SAND2019-4882C

AN EXTENSION OF CONDITIONAL POINT SAMPLING TO MULTI-DIMENSIONAL TRANSPORT

Aaron Olson[1] **and Emily Vu**[2,1]
[1]Sandia National Laboratories
Albuquerque, NM 87185, USA

[2]University of California, Berkeley
Berkeley, CA 94720, USA

aolson@sandia.gov, emilyhvu@berkeley.edu

ABSTRACT

Radiation transport in stochastic media is a challenging problem type relevant for applications such as meteorological modeling, heterogeneous radiation shields, BWR coolant, and pebble-bed reactor fuel. A commonly cited challenge for methods performing transport in stochastic media is to simultaneously be accurate and efficient. Conditional Point Sampling (CoPS), a new method for transport in stochastic media, was recently shown to have accuracy comparable to the most accurate approximate methods for a common 1D benchmark set. In this paper, we use a pseudo-interface-based approach to extend CoPS to application in multi-D for Markovian-mixed media, compare its accuracy with published results for other approximate methods, and examine its accuracy and efficiency as a function of user options. CoPS is found to be the most accurate of the compared methods on the examined benchmark suite for transmittance and comparable in accuracy with the most accurate methods for reflectance and internal flux. Numerical studies examine accuracy and efficiency as a function of user parameters providing insight for effective parameter selection and further method development. Since the authors did not implement any of the other approximate methods, there is not yet a valid comparison for efficiency with the other methods.

KEYWORDS: stochastic media; Monte Carlo transport; Conditional Point Sampling; CoPS

1. INTRODUCTION

For many years, a variety of approaches and algorithms for performing radiation transport in stochastic media have been examined, usually in 1D, and often on a set of problems introduced by Adams, Larsen, and Pomraning [1] and expanded on by Brantley [2]. Recently, Larmier et al. produced a set of benchmark results in 2D and 3D [3] based on the problems in the classic 1D set opening the door for methods to assess accuracy in the multi-D context. Here, we compare the accuracy of a multi-D version of CoPS with those multi-D benchmark results and published approximate method results.

Existing 1D [1,2] and multi-D [3] benchmark results have each been computed by performing transport on each realization of a large ensemble generated according to a statistical rule, i.e.,

Markovian mixing. As long as it is known how to create realizations according to a desired rule, this "brute-force" method is accurate, but very expensive, often requiring thousands or more transport calculations.

One of the earliest and most widely used approximate methods is the "atomic mix" (AM) method in which the materials in the stochastic mixture are "mixed at the atomic level" to form one homogenized material on which transport is performed. This method is straightforward to setup even with existing codes. It tends to lack accuracy, however, with the exception of a few problems, especially those involving finely-mixed, highly-scattering media. Results for the AM approach are found in publications such as Refs. [2,3] (1D) and Refs. [3,4] (multi-D).

Another widely circulated approximate method is the Levermore-Pomraning (LP) closure and derivatives. This closure is a way to solve the stochastic transport equation deterministically [1]. Shortly after its development, Chord Length Sampling (CLS)–in which a particle streams within one material for up to the length of a sampled chord before resampling which material it is in–was shown to be the Monte Carlo equivalent of the LP closure [2,5]. Memory-enhanced versions of CLS, "Algorithm B" and "Algorithm C," were proposed in which one or more local chords are "remembered," yielding successively more accurate transport results [5]. Following the recent publication of multi-D benchmark values, multi-D transport results for CLS [4,6] and an extension of Algorithm B to multi-D called the Local Realization Preserving method (LRP) have been published [4]. Chord Length Sampling showed significant improvements in accuracy on this benchmark set over the atomic mix approach, and LRP showed meaningful improvements in accuracy over CLS.

Larmier et al., recently created and published benchmark results for an approximate method: Poisson Box Sampling (PBS) [7]. In Ref. [8], they observed that Box-Poisson mixing–in which planes are sampled according to Markovian mixing in each of the Cartesian directions leaving a series of parallelepiped tessellations in which the materials are sampled according to their abundance–yielded very similar transport results to Poisson (Markovian) mixing. They then reasoned that sampling boxes, instead of chords, on-the-fly would inherently provide some memory while leveraging the similarity in performance between Poisson and Box-Poisson tessellations. In PBS-1 they sample a "box" at a time on-the-fly and remember it until the particle leaves the box; in PBS-2 they also remember the most recent box the particle left. By retaining some local memory, PBS-1 is analogous to Algorithm B (LRP), and by retaining additional local memory, PBS-2 is analogous to Algorithm C [7]. PBS-1 was shown to be generally more accurate for Markovian mixing than CLS, and PBS-2 results generally improved on PBS-1 results.

