An Efficient Monte Carlo Method for Modeling Radiative Transfer in Protoplanetary Disks

Stacy Kim, SURF Fellow; Neal Turner, Mentor

Monte Carlo methods have been shown to be effective and versatile in modeling radiative transfer processes to calculate model temperature profiles for protoplanetary disks. Temperatures profiles are important for connecting physical structure to observations and for understanding the conditions for planet formation and migration. However, certain areas of the disk such as the optically thick disk interior are under-sampled, or are of particular interest, such as the snow line (where water vapor condenses into ice) and the area surrounding a protoplanet. To improve the sampling, photon packets can be preferentially scattered and reemitted toward the preferred locations at the cost of weighting packet energies to conserve the average energy flux. Here I report on the weighting schemes developed, how they can be applied to various models, and how they affect simulation mechanics and results. We find that improvements in sampling do not always imply similar improvements in temperature accuracies and calculation speeds.

When a protostar is born from a swirling, contracting cloud of dense gas, it also forms a rotating dusty and gaseous disk, produced by accretion of the surrounding interstellar material and the demands of angular momentum conservation. Temperatures throughout the disk are crucial to understanding the physical and chemical processes that occur in the disk and thus the evolution of the disk and the planets that form within. Perhaps of most importance is the role they play in producing spectral energy distributions (SEDs), by which a disk's orientation and density structure and can be inferred. Additionally, they control the distribution and phases of various chemical species due to the temperature dependence of processes such as freeze-out and the migration of gases and solids, which set the conditions for planet formation.

Disk temperature profiles can be calculated with relative ease using the Monte Carlo method due to its adaptability to complex disk shapes (see Figure 1) and the stochastic nature of radiative transfer processes. Despite the success of the Monte Carlo radiative transfer approach in modeling the protoplanetary disk, however, there remain a few conditions under which it performs inefficiently. The most significant of these is caused by the high optical depths of the disk interior. A photon packet trapped in this region can dominate the computational power and extend by orders of magnitude the time to obtain the solution[3]. This can be resolved this in part by assuming the densest regions of the disk interior to be isotropic and making it an exclusion zone where by construction no packets penetrate, thereby reducing the number of interactions and thus the solution time. After temperatures throughout the rest of the grid have been calculated, the temperatures of the deepest few cells for which temperatures were calculated are averaged and extended into the exclusion zone. An alternative is to approximate multiple packet scattering and reemission with a single diffusion step in the dense interior.

A second concern arises from the undersampling of the radiation field in particular cells. If a sufficient number of photon packets do not pass through a given cell and its neighbors or a photon packet remains trapped in one cell for an extended period of time, the resulting temperature profile of the disk can be very noisy. This becomes a particular concern in areas where there is a high number of cells, which requires a greater number of photons to accurately calculate cell temperatures, and again, where optical depths are high, as photons are likely to have escaped from the disk rather than enter the optically thick disk interior. Photons can be preferentially scattered and reemitted into these undersampled regions at the cost of weighting packet energies by implementing photon packet-weighting schemes in the spirit of Juvela[1]. This technique can also be applied to better sample regions of the disk where accurate results are desired. Here I report on the development and implementation of the packet weighting schemes for disks. Despite successfully improving sampling in the disk interior by factors of 6-8, temperatures in the disk interior are incorrect by 27-40% and can more than double the runtime.

I. The Monte Carlo Radiative Transfer Code

The code used in this investigation was developed by N. J. Turner and is similar to RADMC, a Monte Carlo radiative transfer code developed by C. P. Dullemond[2]. Photon packets are followed from their emission at the central protostar through its ensuing scattering, absorption, and reemission in a dusty disk until it escapes to infinity. Assuming that the dust is in local thermodynamic equilibrium, packets are reemitted according to Bjorkman and Wood's immediate reemission procedure[3] in which the packet's energy is added to its surroundings and then reemitted with its original energy but with a frequency chosen from the difference spectrum of the local temperature before and

after the packet is absorbed. Thus as more photons are passed through the disk, the disk temperatures relax to the correct solution, eliminating the need for CPU-costly iteration required by alternate methods. Absorption is treated as a continuous process as prescribed by Lucy[4], which reduces the noisiness inherent to the Monte Carlo approach to radiative transfer and accounts for the possible contribution a packet could have made as it travels through the disk.

