An algorithm for Monte Carlo time-dependent radiation transfer

Tim J. Harries[★]

School of Physics, University of Exeter, Stocker Road, Exeter EX4 4QL

ABSTRACT

A new Monte Carlo algorithm for calculating time-dependent radiative transfer under the assumption of local thermodynamic equilibrium is presented. Unlike flux-limited diffusion, the method is polychromatic, includes scattering, and is able to treat the optically-thick and free-streaming regimes simultaneously. The algorithm is tested on a variety of 1D and 2D problems, and good agreement with benchmark solutions is found. The method is used to calculate the time-varying spectral energy distribution from a circumstellar disc illuminated by a protostar whose accretion luminosity is varying. It is shown that the time-lag between the optical variability and the infrared variability results from a combination of the photon traveltime and the thermal response in the disc, and that the lag is an approximately linear function of wavelength.

Key words: radiative transfer – methods: numerical – protoplanetary discs – stars: pre-main-sequence.

1 INTRODUCTION

Theoretical astrophysics often necessitates computing the energy balance between gas and an ambient radiation field. Most frequently this involves solving the equation of radiative equilibrium, and a wide variety of theoretical tools have been developed to tackle this problem efficiently, ranging from those which make simplifying assumptions both to the geometry (spherical or plane-parallel) and to the underlying microphysics (e.g. the grey approximation) through to a full polychromatic treatment in multiple dimensions.

Although most radiation transfer (RT) codes assume thermal balance, many interesting phenomena occur out of equilibrium, and here a time-dependent approach is required. A spectacular example of this occurs during supernova explosions, and progress is being made in time-dependent modelling of such objects (see e.g. Jack et al. 2009; Kromer & Sim 2009).

Time-dependent methods are also required for radiation hydrodynamics (RHD). Broadly speaking, the transport methods can be split into those concerned with ionizing radiation, which generally adopt ray-casting methods (e.g. Dale, Ercolano & Clarke 2007; Mac Low et al. 2007), and those that deal with transport in dusty media which use either short-characteristic methods, Monte Carlo (MC) techniques or the diffusion approximation (e.g. Turner & Stone 2001; Höfner et al. 2003; Woitke 2006). RHD codes may be further subdivided into those which are applicable to situations in which the radiation transport/thermal equilibrium time-scale for the gas is shorter than the characteristic hydrodynamical time-scale and those in which the hydrodynamical time-scale is shorter. In the former

case, the radiation transport may be conducted as a sequence of pseudo-equilibrium steps (e.g. Höfner et al. 2003; Woitke 2006; Dale, Ercolano & Clarke 2007; Mac Low et al. 2007; Freytag & Höfner 2008; Acreman, Harries & Rundle 2010). However, when the thermal time-scale becomes comparable to, or indeed larger than, the hydrodynamical time-scale, the time-dependent form of the RT equations must be solved (e.g. Krumholz, Stone & Gardiner 2007).

In many situations, RT is implemented assuming energy transport occurs via radiative diffusion of photons. This is a good assumption for optically-thick regions, but breaks down in the optically-thin regime, since the mean free path of photons (and thus the speed of the diffusing radiation field) can become arbitrarily large; a flux limiter is adopted to control this problem. Flux-limited diffusion (FLD) is normally implemented in the grey approximation (Turner & Stone 2001; Whitehouse & Bate 2004; Whitehouse, Bate & Monaghan 2005; Krumholz et al. 2007), although comparisons with a polychromatic treatment demonstrate that this is often quite a poor approximation (Preibisch, Sonnhalter & Yorke 1995). Opaque obstacles also cause problems for FLD, since the radiation field can diffuse around the obstacle when in fact a shadow should be cast.

A modified version of the MC radiative equilibrium method developed by Lucy (1999) offers an attractive route to full time-dependent RT, since it would allow a full polychromatic treatment within which multiple scattering may be incorporated. The method is computationally demanding, but it has the advantage of being extremely efficient to parallelize. Indeed the feasibility of using such a method has already been demonstrated, at least in the case where the thermal time-scales are short (Acreman et al. 2010).

Here I present an algorithm for calculating the matter and radiation field energy densities, as a function of time, for an arbitrary

[★]E-mail: th@astro.ex.ac.uk

distribution photon of sources embedded in an arbitrary distribution of gas under local thermodynamic equilibrium (LTE) conditions. I describe 1D benchmark tests of the method and detail its incorporation into the TORUS RT code. I test the new version of TORUS against a 2D disc benchmark (Pascucci et al. 2004). Finally, I give a simple application of the code demonstrating the effect of a varying accretion luminosity from a protostar on its surrounding dusty circumstellar disc.

2 METHOD

We describe the physical quantities (densities, temperatures, etc.) on a cell-centred grid. For a gas in LTE at temperature T, the rate at which it emits energy is given by

$$\dot{E} = 4\pi \int_0^\infty k_\nu B_\nu \, d\nu, \tag{1}$$

where k_ν is the absorption coefficient and B_ν is the Planck function. The rate at which the same gas absorbs energy is given by

$$\dot{A} = 4\pi \int_0^\infty k_\nu J_\nu \, d\nu, \tag{2}$$

where J_ν is the mean intensity of the radiation field. Clearly, if the gas is in radiative equilibrium, then $\dot{A} = \dot{E}$ and we find

$$T = \left(\frac{\dot{A}}{4\sigma\kappa_P} \right)^{1/4}, \tag{3}$$

where κ_P is the Planck-mean absorption coefficient. However, if we consider gas that is not in radiative equilibrium, then the nett change in the energy density of the gas

$$\dot{u}_g = \dot{A} - \dot{E} \tag{4}$$

and, *mutatis mutandis*, the rate of change in the energy density of the radiation field is

$$\dot{u}_r = \dot{E} - \dot{A} . \tag{5}$$

Now we consider a gas of volume V at time t. Within this volume is a star of luminosity L_*. The luminosity of the gas is given by

$$L_g = \int_V \dot{E} \, dV . \tag{6}$$

We assume that the temperature of the gas is constant over a single time-step Δt. During this time-step, we assume that the gas and the star produce N_g and N_* new photon packets, respectively. The individual photon packet energies are given by

$$\epsilon_g = \frac{L_g \Delta t}{N_g}, \quad \epsilon_* = \frac{L_* \Delta t}{N_*} . \tag{7}$$

