Cell Monte Carlo Radiative Transfer: A Research Note on an Idea for a Speed-up of Monte Carlo Radiative Transfer

David J. Jeffery[1]

ABSTRACT

Cell Monte Carlo radiative transfer (CMC) is a possibly novel speed-up method for Monte Carlo radiative transfer in optically thick (model) atmospheres which are discretized into (atmosphere) cells. The basic idea is that using ordinary Monte Carlo radiative transfer (OMC) with (photon) packets one calculates a set of CMC ingredients (see definitions below) for each interface-direction (see definition below) of the CMC cells which are atmosphere cells that are in general not the thermal-state cells (i.e., cells of homogeneous thermal state properties: i.e., composition, density, temperature, occupation numbers, etc.). The thermal-state cells can be chosen to be subcells of the CMC cells and we usually make use of this feature in our developments. Interface-directions are entities consisting of the interfaces between CMC cells and the directions at the interfaces a packet is being propagated in either by OMC or CMC. In OMC, the packet is just being propagated by OMC. In CMC, the packet is being propagated from an arrival interface-direction by random sampling from a runtime list (the first CMC ingredient which is just a list of exit states) to obtain an exit state for the packet. The exit state includes an exit adjacent interface-direction (which is an exit interface-direction from the combined cell which is the combination of the two CMC cells that surround the arrival interface-direction). The CMC propagation between the arrival and exit adjacent interface-directions is called a CMC step. The exit states are used for further propagation of the packet in OMC or CMC. The second CMC ingredient is the energy deposition array which tells how much energy (relative to packet energy) on average a packet deposits per thermal-state cell, per frequency bin, and per time bin in a CMC step. The deposited energy is used in the thermal-state calculation of a Lambda-iteration.

When the CMC ingredients for an interface-direction are not sufficiently complete for CMC steps, the packets are propagated by OMC and this propagation

[1]Department of Physics & Astronomy, University of Nevada, Las Vegas and the Nevada Center for Astrophysics (NCfA), 4505 S. Maryland Parkway Las Vegas, Nevada 89154, U.S.A. First complete version: 2023 Sep29. Last revision: 2024 Mar10. Presented as a poster paper at NCfA Symposium 2024 Feb28–Mar01, UNLV. For the research note, google “jeffery unlv educationl notes”.

is used to assemble the CMC ingredients. When the CMC ingredients for an interface-direction are judged sufficiently complete, the interface-direction is put in CMC which means that the ingredients are then used for CMC steps described above. Note that the runtime lists more heavily sample in a completely natural way packet exit states that are more important for the radiative transfer. As a CMC radiative transfer occurs, an increasing number of interface-directions are put into CMC and there is a progressive speed-up of the radiative transfer.

For 1-dimensional atmospheres, the CMC cells and thermal-state cells can be chosen to be the same though they do not have to be. For 3-dimensional atmospheres, thermal-state cells will necessarily have 3-dimensional character, but we argue in the text that the CMC cells should be chosen to have a 1-dimensional character, and so, of course, the two kinds of cells will not be the same cells.

In all cases in our discussions, we assume optically thick CMC cells are used for simplicity and uniformity. Since optically thin thermal-state cells must be used sometimes, they must necessarily be sub-cells of CMC cells in our discussions.

In the Lambda-iteration for a self-consistent atmosphere solution (which converges rapidly—in terms of iterations if not computer time—and robustly for Monte Carlo radiative transfer: Lucy 1999a; Kasen et al. 2006), the CMC ingredients must be re-calculated on every radiative transfer phase of the Lambda-iteration cycle. CMC has all generality and accuracy available with the OMC for the Lambda-iteration since one can always make the CMC ingredients more complete, and thus CMC can asymptotically approach the accuracy of OMC when OMC uses as many packets as one deems necessary for accurate convergence in a Lambda-iteration.

The speed-up expected from CMC relative to OMC is orders of magnitude in optimal cases for 1-dimensional atmospheres and is probably significant for 3-dimensional atmospheres, but demonstrating this requires numerical experimentation. CMC requires more overhead in memory and bookkeeping than OMC. However, we do not consider the extra cost of this overhead in this research note. The hope is that that cost is not so bad that it strongly negates the speed-up advantage of CMC.

To emphasize a key point simply, CMC is intended to be general speed-up of Monte Carlo radiative transfer and not just a speed-up for special cases.

Unified Astronomy Thesaurus concepts: Ejecta (453); Monte Carlo methods (2238); Novae (1127); Radiative transfer (1335); Radiative transfer equation (1336); Radiative transfer simulations (1967); Planetary atmospheres (1244);

Stellar atmospheres (1584); Stellar atmospheric opacity (1585); Stellar photospheres (1237); Stellar winds (1636); Supernovae (1668)

1. Introduction

To continue from the abstract, in § 2 we describe OMC and its advantages and disadvantages relative to radiative-transfer-equation radiative transfer (RTE). The advantages all carry over to CMC. In § 3, we describe the basics of CMC including interface-directions and the CMC ingredients (i.e., runtime lists and the energy deposition array). In § 4, we estimate the speed-up obtained from CMC relative to OMC for 1-dimensional atmospheres and discuss the speed-up possible for 3-dimensional atmospheres. A discussion of the further development of the CMC idea is given in § 5. Conclusions are given in the abstract.

2. Ordinary Monte Carlo Radiative Transfer and its Advantages and Disadvantages Relative to Radiative-Transfer-Equation Radiative Transfer

In ordinary Monte Carlo radiative transfer (OMC), in the version known from Lucy (1999a) and Kasen et al. (2006), indestructicle (photon) packets (which imitate real photons, but not in all respects) propagate through a cell depositing and recovering energy both of which are counted by counting the lengths of their (OMC) steps, not by explicit packet events. This counting procedure is a key variance reduction technique (e.g., Wikipedia: Variance reduction) for Monte Carlo calculations. Note the deposition and recovery of energy are built-in to be equal and this ensures energy conservation. The steps end in an event which is either an interaction with matter or the crossing of a cell interface (which may be with empty space). In the former case, the packet direction and frequency are changed by sampling the appropriate probability distributions. In the later case, one changes the cell to which the energy deposition is occuring. If the opacities of the atmosphere were magically correct initially, the first radiative transfer phase of a Lambda-iteration would already give the correct radiation field everywhere and that would solve for the correct thermal state everywhere in the first thermal state solution. By "correct", we mean correct to within the statistical error of Monte Carlo radiative transfer. The essential reasons for the fast (in number of iterations though maybe not computer time) and robust convergence is the built-in energy conservation and, that though opacities for an atmosphere vary wildly with thermal state, it seems they do not vary super wildly: the opacities are ultimately proportional to matter density though many other things too. So it seems the opacities are sufficiently insensitive to thermal state to allow for fast convergence. In general, the more insensitive

an iteration "function" is to the input values, the faster the convergenece.

