SAND2015-11036

SANDIA REPORT

SAND2014-0000
Unlimited Release
Printed August 2014

GenSpec: A Genetic Algorithm for Neutron Energy Spectrum Adjustment

Richard M. Vega and Edward J. Parma

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

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.

Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: http://www.osti.gov/bridge

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd
Springfield, VA 22161

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.fedworld.gov
Online ordering: http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online

SAND2014-0000
Unlimited Release
Printed August 2014

GenSpec: A Genetic Algorithm for Neutron Energy Spectrum Adjustment

Richard M. Vega and Edward J. Parma
Applied Nuclear Technologies Department 1384
Sandia National Laboratories
P.O. Box 5800
Albuquerque, NM 87185
rmvega@sandia.gov and ejparma@sandia.gov

Abstract

Presented in this report is the description of 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 artefacts in the spectral shape as well as a reduced sensitivity to increases in the energy resolution of the derived spectrum. This report presents the adjustment results for various spectrum altering bucket environments in the central cavity of the Annular Core Research Reactor, as well as the adjustment results for the spectrum in the Sandia Pulse Reactor III. In each case, the results are compared to those generated using LSL-M2, which is a code commonly used for the purpose of spectrum adjustment. 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 artefacts in the spectrum, is difficult to quantify but can be clearly seen in the comparison of the final adjustments. Beyond these preliminary results, this report also gives a thorough description of the genetic algorithm and presents instructions for running the code using the graphical user interface. In its present state, the code does not provide uncertainties or correlations for the adjusted spectrum. This capability is currently being added, and will be presented in future work.

Acknowledgment

The authors wish to thank the Annular Core Research Reactor staff and the Radiation Metrology Laboratory staff for their support of this work. In addition the authors would like to thank Patrick Griffin for his advice and support during the code development process, as well as Helmut Katzgraber of Texas A&M University for presenting the genetic algorithm as part of an undergraduate computational physics course. This course was where the idea for this work was originally born.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Prior to introducing the code GenSpec, it is necessary to briefly introduce the concept of spectrum adjustment and why it is needed. This will include a brief description of current methods of spectrum adjustment and their advantages and disadvantages. This section will also aim to point out where the various sources of uncertainty are introduced when calculating and adjusting the neutron energy spectrum of interest.

Neutron energy spectra

When performing neutron irradiation experiments at a specific location in a research reactor, the experimenter needs to know the energy distribution, or spectrum, of the neutron fluence at the location. Typically, this spectrum is desired with a high-resolution with associated uncertainties and correlations. While the calculation of this spectrum is rather straight-forward, the measurement of this spectrum can only be performed with very low, and often insufficient resolution. The spectrum is used with a response function in order to calculate integral quantities such as displacements per atom (DPA), absorbed dose, or fluence above or below a certain threshold energy, which are useful to experimenters. In order to obtain these quantities, the spectrum is "folded" with a response function which can be represented mathematically as

$$I = \int_0^\infty \phi(E)R(E)dE \tag{1.1}$$

where I is the integral quantity, $\phi(E) = d\Phi/dE$ is the differential neutron energy spectrum, Φ is the total fluence, and $R(E)$ is the energy dependent response function. The greater the resolution in each of these functions, especially in energy regions of rapidly varying response functions or spectral shape functions due to cross-section resonances, the more accuracy that can be achieved for the integral quantity. Of course, this accuracy is also effected by the uncertainty in either $\phi(E)$ or $R(E)$. Sometimes the response function is defined such that it has no uncertainty. This is certainly the case when calculating the fluence above or below a certain threshold energy, where the response function is a simple step function. Other times, the response function will be known with high accuracy and high-resolution. It is often the uncertainty and low-resolution in the energy spectrum, given the difficulty of high-resolution measurement, that dominates the uncertainty in the integral quantity.

When using a nuclear reactor as the source of neutrons for an irradiation experiment, the shape of the neutron energy spectrum is assumed fixed. The Annular Core Research Reactor (ACRR) provides experimenters with a variety of spectra to choose from when irradiating samples in its central cavity. This is done by inserting spectrum altering buckets into the central cavity. The spectra as calculated using MCNP[1] for various ACRR central cavity environments along with the spectrum of the Sandia Pulse Reactor III (SPR-III) can be seen in Figure 1.1. These bucket environments work by attenuating the fluence in certain energy ranges. For instance, the polyethylene-lead-graphite (PLG) bucket scatters more neutrons into the lower energy region prior to striking the target by introducing hydrogen in the polyethylene layer. Another purpose of the PLG bucket is to reduce the gamma-ray fluence that is also present in a nuclear reactor environment. This is the purpose of the lead annulus.

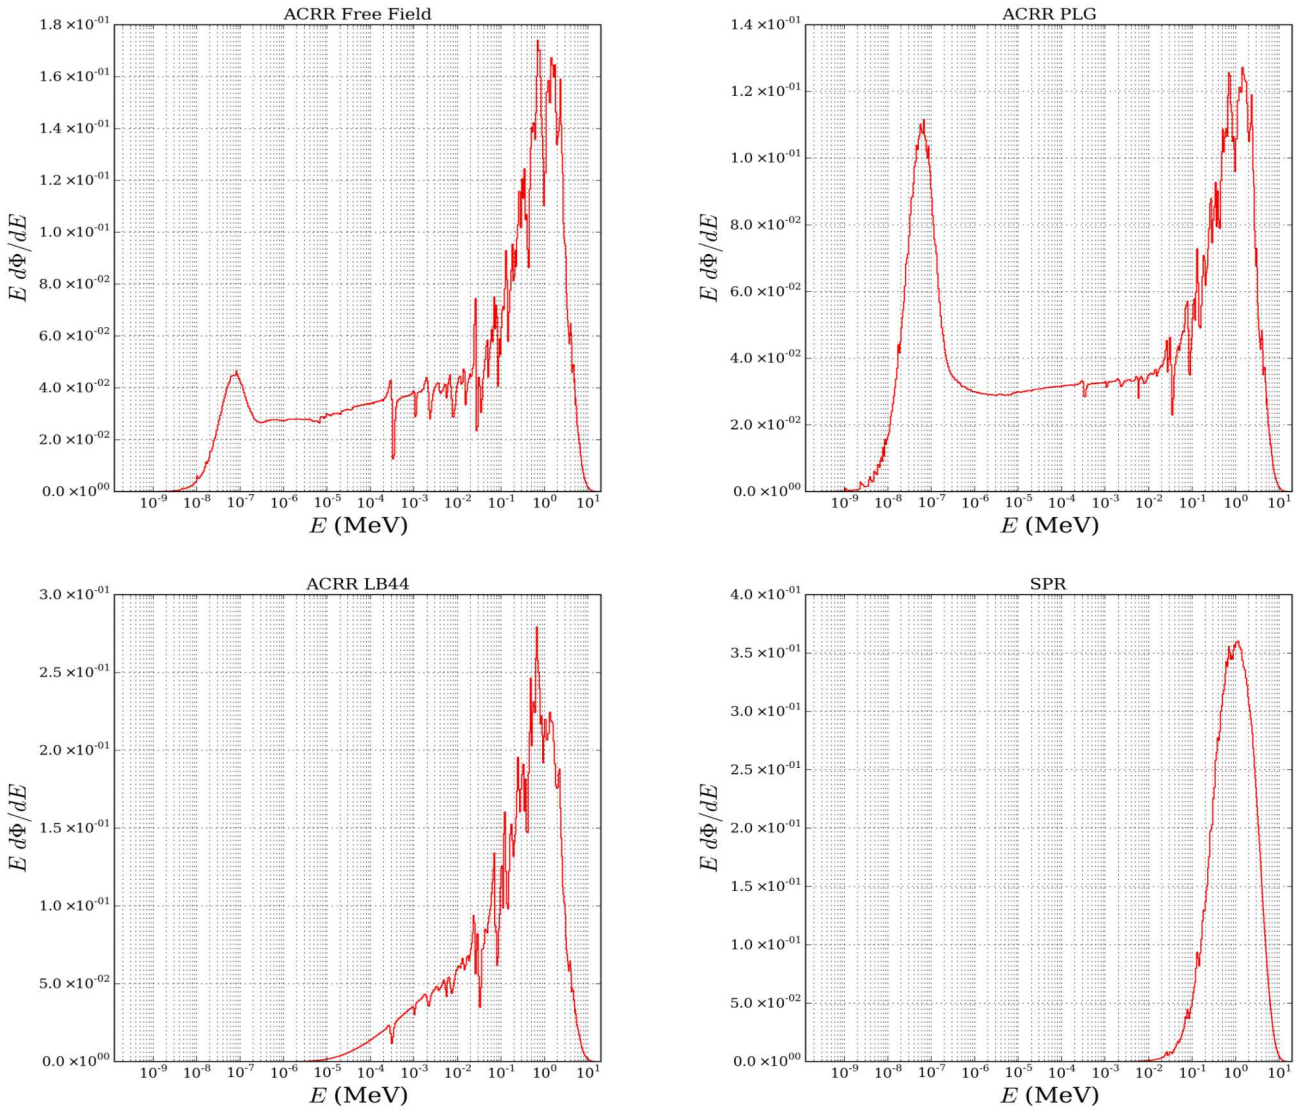

Figure 1.1: Neutron lethargy spectra for bucket environments in the central cavity of the ACRR and SPR-III.

All neutron energy spectra found in a fission reactor like the ACRR or SPR-III are going to have some similar characteristics. The first of these is the fission spectrum, which appears as a peak at the high energy end of the neutron energy spectrum. The shape of this curve is indicative of the distribution of neutron energies born from fission. If there are materials present that scatter neutrons to lower energies, there will also be a noticeable Maxwellian peak in the low-energy end of the spectrum corresponding to the material temperature. These are typically materials rich in hydrogen or other low-Z elements. The region between the two peaks is known as the slowing-down region. If we look at the free-field ACRR spectrum, and then at the 44 inch lead-boron bucket (LB44) and PLG spectra, we see the purpose of each bucket. The PLG bucket increases the low-energy component of the spectrum through hydrogen scattering, whereas the LB44 bucket reduces the this component by absorption in boron, which preferentially absorbs low-energy neutrons.

In addition to the basic shapes present in a fission reactor's spectrum, there are also sharp dips and peaks present throughout the energy range. The dips and peaks are a result of the resonances in the cross-sections for neutrons in reactor materials. For instance, oxygen has strong resonances in its inelastic scattering cross-section at fission neutron energies, and hence scatters neutrons at high rates at nearly discrete energy values near the fission spectrum peak. This is reflected in the spectrum by the sudden dips throughout the fission spectrum peak. The quickly varying behavior of the neutron cross-sections in many materials makes direct measurement of neutron energy spectra in reactors very difficult.

It is helpful in spectrum adjustment to know what the spectrum should look like. This includes explanations for the spectral shape and resonance spikes. In particular, if a spectrum adjustment code predicts a spectrum with a sharp peak where no resonance exists, this might raise questions about the quality of the calculation or integral quantity measurement. In addition, if a spectrum adjustment code predicts a significantly lower Maxwellian peak than what was calculated, it is helpful to understand the reasons for why this might be. For instance, could such an over-estimation of the low-energy fluence, via transport calculations, be simply due to incorrect temperatures in the transport model? Could it be due to a density of scattering material that is incorrect? Building intuition about the effects of model parameters on the neutron energy spectrum can help understand when a spectrum adjustment code is producing unrealistic results.

Spectrum measurement

Attempts at measuring the neutron energy spectrum in a nuclear reactor are hindered by many factors including the complexity of neutron cross-sections as a function of energy and the presence of gamma-rays that accompany fission. The two primary methods of measurement include detectors and neutron activation analysis (NAA). One detector method is the use of Bonner spheres, which are low-energy neutron detectors enclosed in different thicknesses of moderating material. Comparison of the responses from the individual spheres allows the experimenter to estimate the energy spectrum with a resolution depending upon the number of spheres used. One complication of this method is that Bonner spheres can be quite large (12 inches in diameter is not uncommon) depending upon the amount of moderating material enclosing the thermal neutron detectors. For instance, if one

wanted to use Bonner spheres to measure the neutron energy spectrum in the central cavity of the ACRR, the number of detectors would be limited by the constraint that the largest sphere must fit inside the central cavity. With each detector representing a single integral quantity, the spectrum can be "unfolded" to provide an estimate of the neutron energy spectrum.

While much information can be gained from detectors, the number of detectors available will inevitably be limited. A larger number of integral quantities can be obtained from NAA. With this method, a foil made of a specific high-purity material is placed in the irradiation environment for a known amount of time. As the foil material is activated by neutrons, radioactive isotopes are created, and are allowed to decay after removal from the environment. Taking into account both the irradiation and decay time, the activity of the foil can be measured, and this activity can be used to calculate the reaction-rate density that occurred in the foil material during the irradiation. The reaction-rate density will be proportional to an integral quantity as given by equation 1.1, with the response function $R(E)$ replaced by the macroscopic cross-section for the particular reaction that leads to the radioactive daughter being measured. The macroscopic cross-section will be denoted by $\Sigma_j(E)$, where j indicates the reaction type. This reaction-rate density will be proportional to the reaction probability which is defined similarly to the reaction-rate density with the macroscopic cross-section replaced by the microscopic cross-section. These microscopic cross-sections will be denoted by $\sigma_j(E)$, where j indicates the reaction type. Further information regarding NAA can be found in the standard radiation detection textbook by Knoll[2].

There are a wide variety of foil materials that can be used in NAA. GenSpec assumes that all integral quantities are obtained through NAA and comes with several dosimetry cross-section libraries, each consisting of a unique set of cross-sections for specific reactions. The experimenter is not limited only to the number of activation foils used. Often, a single foil material will undergo various reactions that can be inferred from the resulting activity. In addition, activation foils can be used with different cover materials that alter the spectrum that reaches the activation foil. The most common cover material is cadmium, which essentially eliminates all neutrons below 0.5 eV from the neutron fluence. This spectrum modification can be accounted for to give yet another integral quantity. The primary concern in choosing which foils to use is to ensure that the entire energy range of the spectrum is being measured as uniformly as possible. This is determined by the cross-sections of the reactions being investigated. If none of the cross-sections have any low-energy component, then the low-energy range of the spectrum may be adjusted or unfolded arbitrarily because it is having no effect on calculating the integral quantities based on the adjusted spectrum. Indeed, this is the one benefit that Bonner spheres will have over NAA; the energy range of the spectrum can be covered uniformly by choosing the right radii for the spheres.

The problem with trying to infer the energy spectrum from experimental measurements can be stated quite clearly at this point; each measurement is an integral quantity and the differential fluence in each energy bin of the spectrum is a variable. Approximating the integral in equation 1.1 as

$$I = \int_0^\infty \phi(E) R(E) dE \approx \sum_{i=1}^{n} \phi_i \, R_i \, \Delta E_i \tag{1.2}$$

where n is the number of energy groups used to represent the spectrum, ϕ_i is the differential neutron fluence in group i, R_i is the response function in group i, and ΔE_i is the energy bin width of group i, it is clear that each integral quantity provides a linear algebraic equation for the unknowns ϕ_i if the R_i's and ΔE_i's are known. The problem is that the number of activation foils or detector responses is likely an order of magnitude less than the number of energy groups necessary to obtain the resolution desired, leading to a highly under-determined system of linear algebraic equations. The goal of spectrum adjustment is to obtain a spectrum with higher resolution than is possible with the limited amount of experimental data available.

Sources of uncertainty

If we take a closer look at equation 1.2, a few complications arise even for simple low-resolution spectrum unfolding where the number of energy groups equals the number of integral quantities and the system is not under-determined. The integral quantity I that forms the non-homogeneous part of each linear algebraic equation would be an experimentally measured quantity, and hence would be prone to measurement uncertainty. The next thing to note is that the response functions, which are microscopic reaction cross-sections if the integral quantities are reaction probabilities as in NAA, are likely to have uncertainties associated with them as well. The problem of inferring the energy spectrum from experimental data is thus not as simple as solving a system of linear equations, if we are to know anything about the uncertainties in the resulting spectrum.

Instead of measuring the neutron energy spectrum, we could try to calculate it. Calculating the energy spectrum of neutrons at a given position in a reactor is a relatively simple task given the capabilities of modern Monte Carlo transport codes such as MCNP. The problem of course is that even if the transport code modelled the physics perfectly, other parameters such as material compositions, densities, and temperatures must be provided as input, and could never be known with perfect accuracy. Each of these can have a drastic impact on the energy spectrum of neutrons in a reactor. In addition to these input parameters, the geometry of the problem and the transport cross-sections must be provided, each of which is likely to have significant uncertainties. While the user could painstakingly model the geometry to the finest details, the errors in the transport cross-sections are unavoidable. Therefore, if we assume that the transport code implements the physics of neutral particle transport perfectly, all data provided as input to the code must be questioned at some level, and all of these inputs are sources of uncertainty in the calculated spectrum.

The process of spectrum adjustment, which will be discussed in the next section, makes use of a calculated spectrum. It uses this calculated spectrum along with the experimental measurements and reaction cross-sections to arrive at an adjusted spectrum that ideally contains less uncertainty than the calculated spectrum, and when folded with the reaction cross-sections matches the measured reaction probabilities more closely than the calculated spectrum. Therefore, the sources of uncertainty can be split into two categories; calculation uncertainties and adjustment uncertainties. Calculation uncertainties are the uncertainties in the model used to provide the calculated spectrum. Adjustment uncertainties are the uncertainties in the experimental measurements and reaction cross-sections that are used to adjust the calculated spectrum.

Of these two categories, the calculation uncertainties are by far the most difficult to quantify given the number of uncertain parameters and the time required to calculate the spectrum with negligible statistical error given a single set of sampled input parameters. However, these uncertainties must be quantified if the uncertainties in the adjusted spectrum, from which it is derived, are to be quantified. GenSpec in its current state does not attempt to quantify the uncertainties in the adjusted spectrum, and hence does not require uncertainties in the calculated spectrum as input. These capabilities are currently being implemented and will be presented in future work. It is recognized by the authors that without this capability, the code delivers incomplete results.

Spectrum adjustment

The process of obtaining a spectrum using only the integral quantities available in the form of activation foil measurements or detector responses is known as spectrum "unfolding." As mentioned in the previous section, the resolution of the spectrum obtained through this method is directly related to the number of integral quantities available. We could do better by providing more information in forms other than more integral quantities. This additional information comes in the form of an educated guess at the spectrum. Of course, with the capabilities of modern Monte Carlo neutron transport codes, this is likely to be a very good guess if the reactor can be modelled accurately. Thus, where spectrum unfolding aims to reconstruct the spectrum based on experimental data only, spectrum adjustment aims to adjust a "guessed" spectrum that is supplied via computational neutron transport calculations.

