

3D dust radiative transfer¹

Jürgen Steinacker¹

*UJF-Grenoble 1 / CNRS-INSU, Institut de Planétologie et d'Astrophysique de
Grenoble (IPAG), UMR 5274, Grenoble, F-38041, France*

stein@mpia.de

and

Maarten Baes

terrenkundig Observatorium, Universiteit Gent, Krijgslaan 281 S9, B-9000 Gent, Belgium

and

Karl Gordon²

Space Telescope Science Institute, 3700 San Martin Dr., Baltimore, MD 21218, USA

ABSTRACT

Cosmic dust is present in many astrophysical objects, and recent observations across the electromagnetic spectrum have revealed that the dust distribution is often strongly three-dimensional. Dust grains are effective in absorbing and scattering UV/optical radiation, and re-emit the absorbed energy at infrared wavelengths. Understanding the intrinsic properties of these objects, including the dust itself, therefore requires 3D dust radiative transfer calculations. Unfortunately, the 3D dust radiative transfer problem is non-local and non-linear, which makes it one of the hardest challenges in computational astrophysics. Nevertheless, significant progress has been made in the last decade, with an increasing number of codes capable of dealing with the complete 3D dust radiative transfer problem. We discuss the complexity of this problem, describe the two most successful solution techniques (Ray-Tracing and Monte Carlo), and discuss the state of the art in modeling observational data using 3D dust radiative transfer codes. We end with an outlook on the bright future of this field.

Subject headings: scattering – Monte Carlo – ray-tracing – computational astrophysics – numerical algorithms

1. INTRODUCTION

Given the dominant role of radiation in astrophysics, its transport through a medium is one of the most fundamental processes to be considered. Analyzing the radiation received from an object not only provides us with information about its radiation source

but also the medium in between and surrounding the object and the observer.

Among the numerous ways of producing and processing radiation, dust grains mixed in the cosmic gas play a special role. They are efficient at absorbing and scattering ultraviolet (UV) through near-infrared (NIR) photons and then re-radiating the absorbed energy in the infrared and sub-millimeter (submm) wavelength range. Cosmic dust can be found in many astrophysical objects like the solar system (Hoppe et al. 2010), comets and meteoroids (Küppers et al. 2005), sub-stellar atmospheres (Harvey et al. 2012), young stellar objects (Keller et al. 2008), proto-stellar to

¹Max-Planck-Institut für Astronomie, Königstuhl 17, D-69117 Heidelberg, Germany

²Sterrenkundig Observatorium, Universiteit Gent, Krijgslaan 281 S9, B-9000 Gent, Belgium

¹This review was the result of a collaboration of equals, the ordering of the authorship list is not significant.

proto-planetary disks (Watson et al. 2009), evolved stars (Groenewegen et al. 2011), reflection nebulae (Castellanos et al. 2011), supernova remnants (Rho et al. 2009), molecular clouds (Martel, Urban & Evans 2012), the interstellar medium (Zhukovska, Gail & Tieloff 2008), galaxies (Dunne et al. 2011), active galactic nuclei (Haas et al. 2000), and the high-redshift universe (Dwek, Galliano & Jones 2007). As dust grains modify the radiation field in these objects, they should be taken into account for an unbiased analysis of their intrinsic properties. Such an analysis requires radiative transfer (RT) calculations to be performed.

Aside from its importance as a tracer, the physical and chemical processes related to dust itself are of interest. They cover its formation, its cycle in galaxies, the variation in opacity with chemical composition, its growth and destruction processes in cloud cores and circumstellar disks to act as building blocks for planets, its interaction with magnetic fields and the chemistry on the grain surface. For example, the dust RT is important for understanding chemistry in the ISM as photo-dissociation rates are strongly dependent on the UV radiation field that includes a significant amount of photons scattered from dust grains. In this review, we will just address physical properties of the dust where they enter the RT or the modeling of objects, and refer the reader to one of the many published works on dust for the aspects mentioned above (e.g., Draine 2003a; Henning, Grün & Steinacker 2009; Henning 2010).

Many dusty objects have been observed at increasingly higher spatial resolution in the last ten years. These observations cover UV/optical/NIR (e.g., Hubble, GALEX, ground-based telescopes) and infrared/submm wavelengths (e.g., Spitzer, Akari, Herschel, Planck, WISE, ALMA, JCMT, APEX, IRAM). Space telescopes exploring atmospherically absorbed wavelength windows, high-resolution interferometric data, polarization data, or all-sky maps are just a few examples of the rich data set that awaits the RT modeler. The information determined by comparison of dust RT calculations with global and pixel-by-pixel resolved spectral energy distributions (SEDs) of dusty objects include the properties of the illuminating sources (stars, accretion disks, integrated star formation rate, etc.), the distribution of the dust (disk structures, cloud geometries, underlying multi-phase nature, etc.), and properties of the dust grains (size, shape, and composition). A feature commonly seen in high spatial resolution images at all wavelengths is the complex nature of the dust density distribution.

Examples of global 3D geometries include complex arm structures in spiral galaxies (Patrikeev et al. 2006; Fritz et al. 2012), large scale filaments in star forming regions (André et al. 2010; Arzoumanian et al. 2011), and bow-shocked shells around evolved stars (Cox et al. 2012). A prominent example of locally complex 3D geometries is the known fractal nature of the interstellar medium (Beech 1987; Falgarone, Phillips & Walker 1991). Images illustrating complex local and global 3D dust structures are shown in Fig. 1. In addition, the illuminating sources of dust have been long known to have complex distributions from the combination of the anisotropic interstellar radiation field and local neighboring stars to the stellar distribution in galaxies. Both of these issues show that a complete 3D treatment of RT is inevitable and critical to make progress in many fields.

Among the many computational problems in astrophysics, 3D line and dust RT has long been a major challenge and often approximated or neglected. While, for example, 3D (magneto-)hydrodynamics ((M)HD) codes have existed for many years, radiation transport is considered to be one of the four Grand Challenges in Computational Astrophysics². Dust RT is different than line RT in that the dust opacities generally do not depend on the RT solution itself.

The reasons to neglect or approximate 3D dust RT are manifold. A good portion of the difficulties arises from the fact that the underlying physical processes combine, in the stationary case, to a non-local and non-linear 6D problem. Since the radiation field needs to be determined in all directions, at any spatial location, and for each wavelength, the solution vector itself comprises three dimensions more than the variables in (M)HD problems. The RT problem is non-local in space (propagation of the photons within the entire domain), direction (scattering and absorption/re-emission), and wavelength (absorption/re-emission). This non-locality makes it difficult to simplify the problem by neglecting processes or wavelengths. For example, absorption and scattering have roughly the same efficiency from UV to NIR, with strongly anisotropic scattering (Gordon 2004). Modeling far-infrared (FIR) images, a consistent treatment of the dust emission requires the RT to be calculated where the dust absorption happens, at shorter wavelengths.

²see, e.g., the Grand Challenge conference series at the Institute for Pure & Applied Mathematics at the University of California, USA from 2005 to 2007




Fig. 1.— The complex and filamentary structures of ISM dust are clearly seen in both Milky Way star formation regions (top) and external galaxies (bottom). The top image is a color image of the Aquila star formation complex (André et al. 2010). This image was taken as part of the Gould Belt Survey Herschel Key Project, covers ~ 11 deg², and was created from PACS 70 μm (blue), PACS 160 μm (green), and SPIRE 250 μm (red) observations. The bottom image shows the complex structure of the ISM in M31 from HI observations (green, Braun et al. 2009), embedded star formation using Spitzer 24 μm data (red, Gordon et al. 2006), and unobscured young stars from the GALEX far-UV images (blue, Thilker et al. 2005).

Therefore, most of the current 3D dust applications are intrinsically multi-wavelength in nature.

Other difficulties are related to the complexity that is encountered when accessing 3D structures. The underlying grids to resolve the sink and source contributions to the radiation fields are generally discretized and require substantial storage and the RT calculation effort rises with the cell size for a given solution accuracy criterion. Moreover, when modeling the structure, the spatial distribution of the sources and sinks has to be parametrized. Modeling complex structures with a simple spatial distribution model can lead to misleading results. Witt & Gordon (1996) showed that RT through a 3D fractal dust distribution is significantly different than through similar, but smooth distributions. Combined with the runtimes expected for the 3D dust RT code, an exploration or optimization of the parameter space challenges the capabilities of current computers. Finally, more than for simple geometries, applying 3D dust RT also struggles with the loss of information due to projection effects.

As a result of non-local and non-linear effects, the RT equation is an integro-differential equation including a scattering integral; the thermal source term is non-linearly coupled to a double-integral equation, making it difficult to apply common solvers. Moreover, the varying extinction causes changes in the numerical nature of the RT equation. Its character changes from parabolic for the diffusive transport to hyperbolic for freely-streaming photons; to a combination of the two, in the numerically difficult transition region. Solving such a high-dimensional non-local, non-linear problem requires substantial computational resources (both computing power and memory), affecting the solution algorithms and potential limiting the model complexity.

With this review we enter this evolving and dynamic field to report on the significant progress that has been made to tackle this grand challenge problem. Within the last 10 years, the availability of the high resolution images and increase computer speed and storage have triggered an expansion of the dust RT community and the development of new codes capable of dealing with the complete 3D dust RT problem. Many of the techniques used to solve the 3D dust RT problem were developed originally for 1D or 2D geometries. The added computational complexity of solving the 3D problem has emphasized the need for highly efficient techniques, leading to refinements in and use of all possible 1D/2D methods in most 3D

codes. Applications explicitly using 3D RT include models of young stellar objects (Wolf, Fischer & Pfau 1998), proto-stellar to proto-planetary disks (Indebetouw et al. 2006; Niccolini & Alcolea 2006), reflection nebulae (Witt & Gordon 1996), molecular clouds (Steinacker et al. 2005; Pelkonen, Juvela & Padoan 2009), spiral galaxies (Bianchi 2008; Schechtman-Rook, Bershadsky & Wood 2012), interacting and starburst galaxies (Chakrabarti et al. 2007; Hayward et al. 2011), and active galactic nuclei (Schartmann et al. 2008; Stalevski et al. 2012). This expansion motivated the ‘‘Cosmic Dust and Radiative Transfer’’ workshop held in Heidelberg in 2008³. During the workshop, we realized that a review of the various techniques and applications addressing 3D dust RT was needed to communicate common strategies between related fields. We felt that a review of 3D dust RT would be useful for coders and users of dust RT codes and people wishing to enter this field (including writers of line RT and (M)HD codes).

While various solver techniques are in use for RT problems, 3D dust RT is commonly solved using the Monte Carlo (MC) technique, with some applications being solved using the Ray-tracing (RayT) technique. Since modern MC solvers make use of some RayT methods, this review concentrates on describing the spectrum of techniques based on these two approaches.

Other RT solution methods exist but have either not been used to solve the complete 3D dust RT problem in all aspects, or have shown clear disadvantages. One potential solution method is discretizing the RT equation, e.g. with a finite difference approach in spatial Cartesian coordinates and in direction, to create a system of linear equations (Stenholm, Storzner & Wehrse 1991; Steinacker, Bacmann & Henning 2002). The corresponding matrix is extremely difficult to solve even with powerful matrix solvers (van der Vorst 1992). Two more methods are used for RT problems, but have yet to be applied to 3D dust RT. The discretization can be performed on unstructured grids (e.g., a Delaunay grid) and this can provide a very fast solver of RT problems. Currently, such algorithms have been updated to handle freely streaming photon packages and changes in the optical depth, but treatment of scattering has yet to be explored. Finally, the moment method expands the intensity as a function of angle using spherical harmonics as basis functions. It has several numerical advantages both in terms of

solution accuracy and storage requirements, but can exhibit non-physical oscillations. A common variant of the moment method is the Variable Eddington Tensor method (used, e.g., for 2D in the code RADICAL, see description in Pascucci et al. 2004).

This review starts with the mathematical definition of the full 3D dust RT problem. Next, the discretization of the problem in spatial, direction, and wavelength dimensions is presented. The RayT and MC methods of solving the 3D RT problem are described in detail. Challenges in comparing RT models with observations are discussed. A listing of existing 3D dust RT codes is given along with current benchmarking efforts. Finally, the review is concluded with a summary and discussion of the future of 3D RT.

2. THE 3D DUST RADIATIVE TRANSFER PROBLEM

2.1. The radiative transfer equation

The stationary radiation field is described by the specific intensity $I(\mathbf{x}, \mathbf{n}, \lambda)$, where \mathbf{x} gives the location in space, \mathbf{n} is a unit vector indicating the direction of the radiation, and λ its wavelength. The specific intensity represents the amount of energy carried by radiation in a unit wavelength interval, which is transported per unit solid angle and per unit time across an element of unit area perpendicular to \mathbf{n} .⁴ The continuum radiative transfer equation (RTE) describes how the specific intensity varies as a result of interactions with a medium filled with sources and sinks. In its general form, it can be written as (see e.g. Chandrasekhar 1960; Rybicki & Lightman 1979)

$$\mathbf{n} \cdot \nabla I(\mathbf{x}, \mathbf{n}, \lambda) = -\kappa(\mathbf{x}, \lambda) \rho(\mathbf{x}) I(\mathbf{x}, \mathbf{n}, \lambda) + j(\mathbf{x}, \mathbf{n}, \lambda). \quad (1)$$

The left-hand side of this equation represents the change of the intensity over an infinitesimal distance along the path determined by the position \mathbf{x} and the propagation direction \mathbf{n} . The first term on the right-hand side represents the extinction, i.e. the loss of radiant energy, when radiation passes through matter. Here $\kappa(\mathbf{x}, \lambda)$ is the mass extinction coefficient, and $\rho(\mathbf{x})$ is the mass density. The second term on the right-hand side represents the source term, i.e. the new luminosity

⁴The specific intensity can be defined as the intensity per unit of wavelength or per unit of frequency. The convention chosen is typically indicated by a subscript, i.e. as I_λ or I_ν . We adopt the per unit of wavelength convention, and drop the subscript in order not to overload the notations.

³<http://www.mpia.de/RT08/>

released into the medium at \mathbf{x} in the direction \mathbf{n} . The complexity of the RTE depends on the nature of the source and sink terms, i.e. the different physical processes that are responsible for extinction and emission.

An alternative form of the RTE (1) explicitly uses the distance s along the path defined by a position \mathbf{x} and propagation direction \mathbf{n} as a variable. We then obtain

$$\frac{dI}{ds}(s, \lambda) = -\kappa(s, \lambda)\rho(s)I(s, \lambda) + j(s, \lambda). \quad (2)$$

If we assume for now that the source term j does not depend on the intensity I , we can readily solve this differential equation,

$$I(s, \lambda) = \int_{-\infty}^s j(s', \lambda) e^{-\tau(s', s, \lambda)} ds', \quad (3)$$

with the optical depth between two positions defined as

$$\tau(s_1, s_2, \lambda) = \int_{s_1}^{s_2} \kappa(s, \lambda)\rho(s) ds. \quad (4)$$

Expression (3) has a simple physical interpretation: it expresses that the intensity at any position s along a path results from the emission at all anterior points s' along the path, reduced by a factor $e^{-\tau(s', s, \lambda)}$ to account for the extinction by the intervening matter. It is important to stress that the expression (3) is only a *formal* solution of the RTE, and not a very useful solution of the RTE. Indeed, the emissivity generally does not only depend on position, direction and wavelength, but also on the specific intensity itself. The formal solution (3) is then no more than an integral equation version of the radiative transfer equation itself; this is particularly the case for dust radiative transfer. In the following subsections, we gradually build the dust RTE by including the various relevant physical processes.

2.2. Primary emission and absorption

Two important and obvious processes to take into account in a dusty medium are primary emission and absorption. Primary emission accounts for the radiative energy added to the radiation field – this is often stellar emission, but it can include, for example, radiation from an active galactic nucleus, emission line radiation from ionized gas or Bremsstrahlung. In general form, it can be characterized by a function $j_*(\mathbf{x}, \mathbf{n}, \lambda)$. Absorption is the process in which electromagnetic radiation is taken up by dust grains and transformed into the internal energy. It is characterized by the absorption coefficient κ_{abs} ; for a given chemical composition,

size and shape of a dust grain, the absorption coefficient can in principle be determined at any wavelength (Mie 1908; Purcell & Pennypacker 1973; Draine 1988; Mishchenko, Travis & Lacis 2002; Min, Hovenier & de Koter 2005).

