

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

David J. Jeffery¹

ABSTRACT

I've become discouraged with this idea. It is cute and might work for stationary 1-d atmospheres, but making it more general seems cumbersome and inefficient (2026feb06).—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-time (IDT: 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 make that choice when thermal-state cells need to be optically thin for reasons given below. Interface-direction-times (IDTs) are entities consisting of the interfaces between CMC cells, the directions at the interfaces a packet is being propagated in either by OMC or CMC, and, for time-dependent atmospheres, the arrival time of at the interface. In OMC, the packet is just being propagated by OMC. In CMC, the packet is being propagated from an arrival IDT by sequential sampling from a runtime list (the first CMC ingredient which is just a list of exit states for each IDT calculated during the runtime of the atmosphere solution and saved in runtime list array) to obtain an exit state for the packet. The exit state includes an exit adjacent IDT (which is an exit IDT from the combined cell which is the combination of the two CMC cells that surround the arrival interface of the IDT). The CMC propagation between the arrival and exit adjacent IDTs is called a CMC step. The exit states are used for further propagation of the packet in OMC or CMC. The second CMC ingredient

¹Department 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. Presented as a poster paper at NCfA Symposium 2024 Feb28–Mar01, UNLV.

is the CMC energy deposition array which tells how much energy (relative to packet energy) on average a packet deposits per thermal-state cell, per frequency bin, per time bin in a CMC step. The deposited energy is put in a total energy deposition array which is used in the thermal state solution phase of a Lambda-iteration. Note the total energy deposition array is not a CMC ingredient since it is required for the thermal state solution phase of a Lambda-iteration that uses any radiative transfer method for the solution of the radiative transfer solution phase of a Lambda-iteration.

When the CMC ingredients for an IDT are not sufficiently complete for CMC steps, the packets are propagated by OMC and this propagation is used to assemble the CMC ingredients. When the CMC ingredients for an IDT are judged sufficiently complete, the IDT is put in CMC which means that the CMC 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 IDTs are put into CMC and there is a progressive speed-up of the radiative transfer.

One way of looking at the speed-up of CMC relative to OMC is to note that OMC requires calculating a lot of redundant information as packets random walk 1 unit of optical depth on average in each step in a random walk through optically thick thermal-state cells. By calculating the CMC ingredients, CMC is, among other things, a form of machine learning that allows you to avoid calculating much of that redundant information.

For 1-dimensional atmospheres, the thermal-state cells and CMC cells can be chosen to be the same though they do not have to be and as we stated above we make the thermal-state cells subcells of the CMC cells when the thermal-state cells need to be optically thin for reasons given below. 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 (particularly near the surfaces of atmospheres), they must necessarily be subcells 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 always in 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 solution phase of the Lambda-iteration cycle.

CMC has all generality and accuracy available with the OMC for the Lambda-iteration since we 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 we deem necessary for accurate convergence in a Lambda-iteration.

We derive estimation formulae to quantitatively describe and estimate the speed-up of CMC relative to OMC for ideal-case 1-dimensional atmospheres. For these atmospheres, the formulae may be accurate to about 1 order of magnitude and perhaps, with suitable adjustments, may have the same accuracy for general 1-dimensional atmospheres. The speed-up formulae, of course, depend on the number of CMC cells M (which is a calculational parameter) and on an average atmosphere optical depth τ (which is a physical parameter). One of the estimation formulae is for the M value that gives the maximum speed-up (in the jargon of this research note, a speed-up factor (SUF)) for the ideal-case atmospheres for a specified τ . Relying on the estimation formulae, the speed-up expected from CMC relative to OMC is orders of magnitude in optimal cases for 1-dimensional atmospheres. For 3-dimensional atmospheres, we argue for significant speed-up, but probably this speed-up can only be demonstrated and quantitatively estimated by numerical experiments.

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 high that it effectively cancels the speed-up advantage of CMC.

Given that there is an advantage in using CMC, another advantage is that CMC can be supplemented by speed-ups of OMC. For example, the machine learning sampling method of Charles & Chen (2024) that can speed up the sampling of distributions needed in OMC.

Lastly, to emphasize a key point simply, CMC is intended to be a 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 IDTs and the CMC ingredients (i.e., runtime lists and the CMC energy deposition array). In § 4, we derive the estimation formulae used to quantitatively describe and estimate the speed-up obtained from CMC relative to OMC for the aforementioned ideal-case 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. Monte Carlo Radiative Transfer and its Advantages and Disadvantages Relative to Radiative-Transfer-Equation Radiative Transfer

In Monte Carlo radiative transfer (meaning both OMC and CMC hereafter in this section, unless otherwise stated) in the version known from Lucy (1999a) and Kasen et al. (2006), indestructible (photon) packets (which imitate real photons, but not in all respects) propagate through a cell depositing and re-emitting energy both of which are from calculated from the lengths of their (OMC) steps, not by explicit packet events of which there are two kinds: see below. This calculating procedure is a key variance reduction technique (e.g., Wikipedia: Variance reduction) for Monte Carlo calculations. Note the deposition and re-emission of energy are built-in to be equal and this ensures energy conservation. The steps end in an event which is either a packet-matter interaction (scattering or absorption) or the crossing of a cell interface (which may be the cell interface to an atmosphere boundary: see more description below). In the former step-end case, the packet direction, frequency, and, if needed, polarization are changed by sampling the appropriate probability distributions. In the latter step-end case, we change the cell to which the energy deposition is occurring or, if the interface is a boundary, we remove the packet from the atmosphere. If the opacities of the atmosphere were magically correct initially, the first Monte Carlo radiative transfer solution phase of a Lambda-iteration (see description below) would already give the correct radiation field everywhere and that would solve for the correct thermal state everywhere in the first thermal state solution phase of a Lambda-iteration. By “correct”, we mean correct to within the statistical error of Monte Carlo radiative transfer. Thus, the atmosphere would be solved in one radiative transfer solution phase and one thermal state solution phase.

Where do the packets come from and where to they go to in Monte Carlo radiative transfer? The packets are injected from an inner boundary of an atmosphere and/or are injected by creation by interior radioactive energy interior sources or an interior initial-time

thermal energy source as in supernovae (e.g., Kasen et al. 2006, p. 4–5). 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 sources and initial-time thermal energy. 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). Though indestructible, the packets can change energy due to Doppler shifts relative to the comoving frame (i.e., the local rest frame) of the matter in moving atmospheres.

Actually, the case of radioactive energy sources requires some further explication. If we assume local thermodynamic equilibrium (LTE), then the energy from radioactive energy sources is deposited locally and treated just as an extra term in the LTE radiative equilibrium equation that determines the local temperature (e.g., Mihalas 1978, p. 172–173). However, if the atmosphere is treated in non-LTE (NLTE), then the energy deposition from radioactive sources may also need a NLTE treatment: i.e., the amount of radioactive source energy deposited in multiple frequency bins may need to be calculated and this calculation may need to include the calculation of the finite travel of gamma rays and non-thermal massive particles produced by directly by radioactive decay or pair production by gamma rays (i.e., high energy photons). A general approach may be to treat gamma rays just as ordinary photon packets and the non-thermal massive particles as special packets. However, such a general approach may be not be computationally efficient.