We recently proposed a new approximate method, Conditional Point Sampling (CoPS) [9], in which material is defined at sampled points in the domain with full memory until at least one particle history is complete. Streaming is performed using Woodcock tracking, enabling particles to stream without full definition of what they are streaming through. The material type must be defined at Woodcock-tracking pseudo-collision sites. The material type is sampled conditionally on the material designations that have already been made. In Ref. [9], we derived two conditional probability functions in 1D for defining the probability to be sampled against when sampling the material at a new point based on the point-wise material designations already sampled. The first conditional probability function produced reflectance results roughly as accurate as LRP, and transmittance results generally more accurate than Algorithm C. We believe that the second is errorless

in the special case of 1D, Markovian-mixed media and showed numerical results to support this belief. In a companion paper to this one [10], we examine use of a new "Embedded Variance Deconvolution" method [11] with CoPS to enable CoPS to compute not only mean transport results, but the variance caused by the random mixing of the material. In this paper, we examine the accuracy of CoPS in multi-D using recently published benchmark values, compare accuracy with the other approximate methods introduced above, and examine accuracy and efficiency as a function of user-input options.

2. CONDITONAL POINT SAMPLING

CoPS is based on the fact that, when using Woodcock particle tracking [12], it is not necessary to know the exact location of all materials in the domain and the idea that realizations can be sampled on-the-fly in an as-needed basis. Woodcock tracking only requires that material in the domain be specified at pseudo-collision locations: in CoPS we make on-the-fly, point-wise material assignments at these pseudo-collision sites. New material assignments are sampled according to material probabilities: when no other points have been defined these probabilities are simply the material abundances, but when at least one point-wise material assignment has been made, new points are sampled conditionally on previously-made material assignments. In this way, CoPS successively "reveals" a realization, never fully defining it, but defining enough to effect transport, including the ability to take collision tallies at all pseudo-collision sites. In theory, for any type of mixing and any combination of previously defined points, there is a correct conditional probability function from which to sample new material assignments yielding an errorless algorithm; in practice, the true conditional probability function will usually require approximation introducing some bias error. The authors have demonstrated use of two different conditional probability functions in 1D, Markovian-mixed media [9]: the first involves an approximation and yields some bias error and the second we believe to be errorless yielding no bias error. Here, we expand the "pseudo-interface" approach to approximating the true conditional probability function used in Ref. [9] beyond 1D to multi-D.

3. A MULTI-D PSEUDO-INTERFACE-BASED CONDITIONAL PROBABILITY FUNCTION

Media that has Markovian mixing can be generated by sampling a number of $(d-1)$-dimensional hyperplanes of random orientation and randomly filling the resulting cells with materials according to the material abundances [3]. The material abundances can also be called the unconditional material probabilities:

$$\pi(\alpha) = \frac{\Lambda_\alpha}{\Lambda_\alpha + \Lambda_\beta}; \qquad \pi(\beta) = \frac{\Lambda_\beta}{\Lambda_\alpha + \Lambda_\beta}, \tag{1}$$

where Λ_α and Λ_β are the average chord lengths of materials α and β. A classic paper on Markovian mixing [13] shows that the number of hyperplanes crossed between two randomly selected points is Poisson-distributed. Hyperplanes have been called "pseudo-interfaces" in the 1D setting [14]. The frequency of sampling no pseudo-interfaces ($k = 0$) between two points distance r apart is thus

$$f(k=0,r) = \exp\left\{-\frac{r}{\Lambda_C}\right\}, \tag{2}$$

where Λ_C is the correlation length of the random mixing:

$$\Lambda_C = \frac{\Lambda_\alpha \Lambda_\beta}{\Lambda_\alpha + \Lambda_\beta}. \tag{3}$$

In previous work on CoPS [9], we used the Poisson frequency of pseudo-interfaces to derive two conditional probability functions in 1D.

Here, we expand the pseudo-interface derivation of (in general) approximate conditional probability functions [9] to multi-D. Firstly, we assume that previously defined points nearest to a new point most highly affect the conditional probability at the new point. Secondly, we recognize that two previously defined points within a small solid angle relative to a new point will tend to either hold similar information informing the conditional probabilities or that one will "shield" the contribution of the other. Based on these guiding principles, we devise a scheme to select a limited-sized subset of all defined points which highly affect the true conditional probability for a new point and govern our approximation of the true conditional probability. Thirdly, we assume that the selected "governing points" have been chosen sufficiently far apart from one another in solid angle such that their contributions to the conditional probability at the new point can be treated as independent of one another. In other words, we do not account for the probability of a single pseudo-interface hyperplane existing between the new point and more than one governing point. In the following two subsections, we discuss the proposed scheme for selecting governing points and the proposed method for estimating the conditional probability based on the selected governing points.