When we take only primary emission and absorption by dust into account, the RTE (1) becomes

$$\frac{dI}{ds}(\mathbf{x}, \mathbf{n}, \lambda) = -\kappa_{\text{abs}}(\mathbf{x}, \lambda)\rho(\mathbf{x})I(\mathbf{x}, \mathbf{n}, \lambda) + j_*(\mathbf{x}, \mathbf{n}, \lambda). \quad (5)$$

This equation is a simple first-order differential equation and the solution can be found by just integrating along the line of sight, as for the formal solution (3). For general 3D geometries, this integration is done numerically.

2.3. Including scattering

The complexity of the RTE increases substantially when we take scattering into account. Scattering, as absorption, removes radiation from a beam and hence accounts for a second sink term in the RTE, the efficiency of which is quantified by the scattering coefficient κ_{sca} . Rather than converting the radiation to internal energy, it emits the same radiation in a different direction. Scattering hence does not only imply a second sink term, but also a second source term. The scattering phase function $\Phi(\mathbf{n}, \mathbf{n}', \mathbf{x}, \lambda)$ describes the probability that a photon originally propagating in the direction \mathbf{n}' and scattered at the position \mathbf{x} , will have \mathbf{n} as its new propagation direction after the scattering event. Given this definition, the phase function satisfies the normalization

$$\int_{4\pi} \Phi(\mathbf{n}, \mathbf{n}', \mathbf{x}, \lambda) d\Omega' = \int_{4\pi} \Phi(\mathbf{n}, \mathbf{n}', \mathbf{x}, \lambda) d\Omega = 1. \quad (6)$$

Adding the sink and source terms due to scattering to the RTE, we obtain

$$\begin{aligned} \frac{dI}{ds}(\mathbf{x}, \mathbf{n}, \lambda) = & -\kappa_{\text{ext}}(\mathbf{x}, \lambda)\rho(\mathbf{x})I(\mathbf{x}, \mathbf{n}, \lambda) + j_*(\mathbf{x}, \mathbf{n}, \lambda) \\ & + \kappa_{\text{sca}}(\mathbf{x}, \lambda)\rho(\mathbf{x}) \int_{4\pi} \Phi(\mathbf{n}, \mathbf{n}', \mathbf{x}, \lambda) I(\mathbf{x}, \mathbf{n}', \lambda) d\Omega'. \end{aligned} \quad (7)$$

where the extinction coefficient $\kappa_{\text{ext}} = \kappa_{\text{abs}} + \kappa_{\text{sca}}$. Unlike the simple differential equation (5), this equation is an integrodifferential equation in which the radiation fields at all different positions and in all different directions are coupled.

In many RT calculations, scattering is important and adds significant complexity given that dust scattering is anisotropic. This is especially true for UV to NIR wavelengths: observations indicate that the dust scattering albedo at these wavelengths is at least 50%, and that scattering off dust grains is strongly anisotropic (Witt et al. 1992; Calzetti et al. 1995; Burgh, McCandliss & Feldman 2002; Gordon 2004)⁵. Using an isotropic phase function or other approximations such as an isotropic two-stream approximation or an (effective) forward scattering (e.g. Code 1973; Natta & Panagia 1984; Calzetti, Kinney & Storchi-Bergmann 1994) might not always be physically justifiable. Several authors demonstrated that an improper treatment of anisotropic scattering leads to significant errors (e.g. Bruzual, Magris & Calvet 1988; Witt, Thronson & Capuano 1992; Baes & Dejonghe 2001). Even for radiation at MIR wavelengths or longer, where the scattering off common ISM grains is low, an application requiring the heating of the grains to be determined will need to properly handle scattering (Nielbock et al. 2012).

The Henyey-Greenstein (HG) phase function (Henyey & Greenstein 1941) is the most widely used parametrization of the dust phase function and provides a good single parameter approximation. Dust grain models predict small deviations from a HG phase function and other parametrizations or numerical phase functions can be used for increased accuracy (Kattawar 1975; Hong 1985; Draine 2003b).

2.4. Radiative transfer in dust mixtures

In any realistic dust medium, there is a range of different types of dust grains present, with various chemical compositions, sizes, shapes and number densities. Each of these different grain types i is characterized by its own absorption coefficient $\kappa_{\text{abs},i}(\lambda)$, scattering coefficient $\kappa_{\text{sca},i}(\lambda)$ and scattering phase function $\Phi_i(\mathbf{n}, \mathbf{n}', \lambda)$. If we denote the relative contribution of each grain type i at the position \mathbf{x} to the total dust num-

⁵Up-to-date versions of the Gordon (2004) plots of albedo and scattering phase function asymmetry versus wavelength can be found at http://www.stsci.edu/~kgordon/Dust/Scat_Param/scat_data.html.

ber density as $w_i(\mathbf{x})$, the transfer equation becomes

$$\begin{aligned} \frac{dI}{ds}(\mathbf{x}, \mathbf{n}, \lambda) = & - \sum_i w_i(\mathbf{x}) \kappa_{\text{ext},i}(\lambda) \rho(\mathbf{x}) I(\mathbf{x}, \mathbf{n}, \lambda) + j_*(\mathbf{x}, \mathbf{n}, \lambda) \\ & + \sum_i w_i(\mathbf{x}) \kappa_{\text{sca},i}(\lambda) \rho(\mathbf{x}) \int_{4\pi} \Phi_i(\mathbf{n}, \mathbf{n}', \lambda) I(\mathbf{x}, \mathbf{n}', \lambda) d\Omega'. \end{aligned} \quad (8)$$

It is straightforward to see that (8) is formally identical to equation (7) if we define the absorption coefficient, scattering coefficient, extinction coefficient and scattering phase function of a dust mixture as

$$\kappa_{\text{abs}}(\mathbf{x}, \lambda) = \sum_i w_i(\mathbf{x}) \kappa_{\text{abs},i}(\lambda), \quad (9a)$$

$$\kappa_{\text{sca}}(\mathbf{x}, \lambda) = \sum_i w_i(\mathbf{x}) \kappa_{\text{sca},i}(\lambda), \quad (9b)$$

$$\kappa_{\text{ext}}(\mathbf{x}, \lambda) = \sum_i w_i(\mathbf{x}) \kappa_{\text{ext},i}(\lambda), \quad (9c)$$

$$\Phi(\mathbf{n}, \mathbf{n}', \mathbf{x}, \lambda) = \frac{\sum_i w_i(\mathbf{x}) \kappa_{\text{sca},i}(\lambda) \Phi_i(\mathbf{n}, \mathbf{n}', \lambda)}{\sum_i w_i(\mathbf{x}) \kappa_{\text{sca},i}(\lambda)}. \quad (9d)$$

As long as we only consider primary emission, absorption and scattering, radiative transfer in dust mixtures is hence completely identical with radiative transfer in a dust medium with a single average dust grain and no approximations are needed (Martin 1978; Wolf 2003a).

2.5. Including dust emission

Apart from primary emission, absorption and scattering, a fourth physical process to be taken into account in dust RT is the thermal emission by the dust itself. Dust grains that absorb radiation re-emit the acquired radiative energy at wavelengths longwards of about $1 \mu\text{m}$. To account for this astrophysical process, we need to incorporate the third source term $j_d(\mathbf{x}, \lambda)$ in our RTE,

$$\begin{aligned} \frac{dI}{ds}(\mathbf{x}, \mathbf{n}, \lambda) = & -\kappa_{\text{ext}}(\mathbf{x}, \lambda) \rho(\mathbf{x}) I(\mathbf{x}, \mathbf{n}, \lambda) + j_*(\mathbf{x}, \mathbf{n}, \lambda) + j_d(\mathbf{x}, \lambda) \\ & + \kappa_{\text{sca}}(\mathbf{x}, \lambda) \rho(\mathbf{x}) \int_{4\pi} \Phi(\mathbf{n}, \mathbf{n}', \mathbf{x}, \lambda) I(\mathbf{x}, \mathbf{n}', \lambda) d\Omega'. \end{aligned} \quad (10)$$

It might seem as if the dust emissivity term is a simple extra source term similar to the primary stellar emissivity term. Its exact form is strongly dependent on which physical emission processes are important, and

often the dust emissivity term depends in a complicated and non-linear way on the intensity of the radiation field itself.

A common assumption is that the dust grains are in thermal equilibrium with the local interstellar radiation field. In this case, the emissivity of the population of grains of type i can be written as a modified blackbody emission characterized by an equilibrium temperature $T_i(\mathbf{x})$. Summing over all populations, we obtain

$$j_d(\mathbf{x}, \lambda) = \sum_i w_i(\mathbf{x}) \kappa_{\text{abs},i}(\lambda) \rho(\mathbf{x}) B(T_i(\mathbf{x}), \lambda), \quad (11)$$

with $B(T, \lambda)$ being the Planck function. The equilibrium temperature of each type of grain is determined by the balance equation, i.e. the condition that the total amount of energy absorbed equals the total amount of emitted energy,

$$\int_0^\infty \kappa_{\text{abs},i}(\lambda) J(\mathbf{x}, \lambda) d\lambda = \int_0^\infty \kappa_{\text{abs},i}(\lambda) B_\lambda(T_i(\mathbf{x}), \lambda) d\lambda, \quad (12)$$

where $J(\mathbf{x}, \lambda)$ represents the mean intensity of the radiation field,

$$J(\mathbf{x}, \lambda) = \frac{1}{4\pi} \int_{4\pi} I(\mathbf{x}, \mathbf{n}, \lambda) d\Omega. \quad (13)$$

It is important to note that the equilibrium temperature of the dust grains depends explicitly on their size and chemical composition. At the very same location, dust grains with different sizes and/or chemical compositions will obtain different equilibrium temperatures. So far, we have easily combined the absorption and scattering due to different kinds/sizes of dust grains in the RTE without any approximations. The same cannot be done for the thermal re-emission term. One might be inspired by equations (9) and average the temperatures of the different grains to a “mean” temperature. This results in reducing the complexity of the dust mixture at a given position to a single mean grain that will reach a single equilibrium temperature. While this could be useful, sufficient or even necessary for certain applications, it is important to be aware that this is a simplification of the RT problem that is not physically correct (e.g. Wolf 2003a).

While the assumption of thermal equilibrium is useful in some applications, it breaks down in others. Particularly important is the case where the dust medium contains very small dust grains (including polycyclic aromatic hydrocarbons [PAHs]). Large dust grains reach thermal equilibrium and emit as modified blackbodies with an equilibrium temperature. However,

small dust grains have small heat capacities, and the absorption of even a single UV/optical photon can substantially heat the grain. These small grains will not reach an equilibrium temperature but will instead undergo temperature fluctuations that lead to grain emission at temperatures well in excess of the equilibrium temperature. The emission from small grains is necessary to explain the observed mid-infrared emission of many objects (e.g. Sellgren 1984; Boulanger & Perault 1988; Helou et al. 2000; Smith et al. 2007; Draine et al. 2007). When including such transiently heated dust grains, the dust emissivity changes from expression (11) to

$$j_d(\mathbf{x}, \lambda) = \sum_i w_i(\mathbf{x}) \kappa_{\text{abs},i}(\lambda) \rho(\mathbf{x}) \left[\int_0^\infty P_i(T, \mathbf{x}) B(T, \lambda) dT \right]. \quad (14)$$

Here $P_i(T, \mathbf{x})$ is the temperature distribution for dust grains of type i at the location \mathbf{x} . This temperature distribution depends on the chemical composition and size of the dust grains, as well as on the intensity and hardness of the radiation field in which it is embedded. Several methods have been developed to calculate the temperature distribution of small dust grains, using either matrix operations or time averages (Dwek 1986; Desert, Boulanger & Shore 1986; Guhathakurta & Draine 1989; Siebenmorgen, Kruegel & Mathis 1992; Draine & Li 2001; Compiègne et al. 2011). The result is that the dust source term is an intricate, non-linear function of the specific intensity, which adds a significant complexity to the RT problem.

Finally, thermal emission is not the only emission process of dust grains: additional non-thermal processes are extended red emission (Witt & Boroson 1990; Smith & Witt 2002) and blue luminescence (Vijh, Witt & Gordon 2004). These non-thermal processes can account for a substantial fraction of the surface brightness of interstellar clouds at optical wavelengths (e.g. Gordon, Witt & Friedmann 1998; Witt et al. 2008). Both processes can be included as an additional term in the dust source term $j_d(\mathbf{x}, \lambda)$ in the RTE (10).

2.6. Radiative transfer of polarized radiation

The specific intensity $I(\mathbf{x}, \mathbf{n}, \lambda)$ is not a complete description of the radiation field, as it only describes unpolarized light. Scattering of photons from dust grains naturally produces polarized radiation (e.g., Fig. 2 of Gordon et al. 2001). In addition, aligned dust grains also produce polarized radiation (Whitney & Wolff

2002). While alignment of grains has been demonstrated observationally for many decades, the physical mechanism for grain alignment is still matter of debate (Lazarian 2007).

The most common way to describe polarized radiation makes use of the Stokes vector $\mathbf{S} = (I, Q, U, V)$, where I is the total specific intensity, Q and U the linearly polarized intensity in two axes rotated by 45 degrees one from the other, and V the circularly polarized intensity. The four components do not form a preferred basis of this space, but were chosen because they can easily be measured or calculated.

The RT formalism we have described in the previous subsections can be extended to include polarized radiation. Instead of a single RTE, we then obtain a vector RTE, or equivalently, a set of four coupled scalar equations,

$$\begin{aligned} \frac{d\mathbf{S}}{ds}(\mathbf{x}, \mathbf{n}, \lambda) = & -\kappa_{\text{ext}}(\mathbf{x}, \lambda) \rho(\mathbf{x}) \mathbf{S}(\mathbf{x}, \mathbf{n}, \lambda) + \mathbf{j}_*(\mathbf{x}, \mathbf{n}, \lambda) + \mathbf{j}_d(\mathbf{x}, \lambda) \\ & + \kappa_{\text{sca}}(\mathbf{x}, \lambda) \rho(\mathbf{x}) \int_{4\pi} \mathcal{M}(\mathbf{n}, \mathbf{n}', \mathbf{x}, \lambda) \mathbf{S}(\mathbf{x}, \mathbf{n}', \lambda) d\Omega'. \end{aligned} \quad (15)$$

A first complication is the scattering source term, where the phase function $\Phi(\mathbf{n}, \mathbf{n}', \mathbf{x}, \lambda)$ is now replaced by a 4×4 scattering (or Mueller) matrix $\mathcal{M}(\mathbf{n}, \mathbf{n}', \mathbf{x}, \lambda)$ which describes the changes in the Stokes vector when radiation is scattered from propagation direction \mathbf{n}' to a new propagation direction \mathbf{n} . For a full description see e.g. Bohren & Huffman (1983), Fischer, Henning & Yorke (1994) or Code & Whitney (1995). When we consider RT of polarized light through a dust mixture, each type i of grains is characterized by its own Mueller matrix $\mathcal{M}_i(\mathbf{n}, \mathbf{n}', \lambda)$. The radiative transfer equation can then still be written as (15), with

$$\mathcal{M}(\mathbf{n}, \mathbf{n}', \mathbf{x}, \lambda) = \frac{\sum_i w_i(\mathbf{x}) \kappa_{\text{sca},i}(\lambda) \mathcal{M}_i(\mathbf{n}, \mathbf{n}', \lambda)}{\sum_i w_i(\mathbf{x}) \kappa_{\text{sca},i}(\lambda)}. \quad (16)$$

A second complication is the thermal emission term: for thermal emission from aligned grains the full Stokes vector is used in the emission.

3. THE DISCRETE 3D DUST RADIATIVE TRANSFER PROBLEM

A general analytical solution of the stationary 3D dust RT equation is not possible for any of the non-symmetric applications mentioned in the introduction. To apply numerical solution techniques, solvers gener-

ally discretize the solution vector or the physical properties in the RTE. The quantities requiring discretization are the three spatial coordinates, the two directional coordinates, the wavelengths, and/or the dust properties.

A major concern when solving a 6D integro-differential equation is the size of the solution vector. With a resolution of 100 points in each variable, the intensity vector has 10^{12} entries. Handling this intensity vector requires an enormous amount of computer memory and speed. Currently, many solution algorithms avoid this requirement by not storing the full solution vector. There are applications where the full direction-dependent intensity is needed (e.g., when the radiation pressure impacts the gas kinematics). Due to this high dimensionality of 3D RT, the choice of appropriate grids is mandatory to effectively apply existing solution techniques and minimize memory usage and runtime. The solution techniques used does influence the choice of grids (e.g., RayT versus MC).

Another concern is that for a discrete solution vector, the physics is only solved at the grid resolution even if the solver used has no intrinsic error. Thus, if the grid is too coarse, a strong change in the intensity due a change from optical thin to thick may not be resolved and the derived intensity would have large systematic errors. This makes the choice of the grid crucial.