Dust grain opacities are drawn from Preibisch et al.[5], which assumes a mixture of three grain species, amorphous carbon, silicate grains, and dirty ice mantles on silicate grains below 125 K.

II. Importance Weighting

Juvela[1] considers four schemes in which the location of photon packet generation and/or the direction of their motion are weighted to favor the sampling to the optically thick region at the center of a spherically symmetric molecular cloud core. The protoplanetary disk model we are considering differs from Juvela's molecular core model in that there is currently no packet generation within the disk and that the disk is cylindrically symmetrical. However, the basic concept still applies. A photon packet can be preferentially scattered or reemitted towards a particular direction if an appropriate weight is applied to the packet's energy so that, when averaged over many interactions, the energy flux remains isotropic. The weight applied to packet is

$$W = \frac{q(\Omega')}{p(\Omega)}, \tag{1}$$

where $q(\Omega')$ denotes the original probability of being scattered or reemitted into a particular direction and $p(\Omega)$ is the new angular probability distribution applied to the packet (the angles for each distribution may be defined in different reference frames corresponding to Ω and Ω'). Thus in directions where the probability is augmented, more packets will be scattered or reemitted, but each packet will have less energy; in directions where the probability is diminished, less packets will be scattered/reemitted, but each have more energy.

Though the specific weighting scheme used is largely dependent on the nature of the disk model or problem being investigated, we can nevertheless implement weighting schemes with some generality. As the focus of this investigation was on obtaining better results in the disk interior, we focused on weighting in favor of large polar angles. The first scheme

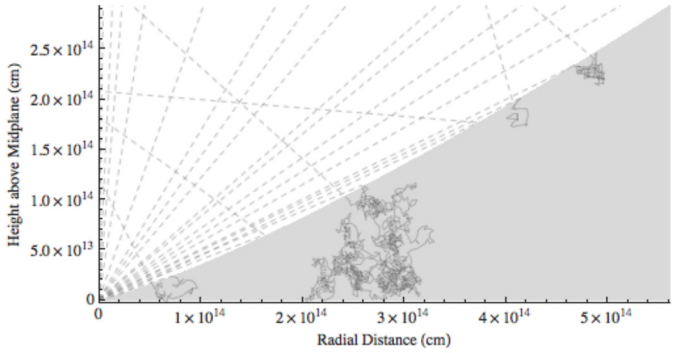

Figure 1 | Packet paths in disk. Sample packet paths in an axi- and midplane symmetric disk, plotted on a linear scale. Step lengths have been exaggerated. The star is in the lower left corner. About 70% of emitted packets do not reach the disk. Packets strike the disk surface at a grazing angle, where about a third are immediately scattered or reemitted into empty space. Only a few reach the deep interior.

(hereafter the "hemisphere" weighting scheme) favors emission in the lower hemisphere ($\pi/2 < \theta < \pi$), but otherwise retains isotropic scattering and reemission, thereby preserving most of the physical scattering/emission process. In detail, given the desired strength of the bias $b \in \Re$ where the quantity $b + 1$ corresponds to the factor that a packet is more likely to enter the lower versus the upper hemisphere, if a randomly generated number $Z \in [0,1]$ exceeds $(b + 2)^{-1}$, then an angle is chosen with isotropic probability in the lower hemisphere; an angle in the upper hemisphere is chosen if the random number is less than or equal to that value. The weighting applied to the packet's energy required to conserve the energy flux is thus

$$\begin{cases} \dfrac{b+2}{2}, & Z < \dfrac{1}{b+2} \\ \dfrac{b+1}{2(b+2)}, & Z > \dfrac{1}{b+2} \end{cases}. \tag{2}$$

A stronger bias towards traveling to the mid-plane can be achieved by specifically favoring the $\theta = \pi$ direction. This is implemented by the second (hereafter "directional") photon weighting scheme, which also has a bias parameter b (again a rational number) corresponding to how heavily the lower pole is favored. The polar angle is chosen according to