The packets are injected from an inner boundary of an atmosphere or are created by radioactive energy deposition as in supernovae (e.g., Kasen et al. 2006). They escape to infinity from atmosphere surface. As aforesaid, the packets are indestructible. This means they are not destroyed in the atmosphere though they can be created by the aforementioned radioactive energy deposition. After injection, the packets are only lost at the outer boundary (to infinity) or, if it exists, at the inner boundary (where they just vanish). The packets can change energy due Doppler shifts relative to the comoving frame (i.e., the local rest frame) of the matter in moving atmospheres.

Monte Carlo radiative transfer has great advantages relative to using the radiative-transfer-equation radiative transfer (RTE): trivially infinitely parallizable (each packet is propagated independently), all complex photon-matter interactions can be easily coded, the mixed boundary conditions of atmospheres is no problem, 3-dimenional atmospheres are easily treated, time dependence is easily treated, moving atmospheres are easily treated, the Lambda-iteration (as aforesaid) converges rapidly (in terms of iterations if not computer time) and robustly (and it does not usually converge practically for optically thick atmospheres using RTE), and has relatively easy coding in general. There are two major disadvantages relative to using RTE. The first is that statistical error in Monte Carlo takes a very large number of packets to reduce to small size: e.g., the emission spectrum from an atmosphere will typically have a relative standard deviation per frequency bin dictated by the Poisson distribution (e.g. Bevington 1969, p. 53), and thus is $\sim 1/\sqrt{N}$, where N the number of packets emitted in that frequency bin. The statistical error can be beaten down by using more packets and various variance reduction techniques (e.g., Wikipedia: Variance reduction), an important one of which we mentioned above. The second major disadvantage is that Monte Carlo radiative transfer (OMC is meant) slows to impracticability for atmosphere with very large optical depths. The essential reason is that packets random walk through an atmosphere and the number of steps scales as τ^2, where τ is an effective average optical depth from inner boundary to the surface (see, e.g., § 4.1 below). Since computer time scales with number of steps, Monte Carlo radiative transfer increases in computation time with optical depth squared. Note that inner cells are usually set to be very optically thick in optically thick atmospheres and the packets in such cells do random walk with each step in optical depth being on average 1, of course. The packets are not going anywhere quickly and are effectively generating a lot very redundant information for the thermal state calculation phase of a Lambda-iteration.

3. The Basics of CMC

CMC is a speed-up method relative to OMC for optically thick atmospheres. It retains all the advantages of OMC discussed in § 2. In § 3.1, we present the basic idea of CMC. In § 3.2, we discuss a few fine points of CMC and in § 3.3, we expand on some details of CMC.

3.1. The Basic Idea of CMC

A CMC radiative transfer starts by using an OMC radiative transfer and then gradually converts to CMC as the CMC cell interface-directions acquire sufficiently complete CMC ingredients (i.e., runtime lists of exit states and energy deposition array) to put said interface-directions into CMC. We explicate the novel terms in last sentence below.

CMC cells are those for which we track the trajectories of packets. They may be identical to the thermal-state cells (i.e., atmospheres cells of homogeneous thermal state: composition, density, temperature, occupation numbers, etc.), but they do not have to be. However, we will assume that all CMC cells are optically thick for simplicity and uniformity in our discussions. Since thermal-state cells sometimes must optically thin (usually near the surfaces of atmospheres in order to capture the emission behavior), such thermal-state cells must necessarily be sub-cells of CMC cells. We discuss optically thin CMC cells in § 3.2.

Another assumption we make is that whether the atmosphere is 1-dimensional or 3-dimensional, we use only CMC cells having a 1-dimensional character. For plane-parallel atmospheres or 3-dimensional atmospheres with infinite extent in a 2 dimensions, we assign infinite planar slab CMC cells. For spherically-symmetric atmospheres or localizable 3-dimensional atmospheres, we assign spherically-symmetric CMC cells that are spherically-symmetric (or approximately spherically-symmetric if that is, in fact, needed for 3-dimensional atmospheres) which in both cases are centered on the point of highest symmetry. One reason for using 1-dimensional character CMC cells for 3-dimensional atmospheres is simplicity in conceptualizing and in coding CMC. A second reason is for better CMC speed-up as we will discuss below in § 4.4.

Note that for 3-dimensional atmospheres, the thermal-state cells in our discussions are necessarily not the same as the CMC cells since they must account for the 3-dimensional behaviors. In fact, for 3-dimensional atmospheres the thermal-state cells are often chosen to be cubes for convenient tiling of the atmosphere space. If cubical thermal-state cells are chosen as the basic thermal-state cells, then the interfaces of exactly spherically-symmetric CMC cells will cut through the cubical thermal-state cells and convert some into odd-shaped thermal-state cells that are sub-cells of the CMC cells. Perhaps, this cutting can be coded

efficiently. On the other hand, perhaps the CMC cells should be made only approximately spherically symmetric (a possibility noted above) to avoid cutting and just make the cubical thermal-state cells as they are sub-cells of the CMC cells. Further consideration of this point to find an optimum solution is beyond the scope of this research note.

Between any two CMC cells, there is an interface with only outward and inward directions (because we assume the CMC cells have a 1-dimensional character). As a shorthand, we define interface-direction to be an entity consisting of an interface and a direction. An interface-direction that has sufficiently complete CMC ingredients is put in CMC and one that does not is still in OMC.

The two CMC cells that interface at an interface, we call the combined cell. Say a combined cell internal interface has at least one interface-direction not in CMC. When a packet (by whatever process, CMC or OMC) arrives at that interface-direction (i.e., that interface and that direction), we propagate it by OMC until it exits the combined cell and record the exit state variables in a runtime list record. Note the packet may cross back and forth over the internal interface multiple times. Also note the two exit adjacent interface-directions have directions pointed away from the internal interface. As long as the arrival interface-direction is not in CMC, its runtime list continues to grow. As mentioned above, when an interface-direction acquires sufficiently complete CMC ingredients, it is put into CMC and packets are propagated from it to one or other of the two exit adjacent interface-directions in CMC steps as described below. The CMC steps give the CMC speed-up. As the radiative transfer in CMC continues, more and more interface-directions are put in CMC and the speed-up increases.

