On the efficiency of variance reduction techniques for Monte Carlo estimates of imaging noise

Diksha Sharma,[1, a)] Josep Sempau,[2] and Aldo Badano[1]

[1)] *Division of Imaging, Diagnostics, and Software Reliability, OSEL/CDRH/USFDA, Maryland, USA*

[2)] *Institute of Energy Technologies, Polytechnic University of Catalonia, Barcelona, Spain*

(Dated: August 28, 2017)

Purpose: Monte Carlo simulations require large number of histories to obtain reliable estimates of the quantity of interest and its associated statistical uncertainty. Numerous variance reduction techniques (VRTs) have been employed to increase computational efficiency by reducing the statistical uncertainty. We investigate the effect of two VRTs for optical transport methods on accuracy and computing time for the estimation of variance (noise) in x-ray imaging detectors.

Methods: We describe two VRTs. In the first, we preferentially alter the direction of the optical photons to increase detection probability. In the second, we follow only a fraction of the total optical photons generated. In both techniques the statistical weight of photons is altered to maintain the signal mean. We use fastDETECT2, an open-source, freely available optical transport routine from the hybridMANTIS package. We simulate VRTs for a variety of detector models and energy sources. The imaging data from the VRTs simulations is then compared to the analog case (no VRT) using pulse height spectra, Swank factor, and the variance of the Swank estimate.

Results: We analyze the effect of VRTs on the statistical uncertainty associated with Swank factors. VRTs increased the relative efficiency by as much as a factor of 9. We demonstrate that we can achieve the same variance of the Swank factor with less computing time. With this approach, the simulations can be stopped when the variance of the variance estimates reaches the desired level of uncertainty.

Conclusion: We implemented analytic estimates of the variance of Swank factor and demonstrated the effect of VRTs on image quality calculations. Our findings indicate that the Swank factor is dominated by the x-ray interaction profile as compared to the additional uncertainty introduced in the optical transport by the use of VRTs. For simulation experiments that aim at reducing the uncertainty in the Swank factor estimate, any of the proposed VRT can be used for increasing the relative efficiency.

Keywords: Variance reduction techniques, Monte Carlo methods, Swank factor.

[a)]Corresponding author: DS, diksha.sharma@fda.hhs.gov.

I. INTRODUCTION

Monte Carlo simulation of radiation transport often requires a large number of histories to ensure the uncertainty of the estimate is below acceptable limits. However, processing such large number of histories can make the simulations computationally expensive. By monitoring the uncertainty of the estimate, the simulation can be terminated when the desired level of uncertainty is reached. In a typical Monte Carlo radiation transport experiment, we estimate a quantity of interest by calculating the mean value over many histories and the associated uncertainty. Variance reduction techniques (VRTs) can be employed to maximize the computational efficiency, that is, to reduce the variance of the mean for a given simulation time[1]. However, an area of current interest is the simulation of cases where the quantity of interest is not the mean but the variance associated with a random physical process and its associated statistical uncertainty.

Star-Lack *et al.*[2] studied how to reduce the number of optical photons used in a scintillator simulation by combining established methods from Fujita[3], Lubberts[4] and Swank[5], to calculate detective quantum efficiency (DQE) rapidly. The authors show that the mean squared error of a simulated noise quantity for estimating image variability, the noise power spectrum (NPS), is inversely proportional to the number of flood images but not dependent upon the input fluence provided it is above a certain threshold. The authors demonstrate methods to further reduce the number of x-ray photons needed to accurately generate NPS by creating a point-spread function from each event instead of a flood field, reducing the time to model DQE to minutes on a single CPU.

In a related effort, Gutierrez *et al.*[6] developed a technique for estimating variance of the variance for energy resolution metrics like the Swank and Fano factors. They used a first-order approximation to derive a closed form estimate of both the Swank and Fano factors. These estimates were computed using the first five moments (zeroth through fourth moments) of the signal distribution. These can be accumulated during the simulation without significant computational overhead.