The energy density is related to the temperature by

$$u_g = \frac{RT\rho}{(\gamma - 1)\mu}, \tag{8}$$

where R is the gas constant, ρ is the mass density, γ is the ratio of specific heats and μ is the mean molecular weight. We follow Lucy (1999) and use the result that the energy density of the radiation field in the interval $(\nu, \nu + d\nu)$ is given by

$$u_{r,\nu} = \frac{4\pi J_\nu \, d\nu}{c}. \tag{9}$$

A photon packet moving between events (scatterings, absorptions, crossing grid-cell boundaries) contributes an energy ϵ_ν for time ℓ/c (where ℓ is the distance between events) to the local energy density. The photon energy density is therefore

$$u_r = \frac{1}{\Delta t} \frac{1}{V} \frac{1}{c} \sum \epsilon_\nu \ell , \tag{10}$$

where the summation is over all photon packets. Now combining equations (2) and (9) with equation (10), we obtain an expression for the energy absorption rate:

$$\dot{A} = \frac{1}{V} \frac{1}{\Delta t} \sum k_\nu \epsilon_\nu \ell . \tag{11}$$

The new energy density of the gas may then be calculated as

$$u_g^{n+1} = u_g^n + (\dot{A} - \dot{E})\Delta t . \tag{12}$$

This explicit integration scheme will require a careful choice of the time-step Δt, which must be short enough to ensure stability while also being long enough that the computation remains tractable. Time-scale considerations are discussed in Section 2.1.

It is worth noting that equation (12) is not the only method for updating the internal energy of the gas. Since we follow the radiation field in detail, we can, for each cell, calculate the energy in photon packets entering and leaving the cell, and hence the change in radiation energy density due to RT ($\Delta u_{r,trans}$). MC estimators exist for u_r for both the start and the end of the time-step, so using

$$u_r^{n+1} = u_r^n + (\dot{E} - \dot{A})\Delta t + \Delta u_{r,trans} \tag{13}$$

we can find $(\dot{E} - \dot{A})\Delta t$ and substitute it into equation (12), bypassing the need to find an estimator for $\dot{A}$. In practice, we found that this method required a larger number of photon packets (in order to improve the estimate of u_r and $\Delta u_{r,trans}$) than that needed using an estimator for $\dot{A}$.

A single time-step encompasses of a loop over photon packets, each with an individual energy ϵ_ν and frequency ν. Information on photon packets that are 'in flight' at the end of a time-step is stored on a stack, to be processed as part of the RT during the subsequent time-step. The total number of photon packets is the sum of the number of packets on the stack (N_s) and the number of new photon packets generated during the time-step N_p. The path of each photon packet is followed as it propagates through the gas, during which time the packet may be scattered multiple times. The random walk of the photon packet ceases when it is either absorbed or leaves the computation domain (in which case the photon packet is destroyed), or when its flight time (ℓ/c) reaches Δt (in which case the photon packet is added to the stack). Once the random walks of all the photons have been calculated, we may then use our MC estimate of the absorption rate (equation 11) to update the gas energy density of each cell via equation (12). One can immediately see that when the photon flight time of a mean free path $\kappa\rho/c \gg \Delta t$, then N_s will dominate over N_p as the calculation proceeds, but eventually if radiative equilibrium is reached, then there will be an approximately constant number of packets at each time-step.

It should be noted that, unlike the Lucy (1999) radiative equilibrium algorithm, energy is not implicitly conserved here, as photon packets are created and destroyed during each time-step. The quality of the energy conservation is controlled by the accuracy of our estimator of the absorption rate (which in turn is dependent on our MC statistics) and by our choice of Δt. Energy conservation is addressed in Section 3.

The algorithm itself is described as a sequence of steps:

(i) The current energy density of the gas (u_g) is used to compute the temperature distribution (equation 8) and thus the gas energy emission rate as given by equation (1). A probability distribution of a photon packet being emitted is calculated over all cells:

$$p_i = \frac{\sum_1^i \dot{E}_i V_i}{\sum_1^{N_{cells}} \dot{E}_i V_i}, \tag{14}$$

where V_i is the volume of the ith cell.

(ii) The probability of a photon being produced in the gas is calculated by

$$p_g = \frac{L_g}{L_* + L_g}. \tag{15}$$

We take η to be the uniform random deviate. If $\eta < p_g$, then the photon packet is emitted within the gas, and its position is found according to the probability distribution given by equation (14); otherwise, the photon is produced by the star.

(iii) If a photon is produced by the gas, then the photon frequency is assigned according to the equation

$$\eta = \int_0^\nu j_\nu \, d\nu \Big/ \int_0^\infty j_\nu \, d\nu, \tag{16}$$

where $j_\nu = k_\nu B_\nu$. If the photon packet is stellar in origin, then its frequency comes from

$$\eta = \int_0^\nu B_\nu \, d\nu \Big/ \int_0^\infty B_\nu \, d\nu, \tag{17}$$

where B_ν is the Planckian at the stellar effective temperature [although other stellar spectral energy distributions (SEDs) may be easily incorporated].

(iv) A photon can propagate an optical depth

$$\tau = (k_\nu + k_{sca,\nu})\rho\ell \tag{18}$$

(where $k_{sca,\nu}$ is the scattering coefficient at frequency ν) before it is either absorbed or scattered. The optical depth is selected randomly:

$$\tau = -\ln(1 - \eta). \tag{19}$$

If τ is greater than the optical depth from the cell boundary, then the photon packet position is moved to the cell boundary and a new optical depth is selected. If the time to travel the optical depth τ is greater than Δt, then the photon packet is moved by the distance $c\Delta t$ and the photon information is placed on the stack. Otherwise, the photon position is updated and the type of interaction is determined by the albedo

$$\eta \leq \alpha = \frac{k_{sca,\nu}}{k_\nu + k_{sca,\nu}}. \tag{20}$$

If the photon packet is scattered, then the new direction vector is chosen randomly from the appropriate phase function. If the photon is absorbed, then the appropriate path-length is stored and the next photon packet is selected.

(v) Once all photon packets have been processed, the new absorption rates are computed and the energy density of the gas is updated using equation (4). Subsequently, the duration of the next time-step is determined (see below).

2.1 Time-scales

Cells in which the emission rate exceeds the absorption rate are cooling, and the time-scale for cooling is calculated as

$$t_{cool} = \frac{u_g}{\dot{E} - \dot{A}}. \tag{21}$$

Conversely, if the energy absorption rate exceeds the emission rate for a given cell, then its temperature is increasing. We may then calculate an approximate radiative equilibrium time-scale by first computing the equilibrium energy density u_e (using equation 3) and then finding

$$t_{eq} = \frac{u_e - u_g}{\dot{A} - \dot{E}}. \tag{22}$$

Cells for which $t_{eq} \leq \Delta t$ have their temperatures set to the radiative equilibrium temperature. Note that this formalism assumes that $\dot{A}$ and $\dot{E}$ are constant over ΔT, and care must be taken to ensure that one does not erroneously set cells to their equilibrium temperature due to a poorly chosen ΔT.