The purposes of this adjustment are two-fold. The first is to bring the reaction probabilities calculated using the spectrum into better agreement with the measured reaction probabilities. This is done because we acknowledge that the calculated spectrum is likely close to the true spectrum, yet ultimately incorrect; however the experimental data does not allow evaluation of the spectrum with the desired resolution. It is therefore necessary to use all information available to arrive at the "most-likely spectrum," even though the problem may mathematically have infinitely many solutions due to its under-determined nature. The second purpose is to produce an adjusted spectrum which contains smaller uncertainties than those associated with the calculated spectrum. After all, if the adjusted spectrum is more uncertain than the calculated spectrum, nothing has really been accomplished. In this case we would have been better off using the calculated spectrum to compute relevant quantities such as absorbed dose during an irradiation experiment.

Although there are many methods in existence today for spectrum adjustment, two of the most common methods are iterative perturbation and statistical least-squares estimates. These two methods differ in their priorities; iterative perturbation is more concerned with matching measured and calculated data, whereas statistical least-squares estimates are primarily concerned with minimizing the uncertainties in the adjusted spectrum. In this sense, the genetic algorithm presented here more closely resembles iterative perturbation techniques. Both of these methods however suffer from some unfortunate side effects such as the production of unrealistic spectral shape artefacts, and in the case of least-squares estimates, a significant deterioration of the method for high-resolution spectra. These known issues remain the driving force behind the development of new methods. The

unfortunate fact remains that spectrum adjustment is an under-determined problem with infinitely many solutions and hence no method can claim to produce the correct spectrum as output.

The term "unrealistic spectral shape artefacts" refers to features of the adjusted spectrum such as wild dips or peaks, that have no scientific explanation and exist only as a result of the adjustment method. This is a possible side-effect in nearly every method for spectrum adjustment in existence today. This side-effect is also one of the main arguments against focusing only on matching measured and calculated data without considering uncertainties. It is possible, through iterative perturbation techniques, to arrive at a wholly unrealistic spectrum consisting of nothing but jagged spikes throughout the entire energy range which happens to predict the measured reaction probabilities with extremely high accuracy[3]. This is particularly the case when the number of iterations is too high, and the calculated spectrum has been perturbed so many times that it no longer resembles its initial form at all. Unfortunately, least-squares methods are also prone to these unrealistic artefacts as can be seen in Figure 1.2 which shows the least-squares adjustment of the spectrum in the PLG bucket environment in the ACRR central cavity. This adjustment was performed with the code LSL-M2 which uses a logarithmic least-squares adjustment.

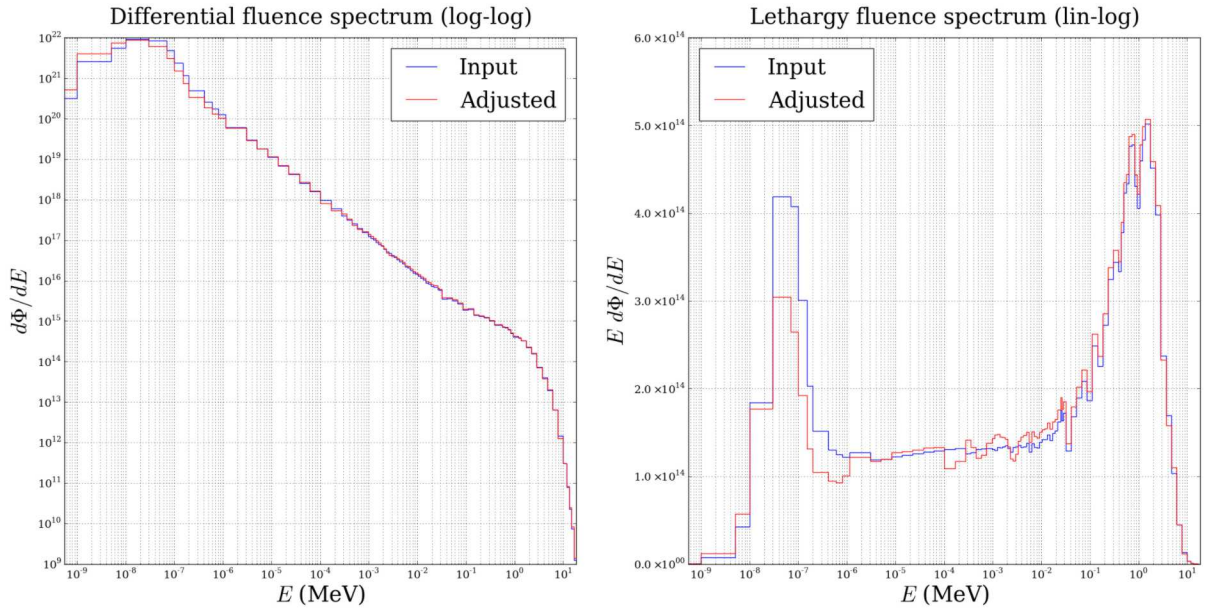

Figure 1.2: LSL-M2 adjustment of the PLG spectrum in the central cavity of the ACRR.

The primary downfall to least-squares methods is their sensitivity to high-resolution spectra. This is a consequence of the purely mathematical solution method which aims to solve the actual under-determined problem by adding more variables and assuming that all input parameters are uncertain and should thus be adjusted in combination with the input spectrum. This requires covariance matrices for the input spectrum and dosimetry cross-sections, that simply are not known with the desired resolution. Re-binning this low-resolution data often leads to singular covariance matrices which renders the resulting equations unsolvable[4]. Typically this will necessitate the use of grids

with fewer than 100 energy groups. For reference, the original iterative perturbation code SAND-II[5] used 621 energy groups, which is more consistent with the resolution desired to capture features of reactor spectra. This is shown in Figure 1.3 which shows the spectrum in the PLG bucket environment in the ACRR central cavity as calculated using MCNP with 89 and 640 energy groups.

Perhaps the biggest pitfall to any method that focuses on the reduction of the uncertainties of the adjusted spectrum is the assumption that the covariance matrix for the calculated spectrum is known with any confidence at all. The fact remains that the uncertainty in the adjusted spectrum must be directly related to the uncertainty in the input spectrum, even though the LSL-M2 user manual claims that "high accuracies [in the input spectrum uncertainties] are not necessary since the results of the adjustment are not very sensitive to changes in the input uncertainties so that relatively crude approximations may be sufficient."[6] This statement is only partially true in that the input spectrum covariance matrix may not significantly effect the uncertainties in the adjusted spectrum, however it will significantly effect the shape of the adjusted spectrum. This is because the input spectrum covariance matrix is used to constrain the adjusted spectrum in conjunction with the covariance matrix for the underlying reaction cross-sections for the activation foils. Thus, the existence and severity of the unrealistic spectral shape artefacts will be directly dependent upon how tightly constrained the adjusted spectrum is by the input spectrum's covariance matrix.

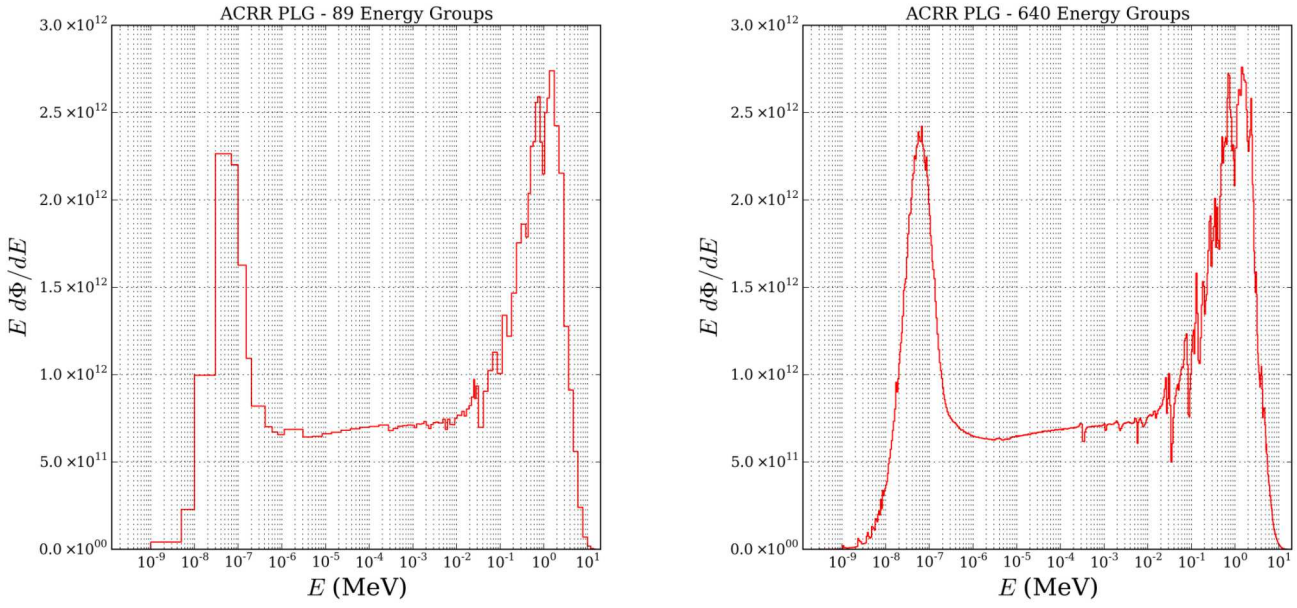

Figure 1.3: Spectrum in the central cavity of the ACRR with the PLG bucket using 89 and 640 energy groups.

Goals for GenSpec

With the discussion above, we can clearly state the reasons for the development of a new method, along with the properties that will be desired. When iterative perturbation techniques were first introduced to match calculated to experimental data, spectral artefacts were an advantage because the input spectra at that time lacked detail such as resonance spikes due to the inability of transport codes to resolve them and the high uncertainty in the transport cross-sections. Iterative perturbation techniques were thus used to introduce detail into a highly approximate input spectrum using reaction cross-sections that were known with much more accuracy. Computational resources now allow the calculation of high-resolution energy spectra and the transport cross-sections used are known with much more accuracy than they were in 1967. Thus, what began as a strength for iterative perturbation techniques to introduce detail that calculation could not provide, is now a downfall, introducing detail in many cases where none should exist. Therefore the first property that is desired of a new method is the preservation of trial spectrum features without the introduction of new ones that have no scientific explanation. In short, this means a smooth adjustment that will reveal if a large portion of the spectrum is overestimated or underestimated without attempting to drastically change the shape of the input spectrum.

The second property that is desired is either a small dependence, or complete independence on the resolution of the input spectrum. Ideally, any new method would retain its accuracy for as high a resolution as possible. It should be noted that there is simply not enough information available in spectrum adjustment problems to justify adjusting a spectrum in a very small energy window independently of the surrounding regions. Thus, the adjustment for a spectrum with high-resolution should be very similar to an adjustment of the same spectrum with a low-resolution. For instance, if the adjustment of a low-resolution spectrum reveals that the input spectrum grossly over-predicted the fluence in the low-energy region, and under-predicted the fluence in the high energy region, this same result should be predicted for the same input spectrum with a higher resolution. In addition, the resolution of the spectrum should not significantly effect the uncertainties in the adjusted spectrum if they are calculated.

The third property that is desired of a new method is that it should produce a spectrum that predicts the measured reaction probabilities at least as well as other existing methods. This comparison to other methods is more important when comparing to least-squares methods since iterative perturbation techniques can produce spectra that are able to predict measurements with arbitrarily high precision, even though they bear no resemblance to a physically realizable spectrum. Thus any new method should compare its precision in predicting measured reaction probabilities to least-squares codes such as STAY-SL or LSL-M2. It should also be noted that results of a new method in this sense should be comparable, and not necessarily better, as closeness to measured data cannot be the only metric judging the success of an adjustment. It may be the most easily quantifiable, yet other properties such as the smoothness of adjustment and the reduction in uncertainties should be taken into account as well.

The final property that is desired is the reduction in uncertainties from the those associated with the input spectrum. Even if the quantification of these uncertainties cannot be performed with high confidence, the adjustment method should act to reduce the variances of the fluence in each

energy group from those associated with the input spectrum. Unless the method is based on the minimization of these uncertainties, it may not be obvious that this property exists. GenSpec, in particular, currently does not have any uncertainty quantification built into the code, yet a detailed examination of the uncertainties of the adjusted spectra will be released in future work. In particular, a parametric uncertainty quantification will be performed which determines the relative contributions to the variances in the adjusted spectrum from each of the sources of uncertainty previously mentioned in this chapter. This will give some insight as to when the input spectrum uncertainties will be important in the final results compared to the uncertainties in the data used to perform the adjustment.

Chapter 2

Theory and Implementation

This section will describe some basic theory about genetic algorithms, as well as the implementation of this algorithm in GenSpec. Any reader who desires to understand how the code works is encouraged to read this section, and those with knowledge of genetic algorithms are encouraged to skip to the implementation section.

Theory

A genetic algorithm is simply an optimization method, often found useful in logistics problems such as the number partitioning problem or the travelling salesman problem. GenSpec is a new spectrum adjustment code that uses a genetic algorithm to optimize functions for the relative adjustment factor to a calculated spectrum as a function of energy over the energy range of interest. This calculated spectrum will be referred to as the trial spectrum. It is the properties of these adjustment functions that are responsible for the smooth adjustment, and other optimization methods such as simulated annealing may be able to perform the same calculations with similar success.

The theory behind any genetic algorithm is relatively simple. The goal is to mimic natural selection to determine the "fittest" possible solution. The first step is to find a way to represent the solution to a problem of interest as a specimen defined by a genotype. This means that solutions should be represented by a number of "genes," with each gene holding a specific value. Once this representation is defined, a population of specimen solutions must be formed either randomly or, as in the case of GenSpec, using some *a priori* information to influence the formation of each initial specimen. The number of specimens in the population is fixed throughout the algorithm.

The next step is to define the fitness function for any specimen solution in the population. Ideally, this would result in a single numerical metric defining how "good" the solution is. For example, in a simple travelling salesman problem where the goal is to minimize the distance travelled while passing through all cities on a map only once, the metric is easily defined as the distance travelled by using a particular route, with each such route represented as a specimen in a genetic algorithm. As optimization problems become more complex, the fitness functions must be more carefully defined, and could measure anything from the smoothness of a function to the difference between calculated and measured data.

Once the solution to the problem of interest has been abstracted to a genotype and the fitness

function has been clearly identified, the genetic algorithm can begin. Genetic algorithms progress in stages called generations. Parents are selected from the population for mating and producing children in such a way that the children inherit qualities of both parents' genotypes. The selection and mating process is continued until the number of children produced equals the initial population size, and these children form the next generation. The keys to success for a genetic algorithm are the selection and mating algorithms. In order to optimize the solution, higher fitness specimens should be chosen for mating more frequently than lower fitness specimens, and the mating algorithm should be designed such that children inherit the qualities of their parents that made them high fitness specimens in the first place.

If a genetic algorithm is successful the minimum, maximum, and average fitnesses of all of the specimens in each generation should increase at a similar rate, and eventually level out indicating convergence. Once the minimum, maximum, and average fitnesses cease to increase appreciably, convergence has been achieved and the specimen with the highest fitness encountered throughout all generations is stored as the optimized solution. One drawback of genetic algorithms is the possibility of convergence to a sub-optimal solution due to several factors which will not be explained in full detail here. It is sufficient to say that this side-effect is a result of insufficient exploration of the solution space. There are two ways to combat this issue, the first being to increase the size of the population. The second is gene-wise mutation.

Gene-wise mutation is the process of altering a specimen's genotype as it is created in the mating process. The idea is to perform this alteration at a pre-defined rate known as the mutation rate, so that each gene inherited by the child in the mating process has a distinct probability of undergoing a mutation. The effect of mutation is a driving force away from convergence and towards solutions that may not yet have been explored. Hence, a high mutation rate along with an effective mutation algorithm ensures that the solution space is being thoroughly explored and reduces the likelihood of convergence to a sub-optimal solution. Of course, there is a balancing act that must be performed since convergence is necessary for completion of the genetic algorithm, and a mutation rate that is too high will not allow for convergence at all.

In truth, the theory behind the genetic algorithm is much more complex than what is presented here. The more interested reader is advised to consult the book by M. D. Vose[7] where the theory and implementation is described with much more detail and mathematical rigour.

Implementation

One of the appealing aspects of genetic algorithms is their flexibility. This flexibility lies in the design of the processes for selection, mating, and mutation. In this section, we describe how each aspect of a genetic algorithm mentioned in the previous section is implemented for the purposes of spectrum adjustment in GenSpec. It should be noted that each process could be performed many different ways, and in most cases the simplest and most trusted methods have been adopted.

Problem abstraction

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 can 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. The number of gene sites N, is chosen by the user, however best practices will be discussed later in this report. 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 energy region of the gene site that it resides at. For instance, a gene value of 1.1 indicates that the specimen should have a 10% increase above the trial spectrum in an energy window around the gene site location. The size of this energy window depends upon the number of gene sites chosen and how close together they are.

The gene values must then be used to define each specimen spectrum. This is done by performing a least-squares polynomial fit through the gene values, although other methods such as cubic spline interpolation are currently being investigated. These polynomials will be called the shift functions. The fluence 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 artefacts in the spectral shape, or at least provides a means of controlling and limiting them. With this abstraction clearly defined, we can begin to describe the various processes that make up the genetic algorithm.

Before discussing the processes that make up the genetic algorithm in GenSpec, it is interesting to note that the abstraction described above effectively decouples the adjustment process from the energy grid structure of the trial spectrum. This decoupling, as will be shown in the Results section, is responsible for the independence of the adjustment to the trial spectrum energy grid.

This is no small feat, as one drawback of existing spectrum adjustment methods is that they are limited to rather coarse energy grid structures to perform a successful adjustment. The reason for this is simple; adding more energy groups without adding more equations in the form of integral quantities makes the problem of spectrum adjustment even more under-determined. This makes any adjustment method based on purely mathematical foundations less useful if a high-resolution spectrum is desired.

