SAND2014-18638C

Development of a Genetic Algorithm for Neutron Energy Spectrum Adjustment

Sandia National Laboratories

Exceptional service in the national interest

Student Information

Name: Richard Vega
School: Texas A&M University
Degree Pursued: Ph.D.
Discipline: Nuclear Engineering
Manager: Ken Reil
Mentor: Ed Parma
Org Name: Applied Nuclear Technologies
Org Number: 1384
Sandia National Laboratories, NM
U.S. Department of Energy

Abstract

Neutron spectrum adjustment, or unfolding as it is commonly called, is an optimization problem. Modern Monte Carlo neutron transport codes can produce high resolution neutron energy spectra. As computing power increases, the statistical error due to the stochastic solution method can be made vanishingly small. Unfortunately, the error incurred by the uncertainty in model parameters and transport cross sections still remains. The goal in spectrum adjustment is to adjust the spectrum produced by a transport code so that it agrees more closely with measured data. Over the past few decades, genetic algorithms have shown the ability to solve optimization problems in logistics such as the traveling salesman and number partitioning problems. The use of a genetic algorithm by NASA in 2006 for antenna design revealed even further potential for this abstract computational method. These and other successes led to the development of a genetic algorithm for spectrum adjustment presented here. The algorithm was used to adjust the spectrum at the center of the ACRR central cavity in the PLG bucket environment, and is compared to the adjustment performed using LSL-M2. Although much work is still to be done, the results are promising.

Introduction

Neutron energy spectrum measurement is complicated by the fact that neutron detectors with high energy resolution simply do not exist. The only information that can be gained experimentally are integral quantities such as the total reaction rate for a given reaction. This information is gained through the use of activation foils, where the activity after irradiation is proportional to the reaction rate. The reaction rate is the integral of the product of the neutron flux and the reaction cross section over energy. As there are typically less than 50 feasible activation foils, each leading to a single equation, this produces an under-determined problem if the desired number of energy groups in the final spectrum, each of which is treated as a variable, is greater than the number of activation foils at the experimenter's disposal. For comparison, codes like MCNP can produce spectra with hundreds of energy groups and very low statistical error. Recently, LSL-M2 has been used to adjust the MCNP produced spectrum at the center of the PLG bucket in the ACRR central cavity. The result is shown in Figure 1. The MCNP spectrum contained 89 energy groups and a total of 37 activation foils were used. The LSL-M2 adjustment shows that the MCNP spectrum over-predicted the flux in the low-energy range and under-predicted the flux in the high-energy range. This could be due to several factors. For instance, the density of HDPE in the MCNP model of the PLG bucket could have been too high, leading to an overestimation of the number of neutrons scattered to low energies. Unfortunately, there is not such a simple explanation for the dips and peaks created by LSL-M2 in the mid-energy range. These unrealistic peaks are one of the main reasons for the development of a new method.