3.1. Spatial grids

For many astrophysical applications, the density values cover orders of magnitudes and are highly structured (e.g., a turbulent gas cloud with filaments and shocks). The same is true for the sources which can be dust grain emission, the distribution of stars, or a layer within a photon dissociation region (PDR) emitting in the MIR. Consequently, the complexity of spatial grids in 3D continuum RT ranges from simple linear Cartesian grids (Stenholm, Stoerzer & Wehrse 1991) to adaptively refined Cartesian grids (Kurosawa & Hillier 2001; Niccolini & Alcolea 2006; Lunttila & Juvela 2012) to multi-wavelength AMR grids (Steinacker, Bacmann & Henning 2002). Complex dust distributions are illustrated for three different cases in Fig. 2 using refined Cartesian grids. In principle, the RTE could be solved on unstructured, dynamic grids like those used in line RT (Petkova & Springel 2011; Paardekooper, Kruip & Icke 2010). Finally, the density or source distribution could be given using an




Fig. 2.— Examples of advanced spatial grids that are currently being used for 3D dust RT calculations, from [Saftly et al. \(2012\)](#), are shown. The examples are 2D plane cuts through octree-based grids used in Monte Carlo simulations; similar grids are applied in ray-tracing techniques. On each row, the left panel shows the dust density, the central panel the grid distribution in a regular octree structure, and the right panel the grid distribution in a barycentric octree structure. The top row represents a regular, analytical model of a double-exponential disc with a three-armed logarithmic spiral density perturbation. The middle row corresponds to a clumpy AGN model, consisting of a large number of optically thick clumps in a smooth density distribution ([Stalevski et al. 2012](#)). The bottom row corresponds to a late-type disc galaxy model resulting from N-body/SPH simulations with strong supernova feedback ([Rahimi & Kawata 2012](#)). In all cases, the grids contain between 3 and 4 million cells.

analytical formula and the 3D dust RT would be a complex function with many parameters.

The optimal grid would be based on changes in the radiation field intensity. As we don't know the radiation field a priori (this is the goal of the RT calculations), it is extremely challenging to generate such an optimal grid. First, the spatial grid is only 3D while the intensity is defined in 6D (spatial, direction, and wavelength). For example, the intensity in a certain direction can remain constant between two positions, while the intensity in another direction can vary drastically between the same positions. The radiation field also depends on the wavelength, so the optimal spatial grid is different for each wavelength. The dust mass density, or the optical depth, could serve as the starting point on which the grid could be based (see e.g. Figure 3, which shows 3D RT octree grids with a refinement criterion based on the total dust mass in a cell). However, the cell optical depth alone is not suitable to define a grid refinement criterion, as the strength of the radiation field can show strong gradients even in regions where the optical depth is small. In summary, the optimal grid should combine the details of the distribution of both the dust and source terms such that it captures the variation of the radiation field intensity.

As a result of the difficulty in creating an optimal spatial grid, various approximations are used. To review the variety of grids and their purpose, it is constructive to distinguish three spatial grid classes. They are used in various combinations by different solvers and in the different astrophysical communities.

Local mean intensity storage grids

This most common class of grids stores the mean intensity with the goal to resolving its variation. In the end, the grid's resolution determines the spatial resolution of the obtained solution.

To achieve a reasonable sampling of the gradients in the mean intensity, it is possible to calculate a series of spatial grids that are refined using the local optical depth averaged over all directions for each wavelength. The computational effort to calculate such grids is negligible compared to the effort of solving the RT equation ([Steinacker, Bacmann & Henning 2002](#)). The drawback of multi-wavelength grids is that a large number of interpolations have to be performed in order to assemble the wavelength dependent radiation field. Such interpolations are time consuming and introduce interpolation errors. As a compromise, grids

optimized for tracking the variation in the local optical depth for the wavelengths that dominate the radiative energy locations have the advantage of resolving the important regions at the expense of excess grid points. An alternative way of building the grid was proposed by [Niccolini & Alcolea \(2006\)](#) where an initial calculation is used to define the explicit location of a number of photons and these locations then used to refine the grid to have a constant absorption rate in each grid cell.

Density and source grids

The grids storing density and source information can originate from a discretization of simple physical models to deep structure grids designed to resolve the kinematic processes. The latter one is often realized as adaptive mesh refinement grid or smoothed particle (SP) distributions used in hydrodynamical simulations (e.g. [Steinacker et al. 2004](#)), typically incorporating several tens of millions of cells and many levels of refinement or SPH particles. While RT calculations can be performed directly using such density grid ([Abel & Wandelt 2002](#)), in most cases the density information is stored on coarser grids to meet the storage and speed limitations of the 6D RT.

Solution grids

The third class of grids is designed to calculate the solution directly at the grid points by advancing from one cell to the next in one step. The refinement criterion is defined to minimize solution errors of a certain order ([Steinacker et al. 2002](#)) or to provide a flexible grid that minimizes solution errors due to the coarse spatial coverage of the physical quantities ([Paardekooper, Kruip & Icke 2010](#)). These grids are well suited for (M)HD codes. Solving the RT equation directly on a grid accumulates discretization errors causing, e.g., a smearing of beams in the case of finite differencing solvers. Additional numerical methods or higher-order corrections are used to compensate for these numerical diffusion errors. It should be noted that in some applications the source function may also contain a smaller number of discrete sources such as stars. These can be considered outside the main spatial grid or using smaller subgrids ([Stamatellos & Whitworth 2005](#)).

3.2. Direction grid

There are two major challenges in defining a fixed direction grid for 3D RT. First, the radiation field can be strongly peaked due to nearby sources, and a smearing of this beam due to insufficient direction resolution will result in remote regions not being illuminated accurately. Since the optimal intensity discretization can vary greatly across the domain, a globally optimal refinement of direction space is not usually possible. Second, even direction grids optimized to be equally-spaced on the unit sphere ([Steinacker, Thamm & Maier 1996](#)) require a large number of points provide good angular resolution. Another possible regular direction grid is given by the HEALPix method which subdivides the unit sphere in pixels of equal size ([Górski et al. 2005](#)). For example, a grid with 600 direction points equally spaced on the unit sphere provides a mean resolution of $\sim 10^\circ$ only, and it takes 10000 grid points to achieve a mean resolution of $\sim 2.5^\circ$. For a protoplanetary disk, this resolution corresponds to a hot dust region with the size of 4!AU placed at a distance of 100 AU from the central star.

It is important to note that the two main RT solution techniques treat the discretization of directions quite differently. In MC, the direction of the photons is not discretized, but sampled from a probability function. For example, in MC the calculation of the scattered intensity is based directly on the scattering phase function, allowing an arbitrarily precise solution. In RayT, a discrete direction grid is used for all calculations. The scattered intensity calculation is done on the direction grid, potentially undersampling the scattering process in the direction space of the solution.

Another issue in direction space is the divergence of rays or photon tracks. The chance to miss physically important parts of the domain increases with the distance from the current point or cell. When tracing the radiation from a single source, this can lead to large errors in the computed radiation field for distant cells, or long runtimes to increase resolution using more photons or rays. Solutions to this issue exist and are discussed in the MC and RayT sections.

3.3. Wavelength and dust grain grids

It is important to consider variations in source spectra and dust opacity when choosing a wavelength grid. The spectrum of the radiation sources should be covered well enough to resolve the overall shape and any important spectral features (e.g., emission lines). The

wavelength grid should resolve variations in the grain properties (e.g., opacities and scattering properties). Where there are features in the dust properties (e.g., 2175 Å bump, MIR aromatic/PAH features), the wavelength grid should resolve the feature, ideally including a point at the maximum of each feature as well as enough points to Nyquist sample the feature.

Beside the discretization of the variables entering the intensity directly, the size distribution of the grains needs to be discretized if not given analytically. The grain size discretization can have a strong impact on the radiation field. The extinction of the radiation is the sum of the extinction of the different species, but the emissivity of individual grains is coupled to the intensity of the incoming radiation field (see Sec. 2.3).

4. THE RAY-TRACING SOLUTION METHOD

Ray-tracing (RayT) is a method widely applied in physics and computer graphics to describe the propagation of electromagnetic waves or particles through a medium with varying properties. Important applications outside astrophysics cover the propagation of radio signals in the ionosphere, the investigation of heating by plasma waves, sound waves in the ocean, optical design of lenses, tomographic reconstruction of the Earth’s interior, and image generation in computer graphics. In radiative transfer, RayT is used to follow the change of intensity in a particular direction which is the basic transport problem described by Eq. (1). The MC solution technique described in the next section is a sophisticated variant where a probabilistic approach is taken to choose the direction of the photon package representing the ray. RayT solvers have been used for a number of 2D continuum RT applications (see e.g. the benchmark comparison in [Pinte et al. 2009](#)) and 3D RayT is based on many techniques first developed in 2D.

In the following subsections, we describe the basic ingredients and challenges when using RayT as a solver of the general 3D continuum RTE. The solution on a single ray under various numerical and physical conditions is given in §4.1. The global solution and the treatment of source terms coupling directions are discussed in §4.2.

4.1. RayT Solution for Single Ray

The elementary RayT operation is to solve the first-order differential equation (2) within a spatial density grid cell along a given direction. The mass density

ρ_0 , the mass extinction coefficient κ_0 , and the source function j_0 are assumed to be constant within the cell for a given wavelength, so that we can determine the intensity $I(s + \Delta s, \lambda)$ from $I(s, \lambda)$ using Eq. (3)

$$I(s + \Delta s, \lambda) = I(s, \lambda)e^{-\tau_0(\lambda)} + \frac{j_0(\lambda)\Delta s}{\tau_0(\lambda)} (1 - e^{-\tau_0(\lambda)}) \quad (17)$$

with $\tau_0(\lambda) = \kappa_0(\lambda)\rho_0\Delta s$. For a ray crossing several cells, the numerical task is to determine the cells that are hit by the ray, calculate the intersection points with the cell borders, and use Eq. (17) to calculate the change in intensity along the ray in each cell. The first step can be time-consuming, since an adaptively refined grid is often stored as a tree, and neighbor-search calculations are required. The error in the intensity is solely defined by the the finite resolution of the underlying spatial grid.

For clarity in the notations used, we note that the solution of the RTE along rays with constant direction is also called the method of *short* or *long* characteristics. Long characteristics refers to calculating the radiation field along a ray through the entire computational domain. Since several rays can cross a certain cell causing redundant calculations, pre-calculated local column density or optical depth values can be used (short characteristics) to interpolate the contributions along a ray. Usually, a sweep of the ordering of the grid is required to ensure that, for a given direction, all information about positions before the currently considered point are given before performing the step described by Eq. (17). The method is less accurate than long characteristics due to the accumulation of interpolation errors. For combined applications of RT and (M)HD, hybrid methods have been proposed combining the advantages of short and long characteristics ([Rijkhorst et al. 2006](#)).

Beyond the Spatial Grid Resolution

There are cases where the best RayT solution is not performed at the density grid resolution. First, the spatial resolution of the density from an Adaptive Mesh Refinement (AMR) MHD calculation with many refinement levels can be too fine for a RT calculation to be done in a reasonable time. One way to deal with this case is to interpolate the density to a coarser grid, but this loses some of the information on the physical structure obtained in the AMR run. Second, the density grid shows strong gradients. One way to soften the gradients is to interpolate the density to a

finer grid. But this can still leave the density changes abrupt. Instead an interpolation scheme can be used to find the density in between grid cells. Finally, the density is given analytically and no step size information is available. Note that in practice analytical density structures are not automatically simpler than AMR density grids. The density used for the RT modeling of Spitzer images of the molecular cloud L183 by [Steinacker et al. \(2010\)](#), e.g., involved 100 3D clumps with Gaussian profiles and 700 free parameters.

The simplest approach to solve the RTE in these cases is to apply an upwinding first-order finite difference scheme. The relation between two intensities located at s and $s + \Delta s$ along the ray is then

$$I(s + \Delta s, \lambda) = I(s, \lambda) [1 - \tau(s, \Delta s, \lambda)] + j(s, \lambda) \Delta s. \quad (18)$$

using the local optical depth $\tau(s, \Delta s, \lambda) = \kappa(s, \lambda) \rho(s) \Delta s$. The step size Δs is chosen to be small enough to stay in the optically thin limit allowing the Taylor expansion of the exponential to be truncated after the first term ($\exp[-\tau] \approx 1 - \tau$). The advantage of this scheme is that it is fast, the disadvantage is that first-order errors can accumulate along the ray. Moreover, the steps become very small in regions of high optical depth, although the radiation field takes a simple form in this limit.

The concept of "trial" steps can improve the accuracy. Using a Runge-Kutta scheme, a step with the size $\Delta s/2$ can be made to calculate a second estimate of the derivative. The scheme then advances to the next point by using a linear combination of both derivatives. The linear factors are chosen by comparison with the Taylor series of $I(s)$ and letting the first-order term vanish.

This improvement can be repeated to achieve a better solution, at the cost of repeated calculations or look-up of the density and source terms. An advanced ray-tracer based on a 5th-order Runge-Kutta accuracy has been proposed by [Press et al. \(2002\)](#). It is coupled with an adaptive step size control using an embedded form of the 4th-order Runge-Kutta formula. As parameters for the error truncation, values determined by [Cash & Karp \(1990\)](#) are used. The tracer steps are chosen adaptively and with full error control. Such a tracer is therefore the first choice for astrophysical RT problems with moderate optical depth variations which are not solved at the resolution of the density grid. This explicit error control is an important characteristic of the RayT solution technique.

High optical depths

There are 3D RT applications with strong gradients in the source function or in the local optical depth. In particular, the case of high optical depth $\tau \gg 1$ occurs in all star formation regions as well as in AGN tori. All algorithms used to solve the RT encounter problems in correctly describing the intensity changes in the optically thick case. Adaptive grid RT solution techniques refine the regions into too many cells. Only moment methods which are applied in conjunction with (M)HD solvers are well-posed to treat the optical thick regime, but in turn fail to describe a strongly peaked radiation field arising in low optical depth regions.

To illustrate the difficulties of high optical depths for the RayT method, we assume a simple modified black body thermal source term for a single dust species $j(s, \lambda) = \kappa(s, \lambda) \rho(s) B(s, \lambda)$ in Eq. (18), which becomes

$$I(s + \Delta s, \lambda) - I(s, \lambda) = \tau(s, \lambda) [B(s, \lambda) - I(s, \lambda)]. \quad (19)$$

When the ray moves from one to the next point, most of the radiation at s is damped, as is the source contribution along the path, so that the local radiation at $s + \Delta s$ is dominated by the source contribution. Hence $B - I$ is small, suggesting that the solver can perform large steps. But $B - I$ is multiplied by the large optical depth τ amplifying any change in the source term that arises from the spatial variation in the temperature. Thus, to meet accuracy requirements, the tracer must perform small steps.

The second reason to modify the ray-tracer is the two exponential functions entering the radiation equation. First, the exponential containing the optical depth has to be calculated precisely along the ray. Given a limited computational range of the computer, an optical depth of 1000 usually exceeds these limits and makes it necessary to renormalize the intensity. Second, the Planck function rises sharply with λ for temperatures in the Wien part $T(x) < hc/\lambda k$.

To solve the two exponential problem, [Steinacker, Bacmann & Henning \(2006\)](#) proposed to use the transformation to the relative intensity $D = I/(I + B)$. For a vanishing source function, it approaches unity, and for the optically thick part with $I(s, \lambda) \approx B(s, \lambda)$, $D(s, \lambda)$ has a value close to 1/2. The authors showed that the transformation avoids numerical problems caused by the exponentials and that the intensity error amplification by the transformation is less than a factor of 2. As criterion for the use of an approximate solver in the

optically thick region $D(s, \lambda) \approx 1/2$, they derived

$$\frac{|B(s_2, \lambda) - B(s_1, \lambda)|}{B(s_1, \lambda)} \frac{1}{\tau(s_1, \lambda)} < \varepsilon \quad (20)$$

where ε is a positive limit for the relative difference $|B - I|/B$. The condition is fulfilled either due to small relative changes in the source term or a large local optical depth. In a pre-calculation along the ray, the regions where the condition is fulfilled can be identified. Then, a fast solution is obtained for these region by performing large steps while assuming $D(\lambda, x) = 1/2$. Applying the scheme to the case of a massive molecular cloud core illuminated by a nearby star, [Steinacker, Bacmann & Henning \(2006\)](#) verified speed-up factors of several hundreds compared to fourth-order Runge-Kutta solvers.

4.2. Ray location and RTE global solution

After discussing the solution methods for a single ray through the computational domain, the next step is to determine how to place the rays to ensure that the radiation field is correctly calculated. This is a critical part of the solution process.