Setting the population

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, however larger population sizes are desired. The limiting factor on the population size will likely be system memory and runtime. At this point, the formation of the specimens that make up the original population can be started. Two options are available for specimen formation. The first is random formation, and the second is variance guided formation. In both cases, each specimen's gene values will be made up of adjustment factors that are randomly perturbed from unity. Variance guided formation is not currently implemented, but will be in future releases. In variance guided formation, the magnitude of the perturbation from unity is influenced by the variance of the trial spectrum fluence in the energy group that the gene site resides in. This variance must be supplied by the user in the form of a covariance matrix for the trial spectrum. Admittedly, finding this covariance matrix is not an easy task, however many spectrum adjustment techniques require such a covariance matrix to perform the adjustment. This is certainly true of any method that uses least-squares.

As GenSpec exists to this date, random formation is used for initializing the gene value of each gene of each specimen in the population. 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. The value of 0.7 is hard-coded in GenSpec and is found to produce initial populations with significant differences from the trial spectrum. Further deviations from the trial spectrum are achieved through mutation. 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 in the problem abstraction.

Calculating the fitness

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 most obvious trait that can be quantified numerically is the agreement between the reaction probabilities calculated using the specimen and appropriate reaction cross-sections, and the measured reaction probabilities. The fitness function is then defined as

$$f = C - \sum_{i=1}^{m} \frac{\left| \left\{ \sum_{j=1}^{n} \sigma_{j,i} \Phi_j \right\} - r_i \right|}{r_i} \tag{2.1}$$

where m is the number of foils used in the NAA, n is the number of energy groups, $\sigma_{j,i}$ is the reaction cross-section for foil/reaction i in energy group j, Φ_j is the total fluence in energy group j of the specimen spectrum, 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.

This fitness function only measures the closeness of the reaction probabilities calculated using the specimen to those measured using NAA. Admittedly, closeness to measured values cannot be the only thing that measures the success of an adjustment. This is because iterative perturbation techniques have shown that spectra can be obtained which match the measured data perfectly, yet are so far from realistic that they must be disregarded. These spectra exhibit wild dips and peaks throughout the energy range of interest with no resemblance to the trial spectrum at all. The goal is then to achieve the highest fitness possible while also preserving attributes of the trail spectrum that are known to exist such as resonance spikes, fission spectra, and a low-energy Maxwellian distribution. This goal is achieved through the use of low order polynomials for the shift function, which constrain the smoothness of the adjustment.

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. GenSpec uses what is known as proportional selection. With this selection method, the probability of specimen j to be chosen for mating is

$$P_j = \frac{f_j - f_{min}}{P_t} \tag{2.2}$$

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

$$P_t = \sum_{i=1}^{s} (f_i - f_{min}) \tag{2.3}$$

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.

Mating

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 2.1 which shows the gene values and shift functions before and after mating. In this figure, the cutting point is shown as a dark vertical line. Other methods were created and implemented with no clear advantage over single point crossover, which is by far the simplest and most widely used mating process for genetic algorithms.

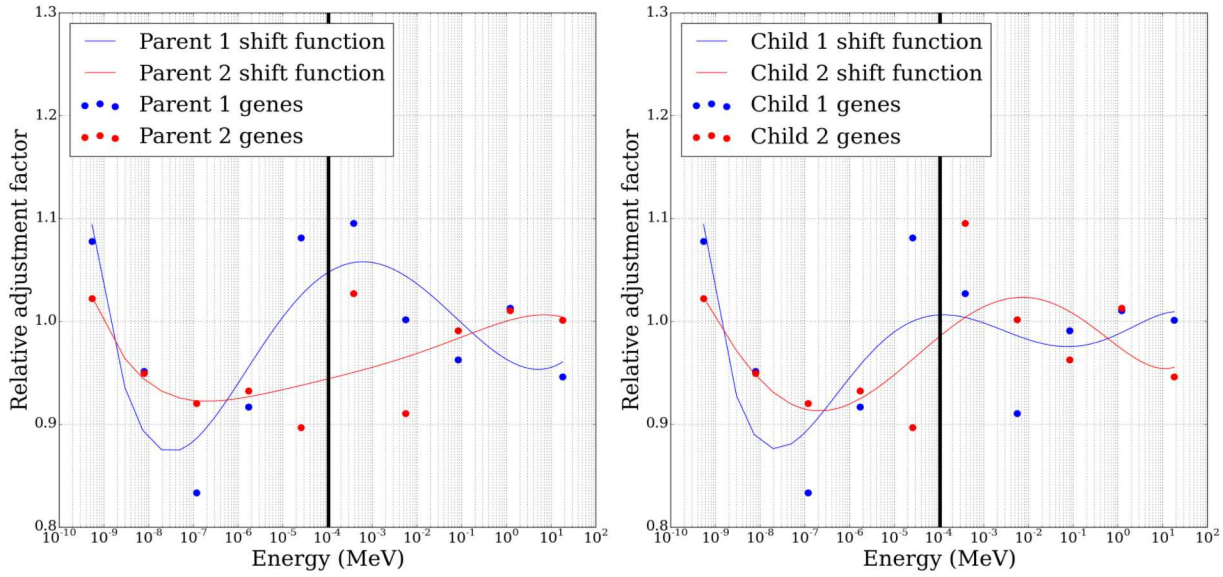

Figure 2.1: Gene values and shift functions of the parents and offspring.

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 genes inherited by the children will undergo some sort of random mutation. Mutations in the algorithm presented here correspond to 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 GenSpec. It should also be noted that mutation is the only way to arrive at gene values that were not attained by any of the parents in the initial population. Hence, mutation is most critical when the population size is small. Figure 2.2 summarizes the steps taken in a typical genetic algorithm, and the sections above describe how each step is implemented in GenSpec.

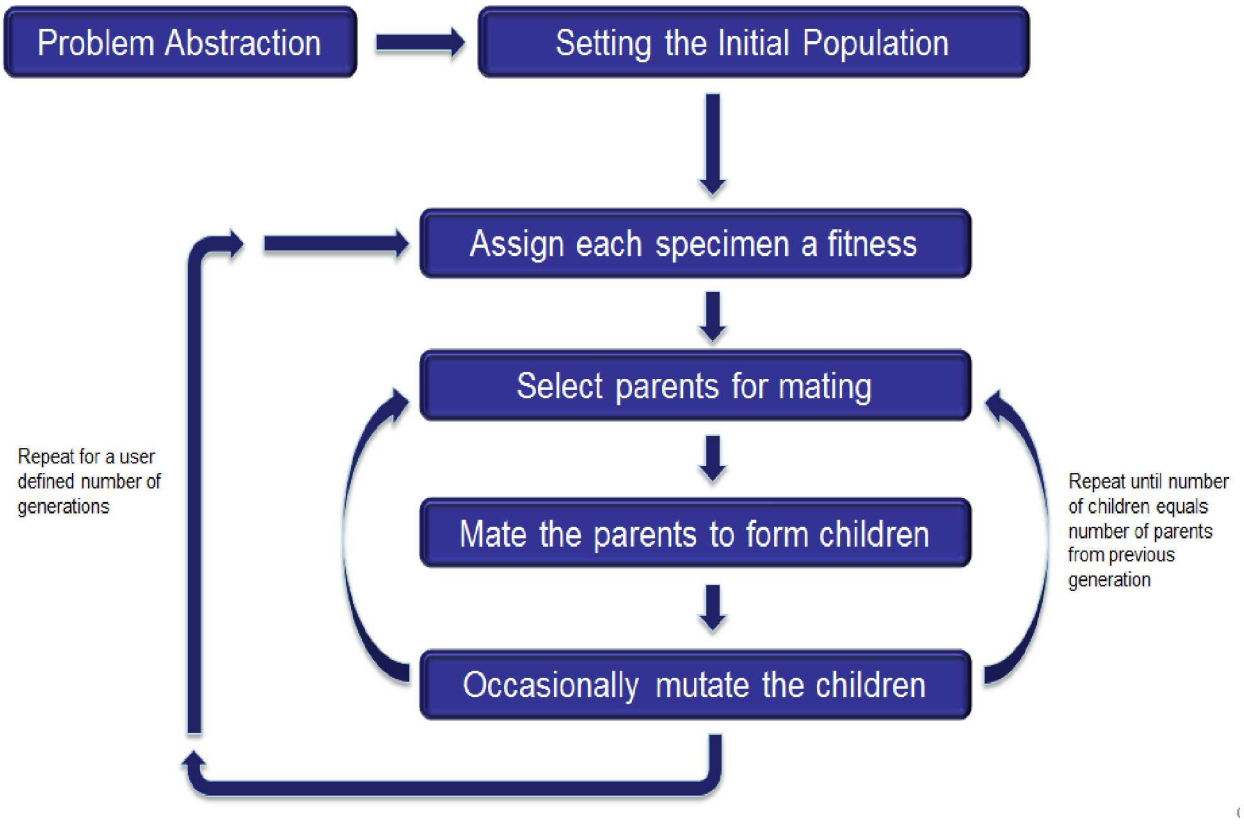

Figure 2.2: Diagram of the steps taken in a typical genetic algorithm.

Chapter 3

Usage

This chapter serves as the user's manual for GenSpec. The code GenSpec was written with user friendliness as a high priority. As such, the code is run through a graphical user interface (GUI), and is currently compatible with Windows, although Mac/Linux versions should follow if there is significant demand. This is because an easy-to-use code leads to more users, which leads to more testing, which finally leads to feedback regarding corrections and improvements. It is hoped that the effort put forth in making the GUI will not go unnoticed or unappreciated. The GUI is written using Python, and hence a working version of Python must be installed on the machine, along with the modules Numpy, Matplotlib, and TraitsUI. Also, a working version of Perl must be installed on the machine as well.

Installing GenSpec is as simple as obtaining the compressed folder from the authors and extracting it to the `C:\` drive. After extraction, the folder `C:\GenSpec` should exist, and should contain all data, source code, and compiled executables required to run the code. In addition, there should be a sample case named `PLG_SAMPLE` which will be discussed in this chapter. Finally, it is convenient to make a desktop shortcut to run GenSpec from the desktop. To do this, navigate to the folder `C:\GenSpec\gui` and right click on the file `GenSpec.bat` and choose Send to $\rightarrow$ Desktop. In the same folder, there is a desktop icon that can be used in place of the standard icon. To replace it, right click on the newly created shortcut and select properties. Then go to the Shortcut tab and select Change icon... at the bottom. Select browse and find `DNA.ico` in the `C:\GenSpec\gui` directory.

Case setup

The directory C:\GenSpec contains the following sub-directories:

- `Act`: contains the source code for the pre-processing code Act which was originally part of the LSL-M2 code package.
- `bin`: contains all compiled executables.
- `Calact`: contains the source code for the pre-processing code Calact which was originally part of the LSL-M2 code package.

- `fluencePRO`: contains the source code for the pre-processing code fluencePRO which was originally part of the LSL-M2 code package.
- `GenSpec`: contains the compilation scripts for GenSpec.
- `gui`: contains the Python files and Perl scripts which make up the GUI.
- `input`: contains GenSpec input files.
- `library`: contains cross-section libraries, trial spectra, energy grid structures, and templates.
- `output`: contains the output files of GenSpec along with the Python script to plot relevant data.
- `self-shield`: an extensive self-shielding library for different cover materials and reactions.
- `src`: contains the source code for GenSpec.
- `work_dir`: contains various transfer and temporary files .

After successfully installing GenSpec and all of its dependencies, double clicking on the desktop icon should bring up the window shown in Figure 3.1. This window gives you the option to either open an existing case, or start a new one. For the purposes of this demonstration, we will choose an existing input file from the folder `C:\GenSpec\input` named `PLG_SAMPLE.pkl`, as shown in the figure. Once the file has been selected, or the New file box has been checked, click the continue button at the bottom of the window.

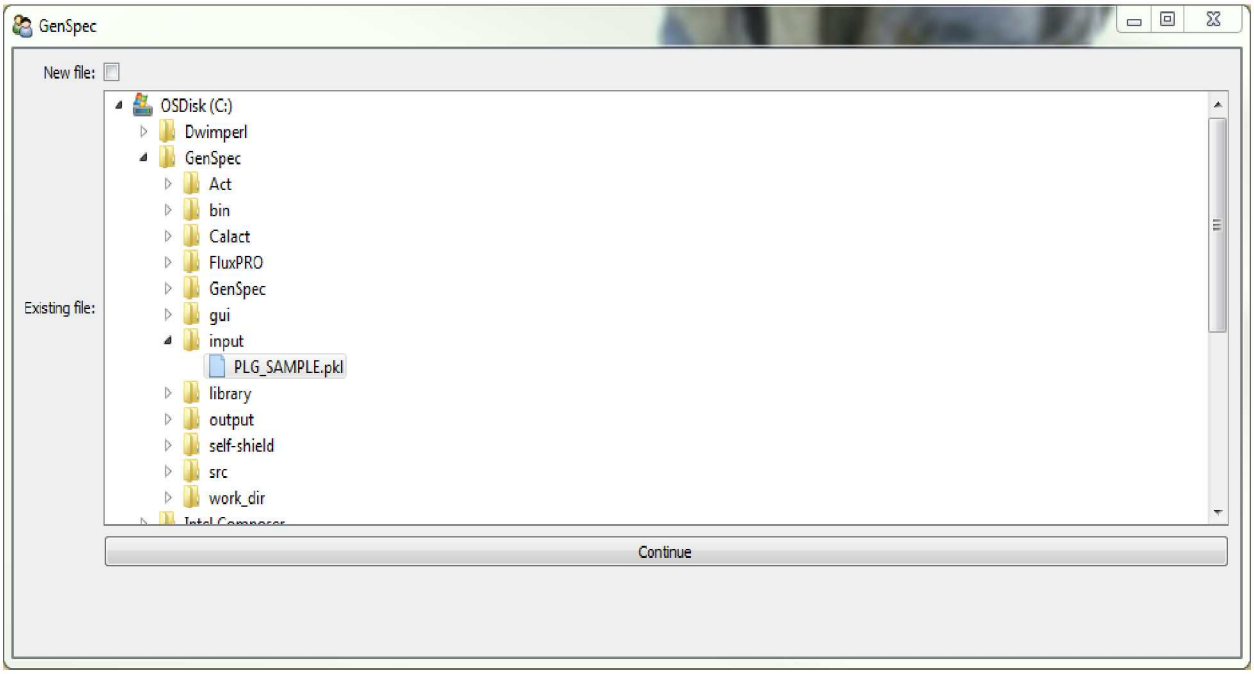

Figure 3.1: Opening window of the GUI.

After clicking continue, the main window of GenSpec should appear. This window has five separate tabs. The first of these tabs is shown in Figure 3.2, which is followed by a list describing each input field.

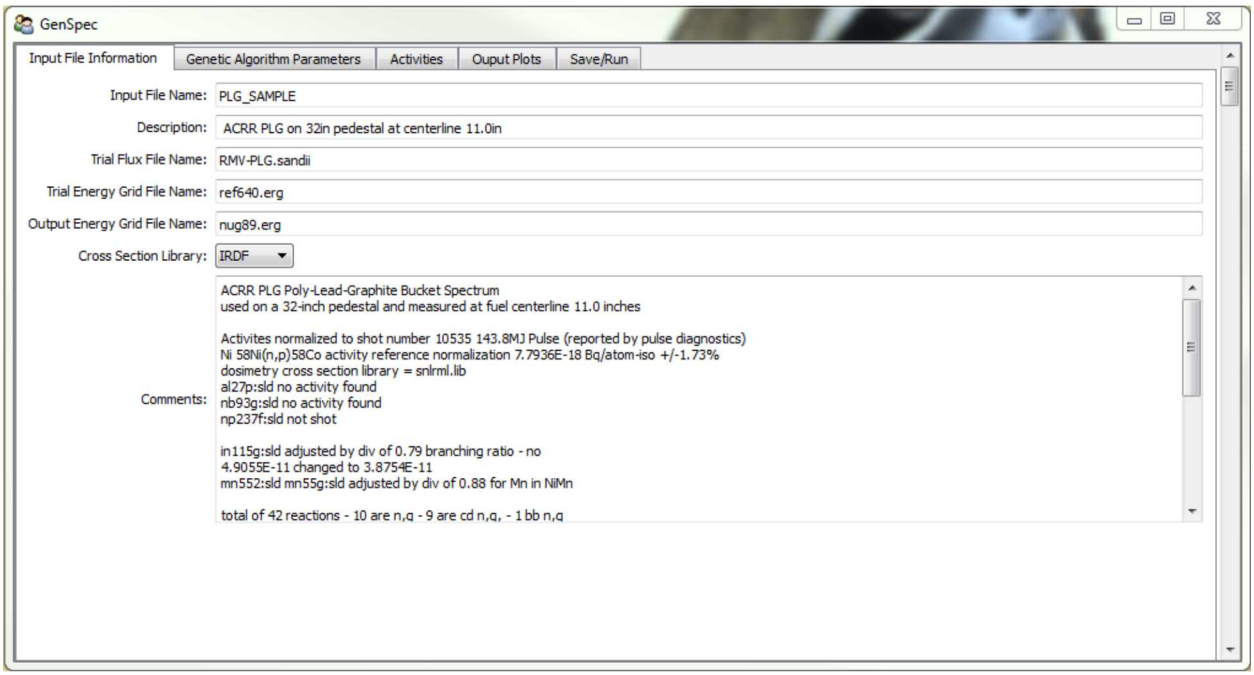

Figure 3.2: Input file information tab of the GUI.

- *Input File Name*: This is the name that GenSpec will save your file as when you save it using the Save/Run tab.
- *Description*: This is a short description of the case. It has no bearing on the calculation.
- *Trial fluence File Name*: This is the name of the file containing the fluence/fluence values of the trial spectrum. It should be located in the `library` folder.
- *Trial Energy Grid File Name*: This is the name of the file containing the energy bin boundaries of the group structure that the fluence values in the *Trial fluence File Name* correspond to. It should be located in the `library` folder.
- *Output Energy Grid File Name*: This is the name of the file containing the energy bin boundaries of the group structure that should be used for the adjustment. It should be located in the `library` folder.
- *cross-section Library*: The cross-section library used for calculating the reaction probabilities given the spectrum. Currently, the only two options are libraries made by the Sandia National Laboratories Radiation Metrology Laboratory (SNLRML) or the International Reactor Dosimetry File (IRDF) produced by the International Atomic Energy Agency (IAEA).

- *Comments*: This is where the user should write comments. These might include suspicions of measured reaction probabilities or dosimetry cross-section data as well as experimental details.

