Estimation Procedures and Error Analysis for Inferring the Total Plutonium Produced by a Graphite-Moderated Reactor

Patrick G. Heasler[a], Tom Burr [c], Bruce Reid[a], Christopher Gesh[a], Charles Bayne[b]

[a]Pacific Northwest National Laboratory, Richland, WA, US;
[b]Oak Ridge National Laboratory, Oak Ridge, TN, US;
[c]Los Alamos National Laboratory, Los Alamos, NM, US
E-mail: patrick.heasler@pnl.gov, bayneck@ornl.gov, tburr@lanl.gov

ABSTRACT

GIRM (Graphite Isotope Ratio Method) is a technique that can determine total Plutonium (Pu) production in a graphite moderated reactor. In the GIRM methodology, total Pu production is estimated by measuring isotopic ratios of trace elements in extracted graphite samples from the target reactor.

Many safeguards problems require an estimate of total Pu production. For example, a declaration of total Pu might need to be verified through GIRM. In some cases, reactor information (such as core dimensions, coolant details, and operating history) are so well documented that reactor computer-models can predict total Pu production. In most cases however, reactor information is imperfectly known or of questionable validity, so a measurement-based method such as GIRM is essential to such a verification strategy.

In this article we concentrate on GIRM's estimation procedure and its associated uncertainty. We describe a simulation strategy to estimate its uncertainty, including the impact of local and global computer codes, and illustrate GIRM for a specific reactor.

Keywords: reactor code errors, simulation, uncertainty in estimated Plutonium

1. INTRODUCTION

In GIRM, samples from the graphite moderator are taken along the fuel channels in the reactor of interest, and isotopic ratios in the samples are measured with mass spectrography (TIMS for Uranium or Pu isotope ratios and SIMS for Boron isotope ratios). These isotopic ratios are converted to local plutonium production rates (i.e. grams of Pu produced per cm of fuel rod) using a lattice physics code (WIMS). Finally, a 3D regression model is used to estimate a Pu fluence field for the reactor which is then integrated over the fuel channels to estimate the total Pu produced. The 3D regression model used in a specific application is determined from a reactor physics code (such as KENO).

The basic scheme has been subjected to several feasibility studies and experimental tests (see [1], [2], and [3]). Two previous error analyses of this methodology have been conducted. The earliest analysis, Reference [1] evaluated uncertainties associated with a "generic" reactor. The other, Reference [3] evaluated actual measurements taken from a British reactor, Trawsfynydd. Because the GIRM methodology has changed substantially in recent years [4] and additional sources of error have been included (i.e. WIMS reactor

code error and 3D fluence model errors), it is timely to re-evaluate the total uncertainty with the GIRM-based estimate of total Pu production. The procedures described in the next section are illustrated by applying them to a "generic" reactor, which is described in [4].

2. METHODOLOGY

GIRM can be organized into three steps: (1) a planning step that assembles and evaluates information about the target reactor; (2) a sampling step that extracts graphite samples from the target reactor, and (3) an analysis/estimation step that measures the isotopic ratios and converts them to a Pu estimate.

In the planning step, information concerning a target reactor is gathered. This includes reactor core dimensions, fuel channel locations, control rod locations, coolant details, fuel rod specifications, and operating history. WIMS produces curves that relate the isotopic ratios to local Pu fluence. WIMS runs also generate a parametric evaluation of the main reactor variables such as temperature, density, and dimensions. A sampling plan is designed that describes where graphite samples will be taken from the target reactor.

The analysis step includes mass spectrographic measurements of the isotopic ratios of each sample, and statistical analysis to combine all the measurements to estimate total Pu production. In the first step in this statistical analysis, the measured isotopic ratios are converted to local Pu fluence (g/cm) estimates using curves calculated by WIMS.

