A Generalized Optimization Methodology for Isotope Management

by

Mark Massie

B.S. Nuclear Engineering, 2008
University of Tennessee, Knoxville

Submitted to the Department of Nuclear Science and Engineering
in Partial Fulfillment of the Requirements for the Degree of

Master of Science in Nuclear Science and Engineering
at the
Massachusetts Institute of Technology

June 2010

Signature of Author: _____
Department of Nuclear Science and Engineering
May 21, 2010

Certified By: _____
Benoit Forget, Ph.D.
Assistant Professor of Nuclear Science and Engineering
Thesis Supervisor

Certified By: _____
Michael J. Driscoll, Ph.D.
Professor Emeritus of Nuclear Science and Engineering
Thesis Reader

Accepted By: _____
Jacquelyn C. Yanch, Ph.D.
Chair, Committee on Graduate Students

A Generalized Optimization Methodology for Isotope Management

by
Mark Massie

Submitted to the Department of Nuclear Science and Engineering on May 21, 2010
in Partial Fulfillment of the Requirements for the Degree of
Master of Science in Nuclear Science and Engineering

ABSTRACT

This research, funded by the Department of Energy's Advanced Fuel Cycle Initiative Fellowship, was focused on developing a new approach to studying the nuclear fuel cycle: instead of using the trial and error approach currently used in actinide management studies in which reactors are designed and then their performance is evaluated, the methodology developed here first identified relevant fuel cycle objectives like minimizing decay heat production in a repository, minimizing Pu-239 content in used fuel, etc. and then used optimization to determine the best way to reach these goals.

The first half of this research was devoted to identifying optimal flux spectra for irradiating used nuclear fuel from light water reactors to meet fuel cycle objectives like those mentioned above. This was accomplished by applying the simulated annealing optimization methodology to a simple matrix exponential depletion code written in Fortran using cross sections generated from the SCALE code system.

Since flux spectra cannot be shaped arbitrarily, the second half of this research applied the same methodology to material composition of fast reactor target assemblies to find optimal designs for minimizing the integrated decay heat production over various timescales. The neutronics calculations were performed using modules from SCALE and ERANOS, a French fast reactor transport code.

The results of this project showed that a thermal flux spectrum is much more effective for transmuting used nuclear fuel. In the spectral optimization study, it was found that a thermal flux spectrum is approximately five times more effective at reducing long-term decay heat production than a fast flux spectrum. This conclusion was reinforced by the results of the target assembly material optimization study, which found that by adding an efficient moderator to a target assembly designed for minor actinide transmutation, the amount of decay heat generated over 10,000 years of cooling can be reduced by over 50% through a single pass in a fast reactor without exceeding standard cladding fluence limits.

Thesis Supervisor: Benoit Forget
Title: Assistant Professor of Nuclear Science and Engineering

Table of Contents

LIST OF FIGURES

LIST OF TABLES

1 INTRODUCTION

As concern about global climate change has risen over the past decade to its present peak, the option of increased reliance on nuclear power as a carbon-free alternative to coal and natural gas has breathed new life into this once floundering industry. While some have called this the renaissance of nuclear power, the extent of growth may largely be determined by how well scientists, engineers, and politicians address concerns over the treatment of used nuclear fuel.

1.1 Current Fuel Cycle

The uranium that fuels today's light water nuclear reactors (LWRs) enters the fuel cycle in the mining process. Before enrichment, the fissile isotope U-235 constitutes only 0.72% of natural uranium isotopes. Since the concentration of U-235 required in fuel for current commercial LWRs is between 3-5%, uranium undergoes an enrichment processes that increases the concentration of U-235. This enriched uranium is then fabricated into fuel rods, which are arranged in fuel assemblies and placed in reactor cores to produced power.

While producing power, the concentration of U-235 and other isotopes changes over time due to fission, neutron capture, and radioactive decay. Eventually, the concentration of U-235 becomes too low to sustain the nuclear chain reaction so the depleted fuel is removed from the core and replaced with fresh fuel. While used nuclear fuel (UNF) is no longer suitable for use in LWRs, it still remains highly radioactive and therefore produces large amounts of heat. This requires storage of UNF in spent fuel pools and dry casks at reactors sites until moved to a permanent repository.