3.1. Selection of governing points

We select governing points by testing each defined point beginning with the point closest to the new point. A point under consideration to be a governing point–a "candidate point"–can be rejected from inclusion as a governing point if it is close to or "shielded by" an already selected governing point. The search for more governing points is terminated either when a certain number of governing points have been chosen or when all points within a prescribed distance from the new point have been tested. Three user input values guide the selection process: the exclusion multiplier M_e, maximum distance multiplier M_d, and maximum number of governing points N_{max}.

Associated with each governing point is a solid angle in which any candidate point is rejected. The size of this solid angle is determined by the correlation length of the material of the governing point and a user-defined parameter, the exclusion multiplier (M_e). We label the distance from the new point to the governing point as having length r. A line is drawn perpendicular to the line connecting the new point and governing point of material type α protruding from the governing point the distance of $\Lambda_\alpha M_e$. A third line is drawn from the end of this one back to the new point forming a right triangle for which two of the sides (r, $\Lambda_\alpha M_e$) and one of the angles ($\pi/2$) is known. The Law of Sines is then used to solve for the exclusion angle θ_g associated with that governing point with respect to the new point:

$$\theta_g = \sin^{-1}\left(\frac{\Lambda_\alpha M_e}{\sqrt{\Lambda_\alpha^2 M_e^2 + r^2}}\right). \tag{4}$$

Candidate points are tested for angle-based exclusion by testing whether they are within the exclusion angle of any governing point. The distance between the new point and governing point is once

again r, the distance between the new point and candidate point is r', and the distance between the governing point and candidate point is d. The Law of Cosines is used to solve for the angle θ_c between the governing point and candidate point with respect to the new point:

$$\theta_c = \cos^{-1}\left(\frac{r^2 + r'^2 - d^2}{2rr'}\right). \tag{5}$$

The candidate point is rejected if θ_c is smaller than θ_g for any governing point.

The selection of governing points is terminated in one of three ways: all existing points have been tested for inclusion, N_{max} points have been selected, or all points within the distance $\Lambda_c M_d$ have been tested. Both the exclusion multiplier M_e and maximum distance multiplier M_d are multipliers to length scales specific to a type of mixing in an attempt to enable the same value of these parameters to seamlessly apply to different stochastic mixing problems in a meaningful way. A sketch of this selection process is provided in Figure 1 in which two points are chosen as governing points (one of each material type), one point is excluded by angle, and six points are excluded by distance.

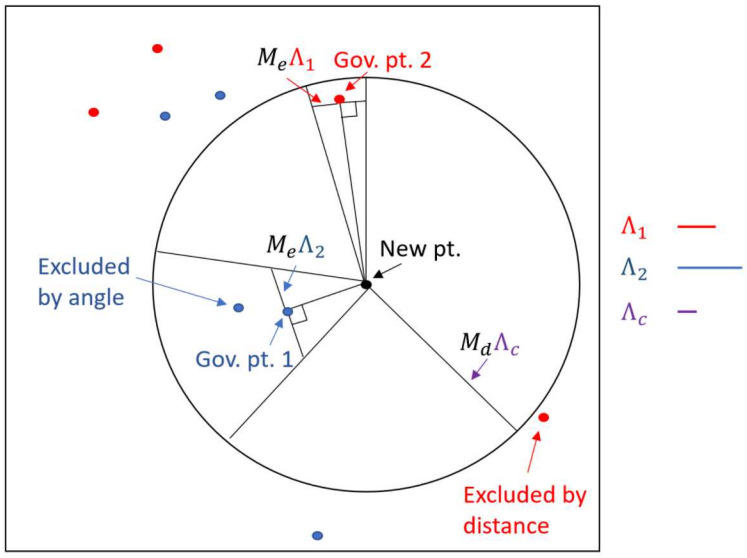

Figure 1: Diagram demonstrating angle- and distance-based exclusion.

3.2. Estimation of conditional probability

Once the governing points have been chosen, they are used to approximate the conditional probability function. We define a probability space comprised of $2^{N_{pt}}$ discrete outcomes including each permutation of there being or not being at least one pseudo-interface between the new point and each of N_{pt} governing points. We define A and B as the sets of permutations for which there are no pseudo-interfaces between the new point and at least one governing point of material type α or β, respectively. The entire space, illustrated in Figure 2, is then uniquely described by the union of the following four non-overlapping scenarios: the new point must be material α ($A \backslash B$), the new point must be material β ($B \backslash A$), the new point must be sampled based on material abundances

($(A \cup B)^C = O$), and the invalid scenario in which there is no pseudo-interface between the new point and at least one point of material α and at least one point of material β ($A \cap B$).

