

History by history statistical estimators in the BEAM code system

B. R. B. Walters,^{a)} I. Kawrakow, and D. W. O. Rogers

Ionizing Radiation Standards, National Research Council of Canada, Ottawa, KIA OR6, Canada

(Received 4 March 2002; accepted for publication 6 September 2002; published 15 November 2002)

A history by history method for estimating uncertainties has been implemented in the BEAMnrc and DOSXYZnrc codes replacing the method of statistical batches. This method groups scored quantities (e.g., dose) by primary history. When phase-space sources are used, this method groups incident particles according to the primary histories that generated them. This necessitated adding markers (negative energy) to phase-space files to indicate the first particle generated by a new primary history. The new method greatly reduces the uncertainty in the uncertainty estimate. The new method eliminates one dimension (which kept the results for each batch) from all scoring arrays, resulting in memory requirement being decreased by a factor of 2. Correlations between particles in phase-space sources are taken into account. The only correlations with any significant impact on uncertainty are those introduced by particle recycling. Failure to account for these correlations can result in a significant underestimate of the uncertainty. The previous method of accounting for correlations due to recycling by placing all recycled particles in the same batch did work. Neither the new method nor the batch method take into account correlations between incident particles when a phase-space source is restarted so one must avoid restarts. [DOI: 10.1118/1.1517611]

I. INTRODUCTION AND THEORY

I.A. The batch approach to uncertainties

The BEAM code system¹ is a widely used Monte Carlo code for simulating radiotherapy beams and calculating dose distributions in patients (see Ref. 2 for a listing of over 150 publications). We have upgraded the BEAM code system to create BEAMnrc,^{3,4} which uses the recently released EGSnrc Monte Carlo code for radiation transport.^{5,6} At the same time we have improved the method for estimating uncertainty in dose and fluence calculated in BEAMnrc and DOSXYZnrc. Previously, calculation of uncertainty in all quantities depended on splitting calculations into statistical batches (usually 10) and then, once the simulation was finished, taking the estimate of the uncertainty in the average of a scored quantity, X , to be:

$$s_{\bar{X}} = \sqrt{\frac{\sum_{i=1}^N (X_i - \bar{X})^2}{N(N-1)}}, \quad (1)$$

where N is the number of batches, X_i is the value of X in batch i , and \bar{X} is the mean value of X evaluated over all batches.

There are three problems with the batch approach. First, unless a large number of statistical batches is used, there are significant fluctuations in the uncertainty itself since the sample size, N , in Eq. (1) is quite small. Second, arbitrarily grouping histories into batches ignores any correlations between incident particles. Incident particles will be correlated when phase-space data from a BEAMnrc simulation of an accelerator are used as a source, especially if variance reduction techniques, such as bremsstrahlung splitting and photon forcing, are used in the accelerator simulation. Finally, the implementation of the batch approach used in most NRC user-codes added an extra dimension (storing the results

from each batch) to all arrays scoring quantities of interest. This last limitation was especially evident in DOSXYZnrc, where the scoring arrays are already large because of the large number of geometrical regions (e.g., 2 million voxels for $128 \times 128 \times 128$ resolution). It should be noted however that this is a reflection of how things were coded in DOSXYZnrc and could have been avoided using other coding techniques.

I.B. The history by history method

In order to eliminate the problems with estimating uncertainty using batches, we have adopted a clever trick (attributed⁷ to Salvat of the University of Barcelona) for efficiently implementing the history by history method for estimating uncertainty. It has been described by Sempau *et al.*⁸ The history by history method is well known and has been used for years in other codes (e.g., since at least 1986 in MCNP⁹). Andreo¹⁰ has also pointed out its advantages. However, the brute force application of this method is inefficient and the inherent improvements did not justify the increase in computation time required to update all scored quantities after each history, especially in calculations with many scoring bins. Use of Salvat's approach removes this increase in computation time. We also found it necessary to make modifications to account for correlations between incident particles when a phase-space source is used. Similar modifications were suggested independently by Sempau *et al.*⁸

Returning to Eq. (1), let X_i be the quantity scored in statistically independent event i (i.e., history i instead of batch i). The equation can be rewritten as

$$s_{\bar{X}} = \sqrt{\frac{1}{N-1} \left(\frac{\sum_{i=1}^N X_i^2}{N} - \left(\frac{\sum_{i=1}^N X_i}{N} \right)^2 \right)}, \quad (2)$$

where N is now the number of independent events, i.e., histories. In BEAMnrc and DOSXYZnrc, when using phase-

space sources, one event or history is defined to be all particle tracks associated with one initial particle (either exiting the accelerator vacuum or from a decaying radioactive source such as ^{60}Co). It should be noted that the X_i may be weighted quantities (if variance reduction techniques, such as bremsstrahlung splitting, are used), while N is an unweighted quantity, always equal to the total number of independent, or primary, histories.

