Monte Carlo Tally Convergence: Runtime Comparisons Between Functional Expansion Tallies and Mesh Tallies

Brycen Wendt,[a] Leslie Kerby[a, b, c]

[a]*Department of Nuclear Engineering, Idaho State University, 921 South 8th Avenue, Pocatello, ID, 83209*
[b]*Center for Advanced Energy Studies (CAES), 995 University Boulevard, Idaho Falls, ID 83401 Reactor Physics*
[c]*Analysis and Design, Idaho National Laboratory, P.O. Box 1625, Idaho Falls, ID 83415*
wendbryc@isu.edu, kerblesl@isu.edu

INTRODUCTION

With the continuing trend in improving computing resources, high-fidelity multiphysics simulations are becoming the ideal approach for design and analysis in many fields; nuclear engineering is no exception. Although some codes are developed as an outright multiphysics simulation environment, many multiphysics codes are comprised of individually-developed applications that are coupled together through an engineered data transfer mechanism. Usually these data transfer mechanisms are specialized and applicable only for the originating application packages.

Recently, the use of functional expansions for coupling has been investigated, yielding a more generalized coupling mechanism [1, 2]. Other functional expansion implementations have yielded results such as accelerated reduced source sampling [3]. For data representation and potential coupling in Monte Carlo (MC) contexts, the function expansion analog is the sample-based functional expansion tally (FET). FETs have a number of desirable properties, including continuous data representation, fast convergence, high data density, and geometric agnosticity limited only by the domains of suitable functional basis sets. In particular, the convergence properties of FETs—with analysis and comparison to mesh tallies—have been rigorously studied previously. It was shown that, under most circumstances, FETs have a faster per-sample convergence rate than conventional tally methods [4]. These characteristics have recently garnered increased interest in FETs.

However, one frequently-voiced concern of using FETs arises from the expensive computational cost caused by evaluating the entire functional basis set for each score to the tally. Do the computational expenses of FETs outweigh the benefits of a reduced sample set? The purpose of this work is to study and quantify the time-to-convergence of FETs. Recently optimized FETs in the MC reactor physics code Serpent will be used [5, 6]. Specifically, it will be shown that FETs have a significant advantage over mesh tallies using computational runtime vs. convergence as a metric. These results will also inherently reconfirm prior sample size convergence studies.

THEORY

Separability was not assumed for this work, meaning fully multivariate functions were used as the basis sets. Often the separability of an FET's dimensional components is imposed, likely due to simplicity of implementation and lower computational cost [2, 4]. However, it is important to show that fully-developed FETs are still computationally advantageous. With an inherent higher-fidelity than univariate (separable) FETs, multivariate (nonseparable) FETs represent the ideal.

Foundational Definitions

FET approaches can be used to approximate a distribution F by a set of expansion coefficients $[\, b \mid b_0, b_1, \ldots, b_N \,]$ corresponding to the functional basis terms $[\, \psi \mid \psi_0, \psi_1, \ldots, \psi_N \,]$.

In Serpent, the collision tally based estimator $\widehat{b}_n$ (the stochastic approximation of b_n) is calculated using:

$$\widehat{b}_n = \frac{c_n}{P} \sum_{p=1}^{P} \sum_{k=1}^{K_p} s_{p,k} \beta_{p,k} \psi_n(\xi_{p,k}) \omega_n(\xi_{p,k}) = \frac{c_n}{P} \sum_{p=1}^{P} G_p \tag{1}$$

$$\frac{1}{c_n} = \|\psi_n\|^2 = \int_\Gamma \psi_n^2 \omega_n \tag{2}$$

where c_n is the orthonormalization constant, P is the number of particle histories, K_p is the number of collisions per particle p, $s_{p,k}$ is the value being tallied, $\beta_{p,k}$ is the stochastic weight, ω_n is the weight function for ψ_n, Γ is the orthogonal domain of ψ, and $\xi_{p,k}$ is the event location mapped from the domain of F to Γ. The condensed term G_p is used in the variance estimator:

$$\widehat{\sigma}^2_{\widehat{b}_n} = \frac{\sum_{p=1}^{P} \left(G_p\right)^2 - \frac{1}{P}\left(\sum_{p=1}^{P} G_p\right)^2}{P\,(P-1)} \tag{3}$$

This approach deviates slightly from the common implementation in literature: the orthonormalization constant c_n is present in Eq. (1) for coefficient estimation rather than the reconstruction of F:

$$\widehat{F}(\xi) = \sum_{n=0}^{N} \widehat{b}_n \psi_n(\xi) \omega_n(\xi) \tag{4}$$

Thus, a notation based on $\widehat{b}_n$ instead of $\widehat{a}_n$ is used herein to avoid confusion. This approach was adopted because the FET coefficient values can be more intuitively understood with respect to the functional basis set.