In this letter, we introduce and study VRTs associated with the production of optical photons. We describe two techniques and analyze their relative efficiency and their effect on image quality.

II. METHODS

In this section, we describe two types of VRTs implemented for this work, and discuss the different detector models simulated and the image quality metrics used for analyzing the performance of our methods.

A. Variance reduction techniques (VRTs)

In the Directed Sampling (DS) technique, a certain fraction of the optical photons are preferentially directed towards the sensor plane while the rest are directed in the opposite direction. This technique is inspired by the fact that the optical photons going towards the sensor plane are more likely to be detected than the ones going in the opposite direction. In order to maintain the signal mean, the statistical weight of the photons is altered appropriately.

We use the fastDETECT2 routine (part of hybridMANTIS package[7]) for transporting optical photons. In the original version, the number of optical photons are sampled using a Poisson distribution with mean equal to the deposited energy at each x-ray photon interaction divided by the mean energy required to generate an optical photon in CsI. Directional cosines for each optical photon are sampled isotropically. For this work, the fastDETECT2 routine was modified to generate optical photons with a non-isotropically distributed initial random z-direction to direct a certain percentage towards the sensor plane, and the remaining in the opposite direction towards the top surface of the detector. If P_d is the new probability of photons going downward, then the statistical weight of each of these photons is given by $W_d = 0.5/P_d$ where 0.5 in the numerator comes from the fact that in analog simulations (no VRT), isotropically emitted photons have equal probability (50%) of going towards or against the sensor plane. Similarly, the statistical weight of the photons going upwards is given by $W_u = 0.5/(1 - P_d)$. The weight (W_d or W_u) of an optical photon is calculated only once based on P_d or P_u depending upon the z-direction of the photon at the time of generation, and any subsequent changes in this during transport do not affect its weight. The directional cosines are randomly calculated based on P_d.

We also implemented a Weighted Poisson Sampling (WPS) technique. Here we generate only a given fraction of the total number of optical photons, as in a Russian roulette

approach. The surviving photons are then transported in an analog manner, with its directional cosines sampled isotropically. The weight of each surviving photon is increased appropriately to maintain the signal mean. As mentioned earlier, the number of optical photons are sampled using a Poisson distribution. For this method, we modify the mean for sampling by multiplying it with the probability of survival, thus the name weighted Poisson sampling. In order to maintain the overall mean, each surviving photon carries the weight equal to the inverse of the survival probability. For instance, if the survival probability is 0.5, each surviving photon will carry a weight of 2.

B. Detector models

We model a 170 μm thick CsI scintillator screen previously used in the validation of hybridMANTIS against MANTIS and experimental data[9]. In addition, we simulate two detector models. First, we implemented a simplified model (*point*, P model) where all the optical photons are generated in the center of a column, each sampled from 25 keV energy deposition (see Ref. 7 for more details). Another detector model we consider is more realistic (*distributed*, D model). In this model, we obtain PENELOPE outputs for the locations and energy depositions for each x-ray photon interactions. At each of these locations within the scintillator, optical photons are sampled based on the energy deposited (more energy deposited yields higher number of optical photons generated). The difference between a point and distributed detector model is the starting location and the number of optical photons generated. For the distributed model, the x-ray source used was a parallel pencil beam with a diameter of 30 μm. We consider three input energy sources which are used to run PENELOPE - monoenergetic 25 keV, and two spectra at 40 and 70 kVp corresponding to a tungsten anode and 1 mm Al internal filtering[9]. The mean photon energy for the 40 and 70 kVp spectra are 25.6 keV and 36.5 keV respectively.

C. Image quality metrics

In a scintillator detector, the x-ray pulses in energy integrating mode are not individually counted and follow some probability distribution based on the distribution of the incident beam energy, the energy absorbed within the detector and the optical pulse distribution

originating from varying paths of the optical photons during the transport. The last of the three - optical pulse distribution, adds to the scintillation noise of the detector and is measured by the Swank factor[5] given by $\mathcal{I} = {m_1}^2/(m_0 m_2)$, where m_i is the i^{th} moment. The higher the Swank factor, the better is the energy resolution of the detector. We use the Swank factor and the variance of Swank estimates as our figures of merit which, for this work, were calculated based on Ref. 6 using the first five moments obtained from pulse height spectrum (PHS). We report the statistical uncertainty as the coefficient of variation (COV), which is the ratio of the square root of variance of Swank factor and the mean of Swank factor (for details, see Ref. 6).