3 TESTS

I have implemented a number of test cases in order to examine the efficacy of the new algorithm.
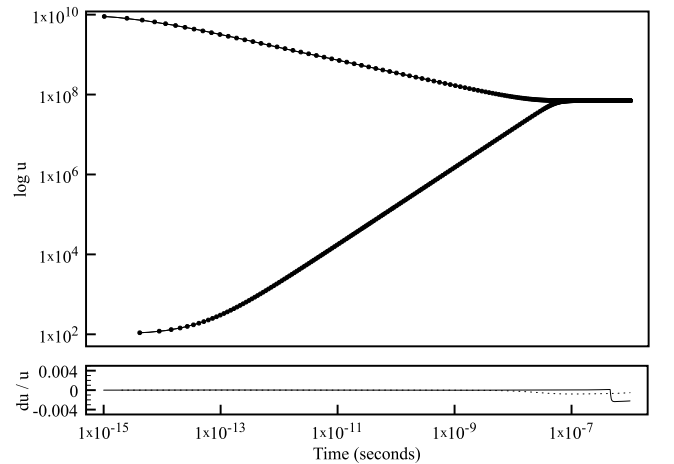
3.1 Radiative equilibrium

If we immerse an absorbing gas in a radiation field of much higher energy density, then the gas will eventually come into equilibrium with the radiation. Since $u_g \ll u_r$, we can assume u_r to be a constant and the evolution of u_g can then be found from the ordinary differential equation

$$\frac{du_g}{dt} = c\kappa u_r - 4\pi\kappa B(u_g). \tag{23}$$

For this test, we adopted the parameters given by Turner & Stone (2001): $\rho = 10^{-7}\,g\,cm^{-3}$, $\kappa = 4 \times 10^{-8}\,cm^{-1}$, $\mu = 0.6$, $\gamma = 5/3$ and $u_r = 10^{12}\,erg\,cm^{-3}$. We ran two separate cases, one with u_g initially well below the equilibrium value ($u_g = 10^2\,erg\,cm^{-3}$) and the other with it well above ($u_g = 10^{10}\,erg\,cm^{-3}$). The results are plotted in Fig. 1 and the agreement with the analytical solution is excellent.

The second test involved the same physical parameters for the box detailed above, but instead of filling the box with a high-density radiation field, we set $u_r = 0$ and $u_g = 10^8$ everywhere. As the RT is followed, the material in the box will cool and emit radiation. Since we have reflective boundary conditions for the photon packets, the gas and radiation field will eventually settle into radiative equilibrium. We set N_p to 1, which means that the gas emits one photon per time-step. In Fig. 2, we plot the evolution of the radiation and gas energy densities with time, and these show good agreement with a simple numerical integration of equations (5) and (4). We find that the total energy is conserved to better than 0.5 per cent throughout the duration of the MC calculation.

Figure 1. The results of the first radiative equilibrium test case described in Section 3.1. In the upper panel, the MC results are shown as dots, while the analytical solution found from equation (23) is displayed as a solid line. The bottom panel shows the fractional difference between the analytical and numerical solutions for the energy density.
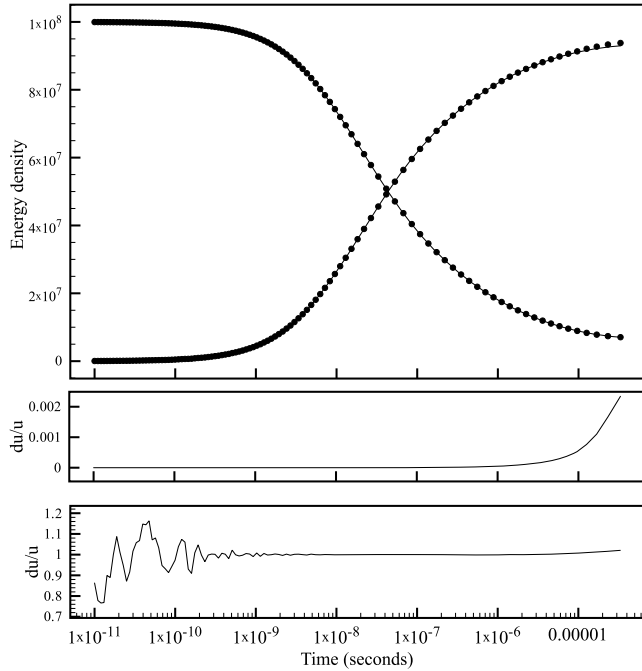
Figure 2. The results of the second radiative equilibrium test case described in Section 3.1. In the upper panel, the MC results are shown as solid lines, while the solutions found from numerical integration of equations (5) and (4) are displayed as dots. The middle and lower panels show the fractional difference between the energy densities from the analytical and MC solutions for equations (5) and (4), respectively.

The above tests suggest that the physics of matter–radiation interaction are adequately captured by the new algorithm, and we now progress to tests of the transport of radiation.

3.2 Diffusion limit

In the optically-thick limit, the RT occurs according to the diffusion equation

$$\frac{\mathrm{d}u_\mathrm{r}}{\mathrm{d}t} = -D\frac{\mathrm{d}^2 u_\mathrm{r}}{\mathrm{d}x^2}, \tag{24}$$

where D is the diffusion coefficient, given by

$$D = \frac{c}{\kappa}. \tag{25}$$

Here we test the MC algorithm on a heat kernel problem, in which energy is deposited at a point at $t = 0\,\mathrm{s}$ and is allowed to diffuse through the medium. The analytical solution to the diffusion equation for this scenario is a Gaussian

$$u(x, t) = \frac{1}{\sqrt{4\pi D t}}\exp\left(-\frac{x^2}{4Dt}\right). \tag{26}$$

We assume a 1-cm 1D box, divided into 101 evenly spaced bins, with reflective boundary conditions for the photons (i.e. adiabatic). We adopt $\kappa = 10^{13}\,\mathrm{cm}^{-1}$ and deposit 10^{10} erg of energy in photons into the central bin at $t = 0\,\mathrm{s}$. (The radiation field immediately comes into radiative equilibrium with the material in the box, but $u_\mathrm{r} \gg u_\mathrm{e}$). We follow the RT using the MC algorithm and also by solving equation (24) using the Crank–Nicolson method.

