

GPU-BASED MONTE CARLO DUST RADIATIVE TRANSFER SCHEME APPLIED TO ACTIVE GALACTIC NUCLEI

FRANK HEYMANN^{1,2,3} AND RALF SIEBENMORGEN²

¹ Department of Physics and Astronomy, University of Kentucky, Lexington, KY 40506-0055, USA; fheyman@pa.uky.edu

² European Southern Observatory (ESO), Karl-Schwarzschild-Strasse 2, 85748 Garching, Germany

³ Astronomisches Institut, Ruhr-University, Bochum, Germany

Received 2011 December 12; accepted 2012 March 11; published 2012 May 1

ABSTRACT

A three-dimensional parallel Monte Carlo (MC) dust radiative transfer code is presented. To overcome the huge computing-time requirements of MC treatments, the computational power of vectorized hardware is used, utilizing either multi-core computer power or graphics processing units. The approach is a self-consistent way to solve the radiative transfer equation in arbitrary dust configurations. The code calculates the equilibrium temperatures of two populations of large grains and stochastic heated polycyclic aromatic hydrocarbons. Anisotropic scattering is treated applying the Heny–Greenstein phase function. The spectral energy distribution (SED) of the object is derived at low spatial resolution by a photon counting procedure and at high spatial resolution by a vectorized ray tracer. The latter allows computation of high signal-to-noise images of the objects at any frequencies and arbitrary viewing angles. We test the robustness of our approach against other radiative transfer codes. The SED and dust temperatures of one- and two-dimensional benchmarks are reproduced at high precision. The parallelization capability of various MC algorithms is analyzed and included in our treatment. We utilize the Lucy algorithm for the optical thin case where the Poisson noise is high, the iteration-free Bjorkman & Wood method to reduce the calculation time, and the Fleck & Canfield diffusion approximation for extreme optical thick cells. The code is applied to model the appearance of active galactic nuclei (AGNs) at optical and infrared wavelengths. The AGN torus is clumpy and includes fluffy composite grains of various sizes made up of silicates and carbon. The dependence of the SED on the number of clumps in the torus and the viewing angle is studied. The appearance of the 10 μm silicate features in absorption or emission is discussed. The SED of the radio-loud quasar 3C 249.1 is fit by the AGN model and a cirrus component to account for the far-infrared emission.

Key words: dust, extinction – galaxies: active – infrared: general – methods: numerical – quasars: individual (3C249.1) – radiative transfer

Online-only material: color figures

1. INTRODUCTION

Dust obscured objects cannot be studied directly in the UV/optical, since the dust shields most of the visible light. To derive the UV/optical component or to constrain the morphological structure from available near-/far-infrared data, a detailed model of the interaction of photons with the dust is required. This necessitates a solution to the radiative transfer (RT) equation. Analytically, this can only be done in some simple configurations, for example by assuming spherical or disk symmetry in which either scattering or absorption is neglected and the wavelength dependency of the dust cross section is strongly simplified, e.g., by a graybody approximation. Nature, however, is usually not well approximated by such assumptions and numerical modeling is the only way to solve this problem.

Dust is detected in the majority of active galactic nuclei (AGNs; Haas et al. 2008). According to the unified scheme (Antonucci & Miller 1985), AGNs are surrounded by a dust obscuring torus. This torus, as argued by Krolik & Begelman (1988), needs a clumpy structure to allow the survival of dust grains in regions where the gas temperatures ($\sim 10^6$ K) are extreme compared to the dust sublimation temperature. Indeed, Tristram et al. (2007) show with VLTI observations of the Circinus AGNs strong evidence for a clumpy or filamentary structure of the nucleus.

In this paper a numerical method is presented, which solves the RT equation in a three-dimensional (3D) geometry, based

on the Monte Carlo (MC) technique (Witt 1977; Section 2). To reduce the required computational effort, different optimization strategies are developed (Lucy 1999; Bjorkman & Wood 2001; Gordon et al. 2001; Misselt et al. 2001; Wolf 2003; Baes 2008; Bianchi 2008; Baes et al. 2011; Section 3). Our numerical solution of the RT equation takes advantage of some of these optimization algorithms and is specifically developed to be vectorized and to run on graphics processing units (GPUs). They are introduced into the original code developed by Krügel (2008). The MC routine handles arbitrary dust distributions in a 3D Cartesian model space at various optical depths. The self-consistent solution provides the dust temperatures; spectral energy distributions (SEDs) and images are computed using a ray tracer (Section 4). The code is tested against existing benchmark results (Section 5). The influence of clumpiness of an AGN dust torus on the SED and the 10 μm silicate feature is discussed and a model of the radio-loud quasar 3C249.1 is presented (Section 6).

2. MONTE CARLO RADIATIVE TRANSFER

In our MC procedure the bolometric luminosity $L = \int_0^\infty L_\nu d\nu$ of the central engine is divided into $N = mn_{\text{zyk}}$ monochromatic photon packets of equal energy $\xi = L/N$, where m is the total number of frequency bins of the heating source and n_{zyk} counts the number of photon packets per bin. The width of the frequency bin with index j between the

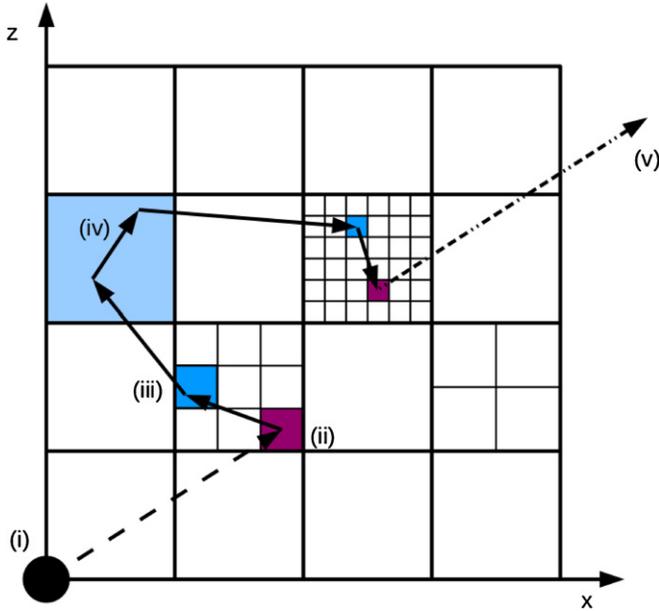


Figure 1. 2D illustration of the three-dimensional grid and the trajectories of the photon packets. The grid is divided into cubes, which can be further divided into sub-cubes. The source emits photons at (i), which interact or not with the dust. By interaction photon packets are either absorbed (ii) or scattered (iii). Multiple interactions may occur in one cell (iv). After absorption the packet will be re-emitted by the dust at a different frequency. Frequency changes are illustrated by different line styles. Finally, (v) the packet escapes the model space.

(A color version of this figure is available in the online journal.)

frequencies ν_j and ν_{j+1} , $\Delta\nu = |\nu_{j+1} - \nu_j|$, is derived from the spectral shape of the radiation source and the energy of a photon packet

$$\xi = \int_{\nu_j}^{\nu_{j+1}} L_\nu d\nu. \quad (1)$$

The primary heating source emits photons, for example, in the optical and UV while the dust emission occurs at IR and submillimeter wavelength. Two frequency grids are considered: the first one is set up by Equation (1) and accounts for the emission of the primary heating source. The second frequency grid is an extended version of the first one and includes additional bins in the IR and submillimeter to scope with the dust emission. The interaction probability between photons and dust particles in the RT problem is computed by the absorption κ_ν^{abs} and scattering κ_ν^{sca} cross sections. The code includes different grain materials and size distributions.

The model space is set up as cubes in an orthogonal Cartesian grid. Each cube is divided into an arbitrary number of sub-cubes of volume V_i and constant density ρ_i (Figure 1), where the index refers to sub-cube i . The subdivision of cubes allows a finer sampling whenever required: for example, close to the dust evaporation zone, in regions of high optical depth or where the spatial gradient of the radiation field is large. One possible trajectory of a photon is illustrated in Figure 1. Photons are emitted by the source at frequency ν_{j_*} , where j_* is the index of the frequency grid of the source. The flight path of the photon throughout the model space is computed.