Table 1 | Two Photon Weighting Schemes

Scheme	Theta	Weighting	Default Preferred Direction	Probability Maps for b=1,2,3
unbiased	$\theta = \cos^{-1}(-1+2Z)$	1 (none)	--	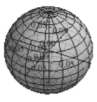
hemisphere	$\theta = \begin{cases} \cos^{-1}(Z_2) & Z_1 < \dfrac{1}{b+2} \\ \cos^{-1}(-Z_2) & Z_1 > \dfrac{1}{b+2} \end{cases}$	$\begin{cases} \dfrac{b+2}{2}, & Z < \dfrac{1}{b+2} \\ \dfrac{b+1}{2(b+2)}, & Z > \dfrac{1}{b+2} \end{cases}$	lower hemisphere	
directional	$\theta = \cos^{-1}[-1+2(1-Z)e^{-bZ}]$	$\left\|(-b+bZ-1)e^{-bZ}\right\|$	θ = π	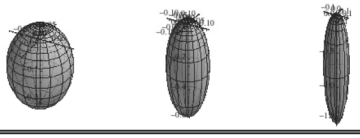

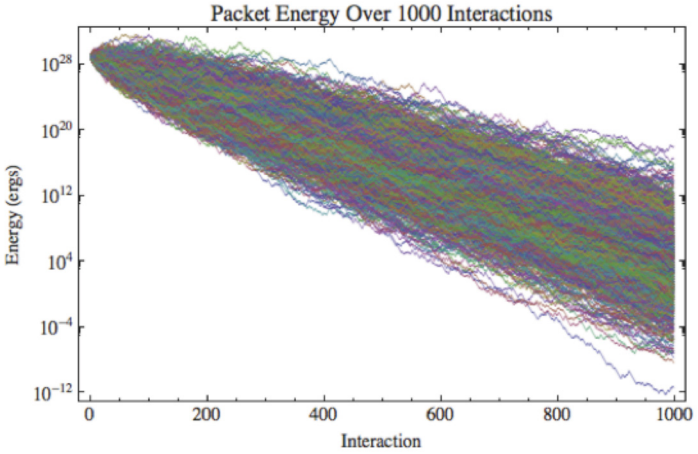

Figure 2 | Packet energy evolution. 1000 realizations of the energy path a packet that undergoes 1000 interactions under the hemisphere weighting scheme with b=1 could take. Note the wide dispersion of energies (often spanning more than 10 orders of magnitude) for a given interaction.

$$\theta = \cos^{-1}[-1 + 2(1 - Z)e^{-bZ}] \quad (3)$$

where all instances of Z denote the same randomly generated number. To derive the corresponding weight required to conserve the averaged energy flux, we take the absolute value (as the weight must be positive) derivative of the weighted [0,1] probability distribution to find the probability per Z:

$$\left|\frac{d}{dZ}[(1 - Z)e^{-bZ}]\right| = \left|(-b + bZ - 1)e^{-bZ}\right|. \quad (4)$$

It was verified that the probability integrated over all solid angles was unity for both schemes. The details of both photon-weighting schemes are summarized in Table 1.

III. Customization

These two weighting schemes are the basic building blocks to a library of more complicated weighting schemes. For instance, a rotation can be applied after generating a new scattering/reemission direction by either of these schemes. This also introduces the ability to bias towards location-dependent directions calculated during runtime, such as radially in/outwards or the direction in which the least sampling has occurred. In general, the hemisphere scheme is useful when one of two general directions is preferred (such as up/down, radially in/out), and the directional scheme, when a specific location (such as the center of a spherical domain) is targeted.

Should the weighting scheme be applied to a packet over a large number of interactions, the packet's energy can drop by several orders of magnitude (see Figure 2), which may cause the packet's contribution to cell temperatures to become negligible. Additionally, should packets be continually biased towards optically thick regions without an easy escape route, such as the optically thick disk interior, packets may become 'trapped' at a high CPU-cost. Two features were implemented to prevent such scenarios. A packet can be flagged for removal and/or the bias reversed (thereby favoring its escape from the disk and preserving the stellar luminosity, one of the main advantages of the immediate reemission procedure) when the packet's energy dips below a given threshold (a separate threshold can be specified for packet removal and bias reversal). In implementation, the bias can be reversed to favor the direction exactly opposite to the original by specifying a negative bias. More sophisticated methods may be used to better gauge when a packet's contribution becomes negligible. Using the hemisphere weighting scheme with a bias strength b=1 for a MMSN disk model with a star of stellar luminosity and an initial packet energy of ~10^{29} ergs, no greater than 2.4% of the total packets were removed or the bias reversed for threshold energies up to 10^{20} ergs (see Table 2).