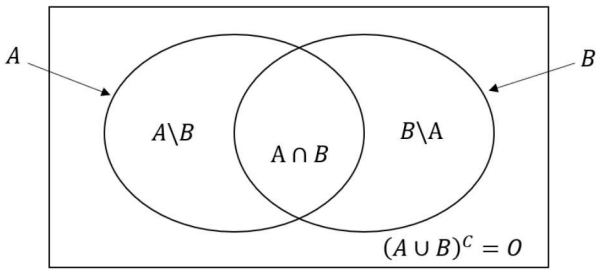

Figure 2: Diagram of four-scenario probability space.

Denoting the distance from the new point to the governing points as $\vec{r} = \{r_1, r_2, \cdots, r_{N_{pt}}\}$ and the material types of the governing points as $\vec{m} = \{m_1, m_2, \cdots, m_{N_{pt}}\}$, the probability of the compliments of sets A and B can be computed as the evaluation of two pi-products:

$$P(A^C) = \prod_{n=1;m_n=\alpha}^{N_{pt}} (1 - f(k=0, r=r_n)) \tag{6a}$$

$$P(B^C) = \prod_{n=1;m_n=\beta}^{N_{pt}} (1 - f(k=0, r=r_n)) \tag{6b}$$

The probabilities of each of the valid scenarios are then solved as a function of these two quantities:

$$P(O) = P((A \cup B)^C) = P(A^C)P(B^C) \tag{7a}$$

$$P(A\backslash B) = P(B^C) - P(O) = P(O)\left(\frac{1}{P(A^C)} - 1\right) \tag{7b}$$

$$P(B\backslash A) = P(A^C) - P(O) = P(O)\left(\frac{1}{P(B^C)} - 1\right) \tag{7c}$$

The probabilities of knowing you have material α, knowing that you have material β, or that you must sample according to material abundance–P_α, P_β, or P_{ind}, respectively–are computed by normalizing according to the probability of having a valid scenario:

$$P_\alpha = \frac{P(A\backslash B)}{P(A\backslash B) + P(B\backslash A) + P(O)} \tag{8a}$$

$$P_\beta = \frac{P(B\backslash A)}{P(A\backslash B) + P(B\backslash A) + P(O)} \tag{8b}$$

$$P_{ind} = \frac{P(O)}{P(A\backslash B) + P(B\backslash A) + P(O)} \tag{8c}$$

Finally, the probabilities of sampling material α or β are computed as the sum of the probability of knowing you have that material and the probability that you both need to sample according to material abundance and sample that material:

$$\pi(\alpha|\vec{r},\vec{m}) = P_\alpha + P_{ind}\pi(\alpha) \tag{9a}$$

$$\pi(\beta|\vec{r},\vec{m}) = P_\beta + P_{ind}\pi(\beta) \tag{9b}$$

Thus, the probability of sampling either material is computed based on a set of governing points and our independence assumption by tallying $P(A^C)$ and $P(B^C)$ on one for-loop and a few arithmetic evaluations.

4. RESULTS

We compute the accuracy of CoPS in 3D using the recently published benchmark set [3]. The problems in this benchmark set involve a geometry of length 10.0 in each direction with an isotropic source and vacuum leakage conditions on the "reflective boundary," vacuum leakage conditions on the opposite "transmissive boundary," and reflecting conditions on the other four. Particles either absorb or scatter isotropically. Two materials are mixed according to Markovian mixing. Three quantities are computed: reflectance, transmittance, and internal flux computed over the whole domain. The nine problems in the benchmark set are described using the unique permutations of case number and case letter outlined in Table 1. We first compare the accuracy of CoPS over the benchmark set with published 3D approximate methods and then examine the accuracy and runtime of CoPS on this problem set as a function of user-input variables.

Table 1: Benchmark Set Parameters

Case Number	$\Sigma_{t,0}$	$\Sigma_{t,1}$	Λ_0	Λ_1	Case Letter	c_0	c_1
1	10/99	100/11	99/100	11/100	a	0.0	1.0
2	10/99	100/11	99/10	11/10	b	1.0	0.0
3	2/101	200/101	101/20	101/20	c	0.9	0.9

4.1. Comparison of accuracy with published approximate methods

We produce the following CoPS results using a maximum of either one governing point (a two-point relationship) and call this CoPS2 or three governing points (a four-point relationship) and call this CoPS4. For both CoPS2 and CoPS4 results, 10^5 histories were used. Benchmark results from Ref. [3] as well as CoPS2 and CoPS4 results are listed in Table 2. In parentheses are uncertainties on the last digit. Both CoPS2 and CoPS4 produce values close to the benchmark values. CoPS4 is generally more accurate than CoPS2, as expected. The relative error of the most long-standing approximate methods in the literature and CoPS4 are plotted in Figure 3. This figure visually demonstrates that in most cases CLS is considerably more accurate than AM and that CoPS4 is generally more accurate than CLS.