The direction of the photon is chosen from a uniform distribution of random numbers ζ_1, ζ_2 from the half-open interval $[0, 1)$, so that $\phi = 2\pi\zeta_1$ and $\cos\theta = 1 - 2\zeta_2$. The distance from the entry point of a photon into a cell i to its exit point along the travel direction is ℓ_i . In the MC method

the interaction of photons with the dust in a sub-cube can be determined with a uniform distributed random number ζ (Witt 1977; Lucy 1999) and the optical depth

$$\tau_i(\nu) = (\kappa_\nu^{\text{abs}} + \kappa_\nu^{\text{sca}})\rho_i \ell_i. \quad (2)$$

Photons leave the cell if $\tau_i(\nu) \leq -\log\zeta$ and otherwise they interact with the dust. If a packet passes through a cell without interacting, it enters a neighboring cell according to its direction of travel. At the border of the model space it eventually escapes. When a photon packet enters a cell, a new random number is chosen to determine the interaction probability of photons with the dust.⁴ The photon packet interacts with the dust after it has traveled a distance

$$\ell'_i = \frac{-\log\zeta}{\rho_i(\kappa_\nu^{\text{abs}} + \kappa_\nu^{\text{sca}})}. \quad (3)$$

The travel distance ℓ'_i defines the point in the cell where interaction takes place and photons are either scattered (Figure 1(ii)) or absorbed (Figure 1(iii)). Depending on its optical depth multiple scattering or absorption events may occur in a cell (Figure 1(iv)). In case of an interaction event the probability of scattering is given by the albedo $A_\nu = \kappa_\nu^{\text{sca}} / (\kappa_\nu^{\text{sca}} + \kappa_\nu^{\text{abs}})$ and therefore the chance of absorption is $1 - A_\nu$.

When a photon is scattered on a dust grain the packet keeps its frequency, but changes its travel direction according to the phase function of the particle. We use the Henyey & Greenstein (1941) phase function to approximate the anisotropic scattering (see Section 4). In the isotropic case the asymmetry factor g_ν equals 0. The index of the cell, the frequency of the scattered photon, and the incoming direction are stored. This allows computing of the source function of the cell i , which will be used in a ray tracer developed to compute the observed image at any frequency (Section 4).

When a photon is absorbed a new one with different frequency and direction is emitted from that position. The frequency of the emitted photon packet is given by the temperature of the absorbing material. The emitted photon has the same energy as the absorbed one. All photon packets contain the same amount of energy ξ , therefore the total number of absorbed photon packets per particle and cell is used to calculate the temperature of each material. After k absorptions of photon packets by dust the cell i has a temperature $T_{i,k}$ calculated from

$$\int \kappa_\nu^{\text{abs}} B_\nu(T_{i,k}) d\nu = \frac{k\xi}{4\pi\rho_i V_i}, \quad (4)$$

where $B_\nu(T_{i,k})$ is the Planck function.

Dust emits a photon packet at the shortest frequency ν' calculated with

$$\int_0^{\nu'} \kappa_\nu^{\text{abs}} B_\nu(T_{i,k}) d\nu \geq \zeta \int_0^\infty \kappa_\nu^{\text{abs}} B_\nu(T_{i,k}) d\nu. \quad (5)$$

In the MC method one follows all N photon packets through the dust cloud until they reach the outer boundary of the model and escape. For each cell the absorption events are counted. When all packets escaped the model space, the number of absorptions in a cell is used (Equation (4)) to calculate its dust temperature. An iterative MC method is started

⁴ It can be confirmed that the interaction probability calculated with a new random number is in agreement with the treatment of a travel distance which is determined by a single random number (Lucy 1999).

with an initial guess of the dust temperatures. This allows computing of the frequencies of packets emitted by the dust (Equation (5)). After all packets escape the model space another N photons are emitted by the source and their flight paths through the model are computed. However, this time using the dust temperatures of the previous run. This procedure continues until the dust temperatures converges, which usually takes about three to five iterations.

2.1. Non-equilibrium Radiative Transfer (PAH)

The enthalpy state of very small grains such as polycyclic aromatic hydrocarbons (PAHs; Tielens 2008) fluctuates strongly when they are hit by photons. In order to solve the RT problem including PAH with MC we treat their emission as a stochastic process. A temperature distribution function $P(T)$ of the PAH is computed using the iterative scheme by Siebenmorgen & Krügel (1992). PAHs are implemented into the MC code in two steps: first, we compute the PAH absorption of photon packages and store position and frequency of these events. Then, after calculating the equilibrium temperature of the large grains (Section 2), the PAH emission is computed. Each cell with PAH absorption becomes a source and emits the same number of photon packages as previously absorbed by the PAH. These photons are traced through the model as described in Section 2. The frequency of the emitted package is calculated according to

$$\int_0^{\nu'} \epsilon_n(\nu) d\nu \geq \xi \int \epsilon_n(\nu) d\nu, \quad (6)$$

with random number ξ and PAH emissivity

$$\epsilon_n(\nu) = K_v^{\text{PAH}} \int B_\nu(T) P_n(T) dT, \quad (7)$$

where K_v^{PAH} is the PAH cross-section per unit mass of dust and $P_n(T)$ is the temperature distribution function as described in Siebenmorgen & Krügel (2010).

Similar to the evaporation of large grains it is important to compute the location where PAHs are able to survive the radiation environment. We treat the photodissociation of PAH as discussed in Siebenmorgen & Krügel (2010). They find that PAH dissociate when, within a cooling time, the total absorbed energy E_{PAH} is larger than a critical energy E_{crit} :

$$E_{\text{PAH}} \geq E_{\text{crit}} = \frac{N_C}{2} \text{ (eV)}, \quad (8)$$

where N_C is the number of carbon atoms per PAH molecule. In cells where the molecules dissociate the absorbed photon packages are treated as if they are emitted from the star. Further details on our and other numerical implementations of PAH in the MC code, a verification against benchmark models and an application to the PAH emission from protoplanetary disks is given in an accompanying paper (Siebenmorgen & Heymann 2012a).

3. OPTIMIZED MONTE CARLO RADIATIVE TRANSFER

In this section we describe various optimization techniques which are introduced into our MC scheme to increase the accuracy and decrease computing-times.

3.1. Small Optical Depth

The MC approach solves the RT problem in a probabilistic way. Dust temperatures become uncertain and have large errors whenever the interaction probability is low. This is the case in cells of low optical depth where many photons flyby without being absorbed or scattered by the dust. The statistical noise in such cells can be reduced significantly applying a procedure developed by Lucy (1999). Contrary to the method described in Section 2, in which only absorption events contribute to the temperature calculation of the cell, the Lucy method considers the absorbed energy of every photon packet traveling through the cell. This technique increases the temperature accuracy of the MC treatment in low optical depth regions. It has no advantage when the optical depth $\tau_i \geq 1$ as interactions between photons and dust are likely.

We follow Section 3.4 of Lucy (1999) and implement a similar optimization technique. It is applied whenever the maximal optical depth of a cell i is smaller than 10^{-2} . The performance of the MC scheme depends critically on this value. It is chosen so that the performance of the code is high while still accounting for reasonable accuracy. For smaller values the optimization is only applied for extreme optical thin cells whereas for higher values the expensive accuracy enhancement is applied to cells which can be accurately computed without the Lucy scheme. With our choice the energy absorbed by the dust E_i^{abs} is a small fraction of the energy of a photon packet. The energy is computed by

$$E_i^{\text{abs}} = (1 - e^{-\tau_i}) \xi, \quad (9)$$

where τ_i is the optical depth from the entry to the exit point of the photon of cell i . Therefore, the frequency of the photon packet remains unchanged.

3.2. High Optical Depth

The MC solution of the RT equation becomes slow for very optical thick regions, where the number of photon interactions with the dust increases exponentially. To avoid photons becoming trapped in cells with very high optical depth we apply a modified random walk procedure. It is based on a diffusion approximation by Fleck & Canfield (1984). The method is tested by Min et al. (2009) and numerical enhancements are given by Robitaille (2010).

