UCLA

Title

Benchmarking the calculation of stochastic heating and emissivity of dust grains in the context of radiative transfer simulations

Permalink

https://escholarship.org/uc/item/2zs5f4fg

Authors

Camps, P
Misselt, K
Bianchi, S
et al.

Peer reviewed

Benchmarking the calculation of stochastic heating
and emissivity of dust grains
in the context of radiative transfer simulations

Peter Camps[1], Karl Misselt[2], Simone Bianchi[3], Tuomas Lunttila[4], Christophe Pinte[5,6,7], Giovanni Natale[8], Mika Juvela[9], Joerg Fischera[10], Michael P. Fitzgerald[11], Karl Gordon[12,1], Maarten Baes[1], and Jürgen Steinacker[6,13]

[1] Sterrenkundig Observatorium, Universiteit Gent, Krijgslaan 281, 9000 Gent, Belgium
    e-mail: peter.camps@ugent.be
[2] Steward Observatory, University of Arizona, Tucson, AZ 85721-0065, USA
[3] INAF-Osservatorio Astrofisico di Arcetri, Largo Enrico Fermi 5, 50125 Firenze, Italy
[4] Chalmers University of Technology, Department of Earth and Space Sciences, Onsala Space Observatory, 439 92 Onsala, Sweden
[5] UMI-FCA, CNRS-INSU, France (UMI 3386) and Universidad de Chile, Cerro Calán, 1058 Santiago, Chile
[6] Univ. Grenoble Alpes, IPAG, 38000 Grenoble, France
[7] CNRS, IPAG, 38000 Grenoble, France
[8] University of Central Lancashire, Jeremiah Horrocks Institute, Leighton building, Preston, PR1 2HE, UK
[9] Department of Physics, PO Box 64, University of Helsinki, 00014 Helsinki, Finland
[10] Max-Planck-Institut für Kernphysik, Saupfercheckweg 1, 69117 Heidelberg, Germany
[11] Department of Physics and Astronomy, 430 Portola Plaza, University of California, Los Angeles, CA 90095-1547, USA
[12] Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218, USA
[13] Max-Planck-Institut für Astronomie, Königstuhl 17, 69117 Heidelberg, Germany

ABSTRACT

Context. Thermal emission by stochastically heated dust grains (SHGs) plays an important role in the radiative transfer (RT) problem for a dusty medium. It is therefore essential to verify that RT codes properly calculate the dust emission before studying the effects of spatial distribution and other model parameters on the simulated observables.
Aims. We define an appropriate problem for benchmarking dust emissivity calculations in the context of RT simulations, specifically including the emission from SHGs. Our aim is to provide a self-contained guide for implementors of such functionality and to offer insight into the effects of the various approximations and heuristics implemented by the participating codes to accelerate the calculations.
Methods. The benchmark problem definition includes the optical and calorimetric material properties and the grain size distributions for a typical astronomical dust mixture with silicate, graphite, and PAH components. It also includes a series of analytically defined radiation fields to which the dust population is to be exposed and instructions for the desired output. We processed this problem using six RT codes participating in this benchmark effort and compared the results to a reference solution computed with the publicly available dust emission code DustEM.
Results. The participating codes implement different heuristics to keep the calculation time at an acceptable level. We study the effects of these mechanisms on the calculated solutions and report on the level of (dis)agreement between the participating codes. For all but the most extreme input fields, we find agreement within 10% across the important wavelength range $3\,\mu m \le \lambda \le 1000\,\mu m$.
Conclusions. We conclude that the relevant modules in RT codes can and do produce fairly consistent results for the emissivity spectra of SHGs. This work can serve as a reference for implementors of dust RT codes, and it will pave the way for a more extensive benchmark effort focusing on the RT aspects of the various codes.

Key words. radiation mechanisms: thermal – dust, extinction – infrared: ISM – radiative transfer – methods: numerical

1. Introduction

Dust substantially affects the radiation emerging from many astrophysical systems. To study the three-dimensional structure of these systems, it is often useful to numerically simulate the transport of radiation through a model that includes a dusty medium with appropriate characteristics. Many authors have described radiative transfer (RT) codes designed to tackle this problem. For an overview, we recommend, for example, Whitney (2011) and Steinacker et al. (2013). In multiwavelength studies, thermal emission by the dust plays an important role. While larger dust grains can often be assumed to be in local thermal equilibrium (LTE) with the surrounding radiation field, this assumption does not typically hold for very small grains (VSGs) or for polycyclic aromatic hydrocarbon molecules (PAHs). The absorption of a single optical photon substantially boosts the internal energy of such a small collection of atoms, causing its emission spectrum to vary over time (see, e.g., Sellgren 1984; Dwek 1986; Boulanger & Perault 1988; Helou et al. 2000; Draine et al. 2007; Smith et al. 2007). We use the term *stochastically*

heated grains (SHGs) to collectively indicate dust grains and PAH molecules that cannot be assumed to be in LTE with the radiation field.

Since SHGs spend a significant amount of time at much higher energy levels than if they are in LTE, they emit at shorter wavelengths, which can have a strong effect on the observed spectrum. It is therefore essential to verify that RT codes properly calculate this emission before studying the effects of spatial distribution and other model parameters on the simulated observables. In this paper we present a benchmark problem for this purpose, including a dust model, a series of input radiation fields, and a reference dust emission spectrum for each input field. The dust model described here has been designed for use in the TRUST benchmarks[1] (Transport of Radiation through a dUSTy medium), which test the actual RT aspect of various codes.

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

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

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

The information provided in this paper and in the accompanying data files[2] is self-contained. Readers can implement the code to calculate the SHG emission, set up the benchmark tests, and verify the results, based solely on the information provided here, without referring to other sources.

In Sects. 2 and 3 we define the dust model and the input radiation fields used in the benchmark problem. Section 4 presents the linear equation method used by the RT codes represented in this paper to calculate the SHG emission spectra, and Sect. 5 briefly introduces each of these codes. In Sect. 6 we compare the

results produced by the RT codes with the reference solution, and we discuss the differences between the methods used by the various codes and how they influence the results. Finally, in Sect. 7 we summarize and conclude.

For ease of reference, Table 1 lists the symbols used in this paper for various physical quantities, with the corresponding SI units.

2. Dust model

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

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

2.1. Optical grain properties

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

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

---

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

Table 1. Symbols used in this paper for various physical quantities, with the corresponding SI units.

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

computed at 1201 wavelength points logarithmically spaced between $0.001\,\mu$m and $10\,000\,\mu$m. The efficiencies thus derived can be cast as either cross sections or mass coefficients:

$$\sigma_k^{\mathrm{abs,sca}}(a, \lambda) = \pi a^2 Q_k^{\mathrm{abs,sca}}(a, \lambda) \tag{1}$$

$$\kappa_k^{\mathrm{abs,sca}}(a, \lambda) = \frac{3 Q_k^{\mathrm{abs,sca}}(a, \lambda)}{4 a \rho_k}. \tag{2}$$

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

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

2.2. Grain size distributions

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

$$\log O(a) = c_0 + b_0 \log\left(\frac{a}{\mu\mathrm{m}}\right) - b_1 \left|\log\left(\frac{a}{a_1}\right)\right|^{m_1} - b_2 \left|\log\left(\frac{a}{a_2}\right)\right|^{m_2}$$
$$- b_3 \left|\frac{a - a_3}{\mu\mathrm{m}}\right|^{m_3} - b_4 \left|\frac{a - a_4}{\mu\mathrm{m}}\right|^{m_4} \tag{3}$$

where a is the grain size and a_i, b_i, c_i, m_i represent constant parameters. In the interest of clarity, the notation in Eq. (3) omits a subscript k indicating the type of material being referenced. For the BARE-GR-S model, we require three sets of parameters, one for each of the materials comprising the dust model. The appropriate values can be found in Table 7 of Zubko et al. (2004) and are reproduced here in Table 2.

By construction, Eq. (3) has been normalized such that

$$\int_{a_k^{\mathrm{min}}}^{a_k^{\mathrm{max}}} O_k(a)\, \mathrm{d}a = 1, \tag{4}$$

Table 2. Parameters for the analytical approximation to the BARE-GR-S size distribution defined in Eqs. (3) through (5) and the bulk densities for the three components in our dust model.

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

Notes. With these parameter values, the grain size a substituted in the equations must be expressed in μm, and the resulting size distribution $\Omega_k(a)$ is expressed in number of dust grains per hydrogen atom and per μm.

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

$$\Omega_k(a) = A_k \, O_k(a), \qquad (5)$$

where A_k is the overall normalization of component k, as listed in Table 2.

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

$$\mathcal{N}_{\mathrm{D}}/\mathcal{N}_{\mathrm{H}} = \sum_k \int_{a_k^{\mathrm{min}}}^{a_k^{\mathrm{max}}} \Omega_k(a)\,\mathrm{d}a \qquad (6)$$

$$\mu = M/\mathcal{N}_{\mathrm{H}} = \sum_k \int_{a_k^{\mathrm{min}}}^{a_k^{\mathrm{max}}} \frac{4}{3}\pi a^3 \rho_k^{\mathrm{bulk}} \Omega_k(a)\,\mathrm{d}a. \qquad (7)$$

We also provide the grain size distributions defined by Eqs. (3) and (5) and Table 2 in tabulated form (see Sect. 2.4), on a size grid with 24, 62, and 63 samples for the PAH, graphite, and silicate component, respectively.

2.3. Calorimetric grain properties

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

$$H_k(a, T) = \frac{4}{3}\pi a^3 \rho_k^{\mathrm{bulk}} h_k(T) \qquad (8)$$

$$h_k(T) = \int_0^T c_k(T')\,\mathrm{d}T' \qquad (9)$$

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

We elected to use the heat capacity functions proposed in Draine & Li (2001) and Li & Draine (2001). To avoid subtle differences between implementations, we provide this information in tabulated form (see Sect. 2.4) rather than expecting each code to implement the equations. One table describes the graphitic components (PAH molecules and graphite grains) and another one describes the silicate component. Each table lists the specific enthalpy and the specific heat capacity at 1000 logarithmically spaced temperature points ranging from 1 K to 2500 K.

Finally, the bulk densities for the three components in our dust model are listed in Table 2.

2.4. Data files

The data files defining the dust properties described above can be downloaded from the web site indicated in footnote 2. They are contained in the `DustModel` directory, which is organized in the following subdirectories:

- `GrainInputs`: optical and calorimetric properties for each of the three grain types k. Optical properties include the absorption efficiencies $Q_k^{\mathrm{abs}}(a, \lambda)$ tabulated on a grid of 1201 wavelengths λ and 28 (for PAH) or 121 (for graphite and silicate) grain sizes a. Calorimetric properties include the specific enthalpy $h(T)$ and the specific heat capacity $c(T)$ of the grain material, tabulated on a grid of 1000 temperatures T. The calorimetry data for graphitic grains should be used for PAH molecules as well.
- `SizeInputs`: tabulated grain size distributions $\Omega_k(a)$ for each grain type k, on a size grid with 24, 62, and 63 samples for the PAH, graphite, and silicate component, respectively. Implementations may choose to compute the size

distribution from the functional form defined in Sect. 2.2, or to load the tabulated data.

