

DEVELOPMENT OF A GENETIC ALGORITHM FOR NEUTRON ENERGY SPECTRUM ADJUSTMENT

Richard M. Vega* and Edward J. Parma

Sandia National Laboratories[†]

1515 Eubank SE, Albuquerque, NM, 87123

rmvega@sandia.gov; ejparma@sandia.gov

ABSTRACT

We describe a new method for neutron energy spectrum adjustment which uses a genetic algorithm to minimize the difference between calculated and measured reaction probabilities. The measured reaction probabilities are found using neutron activation analysis. The method adjusts a trial spectrum provided by the user which is typically calculated using a neutron transport code such as MCNP. Observed benefits of this method over currently existing methods include the reduction in unrealistic artifacts in the spectral shape as well as a reduced sensitivity to increases in the energy resolution of the derived spectrum. The method has thus far been used to perform spectrum adjustments on several spectrum-modifying environments in the central cavity of the Annular Core Research Reactor (ACRR) at Sandia National Laboratories, NM. Presented in this paper are the adjustment results for the polyethylene-lead-graphite (PLG) bucket environment along with a comparison to an adjustment obtained using the code LSL-M2, which uses a logarithmic least squares approach. The genetic algorithm produces spectrum-averaged reaction probabilities with agreement to measured values, and comparable to those resulting from LSL-M2. The true benefit to this method, the reduction of shape artifacts in the spectrum, is difficult to quantify but can be clearly seen in the comparison of the final adjustments.

Key Words: Spectrum Adjustment, Spectrum Unfolding, Genetic Algorithm, Reactor Dosimetry

1 INTRODUCTION

Spectrum adjustment is the process of taking a computed trial spectrum, and adjusting it so that it is more consistent with experimental results. The trial or *a priori* spectrum is typically calculated using neutron transport methods, and hence is prone to modeling errors (e.g. geometry, material composition, temperature, and density), nuclear data uncertainties, and model defects in the physics of the radiation transport. Still, the computed spectrum is often the best initial guess at the true spectrum, as it is typically impossible to measure the neutron energy spectrum in a reactor with high resolution. The experimental data used in spectrum adjustment consists of detector responses or neutron activation analysis (NAA). These experimental measurements are integral quantities of the product of the neutron energy spectrum and an energy dependent response function.

*Richard Vega is currently a graduate student at Texas A&M University

[†]Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Spectrum adjustment is inherently an under-determined problem if a high resolution spectrum is desired as the final product. Treating the flux in each energy group as a variable, the integral quantities that result from detector readings or dosimetry foils each correspond to a single equation. Typically, the number of energy groups desired is in the hundreds, while the number of practical detectors or dosimetry foils is less than 50. The desire to know the spectrum as accurately as possible stems from the desire to know other integral quantities (such as displacements per atom and flux above or below a certain threshold energy) as accurately as possible. While modern Monte Carlo neutron transport codes can take advantage of parallel architectures to calculate neutron energy spectra with extremely high resolution and very small statistical variance in the flux values, model errors such as material temperature and composition uncertainties, are unavoidable.

When performing spectrum adjustments, a solution is being sought to an under-determined problem which by definition has no unique solution. It is therefore desirable to use expert judgment when evaluating the quality of an adjustment by any single method. For example, if an adjustment produces a spectrum consisting of nothing but jagged spikes throughout the entire energy range of interest, yet the predicted reaction probabilities match the experimental data very well, how credible is the adjusted spectrum? Unfortunately, this question is not entirely hypothetical. This exact type of situation arises from iterative perturbation techniques in codes such as SAND-II [1] when too many iterations are performed; the spectrum begins to look more and more unrealistic, yet the predicted reaction probabilities move closer and closer to the measured data [2]. For this reason, closeness to the measured reaction probabilities cannot be the only measurement used to assess the fidelity of a spectrum adjustment.