The runtime list array for all interface-directions is an array of the form $L(k, t_j, i)$ (where the bracketed variables are indices) and the array elements have type record $R(k', t'_j, \hat{n}', \nu', g', \vec{P}')$ (where the bracketed variables are components of the type record). Note the unprimed variables are arrival variables at an interface-direction of a combined cell and the primed variables are the exit variables for exit adjacent interface-directions. The variable specifications are k is the arrival interface-direction, t_j is the arrival time bin j (which is not needed for time independent calculations), i specifies the number of the record in the runtime list (e.g., $i = 1, 2, 3, \ldots, I$) which increments starting from 1 and runs to I (see below), k' is the exit interface-direction, t'_j is the exit time bin j (which is only needed for time dependent calculations), $\hat{n}'$ is the exit direction, ν' is the exit packet frequency, the g' factor is only needed for moving atmospheres (see below), and $\vec{P}'$ is the exit Stokes vector (which is only needed for polarized transfer). Note the maximum value of index i is I the total number of exit states when the runtime list and energy deposition array are sufficiently complete for CMC.

For static atmospheres, packets have identical invariant energy. However, for moving atmospheres, the packet energy changes as packet changes comoving frames. Thus, the packet energy must be kept track of until the atmosphere is completely in CMC when the runtime list does that automatically. The g' factor is the relative shift in packet energy in going from interface-direction k to k'. So the g' factor allows one to update the packet energy when a CMC step is made using the runtime list. Note that packet energies have to be scaled in some way in order to calculate the true energy deposition (e.g., Lucy 1999a; Kasen et al. 2006). Note also for time-dependent atmospheres, the packet time changes as the packet moves and must also be kept track of until the atmosphere is completely in CMC when the runtime list does that automatically.

When an interface-direction is put in CMC, one no longer propagates packets from it by OMC. Instead one samples the runtime list for the interface-direction to effect a CMC step. The sampling is done by randomly choosing an exit state record i from the runtime list as follows. Generate random number R and then set

$$i = \text{int}(RI) + 1 = \text{ceiling}[(1 - R)I] \ , \tag{1}$$

where we rely on the standard property of random number generators that $R \in [0, 1)$ which implies that R can be zero, but never 1. The packet then takes a CMC step which ends with the packet having the exit state properties of exit state record i which include the exit adjacent interface-direction.

We note that the runtime list does not depend on the arrival frequency and direction of packet. This is because that information is encoded implicitly in the runtime list if it is sufficiently complete. We are assuming that frequency and direction distribution of the packets arriving at an interface-direction while it is still in OMC are sufficiently complete. This should be the case if the set of CMC cells in CMC expands progressively through an atmosphere from the innermost CMC cell outward. This will usually be the case since actual photon density will usually decrease outward and the OMC mimics this. However, there is some error in assuming implicit encoding. However, if the criterion for sufficiently complete runtime list is made more strict, the error will decrease just as the overall error in using the runtime list decreases.

The second CMC ingredient, as aforesaid, is energy deposition array. For all interface-directions, there must be an energy deposition array that $E(k, t_j, \ell'', \nu'', t''_j)$ whose elements are the average relative amount of energy deposited by a packet arriving at interface-direction k in time bin t_j per deposition thermal-state cell ℓ'', per deposition frequency bin ν'', and per deposition time t''_j. The energy deposition array is calculated as the runtime list is assembled. In assembling the runtime list, the energy deposition is done as in OMC. In CMC, the energy deposition array is used to do the energy deposition every time a packet makes a CMC step.

How does one judge sufficient completeness of the CMC ingredients. Clearly, one necessary criterion for the sufficient completeness is when energy deposition array elements for arrival interface-direction k and time bin t_j stop changing on average to within some tolerance (e.g., $\epsilon = 0.01$): i.e., it has reached sufficient accuracy. A second criterion is that the calculated probability of packet being transferred by the runtime list to both of the two exit adjacent interface-directions has stopped changing to within some tolerance (e.g., $\epsilon = 0.01$): i.e., it has reached sufficient accuracy. A third criterion is that the distributions of exit t'_j and exit g' factor (when those are needed) are sufficiently well established. This may be when their average values and standard deviations for both of the two exit adjacent interface-directions stop changing to within some tolerance (e.g., $\epsilon = 0.01$): i.e., it has reached sufficient accuracy.

What of the distributions of exit $\hat{n}'$, ν', and $\vec{P}'$ as criteria? In fact, for the interface-directions (except for both those from the 2nd outermost interface which case we discus below), determining these distributions to anything more than very crude accuracy is probably not necessary in most cases. The reason for this is that we are assuming all the CMC cells are optically thick and we here make the extra assumption that they are optically thick in absorption opacity too. This means that throughout a CMC cell with packets being propagated by OMC the packet redistribution on an interaction event will in angle be nearly isotropic, in frequency redistribution determined by the (non-scattering) emissivity, and in polarization will be on average non-polarizing or depolarizing nearly. Thus, the exit $\hat{n}'$, ν', and $\vec{P}'$ at an exit adjacent interface-direction not in CMC have only a small effect on propagation through the combined cell that surrounds that exit adjacent interface-direction. Thus, what redistributions angle, frequency, and polarization happen at exit of the combined cell are relatively unimportant. It is the probability distribution for choosing exit adjacent interface-direction from an arrival interface-direction in CMC which is the important ingredient to the overall radiative transfer. Of course, one makes the distributions of exit $\hat{n}'$, ν', and $\vec{P}'$ better with longer runtime lists, and so those distributions can be improved if needed.

Note for the case where a CMC cell is not optically thick in absorption opacity, but only in scattering opacity, the frequency redistribution at an exit adjacent interface-direction needs to be well determined, and doing this should be a fourth criterion. This special case may not turn up often in realistic atmosphere problems, but must be well treated in testing a CMC code.