The format of the input files which should reside in the `C:\GenSpec\library` folder are important. Energy grid files which should have the extension `.erg` have a very simple format. The first line should contain a short comment and the second line should have the number of bin boundaries in the group structure. On the following lines, the energy bin boundaries should be listed in either increasing or decreasing order, with each energy value on a new line. For example, the 640 group SAND-II grid structure file should look like:

```
Energy bin boundaries of the SAND-II 640 group structure
641
1.00E-04
1.05E-04
1.10E-04
1.15E-04
1.20E-04
1.28E-04
⋮
⋮
2.00E+07
```

Notice that the energy values are in eV, and also that they are bin boundaries, hence the number of entries should be one more than the number of groups in the desired energy grid structure. The fluence files are formatted similarly, with two lines before listing fluence values. The first of these two lines should be `*differential fluence`. The second of these two lines should be a comment line. Notice, the number of entries is not printed as it was for the energy grid file, because it is expected that the number of entries is one less than the number of entries in the accompanying `.erg` file. For example, the input PLG spectrum file should look like:

```
*differential fluence
Diff fluence from MCNP for PLG ACRR environment
9.34388E+03
8.18653E+03
9.76051E+03
1.34731E+04
8.05078E+03
8.42625E+03
5.05863E+03
9.14799E+03
⋮
⋮
1.22950E+17
```

The first line of the fluence/fluence file tells the code that the entries listed are differential fluence/fluence values. For those unfamiliar with this terminology, this simply means that the net fluence in each energy group as determined by a simple fluence tally in MCNP, has been divided by the width of the energy bin corresponding to that fluence/fluence value. Other options for the first line do exist, however all testing thus far has been using input spectra as differential fluence/fluence values and it is recommended that this practice be followed. The normalization of the fluence values is arbitrary and will be adjusted according to the provided measured activities.

Finally, the choice of output energy grid is important. It is advised to make the trial spectrum energy grid and output energy grid the same, but it is understood that due to the time required to compute the spectrum, this may not be feasible in all cases. It is sometimes desirable to reduce the resolution of the trial spectrum prior to adjustment. This is the purpose of the output energy grid file. This change in group structure is performed by fluencePRO, a pre-processing code that is included as part of LSL-M2. Its primary purpose is to take a spectrum and cross-section library of arbitrary group structure, and condense or expand them to the desired output group structure. It is the output group structure that is used in the adjustment process. For instance, in the sample case `PLG_SAMPLE`, the trial spectrum is supplied with a high-resolution 640 group structure, however, the adjustment is performed using only 89 groups. It should also be noted that increasing the resolution by specifying a finer output group structure than the trial spectrum group structure has no meaning. Just because fluencePRO can add more groups to the trial grid structure does not mean that spectrum will magically gain resolution. Thus, fluencePRO should only be used to reduce the spectrum resolution. Thus, it is common practice to calculate all trial spectra with high-resolution, and use fluencePRO to reduce the resolution if necessary.

The second tab of the main window of GenSpec is the genetic algorithm parameters tab pictured in Figure 3.3. The entries here are relatively straightforward and correspond to the parameters described in Chapter 2. For completeness however, they are listed below as well:

- *Population Size*: This is the number of specimen spectra in the population. It is kept fixed throughout the optimization. Higher values reduce the likelihood of convergence to a sub-optimal solution, yet increase the calculation time. Values ≥ 200 are recommended.
- *Number of Generations*: The number of generations to run. GenSpec does not currently check for convergence and stop when convergence is achieved. Instead, the user tells the code how many generations should be simulated, and the user must check for convergence afterwards. This will be problem dependent, yet in all cases observed thus far, convergence is achieved in less than 1,000 generations.
- *Polynomial Order*: The order of the polynomial shift functions used to perform a least-squares fit to the specimen gene values. This should be less than 10 to avoid unrealistic spectral shape artefacts.
- *Number of Gene Sites*: The number of genes equally distributed in base 10 logarithmic space between the minimum and maximum energy values present in the spectrum being adjusted. Ideally, this number should be at least two or three times that of the polynomial order. Larger values can be supplied, but will increase the runtime.

- *Mutation Rate*: The probability that a gene inherited by the child spectrum from the parent spectrum will undergo a random mutation. This should be as high as possible while still allowing for convergence. This will be problem dependent, and will likely need to be found by trial and error. It is recommended to start this trial and error process at 0.15 and increase to just below the value that no longer allows for convergence.

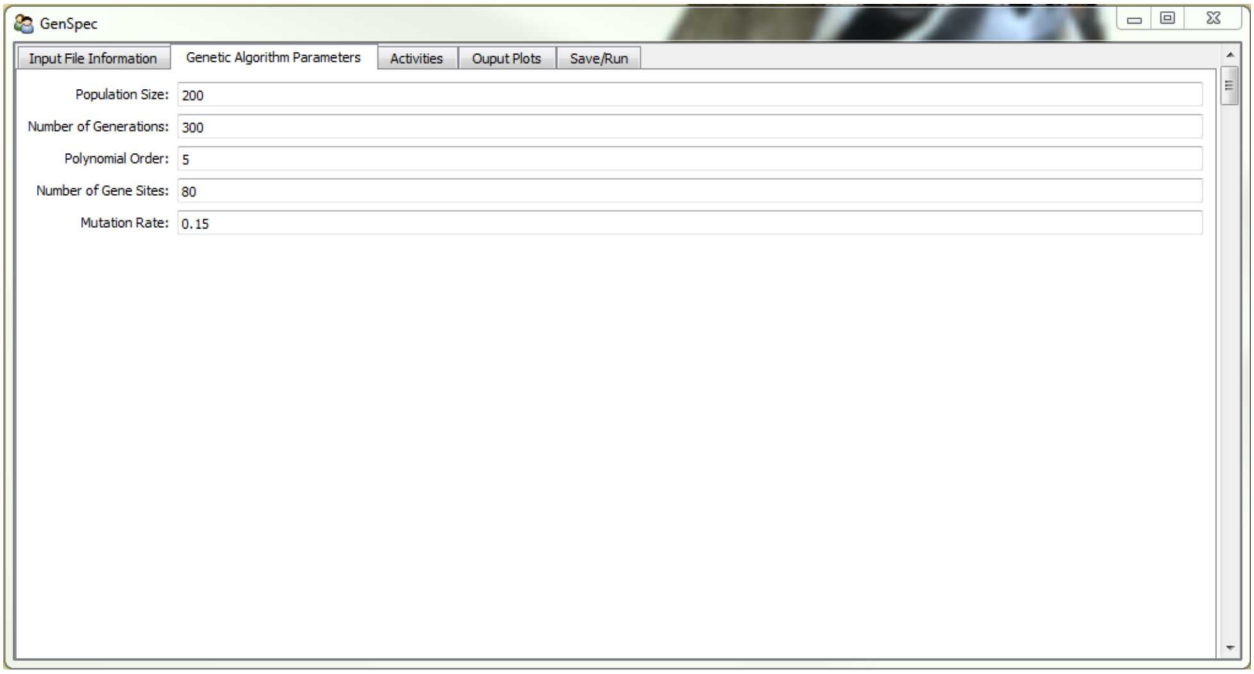

Figure 3.3: Genetic algorithm parameters tab of the GUI.

The next tab in the main window is the activities tab. This is where all experimental measurements are entered. For each reaction existing in the dosimetry cross-section libraries, there is a field in this tab. Multiple foils for the same reaction can be entered with different cover materials, which must exist in the `self-shield` folder. A detailed description of the covers and a discussion of how self-shielding effects are rigorously addressed in Sandia's use of activation measurements to characterize the spectra in our facilities can be found in the IEEE Transactions publication by P. J Griffin and J .G. Kelly[8]. The activities tab is shown in Figure 3.4, and the entries for each foil are listed below.

- *Number of Foils*: The number of different cover materials used with the particular reaction. This can be as high as four, although typically only one cover material is used leading to two entries: bare and covered.
- *Self-shielding/Cover N*: The identity string for the N^{th} cover material.
- *Activity N*: The activity measured for the reaction with the N^{th} cover material in units of Bq/atom of the target isotope. These activities should all correspond to an equivalent integral

reactor energy release. If the foils were not irradiated together, the reported activities should be adjusted so that they all correspond to the same net energy release by the reactor. Of course, the ideal situation would be for them all to be irradiated together, however this may not always be possible.

- *% STD N*: The relative standard deviation written as a decimal (i.e. 0.02 = 2%), for the activity reported previously for the reaction with the N^{th} cover material.
- *Normalize to reaction*: Normalization option for the reaction. This can be selected for only one reaction, and typically is selected for the best known or most trusted activity measurement.

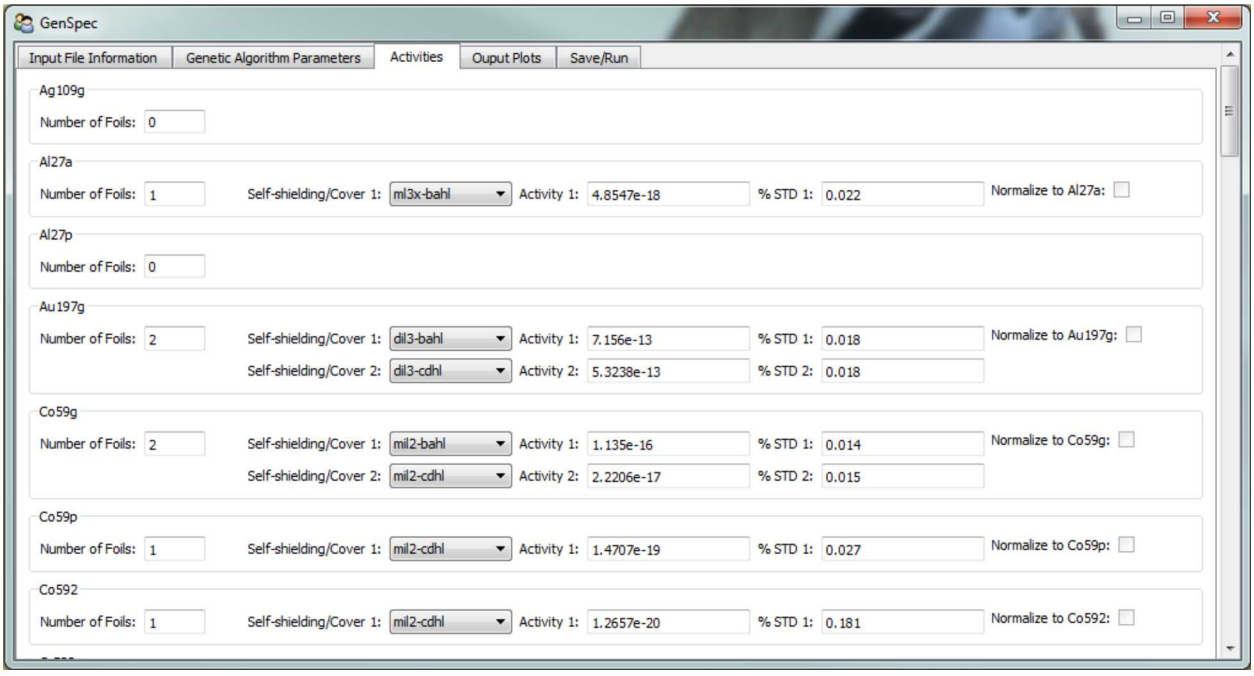

Figure 3.4: Activities tab of the GUI.

The final two tabs are relatively self-explanatory. The output plots tab shown in Figure 3.5 is simply used to choose which output is desired. The output plots that can be chosen here will be described in the next section. The save/run tab does exactly what its name implies, and is shown in Figure 3.6. The save button saves the entries committed thus far in a `.pkl` file which can be opened using the GUI at a later time. The name of the file is determined by the first entry on the first tab. There are then two options for running: run with or without plotting the results. Finally, the run button will start the calculation. Once the calculation is finished the requested plots will appear. The plots are built such that the aspect ratio and label size is ideal when the plot window is maximized on the screen.

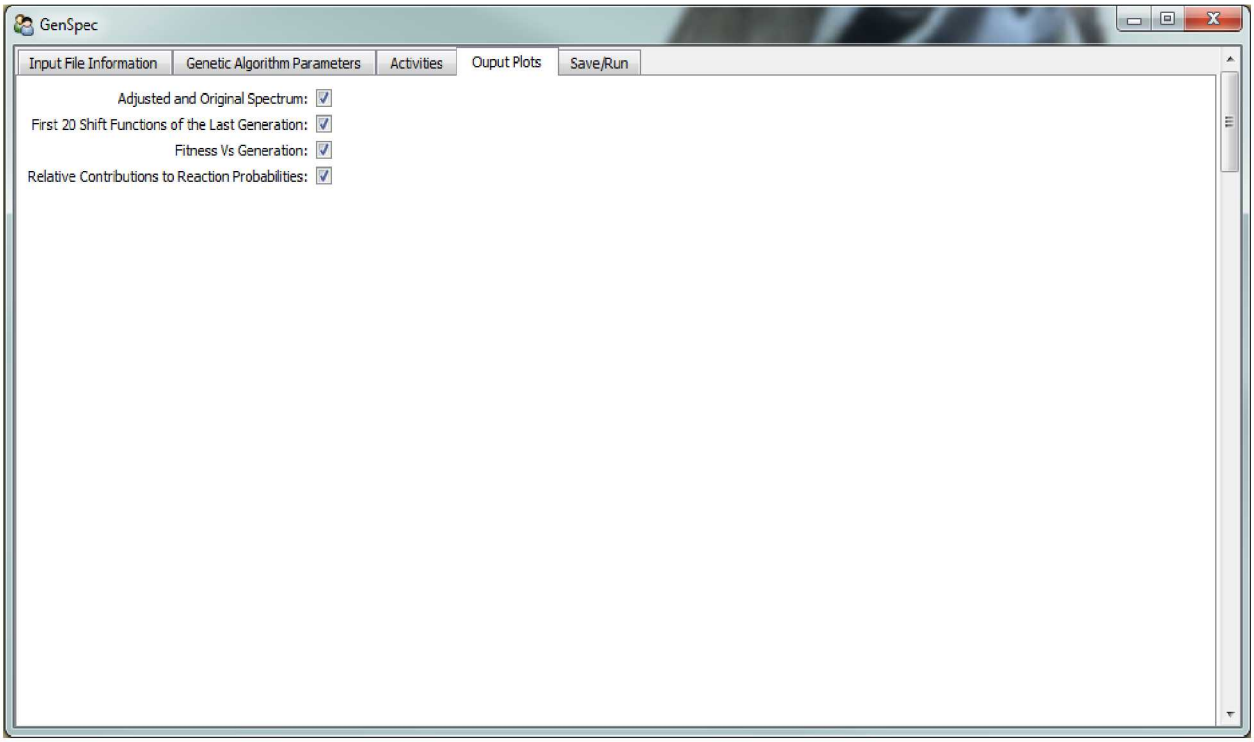

Figure 3.5: Output plots tab of the GUI.

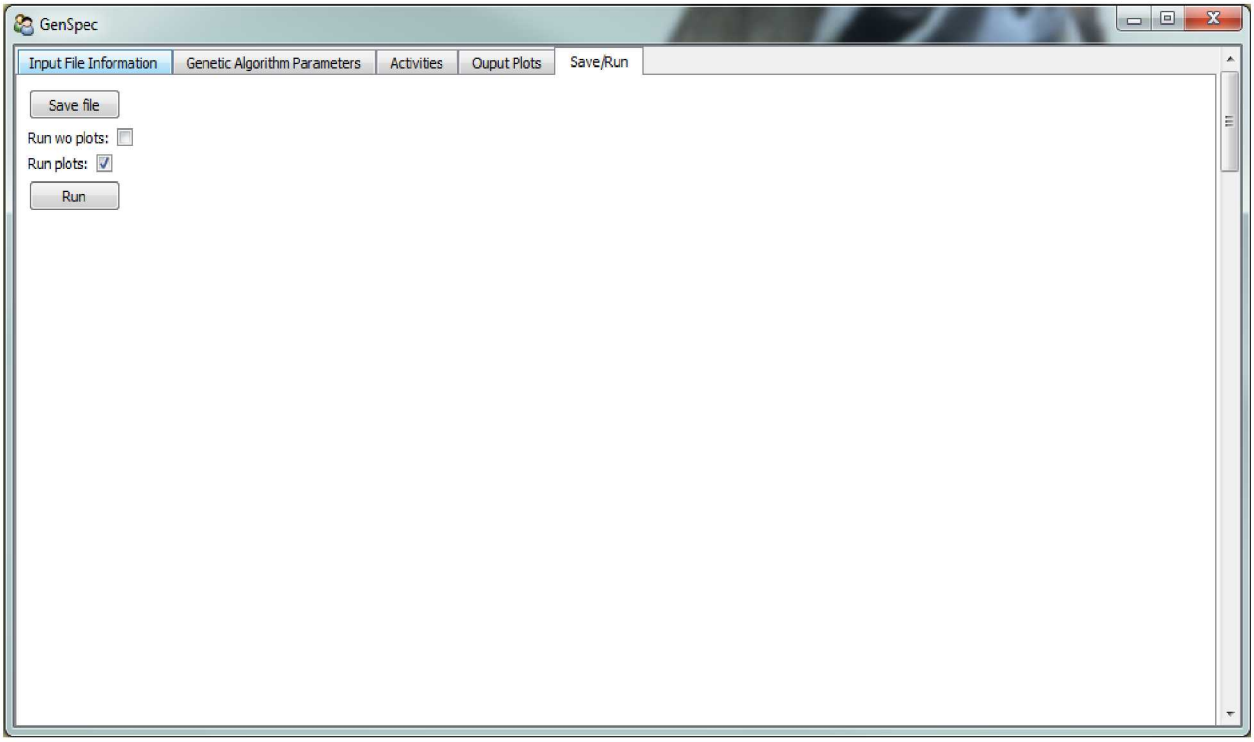

Figure 3.6: Save/run tab of the GUI.

Output description

At present, GenSpec offers four output plots, and a text output file containing all relevant output. The four plots that are offered in the output plots tab of the GUI include the adjusted spectrum compared to the trial spectrum, the first 20 shift functions in the last generation, the fitness as a function of generation number, and the relative contributions to the reaction probabilities as calculated using the adjusted spectrum. Each of these plots will be discussed in this section, and each one serves a different purpose. As the uncertainty quantification capabilities of GenSpec are implemented, the number of plots will likely increase and may include comparisons of trial and adjusted covariance matrices and plots of the variance in each energy group before and after adjustment.