3.3. Iteration-free Monte Carlo

An iteration-free MC scheme is developed by Bjorkman & Wood (2001, hereafter B&W); see Baes et al. (2005) and Krügel (2008) for helpful comments. In the B&W method the dust temperature in a cell is adjusted with respect to the previous absorption event. Contrary to Equation (5), dust re-emission occurs at the shortest frequency ν' for which

$$\int_0^{\nu'} \kappa_\nu^{\text{abs}} \frac{dB_\nu(T_{i,k})}{dT} d\nu \geq \zeta \int_0^\infty \kappa_\nu^{\text{abs}} \frac{dB_\nu(T_{i,k})}{dT} d\nu \quad (10)$$

is valid. Both methods described in Equation (5) and Equation (10) are implemented in the MC code and can be used upon preference.

3.4. Parallelization on Graphics Processing Units

The MC method of solving the RT problem is particularly suited for parallelization because all photon packets are independent of each other (Jonsson 2006; Heymann 2010;

Siebenmorgen et al. 2011). Vectorization of the code allows a number of photon packets to be emitted simultaneously. The number of packets launched at a time depends on the number of threads⁵ available on the computer. The trajectories of the photons can then be calculated in parallel.

In parallel environments an additional requirement is placed on the random number generator (RNG); each thread must have an independent sequence of random numbers. We solve this problem by applying a parallel version of the Mersenne Twister algorithm as given by Matsumoto & Nishimura (2000).

Parallelization is most efficient when all processing units finish their task at about the same time. Then vectorization speeds up the code roughly equal to the number of threads available. Unfortunately, this is not always the case. The number of interactions of photons increases exponentially with the optical depth τ_V and photon packets may get trapped in cells with say $\tau_V \geq 1000$. The workload of these cells is much higher than for cells at much lower optical depth. This results in a rather unbalanced workload over all threads so that the advantage of vectorization may disappear. In our application idle threads are avoided as much as possible. Every time a thread finishes the trajectory of the photon packet within the model space a counter is increased. If the counter reaches 80% of the total number of parallel working threads and if the number of interactions of the photon packet within a cell is larger than 100, the thread pauses. In this case the position and frequency of the photon packet is stored. Close to the end of the simulation when all photons with average processing speed have escaped the model space, the stored photons are resumed. This procedure allows a balanced workload among all threads. In addition, the modified random walk procedure of Robitaille (2010) is implemented (Section 3.2.)

This code utilizes two different parallelization methods. It can be used on shared memory machines using the openMP library to run as many parallel rays as there are processor cores available and on vectorized hardware (GPU) highly optimized for parallelization. The complete MC RT solution is ported to GPU using the Compute Unified Device Architecture (CUDA) developed by NVIDIA (CUDA 2011). This speeds up the entire MC solution in comparison to the method by Jonsson & Primack (2010), which provides only a GPU acceleration when computing the temperature of a grain.

3.5. Numerical Recipes

For vectorized MC RT codes it is recommended to follow some general recipes. In vectorized systems the increased number of processing units decreases the computing-time for numerical operations whereas the time needed for read-and-write operations remains unaltered. This implies a paradigm shift in the programming strategy. In single processing codes the performance is often improved by reading pre-calculated values from memory. Instead in a vectorized code at some point it is faster to calculate such values case by case. For example in our code, and as long as the dust frequency grid has less than 300 bins, it is faster to solve Equation (5) on the fly than reading pre-computed values from memory.

The performance of parallel codes is increased considerably when designing a particular array structure for the particular hardware in use. A read operation from memory is most efficient when it accesses the entire array. For the GPUs used in this

paper the memory bandwidth is 512 bit and we use a 32 bit machine. Therefore in our code the most efficient way to read from memory is by a modulus of 16×32 bit values. The design of a particular array structure is import for problems where the memory access is not predictable, which is unfortunately the case in MC schemes.

In order to improve the performance it is very efficient to code with as less interference between workers and threads as possible. The code should avoid situations where thread A waits until thread B has finished and then thread B waits until thread A has finished. Unfortunately in MC simulations the dependency between workers and threads is often less obvious than in the previous example.

Another challenge with memory interactions is to use local copies of arrays for each thread as often as possible. This decreases the interaction between different threads and usually leads to large performance improvements, of course at the expense of some additional memory needs. When dealing with memory interactions between threads it is helpful to study the special hardware capabilities. In our case we gain a factor of two in the computing-time by using atomic integer operations rather than floating point operations.

It is advantageous to optimize all read-and-write operations to the fastest available source. This can be done by caching small amounts of data from the disk to the global memory, than to the GPU memory and finally to the GPU shared memory. For example in our code we store the wavelength-dependent cross sections in the GPU shared memory. This is about 10 times faster than reading them from the global memory.

For MC simulations the RNG is important and special care must be taken on parallel systems. A detailed study of this topic is beyond the scope of this paper. We tested different RNG available and find that the performance and accuracy as well as thread dependence differs significantly between various RNG implementations. It is therefore important to choose and test an RNG which fits the particular problem best.

3.6. Parallelization and Optimization Algorithms

The run time requirements of the different MC methods are compared. We consider a star at temperature of 2500 K with solar luminosity which heats a spherical and constant density dust envelope with an inner radius of 0.7 AU and an outer radius of 700 AU. The optical depth measured from the star to the outer boundary is $\tau_V = 10$. Parameters of the dust are specified as in Section 5.1. In the models 10^7 photon packets are emitted from the star and 10^6 grid cells are used.

We apply the Lucy (Section 3.1), B&W (Section 3.3.) method in a scalar (single thread), a central processing unit (CPU) version with 8 threads, a GPU version with 256 threads, and in addition the iterative MC scheme of Section 2 on a GPU machine. Initializing times of the various MC methods are identical.

As described in Section 3.1 for the same number of emitted photon packets the Lucy method provides a better temperature estimate than the other methods. In the Lucy method the fraction of a photon packet is considered. This can only be realized by considering floating point operations which are more computer expensive than integer operations. In our test case the scalar version of the Lucy method is run with three iterations and requires a total run time of three hours (Table 1).

For single thread applications the iteration-free MC scheme by B&W (Section 3.3) has the advantage that it speeds up the process by a factor which equals the number of iterations

⁵ A thread is the smallest unit of processing that can be scheduled in a parallel environment.

Table 1
Run Time Requirements of Different MC Methods

Method	Threads	Time (minutes)
Lucy	1	180
B&W	1	60
Lucy	8	45
B&W	8	20
GPU	256	3
GPU+Lucy	256	4
GPU+B&W	256	2
GPU+B&W+Lucy	256	2

required for convergence in the other MC treatments. However, in the B&W method memory interaction between different cells are unavoidable and they slow down the parallelization capabilities of this method. The memory interaction between cells rises with the number of threads and therefore the amount of necessary *atomic memory operations*⁶ increases with parallelization. To minimize the number of atomic memory operations we solve the problem in an iterative MC scheme using Equation (4). This allows parallelization with a huge number of threads. At low budget this can be realized using GPU technology and these MC methods are labeled as such in Column 1 of Table 1. The GPU method may be further optimized in combination with the Lucy (GPU+Lucy) or the B&W method (GPU+B&W). On our conventional computer we use a GPU with 256 threads (eight multiprocessor each with 32 cores) clocked at 1.5 GHz. When compared to a single thread CPU application clocked at 3 GHz a speed up factor due to vectorization of 60 is realized. This is below a theoretical expected speed-up factor of 128 because of additional overheads produced by input/output routines and memory transfers in the GPU machine.

Total run times of the various methods with different numbers of threads are given in Table 1. The vectorized Lucy method scales slightly better with the number of available threads than the other algorithms because fewer atomic memory operations are required. The speed increase by vectorization as compared to scalar MC treatments is proportional to the number of available threads. The GPU method with B&W optimization is a factor of 90 faster than the original Lucy method.

4. RAY TRACER TO COMPUTE SEDs AND IMAGES

Photons which eventually escape the MC model space in a particular solid angle can be counted and converted into a flux density. In principle this method allows computing images of the object. Unfortunately to obtain a moderate signal-to-noise ratio the number of photon counts per solid angle must be high, which is difficult to reach within reasonable computing-times. Here we present a ray tracing method which allows computing of noise-free images, SEDs, and visibilities at high spatial resolution. The ray tracer uses the temperature and scattering events of the cells calculated with the MC code (Section 2). The uncertainty in the derived images is therefore based on the precision of the MC computation.