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

3. Radiation fields

3.1. Basic definitions

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

$$B(\lambda, T) = \frac{2hc^2}{\lambda^5} \frac{1}{\exp\left(\frac{hc}{\lambda kT}\right) - 1} \tag{10}$$

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

We define the solar-neighborhood interstellar radiation field (ISRF) given in Table A3 of Mathis et al. (1983) through the following functional form inspired by (but not identical to[3]) Eq. (31) in Weingartner & Draine (2001b):

$$J^{\mathrm{Mat}}(\lambda) = \begin{cases} 0 \\ \qquad\qquad\qquad\qquad\qquad \lambda < 0.0912\,\mu\mathrm{m} \\ 3069\,\mathrm{W/m^3/sr} \times (\lambda/\mu\mathrm{m})^{3.4172} \\ \qquad\qquad 0.0912\,\mu\mathrm{m} \le \lambda < 0.110\,\mu\mathrm{m} \\ 1.627\,\mathrm{W/m^3/sr} \\ \qquad\qquad 0.110\,\mu\mathrm{m} \le \lambda < 0.134\,\mu\mathrm{m} \\ 0.0566\,\mathrm{W/m^3/sr} \times (\lambda/\mu\mathrm{m})^{-1.6678} \\ \qquad\qquad 0.134\,\mu\mathrm{m} \le \lambda < 0.250\,\mu\mathrm{m} \\ 10^{-14}\,B(\lambda, 7500\,\mathrm{K}) + 10^{-13}\,B(\lambda, 4000\,\mathrm{K}) \\ \qquad\qquad\qquad + 4 \times 10^{-13}\,B(\lambda, 3000\,\mathrm{K}) \\ \qquad\qquad\qquad\qquad\qquad 0.250\,\mu\mathrm{m} \le \lambda. \end{cases} \tag{11}$$

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

$$U = \int_0^\infty J(\lambda)\, \mathrm{d}\lambda \,\Big/ \int_0^\infty J^{\mathrm{Mat}}(\lambda)\, \mathrm{d}\lambda. \tag{12}$$

3.2. Benchmark input fields

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

$$J^{\mathrm{SHG},i}(\lambda) = U_i \times J^{\mathrm{Mat}}(\lambda) \text{ with } U_i = 10^{-4}, 10^{-3}, ..., 10^5, 10^6. \tag{13}$$

---

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

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

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

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

3.3. Calculation and wavelength grid

Using the dust model described in Sect. 2, the codes participating in this benchmark calculate the spectral dust emissivity $\varepsilon(\lambda)$ for each of the input radiation fields specified in Sect. 3, taking stochastic heating of small grains into account. The radiation emitted by the dust itself is ignored with respect to the input field; i.e., it is not our intention to calculate a self-consistent radiation field. The calculations were performed and the results written down using the wavelength grid on which the optical properties have been tabulated. This is a logarithmic grid with 1201 points in the range $0.001\,\mu\mathrm{m} \le \lambda \le 10\,000\,\mu\mathrm{m}$.

4. Dust emission

4.1. Emission from a dust mixture

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

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

$$\varsigma_{k,b}^{\mathrm{abs}}(\lambda) = \int_{a_{k,b}^{\mathrm{min}}}^{a_{k,b}^{\mathrm{max}}} \pi a^2\, Q_k^{\mathrm{abs}}(a, \lambda)\, \Omega_k(a)\, \mathrm{d}a \tag{15}$$

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

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

$$M_{k,b} = \int_{a_{k,b}^{\mathrm{min}}}^{a_{k,b}^{\mathrm{max}}} \rho_k^{\mathrm{bulk}}\, \frac{4\pi}{3}\, a^3\, \Omega_k(a)\, \mathrm{d}a \Big/ \int_{a_{k,b}^{\mathrm{min}}}^{a_{k,b}^{\mathrm{max}}} \Omega_k(a)\, \mathrm{d}a, \tag{16}$$

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

$$H_{k,b}(T) = M_{k,b}\, h_k(T). \tag{17}$$

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

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

$$\varepsilon(\lambda) = \sum_{k,b} \varsigma_{k,b}^{\text{abs}}(\lambda) \int_0^\infty P_{k,b,J}(T) \, B(\lambda, T) \, \mathrm{d}T \tag{18}$$

where $B(\lambda, T)$ is the Planck function defined in Eq. (10), and $P_{k,b,J}(T)$ is the probability of finding the representative grain of bin k, b at temperature T.

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

$$j(\lambda) = \frac{M}{\mu} \, \varepsilon(\lambda) \tag{19}$$

with μ given by Eq. (7). When combining the emission from various dust mixes, it is useful to recall that it is physically meaningful to add cross sections and masses, while mass coefficients, in general, cannot be added meaningfully:

$$\kappa_1 + \kappa_2 = \frac{\varsigma_1}{\mu_1} + \frac{\varsigma_2}{\mu_2} \neq \frac{\varsigma_1 + \varsigma_2}{\mu_1 + \mu_2}. \tag{20}$$

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

In this discussion, we characterize P as a function of grain temperature. The temperature of a grain and its internal energy are related through Eq. (8), so we could equivalently characterize P as a function of internal grain energy.

4.2. Equilibrium heating dust emission

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

$$P_{k,b,J}(T) = \delta(T - T_{k,b,J}^{\text{eq}}), \tag{21}$$

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

$$\int_0^\infty \varsigma_{k,b}^{\text{abs}}(\lambda) \, J(\lambda) \, \mathrm{d}\lambda = \int_0^\infty \varsigma_{k,b}^{\text{abs}}(\lambda) \, B(\lambda, T_{k,b,J}^{\text{eq}}) \, \mathrm{d}\lambda. \tag{22}$$

4.3. Stochastic heating dust emission

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

size bin and a specific radiation field, dropping the indices k, b, and J from the notation.

We select an appropriate temperature grid with N bins T_i, $i = 0, \ldots, N - 1$ (see Sect. 4.1). Our goal is to determine the probabilities $P_i = P(T_i)$ that a grain resides in temperature bin i. We define a transition matrix $A_{f,i}$ that describes the probability per unit time for a grain to transfer from initial temperature bin i to final temperature bin f. The transition matrix elements in the case of heating ($f > i$) are given by

$$A_{f,i} = 4\pi \varsigma^{\text{abs}}(\lambda_{fi}) \, J(\lambda_{fi}) \frac{hc \, \Delta H_f}{\left[H(T_f) - H(T_i) \right]^3}, \tag{23}$$

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

$$\lambda_{fi} = \frac{hc}{H(T_f) - H(T_i)}. \tag{24}$$

We assume that a dust grain cools by radiating photons with an energy that is very small compared to the internal energy of the grain. With this continuous cooling approximation, cooling transitions occur only to the next lower level, so that $A_{f,i} = 0$ for $f < i - 1$ and

$$A_{i-1,i} = \frac{4\pi}{H(T_i) - H(T_{i-1})} \int_0^\infty \varsigma^{\text{abs}}(\lambda) \, B(\lambda, T_i) \, \mathrm{d}\lambda. \tag{25}$$

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

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

$$\sum_{i=0}^{N-1} A_{f,i} P_i = 0 \qquad f = 0, \ldots, N - 1, \tag{26}$$

along with the normalization condition

$$\sum_{i=0}^{N-1} P_i = 1 \tag{27}$$

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

$$B_{N-1,i} = A_{N-1,i} \qquad i = 0, \ldots, N - 2 \tag{28}$$
$$B_{f,i} = B_{f+1,i} + A_{f,i}, \quad f = N - 2, \ldots, 1; \; i = 0, \ldots, f - 1, \tag{29}$$

for the unnormalized probability distribution X_i

$$X_0 = 1 \tag{30}$$

$$X_i = \frac{\sum_{j=0}^{i-1} B_{i,j} X_j}{A_{i-1,i}} \qquad i = 1, \ldots, N - 1, \tag{31}$$

and finally for the normalized probabilities P_i

$$P_i = \frac{X_i}{\sum_{j=0}^{N-1} X_j} \qquad i = 0, \ldots, N - 1. \tag{32}$$

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

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

A third characteristic that may differ between implementations is the discretization of the grain size distribution, as discussed in Sect. 4.1.

5. Reference code and participating codes

5.1. DustEM

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

We produced the reference solutions in this paper with the public version of DustEM (v3.8, dated Spring 2010). To this end, we converted the dust properties defined in Sect. 2 and the input radiation fields defined in Sect. 3 into the data format expected by DustEM. We adjusted the values of the physical constants in the DustEM code, raised the maximum number of grain size and temperature bins to accommodate our input data, and fixed a minor problem in the routine that imports the grain size distribution. More important, to obtain accurate reference solutions, we substantially increased the number of temperature bins and the number of numerical iterations (see Sect. 6.2). Other than this, the DustEM code was used without modifications.

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

The following sections describe each of the participating codes with a focus on the specific heuristics employed for calculating the results presented in this paper. Table 3 offers a (very) concise overview of this information.

5.2. SKIRT

SKIRT (Baes et al. 2003, 2011; Camps & Baes 2015) is a Monte Carlo continuum RT code for simulating the effect of dust on radiation in astrophysical systems. It offers full treatment of absorption and multiple anisotropic scattering by the dust, computes the temperature distribution of the dust and the thermal dust re-emission self-consistently, and supports stochastic heating of small grains using an efficient library approach. The code handles multiple dust mixtures and arbitrary 3D geometries for radiation sources and dust populations, and offers a variety of simulated instruments for measuring the radiation field from any angle. It features a wide range of built-in components that can be configured to construct complex models without changing or adding source code.

SKIRT closely follows the method presented in Sect. 4 to calculate dust emission. The computation time for the emissivity of a stochastically heated dust mixture exposed to a particular radiation field scales roughly with BN^2, where B is the total number of size bins in the dust mixture (for all grain types

Table 3. Overview of the discretization parameters and heuristics used by participating codes.

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

Notes. Refer to Sects. 5.2 through 5.7 for a more extensive description. Also, Fig. 9 further supplements the information in the last column of this table.

combined), and N is the number of temperature bins used in the calculation.

For the results shown in this paper, we use 15 size bins b for each grain type k, distributed logarithmically over the complete size range, so that $B = 45$. The absorption efficiencies loaded from the tables described in Sect. 2.4 are interpolated logarithmically as needed to perform the integrations over grain size presented in Sect. 4.1 on a logarithmic grain size grid with 201 points within each bin.

Because of the N^2 dependency in the computation time, the choice of the temperature grid is fairly crucial. By varying N in a number of experiments, it can be easily shown that, for most input fields, the temperature probability distribution $P(T)$ for very small grains can be calculated accurately on a rather coarse grid. Larger grains require a finer grid because $P(T)$ is narrower and has steep flanks. As discussed in Sect. 4.2, for sufficiently large grains, $P(T)$ approaches a delta function, so that the procedure described in Sect. 4.3 requires an exceedingly refined grid with $N > 5000$ to produce accurate results, which becomes computationally prohibitive. Thus, in the interest of both speed and accuracy, we need to switch from transient to equilibrium calculations for grains that can be considered to be in LTE.

A RT simulation typically calculates the emissivity of a certain (fixed) dust mixture for a large number of radiation fields. This computation can be accelerated substantially by precalculating and storing the elements of the matrix $A_{f,i}$ defined in Sect. 4.3, insofar as they do not depend on the radiation field. The memory requirements scale with BN^2, just like the computation time. More importantly, precalculating these values requires a predefined temperature grid that remains fixed for all emissivity calculations. However, performing all calculations on the finest grid would be very inefficient.

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

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

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

Grid C has an average bin width of 2 K, with a power-law ratio of 3 between the largest and smallest bins. This provides a fine resolution of less than 1 K at low temperatures, while still offering decent resolution at high temperatures. This high-resolution grid is used to calculate the temperature probability distribution for dust grains with a rather narrow temperature range (but not so narrow that they would be considered to be in equilibrium).

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

1. calculate the equilibrium temperature T_{eq} for this grain;
2. use grid A to estimate the temperature range $\Delta T = T_{max} - T_{min}$ in which $P(T)/P_{max} > 10^{-20}$;
3. if $\Delta T < 10$ K or if $T_{max} < T_{eq}$, then calculate the emissivity assuming equilibrium at temperature T_{eq} and exit;
4. calculate $P(T)$ using grid B (if $\Delta T > 200$ K) or grid C (if $\Delta T < 200$ K);
5. update the temperature range $\Delta T = T_{max} - T_{min}$ in which $P(T)/P_{max} > 10^{-20}$ based on the new calculation;
6. if $\Delta T < 10$ K or if $T_{max} < T_{eq}$, then calculate the emissivity assuming equilibrium at temperature T_{eq} and exit;
7. calculate the emissivity using $P(T)$ from step 4 over the range $[T_{min}, T_{max}]$ determined in step 5.

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

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

5.3. DIRTY

DIRTY (Gordon et al. 2001; Misselt et al. 2001) is a Monte Carlo RT code designed to study dust and its effect on radiation in arbitrary astrophysical systems. DIRTY is a fully 3D code allowing for the specification of arbitrary density distributions of both dust and radiation sources. It implements an adaptive mesh, allowing for the efficient allocation of computing

resources amongst regions in the model space depending on the physical characteristics of the system. Dust absorption, temperature distribution, and emission are handled self-consistently and multiple, anisotropic scattering is implemented. The dust heating implementation supports both equilibrium and stochastic processes based on the local radiation field and dust properties at each grid in the model space.

Like other codes presented here, DIRTY follows the approach presented in Sect. 4. Internal to DIRTY, the dust grain size distribution is not further discretized beyond the input discretization of the model. For the benchmark dust model, we computed the heating and emission for all sizes in the input mesh (28 for PAH, 121 for graphitic and silicate components; see Sect. 2.4).

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

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

2. If the time scales computed in step 1 satisfy the inequality $\tau_{abs} < \tau_{rad}$, the grain is considered to be in equilibrium with the local field; it is assigned a temperature of T_{eq} and its emission is calculated following Eq. (18) with $P = \delta(T - T_{eq})$. All grains of the same composition that are larger than the size for which the inequality is first satisfied are by default treated as being in equilibrium with the local field.