The final step in the United States' current once-through fuel cycle is the permanent storage of UNF in a geologic repository. Because UNF remains highly radioactive for hundreds of years and does not reach the radiotoxicity level of natural uranium for thousands of years, regulations require that repositories must be capable of shielding the environment from its contents for 10,000 years. A repository has been partially constructed at the Yucca Mountain site in Nevada but political and regulatory holdups have cast doubt upon its future. Even if it were to begin accepting UNF as planned in 2010, the repository would reach statutory capacity by 2015 because of the already vast reserve of UNF stored at reactors across the country and the approximately 2,100 MT added per year by the nation's fleet of existing LWRs [1].

1.2 Advanced Fuel Cycle

Although the current plan in the United States calls for direct disposal of UNF, this is not because the nuclear material no longer has value; it is because uranium remains inexpensive relative to estimates for reprocessing costs in the U.S. and because of political opposition to reprocessing because of its potential risk in contributing to nuclear weapons proliferation.

Reprocessing is a central part of an advanced fuel cycle as it allows recycling of UNF into new fuel for use in thermal or fast reactors, depending on the type of reprocessing. The inclusion of reprocessing into the fuel cycle results in more efficient use of UNF because after removal from a LWR, it is still composed of over 96% useable nuclear material. If uranium and plutonium are recycled into new fuel, only fission products and minor actinides require permanent disposal, thus reducing volume of material requiring permanent storage by an order of magnitude [1].

The capacity of a repository can be further increased and the radiological hazard of its contents significantly reduced through the use of transmutation, the artificial transformation of unfavorable isotopes into less harmful or shorter-lived nuclides through neutron irradiation.

This already occurs in reactors when plutonium and other elements are created from uranium through neutron capture or neutron-induced fission. Transmutation has also been specifically incorporated into some reactor designs so that they create, or breed, more fuel than they consume through the transmutation of U-238 into Pu-239.

Yet, the use of transmutation does not have to be limited to the production of fuel material; by irradiating used nuclear fuel, one can reduce the amount of long-lived, radiotoxic nuclides sent to a geologic repository and thus decrease the time required for the waste to reach natural levels of radiation. Additionally, by eliminating the isotopes that generate the greatest amount of heat, one could increase the storage capacity of repositories limited by heat considerations. Transmutation could also be used in the same way to make reprocessing easier by lowering the radiotoxicity of used nuclear fuel. The risks of nuclear proliferation can also be reduced by transmuting those isotopes critical to nuclear weapons into their less dangerous counterparts. By studying transmutation, one can determine how these objectives can best be achieved and incorporate these findings in reactor design and studies of the fuel cycle to utilize nuclear material more efficiently and minimize the burden of used nuclear fuel.

1.3 Thesis Objectives and Organization

The purpose of this research is to provide a quantitative optimization of transmutation of used nuclear fuel based on the physics of depletion by identifying optimal flux spectra and material compositions for various fuel cycle applications. The objectives that are important to the nuclear fuel cycle will be identified and then optimization will used to determine the best way to reach those goals.

This thesis is divided into two studies, spectral optimization and material optimization. The first study seeks to identify optimal neutron energy distributions, or spectra, to meet various fuel

cycle objectives such as minimizing decay heat production in a geologic repository to increase storage capacity or reducing the proliferation risks of UNF by minimizing Pu-239 content. This study is presented in chapter 2. The second study focuses on identifying material combinations that enable optimal transmutation of UNF to meet certain fuel cycle objectives. This study is discussed in chapter 3. Chapters 2 and 3 include the methodology used in each study as well as their respective results. Chapter 4 summarizes the results of the spectral and material optimization studies and presents the conclusions drawn from this research.

2 SPECTRAL OPTIMIZATION

2.1 Introduction

The effectiveness of any transmutation scheme is intimately linked with the energy distribution of the neutrons incident upon a target material. This link is derived from the energy dependence of a target material's cross sections, which illustrates how the probability of occurrence for given reaction changes as a function of incident neutron energy. Unfortunately, energy dependent cross sections vary wildly from isotope to isotope, and cross sections for various interactions for a single isotope can be markedly dissimilar. To complicate things even further, the cross sections for a single type of interaction for a given isotope do not vary smoothly with energy; instead, they are marred with sharp peaks and valleys, called resonances, which arise from quantum interactions between the target nuclei and the incident neutrons. When multiple isotopes are present in a sample, the cross sections from one isotope alter the overall neutron energy distribution and therefore change the reaction rates in all of the other isotopes, thus intertwining the energy dependence of all isotopes in a sample.