The first plot, and likely the most important, is the plot of the adjusted spectrum compared to the trial spectrum. This plot can be seen in Figure 3.7 for the `PLG_SAMPLE` case. This plot shows both the differential and lethargy fluence spectra. The lethargy fluence is similar to the differential fluence spectrum discussed earlier, with the exception that the differential fluence in each energy group is multiplied by the midpoint energy of the energy group. The lethargy fluence tends to expose unrealistic spectral shape artefacts better than the differential representation, however, both are commonly found in literature, and hence both are produced here. The true benefit to producing the neutron lethargy fluence plot is that on a lin-log scale, equal areas under the curves correspond to equal total fluences. Thus with the lin-log plot of the lethargy fluence it is easy to approximate what percentage of neutrons lie in a given energy range.

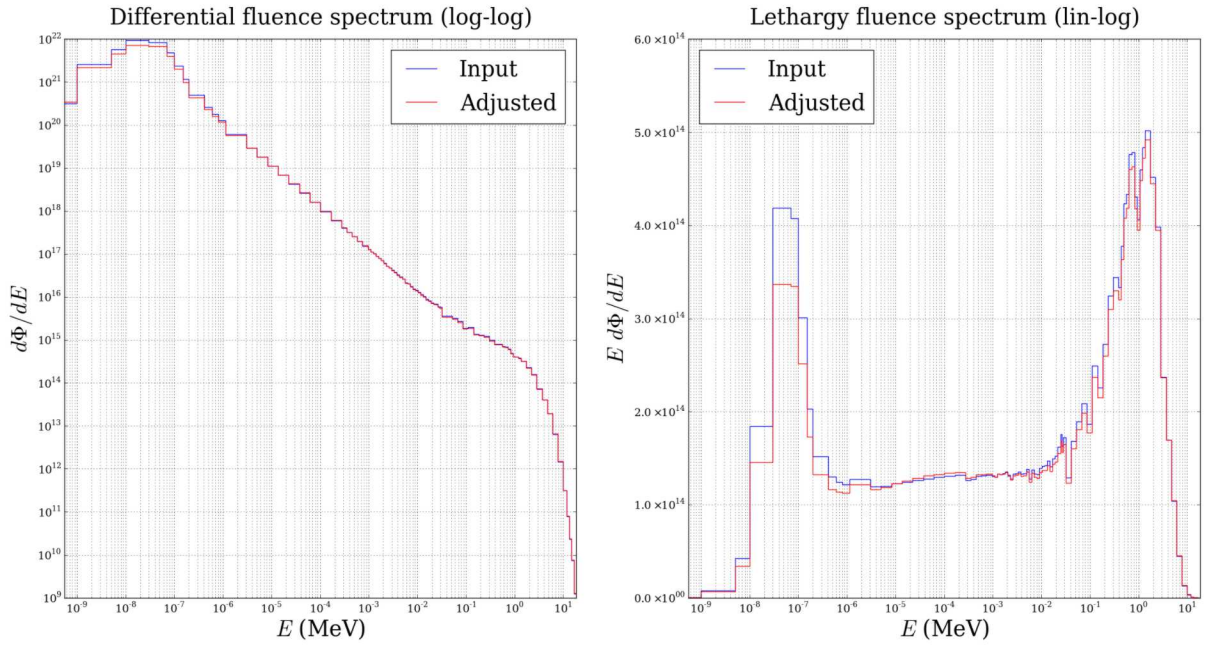

Figure 3.7: Comparison of the adjusted spectrum to the trial spectrum using differential and lethargy fluence representations.

The second figure produced is a plot of the first 20 shift functions of the final generation. Ideally, this should look like a bunch of similar polynomials plotted on the same canvas as shown in Figure 3.8. The purpose of this plot is to show the spread of specimen solutions in the final generation. It should be noted that the polynomials may diverge from each other in regions of low fluence where none of the foil cross-sections have a significant presence. For instance, if the fluence in a certain region is orders of magnitude below the fluence everywhere else, and no foil reaction has a high cross-section in this region, adjusting the spectrum by as high as 300% in this region only is unlikely to change the fitness of the specimen significantly. This is why the shift functions occasionally behave erratically at the minimum and maximum energies considered.

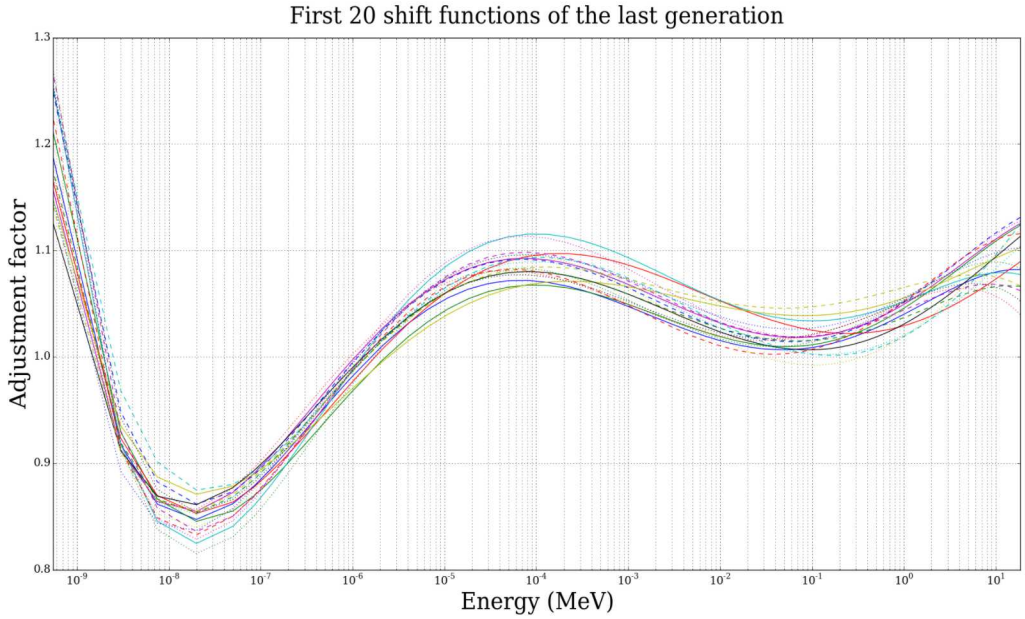

Figure 3.8: Plot of the first 20 shift functions of the population in the last generation.

The third figure produced is a plot of the minimum, average, and maximum encountered fitnesses as a function of generation. The purpose of this plot is to determine whether the method has converged and to look for strange behavior. The fitness plot for the `PLG_SAMPLE` case is shown in Figure 3.9, and shows the ideal behavior. In this plot, all three curves rise initially, and converge nearly at the same time. In this figure, we can see that convergence was actually achieved before 100 generations, and running the code for 300 generations was unnecessary. In most cases, the adjustment for 1,000 generations takes less than a few minutes, so the manual determination of convergence is not a significant hindrance. Even so, future versions of GenSpec will likely have a convergence checking routine whereby the optimization is stopped automatically after convergence is achieved. It should be noted that the minimum and average fitness curves are the minimum and average fitness of each generation. The maximum fitness encountered curve is slightly different in that is gives the maximum fitness found in all previous generations up to the generation on the abscissa. This is done because the specimen with the highest fitness is stored as the optimal solution

until a specimen with an even higher fitness is found. Thus, several generations may pass without the maximum fitness specimen being unseated.

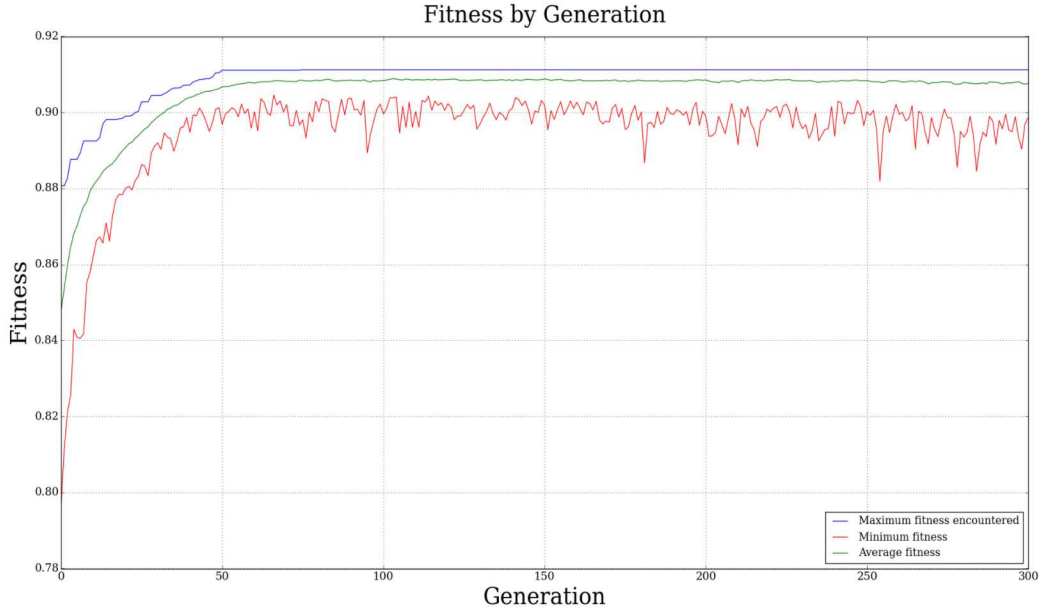

Figure 3.9: Minimum, average, and maximum encountered fitnesses as a function of generation number showing convergence.

The final figure produced is the relative contributions to the reaction probabilities for each reaction, and Figure 3.10 shows the results for the `PLG_SAMPLE` case. The purpose of this plot is twofold. The first, and most important is to show how uniformly the energy range of the spectrum is being weighted during the adjustment. The second is to see which foils, if any are in significant error with the measured quantities. Mathematically, the relative contribution for energy group j and foil k can be written as

$$R_{j,k} = \frac{\sigma_{j,k}\Phi_j}{\sum_{i=1}^{N} \sigma_{i,k}\Phi_i} \tag{3.1}$$

where $\sigma_{j,k}$ is the microscopic cross-section for energy group j and reaction k and Φ_j is the total fluence in energy group j. By plotting this relative contribution for each energy group and each reaction/cover combination, we can see which energy ranges are being weighted more heavily than others. Ideally, the entire energy range would be weighted equally, but this is simply not possible when using NAA as the measured data due to the chaotic nature of neutron reaction cross-sections. In particular, in Figure 3.10, it is clear that the energy range from roughly 0.01 to 1 MeV is being inadequately covered by the activation foils used. Possible remedies to this would be to find a

particular reaction that has a large cross-section in that region which could be feasibly measured using NAA, or to incorporate detector methods such as the use of Bonner spheres to sample that region specifically.

The second purpose of the figure is accomplished by the coloring of each curve. Curves with colors closer to cyan, which are most of them in the figure, have less error between measured and calculated reaction probabilities than those with colors closer to magenta. Ideally, all curves would have colors as close as possible to cyan, however as mentioned in previous chapters, the difference between measured and calculated data cannot be the only metric by which an adjustment is judged. The true purpose of coloring the curves according to error is to find outliers, such as mn552#-mil-2-bahl in Figure 3.10. If a single reaction is producing a much larger error than foils with similar energy coverage, this may be a sign of faulty measurement, or large errors in the cross-section data for that reaction. In this case, the foil should be irradiated again to check for experimental errors.

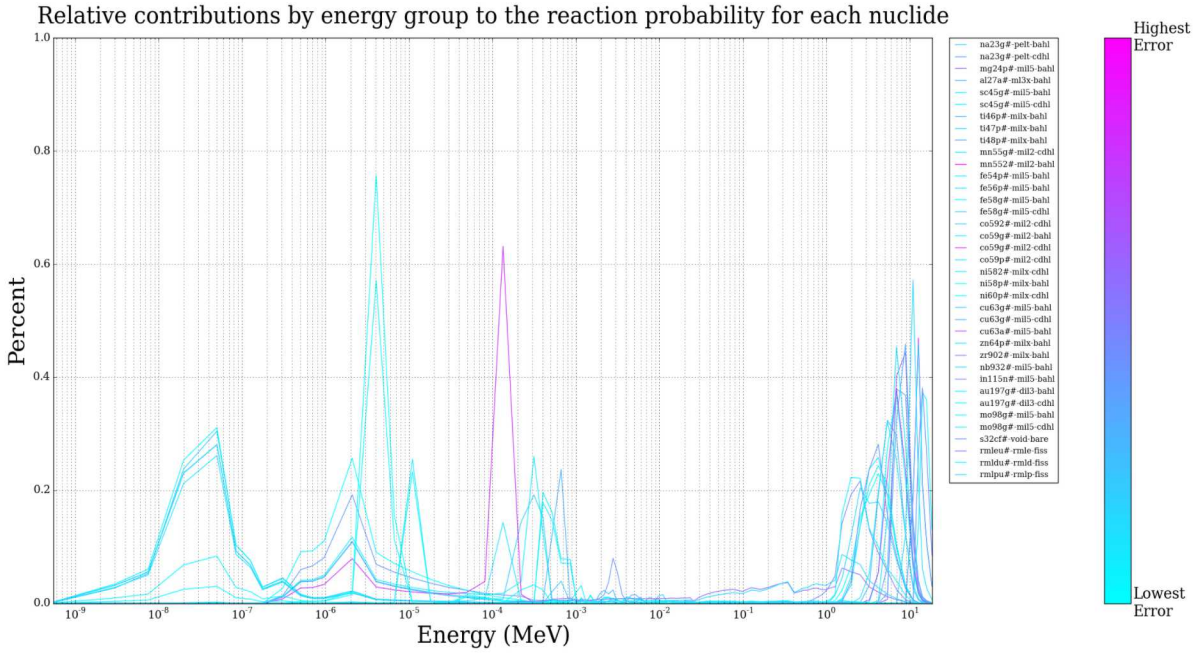

Figure 3.10: Relative contributions to the reaction probabilities for each foil/cover combination showing energy span coverage and error between calculated and measured reaction probabilities.

It should finally be noted that each of the figures produced by GenSpec are interactive. They can be zoomed in or out, line colors and line thicknesses can be changed, and the figures can be saved once all desired changes have been made in various image formats. The authors welcome suggestions on additional plots that may be useful.

Chapter 4

Results

This section will present the results obtained using GenSpec when applied to various neutron irradiation environments here at Sandia National Laboratories, NM. These include the central cavity of the ACRR using the LB44 bucket environment, the PLG bucket environment, and the free field environment. In addition, adjustment results for the SPR-III reactor will also be presented, even though this reactor is no longer in operation. In each case, the results obtained by GenSpec will be compared to those obtained using LSL-M2. The statistical least-squares estimate code LSL-M2 was chosen for comparison because it is widely used and trusted in the field of spectrum adjustment. In addition, each of the ACRR environments have recently been characterized, and LSL-M2 was used for the neutron energy spectrum adjustment portions of those characterizations. Results will be presented here while deferring discussion of said results to Chapter 5.

It should be remembered throughout the comparisons presented in this chapter that LSL-M2 is a least-squares adjustment method and as such is primarily concerned with minimizing the uncertainty in the adjusted spectrum, and only secondarily concerned with matching calculated to measured data. Within this chapter, the calculated reaction probabilities for the adjusted spectra from both GenSpec and LSL-M2 are compared, but this comparison does not intend to show superiority of either method. The important aspect in these comparisons is that GenSpec is producing a spectrum that produces reaction probabilities with similar accuracy to those produced by LSL-M2.

The two most beneficial qualities of GenSpec are difficult to quantify. The first of these is the reduction in spectral shape artefacts which will be clearly seen in the comparisons to adjustments performed using LSL-M2. Admittedly, this is a qualitative analysis and thus is less important than a reduction in uncertainties. The issue of course is that, even if we were given all of the uncertainties of the adjusted spectrum, there may be features that are unrealistic that could not possibly exist. The question then becomes whether to use the spectrum with the minimized uncertainties or a spectrum that predicts similar reaction probabilities and possesses similar properties, yet looks more realistic.

The second beneficial quality of GenSpec is its independence to the trial spectrum resolution, allowing for high-resolution adjustment. This is a quality that is very difficult to achieve with least-squares methods due to their purely mathematical derivation and the singular nature of the cross-section covariance matrices for high-resolution energy group structures. To see this independence, each spectrum will be adjusted with an 89 energy group bin structure for comparison with LSL-M2, followed by an adjustment using a 640 group structure to compare to the 89 group adjustment. Ideally, the two adjustments using GenSpec should not have significant differences.

Free-field environment

The ACRR is the successor of the Annular Core Pulse Reactor (ACPR) which used standard TRIGA fuel in a hexagonal lattice surrounding a large dry cavity at the center of the reactor. The only difference between the two reactors is the fuel material. The ACRR uses UO_2BeO fuel which allows for a larger heat capacity so that it can sustain larger pulses. Figure 4.1 shows the ACRR lattice as modelled in MCNP. It can operate at steady state with power levels up to 2 MW, and can achieve a maximum pulse of 250 MJ with a full-width half-maximum of 6 ms. The ACRR was designed to have an epithermal spectrum which is a compromise between the thermal spectra of light water reactors (LWRs) and the fast spectra of fast reactors. It was designed this way so that different buckets could be placed in the central cavity to reproduce thermal or fast spectra depending upon the bucket materials. This section performs an adjustment on the free field spectrum, which is to say, the spectrum in the central cavity with no bucket used to alter the spectrum. Further details on the ACRR can be found in SAND2006-3067[9]. All LSL-M2 adjustment results used here for comparison to GenSpec were recently obtained as part of the free field characterization in SAND-20XX-XXXX[10]. The adjustments resulting from using the same trial spectrum and experimental measurements in LSL-M2 and GenSpec with an 89 group energy grid are shown in Figure 4.2. The adjustments performed using GenSpec for the 89 and 640 group energy grids are shown in Figure 4.3. Table 4.1 shows the measured and calculated reaction probabilities for both GenSpec and LSL-M2.