In the algorithm, rays are traced through the model space from an observer plane of arbitrary orientation. The ray tracer together with the MC method allows us to calculate the flux

received on each pixel of an image in the plane of the observer. The observed image is located at distance D from the object which is at the center of the cube. The orientation of the image plane is defined by its surface normal \mathbf{e}'_z . The axis of the image plane $\mathbf{e}'_x, \mathbf{e}'_y$ is perpendicular to \mathbf{e}'_z . The coordinates (x, y, z) of the 3D model space are transformed by a parallel projection into the coordinates (x', y') of the two-dimensional (2D) image in the observers plane. The image consists of $n_x \times n_y$ pixels, with a pixel size chosen so that the complete model fits in the projection. The center of the image \mathbf{i}_0 is located at position $[x_0, y_0, z_0]$ and has image coordinates $[x'_0, y'_0]$. The projected coordinates $[x', y']$ of the image are transformed into MC cube coordinates by

$$\mathbf{r}(x', y') = \mathbf{i}_0 + (x' - x'_0)\mathbf{e}'_x + (y' - y'_0)\mathbf{e}'_y, \quad (11)$$

where $\mathbf{r} = [x, y, z]$. The ray tracer follows the line of sight from each detector pixel in direction \mathbf{e}'_z through the MC model. Adding the contribution of emission and scattering from each cell i along the line of sight results in the observed intensity. The contribution of cell i to the total intensity is

$$I_{v,i} = (I_{v,i}^e + I_{v,i}^s)e^{-\tau_{v,i}}, \quad (12)$$

where $I_{v,i}^e$ refers to the intensity of the dust emission and $I_{v,i}^s$ to scattering; $\tau_{v,i}$ is the optical depth at frequency ν of cell i . For convenience we drop the index i in the following. The optical depth is

$$\tau_\nu = K_\nu^{\text{abs}} \int_0^\ell \rho(\mathbf{r}) ds, \quad (13)$$

where K_ν^{abs} is the absorption cross section per unit dust mass and ℓ is the path lengths. We call the optical depth measured along the ray from the border of the cell to the observer τ_ν , and the one measured face to face from one side of the cell to the other τ_ν^{cell} . The emission is

$$I_\nu^e = [1 - \exp(-\tau_\nu^{\text{cell}})]B_\nu(T). \quad (14)$$

For scattering we implement either a g -factor approximation following the notation of Krügel (2008) or consider non-isotropic scattering applying the Henyey–Greenstein phase function P_{HG} (Henyey & Greenstein 1941) defined as

$$P_{\text{HG}}(\beta) = \frac{1}{4\pi} \frac{(1 - g^2)}{[1 + g^2 - 2g\cos(\beta)]^{3/2}}, \quad (15)$$

where g is the anisotropy parameter. If $g = 0$ the scattering is isotropic. The scattered light intensity is given by

$$I_\nu^s = \sum_\beta P_{\text{HG}}(\beta) \frac{N(\beta, \nu)^{\text{sca}} \epsilon}{\Delta\nu}, \quad (16)$$

where $\Delta\nu$ is the width of the frequency bin and $N(\beta, \nu)^{\text{sca}}$ is the total number of scattering events with frequency ν and angle β between the observer and the original scattering direction. This method to calculate the scattered intensity is similar to the peel-off technique proposed by Yusef-Zadeh et al. (1984). The difference lies in the time of the calculation of the scattering intensity. The peel-off technique calculates the intensity during the MC run. Whenever a scattering event occurs the scattered intensity of the event is added to the observed images. Therefore, in the peel-off technique the location of the observer has to be known before the MC run. This information is not required

⁶ Atomic operations are operations which are performed without interference from any other threads. Atomic operations are often used to prevent conditions which are common problems in multi-thread applications.

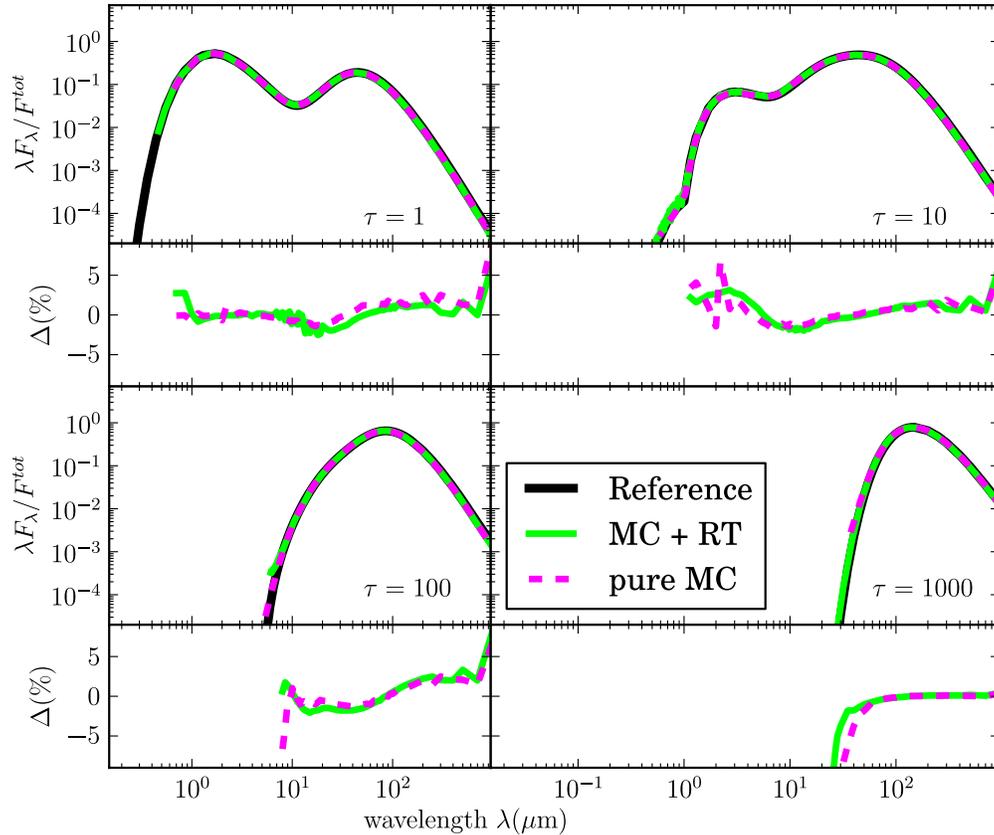


Figure 2. Comparison of SEDs of dust spheres at optical depths between $\tau_v = 1$ and 1000 computed with MC and ray tracer (green) or photon counting (dashed), benchmark results in black (Ivezic et al. 1997). Lower panels give the difference between benchmark and the procedures of this work. (A color version of this figure is available in the online journal.)

in our MC scheme which stores only the scattering events. Then, in post-processing using the ray tracer, the contribution of the scattered light flux can be computed for arbitrary observer orientations. The flux density is calculated summing up the contribution of each cell given by

$$F_{v,i} = \frac{A_{\text{pix}}}{D^2} [I_v^e + I^s(\theta)] \exp(-\tau_v), \quad (17)$$

where τ_v is the optical depth from the border of the cell to the observer and A_{pix} is the area of the detector pixel. The ray tracer calculates emission and scattering images at high signal to noise. It cannot remove the noise which is introduced by the MC simulation where the temperature distribution and scattering events are calculated with a stochastic sampling. However, for a known distribution of the temperature and scattering events the ray tracer provides noise-free images.

5. BENCHMARK

The vectorized MC method is tested against the 1D ray tracing code by Krügel (2008). Further verifications of the MC and comparison with benchmark tests are described below.

5.1. Spherical Symmetry (1D)

Benchmark models in one dimension are provided by Ivezic et al. (1997). They consider a spherical dust envelope which is centrally heated by a 2500 K star with luminosity of $1 L_\odot$. For this test the dust absorption and scattering coefficient for wavelength below $\lambda_0 = 1 \mu\text{m}$ are $q_{\text{abs}} = q_{\text{sca}} = 1$ and are $q_{\text{sca}} = \lambda_0/\lambda$ and $q_{\text{abs}} = (\lambda_0/\lambda)^{-4}$ at longer wavelengths.