An important advantage of Monte Carlo radiative transfer over RTE is that it gives fast (in number of iterations though not necessarily in computer time) and robust Lambda-iteration convergence in the solution of optically thick atmospheres (Lucy 1999a; Kasen et al. 2006). RTE on the other hand, does not effectively converge if the atmosphere becomes too optically thick (e.g., with some extrapolation, Mihalas 1978, p. 147–150; Hubeny & Mihalas 2015, p. 381–385). The Lambda-iteration is an alternation of radiative transfer solution phases and thermal state solution phases. Each phase is the initial condition for the alternate phase. The alternation of phases continues to convergence: i.e., when both solutions become unchanging to within some relative tolerance.

The essential two factors for the fast and robust Lambda-iteration convergence with Monte Carlo radiative transfer that do not hold for RTE are the indestructibility of the packets and the built-in energy conservation of the indestructible packets. The first factor means that the packets can convey information about the changes in the thermal state of a particular thermal-state cell after a Lambda-iteration thermal state solution phase to all other thermal-state cells in the succeeding radiative transfer solution phase. In the Lambda-iteration using RTE, such information can only propagate to nearest neighbor thermal-state cells if the thermal-state are cells optically thick, and this slows down convergence to the

point of effective nonconvergence if the atmosphere is sufficiently optically thick (e.g., with some extrapolation, Mihalas 1978, p. 149–150). The second factor prevents the unphysical nonconservation of energy that happens in the Lambda-iteration steps using RTE: energy conservation is only reached asymptotically on convergence (Kasen et al. 2006, p. 10). The built-in energy conservation of the indestructible packets keeps Monte Carlo radiative transfer closer to physical realism than RTE and this helps speed convergence. Of course, this second factor is essential to the effectiveness of the first factor for fast convergence.

There is another factor essential to Lambda-iteration convergence that is general to Monte Carlo radiative transfer and RTE. This is the sufficient insensitivity of the opacities to the thermal state. The sufficient insensitivity makes the Lambda-iteration adhere sufficiently to the general rule the more insensitive an iteration “function” is to the input values, the faster the convergence. Note the opacities for an atmosphere can vary wildly with thermal state, but it seems not so wildly that Lambda-iteration convergence fails because of that wild variation as proven by convergence with Monte Carlo radiative transfer. One obvious reason for the sufficient insensitivity of opacities is that opacities are ultimately proportional to matter density even though the thermal state can superimpose the aforesaid wild variation.

Other important advantages that Monte Carlo radiative transfer has relative to RTE are that it is trivially infinitely parallelizable (since each packet is propagated independently), all complex photon-matter interactions can be easily coded, there is no problem in the with mixed boundary conditions (i.e., photon inflow to an atmosphere is specified, but not outflow), 3-dimensional atmospheres are easily treated, time dependence is easily treated, moving atmospheres are easily treated, and there is relatively easy coding in general.

However, there are two major disadvantages that Monte Carlo radiative transfer has relative to using RTE. The first is that statistical error (which does not occur at all with RTE) in Monte Carlo radiative transfer takes a very large number of packets to reduce it to small size: e.g., the emission spectrum from an atmosphere calculated just by counting packets will typically have a relative standard deviation (per frequency bin) dictated by the Poisson distribution (e.g., Bevington 1969, p. 53), and thus the relative standard deviation has the formula $\sim 1/\sqrt{N}$, where N is the number of packets emitted in a frequency bin. The statistical error can be beaten down by using more packets, but this rapidly becomes computationally expensive: e.g., to decrease the relative standard deviation given by $\sim 1/\sqrt{N}$ by a factor of order 10 requires using of order 100 more packets, and therefore requires 100 times more computer time. Statistical error can also be reduced by various variance reduction techniques (e.g., Wikipedia: Variance reduction), a key one of which we mentioned above (i.e., calculating energy deposition and re-emission from OMC step lengths rather than by packet events).

The second major disadvantage is that Monte Carlo radiative transfer (OMC is meant here) 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 an inner boundary and/or interior sources to the surface (see, e.g., below § 4.1). Since computer time scales with number of steps, Monte Carlo radiative transfer increases in computation time with optical depth squared (i.e., τ^2). 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 solution phase of a Lambda-iteration. As stated above in the abstract, CMC is, among other things, a form of machine learning that allows you to avoid calculating much of that redundant information.

3. The Basics of CMC

CMC is a speed-up method relative to OMC for optically thick atmospheres. It retains all the advantages of (general) Monte Carlo radiative transfer 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

Consider an atmosphere (i.e., an atmosphere model) that may be semi-infinite atmosphere (e.g., Mihalas 1978, p. 36) or extended atmosphere (i.e., finite localizable atmosphere: e.g., Mihalas 1978, p. 243). We only consider atmospheres where the outer boundary is empty space and packets exiting the outer boundary escape to infinity. Other kinds outer-boundary conditions can be treated in CMC, but for simplicity in this research note, we do not consider them further. Semi-infinite atmospheres must have an inner boundary and packets can be injected into the atmosphere from there. Localizable atmospheres may or may not have an inner boundary. Whether there is an inner boundary or not, sources of packets can be embedded in the atmosphere. The packets from embedded sources are created by a distribution of radioactive energy sources and/or from initial-time thermal energy.

The 1-dimensional case of semi-infinite atmospheres are plane-parallel atmospheres and of extended atmospheres are spherically-symmetric atmospheres.

For CMC, we divide the atmospheres into I CMC cells that have a 1-dimensional char-

acter: i.e., for semi-infinite atmospheres the CMC cells are parallel slabs perpendicular to the radial direction and for extended atmospheres they are spherical shells centered on the point of highest symmetry (which for spherically-symmetric atmospheres is the geometrical center). The CMC cells and their outer interfaces are numbered $i = 1, 2, 3, \dots, I$. If there is an inner boundary, it is interface $i = 0$. The CMC cells are used for assembling the CMC ingredients: i.e., runtime lists and the CMC energy deposition array both of which we will explicate below.

There are also thermal-state cells (i.e., cells of homogeneous thermal state properties: i.e., composition, density, temperature, occupation numbers, etc.). The thermal-state cells are used for the thermal state calculation of a Lambda-iteration. For 3-dimensional atmospheres, the thermal-state cells will in general will be subcells of the CMC cells. For 1-dimensional, the thermal-state cells can be coincident with the CMC cells or can be subcells.

For uniformity in our discussion, we will consider almost always??? only CMC cells that are optically thick. 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 subcells of CMC cells. We discuss optically thin CMC cells in § 3.2.???