Convergence

Like all MC-based tallying methods, FETs are subject to both truncation and statistical errors. Another important realization is that some terms may be orthogonal—or nearly orthogonal—to the sampled distribution F. Accordingly, the variance estimators for these terms can be statistically large even though the coefficient estimators may be small or zero. The challenge is to accurately represent the total uncertainty without an unreasonable bias toward less-meaningful terms.

A cost-to-benefit ratio has been proposed to evaluate the worth of each coefficient (note the use of $\widehat{a}_n$ instead of $\widehat{b}_n$) [4]:

$$R_n = \frac{\widehat{\sigma}^2_{\widehat{a}_n} c_n}{\widehat{a}_n^2} \tag{5}$$

The accompanying recommendation was to accept coefficients with $R_n \ll 1$ and reject those with $R_n \gg 1$. Unfortunately, the fate of coefficients with $R_n \approx 1$ is unsettled. One solution is to set an arbitrary limit, e.g. $R_n = 2.0$, as a cutoff. Without detailed insight into the sampled distribution F itself, though, this cutoff approach can be unsatisfactory as it does not optimize the balance between information and uncertainty contributions. Prior searches for an optimal cutoff are situation-specific [2, 4], implying that prescribing a generalized cutoff value is highly speculative. Other possible heuristics are still being investigated that may provide more generalized filtering recommendations than a manually-fixed cutoff value.

The need for an unbiased uncertainty measure—a quantitative metric of the data quality vs. uncertainty significance—induced development of a new approach to managing the statistical uncertainty analysis of an FET. First, it was decided that all evaluated values should be included in the analysis. Next, an equivalent measure must exist for mesh tallies. Finally, it must be evaluated in ℓ^2 space to capture the effects of non-zeroth (higher order) terms in orthogonal functional bases.

The proposed metric is a total relative variance U, which measures the uncertainty contribution $\|\sigma_i\|^2$ of each term with respect to the aggregated information $\sum\|v_i\|^2$ provided by the values themselves:

$$U_i = \frac{\|\sigma_i\|^2}{\sum\|v_i\|^2} \tag{6}$$

$$U = \sum U_i = \frac{\sum\|\sigma_i\|^2}{\sum\|v_i\|^2} \tag{7}$$

This metric is not a "true" uncertainty, but a sensible measure found to be very useful for this time-to-convergence analysis.

U_{FET} Derivation The uncertainty and value contributions, with the resulting total relative variance, are:

$$\|\sigma_i\|^2 = \|\widehat{\sigma}_{\widehat{b}_n}\|^2 = \int_\Gamma \left(\widehat{\sigma}_{\widehat{b}_n}\psi_n\right)^2 \omega_n = \frac{\widehat{\sigma}^2_{\widehat{b}_n}}{c_n} \tag{8}$$

$$\|v_i\|^2 = \|\widehat{b}_n\|^2 = \int_\Gamma \left(\widehat{b}_n\psi_n\right)^2 \omega_n = \frac{\widehat{b}_n^2}{c_n} \tag{9}$$

$$U_{\text{FET}} = \frac{\sum \frac{\widehat{\sigma}^2_{\widehat{b}_n}}{c_n}}{\sum \frac{\widehat{b}_n^2}{c_n}} \tag{10}$$

Note that the bias of $\sigma \approx \sqrt{\sigma^2}$, a variance to standard deviation relation used in Eq. (8), is less than 0.03 % when the sample size is greater than 1000 (for most MC-derived distributions).

U_{mesh} Derivation The equivalent mesh tally form is:

$$\|\sigma_i\|^2 = \int_b \sigma_b^2 dV = \sigma_b^2 V_b \tag{11}$$

$$\|v_i\|^2 = \int_b v_b^2 dV = v_b^2 V_b \tag{12}$$

$$U_{\text{mesh}} = \frac{\sum \sigma_b^2 V_b}{\sum v_b^2 V_b} \xrightarrow{\text{uniform } V_b} U_{\text{mesh}} = \frac{\sum \sigma_b^2}{\sum v_b^2} \tag{13}$$

for bin b and volume V.

Figure of Merit (FOM) A FOM is appropriate for this context, in which the effectiveness of an approach is measured against the runtime. Using the total relative variance U as the uncertainty measure, the time-to-convergence FOM is:

$$\text{FOM} = \frac{1}{UT} \tag{14}$$

where T is the trial's runtime.