There are three special cases of interfaces. First is the inner-boundary interface which only exists if there is an inner boundary for the atmosphere. The inner-boundary has ouly an outward-going interface-direction and the photon packet injection properties (injection time, direction, frequency, g' factor, and polarization) are determined by the boundary conditions.

Assuming (as we do) that the innermost CMC cell is optically thick, most injected packets will not transfer to the outward-going exit adjacent interface-direction (i.e., will not escape), but will exit at the inner boundary itself. For the inner-boundary outward-going interface-direction, it therefore makes sense to count as one trajectory for the runtime list all the trajectories between escapes: i.e., only an escape ends a trajectory for the runtime list. Such trajectories for time-dependent atmospheres might have time discontinuities in them. But this is no problem in using them in the runtime list given our treatment of CMC steps (described above).

The second special case is the second outermost interface. Packets from the outward-going interface-direction for this interface can escape through the outer boundary. If one wants to use these escaping packets to calculate the emergent spectrum (rather than use formal integration and the source functions: see § 3.2), then the distributions for direction, frequency, and polarization for the outward-going interface-direction have be calculated for sufficient accuracy for the emergent spectrum rather than for the crude accuracy that most often is needed for deeper interface-directions (as explained above). The inward-going interface-direction needs no CMC ingredients since no packets arrive there from the outer-boundary interface as we discuss below for the third special case. All packets that arrive at the inward-going interface-direction are on trajectories that started from the the outward-going interface-direction and those trajectories do not end until the packet has exited the combined cell surrounding the second outermost interface.

The third special case is the outer-boundary interface itself. This case is simple since no CMC ingredients are needed. Packets that arrive its outward-going interface-direction just escape to infinity and no packets arrive at its inward-going interface-direction.

3.2. A Few Fine Points of CMC

A few fine points of CMC should be discussed:

1. For most atmospheres the transition of CMC cells from OMC to CMC will propagate from the deepest layer to the surface. We call this natural transition propagation. This is just because the density of packets will usually increase inward, and so the CMC ingredients should become sufficiently complete going outward. Natural transition propagation should be beneficial since it probably ensures that runtime list distributions do encode sufficiently complete distributions of arrival states at an interface-direction. With natural transition propagation, one does not have randomly placed interface-directions going into CMC with inadequate encoding of arrival states. Of course, there

are atmospheres for which natural transition propagation will not hold. However, this should not be a special problem since conditions for sufficiency can in always be increased in an atmosphere until CMC radiative transfer converges to full OMC radiative transfer. Convergence might take longer without natural transition propagation.

To explicate natural transition propagation a bit further, we consider two important atmosphere cases. First, if there is an inner boundary surface, natural transition propagation is plausible given that packet interactions generally decrease going outward as most packets return to the inner boundary and only a decreasing fraction leak outward with relatively few escaping at the outer boundary. Second, for localizable atmospheres, spherically symmetric or not, if there is no inner boundary, but the packet creation is still mainly in a compact deep interior (as for those supernova model atmospheres where the packets are created by a distribution of radioactive sources concentrated in a compact interior), the CMC cells are still probably mostly put into CMC going outward from the deep interior since that is where their packet density will be highest just by geometric concentration.

2. CMC radiative transfer gives the same Monte Carlo statistical behavior as (full) OMC radiative transfer in the asymptotic limit of sufficient completeness of CMC ingredients. Of course, the Monte Carlo statistical noise behavior between CMC radiative transfer and OMC radiative transfer will always be distinct since, among other reasons, the runtime list sampling does not depend on particular arrival states, but on only encoded distributions.

3. Calculating the emergent spectrum from escaping packets is a straightforward approach, but will give rise to statistical noise that takes large numbers of packets to strongly reduce. One can just smooth the emergent spectrum calculated this way, but that introduces some uncontrolled error. Another approach (referred to in § 3.1) is to use the source functions for CMC cells near the surface of the atmosphere to do a formal integration for the emergent spectrum which gives a physically-motivated averaging away of the statistical noise (e.g., Lucy 1999b). For the formal-integration spectrum, one needs to explicitly calculate the source functions (including the scattering source functions with, if needed for polarized transfer, Stokes vector) for the outermost CMC cells: only the outermost CMC cell if that is optically thick in all frequency bins. However, the formal-integration spectrum only needs to be calculated once for the solved atmosphere given by the Lambda-iteration. Calculating the formal-integration spectrum is a variance reduction technique (e.g., Wikipedia: Variance reduction) for Monte Carlo radiative transfer.

4. As the interface-directions in CMC accumulate in a CMC radiative transfer calculation,

there is an increasing speed-up of the radiative transfer through the whole atmosphere. Eventually, all interface-directions that one wishes to treat in CMC are in CMC.

5. Should all interface-directions be put in CMC? If one elects to have optically thin CMC cells, most optically thin CMC cells should have their inward-to-the-CMC-cell-pointing interface-directions left in OMC since ordinary OMC allows many packets to just fly through such CMC cells without interacting whereas being in CMC would cause propagation through such CMC cells to be a CMC step. Thus, for such CMC cells, CMC is a slow-down not a speed-up. However, we suggest not having any optically thin CMC cells so that the radiative transfer has a uniform treatment. However, optically thin thermal-state cells are often needed near the surface of atmosphere to obtain the correct emergent spectrum. However, as we suggested and assumed in § 3.1, one can make optically thin thermal-state cells sub-cells of optically thick CMC cells. We recommend this in order to have a uniform treatment as aforesaid.

6. There will very probably be cases where the radiative transfer becomes complete (as judged by some criterion) before all interface-directions are put in CMC and then the interface-directions not in CMC can be left not in CMC.

7. In a Lambda-iteration, the completeness of the CMC ingredients can be increased from one radiative transfer phase to the next as overall convergence to the exact atmosphere solution is approached, and therefore one can approach the exact atmosphere solution as closely as one likes just as with OMC. Of course, in both OMC and CMC, Monte Carlo statistical error cannot be reduced to zero in finite computer time. One runs the calculation to the desired accuracy in both cases.

8. A generalization of CMC for very optically thick atmospheres may be necessary if some CMC cells need to be very optically thick and determining the CMC ingredients for their interface-directions becomes a very slow procedure. In this case, CMC sub-cells with their own interface-directions could implemented. The CMC sub-cell interface-directions are then put into CMC when they have sufficiently complete CMC ingredients and they can then be used to accumulate the CMC ingredients to put the CMC cell interface-directions in CMC. Perhaps even a hierarchy of CMC sub-cells might be needed. Further consideration of CMC sub-cells and hierarchies of CMC sub-cells is beyond the scope of this research note.