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

4. Depending on the results of step 3, the algorithm proceeds in one of two directions:
   (a) If $P(T)$ is not below a specified tolerance at the endpoints of the temperature grid ($P(T_{min}), P(T_{max}) < 10^{-15}$), the temperature limits on the grid are expanded by 50%, and we return to step 3 with the new, expanded temperature grid.
   (b) If $P(T)$ is below the tolerance at the endpoints of the temperature grid, we compute the energy emitted by the stochastically heated grain using Eq. (18). If the emitted energy is within 1% of the energy absorbed from the radiation field by the grain, we consider the calculation converged and return the emitted energy spectrum. Otherwise, proceed to step 5.

5. The grid is now refined by a series of moves that refine the temperature limits and increase the number of temperature samples if necessary.
   (a) Remove all points on the ends of the temperature range that have low probabilities ($P(T) < 10^{-15}$).
   (b) Recompute the probabilities on the smaller grid with the same number of samples.
   (c) Compute the emission and emitted energy. If the emitted energy matches the absorbed energy to 1%, consider the calculation converged and return the emitted energy spectrum. Otherwise.
   (d) Keeping the same temperature limits, increase the number of samples by 50% and recompute the probabilities. Trim temperature endpoints with $P(T) < 10^{-15}$ and return to step 5b. If this step would result in the number of bins exceeding a predefined maximum (1000), record the

failure of the stochastic heating algorithm and proceed to the next grid in the model space.

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

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

5.4. TRADING

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

TRADING computes stochastic heating following the method described in Sect. 4. A single temperature grid is used to precompute the field-independent terms of the matrix elements (Eq. (25) and most factors in Eq. (23)). Since the RT models of clumpy galactic disks in Bianchi (2008) need a few million dust cells, and each cell requires the calculation of dust heating for the full grain size distribution, Bianchi (2008) uses a limited temperature grid of 80 logarithmically spaced bins between 2.7 and 2000 K. As mentioned in Sect. 4.3, faster thermal equilibrium calculations were performed for grains larger than $a_{trans} = 0.01 \mu$m. The original setup resulted in SEDs that were estimated to be within 10% of full solutions for wavelengths up to 1000 μm.

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

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

For $a_{trans} = 0.01 \mu$m, only about half of the size bins pass through the stochastic heating calculation. For $a_{trans} = 0.05 \mu$m, this goes up to 75%: all of the PAH bins and 14 of the 20 graphite and silicate bins. If a grain with $a < a_{trans}$ attains the condition of thermal equilibrium, the adopted temperature grid might fail to compute the resulting narrow probability function. Thus, in addition to the grain size cut-off, we chose to assume thermal equilibrium whenever the range of the computed temperature distribution defined by $P(T) > 10^{-15}$ is smaller than the equilibrium temperature.

5.5. CRT

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

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

To select the temperature grid that is used for calculating the emission for a grain size and type in a given input field, we calculate the absorption time scale τ_{abs} and the mean energy of the absorbed photon $\langle E_{abs} \rangle$. We choose the predefined grid that has been built for the reference field most like the current input field. In the calculations presented in this paper, the selected grid is simply the one for whose reference field the mean absorbed power per grain $\mathcal{P}_{abs} = \langle E_{abs} \rangle / \tau_{abs}$ is closest to the values calculated for the input field. Although the selected temperature grid is not necessarily optimal, it is good enough to allow accurate results using a modest number of energy bins.

The computation of emission from stochastically heated grains in CRT differs slightly from the description given in Sect. 4.3. Instead of using Eq. (23) for calculating the upward transition rates, we apply Eqs. (15)–(25) and (28) from Draine & Li (2001), which include corrections for the finite size of an enthalpy bin. Similarly, instead of Eq. (25), we use Eq. (41) from Draine & Li (2001) for calculating the cooling part of the transition matrix. Including these finite bin size corrections allows for using a lower resolution grid, which substantially benefits the computation time.

The cooling matrix elements ($f < i$) are independent of the radiation field and they are precomputed. The heating elements ($f > i$) depend on the radiation field and must be calculated separately for each input field. Moreover, when using the equations from Draine & Li (2001), instead of using the radiation field strength at a single wavelength as in Eq. (23), we must integrate numerically over the wavelength grid. We use precomputed integration weights $w_{f,i}$ corresponding to each grid point of the wavelength grid. The integral in Eq. (15) in Draine & Li (2001) can then be evaluated as $A_{f,i} = \sum_{k=1}^{N_\lambda} w_{f,i}(k) J(\lambda_k)$. If there are a lot of points in the wavelength grid, calculating the full sum is slow. However, for given energy bins i and f, radiation within only a narrow range of wavelengths can induce transitions $i \rightarrow f$. Therefore, $w_{f,i}(k)$ is non-zero only for a few k, and only non-zero integration weights are stored and used in the summation.

Discretization of the particle size distribution can be defined by the user. In this paper we employ the same discretization as SKIRT: 15 logarithmic bins for each of the three dust components. Dust properties for each size bin are calculated according to Eqs. (15)–(17) using numerical integration with 256 grid points.

5.6. MCFOST

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

MCFOST computes the stochastic heating of small dust grains following the method presented in Sect. 4, with a few refinements to ensure numerical stability and speed. We first compute the time between two successive absorptions of a photon and compare it to the cooling time of the grain, following the method described in Draine & Li (2001). For dust grains where the time between two absorptions is shorter than the cooling time, only the equilibrium temperature is calculated.

For those grains that have a shorter cooling time, we compute the full temperature probability distribution using a fixed temperature grid with 300 points logarithmically distributed between 1 K and 3000 K. The cooling terms of the transition matrix (Eq. (25)), which are independent of the radiation field, are precomputed, as are most of the heating terms factors (Eq. (23)). We estimate the specific intensity $J(\lambda_{fi})$ for each term in Eq. (23) by interpolating the radiation field computed by the Monte Carlo run. Since the interpolation coefficients are identical for every cell in the model, they are also precomputed to speed the calculations up.

For dust in radiative equilibrium, MCFOST solves the problem of self-consistent dust heating and re-emission using the immediate re-emission concept of Bjorkman & Wood (2001). This methods eliminates the need for iteration and ensures a perfect conservation of energy. For non-equilibrium dust grains, however, this procedure is prohibitive because it requires a temperature calculation at each absorption or re-emission. Instead, we use the classical iterative scheme where we store the energy absorbed by the dust, compute the temperature probability distribution in all cells once all the packets have been propagated, and re-emit the absorbed energy via new packets according to the new temperature probability. The procedure is iterated until a desired convergence on the temperature or energy is reached. Because only the fraction of radiation that is absorbed by the non-equilibrium grains needs to be re-emitted in an iterative way, convergence is usually reached after only a few

iterations. In practice, when a packet is absorbed inside a cell, its energy is split: the fraction absorbed by dust grains in equilibrium is immediately re-emitted, while the fraction absorbed by non-equilibrium grains is stored to be re-emitted during the next iteration. For the benchmark presented in this paper, the input radiation field is fixed so no iteration is required.

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

5.7. DART-Ray

DART-Ray is a ray-tracing 3D dust RT code that implements the RT algorithm described in Natale et al. (2014). It can be used to derive radiation-field energy-density distributions and outgoing radiation surface brightness maps for arbitrary 3D distributions of dust mass and stellar emission. It includes treatment of both absorption and anisotropic scattering. For the dust emission calculations, DART-Ray uses the prescription initially incorporated into the 2D RT model of Popescu et al. (2000) and later updated in Popescu et al. (2011). However, unlike the 2D models where the stochastically heated dust emission could be explicitly computed for each individual position, the calculations for the 3D models, often containing millions of cells, can be accelerated by using an adaptive SED library approach (see Natale et al. 2015).

For a given radiation field intensity spectrum found for a particular model cell, the stochastically heated dust emission is derived following Voit (1991). The method used to determine the probability distribution $P(T)$ combines the numerical integration of Guhathakurta & Draine (1989) with a step-wise analytical solution. The algorithm provides accurate and swift results on a relatively coarse grid. This is particularly useful for larger grain sizes where the probability distribution $P(T)$ converges to a narrow distribution around the equilibrium temperature. In the case of the first-order integration of Guhathakurta & Draine (1989), the width of the energy bins needs to be considerably smaller than the mean deposited energy to preserve energy balance. An increasing number of energy bins would be required to avoid energy losses in the heating process, which would make the calculation inefficient (see Fischera 2000, and Sect. 4.3).

The absorption of CMB photons is assumed to provide a continuous heating source. It is taken into consideration by subtracting the heating rate related to the CMB from the cooling rate (Eq. (25)), which limits the temperatures to values not lower than the CMB temperature. As further discussed in Sects. 6.3 and 6.6, the cooling assumed in the dust model is only valid as long as the emitted photon energy of the modified black body is low relative to the enthalpy of the grain. This assumption no longer applies at very low temperatures of small dust grains or PAH molecules (Fischera 2000). While negligible for most cases, we find that the CMB becomes a considerable heating source for the very low radiation field strength. For $U = 10^{-4}$, the heating of silicate grains by CMB photons is for all sizes larger than 10%. For graphites the contribution is only a few percentage points and it still negligible for PAH molecules.

The temperature distribution $P(T)$ for each grain size of a given composition is obtained consecutively by starting with the largest grain size and utilizing the basic characteristic that the distributions broaden with decreasing grain size. The probability distribution is only considered for all grains below a certain size where the stochastic heating leads to a considerable distribution of the dust temperatures. Above this critical grain size, the grains are assumed to radiate at the equilibrium temperature T_{eq}.

To estimate the transition from equilibrium to non-equilibrium, we apply the Gaussian approximation in the limit of large grains as derived by Voit (1991). We consider the grains to be stochastically heated if $2\sigma_T/T_{eq} > 0.05$. For $0.05 < 2\sigma_T/T_{eq} < 0.1$, we apply the Gaussian approximation, and the full $P(T)$ distributions are derived for $2\sigma_T/T_{eq} > 0.1$.

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

For the results presented in this paper, we derived the dust emission for each grain species at each grain size of the tabulated size distribution described in Sect. 2.4. To calculate the total emission, we then integrated over the size distribution and summed the contributions from different grain species. By comparing the total dust emission with the total absorbed energy, we ascertained that the energy balance for every grain species is fulfilled with an accuracy better than a few percentage points. The largest discrepancies (3−4%) are found for the most extremely scaled Mathis fields considered in the benchmark ($U = 10^{-4}$ and $U = 10^6$).

6. Results and discussion

6.1. Data files

The data files representing the benchmark results can be downloaded from the web site indicated in footnote 2. For each participating code, and for each input radiation field, the calculated solution is stored in a separate text file with columns specifying the wavelength λ (in μm); the mean intensity J_λ of the input field (in W m^{-3} sr^{-1}); and the silicate, graphite, and PAH emissivities $\lambda \varepsilon_\lambda$ (in W sr^{-1} H^{-1}), in that order. The file naming scheme and the precise file format are described on the web site.

6.2. Reference solutions

As mentioned in Sect. 5.1, we compare the results below from the codes participating in this benchmark against reference solutions generated with DustEM. Figure 1 shows these solutions for a selection of the input fields defined in Sect. 3.

To ensure proper accuracy, we increased the number of temperature bins and the number of iterations in the DustEM integral equation solver until the calculated emission converged to a stable solution; see Fig. 2. Specifically, the number of temperature bins was raised from 200 to 3500 and the number of iterations from 80 to 250. These changes dramatically increased the computation time, however this is acceptable for calculating a reference solution.

