucke and De Rijcke submitted). In elk geval kunnen deze codes "foto's" van het gesimuleerde systeem op bepaalde tijdstippen in de dynamische evolutie wegschrijven in een gegevensbestand. De structuur van deze informatie hangt uiteraard af van de gebruikte techniek, en is dikwijls ook specifiek voor elke code.

We hebben modules geïmplementeerd in SKIRT om de relevante informatie uit dergelijke foto's (of "snapshots") in te lezen en te gebruiken als stellaire of stofcomponent in de stralingsoverdacht-simulatie. We hebben ervoor gekozen om voor elke hydrodynamische techniek (SPH, AMR, of Voronoi) een eenvoudig tekstbestandsformaat vast te leggen dat gemakkelijk kan aangemaakt worden via bijvoorbeeld een Python script, alhoewel we soms ook het binaire formaat van de simulatiecode ondersteunen.

Voor de stellaire component moet naast de ruimtelijke verdeling ook het emissie-spectrum op elke positie vastgelegd worden (zie Fig. 8.1). Die informatie kan uit de snapshot gehaald worden door de eigenschappen van elk deeltje of elke cel te gebrui-ken als parameters om het spectrum te bepalen. SKIRT heeft daartoe ingebouwde sjablonen voor sterpopulaties (Bruzual and Charlot 2003) en voor stervormingsgebie-den (Groves et al. 2008).

Bij heel wat hydrodynamische simulaties wordt de stofcomponent in het interstellair medium niet apart gevolgd, en moet de stofverdeling dus afgeleid worden uit de verdeling van het koude gas. Daartoe voorziet SKIRT een eenvoudige heuristiek die ervan uitgaat dat een constante fractie van de metalen in het interstellair medium in stofdeeltjes opgesloten zit (Dwek 1998; James et al. 2002; Brinchmann et al. 2013; Zafar and Watson 2013). Daarnaast kan de gebruiker één van de ingebouwde stofmodellen kiezen om te stofeigenschappen te bepalen, die worden verondersteld overal dezelfde te zijn.

We mogen besluiten dat we hiermee een behoorlijk flexibel raamwerk gebouwd hebben om informatie van om het even welk type hydrodynamische simulatie te verwerken met minimale extra inspanningen.

| 8.6 | Stof in de EAGLE galaxieën |

Kosmologische simulaties zoals het EAGLE project (Schaye et al. 2015) kunnen ons helpen om de vorming en de evolutie van galaxieën beter te begrijpen door de overeenkomsten en verschillen te bestuderen tussen de simulatieresultaten en waar-nemingen. Voor dit proefschrift hebben we een methode ontwikkeld en gebruikt om kunstmatige "waarnemingen" te doen voor de galaxieën uit de EAGLE simulaties, rekening houdend met de effecten van interstellair stof, zodat de simulatieresultaten op een correcte manier kunnen vergeleken worden met sterrenkundige waarnemin-gen (Hayward and Smith 2015; Guidi et al. 2015). Om onze methode te testen en te ijken hebben we een aantal gesimuleerde EAGLE galaxieën gekozen met eigen-schappen vergelijkbaar met die van de galaxieën in de *Herschel* Reference Survey (HRS) (Boselli et al. 2010). Specifiek kozen we 282 willekeurige hedendaagse EAGLE galaxieën van hetzelfde type (spiraalvormig of elliptisch) en met dezelfde *K*-band lichtsterkte-verdeling als de 282 HRS galaxieën bestudeerd door Cortese et al. (2012), waarvoor zowel optische als sub-millimeter waarnemingen beschikbaar zijn.

Figuur 8.4 toont een overzicht van de procedure die we ontwikkelden om een EAGLE galaxie klaar te maken voor de SKIRT stralingsoverdracht-simulatie. Een EAGLE ga-laxie wordt voorgesteld door een verzameling door zwaartekracht gebonden deeltjes met uitgesmeerde dichtheid voor de sterren (a) en voor het gas (b). Tijdens de gesi-muleerde evolutie van een galaxie, wordt van tijd tot tijd een gasdeeltje omgevormd tot een sterdeeltje. Aangezien elk deeltje typisch miljoenen zonsmassa's voorstelt, zijn er op elk moment slechts een beperkt aantal jonge sterdeeltjes aanwezig, terwijl zo'n stervormingsgebied toch een specifieke voetafdruk nalaat in de uitgezonden
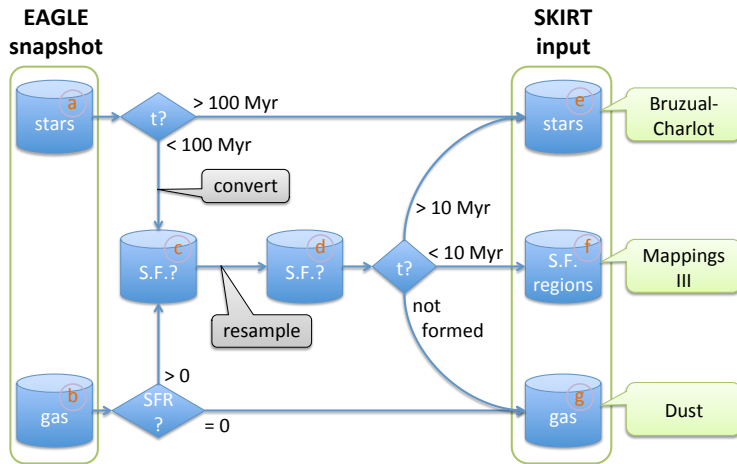
Figuur 8.4: Schematisch overzicht van de procedure om een EAGLE galaxie klaar te maken voor de SKIRT stralingsoverdracht-simulatie.

straling. Om onze kunstmatige waarnemingen realistischer te maken, geven we een speciale behandeling aan de net gevormde sterdeeltjes en de gasdeeltjes die op het punt staan om eventueel omgevormd te worden (c). Die deeltjes worden in onze methode opgedeeld in een aantal kleinere deeltjes (d) met een massa bemonsterd uit de waargenomen massa-functie van moleculaire wolken in onze Melkweg (Heyer et al. 2001), en met een willekeurige leeftijd rekening houdend met de stervormingsgraad van het moederdeeltje. De jongste van deze kleinere deeltjes brengen we onder in een nieuwe verzameling van stervormingsgebieden (f); de iets oudere deeltjes voegen we terug bij de sterdeeltjes (e), en de deeltjes die nog niet aan stervorming toegekomen zijn voegen we terug bij de gasdeeltjes (g).

Het emissiespectrum voor de geëvolueerde sterpopulaties (e) bepalen we aan de hand van de Bruzual and Charlot (2003) sjablonen, en voor de stervormingsgebieden (f) aan de hand van de Groves et al. (2008) sjablonen. In beide gevallen maken we gebruik van de eigenschappen van elk deeltje om de parameters van het sjabloon te bepalen, behalve voor de stofdekkingsgraad van de stervormingsgebieden f_{PDR} (zie verder). De stofverdeling leiden we af uit de eigenschappen van de gasdeeltjes (g), waarbij we ervan uitgaan dat de hoeveelheid stof een constante fractie f_{dust} is van de metalen in het gas, behalve voor het hete gas waarin zich geen stof kan vormen, en dat zich vooral aan de buitenkant van een galaxie bevindt. We gebruiken het Zubko et al. (2004) model voor de stralingsoverdracht-eigenschappen van het stof. Er blijven bij deze methode dus twee vrije parameters over waarvoor we zelf een vaste, optimale waarde hebben bepaald op basis van een uitgebreide parameterstudie: de stoffractie $f_{dust} = 0.3$ en de stofdekkingsgraad $f_{PDR} = 0.1$. Deze waarden zijn verenigbaar met de waarnemingen (Dwek 1998; Brinchmann et al. 2013; Zafar and Watson 2013) en met de waarden voorgesteld door anderen (Jonsson et al. 2010).

Vooraleer onze resultaten te berekenen, gingen we uiteraard ook na dat de verdere configuratie van de SKIRT-simulaties aan de eisen voldoet, zoals bijvoorbeeld het
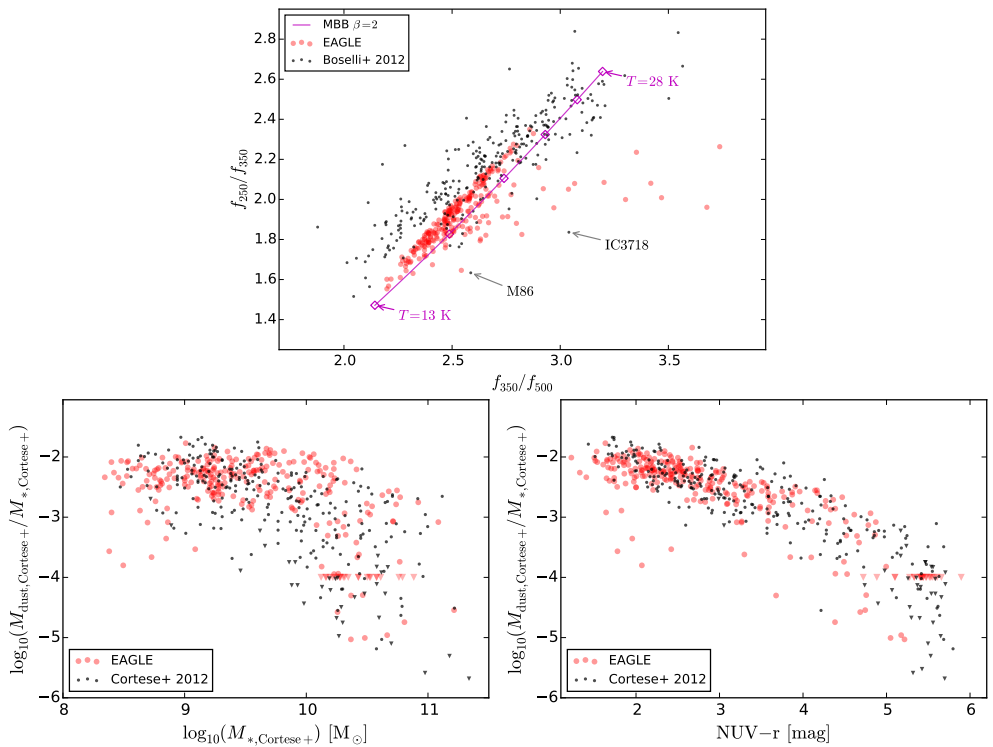
Figuur 8.5: Kunstmatige waarnemingen van onze EAGLE galaxieën (rood) vergeleken met HRS waarnemingen (zwart) uit Boselli et al. (2012) en Cortese et al. (2012). *Boven*: sub-millimeter fluxverhoudingen f_{250}/f_{350} versus f_{350}/f_{500}; de magenta ruiten geven de waarden aan voor een "grijze" straler met $\beta = 2$ voor temperaturen van 13 K tot 28 K. *Linksonder*: stof-tot-stermassaverhouding versus stermassa. *Rechtsonder*: stof-tot-stermassaverhouding versus NUV $- r$ kleur.

aantal gebruikte golflengten en de constructie van het stofrooster. In een volgende stap bevestigden we dat een aantal intrinsieke eigenschappen van onze EAGLE galaxieën, die kunnen berekend worden door de eigenschappen van de individuele deeltjes bij elkaar op te tellen, voldoen aan de verwachtingen voor het soort galaxieën dat we gekozen hebben. Hier keken we naar de stellaire massa, de stervormingsgraad, de stofmassa, en de metalliciteit.

Bij het interpreteren van de resultaten vergeleken we eerst de intrinsieke ster- en stofmassa in onze EAGLE galaxieën met dezelfde eigenschappen afgeleid uit de kunstmatige SKIRT waarnemingen met de technieken gebruikt door Cortese et al. (2012) voor de HRS galaxieën. In beide gevallen vonden we een systematisch verschil van ongeveer een factor twee. Voor de stermassa komt dit overeen met de verschillen tussen de schattingstechnieken die men in de literatuur kan terugvinden (Gallazzi and Bell 2009; Zibetti et al. 2009; Baldry et al. 2012). Voor de stofmassa heeft het verschil vooral te maken met de opaciteit die we toekennen aan het stof in het sub-mm gebied, en waarover ook onzekerheid bestaat (Zubko et al. 2004; Cortese et al. 2012). Zolang we geen uitspraken doen over de "absolute" ster- of stofmassa in onze galaxieën,

vormen deze verschillen echter geen probleem voor de verdere interpretatie van onze resultaten, aangezien we waargenomen eigenschappen enkel vergelijken met kunstmatige waarnemingen die op dezelfde manier tot stand kwamen.

Onze belangrijkste resultaten worden getoond in Fig. 8.5. Voor deze figuren gebruikten we onze "standaard" parameter waarden $f_{\text{dust}} = 0.3$ en $f_{\text{PDR}} = 0.1$. We zien dat de EAGLE galaxieën de waargenomen relaties behoorlijk benaderen. Bijvoorbeeld, de overeenkomst voor de relatie stof-tot-stermassaverhouding versus $NUV - r$ kleur (Fig. 8.5, rechtsonder) blijkt uitzonderlijk goed te zijn. Tegelijkertijd zien we dat de sub-millimeter fluxverhoudingen (Fig. 8.5, bovenaan) over het algemeen kleiner zijn dan de waargenomen verhoudingen, wat wijst op een vlakkere helling van het stofemissiespectrum bij die golflengten, wat op zijn beurt de aanwezigheid van een grotere hoeveelheid heel koud stof verraadt (zie ook de magenta temperatuurindicaties in de figuur). Volgens onze analyse wordt deze afwijking vooral verklaard door het feit dat de EAGLE simulaties de koude fase van het interstellair medium niet apart volgen, een beperking die we niet volledig kunnen corrigeren door de her-bemonstering van de stervormingsgebieden zoals hiervoor beschreven.