If we keep track of $\sum_{i=1}^N X_i^2$ and $\sum_{i=1}^N X_i$ on the fly, then we can calculate the uncertainty at the end of the simulation without the need to store the scored quantity in batches. The problem is that when there are a large number of quantities being scored, it can be very computationally inefficient to evaluate $\sum_{i=1}^N X_i^2$ at the end of each history. To overcome this problem, Sempau *et al.*⁸ outlined the following algorithm for quantity X :

```
IF(nhist=X_last) THEN
  X_tmp=X_tmp+delta
ELSE
  X=X+X_tmp
  X2=X2+(X_tmp)**2
  X_tmp=delta
  X_last=nhist
ENDIF
```

where X stores $\sum_{i=1}^N X_i$ during the run, but, after analysis, will store the quantity \bar{X} , `nhist` is the current history number, `X_last` is the number of the last history that contributed to X , `X_tmp` stores the sum of the contributions to X during the current history, `delta` is a contribution to X during the current step, and `X2` stores $\sum_{i=1}^N X_i^2$.

Using the algorithm outlined above together with Eq. (2), the three main problems with using the batch method are eliminated. The problem of small sample size is eliminated, since N is now the number of histories and is usually large for a calculation with reasonable statistics. Also, if phase-space data from BEAMnrc are used as a source then, by ensuring that N and `nhist` only count primary histories (i.e., histories from the original nonphase-space source), we properly take into account correlations between incident particles. Finally, the additional 10 or more dimensions required to store the value of X in each batch have been eliminated and three new scoring arrays, `X_tmp`, `X2`, and `X_last`, have been introduced (X is retained from the batch method, but with the batch dimension eliminated). However, X and `X2` have been made double precision to avoid any potential round-off errors and taking into account the other large arrays, the memory requirement for large arrays is only reduced by a factor of 2 compared to when 10 batches were used.

Some quantities output by BEAMnrc are actually ratios of correlated quantities. An example of this is the average energy of photons in a scoring zone, which is given by the total photon energy crossing the zone divided by the total number of photons crossing the zone. These quantities are correlated with each other, and, if output separately, would each have their own uncertainty. In order to estimate uncertainty on

these ratios using the history by history method, we use the equation for the fractional uncertainty on a ratio of correlated quantities, $C=X/Y$:

$$\frac{s_{\bar{C}}}{\bar{C}} = \sqrt{\left(\frac{s_{\bar{X}}}{\bar{X}}\right)^2 + \left(\frac{s_{\bar{Y}}}{\bar{Y}}\right)^2 - \frac{2\text{cov}(X,Y)}{(N-1)(\bar{X}\bar{Y})}}, \quad (3)$$

where $s_{\bar{X}}$ and $s_{\bar{Y}}$ are the uncertainties on \bar{X} and \bar{Y} estimated using the history by history method outlined above, and $\text{cov}(X,Y)$ is the covariance of X and Y , given by:

$$\text{cov}(X,Y) = \frac{\sum_{i=1}^N X_i Y_i}{N} - \frac{\sum_{i=1}^N X_i \sum_{i=1}^N Y_i}{N^2}. \quad (4)$$

In order to calculate $\text{cov}(X,Y)$, we need to keep track of $\sum_{i=1}^N X_i Y_i$ on a history by history basis. This is done on the fly using an algorithm similar to the one given above used to keep track of $\sum_{i=1}^N X_i^2$ and $\sum_{i=1}^N X_i$. Keeping track of $\sum_{i=1}^N X_i Y_i$ requires an additional REAL*8 variable for each ratio scored.

Note that keeping track of primary histories in a phase-space source has made it necessary to modify the format of BEAMnrc phase-space files slightly. We now mark the first particle scored in the phase-space file from each primary history by setting the particle energy negative. Then, when the phase-space file is used as a source, we increment N and `nhist` only when a negative energy is read.

Ma *et al.*¹¹ have used a similar method to analyze uncertainty in their Monte Carlo dose calculation code, MCDOSE. However, in their approach, the quantity of interest (in this case, energy deposited) is not grouped according to primary history, but according to each energy deposition event. Thus, if a primary history gives rise to two or more charged particles (through interactions and/or variance reduction techniques, such as bremsstrahlung splitting), their method would put the energy deposited by each charged particle in a different group, whereas our technique would put the energy deposited by all resultant charged particles in the same group. It will be seen below that the event-by-event method can give rise to errors in the uncertainty estimate, especially if particles in a phase-space source are recycled.

I.C. Latent variance of a phase-space file

Sempau *et al.*⁸ have introduced the term ‘‘latent variance of a phase-space file’’ to distinguish between the uncertainty in a dose calculation due to the random nature of the transport in the phantom versus that due to the statistical fluctuations in the phase-space file. Thus, a Monte Carlo calculation of the dose distribution generated by a 20 MeV pencil beam of electrons has a statistical uncertainty despite the fact that there are no fluctuations in the incident phase-space of the source. This is the inherent uncertainty of the dose calculation. This uncertainty will approach zero if the number of histories is increased sufficiently. However, when a finite phase-space file is used as a source, the statistical uncertainty on the calculated doses approaches a finite value, independent of how often the phase-space file is reused. This value represents the latent variance of the phase-space file.