METHODOLOGY

Benchmarks PU-COMP-MIXED-001 case 5 (unreflected) and PU-COMP-MIXED-002 case 23 (reflected) were selected from the International Criticality Safety Benchmark Evaluation Project (ICSBEP) Handbook as the base models for testing. These were selected due to: 1) plain models, 2) reflected and unreflected configurations, and 3) common dimensions.

Testing

The benchmark trials were run on a Beowulf cluster, constructed of nodes containing eight CPU cores clocked at 2.0 GHz and 16 GB RAM. Each trial was run on a separate node from the others to isolated each trial's performance and ensure the highest level of computational environment similarity. Serpent was run with 8 OpenMP threads to fully utilize each node's CPU resources, and its self-reported transport cycle duration was collected as the runtime metric.

All tallies were generated using Serpent's detector functionality. Each trial contained 1100 generations with the first 100 discarded; the Shannon entropy initially verified to ensure source term convergence before the active cycles were begun. The FETs were based on a convolution of three Legendre polynomial series (one for each dimension) to create a fully multivariate 3D functional basis. The mesh tallies were uniformly divided so that the bin volumes were equal. Well-resolved high-resolution "truth" mesh tallies were also generated to provide a sensible and accurate approximation of the underlying sampled distribution F. These truth tallies were utilized to provide an additional comparative measurement of the trial tallies' overall statistical and truncation error.

The truth tallies were recorded with 1×10^{10} total particles during the active cycle. Coarse and fine trial tally granularities were defined as shown in Table I; the truth tally specifications are also provided. Since the objective was to research the time-based convergence of each tally type, short, medium, and long trial runtimes were performed with respective total particle sources of 1×10^6, 1×10^7, and 1×10^8.

Comparisons and Analysis

All tallies were normalized to a count density of 1 over the detector volume. For the FETs this required a transformation from the functional domain to the physical model. The total relative uncertainties were calculated using Eqs. (10) and (13).

The trial tally results were then compared to the truth tally to individual accuracies. A full 3D comparison was not performed; instead, 2D data slices were evaluated using perpendicular planes passing though 29.5 %, 45.5 %, 79.5 %, and 95.5 % of each dimension's length, resulting in twelve slices

TABLE I. The Granularities Used, and the Total Number of Values Tracked, by Each Tally Type.

Type	Specification	Total Values
FET	Order	
Coarse	$4 \times 4 \times 4$	125
Fine	$8 \times 8 \times 8$	729
Mesh Tally	Bins	
Coarse	$8 \times 8 \times 8$	512
Fine	$32 \times 32 \times 32$	32 768
Truth	$100 \times 100 \times 100$	1 000 000

Note: the Cartesian domain correlation of the tally specification is: $x \times y \times z$

per trial mesh comparison. These positions were chosen from the nearest common bin centroids shared among the three mesh tally granularities, enabling the most accurate comparison between the mesh tallies. Each FET's error was found by numerically evaluating the expansion using Eq. (4) over a fine quadrature, calculating average values over truth tally bin ranges, then accumulating the differences with the corresponding truth tally bin values. Each mesh tally's errors were evaluated at a resolution constructed as the union of the bin-edge sets from itself and the truth tally. The bin values were projected onto this common mesh and the errors accumulated as the volume-weighted difference with the collocated truth bin. The standard deviation of these accumulations represented the degree to which a trial tally deviated from the truth; therefore, it was used as the measure of overall relative error.

RESULTS

The eigenvalues reported by Serpent were in agreement with the benchmarks, indicating that the models were correctly defined in the input files. The analysis results of PU-COMP-MIXED-002 Case 23 are shown in Table II; a complete dataset is found in a publicly-viewable Google Sheet. Also, Fig. 1 compares two fine tally trials—the FET medium-runtime and mesh tally long-runtime—to the truth tally. The full analysis for both benchmarks plotted in Figs. 2 and 3; the optimal performance region is in the lower-left corner of each graph.

First, and foundationally requisite to these results and analyses, Fig. 2 and Table II exhibit a near-perfect $U \propto \frac{1}{t}$ relation. This indicates that U is a sensible measure for evaluating the statistical convergence of both FETs and mesh tallies. This also means that U is an appropriate uncertainty measure for use in calculating the FOM according to Eq. (14). The final FOM for each tally type, shown in Table III, was calculated as the average of the individual FOMs from each trial duration. The FET FOMs are consistently higher than the mesh FOMs, somtimes by more than a factor of ten.