As an extra sanity check, we verified that the emissivities calculated by DustEM (using the same number of temperature bins and iterations as for the reference solutions) indeed
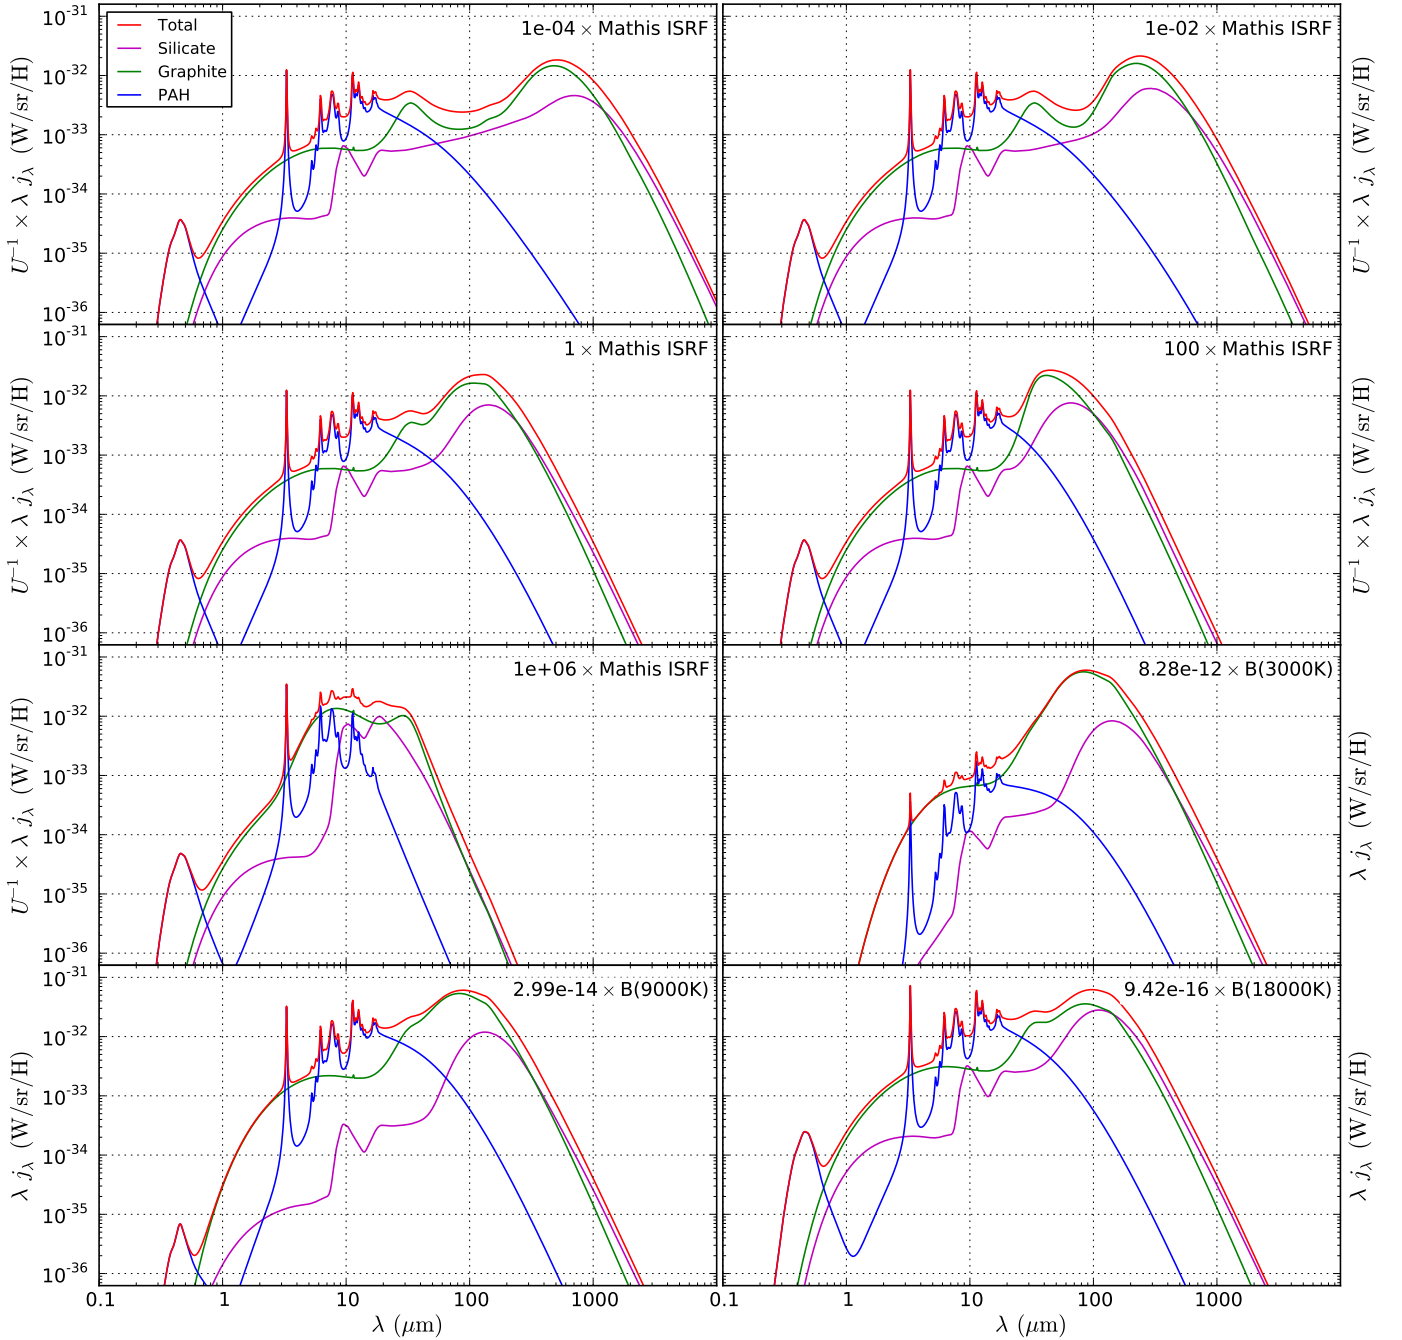
Fig. 1. Reference solutions generated with the public version of DustEM (see Sect. 5.1) using 3500 temperature bins and 250 iterations in the integral equation solver. The panels show the calculated dust emissivity for a selection of the input fields defined in Sect. 3. In each panel, the red curve represents the total emissivity and the other curves represent the portion of the emissivity for each grain type, silicate (magenta), graphite (green), and PAHs (blue). For the scaled Mathis input fields, the emissivity is divided by the input field strength U to allow identical axis ranges for all plots.

converge to the corresponding equilibrium emissivities. Figure 3 shows this comparison for $0.05\,\mu$m grains exposed to radiation fields ranging from extremely weak (left) to strong (right). For a strong field, where we expect the grains to be in equilibrium, the DustEM solutions indeed match the equilibrium emissivities. For a weak field, the solutions differ since the grains are no longer completely in equilibrium. This shows the importance of performing the full stochastic calculation in the presence of extremely weak fields, even for grain sizes up to $0.05\,\mu$m. Codes transitioning from stochastic to equilibrium calculation at a fixed grain size should thus set a sufficiently high value of $a_{\rm trans}$ (see Sects. 4.3 and 5.4).

6.3. Benchmark solutions

Figure 4 compares the total emissivities calculated by each of the codes participating in this benchmark to the corresponding reference solutions for a selection of the input fields defined in Sect. 3. Subsequent figures zoom in on the emissivities for each dust component separately: silicate (Fig. 5), graphite (Fig. 6), and PAH (Fig. 7) grains.

In these figures we limited the displayed wavelength range to $1\,\mu$m $\leq \lambda \leq 1000\,\mu$m. Outside of this range, other sources or processes usually dominate the radiation emanating from astrophysical objects, so it is not relevant to evaluate
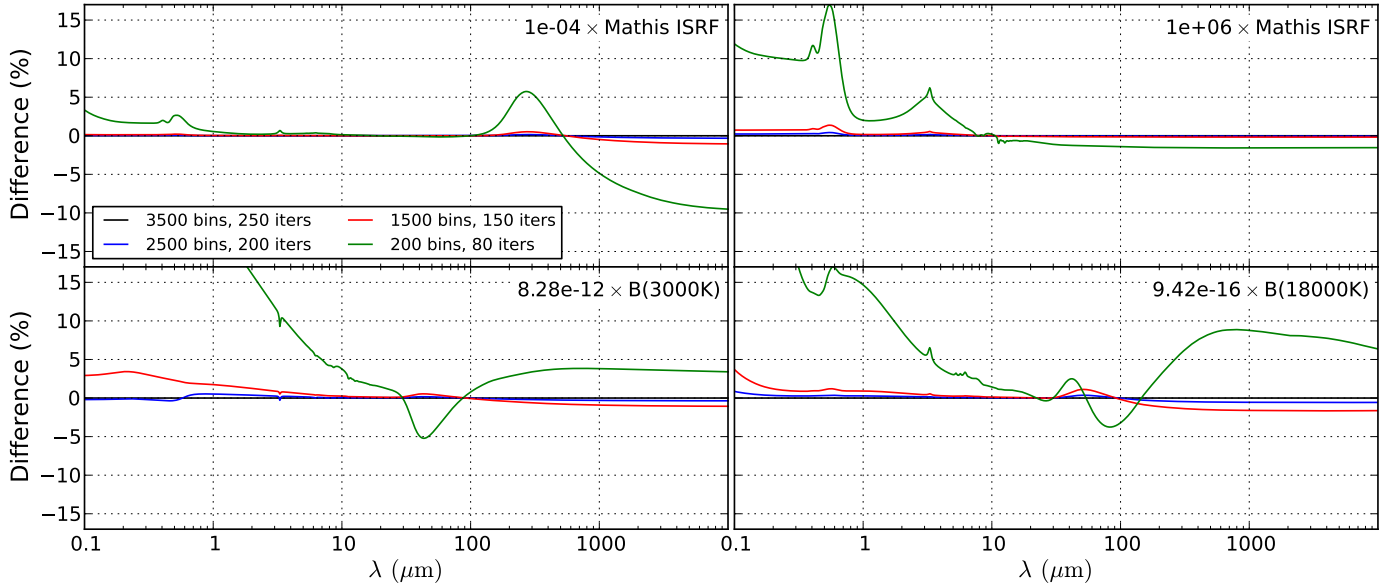
Fig. 2. Comparison of DustEM solutions for the most extreme input fields defined in Sect. 3, calculated with a varying number of temperature bins and iterations in the DustEM integral equation solver. The solutions employed as a reference for our benchmark are calculated with 3500 temperature bins and 250 iterations; these solutions are represented in this figure by the zero lines. The solutions calculated with the standard DustEM values of 200 temperature bins and 80 iterations are represented by the green curve. For these extreme fields, the standard solution deviates by up to 20% (and even more for wavelengths shorter than $1\,\mu m$). The solutions using 2500 temperature bins and 200 iterations (the blue curve) differ by less than 1% from the reference solution, indicating numerical convergence at these parameter values. The contribution of each grain type separately has a similar convergence behavior (not shown).
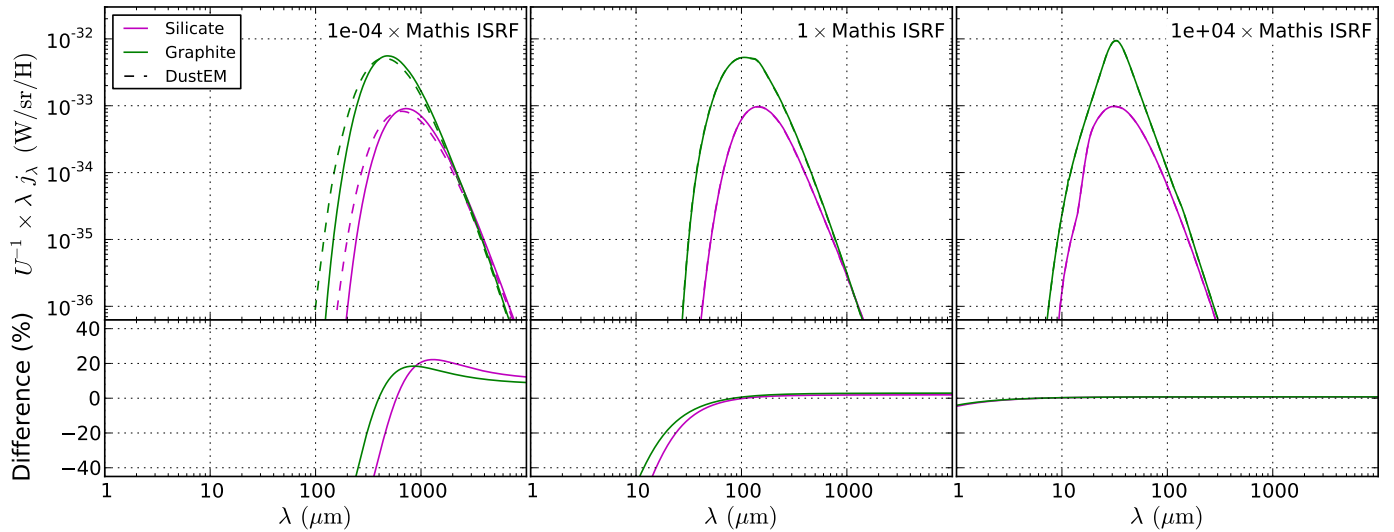
Fig. 3. Comparison of the emissivities calculated by DustEM (using 3500 temperature bins and 250 iterations) for single-size, near-LTE grain populations to the corresponding equilibrium emissivities. The panels show the comparison for input fields ranging from extremely weak (*left*) to strong (*right*). The emissivity is divided by the input field strength U to allow identical axis ranges for all plots. We used a dust mixture consisting of $0.05\,\mu m$ silicate (magenta) or graphite (green) grains with a total dust mass per hydrogen atom of 10^{-30} kg/H for each grain type. The solid curves represent the emissivities calculated by one of our codes (SKIRT) under the assumption of LTE. The dashed curves represent the solutions calculated by DustEM without any LTE assumptions. The *lower panels* show the deviation of the equilibrium solutions from the corresponding full solutions. In a strong field, where we expect the grains to be in equilibrium, the solutions are indeed virtually identical. In a weaker field, the solutions differ since the grains are no longer completely in equilibrium.

the results of a dust emissivity calculation in that spectral range. Moreover, some of the assumptions underlying the computations are no longer valid, rendering the results physically meaningless.