Models are computed at optical depths of $\tau_v = 1, 10, 100,$ and 1000 with inner radius $r_{\text{in}} = 3.14, 3.15, 3.20,$ and 3.52 AU, respectively. The dust density is constant and the outer radius is $r_{\text{out}} = 1000 \times r_{\text{in}}$. The models are run on a grid with 10^6 cells and 2×10^8 photon packets per iteration are launched from the star. The number of MC iterations are 3, 4, and 5 for $\tau_v \leq 10, 100,$ and $1000,$ respectively.

MC computed temperatures are compared to the one calculated by DUSTY (Ivezic & Elitzur 1997) and Krügel (2008); they agree within $\lesssim 1\%$. The SEDs are shown in Figure 2 in which each quadrant represents models at different optical depths. In the quadrants, the SEDs are shown in the top and the flux difference between MC and benchmark in the bottom panel. The SEDs of the MC models are either computed by photon counting or using the ray tracer (Section 4). The difference of the SED computed by MC and DUSTY is typically better than a few percent. It becomes larger at faint fluxes where the ray tracing is more accurate than the packet counting procedure because of photon noise. For the $\tau_v = 1000$ model the counting method is starved at fluxes below 0.2% from the peak flux. We encountered numerical inconsistencies in the DUSTY code at flux levels $\lesssim 10^{-8}$ from the peak when compared against the RT code by Krügel (2008).

The differences between the reference codes and our MC program are mainly caused by Poisson noise. The photon noise is large at flux levels which are low compared to the peak flux. The wavelength range where the frequency grid of the source and that of the dust overlap is an additional source of systematic error. The source grid is build so that the energy of the photon packets of each frequency bin is identical. This results in a

Table 2
Parameters of 2D Benchmark by Pascucci et al. (2004)

Symbol	Parameter	Value
M_*	Stellar mass	$1 M_\odot$
R_*	Stellar radius	$1 R_\odot$
T_*	Stellar effective temperature	5800 K
r_{in}	Inner disk radius	1 AU
r_{out}	Outer disk radius	1000 AU
r_d	Half of outer radius	500 AU
z_d	Disk height	125 AU
a	Grain radius	$0.12 \mu\text{m}$
ρ_0	Density for $\tau_V = 1$	$8.45 \times 10^{-22} \text{ g cm}^{-3}$
τ_V	Optical depth at 550 nm	1, 10, 100

larger bin size at the Rayleigh–Jeans tail of the star where, on the other hand, the dust grid has a fine sampling because of the silicate and PAH features. To reduce this sampling problem it is necessary to increase the number of frequency bins of the star. Unfortunately, the total number of photons emitted from the star must be increased to achieve a certain signal to noise for each bin. The number of photons emitted from the star in each frequency bin must be large enough to reduce the Poisson noise to acceptable levels. We use 1000 frequency bins for the star and emit in each bin 10^5 photons. This choice is based on a run time optimization for the achieved accuracy. The third source of systematic differences between the reference codes and our MC scheme is the spatial grid. While the MC code uses a Cartesian grid in three dimensions with one level of refinement to describe the spatial dust distribution, the reference codes, and in particular that of Krügel (2008), use an extremely optimized grid to account for the spherical symmetry. This systematic effect could be reduced by an increase of the number of MC grid cells. Current hardware limits the number of grid cells because of memory and speed constraints.

5.2. Disk Geometry (2D)

Different methods to compute the RT in a 2D dust configuration are compared by Pascucci et al. (2004). They consider a star which heats a dust disk. The inner region $r < r_{in}$ is dust free and the dust density distribution is similar to those described by Chiang & Goldreich (1997, 1999):

$$\rho(r, z) = \rho_0 \left(\frac{r_d}{r} \right) e^{-(0.25\pi(z/h)^2)} \quad (18)$$

with

$$h = z_d \left(\frac{r}{r_d} \right)^{1.125}, \quad (19)$$

where r is the distance from the central star in the midplane of the disk and z is the height above the midplane. The parameter ρ_0 is chosen such that the optical depth at 550 nm along the midplane is 1, 10, and 100, which leads to densities $\rho_0 = 8.45 \times 10^{-22}$, 8.45×10^{-21} , and $8.45 \times 10^{-20} \text{ g cm}^{-3}$. Additional parameters are specified in Table 2. The absorption and scattering cross-sections are that of a $0.12 \mu\text{m}$ silicate grain (optical data⁷ are taken from Draine & Lee 1984).

The SEDs calculated by the various algorithms and treatments used by Pascucci et al. (2004) agree to better than 20%. The RT in the disk is solved by our GPU code and the derived SEDs are compared to an averaged SED of the results given by Pascucci et al. (2004).

⁷ Download from <http://www.mpa.de/PSF/PSFpages/RT/benchmark.html>.

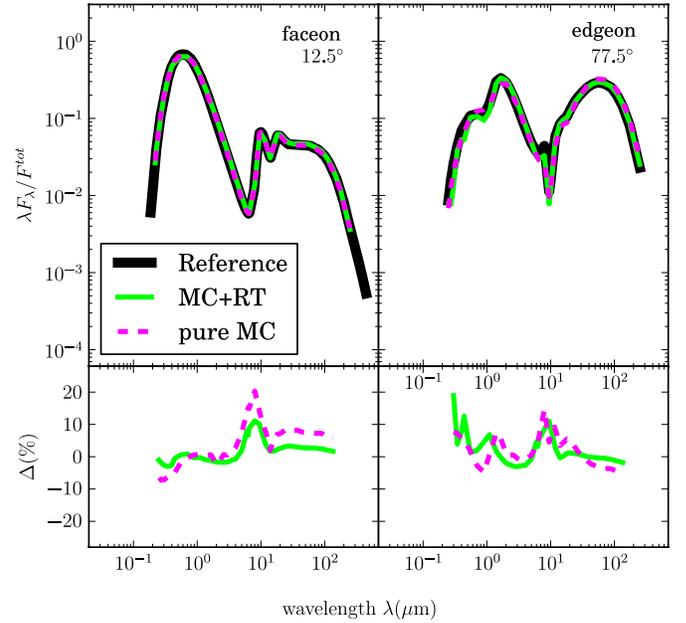


Figure 3. Comparison of SEDs of a dust disk at optical depths along the midplane of $\tau_V = 100$, face-on view (12.5°) is shown left and edge-on view (77.5°) right. SEDs are computed with MC and ray tracer (green) or photon counting (dashed). The mean SED of the benchmark results by Pascucci et al. (2004) is in black. Lower panels give the difference between the mean SED of the benchmark and the procedures of this work.

(A color version of this figure is available in the online journal.)

In the MC program we use 3×10^6 cells and run models with 2×10^8 photon packets per iteration.

For the optical depth along the midplane of the disk with $\tau_V \leq 10$ we use three iterations and found in the SED an overall agreement to the benchmark results to within a few percent. In the $\tau_V = 100$ case four iterations are used and the SED and differences to a mean SED over those computed in the reference models are shown in Figure 3. Two different inclinations 12.5° (face-on) and 77.5° (edge-on) are displayed. The residual between SED computed by us and the mean SED of the benchmarks is typically a few percent and larger deviation is found at fluxes below 1% of the peak flux. This is visible at the border of the wavelength grid and in the $10 \mu\text{m}$ silicate absorption band where variations from one SED to another in the reference codes are also larger.

The differences in the 2D benchmark and our code are twofold. There is the Poisson noise which is present in our as well as in the reference. Note that the reference SED in Pascucci et al. (2004) is computed as an average of various SEDs computed by different codes. The second source of systematic error is the grid. We use a 3D Cartesian grid with one level of refinement. The reference codes calculate the benchmark in a 2D grid optimized to account for the disk symmetry. It is possible to improve our solution by increasing the number of grid cells as well as the number of photon packages. However, the current hardware gives an upper limit to run time and memory requirements.