1.1 Desired Properties of a New Method

Due to the fact that spectrum adjustment is an under-determined problem, many possible adjustments will exist for any given set of input data, making all results somewhat subjective. By definition, any method should adjust the trial spectrum so that when folded with the appropriate dosimetry cross sections, it predicts the measured integral quantities more accurately than the trial spectrum. We can however list general qualities that would be desirable of a new method:

1. The adjustment should preserve physically meaningful features in the trial spectrum without introducing new ones that have no physical foundation.
2. The adjusted spectrum should have smaller energy-dependent uncertainties than the trial spectrum, and the uncertainty in the adjusted spectrum should be quantifiable.
3. The adjustment method should contain as few "knobs", i. e. arbitrary adjustable parameters, as possible.
4. The adjustment method should be independent of the energy grid structure and the quality/credibility of the adjustment should not deteriorate as the number of energy groups increases.

The first of these qualities has historically been the hardest to achieve. Even methods that are not iterative, and hence are not subject to the spikes that result from over iteration, can still produce unrealistic dips and peaks in their adjusted spectra due to the structure in the underlying cross sections for the measured activities. This occurs if the correlation in the energy-dependent spectrum is not properly represented in the covariance matrix for the trial spectrum. For instance, Figure 1 shows two spectrum adjustments recently performed using LSL-M2 at the Annular Core Research Reactor (ACRR). The spikes in the trial spectrum are evidence of resonances in the transport cross sections. For example, a large dip in the spectrum corresponds to the presence of a nuclide in the model with a strong absorption resonance at that energy. The peaks and dips in the adjusted spectrum that do not align with any such physical features that relate back to the radiation transport modeling of the trial spectrum are unrealistic. These features were introduced by the adjustment method solely for the purpose of bringing the predicted reaction probabilities into better agreement with the measured data while reflecting the structure in the cross sections supporting the activity measurements and the stiffness or correlation assumed in the trial spectrum. So the question is can we get the same agreement without introducing these artifacts?




Figure 1. (a) Spectrum adjustment for the PLG bucket environment in the central cavity of the ACRR using LSL-M2. (b) Spectrum adjustment for the free field environment in the central cavity of the ACRR using LSL-M2.

The second quality is not terribly difficult to achieve, yet it can be very difficult to reduce an uncertainty comparison of the spectral shape to a simple single value metric. While uncertainty in the energy-dependent spectral shape can be readily represented in the form of a covariance matrix, it requires many more computational resources to sample all the sources of uncertainty in the adjustment process and it is even more difficult to derive the covariance in the trial spectrum in a manner consistent with its derivation, i.e. via a radiation transport calculation. Performing a detailed uncertainty quantification of a model with so many uncertain parameters can be a very time consuming process especially when a single run of a Monte Carlo based radiation transport code with high energy-resolution can take days to reduce the statistical sampling errors to a reasonable level ($< 1\%$). Nonetheless many methods, such as least squares, use an estimation of the covariance

matrix of the trial spectrum to perform the adjustment. Due to the need to capture constraints in the input trial spectrum, physics-based approximation techniques are often used to estimate the covariance matrix for the trial spectrum [3].

The third quality is particularly important given the skepticism that any adjustment code is likely to face. If it is believed that a code can be made to give a better solution by simply tweaking some parameters, it will likely not be looked upon favorably. Note that this may not be the case for problems that have a unique solution such as the Boltzmann solution to the neutron transport equation, where increasing the grid resolution in any of the independent variables is expected to give a better solution. This is indeed a common criticism of iterative perturbation techniques where expert judgment must be used to determine how many iterations is enough to balance the agreement with experimental data and the credibility of the adjusted spectrum. If there are parameters that can be adjusted, they should have an optimal value that consistently produces the best result or they should be analogous to grid refinement or particle histories in Monte Carlo, where a larger value always produces a better result.

Finally, the fourth quality is difficult to achieve for reasons that are very easy to understand. Adding energy groups adds variables to an under-determined problem without adding any equations, making it even more under-determined. Least squares approaches typically address the added degrees of freedom while providing a unique solution by considering the energy-dependent constraints provided by the representation of the trial spectrum, i.e. its covariance matrix, and by the uncertainty in the cross sections for the activity metrics. Any method that uses least squares is going to be sensitive to the energy-group structure. The user manual for LSL-M2 warns that for high resolution energy grids, the results will show larger uncertainties [4]. The integral constraints provided by the activity metrics provide limits on the possible spectrum deviations over larger energy intervals but will permit large deviations within narrow energy regions unless constrained by the energy-dependent covariance in the trial spectrum. Least squares methods can also exhibit numerical instabilities from the necessity to invert large dimensioned matrices associated with high resolution energy grids. Thus, it is desirable for any new method to be less sensitive to increased energy grid resolution.