We initially assume pure scattering ($\alpha = 1$), in which case the MC algorithm will always perfectly preserve energy since no photons are created or destroyed and matter/radiation energy transfer terms
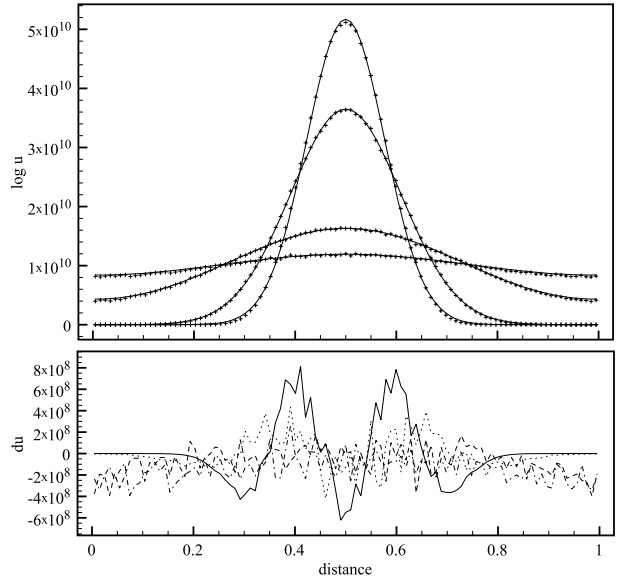
Figure 3. Heat kernel test for the pure scattering case. In the upper panel, the solution to the diffusion equation is plotted (solid lines) for times $t = 10^{-11}$, 2×10^{-11}, 10^{-10} and 2×10^{-10} s, while the radiation energy densities for the same times found using the MC algorithm are plotted as the crosses. The lower panel shows the absolute differences between the analytical and MC-based energy densities at the same time-steps (solid, dotted, dashed and dot–dashed lines, respectively).

are by definition zero. This scenario tests the random walk and photon flight time limit section of the MC algorithm, and we find good agreement between the diffusion approximation and the MC algorithm for this test case (Fig. 3).

The purely absorptive case ($\alpha = 0$) is a much more challenging proposition. Both the emissivity and absorption rates are very high, and MC sampling errors can potentially lead to deviations from energy conservation. However, we find good agreement (Fig. 4) with the diffusion approximation, and energy is conserved to within 2 per cent even after more than 4000 time-steps.

3.3 Time-varying source

It appears that the diffusion limit is well modelled using the MC algorithm, but such conditions are naturally more efficiently treated using the diffusion approximation. However, many astrophysical problems involve transport through substantial regions of optically-thin material, and here we conduct a test of the MC algorithm that incorporates both optically-thin and opaque materials. We also include a time-varying source of photons which allows us to test both how well the algorithm retains the coherence of the radiation field and how well it captures the heating and cooling of material.

We adopt a 1-cm box with a density of $2 \times 10^{-5}\,\mathrm{g\,cm}^{-3}$ and

$$\kappa(x) = \begin{cases} 0 & \text{if } x < 0.5 \\ 10^6 & \text{if } x \geq 0.5 \,. \end{cases} \tag{27}$$

Photons are injected into the left-hand boundary of the box from a source of luminosity that varies sinusoidally between zero and $10^{20}\,\mathrm{erg\,s}^{-1}$ on a period of 10^{-11} s. Photons propagate through the optically-thin part of the box until they encounter the optically-thick material and are absorbed, thus heating the right-hand half of the box. This heated material emits radiation in an attempt to come into thermal balance; since the optical depth is smaller in the $-x$ direction, these photons are preferentially emitted in that direction
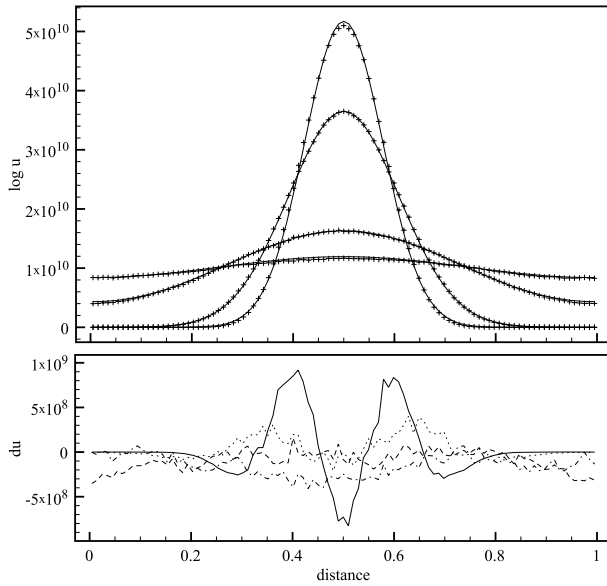
Figure 4. Heat kernel test for the pure absorption case. In the upper panel, the solution to the diffusion equation is plotted (solid lines) for times $t = 10^{-11}$, 2×10^{-11}, 10^{-10} and 2×10^{-10} s, while the radiation energy densities for the same times found using the MC algorithm are plotted as the crosses. The lower panel shows the absolute differences between the analytical and MC-based energy densities at the same time-steps (solid, dotted, dashed and dot–dashed lines, respectively).

(accompanied by a transport by diffusion in the $+x$ direction). Outflow boundary conditions are imposed at either end of the box.

Several time-steps from this test calculation are presented in Fig. 5. Initially, the sinusoidally varying radiation field can be seen to propagate at speed c through the left-hand half of the box. The optically-thick material is heated and re-emits radiation which can then be seen propagating to the left-hand side with a sinusoidally varying flux (the amplitude of the variability is lower than that of the impinging radiation field, never reaching zero, due to the thermal capacity of the optically-thick material). The diffusion of radiation through the optically-thick material is also apparent.

We have checked the energy conservation of the algorithm as the source radiation is transported, absorbed and re-emitted. Fig. 6 shows the integrated thermal and radiation energy presented in the box at each time-step (prior to any radiation escaping the box at the boundaries). The total energy is conserved to better than 1 per cent at all time-steps.

3.4 Implementation

The above 1D, grey tests indicate that the algorithm is fundamentally sound and can reproduce analytical results. We therefore implemented a version of the algorithm as a module within the TORUS radiative transfer code (Harries 2000; Kurosawa et al. 2004; Pinte et al. 2009). The code is written in FORTRAN 90 and follows the radiative transfer on an adaptive mesh stored as a quadtree (2D) or octree (3D).