Niettegenstaande deze beperking mogen we besluiten dat we er in geslaagd zijn om kunstmatige waarnemingen te doen van de EAGLE galaxieën die behoorlijk overeenkomen met de waarnemingen. Bovendien heeft onze parameterstudie op basis van de infrarood en sub-millimeter fluxen de mogelijke waarden voor de vrije parameters in onze methode voldoende kunnen beperken, zodat we correcte stofhoeveelheden en -eigenschappen aan de EAGLE galaxieën kunnen toekennen. Dit laat ons toe om een correcte demping door stof te berekenen in het UV en optische bereik, zoals dat door Trayford et al. (in prep) ondernomen wordt.

8.7	Besluit

Voor dit proefschrift hebben we een algemeen bruikbaar raamwerk ontwikkeld om kunstmatige waarnemingen te doen voor de resultaten van om het even welke hydrodynamische simulatie. We hebben dit raamwerk met succes toegepast om de stof-gerelateerde infrarood en sub-millimeter eigenschappen van een aantal EAGLE galaxieën te vergelijken met sterrenkundige waarnemingen. Op die manier hebben we onze methode kunnen ijken aan de realiteit.

Tijdens dit werk zijn er ook een aantal domeinen aan het licht gekomen waar verder onderzoek nuttig zou zijn. Bijvoorbeeld, we kunnen de stofroosters in stralingsoverdracht-simulatie codes nog verder verbeteren, en de correcte berekening van de emissie van kleine stofdeeltjes uitbreiden naar meer extreme stralingsomgevingen. Ons raamwerk om kunstmatige waarnemingen te doen kan nog verder gestroomlijnd worden, en toegepast worden om bijvoorbeeld de eigenschappen van de EAGLE galaxieën op hoge roodverschuiving te bestuderen. Op langere termijn zullen er ongetwijfeld simulatietechnieken ontwikkeld worden die de fysische processen nog beter vatten, en we hopen dat dit werk daartoe een bijdrage kan leveren.

Publications

WRITING IS NOT NECESSARILY SOMETHING TO BE ASHAMED OF
− BUT DO IT IN PRIVATE AND WASH YOUR HANDS AFTERWARDS.

In the following tables, the chapter references indicate the closest relationship of each publication with the contents of this thesis.

First-author publications

Subject	Ch.	Authors
Voronoi grids	3	Camps, Baes, and Saftly (2013b)
Code design	2	Camps and Baes (2015)
Stochastic heating	4	Camps, Misselt, Bianchi, Lunttila, Pinte, Natale, Juvela, Fischera, Fitzgerald, Gordon, Baes, and Steinacker (2015)
Dust in EAGLE	6	Camps, Trayford, Baes, Theuns, Schaller, and Schaye (submitted)

Co-authored publications

Subject	Ch.	Authors
FitSKIRT	1	De Geyter, Baes, Fritz, and Camps (2013)
Hierarchical grids	3	Saftly, Camps, Baes, Gordon, Vandewoude, Rahimi, and Stalevski (2013)
Hierarchical grids	3	Saftly, Baes, and Camps (2014)
CALIFA galaxies	1	De Geyter, Baes, Camps, Fritz, De Looze, Hughes, Viaene, and Gentile (2014)
M51 spiral	1	De Looze, Fritz, Baes, Bendo, Cortese, Boquien, Boselli, Camps, Cooray, Cormier, Davies, De Geyter, Hughes, Jones, Karczewski, Lebouteiller, Lu, Madden, Rémy-Ruyer, Spinoglio, Smith, Viaene, and Wilson (2014)

Subject	Ch.	Authors
EAGLE simulation	6	Schaye, Crain, Bower, Furlong, Schaller, Theuns, Dalla Vecchia, Frenk, McCarthy, Helly, Jenkins, Rosas-Guevara, White, Baes, Booth, Camps, Navarro, Qu, Rahmati, Sawala, Thomas, and Trayford (2015)
Molecular clouds	5	Hendrix, Keppens, and Camps (2015)
Energy balance	5	Saftly, Baes, De Geyter, Camps, Renaud, Guedes, and De Looze (2015)
Algol binaries	5	Deschamps, Braun, Jorissen, Siess, Baes, and Camps (2015)
Energy balance	1	De Geyter, Baes, De Looze, Bendo, Bourne, Camps, Cooray, De Zotti, Dunne, Dye, Eales, Fritz, Furlanetto, Gentile, Hughes, Ivison, Maddox, Michałowski, Smith, Valiante, and Viaene (2015)
Decorators	2	Baes and Camps (2015)
EAGLE database	6	McAlpine, Helly, Schaller, Trayford, Qu, Furlong, Bower, Crain, Schaye, Theuns, Dalla Vecchia, Frenk, McCarthy, Jenkins, Rosas-Guevara, White, Baes, Camps, and Lemson (2016)
Composite biasing	2	Baes, Gordon, Lunttila, Bianchi, Camps, Juvela, and Kuiper (2016)
IC 2531 edge-on	1	Mosenkov, Allaert, Baes, Bianchi, Camps, De Geyter, De Looze, Fritz, Gentile, Hughes, Lewis, Verstappen, Verstocken, and Viaene (in press)
Wolf-Rayet binaries	5	Keppens et al. (in press)
Optical EAGLE	6	Trayford et al. (in prep)
M31 spiral	1	Viaene et al. (in prep)
Parallelization	2	Verstocken et al. (in prep)
Polarization	2	Peest et al. (in prep)

Conference proceedings

Subject	Ch.	Authors
SKIRT introduction	1	Camps, Baes, De Geyter, Saftly, and De Looze (2013a)

Bibliography

D. M. Acreman, T. J. Harries, and D. A. Rundle. Modelling circumstellar discs with three-dimensional radiation hydrodynamics. *MNRAS*, 403:1143–1155, April 2010. doi: 10.1111/j.1365-2966.2009.16199.x.

O. Agertz, B. Moore, J. Stadel, D. Potter, F. Miniati, J. Read, L. Mayer, A. Gawryszczak, A. Kravtsov, Å. Nordlund, F. Pearce, V. Quilis, D. Rudd, V. Springel, J. Stone, E. Tasker, R. Teyssier, J. Wadsley, and R. Walder. Fundamental differences between SPH and grid methods. *MNRAS*, 380:963–978, September 2007. doi: 10.1111/j.1365-2966.2007.12183.x.

C. Allende Prieto, D. L. Lambert, and M. Asplund. The Forbidden Abundance of Oxygen in the Sun. *ApJ*, 556:L63–L66, July 2001. doi: 10.1086/322874.

G. Altay and T. Theuns. URCHIN: a reverse ray tracer for astrophysical applications. *MNRAS*, 434:748–764, September 2013. doi: 10.1093/mnras/stt1067.

G. Aniano, B. T. Draine, D. Calzetti, D. A. Dale, C. W. Engelbracht, K. D. Gordon, L. K. Hunt, R. C. Kennicutt, O. Krause, A. K. Leroy, H.-W. Rix, H. Roussel, K. Sandstrom, M. Sauvage, F. Walter, L. Armus, A. D. Bolatto, A. Crocker, J. Donovan Meyer, M. Galametz, G. Helou, J. Hinz, B. D. Johnson, J. Koda, E. Montiel, E. J. Murphy, R. Skibba, J.-D. T. Smith, and M. G. Wolfire. Modeling Dust and Starlight in Galaxies Observed by Spitzer and Herschel: NGC 628 and NGC 6946. *ApJ*, 756:138, September 2012. doi: 10.1088/0004-637X/756/2/138.

M. Baes and P. Camps. SKIRT: The design of a suite of input models for Monte Carlo radiative transfer simulations. *Astronomy and Computing*, 12:33–44, September 2015. doi: 10.1016/j.ascom.2015.05.006.

M. Baes and H. Dejonghe. Radiative transfer in disc galaxies - II. The influence of scattering and geometry on the attenuation curve. *MNRAS*, 326:733–744, September 2001. doi: 10.1046/j.1365-8711.2001.04626.x.

M. Baes, J. I. Davies, H. Dejonghe, S. Sabatini, S. Roberts, R. Evans, S. M. Linder, R. M. Smith, and W. J. G. de Blok. Radiative transfer in disc galaxies - III. The observed kinematics of dusty disc galaxies. *MNRAS*, 343:1081–1094, August 2003. doi: 10.1046/j.1365-8711.2003.06770.x.

M. Baes, J. Fritz, D. A. Gadotti, D. J. B. Smith, L. Dunne, E. da Cunha, A. Amblard, R. Auld, G. J. Bendo, D. Bonfield, D. Burgarella, S. Buttiglione, A. Cava, D. Clements, A. Cooray, A. Dariush, G. de Zotti, S. Dye, S. Eales, D. Frayer, J. Gonzalez-Nuevo, D. Herranz, E. Ibar, R. Ivison, G. Lagache, L. Leeuw, M. Lopez-Caniego, M. Jarvis, S. Maddox, M. Negrello, M. Michałowski, E. Pascale, M. Pohlen, E. Rigby, G. Rodighiero, S. Samui, S. Serjeant, P. Temi, M. Thompson, P. van der Werf, A. Verma, and C. Vlahakis. Herschel-ATLAS: The dust energy balance in the edge-on spiral galaxy UGC 4754. *A&A*, 518:L39, July 2010. doi: 10.1051/0004-6361/201014644.

M. Baes, J. Verstappen, I. De Looze, J. Fritz, W. Saftly, E. Vidal Pérez, M. Stalevski, and S. Valcke. Efficient Three-dimensional NLTE Dust Radiative Transfer with SKIRT. *ApJS*, 196:22, October 2011. doi: 10.1088/0067-0049/196/2/22.

M. Baes, K. D. Gordon, T. Lunttila, S. Bianchi, P. Camps, M. Juvela, and R. Kuiper. Composite biasing in Monte Carlo radiative transfer. *A&A*, 590:A55, May 2016. doi: 10.1051/0004-6361/201528063.

Y. M. Bahé, R. A. Crain, G. Kauffmann, R. G. Bower, J. Schaye, M. Furlong, C. Lagos, M. Schaller, J. W. Trayford, C. Dalla Vecchia, and T. Theuns. The distribution of atomic hydrogen in EAGLE galaxies: morphologies, profiles, and H I holes. *MNRAS*, 456:1115–1136, February 2016. doi: 10.1093/mnras/stv2674.

I. K. Baldry, S. P. Driver, J. Loveday, E. N. Taylor, L. S. Kelvin, J. Liske, P. Norberg, A. S. G. Robotham, S. Brough, A. M. Hopkins, S. P. Bamford, J. A. Peacock, J. Bland-Hawthorn, C. J. Conselice, S. M. Croom, D. H. Jones, H. R. Parkinson, C. C. Popescu, M. Prescott, R. G. Sharp, and R. J. Tuffs. Galaxy And Mass Assembly (GAMA): the galaxy stellar mass function at z ¡ 0.06. *MNRAS*, 421:621–634, March 2012. doi: 10.1111/j.1365-2966.2012.20340.x.

A. Bauer and V. Springel. Subsonic turbulence in smoothed particle hydrodynamics and moving-mesh simulations. *MNRAS*, 423:2558–2578, July 2012. doi: 10.1111/j.1365-2966.2012.21058.x.

G. J. Bendo, M. Baes, S. Bianchi, M. Boquien, A. Boselli, A. Cooray, L. Cortese, I. De Looze, S. di Serego Alighieri, J. Fritz, G. Gentile, T. M. Hughes, N. Lu, C. Pappalardo, M. W. L. Smith, L. Spinoglio, S. Viaene, and C. Vlahakis. The identification of dust heating mechanisms in nearby galaxies using Herschel 160/250 and 250/350 μm surface brightness ratios. *MNRAS*, 448:135–167, March 2015. doi: 10.1093/mnras/stu1841.

T. J. Bethell, E. G. Zweibel, F. Heitsch, and J. S. Mathis. Dust Heating by the Interstellar Radiation Field in Models of Turbulent Molecular Clouds. *ApJ*, 610:801–812, August 2004. doi: 10.1086/421902.

S. Bianchi. Dust extinction and emission in a clumpy galactic disk. An application of the radiative transfer code TRADING. *A&A*, 490:461–475, October 2008. doi: 10.1051/0004-6361:200810027.

S. Bianchi, A. Ferrara, J. I. Davies, and P. B. Alton. Effects of clumping on the observed properties of dusty galaxies. *MNRAS*, 311:601–610, January 2000. doi: 10.1046/j.1365-8711.2000.03113.x.

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

A. Boselli, S. Eales, L. Cortese, G. Bendo, P. Chanial, V. Buat, J. Davies, R. Auld, E. Rigby, M. Baes, M. Barlow, J. Bock, M. Bradford, N. Castro-Rodriguez, S. Charlot, D. Clements, D. Cormier, E. Dwek, D. Elbaz, M. Galametz, F. Galliano, W. Gear, J. Glenn, H. Gomez, M. Griffin, S. Hony, K. Isaak, L. Levenson, N. Lu, S. Madden, B. O'Halloran, K. Okamura, S. Oliver, M. Page, P. Panuzzo, A. Papageorgiou, T. Parkin, I. Perez-Fournon, M. Pohlen, N. Rangwala, H. Roussel, A. Rykala, N. Sacchi, M. Sauvage, B. Schulz, M. Schirm, M. W. L. Smith, L. Spinoglio, J. Stevens, M. Symeonidis, M. Vaccari, L. Vigroux, C. Wilson, H. Wozniak, G. Wright, and W. Zeilinger. The Herschel Reference Survey. *PASP*, 122:261–287, March 2010. doi: 10.1086/651535.