II. RESULTS AND DISCUSSION

II.A. New vs old methods without recycling

To see the effect of the improved uncertainty analysis, we simulated an 18 MeV electron beam (field size = 20×20 cm at SSD = 100 cm) from a Clinac 2100C accelerator and examined the uncertainty in the calculated dose in a water phantom. This was a two-stage process. The accelerator was simulated using BEAMnrc, and a phase-space file was scored at the SSD. This phase-space file was then used as a source in DOSXYZnrc simulations of a water phantom.

In the BEAMnrc accelerator simulation, 56 million primary histories were used to generate a phase-space source containing 51 million particles (including 34 million photons). Range rejection for particles below 3 MeV was used in the accelerator simulation, however, no variance reduction techniques were used which could have led to multiple particles in the same phase-space file from the same primary history (e.g., bremsstrahlung splitting). Nonetheless, there can be multiple particles in the phase-space file for a given primary history, e.g., the primary electron and bremsstrahlung photons and/or knock-on electrons.

In the DOSXYZnrc simulation of the water phantom 50 mil incident particles (of all types) from the phase-space source were used. The simulated water phantom itself had dimensions of $20 \times 20 \times 15$ cm, and dose was scored in $1 \times 1 \times 0.5$ cm voxels down the central axis. Region-by-region range rejection for particles below 5 MeV was used in the DOSXYZnrc simulation. This means that a charged particle history was terminated and its energy deposited locally if the particle energy was ≤ 5 MeV and the particle could not make it to the nearest voxel boundary with energy $> ECUT$, the low-energy threshold for particle transport (700 keV in this case). This range-rejection scheme saves CPU time by cutting short unnecessary transport in the large, off-axis voxels.

Figure 1 shows fractional uncertainty in dose versus depth estimated using the new method (grouping energy deposited according to primary history) and using the old method with 10 and 40 batches. A scaled depth-dose curve is also shown for reference. Incident particles from the phase-space file were not recycled in these simulations.

It is clear from Fig. 1 that the new method estimates a much smoother uncertainty versus depth curve than the batch method, indicating a much lower uncertainty on the uncertainty estimate. It is also clear that increasing the number of arbitrary batches from 10 to 40 results in reduced fluctuations in the uncertainty, albeit at a cost of increasing the memory requirement for the scoring arrays by a factor of 4. Despite fluctuations, the mean uncertainty in all doses $> 0.5D_{\max}$ estimated using the batch method (0.55% using 10 batches and 0.53% using 40 batches) is in good agreement with that estimated using the new method (0.54%). The overall agreement between the three methods (aside from fluctuations) implies that any correlation between multiple particles scored in the same primary history, which were placed in separate batches in the old technique, is negligible.

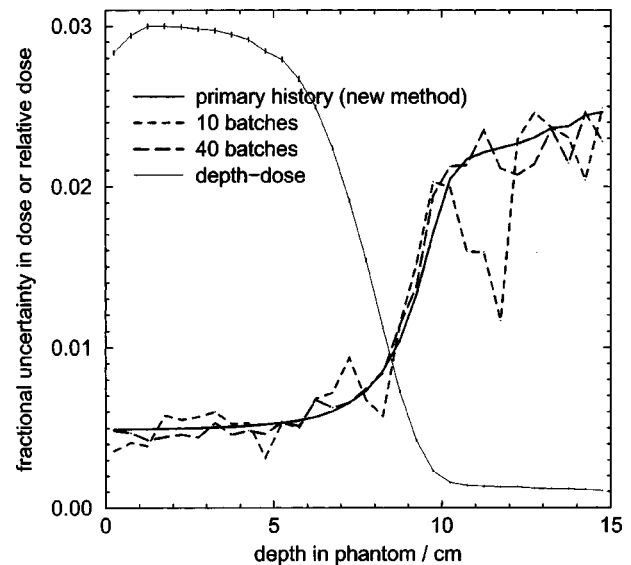


Fig. 1. Fractional uncertainty in dose calculated on the central axis vs depth for a simulated 18 MeV beam from a Clinac 2100C (20×20 cm field at SSD = 100 cm) in a simulated water phantom. Fractional uncertainties are estimated using the new method (grouping by primary history) and using the old method with 10 and 40 batches. A scaled depth-dose curve is also shown for reference (with uncertainties estimated using the history by history method). Dose was scored in $1 \times 1 \times 0.5$ cm voxels on the central axis.

II.B. Effects of recycling

In many cases, the phase-space data available at the bottom of a simulated accelerator may be relatively sparse. In such cases, reducing the inherent uncertainty in the dose calculation to an acceptable level can only be achieved by recycling each particle in the phase-space source before moving on to the next one. In order to investigate the effects of particle recycling on the uncertainty, the water phantom simulation was repeated with the same number of incident particles (50 million), but this number was achieved by using only 12.5 million particles from the phase-space file with each particle recycled 3 times (each particle used a total of 4 times—the recommended maximum for electron beams) and then by using only 1.8 million particles from the phase-space file with each particle recycled 27 times (each particle used a total of 28 times). A DOSXYZnrc technique called smoothing, in which incident particles are reflected about the X axis and the Y axis before being re-used in the simulation, was also used, but this has little impact on the central axis (see Sec. II.E). Note that when a particle is recycled, the number of primary histories (n_{hist} in the new algorithm described in Sec. I) is not incremented. Thus, energy deposited by all occurrences of a particle is grouped into the same primary history.