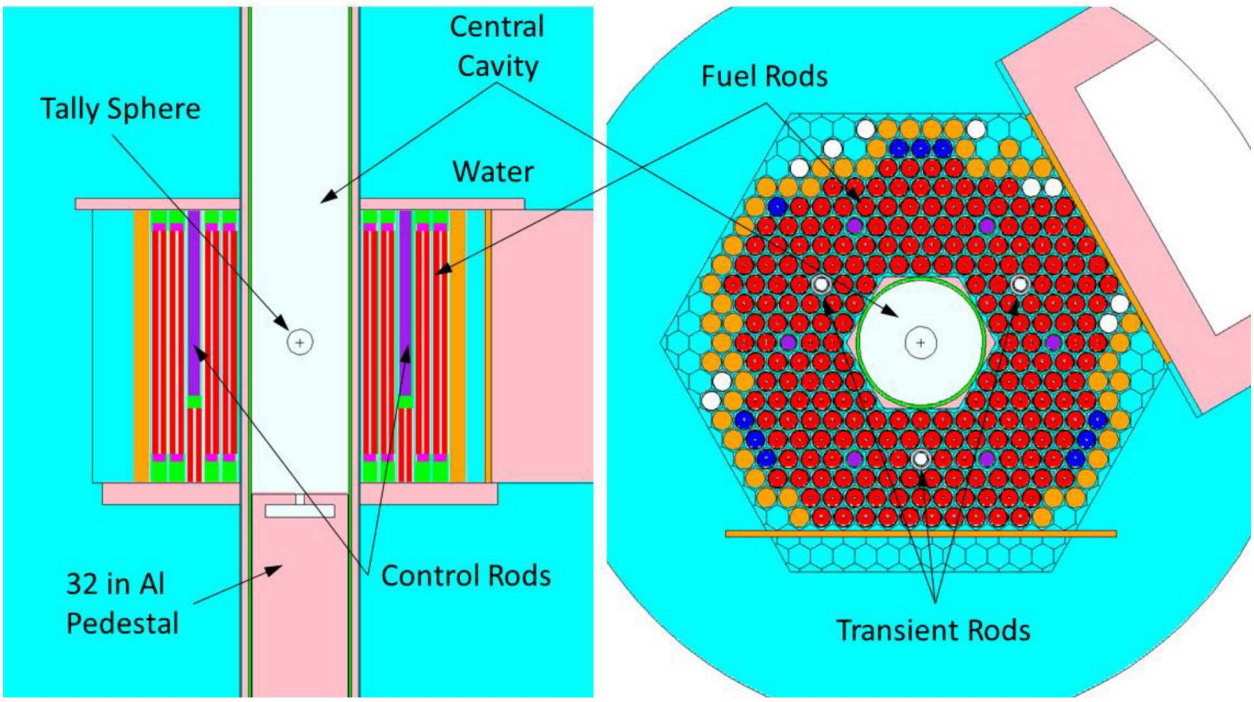

Figure 4.1: Side view (left) and top view (right) of the ACRR as modelled using MCNP showing the hexagonal fuel lattice and the central irradiation cavity.

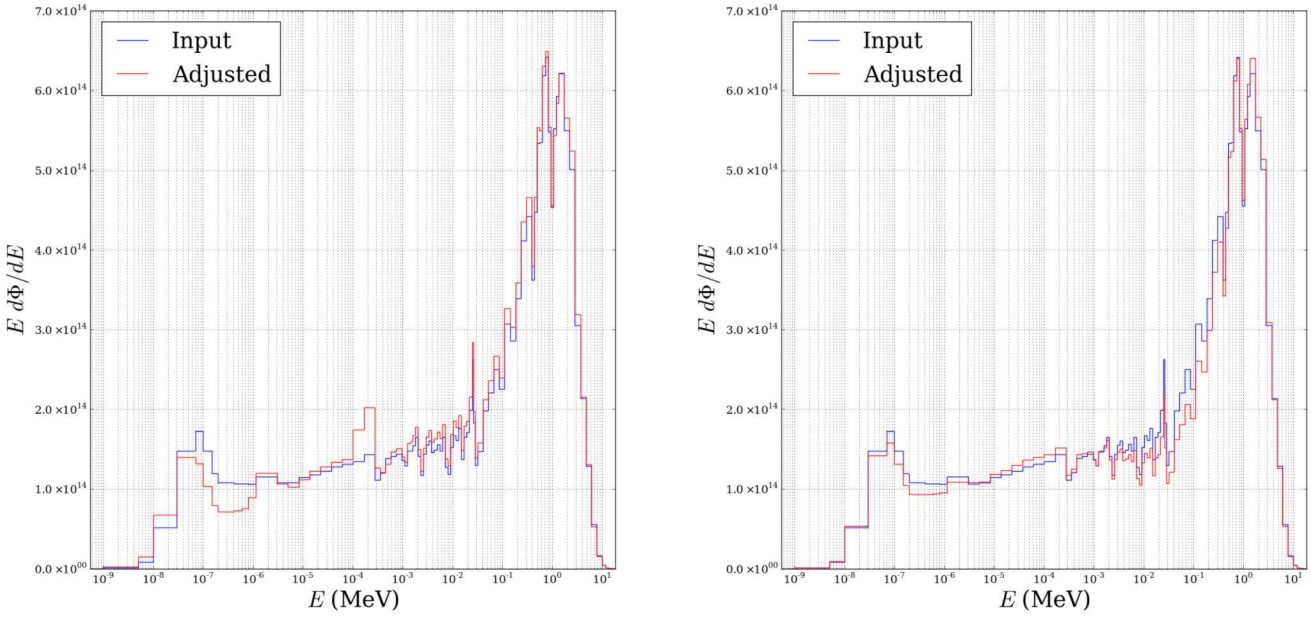

Figure 4.2: Comparison of the adjustments performed using the 89 energy group structure for LSL-M2 (left) and GenSpec (right) for the free field environment in the central cavity of the ACRR.

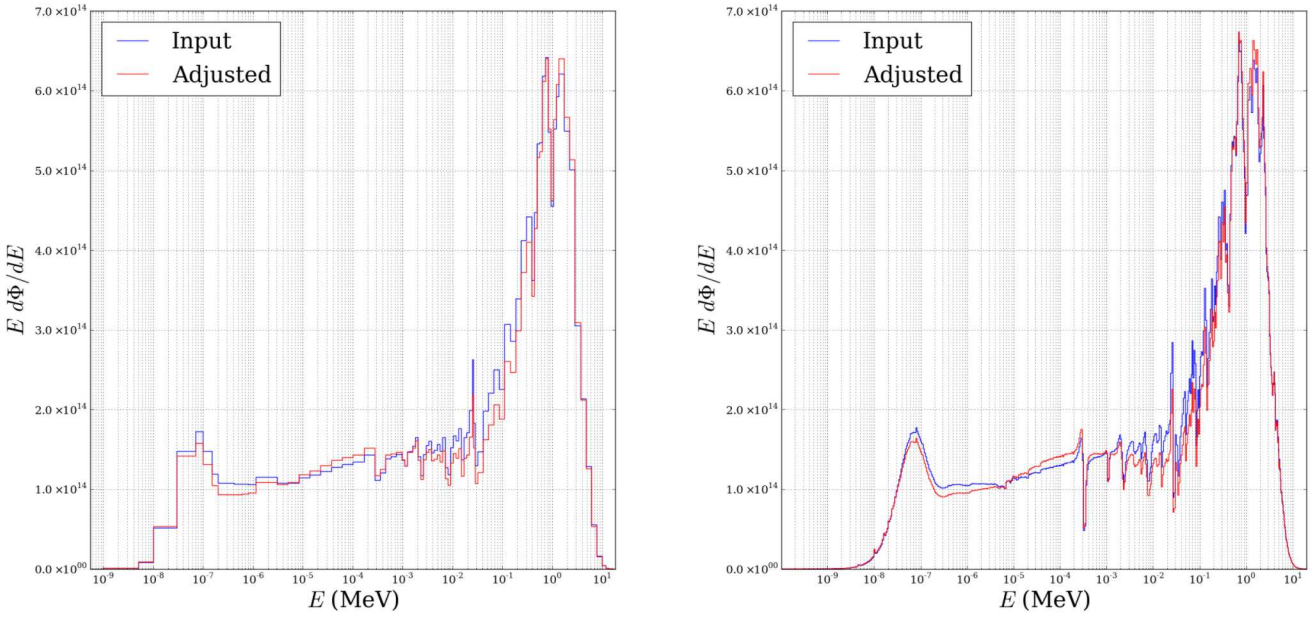

Figure 4.3: Comparison of the adjustments performed using GenSpec for the 89 group (left) and 640 group (right) structures for the free field environment in the central cavity of the ACRR.

Table 4.1: Comparison of the reaction probabilities predicted by LSL-M2 and GenSpec for the free field ACRR environment.

Reaction-Cover	LSL	GenSpec	Measured	LSL % diff	GenSpec % diff
na23g#-pelt-bahl	1.534E-010	1.544E-010	1.445E-010	6.157	6.818
na23g#-pelt-cdhl	3.327E-011	3.307E-011	3.338E-011	0.338	0.930
mg24p#-mil5-bahl	9.993E-013	1.022E-012	1.002E-012	0.270	1.949
al27a#-ml3x-bahl	4.698E-013	4.777E-013	4.787E-013	1.853	0.205
sc45g#-mil5-bahl	7.950E-009	7.966E-009	7.641E-009	4.039	4.250
sc45g#-mil5-cdhl	1.188E-009	1.226E-009	1.195E-009	0.568	2.569
ti46p#-milx-bahl	8.144E-012	8.198E-012	8.200E-012	0.689	0.026
ti47p#-milx-bahl	1.550E-011	1.554E-011	1.589E-011	2.480	2.204
ti48p#-milx-bahl	2.031E-013	2.064E-013	2.006E-013	1.253	2.891
mn55g#-mil2-bahl	4.586E-009	4.559E-009	4.647E-009	1.303	1.895
mn55g#-mil2-cdhl	1.243E-009	1.231E-009	1.229E-009	1.121	0.132
fe54p#-mil5-bahl	6.328E-011	6.317E-011	6.597E-011	4.071	4.242
fe56p#-mil5-bahl	7.432E-013	7.566E-013	7.803E-013	4.760	3.041
fe58g#-mil5-bahl	4.684E-010	4.584E-010	4.618E-010	1.434	0.737
fe58g#-mil5-cdhl	1.376E-010	1.293E-010	1.398E-010	1.599	7.511
co59p#-mil2-cdhl	1.005E-012	1.009E-012	9.964E-013	0.862	1.277
ni58p#-milx-bahl	8.633E-011	8.633E-011	8.633E-011	0.000	0.000
ni60p#-milx-bahl	1.504E-012	1.521E-012	1.517E-012	0.856	0.287
zn64p#-milx-bahl	3.041E-011	3.033E-011	3.165E-011	3.918	4.161
zr902#-milx-bahl	6.961E-014	6.968E-014	6.964E-014	0.040	0.056
nb932#-mil5-bahl	3.151E-013	3.050E-013	3.230E-013	2.431	5.567
w186g#-mil6-bahl	2.732E-008	2.723E-008	2.724E-008	0.276	0.026
au197g#-dil3-bahl	1.851E-007	1.875E-007	1.899E-007	2.548	1.281
au197g#-dil3-cdhl	1.588E-007	1.613E-007	1.590E-007	0.110	1.421
np237f#-void-fisa	1.574E-009	1.598E-009	1.581E-009	0.463	1.089
mo98g#-mil5-bahl	8.416E-010	8.526E-010	8.512E-010	1.127	0.160
mo98g#-mil5-cdhl	8.106E-010	8.218E-010	8.049E-010	0.702	2.095
s32cf#-void-bare	7.431E-002	7.411E-002	7.041E-002	5.540	5.259
rmleu#-rmle-fiss	2.899E-009	2.679E-009	3.140E-009	7.660	14.692
rmldu#-rmld-fiss	2.790E-010	2.817E-010	2.853E-010	2.202	1.278
rmlpu#-rmlp-fiss	3.333E-009	3.167E-009	3.156E-009	5.602	0.341

LB44 environment

The purpose of the LB44 bucket is to filter out low-energy neutrons via absorption in boron. Because of this absorption, the bucket has a reactivity worth of -$6.07 compared to the free field environment. In addition to filtering out low-energy neutrons from the spectrum, the lead layer is able to attenuate the gamma-ray fluence. Figure 4.4 shows the details of the LB44 bucket. Further details on the LB44 bucket can be found in SAND2013-3406[11]. All LSL-M2 adjustment results used for comparison to GenSpec were obtained as part of the LB44 bucket characterization in SAND-2013-3406 as well. The adjustments resulting from using the same trial spectrum and experimental measurements in LSL-M2 and GenSpec with an 89 group energy grid are shown in Figure 4.5. The adjustments performed using GenSpec for the 89 and 640 group energy grids are shown in Figure 4.6. Table 4.2 shows the measured and calculated reaction probabilities for both GenSpec and LSL-M2.

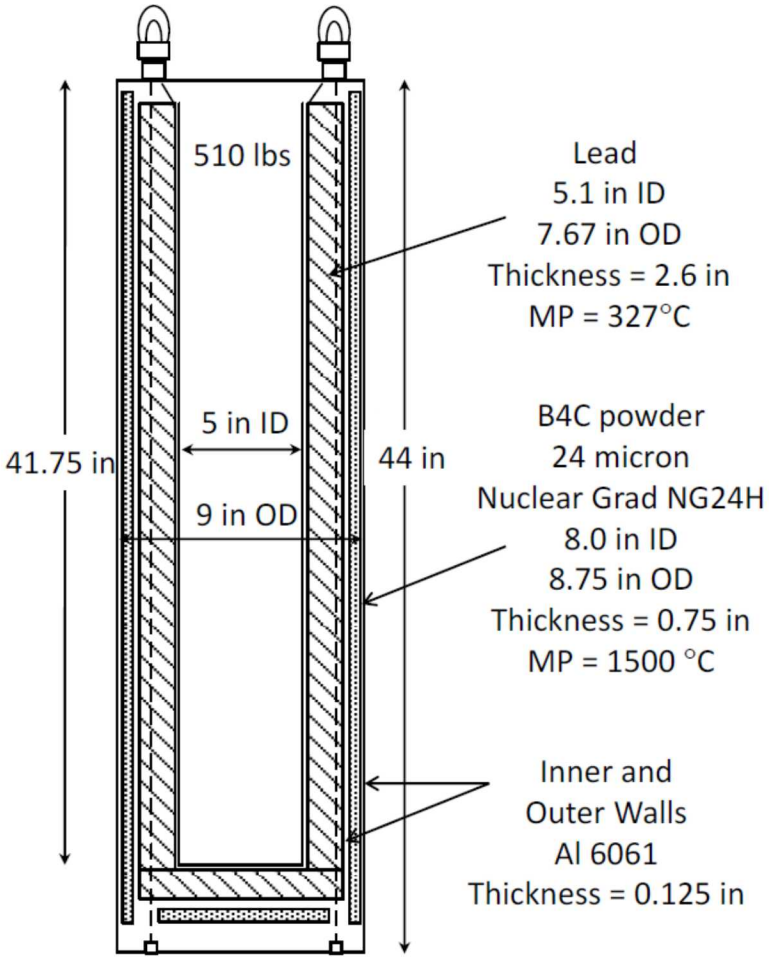

Figure 4.4: Details of the LB44 bucket.

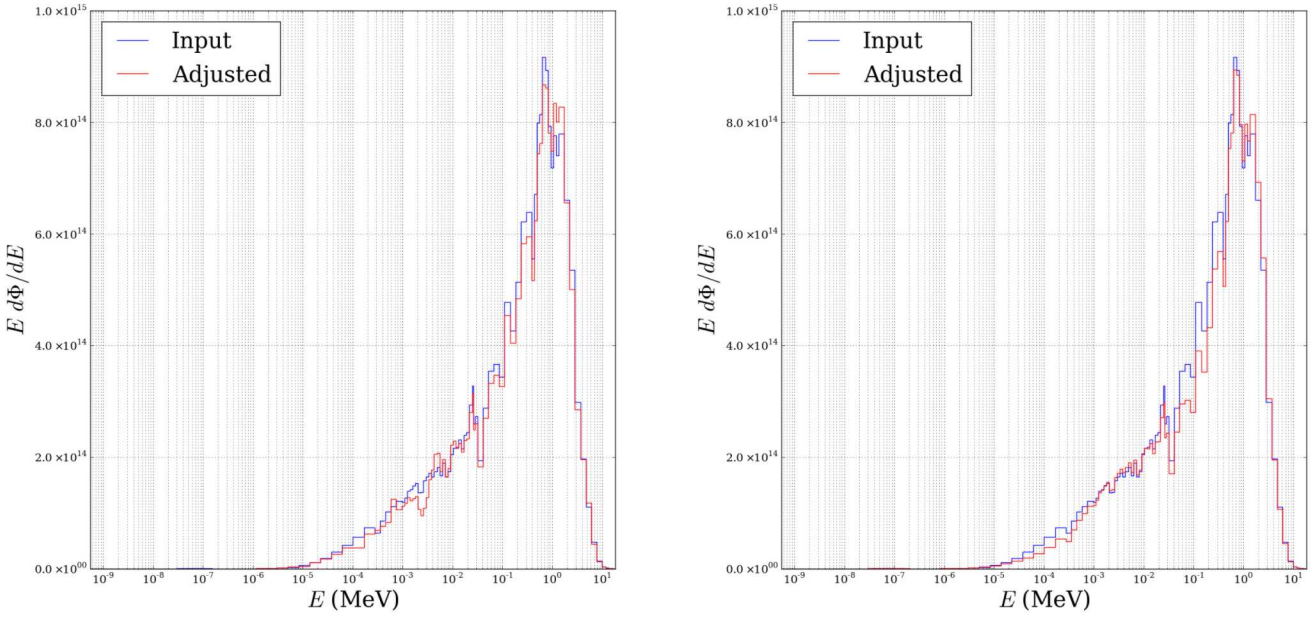

Figure 4.5: Comparison of the adjustments performed using the 89 energy group structure for LSL-M2 (left) and GenSpec (right) for the LB44 environment in the central cavity of the ACRR.