9. There may be a general and efficient way of using growing packet trajectories in runtime lists that start in small CMC cells, much smaller than the thermal-state cells. The runtime lists first become sufficiently complete for nearest neighbor CMC cells and CMC steps are used to transverse these small CMC cells. But the trajectories from

different CMC cells are progressively concatenated and there is a gradual general speed-up of the radiative transfer that can efficiently deal both with very optically thick and thin thermal-state cells. This hypothetical trajectory concatenation method needs further thought which is beyond the scope of this research note.

10. CMC requires more overhead in memory and bookkeeping than OMC. We do not consider further the extra cost of this overhead in this research note. Hopefully, that cost is not so bad that it strongly negates the speed-up advantage of CMC.

3.3. Some Details of CMC

Now we expand on some details of CMC.

As a packet propagates through the combined cell, one sums all relative energy deposition in a temporary energy deposition array with bins indexed by the thermal-state cell (which is not necessarily the same the CMC cell as discussed above in § 3.1), frequency, and time. These temporary energy deposition bins are consistent in cell/frequency/time discretation with the energy deposition array, and with the bins needed for total energy deposition array for a radiative transfer run. Note the packet can change time bin one or more times while it propagates through the combined cell. For LTE without requiring scattering source functions, there is no frequency index for the temporary energy deposition array, the energy deposition array, and the total energy deposition array since there is only one energy bin for all depositied energy. For NLTE and/or if one requires scattering source functions, a frequency index is required. As discussed in § 3.2, the scattering source functions are needed at least for CMC cells near the surface of the atmosphere if one wants to calculate a formal-integration spectrum for the atmosphere rather than rely on the spectrum derived from the emergent packets (e.g., Lucy 1999b). Of course, for a pure scattering atmosphere and without needing the scattering source functions, one does not need energy deposition at all.

At the completion of the runtime list for an interface-direction, the temporary deposition array is converted to the energy deposition array (which is the average energy depositon per packet, per thermal-state cell, per frequency bin, and per time bin). Once an interface-direction is put in CMC, the energy deposition array is used every time a packet makes a CMC step to update to the total energy deposition array: i.e, the energy deposition array is added to the total energy deposition array. The total energy deposition array is used for the next thermal state calculation of the Lambda-iteration. As mentioned in § 3.1, packet energies have to be scaled in some way in order to calculate the true energy deposition (e.g.,

Lucy 1999a; Kasen et al. 2006).

We can expand here on one criterion for the sufficient completeness of the CMC ingredients mentioned in § 3.1. This criterion is having sufficiently accurate exit probability distribution for the exit adjacent interface-directions from the combined cell that surrounds an arrival interface-direction. The exit probability distribution (except for an inner-boundary interface-direction: see below) is, in fact, a binomial distribution (e.g. Bevington 1969, p. 53) with relative standard deviation σ/μ (where σ is standard deviation and μ is the mean for count of packet exits) and its estimate relerr after n trials (i.e., n starts of packet from an arrival interface-direction). The formulae for σ/μ and relerr are, respectively,

$$\frac{\sigma}{\mu} = \begin{cases} \sqrt{\dfrac{1-p}{np}} & \text{in general;} \\ \dfrac{1}{\sqrt{n}} & \text{for } p = 1/2\text{, the equal probability for the two events;} \\ \dfrac{1}{\sqrt{np}} & \text{for } p << 1\text{, the Poisson limit case} \end{cases}$$

and

$$\text{relerr} \approx \begin{cases} \sqrt{\dfrac{1-N/n}{N}} & \text{in general;} \\ \dfrac{1}{\sqrt{n}} \approx \dfrac{1}{\sqrt{2N}} & \text{for } p = 1/2\text{, the equal probability for the two events with the first formula being actually exact;} \\ \dfrac{1}{\sqrt{N}} & \text{for } p << 1\text{, the Poisson distribution limit case,} \end{cases} \tag{2}$$

where p is the probability of exit event E (which is either of the two events) and N is the number of exit events E after an actual run of n trials. The estimate relerr is the metric for goodness. When relerr is sufficiently small and other criteria are met, the arrival interface-direction can be put in CMC. For example, one could require relerr $\leq 10^{-2}$ for the both of the possible events E. In this case, N would have to be of order 10^4.

The inner-boundary interface-direction is a special case since here the exit probability distribution has only 1 possible exit event with probability 1 given how we treat the packet with discontinuous trajectories (see § 3.1 above). So here, one can use a temporary exit probability distribution which is a binomal distribution based on the events of the packet returning to the inner boundary (i.e., returns) or escaping outward from the top of the

innermost CMC cell without returns. The metric relerr used for the ordinary exit probability distributions given above can then be applied to the temporary exit probability distribution.

4. Estimanting the CMC Speed-Up

In this section, we review from some general random walk results for 1-dimensional atmospheres with the simplifying assumptions of having a single frequency and constant pure-scattering opacity (see § 4.1), estimate the CMC speed-up relative to OMC for 1-dimensional atmospheres (see § 4.2), for reference give all interesting cases of the computer time function defined in § 4.2 (see § 4.3), and discuss the CMC speed-up obtainable for 3-dimensional atmospheres (see § 4.4).

4.1. Some General Random Walk Results

For §§ 4.2 and 4.4, we need to review some general random walk results for 1-dimensional atmospheres with the simplifying assumptions of having a single frequency and constant pure-scattering opacity. The radial outward variance of distance traveled from an arrival surface for a random walk after n steps is given by

$$\begin{aligned}
\langle r_n^2 \rangle &= \langle (r_{n-1} + \Delta r)^2 \rangle && \text{where } \Delta r \text{ is a random walk step.} \\
\langle r_n^2 \rangle &= \langle r_{n-1}^2 \rangle + 2\langle r_{n-1} \Delta r \rangle + \langle \Delta r^2 \rangle \\
\langle r_n^2 \rangle &= \langle r_{n-1}^2 \rangle + \langle \Delta r^2 \rangle && \text{assuming } r_{n-1} \text{ and } \Delta r \text{ are uncorrelated.} \\
\langle r_n^2 \rangle &= n\ell^2 \times \text{constant} && \text{where } \ell \text{ is mean free path and the} \\
& && \text{constant is determined as specified just below.}
\end{aligned}$$

$$\langle r_n^2 \rangle = n\ell^2 \times \begin{cases} \dfrac{2}{3} & \text{Case 1: for the distance } s \text{ traveled in a beam path} \\ & \text{in any direction with scattering} \\ & \text{probability density } e^{-s/\ell}/\ell \\ & \text{and isotropic scattering. This case is how ordinary photons} \\ & \text{propagate, and so is how OMC behaves for packets.} \\ \dfrac{1}{3} & \text{Case 2: for constant distance } s = \ell \text{ traveled} \\ & \text{in a beam path in any direction and isotropic scattering.} \\ 1 & \text{Case 3: for constant distance traveled } s = \ell \text{ with only} \\ & \text{radially outward and inward scattering with} \\ & \text{equal probability.} \end{cases} \tag{3}$$