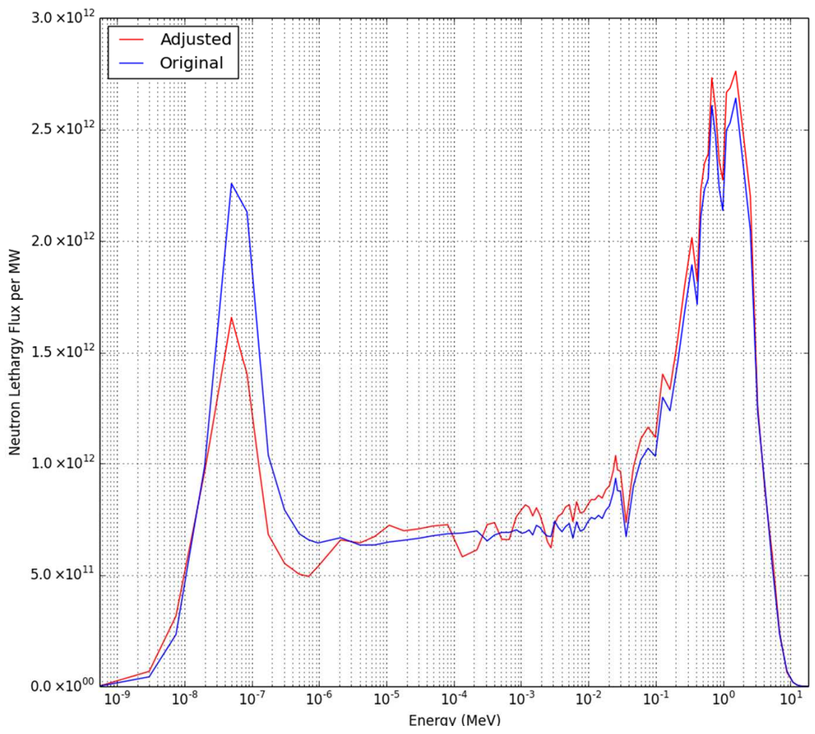

Figure 1. LSL-M2 adjustment of the PLG spectrum

Methods

Genetic algorithms attempt to solve an optimization problem by mimicking the theory of natural selection. A population of possible solutions to the problem is formed, and each specimen is assigned a fitness value based on the quantity being optimized. To produce children, specimens are then chosen for mating based on their fitness values; high fitness specimens mate more frequently than low fitness specimens. In addition, the genes of each child are subject to a random mutation with a pre-defined probability. The children produced should inherit the good qualities of their parents, carrying them into further generations. Mating is performed until the number of children equals the number of specimens in the original population. This process is then repeated for a pre-defined number of generations. Thus, for a genetic algorithm to work, techniques for setting the population, assigning fitness values, specimen mating, and gene mutation must be properly chosen.

Setting the population: In light of the unrealistic peaks produced by other adjustment methods, it was desired to produce a smooth adjustment. A total of 13 points in the energy domain of interest were chosen. These points will be referred to as sites. For each site, a Gaussian distributed random number with a mean of zero and a standard deviation of 0.07 was added to unity and then a least squares polynomial fit was generated, resulting in what will be referred to as the shift function. Each specimen is characterized by the perturbation at each site and the associated shift function. Thus, each specimen has 13 genes and the shape of each specimen is determined by a polynomial fit through these 13 values. The shift functions of the initial population along with the spectra themselves are shown in Figure 2. The results of Figure 2 are for a 5th order polynomial fit, however the polynomial order is flexible in the code.

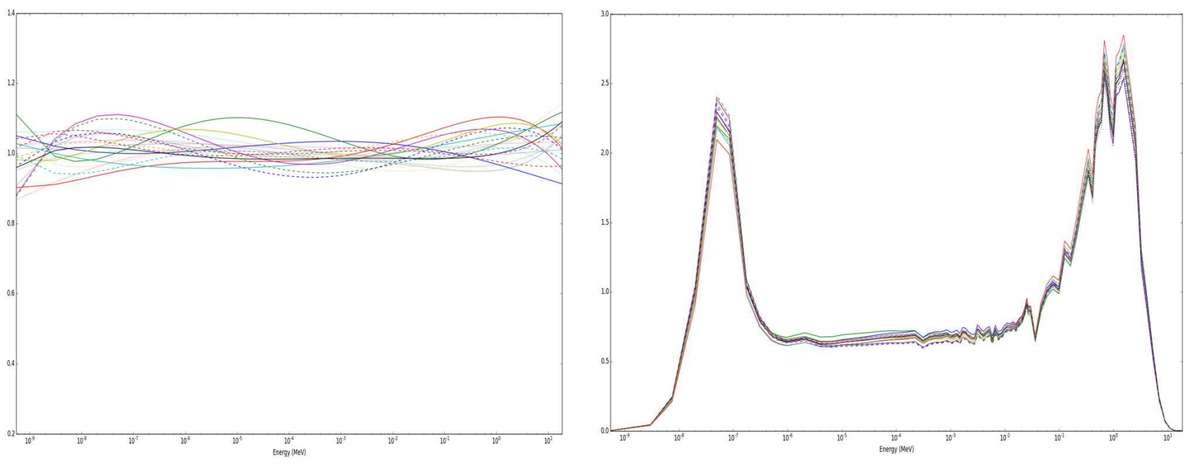

Figure 2. (left) Shift functions for the first 20 specimens in the initial population. (right) The associated specimen's spectral shape functions (arbitrary magnitude).