The second step in statistical analysis fits a 3D linear regression model to the local Pu fluence estimates produced by the first step. The 3D regression model (called the global regression model), is produced from 3D reactor-physics solutions. After the global regression model has been fit to the local fluence estimates, it is easy to determine the total Pu production by integration. Standard propagation-of-error (POE) formulas also produce the uncertainty associated with this estimate in the form of a standard error. Uncertainties of the entire estimation procedure are verified using Monte-Carlo simulation. Because of deficiencies in the POE methods, we presently rely on simulation to produce the most comprehensive estimate of total uncertainty.

2.1. Pu-fluence from U/Pu Data and B data

At a given sample location, the TIMS mass spectrograph produces a vector of 5 isotopic ratios ($^{236}U/^{238}U$, $^{236}U/^{235}U$, $^{240}Pu/^{239}Pu$, $^{241}Pu/^{239}Pu$, $^{242}Pu/^{239}Pu$), while SIMS produces a single isotopic ratio for $^{10}B/^{11}B$. All measured isotopic ratios have an associated uncertainty, supplied by the chemical analysis. It should be noted that Boron is particularly suited to a low burn-up reactor. For a high burn-up reactor, the Boron measurements would be replaced by Titanium.

These isotopic ratios are used to estimate the local Pu-fluence, ϕ_{pu} at each sample location, and produce $se(\hat{\phi}_{pu})$. To accomplish this, a nonlinear regression model is used to relate Pu-fluence ϕ_i to the measured isotopic ratios, R_{ij}. The model is;

$$R_{ij} = H_j(\phi_i) + e_{ij} \tag{1}$$

where the index i identifies the sample location, and the index j the specific isotopic ratio measured. The function $H_j(\phi)$ represents the isotopic ratio curves produced by WIMS runs for the target reactor and the error term, e_{ij}, represents error. The least-squares estimator for the vector of ϕ_i is the value that minimizes the quadratic form[*];

$$SSE_i = [R_i - H(\phi_i)]^T Cov(e_i)^{-1}[R_i - H(\phi_i)] \tag{2}$$

while its standard error is given by an asymptotic approximation used for non-linear regression (i.e., the inverse of Fisher information matrix).

The covariance of the regression uncertainties includes measurement error, and reactor physics error (RPE). In other words, the covariance matrix used by the regression is $Cov(e) = Cov(\text{Meas. Error}) + Cov(\text{Reactor Phy. Error})$. Including RPE in the regression weights improves the total Pu estimates by 50% because it is the dominant source of uncertainty in the data. RPE describes the effect that uncertainties in important reactor parameters (such as graphite temperature) have on this regression problem.

We note here that the se's produced by this algorithm do not account for calibration bias. Although the se of the result does include RPE, correlation between local estimates is not produced, which also can have an important effect on total Pu uncertainty and is the chief reason that Monte-Carlo generated uncertainties are better than POE uncertainties.

2.2. Global Regression and Estimate of Total Pu

Linear regression is used to fit a 3D fluence field to the local fluence estimates produced by the previous local estimation step. Let $\Phi(x; \beta)$ represent the fluence field model, with x representing a location in the reactor core, and β a set of unknown parameters that determines the shape of the fluence field. Linear regression is used to determine parameter estimates that produce a fluence field that is as close as possible to the local fluence values.

We call this "global" regression because it transforms the local Pu estimates into a fluence field that is defined at any point within the entire reactor core. The global regression model is described by; $\hat{\phi}_i = \Phi(x_i; \beta) + e_i$ with weight W_i associated with location x_i. The weight is determined by the se assigned to $\hat{\phi}_{pu}(x)$ from the local Pu estimation step, $W_i = se(\hat{\phi}_{pu}(x))^{-2}$. Weighted regression is then performed to produce estimates, $\hat{\beta}$, of the unknown parameter vector β, and its covariance matrix, $Cov(\hat{\beta})$. The form used for the 3D fluence field model is

$$\Phi(x; \beta) = \sum_{k=1}^{p} \beta_k \Psi_k(x). \tag{3}$$

The $\Psi_k(x)$ must be known functions, that are relatively good at approximating the target reactor fluence field. An estimate of the total Plutonium production, $\hat{T}_{pu}$, is then determined by integrating the estimated fluence field over all fuel channels in the reactor. The integration result is a linear combination of the unknown beta parameters; consequently, it is simple to calculate the standard error of the estimate, $\hat{T}_{pu}$, from $Cov(\hat{\beta})$.