First considering the shorter wavelength range, a black body with peak emission at $\lambda = 1\,\mu m$ has a temperature of $T \approx 2900\,K$. The sublimation temperature of a dust grain is estimated at 1200 K for silicates and at 2100 K for graphites

(Kobayashi et al. 2009). Evaporation rates rise roughly exponentially with increasing grain temperature (Guhathakurta & Draine 1989; Kobayashi et al. 2009), i.e. with decreasing emission wavelength. It is clear that a relevant portion of the dust that would emit at $\lambda \lesssim 1\,\mu m$ is destroyed by evaporation, so that the grain size distribution in our dust model is no longer valid under these conditions. Consequently, the calculation would substantially overestimate the resulting dust emission.
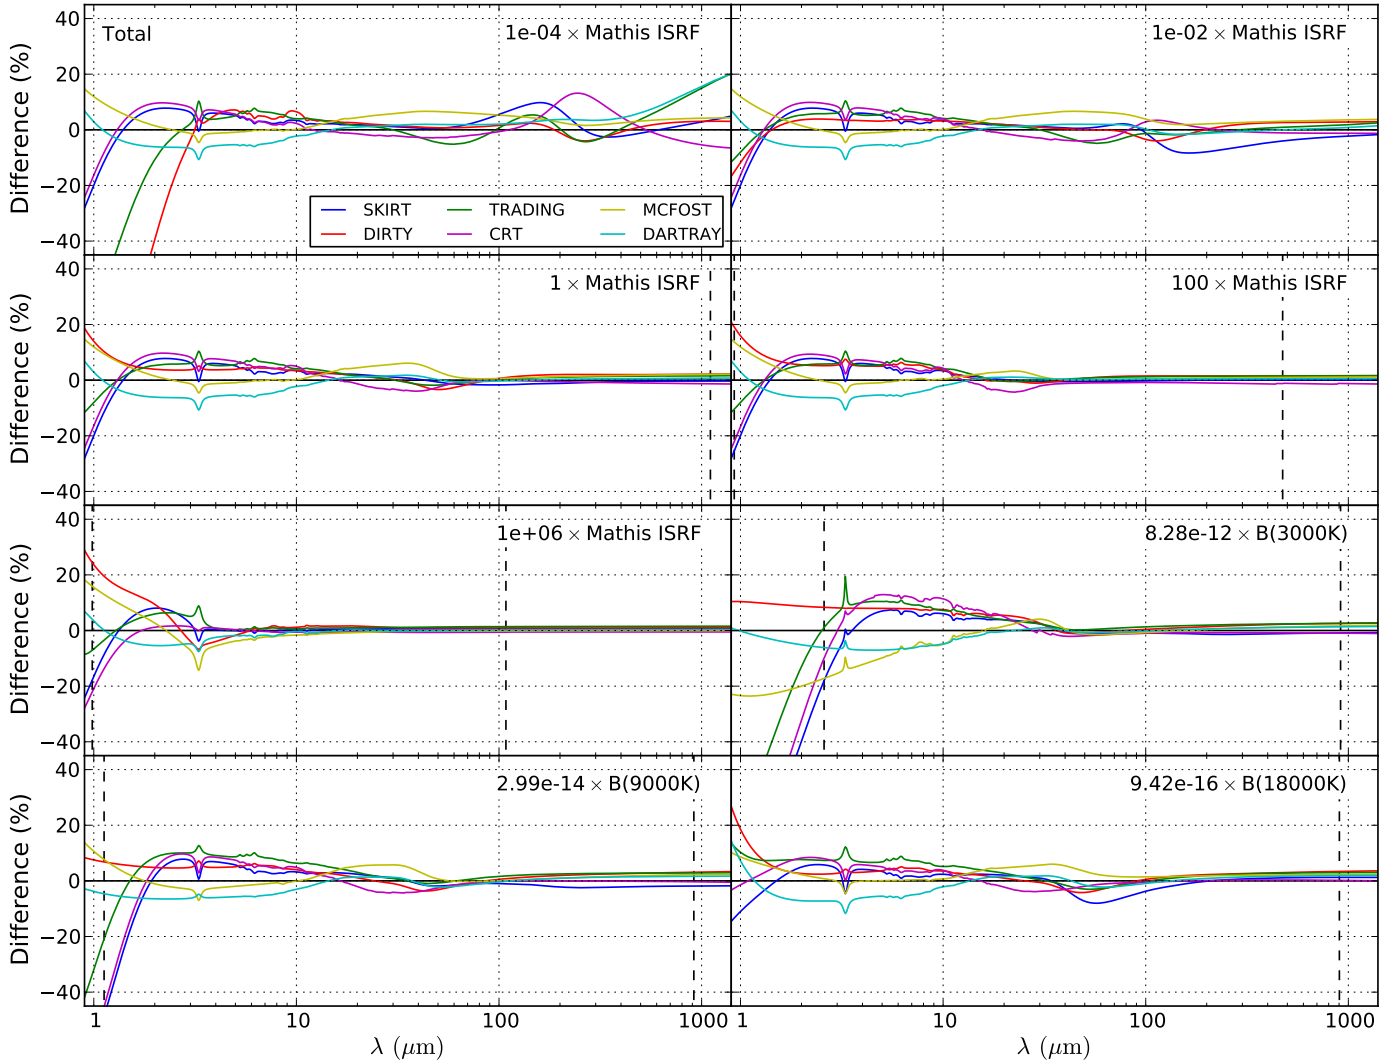
Fig. 4. Relative differences between the total emissivities calculated by each of the codes participating in this benchmark and the corresponding reference solutions. The panels show the results for a selection of the input fields defined in Sect. 3. In each panel, the reference solution is represented by the zero line. Positive percentages indicate results above the reference solution. The vertical dashed lines indicate where the reference solution becomes three orders of magnitude smaller than its peak value.

We now consider the longer wavelength range. For a sufficiently strong input field, say $U \gtrsim 1$, we can expect the calculated emissivity results to be correct because most of the grains emitting at wavelengths longer than $1000\,\mu$m are in LTE. However, the emissivity peaks at much shorter wavelengths, so that the level at $1000\,\mu$m is already several orders of magnitude below the peak level (see all panels in Fig. 1 except for the first two). The situation is different for extremely weak input fields approaching the level of the cosmic microwave background (CMB). A black body with peak emission at $\lambda = 1000\,\mu$m has a temperature of $T \approx 2.9$ K, just above the temperature of the CMB. For small dust grains at such low energies, the continuous cooling assumption no longer holds,[4] meaning that the method presented in Sect. 4.3 does not necessarily yield correct results. In conclusion, the dust emission at wavelengths $\lambda \gtrsim 1000\,\mu$m is either calculated assuming equilibrium conditions (which renders comparison uninteresting) or calculated improperly (which renders comparison meaningless).

6.4. Evaluation of benchmark results

In the wavelength range $3\,\mu\text{m} \le \lambda \le 1000\,\mu\text{m}$, all participating codes reproduce the total dust emissivity within 20% of the reference solution for all input fields used in this benchmark (see Fig. 4). Excluding the weakest ($U \lesssim 10^{-4}$) and the softest ($T \lesssim 3000$ K) fields, the correspondence in the same wavelength range is within 10%. The larger relative deviations at wavelengths shorter than $3\,\mu$m are caused in part by the much lower absolute emissivity values in that range (two to three orders of magnitude below peak values; see Figs. 1 and 4).

The emissivities calculated for the silicate and graphite components (Figs. 5 and 6) show a similar pattern, although the deviations in the individual components are sometimes slightly larger. The emissivities calculated for the PAH component (Fig. 7) show the largest deviations. If we restrict the analysis to the wavelength range in which the emissivity of the reference solution is within three orders of magnitude of its peak value, the correspondence is still within 40% for all codes and for all input fields, and often a lot better.

The larger discrepancies between the various codes for the PAH component can be traced to the PAH molecules generally
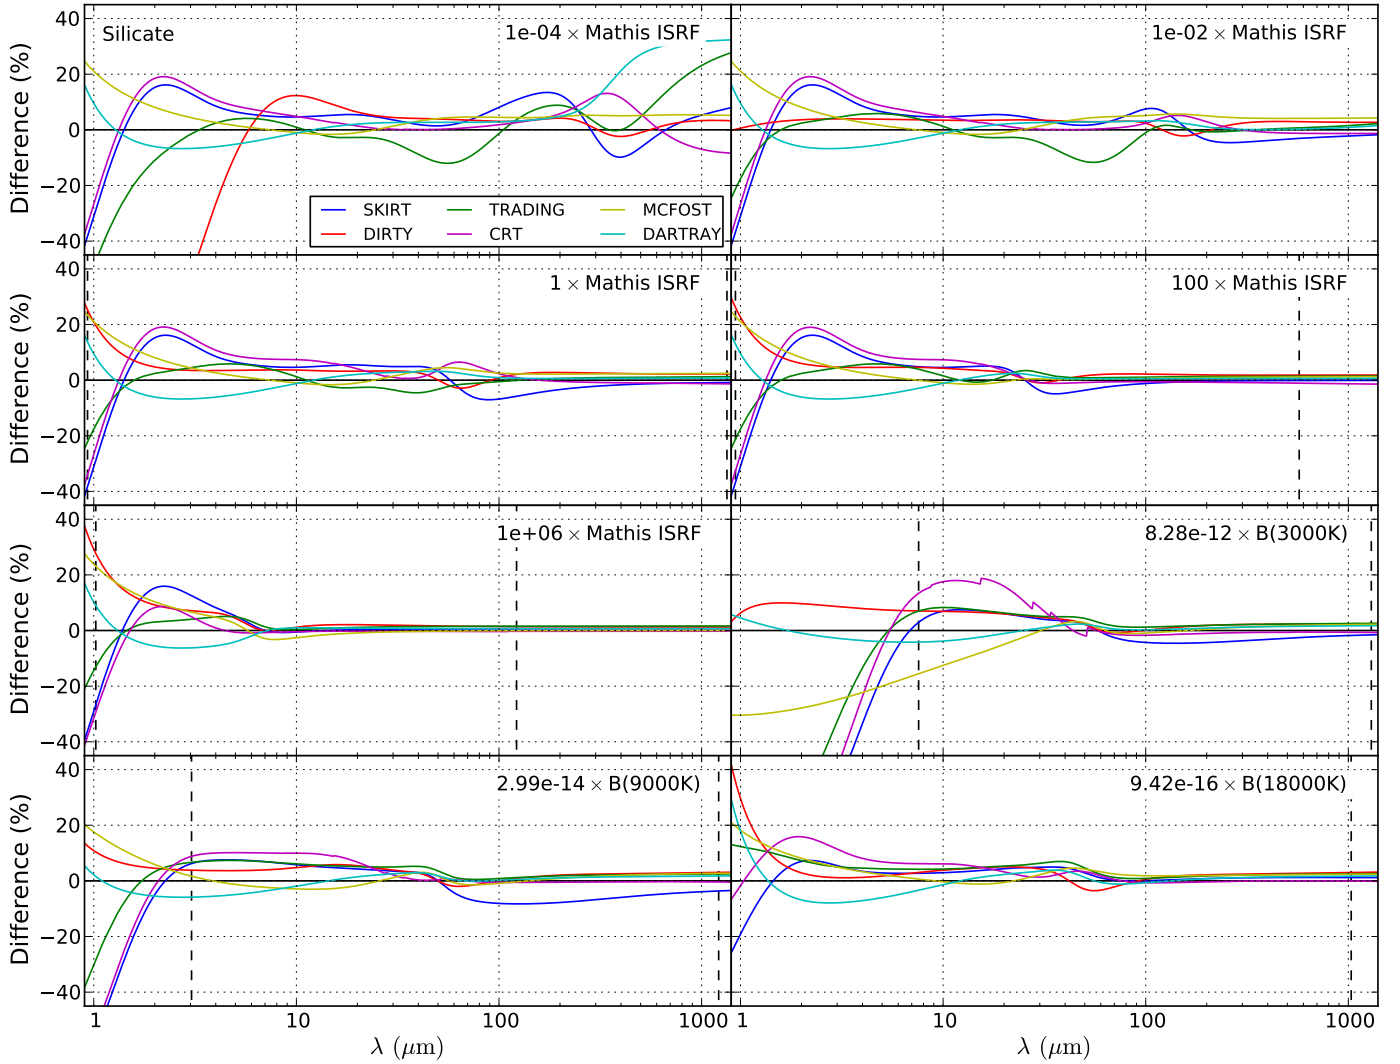
---

[4] With the properties of our dust model, the internal energy of a $0.007\,\mu$m graphite grain at 2.9 K is insufficient to emit a $1000\,\mu$m photon.