1.2 The Genetic Algorithm

Genetic algorithms have been used in recent decades to solve logistics problems such as the traveling salesman, word matching, and number partitioning problems. This abstract stochastic method has been gaining popularity steadily in the past decade and has made its way into the world of engineering as evidenced by NASA's use of a genetic algorithm for antenna design [5]. If a problem can be abstracted to a set of genes with a clear fitness function, a genetic algorithm will probably be applicable.

The genetic algorithm is essentially just a simulation designed to mimic natural selection. The simulation starts off with a population of possible solutions, or specimens, and each specimen is assigned a fitness value based on what makes any single specimen favorable over any other. The specimens are then chosen for mating in such a way that higher fitness specimens are selected for

mating more frequently. The mating of these specimens produces children specimens which carry onto the next generations and hopefully carry the favorable characteristics of their parents with them. The process is then repeated for a set number of generations.

If a genetic algorithm is successful, the fitness of the population as a function of generations will increase in its minimum, maximum, and average. At a certain point, convergence will be achieved, and the fitness levels of the population will cease to rise any further. If a genetic algorithm is to show these symptoms of success, the individual processes required to initiate and propagate the specimens through the generations must be chosen wisely. These processes include:

- Representing the solution as a set of genes
- Setting the initial population
- Designing the fitness function
- Selecting the parents
- Mating the parents
- Mutating the children

The next section will describe in detail how these processes are chosen. Termination of the algorithm occurs either when the user defined number of generations have completed, or the relative difference between the maximum fitness in consecutive generations is below a user defined tolerance. The flexibility in creating the processes listed above and the ability to map the process into the metrics important to the analysis is one of the things that makes the genetic algorithm so appealing from a computational point of view. For the genetic algorithm presented here, several different methods for each step listed above were considered before a successful formula was found.

2 ALGORITHM DESCRIPTION

The first step in creating a genetic algorithm is finding a way to represent the population of solutions as specimens in a gene pool. This requires defining each specimen as a collection of genes, which have specific values at locations known as gene sites. For the genetic algorithm presented here, this means representing each possible spectrum in terms of a unique set of genes. Once this is defined, the selection, mating, and mutation processes must be defined, along with a way to determine the fitness of any individual specimen.

To perform this abstraction, the gene sites are chosen to be discrete energy values in the range of the trial spectrum being adjusted. As the energy spectrum in any reactor typically spans many orders of magnitude, the algorithm finds the minimum and maximum energies of the trial spectrum and distributes N points in the energy domain, equidistant in base 10 logarithmic space, starting at the minimum and ending at the maximum. At each gene site, there must be a gene value. For the

algorithm considered here, this value can be thought of as a relative adjustment factor to the trial spectrum in the region of the gene site that it resides at.

The gene values must then be used to define the adjusted spectrum. This is done by performing a least squares polynomial fit through the gene values. The polynomial will be called the shift function. The flux in each group of the trial spectrum is then multiplied by the value of the shift function at the midpoint of its energy group. In this way, the relative magnitude of the trial spectrum is adjusted differently in different regions of the energy domain, however the polynomial nature of the shift function assures that this adjustment is smooth for low order polynomials. This reduces the likelihood of introducing artifacts in the spectral shape. With this abstraction clearly defined, we can begin to describe the various processes that make up the genetic algorithm.

2.1 Initial Population Formation

To set the initial population of possible solution spectra, the population size is first selected by the user. This population size should be in the hundreds to fully explore the solution space. For each specimen in the initial population, the gene value at each site is chosen to be randomly perturbed from a baseline value of unity. To do this, a normal random variable with a mean of zero and a relative standard deviation of 0.07 is sampled, and then added to unity to obtain the gene value. Once each specimen has been assigned gene values for all gene sites, the shift function for each specimen is obtained, and the specimen spectrum is calculated as described earlier in this section.