---

[*]$R_i = (R_{11}, R_{12}, ...R_{1n})$ and similarly for $H(\phi)$, and e_i.

Let $(x_{1\ell}, x_{2\ell})$ represent the horizontal coordinates of the ℓ'th fuel channel, so that $\hat{T}_{pu}$ is given by;

$$\hat{T}_{pu} = \sum_\ell \int_{Z_{bottom}}^{Z_{top}} \Phi(x_{1,\ell}, x_{2\ell}, z_3; \hat{\beta}) dz_3 = \sum_{k=1}^{p} \hat{\beta}_k \sum_\ell \int_{Z_{bottom}}^{Z_{top}} \Psi_k(x_{1,\ell}, x_{2\ell}, z_3) dz_3 = \sum_{k=1}^{p} \hat{\beta}_k \omega_k \tag{4}$$

with the integration constants ω_k defined by $\omega_k = \sum_\ell \int_{Z_{bottom}}^{Z_{top}} \Psi_k(x_{1,\ell}, x_{2,\ell}, z_3) dz_3$.

An estimate for the total Pu can therefore be found by multiplying the beta estimates by the integration constants ω_k. In vector notation this reduces to; $\hat{T}_{pu} = \omega^T \hat{\beta}$ and the estimate of total Pu has $se(\hat{T}_{pu})^2 = \omega^T Cov(\hat{\beta}) \omega$. These formulas produce an estimate of T_{pu}, and its se, which is the final objective of this statistical procedure. Incidentally, if the regression model $\Phi(x, \beta)$ can fit the true fluence field without any error, then the estimate is unbiased and the true total fluence, T_{pu} can be expressed as a linear combination of the ω_k and true β_k that resembles Eq. (4).

The adequacy of the global regression model depends upon the proper selection of the set of "basis functions", $\Psi_k(x)$. 3D reactor physics models are used to produce a "basis set" of reactor fluence profiles, based on reactor operating history. If one is fairly certain about the operating history, this set will be small, if one is less certain, the set will contain more profiles and uncertainty related to global regression will be correspondingly greater.

In the generic example presented in this article, the KENO code has been used to produce a "best estimate" for fluence, which is included in the basis set, along with 44 "eigen-function" profiles from an homogeneous-core solution. The homogeneous-core "eigen-function" profiles are the functions we would use when little reliable information is available for the target reactor. Thus the example global regression model is constructed from generic information and one "best estimate" profile.

2.3. Monte Carlo Simulation for Error Analysis

The steps for our simulation are: (1) Calculate a "true state" for the target reactor (the global Pu-fluence, $\Phi(x)$ at locations x_i) (see the next paragraph); (2) Add error to locations x_i that arises due to limitations of the sampler; (3) Sample the reactor parameters (such as coolant and fuel temperature, etc.) from distributions centered on true values with standard deviations estimated using linear approximations to results of auxiliary WIMS runs. The "true" isotopic/fluence curves are generated WIMS using these sampled reactor parameters; (4) Add error to the true isotopic ratios to create measured values at samples (with sample contamination, calibration, and random errors), (5) Run the simulated measurements through the Pu estimation algorithm, and finally (6) compare the estimated result to the "true" Pu value, and repeat many times.

The "true" global fluence is simulated by random selection from a set of "representative" 3D fluence fields. This set is related to, but not the same as the "basis" set used to construct the global regression model. The set of 3D fluence fields contains plausible (best-estimate) and extreme fluence shapes, calculated from what is known about the target reactor's operating history. For the generic reactor used here, this set consists of seven fields calculated by KENO. Of these seven fields, three represent best-estimates

versions of reactor fluence, while the other four represent extreme fields. The extreme fields are calculated with unusual control rod configurations which are possible, but it is unlikely that the reactor would be operated in this configuration for any length of time.