A. Boselli, L. Ciesla, L. Cortese, V. Buat, M. Boquien, G. J. Bendo, S. Boissier, S. Eales, G. Gavazzi,

T. M. Hughes, M. Pohlen, M. W. L. Smith, M. Baes, S. Bianchi, D. L. Clements, A. Cooray, J. Davies, W. Gear, S. Madden, L. Magrini, P. Panuzzo, A. Remy, L. Spinoglio, and S. Zibetti. Far-infrared colours of nearby late-type galaxies in the Herschel Reference Survey. *A&A*, 540:A54, April 2012. doi: 10.1051/0004-6361/201118602.

F. Boulanger and M. Perault. Diffuse infrared emission from the galaxy. I - Solar neighborhood. *ApJ*, 330:964–985, July 1988. doi: 10.1086/166526.

N. Bourne, S. J. Maddox, L. Dunne, R. Auld, M. Baes, I. K. Baldry, D. G. Bonfield, A. Cooray, S. M. Croom, A. Dariush, G. de Zotti, S. P. Driver, S. Dye, S. Eales, H. L. Gomez, J. González-Nuevo, A. M. Hopkins, E. Ibar, M. J. Jarvis, A. Lapi, B. Madore, M. J. Michałowski, M. Pohlen, C. C. Popescu, E. E. Rigby, M. Seibert, D. J. B. Smith, R. J. Tuffs, P. van der Werf, S. Brough, S. Buttiglione, A. Cava, D. L. Clements, C. J. Conselice, J. Fritz, R. Hopwood, R. J. Ivison, D. H. Jones, L. S. Kelvin, J. Liske, J. Loveday, P. Norberg, A. S. G. Robotham, G. Rodighiero, and P. Temi. Herschel-ATLAS/GAMA: a census of dust in optically selected galaxies from stacking at submillimetre wavelengths. *MNRAS*, 421:3027–3059, April 2012. doi: 10.1111/j.1365-2966.2012.20528.x.

C. Brinch and M. R. Hogerheijde. LIME - a flexible, non-LTE line excitation and radiation transfer method for millimeter and far-infrared wavelengths. *A&A*, 523:A25, November 2010. doi: 10.1051/0004-6361/201015333.

J. Brinchmann, S. Charlot, G. Kauffmann, T. Heckman, S. D. M. White, and C. Tremonti. Estimating gas masses and dust-to-gas ratios from optical spectroscopy. *MNRAS*, 432: 2112–2140, July 2013. doi: 10.1093/mnras/stt551.

T. M. Brown, H. C. Ferguson, A. F. Davidsen, and B. Dorman. A Far-Ultraviolet Analysis of the Stellar Populations in Six Elliptical and S0 Galaxies. *ApJ*, 482:685–707, June 1997.

T. M. Brown, H. C. Ferguson, E. Smith, C. W. Bowers, R. A. Kimble, A. Renzini, and R. M. Rich. Far-Ultraviolet Emission from Elliptical Galaxies at z = 0.33. *ApJ*, 584:L69–L72, February 2003. doi: 10.1086/374035.

G. Bruzual and S. Charlot. Stellar population synthesis at the resolution of 2003. *MNRAS*, 344: 1000–1028, October 2003. doi: 10.1046/j.1365-8711.2003.06897.x.

G. L. Bryan, M. L. Norman, B. W. O'Shea, T. Abel, J. H. Wise, M. J. Turk, D. R. Reynolds, D. C. Collins, P. Wang, S. W. Skillman, B. Smith, R. P. Harkness, J. Bordner, J.-h. Kim, M. Kuhlen, H. Xu, N. Goldbaum, C. Hummels, A. G. Kritsuk, E. Tasker, S. Skory, C. M. Simpson, O. Hahn, J. S. Oishi, G. C. So, F. Zhao, R. Cen, Y. Li, and Enzo Collaboration. ENZO: An Adaptive Mesh Refinement Code for Astrophysics. *ApJS*, 211:19, April 2014. doi: 10.1088/0067-0049/211/2/19.

Y. I. Byun, K. C. Freeman, and N. D. Kylafis. Diagnostics of dust content in spiral galaxies: Numerical simulations of radiative transfer. *ApJ*, 432:114–127, September 1994. doi: 10.1086/174553.

P. Camps and M. Baes. SKIRT: An advanced dust radiative transfer code with a user-friendly architecture. *Astronomy and Computing*, 9:20–33, March 2015. doi: 10.1016/j.ascom.2014.10.004.

P. Camps, M. Baes, G. De Geyter, W. Saftly, and I. De Looze. Three-dimensional Continuum Radiative Transfer Simulations of Dusty Systems. In N. V. Pogorelov, E. Audit, and G. P.

Zank, editors, *Numerical Modeling of Space Plasma Flows (ASTRONUM2012)*, volume 474 of *Astronomical Society of the Pacific Conference Series*, page 53, April 2013a.

P. Camps, M. Baes, and W. Saftly. Using 3D Voronoi grids in radiative transfer simulations. *A&A*, 560:A35, December 2013b. doi: 10.1051/0004-6361/201322281.

P. Camps, K. Misselt, S. Bianchi, T. Lunttila, C. Pinte, G. Natale, M. Juvela, J. Fischera, M. P. Fitzgerald, K. Gordon, M. Baes, and J. Steinacker. Benchmarking the calculation of stochastic heating and emissivity of dust grains in the context of radiative transfer simulations. *A&A*, 580:A87, August 2015. doi: 10.1051/0004-6361/201525998.

Peter Camps, James W. Trayford, M. Baes, Tom Theuns, Matthieu Schaller, and Joop Schaye. Far-infrared and dust properties of present-day galaxies in the eagle simulations. *MNRAS*, submitted.

M. Cappellari. Efficient multi-Gaussian expansion of galaxies. *MNRAS*, 333:400–410, June 2002. doi: 10.1046/j.1365-8711.2002.05412.x.

E. D. Cashwell and C. J. Everett. *A practical manual on the Monte Carlo method for random walk problems*. International tracts in computer science and technology and their application. Oxford: Pergamon Press, 1959.

B. Catinella, D. Schiminovich, G. Kauffmann, S. Fabello, J. Wang, C. Hummels, J. Lemonias, S. M. Moran, R. Wu, R. Giovanelli, M. P. Haynes, T. M. Heckman, A. R. Basu-Zych, M. R. Blanton, J. Brinchmann, T. Budavári, T. Gonçalves, B. D. Johnson, R. C. Kennicutt, B. F. Madore, C. D. Martin, M. R. Rich, L. J. Tacconi, D. A. Thilker, V. Wild, and T. K. Wyder. The GALEX Arecibo SDSS Survey - I. Gas fraction scaling relations of massive galaxies and first data release. *MNRAS*, 403:683–708, April 2010. doi: 10.1111/j.1365-2966.2009.16180.x.

G. Chabrier. The Galactic Disk Mass Function: Reconciliation of the Hubble Space Telescope and Nearby Determinations. *ApJ*, 586:L133–L136, April 2003. doi: 10.1086/374879.

S. Chakrabarti, T. J. Cox, L. Hernquist, P. F. Hopkins, B. Robertson, and T. Di Matteo. Feedback-driven Evolution of the Far-Infrared Spectral Energy Distributions of Luminous and Ultraluminous Infrared Galaxies. *ApJ*, 658:840–850, April 2007. doi: 10.1086/510113.

Su-shing Chen. *Mapping and Spatial Modelling for Navigation*, chapter A Spherical Model for Navigation and Spatial Reasoning, pages 59–72. Springer Berlin Heidelberg, Berlin, Heidelberg, 1990. ISBN 978-3-642-84215-3. doi: 10.1007/978-3-642-84215-3_4. URL http://dx.doi.org/10.1007/978-3-642-84215-3_4.

L. Ciesla, A. Boselli, M. W. L. Smith, G. J. Bendo, L. Cortese, S. Eales, S. Bianchi, M. Boquien, V. Buat, J. Davies, M. Pohlen, S. Zibetti, M. Baes, A. Cooray, I. De Looze, S. di Serego Alighieri, M. Galametz, H. L. Gomez, V. Lebouteiller, S. C. Madden, C. Pappalardo, A. Remy, L. Spinoglio, M. Vaccari, R. Auld, and D. L. Clements. Submillimetre photometry of 323 nearby galaxies from the Herschel Reference Survey. *A&A*, 543:A161, July 2012. doi: 10.1051/0004-6361/201219216.

L. Ciotti and G. Bertin. Analytical properties of the $R^{1/m}$ law. *A&A*, 352:447–451, December 1999.

M. Cohen, W. A. Wheaton, and S. T. Megeath. Spectral Irradiance Calibration in the Infrared. XIV. The Absolute Calibration of 2MASS. *AJ*, 126:1090–1096, August 2003. doi: 10.1086/376474.

D. C. Collins, H. Xu, M. L. Norman, H. Li, and S. Li. Cosmological Adaptive Mesh Refinement Magnetohydrodynamics with Enzo. *ApJS*, 186:308–333, February 2010. doi: 10.1088/0067-0049/186/2/308.

M. Compiègne, L. Verstraete, A. Jones, J.-P. Bernard, F. Boulanger, N. Flagey, J. Le Bourlot, D. Paradis, and N. Ysard. The global dust SED: tracing the nature and evolution of dust with DustEM. *A&A*, 525:A103, January 2011. doi: 10.1051/0004-6361/201015292.

L. Cortese, B. Catinella, S. Boissier, A. Boselli, and S. Heinis. The effect of the environment on the H I scaling relations. *MNRAS*, 415:1797–1806, August 2011. doi: 10.1111/j.1365-2966.2011.18822.x.

L. Cortese, L. Ciesla, A. Boselli, S. Bianchi, H. Gomez, M. W. L. Smith, G. J. Bendo, S. Eales, M. Pohlen, M. Baes, E. Corbelli, J. I. Davies, T. M. Hughes, L. K. Hunt, S. C. Madden, D. Pierini, S. di Serego Alighieri, S. Zibetti, M. Boquien, D. L. Clements, A. Cooray, M. Galametz, L. Magrini, C. Pappalardo, L. Spinoglio, and C. Vlahakis. The dust scaling relations of the Herschel Reference Survey. *A&A*, 540:A52, April 2012. doi: 10.1051/0004-6361/201118499.

R. A. Crain, J. Schaye, R. G. Bower, M. Furlong, M. Schaller, T. Theuns, C. Dalla Vecchia, C. S. Frenk, I. G. McCarthy, J. C. Helly, A. Jenkins, Y. M. Rosas-Guevara, S. D. M. White, and J. W. Trayford. The EAGLE simulations of galaxy formation: calibration of subgrid physics and model variations. *MNRAS*, 450:1937–1961, June 2015. doi: 10.1093/mnras/stv725.

C. Dalla Vecchia and J. Schaye. Simulating galactic outflows with thermal supernova feedback. *MNRAS*, 426:140–158, October 2012. doi: 10.1111/j.1365-2966.2012.21704.x.

G. De Geyter, M. Baes, J. Fritz, and P. Camps. FitSKIRT: genetic algorithms to automatically fit dusty galaxies with a Monte Carlo radiative transfer code. *A&A*, 550:A74, February 2013. doi: 10.1051/0004-6361/201220126.

G. De Geyter, M. Baes, P. Camps, J. Fritz, I. De Looze, T. M. Hughes, S. Viaene, and G. Gentile. The distribution of interstellar dust in CALIFA edge-on galaxies via oligochromatic radiative transfer fitting. *MNRAS*, 441:869–885, June 2014. doi: 10.1093/mnras/stu612.

G. De Geyter, M. Baes, I. De Looze, G. J. Bendo, N. Bourne, P. Camps, A. Cooray, G. De Zotti, L. Dunne, S. Dye, S. A. Eales, J. Fritz, C. Furlanetto, G. Gentile, T. M. Hughes, R. J. Ivison, S. J. Maddox, M. J. Michałowski, M. W. L. Smith, E. Valiante, and S. Viaene. Dust energy balance study of two edge-on spiral galaxies in the Herschel-ATLAS survey. *MNRAS*, 451: 1728–1739, August 2015. doi: 10.1093/mnras/stv1104.

I. De Looze, M. Baes, G. J. Bendo, L. Ciesla, L. Cortese, G. de Geyter, B. Groves, M. Boquien, A. Boselli, L. Brondeel, A. Cooray, S. Eales, J. Fritz, F. Galliano, G. Gentile, K. D. Gordon, S. Hony, K.-H. Law, S. C. Madden, M. Sauvage, M. W. L. Smith, L. Spinoglio, and J. Verstappen. The dust energy balance in the edge-on spiral galaxy NGC 4565. *MNRAS*, 427: 2797–2811, December 2012a. doi: 10.1111/j.1365-2966.2012.22045.x.

I. De Looze, M. Baes, J. Fritz, and J. Verstappen. Panchromatic radiative transfer modelling of stars and dust in the Sombrero galaxy. *MNRAS*, 419:895–903, January 2012b. doi: 10.1111/j.1365-2966.2011.19759.x.

I. De Looze, J. Fritz, M. Baes, G. J. Bendo, L. Cortese, M. Boquien, A. Boselli, P. Camps, A. Cooray, D. Cormier, J. I. Davies, G. De Geyter, T. M. Hughes, A. P. Jones, O. Ł. Karczewski, V. Lebouteiller, N. Lu, S. C. Madden, A. Rémy-Ruyer, L. Spinoglio, M. W. L. Smith, S. Viaene,

and C. D. Wilson. High-resolution, 3D radiative transfer modeling. I. The grand-design spiral galaxy M 51. *A&A*, 571:A69, November 2014. doi: 10.1051/0004-6361/201424747.