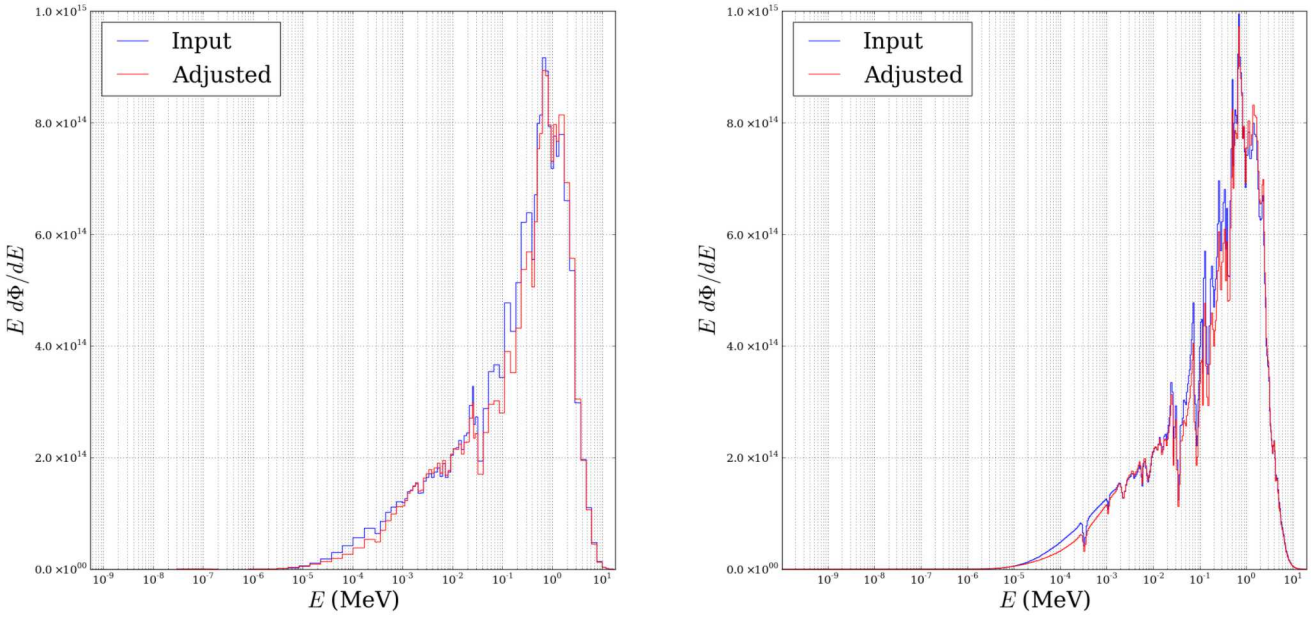

Figure 4.6: Comparison of the adjustments performed using GenSpec for the 89 group (left) and 640 group (right) structures for the LB44 environment in the central cavity of the ACRR.

Table 4.2: Comparison of the reaction probabilities predicted by LSL-M2 and GenSpec for the LB44 environment in the ACRR central cavity.

Reaction-Cover	LSL	GenSpec	Measured	LSL % diff	GenSpec % diff
na23g#-pelt-bare	1.080E-011	1.284E-011	9.970E-012	8.293	28.757
al27p#-void-bare	2.613E-012	2.493E-012	2.822E-012	7.406	11.651
al27a#-void-bare	4.127E-013	4.052E-013	4.096E-013	0.762	1.075
sc45g#-mil5-bare	1.784E-010	1.730E-010	1.749E-010	2.029	1.109
ti46p#-void-bare	7.518E-012	7.166E-012	7.019E-012	7.102	2.100
ti47p#-void-bare	1.568E-011	1.586E-011	1.508E-011	4.004	5.189
ti48p#-void-bare	1.788E-013	1.749E-013	1.760E-013	1.604	0.617
mn55g#-wcu2-bare	5.623E-010	4.870E-010	5.444E-010	3.280	10.540
mn552#-void-bare	1.204E-013	1.221E-013	1.362E-013	11.624	10.373
fe54p#-void-bare	6.095E-011	6.077E-011	5.821E-011	4.702	4.397
fe56p#-void-bare	6.638E-013	6.434E-013	6.434E-013	3.175	0.007
fe58g#-mil5-bare	6.584E-011	5.614E-011	6.901E-011	4.589	18.648
fe58g#-void-fiss	1.704E-011	1.577E-011	1.546E-011	10.213	2.031
co592#-void-bare	1.179E-013	1.196E-013	1.106E-013	6.567	8.098
co59g#-mil2-bare	1.542E-009	1.515E-009	1.499E-009	2.882	1.045
co59g#-void-fiss	5.141E-011	4.907E-011	5.037E-011	2.073	2.575
co59p#-void-bare	9.400E-013	8.987E-013	9.141E-013	2.828	1.687
ni58p#-void-bare	8.470E-011	8.470E-011	8.470E-011	0.000	0.000
ni60p#-void-bare	1.368E-012	1.308E-012	1.296E-012	5.568	0.904
cu63g#-void-bare	4.304E-010	3.977E-010	4.392E-010	1.996	9.448
nb932#-void-bare	2.613E-013	2.563E-013	2.563E-013	1.941	0.006
in115n#-void-bare	2.187E-010	2.215E-010	2.291E-010	4.542	3.322
in115g#-mil5-bare	2.429E-009	2.314E-009	2.096E-009	15.907	10.414
au197g#-dil5-bare	8.552E-009	6.849E-009	8.161E-009	4.797	16.076
au197g#-void-fiss	8.613E-010	8.326E-010	9.122E-010	5.577	8.729
u235f#-void-fisa	3.956E-009	3.809E-009	3.804E-009	4.003	0.124
np237f#-void-fisa	1.984E-009	1.976E-009	1.975E-009	0.454	0.026
pu239f#-void-fisa	4.588E-009	4.437E-009	4.547E-009	0.912	2.426
mo98g#-mil5-bare	5.970E-010	5.750E-010	5.644E-010	5.772	1.882
mo98g#-void-fiss	1.629E-010	1.574E-010	1.363E-010	19.484	15.500
u238f#-void-fisa	2.956E-010	3.014E-010	2.886E-010	2.411	4.425

PLG environment

The purpose of the PLG bucket is to produce more low-energy neutrons via scattering with hydrogen in the high-density polyethylene (HDPE) layer. In addition to increasing this low-energy component, the lead layer is able to attenuate the gamma-ray fluence. Figure 4.7 shows the details of the PLG bucket. Further details on the PLG bucket can be found in SAND2015-4844[12]. All LSL-M2 adjustment results used for comparison to GenSpec were obtained as part of the PLG bucket characterization in SAND-2015-4844 as well. The adjustments resulting from using the same trial spectrum and experimental measurements in LSL-M2 and GenSpec with an 89 group energy grid are shown in Figure 4.8. The adjustments performed using GenSpec for the 89 and 640 group energy grids are shown in Figure 4.9. Table 4.3 shows the measured and calculated reaction probabilities for both GenSpec and LSL-M2.

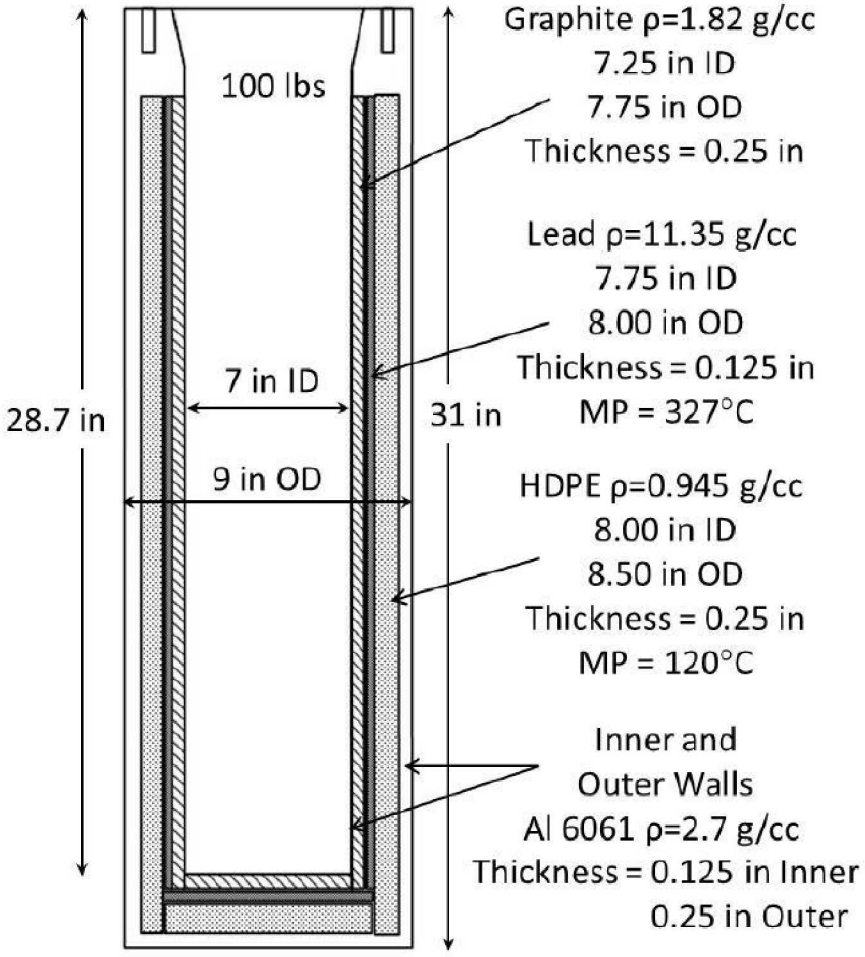

Figure 4.7: Details of the PLG bucket.

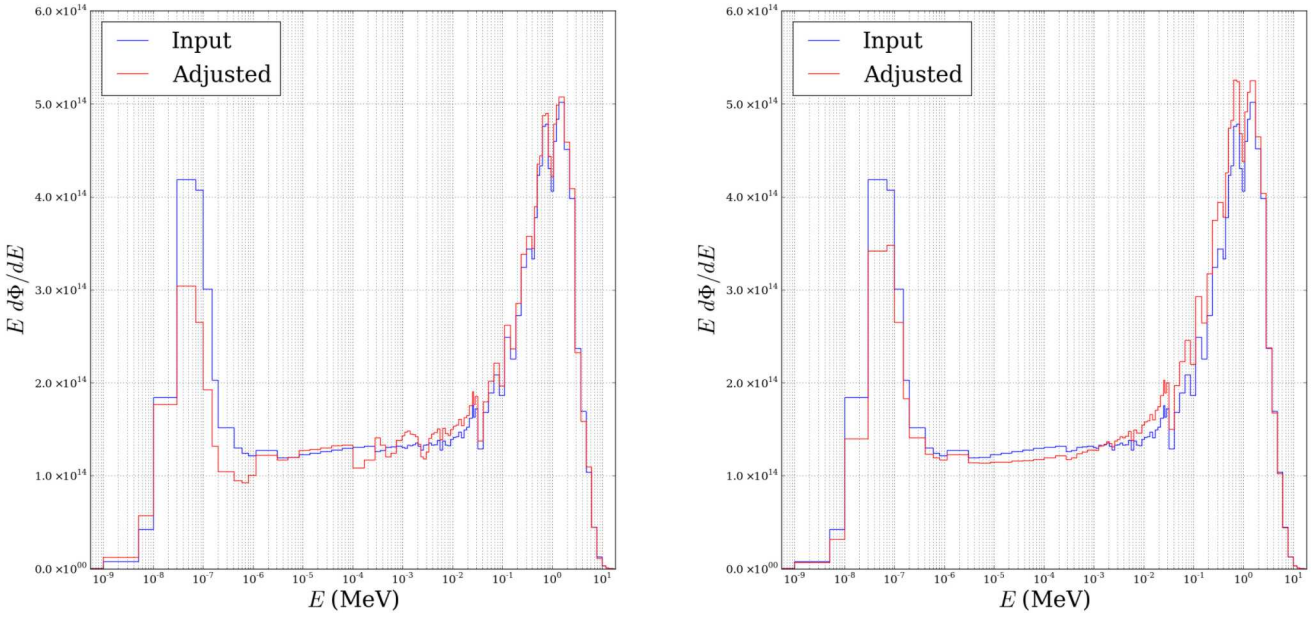

Figure 4.8: Comparison of the adjustments performed using the 89 energy group structure for LSL-M2 (left) and GenSpec (right) for the PLG environment in the central cavity of the ACRR.

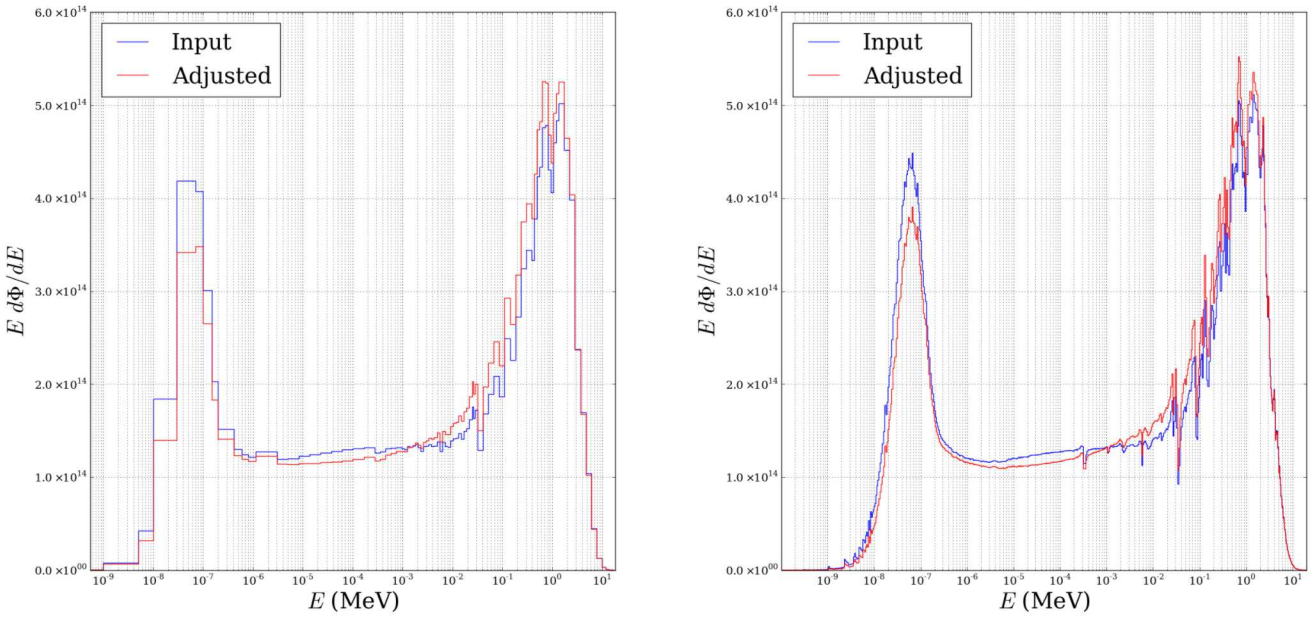

Figure 4.9: Comparison of the adjustments performed using GenSpec for the 89 group (left) and 640 group (right) structures for the PLG environment in the central cavity of the ACRR.

Table 4.3: Comparison of the reaction probabilities predicted by LSL-M2 and GenSpec for the PLG environment in the ACRR central cavity.

Reaction-Cover	LSL	GenSpec	Measured	LSL % diff	GenSpec % diff
na23g#-pelt-bahl	3.243E-010	3.173E-010	3.057E-010	6.096	3.800
na23g#-pelt-cdhl	3.359E-011	3.500E-011	3.245E-011	3.509	7.855
mg24p#-mil5-bahl	8.016E-013	8.373E-013	7.661E-013	4.640	9.292
al27a#-ml3x-bahl	3.736E-013	3.910E-013	3.774E-013	1.003	3.612
sc45g#-mil5-bahl	1.801E-008	1.729E-008	1.731E-008	4.020	0.099
sc45g#-mil5-cdhl	1.297E-009	1.351E-009	1.352E-009	4.104	0.049
ti46p#-milx-bahl	6.693E-012	6.641E-012	6.400E-012	4.575	3.773
ti47p#-milx-bahl	1.238E-011	1.240E-011	1.265E-011	2.170	1.963
ti48p#-milx-bahl	1.638E-013	1.686E-013	1.625E-013	0.791	3.724
mn55g#-mil2-cdhl	1.489E-009	1.415E-009	1.487E-009	0.132	4.839
mn552#-mil2-bahl	1.166E-013	1.106E-013	1.375E-013	15.210	19.572
fe54p#-mil5-bahl	5.031E-011	5.026E-011	4.970E-011	1.235	1.137
fe56p#-mil5-bahl	6.116E-013	6.182E-013	6.200E-013	1.359	0.288
fe58g#-mil5-bahl	9.772E-010	9.330E-010	9.335E-010	4.686	0.050
fe58g#-mil5-cdhl	1.545E-010	1.504E-010	1.493E-010	3.495	0.728
co592#-mil2-cdhl	1.126E-013	1.068E-013	1.117E-013	0.776	4.378
co59g#-mil2-bahl	2.796E-008	2.731E-008	2.724E-008	2.636	0.264
co59g#-mil2-cdhl	5.386E-009	5.727E-009	5.329E-009	1.069	7.473
co59p#-mil2-cdhl	8.208E-013	8.156E-013	8.157E-013	0.620	0.014
ni582#-milx-cdhl	2.318E-015	2.157E-015	2.152E-015	7.732	0.215
ni58p#-milx-bahl	6.879E-011	6.879E-011	6.879E-011	0.000	0.000
ni60p#-milx-cdhl	1.228E-012	1.225E-012	1.254E-012	2.066	2.283
cu63g#-mil5-bahl	3.305E-009	3.186E-009	3.305E-009	0.012	3.587
cu63g#-mil5-cdhl	5.111E-010	5.270E-010	4.975E-010	2.738	5.928
cu63a#-mil5-cdhl	2.972E-013	2.983E-013	3.596E-013	17.351	17.034
zn64p#-milx-bahl	2.403E-011	2.408E-011	2.451E-011	1.942	1.756
zr902#-milx-bahl	5.655E-014	5.269E-014	5.748E-014	1.622	8.338
nb932#-mil5-bahl	2.437E-013	2.438E-013	2.411E-013	1.089	1.117
in115n#-mil5-bahl	1.546E-010	1.561E-010	1.651E-010	6.386	5.446
au197g#-dil3-bahl	2.438E-007	2.341E-007	2.404E-007	1.416	2.601
au197g#-dil3-cdhl	1.802E-007	1.733E-007	1.788E-007	0.760	3.088
mo98g#-mil5-bahl	8.836E-010	8.459E-010	8.667E-010	1.950	2.404
mo98g#-mil5-cdhl	8.037E-010	7.699E-010	7.982E-010	0.692	3.544
s32cf#-void-bare	5.847E-002	5.875E-002	5.437E-002	7.537	8.052
rmleu#-rmle-fiss	2.461E-009	2.572E-009	2.573E-009	4.362	0.053
rmldu#-rmld-fiss	2.252E-010	2.270E-010	2.194E-010	2.631	3.473
rmlpu#-rmlp-fiss	2.792E-009	2.916E-009	2.570E-009	8.627	13.482

SPR-III environment

SPR-III is an advanced fast-burst Godiva-type reactor with a large 16.5 cm central cavity. Although it was developed primarily for the radiation testing of electronic parts and systems, it has been used in a wide variety of research activities. It is positioned in the center of an air-filled shield building called a Kiva. SPR-III has a neutron absorbing shroud that serves to decouple it from room return neutrons. Experiments are conducted not only in the central cavity but also outside the core at distances between 0.3 and 3.0 meters from the core axis. There are also ports in the shield wall for fielding of experiments that require collimated beam geometries. The reactor can be operated in steady-state (up to 10 kW power) or pulsed mode (10 MJ in an 80 μs FWHM pulse that yields approximately 5×10^{14} n/cm^2 in the central cavity). Figure 4.10 shows SPR-III as modelled in MCNP. The adjustments resulting from using the same trial spectrum and experimental measurements in LSL-M2 and GenSpec with an 89 group energy grid are shown in Figure 4.11. The adjustments performed using GenSpec for the 89 and 640 group energy grids are shown in Figure 4.12. Table 4.4 shows the measured and calculated reaction probabilities for both GenSpec and LSL-M2.