Since each MC photon packet is essentially an independent event, the code is straightforwardly parallelized under the Message Passing Interface (MPI). Each thread holds a copy of the grid in memory, and the work of the photon loop is distributed over all the MPI threads. At the end of the photon packet loop, the summations in equations (10) and (11) are made over all threads, and thus estimates for the photon energy density and absorption rate are found for each cell
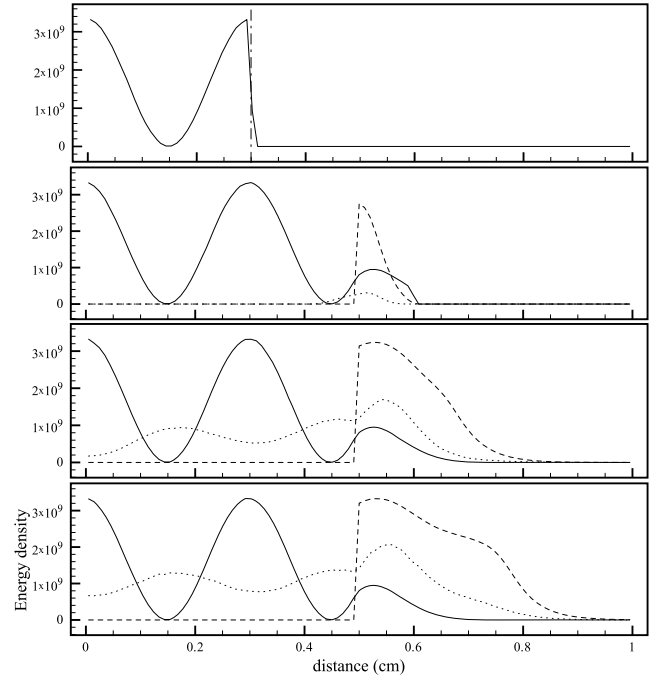
Figure 5. An optically-thin/thick radiation propagation test with a varying source. The panels (top to bottom) are snapshots at $t = 10^{-11}$, 2×10^{-11}, 4×10^{-11} and 6×10^{-11} s. Plotted are the radiation energy density of the source photons (solid line), the thermal energy density of the material (dashed line) and the radiation energy density of photons emitted by the material (dotted line). The expected maximum extent of the radiation field from the source at $t = 10^{-11}$ s is plotted as a dot–dashed line in the upper panel.

in the grid, and finally the results are distributed back to all threads. The communication overhead for the last step is minimal compared to the calculations for the photon loop, and the calculation is thus CPU rather than thread-communication limited.

Example CPU times are susceptible to rapid erosion in usefulness due to Moore's law, but for the benefit of contemporary readers the science grade calculation described in Section 4 was run on 16 theads on the University of Exeter's SGI Altix ICE 8200 supercomputer and each of the later time-steps (when the photon stack size had maximized) took approximately 4 min, of which 20 s was inter-thread communication. It was found that the calculation scaled almost perfectly up to 64 threads (the maximum that we tested).

3.5 Benchmark protostellar disc

As an initial test, we used the 2D RT benchmark disc described by Pascucci et al. (2004). By following the radiation transport over a sufficiently long duration, the disc (initially at $T = 0$ K) will come into radiative equilibrium with the central star. This represents a rigourous test of the new algorithm since the circumstellar material contains both optically-thin and optically-thick regions and thus simultaneously incorporates both free-streaming radiation and diffusing photons. Furthermore, the polychromatic nature of the algorithm is tested with excess luminosity from the photosphere at short wavelengths, heating the disc and being re-emitted in the mid-infrared (mid-IR) and submillimetre regimes.

We adopted an initial time-step of 100 s and used 10^6 photons per time-step. In Fig. 7, the gradual heating of the disc is shown. The rarified regions above and below the disc quickly reach thermal equilibrium, while the disc mid-plane, which has a characteristic
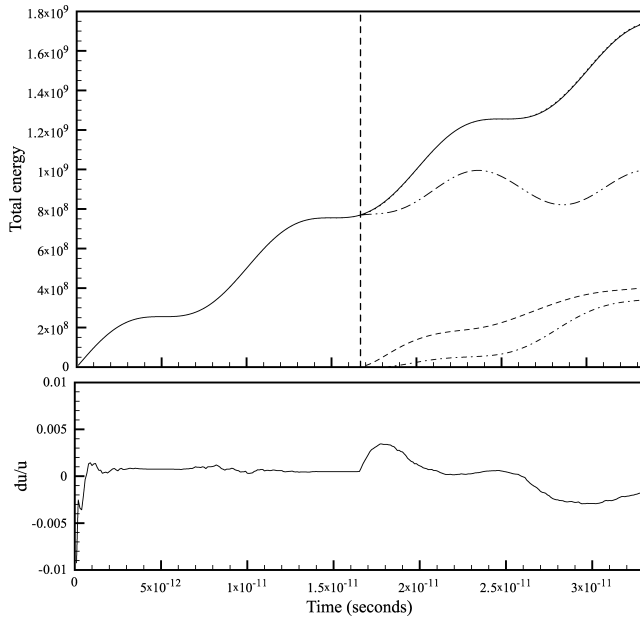
Figure 6. Total energies for the optically-thin/thick radiation propagation test with a varying source. In the upper panel, the expected total energy in the box is plotted as a solid line, along with the total energy calculated from the MC algorithm (dotted line, very close to the solid line), the energy in the radiation field (dot–dot–dashed line), the thermal energy (dashed line) and the radiation density emitted from the heated material (dot–dashed line). The vertical dashed line indicates the expected time in which the source radiation field enters the optically-thick material (1.67×10^{-11} s). The lower panel shows the fraction difference between the expected total energy in the box and that derived from the MC method.

radial optical depth of ~ 100 at 5500 Å, approaches equilibrium much more slowly. After $\sim 3 \times 10^{10}$ s, the entire disc has reached a steady state.

A quantitative comparison with the benchmark solution is shown in Fig. 8. Good agreement with the published benchmark mid-plane temperature distribution is found after $\sim 10^{10}$ s.

4 APPLICATION: PROTOSTELLAR DISC VARIABILITY

The accretion rate on to protostars is inherently variable on a wide variety of time-scales from hours to years (e.g. Bouvier et al. 2007; Nguyen et al. 2009). This variable accretion flux (emitted primarily at wavelengths less than 4000 Å) is scattered and reprocessed in the disc to be emitted in the near- and mid-IR. There will be an intrinsic delay between the optical variability and the disc's IR response, in part due to the light traveltime from the star to the disc and due to thermal lag within the disc itself. Since the temperature in the disc decreases radially, one expects the time-lag (with respect to the optical flux variability) to increase with wavelength. Hence, by obtaining simultaneous time-series photometry in the optical and IR, one may attempt to map the disc emission using the wavelength-dependent lag, in much the same way as reverberation mapping is used to map the broad-line region around active galactic nuclei from the emission-line/ultraviolet continuum variability (e.g. Denney et al. 2009). Of course, there are a host of other factors that might complicate this rather naive interpretation of the variability, such as changes to the disc structure, and rotational modulation of either an

azimuthally structured inner edge (warps) or an anisotropic radiation field (hotspots), but it is none the less interesting to examine the expected time-scales from a simple thermal response model.