For our variance analysis, we collect from the simulations the point response function image (PRF), number of detected optical photons per primary (DPP), as well as optical transport statistics. In order to maintain the overall image mean, we add the statistical weight of each optical photon to the corresponding PRF and DPP elements.

Monitoring the uncertainty of these estimates while running the simulation enables us to terminate it when the estimate reaches the desired uncertainty. For this work, we ran a fixed number of histories and collected the estimates every 10 histories because the focus was to compare the effect of various VRT on these estimates. We use root mean square error (RMSE) to analyze the effect of VRTs on the PRFs.

D. Monte Carlo engine

We use hybridMANTIS, a package specifically designed for running coupled x-ray/electron and optical photon transport. It uses PENELOPE[8] for running the x-ray/electron transport and fastDETECT2 for transporting optical photons. For this work, we decoupled the two types of transport - the x-ray transport records the location, energy, and the history number of the x rays depositing energy which is then used as an input for the fastDETECT2 optical transport routine which was modified to incorporate the variance reduction algorithms.

III. RESULTS

Based on the detector models and energy sources, we simulated the following four models: point at 25 keV (P25), distributed at 25 keV (D25), distributed at 40 kVp (D40) and

distributed at 70 kVp (D70). For each model we obtained six sub-cases: analog (no VRT, 50% optical photons going towards and away from the sensor plane), DS with 20% and 80% directed towards the sensor plane (we will refer to them as DS 0.2 and 0.8), and WPS with 10%, 30%, and 50% of optical photons surviving and being transported (referred to hereafter as WPS 0.1, 0.3 and 0.5). All cases used 10,000 x-ray histories. For the distributed models, our optical transport routine fastDETECT2 uses energy deposition events information generated by PENELOPE. For obtaining the PHS and Swank factor plots (Fig. 1) we only accounted for x rays which deposited energy in the scintillator. However, for calculating the variance of Swank factor and COV (Fig. 2) we considered all x rays including those which did not deposit any energy.

The Swank factor for P25 and D25 cases is close to 1 (not shown in the paper) which matches our results from Ref. 9 and drops slightly for the D40 and D70 cases (Fig. 1). This is due to loss of x-ray energy during Compton and fluorescence events. The Swank factor stabilizes around 1,000 histories. The variance of the Swank factor is low - in the range 10^{-7} to 10^{-10} for P25 at 10,000th history and increases for distributed models (Fig. 2). Note that the Swank factor depends more strongly on x-ray energy and detector model than on optical transport variability. Fig. 1 indicates that the Swank factor is mainly dominated by the correlation of x-ray interactions as opposed to by the randomness of the optical transport.

From Fig. 2 we observe noticeable differences between the various VRTs for P25 which is not so for the D25 detector model (we are not showing data for D40 and D70 models because they show trends similar to the D25 model). This is because PHS for the P25 detector model has a single peak corresponding to the full energy absorption events with a much lower variance as compared to the other detector models. In this case, slight differences in the variance of the Swank factor appear more pronounced as compared to other detector models.

We also compared the PRFs for different VRTs with the analog case using RMSE to gauge the effect of a VRT but did not find a noticeable difference. The RMSE values were in the order of 10^{-4} with highest being WPS 0.1 with 10^{-3}. Fig. 3 shows the line spread function (LSF) and modulation transfer function (MTF) for WPS 0.1 for D25 and D70 models. Here we show data for only one VRT to depict the effect of the detector models on the MTF.