Figure 2 shows the effect of particle recycling on the estimated uncertainty on the central axis. The “no recycling” curves are the same curves as shown for the new method and using 10 batches in Fig. 1. The uncertainty estimated using the new method clearly indicates that uncertainty increases with the number of times each particle is recycled. The increase is greatest at the phantom surface where the fluctuations reflect the statistical fluctuations of the initial particles

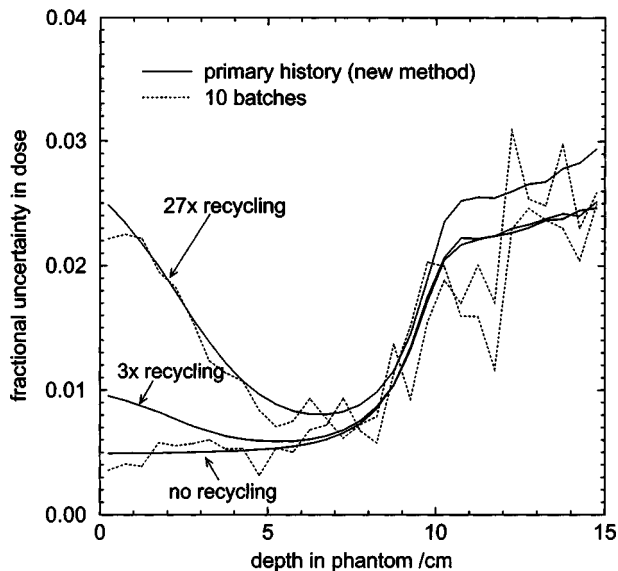


FIG. 2. Fractional uncertainty in central-axis dose vs depth for a simulated 18 MeV beam from a Clinac 2100C (20×20 cm field at SSD = 100 cm) in a simulated water phantom. Fractional uncertainties estimated using the new method are shown for cases in which incident particles were recycled 27 times, 3 times (the recommended maximum), and not at all. Fractional uncertainties estimated using 10 batches (old method) are shown for the cases in which particles were recycled 27 times and not at all. The total number of particle tracks simulated was 50 million in all cases. Dose was scored in $1 \times 1 \times 0.5$ cm voxels on the central axis.

hitting just that voxel whereas at depth, the particles from a wider initial area are involved and the in-phantom statistical variations in repeated histories play a role, thereby decreasing the uncertainty.

Recycling incident particles 3 times is equivalent to decreasing the number of primary histories by a factor of 4 (recall the number of primary histories is not incremented when a particle is recycled), and, thus, in an electron beam the uncertainty at the surface of the phantom is expected to increase by a factor of $\sqrt{4} = 2$ because the uncertainty is dominated by the latent variance of the phase-space file. This matches the increase at the surface shown in Fig. 2. Similarly, when particles are recycled 27 times, the uncertainty at the surface is expected to increase by a factor of $\sqrt{28} \sim 5$, which matches the increase shown in the figure. The uncertainty estimated using 10 batches in the case of $27 \times$ recycling shows expected fluctuations, however, similar to the “no recycling” case, the mean uncertainty is in fairly good agreement with that estimated using the new method (mean uncertainty on all doses $> 0.5D_{\max}$ is 1.30% using 10 batches and 1.38% using the new method). This agreement validates our previous scheme to account for correlations between recycled particles by putting them in the same statistical batch (accomplished by having the recycling loop inside the statistical batch loop in the code).

When recycling photon sources, the effects on the uncertainty are much less dramatic since, except when highly recycled, photons are not likely to interact in the same voxel and therefore the recycling has little effect. For example, recycling a 13 MV beam 3 times had no effect on the central-

axis dose uncertainties and recycling 27 times only increased the uncertainty by about 40% rather than the factor of 5 seen at the surface in the electron beam (Fig. 2).

The above-noted results indicate that both the new and old methods of statistical analysis took the latent variance of the phase-space file (as discussed in Sec. I) into account properly. Thus no matter how often a given particle was recycled, the uncertainty would only decrease to a fixed value which reflected the latent variance of the phase-space file. This is not the case if the phase-space file is restarted rather than each particle being recycled.

II.C. Effects of restarting

Problems arise in the new and old methods of estimating uncertainty when a phase-space file is restarted. This happens automatically upon reaching the end of the source file. A particle that is re-used because of a restart will not be grouped into the same primary history as it was on the previous pass through the phase-space file, correlations between re-used particles will be ignored, and the uncertainty will be underestimated, as shown above. Thus, for the purpose of estimating uncertainty, it is recommended that particles be recycled enough times so that the phase-space file is only used once. To make this feasible BEAMnrc and DOSXYZnrc offer an option in which the number of times to recycle a particle is automatically calculated based on the number of particles in the phase-space source and the number of requested histories. This option may underestimate the number of times to recycle a particle because it cannot take into account particles that do not get used because they are outside the geometry, do not have the correct LATCH value, etc., in which case the phase-space source will be restarted. If the source is restarted only once and only a small fraction of it is re-used on the second pass, this is unlikely to have a significant effect on the uncertainty. However, if most of the source is re-used on the second pass, or if it is restarted more than once, then a manually-calculated value for the number of times to recycle particles (based on data from the previous run in which the value was automatically calculated) is recommended.