We now define entities we call interface-time-locations (ITLs) which consist of an interface at time bin and, for 3-dimensional atmospheres, at a given location bin on the interface which we symbolize by Ω (which has an natural interpretation as angular position on spherical symmetric shell for extended atmospheres: i.e., those localizable in 3-dimensional space). There is no ITL at interface $i = I$ since packets that arrive there only arrive from below, and so always escape the atmosphere and can be used to calculate the emergent spectrum by the direct counting technique (DCT) on the last radiative transfer calculation of a Lambda-iteration.

The runtime lists are created for all ITLs. In OMC, a runtime list for an ITL is sampled in order to propagate a pack from a first ITL to a second ITL. The second ITL is one of the two adjacent interfaces to the interface of the first ITL, except if the first ITL has interface $i = 0$ for an atmosphere with an inner boundary or $i = 1$ for an atmosphere with no inner boundary. The first exception requires special treatment which we discuss below. The second exception just has one adjacent interface for a packet to be propagated to.

Before an OMC radiative transfer calculation, one must determine the runtime list of packet trajectories for each ITL. The runtime lists are collected in a runtime list array $R(i, t_j, \Omega_k, m)$ where for ijk are the indices for the ITL and m is the index for the runtime list for the ITL. Say a packet arrives at a first ITL specified by ijk , then one propagates

the packet to a second ITL specified by $i'j'k'$ by running sequentially $m = 1, 2, 3, \dots, M$, where M is the maximum number of trajectories in the runtime list for ITL ijk . One does not need to select from the M trajectories randomly since the trajectories were calculated randomly as described below.

The runtime list array elements have a type (a Fortran-95 type), which we specify by $RE(i', t_{j'}, \Omega_{k'}, g')$, where the g' factor is the relative change in energy of the packet going from ITL ijk to $i'j'k'$. The g' factor is only needed for moving atmospheres since there packet energies change due Doppler shifts.

In assembling the runtime list array elements, the element set for each ITL is found by propagating M packets from the ITL by OMC through what we call the combined cell. The combined cell consists of the two CMC cells adjacent to the interface, except for the special case of interfaces $i = 0$. For interface $i = 0$ (which does not exist if the atmosphere has not inner boundary), the combined cell is just the CMC cell $i = 1$. The trajectory for packet ends when it leaves the combined cell. The case with an interface $i = 0$ requires special treatment which we describe below.

The OMC propagation what ITL how is the initial frequency and direction to be chosen????

The frequency and direction of the

For very optically thick absorption dominated CMC cells, the frequency bins can be neglected since the particular frequency of a packet injected into a combined cell has a negligible effect compared to the frequency redistributions a packet undergoes propagating through the combined cell and these redistributions are determined by the lagged thermal state of the thermal-state cells. However, as the absorption optical depth is decreased (either because optical depth is decreasing generally or because absorption optical depth is decreasing even as total optical depth stays high), then frequency bins will be needed, the number depending on the case. For time-independent calculations, the time bins are not used. For 1-dimensional atmospheres, location bins are not used.

What of direction bins? They do not need to be considered as long as the CMC cells are very optically thick since the particular direction of a packet injected into a combined cell has a negligible effect compared to the directions redistributions a packet undergoes propagating through the CMC cell and these redistributions are determined by the lagged thermal state of the thermal-state cells. However, giving the packets an isotropic direction distribution for all injections seems best since that is the fairest simple distribution for optically thick conditions. Of course, if the optical depths of the CMC cells are made less than optically thick, then a direction distribution would need to be added to the ITLs: i.e., one would need

interface-frequency-time-directions (IFTDs). We will not consider further cases where the optical depths of the CMC cells are made less than optically thick.

CMC replaces OMC by propagating packets from one IFT to another one located an adjacent interface (i.e., one of the two that enclose the combined cell) using runtime lists of trajectories. Say a packet arrives at an IFT specified by indices i, ν_j, t_k , then one selects and propagates the packet by running sequentially $m = 1, 2, 3, \dots, M$ through the runtime list array $R(i, \nu_j, t_k, m)$ for i, ν_j, t_k . The runtime list array elements have a type (a Fortran-95 type), which we specify by $RE(i', \nu_{j'}, t_{k'}, g')$, where i' specifies the adjacent interface to which the packet is propagated, $\nu_{j'}$ specifies the arrival frequency bin interface i' (which as described above is not needed for very optically absorption dominate CMC cells), $t_{k'}$ is the arrival time bin at interface i' . and the g' factor is the relative change in energy of the packet going from interface i to i' .

(it can only be one of interfaces $i - 1$ or $i + 1$, except

(where the bracketed variables are the type components of the elements). Note the unprimed variables are arrival variables at an IDT of a combined cell and the primed variables are the exit variables for exit adjacent IDTs. The variable (index and component) specifications are i is the arrival interface (and also the CMC cell number), d is the arrival direction (which can only be inward -1 or outward $+1$), t_j is the arrival time bin j (which is not needed for time independent calculations), k specifies the element number in the runtime list (e.g., $k = 1, 2, 3, \dots, K(i, d, t_j)$) which increments starting from 1 and runs to $K(i, d, t_j)$ (see below), i' is the exit interface, d' is the exit direction, $t'_{j'}$ is the exit time bin j' (which is only needed for time dependent calculations), \hat{n}' is the exit unit vector that specifies the exit direction (here meaning in 3-dimensional space), 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). For simplicity in our discussion, we usually just say runtime list when we mean the set of $R(i, d, t_j, k)$ elements for a fixed IDT (i.e., for fixed i, d , and, for time-dependent atmospheres, t_j).

A CMC radiative transfer calculation starts by using OMC radiative transfer and then gradually converts to CMC radiative transfer as the CMC cell IDTs individually acquire sufficiently complete CMC ingredients (i.e., for each IDT, a sufficiently complete runtime list of exit states and the CMC energy deposition array element) to put said IDTs 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 subcells 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 (which are 1-dimensional atmospheres) or 3-dimensional atmospheres with infinite extent in a 2 dimensions, we assign infinite planar slab CMC cells. For spherically-symmetric atmospheres (which are 1-dimensional atmospheres) or localizable (and therefore finite) 3-dimensional atmospheres, we assign spherically-symmetric CMC cells (which would be spherically-symmetric shells) 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. For 3-dimensional atmospheres with infinite extent in 2 dimensions, the thermal-state cells could be chosen to be cubes that in each CMC cell have the same size. For localizable 3-dimensional atmospheres, the thermal-state cells could be chosen to be spherical-coordinate system “boxes” for each spherically-symmetric CMC cell. Other choices for the shape of the thermal-state cells are possible (e.g., Lauwers et al. 2024). A thermal-state cell should be in only one CMC cell, and so the thermal-state cells should never be cut by a CMC cell boundary. Further consideration of the shape of thermal-state cells 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 in all cases as discussed above). As a shorthand, we define IDT to be an entity consisting of an interface, a direction, and, for time-dependent atmospheres, a time. The direction and time are needed since when a packet arrives at an interface it does so in particular direction (here meaning only either inward or outward) and, for time-dependent atmospheres, at a particular time. The direction and time are needed in specifying the CMC ingredients. As a CMC radiative transfer calculation proceeds, an IDT that acquires sufficiently complete CMC ingredients is put in CMC and one that has not acquired sufficiently complete CMC ingredients 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 some IDTs not in CMC. When a packet (by whatever