2.2 Fitness Function

The fitness function should quantify what makes any particular specimen better than any other. Specimens with high fitness values will be more likely to participate in mating, thereby passing their genes to the next generation. In the case of spectrum adjustment, the only obvious trait that can be quantified numerically is the agreement between the reaction probabilities calculated using the specimen and appropriate dosimetry cross sections, and the measured reaction probabilities. The fitness function is then defined as

$$f = C - \sum_{i=1}^m \frac{\left| \left\{ \sum_{j=1}^n \sigma_{j,i} \phi_j \Delta E_j \right\} - r_i \right|}{r_i} \quad (1)$$

where m is the number of foils used in the NAA, n is the number of energy groups, $\sigma_{j,i}$ is the dosimetry cross section for foil/reaction i in energy group j , ϕ_j is the flux in energy group j of the specimen spectrum, ΔE_j is the energy bin width of energy group j , r_i is the measured reaction probability for foil/reaction i , C is an arbitrary constant, and f is the fitness of the specimen. The sum of relative differences between calculated and measured reaction probabilities is subtracted from a constant so that when the difference is small, the fitness is large.

2.3 Parent Selection

In order to produce the next generation, specimens from the current generation must be chosen for mating to produce children until the number of children equals the population size, which is held constant throughout all generations. The selection of parents must be performed such that high fitness specimens are chosen more frequently for mating than lower fitness specimens. This is done so that desirable features are more likely to propagate to the next generation. The algorithm uses what is known as proportional selection. With this selection method, the probability of specimen j to be chosen for mating is

$$P_j = \frac{f_j - f_{min}}{P_t} \quad (2)$$

where f_j is the fitness of the specimen, f_{min} is the minimum fitness of all specimens in the current generation, and P_t is defined by

$$P_t = \sum_{i=1}^s (f_i - f_{min}) \quad (3)$$

where s is the number of specimens in a generation. In this way, the specimen with the lowest fitness will never be selected for mating, and the specimen with the highest fitness will have the highest probability of being selected for mating, and all other specimens will have probabilities of being selected linearly between these two extremes.

2.4 Mating Algorithm

In order to produce children, the chosen parents must mate in a way that passes on their characteristics to the next generation. There are many ways that this could be performed that vary in the number of parents that participate, and the number of offspring formed. The algorithm uses what is known as single point crossover. This method takes two parents and produces two children. Single point crossover is performed by choosing a cutting point between two consecutive gene sites at random, and exchanging all genes to the right of this cutting point between the two parents. This is illustrated in Figure 2 which shows the gene values and shift functions before and after mating.

2.5 Mutation

It is not uncommon for genetic algorithms to stagnate on a solution that has a high fitness, but not the highest fitness possible given the constraints of the method. For this reason, it is necessary to give the specimens a random kick once in a while to ensure that the entire solution space is being explored. This is known as mutation, and it occurs at a rate set by the user. For instance, a mutation rate of 0.1 means that one out of every ten children produced will undergo some sort of random




Figure 2. (a) Gene values and shift functions of the parents. (b) Gene values and shift functions of the offspring.

mutation. Mutations in the algorithm presented here correspond to selecting a gene site at random and adding a normal random variable with a mean of zero and a relative standard deviation of 0.02 to the value of the gene at that site. Care should be taken that the mutation rate is not set so high that convergence is never achieved. Values as high as 0.15 have been used successfully in the algorithm presented here.

3 CONCLUSIONS

The results presented here are for the polyethylene-lead-graphite (PLG) bucket environment in the ACRR central cavity, although adjustments have also been performed for the lead-boron (LB44) bucket environment as well as the free field environment with similar results to be published at a later date. A total of 37 reactions were considered using a total of 20 activation foils, some with different cover combinations, in the NAA. As mentioned in the introduction, the quality of an adjustment should not be judged solely on how well the adjusted spectrum predicts the measured reaction probabilities, but also on the physical meaningfulness of the adjusted spectrum. A large dip or peak in the adjusted spectrum in a region where no transport cross section resonance exists in the nearby reactor materials and no structural feature is thought to exist in the neutron source term, ^{235}U here, should raise questions about whether the structure in the spectrum can be physically meaningful. Figure 3 shows a comparison of the adjustments for the PLG bucket environment using LSL-M2 and the genetic algorithm.