II.D. Effects of correlations

To investigate the effects of correlations, the uncertainty has also been estimated by grouping deposited energy according to incident particle instead of primary history. In this case, *nhist* in the algorithm described in Sec. I is always incremented, even if the particle is being recycled. This method is similar to the new method but ignores correlations between incident particles when a phase-space source is used. When this method was applied to the simulation of the electron beam in water then, regardless of how many times incident particles were recycled, uncertainties were found to agree exactly with those estimated using the new method with no recycling (shown in Fig. 2). The exact agreement in the case of no recycling indicates that correlations between particles in the phase-space file did not affect uncertainty in this electron beam simulation. More importantly, the distinct

differences in the cases of recycling indicate that unless correlations introduced by recycling are taken into account, the uncertainty is significantly underestimated because the latent uncertainty of the phase-space data is not accounted for.

In electron beams one does not expect significant correlations between particles. To investigate the effects on estimated uncertainties of correlations between incident particles (other than those introduced by particle recycling) we used BEAMnrc to simulate a 13.5 MV photon beam (10×10 cm field at SSD=100 cm) from a Siemens KD2 accelerator and used selective bremsstrahlung splitting (maximum splitting number of 400, minimum splitting number of 40) to enhance photon output. When bremsstrahlung splitting is used, each primary history can potentially generate a large number of photons, all of which are correlated, at the bottom of the accelerator. For this simulation, 25.5 million primary histories were run to generate a phase-space file containing 63.5 million particles (63 million of which were photons).

The phase-space file generated at the bottom of the photon accelerator was then used as a source in a DOSXYZnrc simulation of a water phantom (dimensions $20 \times 20 \times 60$ cm). Range rejection ($E_{\text{SAVE}} = 5$ MeV) was used in all DOSXYZnrc simulations. Initially, we ran 50 million incident particles from the phase-space source (no particle recycling) and examined uncertainty on the dose in $2 \times 2 \times 0.5$ cm voxels down the central axis of the phantom. It was found that this did not include enough of the incident beam field to show any effects of correlations between incident particles. By increasing the size of the central-axis voxels to $10 \times 10 \times 0.5$ cm, with a corresponding decrease in the number of incident particles to 1.5 million (no recycling), we were able to see the effects of correlations.

Figure 3 shows the estimated uncertainty in dose versus depth for the photon beam simulation. Uncertainties were estimated by grouping deposited energy according to primary history (the new method), by splitting the run into 40 statistical batches (the old method), and by grouping energy deposited according to incident particle (similar to new method but ignoring correlations). A scaled depth-dose curve is also shown in Fig. 3 for reference. It is interesting to note that the 1.5 million incident particles used in this simulation were found to represent only 209 196 primary histories, indicating a fairly high degree of correlation between photons in the phase-space source.

From Fig. 3 it is clear that, similar to the electron beam example, the new method of estimating uncertainty yields a much smoother uncertainty versus depth curve than when 40 batches are used, indicating a much lower uncertainty on the uncertainty estimate. Figure 3 also shows that ignoring correlations between incident photons by grouping according to incident particle results in a consistent underestimate of the uncertainty (although the curve is still smooth). The batch method also ignores correlations between particles and in Fig. 3 it can be seen that it also tends to underestimate uncertainty. The mean uncertainty on all doses $> 0.5D_{\text{max}}$ estimated using 40 batches is 0.64%, which is in better agreement with the mean of 0.63% estimated by grouping

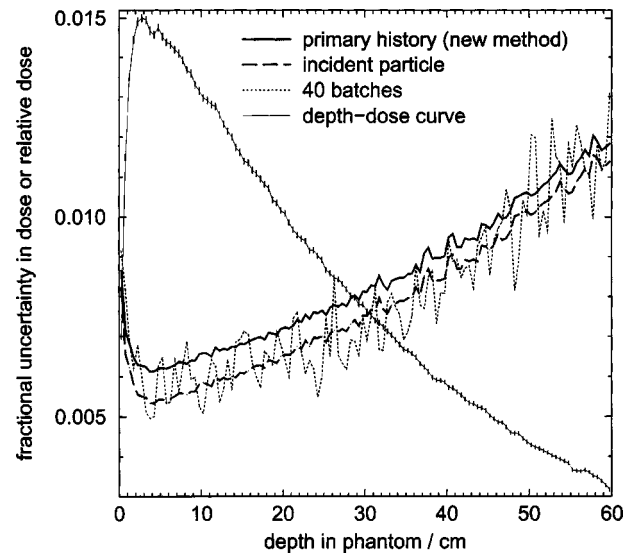


Fig. 3. Fractional uncertainty in central-axis dose vs depth for a simulated 13.5 MV photon beam from a Siemens KD2 accelerator (10×10 cm field at SSD=100 cm) in a simulated water phantom ($20 \times 20 \times 60$ cm). Uncertainties were estimated three ways: using the new method (grouping by primary history); using 40 statistical batches; and using a method similar to the new method but which ignores correlations between incident particles (grouping by incident particle). A scaled depth-dose curve is also shown for reference (with uncertainties estimated using the history by history method). Dose was scored in $10 \times 10 \times 0.5$ cm voxels on the central axis.

according to incident particle than it is with the mean of 0.70% estimated using the new method.