The number of computer operations and therefore the computer time needed to complete a radiative transfer calculation scales approximately as the n needed for the variance divided by ℓ^2 to increase to the square of the radial optical depth of the atmosphere $\tau^2 = (r/\ell)^2$, where r is the physical thickness of the atmosphere. Thus,

$$\text{approximate computer time} \propto n = \tau^2 \times \begin{cases} \frac{3}{2} & \text{Case 1.} \\ 3 & \text{Case 2.} \\ 1 & \text{Case 3.} \end{cases} \tag{4}$$

We make use of the three cases below in §§ 4.2 and 4.4.

4.2. 1-Dimensional Atmospheres

First, we divide the (1-dimensional) atmosphere into M CMC cells which are layers perpendicular to the outward direction. The optical depth of the atmosphere is τ in some kind of average opacity. Since we are only interested in optically thick atmospheres $\tau >> 1$. We assume the CMC cells have equal optical depth (in the average opacity), and so each CMC cell has optical depth τ/M. We also assume the CMC cells individually can be treated as optically thick (i.e., as has having $\tau/M >> 1$), but this assumption may not actually hold in cases of very large M. We do not worry about this complication since we are just doing an estimate of ideal speed-up where the M values of practical interest to not get so large.

For ordinary OMC, Case 1 (see § 4.1) applies and the computer time period to establish the radiative transfer through a CMC cell scales as $(3/2)(\tau/M)^2$ as follows from the discussion in § 4.1. However, this computer time period is not a single block of computer time in general, but is spread out into snippets of various durations in computer time of the whole radiative transfer calculation. But at our level of approximation, this indeterminacy of computer time period is of no importance. We add up the scaling terms $(3/2)(\tau/M)^2$ without regard to the indeterminacy in ordering of computer time snippets to get a scaling proportional to be putting all CMC cells in CMC. (Putting a CMC cell in CMC is a shorthand for putting its interface-directions in CMC as described above in § 3.1.) Note once a CMC cell is put in CMC, the computer time to propagate a packet across it (i.e., a CMC step) is approximated as the same for a single OMC step.

Once the atmosphere is all in CMC, the CMC cells are unit "optical depths" for a random walk with only radial steps through an atmosphere of "optical depth" M. As simplification, we assume an equal probability for radially outward and inward CMC steps, and so Case 3 applies (see § 4.1). Thus, the computer time for the radiative transfer through the whole atmosphere in CMC scales as M^2. However, we emphasize in actual cases using CMC,

the radiative transfer through the atmosphere is being done as the CMC cells are being progressively put into CMC, and so the computer time for the radiative transfer through the whole atmosphere is overestimated to some degree by using M^2 for the scaling.

We now define as a metric for CMC what we call the computer time function CTF by

$$\mathrm{CTF}(M) = M^2 - 1 + M\left(\frac{3}{2}\right)\left(\frac{\tau}{M}\right)^2 = M^2 - 1 + \left(\frac{3}{2}\right)\frac{\tau^2}{M} \ , \tag{5}$$

where we have just summed the contributions of the scalings discussed above. Note we have added a -1 term to the CTF to make the 1 cell limit for CMC equal to the analogue computer time function for OMC where computer time scales as $(3/2)\tau^2$ as follows from the discussion in § 4.1. Building in this equality seems reasonable so that the $M \to 1$ limit of the CTF is the OMC analogue CTF.

The CTF allows us to estimate the speed-up of CMC. Of course, the CTF decreases as speed-up increases. However, it seems most intelligible to analyze in terms of CTF in the first instance and then derive the time reduction factor (TRF) and the speed-up factor (SUF) which factors we define below.

Note we are neglecting any overall factor for the CTF (i.e., Equation (5)), and so the CTF can only be used for determining relative speed-ups for specified atmospheres: i.e., specified plane-parallel atmopsheres or spherically symmetric atmospheres. We are also neglecting coefficients for the terms in CTF which could be order of magnitude or more, but we are just trying to get a rough estimate of how the computer time scales with M and τ, and in any case those coefficients can only be determined for specified atmospheres and we are trying to be general. Note also that the CTF, following from a point we made above, will to some degree overestimate computer time scaling since establishing the radiative transfer through the atmosphere makes use of the OMC through CMC cells before they are put into CMC. For 1-dimensional atmospheres, the overestimate is probably relatively unimportant, but for 3-dimensional atmospheres we believe it is important as we discuss below in § 4.4 where we consider the speed-up for 3-dimensional atmospheres.

To explicate the behavior of the CTF, we will treat M as a continuous variable. Obviously, and with no computational relevance, $\mathrm{CTF}(M = 0) = \infty$. Also obviously, $\mathrm{CTF}(M = 1) = (3/2)\tau^2$ which is the OMC limit as discussed above. In fact, as M increases above 0, the CTF strictly decreases until it reaches a minimum $\mathrm{CTF}(M_{\min})$ at $M_{\min}$ determined as follows:

$$\frac{d\mathrm{CTF}(M)}{dM} \ = \ 2M - \left(\frac{3}{2}\right)\frac{\tau^2}{M^2} = 0 \ , \qquad M^3 = \left(\frac{3}{4}\right)\tau^2 \ , \qquad M_{\min} = \left(\frac{3}{4}\right)^{1/3}\tau^{2/3} \ ,$$