Table 2: Selected Transport Results

Case		Reflectance			Transmittance			Flux		
		Bench [3]	CoPS2	CoPS4	Bench [3]	CoPS2	CoPS4	Bench [3]	CoPS2	CoPS4
1	a	0.4065(4)	0.408(2)	0406(2)	0.0162(1)	0.0167(4)	0.0164(4)	6.318(8)	6.33(2)	6.31(2)
	b	0.0376(2)	0.0376(6)	0.0358(6)	0.00085(3)	0.0009(1)	0.0009(1)	1.920(3)	1.921(6)	1.915(6)
	c	0.4036(4)	0.407(2)	0.409(2)	0.0164(1)	0.0175(4)	0.0175(4)	6.2960(8)	6.36(2)	6.36(2)
2	a	0.223(2)	0.217(1)	0.227(1)	0.0935(8)	0.0995(9)	0.0941(9)	7.55(2)	7.47(2)	7.41(2)
	b	0.161(2)	0.143(1)	0.152(1)	0.119(2)	0.111(1)	0.113(1)	7.76(7)	7.32(2)	7.46(2)
	c	0.3438(6)	0.314(1)	0.347(2)	0.1650(2)	0.155(1)	0.162(1)	10.76(6)	10.08(3)	10.74(6)
3	a	0.670(4)	0.654(2)	0.666(1)	0.169(3)	0.185(1)	0.170(1)	16.35(8)	15.96(6)	15.71(6)
	b	0.0167(6)	0.0129(4)	0.0143(4)	0.045(3)	0.0416(6)	0.0410(6)	3.49(8)	3.26(1)	3.31(1)
	c	0.395(1)	0.374(2)	0.398(2)	0.085(3)	0.0817(9)	0.0852(8)	7.9(1)	7.53(3)	7.82(3)

We compare CoPS accuracy with that of the other approximate methods that have published multi-D results in Table 3 using the root mean squared (RMS) relative error, the mean absolute relative error, and the maximum absolute relative error across the benchmark set as in Ref. [4]. The atomic mix (AM) results shown here were generated by homogenizing the medium and running CoPS (the results agree within statistics to other published multi-D AM results [3,4]). We use CLS results from Ref. [6] instead of Ref. [4] since the former provides results for internal flux but note that results from the two agree within statistics. Likewise, LRP results are only given in Ref. [4] and do not contain internal flux results such that we cannot compute error metrics for this quantity.

We first note that each approximate method compared here outperforms AM by a wide margin. We secondly note that each higher-fidelity version of an approximate method shows improvement over its lower-fidelity counterpart in accuracy in nearly every metric (i.e., LRP improves on CLS, PBS-2 improves on PBS-1, and CoPS4 improves on CoPS2). PBS-1 and CoPS2 each outperform CLS; PBS-1 is the most accurate of these three in reflectance and flux while CoPS2 is the most accurate in transmittance. LRP and CoPS4 show comparable accuracy in reflectance while PBS-2 is the most accurate of the methods compared here. LRP and PBS-2 show comparable accuracy in transmittance while CoPS4 is the most accurate of the methods compared here. PBS-2 is somewhat more accurate than CoPS4 for internal flux. We note that the statistical uncertainties on our CoPS results are of a similar order as the statistical uncertainty on the benchmark values such that running CoPS with more histories may somewhat decrease the value of the error metrics, but significantly better resolution for the error metrics would require better-resolved benchmark values.

4.2. Accuracy and runtime as a function of user-input variables

The previous results in this paper were generated using a maximum distance multiplier of $M_d = \infty$, an exclusion multiplier of $M_e = 1.0$, and a maximum number of governing points of either $N_{max} = 1$ or $N_{max} = 3$. Table 4 and Figure 4 present error and runtime results numerically generated by successively varying each of these parameters.