L. Decin, N. L. J. Cox, P. Royer, A. J. Van Marle, B. Vandenbussche, D. Ladjal, F. Kerschbaum, R. Ottensamer, M. J. Barlow, J. A. D. L. Blommaert, H. L. Gomez, M. A. T. Groenewegen, T. Lim, B. M. Swinyard, C. Waelkens, and A. G. G. M. Tielens. The enigmatic nature of the circumstellar envelope and bow shock surrounding Betelgeuse as revealed by Herschel. I. Evidence of clumps, multiple arcs, and a linear bar-like structure. *A&A*, 548:A113, December 2012. doi: 10.1051/0004-6361/201219792.

W. Dehnen. A Family of Potential-Density Pairs for Spherical Galaxies and Bulges. *MNRAS*, 265:250, November 1993.

H. Dejonghe. A completely analytical family of anisotropic Plummer models. *MNRAS*, 224:13–39, January 1987.

Boris Delaunay. Sur la sphère vide (a la mémoire de georges voronoi). *Classe des Sciences Mathématiques et Naturelles*, 7(6):793–800, 1934.

Mathieu Desbrun and Marie-Paule Gascuel. Smoothed particles: A new paradigm for animating highly deformable bodies. In *Proceedings of the Eurographics Workshop on Computer Animation and Simulation '96*, pages 61–76, New York, NY, USA, 1996. Springer-Verlag New York, Inc. ISBN 3-211-82885-0. URL http://dl.acm.org/citation.cfm?id=274976.274981.

R. Deschamps, K. Braun, A. Jorissen, L. Siess, M. Baes, and P. Camps. Non-conservative evolution in Algols: where is the matter? *A&A*, 577:A55, May 2015. doi: 10.1051/0004-6361/201424772.

F. X. Desert, F. Boulanger, and S. N. Shore. Grain temperature fluctuations - A key to infrared spectra. *A&A*, 160:295–300, May 1986.

S. di Serego Alighieri, S. Bianchi, C. Pappalardo, S. Zibetti, R. Auld, M. Baes, G. Bendo, E. Corbelli, J. I. Davies, T. Davis, I. De Looze, J. Fritz, G. Gavazzi, C. Giovanardi, M. Grossi, L. K. Hunt, L. Magrini, D. Pierini, and E. M. Xilouris. The Herschel Virgo Cluster Survey. XIII. Dust in early-type galaxies. *A&A*, 552:A8, April 2013. doi: 10.1051/0004-6361/201220551.

Lejeune Dirichlet. Über die reduction der positiven quadratischen formen mit drei unbestimmten ganzen zahlen. *Journal für die reine und angewandte Mathematik*, 40:209–227, 1850.

M. Disney, J. Davies, and S. Phillipps. Are galaxy discs optically thick? *MNRAS*, 239:939–976, August 1989.

M. Doi, M. Tanaka, M. Fukugita, J. E. Gunn, N. Yasuda, Ž. Ivezić, J. Brinkmann, E. de Haars, S. J. Kleinman, J. Krzesinski, and R. French Leger. Photometric Response Functions of the Sloan Digital Sky Survey Imager. *AJ*, 139:1628–1648, April 2010. doi: 10.1088/0004-6256/139/4/1628.

K. Dolag and F. Stasyszyn. An MHD GADGET for cosmological simulations. *MNRAS*, 398:1678–1697, October 2009. doi: 10.1111/j.1365-2966.2009.15181.x.

K. Dolag, S. Borgani, G. Murante, and V. Springel. Substructures in hydrodynamical cluster simulations. *MNRAS*, 399:497–514, October 2009. doi: 10.1111/j.1365-2966.2009.15034.x.

S. D. Doty, R. A. Metzler, and M. L. Palotti. Effects of clumping on temperature - I. Externally heated clouds. *MNRAS*, 362:737–747, September 2005. doi: 10.1111/j.1365-2966.2005.09363.x.

B. T. Draine. The discrete-dipole approximation and its application to interstellar graphite grains. *ApJ*, 333:848–872, October 1988. doi: 10.1086/166795.

B. T. Draine. Scattering by Interstellar Dust Grains. II. X-Rays. *ApJ*, 598:1026–1037, December 2003. doi: 10.1086/379123.

B. T. Draine and H. M. Lee. Optical properties of interstellar graphite and silicate grains. *ApJ*, 285:89–108, October 1984. doi: 10.1086/162480.

B. T. Draine and A. Li. Infrared Emission from Interstellar Dust. I. Stochastic Heating of Small Grains. *ApJ*, 551:807–824, April 2001. doi: 10.1086/320227.

B. T. Draine and A. Li. Infrared Emission from Interstellar Dust. IV. The Silicate-Graphite-PAH Model in the Post-Spitzer Era. *ApJ*, 657:810–837, March 2007. doi: 10.1086/511055.

B. T. Draine, D. A. Dale, G. Bendo, K. D. Gordon, J. D. T. Smith, L. Armus, C. W. Engelbracht, G. Helou, R. C. Kennicutt, Jr., A. Li, H. Roussel, F. Walter, D. Calzetti, J. Moustakas, E. J. Murphy, G. H. Rieke, C. Bot, D. J. Hollenbach, K. Sheth, and H. I. Teplitz. Dust Masses, PAH Abundances, and Starlight Intensities in the SINGS Galaxy Sample. *ApJ*, 663:866–894, July 2007. doi: 10.1086/518306.

P. C. Duffell and A. I. MacFadyen. TESS: A Relativistic Hydrodynamics Code on a Moving Voronoi Mesh. *ApJS*, 197:15, December 2011. doi: 10.1088/0067-0049/197/2/15.

E. Dwek. Temperature fluctuations and infrared emission from dust particles in a hot gas. *ApJ*, 302:363–370, March 1986. doi: 10.1086/163995.

E. Dwek. The Evolution of the Elemental Abundances in the Gas and Dust Phases of the Galaxy. *ApJ*, 501:643, July 1998. doi: 10.1086/305829.

H. Edelsbrunner, F.P. Preparata, and D.B. West. Tetrahedrizing point sets in three dimensions. *Journal of Symbolic Computation*, 10(3–4):335 – 347, 1990. ISSN 0747-7171. doi: http://dx.doi.org/10.1016/S0747-7171(08)80068-5. URL http://www.sciencedirect.com/science/article/pii/S0747717108800685.

J. Einasto. On the Construction of a Composite Model for the Galaxy and on the Determination of the System of Galactic Parameters. *Trudy Astrofizicheskogo Instituta Alma-Ata*, 5:87–100, 1965.

A. Einstein. Die Grundlage der allgemeinen Relativitätstheorie. *Annalen der Physik*, 354:769–822, 1916. doi: 10.1002/andp.19163540702.

E. Emsellem, G. Monnet, and R. Bacon. The multi-gaussian expansion method: a tool for building realistic photometric and kinematical models of stellar systems I. The formalism. *A&A*, 285:723–738, May 1994.

S. Fabello, B. Catinella, R. Giovanelli, G. Kauffmann, M. P. Haynes, T. M. Heckman, and D. Schiminovich. ALFALFA H I data stacking - I. Does the bulge quench ongoing star formation in early-type galaxies? *MNRAS*, 411:993–1012, February 2011. doi: 10.1111/j.1365-2966.2010.17742.x.

D. Fabian, T. Henning, C. Jäger, H. Mutschke, J. Dorschner, and O. Wehrhan. Steps toward interstellar silicate mineralogy. VI. Dependence of crystalline olivine IR spectra on iron content and particle shape. *A&A*, 378:228–238, October 2001. doi: 10.1051/0004-6361:20011196.

C. Fallscheer, M. A. Reid, J. Di Francesco, P. G. Martin, T. Hill, M. Hennemann, Q. Nguyen-Luong, F. Motte, A. Men'shchikov, P. André, D. Ward-Thompson, M. Griffin, J. Kirk, V. Konyves, K. L. J. Rygl, S. Sadavoy, M. Sauvage, N. Schneider, L. D. Anderson, M. Benedettini, J.-P. Bernard, S. Bontemps, A. Ginsburg, S. Molinari, D. Polychroni, A. Rivera-Ingraham, H. Roussel, L. Testi, G. White, J. P. Williams, C. D. Wilson, M. Wong, and A. Zavagno. Herschel Reveals Massive Cold Clumps in NGC 7538. *ApJ*, 773:102, August 2013. doi: 10.1088/0004-637X/773/2/102.

L. Fanciullo, V. Guillet, G. Aniano, A. P. Jones, N. Ysard, M.-A. Miville-Deschênes, F. Boulanger, and M. Köhler. Dust models post-Planck: constraining the far-infrared opacity of dust in the diffuse interstellar medium. *A&A*, 580:A136, August 2015. doi: 10.1051/0004-6361/201525677.

G. Fasano and A. Franceschini. A multidimensional version of the Kolmogorov-Smirnov test. *MNRAS*, 225:155–170, March 1987. doi: 10.1093/mnras/225.1.155.

R. Feldmann and L. Mayer. The Argo simulation - I. Quenching of massive galaxies at high redshift as a result of cosmological starvation. *MNRAS*, 446:1939–1956, January 2015. doi: 10.1093/mnras/stu2207.

M. Fioc and B. Rocca-Volmerange. PEGASE: a UV to NIR spectral evolution model of galaxies. Application to the calibration of bright galaxy counts. *A&A*, 326:950–962, October 1997.

Joerg Fischera. *Infrarotemission der SN 1987A nach 11 Jahren*. PhD thesis, Heidelberg, 2000.

Jerome H. Friedman, Jon Louis Bentley, and Raphael Ari Finkel. An algorithm for finding best matches in logarithmic expected time. *ACM Trans. Math. Softw.*, 3(3):209–226, September 1977. ISSN 0098-3500. doi: 10.1145/355744.355745. URL http://doi.acm.org/10.1145/355744.355745.

J. Fritz, G. Gentile, M. W. L. Smith, W. K. Gear, R. Braun, J. R. Duval, G. J. Bendo, M. Baes, S. A. Eales, J. Verstappen, J. A. D. L. Blommaert, M. Boquien, A. Boselli, D. Clements, A. R. Cooray, L. Cortese, I. De Looze, G. P. Ford, F. Galliano, H. L. Gomez, K. D. Gordon, V. Lebouteiller, B. O'Halloran, J. Kirk, S. C. Madden, M. J. Page, A. Remy, H. Roussel, L. Spinoglio, D. Thilker, M. Vaccari, C. D. Wilson, and C. Waelkens. The Herschel Exploitation of Local Galaxy Andromeda (HELGA). I. Global far-infrared and sub-mm morphology. *A&A*, 546:A34, October 2012. doi: 10.1051/0004-6361/201118619.

S. Fromang, P. Hennebelle, and R. Teyssier. A high order Godunov scheme with constrained transport and adaptive mesh refinement for astrophysical magnetohydrodynamics. *A&A*, 457:371–384, October 2006. doi: 10.1051/0004-6361:20065371.

B. Fryxell, K. Olson, P. Ricker, F. X. Timmes, M. Zingale, D. Q. Lamb, P. MacNeice, R. Rosner, J. W. Truran, and H. Tufo. Flash: An adaptive mesh hydrodynamics code for modeling astrophysical thermonuclear flashes. *The Astrophysical Journal Supplement Series*, 131(1):273, 2000. URL http://stacks.iop.org/0067-0049/131/i=1/a=273.

M. Furlong, R. G. Bower, T. Theuns, J. Schaye, R. A. Crain, M. Schaller, C. Dalla Vecchia, C. S. Frenk, I. G. McCarthy, J. Helly, A. Jenkins, and Y. M. Rosas-Guevara. Evolution of galaxy stellar masses and star formation rates in the EAGLE simulations. *MNRAS*, 450:4486–4504, July 2015. doi: 10.1093/mnras/stv852.

A. Gallazzi and E. F. Bell. Stellar Mass-to-Light Ratios from Galaxy Spectra: How Accurate Can They Be? *ApJS*, 185:253–272, December 2009. doi: 10.1088/0067-0049/185/2/253.

Erich Gamma, Richard Helm, Ralph Johnson, and John Vlissides. *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley Professional, 1st edition, 11 1994.

P. F. Goldsmith, M. Heyer, G. Narayanan, R. Snell, D. Li, and C. Brunt. Large-Scale Structure of the Molecular Gas in Taurus Revealed by High Linear Dynamic Range Spectral Line Mapping. *ApJ*, 680:428–445, June 2008. doi: 10.1086/587166.

H. L. Gomez, M. Baes, L. Cortese, M. W. L. Smith, A. Boselli, L. Ciesla, G. J. Bendo, M. Pohlen, S. di Serego Alighieri, R. Auld, M. J. Barlow, J. J. Bock, M. Bradford, V. Buat, N. Castro-Rodriguez, P. Chanial, S. Charlot, D. L. Clements, A. Cooray, D. Cormier, J. I. Davies, E. Dwek, S. Eales, D. Elbaz, M. Galametz, F. Galliano, W. K. Gear, J. Glenn, M. Griffin, S. Hony, K. G. Isaak, L. R. Levenson, N. Lu, S. Madden, B. O'Halloran, K. Okumura, S. Oliver, M. J. Page, P. Panuzzo, A. Papageorgiou, T. J. Parkin, I. Perez-Fournon, N. Rangwala, E. E. Rigby, H. Roussel, A. Rykala, N. Sacchi, M. Sauvage, M. R. P. Schirm, B. Schulz, L. Spinoglio, S. Srinivasan, J. A. Stevens, M. Symeonidis, M. Trichas, M. Vaccari, L. Vigroux, C. D. Wilson, H. Wozniak, G. S. Wright, and W. W. Zeilinger. The dust morphology of the elliptical Galaxy M 86 with SPIRE. *A&A*, 518:L45, July 2010. doi: 10.1051/0004-6361/201014530.

K. D. Gordon, K. A. Misselt, A. N. Witt, and G. C. Clayton. The DIRTY Model. I. Monte Carlo Radiative Transfer through Dust. *ApJ*, 551:269–276, April 2001. doi: 10.1086/320082.