$$\mathrm{CTF}(M_{\min}) = \left(\frac{3}{4}\right)^{2/3}\tau^{4/3} - 1 + \left(\frac{3}{2}\right)\frac{\tau^2}{(3/4)^{1/3}\tau^{2/3}} = 3\left(\frac{3}{4}\right)^{2/3}\tau^{4/3} - 1 \ , \tag{6}$$

where recall we have assumed $\tau >> 1$. For $M > M_{\min}$, the CTF strictly increases as $M \to \infty$. Thus, there is only a range of M between points where $\mathrm{CTF}(M) = (3/2)\tau^2$ where we can obtain speed-ups relative to OMC. We solve for the 3 points (which we call M_1, M_2 and M_3) where $\mathrm{CTF}(M) = (3/2)\tau^2$ as follows:

$$\begin{aligned}
\frac{3}{2}\tau^2 &= M^2 - 1 + \left(\frac{3}{2}\right)\frac{\tau^2}{M} \\
0 &= M^3 - M\left(1 + \frac{3}{2}\tau^2\right) + \frac{3}{2}\tau^2 \\
0 &= (M-1)\left(M^2 + M - \frac{3}{2}\tau^2\right) \\
M &= \begin{cases}
M_1 = 1 & \text{which solution we already knew and is the lower bound of the speed-up range.} \\
M_2 = \dfrac{-1+\sqrt{1+6\tau^2}}{2} & \text{which is the upper bound of the speed-up range.} \\
M_3 = \dfrac{-1-\sqrt{1+6\tau^2}}{2} & \text{which is irrelevant to CMC since it is always negative.}
\end{cases}
\end{aligned} \tag{7}$$

The speed-up range is thus between M_1 and M_2.

Given the formalism above, it is best to choose $M = M_{\min}$. However, for actual atmosphere solutions, the CMC cells will usually not have equal optical depth in an average opacity since, among other things, one usually does not know any of the opacities before one starts calculating the solution. One can only determine an approximate $M_{\min}$ analogue for most actual cases and it is useful in determining the $M_{\min}$ analogue to have approximate analogues for M_1 and M_2 in order to stay far away from them, but still between them, of course.

However, assuming the CTF is exact, there is obviously a considerable speed-up for large optical depths τ by choosing $M = M_{\min}$ since $\mathrm{CTF}(M_1 = 1) \propto \tau^2$ (see Equation (5) above) and $\mathrm{CTF}(M_{\min}) \propto \tau^{4/3}$ (see Equation (6) above). We now define the time reduction factor (TRF) and its inverse the speed-up factor (SUF) by, respectively,

$$\mathrm{TRF}(\tau) = \frac{\mathrm{CTF}(M_{\min})}{\mathrm{CTF}(M=1)} = \frac{3(3/4)^{2/3}\tau^{4/3} - 1}{(3/2)\tau^2} \approx 2\left(\frac{3}{4}\right)^{2/3}\tau^{-2/3} \tag{8}$$

and

$$\mathrm{SUF}(\tau) \;=\; \frac{(3/2)\tau^2}{3(3/4)^{2/3}\tau^{4/3}-1} \approx \frac{1}{2}\left(\frac{4}{3}\right)^{2/3}\tau^{2/3}\;, \tag{9}$$

where the approximations are highly accurate for $\tau >> 1$ which we have assumed. We see that the speed-up factor increases with τ as $\tau^{2/3}$ and will be large for $\tau >> 1$.

To illustrate the speed-ups possible for 1-dimensional atmospheres in CMC, we present in Table 1 (given at the end of this research note) the speed-up factors (SUFs), time reduction factors (TRFs), and other results for two atmospheres specified by optical τ. As Table 1 shows, for $\tau = 10^2$, the SUF is about 13 and for $\tau = 10^6$, the SUF is about 6×10^3.

To conclude this subsection, we need to reiterate the point made at the end of § 3.2: CMC requires more overhead in memory and bookkeeping than OMC. These overheads are not considered in the above estimation formulae and their results (as above and in Table 1), and so actual speed-ups possible for 1-dimensional atmospheres will probably be less, maybe much less.

4.3. For Reference, the Computer Time Function (CTF) for All Interesting Cases

For reference, we display the computer time function (CTF) from § 4.2 for all interesting cases:

$$
\mathrm{CTF}(M) = \begin{cases}
M^2 - 1 + M\left(\frac{3}{2}\right)\left(\frac{\tau}{M}\right)^2 = M^2 - 1 + \left(\frac{3}{2}\right)\frac{\tau^2}{M} & \text{in general for CMC.} \\
\left(\frac{3}{2}\right)\tau^2 & \text{for } M = M_3 \text{ which occurs for} \\
\qquad M_3 = \dfrac{-1-\sqrt{1+6\tau^2}}{2} & \text{which is irrelevant to CMC} \\
& \text{since it is always negative.} \\
\qquad\quad = -\sqrt{\frac{3}{2}}\tau\sqrt{1+\frac{1}{6\tau^2}} - \frac{1}{2} & \\
\qquad\quad = -\sqrt{\frac{3}{2}}\tau - \frac{1}{2} - \sqrt{\frac{3}{2}}\frac{1}{12\tau} + \ldots & \\
\infty & \text{for } M = 0. \\
\left(\frac{3}{2}\right)\tau^2 & \text{for } M_1 = 1 \text{ which} \\
& \text{is the lower bound} \\
& \text{of the speed-up range.} \\
3\left(\frac{3}{4}\right)^{2/3}\tau^{4/3} - 1 = (2.476\ldots)\times\tau^{4/3} - 1 & \text{is the minimum value} \\
& \text{which occurs for} \\
\qquad M_{\min} = \left(\frac{3}{4}\right)^{1/3}\tau^{2/3} = (0.90856\ldots)\times\tau^{2/3} & \\
\left(\frac{3}{2}\right)\tau^2 & \text{for } M = M_2 \text{ which occurs for} \\
\qquad M_2 = \dfrac{-1+\sqrt{1+6\tau^2}}{2} & \text{which is the upper bound of} \\
& \text{the speed-up range.} \\
\qquad\quad = \sqrt{\frac{3}{2}}\tau\sqrt{1+\frac{1}{6\tau^2}} - \frac{1}{2} & \text{For } M > M_2, \text{ one gets a} \\
\qquad\quad = \sqrt{\frac{3}{2}}\tau - \frac{1}{2} + \sqrt{\frac{3}{2}}\frac{1}{12\tau} + \ldots & \text{slow-down, not a speed-up.} \\
M^2 & \text{for } M >> M_2 \approx \sqrt{3/2}\,\tau \text{ which} \\
& \text{is far in the slow-down region.} \\
\left(\frac{3}{2}\right)\frac{\tau^2}{M} & \text{for fixed } M \text{ with } \tau \\
& \text{sufficiently large.}
\end{cases}
\tag{10}
$$

The last case in Equation (10) shows that for sufficiently large τ, there is always a time reduction (i.e., speed-up) for CMC with $M \geq 2$, but the CTF grows in this case as τ^2 rather than as $\tau^{4/3}$ for CTF($M = M_{\min}$).