One surprising result from Fig. 3 is the similarity between the shape of the uncertainty versus depth curve estimated using the new method and that estimated by grouping according to incident particle, even the small fluctuations. In fact, they appear to be the same curve, just translated in the Y direction. This similarity was also found to exist in the absolute uncertainties estimated using these two methods (i.e., once fluctuations in dose had been factored out). This indicates that there are very few cases in which multiple photons originating from the same primary history lead to energy deposited in the same voxel. Thus, differences between uncertainty estimated using the new method and by grouping according to incident particle are mainly due to differences in the number N in Eq. (2). Using the new method, N is the number of primary histories, while grouping according to incident particle, N is the number of incident particles.

II.E. Effects of smoothing on uncertainties

When recycling or restarting a phase-space source in DOSXYZnrc, one has the option of using a routine which makes use of the symmetry in many beams to redistribute the phase-space particles to 3 symmetrical positions $[(x,y)$ with direction cosines (u,v) goes to $(-x,y)$ with $(-u,v)$ etc.]. This means that, away from the axes, one can effectively gain a factor of 4 increase in the number of different initial particles incident on a particular region since one gets them from all four quadrants. Unfortunately, close to the central axis this has little value since all four locations are very close

to each other. This is clear in Fig. 2 where the uncertainty on the central-axis dose increased by a factor of 2 near the surface, even with smoothing turned on. At greater depths, the recycling causes only a small increase in the uncertainty, but the accelerator simulation required four times fewer histories. We find that for a voxel element away from the axes, with smoothing on, the uncertainty estimate for the 3 times recycling case is the same as in the no recycle case. Hence, for electron beams the use of the smoothing option with the recycling option clearly has a positive benefit in much of the volume, but has no benefit on the central axis at the surface, and also a reduced benefit along the x and y axes. In any case, the uncertainty estimates are accurate.

For photon beams the effects of smoothing are not so dramatic since the effects of recycling are not as dramatic as discussed earlier. Nonetheless, the trends are similar to the effects with electron beams, namely smoothing is most effective away from the axes.

II.F. Ratios of correlated quantities

In order to determine how well the history by history method is able to estimate uncertainties in ratios of correlated quantities [using Eqs. (3) and (4)], we examined the χ^2 per degree of freedom for average photon energies at the bottom of a simulated accelerator. As mentioned in Sec. I.B. average photon energy is a ratio of the correlated quantities total photon energy and number of photons.

We performed 20 separate simulations (500 000 histories each) of a generic 16 MV photon accelerator (field size = 10×10 cm at SSD = 100 cm) and calculated the average photon energies (with uncertainties) for photons crossing 10 scoring zones at the SSD. The scoring zones were square “rings” centered on the beam axis with midpoints at a distance 0, 1.5, 2.5, 3.5, 4.5, 5.5, 6.5, 7.5, 8.5, and 9.5 cm from the beam axis. We then estimated the χ^2 per degree of freedom in each scoring zone using:

$$\frac{\chi^2}{df} = \frac{1}{N-1} \sum_{i=1}^N \frac{(E_i - \bar{E})^2}{s_{E_i}^2}, \quad (5)$$

where N is the number of simulations (20 in this case), E_i is the average photon energy crossing the scoring zone in simulation i , s_{E_i} is the uncertainty on E_i , and \bar{E} is the photon energy crossing the scoring zone averaged over all N simulations.

The results are shown in Fig. 4. If the uncertainty is estimated accurately then we expect χ^2/df to be ≈ 1 . A $\chi^2/df \ll 1$ indicates that the uncertainty has been overestimated, and a $\chi^2/df \gg 1$ indicates an underestimate of the uncertainty. Figure 4 shows that χ^2/df is ≈ 1 for most scoring zones.

We have found that for a simulation with extreme weight variations, the χ^2/df was about 1 except for one or two scoring zones where it was about 2, caused by one or two outliers in the 20 calculations. We take this as an indicator of using too much biasing, causing inconsistent results due to “fat” particles with very large weights relative to most par-




Fig. 4. χ^2 per degree of freedom for average photon energy at the SSD (100 cm) of a generic 16 MV photon accelerator evaluated over 20 separate simulations (500 000 histories each). Scoring zones in which average photon energies (and their uncertainties) were calculated were square “rings” centered on the beam axis. Midpoints of the square rings were at 0, 1.5, 2.5, 3.5, 4.5, 5.5, 6.5, 7.5, 8.5, and 9.5 cm from the beam axis. Field size was 10×10 cm at the SSD.

ticles. This could possibly even be developed into a diagnostic tool for the variance reduction techniques.