Fig. 5. Relative differences between the emissivities of the silicate component calculated by each of the codes participating in this benchmark and the corresponding reference solutions. The panels show the results for a selection of the input fields defined in Sect. 3. In each panel, the reference solution is represented by the zero line. Positive percentages indicate results above the reference solution. The vertical dashed lines indicate where the reference solution becomes three orders of magnitude smaller than its peak value.

being substantially smaller than silicate and graphite grains: see the a^{max} values in Table 2. As noted in Sect. 4.3 and further discussed in Sect. 6.5, smaller grains are more likely to remain in the stochastic regime, requiring complex calculations that are more sensitive to differences in discretization (choice of grids) and concrete implementation (even if the same overall method is employed), as compared to equilibrium calculations.

Specifically, the PAH emissivities calculated by CRT (magenta curves in Fig. 7) deviate from the other codes because, as described in Sect. 5.5, CRT implements the Draine & Li (2001) equations rather than Eqs. (23) and (25) for calculating the heating and cooling transition rates. This method does not allow the emission of photons with an energy higher than the enthalpy content of the emitting grain (see Eq. (56) in Draine & Li 2001). The upward jumps in the emission spectrum appear when, going toward longer wavelengths (lower photon energy), a new enthalpy bin enters the emission calculation (see also Figs. 14 and 15 in Draine & Li 2001). These discontinuities appear only in the wavelength range where emission is largely from grains with very low enthalpy, for which the continuous cooling approximation is not valid, as described in Sect. 6.3.

Figure 8 shows the contributions to the total emissivity calculated by one of our codes (DIRTY) in the stochastic and equilibrium regimes for each of the three grain types and for a number of input fields[5]. While the silicate and graphite components show a significant equilibrium contribution for all fields, the PAHs remain in the stochastic regime for all but the strongest fields.

With respect to our conclusion at the end of Sect. 6.3, it is worth noting in Fig. 8 that, at $\lambda = 1000\,\mu\mathrm{m}$, the equilibrium contributions of the silicate and graphite components dominate the total stochastic contribution for all input fields. Only for the weakest fields ($U \approx 10^{-4}$) does the stochastic contribution at that wavelength become a noticeable fraction of the total emission.

Considering the contributions of the different grain types to the total spectrum (Fig. 1), it turns out that, for our dust model, the graphite emission dominates in most of the wavelength range for most input fields. The PAHs dominate in a very small region (their peak) and the silicates only at the longest wavelengths. Since the agreement between the different codes is significantly

---

[5] The precise form of these respective contributions varies between codes because of differences in transitioning from one regime to the other, but the general trend is similar.
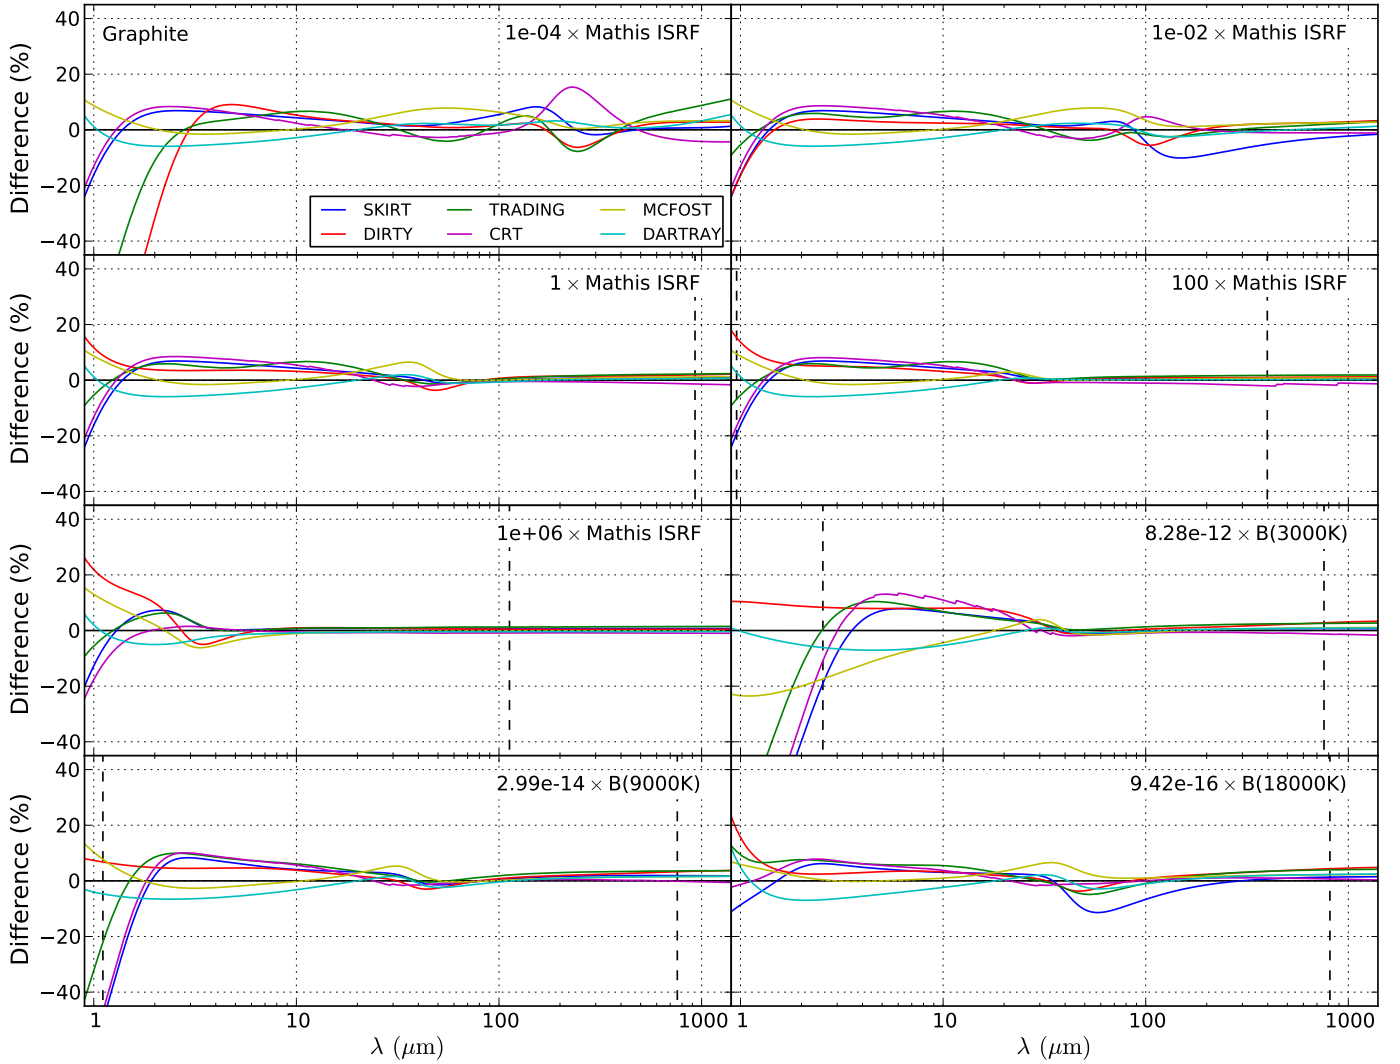
Fig. 6. Relative differences between the emissivities of the graphite component calculated by each of the codes participating in this benchmark and the corresponding reference solutions. The panels show the results for a selection of the input fields defined in Sect. 3. In each panel, the reference solution is represented by the zero line. Positive percentages indicate results above the reference solution. The vertical dashed lines indicate where the reference solution becomes three orders of magnitude smaller than its peak value.

better for graphite and silicates than for the PAHs, our choice of dust mixture (unintentionally) benefits the global agreement between the different codes.

6.5. Transition to equilibrium

As introduced in Sect. 4.3, and further elaborated upon in the code descriptions in Sect. 5, each code handles the transition from the stochastic to the equilibrium calculation regime in its own way. Figure 9 shows the grain size a_{trans} for which the participating codes transition from the stochastic to the equilibrium calculation regime, for each grain type, and for each of the input fields defined in Sect. 3. While the details differ between codes, as a general trend, small grains (e.g. $a < 0.01\,\mu$m) are considered to be in equilibrium only when exposed to the strongest fields ($U > 10^2$).

Because most of the PAHs remain in the stochastic regime (Fig. 8 and right panel of Fig. 9), the transition differences are most easily seen in the results for the silicate and graphite grains (Figs. 5 and 6). As a general trend, the results converge for longer wavelengths, because the equilibrium emission dominating in that range is calculated in the same straightforward manner

across all codes. The larger discrepancies are found at shorter wavelengths, where the stochastic regime dominates. Depending on the input field, the transition point shifts on the wavelength scale. Interestingly, for some fields, the discrepancies show extra "wiggles" near the transition points, most likely caused by differences in handling the transition. This is particularly evident in the silicate emissivities for the scaled Mathis fields with strengths $U = 10^{-2}$ to 10^2 (see the top half of Fig. 5).

The total emissivities are influenced by these transition differences mostly in the wavelength range just shortward of the large sub-mm emission peak, which is dominated by LTE emission from silicate and graphite grains (see Fig. 4). The position of this peak is determined by the temperatures of the dust grains, and thus depends on the input field. For the strongest field in our benchmark ($U = 10^6$), the broad peak even overlaps the PAH features in the 3 to 30 μm wavelength range (see the third panel in the leftmost column of Fig. 4).