4.4. 3-Dimensional Atmospheres

Recall from § 3.1, that we assume CMC cells with a 1-dimensional character for 3-dimensional atmospheres. So it is only the thermal-state cells that have a 3-dimensional character for 3-dimensional atmospheres.

As stated in § 3.1, one reason for using 1-dimensional character CMC cells for 3-dimensional atmospheres is simplicity in conceptualizing and in coding CMC. A second reason is for better CMC speed-up since the thermal-state cells that are subcells of a 1-dimensional CMC cells will have an average behavior that can be encoded faster than their individual detailed behavior. That average behavior will probably be adequate to converge a Lambda-iteration, except maybe near final convergence.

To recapitulate from § 3.1, there are two cases to consider. First, for 3-dimensional atmospheres that are essentially planar (i.e., have an infinite extent in 2 dimensions), we use planar slab CMC cells. Second, for 3-dimensional atmospheres that are localizable, we use spherically symmetric (or approximately spherically symmetric) shell CMC cells centered on the point of highest symmetry of the 3-dimenstional atmospheres.

For the case of the essentially planar 3-dimensional atmospheres, we assume periodic boundary conditions so that planar slab CMC cells are actually finite in extent. We now assume the 3-dimensional atmosphere computer time function

$$\mathrm{CTF}_{\text{3-d}}(M) = M^2 - 1 + CM\left(\frac{3}{2}\right)\left(\frac{\tau}{M}\right)^2 = M^2 - 1 + C\left(\frac{3}{2}\right)\frac{\tau^2}{M} \tag{11}$$

which is almost the same as our original 1-dimensional atmosphere computer time function (CTF) Equation (5) given in § 4.2. The difference between the two computer time functions is the factor of C which is of order the number of thermal-state cells in each planar slab CMC cell. The extra computer time taken for the OMC phase of CMC to adequately encode the radiative transfer behavior of the thermal-state cells is the motivation of the factor C. Since we can have $C >> 1$, we find that we can have $\mathrm{CTF}_{\text{3-d}} >>$ CTF. We will not attempt a detailed estimate of speed-up for essentially planar 3-dimensional atmosphere case, but clearly it will be less, maybe much less than, than for 1-dimensional atmsopheres. However, we still expect a significant speed-up for CMC for two reasons. The first reason is that as the atmosphere solution proceeds, CMC cells convert from OMC to CMC, and therefore there

must be a speed-up provided the computational overhead of using CMC is not too large. The second reason is just the second reason for using 1-dimensional CMC cells given above.

For the case of localizable 3-dimensional atmospheres, we assume that to capture the behavior of the thermal-state cells their number will have to be of order the cube of the number of CMC cells (i.e., M^3) so that they tile 3-dimensional space adequately. This leads to a 3-dimensional atmosphere computer time function of the form

$$\mathrm{CTF}_{\text{3-d}}(M) = M^2 - 1 + M^3 \left(\frac{3}{2}\right) \left(\frac{\tau}{M}\right)^2 = M^2 - 1 + M \left(\frac{3}{2}\right) \tau^2 \tag{12}$$

which, like Equation (11) above, is almost the same as our original 1-dimensional atmosphere computer time function (CTF) Equation (5) given in § 4.2. The difference between the two computer time functions is, of course, the factor of M^3 (in the first of form of Equation (12)) which is as aforesaid of order the number of thermal-state cells needed to tile the 3-dimensional space adequately, and so is the factor of increase needed in encode the radiative transfer behavior of the thermal-state cells. Equation (12) increases strictly with M for the CMC-relevant $M \geq 1$ region, and so it seems there is always a slow-down not a speed-up for using CMC with $M > 1$. However, we still expect a significant speed-up for CMC for the same two reasons cited above for the essentially planar 3-dimensional atmosphere case.

The upshot of our analysis of the two 3-dimensional atmosphere cases is that we expect a significant speed-up for CMC for 3-dimensional atmospheres. However, it will take further analysis to obtain a quantitative estimate of the speed-up. In fact, perhaps the speed-up can only be quantitatively estimated by numerical experiments.

5. Discussion

In order to further develop the CMC idea, numerical experiments with CMC are required. A first experiment would be applying CMC for the solution of the plane-parallel grey atmosphere (e.g., Mihalas 1978, p. 53–76; Hubeny & Mihalas 2015, p. 569–588).

Recall, conclusions are given in the abstract.

Support for this work has been provided the Department of Physics & Astronomy of the University of Nevada, Las Vegas and the Nevada Center for Astrophysics (NCfA).

REFERENCES

Bevington, P. R. 1969, Data Reduction and Error Analysis for the Physical Sciences (New York: McGraw-Hill Book Company)

Hubeny, I., & Mihalas, D. 2015, Theory of Stellar Atmospheres (Princeton, New Jersey: Princeton University Press)

Kasen, D., Thomas, R.C., & Nugent, P. 2006, ApJ, 651, 366

Lucy, L. B. 1999a, A&A, 344, 282

Lucy, L. B. 1999b, A&A, 345, 211

Mihalas, D. 1978, Stellar Atmospheres, (San Francisco: W. H. Freeman and Company)

Table 1. Speed-Up Factors (SUFs), Time Reduction Factors (TRFs), Etc. for 1-Dimensional Atmospheres Solved Using CMC

τ	$M_{\min}$	$\frac{\tau}{M_{\min}}$	CTF($M = 1$) $= \frac{3}{2}\tau^2$	CTF($M_{\min}$)	SUF TRF	M_2
10^2	19.6	5.11	1.50×10^4	1.15×10^3	13.1 0.0766	122
10^6	9.09×10^3	110.0	1.50×10^{12}	2.48×10^8	6.06×10^3 1.65×10^{-4}	1.22×10^6