process, CMC or OMC) arrives at one of those IDTs, we propagate it by OMC until it exits the combined cell and record the exit state in a runtime list (which is saved in runtime list array: see below). Note the packet may cross back and forth over the internal interface multiple times before making an exit. Also note the exit adjacent IDTs (of the combined cell) have directions pointed away from the internal interface. As long as the arrival IDT is not in CMC, its CMC ingredients continue to evolve to completion. As mentioned above, when an IDT acquires sufficiently complete CMC ingredients, it is put into CMC and packets are propagated from it to one of the exit adjacent IDTs in CMC steps as described below. The CMC steps give the CMC speed-up since each one replaces many OMC steps. As the CMC radiative transfer calculation continues, more and more IDTs are put in CMC and the CMC speed-up increases.

The runtime list array for all IDTs is an array of the form $R(i, d, t_j, i)$ (where the bracketed variables are indices) and the runtime list array elements have a type (a Fortran-95 type) which we specify by $RE(i', d', t'_{j'}, \hat{n}', \nu', g', \vec{P}')$ (where the bracketed variables are the type components of the elements). Note the unprimed variables are arrival variables at an IDT of a combined cell and the primed variables are the exit variables for exit adjacent IDTs. The variable (index and component) specifications are i is the arrival interface (and also the CMC cell number), d is the arrival direction (which can only be inward -1 or outward $+1$), t_j is the arrival time bin j (which is not needed for time independent calculations), k specifies the element number in the runtime list (e.g., $k = 1, 2, 3, \dots, K(i, d, t_j)$) which increments starting from 1 and runs to $K(i, d, t_j)$ (see below), i' is the exit interface, d' is the exit direction, $t'_{j'}$ is the exit time bin j' (which is only needed for time dependent calculations), \hat{n}' is the exit unit vector that specifies the exit direction (here meaning in 3-dimensional space), 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). For simplicity in our discussion, we usually just say runtime list when we mean the set of $R(i, d, t_j, k)$ elements for a fixed IDT (i.e., for fixed i , d , and, for time-dependent atmospheres, t_j).

Note the maximum value of index k is $K(i, d, t_j)$ which is the total number of exit states for an IDT when its runtime list and the CMC energy deposition array elements are sufficiently complete to put the IDT in CMC. As indicated by its arguments, $K(i, d, t_j)$ depends, of course, on the specification of the IDT.

Three points need to be made about packets. First, the 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). Second, for time-dependent atmospheres, the packet time must be kept track of as the packet moves. In OMC, the packet time needs to be updated on every OMC step.

In CMC, the runtime list element updates the packet time whenever a CMC step is made using the time bins t_j and t'_j . Third, for static atmospheres, packets have identical invariant energy. However, for moving atmospheres, the packet energy changes as the packet changes comoving frames. Thus, for moving atmospheres the packet energy must be kept track of. In OMC, this means packet energy must be updated for every OMC step. In CMC, the runtime list element updates the packet energy using the g' factor whenever a CMC step is made. The g' factor is the relative shift in packet energy in going from arrival IDT to exit IDT.

When an IDT is put in CMC, we no longer propagate packets from it by OMC. Instead we sequentially sample the runtime list for the IDT to choose a CMC step: i.e., we sample the runtime list (possible CMC steps) in order $k = 1, 2, 3, \dots K(i, d, t_j)$ and then start over at $k = 1$ after element $K(i, d, t_j)$ has been sampled. If the runtime list is sufficiently complete, this is a valid sampling procedure since each element is independent of the others. A random sampling of the runtime lists seems unnecessary to prevent favoring the first parts of the runtime lists since such favoring is uncorrelated between the runtime lists and it is their overall behavior that goes into the convergence of the overall radiative transfer solution. After the sampling of a runtime list, the packet takes the chosen CMC step (i.e., element k) which ends with the packet having the exit state properties of element k which include, of course, the exit adjacent IDT.

We note that the runtime list for an IDT does not depend on the arrival state of a packet. This is because the arrival state distribution is already encoded implicitly in the runtime list if the runtime list is itself sufficiently complete. Obviously, there is some error in assuming sufficient implicit encoding. However, if the criteria for sufficiently complete runtime list are made more strict, that error will decrease just as the overall error in using the runtime list decreases.

The second CMC ingredient, as aforesaid, is the CMC energy deposition array. The form of the CMC energy deposition array is $E(i, d, t_j, m'', \nu'', t''_{j''})$ and the elements are the average relative amount of energy deposited by a packet arriving at an IDT (specified by i, d, t_j) per deposition thermal-state cell m'' , per deposition frequency bin ν'' , and per deposition time $t''_{j''}$. The CMC energy deposition array elements for the arrival IDT are calculated as the runtime list for the IDT is assembled. While assembling the runtime list, the energy deposition is done as in OMC with the energy deposition going into a total energy deposition array which has just 3 indices m'' , ν'' , and $t''_{j''}$ which are the same as the last 3 of the CMC energy deposition array. In CMC, the CMC energy deposition array is used to do the energy deposition into the total energy deposition array every time a packet makes a CMC step. Thus, the total energy deposition array gets contributions from both the OMC and CMC

phases of the radiative transfer solution phase of the Lambda-iteration. The total energy deposition array is used to calculate the thermal state in the thermal state solution phase of the Lambda-iteration. To be more specific, the total energy deposition array is used to calculate the thermal state of each thermal-state cell m'' and for a time bin $t''_{j''}$ in the thermal state solution phase of the Lambda-iteration. Note the total energy deposition array is not a CMC ingredient since it used for the thermal state solution phase for a Lambda-iteration using any method of radiative transfer.

How do we judge sufficient completeness of the CMC ingredients for an IDT? Clearly, a first and necessary criterion for the sufficient completeness is when the calculated probabilities of packet being transferred by the runtime list for the IDT to the exit adjacent IDTs have stopped changing to within some relevant tolerance as the runtime list grows. One suggestion for this criterion is that the largest transfer probabilities to each of the two outer interfaces of the combined cell have stopped changing to within relative tolerance ϵ (e.g., $\epsilon = 0.01$). Additional criteria for the runtime list can easily be thought of. As for the CMC energy deposition array elements for the IDT, one criterion for completeness is when the larger CMC energy deposition array elements for the IDT stop changing on average to within some relative tolerance ϵ (e.g., $\epsilon = 0.01$). Additional criteria for the CMC energy deposition array elements for the IDT can easily be thought of.