4.1 Model parameters

We adopt a central protostar of radius $R_* = 2\,R_\odot$ and a blackbody photospheric flux distribution at $T_{\mathrm{eff}} = 4000$ K. The recent study by Nguyen et al. (2009) showed that the accretion rate on to a typical protostar varies by a factor of ~ 2 on a time-scale of days to weeks. We assume the star is accreting at a rate that varies sinusoidally between 5×10^{-8} and $1 \times 10^{-7}\,M_\odot\,\mathrm{yr}^{-1}$ with a period of 1 h. This period is comparable with the flushing time-scale of the magnetosphere and line profile variability on such time-scales has been observed (see e.g. Smith et al. 1999; Kurosawa et al. 2005). Such rapid variability corresponds to a change on a much shorter time-scale than the canonical rotation period, making it easier observationally to distinguish disc reprocessing effects from rotationally modulated disc-structure effects.

In order to mimic the additional flux associated with the accretion, we add the accretion luminosity as a blackbloby with a characteristic temperature found by assuming that the accretion power is emitted from an area equivalent to 5 per cent of the stellar photosphere.

A simple flared structure is adopted for the disc, namely

$$\rho(r, z) = \rho_0 r^{-\alpha} \exp\left[-\frac{1}{2} \frac{z^2}{h(r)^2} \right], \tag{28}$$

where

$$h = h_0 \left(\frac{r}{r_0} \right)^\beta \tag{29}$$

with $\alpha = 2.25$ and $\beta = 1.25$. The disc scaleheight is set to $h_0 = 125$ au at $r_0 = 100$ au. In order to simulate the truncation of the inner disc by the magnetosphere, we assume an inner hole of $10R_*$, while the outer radius is 300 au. We fix ρ_0 by assuming a disc mass of $0.01\,M_\odot$. The disc is assumed to contain the dust size distribution and chemistry described by Kurosawa et al. (2004), and for the purposes of this test, we assume isotropic scattering.

4.2 Method

The disc is initially brought into radiative equilibrium with the central object (assuming a constant accretion rate of $5 \times 10^{-8}\,M_\odot\,\mathrm{yr}^{-1}$) by using a time-independent algorithm. The time-dependent method is then turned on (which defines $t = 0$ s) and the RT is followed for 10 periods at a time-step of 18 s. This short time-step allows us to adequately resolve the shortest time-scales that are of interest, which is basically the light-crossing time from the central object to the inner disc (47 s).

Since the photon packet estimate of the radiation field is only defined within a sphere of radius ct around the central object, it is only cells within this radius whose thermal properties are changed at each time-step, with the rest of the disc held in thermal equilibrium. Free-streaming photon packets that have a negligible probability of interacting with the disc (found by integrating the optical depth along the packet's path) are deleted from the stack at the end of each time-step. This has the advantage of significantly reducing the memory requirement of the photon stack.

The shape of the cross-correlation function (CCF) should be a function of the system's viewing angle (inclination) as different
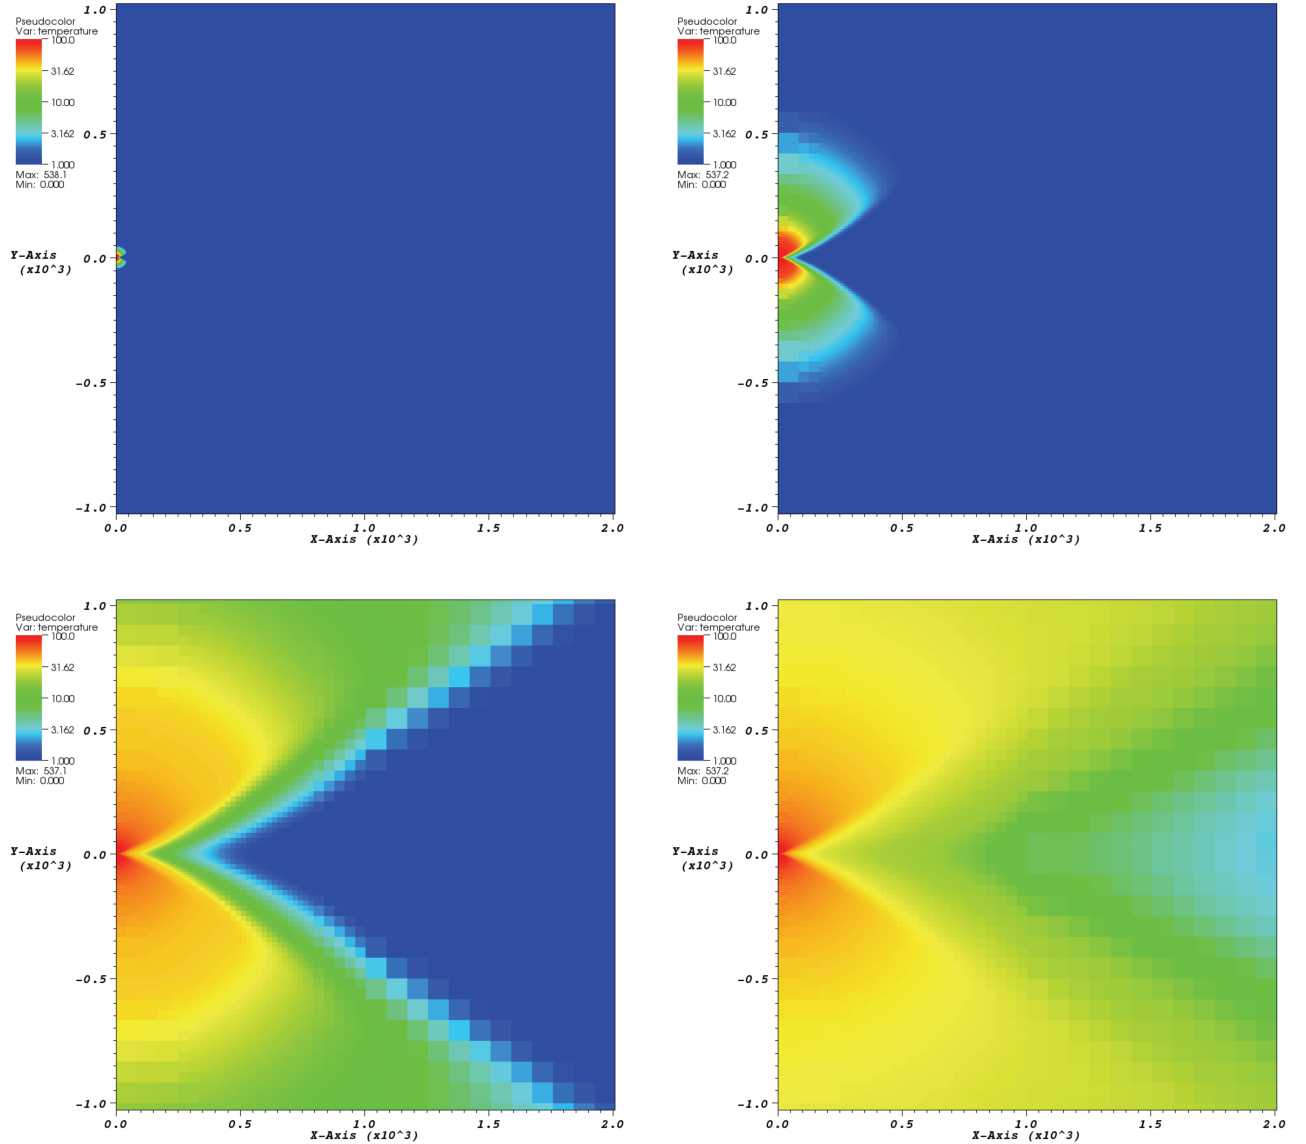
Figure 7. The Pascucci benchmark disc heated from $T = 0$ K to thermal equilibrium using the MC algorithm. The temperature of the disc is plotted as a logarithmic colour-scale scaled between 1 and 100 K. Distances are in au. Four snapshots are shown: $t = 3.1 \times 10^4$ s (top left-hand panel), $t = 1.02 \times 10^6$ s (top right-hand panel), $t = 3.3 \times 10^7$ s (bottom left-hand panel) and $t = 1.05 \times 10^9$ s (bottom right-hand panel).