Thermal Emission

We describe how to place the rays in order to calculate the intensity of the radiation field when the source term is dominated by thermal emission from dust grains. The various ray patterns used in RayT are illustrated in Fig. 3. For what follows, the term "placing a ray" will be used for the basic RayT step that includes defining the ray by a point in space and a direction, solving the intensity along the ray as previously described, and storing the absorbed energy in the cells that are crossed. According to Eq. (17), the intensity of the absorbed radiation per cell is

$$I(s, \Delta s, \lambda) = I(s, \lambda) (1 - e^{-\tau_0(\lambda)}) + j_0(\lambda) \Delta s \left[1 - \frac{1}{\tau_0(\lambda)} (1 - e^{-\tau_0(\lambda)}) \right] \quad (21)$$

Its contribution to the local radiation field then is calculated from the formula for the mean intensity (13) and the balance equation (12). In each RayT step, the local source term (11) contains the thermal energy of the currently crossed cell; this is updated with each ray crossing.

In RayT, pre-calculation steps are done to analyze the specific RT problem and accelerate the computations. The optical depth affects the transport within

the domain, the thermalization of the radiation, and the appearance of the object on the $\tau(\lambda) \approx 1$ -surface for the observer. Therefore a *penetration depth analysis* is performed at all wavelengths determining the optical depth distribution with respect to all discrete sources as well as to the observer (see, e.g., [Steinacker, Bacmann & Henning 2002](#)). For this calculation, rays are placed from the sources to all grid cells (Fig. 3a,b). In addition, we calculate the optical depth from the source to the cell using

$$\tau_{\text{sou}}(s_1, s_2, \lambda) = \int_{s_1}^{s_2} \kappa(s, \lambda) \rho(s) ds = \sum_{i=1}^{N_c} \kappa_i(\lambda) \rho_i \Delta s_i \quad (22)$$

with the ray crossing N_c cells with their individual κ_i , ρ_i , and crossing length Δs_i . $\tau_{\text{sou}} \approx 1$ marks the region where most of the source radiation is reprocessed.

The solver also calculates the optical depth from each cell to the observer τ_{obs} as described above (Fig. 3f). This information is used to understand which regions are shielded at which wavelengths from the observer, and to calculate the final images once the main RayT has been performed.

The second pre-calculation is to propagate the initially deposited discrete source energy and the radiation field at the domain boundary through the domain. The calculation is done on a regular grid of rays like the one shown in Fig. 3e. This calculation also determines the *optical depth between cells*, τ_{cc} , on a coarse spatial grid at all wavelengths, in this way providing information on which regions in the domain exchange significant amounts of radiation. In star formation applications, e.g., regions often do not connect significantly at UV wavelengths.

If the source of radiation is the interstellar radiation field, the pre-calculation can be very time-consuming, as it comprises rays from all directions to all cells at all wavelengths for which the optical depth can be ≥ 1 . However, there are many cases where this pre-calculation is less time-consuming. For example, when the dust properties are constant in the domain as then the optical depth is $\tau_{\text{sou}} = \kappa N_{\text{col}}$ and this is an integral only over the density.

The main RayT calculations are performed separately for each wavelength. Rays are placed through the grid points in the $\tau_{\text{sou}} \approx 1$ -layers covering all directions. In this way additional resolution is provided where the largest changes in the radiation field are expected. Moreover, additional rays from discrete sources are placed to help resolving the illumination

of the $\tau_{\text{sou}} \approx 1$ layers. This can be important, e.g., in the case of an accretion disk atmosphere that is illuminated by a protostar through a narrow solid angle. It is important to note that the order in which the computation is done for each wavelength influences the convergence. In RayT, information has already been transported from the sources to the cells during the pre-calculation step. Therefore, starting with wavelengths covering the peak of the re-emitted photon energy (e.g., in the MIR) can speed up the information transport in thermally dominated problems. For problems dominated by low optical depths in the UV/optical, the order is less important as the dust self-heating is small compared to the dust heating from UV/optical photons.

In practice, the placing of rays is controlled by the maximal numbers of rays per cell N_l and the "width" of the $\tau_{\text{sou}} \approx 1$ layers $\Delta\tau_l$ with $|\tau_{\text{sou}} - 1| < \Delta\tau_l$. The placing and iteration over wavelength is stopped when the change in the energy deposited in all cells drops below a chosen change limit and the mean field distribution has been determined. If the energy does not converge after placing N_l rays in all cells within the $\tau_{\text{sou}} \approx 1$ layers, a second run with all rays is performed but using the updated energy information from the first run. If still no convergence is obtained in this run, N_l is increased to improve the resolution. The number of rays for typical 3D dust RT applications can easily exceed 10^6 .

The final step is to calculate the radiation that is received by the observer from each cell. This calculation makes use of the pre-calculated τ_{obs} . Moreover, radiation directly received from the inner and outer discrete sources as well as from the background radiation that is inside the field of view is calculated by placing rays from the sources and the background to the observer.

To illustrate how the rays are placed in various situations, we present a few simple examples. The first is a molecular cloud core with a gas mass of $1 M_{\odot}$ that is illuminated by a strong MIR radiation field that dominates its thermal budget. The core has no internal source of radiation except the thermally emitting dust. Furthermore, the self-heating of the dust by neighboring dust within the core can be neglected. The pre-calculations will not yield any τ -surfaces since the optical depth is too low in the MIR. The main calculation will therefore be to propagate the external anisotropic radiation field through the core with little need for refinement by placing additional rays. Correspondingly, the rays can be placed dense enough in

direction space to resolve all features of the extended illumination source (e.g., a nearby photodissociation region, see [Steinacker et al. 2005](#)).

The second example is a binary star surrounded by a circumstellar disk. The pre-calculation will identify the inner disk and the disk surface as the $\tau \approx 1$ -zone for the NIR wavelengths. The ray pattern will therefore be dense at the inner rim, and in the zone above and below the disk. At MIR or longer wavelengths, the ray pattern will be less dense in the atmosphere since the $\tau \approx 1$ -range will move deeper into the disk as the thermally emitting inner dust is important. This application has an additional complication in that it requires the scattered light to be calculated (see below).

More complex structures being illuminated by a central source are a problem for all RT solvers: a proper resolution requires many spatial cells and therefore many rays from the star to the cell and from the cell into the surrounding regions. For RayT, if a fixed direction grid is used, more and more cells are overlooked when the diverging rays reach the outer parts. For such applications, beam-splitting can be used to split a ray in order to sample several neighboring cells ([Abel & Wandelt 2002](#)).

Including scattered radiation

While the thermal source contribution can be calculated using the mean intensity $J(x, \lambda)$ per cell, the scattered light intensity depends on the direction of the incoming radiation, the optical depth for scattering, and the phase function of the grains. The RayT scheme that includes thermal re-radiation and scattering must therefore store J and the intensity and direction of the incoming radiation for each cell and wavelength.

Every RayT step is expanded to include "directly scattered light" by calculating the amount of radiation that is scattered towards the observer using the pre-calculated τ_{obs} . In many applications, this singly scattered radiation is a good fraction of the total scattered light in the computational domain. Including multiple scattering in RayT methods significantly increases the computational requirements. Applications using 3D dust RT based on RayT with scattering sources have been presented for massive disk candidates ([Steinacker et al. 2006](#)) and grain growth in molecular cloud cores ([Steinacker et al. 2010](#)). Further codes using ray-tracing solvers and also alternative techniques to deal with high-optical depth are described in [Pinte et al. \(2009\)](#) and the references within. Fortunately, there

are specific astrophysical regimes where multiple scattering is not crucial or where reasonable approximations can be made to limit the computational requirements.

The standard ISM dust grain distribution can be approximated as a mean grain with a $0.1 \mu\text{m}$ grain size. Such grains have very low scattering efficiency in the MIR and, thus, scattering is often neglected if calculations are done at MIR or longer wavelengths. But caution should be exercised as for coagulated grains in dense molecular clouds scattering can be important even at MIR wavelengths (Steinacker et al. 2010). In addition, dust grains have a strongly forward peaked scattering phase function and a portion of the scattered light can be transported along the already calculated rays (see, e.g., Steinacker et al. 2003). For the wavelengths where dust shows a more isotropic scattering phase function (e.g., MIR), additional rays can be added to carry the scattered radiation. An example of where adding scattering is fairly straightforward is in modeling star formation regions. Models of these regions will have a high number of rays in cells that are identified using the $\tau \approx 1$ -search algorithm, allowing scattering to be calculated along existing rays. Finally, the high albedo of dust grains and the distribution of the scattered intensity over the unit sphere naturally reduce the intensity of the scattered radiation with subsequent scatterings allowing an intensity threshold to be used to limit the number of scatterings needing to be calculated.

In astrophysical objects where the optical depth is very small, or very large, or where the $\tau \approx 1$ layer is not resolved, calculating multiple scattering can be ignored. If $\tau \ll 1$ for an object (e.g., the diffuse ISM or the Zodiacal light), then single scattering completely dominates the scattered intensity and multiple scattering can be ignored. In models of $\tau \gg 1$, the main modeling goal is usually to calculate the thermal dust emission. Since scattering is generally important at optical and shorter wavelengths, most of energy at these wavelengths is absorbed in the $\tau \approx 1$ regions of the object. Thus, not calculating the scattered light properly will not have a large impact on the dust emission results. For very embedded sources like massive young stellar objects, the stellar energy is completely converted to radiation at MIR or longer wavelengths before it leaves the core and therefore scattering does not influence the outer radiation field. Finally, if a model does not resolve optically thick regions (e.g., in some extragalactic applications) and thus not having $\tau \approx 1$ regions,

then all the rays can be used to determine the global field and transport the single scattering.

4.3. RayT Error analysis

The precision of a solution of the RT equation along a ray through a spatial grid on which the opacity and source terms are described analytically is limited by machine precision.

Solving the RT equation along a ray through a spatial grid on which the opacity and source terms are analytically described can be done with only machine precision errors. For example, the previously mentioned Runge-Kutta solver with adaptive step size control provides a good compromise between accuracy and computational cost while providing explicit error control.

The main source of error for solutions on a single ray is the interpolation error of the density and source function from the underlying grid. It can accumulate if the density grid shows strong (and partially unresolved) gradients. The pre-calculation of $\tau \approx 1$ -regions and shielded areas helps to characterize how well the chosen grid describes the underlying physical problem, and refining the grid with this information can reduce the interpolation errors in the intensity distribution. Unlike solvers calculating the intensity on a grid (e.g. using finite differencing), RayT solvers create no diffusive errors (beam-smearing).

There is no general formula for the global error in the achieved intensity distribution, because deviations caused by coarse resolution due to the placement of rays are hard to quantify. A good test for the global accuracy of the thermal conversion of radiation is to calculate the energy output of the source term integrated over the domain, and to compare it to the energy in the radiation leaving the domain. Another test for the overall resolution of the important regions in the domain is based on the number of crossing rays per cell. All the cells that contribute significantly to the radiation field overall should be crossed by many rays. Having important cells with a small number of rays is an indication of underresolving the grid.

A practical test to understand and measure the global error is to increase the spatial or wavelength resolution by inserting more rays or adding wavelength grid points, and testing the stability of the solution.

5. THE MONTE CARLO SOLUTION METHOD

The Monte Carlo (MC) method is a general computational technique that is widely used in many different areas, including numerical mathematics, physical sciences, finance, and medicine. It is a joint name for a variety of stochastic or probabilistic techniques, which all have in common that they solve equations by sampling random numbers. MC methods are particularly interesting for transport systems, which was the motivation of their first application in the 1940s. For a general overview of MC as a tool for transport problems, see e.g. Dupree & Fraley (2002), Kalos & Whitlock (2009) or Whitney (2011). Its application to dust RT problems in an astrophysical context has a history of more than 40 years (e.g. Mattila 1970; Roark, Roark & Collins 1974; Witt & Stephens 1974; Witt 1977). In the last four decades, it has become a mainstream method for 3D dust RT calculations.

The basis of MC RT is to treat the radiation field as the flow of a large but finite number of photon packages (often called photons). Each individual photon is followed along its journey through the dusty medium. At every stage in its journey, the characteristics that determine the path of each photon (such as its birth location, initial propagation direction, or the distance along the path until the next interaction with a dust grain) are determined in a probabilistic way by generating random numbers from an appropriate probability density functions (PDF). At the end of the simulation, the radiation field is recovered from a statistical analysis of the photon paths. Hence, the MC technique simulates the RT instead of explicitly solving the RT equation.

Central to all MC techniques is the process of generating random numbers from a given PDF $p(x) dx$. Thus an algorithm is needed that returns a set of numbers X such that the probability that X lies in the infinitesimal interval between x and $x + dx$ is equal to $p(x) dx$. The starting point for such algorithms, which are key to the MC process, is a (pseudo-)random number generator. This is a code that generates uniform deviates (random numbers with an equal probability to be chosen in the unit interval between 0 and 1). In order to generate random numbers from another, arbitrary PDF, one almost always applies appropriate operations on one or more uniform deviates. The most popular methods are the so-called transformation and rejection methods, details of which can be found in Devroye (1986) or Press et al. (2002, Ch. 7).

In the following subsections, we will gradually in-

troduce the ingredients and techniques that are combined to develop a 3D MC dust RT code. We start by describing the simplest MC RT technique, as it is the basis for all modern MC RT, in §5.1. These simple techniques are sufficient for MC calculations in geometries with a large degree of symmetry, such as 1D spherical or slab geometries. For 3D applications, however, they would result in codes that would be very inefficient. Fortunately, there are various weighting schemes that improve the performance of this technique making modern MC RT quite efficient. Some of these weighting schemes have a strong heritage in RayT methods, making most modern MC RT codes hybrids between the classical MC and RayT techniques. A number of these weighting schemes were developed for 2D geometries, especially cylindrical geometries. The use of weighting schemes is critical for 3D MC RT; in §5.2 we discuss several of these techniques.

5.1. Simple MC RT

The simplest MC calculation consists of considering the RT problem at a single wavelength λ . We consider a source of photons, characterized by the source term $j_*(\mathbf{x}, \mathbf{n}, \lambda)$, and a distribution of dust, characterized by the dust density $\rho(\mathbf{x})$. Throughout the calculation, the state of each photon is tabulated by its energy, position, direction of travel, and polarization state. For 3D RT, the Cartesian coordinate system is usually used resulting in the photon's position being stored as $\mathbf{x} = (x, y, z)$ and the direction using direction cosines as $\mathbf{n} = (n_x, n_y, n_z)$. The polarization state is stored using the Stokes vector, $S = (I, Q, U, V)$.

Step 1: birth

The first step in a photon's life cycle is its birth, i.e. its injection into the computational domain. If N is the number of photons in the model run and $L_{\text{tot}}(\lambda)$ is the total luminosity of the source, the luminosity carried by each photon is $L = L_{\text{tot}}(\lambda)/N$. The initial position \mathbf{x} and propagation direction \mathbf{n} are to be chosen randomly according to the source term $j_*(\mathbf{x}, \mathbf{n}, \lambda)$, which means that they need to be sampled from the PDF

$$p(\mathbf{x}, \mathbf{n}) d\mathbf{x} d\mathbf{n} = \frac{j_*(\mathbf{x}, \mathbf{n}, \lambda) d\mathbf{x} d\mathbf{n}}{\iint j_*(\mathbf{x}, \mathbf{n}, \lambda) d\mathbf{x} d\mathbf{n}} = \frac{j_*(\mathbf{x}, \mathbf{n}, \lambda) d\mathbf{x} d\mathbf{n}}{L_{\text{tot}}(\lambda)}. \quad (23)$$

In many cases, e.g. for emission by stars or thermal emission by dust grains, the emission is isotropic, which implies that the initial propagation direction can

be chosen randomly from the unit sphere,

$$p(\mathbf{n}) d\mathbf{n} = \frac{d\mathbf{n}}{4\pi} = \frac{\sin\theta d\theta d\phi}{4\pi}. \quad (24)$$

Generating random values for θ and ϕ from this distribution can easily be done using the transformation method. In other cases, e.g. when we deal with external emission illuminating an interstellar cloud or anisotropic emission from an AGN accretion disc, different PDFs for the initial propagation direction need to be considered (e.g. Niccolini, Woitke & Lopez 2003; Stalevski et al. 2012). For most cases, the photon is assumed to be emitted unpolarized, i.e. $S = (1, 0, 0, 0)$.