What of the accuracies of exit distributions of \hat{n}' , ν' , and \vec{P}' as criteria of completeness for the runtime list for an IDT? In fact, for the IDTs (except for those from the 2nd outermost interface which case we discuss below), determining these exit distributions to anything more than very crude accuracy is probably not necessary in most cases, and so using their accuracy as a criteria is probably not necessary. The reason why only very crude accuracy is needed 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 the peculiar exit values of \hat{n}' and ν' have very little effect on the overall radiative transfer since, viewed as redistributions in angle and frequency, they are only a tiny fraction of the redistributions in angle and frequency that happen on packet-matter interactions inside the CMC cells and which are encoded in the runtime lists and in the CMC energy deposition array. The peculiar exit values will usually be statistically negligible. Note the exit values of \hat{n}' and ν' will also have only negligible dependence on the arrival state of a packet at an arrival IDT. As for polarization in optically thick CMC cells, it is nearly zero because of the large number randomly polarizing and/or depolarizing packet-matter interactions. Thus, the exit polarization will be nearly zero in all cases and of little importance in any case. Of course, we make the exit distributions of \hat{n}' , ν' , and \vec{P}' better with longer runtime lists, and so those exit distributions can be improved if needed. Note it is the probability distribution for choosing the exit adjacent IDT for an arrival IDT in CMC which is the important ingredient

to the overall radiative transfer.

What of the case where CMC cells are optically thick in scattering opacity, but not in absorption opacity? The exit frequency distribution at an exit adjacent IDT needs to be well determined in this case since that distribution strongly effects the subsequent propagation. Thus, ensuring the exit frequency distributions are high accuracy should be another criterion for the completeness of the CMC ingredients. Of course, the frequency distribution at an arrival IDT is encoded in the runtime list for the IDT if that runtime list is sufficiently complete. Note that exit angle and polarization distributions need only crude accuracy for optically thick CMC cells whether in absorption or scattering opacity or any combination since the arguments of the last paragraph apply in any case for these distributions. The case of optically thick CMC cells in scattering, but not absorption, opacity, may not turn up often in realistic atmosphere problems, but must be well treated in testing a general 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 only outward-going IDTs (i.e., there are no inward-going IDTs at the inner-boundary), the inward-going exit adjacent ITDs are to the inner boundary itself (which means the combined cell in this case is just the innermost CMC cell itself), and the injected packets (i.e., those injected at the inner boundary) have properties (direction, time, frequency, g' factor (which is probably chosen to be 1 in all cases), and polarization (which is probably chosen to be zero in all cases) that are determined by the inner-boundary boundary conditions. Assuming (as we do) that the innermost CMC cell is optically thick, most injected packets in OMC will not transfer to the outward-going exit adjacent IDTs (i.e., will not escape from the top of the innermost CMC cell), but will exit at the inner-boundary interface itself (i.e., the inward-going exit adjacent interface). However, no matter how optically thick the innermost CMC cell, the runtime lists for IDTs for the inner boundary should not include the trajectories of packets that exit at the inner boundary (i.e. non-escaping packets), but only the trajectories of inner-most-CMC-cell escaping packets. Thus, in CMC all injected packets are inner-most-CMC-cell escaping packets. To account correctly for the energy deposition in CMC, all the energy deposited by non-escaping packets in OMC should be included in the averaging process to determine the CMC energy deposition array. Does there have to be a decrease in the number of injected packets in going from OMC to CMC? If the energy scaling of the packets is done using those that escape the atmosphere (i.e., atmosphere-escaping packets), no: the same number on average escape the atmosphere in OMC and CMC, and the CMC energy deposition array determined as we have described gives the same energy deposition to the innermost CMC in both cases. However, if inner-boundary outward-going flux is a boundary condition and used to scale the energy of the packets, then to keep the flux

the same between OMC and CMC, one would have to reduce rate of packet injection for an inner-boundary IDT by a factor K/K^* where K^* is the number of packets injected in OMC and K is number of escaping packets and also the length of runtime list. Note for a time-independent atmosphere, there is effectively a singular time bin.

The second special case is the second outermost interface. Packets from the outward-going IDTs for this interface can escape through the outer boundary and can be used calculate the emergent spectrum after the convergence of a Lambda-iteration (see below § 3.2). However, one can also use formal integration and source functions to calculate emergent spectrum after convergence (see below § 3.2). The inward-going IDTs (of the second outermost interface) need no CMC ingredients since no packets arrive at it from the outer-boundary interface as we discuss below for the third special case. All packets that arrive at the inward-going IDTs (of the second outermost interface) are on trajectories that started from the the outward-going IDTs (of the second outermost interface) 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 for its IDTs. Packets that arrive its outward-going IDTs just escape to infinity and no packets arrive at its inward-going IDTs. As mentioned above, to calculate the emergent spectrum, we can use the escaping packets or we can use formal integration and the source functions (see below § 3.2).

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 (i.e., their IDTs) 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.

To explicate natural transition propagation a bit further, we consider two important atmosphere cases. First, if there is an atmosphere inner boundary, 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 atmosphere outer boundary. Second, for localizable (and therefore finite) 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 energy sources and initial-time thermal energy 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 the atmosphere 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 all the individual arrival states of OMC (i.e., full OMC), but on only the distributions of arrival states encoded in the runtime lists.
3. Calculating the emergent spectrum from escaping packets is a straightforward approach. Of course, only needs to that radiative transfer calculation once after the Lambda-iteration has converged. This post-convergence radiative transfer calculation should probably use far more packets and longer runtime lists than during the Lambda-iteration itself to give the emergent spectrum higher accuracy. In particular, the distributions for direction, frequency, and polarization for the outward-going IDTs from second outermost interface (and deeper ones if the CMC cells near the surface of the atmosphere are not optically thick) have be calculated for sufficient accuracy for the emergent spectrum rather than for the crude accuracy that most often is needed for IDTs as explained in § 3.1).
4. 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 parenthetically 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, we need to explicitly calculate the source functions (including the scattering source functions with, if doing polarized transfer, Stokes vector formalism) 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.
5. As the IDTs in CMC accumulate in a CMC radiative transfer solution phase, there is

an increasing speed-up of the radiative transfer through the whole atmosphere. Eventually, all IDTs that we wish to treat in CMC are in CMC.

6. Should all IDTs be put in CMC? If we elect to have optically thin CMC cells, then most of the optically thin CMC cells should have their into-the-CMC-cell-pointing IDTs 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 most optically-thin 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, we can make optically thin thermal-state cells subcells of optically thick CMC cells. We recommend this in order to have a uniform treatment as aforesaid.
7. There will very probably be cases where the radiative transfer phase of an Lambda-iteration iteration becomes complete (as judged by some criterion) before all IDTs are put in CMC and then the IDTs not in CMC can be left not in CMC.
8. In a Lambda-iteration, the completeness of the CMC ingredients can be increased from one radiative transfer solution phase to the next as overall convergence to the exact atmosphere solution is approached, and therefore we can approach the exact atmosphere solution as closely as we like 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.
9. 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 IDTs becomes itself a very slow procedure. In this case, CMC subcells with their own IDTs could be implemented. The CMC subcell IDTs 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 IDTs in CMC. Perhaps even a hierarchy of CMC subcells might be needed. Further consideration of CMC subcells and hierarchies of CMC subcells 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 effectively cancels the estimated speed-up of CMC (see below § 4).