parts of the disc have different star–disc–observer path lengths. In the following, we adopt an inclination of 60° when calculating the SEDs.

Initially, a set of photon packets were generated from the volume of grid outside the *ct* sphere defined by the length of the time-series. Since this volume is at constant temperature, it was possible to calculate this contribution to the SED in a time-independent manner. Subsequently, at each time-step, 10^6 new photon packets (photospheric or thermal disc) were generated, and these photons were tagged by their generation time. A 'peel-off' technique was then employed: the light traveltime and optical depths to the observer were calculated, along with the probability that the photon packet was emitted towards the observer direction. The photon packets were then binned in the observer's frame, both spectrally and temporally. The new photon packets, plus the photon packet stack, were then followed for a single time-step, with packet peel-off and binning occurring at each scattering event.
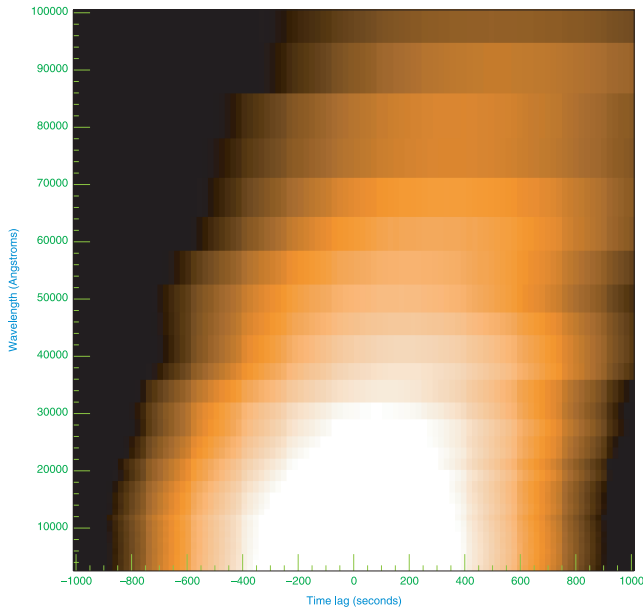
4.3 Results

We have computed the CCF of monochromatic time-series from 3000 Å to 10 μm against a fiducial time-series at 3000 Å (see Fig. 9). It is immediately apparent that the lag increases with wavelength due to a combination of the light traveltime to the disc and the thermal lag as the increased radiation field heats the local disc material. The maximum value of the CCF decreases at longer wavelengths as the thermal time-scale of the disc begins to dominate over the variability time-scale.

I have quantified the lag by fitting a Gaussian to each CCF peak as a function of wavelength. The central positions of these Gaussians are plotted in Fig. 10. The lags at blue wavelengths ($\ll 2$ μm) are short, in fact less than the light traveltime between the star and the disc; this is principally because there is significant direct photospheric emission at these wavelengths (with zero lag) in addition to the scattered and reprocessed components. Redwards of ~ 2 μm,
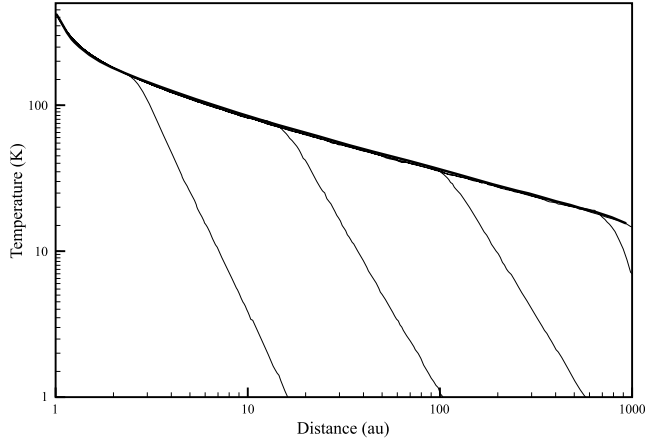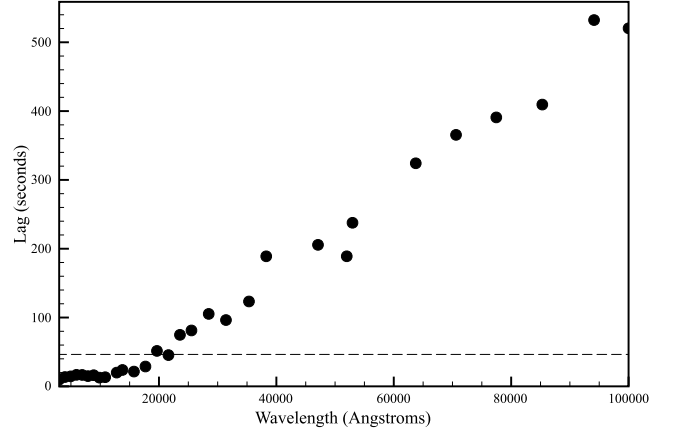
Figure 8. The disc mid-plane temperature for the Pascucci benchmark disc. The heavy line shows the benchmark temperature distribution, while the other lines show (from the left-hand side) snapshots at $t = 3.1 \times 10^4$, 1.02×10^6, 3.3×10^7, 1.05×10^9 and 3.3×10^{10} s.