The domain in which the hemisphere weighting schemes are applied can also be restricted. For instance, due to the difficulties caused by low packet energies in the optically thin disk interior, one could apply the bias in the upper regions of the disk but revert to unbiased scattering and reemission in the interior, thereby preserving the packet's energy upon entering the inner disk. One could also imagine applying the bias to packets emitted from the star to direct packets into the disk (~80% of packets undergo 0 or 1 interaction).

The ways in which the biasing schemes could be applied are myriad, providing flexibility and customizability depending on the nature of the disk model and the problem to be solved. The different applications suggested here can also be used in combination. For instance, reverting to unbiased scattering and reemission at high optical depths (as suggested above) reintroduces the concern of high CPU costs at high optical depths; the total number of interactions a packet undergoes in the disk interior could be monitored and the packet biased upwards when it has undergone a large number of interactions.

IV. Tests for Energy Conservation

To confirm that the packet weighting scheme did indeed conserve energy, sample sequences of 1, 10, and 100 interactions were generated using the hemisphere weighting scheme with b=1. For simplicity, we assumed that the photon travels in a medium with unit density and opacity. The final energies and number of packets were binned according to the direction of the final positions of the packets relative to their common source of origin. Due to the sheer number of packets required to correctly assess the energy distribution at large numbers of interactions, we focused on the results of 1 and 10 interactions with many packets (10^6 and 10^7). The rms fractional difference from the mean of the energy received by each bin (constructed so that each of the 10,000 bins covers an equal solid angle) was 3.18%. The energies converged to those found in the unbiased case, with a biased/unbiased energy ratio averaged over all cells of 1.0012 and with an rms fractional difference of 5.97%. As expected, errors were higher at the upper pole due to the larger role of small number statistics.

V. Tests with Monte Carlo: The Models

The subroutine implementing the weighting schemes was integrated into the main Monte Carlo code and tests were run on two 2D axisymmetric models symmetric above and below the midplane, a disk model and a simple cylinder model with constant density. In the following analyses, the density distribution was kept fixed. The models are described below.

Disk Model. The disk model reflects the minimum mass solar nebula (MMSN), a lower bound on the mass required to form the planets in the solar system, distributed with a surface density given by 1700 $(r/AU)^{-3/2}$ g·cm^{-2}. The underlying grid spans 6 scale heights in the vertical direction, divided into 60 cells equally spaced in each column, and covers ~0.04 AU to ~40 AU in the radial direction, divided into 120 logarithmically spaced cells. Packets were bounced out of regions where the mass column, defined by

Table 2 | Parameters and results of test runs of the MMSN disk model, hemisphere weighting scheme

Test No.	Total Packets	Bias Strength	Threshold Energies		% removed	Time (cpu-sec)	Packet E_0*	Upper Disk Interior ($1 < m(z) < 20$)			Disk Exterior ($m(z) < 1$)		
			Discard (ergs)	Reverse Bias (ergs)				Sampling (b/b_0)**	RMS Frac. Diff.***	dA /packet (b/b_0)**	Sampling (b/b_0)**	RMS Frac. Diff.***	dA /packet (b/b_0)**
30	1e4	0	0	0	0	90.4789	7.8e29	–	–	–	–	–	–
01a	1e5	0	0	0	0	392.722	7.8e28	–	–	–	–	–	–
03a	1e6	0	0	0	0	3321.68	7.8e27	–	–	–	–	–	–
17	1e4	1	1	1	0.00570	88.9478	7.8e29						
38	1e5	1	1	1	0.01490	734.471	7.8e28	6.89968	0.2734831	0.277542	1.16224	0.037601	0.873414
39	1e6	1	1	1	0.01760	7324.57	7.8e27						
19	1e5	1	1e5	–	0.01589	628.860	7.8e28						
20	1e5	1	1e10	–	0.01639	512.451	7.8e28						
21	1e5	1	1e15	–	0.02019	437.465	7.8e28						
22	1e5	1	1e20	–	0.02355	337.303	7.8e28						
26	1e5	1	1	1e5	0.00835	633.475	7.8e28						
27	1e5	1	1	1e10	0.00210	649.601	7.8e28	5.56922	0.273453	0.266874	1.19491	0.055996	0.873879
28	1e5	1	1	1e15	0.00014	636.268	7.8e28	6.95861	0.279865	0.288188	1.20125	0.037621	0.847911
29	1e5	1	1	1e20	0.00002	479.721	7.8e28	4.68488	0.279081	0.264277	1.20408	0.038486	0.848605
31	1e5	1	1	1e25	0	379.856	7.8e28	3.79296	0.276572	0.303355	1.20404	0.039506	0.855356
32	1e5	2	1e10	–	0.03266	410.480	7.8e28	7.85723	0.334508	0.211585	1.24540	0.048401	0.852647
33	1e5	3	1e10	–	0.04382	359.441	7.8e28	7.19205	0.361656	0.153921	1.27953	0.052018	0.838428
34	1e5	4	1e10	–	0.05545	344.475	7.8e28	7.70680	0.379615	0.168746	1.28735	0.055529	0.829331
35	1e5	5	1e10	–	0.06142	340.662	7.8e28	7.36734	0.390004	0.127478	1.30040	0.063766	0.850288