3.3. Some Details of CMC

Now we expand on some details of CMC.

As a packet propagates through a combined cell in OMC, we simultaneously build up the CMC energy deposition array needed for CMC and the total energy deposition array needed for the next Lambda-iteration thermal state solution phase. Note the energy deposition in both arrays will require scaling for a thermal state solution phase (e.g., Lucy 1999a; Kasen et al. 2006), unless the absolute scale of the packets is set initially which can be done when the packets are created by radioactive energy sources or initial-time thermal energy. The CMC energy deposition array and the total energy deposition array have their energy bins indexed by thermal-state cell (which is not necessarily the same the CMC cell as discussed above in § 3.1), frequency, and time. 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 CMC energy deposition array and the total energy deposition array since there is only one energy bin for all deposited energy. For NLTE and/or if we require 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 we want 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, we do not need energy deposition at all.

At the completion of the runtime list for an IDT and the CMC energy deposition array elements for an IDT (which to complete requires averaging the energy deposition per packet), the IDT is put in CMC. Once the IDT is put in CMC, the CMC energy deposition array is used every time a packet makes a CMC step to update to the total energy deposition array: i.e., the CMC energy deposition array is added to the total energy deposition array. As aforesaid, the total energy deposition array is used for the next thermal state solution phase of the Lambda-iteration.

We will here explore one of the criteria for the sufficient completeness of the CMC ingredients mentioned in § 3.1. This criterion is having a sufficiently accurate exit probability distribution for the exit adjacent IDTs from the combined cell that surrounds an arrival IDT. We will assume that there is no time dependence as a simplification for this exploration. This means there are only two exit adjacent IDTs Their directions both point away from the arrival IDT. Recall, we assume the CMC cells have a 1-dimensional character even for 3-dimensional atmospheres, and so the exit probability distribution is, in fact, a binomial distribution (e.g., Bevington 1969, p. 53) that describes the exit behavior at the 2 possible exit adjacent IDTs. For the exits from one of the 2 IDTs, the relative standard deviation

is σ/μ (where σ is the standard deviation and μ is the mean of the counts of packet exits) after n trials (i.e., n starts of packet from an arrival IDT). We use the word *relerr* to be the estimate of σ/μ . The formulae for σ/μ and *relerr* are, respectively,

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

and

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

where p is the probability of exit event E (which is either of the two events), N is the number of exit events E after an actual run of n trials (and not an average number of exit events $\mu = np$ after an infinite number of runs of n trials), and N/n is the estimate of p after the aforesaid actual run of n trials. The estimate *relerr* is the metric for goodness. When *relerr* is sufficiently small and other criteria for CMC are met, the arrival IDT can be put in CMC. For example, we could require $\text{relerr} \leq 0.01 = 10^{-2}$ for the both of the possible events E assuming about equal probability (i.e., $p \approx 1/2$) for the two exit adjacent IDTs which could be roughly true for most arrival IDTs (but not an inner-boundary arrival IDT). In this case, N for both exit events E would have to be of order 5×10^3 . However, for the inner-boundary arrival IDT, one exit adjacent interface is the inner boundary itself and most packets would leave the CMC at that exit adjacent interface (as discussed in § 3.1). In this case, $p \ll 1$ for the outward-going events E and the limit Poisson distribution applies. So we would require N for the outward-going events E to be of order 10^4 to give $\text{relerr} \leq 0.01 = 10^{-2}$ for the outward-going events E. In the two cases just discussed, the required *relerr* is essentially the same condition as the relative tolerance $\epsilon = 0.01$ we suggested in § 3.1 for the more general case where the atmosphere was allowed to be time dependent.

4. Estimating the CMC Speed-Up

In this section, we review 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), derive formulae to estimate the CMC speed-up relative to OMC for an ideal-case 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 below, 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 point source for a random walk after n steps $\langle r_n^2 \rangle$ is given by

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

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

It follows from Equation (3), that the number of computer operations and therefore the computer time needed for packets to spread out from a point or surface source to having a variance equal to r^2 (r being the distance from the point or surface source to some surface)

scales to some approximation as n (the number of random walk steps for the spreading out) and therefore to some approximation as $C\tau^2 = C(r^2/\ell^2)$, where $\tau = r/\ell$ is the optical depth corresponding to r and C is a constant obtained from Equation (3). It then follows that the total computer time for a whole radiative transfer calculation through a layer of optical depth τ must scale approximately as the computer time mentioned in the last sentence. Thus, for this total computer time, we have the result

$$\text{approximate total computer time} \propto \tau^2 \times \begin{cases} \frac{3}{2} & \text{Case 1;} \\ 3 & \text{Case 2;} \\ 1 & \text{Case 3,} \end{cases} \quad (4)$$

where the three cases are those specified in Equation (3). We make use of Equation (4) below in §§ 4.2 and 4.4.

4.2. 1-Dimensional Atmospheres

In this subsection, we derive estimation formulae for quantitatively describing and estimating the speed-up of CMC relative to OMC for an ideal-case 1-dimensional atmosphere.

First, we divide the atmosphere into M CMC cells which are layers perpendicular to the outward direction. The optical depth of the atmosphere is τ in an average opacity of some kind. Note M is a calculational parameter under the control of the calculator and τ is a physical parameter which in realistic cases is not under the control of the calculator. Thus, one goal is to find the formula for M that gives the maximum speed-up for a specified τ .

Since we are only interested in optically thick atmospheres, we assume $\tau \gg 1$. 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 are optically thick (i.e., as has having $\tau/M \gg 1$). This assumption will not hold if M gets sufficiently large. But 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 do not get so large.

For ordinary OMC, Case 1 (see above § 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 (see especially Equation (4) 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 additional roughness in estimating the computer time period is of no importance. We add up the scaling terms $(3/2)(\tau/M)^2$ without regard to the additional

roughness to get a scaling $M(3/2)(\tau/M)^2 = (3/2)(\tau^2/M)$ proportional to the computer time needed to put all CMC cells in CMC. (Putting a CMC cell in CMC is a shorthand for putting its IDTs 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 as per the discussion in § 4.1 (see especially Equation (4) in § 4.1). Thus, the computer time for a pure CMC radiative transfer through the whole atmosphere scales as M^2 . However, we emphasize in actual cases using CMC, the radiative transfer through the atmosphere is being done in part by OMC as the CMC cells are being progressively put into CMC. Therefore, estimating the computer time for the pure CMC radiative transfer part of the radiative transfer as scaling as M^2 (which is as if all the radiative transfer as being done by pure CMC radiative transfer) will tend to be an overestimate.