Step 2: determination of the interaction point

Once the photon is launched into the dusty medium, the next step consists of randomly determining whether it will interact with a dust grain, and if so, where this interaction will take place. The PDF that describes the free pathlength before an interaction is most conveniently described in optical depth space, where it has an exponential distribution, $p(\tau) d\tau = e^{-\tau} d\tau$. The optical depth τ to which a particular photon travels along its path before it interacts with a dust grain, is drawn from from this exponential distribution. This is easily done using the transformation method: we simply pick a uniform deviate ξ and solve the equation

$$\xi = \int_0^\tau e^{-\tau'} d\tau' \quad (25)$$

for τ . Integrating and solving gives $\tau = -\ln\xi$ remembering that the distributions of ξ and $1 - \xi$ are equivalent. If τ is greater than the optical depth τ_{path} to the surface of the system in the direction the photon is traveling, the photon escapes the system, and the life cycle of this particular photon is over. Otherwise, the photon will interact with the dust medium at the location along the path corresponding to the traveled optical depth τ .

The next step is converting the traveled optical depth τ to a physical path length s , such that we can move the photon to the interaction site. This means that we have to integrate along the path and solve the integral equation

$$\int_0^s \kappa_{\text{ext}}(s', \lambda) \rho(s') ds' = \tau \quad (26)$$

for the path length s . Comparison between this equation and Eq. (22) highlights the intimate link between

the MC solution technique and the RayT technique: a large fraction of the calculations in MC simulations are pure ray-tracing operations.

In practice, MC codes virtually always use a grid structure of dust cells on which the dust density and the optical properties are discretized. The integral in equation (26) then reduces to a sum over the consecutive grid cells along the path, and the integral equation comes down to summing the optical depth along the path until τ is reached. This calculation is often one of the more computationally intensive portions of MC RT and is a strong driver for choosing a grid optimized for the particular astrophysical object of interest. For a Cartesian model grid, the distance traveled in each cell is easy to calculate, as is the next grid cell along the path. For hierarchical grids, more complex neighbor search algorithms may be required.

Step 3: absorption and scattering

Once the path length s has been calculated, the photon moves from its old location \mathbf{x} to its updated location, i.e. the interaction site $\mathbf{x} + s\mathbf{n}$. At this location, the photon can either be absorbed or scattered; the appropriate PDF is hence not a continuous, but a discrete function with only two possible values. The probability that the interaction is a scattering event is equal to the dust albedo $a = \kappa_{\text{sca}}/\kappa_{\text{ext}}$. Using a uniform deviate ξ , the nature of the interaction is easily determined: if $\xi \leq a$ we have a scattering event, if $\xi > a$ an absorption event.

In the case of an absorption event, this is the end of the photon's life cycle. If dust emission is to be calculated in the simulation, the absorbed photon luminosity is stored in the interaction cell. This absorbed luminosity will be used at a later stage to calculate the dust emission spectrum, which can then be used as the source function for another MC cycle.

If the interaction is a scattering event, the next step consists of determining the new propagation direction. In the case of isotropic scattering, this just comes down to generating a new random point from the unit sphere. In the case of anisotropic scattering, the new propagation direction \mathbf{n} is chosen according to the PDF

$$p(\mathbf{n}) d\mathbf{n} = \frac{\Phi(\mathbf{n}, \mathbf{n}', \mathbf{x}, \lambda) d\mathbf{n}}{4\pi}, \quad (27)$$

where $\Phi(\mathbf{n}, \mathbf{n}', \mathbf{x}, \lambda)$ is the scattering phase function and \mathbf{n}' is the original propagation direction before the scattering event. For spherical dust grains, the scat-

tering phase function is dependent only on the scattering angle between \mathbf{n} and \mathbf{n}' . For the HG phase function, the most popular approximation to the real phase function, the generation of a random scattering angle and hence the calculation of the new propagation direction, can be done analytically (Witt 1977). For polarized RT, the scattering process is more complex, as the scattering phase function is dependent on the polarization of the photon and, in addition, scattering changes the polarization state of the photon. A detailed description of scattering events in the case of polarized radiation can be found in e.g. Fischer, Henning & Yorke (1994), Code & Whitney (1995), or Goosmann & Gaskell (2007).

With its new propagation direction determined, the photon can continue its journey through the dusty medium. This means that the second and third step are repeated until the photon either escapes from the system, or is absorbed by a dust grain.⁶ This entire process is repeated for all the photons until the last photon has left the dusty medium. The RT at different wavelengths can be independently calculated given there is no explicit wavelength coupling in the RT. There is an implicit coupling when dust emission is included in the modeling. In this case, the absorbed luminosity at every wavelength is stored in every dust cell in the spatial grid. After finishing the simulation at all the wavelengths, the absorbed luminosity is used to compute the mean intensity $J(\mathbf{x}, \lambda)$ of the radiation field and subsequently the dust source term $j_d(\mathbf{x}, \lambda)$. It is common to include the dust emission as a second source of photons by first computing the RT from the primary sources, computing the RT for the dust emission, recomputing the dust emission with the new radiation field, and iterating until a set convergence level is reached.

In simple MC RT, the output desired from the calculation is usually the view of the system from a particular observer location. Images of the system are constructed by projecting the position of each photon that escapes in the direction of the observer within some angular tolerance onto the plane of the sky. This plane

⁶It is possible to have cases where the optical depth is very large and the number of scatterings is correspondingly large. It is common to impose a limit on the number of scatterings to calculate as a result. This has the consequence that high optical depths are not sampled well with potential systematic errors in the calculation. For most cases, the systematic errors from imposing a maximum number of scatterings are negligible.

is divided into pixels and the images are built from those photons. A large number of the photons will escape the system in directions other than the observer and are not counted, unless symmetries (e.g., circular or cylindrical) can be exploited. Special care must be applied when constructing the images which give the polarization state of the scattered flux, specifically the Q and U images. As part of the construction of these images, the polarization vector of the photon must be rotated so that it is referenced correctly in the image plane.

5.2. Weighted MC RT

Simple MC RT is the easiest to understand, but it would be extremely inefficient for full 3D calculations. For 1D and 2D MC RT calculations, there are symmetries, such as spherical and cylindrical, that can be exploited and thus the number of photons needed to achieve accurate results is relatively small. In the case of 3D dust RT, there are no symmetries by definition; this has been one of the motivations for a number of acceleration methods. Most of these acceleration methods are well established and validated, while others still have a more experimental character. Almost all of these methods were first developed for 2D RT calculations. The benefits of weighted MC techniques as compared to simple MC RT methods are illustrated in Fig. 4.

The basis of all acceleration methods is to assign a weight W to each photon and modify this weight. The weight of each photon is equivalent to the fraction of the luminosity of the emission source carried by that photon i.e. the number of photons in each photon package. Several acceleration techniques use the idea of biasing, i.e. generating random numbers from a PDF $q(x) dx$ rather than from the appropriate PDF $p(x) dx$. This biased behavior is corrected for by assigning the weight $W = p(x)/q(x)$ to the photon. The biasing technique is used in many MC applications, and can be a very effective way of reducing the variance (Dupree & Fraley 2002). It is, for example, common practice in MC numerical integration where it is known as importance sampling.

Biased emission

A direct application of the biasing technique is the so-called biased emission. The initial propagation direction of the photons launched into the dust medium is usually determined from the angular part of the

source term $j(\mathbf{x}, \mathbf{n}, \lambda)$. There are cases, however, where an increased level of emission in particular directions is desired, for example to increase the signal-to-noise in particularly interesting directions, such as the polar regions of a star with an accretion disk (Yusef-Zadeh, Morris & White 1984). In this case, the emission of photons is biased towards the directions of interest, and the initial weight of the photon is determined using the standard biasing weight formula. The same technique can also be applied to the spatial part of the source term to increase the number of photons emitted from regions with a low emission rate (Juvola 2005).

This technique of biased emission has the potential of strongly increasing the efficiency of a MC simulation. On the other hand, it has the drawback that it is very model-dependent, and therefore requires significant manual interaction. It is, in a sense, comparable to the placement of the rays in the RayT method.

Absorption-scattering split

This acceleration method allows for a photon to contribute to both absorption and scattering at each interaction site. Instead of randomly choosing the nature of the interaction, the photon is split in two parts: one that is absorbed and one that scatters. The fraction absorbed is equal to $(1 - a)$ times the current weight of the photon. For the scattered part that continues its life cycle through the dust, the weight is multiplied by a factor a .

Forced scattering

Instead of having each photon either scatter or escape the system, the scattering can be forced to occur every time (Cashwell & Everett 1959). In the simple MC routine, the randomly generated optical depth is compared to the total optical τ_{path} along the photon's path. This approach has a problem for regions with low optical depth: in those regions, many photons just leave the system without interacting with the dust, resulting in a low efficiency of dust scattering and heating. A way to avoid this low efficiency is the technique of forced scattering, which limits the values of the randomly chosen optical depths to the range between 0 and τ_{path} . This can be achieved by biasing the PDF from which the optical depth is generated. Instead of sampling from the actual exponential PDF, we consider an exponential distribution cut off at $\tau = \tau_{\text{path}}$.

Properly normalized, this PDF reads

$$q(\tau) d\tau = \begin{cases} \frac{e^{-\tau} d\tau}{1 - e^{-\tau_{\text{path}}}} & \tau < \tau_{\text{path}}, \\ 0 & \tau \geq \tau_{\text{path}}. \end{cases} \quad (28)$$

Generating a random τ from this distribution is straightforward, and guaranteed to produce an interaction before the photon has left the system. The compensation to be applied to the weight of the photon is a factor

$$W_{\text{fs}} = \frac{q(\tau)}{p(\tau)} = 1 - e^{-\tau_{\text{path}}}. \quad (29)$$

One issue with forcing every scattering, when combined with the absorption-scattering split, is that there is no longer a natural stopping criterion for the scattering calculation. In the original MC cycle, photons end their journey when they are either absorbed by the dust or leave the dusty system. The common solution is to set a minimum weight for a photon, below which the photon's life cycle is terminated. The value of this termination weight is usually set to be very low, after the equivalent of many scatterings. An alternative solution is to only force the first scattering and revert to the standard scattering calculation for subsequent scatterings.

Peel-off technique

For 3D cases, it is usually desired to calculate the appearance of a system for an observer at a particular orientation. Simple Monte Carlo RT is particularly inefficient in building up such an image as only the photons that are emitted from the system in the direction of the observer contribute to the output appearance. In addition, some blurring of the image is inherent as photons that are within some tolerance of the desired direction are used. This inefficiency can be eliminated by requiring that all photons directly contribute to the output images by calculating the portion of the photon that is emitted from sources and scattered at every interaction point in the observer's direction (Yusef-Zadeh, Morris & White 1984). The weight factor of a photon in the direction of the observer is

$$W_{\text{po}} = p(\mathbf{n}_{\text{obs}}) e^{-\tau_{\text{obs}}} \quad (30)$$

where τ_{obs} is the optical depth from the position of the emission or scattering event, and $p(\mathbf{n}_{\text{obs}})$ is the probability that the photon would be directed towards the observer. For example, for isotropic emission, $p(\mathbf{n}_{\text{obs}}) = 1$, and after a scattering event we have $p(\mathbf{n}_{\text{obs}}) = \Phi(\mathbf{n}, \mathbf{n}_{\text{obs}}, \mathbf{x}, \lambda)$.

Peel-off is a pure application of RayT, again highlighting the connection between both approaches. Peel-off is probably the most important acceleration technique for 3D MC RT simulation. It has a significant computational cost, however, as it requires a calculation of the optical depth from the emission or scattering location towards the observer after every emission or scattering event. One way to alleviate this computational burden is by precalculating the optical depth towards the observer for each cell. Another possibility is to store the information per cell and create the images at the end of the simulation (e.g. [Dullemond & Turolla 2000](#); [Pinte et al. 2006](#); [Min et al. 2009](#)). Using a precomputed τ_{obs} or using stored information to compute the image does come with a price, loss of subgrid resolution in the model images. The benefit of this approximation in decreased computation time has to be weighted against the loss of subgrid resolution for the particular modeled astrophysical object.

Continuous absorption

The accuracy of the dust re-emission of the absorbed energy can be enhanced by absorbing not just at the interaction site, but along the path the photon travels. Depending on the implementation, this can be done up to the location of the scattering event ([Lucy 1999](#)), or along the entire path ([Niccolini, Woitke & Lopez 2003](#); [Baes et al. 2011](#)). In the latter scenario, the photon is effectively split in $N + 2$ different parts: one part W_{esc} that leaves the system (and is hence not accounted for anymore), one part W_{sca} that is scattered at the interaction location (determined stochastically), and N parts $W_{\text{abs},j}$ that represent the fraction absorbed in each of the N cells along the photon's path. These different fractions are

$$W_{\text{esc}} = e^{-\tau_{\text{path}}}, \quad (31)$$

$$W_{\text{sca}} = a(1 - e^{-\tau_{\text{path}}}), \quad (32)$$

$$W_{\text{abs},j} = (1 - a)(e^{-\tau_{j-1}} - e^{-\tau_j}), \quad (33)$$

where τ_j is the optical depth measured from the photon's location to the surface of the j 'th cell along the path. An alternative interpretation of this continuous absorption approach is that it estimates the mean intensity of the radiation field not by means of the absorbed luminosity, but by means of counting the luminosity that passes through each cell. The strength of the technique lies in the fact that all photons contribute to the calculation of the absorption rate of each cell they pass

through, and not only of those cells with which they interact. This is particularly useful for the optically thin regime, which has very few absorptions in the simple MC approach.

Instantaneous dust emission

The traditional method in computing the self-consistent dust emission in a MC RT simulation consists of running an independent MC simulation at every individual wavelength and storing the absorbed luminosity over the computational grid. In a second stage, the dust emission spectrum is calculated in each dust cell and used as secondary source term. This approach inevitably leads to iteration, as the dust emission itself affects the radiation field. It is possible to compute the output dust emission spectrum from a multiwavelength model without iterating. The instantaneous dust emission technique or frequency distribution adjustment technique emits dust thermal photons immediately after each absorption event in a specific grid cell, with the wavelength of the emitted thermal dust photon carefully chosen to retain thermal equilibrium ([Bjorkman & Wood 2001](#), [Baes et al. 2005](#)).

Besides eliminating the need for iteration in the computation of the dust emission spectrum, the instantaneous dust emission technique also has the advantage that photon packages are emitted at the exact position where the absorption event took place, so subgrid resolution is achieved. One disadvantage is that the absorption/re-emission event takes place in a single cell, which results in a poor convergence rate of the radiation field in cells with a low dust absorption rate. This makes the classical iterative technique with continuous absorption more efficient than the instantaneous dust emission technique, at least when applied in its original form, for 3D simulations (e.g. [Chakrabarti & Whitney 2009](#)). This problem can be alleviated by applying a combination of the instantaneous dust emission and the continuous absorption techniques. In this hybrid method, the photon packages are followed through the domain using the instantaneous dust emission technique, but the final dust emission spectrum of the cells, used to create images and SEDs, is calculated based on the continuous absorption approach ([Pinte et al. 2006](#)). A second problem for the instantaneous dust emission technique is that it was originally designed to work with equilibrium dust emission. Several ways have been explored to adapt the instantaneous dust emission technique for transiently heated

grains (Krügel 2007, Wood et al. 2008, Heymann & Siebenmorgen 2012).

High optical depths

In regions of high optical depth, the simple MC routine becomes very inefficient, as photons can be trapped in a virtually never-ending loop of scattering events. This problem is largely solved when the absorption-scattering split is applied, as the weight of the photon then decreases with every scattering event terminating when a very low weight is reached. However, in regions with extreme optical depths (such as in the midplane of circumstellar discs around protostars) or at wavelengths where scattering largely dominates absorption (such as the far-UV), this can still imply a significant computation burden. A solution to this problem is to mirror or reflect photons from high optical depth regions and use the diffusion approximation to find the RT solution in these regions. The diffusion approximation allows for multiple interaction steps to be calculated in a single computation. An elegant solution that is well adapted to the MC method is to solve for the RT in optically thick cells using a modified random walk technique that also uses the diffusion approximation (Min et al. 2009; Robitaille 2010).

Polychromatism

For multi-wavelength RT, it is possible to significantly speed up the calculation by considering photon packages that consist of photons of all wavelengths. The advantage of this technique is that a MC run is simultaneously solved at all wavelengths, instead of a run for each wavelength. The difficulty in this approach is that many of the PDFs that describe the life cycle of a photon are wavelength dependent, such as the path length distribution or the scattering phase function. One solution is to consider partly polychromatic photon packages, which shift to monochromatism as soon as wavelength-dependent PDFs are involved (Baes, Dejonghe & Davies 2005). A more advanced option is to perform the calculations at one reference wavelength λ_{ref} and use the biasing technique to adjust for the wavelength-dependent PDFs (Juvella 2005; Jonsson 2006). The efficiency gains of full polychromatic RT are large given that every photon calculated would contribute to the output images and radiation field density at all wavelengths instead of a single wavelength. But there is a known significant complication – the biasing factors can be very large and, as

a result, can dominate the results at a particular wavelength. This has a systematic effect on the results that is hard to control and therefore, this method should be considered experimental.