Fitness function: The ultimate goal of spectrum adjustment is to produce a spectrum that agrees more closely with measured data. In addition, the fitness is typically a quantity to be maximized. Thus, a large fitness should correspond to a small difference between measured and calculated reaction rates. This leads to the following fitness function:

$$f = C - \sum_{foils} \left(\frac{\left| \{\sum \sigma\varphi\Delta E\}_{calc} - \{\sum \sigma\varphi\Delta E\}_{meas} \right|}{\{\sum \sigma\varphi\Delta E\}_{meas}} \right)$$

where f is the fitness, C is a constant, σ is the reaction cross section, φ is the neutron flux, and ΔE is the bin width of the energy group.

Mating: In order to produce new solutions from the population of possible solutions, gene crossover is chosen as the mating algorithm. This is performed by choosing a cutting point between the 13 genes and exchanging the genes of each parent to produce two new solutions. This process is shown schematically in Figure 3. In the genetic algorithm presented here, each gene is a value perturbed from unity through which a polynomial fit is generated. The product of the polynomial fit and the trial spectrum is the specimen spectrum. This crossover guarantees that if two specimens of relatively high fitness are chosen for mating, the children will be likely to inherit the best qualities of both. Although multiple crossover points can be used, the algorithm presented here uses single point crossover.

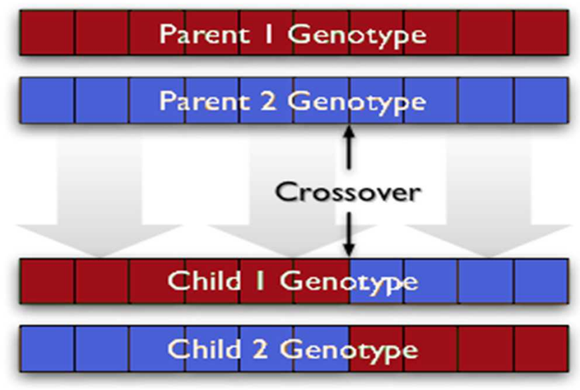

Figure 3. Depiction of the crossover method.

Mutation: It is often the case that several maxima exist for a given problem. To avoid having the population converge on a fitness that is not the absolute maximum fitness, there has to be competition to said convergence. This is the purpose of mutation. Even after a solution has been converged upon, mutation ensures that other areas of the solution space are being explored. The genetic algorithm presented here uses gene-wise mutation. This is accomplished by first setting a mutation probability m. During the mating process, each gene of each of the children will have a set probability m of being perturbed randomly from its inherited value. The perturbation is a Gaussian distributed random number that is added to the previous value of the gene site.

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.
SAND No. 2011-XXXXP

Results

Effect of polynomial order: By far the most significant variable in the method presented here is the order of the polynomial used when determining the shift function for each specimen. As previously mentioned, to obtain a smooth adjustment and to not introduce unrealistic peaks in the final result, a polynomial shift function is fitted to the perturbation values that make up the genes of each specimen. If the method converges on a maximum fitness, we would expect for the shift functions of all specimens in the population to migrate towards that which produces the maximum fitness. In fact, this is exactly what is observed as can be seen in Figure 4. Each panel in Figure 4 represents the shift functions of the population at the end of 500 generations.