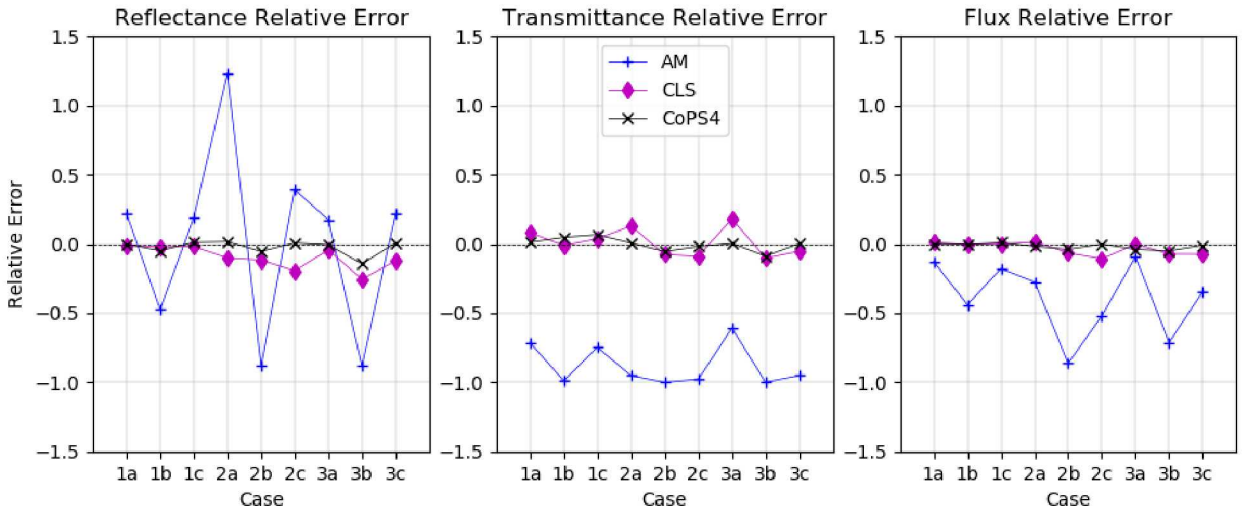

Figure 3: Relative error for selected methods on benchmark problems.

Root mean squared error across the benchmark set decreases as the maximum distance multiplier increases until it roughly plateaus. We think the plateau is caused by the level of uncertainty in the benchmarking and expect that the statistically resolved error would likely decrease asymptotically towards a finite error value. Runtime increases monotonically as the maximum distance multiplier increases. Based on this data, we choose $M_d = 3$ for subsequent numerical studies since this value minimizes runtime without measurably increasing error.

The behavior of CoPS as a function of the exclusion multipier is somewhat more complicated. Firstly, we compare having no angle-based exclusion ($M_e = 0.0$) with a small exclusion angle ($M_e = 0.001$). The error and runtime values for $M_e = 0.0$ are plotted at $M_e = 0.0005$ in Figure 4 so that they can be plotted for visual comparison on a log plot. This comparison shows that even a small amount of angle-based exclusion significantly improves accuracy. Secondly, we characterize accuracy and runtime as a function of non-zero exclusion angles. Our data suggests that the most interesting behavior likely occurs between $M_e = 0.001$ and $M_e = 0.25$, a span over which the error significantly diminishes. We think that error is largely swamped by statistical uncertainty beyond $M_e = 0.25$, though there may be a true increase in error for larger exclusion angles. With the exception of $M_e = 2.0$, runtime appears to decreases at first steeply, then gradually with larger M_e–the nonconformity of the $M_e = 2.0$ data point may be caused by cpu runtime fluctuations. Based on these results, we chose $M_e = 0.5$ for the next numerical study.

Finally, we compute error and runtime as a function of the maximum number of governing points. Error decreases sharply, then plateaus. The reflectance error appears to reach a local minimum around $N_{max} = 3$, then increase. This behavior may represent lack of statistical resolution or be a real effect perhaps caused by having chosen a relatively small M_e, which may affect reflectance more than the other quantities. Runtime monotonically increases as the maximum number of points increases. After a steep initial increase the runtime appears to asymptotically approach a maximum value. Non-negligible distance- and angle-based exclusion tends to limit the number of governing points chosen even when N_{max} is large and may cause the runtime plateauing.

Table 3: Transport Results Error Metrics

Leakage	Error	AM	CLS [6]	LRP [4]	PBS-1 [7]	PBS-2 [7]	CoPS2	CoPS4
Refl.	RMS E_R	1.900	0.380	0.160	0.221	0.087	0.277	0.164
	Mean $\lvert E_R \rvert$	0.517	0.098	0.042	0.052	0.021	0.061	0.033
	Max $\lvert E_R \rvert$	1.232	0.254	0.096	0.163	0.056	0.228	0.144
Trans.	RMS E_R	2.678	0.290	0.235	0.275	0.247	0.197	0.134
	Mean $\lvert E_R \rvert$	0.881	0.084	0.073	0.066	0.060	0.063	0.033
	Max $\lvert E_R \rvert$	1.000	0.181	0.122	0.206	0.209	0.097	0.088
Flux	RMS E_R	1.399	0.163	N/A	0.095	0.044	0.119	0.078
	Mean $\lvert E_R \rvert$	0.394	0.039	N/A	0.022	0.010	0.031	0.019
	Max $\lvert E_R \rvert$	0.861	0.107	N/A	0.058	0.032	0.065	0.050