The local RPE is the result of errors in certain reactor parameters required to complete a WIMs calculation. The WIMs code result can be mathematically described by the function $R = H(\phi; \alpha)$, where R is a vector whose components represent the ratios of all of the measured isotopes in graphite, while ϕ represents the corresponding Pu-fluence. $H(\cdot)$ is also a function of reactor parameters, as represented by the vector α. There is some error associated with the best-estimate value for α used in the estimation procedure, and to account for this, α is considered a random vector, with a mean centered around the "best estimate" and standard deviation representative of the uncertainty associated with these parameters. In the Monte-Carlo, a set of "true" fluence curves is simulated by sampling α from this distribution and computing $H(\phi, \alpha_{sampled})$.

For the example reactor, the reactor parameter vector α represents the following; fuel pin radius, fuel temperature, graphite density, graphite temperature, graphite equivalent boron concentration, and WIMS code uncertainty. "WIMS code uncertainty" represents numerical errors produced in WIMS, estimated by comparing WIMS results to results from another code. It is supposed to represent the difference between the answer produced by WIMS in GIRM, and a computer code that could solve the problem without error.

As a final step, the Monte Carlo simulates SIMS and TIMS measurement errors. These consist of random, calibration, and contamination errors, which are added to the true isotopic ratios to produce a measured value. These operations are summarized by the formula $R_{meas} = (1-C)R_{irr} + CR_{nat} + e_{cal} + e_{ran}$ where, for example, R_{irr} is the true (irradiated) ratio of Boron in the sample, which is contaminated by $C\%$ of natural Boron (that has ratio R_{nat}). A calibration error of e_{cal} and a random error of e_{ran} is then added to this result to obtain the measured ratio, R_{meas}.

It is important to note that the simulated measurements are more complicated than those assumed for the regression model (Equation 1) used in the estimation step. The contamination can cause bias, while the calibration errors can cause correlations, and neither of these effects are accounted for in the regression model.

2.4. Sampling as an Optimization Problem

A sampling plan consists of a set of specified locations in a reactor. These locations are represented by $x_1, x_2, ...x_n$, with each x_i a 3D vector representing the coordinates of a sampling location. The Monte Carlo error analysis methodology can be used to select the best sampling plan from a limited set of candidate plans. However, Monte Carlo evaluation is unsuited (too slow) to be used as part of a sample optimization scheme.

To find the optimal sample design, we have utilized the global regression model described in the estimation step. This global regression model can be quickly evaluated to produce an approximate RMSE for a particular to a sample design. These results are used by an optimization algorithm to find the sample design with a small root mean square error. The optimization is typically done with constraints; Samples can only be taken in certain allowed channels, and the sample locations within a channel are fixed.

Our optimization evaluations for the example reactor resulted in a design with 100 sample locations, taken in 10 channels. Roughly, the design consists of three "half-planes" in three radial directions. An experimental design tool (ECHIP, available at echip.com) was also employed to evaluate the designs and identify "high-bias" sample locations, which were eliminated. This optimized design was used in the example evaluations presented in this article.

3. EXAMPLE RESULTS FOR A LOW-BURNUP REACTOR

This example provides an upper-bound case for uncertainties when GIRM is applied to a low-burnup reactor. This means the uncertainties in RPE, contamination, instrument error, etc. are large, but not unreasonably large.

3.1. Estimation Uncertainty for the Example

To determine Pu estimation error for the example, 1400 sets of measurements were simulated, resulting in 1400 "total Pu estimates," which were compared to the "true Pu" (250 Kg for the generic case). Table 1 illustrates a few of these results.