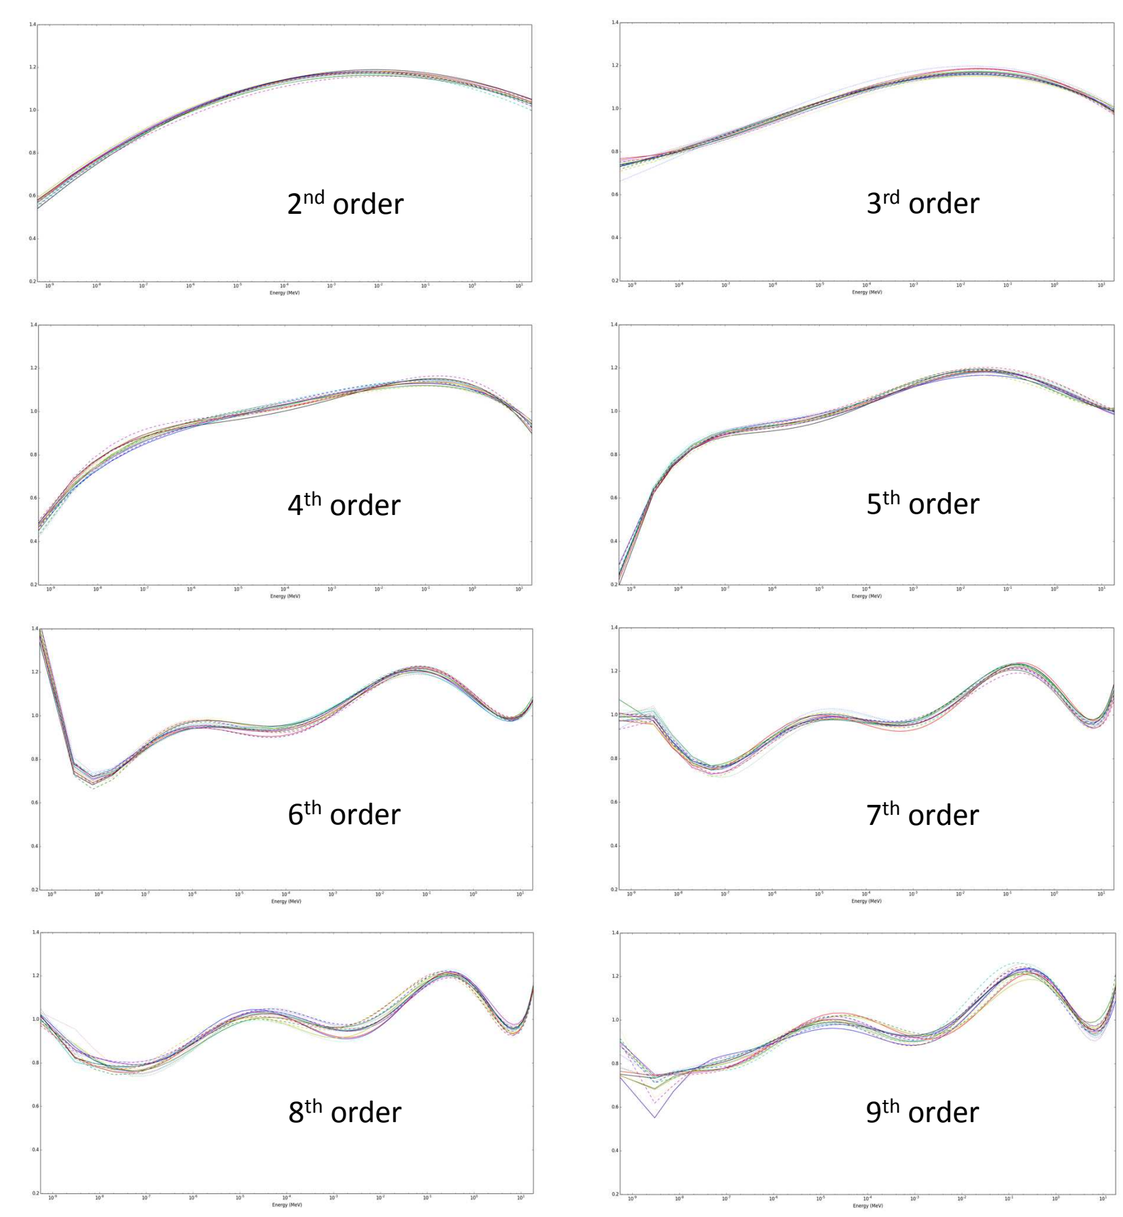

Figure 4. Shift functions for the first 20 specimens in the final population for various polynomial orders.