Figure 3 shows the true benefit of the genetic algorithm over conventional least squares approaches. The first thing to notice from this figure is that both adjustment methods appear to be conveying the same qualitative result; the trial spectrum is over-predicting the thermal flux, and under-predicting the fast flux. This could be due to a number of things such as incorrect

temperatures or densities in the transport model. The two methods only differ significantly in the mid-energy region. Whereas LSL-M2 predicts oscillations in the energy range $10^{-6} - 10^{-1}$ MeV, the genetic algorithm does not.




Figure 3. (a) PLG spectrum adjustment using LSL-M2. (b) PLG spectrum adjustment using the genetic algorithm.

While it is unfortunate that the true benefit of the genetic algorithm is difficult to quantify, it is possible to examine whether it predicts some metrics, such as the reaction probabilities, as well as LSL-M2 and other existing methods. LSL-M2 was chosen for comparison because it is not an iterative perturbation method like the commonly used SAND-II code [1], and hence it is not sensitive to the number of iterations performed, but also because it is widely used and belongs to a large family of least squares adjustment methods that are well tested and trusted by the radiation effects community. Table I shows the measured reaction probabilities alongside the reaction probabilities predicted by the adjusted spectra from LSL-M2 and the genetic algorithm, as well as their percent differences.

The results given in Table I show that the genetic algorithm predicts reaction probabilities closer to the measured value for 20 out of the 37 reactions considered in this study. Of course, the method would not have been considered unsuccessful if it had predicted less reaction probabilities closer to the measured values, as long as the results were comparable to those found using LSL-M2. Perhaps one of the biggest advantages in the genetic algorithm is that the method is insensitive to the number of energy groups used, and does not exhibit numerical instabilities when a large number of energy groups are used. The adjustments shown in Figure 3 used an 89 group structure, however the same adjustment is achieved using 640 energy groups as shown in Figure 4. This independence is due to the fact that the shift function for each specimen in the genetic algorithm is decoupled from the energy group structure of the trial spectrum, and hence produces a nearly identical polynomial in most cases at higher resolution.

Table I. Comparison of the reaction probabilities predicted by LSL-M2 and the genetic algorithm. The reference reaction $^{58}\text{Ni}(n,p)^{58}\text{Co}$ shows no difference because it is used to normalize both the trial and adjusted spectra.