*The initial packet energy is calculated according to L_{star}*t/n , where t is the length of time and n is the number of packets modeled by the simulation

**fraction of the value in the unbiased case with the same number of packets

***RMS errors in temperature, using test01 (unbiased, 10^5 packets) as reference

$$m(z) = \int_z^\infty \rho(z')dz' \quad (5)$$

exceeded 20 g·cm^{-2} and the resulting exclusion zone where no packets reach was treated isothermally. As the cells in this grid are generally radially much longer than they are tall, a packet may cause many interactions to occur in one cell while the cells radially adjacent may have no packets, producing noisy results. The hemisphere scheme with a bias toward the lower hemisphere was thus used, which retains the radial sampling of the original isotropic model. Four different sets of tests were performed, one varying the number of packets (and thus initial packet energies), two varying the threshold energy at which photons were either discarded or the bias reversed, and the last varied the strength of the bias. Results were compared against temperature profiles obtained without using weighting schemes but with the same number of packets. The disk tests and their results are summarized in Table 2.

Cylinder Model. To simplify the geometry, we also constructed a simple cylindrical model of constant density and gray atmosphere with a star of solar luminosity. The inner 0.1 AU was evacuated. A coarse square 10 x 10 grid starting from 0.1 AU and extending to 10.1 AU in the radial direction and stretching 10 AU high from the midplane was used to simplify the analysis and to dampen the noise in each cell. A large optical depth of $\tau = 1000$ was used to obtain better spatial sampling throughout the grid. In contrast to the disk model, the least-sampled region of this model is the upper region of the cylinder, and so the hemisphere scheme favoring the upward direction was applied with varying bias strengths. Weaker biases than in the disk tests were applied to prevent excess energy loss. For each bias strength, tests using 10^4 and 10^5 packets were run. Again, results were compared against runs that did not use a weighting scheme but sampled the disk using the same number of packets. The parameters and results of the cylinder model tests are summarized in Table 3a and 3b, respectively.

As the regions of the cylinder near its surface are under-sampled compared to its midplane, the upper regions of the cylinder draw comparison to the optically thick disk interior. It is important to note the differences between the two models, however; in the disk model, packets can be biased continually towards the midplane due to the bounce layer, and thus become 'trapped' in the disk interior, whereas in the cylinder model, packets can immediately escape the disk once reaching the top surface of the cylinder. Though this affects runtimes, the distribution of packet energies, and temperature accuracies, among others, the simpler cylinder model can provide better insight on general trends observed in both models and highlight model-dependent results.

VI. Sampling

A key measure of the accuracy in sampling the radiation field is the number of times the grid cell is crossed by packets. We compare this sampling measure between runs with and without biasing, having the same total number of packets. For both the disk and cylindrical models, the sampling in the targeted regions was improved. In the disk interior, sampling was improved by a