G. L. Granato and L. Danese. Thick Tori around Active Galactic Nuclei - a Comparison of Model Predictions with Observations of the Infrared Continuum and Silicate Features. *MNRAS*, 268:235, May 1994.

T. H. Greif, V. Springel, S. D. M. White, S. C. O. Glover, P. C. Clark, R. J. Smith, R. S. Klessen, and V. Bromm. Simulations on a Moving Mesh: The Clustered Formation of Population III Protostars. *ApJ*, 737:75, August 2011. doi: 10.1088/0004-637X/737/2/75.

M. J. Griffin, A. Abergel, A. Abreu, P. A. R. Ade, P. André, J.-L. Augueres, T. Babbedge, Y. Bae, T. Baillie, J.-P. Baluteau, M. J. Barlow, G. Bendo, D. Benielli, J. J. Bock, P. Bonhomme, D. Brisbin, C. Brockley-Blatt, M. Caldwell, C. Cara, N. Castro-Rodriguez, R. Cerulli, P. Chanial, S. Chen, E. Clark, D. L. Clements, L. Clerc, J. Coker, D. Communal, L. Conversi, P. Cox, D. Crumb, C. Cunningham, F. Daly, G. R. Davis, P. de Antoni, J. Delderfield, N. Devin, A. di Giorgio, I. Didschuns, K. Dohlen, M. Donati, A. Dowell, C. D. Dowell, L. Duband, L. Dumaye, R. J. Emery, D. Ferlet, D. Ferrand, J. Fontignie, M. Fox, A. Franceschini, M. Frerking, T. Fulton, J. Garcia, R. Gastaud, W. K. Gear, J. Glenn, A. Goizel, D. K. Griffin, T. Grundy, S. Guest, L. Guillemet, P. C. Hargrave, M. Harwit, P. Hastings, E. Hatziminaoglou, M. Herman, B. Hinde, V. Hristov, M. Huang, P. Imhof, K. J. Isaak, U. Israelsson, R. J. Ivison, D. Jennings, B. Kiernan, K. J. King, A. E. Lange, W. Latter, G. Laurent, P. Laurent, S. J. Leeks, E. Lellouch, L. Levenson, B. Li, J. Li, J. Lilienthal, T. Lim, S. J. Liu, N. Lu, S. Madden, G. Mainetti, P. Marliani, D. McKay, K. Mercier, S. Molinari, H. Morris, H. Moseley, J. Mulder, M. Mur, D. A. Naylor, H. Nguyen, B. O'Halloran, S. Oliver, G. Olofsson, H.-G. Olofsson, R. Orfei, M. J. Page, I. Pain, P. Panuzzo, A. Papageorgiou, G. Parks, P. Parr-Burman, A. Pearce, C. Pearson, I. Pérez-Fournon, F. Pinsard, G. Pisano, J. Podosek, M. Pohlen, E. T. Polehampton, D. Pouliquen, D. Rigopoulou, D. Rizzo, I. G. Roseboom, H. Roussel, M. Rowan-Robinson, B. Rownd, P. Saraceno, M. Sauvage, R. Savage, G. Savini, E. Sawyer, C. Scharmberg, D. Schmitt, N. Schneider, B. Schulz, A. Schwartz, R. Shafer, D. L. Shupe, B. Sibthorpe, S. Sidher, A. Smith, A. J. Smith, D. Smith, L. Spencer, B. Stobie, R. Sudiwala, K. Sukhatme, C. Surace, J. A. Stevens, B. M. Swinyard, M. Trichas, T. Tourette, H. Triou, S. Tseng, C. Tucker, A. Turner, M. Vaccari, I. Valtchanov, L. Vigroux, E. Virique, G. Voellmer, H. Walker, R. Ward, T. Waskett, M. Weilert, R. Wesson, G. J. White, N. Whitehouse, C. D.

Wilson, B. Winter, A. L. Woodcraft, G. S. Wright, C. K. Xu, A. Zavagno, M. Zemcov, L. Zhang, and E. Zonca. The Herschel-SPIRE instrument and its in-flight performance. *A&A*, 518:L3, July 2010. doi: 10.1051/0004-6361/201014519.

B. Groves, M. A. Dopita, R. S. Sutherland, L. J. Kewley, J. Fischera, C. Leitherer, B. Brandl, and W. van Breugel. Modeling the Pan-Spectral Energy Distribution of Starburst Galaxies. IV. The Controlling Parameters of the Starburst SED. *ApJS*, 176:438–456, June 2008. doi: 10.1086/528711.

J. Guedes, S. Callegari, P. Madau, and L. Mayer. Forming Realistic Late-type Spirals in a ΛCDM Universe: The Eris Simulation. *ApJ*, 742:76, 2011.

P. Guhathakurta and B. T. Draine. Temperature fluctuations in interstellar grains. I - Computational method and sublimation of small grains. *ApJ*, 345:230–244, October 1989. doi: 10.1086/167899.

G. Guidi, C. Scannapieco, and C. J. Walcher. Biases and systematics in the observational derivation of galaxy properties: comparing different techniques on synthetic observations of simulated galaxies. *MNRAS*, 454:2381–2400, December 2015. doi: 10.1093/mnras/stv2050.

Antonin Guttman. R-trees: a dynamic index structure for spatial searching. *SIGMOD Rec.*, 14 (2):47–57, June 1984. ISSN 0163-5808. doi: 10.1145/971697.602266. URL http://doi.acm.org/10.1145/971697.602266.

F. Haardt and P. Madau. Modelling the UV/X-ray cosmic background with CUBA. In D. M. Neumann and J. T. V. Tran, editors, *Clusters of Galaxies and the High Redshift Universe Observed in X-rays*, 2001.

C.-N. Hao, R. C. Kennicutt, B. D. Johnson, D. Calzetti, D. A. Dale, and J. Moustakas. Dust-corrected Star Formation Rates of Galaxies. II. Combinations of Ultraviolet and Infrared Tracers. *ApJ*, 741:124, November 2011. doi: 10.1088/0004-637X/741/2/124.

T. J. Harries, J. D. Monnier, N. H. Symington, and R. Kurosawa. Three-dimensional dust radiative-transfer models: the Pinwheel Nebula of WR 104. *MNRAS*, 350:565–574, May 2004. doi: 10.1111/j.1365-2966.2004.07668.x.

C. C. Hayward and D. J. B. Smith. Should we believe the results of ultraviolet-millimetre galaxy spectral energy distribution modelling? *MNRAS*, 446:1512–1535, January 2015. doi: 10.1093/mnras/stu2195.

C. C. Hayward, D. Kereš, P. Jonsson, D. Narayanan, T. J. Cox, and L. Hernquist. What Does a Submillimeter Galaxy Selection Actually Select? The Dependence of Submillimeter Flux Density on Star Formation Rate and Dust Mass. *ApJ*, 743:159, December 2011. doi: 10.1088/0004-637X/743/2/159.

G. Helou, N. Y. Lu, M. W. Werner, S. Malhotra, and N. Silbermann. The Mid-Infrared Spectra of Normal Galaxies. *ApJ*, 532:L21–L24, March 2000. doi: 10.1086/312549.

T. Hendrix, R. Keppens, and P. Camps. Modelling ripples in Orion with coupled dust dynamics and radiative transfer. *A&A*, 575:A110, March 2015. doi: 10.1051/0004-6361/201425498.

M. H. Heyer, J. M. Carpenter, and R. L. Snell. The Equilibrium State of Molecular Regions in the Outer Galaxy. *ApJ*, 551:852–866, April 2001. doi: 10.1086/320218.

F. Heymann and R. Siebenmorgen. GPU-based Monte Carlo Dust Radiative Transfer Scheme Applied to Active Galactic Nuclei. *ApJ*, 751:27, May 2012. doi: 10.1088/0004-637X/751/1/27.

P. F. Hopkins, D. Kereš, J. Oñorbe, C.-A. Faucher-Giguère, E. Quataert, N. Murray, and J. S. Bullock. Galaxies on FIRE (Feedback In Realistic Environments): stellar feedback explains cosmologically inefficient star formation. *MNRAS*, 445:581–603, November 2014. doi: 10.1093/mnras/stu1738.

D. A. Hubber, C. P. Batty, A. McLeod, and A. P. Whitworth. SEREN - a new SPH code for star and planet formation simulations. Algorithms and tests. *A&A*, 529:A27, May 2011. doi: 10.1051/0004-6361/201014949.

E. P. Hubble. NGC 6822, a remote stellar system. *ApJ*, 62, December 1925. doi: 10.1086/142943.

T. M. Hughes, L. Cortese, A. Boselli, G. Gavazzi, and J. I. Davies. The role of cold gas and environment on the stellar mass-metallicity relation of nearby galaxies. *A&A*, 550:A115, February 2013. doi: 10.1051/0004-6361/201218822.

L. K. Hunt, B. T. Draine, S. Bianchi, K. D. Gordon, G. Aniano, D. Calzetti, D. A. Dale, G. Helou, J. L. Hinz, R. C. Kennicutt, H. Roussel, C. D. Wilson, A. Bolatto, M. Boquien, K. V. Croxall, M. Galametz, A. Gil de Paz, J. Koda, J. C. Muñoz-Mateos, K. M. Sandstrom, M. Sauvage, L. Vigroux, and S. Zibetti. Cool dust heating and temperature mixing in nearby star-forming galaxies. *A&A*, 576:A33, April 2015. doi: 10.1051/0004-6361/201424734.

G. D. Illingworth, D. Magee, P. A. Oesch, R. J. Bouwens, I. Labbé, M. Stiavelli, P. G. van Dokkum, M. Franx, M. Trenti, C. M. Carollo, and V. Gonzalez. The HST eXtreme Deep Field (XDF): Combining All ACS and WFC3/IR Data on the HUDF Region into the Deepest Field Ever. *ApJS*, 209:6, November 2013. doi: 10.1088/0067-0049/209/1/6.

R. Indebetouw, B. A. Whitney, K. E. Johnson, and K. Wood. Three-dimensional Models of Embedded High-Mass Stars: Effects of a Clumpy Circumstellar Medium. *ApJ*, 636:362–380, January 2006. doi: 10.1086/497886.

C. Jaeger, F. J. Molster, J. Dorschner, T. Henning, H. Mutschke, and L. B. F. M. Waters. Steps toward interstellar silicate mineralogy. IV. The crystalline revolution. *A&A*, 339:904–916, November 1998.

A. James, L. Dunne, S. Eales, and M. G. Edmunds. SCUBA observations of galaxies with metallicity measurements: a new method for determining the relation between submillimetre luminosity and dust mass. *MNRAS*, 335:753–761, September 2002. doi: 10.1046/j.1365-8711.2002.05660.x.

A. P. Jones, L. Fanciullo, M. Köhler, L. Verstraete, V. Guillet, M. Bocchio, and N. Ysard. The evolution of amorphous hydrocarbons in the ISM: dust modelling from a new vantage point. *A&A*, 558:A62, October 2013. doi: 10.1051/0004-6361/201321686.

P. Jonsson. SUNRISE: polychromatic dust radiative transfer in arbitrary geometries. *MNRAS*, 372:2–20, October 2006. doi: 10.1111/j.1365-2966.2006.10884.x.

P. Jonsson, B. A. Groves, and T. J. Cox. High-resolution panchromatic spectral models of galaxies including photoionization and dust. *MNRAS*, 403:17–44, March 2010. doi: 10.1111/j.1365-2966.2009.16087.x.

M. Juvela. Efficient Monte Carlo methods for continuum radiative transfer. *A&A*, 440:531–546, September 2005. doi: 10.1051/0004-6361:20042615.

M. Juvela and P. Padoan. Dust emission from inhomogeneous interstellar clouds: Radiative

transfer in 3D with transiently heated particles. *A&A*, 397:201–212, January 2003. doi: 10.1051/0004-6361:20021433.

M. Juvela, J. Malinen, and T. Lunttila. Profiles of interstellar cloud filaments. Observational effects in synthetic sub-millimetre observations. *A&A*, 544:A141, August 2012. doi: 10.1051/0004-6361/201219558.

R. C. Kennicutt and N. J. Evans. Star Formation in the Milky Way and Nearby Galaxies. *ARA&A*, 50:531–608, September 2012. doi: 10.1146/annurev-astro-081811-125610.

R. C. Kennicutt, Jr. The Global Schmidt Law in Star-forming Galaxies. *ApJ*, 498:541–552, May 1998. doi: 10.1086/305588.

R. Keppens, Z. Meliani, A.J. van Marle, P. Delmont, A. Vlasis, and B. van der Holst. Parallel, grid-adaptive approaches for relativistic hydro and magnetohydrodynamics. *Journal of Computational Physics*, 231(3):718 – 744, 2012. ISSN 0021-9991. doi: http://dx.doi.org/10.1016/j.jcp.2011.01.020. URL http://www.sciencedirect.com/science/article/pii/S0021999111000386.

Keppens et al. Pinwheels in the sky, with dust: 3d modeling of the wolf-rayet 98a environment. *MNRAS*, in press.

D. Kereš, M. Vogelsberger, D. Sijacki, V. Springel, and L. Hernquist. Moving-mesh cosmology: characteristics of galaxies and haloes. *MNRAS*, 425:2027–2048, September 2012. doi: 10.1111/j.1365-2966.2012.21548.x.

H. Kobayashi, S.-i. Watanabe, H. Kimura, and T. Yamamoto. Dust ring formation due to sublimation of dust grains drifting radially inward by the Poynting-Robertson drag: An analytical model. *Icarus*, 201:395–405, May 2009. doi: 10.1016/j.icarus.2009.01.002.