From Figure 4, it appears that not only does the shift function for the entire population converge to that of the highest fitness, it also appears that the shape of the maximum fitness shift function converges for polynomial orders greater than 6. However, as the polynomial order approaches the number of sites, it is expected that the polynomials will swing wildly in an attempt to pass through all 13 values. This is expected to impact the maximum fitness negatively, which can be seen in Figure 5.

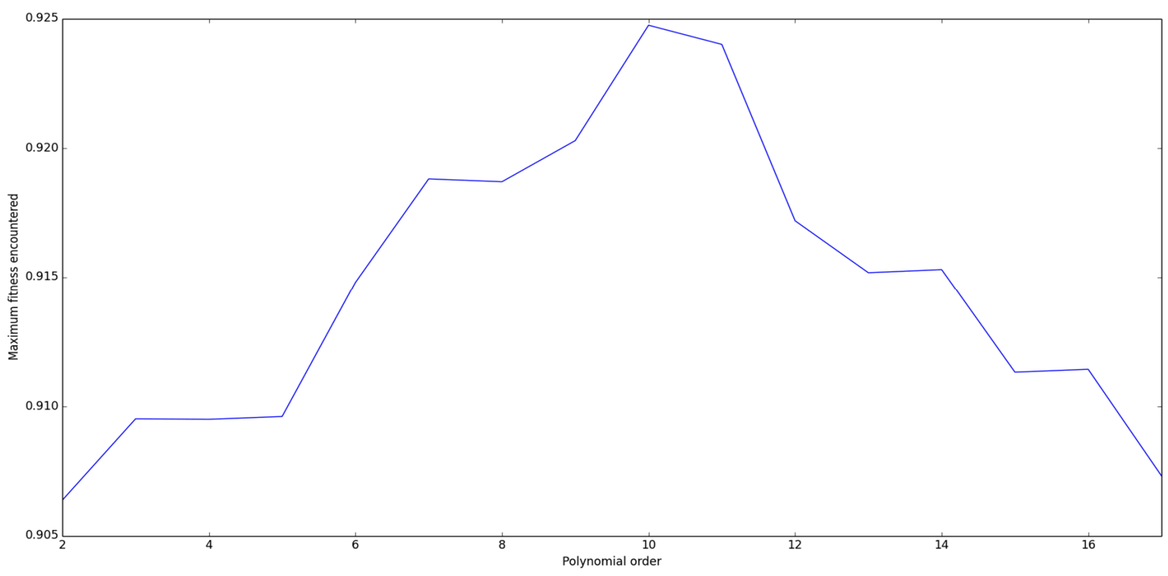

Figure 5. Fitness of best solution as a function of polynomial order.

Convergence speed: In all cases considered, the algorithm was shown to converge in less than 100 generations. In addition, the convergence showed the characteristics of a successful genetic algorithm; namely that the minimum, average, and maximum fitness increased with increasing number of generation at roughly the same rate and converged at nearly the same time. This general trend can be seen in Figure 6.

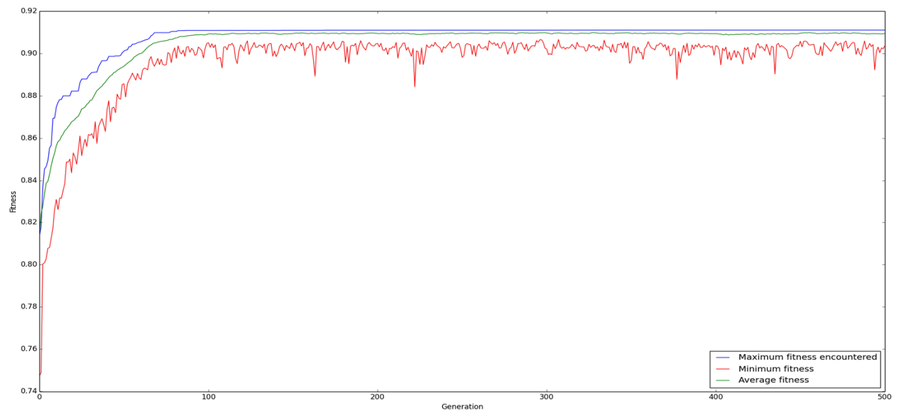

Figure 6. Minimum, average, and maximum fitness as a function of generation.