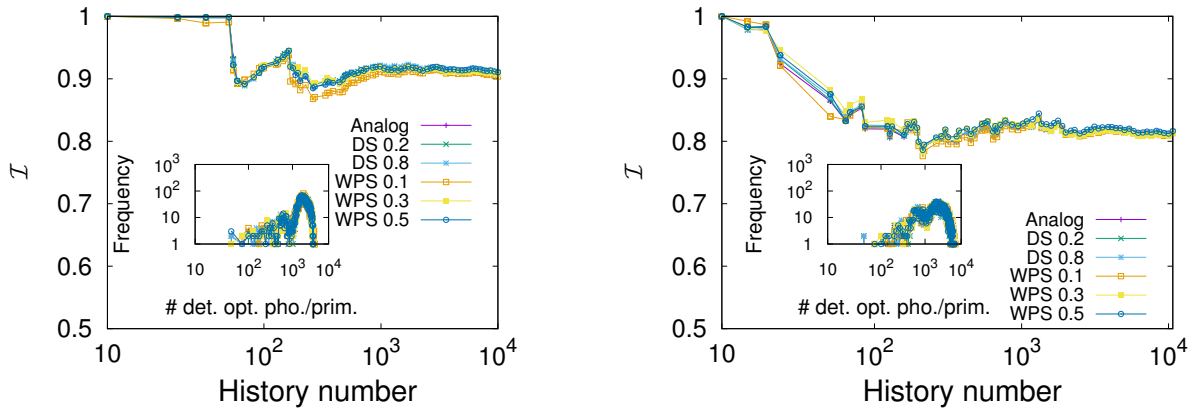

(a)D40 (b)D70

Figure 1. Swank factor (with PHS as inset) for the D40 and D70 detector models.

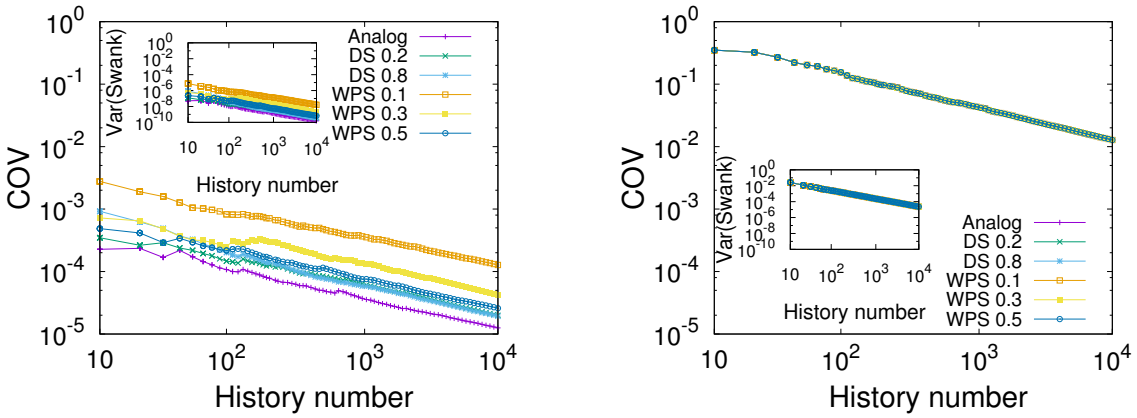

(a)P25 (b)D25

Figure 2. Coefficient of variation (COV) (with variance of the Swank factor as inset) for P25 and D25 detector models.

Table I shows the optical transport statistics for various VRTs for the D25 model. We observe that the number of generated optical photons per primary for DS remains the same as analog while for WPS it was reduced by the fraction of photons killed. For instance, WPS 0.1 would have only 10% optical photons generated as compared to the analog case. The number of detected optical photons per primary also follow a similar pattern. We observe that the WPS 0.1 is about 10 times faster than analog. DS 0.2 increases the average path length because 80% of optical photons are forced to travel away from the sensor and reverse is observed for DS 0.8, as expected. For WPS, this parameter remains constant because the transport of the photons is not affected.