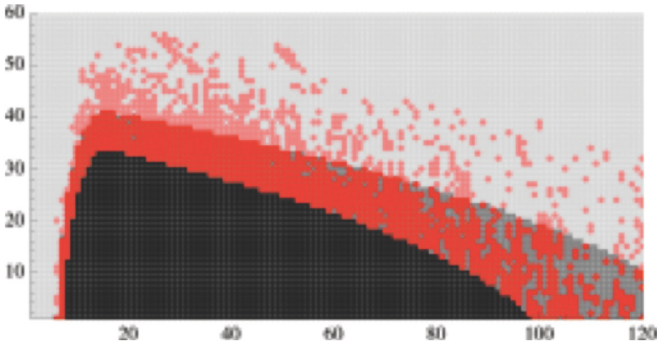

Figure 3 | Disk Regions. Each pixel corresponds to a cell in the grid (the axes are labeled by cell index). The exclusion zone is marked in black, the region where the mass column $20<m(z)<1$g·cm^{-2}, gray, and the region where $m(z)<1$g·cm^{-2}, light gray. Points shaded red denote cells where temperatures in test 38 were higher or lower by more than 5% of that in test 01a.

factor of 7, and could be improved to nearly a factor of 8 with stronger biases (though continually increasing the bias did not consistently produce better sampling). In the top surface of the cylinder, sampling improved by a factor of 2.

Sampling was not necessarily compromised in untargeted regions. While in the cylinder model, the sampling did decrease by a factor of 3 to 8 in the midplane, sampling increased by 20% in the disk exterior. This is not surprising, as packets must pass through the disk exterior in order to reach the disk interior, whereas the cylinder midplane is not located in the path to the targeted top surface regions.

In the cylinder model, sampling improves roughly linearly with bias strength; a decrease in bias strength by a factor of 1/4 from b = 1/4 to 1/16 decreased sampling by about the same factor in the top surface, whereas sampling increased in the midplane by a factor of 4. However, in the disk model, a bias of b=1 increased the sampling in the disk interior by a factor of 6.9, more than three times greater than expected, while greater increases in bias from b=2 to b=5 produced a maximum sampling increase of 14% (as compared to b=1). This could again be attributed to the bounce layer and the lack of an escape channel.

VII. Temperature Accuracy

In the basic run (b=1, threshold energies for packet removal and bias reversal = 1 erg), temperatures in the disk atmosphere were accurate to within 4% on average, while temperatures in the disk interior were off by 27% on average, though individual cell temperatures could be off by as much as ~90%. Thus despite successfully increasing the number of packets reaching the disk interior, there were significantly large departures in temperature.

Similarly, in the cylinder model, though overall temperature profiles were qualitatively similar to that of the unbiased case, actual temperatures deviated significantly from those of unbiased tests. Average rms fractional differences (using results from an unbiased run with 10^6 packets as reference) ranged from 32-55% along the outer edges of the grid for runs with 10^5 packets (errors were greater with fewer packets). Temperatures along the vertical inner surface were the most accurate, and those along the vertical outer surface, most inaccurate (see Table 3b).

Packet removal and bias reversal. Bias reversal produces marginally (~0.1-0.2%) more accurate results than packet removal (in keeping with greater accuracy expected by preserving stellar luminosity).

Table 3a | Parameters of test runs of the cylinder model, hemisphere scheme

Test No.	Total Packets	Bias Strength	Discard Packet E	% removed	Time (cpu-sec)	Packet E_0*
37h	1e4	0	0	0	2 826	7.8e29
40j	1e4	-0.25	1	0	781	7.8e29
40k	1e4	-0.125	1	0	1 285	7.8e29
40l	1e4	-0.0625	1	0	1 851	7.8e29
37m	1e5	0	0	0	27 734	7.8e28
40o	1e5	-0.25	1	0	7 190	7.8e28
40p	1e5	-0.125	1	0	12 664	7.8e28
40q	1e5	-0.0625	1	0	18 446	7.8e28
37r	1e6	0	0	0	276 519	7.8e27

*The initial packet energy is calculated according to L_{star}*t/n, where t is the length of time and n is the number of packets modeled by the simulation

VIII. Calculation Times

For the disk model, stronger biases produce shorter runtimes. However, at bias strengths as low as 1, the runtime can more than double compared to the unbiased runs with the same number of packets; it is only at bias strengths of 3 or greater that improvements in runtime are made, perhaps because the greater weights applied at high biases decrease packet energies more quickly, and thus not only are packets are removed earlier, but more of them are removed as well. These improvements, however, come with a loss in accuracy. Similar improvements, but without as great a loss in accuracy (<1%), can be made by increasing the threshold energy for packet removal to 10^{20} ergs or for bias reversal to 10^{25} ergs.