5.3. Uncertainties for Monte Carlo

In the simple (unweighted) MC RT, the noise in the output quantities (i.e., scattered intensity, polarization, etc.) scales as $N^{-1/2}$, where N is the number of photons. In the case of weighted Monte Carlo, the uncertainties in the output quantities does not scale directly with $N^{-1/2}$.

The uncertainties can be calculated by using the dispersion in the average properties of the photons used to determine an integrated quantity (Gordon et al. 2001). If the integrated quantity is X , then

$$X = \sum_{j=1}^N x_j = N \bar{x}, \quad (34)$$

where x_j is the contribution of the j 'th photon to X , N is the total number of photons in the model run, and \bar{x} is the average contribution each photon makes to X . The uncertainty in X is then $\sigma_X = X(\sigma_x/\bar{x})$, where σ_x is the standard deviation of \bar{x} , calculated using

$$\sigma_x^2 = \frac{1}{N(N-1)} \sum_{j=1}^N (x_j - \bar{x})^2 = \frac{1}{N-1} (\overline{x^2} - \bar{x}^2). \quad (35)$$

The equations for the uncertainties in output quantities given above can be used to explicitly enforce a particular uncertainty level in the final results of a model run. The uncertainties in the output quantities of interest can be checked during the model run and the number of photons adjusted dynamically to achieve the desired accuracy. For example, when calculating a multi-wavelength model more photons can be put at the wavelengths with higher optical depths compared to wavelengths with lower optical depths. As the number of photons is determined from the statistics of the particular run itself, it automatically takes into account the full RT solution, including the locations of the photon sources, dust, etc. The output quantities to be used for uncertainty control can range from the global flux, to pixels in resolved images, to the radiation field density in each cell. Additionally, it is possible to improve the convergence to the needed accuracy by identifying areas of the model with high noise (e.g., particular cells with few photon interactions) and dynamically adding photons to those areas (e.g., though bi-

ased emission). The use of dynamically determine uncertainties to improve the convergence of a model has been done in limited cases and is clearly an area for future improvement.

Uncertainties associated with the model setup itself are more difficult to quantify. Evaluating the systematic uncertainties due to specific model choices is usually done by trying other parametrizations or increasing the grid resolution and evaluating how the output quantities change. Example model setup choices that are prone to systematic uncertainties are the dust density grid, specific parametrizations of the photon sources, and wavelength grid. These kinds of uncertainties are clearly the most difficult to diagnose and generally rely on the expertise of the coder and user.

6. CHALLENGES IN MODELING OBSERVATIONS

The uses of 3D dust RT codes are many, from understanding the impact of locally clumpy dust on dust RT (Witt & Gordon 1996, 2000; Bianchi 2008) to modeling observed images of objects to derive the source and dust distributions (De Looze et al. 2012b; Steinacker et al. 2005; Gordon et al. 1994) to predicting the appearance of an object that has been modeled with a HD code (Steinacker et al. 2004; Bethell et al. 2004; Jonsson, Groves & Cox 2010; Robitaille 2011) to investigating the ensemble behavior of objects to derive average source and dust properties (Law, Gordon & Misselt 2011) to test observability of objects for instrument and observation planning (Wolf 2008; Gonzalez et al. 2012). To illustrate the challenges in modeling observations, we focus on the modeling of objects to derive their physical properties from observational data. The common approach is to perform some steps of the following scheme:

1. Choose a parametrized model for the dust density, radiation sources, and dust optical properties.
2. Discretize the problem with the choice of grids and use the RT code to derive a SED and/or images for this object model.
3. Compare simulated and observed SED or images.
4. Evaluate the differences in the SED or images and change the model parameters and/or the

model assumptions to minimized these differences. Repeat step 2.

5. Find the parameter sets that come closest to describing the observed data.
6. Evaluate the observational and theoretical errors that are important for the comparison.

Although the scheme seems straightforward, the full set of steps has rarely been carried out for existing, published 3D dust RT models as each of these steps provides challenges for 3D dust RT.

Model choice

The aim is to choose an appropriate model that allows meaningful physical information to be derived from the given observational data. This is a general topic, but 3D RT modeling has special features which are important to consider. The data are usually SEDs and/or images and we can investigate the number of independent information bits a priori, e.g. by counting the number of wavelengths for which the flux density values have been measured. This is then compared to the number of model parameters.

Most 3D models include several 10 to hundreds of free parameters depending on the complexity of the spatial model and the assumed dust properties. Even for attempting only to reproduce the global SED of an object, it is difficult to stay below 10 parameters since modeling a dusty 3D structure requires a few length scales and/or power laws to describe density, dust properties, and viewing angles. Fig. 5 illustrates the large number of free parameters needed to reproduce the 3D dust distribution of a molecular cloud. Fig. 6 shows four examples of state-of-the-art RT fits to dusty galaxy images. Such RT model fits are designed to determine the parameters that describe the intrinsic 3D distribution of stars and dust both large-scale geometric parameters (e.g., stellar and dust scale lengths and heights) and measures of the small scale inhomogeneity or clumpiness of the dust.

Starting with more free model parameters than data points is often considered to produce meaningless parameter values. However, a careful exploration of the parameter space can detect ambiguities and reveal redundant parameters. As long as the parameter space exploration can be afforded computationally, starting with a detailed model and analyzing the parameter ambiguities is usually the most accurate way to model an object.

Gridding

As discussed in §3, there are a variety of options for discretizing space, direction, wavelength, and the dust model. If the grid fails to resolve structures that are important at a particular wavelength, the corresponding model image will be correspondingly inaccurate. A good example is the surface layer of an accretion disk illuminated by a central star. If the layer is not well-resolved, the description of the radiation field that is scattered in the layer will be poor, and the scattered light images inaccurate. Therefore, the discretization step often includes running the RT code before the start of the modeling to verify that expected features are present in the image or SED, and that changes/refinement do not influence the overall appearance of the object.

Comparison of model and data

Once the model images or SED have been calculated, the results should be convolved with the beam of the observed instruments/telescopes and the sampling should be made equal (pixel size for images). It is important to apply detection limits, especially when studying faint structures. For interferometric observations, the incomplete coverage of all spatial scales implies that the comparison should be performed in the (u,v) plane rather than in spatial plane to achieve the highest fidelity. Calibration uncertainties and correlated noise properties in the observations should be understood and included in the model versus observations comparison. In addition, foreground emission (e.g. caused by the zodiacal light) should be carefully removed or modeled, especially for on-off chopping observing modes that remove this emission as part of the data reduction process. Generally, modeling can benefit from good communication with experienced observers. Another potential source of error in comparing model results and data can be a displacement of structures due to observational uncertainties. Additional translation parameters can correct this and provide physical insight by deriving improved positioning.

Exploration of the parameter space

The parameter space of 3D dust distributions is large: a uniform coverage of a 10-dimensional parameter space with 5 grid points in each parameter, would imply about 10 million individual RT calculations. Therefore, almost all 3D RT modeling of data

has been done "by hand," that is starting from a point in the parameter space and then varying just one parameter to explore the variations in the results. This drastically reduces the coverage in parameter space at the expense of not exploring correlations between parameters. In structures with strong gradients, inhomogeneously distributed radiation sources or varying optical depth, the resulting radiation field may strongly react to changes in the model parameter. In this case, low-coverage "by hand" explorations are likely quite unreliable.

To evaluate the model and the variation in the model parameter, the difference between the model data and the observed data is defined (often using a χ^2 metric) and the fitting procedure is then to minimize this difference by varying the parameters. There are standard optimization methods that have been applied in astrophysics such as gradient descent, Newton's method, Metropolis optimization, or genetic algorithms. The latter two are able to leave local minima with the goal of ending in the deepest minimum providing the best-fitting model parameters.

There are advancements in the application of automated fitting techniques that are already in use for 2D RT calculations and are or should be used for 3D RT calculations as well. One such technique is the Metropolis algorithm that has been applied to 3D dust RT in the form of simulated annealing (Steinacker et al. 2005). Precomputing a large grid of models is a promising technique. Robitaille et al. (2007) used a parametrized circumstellar disk model to create just such a large grid of precalculated SEDs for a 2D configuration and then used the grid to fit the disk parameters and to characterize uncertainties in fit parameters. Other promising techniques are 2D RT fitting techniques based on the Levenberg-Marquardt algorithm (e.g. Xilouris et al. 1999), downhill-simplex method (Bianchi 2007), and genetic algorithms (De Geyter et al. 2012). Some of these techniques have already been applied to 3D structures, albeit in limited parameter spaces (e.g. Witt & Gordon 2000; Schechtman-Rook, Bershadsky & Wood 2012).

Quantifying the ambiguity of the derived parameters is also important. The 2D circumstellar disk modeling of SEDs in Robitaille et al. (2007) provides a template for determining the ambiguity when a complex model is applied to only a few SED points. A general strategy to assess ambiguity is to explore the variation of the fitting metric (e.g., χ^2) in the vicinity of the best fit parameters, as well as to characterize

the overall degree of variation by random or grid-based parameter space exploration.

Error analysis

Each of the RT solvers has its own sources of errors due to approximations, the underlying grid, undetected regions of the computational domains, etc. Tracing these errors, measuring them, and estimating their importance is non-trivial. RayT provides error control for the solution along one ray and MC can provide error control on the number of photons needed, but whether the rays are sent through the right parts of the domain is not well quantified. For modeling, a comparison of these errors with the solution and observational errors determines the sensitivity of the observed data to the specific science question of interest.

Inverse RT

The true challenge of RT modeling is to invert it: determine the 3D density structure of the dust and the sources, and the dust properties from a set of images taken at different wavelengths. Direct inversion is computationally impossible in 3D with current or expected computing capabilities. In 1D an analytic inversion can be performed under special assumptions (Steinacker, Michel & Bacmann 2002), and numerical 1D inverse RT modeling was used in Doty et al. (2010). A forward RT method using many iterations provides a good solution, but is limited to fairly simple RT applications. A prominent difficulty is the loss of information due to the line-of-sight integration inherent in the observations. Multiwavelength observations can help disentangle the line-of-sight integration. For example, points in the center of a molecular core are better shielded and cold compared to points in the outer parts, so mm observations can be used to constrain the core center and shorter wavelength observations the outer parts. Such molecular cores are simple enough to enable inverse RT to be done. For example, Steinacker et al. (2005) used a 3D background radiation field illuminating a 3D model core and fitted the MIR and mm images of the core Rho Oph D, deriving the density and temperature structure with assumed dust properties.

7. CODES AND BENCHMARKS

7.1. Available 3D codes

Until the mid-1990s, few RT codes could handle the absorption, scattering, and thermal emission by interstellar dust in general 3D geometries. The available codes were either limited to 1D or 2D geometries, or were not able to calculate full 3D problem (e.g., missing either scattering or thermal emission). The spectacular increase in computational capabilities during the past two decades, as well as the development of more powerful techniques to solve the RT problem, has led to the creation of a number of 3D RT codes. There are now (to our knowledge) almost 30 codes operational that can handle the full dust RT problem, i.e. including absorption and multiple anisotropic scattering, in a general 3D geometry. Table 1 gives a list of published 3D dust RT codes that have been or are being used for astronomical applications. The application fields of the different codes varies widely, ranging from prestellar cores and circumstellar discs to active galactic nuclei and galaxies.

With many different codes and several techniques available, it might be difficult for a potential user or future developer to decide which one to choose. The choice for a given code or solution technique should primarily be driven by the specific nature of the problem that is to be solved. While most of the 3D codes in Table 1 are applicable for a wide range of RT problems, virtually all of them been developed with a particular application in mind, and hence have been optimized for that particular goal.

The programming language can also influence the choice; most existing 3D RT codes have been coded in FORTRAN, C, or C++. There are no features in the 3D RT problem itself or the two solution methods presented in this review that make certain languages preferential to others. The main driver is speed: since the 3D RT problem is computationally very challenging, all codes should be developed in a language suitable for high-performance computing. Moreover, parallelization is becoming increasingly important as a means to increase computational speed and memory availability. Several codes are designed to work on shared-memory or distributed-memory clusters and/or adopt graphical processing units (e.g. Jonsson 2006; Jonsson & Primack 2010; Baes et al. 2011; Robitaille 2011; Heymann & Siebenmorgen 2012).

The choice of a given code or method primarily de-

Table 1: List of published 3D dust RT codes that have been or are being used for astronomical applications. The codes are ordered alphabetically as it is difficult to establish true code development dates given that many codes are first used for specific investigations with the full code details (references in table) published at a later date.

name	type ^a	reference	main application
SKIRT	MC	Baes+ 2003; 2011	galaxies, AGNs
(no name)	MC	Bethell+ 2004; 2007	SF clouds
TRADING	MC	Bianchi+ 1996, Bianchi 2008	galaxies
RADISHE	MC	Chakrabarti+ 2007; 2009	galaxies
(no name)	MC	Doty+ 2005	SF clouds
RADMC-3D	MC	Dullemond (in prep.)	SF disks
MOCASSIN	MC	Ercolano+ 2005	photoionized regions
(no name)	MC	Fischer+ 1994	SF disks
(no name)	MC	Gonçalves+ 2004	SF clouds
STOKES	MC	Goosmann+ 2007	AGNs
DIRTY	MC	Gordon+ 2001, Misselt+ 2001	galaxies, nebulae
TORUS	MC	Harries 2000, Harries+ 2004	SF disks
(no name)	MC	Heymann+ 2012	SF disks, AGNs
SUNRISE	MC	Jonsson 2006, Jonsson+ 2010	galaxies
CRT	MC	Juvela+ 2003, Juvela 2005	SF clouds
(no name)	MC	Lucy 1999; 2005	supernovae
MCMax	MC	Min+ 2009; 2011	SF disks
STSH	MC	Murakawa+ 2008	SF disks
MCTRANSF	MC	Niccolini+ 2003; 2006	SF disks
mcsim mpi	MC	Ohnaka+ 2006	carbon stars
MCFOST	MC	Pinte+ 2006	SF disks
HYPERION	MC	Robitaille 2011	SF clouds
PHAETHON	MC	Stamatellos+ 2004; 2005	SF cores
STEINRAY	FD	Steinacker+ 2003	SF disks
	RayT	Steinacker+ 2006	SF cores
(no name)	FD	Stenholm+ 1991	SF disks, AGNs
HO-CHUNK	MC	Whitney+ 2002	SF disks
MC3D	MC	Wolf+ 2000, Wolf 2003b	SF disks, SF cores, AGNs
(no name)	MC	Wood+ 1999, Bjorkman+ 2001	SF disks, galaxies

^a finite differencing (FD), Monte Carlo (MC) or ray-tracing (RayT)

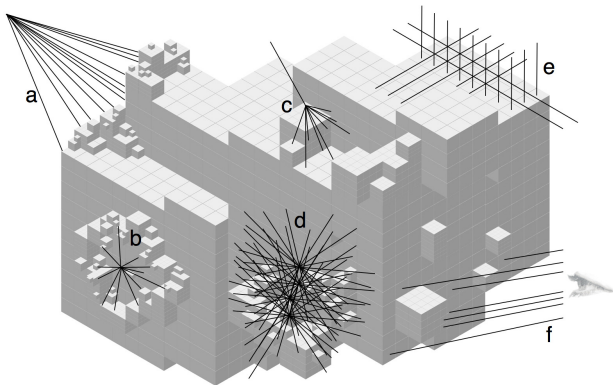


Fig. 3.— Illustration of the various types of rays used in 3D dust RT calculations performed within an adaptively refined density grid. Only a small fraction of the typical number of rays are shown for (a) an outer radiation source, (b) an inner radiation source, (c) a scattering event, (d) a region with an optical depth near 1, (e) a coarse regular outer grid, and (f) rays to the observer.

depends on the specific needs of the application and/or the personal preferences of the user or developer. Nevertheless, the different approaches have some general strengths and weaknesses. These need to be considered as rough guidelines only, as there are significant differences among codes based on the same approach. For example, simple MC codes based on the naive techniques explained in §5.1 are many orders of magnitude less efficient than modern MC codes that use weighting schemes.