A. N. Kolmogorov. Sulla determinazione empirica di una legge di distribuzione. *Giornale dell'Istituto Italiano degli Attuari*, 4:83–91, 1933.

R. Kurosawa and D. J. Hillier. Tree-structured grid model of line and polarization variability from massive binaries. *A&A*, 379:336–346, November 2001. doi: 10.1051/0004-6361:20011233.

R. L. Kurucz. Model Atmospheres (Kurucz, 1979). *VizieR Online Data Catalog*, 6039:0, October 1993.

A. Laor and B. T. Draine. Spectroscopic constraints on the properties of dust in active galactic nuclei. *ApJ*, 402:441–468, January 1993. doi: 10.1086/172149.

P. Laursen, A. O. Razoumov, and J. Sommer-Larsen. Lyα Radiative Transfer in Cosmological Simulations Using Adaptive Mesh Refinement. *ApJ*, 696:853–869, May 2009. doi: 10.1088/0004-637X/696/1/853.

A. M. C. Le Brun, I. G. McCarthy, J. Schaye, and T. J. Ponman. Towards a realistic population of simulated galaxy groups and clusters. *MNRAS*, 441:1270–1290, June 2014. doi: 10.1093/mnras/stu608.

C. Leitherer, D. Schaerer, J. D. Goldader, R. M. G. Delgado, C. Robert, D. F. Kune, D. F. de Mello, D. Devost, and T. M. Heckman. Starburst99: Synthesis Models for Galaxies with Active Star Formation. *ApJS*, 123:3–40, July 1999. doi: 10.1086/313233.

A. Li and B. T. Draine. Infrared Emission from Interstellar Dust. II. The Diffuse Interstellar Medium. *ApJ*, 554:778–802, June 2001. doi: 10.1086/323147.

S.H. Lo. Parallel delaunay triangulation in three dimensions. *Computer Methods in Applied Mechanics and Engineering*, 237–240(0):88 – 106, 2012. ISSN 0045-7825. doi: http://dx.doi.org/10.1016/j.cma.2012.05.009. URL http://www.sciencedirect.com/science/article/pii/S0045782512001570.

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

T. Lunttila and M. Juvela. Radiative transfer on hierarchial grids. *A&A*, 544:A52, August 2012. doi: 10.1051/0004-6361/201219220.

V. Manske and T. Henning. Two-dimensional radiative transfer with transiently heated particles: methods and applications. *A&A*, 337:85–95, September 1998.

V. Manske and T. Henning. 2D radiative transfer with transiently heated particles for the circumstellar environment of Herbig Ae/Be stars. *A&A*, 349:907–911, September 1999.

C. Maraston. Evolutionary synthesis of stellar populations: a modular tool. *MNRAS*, 300: 872–892, November 1998. doi: 10.1046/j.1365-8711.1998.01947.x.

F. Marinacci, R. Pakmor, and V. Springel. The formation of disc galaxies in high-resolution moving-mesh cosmological simulations. *MNRAS*, 437:1750–1775, January 2014. doi: 10.1093/mnras/stt2003.

J. S. Mathis, W. Rumpl, and K. H. Nordsieck. The size distribution of interstellar grains. *ApJ*, 217:425–433, October 1977. doi: 10.1086/155591.

J. S. Mathis, P. G. Mezger, and N. Panagia. Interstellar radiation field and dust temperatures in the diffuse interstellar matter and in giant molecular clouds. *A&A*, 128:212–229, November 1983.

D. J. Mavriplis. Unstructured Grid Techniques. *Annual Review of Fluid Mechanics*, 29:473–514, 1997. doi: 10.1146/annurev.fluid.29.1.473.

Nelson Max. Consistent subdivision of convex polyhedra into tetrahedra. *Journal of Graphics Tools*, 6(3):29–36, 2001. doi: 10.1080/10867651.2001.10487543. URL http://dx.doi.org/10.1080/10867651.2001.10487543.

S. McAlpine, J. C. Helly, M. Schaller, J. W. Trayford, Y. Qu, M. Furlong, R. G. Bower, R. A. Crain, J. Schaye, T. Theuns, C. Dalla Vecchia, C. S. Frenk, I. G. McCarthy, A. Jenkins, Y. Rosas-Guevara, S. D. M. White, M. Baes, P. Camps, and G. Lemson. The EAGLE simulations of galaxy formation: Public release of halo and galaxy catalogues. *Astronomy and Computing*, 15:72–89, April 2016. doi: 10.1016/j.ascom.2016.02.004.

R. McKinnon, P. Torrey, and M. Vogelsberger. Dust Formation in Milky Way-like Galaxies. *MNRAS*, February 2016. doi: 10.1093/mnras/stw253.

M. Min, J. W. Hovenier, and A. de Koter. Modeling optical properties of cosmic dust grains using a distribution of hollow spheres. *A&A*, 432:909–920, March 2005. doi: 10.1051/0004-6361:20041920.

A. Misiriotis, N. D. Kylafis, J. Papamastorakis, and E. M. Xilouris. Is the exponential distribution a good approximation of dusty galactic disks? *A&A*, 353:117–123, January 2000.

K. A. Misselt, K. D. Gordon, G. C. Clayton, and M. J. Wolff. The DIRTY Model. II. Self-consistent

Treatment of Dust Heating and Emission in a Three-dimensional Radiative Transfer Code. *ApJ*, 551:277–293, April 2001. doi: 10.1086/320083.

J. J. Monaghan and J. C. Lattanzio. A refined particle method for astrophysical problems. *A&A*, 149:135–143, August 1985.

P. Morrissey, T. Conrow, T. A. Barlow, T. Small, M. Seibert, T. K. Wyder, T. Budavári, S. Arnouts, P. G. Friedman, K. Forster, D. C. Martin, S. G. Neff, D. Schiminovich, L. Bianchi, J. Donas, T. M. Heckman, Y.-W. Lee, B. F. Madore, B. Milliard, R. M. Rich, A. S. Szalay, B. Y. Welsh, and S. K. Yi. The Calibration and Data Products of GALEX. *ApJS*, 173:682–697, December 2007. doi: 10.1086/520512.

G. M. Morton. *A computer oriented geodetic data base and a new technique in file sequencing*. Ottawa, Canada: International Business Machines Co., 1966.

A. V. Mosenkov, F. Allaert, M. Baes, S. Bianchi, P. Camps, G. De Geyter, I. De Looze, J. Fritz, G. Gentile, T. M. Hughes, F. Lewis, J. Verstappen, S. Verstocken, and S. Viaene. HERschel Observations of Edge-on Spirals (HEROES). III. Dust energy balance study of IC 2531. *A&A*, May in press.

Matthias Müller, David Charypar, and Markus Gross. Particle-based fluid simulation for interactive applications. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, SCA '03, pages 154–159, Aire-la-Ville, Switzerland, Switzerland, 2003. Eurographics Association. ISBN 1-58113-659-5. URL http://dl.acm.org/citation.cfm?id=846276.846298.

E. J. Murphy, J. J. Condon, E. Schinnerer, R. C. Kennicutt, D. Calzetti, L. Armus, G. Helou, J. L. Turner, G. Aniano, P. Beirão, A. D. Bolatto, B. R. Brandl, K. V. Croxall, D. A. Dale, J. L. Donovan Meyer, B. T. Draine, C. Engelbracht, L. K. Hunt, C.-N. Hao, J. Koda, H. Roussel, R. Skibba, and J.-D. T. Smith. Calibrating Extinction-free Star Formation Rate Diagnostics with 33 GHz Free-free Emission in NGC 6946. *ApJ*, 737:67, August 2011. doi: 10.1088/0004-637X/737/2/67.

G. Natale, C. C. Popescu, R. J. Tuffs, and D. Semionov. DART-RAY: a 3D ray-tracing radiative transfer code for calculating the propagation of light in dusty galaxies. *MNRAS*, 438:3137–3162, March 2014. doi: 10.1093/mnras/stt2418.

D. Nelson, M. Vogelsberger, S. Genel, D. Sijacki, D. Kereš, V. Springel, and L. Hernquist. Moving mesh cosmology: tracing cosmological gas accretion. *MNRAS*, 429:3353–3370, March 2013. doi: 10.1093/mnras/sts595.

H. Netzer. Quasar discs. II - A composite model for the broad-line region. *MNRAS*, 225:55–72, March 1987.

H. T. Nguyen, B. Schulz, L. Levenson, A. Amblard, V. Arumugam, H. Aussel, T. Babbedge, A. Blain, J. Bock, A. Boselli, V. Buat, N. Castro-Rodriguez, A. Cava, P. Chanial, E. Chapin, D. L. Clements, A. Conley, L. Conversi, A. Cooray, C. D. Dowell, E. Dwek, S. Eales, D. Elbaz, M. Fox, A. Franceschini, W. Gear, J. Glenn, M. Griffin, M. Halpern, E. Hatziminaoglou, E. Ibar, K. Isaak, R. J. Ivison, G. Lagache, N. Lu, S. Madden, B. Maffei, G. Mainetti, L. Marchetti, G. Marsden, J. Marshall, B. O'Halloran, S. J. Oliver, A. Omont, M. J. Page, P. Panuzzo, A. Papageorgiou, C. P. Pearson, I. Perez Fournon, M. Pohlen, N. Rangwala, D. Rigopoulou, D. Rizzo, I. G. Roseboom, M. Rowan-Robinson, D. Scott, N. Seymour, D. L. Shupe, A. J. Smith, J. A. Stevens, M. Symeonidis, M. Trichas, K. E. Tugwell, M. Vaccari, I. Valtchanov,

L. Vigroux, L. Wang, R. Ward, D. Wiebe, G. Wright, C. K. Xu, and M. Zemcov. HerMES: The SPIRE confusion limit. *A&A*, 518:L5, July 2010. doi: 10.1051/0004-6361/201014680.

G. Niccolini and J. Alcolea. 3D continuum radiative transfer. A new adaptive grid construction algorithm based on the Monte Carlo method. *A&A*, 456:1–12, September 2006. doi: 10.1051/0004-6361:20054410.

G. Niccolini, P. Woitke, and B. Lopez. High precision Monte Carlo radiative transfer in dusty media. *A&A*, 399:703–716, February 2003. doi: 10.1051/0004-6361:20021761.

B. D. Oppenheimer, R. A. Crain, J. Schaye, A. Rahmati, A. J. Richings, J. W. Trayford, J. Tumlinson, R. G. Bower, M. Schaller, and T. Theuns. Bimodality of low-redshift circumgalactic O VI in non-equilibrium EAGLE zoom simulations. *MNRAS*, May 2016. doi: 10.1093/mnras/stw1066.

J.-P. Paardekooper, C. J. H. Kruip, and V. Icke. SimpleX2: radiative transfer on an unstructured, dynamic grid. *A&A*, 515:A79, June 2010. doi: 10.1051/0004-6361/200913821.

R. Pakmor, P. Edelmann, F. K. Röpke, and W. Hillebrandt. Stellar GADGET: a smoothed particle hydrodynamics code for stellar astrophysics and its application to Type Ia supernovae from white dwarf mergers. *MNRAS*, 424:2222–2231, August 2012. doi: 10.1111/j.1365-2966.2012.21383.x.

T. Park, S.-H. Lee, and C.-H. Kim. S-octree: An extension to spherical coordinates. *Journal of Korea Multimedia Society*, 13(12):12, 12 2010.

S. Paron, W. Weidmann, M. E. Ortega, J. F. Albacete Colombo, and A. Pichel. The interstellar medium and the massive stellar content towards the SNR G18.1-0.1 and neighbouring H II regions. *MNRAS*, 433:1619–1627, August 2013. doi: 10.1093/mnras/stt837.

I. Pascucci, S. Wolf, J. Steinacker, C. P. Dullemond, T. Henning, G. Niccolini, P. Woitke, and B. Lopez. The 2D continuum radiative transfer problem. Benchmark results for disk configurations. *A&A*, 417:793–805, April 2004. doi: 10.1051/0004-6361:20040017.

J. A. Peacock. Two-dimensional goodness-of-fit testing in astronomy. *MNRAS*, 202:615–627, February 1983. doi: 10.1093/mnras/202.3.615.

Peest et al. Polarization Signatures in 3D Radiative Transfer Simulations. *A&A*, in prep.

V.-M. Pelkonen, M. Juvela, and P. Padoan. Predictions of polarized dust emission from interstellar clouds: spatial variations in the efficiency of radiative torque alignment. *A&A*, 502:833–844, August 2009. doi: 10.1051/0004-6361/200811549.

C. Pinte, F. Ménard, G. Duchêne, and P. Bastien. Monte Carlo radiative transfer in protoplanetary disks. *A&A*, 459:797–804, December 2006. doi: 10.1051/0004-6361:20053275.

C. Pinte, T. J. Harries, M. Min, A. M. Watson, C. P. Dullemond, P. Woitke, F. Ménard, and M. C. Durán-Rojas. Benchmark problems for continuum radiative transfer. High optical depths, anisotropic scattering, and polarisation. *A&A*, 498:967–980, May 2009. doi: 10.1051/0004-6361/200811555.