Two of these energy-dependent interactions, neutron capture and neutron-induced fission, compete with radioactive decay to determine the time-dependent concentration of nuclides in a sample. Without an incident neutron flux, a sample's isotopic composition is determined solely by radioactive decay but if the sample is exposed to neutron radiation, neutron capture and fission begin competing with radioactive decay to determine time-dependent nuclide concentrations within the material. This competition is what introduces the energy dependence into the rate of transmutation. Since the rate at which neutron capture and fission occur is determine by their respective energy-dependent cross sections, it is ultimately the energy distribution, or spectrum, of the incident neutrons that determines which interaction controls

the time-dependent composition of a given sample. In turn, the time-dependent concentration of nuclides within a sample is what determines the many physical characteristics that dictate important fuel cycle parameters such as decay heat production, radiological hazard, ease of reprocessing, and proliferation resistance. It is because of this fundamental link between neutron energy distribution and time-dependent nuclide composition and associated physical characteristics that this project begins with a study on determining optimal energy spectra for transmuting used nuclear fuel to meet various fuel cycle goals.

2.2 General Methodology

The purpose of this section of the project is to create and implement a quantitative optimization routine that determines the optimal flux spectrum for various fuel cycle applications. In order to accomplish this, the first step is to define objectives that are important to the nuclear fuel cycle. From there, a neutronics code needs to be developed that evaluates performance for each fuel cycle objective based on the incident neutron spectrum. The final step is to apply an optimization scheme that iterates over the neutronics calculations to determine an optimal energy spectrum for each fuel cycle objective. These steps will be discussed in detail in the following sections, but before explaining the solution methodology, the problem must be given some context.

2.3 Case Study Description

Although the general methodology described here could be applied to any transmutation problem, this project focuses on the transmutation of light water reactor (LWR) used nuclear fuel (UNF) because of its abundance and lack of permanent disposal solution. In accordance

with the goals of the Advanced Fuel Cycle Initiative, this study investigates how the United States' supply of UNF can be transmuted to reduce the long-term radiological impact of waste, improve geological repository efficiency, and reduce proliferation risks [2]. Therefore, the case study will be the irradiation of typical LWR UNF in a reactor with flux spectrum constraints similar to those of current reactor designs (both fast and thermal reactors) for a residence time similar to that of an LWR fuel assembly and a permanent storage time on the same timescale as proposed geologic repositories such as Yucca Mountain. The exact case study parameters are discussed in more detail below.

2.3.1 Used Nuclear Fuel Description

The material to be transmuted in this case study is LWR UNF. The UNF isotopic vector was calculated by the ORIGEN-ARP graphical user interface, which is a SCALE analytical sequence that solves for time-dependent material concentrations using the ORIGEN-S depletion code using pre-computed cross section sets for common reactor designs [3]. In this case, a Westinghouse 17x17 assembly normalized to 1 metric ton of uranium with an initial enrichment of 3% was depleted to 50 GWd/MTHM and the isotopic concentrations from the ORIGEN output file were used to calculate the weight percent (wt%) for each isotope in the used fuel. This vector, which is shown in TABLE 2.1 below, is later used as the initial relative isotopic concentrations in the UNF that is to be transmuted.

TABLE 2.1 - Used Nuclear Fuel Isotopic Composition

Isotope	wt %
U234	1.84E-02
U235	7.46E-01
U236	6.05E-01
U238	9.73E+01
Np237	7.59E-02
Pu236	1.00E-10
Pu238	3.50E-02
Pu239	6.33E-01
Pu240	3.10E-01
Pu241	1.41E-01
Pu242	9.61E-02
Am241	4.50E-02
Am242	1.38E-04
Am243	2.61E-02
Cm242	1.41E-06
Cm243	7.40E-05
Cm244	8.80E-03
Cm245	5.23E-04
Cm246	6.76E-05
Cm247	1.07E-06
Cm248	7.74E-08
Bk249	1.00E-10
Cf249	1.08E-09
Cf250	3.51E-10
Cf251	1.85E-10
Cf252	3.41E-11

The model used in this spectral optimization study simulates the neutron irradiation of 1 metric ton of UNF, but it should be noted that the mass of the UNF sample does not affect any results because the depletion calculations that are used later are zero-dimensional and all cost functions are normalized to the starting mass of UNF.