6.6. Weak fields

As discussed in Sect. 4.3, the transition size a_{trans} above which a grain can be considered to be in LTE depends on the radiation
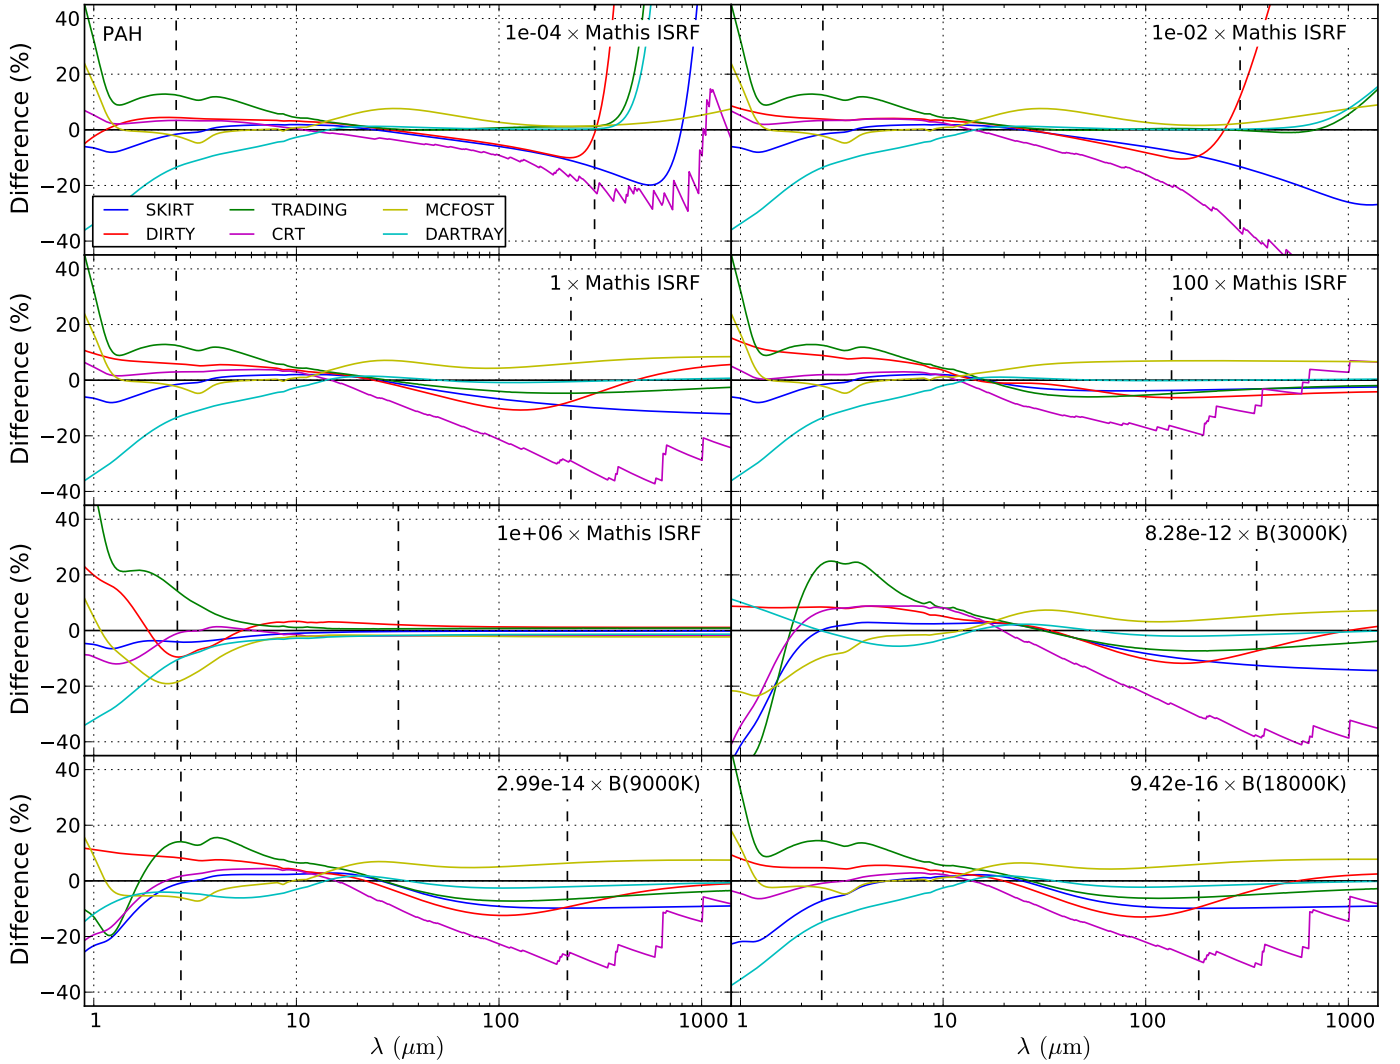
Fig. 7. Relative differences between the emissivities of the PAH component calculated by each of the codes participating in this benchmark and the corresponding reference solutions. The panels show the results for a selection of the input fields defined in Sect. 3. In each panel, the reference solution is represented by the zero line. Positive percentages indicate results above the reference solution. The vertical dashed lines indicate where the reference solution becomes three orders of magnitude smaller than its peak value.

field to which the grain is exposed. For a given grain composition, a_{trans} tends to be higher for weaker fields because photon interactions are less frequent, which keeps larger grains in the stochastic regime. Consequently, the dust emission calculations are more complex for weaker fields, especially for small grains in weaker fields. In addition to the computation time, this complexity affects the accuracy of the results, which explains why the largest discrepancies between the various solutions occur for the weakest field in the benchmark (see the top left panels in Figs. 4 through 7).

In fact, the weakest field in our benchmark ($U = 10^{-4}$) may be unrealistically weak, since its peak intensity is below the peak of the CMB, albeit in a different wavelength regime (UV versus mm wavelengths). To evaluate the effect of neglecting the CMB, we added the CMB to the $10^{-4} \times J^{Mat}$ field and used DustEM to recalculate the emissivity of our dust model exposed to this new input field. For wavelengths $\lambda \leq 100\,\mu m$, the results are essentially identical to those shown in the top lefthand panel of Fig. 1. The sub-mm peak is a notch higher and slightly shifted to longer wavelengths, causing an emissivity increase of about 35% at $\lambda = 1000\,\mu m$. While this effect may not be negligible, it does not invalidate the benchmark test.

However, as argued in Sect. 6.3, our computations may no longer be physically founded for these weak fields, especially for small grains with internal energies comparable to those of the CMB photons. In conclusion, the $10^{-4} \times J^{Mat}$ input field is benchmarking the various codes properly, but the calculations may be collectively incorrect because the continuous cooling approximation is inappropriate in this regime.

6.7. Temperature discretization

The participating codes implement various ways to discretize the grain temperature (or equivalently, the grain enthalpy), as described in Sect. 5. The different schemes are mostly driven by the aim to increase performance while preserving accuracy. Here we discuss the impact of the minimum and maximum temperature values allowed on the grid.

DustEM, which was used to calculate our reference solutions, does not impose a lower temperature limit other than the zero point. Indeed, under the continuous cooling assumption, a small dust grain does not have to be in equilibrium with the CMB, and thus there seems to be no reason the grain should not
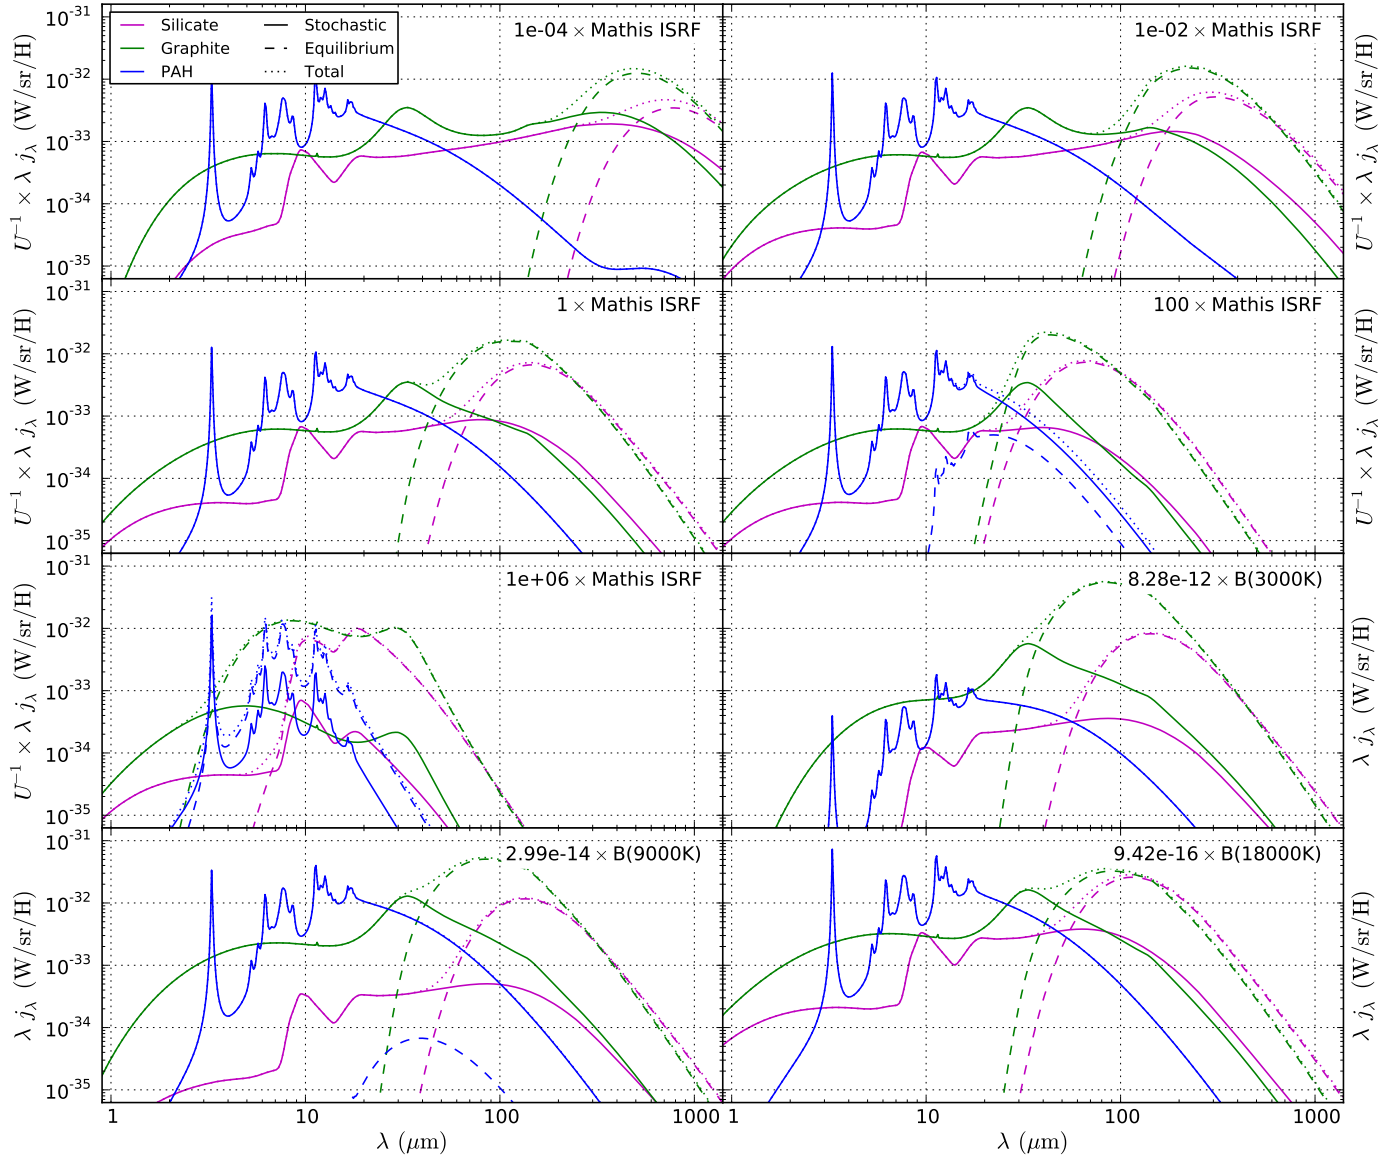
Fig. 8. Contributions calculated in the stochastic (solid) and equilibrium (dashed) regimes to the total emissivity (dotted) for each of the grain types silicate (magenta), graphite (green), and PAHs (blue), by one of our codes (DIRTY). The panels show the results for a selection of the input fields defined in Sect. 3. For the scaled Mathis input fields, the emissivity is divided by the input field strength U to allow identical axis ranges for all plots. The relative contributions for each regime may vary between codes because of the differences in the schemes to transition from one regime to the other.
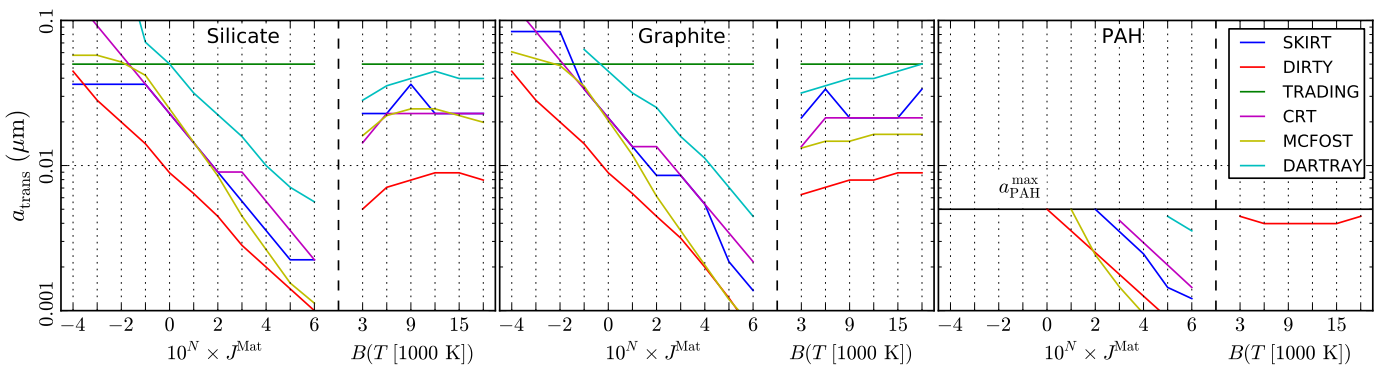
Fig. 9. Grain size a_{trans} for which the participating codes transition from the stochastic to the equilibrium calculation regime. The panels show the smallest grain size that was considered to be in equilibrium for silicate (*left*), graphite (*middle*), and PAH (*right*) grains, for each of the codes, across all of the input fields defined in Sect. 3 (within each panel, the scaled Mathis ISRF fields to the left, and the diluted black body fields to the right).

have, at any given moment in time, a temperature below 2.73 K. In other words, we need to calculate the temperature probability distribution as usual. As argued in Sects. 6.3 and 6.6, this line of reasoning breaks down for small grains at very low energies, since the continuous cooling assumption no longer holds.