Reaction-Cover	LSL	Genetic	Measured	LSL % diff	Genetic % diff
$^{27}\text{Al}(n,\alpha)^{24}\text{Na}$ -Bare	3.736E-13	3.768E-13	3.774E-13	1.008	0.171
$^{197}\text{Au}(n,\gamma)^{198}\text{Au}$ -Bare	2.439E-07	2.386E-07	2.404E-07	1.467	0.756
$^{197}\text{Au}(n,\gamma)^{198}\text{Au}$ -Cd	1.803E-07	1.777E-07	1.788E-07	0.841	0.630
$^{59}\text{Co}(n,\gamma)^{60}\text{Co}$ -Bare	2.797E-08	2.720E-08	2.724E-08	2.694	0.140
$^{59}\text{Co}(n,\gamma)^{60}\text{Co}$ -Cd	5.391E-09	5.591E-09	5.329E-09	1.165	4.918
$^{59}\text{Co}(n,p)^{59}\text{Fe}$ -Cd	8.216E-13	8.014E-13	8.157E-13	0.723	1.753
$^{59}\text{Co}(n,2n)^{58}\text{Co}$ -Cd	1.128E-13	1.112E-13	1.117E-13	0.989	0.425
$^{63}\text{Cu}(n,\alpha)^{60}\text{Co}$ -Bare	3.014E-13	2.933E-13	3.596E-13	16.182	18.426
$^{63}\text{Cu}(n,\gamma)^{64}\text{Cu}$ -Bare	3.310E-09	3.220E-09	3.305E-09	0.162	2.577
$^{63}\text{Cu}(n,\gamma)^{64}\text{Cu}$ -Cd	5.135E-10	5.490E-10	4.975E-10	3.216	10.353
$^{54}\text{Fe}(n,p)^{54}\text{Mn}$ -Bare	5.032E-11	5.024E-11	4.970E-11	1.249	1.089
$^{56}\text{Fe}(n,p)^{56}\text{Mn}$ -Bare	6.122E-13	5.970E-13	6.200E-13	1.255	3.714
$^{58}\text{Fe}(n,\gamma)^{59}\text{Fe}$ -Bare	9.783E-10	9.375E-10	9.335E-10	4.804	0.426
$^{58}\text{Fe}(n,\gamma)^{59}\text{Fe}$ -Cd	1.548E-10	1.512E-10	1.493E-10	3.667	1.304
$^{115}\text{In}(n,n')^{115m}\text{In}$ -Bare	1.545E-10	1.564E-10	1.651E-10	6.434	5.274
$^{24}\text{Mg}(n,p)^{24}\text{Na}$ -Bare	8.013E-13	8.042E-13	7.661E-13	4.592	4.973
$^{55}\text{Mn}(n,\gamma)^{56}\text{Mn}$ -Cd	1.490E-09	1.436E-09	1.487E-09	0.209	3.400
$^{55}\text{Mn}(n,2n)^{54}\text{Mn}$ -Bare	1.168E-13	1.151E-13	1.375E-13	15.045	16.271
$^{98}\text{Mo}(n,\gamma)^{99}\text{Mo}$ -Bare	8.740E-10	8.642E-10	8.667E-10	0.842	0.284
$^{98}\text{Mo}(n,\gamma)^{99}\text{Mo}$ -Cd	7.940E-10	7.881E-10	7.982E-10	0.528	1.271
$^{23}\text{Na}(n,\gamma)^{24}\text{Na}$ -Bare	3.242E-10	3.155E-10	3.057E-10	6.064	3.219
$^{23}\text{Na}(n,\gamma)^{24}\text{Na}$ -Cd	3.358E-11	3.547E-11	3.245E-11	3.469	9.305
$^{93}\text{Nb}(n,2n)^{92m}\text{Nb}$ -Bare	2.438E-13	2.435E-13	2.411E-13	1.124	0.976
$^{58}\text{Ni}(n,p)^{58}\text{Co}$ -Bare	6.879E-11	6.879E-11	6.879E-11	0.000	0.000
$^{58}\text{Ni}(n,2n)^{57}\text{Ni}$ -Cd	2.293E-15	2.332E-15	2.152E-15	6.569	8.382
$^{60}\text{Ni}(n,p)^{60}\text{Co}$ -Cd	1.228E-12	1.189E-12	1.254E-12	2.040	5.168
$^{238}\text{U}(n,f)\text{FF-B}$	2.250E-10	2.281E-10	2.194E-10	2.546	3.958
$^{235}\text{U}(n,f)\text{FF-B}$	2.463E-09	2.495E-09	2.573E-09	4.257	3.040
$^{239}\text{Pu}(n,f)\text{FF-B}$	2.796E-09	2.803E-09	2.570E-09	8.776	9.072
$^{32}\text{S}(n,p)^{32}\text{P}$ -Bare	5.845E-02	5.886E-02	5.437E-02	7.501	8.262
$^{45}\text{Sc}(n,\gamma)^{46}\text{Sc}$ -Bare	1.802E-08	1.730E-08	1.731E-08	4.087	0.057
$^{45}\text{Sc}(n,\gamma)^{46}\text{Sc}$ -Cd	1.295E-09	1.314E-09	1.352E-09	4.208	2.791
$^{46}\text{Ti}(n,p)^{46}\text{Sc}$ -Bare	6.700E-12	6.509E-12	6.400E-12	4.684	1.698
$^{47}\text{Ti}(n,p)^{47}\text{Sc}$ -Bare	1.237E-11	1.243E-11	1.265E-11	2.188	1.723
$^{48}\text{Ti}(n,p)^{48}\text{Sc}$ -Bare	1.639E-13	1.628E-13	1.625E-13	0.839	0.183
$^{64}\text{Zn}(n,p)^{64}\text{Cu}$ -Bare	2.403E-11	2.408E-11	2.451E-11	1.940	1.735
$^{90}\text{Zr}(n,2n)^{89}\text{Zr}$ -Bare	5.670E-14	5.689E-14	5.748E-14	1.353	1.023




Figure 4. (a) PLG spectrum adjustment using the genetic algorithm with 89 energy groups. (b) PLG spectrum adjustment using the genetic algorithm with 640 energy groups.