5. CONCLUSIONS AND FUTURE WORK

An extension of Conditional Point Sampling (CoPS) to multi-D for Markovian-mixed media is presented and transport results are compared with published benchmark and approximate method results. The ability to improve accuracy by using a more accurate conditional probability function for sampling new points is demonstrated by improving accuracy when using more governing points in the conditional probability function. A basic version of CoPS (CoPS2) performs somewhat better than Chord Length Sampling (CLS) and roughly as accurately as the basic version of Poisson Box Sampling (PBS-1). An improved version of CoPS (CoPS4) performs roughly as well as the Local Realization Preserving (LRP) method and an improved version of PBS (PBS-2). CoPS4 is the most accurate of the examined methods for transmittance while PBS-2 is the most accurate for reflectance and internal flux. It is not known how LRP compares to the other methods for internal flux due to lack of data in the literature. Additionally, several numerical studies are performed examining the accuracy and runtime of CoPS as a function of three user-input parameters.

Future work may include resolving benchmark and/or CoPS results to have less uncertainty. Similar numerical results could be generated for 2D problems. Different conditional probability functions, i.e., geared for Poisson-Box mixing or using an approach other than the pseudo-interface approach, could be developed for application to different types of material mixing or for improved accuracy. The Embedded Variance Deconvolution approach demonstrated with CoPS in 1D [10] could be implemented in multi-D to yield not only mean but variance of output quantities. Other methods could be implemented in a similar code framework to enable valid runtime comparisons between CoPS and the other methods. The authors would like to investigate use of biased Woodcock tracking [12] for possible improvements in efficiency and/or accuracy.

6. ACKNOWLEDGMENTS

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell

Table 4: Error and Runtime as a Function of User Variables

Maximum Distance Multiplier, M_d ($M_e = 1.0$, $N_{max} = 3$)								
M_d	0.5	1.0	2.0	3.0	3.75	4.5	6.0	17.3
Refl. RMS E_R	0.464	0.243	0.148	0.109	0.167	0.156	0.159	0.164
Trans. RMS E_R	1.284	0.665	0.219	0.147	0.138	0.157	0.113	0.134
Flux. RMS E_R	0.343	0.149	0.079	0.078	0.073	0.073	0.069	0.078
Runtime [min]	246	306	401	479	526	574	633	813
Exclusion Multiplier, M_e ($M_d = 3.0$, $N_{max} = 3$)								
M_e	0.0	0.001	0.25	0.5	0.75	1.0	2.0	1E+9
Refl. RMS E_R	2.769	0.305	0.134	0.096	0.133	0.141	0.131	0.187
Trans. RMS E_R	15.22	0.628	0.186	0.138	0.217	0.117	0.187	0.150
Flux. RMS E_R	2.049	0.855	0.041	0.068	0.066	0.081	0.072	0.084
Runtime [min]	1011	1579	426	419	416	390	477	381
Maximum Number of Points, N_{max} ($M_d = 3.0$, $M_e = 0.5$)								
N_{max}		0	1	2	3	4	5	
Refl. RMS E_R		1.898	0.273	0.139	0.096	0.113	0.186	
Trans. RMS E_R		2.691	0.167	0.109	0.138	0.144	0.147	
Flux. RMS E_R		1.398	0.125	0.070	0.068	0.070	0.066	
Runtime [min]		45	136	360	392	396	403	

International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

REFERENCES

[1] M. L. Adams, E. W. Larsen, and G. C. Pomraning. "Benchmark results for particle transport in a binary Markov statistical medium." *J Quant Spectrosc and Rad Transfer*, **volume 42**(4), pp. 253–266 (1989).

[2] P. S. Brantley. "A benchmark comparison of Monte Carlo particle transport algorithms for binary stochastic mixtures." *J Quant Spectrosc and Rad Transfer*, **volume 112**, pp. 599–618 (2011).

[3] C. Larmier, F. Hugot, F. Malvagi, A. Mazzolo, and A. Zoia. "Benchmark solutions for transport in d-dimensional Markov binary mixtures." *J Quant Spectrosc and Rad Transfer* (2017).

[4] P. S. Brantley and G. B. Zimmerman. "Benchmark comparison of Monte Carlo algorithms for three-dimensional binary stochastic media." *Trans Am Nucl Soc*, **volume 117**, pp. 765–768 (2017).