Second, Fig. 2 reconfirms that the time-based FET performance is consistently better than the mesh tallies as indicated by the FOMs. This is apparent when comparing the short runtime FETs to the long runtime mesh tallies: the short FETs require only minutes of runtime to surpass the quality of data provided by the mesh tallies running for over two hours. Figure 1 is an especially poignant demonstration, in which the fine FET tally shows remarkable similarity to the truth tally, but requires only a fraction of the less-precise fine mesh tally's

TABLE II. Raw Data for the Reflected Benchmarks; Plots are Presented on the Right Sides of Figs. 2 and 3.

Tally	Time [h]	U	Error [%]
Truth	285.4	—	—
Coarse FET	0.041 94	6.23×10^{-5}	2.13
	0.3842	6.21×10^{-6}	1.65
	3.812	6.20×10^{-7}	1.50
Fine FET	0.089 17	2.37×10^{-4}	2.39
	0.8647	2.37×10^{-5}	1.02
	8.589	2.36×10^{-6}	0.58
Coarse Mesh	0.030 83	2.50×10^{-4}	7.52
	0.2844	2.51×10^{-5}	7.25
	2.749	2.48×10^{-6}	7.19
Fine Mesh	0.038 88	6.47×10^{-3}	9.21
	0.2917	6.49×10^{-4}	3.41
	2.771	6.48×10^{-5}	2.04

TABLE III. Final FOMs of the Tally Types for each Benchmark Model, Organized by Granularity

Granularity	Type	Unreflected	Reflected
Coarse	FET	3.57×10^{6}	9.80×10^{6}
	Mesh	6.70×10^{5}	3.33×10^{6}
Fine	FET	6.92×10^{5}	1.16×10^{6}
	Mesh	2.99×10^{4}	1.19×10^{5}

longer runtime. This is a strong confirmation of the hypothesis: despite a heavier per-sample computational footprint, FETs have a faster time-to-convergence rate than mesh tallies.

Third, inspection of Fig. 3 reveals that the coarser tallies have a lower U but a higher overall relative error due to truncation effects. Interestingly enough, Fig. 2 shows that, for the unreflected case, the coarse mesh tally and fine FET have exactly the same convergence rate; however, when checking the overall relative errors in Fig. 3 it is apparent that the FET provides a much better representation. The fine FET is tracking 729 coefficient values, compared to the coarse mesh tally's 512 bins. This places the FET at a statistical disadvantage, making the outcome in favor of the FET even more impressive.

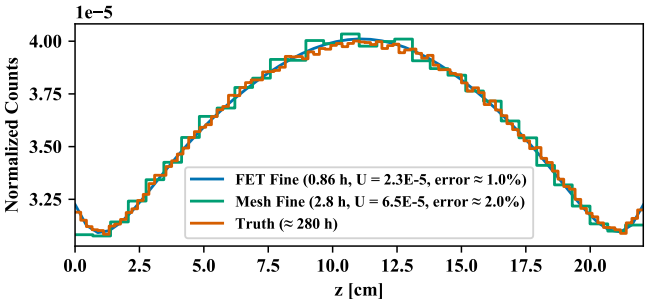

Fig. 1. Comparison of fine tally methods for the reflected model, with the 3D Data Sliced at 45.5% of the X and Y Axes to Produce a 1D Plot in Z.

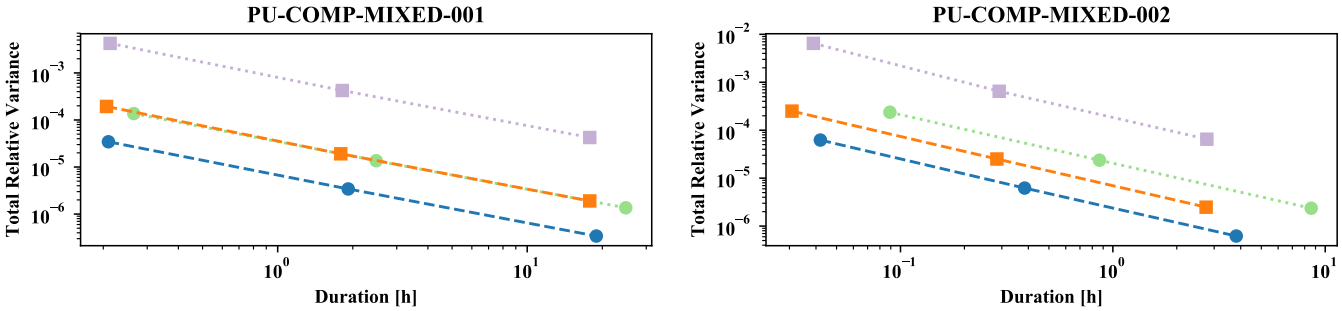

Fig. 2. Comparison of the total relative variances as a function of trial runtime (both axes are logarithmically scaled).