6. CLUMPY AGN TORUS MODEL

The unified scheme of AGNs (Antonucci & Miller 1985; Padovani & Urry 1990) predicts that the central engine is surrounded by dust so that the obscuration depends on the viewing angle. The detailed geometrical configuration of dust

around the AGN is still a matter of debate (Stalevski et al. 2012). Observations are not able to resolve the inner part of the objects. Theoretical considerations favor a toroidal and clumpy geometry (Nenkova et al. 2002; Schartmann et al. 2005; Höning et al. 2006; Nenkova et al. 2008; Schartmann et al. 2008, 2010). Observational hints of a clumpy structure in which optically thick dust clouds surround the accretion disk of the central black hole are provided by the $10\mu\text{m}$ silicate band. This band is detected in absorption and emission. The strength of the feature is much weaker than predicted by homogeneous torus models (Pier & Krolik 1992; Granato & Danese 1994; Efstathiou & Rowan-Robinson 1995; van Bemmel & Dullemond 2003). A silicate emission feature is also seen in obscured AGNs, where the broad emission lines are hidden (Sturm et al. 2006; Mason et al. 2009). Observations are explained by postulating a clumpy torus structure.

For the dust around AGNs we consider fluffy agglomerates of silicate and carbon as sub-particles. We use a power-law size distribution: $n(a) \propto a^{-3.5}$ with particle radii between $a_- \leq a \leq a_+$ and $a_- = 160 \text{ \AA}$, $a_+ = 0.13 \mu\text{m}$. The bulk density of both materials is 2.5 g cm^{-3} . Dust abundances (ppm) are 31 [Si]/[H] and 200 [C]/[H], which agree with cosmic abundance constraints (Asplund et al. 2009). This gives a gas-to-dust mass ratio of 130. We take the relative volume fractions in composite grain to be 34% silicates, 16% carbon, and 50% vacuum, which translates, with abundances as above, to a relative mass fraction of 68% silicates and 32% amorphous carbon. Absorption and scattering cross sections and the scattering asymmetry factor are computed with Mie theory and the Bruggeman mixing rule. Optical constants for silicates are taken from Draine (2003) and for carbon we use the ACH2 mixture from Zubko et al. (1996). This gives a total mass extinction cross section in the optical ($0.55 \mu\text{m}$) of $K_V^{\text{ext}} = 35000$. The wavelength dependence of K^{ext} is displayed in Kruegel & Siebenmorgen (1994).

The hard (UV/optical) AGN radiation emerges from the accretion disk around the massive black hole. Its spectral shape is suggested to follow a broken power law (Rowan-Robinson 1995):

$$\lambda F_\lambda \propto \begin{cases} \lambda^{1.2} & \text{for } -\infty < \lambda < 0.01 \mu\text{m} \\ \text{const} & \text{for } 0.01 < \lambda < 0.10 \mu\text{m} \\ \lambda^{-0.5} & \text{for } 0.1 < \lambda < 1.0 \mu\text{m} \\ \lambda^{-3} & \text{for } 1.0 < \lambda < +\infty \end{cases} \quad (20)$$

In the following we consider a clumpy AGN dust torus of total luminosity of $L_{45} = 10^{45} \text{ erg s}^{-1}$. The inner radius of the torus r_{in} is set by dust evaporation, which is approximated by $r_{\text{in}} = 0.4(L_{45})^{0.5} \text{ pc}$. A clump is represented as a sphere of constant density and radius of 0.5 pc . The optical depth through the center of the clump is $\tau_V = 30$. The clumps are randomly distributed within a half-opening angle of $\theta = 45^\circ$ between an inner radius of $r_{\text{in}} = 0.4 \text{ pc}$ and outer radius of $r_{\text{out}} = 6 \text{ pc}$. For random numbers $\zeta_1, \zeta_2, \zeta_3$ the position of the center of the clouds is computed by

$$\begin{aligned} r &= \zeta_1(r_{\text{out}} - r_{\text{in}}) \\ \phi &= 2\pi\zeta_2 \\ \theta &= \frac{\pi}{2} + (2\zeta_3 - 1)\theta_1. \end{aligned} \quad (21)$$

If clumps overlap the density is constant and the same as for a single clump.

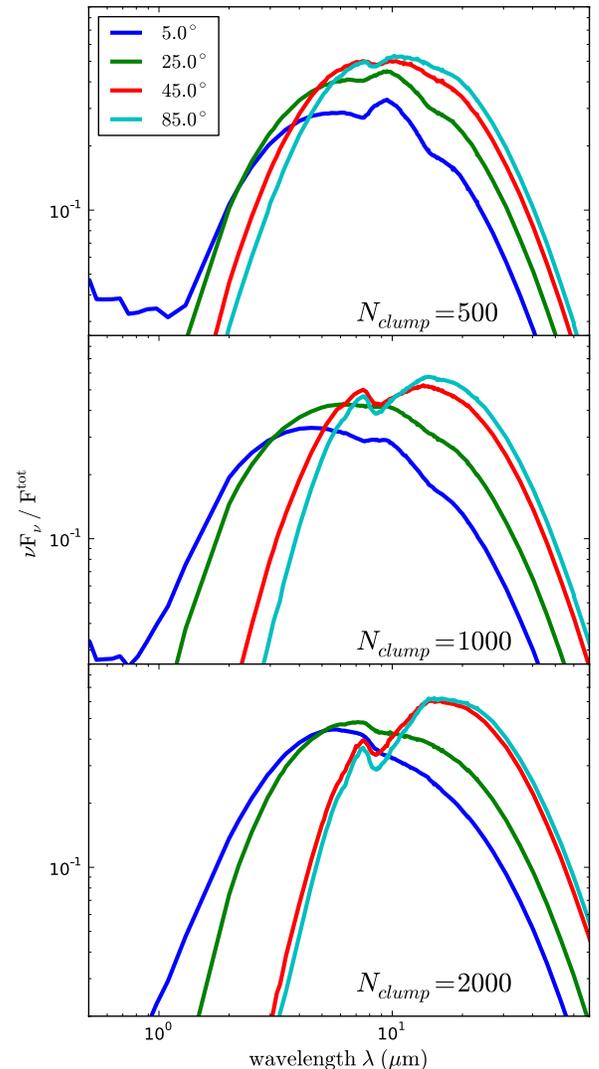


Figure 4. Spectral energy distribution of clumpy AGN torus models at viewing angles between 5° and 85° and number of clouds of $N_{\text{clump}} = 500, 1000,$ and 2000 .

(A color version of this figure is available in the online journal.)

6.1. Influence on Clumps on SED

The AGN torus model is used to study the influence of clumps on the SED. In the models the number of clouds is varied using $N_{\text{clump}} = 500, 1000,$ and 2000 . This gives a total optical depth along the midplane of the torus of $\tau_V \sim 80, 140,$ and 200 . The SED is displayed in Figure 4; at wavelengths $\lesssim 1 \mu\text{m}$, the flux is emitted by the central source and between $1 \mu\text{m}$ and $50 \mu\text{m}$ the SED is dominated by dust emission. The dust temperatures are between 100 and 1500 K. At short wavelengths of the SED the AGN is visible only in the face-on view and becomes faint at higher inclination angles. At longer wavelengths the flux is stronger in the edge-on view as compared to observations at smaller inclination angles. The optical depth of a single clump is high and already one or two of such intervening clouds define most of the spectral behavior of the torus. The SEDs at viewing angles between 45° and 85° are similar. At these angles a few clouds are always in the line of sight and dominate the spectrum (Figure 4).