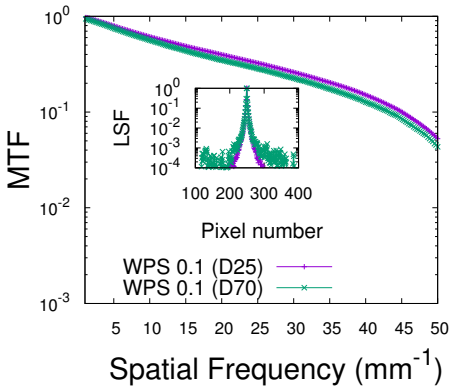

Figure 3. MTF (with LSF as inset) for WPS 0.1 VRT for D25 and D70 models. There was not a significant difference observed between various VRTs for a detector model.

Table I. Optical transport statistics for D25 case for simulating 10,000 primaries.

	Analog	DS 0.2	DS 0.8	WPS 0.1	WPS 0.3	WPS 0.5
Optical photons generated per primary	520	520	520	52	156	260
Optical photons detected per primary	432	416	447	42	130	216
Average path length (μm)	887	1008	767	885	888	889
Total execution time (s) [a]	421	486	362	43	128	212
Speed (x-rays/prim) (hist/s)	24	21	28	230	78	47
Efficiency (ε) (s^{-1})	5850	5105	6754	52854	18780	11792
Efficiency improvement	-	0.87	1.15	9.03	3.21	2.01

[a] Using one six-core AMD® Opteron Processor 2427 with 16 GB RAM[10]

We report relative efficiency as

$$\varepsilon = 1/(\mathrm{var}(\mathcal{I})t), \tag{1}$$

where t is the total computing time. Since the variance of Swank estimates are very close for different VRTs, ε follows the reverse trend as time. WPS 0.1 is about 9x more efficient than analog and 8x as compared to DS 0.8. We observe that WPS is much more efficient in increasing computational speed and efficiency as compared to DS.

In Monte Carlo simulations, it is essential to have an indication of how much computing time would be necessary to achieve a certain level of precision. Fig. 4 shows how the variance

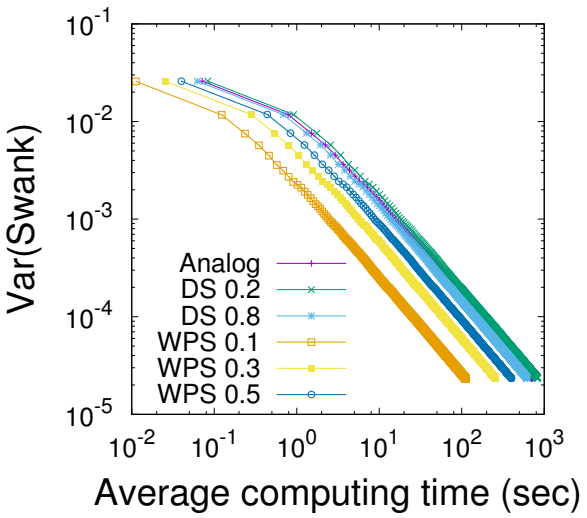

Figure 4. Variance of the Swank factor vs. computing time for various VRTs for the D25 model with 10,000 histories.

of the Swank factor reduces with time for different VRTs for the D25 model. Variance of the Swank factor was taken from Fig. 2 and the average computing time was obtained by dividing the history number with the average speed of the model. This assumes that all the histories take approximately the same time. This important figure shows that VRTs can attain the same variance of the Swank factor as the analog simulation albeit faster. For instance, a variance of the Swank factor of 10^{-4} can be reached in 40 seconds using WPS 0.1 as compared to 150 seconds with analog simulation. By monitoring the variance of the variance estimate during the simulation one can determine when to stop the simulation when the desired level of uncertainty is reached.

IV. DISCUSSION

Fig. 1 shows that for the D40 and D70 detector models although the Swank factor is noisier in the beginning, it starts to converge between 1000 and 10,000 histories. This is an important finding which indicates that if the aim of a simulation experiment is to reduce the variance of the Swank estimate, any of the proposed VRT can be used to obtain a better estimate of the Swank factor with fewer histories thus reducing the computational bottleneck of the optical transport.