Improvements for the cylindrical model were more significant. Biasing improved calculation times by 25%-65%, and more significant improvements were obtained with stronger biasing, as is expected when biasing towards an escape channel. Again, however, stronger biases also produced greater inaccuracies.

IX. Conclusions

Two importance-weighting schemes for scattering and reemission were developed and implemented into a Monte Carlo radiative transfer code for disks.

Many variations on a basic weighting scheme can be implemented with relative ease, providing adaptability to a variety of different models.

Sampling improvements are obtained in the targeted region in all tests. The degree of improvement and the degree of compromise (if any) elsewhere in the model are dependent on the

Table 3b | Results of test runs of the cylinder model, hemisphere scheme

Test No.	Top Surface			Inner vertical surface			Midplane			Outer vertical surface		
	Sampling (b/b_0)*	RMS Frac. Diff. in T**	dA /packet (b/b_0)*	Sampling (b/b_0)*	RMS Frac. Diff.**	dA /packet (b/b_0)*	Sampling (b/b_0)*	RMS Frac. Diff.**	dA /packet (b/b_0)*	Sampling (b/b_0)*	RMS Frac. Diff.**	dA /packet (b/b_0)*
37h	1	-	1	1	-	1	1	-	1	1	-	1
40j	2.00069	0.636705	0.010872	1.33985	0.389825	0.418756	0.082640	0.494738	8.45131	0.143313	0.672621	2.21687
40k	*2.15268*	0.218291	0.211629	1.29821	0.102949	0.736781	0.166375	0.197286	5.91089	0.467242	0.266671	2.43391
40l	1.91692	0.044022	0.454253	1.23591	0.063650	0.829071	0.334203	0.047039	3.13419	0.772488	0.063650	1.82104
37m	1	-	1	1	-	1	1	-	1	1	-	1
40o	2.04717	0.467216	0.627511	1.36217	0.318828	0.547011	0.081988	0.338024	36.3669	0.134109	0.545028	19.8787
40p	*2.19398*	0.177916	0.749441	1.30012	0.056730	0.856637	0.164742	0.138925	8.39781	0.453154	0.199072	5.01318
40q	1.93090	0.023391	0.483339	1.22339	0.017048	0.865888	0.338035	0.020446	3.45766	0.753760	0.036610	1.92746
37r	1	0	1	1	0	1	1	0	1	1	0	1

Each of the 4 regions denotes the outermost single column or row in the 2D grid. Values are averaged over all cells in the row or column. Reference unbiased tests are colored gray.

*fraction of the value in the unbiased case with the same number of packets

**RMS errors in temperature, using test37r (unbiased, 10^6 packets) as reference

location of the targeted region and whether an escape channel is nearby.

Improvements in sampling, however, are not clearly correlated with improvements in temperature accuracy and are sometimes even anti-correlated. Why could this be so? Figure 4, which depicts the energy of 1000 packets that have been restricted to move one unit up or down after each interaction gives a conservative estimate of the wide range of packet energies that exist for a given height (which can span 4–5 magnitudes at heights as low as -30 to -40 with variable step size). It is therefore possible that despite the large improvement in sampling in the disk interior, effectively only a few packets with substantially high energies (compared with other packets that contributed to the cell temperature), causing high errors. This may explain why temperatures in the midplane and inner vertical surface of the cylinder model were the most accurate, even despite receiving 8-34% of the packets in the unbiased case in the midplane—as the anti-targeted region, the midplane should mostly receive packets with high energies, and the inner vertical surface intercepts all the high energy packets that enter the cylinder.

Improvements in runtime can be obtained with the existence of an escape channel and at strong biases. However, more accurate results are obtained at more CPU-costly weaker biases.

X. Future Work

The application of weighting schemes is not necessarily limited to the spatial domain, but to any quantity that is randomly selected during the simulation. This suggests weighting the wavelength at which a packet is either emitted from the star or reemitted by a dust particle in order to achieve better results at a particular wavelength or wavelength region of interest.

On a broader scope, we can sophisticate the code in primarily two ways; we can more accurately treat the radiative transfer process and/or better treat the optically thick disk interior by applying different approximations to sample the disk interior without paying too high a price in CPU costs. With regards to the former, instead of scattering packets isotropically, new scattering directions can be chosen from the scattering phase function.