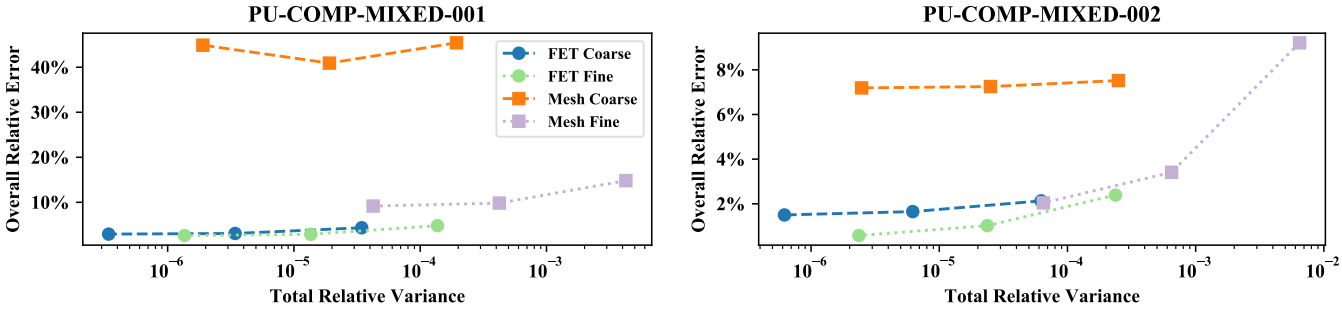

Fig. 3. Relationship between the total relative variances and the overall relative errors for the different trials (the horizontal axis is logarithmically scaled; the legend is common for Figs. 2 and 3).

Fourth, the reflected trials have better overall results. This is primarily because the reflection causes the neutron profile to be less varying throughout the detector volumes. Comparing the left and right sides of Fig. 3 demonstrates how both tally methods benefit from this effect.

Finally, the trial tallies appear to asymptotically approach a minimum error-to-truth value for longer computational times, especially for the unreflected benchmark (PU-COMP-MIXED-001). It is inferred via Fig. 3 that all the mesh tallies—with the exception of the reflected fine trial—are unlikely to reach the minimal error values of the FETs; finer resolutions (needing a significantly longer runtime) would be required.

CONCLUSIONS

FETs have been verified be computationally advantageous due to superior convergence qualities, despite an algorithmically heavier tallying methodology. They showed a performance speedup of several times over a fine mesh—and hundreds over a truth tally (see Table II)—while consistently providing a higher-quality representation than a coarse mesh. Low-order FETs are suitable to capture high-fidelity information of simple distributions with very short runtimes. Higher-order FETs are required when capturing more complex distributions, yet can still provide suitable results with medium-length runtimes. In high-performance computing contexts, gains of 10 % to 20 % are significant; the performance advantages of FETs are monumental: often several times faster (or more) than mesh tallies providing a comparable quality. Consequently, FETs have an impressive case for consideration as the primary MC tally mechanism. Finally, these results firmly support the role of FET methods in MC-coupled multiphysics simulations seeking reduced computational time and statistical error while insuring or improving fidelity.

REFERENCES

1. L. KERBY, A. TUMULAK, J. LEPPÄNEN, and V. VALTAVIRTA, "Preliminary Serpent–MOOSE Coupling and Implementation of Functional Expansion Tallies in Serpent," in "International Conference on Mathematics & Computational Methods Applied to Nuclear Science and Engineering (M&C 2017)," (2017).
2. M. ELLIS, *Methods for Including Multiphysics Feedback in Monte Carlo Reactor Physics Calculations*, Ph.D. thesis, Massachusetts Institute of Technology (2017).
3. J. FAVORITE and H. LICHTENSTEIN, "Exponential Monte Carlo Convergence of a Three-Dimensional Discrete Ordinates Solution," *Transactions of the American Nuclear Society*, **81**, 147–148 (1999).
4. D. GRIESHEIMER, *Functional Expansion Tallies for Monte Carlo Simulations*, Ph.D. thesis, University of Michigan (2005).
5. B. WENDT, L. KERBY, A. TUMULAK, and J. LEPPÄNEN, "Advancement of Functional Expansion Capabilities: Implementation and Optimization in Serpent 2," *Nuclear Engineering and Design* (2018), submitted.
6. J. LEPPÄNEN, M. PUSA, T. VIITANEN, V. VALTAVIRTA, and T. KALTIAISENAHO, "The Serpent Monte Carlo code: Status, development and applications in 2013," *Annals of Nuclear Energy*, **82**, 142–150 (2015).