II.G. A note on parallel runs

If one has multiple machines running on a network, then it is possible to split BEAMnrc and DOSXYZnrc simulation up into a number of parallel jobs.⁴ When a simulation that uses a phase-space source is split up into n multiple jobs, then the phase-space source is automatically divided into n equal partitions. Each job uses a separate partition. Thus, job i uses particle numbers that fall in the range:

$$(i-1) \times \frac{nshist}{n} < nnpshp \leq i \times \frac{nshist}{n}, \quad (6)$$

where $nshist$ is the total number of particles in the phase-space source and $nnpshp$ is the number of the particle used. This partitioning scheme was adopted to ensure that the entire phase-space source is adequately sampled over all of the parallel jobs.

Partitioning of phase-space sources potentially interferes with the new method of estimating uncertainties because it may split up particles generated by the same primary history (correlated) and use them in different jobs. When the jobs are recombined for final analysis, any quantity scored by these particles will be grouped as if it originated in two or more different primary histories, instead of just one. This introduces the possibility of some correlations being ignored and an underestimate of the uncertainty as discussed earlier.

In order to observe the effects of partitioning a phase-space source on uncertainty, we split the simulation of the 13.5 MV photon beam in water (described earlier) into 50

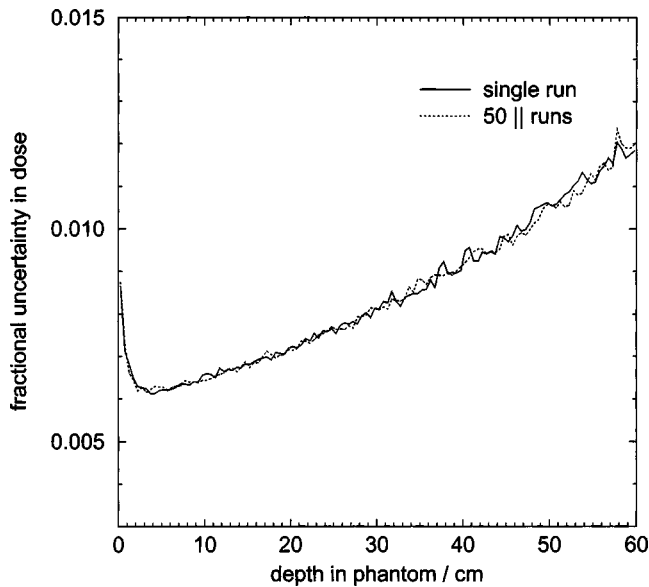


FIG. 5. Fractional uncertainty in central-axis dose vs depth for a simulated 13.5 MV photon beam from a Siemens KD2 accelerator (10×10 cm field at SSD = 100 cm) in a simulated water phantom ($20 \times 20 \times 60$ cm). Uncertainties are shown estimated using the new method when the simulation was performed in a single run (1.5 million incident particles) and when the simulation was divided into 50 parallel jobs (30 000 incident particles each). Dose was scored in $10 \times 10 \times 0.5$ cm voxels on the central axis.

parallel jobs. Each job simulated 30 000 histories ($1.5 \times 10^6/50$). Division into 50 jobs represents an extreme case, and increases the chances that correlated groups of photons in the phase-space source will be broken up by partitioning the file.

Figure 5 shows the fractional uncertainty in dose versus depth estimated after recombining the 50 parallel jobs, along with the uncertainty estimated when the simulation was performed in a single run (same as the solid line in Fig. 3). Figure 5 shows no significant difference between the uncertainties in the two cases. The mean uncertainty of all doses $> 0.5D_{\max}$ estimated after recombining the 50 runs is 0.697%, which is in good agreement with the mean uncertainty of 0.700% estimated in the single run. Thus, we conclude that, in general, partitioning of a phase-space source has no effect on the estimated uncertainty.

III. CONCLUSIONS

A new method for better estimates of uncertainty in the BEAMnrc and DOSXYZnrc codes has been introduced. Scored quantities are now grouped according to primary history and not statistical batches as before. This new method eliminates the problems of fluctuations in estimated uncertainty due to small sample size (i.e., small number of batches). It also eliminates one dimension from the arrays that score quantities of interest, resulting in a decrease in the memory required by the scoring arrays by a factor of about 2. This new method can easily be adapted to estimate the uncertainty on ratios of correlated quantities as well.

The new method also accounts for correlations between incident particles in a phase-space source by grouping inci-

dent particles according to the primary histories that generated them. This has been accomplished by changing the format of phase-space files slightly, so that the first particle scored from a new primary history is marked by setting its energy negative.

Recycling incident particles (i.e., using the same particle as an incident particle many times immediately after reading it), which is often necessary when phase-space data are sparse, introduces the correlations with the most significant effect on uncertainty. Failure to take into account these correlations can lead to a significant underestimate of the uncertainty, basically because it ignores the statistical fluctuations inherent in the phase-space file. On the other hand, when correlations are taken into account then, even if particles are recycled many times, the uncertainty will always reflect the uncertainty in the phase-space data itself. It was found that the batch method also successfully accounted for correlations introduced with recycling by placing all recurrences of a particle in the same statistical batch. Problems occur with both the new method and the batch method of estimating uncertainties if a phase-space source is restarted. In the case of the new method, a particle that gets re-read after the phase-space file has been restarted will not be grouped in the same primary history as it was on the previous pass. In the case of the batch method, there is no guarantee that the particle re-read after a restart will be put in the same batch as it was on the previous pass. Thus, restarting can cause uncertainty to be underestimated, and we recommend recycling particles enough times to avoid restarting a phase-space source (as well as reducing needless file re-reading).