We now specify an (unscaled) estimate of CMC computer time the computer time function (CTF) by

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

where we have just summed the contributions of the time scalings discussed above. Note we have added a -1 term to the CTF to make the 1 cell limit for CMC equal to the computer time function for OMC where computer time scales as $(3/2)\tau^2$ as follows from the discussion in § 4.1 (see especially Equation (4) in § 4.1). If $M = 1$, then an in ideal sense at least a CMC radiative transfer calculation is solved by OMC using many thermal state cells and one giant CMC cell which is never put into CMC: i.e., the calculation is just an OMC calculation and the CTF with $M = 1$ is the OMC limit of the CTF. Of course, the CTF is a rough estimate for both CMC and OMC.

Another point to make about the CTF (i.e., Equation (5)) is that we have set the coefficients for the terms in the CTF to 1. Realistic coefficients could differ from 1 by an order of magnitude or more. But since the CTF only allows a rough estimate of how the computer time scales with M and τ in any case, setting the coefficients to 1 is the only reasonable choice. Also any case, the coefficients can only be determined for specified atmospheres and we are trying to be general.

The CTF allows us to investigate the speed-ups obtainable from CMC. However, as noted above, the CTF is an unscaled estimate of computer time, and so it can only be used

for determining relative speed-ups for specified atmospheres: i.e., specified plane-parallel atmospheres or spherically symmetric atmospheres. Note also that the CTF, following from a point we made above, will tend to overestimate computer time for CMC radiative transfer since the term M^2 tends to overestimate the computer time for pure CMC radiative transfer part of the radiative transfer. For 1-dimensional atmospheres, the tendency to overestimate is probably relatively unimportant, but for 3-dimensional atmospheres, we believe the tendency will lead to a real and significant overestimate as we discuss below in § 4.4 where we consider the speed-up for 3-dimensional atmospheres.

The CTF allows us to estimate the maximum speed-up of CMC as a function of M which we do below. The maximum speed-up is, of course, of most interest. Note that the speed-up increases as the CTF decreases, and so maximum speed-up is given by the minimum CTF. However, more useful than the CTF itself in describing the maximum speed-up is what we call the time reduction factor (TRF) and the speed-up factor (SUF) which factors we define below. To obtain the TRF and SUF, we first find the minimum of the CTF as a function of M which we treat as a continuous variable. Obviously, and with no computational relevance, $\text{CTF}(M = 0) = \infty$. Also obviously, $\text{CTF}(M = 1) = (3/2)\tau^2$ which is the OMC limit of the CTF as discussed above. In fact, as M increases above 0, the CTF strictly decreases until it reaches a minimum $\text{CTF}(M_{\min})$ at M_{\min} determined as follows:

$$\begin{aligned} \left. \frac{d\text{CTF}(M)}{dM} \right|_{M_{\min}} &= \left[2M - \left(\frac{3}{2} \right) \frac{\tau^2}{M^2} \right] \Big|_{M_{\min}} = 0 , \\ M_{M_{\min}}^3 &= \left(\frac{3}{4} \right) \tau^2 , \quad M_{\min} = \left(\frac{3}{4} \right)^{1/3} \tau^{2/3} = (0.908560\dots) \times \tau^{2/3} , \\ \text{CTF}(M_{\min}) &= \left(\frac{3}{4} \right)^{2/3} \tau^{4/3} - 1 + \left(\frac{3}{2} \right) \frac{\tau^2}{(3/4)^{1/3} \tau^{2/3}} , \\ \text{CTF}(M_{\min}) &= 3 \left(\frac{3}{4} \right)^{2/3} \tau^{4/3} - 1 = (2.476445\dots) \times \tau^{4/3} - 1 . \end{aligned} \tag{6}$$

Note for computationally relevant M_{\min} , we demand

$$M_{\min} > 1 \quad \text{implying} \quad \tau > \sqrt{\frac{4}{3}} . \tag{7}$$

Since we have assumed $\tau \gg 1$, the condition of Equation (7) is always satisfied.

Now for $M \in [1, M_{\min})$, the CTF strictly decreases and for $M \in (M_{\min}, \infty]$, the CTF strictly increases. Thus, there is only a range of M values between two special M values giving $\text{CTF}(M) = (3/2)\tau^2$ where we can obtain speed-ups relative to OMC. This range is the speed-up range. In fact, there are 3 points (which we call M_1, M_2, M_3) where $\text{CTF}(M) =$

$(3/2)\tau^2$ and we solve for them as follows:

$$\begin{aligned}
 \frac{3}{2}\tau^2 &= M^2 - 1 + \left(\frac{3}{2}\right) \frac{\tau^2}{M} \\
 0 &= M^3 - M \left(1 + \frac{3}{2}\tau^2\right) + \frac{3}{2}\tau^2 \\
 0 &= (M - 1) \left(M^2 + M - \frac{3}{2}\tau^2\right)
 \end{aligned}$$

$$M = \begin{cases} M_1 = 1 & \text{which solution we already knew} \\ & \text{and is the lower bound of} \\ & \text{the speed-up range.} \\ M_2 = \frac{-1 + \sqrt{1 + 6\tau^2}}{2} & \text{which is the upper bound of} \\ & \text{the speed-up range.} \\ M_3 = \frac{-1 - \sqrt{1 + 6\tau^2}}{2} & \text{which is irrelevant to CMC} \\ & \text{since it is always negative.} \end{cases} \quad (8)$$

The main conclusion from Equation (8) is that the speed-up range is (M_1, 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, we usually do not know any of the opacities before we start calculating the solution. We 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 be 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 $\text{CTF}(M_1 = 1) \propto \tau^2$ (see Equation (5) above) and $\text{CTF}(M_{\min}) \propto \tau^{4/3}$ (see Equation (6) above). To analyze this speed-up, we now define the time reduction factor (TRF) and its inverse the speed-up factor (SUF) by, respectively,

$$\text{TRF}(\tau) = \frac{\text{CTF}(M_{\min})}{\text{CTF}(M_1 = 1)} = \frac{3(3/4)^{2/3}\tau^{4/3} - 1}{(3/2)\tau^2} \approx 2 \left(\frac{3}{4}\right)^{2/3} \tau^{-2/3} \quad (9)$$

and

$$\text{SUF}(\tau) = \frac{(3/2)\tau^2}{3(3/4)^{2/3}\tau^{4/3} - 1} \approx \frac{1}{2} \left(\frac{4}{3}\right)^{2/3} \tau^{2/3}, \quad (10)$$

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

To illustrate the speed-ups possible for 1-dimensional atmospheres solved using CMC, we present in Table 1 (given at the end of this research note) time reduction factors (TRFs), speed-up factors (SUFs), and other results for two ideal-case 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 .

The estimation formulae derived above may be order of magnitude accurate for quantitatively describing and estimating the speed-up of CMC relative to OMC for ideal-case 1-dimensional atmospheres. If so, one can see that for large τ (e.g., $\tau = 10^6$), very large speed-ups may be obtainable for ideal-case atmospheres. With suitable adjustments, the estimation formulae may have the same accuracy for general 1-dimensional atmospheres.