Generally speaking, setting up a reasonably efficient 3D MC RT code fits within the lifetime of a PhD project. Most MC RT codes are intrinsically 3D codes, and hence the shift from 1D or 2D codes to a full 3D geometry is fairly straightforward. For RayT codes, on the other hand, the increase in complexity when moving from 1D and 2D to 3D is much steeper. These differences explain the relative scarcity of general 3D RayT codes compared to MC codes in Table 1; a similar comparison of 1D or 2D codes would result in a table with a much larger fraction of RayT codes. A big part of this complexity lies in the placements of the rays, which needs manual adjustment in RayT codes, but is done automatically in MC codes. The RayT pre-

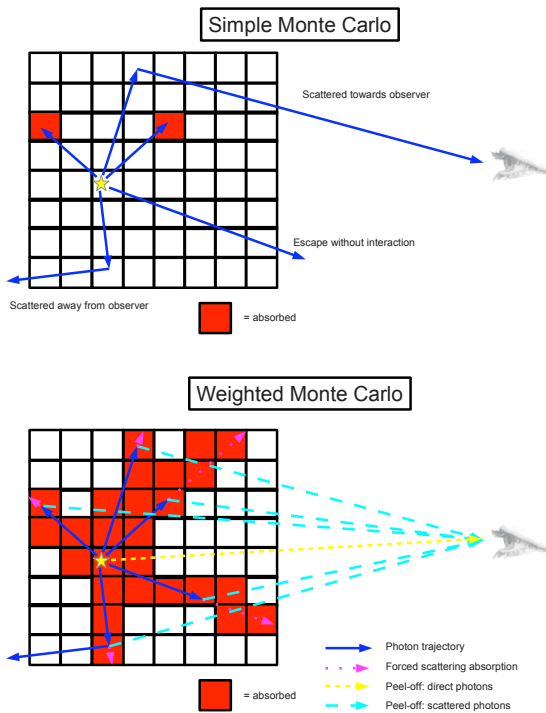


Fig. 4.— Both simple and weighted Monte Carlo are illustrated graphically. This example includes 5 photons resulting in one scattering photon reaching the observer and two absorption events for simple MC RT. This same set of 5 photons produces 5 scattered photons towards the observer and 28 absorption events for weighted MC. The improvement in computational efficiency of weighted Monte Carlo is clearly seen.

calculation step needed for the manual placement of rays reveals an advantage for this solver as it identifies the critical locations in the model that dominate the appearance of the object. Other advantages of the MC method are that it needs less storage than RayT codes when scattering is included, and that it is widely used with many coders in the community improving its application to astrophysics with new algorithms, whereas RayT methods are mainly improved outside astrophysics. One strength of the RayT codes is the stronger and explicit error control. For example, the precalculation step can identify areas that are potentially under-resolved allowing for changes to the grid to provide adequate resolution for the full calculation. At this point, MHD codes tend to use RayT rather than MC techniques (e.g. Heinemann et al. 2006; Kuiper et al. 2012).

7.2. Benchmark efforts

Probably the most objective way to compare the strengths and weaknesses of the different codes is using benchmark problems. In the past few years, there have been benchmark efforts in many different computational astrophysics areas, including molecular line transfer (van Zadelhoff et al. 2002), halo and void finder algorithms (Colberg et al. 2008; Knebe et al. 2011), astrophysical hydrodynamics (Agertz et al. 2007), cosmological hydrodynamical simulations (Frenk et al. 1999; O’Shea et al. 2005) and cosmological radiation hydrodynamics (Iliev et al. 2006, 2009). At the moment, no code validation or benchmark project exists for 3D dust RT. The most advanced dust RT code validation projects are 2D benchmarking efforts.

Pascucci et al. (2004) presented a benchmark test for two-dimensional equilibrium RT problems. Their system consists of a single star surrounded by a flared axisymmetric accretion disc. The optical depth through the disc varies from $\tau_V = 0.1$ to $\tau_V = 100$. The accretion disc contains strong density gradients, which makes it an ambitious benchmark problem (unfortunately, anisotropic scattering is not taken into account). The authors compare the temperature maps and spectral energy distributions for five 2D dust RT codes (two grid-based codes and three MC codes). Differences between the various codes in the temperature maps are smaller than 1% for the most optically thin model, but reach up to some 15% in the most optically thick system. For the emerging spectra, the differences range from a few percent for the optically

thin models to more than 20% for the most optically thick models.

A first extension of the [Pascucci et al. \(2004\)](#) benchmark test was presented by [Pascucci et al. \(2003\)](#)⁷, with two of the five codes from the original benchmark participating (one MC code and one RayT code, both intrinsically 3D codes). They considered a similar disc as in the [Pascucci et al. \(2004\)](#) benchmark, with an azimuthal ring added as a simplified model for a spiral density distortion. The main advancements were the addition anisotropic scattering and that images and visibilities, in addition to SEDs and temperature maps, were compared. The difference in flux in the entire image was smaller than 20%, but the visibility curves showed substantially larger discrepancies in their overall shape.

[Pinte et al. \(2009\)](#) go another step further, in what is the most advanced dust RT benchmark to date. They start from a similar circumstellar disc model, but consider optical depths up to $\tau_V = 10^6$, use anisotropic scattering, and compare images and polarization maps. The four different codes that participate in the benchmark are all MC codes. The agreement in the temperature distributions is very good, with differences almost always smaller than 10%. Differences in the SEDs remain smaller than 15% for models with $\tau_V = 1000$ and agree within 20% over almost the entire wavelength range for the most optical thick cases. Pixel-to-pixel differences in high-resolution scattered light images remain limited to 10% and the polarization maps do not differ by more than 5° in regions where the polarization can be effectively measured by observations (Figure 7).

8. THE FUTURE OF THE FIELD

Present status

3D dust RT is a rich and diverse field, with applications across a broad range of astrophysical topics from dust near stars to entire galaxies. Correctly modeling the effects of dust on the transfer of radiation is critical to studying many astrophysical objects, including the dust itself. Recent years have seen an impressive improvement in observational capabilities across the electromagnetic spectrum and this has shown that the dust distribution in many regions is strongly 3D. This requires methods to compute the dust radiative trans-

fer that can handle 3D structures and return solutions in a reasonable amount of time. The most common 3D dust RT solver is based on MC techniques, with RayT features in its modern accelerated form. A few applications have used pure RayT solvers. Both methods face the challenges of grid discretization, determination of uncertainties in solutions, and accurate comparison between observations and the model calculations. Almost 30 codes are currently able to deal with the full 3D dust RT with code variations arising from the prime field of application. There is no 3D dust RT benchmark, currently code comparisons are done using 2D benchmarks.

General trends

Several trends indicate that the future of 3D dust RT is bright. The number of people actively involved in 3D dust RT is growing and the number of new published codes has increased significantly in recent years. A 3D approach to modeling complex distributions is becoming common in many fields featuring 3D dust distributions. The continuing increase in available computing power and storage will support this trend, allowing a full transition from 2D to 3D dust RT for all objects showing 3D signatures. A prominent example of this trend is circumstellar disks with (proto-)planets, where the (M)HD simulations have been 3D for years, dust RT modeling often was 2D, and observations are now reaching the resolution to identify the 3D signatures of disk deformation due to a planet. In addition, modern online tools are expected to support the access to the codes by users through sophisticated interfaces.

Future benchmarks

For progress in 3D dust RT to continue, 3D dust RT benchmarks need to be established. Given the complexity of the codes, the increasing number of acceleration algorithms, and large number of specific applications, it is critical to provide a quantitative comparison between codes. Experience with existing dust RT benchmarks and similar efforts in other areas indicate a suite of 3D dust RT benchmarks is needed. Ideally, each benchmark would focus on a particular part of the RT solution (e.g., scattering, polarization, equilibrium dust emission, or non-equilibrium dust emission) in a 3D geometry. This would provide a clear test of different aspects of 3D dust RT and support the participation of all codes in at least part of the suite.

⁷This paper appeared in print before the [Pascucci et al. \(2004\)](#) paper, but was actually a follow-up project.

Data modeling future

Given the impressive flow of new data now and coming years from ground- and space-based observatories, it is clear that the demand to accurately model 3D dusty structures will rise strongly. Interfaces that can simulate observations with different telescope properties will become necessary to perform modeling. We expect a rise of 3D dust RT modeling efforts that rely on automated fitting processes rather than by-hand explorations of the model parameter space. Since the number of sources of multi-wavelength data will rise, collaborations crossing borders between observers and modelers will become more frequent. The ultimate goal of 3D dust calculations is to modeling multi-wavelength images and derive quantitative and statistically sound information about 3D structures, embedded sources, and the dust itself.

Future connections to non-dust RT codes

Another future direction is the coupling of 3D dust RT codes with codes describing other physical effects in astrophysical objects. This trend is already happening with 2D dust RT codes and the extension to 3D dust RT codes is clearly the next step. A variant of this type of connection is already happening where 3D dust RT is used to calculate the radiation field in a dust distribution generated with a (M)HD code. Furthermore, (M)HD codes that make use of simple dust RT could be tested or the simple algorithms improved by comparison to full 3D dust RT solutions. Chemical network calculations could be based on a more realistic estimate of the incoming radiation calculated from 3D dust codes. Finally, a combined calculation for 3D line and dust RT would enable line and continuum data to be simultaneously investigated using the same underlying physical model.

Future algorithms

In the past, conferences and keyword related publication searches have often been used to improve the unfortunately rare communication of new numerical algorithms from applied mathematics to astrophysics. The basic issue is the sheer flow of new findings and the different language of the two communities. Recent MC improvements have been mainly developed by coders working in the field, and additional efforts should be made to enable community-crossing exchange on algorithms and error control. As a result of communications between coders preferring differ-

ent solvers, we expect hybrid solvers making use of the advantages of the various approaches to appear more frequently. Given the increase in complexity in the modeled objects, we expect future activities to establish grid generation algorithms that are optimized for 3D dust RT; besides the octree or AMR-style grids that are now routinely implemented in 3D dust RT codes, unstructured grids as used in line RT (Paardekooper, Kruip & Icke 2010) and (M)HD codes (Springel 2010) are an interesting alternative. The inclusion of time dependence in the 3D RT problem, which could be important in the context of star formation or episodic accretion, will also need to be tackled with new algorithms (e.g. Harries 2011). The increasing availability of massively parallel machines will support algorithms that are optimized to run on many processors.

Input physics improvements

The improvement of the solvers is not restricted to developing algorithms that provide accelerated solutions. The interaction of radiation with cosmic dust is still not fully understood, and the variation of the dust properties with environment is an area of active research. The various continuum radiation sources like stars, PDRs, AGN accretion discs, and the ISRF are areas of vigorous investigation. For example, efforts based on existing and upcoming large scale surveys are being used to update the 3D structure of the stars in the Milky Way. Consequently, we expect to achieve a better understanding of the observed radiation from future research on the optical properties of dust, and improved data on the stellar and non-stellar sources that enter the 3D dust RT equation.

Challenges

A major challenge in 3D dust RT that emerges from this review is how to account for and mitigate systematic uncertainties in the dust RT solution. They arise from under-resolving grids, from not propagating rays/photons to important cells, and/or from uncertainties in the underlying dust grain models. As under-resolving of the dust and radiation field grid is often a result of constraints on computer memory and speed and improvements in algorithms to implicitly handle optimal grids are clearly needed. The pre-processing steps necessary for the RayT solver address some of these issues, but need further automating. The issue of not propagating enough rays/photons into particular cells has been solved for both RayT (placement of

rays) and MC (biased emission), but in both cases currently requires hand-tuning. An algorithm to automatically add additional rays/photons similar to that used for adaptive mesh refinement would clearly be useful. Finally, uncertainties in the assumed dust grain model provide a systematic uncertainty in the dust RT modeling that is difficult to quantify. Different dust grains models can be used to provide an estimate of this uncertainty, but the best way to reduce this uncertainty is to support the improvement of dust grain models through the use of improved laboratory and observational data.

ACKNOWLEDGEMENT

The authors acknowledge the support of the Ghent University for two excellent week long meetings in Ghent, Belgium where a large portion of the work on this review was done. We thank Simon Bruderer, Jacopo Fritz, Gianfranco Gentile, Michiel Min, Kirill Tchernyshyov, Ewine van Dishoeck, and Adolf Witt for providing comments on this review that significantly improved it.

9. LITERATURE CITED

REFERENCES

- Abel T, Wandelt BD. 2002. *MNRAS* 330:L53–L56
- Agertz O, Moore B, Stadel J, Potter D, Miniati F, et al. 2007. *MNRAS* 380:963–978
- André P, et al. 2010. *A&A* 518:L102
- Arzoumanian D, André P, Didelon P, Könyves V, Schneider N, et al. 2011. *A&A* 529:L6
- Baes M, Davies JI, Dejonghe H, Sabatini S, Roberts S, et al. 2003. *MNRAS* 343:1081–1094
- Baes M, Dejonghe H. 2001. *MNRAS* 326:733–744
- Baes M, Dejonghe H, Davies JI. 2005. In *The Spectral Energy Distributions of Gas-Rich Galaxies: Confronting Models with Data*, eds. CC Popescu, RJ Tuffs, vol. 761 of *American Institute of Physics Conference Series*
- Baes M, Stamatellos D, Davies JI, Whitworth AP, Sabatini S, et al. 2005. *New A* 10:523–533
- Baes M, Verstappen J, De Looze I, Fritz J, Saftly W, et al. 2011. *ApJS* 196:22
- Beech M. 1987. *Ap&SS* 133:193–195
- Bethell TJ, Chepurnov A, Lazarian A, Kim J. 2007. *ApJ* 663:1055–1068
- Bethell TJ, Zweibel EG, Heitsch F, Mathis JS. 2004. *ApJ* 610:801–812
- Bianchi S. 2007. *A&A* 471:765–773
- Bianchi S. 2008. *A&A* 490:461–475
- Bianchi S, Ferrara A, Giovanardi C. 1996. *ApJ* 465:127
- Bjorkman JE, Wood K. 2001. *ApJ* 554:615–623
- Bohren CF, Huffman DR. 1983. *Absorption and scattering of light by small particles*. Wiley
- Boulanger F, Perault M. 1988. *ApJ* 330:964–985
- Braun R, Thilker DA, Waltherbos RAM, Corbelli E. 2009. *ApJ* 695:937–953
- Bruzual GA, Magris G, Calvet N. 1988. *ApJ* 333:673–688
- Burgh EB, McCandliss SR, Feldman PD. 2002. *ApJ* 575:240–249
- Calzetti D, Bohlin RC, Gordon KD, Witt AN, Bianchi L. 1995. *ApJ* 446:L97
- Calzetti D, Kinney AL, Storchi-Bergmann T. 1994. *ApJ* 429:582–601
- Cash JR, Karp AH. 1990. *ACM Trans. Math. Software* 16:201–
- Cashwell E, Everett C. 1959. *A practical manual on the Monte Carlo method for random walk problems*. Pergamon Press
- Castellanos P, Casassus S, Dickinson C, Vidal M, Paladini R, et al. 2011. *MNRAS* 411:1137–1150
- Chakrabarti S, Cox TJ, Hernquist L, Hopkins PF, Robertson B, Di Matteo T. 2007. *ApJ* 658:840–850
- Chakrabarti S, Whitney BA. 2009. *ApJ* 690:1432–1451
- Chandrasekhar S. 1960. *Radiative transfer*. Dover Publications Inc.