Table 1. Simulations for Generic Case

| | Estimation Results | | | | Est Error |
| | (Kg) | (Kg) | (Kg) | | (Kg) |
Run	$\hat{T}_{Pu}$	True T_{Pu}	$se(\hat{T}_{Pu})$	GOF	$\hat{T}_{Pu} - T_{Pu}$
1	253.80	250.00	1.22	4.46	3.80
2	247.66	250.00	1.40	3.88	-2.33
3	252.76	250.00	1.00	2.54	2.76
.	.	.	.	.	.
1400	252.08	250.00	1.59	3.27	2.08

$RMSE = 4.04 Kg$ %RMSE= 1.62

Mean $se(\hat{T}_{pu})$=1.11Kg   (Prop. Error SE)

In Table 1, each row summarizes one simulation. The total Pu estimate is in column 2 and the POE-based estimate of $(se(\hat{T}_{Pu}))$ is in column 3. Column 4 is a "Goodness of Fit" statistic computed by the fluence regression. When the data fits the statistical models presented in the previous sections, the GOF statistic should be approximately 1. Theoretically a value above 1.5 indicates significant lack of fit. Note that GOF statistic is typically higher than it should be. This is because the simulated data contains non-independent errors (i.e. contamination, RPE, measurement calibration errors) that are being detected by this statistic. The GOF statistic measures how severely the data deviates from the assumed model, which assumes independent, zero-mean errors. We expect the GOF statistic to be larger than its theoretical value for real data, because real data will have calibration and contamination problems to some degree. At the bottom of the table are the average RMSE and %RMSE over the 1400 simulations. As one can see from

the example, the POE-based se is too small; actual errors are more than 3 times larger in this example (average RMSE of 404 compared to average POE-based se of 111).

Table 2 decomposes the generic scenario RMSE into the three main components. Bias describes a general bias that may be due to any combination of global regression model misfit, measurement contamination, calibration bias, or RPE. Although significant, this bias is relatively small. Because bias is influenced by many inputs, it can be altered significantly by changing the above-mentioned inputs.

Table 2. Decomposition of RMSE for Generic Scenario

Source of Variation	Absolute (g)	Relative % of T_{pu}
Bias: Due to various sources.	74	0.30
Between Model SE: Global Model Uncertainty	111	0.44
Within Model SE: All Sources Except Global Modal Uncertainty	381	1.52
Decomposition of Within SE:		
–Random Measurement Error:	50	0.20
–Measurement Calibration Error:	30	0.12
–Contamination Error:	25	0.10
–Location Error	0	0
–Reactor Physics Error	375	1.50
Total RMSE*	404	1.62

$$* \ RMSE = \sqrt{\text{Bias}^2 + \text{SE}^2}$$

The between model error describes the bias in the global regression model when applied to the population of seven global fluence models used as truth. This error is the second largest component of error, but still relatively small. The major difference in these results is not in the se's (which are all approximately 1.50%), but in the bias. Biases range from -0.18 to 0.81 for the seven fluence fields. These seven biases correspond to a between model se of 0.44%, which is about half of the maximum, and describes the error in the population of "true state" fields.

The third and largest source of error, "within model" error, comprises all other sources of error in the data (contamination, measurement error, calibration error and local RPE). By switching specific sources of error off and re-running the Monte Carlo program, one can decompose this within model error further, into its major sources. The results of the Monte Carlo runs are also listed in Table 2. The decomposition of within model error shows that local RPE is the largest contributor to Total Pu error, producing an error of 1.50%. The next largest contributor is SIMS and TIMS random measurement error, producing a contribution of 0.20%. SIMs and TIMs calibration error produce a 0.12% error, which roughly corresponds to the magnitude of calibration error with respect to random error. The assumed contamination for SIMS and TIMS measurements produces a 0.10% error, which is the smallest of the measurement error effects. Finally, sampler location error is less than 0.01% and we have therefore set it to zero. Even a unrealistically

large location error of 2cm only increases the se in total Pu to 0.03%, so it is safe to eliminate location error as a serious source of estimation error.