Final adjustment and LSL-M2 comparison: The final results of the genetic algorithm can be seen in Figure 7 and Table 1. Figure 7 should be compared to the LSL-M2 adjustment of Figure 1. Table 1 shows the reaction rates predicted by the LSL-M2 and genetic algorithm spectra, as well as the measured reaction rates and percent error for each. The foil labeled ** is the reference foil.

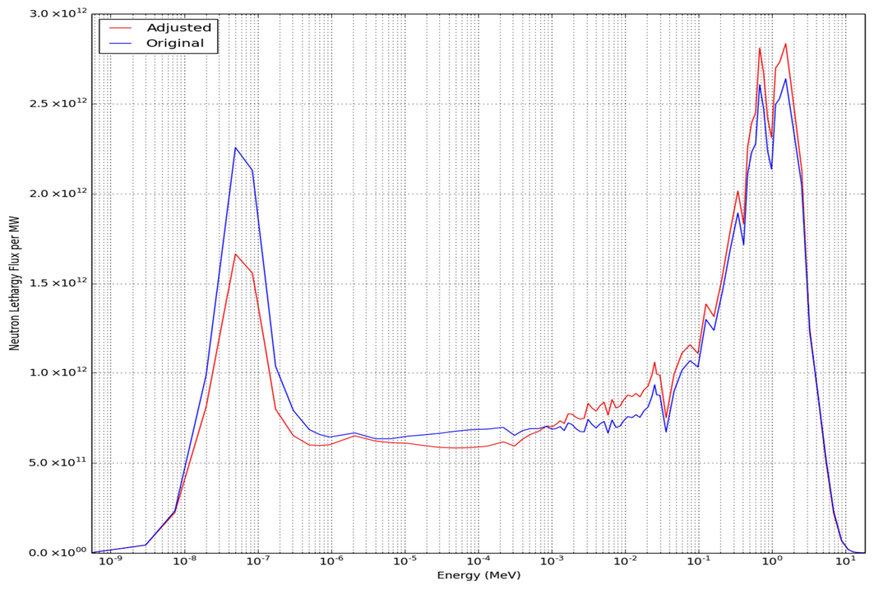

Figure 7. Final adjustment by the genetic algorithm of the PLG Spectrum

Table1. Comparison of reaction probabilities for LSL-M2 and the genetic algorithm.