Secondly, instead of assuming that the disk interior is isothermal, we can more accurately treat this region and minimize computational costs by applying approximations. A modified random walk algorithm described by Min et al.[6] and improved by Robitaille[7] has been implemented into the code, which approximate photon scattering and reemission as diffusion process, thereby replacing a large number of interactions and packet travel with a single approximation step, if the interaction depth to the nearest cell wall passes a parameterized threshold, one can. Reductions in calculation times of about an order of magnitude at the cost of small errors on the order of a few percent were achieved. A second method to improve the accuracy in optically thick regions where photon counts are low is to use a partial diffusion approximation[6].

All these features have been added piecemeal and lack integration, i.e. they have not been made compatible with each other. This is another area that requires further work.

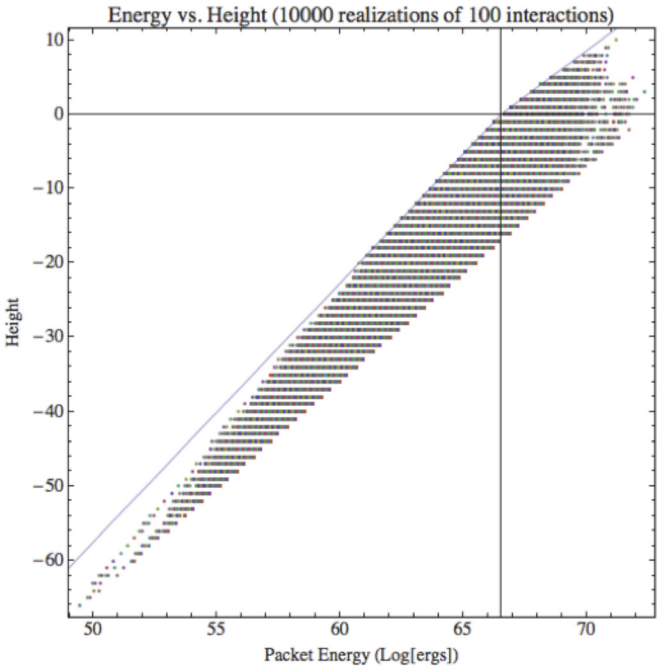

Figure 4 | Packet Energy Distribution. Packet energies after every interaction (up to 100 interactions) under the hemisphere weighting scheme favoring the lower hemisphere with b=1. Packets begin at zero height with 7.8 x 10^{28} ergs and move one unit up or down along a line depending on the randomly chosen hemisphere. The solid lines denote the energy path a packet continually biased down or continually biased up would follow. The possible energies for a given number of interactions lie along a straight line intersecting the two solid lines. Packet energies fall away from the solid lines due to the small probability (on the order of 2^N, where N is the number of interactions) of being directed up/down for a majority of the time. The expected frequency of each point is ${}_N\mathrm{C}_n(2/3)^n(1/3)^{N-n}$ where n denotes the number of times the packet went down.

1. Juvela, M. Efficient Monte Carlo Methods for Continuum Radiative Transfer. *Astron. Astrophys.* **440**, 531-546 (2006).
2. Dullemond, C. P. and Dominick, C. Flaring vs. Self-Shadowed Disks: The SEDs of Herbig Ae/Be Stars. *Astron. Astrophys.* **417**, 159-168 (2004).
3. Bjorkman, J. E., Wood, K. Radiative Equilbrium and Temperature Correction in Monte Carlo Radiative Transfer. *Astrophys. J.* **554**, 615-623 (2001).
4. Lucy, L. B. Computing Radiative Equilibria with Monte Carlo Techniques. *Astron. Astrophys.* **344**, 282-288 (1999).
5. Preibisch, T., Ossenkopf, V., Yorke, H. W., Henning, T. The Influence of Ice-Coated Grains on Protostellar Spectra. *Astron. Astrophys.* **279**, 577-588 (1993).
6. Min, M. *et al.* Radiative Transfer in Very Optically Thick Circumstellar Disks. *Astron. Astrophys*. **497**, 155-166 (2009).
7. Robitaille, T. P. On the Modified Random Walk Algorithm for Monte-Carlo Radiation Transfer. *Astron. Astrophys.* **520**, A70 (2010).

Acknowledgements This work was carried out at the Jet Propulsion Laboratory, California Institute of Technology, and was sponsored by the Summer Undergraduate Research Fellowship (SURF) and the National Aeronautics and Space Administration (NASA).