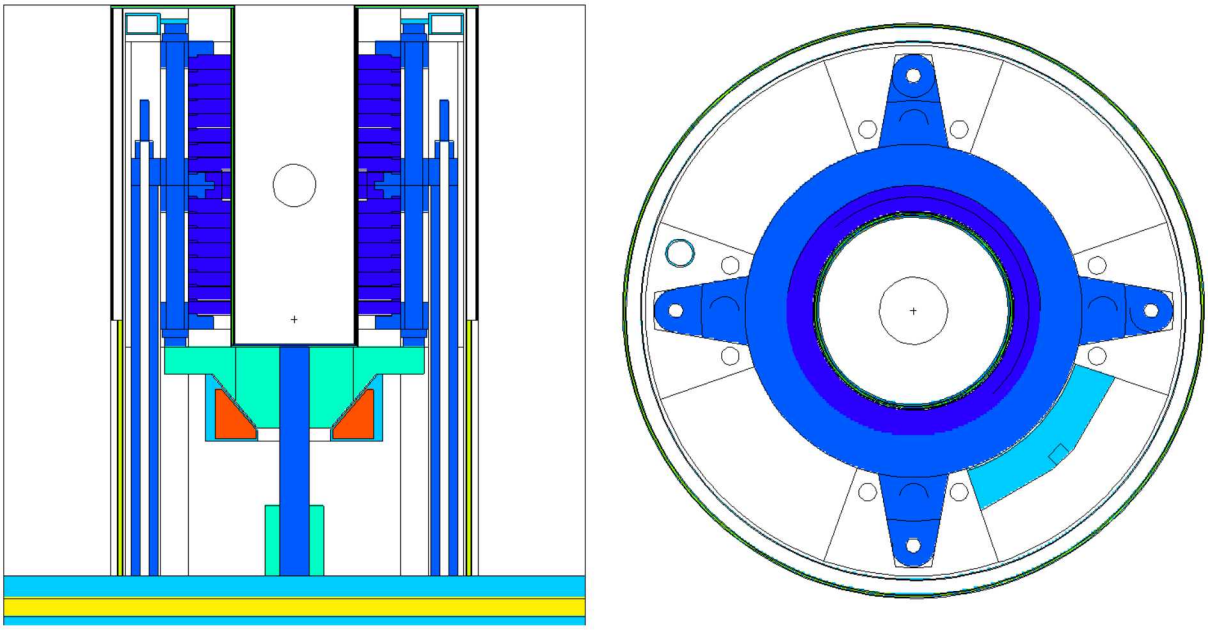

Figure 4.10: Side view (left) and top view (right) of SPR-III as modelled using MCNP.

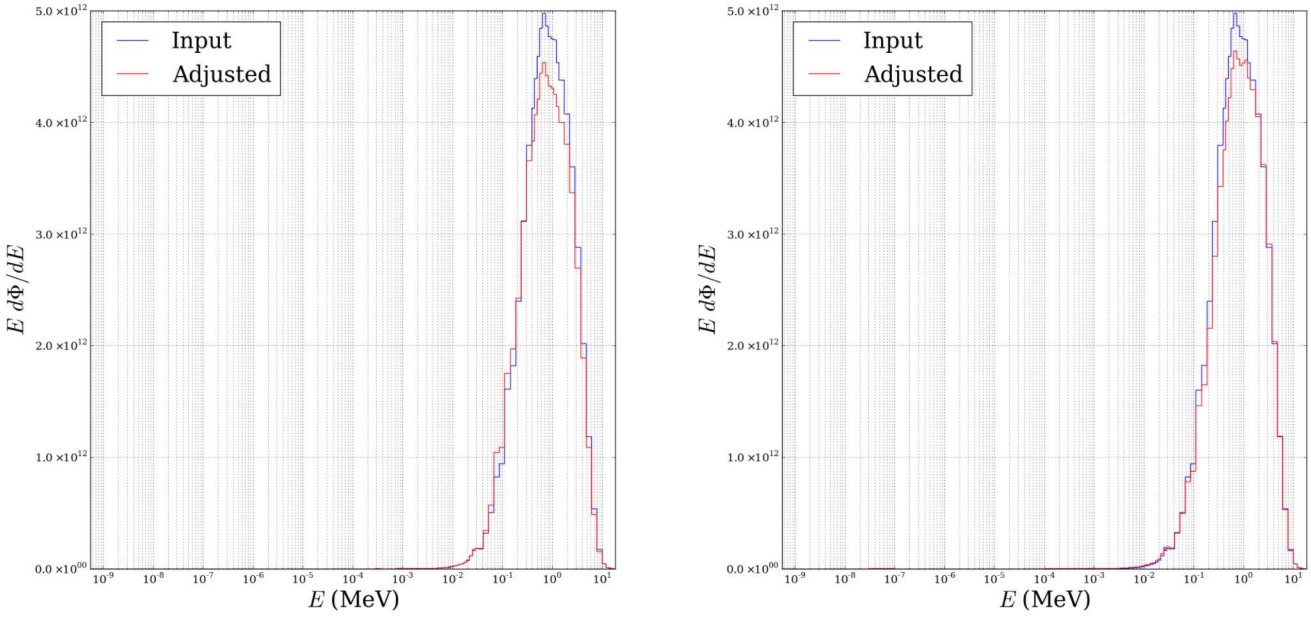

Figure 4.11: Comparison of the adjustments performed using the 89 energy group structure for LSL-M2 (left) and GenSpec (right) for the spectrum in the SPR-III central cavity.

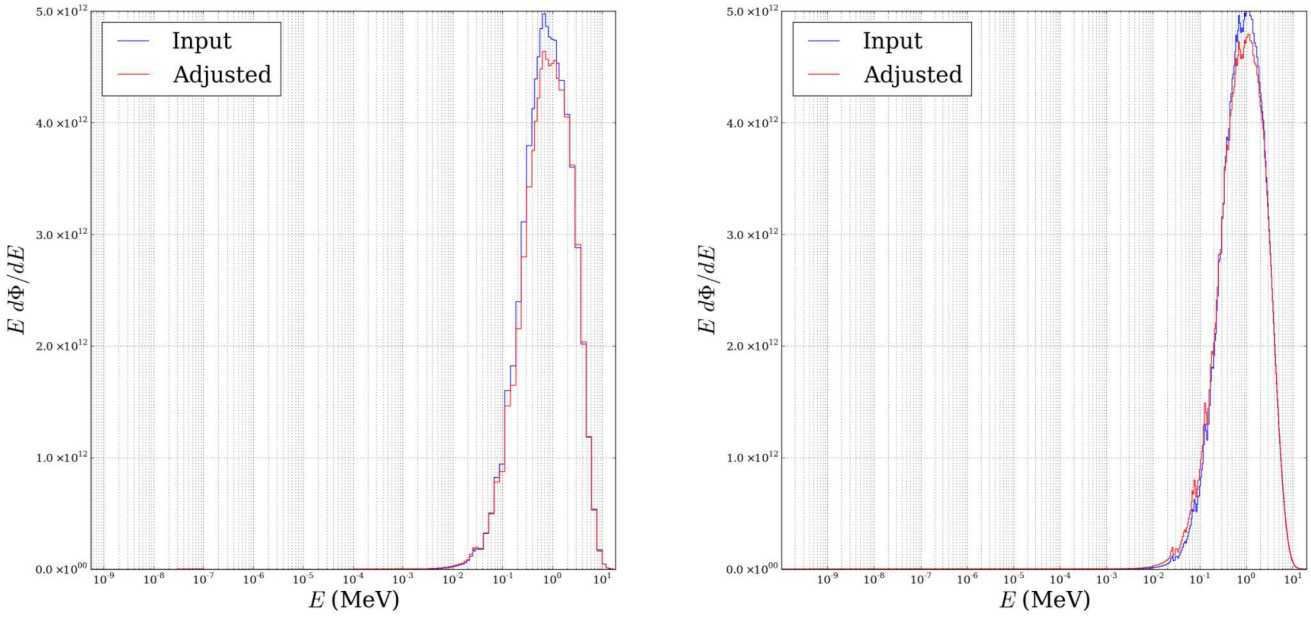

Figure 4.12: Comparison of the adjustments performed using GenSpec for the 89 group (left) and 640 group (right) structures for the spectrum in the SPR-III central cavity.

Table 4.4: Comparison of the reaction probabilities predicted by LSL-M2 and GenSpec for the spectrum in the SPR-III central cavity.

Reaction-Cover	LSL	GenSpec	Measured	LSL % diff	GenSpec % diff
na23g#-pelt-cdna	6.183E-015	5.576E-015	5.573E-015	10.943	0.053
na23g#-pelt-bare	6.445E-015	5.857E-015	5.509E-015	16.986	6.308
na23g#-pelt-fiss	5.681E-015	5.001E-015	4.635E-015	22.577	7.886
mg24p#-mil5-cdnm	1.072E-014	1.036E-014	9.654E-015	11.036	7.285
al27p#-ml3x-cdnm	2.771E-014	2.763E-014	2.437E-014	13.713	13.383
al27a#-ml3x-cdnm	5.049E-015	4.846E-015	4.769E-015	5.880	1.617
sc45g#-mil5-cdnm	1.509E-013	1.332E-013	1.433E-013	5.334	7.060
sc45g#-mil5-fiss	1.379E-013	1.184E-013	1.245E-013	10.747	4.860
ti46p#-milx-cdnm	7.986E-014	7.958E-014	7.650E-014	4.391	4.033
ti47p#-milx-cdnm	1.333E-013	1.330E-013	1.203E-013	10.828	10.529
ti48p#-milx-cdnm	2.148E-015	2.073E-015	1.970E-015	9.059	5.250
mn55g#-wcu2-cdnm	7.799E-014	7.283E-014	8.883E-014	12.205	18.014
mn55g#-wcu2-fiss	6.940E-014	6.064E-014	6.117E-014	13.458	0.871
fe54p#-mil5-cdnm	5.722E-013	5.727E-013	5.447E-013	5.054	5.145
fe56p#-mil5-cdnm	7.689E-015	7.527E-015	7.259E-015	5.925	3.693
ni58p#-mil5-cdnm	7.725E-013	7.725E-013	7.725E-013	0.000	0.000
cu63g#-mil5-cdnm	2.344E-013	2.101E-013	2.179E-013	7.591	3.566
zn64p#-milx-cdnm	2.768E-013	2.770E-013	2.562E-013	8.036	8.118
zr902#-mil5-cdnm	7.183E-016	6.189E-016	6.580E-016	9.160	5.942
in115n#-mil5-bare	1.548E-012	1.532E-012	1.447E-012	6.977	5.890
in115g#-mil5-cdnm	2.665E-012	2.459E-012	2.330E-012	14.387	5.541
au197g#-dil5-bare	2.274E-012	2.317E-012	2.491E-012	8.709	6.985
au197g#-dil5-cdnm	2.208E-012	2.245E-012	2.208E-012	0.016	1.692
u235f#-void-cdtk	1.869E-011	1.705E-011	1.755E-011	6.473	2.824
u235f#-void-fiss	1.664E-011	1.518E-011	1.500E-011	10.913	1.208
np237f#-void-cdtk	1.375E-011	1.333E-011	1.234E-011	11.457	8.030
np237f#-void-fiss	1.298E-011	1.259E-011	1.182E-011	9.833	6.519
pu239f#-void-cdtk	2.517E-011	2.321E-011	2.233E-011	12.740	3.950
s32cf#-sulf-bare	5.960E-004	5.964E-004	5.991E-004	0.518	0.448
u238f#-void-cdnm	2.409E-012	2.392E-012	2.317E-012	3.955	3.228
u238f#-bvod-fiss	2.141E-012	2.127E-012	2.223E-012	3.694	4.328

Chapter 5

Conclusions

This work has aimed to develop a spectrum adjustment method that is less prone to unrealistic spectral shape artefacts and less hindered by high-resolution input spectra, than currently existing methods. This has been accomplished through the use of polynomial adjustment functions that are optimized through the use of a genetic algorithm. The final product is a spectrum adjustment code named GenSpec that produces adjusted spectra that closely resemble those obtained using other spectrum adjustment methods and which predict the measured reaction probabilities with similar accuracy. GenSpec does not currently calculate the uncertainties in the adjusted spectrum, but this capability is currently being implemented. This report has introduced the concept of spectrum adjustment and described general features of a genetic algorithm in Chapters 1 and 2 respectively. Chapter 3 serves as the user's manual for GenSpec. Chapter 4 compared the results of various spectrum adjustments performed using GenSpec to the same adjustments performed using LSL-M2 which uses least-squares estimates.

While the results shown in Chapter 4 are promising, no spectrum adjustment method can claim to produce the correct spectrum. This is an unfortunate side effect of trying to solve a problem which mathematically has infinitely many solutions. The best course of action when characterizing a neutron irradiation environment may be to use several adjustment techniques to gather appropriate information from each. For instance, if several methods agree that the thermal component of the spectrum is too high, as was the case for the free-field and PLG environments in the ACRR from the previous chapter, this should lead the user to re-examine the MCNP input deck for model properties that could be responsible for this. If two methods predict significantly different reaction probabilities for a specific foil/cover combination, several factors including measurement error, high cross-section uncertainty, or inaccurate self-shielding corrections could be responsible. In this case, the foil should be re-fielded in the irradiation environment to rule out experimental error.

References

[1] F. B. Brown et al., *MCNP6 User's Manual, Version 1.0, LA-CP-13-00634*, Los Alamos National Laboratories, Los Alamos, NM (2004).

[2] G. Knoll, *Radiation Detection and Measurement*, John Wiley & Sons (2010).

[3] W. Zijp, "Comparison of Neutron Spectrum Unfolding Codes," 1978, IAEA Report: IAEA-R-1811-F.

[4] F. G. Perey, "Spectrum Unfolding by the Least-squares Method," Oak Ridge National Laboratory (1979).

[5] 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).

[6] F. W. Stallmann, "LSL-M2: A Computer Program For Least-Squares Logarithmic Adjustment of Neutron Spectra," 1986, NUREG/CR-4349, ORNL/TM-9933.

[7] M. D. Vose, *The Simple Genetic Algorithm: Foundations and Theory*, MIT Press, Cambridge, MA, USA (1998).

[8] P. Griffin and J. Kelly, "A rigorous treatment of self-shielding and covers in neutron spectra determinations," *IEEE Trans. Nucl. Sci.*, **42**, *6*, pp. 1878–1885 (1995).

[9] K. R. DePriest, P. J. Cooper, and E. J. Parma, *MCNP/MCNPX model of the Annular Core Research Reactor, SAND2006-3067*, Sandia National Laboratories, Albuquerque, NM (2006).

[10] E. J. Parma et al., *Radiation Characterization Summary: ACRR Central Cavity Free-Field Environment with the 32-Inch Pedestal at the Core Centerline (ACRR-FF-CC-32-cl), SAND2015-XXXX (to be published in 2015).*, Sandia National Laboratories, Albuquerque, NM (2015).

[11] E. J. Parma et al., *Radiation Characterization Summary: ACRR 44-Inch Lead-Boron Bucket Located in the Central Cavity on the 32-Inch Pedestal at the Core Centerline (ACRR-LB44-CC-32-cl), SAND2013-3406*, Sandia National Laboratories, Albuquerque, NM (2015).

[12] E. J. Parma et al., *Radiation Characterization Summary: ACRR Polyethylene-Lead-Graphite (PLG) Bucket Located in the Central Cavity on the 32-Inch Pedestal at the Core Centerline (ACRR-PLG-CC-32-cl), SAND2015-4844*, Sandia National Laboratories, Albuquerque, NM (2015).

DISTRIBUTION:

1	MS 1115	Warren Strong, 4248 (electronic copy)
1	MS 1115	Randell Salyer, 4248 (electronic copy)
1	MS 1136	Billy Martin, 6221 (electronic copy)
1	MS 1141	Ron Knief, 1382 (electronic copy)
1	MS 1141	Dave Wheeler, 1382 (electronic copy)
1	MS 1141	Danielle Redhouse, 1382
1	MS 1141	John Miller, 1383 (electronic copy)
1	MS 1141	Peter Subaiya, 1383 (electronic copy)
1	MS 1143	Thomas Quirk, 1384 (electronic copy)
1	MS 1146	Ed Parma, 1384
1	MS 1146	Taylor Lane, 1384 (electronic copy)
1	MS 1146	Pat Griffin, 1300 (electronic copy)
1	MS 1146	Ken Reil, 1384 (electronic copy)
1	MS 1146	Russell DePriest, 1384 (electronic copy)
1	MS 1146	Richard Vega, 1384
1	MS 1146	David Vehar, 1384 (electronic copy)
1	MS 1146	Dimitrios Michaelides, 1384 (electronic copy)
1	MS 0899	Technical Library, 9536 (electronic copy)

Sandia National Laboratories