- Code AD. 1973. In *Interstellar Dust and Related Topics*, ed. J. M. Greenberg & H. C. van de Hulst, vol. 52 of *IAU Symposium*
- Code AD, Whitney BA. 1995. *ApJ* 441:400–407
- Colberg JM, et al. 2008. *MNRAS* 387:933–944
- Compiègne M, Verstraete L, Jones A, Bernard JP, Boulanger F, et al. 2011. *A&A* 525:A103
- Cox NLJ, Kerschbaum F, van Marle AJ, Decin L, Ladjal D, et al. 2012. *A&A* 537:A35
- De Geyter G, Baes M, Fritz J, Camps P. 2012. *ArXiv e-prints*
- De Looze I, Baes M, Bendo GJ, Ciesla L, Cortese L, et al. 2012a. *MNRAS* 427:2797–2811
- De Looze I, Baes M, Fritz J, Verstappen J. 2012b. *MNRAS* 419:895–903
- Desert FX, Boulanger F, Shore SN. 1986. *A&A* 160:295–300
- Devroye L. 1986. *Non-uniform random variate generation*. Springer-Verlag
- Doty SD, Metzler RA, Palotti ML. 2005. *MNRAS* 362:737–747
- Doty SD, Tidman R, Shirley Y, Jackson A. 2010. *MNRAS* 406:1190–1200
- Draine BT. 1988. *ApJ* 333:848–872
- Draine BT. 2003a. *ARA&A* 41:241–289
- Draine BT. 2003b. *ApJ* 598:1017–1025
- Draine BT, Li A. 2001. *ApJ* 551:807–824
- Draine BT, et al. 2007. *ApJ* 663:866–894
- Dullemond CP, Turolla R. 2000. *A&A* 360:1187–1202
- Dunne L, Gomez HL, da Cunha E, Charlot S, Dye S, et al. 2011. *MNRAS* 417:1510–1533
- Dupree S, Fraley S. 2002. *A Monte Carlo Primer: A Practical Approach to Radiation Transport*. Kluwer Academic/Plenum
- Dwek E. 1986. *ApJ* 302:363–370
- Dwek E, Galliano F, Jones AP. 2007. *ApJ* 662:927–939
- Ercolano B, Barlow MJ, Storey PJ. 2005. *MNRAS* 362:1038–1046
- Falgarone E, Phillips TG, Walker CK. 1991. *ApJ* 378:186–201
- Fischer O, Henning T, Yorke HW. 1994. *A&A* 284:187–209
- Frenk CS, et al. 1999. *ApJ* 525:554–582
- Fritz J, Gentile G, Smith MWL, Gear WK, Braun R, et al. 2012. *A&A* 546:A34
- Gonçalves J, Galli D, Walmsley M. 2004. *A&A* 415:617–625
- Gonzalez JF, Pinte C, Maddison ST, Ménard F, Fouchet L. 2012. *ArXiv e-prints*
- Goosmann RW, Gaskell CM. 2007. *A&A* 465:129–145
- Gordon KD. 2004. In *Astrophysics of Dust*, eds. AN Witt, GC Clayton, BT Draine, vol. 309 of *Astronomical Society of the Pacific Conference Series*
- Gordon KD, Bailin J, Engelbracht CW, Rieke GH, Misselt KA, et al. 2006. *ApJ* 638:L87–L92
- Gordon KD, Misselt KA, Witt AN, Clayton GC. 2001. *ApJ* 551:269–276
- Gordon KD, Witt AN, Carruthers GR, Christensen SA, Dohne BC. 1994. *ApJ* 432:641–647
- Gordon KD, Witt AN, Friedmann BC. 1998. *ApJ* 498:522
- Górski KM, Hivon E, Banday AJ, Wandelt BD, Hansen FK, et al. 2005. *ApJ* 622:759–771
- Groenewegen MAT, Waelkens C, Barlow MJ, Kerschbaum F, Garcia-Lario P, et al. 2011. *A&A* 526:A162
- Guhathakurta P, Draine BT. 1989. *ApJ* 345:230–244
- Haas M, Müller SAH, Chini R, Meisenheimer K, Klaas U, et al. 2000. *A&A* 354:453–466
- Harries TJ. 2000. *MNRAS* 315:722–734
- Harries TJ. 2011. *MNRAS* 416:1500–1508
- Harries TJ, Monnier JD, Symington NH, Kurosawa R. 2004. *MNRAS* 350:565–574

- Harvey PM, Henning T, Liu Y, Ménard F, Pinte C, et al. 2012. *ApJ* 755:67
- Hayward CC, Kereš D, Jonsson P, Narayanan D, Cox TJ, Hernquist L. 2011. *ApJ* 743:159
- Heinemann T, Dobler W, Nordlund Å, Brandenburg A. 2006. *A&A* 448:731–737
- Helou G, Lu NY, Werner MW, Malhotra S, Silbermann N. 2000. *ApJ* 532:L21–L24
- Henning T. 2010. *ARA&A* 48:21–46
- Henning T, Grün E, Steinacker J, eds. 2009. *Cosmic Dust - Near and Far*, vol. 414 of *Astronomical Society of the Pacific Conference Series*
- Heney LG, Greenstein JL. 1941. *ApJ* 93:70–83
- Heymann F, Siebenmorgen R. 2012. *ApJ* 751:27
- Hong SS. 1985. *A&A* 146:67–75
- Hoppe P, Leitner J, Gröner E, Marhas KK, Meyer BS, Amari S. 2010. *ApJ* 719:1370–1384
- Iliev IT, Ciardi B, Alvarez MA, Maselli A, Ferrara A, et al. 2006. *MNRAS* 371:1057–1086
- Iliev IT, Whalen D, Mellema G, Ahn K, Baek S, et al. 2009. *MNRAS* 400:1283–1316
- Indebetouw R, Whitney BA, Johnson KE, Wood K. 2006. *ApJ* 636:362–380
- Jonsson P. 2006. *MNRAS* 372:2–20
- Jonsson P, Groves BA, Cox TJ. 2010. *MNRAS* 403:17–44
- Jonsson P, Primack JR. 2010. *New A* 15:509–514
- Juvela M. 2005. *A&A* 440:531–546
- Juvela M, Padoan P. 2003. *A&A* 397:201–212
- Kalos M, Whitlock P. 2009. *Monte Carlo Methods*. John Wiley & Sons
- Kattawar GW. 1975. *J. Quant. Spec. Radiat. Transf.* 15:839–849
- Keller LD, Sloan GC, Forrest WJ, Ayala S, D’Alessio P, et al. 2008. *ApJ* 684:411–429
- Knebe A, et al. 2011. *MNRAS* 415:2293–2318
- Krügel E. 2007. *An introduction to the physics of interstellar dust*. Taylor & Francis
- Kuiper R, Klahr H, Beuther H, Henning T. 2012. *A&A* 537:A122
- Küppers M, Bertini I, Fornasier S, Gutierrez PJ, Hviid SF, et al. 2005. *Nature* 437:987–990
- Kurosawa R, Hillier DJ. 2001. *A&A* 379:336–346
- Law KH, Gordon KD, Misselt KA. 2011. *ApJ* 738:124
- Lazarian A. 2007. *J. Quant. Spec. Radiat. Transf.* 106:225–256
- Lucy LB. 1999. *A&A* 344:282–288
- Lucy LB. 2005. *A&A* 429:19–30
- Lunttila T, Juvela M. 2012. *A&A* 544:A52
- Martel H, Urban A, Evans II NJ. 2012. *ApJ* 757:59
- Martin PG. 1978. *Cosmic dust. Its impact on astronomy*. Clarendon Press
- Matthews LD, Wood K. 2001. *ApJ* 548:150–171
- Mattila K. 1970. *A&A* 9:53–63
- Mie G. 1908. *Annalen der Physik* 330:377–445
- Min M, Dullemond CP, Dominik C, de Koter A, Hovenier JW. 2009. *A&A* 497:155–166
- Min M, Dullemond CP, Kama M, Dominik C. 2011. *Icarus* 212:416–426
- Min M, Hovenier JW, de Koter A. 2005. *A&A* 432:909–920
- Mishchenko MI, Travis LD, Lacis AA. 2002. *Scattering, absorption, and emission of light by small particles*. Cambridge University Press
- Misselt KA, Gordon KD, Clayton GC, Wolff MJ. 2001. *ApJ* 551:277–293
- Murakawa K, Preibisch T, Kraus S, Weigelt G. 2008. *A&A* 490:673–684
- Natta A, Panagia N. 1984. *ApJ* 287:228–237
- Niccolini G, Alcolea J. 2006. *A&A* 456:1–12
- Niccolini G, Woitke P, Lopez B. 2003. *A&A* 399:703–716

- Nielbock M, Launhardt R, Steinacker J, Stutz AM, Ba-log Z, et al. 2012. *ArXiv e-prints*
- Ohnaka K, Driebe T, Hofmann KH, Leinert C, Morel S, et al. 2006. *A&A* 445:1015–1029
- O’Shea BW, Nagamine K, Springel V, Hernquist L, Norman ML. 2005. *ApJS* 160:1–27
- Paardekooper J, Kruip CJH, Icke V. 2010. *AAP* 515:A79
- Pascucci I, Henning T, Steinacker J, Wolf S. 2003. *Ap&SS* 286:113–118
- Pascucci I, Wolf S, Steinacker J, Dullemond CP, Henning T, et al. 2004. *A&A* 417:793–805
- Patrikeev I, Fletcher A, Stepanov R, Beck R, Berkhuijsen EM, et al. 2006. *A&A* 458:441–452
- Pelkonen VM, Juvela M, Padoan P. 2009. *A&A* 502:833–844
- Petkova M, Springel V. 2011. *MNRAS* 412:935–946
- Pinte C, Harries TJ, Min M, Watson AM, Dullemond CP, et al. 2009. *A&A* 498:967–980
- Pinte C, Ménard F, Duchêne G, Bastien P. 2006. *A&A* 459:797–804
- Press WH, Teukolsky SA, Vetterling WT, Flannery BP. 2002. *Numerical recipes in C++ : the art of scientific computing*. Cambridge University Press
- Purcell EM, Pennypacker CR. 1973. *ApJ* 186:705–714
- Rahimi A, Kawata D. 2012. *MNRAS* 422:2609–2619
- Rho J, Reach WT, Tappe A, Hwang U, Slavin JD, et al. 2009. *ApJ* 700:579–596
- Rijkhorst EJ, Plewa T, Dubey A, Mellema G. 2006. *A&A* 452:907–920
- Roark T, Roark B, Collins II GW. 1974. *ApJ* 190:67–72
- Robitaille TP. 2010. *A&A* 520:A70
- Robitaille TP. 2011. *A&A* 536:A79
- Robitaille TP, Whitney BA, Indebetouw R, Wood K. 2007. *ApJS* 169:328–352
- Rybicki GB, Lightman AP. 1979. *Radiative processes in astrophysics*. Wiley
- Saftly W, Camps P, Baes M, Gordon, K. D.Vandewoude S, et al. 2012. *A&A* submitted
- Schartmann M, Meisenheimer K, Camenzind M, Wolf S, Tristram KRW, Henning T. 2008. *A&A* 482:67–80
- Schechtman-Rook A, Bershady MA, Wood K. 2012. *ApJ* 746:70
- Sellgren K. 1984. *ApJ* 277:623–633
- Siebenmorgen R, Kruegel E, Mathis JS. 1992. *A&A* 266:501–512
- Smith JDT, et al. 2007. *ApJ* 656:770–791
- Smith TL, Witt AN. 2002. *ApJ* 565:304–318
- Springel V. 2010. *MNRAS* 401:791–851
- Stalevski M, Fritz J, Baes M, Nakos T, Popović LČ. 2012. *MNRAS* 420:2756–2772
- Stamatellos D, Whitworth AP. 2005. *A&A* 439:153–158
- Stamatellos D, Whitworth AP, André P, Ward-Thompson D. 2004. *A&A* 420:1009–1023
- Stamatellos D, Whitworth AP, Boyd DFA, Goodwin SP. 2005. *A&A* 439:159–169
- Steinacker J, Bacmann A, Henning T. 2002. *JQSRT* 75:765–786
- Steinacker J, Bacmann A, Henning T. 2006. *ApJ* 645:920–927
- Steinacker J, Bacmann A, Henning T, Klessen R, Stickel M. 2005. *A&A* 434:167–180
- Steinacker J, Chini R, Nielbock M, Nürnberger D, Hoffmeister V, et al. 2006. *A&A* 456:1013–1026
- Steinacker J, Hackert R, Steinacker A, Bacmann A. 2002. *J. Quant. Spec. Radiat. Transf.* 73:557–569
- Steinacker J, Henning T, Bacmann A, Semenov D. 2003. *A&A* 401:405–418
- Steinacker J, Lang B, Burkert A, Bacmann A, Henning T. 2004. *ApJ* 615:L157–L160
- Steinacker J, Michel B, Bacmann A. 2002. *J. Quant. Spec. Radiat. Transf.* 74:183–193

- Steinacker J, Pagani L, Bacmann A, Guieu S. 2010. *A&A* 511:A9
- Steinacker J, Thamm E, Maier U. 1996. *J. Quant. Spec. Radiat. Transf.* 56:97–107
- Stenholm LG, Stoerzer H, Wehrse R. 1991. *JQSRT* 45:47–56
- Thilker DA, Hoopes CG, Bianchi L, Boissier S, Rich RM, et al. 2005. *ApJ* 619:L67–L70
- van der Vorst HA. 1992. *SIAM JSSC* 13:631–644
- van Zadelhoff GJ, Dullemond CP, van der Tak FFS, Yates JA, Doty SD, et al. 2002. *A&A* 395:373–384
- Vijh UP, Witt AN, Gordon KD. 2004. *ApJ* 606:L65–L68
- Watson AM, Henney WJ. 2001. *Rev. Mexicana Astron. Astrofis.* 37:221–236
- Watson DM, Leisenring JM, Furlan E, Bohac CJ, Sargent B, et al. 2009. *ApJS* 180:84–101
- Whitney BA. 2011. *Bulletin of the Astronomical Society of India* 39:101–127
- Whitney BA, Wolff MJ. 2002. *ApJ* 574:205–231
- Witt AN. 1977. *ApJS* 35:1–6
- Witt AN, Boroson TA. 1990. *ApJ* 355:182–189
- Witt AN, Gordon KD. 1996. *ApJ* 463:681
- Witt AN, Gordon KD. 2000. *ApJ* 528:799–816
- Witt AN, Mandel S, Sell PH, Dixon T, Vijh UP. 2008. *ApJ* 679:497–511
- Witt AN, Petersohn JK, Bohlin RC, O’Connell RW, Roberts MS, et al. 1992. *ApJ* 395:L5–L8
- Witt AN, Stephens TC. 1974. *AJ* 79:948
- Witt AN, Thronson Jr. HA, Capuano Jr. JM. 1992. *ApJ* 393:611–630
- Wolf S. 2003a. *ApJ* 582:859–868
- Wolf S. 2003b. *Computer Physics Communications* 150:99–115
- Wolf S. 2008. *Ap&SS* 313:109–112
- Wolf S, Fischer O, Pfau W. 1998. *A&A* 340:103–116
- Wolf S, Henning T. 2000. *Computer Physics Communications* 132:166–188
- Wood K, Reynolds RJ. 1999. *ApJ* 525:799–807
- Wood K, Whitney BA, Robitaille T, Draine BT. 2008. *ApJ* 688:1118–1123
- Xilouris EM, Byun YI, Kylafis ND, Paleologou EV, Papamastorakis J. 1999. *A&A* 344:868–878
- Yusef-Zadeh F, Morris M, White RL. 1984. *ApJ* 278:186–194
- Zhukovska S, Gail HP, Tieloff M. 2008. *A&A* 479:453–480




Fig. 5.— Left: Spitzer image of the molecular cloud L183 at $3.6 \mu\text{m}$ revealing "coreshine" which is scattered light from the densest part of the cloud. Right: Spatial modeling based on basis functions which have Gaussian density structure in all three coordinates (Steinacker et al. 2010). The ellipses give the FWHM of the various Gaussians in the plane of sky.

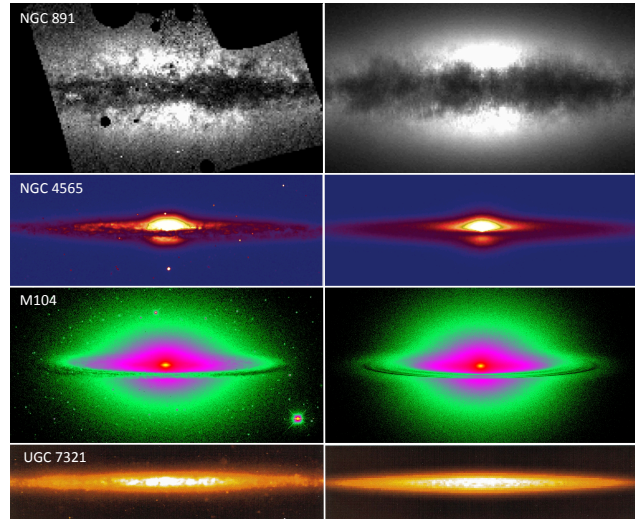


Fig. 6.— Examples illustrating the current state of the art in fitting dust RT models to images of galaxies. The left panels show the observed images, the corresponding panels on the right are the fits to these images. From top to bottom: (a) a clumpy 3D spiral galaxy model fit to an HST B-band image of the prototypical edge-on spiral galaxy NGC 891 by (Schechtman-Rook, Bershady & Wood 2012); (b) a 2D disc galaxy model fit to an SDSS g band image of NGC 4565 by De Looze et al. (2012a) using a fully automatic fitting based on genetic algorithms (De Geyster et al. 2012); (c) a detailed 2D model for the Sombrero Galaxy, fit to a V-band image (De Looze et al. 2012b); (d) a 3D clumpy disc model fit to an R-band image of the edge-on low surface brightness galaxy UGC 7321 (Matthews & Wood 2001).




Fig. 7.— An example of the [Pinte et al. \(2009\)](#) 2D dust RT benchmark, for a flared circumstellar disc with a V-band optical depth in the midplane of $\tau = 10^6$. The image is a scattered light image, the three panels on the right show brightness profiles along the cuts plotted in the left panel, and the differences among the four different models used in the benchmark. The different line styles and colors indicate if the model was Pinball ([Watson & Henney 2001](#)), MCFROST ([Pinte et al. 2006](#)), MCMMax ([Min et al. 2009](#)), or TORUS ([Harries et al. 2004](#)).