Depending on the viewing angle and the total optical depth and dust temperatures in the AGN, the $10 \mu\text{m}$ silicate band is

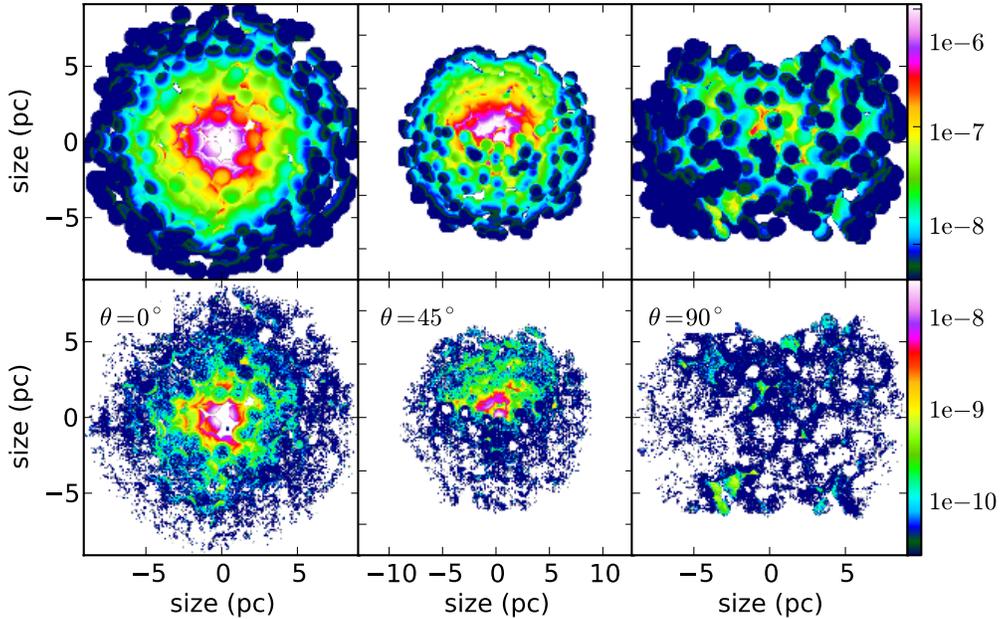


Figure 5. Images at $10\ \mu\text{m}$ (top) and $0.55\ \mu\text{m}$ (bottom) of the clumpy AGN torus with $N_{\text{clump}} = 1000$ and viewing angles (from left to right) of 0° , 45° , and 90° . Color bar gives flux in (Jy arcsec^{-2}). The source with luminosity $L = 10^{45}\ \text{erg s}^{-1}$ is at a distance of 15 Mpc. (A color version of this figure is available in the online journal.)

observed either in emission or in absorption. In AGNs viewed face-on (type I sources) there is a direct view on the inner region of the torus and a weak $10\ \mu\text{m}$ silicate emission feature is observed (Siebenmorgen et al. 2005; Hao et al. 2005; Sturm et al. 2005). In edge-on sources (type II) the $10\ \mu\text{m}$ silicate band is seen in absorption.

The silicate feature traces the amount of dust for a given viewing direction and constrains the distribution of the clumps in the torus (Nikutta et al. 2009). Homogeneous AGN torus models without a clumpy structure overpredict the strength of the silicate feature (Efstathiou & Rowan-Robinson 1995). In the clumpy torus models presented in Figure 4 the $10\ \mu\text{m}$ feature is in emission for face-on views as long as $N_{\text{clump}} \lesssim 1000$. Individual clouds are optically thick at $10\ \mu\text{m}$ so that a single cloud already obscures the warm and optical thin region required to observe the silicate band in emission. This is the case in the face-on view of the model with 2000 clumps where the emission and self-absorption of the silicate grains cancel out and the spectrum becomes featureless. For inclination angles $\theta \geq 45^\circ$ the silicate features are in absorption (Roche et al. 1991). The strength of the absorption feature depends on the optical depth and therefore increases with increasing number of clumps.

6.2. AGN Images

The MC combined with the ray tracer (Section 4) allows computing of images of the object at a given wavelength. In Figure 5 we present a $10\ \mu\text{m}$ and $0.55\ \mu\text{m}$ image of the clumpy AGN torus using $N_{\text{clump}} = 1000$. The mid-IR image is dominated by dust emission and the optical image by scattered light. The face-on image in the mid-IR provides an unobscured view of the emission from hot dust in the inner torus close to the AGN. In this image, clouds near the center dominate the emission. Further out, clumps become cooler and the contribution to the emission decreases. In the tilted view, at 45° , intervening clumps obscure part of the emission from the central region and the contribution of cooler dust becomes more important than in the face-on view. In the edge-on view most

of the central region is no longer visible and the emission is dominated by clouds located close to the border of the torus. However, close to the edge-on view and because of the clumpy nature of the medium there are still a few lines of sight penetrating to the inner torus. This is not the case in AGN models using a homogeneous dust distribution. For all viewing angles and clump configurations the scattered light images appear similar to the emission images (Figure 5). However, the flux in the optical is two orders of magnitude fainter than at $10\ \mu\text{m}$. The scattering light image in the edge-on view becomes fuzzy. The assumption of isotropic scattering likely over predicts the scattered light in the face-on direction. However, the clumpy structure visible in the optical image is preserved.

6.3. Quasar 3C 249.1

The clumpy AGN dust torus model is applied to fit the SED of the quasar 3C249.1. The luminosity of the object is $7 \times 10^{45}\ \text{erg s}^{-1}$ at redshift of $z = 0.311$, which translates, using standard cosmological parameters, to a luminosity distance of 1600 Mpc. The torus inner edge is set by the sublimation temperature and leads to an inner radius of $r_{\text{in}} = 1\ \text{pc}$. In the inner region of the model at $r \leq 25 \times r_{\text{in}}$ the half-opening angle of the torus is $\theta_1 = 25^\circ$ and in the outer region up to $r_{\text{out}} = 50 \times r_{\text{in}}$ it is 55° . The torus includes 5500 clouds. The number of clumps per volume element decreases in the inner part and increases in the outer part of the torus; other parameters remain the same.

In order to fit data at $\lambda > 40\ \mu\text{m}$ we increased the outer radius of the torus r_{out} and considered a two phase medium in which the density in between the clumps is varied. However, the far-IR emission of the quasar cannot be explained by the AGN torus. The far-IR is dominated by emission of cold dust located further out in the host galaxy. For simplicity, we approximate this cirrus component by a modified blackbody, $\kappa_{\nu}^{\text{abs}} B_{\nu}(T)$ at 50 K. The dust torus model fits the silicate emission feature as observed in the *Spitzer* spectrum and together with the cirrus component a good fit to the overall SED is achieved (Figure 6).

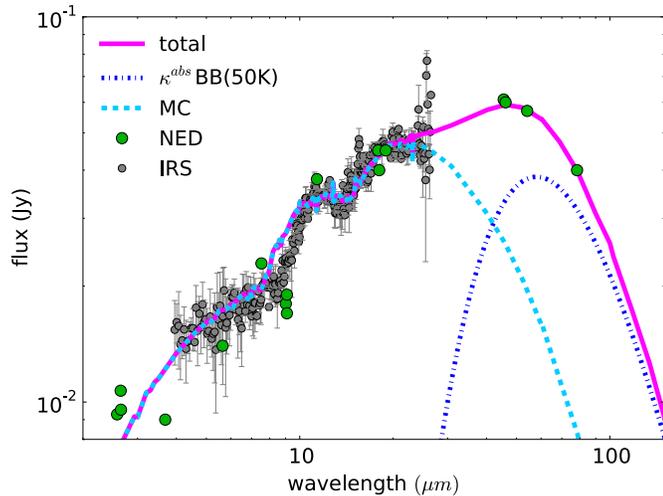


Figure 6. IR emission of the quasar 3C249.1. Photometric data by NED and *Spitzer* spectrum by Siebenmorgen et al. (2005). The total flux (full line) is given as sum of the clumpy AGN torus model (dashed) and a cirrus component (dashed dotted).

(A color version of this figure is available in the online journal.)

7. CONCLUSIONS

A parallel 3D MC method is presented to solve the RT problem of dust obscured objects in arbitrary geometries. The code uses an adaptive 3D Cartesian grid. It utilizes the advantage of different optimization algorithms: in optical thin cells the Lucy algorithm and in cells at very high optical depths a modified random walk procedure by Fleck & Canfield (1984) is included. The temperature of the dust is computed in an iterative MC method or with the iteration-free Bjorkman & Wood (2001) algorithm.

The SED of the objects is derived by a photon counting procedure or a ray tracing routine. Photons which eventually escape the MC model space are counted and converted into a flux density. This method computes the SED of an object very quickly when a large aperture of several degree is used. Unfortunately, for high spatial resolution the method is starved by the low count rate. The SED of an object observed with a pencil beam is computed by a ray tracer which uses the dust scattering and absorption events of the MC cells as input. The ray tracer allows computing of noise-free images of the scattered and emitted radiation at any frequency.