RPE is the largest error source (1.50%), and the contributions of the key components of RPE are fuel pin radius (1.05%), fuel temperature (0.30%), graphite density (0.24%), equivalent boron concentration (0.16%), graphite temperature (0.63%), and specific power (0.17%), and WIMs Code (0.25%).

Figure 1 illustrates the effect of different sources of error. In the first five plots, the local errors are illustrated for each of the measured isotopic ratios. In these plots, the dashed line represents the "true" relationship between fluence and the ratio, while the solid line represents the relationship used in the local estimation procedure. The difference between the two curves is local RPE. The points scattered about the dashed line represent measurement error. The local estimation procedure uses the solid curves to produce a best estimate for Pu-fluence at each sampling location.

The last plot in Fig. 1 represents the results of the global regression procedure. In this plot the estimated fluence is plotted against true fluence at each sample location. Heasler et al. [4] includes more details, including some additional worst-case results in which all error sources were increased to extremely large values, and and in one example, the RPEs were forced to vary in the same direction. So, for example, both fuel pin radius and temperature would be forced to have positive errors. These addition results resulted in RMSEs of 2% to 6% (the 6% is arguably worse than worst-case).

4. DISCUSSION

We presented GIRM and a simulation method for assessing its RMSE. Local and global RPEs have been considered, with local RPEs appearing to be the largest error source. Although we focused on estimating Pu production for a specific example reactor, clearly the approach can be applied to almost any graphite reactor.

Although some basic knowledge of the reactor operating parameters should be available, the methodology can be applied to reactors for which no detailed operating history is available. If a sufficient number of samples (say 100 or so) can be taken, total Pu production can be estimated to within a few percent. If, on the other hand, a basic shape of the fluence profile can be calculated from operating history, a better estimate of total Pu production can be obtained with fewer samples. For example, in the Trawsfynydd exercise [3], estimates were produced from approximately 30 samples.

GIRM is a unique blend of reactor physics models and empirical data. Any error associated with the local (WIMS) model or global 3D model has important consequences for the resulting total Pu estimate. The simulation quantifies the effects of error sources, and allows us to evaluate sampling schemes. Benchmark comparisons [3] add to our confidence in the performance claims here. However, RPE remain the dominant error source and it is important to understand potential pitfalls.

The local RPE considered here are due to improperly specified parameters; implicitly we have assumed that the model is truth, provided the parameter values (such as fuel pin radius) are accurately specified and that numerical accuracy is good. In experimental