Other correlations between particles in a phase-space source include those that occur “naturally” (e.g., a primary history that undergoes interactions leading to more than one particle in the phase-space source) and those that occur due to variance reduction techniques, such as bremsstrahlung splitting. These correlations are taken into account by the new method, but not by the batch method of estimating uncertainty. We examined the case of a photon beam generated by splitting bremsstrahlung photons up to 400 times (using selective bremsstrahlung splitting) in water and found that the batch method did, in fact, underestimate uncertainty. However, the large bremsstrahlung splitting number and the fact that significant differences were not noticeable until the voxels in the water phantom had been enlarged (to $10 \times 10 \times 0.5$ cm) to encompass most of the incident beam leads to the conclusion that, in most cases, these correlations do not play a role in the uncertainty estimate.

It should be noted that the event-by-event technique used by Ma *et al.*¹¹ and mentioned in Sec. I will not take into account correlations between incident particles in a phase-space source. This is because this technique does not trace energy deposited back to primary histories. It is even more extreme than the case discussed in Sec. II.D in which energy deposition was grouped according to particle incident on the phantom rather than by primary history. As shown there, the latter case is not likely to be a problem when correlations are due to interactions and/or variance reduction techniques. However, if correlations are introduced by recycling par-

ticles, then grouping by particles incident on the phantom is expected to significantly underestimate uncertainty since it ignores the latent variance of the phase-space file. The techniques used by Ma *et al.* will suffer this same problem.

One closing observation is that issues about the latent variance of a phase-space file are overcome, or at least highly mitigated, by using beam characterization models.¹² These models do not need to recycle particles and hence avoid the associated problems. Of course there is still some uncertainty to be associated with the models themselves, but as shown elsewhere, the use of the model greatly reduces these effects.¹²

The history by history method has also recently been added to the standard EGSnrc user-codes which are distributed by NRC.¹³

³Electronic mail: bwalters@irs.phy.nrc.ca

¹D. W. O. Rogers, B. A. Faddegon, G. X. Ding, C.-M. Ma, J. Wei, and T. R. Mackie, "BEAM: A Monte Carlo code to simulate radiotherapy treatment units," *Med. Phys.* **22**, 503–524 (1995).

²http://www.irs.inms.nrc.ca/inms/irs/BEAM/bibliog/omega_pubs.html.

³C.-M. Ma, D. W. O. Rogers, and B. Walters, "DOSXYZnrc Users Manual," NRC Report No. PIRS 509b(revF), 2001.

⁴D. W. O. Rogers, C.-M. Ma, G. X. Ding, B. Walters, D. Sheikh-Bagheri, and G. G. Zhang, "BEAMnrc Users Manual," NRC Report No. PIRS 509(a)revF, 2001.

⁵I. Kawrakow, "Accurate condensed history Monte Carlo simulation of electron transport. I. EGSnrc, the new EGS4 version," *Med. Phys.* **27**, 485–498 (2000).

⁶I. Kawrakow and D. W. O. Rogers, "The EGSnrc code system: Monte Carlo simulation of electron and photon transport," Technical Report No. PIRS-701, National Research Council of Canada, Ottawa, Canada, 2000.

⁷J. Sempau, S. J. Wilderman, and A. F. Bielajew, "DPM, a fast, accurate Monte Carlo code optimized for photon and electron radiotherapy treatment planning dose calculations," *Phys. Med. Biol.* **45**, 2263–2291 (2000).

⁸J. Sempau, A. Sánchez-Reyes, F. Salvat, H. Oulad ben Tahar, S. B. Jiang, and J. M. Fernández-Varea, "Monte Carlo simulation of electron beams from an accelerator head using PENELOPE," *Phys. Med. Biol.* **46**, 1163–1186 (2001).

⁹J. Briesmeister, "MCNP—A general purpose Monte Carlo code for neutron and photon transport, Version 3A," Los Alamos National Laboratory Report No. LA-7396-M, Los Alamos, NM, 1986.

¹⁰P. Andreo, "Monte Carlo techniques in medical radiation physics," *Phys. Med. Biol.* **36**, 861–920 (1991).

¹¹C.-M. Ma, J. S. Li, T. Pawlicki, S. B. Jiang, J. Deng, M. C. Lee, T. Koumrian, M. Luxton, and S. Brain, "A Monte Carlo dose calculation tool for radiotherapy treatment planning," *Phys. Med. Biol.* **47**, 1671–1690 (2002).

¹²C.-M. Ma, B. A. Faddegon, D. W. O. Rogers, and T. R. Mackie, "Accurate characterization of Monte Carlo calculated electron beams for radiotherapy," *Med. Phys.* **24**, 401–416 (1997).

¹³D. W. O. Rogers, I. Kawrakow, J. P. Seuntjens, and B. R. B. Walters, "NRC User Codes for EGSnrc," Technical Report No. PIRS-702, National Research Council of Canada, Ottawa, Canada, 2000.