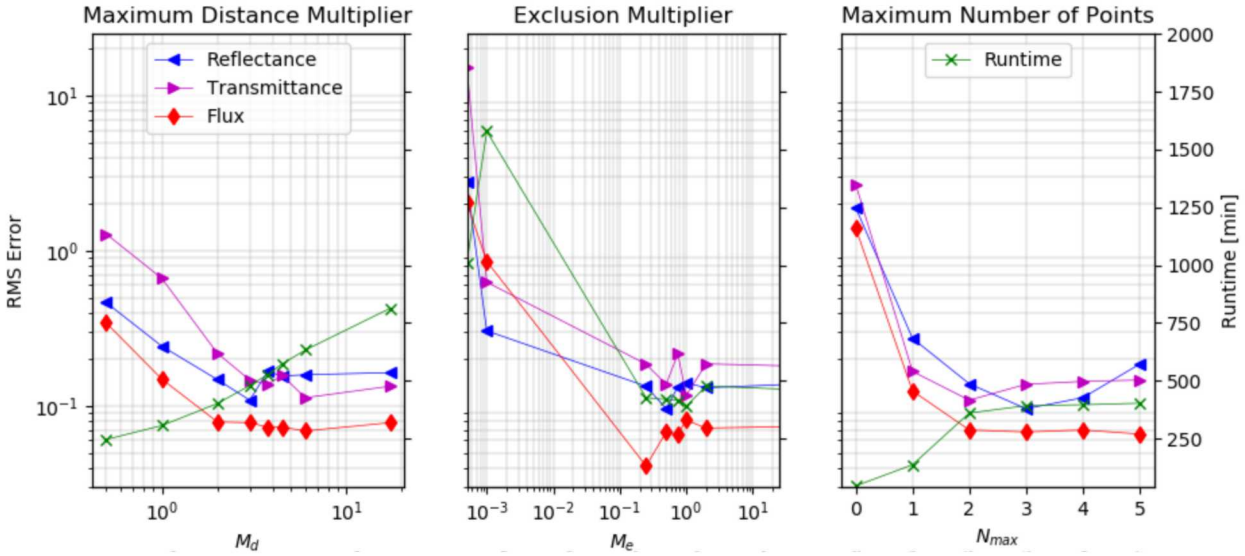

Figure 4: Root mean squared error and runtime as a function of user parameters.

[5] G. B. Zimmerman and M. L. Adams. "Algorithms for Monte-Carlo Particle Transport in Binary Statistical Mixtures." *Trans Am Nucl Soc*, **volume 63**, p. 287 (1991).

[6] C. Larmier, A. Lam, P. Brantley, F. Mavagi, T. Palmer, and A. Zoia. "Monte Carlo chord length sampling for d-dimensional Markov binary mixtures." *J Quant Spectrosc and Rad Transfer* (2018).

[7] C. Larmier, A. Zoia, F. Malvagi, E. Dumonteil, and A. Mazzolo. "Poisson-Box Sampling algorithms for three-dimensional Markov binary mixtures." *J Quant Spectrosc and Rad Transfer* (2018).

[8] C. Larmier, A. Zoia, F. Malvagi, E. Dumonteil, and A. Mazzolo. "Monte Carlo particle transport in random media: The effects of mixing statistics." *J Quant Spectrosc and Rad Transfer* (2017).

[9] E. H. Vu and A. J. Olson. "Conditional Point Sampling: A novel Monte Carlo method for radiation transport in stochastic media." *Trans Am Nucl Soc*, **volume 120** (2019, submitted).

[10] E. H. Vu and A. J. Olson. "An extension of Conditional Point Sampling to quantify uncertainty due to material mixing randomness." In *M&C2019*. American Nuclear Society, Portland, OR (2019, submitted).

[11] A. J. Olson. "Calculation of parametric variance using variance deconvolution." *Trans Am Nucl Soc*, **volume 120** (2019, submitted).

[12] B. Molnar, G. Tolnai, and D. Legrady. "Variance reduction and optimization strategies in a biased Woodcock particle tracking framework." *Nucl Sci Eng*, **volume 190**(1), pp. 56–72 (2018).

[13] P. Switzer. "A random set process in the plane with a Markovian property." *Annals of Mathematical Statistics*, **volume 36**(6), pp. 1859–1863 (1965).

[14] S. D. Pautz, B. C. Franke, A. K. Prinja, and A. J. Olson. "Solution of stochastic media transport problems using a numerical quadrature-based method." In *M&C 2013*. American Nuclear Society, Sun Valley, ID (2013).