Foil identification	LSL	Genetic	Measured	LSL % diff	Genetic % diff
al27a#-ml3x-bahl	3.736E-13	3.768E-13	3.774E-13	1.008%	0.171%
au197g#-dil3-bahl	2.439E-07	2.386E-07	2.404E-07	1.467%	0.756%
au197g#-dil3-cdhl	1.803E-07	1.777E-07	1.788E-07	0.841%	0.630%
co59g#-mil2-bahl	2.797E-08	2.720E-08	2.724E-08	2.694%	0.140%
co59g#-mil2-cdhl	5.391E-09	5.591E-09	5.329E-09	1.165%	4.918%
co59p#-mil2-cdhl	8.216E-13	8.014E-13	8.157E-13	0.723%	1.753%
co592#-mil2-cdhl	1.128E-13	1.112E-13	1.117E-13	0.989%	0.425%
cu63a#-mil5-bahl	3.014E-13	2.933E-13	3.596E-13	16.182%	18.426%
cu63g#-mil5-bahl	3.310E-09	3.220E-09	3.305E-09	0.162%	2.577%
cu63g#-mil5-cdhl	5.135E-10	5.490E-10	4.975E-10	3.216%	10.353%
fe54p#-mil5-bahl	5.032E-11	5.024E-11	4.970E-11	1.249%	1.089%
fe56p#-mil5-bahl	6.122E-13	5.970E-13	6.200E-13	1.255%	3.714%
fe58g#-mil5-bahl	9.783E-10	9.375E-10	9.335E-10	4.804%	0.426%
fe58g#-mil5-cdhl	1.548E-10	1.512E-10	1.493E-10	3.667%	1.304%
in115n#-mil5-bahl	1.545E-10	1.564E-10	1.651E-10	6.434%	5.274%
mg24p#-mil5-bahl	8.013E-13	8.042E-13	7.661E-13	4.592%	4.973%
mn55g#-mil2-cdhl	1.490E-09	1.436E-09	1.487E-09	0.209%	3.400%
mn552#-mil2-bahl	1.168E-13	1.151E-13	1.375E-13	15.045%	16.271%
mo98g#-mil5-bahl	8.740E-10	8.642E-10	8.667E-10	0.842%	0.284%
mo98g#-mil5-cdhl	7.940E-10	7.881E-10	7.982E-10	0.528%	1.271%
na23g#-pelt-bahl	3.242E-10	3.155E-10	3.057E-10	6.064%	3.219%
na23g#-pelt-cdhl	3.358E-11	3.547E-11	3.245E-11	3.469%	9.305%
nb932#-mil5-bahl	2.438E-13	2.435E-13	2.411E-13	1.124%	0.976%
ni58p#-milx-bahl **	6.879E-11	6.879E-11	6.879E-11	0.000%	0.000%
ni582#-milx-cdhl	2.293E-15	2.332E-15	2.152E-15	6.569%	8.382%
ni60p#-milx-cdhl	1.228E-12	1.189E-12	1.254E-12	2.040%	5.168%
rmldu#-rmld-fiss	2.250E-10	2.281E-10	2.194E-10	2.546%	3.958%
rmleu#-rmle-fiss	2.463E-09	2.495E-09	2.573E-09	4.257%	3.040%
rmlpu#-rmlp-fiss	2.796E-09	2.803E-09	2.570E-09	8.776%	9.072%
s32cf#-void-bare	5.845E-02	5.886E-02	5.437E-02	7.501%	8.262%
sc45g#-mil5-bahl	1.802E-08	1.730E-08	1.731E-08	4.087%	0.057%
sc45g#-mil5-cdhl	1.295E-09	1.314E-09	1.352E-09	4.208%	2.791%
ti46p#-milx-bahl	6.700E-12	6.509E-12	6.400E-12	4.684%	1.698%
ti47p#-milx-bahl	1.237E-11	1.243E-11	1.265E-11	2.188%	1.723%
ti48p#-milx-bahl	1.639E-13	1.628E-13	1.625E-13	0.839%	0.183%
zn64p#-milx-bahl	2.403E-11	2.408E-11	2.451E-11	1.940%	1.735%
zr902#-milx-bahl	5.670E-14	5.689E-14	5.748E-14	1.353%	1.023%

Discussion and Future Work

Despite the promising results shown above, the genetic algorithm presented here lacks one very important quality; it does not give any information about the uncertainty in the adjusted spectrum as LSL-M2 does. Unfortunately, the only way to get around this downfall would be to perform many adjustments given different input data, reflecting the uncertainty in the transport cross sections and other model parameters. A full run of the algorithm presented here of 500 generations takes less than 10 seconds in serial. Parallelization may make it feasible to perform several thousand adjustments; however, the largest computational cost would be in generating a large number of input spectra using MCNP with different cross section libraries. Quantities that can be varied easily within the code and without generating new input spectra are the measured activities and dosimetry cross sections, both of which have known uncertainties. Aside from parallelization, further work in code verification will need to be done. This will include examining the effect that the input ACRR trial spectrum has on the final adjusted spectrum. For instance, changing the HDPE density in the MCNP model would significantly change the trial spectrum. It is hoped that the genetic algorithm will arrive at the same adjusted spectrum shown in Figure 7. Code verification using known standards in radiation metrology will also be performed.

Acknowledgements

I would like to thank my manager Ken Reil, mentor Ed Parma, and Patrick Griffin for their support in this project. A great deal of freedom was given to explore this abstract computational method, and there were surely doubts along the way. I would also like to thank my previous organization, 1386: Nuclear Materials Management, for giving me my first opportunity to work at the lab. Finally, I would like to thank my computational physics professor at Texas A&M University, Dr. Helmut Katzgraber, for introducing me to the genetic algorithm (and other abstract computational methods that you just don't learn in undergraduate engineering courses). It has truly been a pleasure working at SNL for the past year.