We observe that the Swank factor depends more on the input energy and the choice of

detector model than on the uncertainty introduced by the optical transport method. From Fig. 1 we observe that the value of the Swank factor is mainly dominated by the x-ray interactions in the scintillator as opposed to by the randomness introduced by the optical transport, suggesting the need to consider VRTs for x-ray and electron transport.

In order to better understand the correlation between the Swank factors for various VRTs, we performed another experiment (data not shown in this paper) using a different x-ray interaction profile. We observed that the Swank factors were not correlated with each other as seen in Fig. 1 which again confirms that their value is dominated by the x-ray interactions. We also ran the same cases for a thicker CsI screen (450 μm) and observed similar trends suggesting that the thickness of the screen does not significantly impact either the Swank factor or the variance of the Swank estimates.

Our results suggest that a WPS technique is effective in speeding up optical transport simulations. However, it might be possible that a WPS with a smaller survival probability than the ones studied in this work might have significant effects on other quantities including MTF. In addition, an effective VRT for optical transport might make the most significant bottleneck for the imaging simulations to be not in the optical transport as it is typically the case, but in the x-ray and electron transport component. More research is needed to determine optimal settings for using VRT on imaging simulations.

V. CONCLUSION

More efficient simulations of indirect imaging detectors where optical transport is typically the slowest component in the imaging pipeline are needed. Here, we implemented a previously described analytic method to estimate the variance of the Swank factor and demonstrated the effect of VRTs on the image quality parameters. Our findings suggest that VRTs are useful in reducing the computational bottleneck of the optical transport in Monte Carlo simulations. We analyzed various VRTs and found that x-ray interactions are the dominant factor contributing to the Swank factor as compared to randomness in the optical transport. VRTs increased the relative efficiency by as much as a factor of 9. Monitoring the uncertainty of the variance of the variance estimates helps determine when to stop the simulation for a desired level of precision.

REFERENCES

[1]T. M. Jenkins, W. R. Nelson, A. Rindi. Monte Carlo Transport of Electrons and Photons Chapter 18, *Plenum Press*, 407-419, 1988.

[2]J. Star-Lack, M. Sun, A. Meyer, D. Morf, D. Constantin, R. Fahrig, E. Abel. Rapid Monte Carlo simulation of detector DQE(f). *Medical Physics*, 41(3):031916, 2014.

[3]H. Fujita, D. Tsai, T. Itoh, K. Doi, J. Morishita, K. Ueda, A. Ohtsuka. A simple method for determining the modulation transfer function in digital radiography. *IEEE Trans. Med. Imaging*, 11(1):34-39, 1992.

[4]G. Lubberts. Random noise produced by x-ray fluorescent screens. *J. Opt. Soc. Am.*, 58(11):1475-1482, 1968.

[5]R. Swank. Absorption and noise in x-ray phosphors. *Jour. Applied Physics*, 44:4199-4203, 1973.

[6]B. Gutierrez, A. Badano, F. Samuelson. Analytic variance estimates of Swank and Fano factors. *Medical Physics*, 41:072102, 2014.

[7]D. Sharma, A. Badal, A. Badano. hybridMANTIS - a CPU-GPU Monte Carlo method for modeling indirect x-ray detectors with columnar scintillators. *Phys. in Med. Biol.*, 57:2357-2372, 2012.

[8]J. Sempau, A. Badal, L. Brualla. A PENELOPE-based system for the automated Monte Carlo simulation of clinacs and voxelized geometries - application to far-from-axis fields. *Medical Physics*, 38:5887-5895, 2011.

[9]D. Sharma, A. Badano. Validation of columnar CsI x-ray detector responses obtained with hybridMANTIS, a CPU-GPU Monte Carlo code for coupled x-ray, electron, and optical transport. *Medical Physics*, 40(3):031907, 2013.

[10] The mention of commercial products herein is not to be construed as either an actual or implied endorsement of such products by the Department of Health and Human Services. This is a contribution of the Food and Drug Administration and is not subject to copyright.