Figure 1. Local and Global Estimation Results
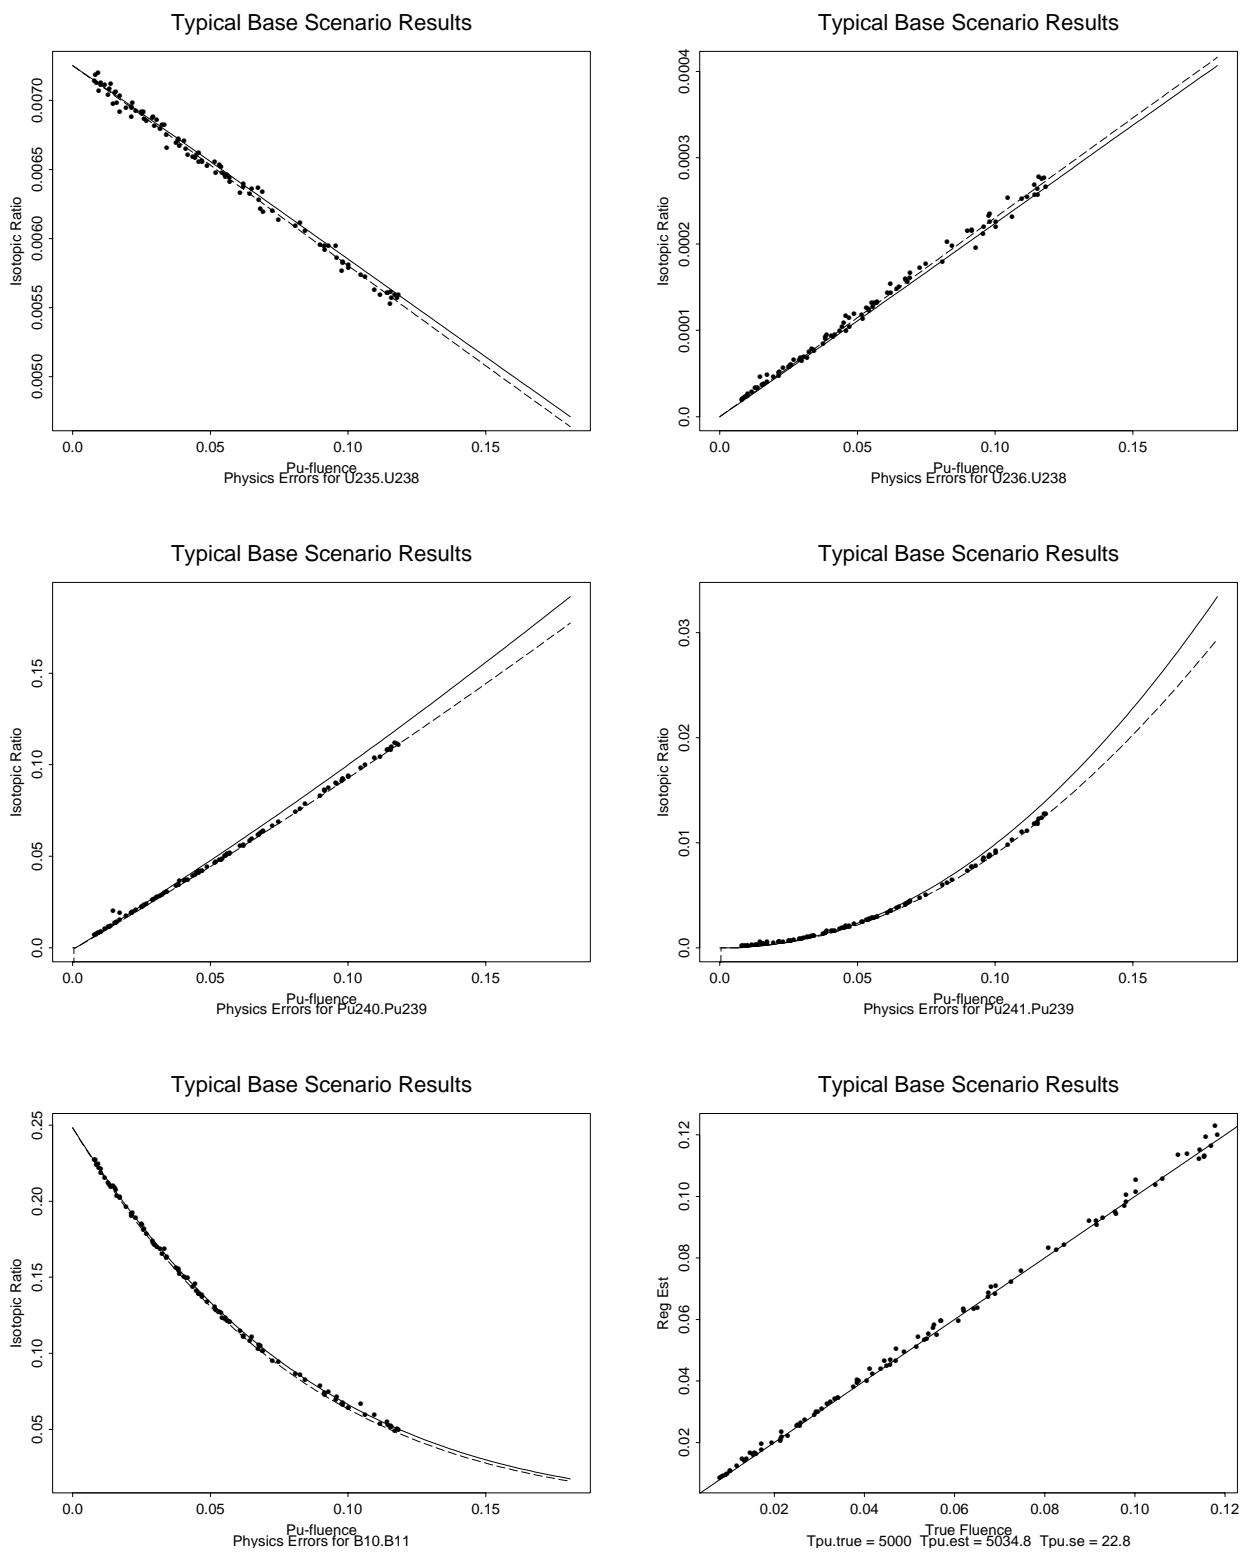
data, any outlying ratios in curves such as in Fig. 1 are assumed to be due to sample contamination or other measurement problems. The fact that other models give similar results provides some assurance that pure model error is minimal, but the main assurance is the experimental data such as in [3].