H. C. Plummer. On the problem of distribution in globular star clusters. *MNRAS*, 71:460–470, March 1911.

A. Poglitsch, C. Waelkens, N. Geis, H. Feuchtgruber, B. Vandenbussche, L. Rodriguez, O. Krause, E. Renotte, C. van Hoof, P. Saraceno, J. Cepa, F. Kerschbaum, P. Agnèse, B. Ali, B. Altieri, P. An-

dreani, J.-L. Augueres, Z. Balog, L. Barl, O. H. Bauer, N. Belbachir, M. Benedettini, N. Billot, O. Boulade, H. Bischof, J. Blommaert, E. Callut, C. Cara, R. Cerulli, D. Cesarsky, A. Contursi, Y. Creten, W. De Meester, V. Doublier, E. Doumayrou, L. Duband, K. Exter, R. Genzel, J.-M. Gillis, U. Grözinger, T. Henning, J. Herreros, R. Huygen, M. Inguscio, G. Jakob, C. Jamar, C. Jean, J. de Jong, R. Katterloher, C. Kiss, U. Klaas, D. Lemke, D. Lutz, S. Madden, B. Marquet, J. Martignac, A. Mazy, P. Merken, F. Montfort, L. Morbidelli, T. Müller, M. Nielbock, K. Okumura, R. Orfei, R. Ottensamer, S. Pezzuto, P. Popesso, J. Putzeys, S. Regibo, V. Reveret, P. Royer, M. Sauvage, J. Schreiber, J. Stegmaier, D. Schmitt, J. Schubert, E. Sturm, M. Thiel, G. Tofani, R. Vavrek, M. Wetzstein, E. Wieprecht, and E. Wiezorrek. The Photodetector Array Camera and Spectrometer (PACS) on the Herschel Space Observatory. *A&A*, 518:L2, July 2010. doi: 10.1051/0004-6361/201014535.

C. C. Popescu, A. Misiriotis, N. D. Kylafis, R. J. Tuffs, and J. Fischera. Modelling the spectral energy distribution of galaxies. I. Radiation fields and grain heating in the edge-on spiral NGC 891. *A&A*, 362:138–150, October 2000.

C. C. Popescu, R. J. Tuffs, M. A. Dopita, J. Fischera, N. D. Kylafis, and B. F. Madore. Modelling the spectral energy distribution of galaxies. V. The dust and PAH emission SEDs of disk galaxies. *A&A*, 527:A109, March 2011. doi: 10.1051/0004-6361/201015217.

L. Č. Popović, P. Jovanović, M. Stalevski, S. Anton, A. H. Andrei, J. Kovačević, and M. Baes. Photocentric variability of quasars caused by variations in their inner structure: consequences for Gaia measurements. *A&A*, 538:A107, February 2012. doi: 10.1051/0004-6361/201117245.

W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical recipes in C++ : the art of scientific computing*. Cambridge University Press, third edition, 2007.

F. Renaud, F. Bournaud, E. Emsellem, B. Elmegreen, R. Teyssier, J. Alves, D. Chapon, F. Combes, A. Dekel, J. Gabor, P. Hennebelle, and K. Kraljic. A sub-parsec resolution simulation of the Milky Way: global structure of the interstellar medium and properties of molecular clouds. *MNRAS*, 436:1836, 2013.

E. Retana-Montenegro, E. van Hese, G. Gentile, M. Baes, and F. Frutos-Alfaro. Analytical properties of Einasto dark matter haloes. *A&A*, 540:A70, April 2012. doi: 10.1051/0004-6361/201118543.

G. H. Rieke, E. T. Young, C. W. Engelbracht, D. M. Kelly, F. J. Low, E. E. Haller, J. W. Beeman, K. D. Gordon, J. A. Stansberry, K. A. Misselt, J. Cadien, J. E. Morrison, G. Rivlis, W. B. Latter, A. Noriega-Crespo, D. L. Padgett, K. R. Stapelfeldt, D. C. Hines, E. Egami, J. Muzerolle, A. Alonso-Herrero, M. Blaylock, H. Dole, J. L. Hinz, E. Le Floc'h, C. Papovich, P. G. Pérez-González, P. S. Smith, K. Y. L. Su, L. Bennett, D. T. Frayer, D. Henderson, N. Lu, F. Masci, M. Pesenson, L. Rebull, J. Rho, J. Keene, S. Stolovy, S. Wachter, W. Wheaton, M. W. Werner, and P. L. Richards. The Multiband Imaging Photometer for Spitzer (MIPS). *ApJS*, 154:25–29, September 2004. doi: 10.1086/422717.

G. H. Rieke, A. Alonso-Herrero, B. J. Weiner, P. G. Pérez-González, M. Blaylock, J. L. Donley, and D. Marcillac. Determining Star Formation Rates for Infrared Galaxies. *ApJ*, 692:556–573, February 2009. doi: 10.1088/0004-637X/692/1/556.

T. P. Robitaille. HYPERION: an open-source parallelized three-dimensional dust continuum radiative transfer code. *A&A*, 536:A79, December 2011. doi: 10.1051/0004-6361/201117150.

C. Rocchini and P. Cignoni. Generating random points in a tetrahedron. *Journal of Graphics*

Tools, 5(4):9–12, 2000. doi: 10.1080/10867651.2000.10487528. URL http://dx.doi.org/10.1080/10867651.2000.10487528.

Y. M. Rosas-Guevara, R. G. Bower, J. Schaye, M. Furlong, C. S. Frenk, C. M. Booth, R. A. Crain, C. Dalla Vecchia, M. Schaller, and T. Theuns. The impact of angular momentum on black hole accretion rates in simulations of galaxy formation. *MNRAS*, 454:1038–1057, November 2015. doi: 10.1093/mnras/stv2056.

Chris H. Rycroft. Voro++: A three-dimensional voronoi cell library in c++. *Chaos*, 19:041111; doi: 10.1063/1.3215722, 2009.

W. Saftly, P. Camps, M. Baes, K. D. Gordon, S. Vandewoude, A. Rahimi, and M. Stalevski. Using hierarchical octrees in Monte Carlo radiative transfer simulations. *A&A*, 554:A10, June 2013. doi: 10.1051/0004-6361/201220854.

W. Saftly, M. Baes, and P. Camps. Hierarchical octree and k-d tree grids for 3D radiative transfer simulations. *A&A*, 561:A77, January 2014. doi: 10.1051/0004-6361/201322593.

W. Saftly, M. Baes, G. De Geyter, P. Camps, F. Renaud, J. Guedes, and I. De Looze. Large and small-scale structures and the dust energy balance problem in spiral galaxies. *A&A*, 576:A31, April 2015. doi: 10.1051/0004-6361/201425445.

T. Sawala, C. S. Frenk, A. Fattahi, J. F. Navarro, R. G. Bower, R. A. Crain, C. Dalla Vecchia, M. Furlong, J. C. Helly, A. Jenkins, K. A. Oman, M. Schaller, J. Schaye, T. Theuns, J. Trayford, and S. D. M. White. The APOSTLE simulations: solutions to the Local Group's cosmic puzzles. *MNRAS*, 457:1931–1943, April 2016. doi: 10.1093/mnras/stw145.

M. Schaller, C. Dalla Vecchia, J. Schaye, R. G. Bower, T. Theuns, R. A. Crain, M. Furlong, and I. G. McCarthy. The EAGLE simulations of galaxy formation: the importance of the hydrodynamics scheme. *MNRAS*, 454:2277–2291, December 2015. doi: 10.1093/mnras/stv2169.

M. Schartmann, K. Meisenheimer, M. Camenzind, S. Wolf, and T. Henning. Towards a physical model of dust tori in Active Galactic Nuclei. Radiative transfer calculations for a hydrostatic torus model. *A&A*, 437:861–881, July 2005. doi: 10.1051/0004-6361:20042363.

M. Schartmann, K. Meisenheimer, M. Camenzind, S. Wolf, K. R. W. Tristram, and T. Henning. Three-dimensional radiative transfer models of clumpy tori in Seyfert galaxies. *A&A*, 482:67–80, April 2008. doi: 10.1051/0004-6361:20078907.

J. Schaye. Star Formation Thresholds and Galaxy Edges: Why and Where. *ApJ*, 609:667–682, July 2004. doi: 10.1086/421232.

J. Schaye and C. Dalla Vecchia. On the relation between the Schmidt and Kennicutt-Schmidt star formation laws and its implications for numerical simulations. *MNRAS*, 383:1210–1222, January 2008. doi: 10.1111/j.1365-2966.2007.12639.x.

J. Schaye, C. Dalla Vecchia, C. M. Booth, R. P. C. Wiersma, T. Theuns, M. R. Haas, S. Bertone, A. R. Duffy, I. G. McCarthy, and F. van de Voort. The physics driving the cosmic star formation history. *MNRAS*, 402:1536–1560, March 2010. doi: 10.1111/j.1365-2966.2009.16029.x.

J. Schaye, R. A. Crain, R. G. Bower, M. Furlong, M. Schaller, T. Theuns, C. Dalla Vecchia, C. S. Frenk, I. G. McCarthy, J. C. Helly, A. Jenkins, Y. M. Rosas-Guevara, S. D. M. White, M. Baes, C. M. Booth, P. Camps, J. F. Navarro, Y. Qu, A. Rahmati, T. Sawala, P. A. Thomas, and

J. Trayford. The EAGLE project: simulating the evolution and assembly of galaxies and their environments. *MNRAS*, 446:521–554, January 2015. doi: 10.1093/mnras/stu2058.

A. Schechtman-Rook, M. A. Bershady, and K. Wood. The Three-dimensional Distribution of Dust in NGC 891. *ApJ*, 746:70, February 2012. doi: 10.1088/0004-637X/746/1/70.

D. Schiminovich, T. K. Wyder, D. C. Martin, B. D. Johnson, S. Salim, M. Seibert, M. A. Treyer, T. Budavári, C. Hoopes, M. Zamojski, T. A. Barlow, K. G. Forster, P. G. Friedman, P. Morrissey, S. G. Neff, T. A. Small, L. Bianchi, J. Donas, T. M. Heckman, Y.-W. Lee, B. F. Madore, B. Milliard, R. M. Rich, A. S. Szalay, B. Y. Welsh, and S. Yi. The UV-Optical Color Magnitude Diagram. II. Physical Properties and Morphological Evolution On and Off of a Star-forming Sequence. *ApJS*, 173:315–341, December 2007. doi: 10.1086/524659.

K. Sellgren. The near-infrared continuum emission of visual reflection nebulae. *ApJ*, 277: 623–633, February 1984. doi: 10.1086/161733.

J. L. Sérsic. Influence of the atmospheric and instrumental dispersion on the brightness distribution in a galaxy. *Boletin de la Asociacion Argentina de Astronomia La Plata Argentina*, 6: 41, 1963.

R. Siebenmorgen, E. Kruegel, and J. S. Mathis. Radiative transfer for transiently heated particles. *A&A*, 266:501–512, December 1992.

D. Sijacki, M. Vogelsberger, D. Kereš, V. Springel, and L. Hernquist. Moving mesh cosmology: the hydrodynamics of galaxy formation. *MNRAS*, 424:2999–3027, August 2012. doi: 10.1111/ j.1365-2966.2012.21466.x.

M. F. Skrutskie, R. M. Cutri, R. Stiening, M. D. Weinberg, S. Schneider, J. M. Carpenter, C. Beichman, R. Capps, T. Chester, J. Elias, J. Huchra, J. Liebert, C. Lonsdale, D. G. Monet, S. Price, P. Seitzer, T. Jarrett, J. D. Kirkpatrick, J. E. Gizis, E. Howard, T. Evans, J. Fowler, L. Fullmer, R. Hurt, R. Light, E. L. Kopan, K. A. Marsh, H. L. McCallon, R. Tam, S. Van Dyk, and S. Wheelock. The Two Micron All Sky Survey (2MASS). *AJ*, 131:1163–1183, February 2006. doi: 10.1086/498708.

N. Smirnov. Table for estimating the goodness of fit of empirical distributions. *Ann. Math. Statist.*, 19(2):279–281, 06 1948. doi: 10.1214/aoms/1177730256. URL http://dx.doi.org/ 10.1214/aoms/1177730256.

J. D. T. Smith, B. T. Draine, D. A. Dale, J. Moustakas, R. C. Kennicutt, Jr., G. Helou, L. Armus, H. Roussel, K. Sheth, G. J. Bendo, B. A. Buckalew, D. Calzetti, C. W. Engelbracht, K. D. Gordon, D. J. Hollenbach, A. Li, S. Malhotra, E. J. Murphy, and F. Walter. The Mid-Infrared Spectrum of Star-forming Galaxies: Global Properties of Polycyclic Aromatic Hydrocarbon Emission. *ApJ*, 656:770–791, February 2007. doi: 10.1086/510549.

V. Springel. The cosmological simulation code GADGET-2. *MNRAS*, 364:1105–1134, December 2005. doi: 10.1111/j.1365-2966.2005.09655.x.

V. Springel. E pur si muove: Galilean-invariant cosmological hydrodynamical simulations on a moving mesh. *MNRAS*, 401:791–851, January 2010. doi: 10.1111/j.1365-2966.2009.15715.x.

V. Springel, S. D. M. White, G. Tormen, and G. Kauffmann. Populating a cluster of galaxies - I. Results at [formmu2]z=0. *MNRAS*, 328:726–750, December 2001. doi: 10.1046/j.1365-8711. 2001.04912.x.

V. Springel, T. Di Matteo, and L. Hernquist. Modelling feedback from stars and black holes in galaxy mergers. *MNRAS*, 361:776–794, August 2005. doi: 10.1111/j.1365-2966.2005.09238.x.

M. Stalevski. SKIRTOR - database of modelled AGN dusty torus SEDs. *Bulgarian Astronomical Journal*, 18(3):030000, 2012.

M. Stalevski, J. Fritz, M. Baes, T. Nakos, and L. Č. Popović. 3D radiative transfer modelling of the dusty tori around active galactic nuclei as a clumpy two-phase medium. *MNRAS*, 420: 2756–2772, March 2012a. doi: 10.1111/j.1365-2966.2011.19775.x.

M. Stalevski, P. Jovanović, L. Č. Popović, and M. Baes. Gravitational microlensing of active galactic nuclei dusty tori. *MNRAS*, 425:1576–1584, September 2012b. doi: 10.1111/j.1365-2966. 2012.21611.x.

M. Stalevski, C. Ricci, Y. Ueda, P. Lira, J. Fritz, and M. Baes. The dust covering factor in active galactic nuclei. *ArXiv e-prints*, February 2016.