With the results presented above, a few remarks regarding the method are in order. First, and most importantly, it should be noted that the polynomial shift functions are responsible for the smooth adjustment, and not the genetic algorithm itself. The genetic algorithm serves as the optimization technique to converge on the correct shape of the polynomial shift function. It should be acknowledged that other methods could be used to perform this optimization. In addition, the shift function serves the purpose of a relative adjustment factor as a function of neutron energy. Polynomials were chosen out of convenience, but it is recognized that this constraint is not physics based. The value of the shift function at the gene sites however, is physics based via the fitness function in the genetic algorithm, while the polynomial serves as a means of curve fitting between the gene sites. It should also be noted that least squares methods can include smoothing constraints, and LSL-M2 does include such a constraint in a mathematically defensible manner by using an “a priori” covariance matrix on the trial spectrum. This trial spectrum constraint is used in conjunction with the covariance in the underlying dosimetry cross sections which characterize the energy-dependence of the response functions of the activation reactions. The reason the peaks still exist despite this smoothing mechanism is that the cross section constraints are much tighter than those imposed on the trial spectrum, as is often the case. External smoothing mechanisms applied to the adjusted spectrum itself would unfortunately smooth out details that actually do exist in the spectrum such as sharp dips at resonances.

One shortcoming of the genetic algorithm presented here is that it does not provide uncertainties or correlations for the energy-dependent flux as do most least squares methods. In fact, many experts in the field of spectrum adjustment would state that the adjusted spectrum is not as significant as the determination of the uncertainties associated with it [6]. Work is currently being performed to remedy this shortcoming. The sources of uncertainty in the adjustment process, namely the uncertainty in the measured reaction probabilities and dosimetry cross sections used to compute

them, are currently being accounted for and their effects on the adjusted spectrum are being studied. The uncertainties in the trial spectrum are much more difficult to quantify. This difficulty to quantify stems from the inability to generate a significant number of continuous energy Monte Carlo transport runs to rigorously sample different temperatures, densities, dimensions, and material compositions, not to mention the difficulty in analyzing the energy-dependent uncertainty in the transport cross sections. Another critical issue with the use of forward radiation transport to calculate the uncertainty in the spectrum is “model defect.” Many of the uncertainty terms in the calculation are not directly traceable to input parameters but are attributed to defects in the models used. An example is the shape of the fission neutron spectrum used. While a detailed uncertainty quantification of the forward problem may be easier to handle using deterministic methods, continuous energy calculations are preferred to capture the effect of resonance structures in the transport cross sections. While the uncertainty quantification of the forward problem using continuous energy Monte Carlo transport is doable, it is difficult to say the least.

4 ACKNOWLEDGMENTS

We would like to thank Patrick Griffin for his helpful feedback and review during the code development process. His expertise in the area of spectrum adjustment has been invaluable. In addition, we would like to thank Kenneth Reil for his support. Finally we would like to thank Sandia National Laboratories.

5 REFERENCES

- [1] W. McElroy, S. Berg, T. Crockett, R. Hawkins, and A. I. C. P. CA., *A Computer-automated Iterative Method for Neutron Flux Spectra Determination by Foil Activation. Volume 1. a Study of the Iterative Method*, Defense Technical Information Center (1967).
- [2] W. Zijp, “Comparison of Neutron Spectrum Unfolding Codes,” 1978, IAEA Report: IAEA-R-1811-F.
- [3] J. G. Williams et al., “Simultaneous Evaluation of Neutron Spectra and 1-MeV-Equivalent (Si) Fluences at SPR-III and ACRR,” *IEEE Transactions on Nuclear Science*, **54**, pp. 2296–2302 (2007).
- [4] F. W. Stallmann, “LSL-M2: A Computer Program For Least-Squares Logarithmic Adjustment of Neutron Spectra,” 1986, NUREG/CR-4349, ORNL/TM-9933.
- [5] G. S. Hornby, J. D. Lohn, and D. S. Linden, “Computer-Automated Evolution of an X-Band Antenna for NASA’s Space Technology 5 Mission,” *MIT Press*, **19**, pp. 1–23 (2011).
- [6] M. Reginatto, “Overview of Spectral Unfolding Techniques and Uncertainty Estimation,” *Radiation Measurements*, **45**, pp. 1323–1329 (2010).