Figure 10. The time-lag with respect to the 3000-Å light curve plotted as a function of wavelength (filled circles). The light traveltime between the stellar centre and the inner disc edge is plotted as a dashed line.

Figure 9. A cross-correlation image of the monochromatic disc light curve at 3000 Å against wavelengths from 3000 Å to 10 μm. The image is a linear grey-scale of the CCF scaled between 0.8 (white) and 1 (black). The bottom row of the image is the autocorrelation function of the 3000-Å light curve.

the flux is dominated by the thermal emission from the disc, and we see that the lag increases monotonically (within the errors on the CCF Gaussian fits) with wavelength (although the strength of the correlation is decreasing).

Observing this correlated phenomenon would be a challenging proposition, necessitating high-cadence temporal sampling in the optical and near-IR simultaneously. Optimally, one would select ground-based *UBV* photometry combined with *Spitzer* IRAC photometry. Although it may be possible to conduct the IR observations into the mid-IR, it is clear that the correlation is dropping rapidly at these wavelengths, and it is not at all clear whether the necessary temporal sampling could be achieved using ground-based facilities. It may be that systems with a larger inner hole, such as Herbig AeBe stars (Monnier et al. 2005), would be superior targets to Classical

T Tauri stars, allowing light-traveltime effects to remain dominant at longer time-scales.

Of course, in reality, the underlying accretion luminosity is unlikely to vary periodically, but more stochastic variations could still be investigated. It may be possible to distinguish between disc thermal responses and scattering by observing polarized light, which will be dominated by the scattered light contribution.

5 CONCLUSIONS

I have presented a new method for computing time-dependent radiation transport for an arbitrary distribution of sources embedded in an arbitrary distribution of absorbing, emitting and scattering material. The new algorithm, based on the MC radiative equilibrium method of Lucy (1999), can be used in 3D and scales almost perfectly under parallelization. It has advantages over the FLD approximation in that it is polychromatic and correctly treats the directionality and flux of the radiation field in the optically-thin limit. Note that although the algorithm as presented here is only applicable under conditions of LTE, the method may be straightforwardly extended to the non-LTE regime (see e.g. Carciofi & Bjorkman 2006).

I have applied the new method to the problem of a circumstellar disc illuminated by a protostar with a periodic time-variable accretion rate. I have shown that the lag between the blue continuum resulting from the accretion hotspots and the reprocessed IR radiation from the disc is a strong function of wavelength. It appears that photometric time-series data in the blue part of the optical spectrum, combined with equally-intensive near-IR time-series, could be used to probe the geometrical and thermal structure of the disc, although it is likely that complications arising from stochastic variations in the disc structure close to the inner edge could mask the correlation calculated here.

In future, the time-dependent method will be incorporated in RHD calculations of protostellar disc fragmentation, extending the time-independent transfer simulations of Acreman et al. (2010). It is well established that the thermal properties of the disc strongly affect the fragmentation (Boss 2008; Stamatellos & Whitworth 2008), and it is also clear that the FLD approximation does not capture the full physics of the transport in the disc itself.

ACKNOWLEDGMENTS

REFERENCES

Acreman D. M., Harries T. J., Rundle D. A., 2010, MNRAS, 403, 1143
Boss A. P., 2008, ApJ, 677, 607
Bouvier J., Alencar S. H. P., Harries T. J., Johns-Krull C. M., Romanova M. M., 2007, in Reipurth B., Jewitt D., Keil K., eds, Protostars and Planets V. Univ. Arizona Press, Tucson, p. 479
Carciofi A. C., Bjorkman J. E., 2006, ApJ, 639, 1081
Dale J. E., Ercolano B., Clarke C. J., 2007, MNRAS, 382, 1759
Denney K. D. et al., 2009, ApJ, 704, L80
Freytag B., Höfner S., 2008, A&A, 483, 571
Harries T. J., 2000, MNRAS, 315, 722
Höfner S., Gautschy-Loidl R., Aringer B., Jørgensen U. G., 2003, A&A, 399, 589
Jack D., Hauschildt P. H., Baron E., 2009, A&A, 502, 1043
Kromer M., Sim S. A., 2009, MNRAS, 398, 1809
Krumholz M. R., Klein R. I., McKee C. F., Bolstad J., 2007, ApJ, 667, 626
Krumholz M. R., Stone J. M., Gardiner T. A., 2007, ApJ, 671, 518

Kurosawa R., Harries T. J., Bate M. R., Symington N. H., 2004, MNRAS, 351, 1134
Kurosawa R., Harries T. J., Symington N. H., 2005, MNRAS, 358, 671
Lucy L. B., 1999, A&A, 344, 282
Mac Low M.-M., Toraskar J., Oishi J. S., Abel T., 2007, ApJ, 668, 980
Monnier J. D. et al., 2005, ApJ, 624, 832
Nguyen D. C., Scholz A., van Kerkwijk M. H., Jayawardhana R., Brandeker A., 2009, ApJ, 694, L153
Pascucci I., Wolf S., Steinacker J., Dullemond C. P., Henning T., Niccolini G., Woitke P., Lopez B., 2004, A&A, 417, 793
Pinte C., Harries T. J., Min M., Watson A. M., Dullemond C. P., Woitke P., Ménard F., Durán-Rojas M. C., 2009, A&A, 498, 967
Preibisch T., Sonnhalter C., Yorke H. W., 1995, A&A, 299, 144
Smith K. W., Lewis G. F., Bonnell I. A., Bunclark P. S., Emerson J. P., 1999, MNRAS, 304, 367
Stamatellos D., Whitworth A. P., 2008, A&A, 480, 879
Turner N. J., Stone J. M., 2001, ApJS, 135, 95
Whitehouse S. C., Bate M. R., 2004, MNRAS, 353, 1078
Whitehouse S. C., Bate M. R., Monaghan J. J., 2005, MNRAS, 364, 1367
Woitke P., 2006, A&A, 452, 537

This paper has been typeset from a TEX/LATEX file prepared by the author.