D. Stamatellos and A. P. Whitworth. Monte Carlo radiative transfer in SPH density fields. *A&A*, 439:153–158, August 2005. doi: 10.1051/0004-6361:20052951.

J. Steinacker, A. Bacmann, T. Henning, R. Klessen, and M. Stickel. 3D continuum radiative transfer in complex dust configurations. II. 3D structure of the dense molecular cloud core ρ Oph D. *A&A*, 434:167–180, April 2005. doi: 10.1051/0004-6361:20041978.

J. Steinacker, M. Baes, and K. D. Gordon. Three-Dimensional Dust Radiative Transfer. *ARA&A*, 51:63–104, August 2013. doi: 10.1146/annurev-astro-082812-141042.

Jürgen Steinacker, Aurore Bacmann, and Thomas Henning. Application of adaptive multi-frequency grids to three-dimensional astrophysical radiative transfer. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 75(6):765 – 786, 2002. ISSN 0022-4073. doi: http://dx. doi.org/10.1016/S0022-4073(02)00042-0. URL http://www.sciencedirect.com/science/article/pii/S0022407302000420.

J. M. Stone and M. L. Norman. ZEUS-2D: A radiation magnetohydrodynamics code for astrophysical flows in two space dimensions. I - The hydrodynamic algorithms and tests. *ApJS*, 80:753–790, June 1992. doi: 10.1086/191680.

H. Suto, H. Sogawa, S. Tachibana, C. Koike, H. Karoji, A. Tsuchiyama, H. Chihara, K. Mizutani, J. Akedo, K. Ogiso, T. Fukui, and S. Ohara. Low-temperature single crystal reflection spectra of forsterite. *MNRAS*, 370:1599–1606, August 2006. doi: 10.1111/j.1365-2966.2006.10594.x.

E. N. Taylor, M. Franx, K. Glazebrook, J. Brinchmann, A. van der Wel, and P. G. van Dokkum. On the Dearth of Compact, Massive, Red Sequence Galaxies in the Local Universe. *ApJ*, 720: 723–741, September 2010. doi: 10.1088/0004-637X/720/1/723.

E. N. Taylor, A. M. Hopkins, I. K. Baldry, M. J. I. Brown, S. P. Driver, L. S. Kelvin, D. T. Hill, A. S. G. Robotham, J. Bland-Hawthorn, D. H. Jones, R. G. Sharp, D. Thomas, J. Liske, J. Loveday, P. Norberg, J. A. Peacock, S. P. Bamford, S. Brough, M. Colless, E. Cameron, C. J. Conselice, S. M. Croom, C. S. Frenk, M. Gunawardhana, K. Kuijken, R. C. Nichol, H. R. Parkinson, S. Phillipps, K. A. Pimbblet, C. C. Popescu, M. Prescott, W. J. Sutherland, R. J. Tuffs, E. van Kampen, and D. Wijesinghe. Galaxy And Mass Assembly (GAMA): stellar mass estimates. *MNRAS*, 418:1587–1620, December 2011. doi: 10.1111/j.1365-2966.2011.19536.x.

R. Teyssier. Cosmological hydrodynamics with adaptive mesh refinement. A new high resolution code called RAMSES. *A&A*, 385:337–364, April 2002. doi: 10.1051/0004-6361:20011817.

A. T. Tokunaga and W. D. Vacca. "The Mauna Kea Observatories Near-Infrared Filter Set. III. Isophotal Wavelengths and Absolute Calibration" (¡A href="/abs/2005PASP.117..421T"¿PASP, 117, 421 [2005]¡/A¿). *PASP*, 117:1459–1459, December 2005. doi: 10.1086/499029.

P. Torrey, M. Vogelsberger, D. Sijacki, V. Springel, and L. Hernquist. Moving-mesh cosmology: properties of gas discs. *MNRAS*, 427:2224–2238, December 2012. doi: 10.1111/j.1365-2966.2012.22082.x.

J. W. Trayford, T. Theuns, R. G. Bower, J. Schaye, M. Furlong, M. Schaller, C. S. Frenk, R. A. Crain, C. Dalla Vecchia, and I. G. McCarthy. Colours and luminosities of z = 0.1 galaxies in the EAGLE simulation. *MNRAS*, 452:2879–2896, September 2015. doi: 10.1093/mnras/stv1461.

Trayford et al. SKIRTing the issue: investigating observable properties of EAGLE galaxies using 3D radiative transfer. *MNRAS*, in prep.

S. Tremaine, D. O. Richstone, Y.-I. Byun, A. Dressler, S. M. Faber, C. Grillmair, J. Kormendy, and T. R. Lauer. A family of models for spherical stellar systems. *AJ*, 107:634–644, February 1994. doi: 10.1086/116883.

C. A. Tremonti, T. M. Heckman, G. Kauffmann, J. Brinchmann, S. Charlot, S. D. M. White, M. Seibert, E. W. Peng, D. J. Schlegel, A. Uomoto, M. Fukugita, and J. Brinkmann. The Origin of the Mass-Metallicity Relation: Insights from 53,000 Star-forming Galaxies in the Sloan Digital Sky Survey. *ApJ*, 613:898–913, October 2004. doi: 10.1086/423264.

R. van de Weygaert. Fragmenting the Universe. 3: The constructions and statistics of 3-D Voronoi tessellations. *A&A*, 283:361–406, March 1994.

P. C. van der Kruit. Surface photometry of edge-on spiral galaxies. V - The distribution of luminosity in the disk of the galaxy derived from the Pioneer 10 background experiment. *A&A*, 157:230–244, March 1986.

B. Vandenbroucke and S. De Rijcke. Shadowfax: hydrodynamics on a (co-)moving mesh. *Astronomy and Computing*, submitted.

J. Verstappen, J. Fritz, M. Baes, M. W. L. Smith, F. Allaert, S. Bianchi, J. A. D. L. Blommaert, G. De Geyter, I. De Looze, G. Gentile, K. D. Gordon, B. W. Holwerda, S. Viaene, and E. M. Xilouris. HERschel Observations of Edge-on Spirals (HEROES). I. Far-infrared morphology and dust mass determination. *A&A*, 556:A54, August 2013. doi: 10.1051/0004-6361/201220733.

Verstocken et al. SKIRT: the design of hybrid parallelization for Monte Carlo radiative transfer simulations. *Astronomy and Computing*, in prep.

S. Viaene, M. Baes, G. Bendo, M. Boquien, A. Boselli, L. Ciesla, L. Cortese, I. De Looze, S. Eales, J. Fritz, O. Ł. Karczewski, S. Madden, M. W. L. Smith, and L. Spinoglio. The bolometric and UV attenuation in normal spiral galaxies of the Herschel Reference Survey. *A&A*, 586:A13, February 2016. doi: 10.1051/0004-6361/201527586.

Viaene et al. The herschel exploitation of local galaxy andromeda (helga) vii. a skirt radiative transfer model and insights on dust heating. *MNRAS*, in prep.

M. Vogelsberger, D. Sijacki, D. Kereš, V. Springel, and L. Hernquist. Moving mesh cosmology: numerical techniques and global statistics. *MNRAS*, 425:3024–3057, October 2012. doi: 10.1111/j.1365-2966.2012.21590.x.

M. Vogelsberger, S. Genel, V. Springel, P. Torrey, D. Sijacki, D. Xu, G. Snyder, D. Nelson, and L. Hernquist. Introducing the Illustris Project: simulating the coevolution of dark and visible matter in the Universe. *MNRAS*, 444:1518–1547, October 2014. doi: 10.1093/mnras/stu1536.

G. M. Voit. X-ray irradiation of interstellar grains in active galaxies - Evaporation and infrared spectra. *ApJ*, 379:122–140, September 1991. doi: 10.1086/170490.

George Voronoi. Nouvelles applications des paramètres continus à la théorie de formes quadratiques. deuxième mémoire. recherches sur les parallélloèdres primitifs. *Journal für die reine und angewandte Mathematik*, 134:198–287, 1908.

J. W. Wadsley, J. Stadel, and T. Quinn. Gasoline: a flexible, parallel implementation of TreeSPH. *New A*, 9:137–158, February 2004. doi: 10.1016/j.newast.2003.08.004.

M. Wall. GAlib: A C++ library of genetic algorithm components. *PhD thesis, Mechanical Engineering Department, Massachusetts Institute of Technology*, 1996.

L. Wang, A. A. Dutton, G. S. Stinson, A. V. Macciò, C. Penzo, X. Kang, B. W. Keller, and J. Wadsley. NIHAO project - I. Reproducing the inefficiency of galaxy formation across cosmic time with a large sample of cosmological hydrodynamical simulations. *MNRAS*, 454: 83–94, November 2015. doi: 10.1093/mnras/stv1937.

Z. Wang, D. L. Kaplan, P. Slane, N. Morrell, and V. M. Kaspi. Serendipitous Discovery of an Infrared Bow Shock near PSR J1549-4848 with Spitzer. *ApJ*, 769:122, June 2013. doi: 10.1088/0004-637X/769/2/122.

J. C. Weingartner and B. T. Draine. Dust Grain-Size Distributions and Extinction in the Milky Way, Large Magellanic Cloud, and Small Magellanic Cloud. *ApJ*, 548:296–309, February 2001a. doi: 10.1086/318651.

J. C. Weingartner and B. T. Draine. Photoelectric Emission from Interstellar Dust: Grain Charging and Gas Heating. *ApJS*, 134:263–281, June 2001b. doi: 10.1086/320852.

B. A. Whitney. Monte Carlo radiative transfer. *Bulletin of the Astronomical Society of India*, 39: 101–127, March 2011.

R. P. C. Wiersma, J. Schaye, and B. D. Smith. The effect of photoionization on the cooling rates of enriched, astrophysical plasmas. *MNRAS*, 393:99–107, February 2009a. doi: 10.1111/j. 1365-2966.2008.14191.x.

R. P. C. Wiersma, J. Schaye, T. Theuns, C. Dalla Vecchia, and L. Tornatore. Chemical enrichment in cosmological, smoothed particle hydrodynamics simulations. *MNRAS*, 399:574–600, October 2009b. doi: 10.1111/j.1365-2966.2009.15331.x.

A. N. Witt and K. D. Gordon. Multiple Scattering in Clumpy Media. I. Escape of Stellar Radiation from a Clumpy Scattering Environment. *ApJ*, 463:681, June 1996. doi: 10.1086/ 177282.

A. N. Witt and K. D. Gordon. Multiple Scattering in Clumpy Media. II. Galactic Environments. *ApJ*, 528:799–816, January 2000. doi: 10.1086/308197.

A. N. Witt, H. A. Thronson, Jr., and J. M. Capuano, Jr. Dust and the transfer of stellar radiation within galaxies. *ApJ*, 393:611–630, July 1992. doi: 10.1086/171530.

S. Wolf. Efficient Radiative Transfer in Dust Grain Mixtures. *ApJ*, 582:859–868, January 2003a. doi: 10.1086/344710.

S. Wolf. MC3D-3D continuum radiative transfer, Version 2. *Computer Physics Communications*, 150:99–115, February 2003b. doi: 10.1016/S0010-4655(02)00675-6.

S. Wolf, O. Fischer, and W. Pfau. Radiative transfer in the clumpy environment of young stellar objects. *A&A*, 340:103–116, December 1998.

E. L. Wright, P. R. M. Eisenhardt, A. K. Mainzer, M. E. Ressler, R. M. Cutri, T. Jarrett, J. D. Kirkpatrick, D. Padgett, R. S. McMillan, M. Skrutskie, S. A. Stanford, M. Cohen, R. G. Walker, J. C. Mather, D. Leisawitz, T. N. Gautier, III, I. McLean, D. Benford, C. J. Lonsdale, A. Blain, B. Mendez, W. R. Irace, V. Duval, F. Liu, D. Royer, I. Heinrichsen, J. Howard, M. Shannon, M. Kendall, A. L. Walsh, M. Larsen, J. G. Cardon, S. Schick, M. Schwalm, M. Abid, B. Fabinsky, L. Naes, and C.-W. Tsai. The Wide-field Infrared Survey Explorer (WISE): Mission Description and Initial On-orbit Performance. *AJ*, 140:1868, December 2010. doi: 10.1088/0004-6256/140/6/1868.

N. Ysard, M. Juvela, and L. Verstraete. Modelling the spinning dust emission from dense interstellar clouds. *A&A*, 535:A89, November 2011. doi: 10.1051/0004-6361/201117394.

F. Yusef-Zadeh, M. Morris, and R. L. White. Bipolar reflection nebulae - Monte Carlo simulations. *ApJ*, 278:186–194, March 1984. doi: 10.1086/161780.

T. Zafar and D. Watson. The metals-to-dust ratio to very low metallicities using GRB and QSO absorbers; extremely rapid dust formation. *A&A*, 560:A26, December 2013. doi: 10.1051/0004-6361/201321413.

H. J. Zahid, G. I. Dima, R.-P. Kudritzki, L. J. Kewley, M. J. Geller, H. S. Hwang, J. D. Silverman, and D. Kashino. The Universal Relation of Galactic Chemical Evolution: The Origin of the Mass-Metallicity Relation. *ApJ*, 791:130, August 2014. doi: 10.1088/0004-637X/791/2/130.

S. Zibetti, S. Charlot, and H.-W. Rix. Resolved stellar mass maps of galaxies - I. Method and implications for global mass estimates. *MNRAS*, 400:1181–1198, December 2009. doi: 10.1111/j.1365-2966.2009.15528.x.

V. Zubko, E. Dwek, and R. G. Arendt. Interstellar Dust Models Consistent with Extinction, Emission, and Abundance Constraints. *ApJS*, 152:211–249, June 2004. doi: 10.1086/382351.