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 for the specified optical depths τ (as above and in Table 1), and so actual speed-ups possible for 1-dimensional atmospheres will probably be less, maybe much less, than those suggested by those estimation formulae and those results.

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:

$$\text{CTF}(M) = \left\{ \begin{array}{l}
 M^2 - 1 + M \left(\frac{3}{2}\right) \left(\frac{\tau}{M}\right)^2 = M^2 - 1 + \left(\frac{3}{2}\right) \frac{\tau^2}{M} \\
 \left(\frac{3}{2}\right) \tau^2 \\
 M_3 = \frac{-1 - \sqrt{1 + 6\tau^2}}{2} \\
 = -\sqrt{\frac{3}{2}}\tau \sqrt{1 + \frac{1}{6\tau^2}} - \frac{1}{2} \\
 = -\sqrt{\frac{3}{2}}\tau - \frac{1}{2} - \sqrt{\frac{3}{2}} \frac{1}{12\tau} + \dots \\
 \infty \\
 \left(\frac{3}{2}\right) \tau^2 \\
 3 \left(\frac{3}{4}\right)^{2/3} \tau^{4/3} - 1 = (2.476445\dots) \times \tau^{4/3} - 1 \\
 M_{\min} = \left(\frac{3}{4}\right)^{1/3} \tau^{2/3} = (0.908560\dots) \times \tau^{2/3} \\
 \left(\frac{3}{2}\right) \tau^2 \\
 M_2 = \frac{-1 + \sqrt{1 + 6\tau^2}}{2} \\
 = \sqrt{\frac{3}{2}}\tau \sqrt{1 + \frac{1}{6\tau^2}} - \frac{1}{2} \\
 = \sqrt{\frac{3}{2}}\tau - \frac{1}{2} + \sqrt{\frac{3}{2}} \frac{1}{12\tau} + \dots \\
 M^2 \\
 \left(\frac{3}{2}\right) \frac{\tau^2}{M}
 \end{array} \right.$$

in general for CMC.

for $M = M_3$, where

The M_3 case is irrelevant to CMC since M_3 is always negative.

for $M = 0$.

for $M_1 = 1$ which is the lower bound of the speed-up range.

is the minimum value, where M_{\min} is given by

for $M = M_2$, where

The M_2 case is the upper bound of the speed-up range.

For $M > M_2$, we gets a slow-down, not a speed-up.

for $M \gg M_2 \approx \sqrt{3/2}\tau$ which is far in the slow-down region.

for fixed M with τ sufficiently large.

(11)

The last case in Equation (11) 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 $\text{CTF}(M = M_{\min})$. However, Equation (11) suggests that even if there is difficulty in finding the M values for the maximum speed-ups for general 1-dimensional atmospheres (meaning not just our ideal-case 1-dimensional atmospheres), we can still obtain useful speed-ups by just experimentally adjusting M starting perhaps from small values.

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. To explicate, since the thermal-state cells are subcells of 1-dimensional CMC cells, they will have an average behavior that will be encoded faster in the CMC ingredients than their individual detailed behavior. That encoded average behavior will probably be adequate to speed up the convergence of the Lambda-iteration, except maybe near final convergence.

To recapitulate from § 3.1, there are two cases of 3-dimensional atmospheres 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 (and therefore finite), we use spherically symmetric (or approximately spherically symmetric) shell CMC cells centered on the point of highest symmetry of the 3-dimensional atmospheres.

To elucidate the behavior of both cases of 3-dimensional atmospheres, we consider ideal-case 3-dimensional atmospheres, analogous to the ideal-case 1-dimensional atmosphere considered in § 4.2.

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 specify a 3-dimensional atmosphere computer time function ($\text{CTF}_{3\text{-d}}$) by

$$\text{CTF}_{3\text{-d}}(M) = M^2 - 1 + CM \left(\frac{3}{2}\right) \left(\frac{\tau}{M}\right)^2 = M^2 - 1 + C \left(\frac{3}{2}\right) \frac{\tau^2}{M}, \quad (12)$$

where τ is again the optical thickness of the atmosphere. Equation (12) 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 average number of thermal-state cells in each planar slab CMC cell. The extra computer time taken for the OMC phase of a CMC radiative transfer to adequately encode the radiative transfer behavior of the thermal-state cells is the motivation of the factor C . Since we can have $C \gg 1$, we find that we can have $\text{CTF}_{3\text{-d}} \gg \text{CTF}$. 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 atmospheres. However, we still expect a significant speed-up for CMC for two reasons A and B. The reason A 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 reason B is just the second reason for using 1-dimensional CMC cells given above in this subsection.

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

$$\text{CTF}_{3\text{-d}}(M) = M^2 - 1 + M^3 \left(\frac{3}{2}\right) \left(\frac{\tau}{M}\right)^2 = M^2 - 1 + M \left(\frac{3}{2}\right) \tau^2, \quad (13)$$

where τ is again the optical thickness of the atmosphere. Equation (13), like Equation (12) 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 (13)) 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 in computer time needed to encode the radiative transfer behavior of the thermal-state cells. Equation (13) 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 reasons A and B cited above for the essentially planar 3-dimensional atmosphere case.

The upshot of our analysis of the two (ideal-case) 3-dimensional atmospheres is that we expect a significant speed-up for CMC for 3-dimensional atmospheres in general, generalizing from those atmospheres. However, it will take further analysis to demonstrate and quantitatively estimate of the speed-up. In fact, probably the speed-up can only be demonstrated and 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).

On another discussion front, an advantage of CMC is that CMC can be supplemented by speed-ups of OMC. For example, the machine learning sampling method of Charles & Chen (2024) that can speed up the sampling of distributions needed in OMC: e.g., those needed for the complex frequency and angular redistribution in some kinds of scattering: e.g., electron scattering treated in detail (e.g., Mihalas 1978, p. 420).

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). I thank colleagues for their comments on this research note.

REFERENCES

- Bevington, P. R. 1969, *Data Reduction and Error Analysis for the Physical Sciences* (New York: McGraw-Hill Book Company)
- Charles, W., & Chen, A.Y. 2024, arXiv:2406.19385
- 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
- Lauwers, A., Baes, M., Camps, P., & Vander Meulen, B. 2024, *A&A*, in press
- 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. Time Reduction Factors (TRFs), Speed-Up Factors (SUFs), Etc. for Ideal-Case 1-Dimensional Atmospheres Solved Using CMC

τ	M_{\min}	$\frac{\tau}{M_{\min}}$	CTF($M_1 = 1$) $= \frac{3}{2}\tau^2$	CTF(M_{\min})	TRF SUF	M_2
10^2	19.6	5.11	1.50×10^4	1.15×10^3	0.0766 13.1	122
10^6	9.09×10^3	110.0	1.50×10^{12}	2.48×10^8	1.65×10^{-4} 6.06×10^3	1.22×10^6