MC schemes are known to be computing-time expensive. Therefore, we apply particular attention to solve this problem. We take advantage of the independence of the trajectories of each photon packet. This allows a highly vectorized design of the MC program which is not realized in previous work. With the vectorized code a speed improvement of about 100 is reached on a low budget computer when compared to a scalar, non-parallel version of the same program. This gain in the speed performance of the computations is reached by applying either a multi-core application using conventional CPUs or the recent technological development of graphics power units (GPU). The parallelization capabilities of various MC RT algorithms are analyzed. By combining different MC algorithms we report a linear scaling of the computing-time with the number of available threads (processor cores). The ray tracer to compute SED and images is another computer time expensive procedure and is developed, similar to the MC code, in vectorized form.

The 3D MC is tested against the ray tracer solution of the RT problem in spherical symmetry by Krügel (2008). Further we verified our procedure against 1D and 2D benchmarks of the literature. The code reproduces the spectral energy and dust temperature distributions of the test cases at high accuracy. As a first astrophysical application we use the MC code to investigate the appearance of a dusty AGN at optical and IR wavelengths; a second application on protoplanetary disks is presented in an accompanying paper (Siebenmorgen & Heymann 2012b).

The dust around the AGN is geometrically configured in a clumpy torus structure and is represented by fluffy composite grains of various sizes made of silicate and carbon. The influence of the number of clumps in the torus on the SED is studied. We find that the clumpy torus explains the observed “weak” silicate absorption and emission feature observed in AGNs. Images of the AGN in the optical and the mid-IR are presented by viewing the torus from different angles. The SED of the radio-loud quasar 3C249.1 is fit by the torus model including a cirrus component for the far-IR and submillimeter emission.

We are grateful to Endrik Krügel for discussions and helpful suggestions. This work is based in part on observations made with the *Spitzer Space Telescope*, which is operated by the Jet Propulsion Laboratory, California Institute of Technology under a contract with NASA. This work was partially supported by NASA and NSF.

REFERENCES

- Antonucci, R. R. J., & Miller, J. S. 1985, *ApJ*, **297**, 621
 Asplund, M., Grevesse, N., Sauval, A. J., & Scott, P. 2009, *ARA&A*, **47**, 481
 Baes, M. 2008, *MNRAS*, **391**, 617
 Baes, M., Stamatellos, D., Davies, J. I., et al. 2005, *New Astron.*, **10**, 523
 Baes, M., Verstappen, J., De Looze, I., et al. 2011, *ApJS*, **196**, 22
 Bianchi, S. 2008, *A&A*, **490**, 461
 Bjorkman, J. E., & Wood, K. 2001, *ApJ*, **554**, 615
 Chiang, E. I., & Goldreich, P. 1997, *ApJ*, **490**, 368
 Chiang, E. I., & Goldreich, P. 1999, *ApJ*, **519**, 279
 CUDA 2011, *NVIDIA C Programming Guide* (Santa Clara: NVIDIA)
 Draine, B. T. 2003, *ApJ*, **598**, 1017
 Draine, B. T., & Lee, H. M. 1984, *ApJ*, **285**, 89
 Efstathiou, A., & Rowan-Robinson, M. 1995, *MNRAS*, **273**, 649
 Fleck, J. J., & Canfield, E. 1984, *J. Comput. Phys.*, **54**, 508
 Gordon, K. D., Misselt, K. A., Witt, A. N., & Clayton, G. C. 2001, *ApJ*, **551**, 269
 Granato, G. L., & Danese, L. 1994, *MNRAS*, **268**, 235
 Haas, M., Willner, S. P., Heymann, F., et al. 2008, *ApJ*, **688**, 122
 Hao, L., Spoon, H. W. W., Sloan, G. C., et al. 2005, *ApJ*, **625**, L75
 Henyey, L. G., & Greenstein, J. L. 1941, *ApJ*, **93**, 70
 Heymann, F. 2010, PhD thesis, Univ. Bochum
 Hönig, S. F., Beckert, T., Ohnaka, K., & Weigelt, G. 2006, *A&A*, **452**, 459
 Ivezić, Z., & Elitzur, M. 1997, *MNRAS*, **287**, 799
 Ivezić, Z., Groenewegen, M. A. T., Men’shchikov, A., & Szczerba, R. 1997, *MNRAS*, **291**, 121
 Jonsson, P. 2006, *MNRAS*, **372**, 2
 Jonsson, P., & Primack, J. R. 2010, *New Astron.*, **15**, 509
 Krolik, J. H., & Begelman, M. C. 1988, *ApJ*, **329**, 702
 Kruegel, E., & Siebenmorgen, R. 1994, *A&A*, **288**, 929
 Krügel, E. 2008, *An Introduction to the Physics of Interstellar Dust* (New York: Taylor & Francis)
 Lucy, L. B. 1999, *A&A*, **344**, 282
 Mason, R. E., Levenson, N. A., Shi, Y., et al. 2009, *ApJ*, **693**, L136
 Matsumoto, M., & Nishimura, T. 2000, in *Monte Carlo and Quasi-Monte Carlo Methods 1998*, ed. H. Niederreiter & J. Spanier (Berlin: Springer-Verlag), 59
 Min, M., Dullemond, C. P., Dominik, C., de Koter, A., & Hovenier, J. W. 2009, *A&A*, **497**, 155
 Misselt, K. A., Gordon, K. D., Clayton, G. C., & Wolff, M. J. 2001, *ApJ*, **551**, 277
 Nenkova, M., Ivezić, Ž., & Elitzur, M. 2002, *ApJ*, **570**, L9

- Nenkova, M., Sirocky, M. M., Ivezić, Ž., & Elitzur, M. 2008, *ApJ*, **685**, 147
- Nikutta, R., Elitzur, M., & Lacy, M. 2009, *ApJ*, **707**, 1550
- Padovani, P., & Urry, C. M. 1990, *ApJ*, **356**, 75
- Pascucci, I., Wolf, S., Steinacker, J., et al. 2004, *A&A*, **417**, 793
- Pier, E. A., & Krolik, J. H. 1992, *ApJ*, **399**, L23
- Robitaille, T. P. 2010, *A&A*, **520**, A70
- Roche, P. F., Aitken, D. K., Smith, C. H., & Ward, M. J. 1991, *MNRAS*, **248**, 606
- Rowan-Robinson, M. 1995, *MNRAS*, **272**, 737
- Schartmann, M., Burkert, A., Krause, M., et al. 2010, *MNRAS*, **403**, 1801
- Schartmann, M., Meisenheimer, K., Camenzind, M., Wolf, S., & Henning, T. 2005, *A&A*, **437**, 861
- Schartmann, M., Meisenheimer, K., Camenzind, M., et al. 2008, *A&A*, **482**, 67
- Siebenmorgen, R., Haas, M., Krügel, E., & Schulz, B. 2005, *A&A*, **436**, L5
- Siebenmorgen, R., & Heymann, F. 2012a, *A&A*, **539**, A20
- Siebenmorgen, R., & Heymann, F. 2012b, *A&A*, **539**, A20
- Siebenmorgen, R., Heymann, F., & Krügel, E. 2011, in PAHs and the Universe, ed. C. Joblin & A. G. G. M. Tielens (EAS Publication Series, Vol. 46; Les Ulis, France: EDP Sciences), 285
- Siebenmorgen, R., & Krügel, E. 1992, *A&A*, **259**, 614
- Siebenmorgen, R., & Krügel, E. 2010, *A&A*, **511**, A6
- Stalewski, M., Fritz, J., Baes, M., Nakos, T., & Popovic, L. C. 2012, *MNRAS*, **420**, 2756
- Sturm, E., Hasinger, G., Lehmann, I., et al. 2006, *ApJ*, **642**, 81
- Sturm, E., Schweitzer, M., Lutz, D., et al. 2005, *ApJ*, **629**, L21
- Tielens, A. G. G. M. 2008, *ARA&A*, **46**, 289
- Tristram, K. R. W., Meisenheimer, K., Jaffe, W., et al. 2007, *A&A*, **474**, 837
- van Bemmell, I. M., & Dullemond, C. P. 2003, *A&A*, **404**, 1
- Witt, A. N. 1977, *ApJS*, **35**, 1
- Wolf, S. 2003, *Comput. Phys. Commun.*, **150**, 99
- Yusef-Zadeh, F., Morris, M., & White, R. L. 1984, *ApJ*, **278**, 186
- Zubko, V. G., Mennella, V., Colangeli, L., & Bussoletti, E. 1996, *MNRAS*, **282**, 1321