The impact of global (KENO) model error in predicting the 3D fluence field depends critically on the GIRM strategy involving fitting the local Pu estimates using a special set of basis functions, and sample size/location choice. Because there are constraints on the 3D shape of the flux (it can have spikes near control rods, but generally it is relatively smooth), it is reasonable to assume that not too large a basis set is required for a good fit (allowing accurate integration of the local Pu estimates). Also because of known constraints on flux shapes, we did not allow arbitrary true flux shapes. Instead, in the simulations, the true flux was randomly selected to be one of the seven shapes described. All seven shapes were relatively smooth, and it is possible in a new application (although highly unlikely) that the true shape would have sharp peaks and valleys, leading to understatement of the 3D model error. The use of basis functions to fit the 3D fluence field has several advantages compared to using the single best-guess field shape. For example, in the evaluation presented here if the true field shape were assumed to vary randomly around the best-guess field shape then the estimated impact of global model error would most likely have been too low. In addition, the sampling was optimized under the assumption that the true shape was among the seven choices. However, a linear combination of the basis functions can accurately approximate an arbitrary field if enough terms are used. In the example provided, because sufficient detail was provided for the target reactor to have a "best guess" shape that is close to the true shape, it is unlikely that the 3D shape could be even as misspecified as assumed in the simulation, so if anything, experts believe that Table 2 overstates, but not drastically, the impact of global model uncertainty. Therefore, performance claims have not yet been attempted for the situation where initial sample results indicated an extremely rough 3D flux field, leading to a request for more samples at key locations. So, performance claims for the situation in which test samples not used for fitting indicated a problem with the 3D fit would be desired on an as-needed basis.

REFERENCES

1. R. J. Talbert, B. D. Reid, P. G. Heasler, D. C. Gerlach, C. M. Heeb, and K. A. Pauley, Accuracy of Plutonium Production Estimates from Isotopic Ratios in Graphite Reactors, PNL-RTC-0693-REV1-DEL, Feb. 1995.
2. B. D. Reid, Morgan, E. F. Love, D. C. Gerlach, S. L. Peterson, Livingston, L. R. Greenwood, and McNeece, Graphite Isotope Ratio Method Development Report: Irradiation Test Demonstration of Uranium as a Low Fluence Indicator," PNNL-13056, Sept. 1999.
3. B. D. Reid, D. C. Gerlach, P. G. Heasler, and J. V. Livingston, "Trawsfynydd Plutonium Estimate," PNNL, Sept. 1997.
4. P. G. Heasler, B. D. Reid, C. J. Gesh, D. C. Gerlach, S. L. Petersen, C. K. Bayne, and T. L. Burr, Estimation Procedures and Error Analysis for GIRM, PNNL-14332, July 2003.