This is why most codes participating in this benchmark impose a lower temperature limit of 2.73 K, or even rather arbitrarily, 1, 2, or 3 K. A limit of 2.73 K, for example, causes a bump in the PAH emission peaking at $\lambda \approx 1060\,\mu m$ (the peak CMB wavelength) because all the probabilities for lower temperatures are bunched together in the 2.73 K temperature bin. This effect is to some extent responsible for the discrepancies between the codes and the reference solution seen in the top lefthand panel of Fig. 7, in the wavelength range to the right of the dashed line, where the absolute value of the emissivity has become small anyway. The effect is negligible for all but the weakest fields.

At the other end of the scale, all codes in this benchmark use the complete temperature range for which the dust properties are defined, i.e. up to 2500 K for our dust model. As described in Sect. 6.3, this is well above the sublimation temperature of the dust material, although a fraction of the grains may survive at these temperatures for some time. Because the method used for this benchmark ignores dust grain destruction, we expect it to overestimate the emissivity for shorter wavelengths. To evaluate this effect, we reran the benchmark calculations with one of the codes (SKIRT) using a maximum grid temperature of 2250 K instead of 2500 K. As expected, the emissivities in the range $\lambda > 3\,\mu m$ are essentially unaffected by this change. At wavelengths shorter than $3\,\mu m$, the total emissivity for the hardest black body input fields ($T \gtrsim 15\,000$ K) decreases noticeably (by about 30% at $1\,\mu m$), while there is no perceptible change down to $1\,\mu m$ for the softer fields or for the scaled Mathis fields. The silicate and graphite components behave similarly. Interestingly, the emissivity of the PAH component shows a noticeable decrease for all but the softest black body fields ($T \lesssim 6000$ K) and for all scaled Mathis fields. This is can be understood by recalling that the PAH particles are, on average, a lot smaller than those in the other components, so that they are more easily propelled to higher temperatures.

6.8. Wavelength discretization

The method described in Sect. 4 for calculating dust emission involves wavelength discretization in several distinct areas. The optical dust properties are tabulated on some predefined wavelength grid. We also need to configure the wavelength range and sampling resolution for the input fields and for the output emissivity. And finally, the calculation of the cooling coefficients defined in Eq. (25) requires an integration of the optical properties over wavelength, which implies a grid as well. As long as proper interpolation procedures are in place, these wavelength grids do not need to be identical.

To keep matters simple for the benchmarks presented in this paper, Sect. 3.3 specifies the same wavelength grid for the optical properties as for the calculated emissivities, i.e., a logarithmic grid with 1201 points in the range $0.001\,\mu m \leq \lambda \leq 10\,000\,\mu m$. In this section we discuss the impact of the resolution and of the lower and upper limits of the wavelength grid on the emissivity calculations. We reran the benchmark calculations with one of the codes (SKIRT) using some wavelength grid variations as reported below. SKIRT employs a single (configurable) wavelength grid for all aspects of the calculation. The optical dust properties are interpolated to this grid, and the grid is

subsequently used for the input and output fields and for the integration to obtain the cooling coefficients.

In practice, the lower wavelength limit is determined by the need to properly capture the input field in calculating the heating coefficients (Eq. (23)) and the equilibrium temperature (Eq. (22)). For a scaled Mathis field, the lower limit can be increased to $0.09\,\mu m$ (see Eq. (11)). For the hardest black body fields ($T \gtrsim 12\,000$ K), the limit should be lower. With a lower limit of $0.01\,\mu m$, there is no noticeable difference in the calculated emissivities even for the hardest field in this benchmark ($T = 18\,000$ K).

The upper wavelength limit is mostly determined by the need to calculate the emissivity up to mm wavelengths. The upper limit also affects the calculation of the heating coefficients and the equilibrium temperature, but this effect is smaller because the absorption coefficients for the grain material are much lower at longer wavelengths. With an upper limit of $2000\,\mu m$ instead of $10\,000\,\mu m$, there is no noticeable difference in the calculated emissivities for any of the input fields. With an upper limit of $1000\,\mu m$, the sub-mm emissivity peak is overestimated by 20% for the weakest scaled Mathis field ($U = 10^{-4}$). The emissivity peak for the $U = 10^{-3}$ field shows a similar but much smaller effect. For all other input fields there is no noticeable difference.

Finally, we reran the benchmark calculations for logarithmically distributed wavelength grids with successively lower resolution, always using a range of $0.01\,\mu m \leq \lambda \leq 2000\,\mu m$. First of all, lowering the wavelength resolution affects the shape of the sharp PAH-dominated emissivity peaks in the 3 to $30\,\mu m$ wavelength range; if there is no wavelength point at the center of a peak, the peak cannot be resolved. However, this does not affect the accuracy of the emissivity at other wavelength points unless the resolution becomes too low, as described in what follows. Other than this peak resolution effect, using 601 wavelength points instead of 1201 does not noticeably influence the results for any of the input fields. Lowering the resolution to 301 wavelength points causes minor deviations in the calculated PAH emissivities for wavelengths $\lambda < 3\,\mu m$, which however do not noticeably affect the total emissivities down to $\lambda \geq 1\,\mu m$. With only 151 wavelength points, the deviation in the total emissivities is still limited (a few percent at $1\,\mu m$), but the PAH features are now clearly under-resolved and fairly smoothed out. This can be improved by concentrating more grid points in the wavelength range of the PAH features. For example, a specialty grid with a total of 151 points, 61 of which are concentrated in the 3 to $30\,\mu m$ wavelength range, seems to provide an acceptable compromise.

6.9. Grain size discretization

The total calculation time for the emissivity of a dust population is roughly proportional to the number of grain size bins used to represent the population (see Sect. 4.1). Therefore, some of the participating codes recompute the optical grain properties on an internally defined grain size grid rather than using the size bins tabulated in the dust model data. We used two of the codes (SKIRT and CRT) to investigate the effect of the number of grain size bins on the calculated emission. As mentioned in Sects. 5.2 and 5.5, these codes were configured with 15 size bins per grain type to calculate the benchmark results presented in Figs. 4 through 7.

The predominant effect of changing the number of grain size bins appears for wavelengths $\lambda < 10\,\mu m$ and increases for shorter wavelengths. This is to be expected, because with a coarse grid, the effective size of grains in the smallest bin

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

A secondary effect occurs for longer wavelengths because the rebinning influences the heuristic for transitioning between the stochastic and equilibrium calculation regimes (Sect. 6.5). This effect seems to be somewhat random in nature, causing deviations that remain within the accuracy limits described in Sect. 6.4.

6.10. Calculation time

A typical 3D RT simulation calculates the dust emission spectra for a large number of dust cells. In cases where dust self-absorption is a relevant factor, this calculation is repeated for each iteration of the loop that self-consistently determines dust heating and re-emission (see, e.g., Sect. 5.6). The time spent on calculating dust emission might thus become a significant or even dominant fraction of the total RT simulation time. The aim of reducing the dust emission calculation time has guided many of the choices in the implementations of the RT codes participating in this benchmark. Most fundamentally, all codes adopt the continuous cooling approximation. In addition, all codes select specific discretization schemes, most codes employ heuristics to transition between stochastic and equilibrium regimes, and some precompute field-independent data. Often these choices affect not only the calculation time, but also the accuracy of the results. In principle at least, the results can be made to match perfectly by increasing grid sizes and removing the heuristics, at the expense of calculation time.

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

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

7. Conclusions

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

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

We reported in detail on the similarities and differences between the results from the participating codes and a reference solution. In the important wavelength range $3\,\mu m \leq \lambda \leq 1000\,\mu m$, all participating codes reproduce the total dust emissivity within 20% of the reference solution for all input fields used in this benchmark. Excluding the weakest and the softest input fields, the agreement in the same wavelength range is within 10%.

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

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

References

Aniano, G., Draine, B. T., Calzetti, D., et al. 2012, ApJ, 756, 138
Baes, M., Davies, J. I., Dejonghe, H., et al. 2003, MNRAS, 343, 1081
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
Boulanger, F., & Perault, M. 1988, ApJ, 330, 964
Camps, P., & Baes, M. 2015, Astron. Comput., 9, 20

Compiègne, M., Verstraete, L., Jones, A., et al. 2011, A&A, 525, A103
Desert, F. X., Boulanger, F., & Shore, S. N. 1986, A&A, 160, 295
Draine, B. T. 1988, ApJ, 333, 848
Draine, B. T. 2003, ApJ, 598, 1026
Draine, B. T., & Lee, H. M. 1984, ApJ, 285, 89
Draine, B. T., & Li, A. 2001, ApJ, 551, 807
Draine, B. T., & Li, A. 2007, ApJ, 657, 810
Draine, B. T., Dale, D. A., Bendo, G., et al. 2007, ApJ, 663, 866
Dwek, E. 1986, ApJ, 302, 363
Fischera, J. 2000, Ph.D. Thesis, Heidelberg
Gordon, K. D., Misselt, K. A., Witt, A. N., & Clayton, G. C. 2001, ApJ, 551, 269
Guhathakurta, P., & Draine, B. T. 1989, ApJ, 345, 230
Helou, G., Lu, N. Y., Werner, M. W., Malhotra, S., & Silbermann, N. 2000, ApJ, 532, L21
Hunt, L. K., Draine, B. T., Bianchi, S., et al. 2015, A&A, 576, A33
Jones, A. P., Fanciullo, L., Köhler, M., et al. 2013, A&A, 558, A62
Juvela, M. 2005, A&A, 440, 531
Juvela, M., & Padoan, P. 2003, A&A, 397, 201
Kobayashi, H., Watanabe, S.-I., Kimura, H., & Yamamoto, T. 2009, Icarus, 201, 395
Laor, A., & Draine, B. T. 1993, ApJ, 402, 441
Li, A., & Draine, B. T. 2001, ApJ, 554, 778
Lunttila, T., & Juvela, M. 2012, A&A, 544, A52

Manske, V., & Henning, T. 1998, A&A, 337, 85
Manske, V., & Henning, T. 1999, A&A, 349, 907
Mathis, J. S., Mezger, P. G., & Panagia, N. 1983, A&A, 128, 212
Misselt, K. A., Gordon, K. D., Clayton, G. C., & Wolff, M. J. 2001, ApJ, 551, 277
Natale, G., Popescu, C. C., Tuffs, R. J., & Semionov, D. 2014, MNRAS, 438, 3137
Natale, G., Popescu, C. C., Tuffs, R. J., et al. 2015, MNRAS, 449, 243
Pinte, C., Ménard, F., Duchêne, G., & Bastien, P. 2006, A&A, 459, 797
Pinte, C., Harries, T. J., Min, M., et al. 2009, A&A, 498, 967
Popescu, C. C., Misiriotis, A., Kylafis, N. D., Tuffs, R. J., & Fischera, J. 2000, A&A, 362, 138
Popescu, C. C., Tuffs, R. J., Dopita, M. A., et al. 2011, A&A, 527, A109
Sellgren, K. 1984, ApJ, 277, 623
Siebenmorgen, R., Kruegel, E., & Mathis, J. S. 1992, A&A, 266, 501
Smith, J. D. T., Draine, B. T., Dale, D. A., et al. 2007, ApJ, 656, 770
Steinacker, J., Baes, M., & Gordon, K. D. 2013, ARA&A, 51, 63
Voit, G. M. 1991, ApJ, 379, 122
Weingartner, J. C., & Draine, B. T. 2001a, ApJ, 548, 296
Weingartner, J. C., & Draine, B. T. 2001b, ApJS, 134, 263
Whitney, B. A. 2011, Bull. Astron. Soc. India, 39, 101
Wolf, S. 2003, ApJ, 582, 859
Ysard, N., Juvela, M., & Verstraete, L. 2011, A&A, 535, A89
Zubko, V., Dwek, E., & Arendt, R. G. 2004, ApJS, 152, 211