2.3.2 Constraints

Constraints were placed on several of the transmutation parameters in order to make this a more realistic simulation. Upper and lower bounds were placed on the allowable flux values for each energy group of incident neutron radiation to limit the search space to flux levels that could reasonably be achieved in current reactor designs. The maximum flux for a given energy group was set 10^{14} n/cm^2-sec and the minimum was set at 10^8 n/cm^2-sec. These values were chosen by examining typical energy spectra for fast and thermal reactors, which are shown in FIGURE 2.1 below.

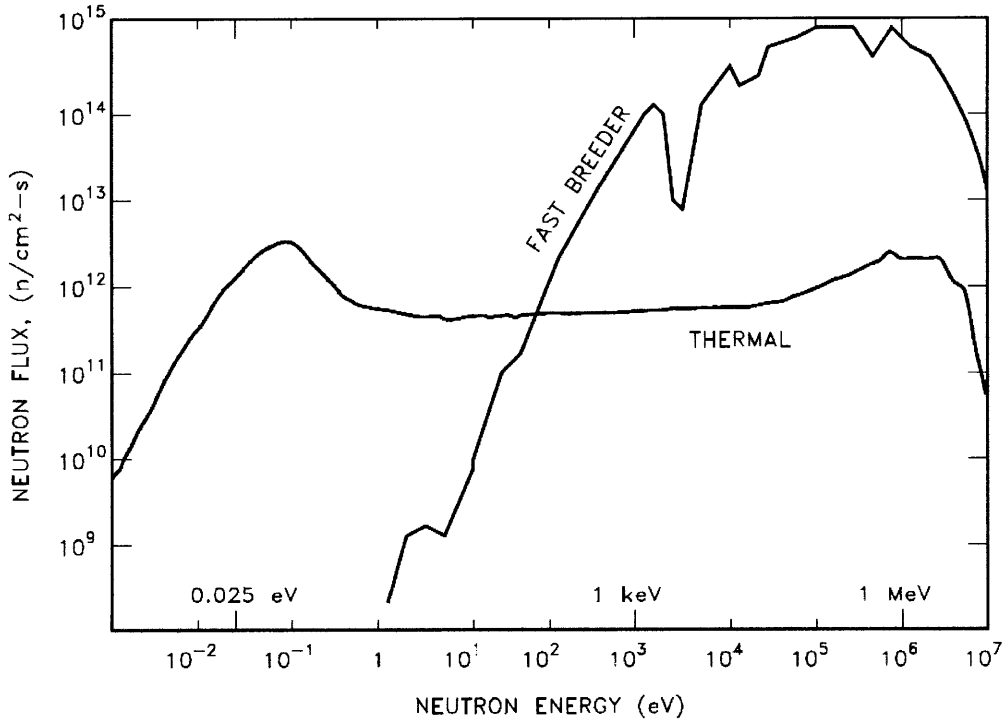

FIGURE 2.1 – Neutron energy spectra for fast and thermal reactors [4]

No constraints were placed on the total flux because the objective of this study was to find optimal flux spectra independent of reactor design and if the total flux were set to a specific

value, large peaks in the spectrum would cause depressions in other energy groups that may not be beneficial to minimizing the given cost function. Instead, by limiting group flux values, the search space is constrained to that of current reactor designs while still being able to produce optimal flux spectra with unnecessary constraints.

In this study, the UNF sample is irradiated for 1,000 days. Although the method developed here for finding optimal flux spectra can be applied to any length of irradiation time, an exposure time of 1,000 days was chosen. This is along the same order of magnitude of typical residence times of assemblies in a light water reactor.

After being irradiated for 1,000 days, a decay-only depletion case is run on the transmuted UNF to calculate nuclide concentrations over the course of 10,000 years. This 10,000 year period is designed to simulate storage of UNF in a permanent geologic repository. This length of cooling time was chosen because the Environmental Protection Agency's (EPA) radiation protection standards for the Yucca Mountain repository were originally to be in effect for 10,000 years after the site was closed, but even though those standards are now extended to 100,000 years after site closing, essentially all measured parameters in this study have converged by 10,000 years and so extending the depletion calculation through 100,000 years would only increase computation time without yielding any additional results [6]. Additionally, the Nuclear Regulatory Commission's (NRC) regulations on Yucca Mountain are only applicable through 10,000 years after repository closure [7].

2.4 Neutronic Methodology

There are two basic neutronics calculations that need to be preformed for this spectral optimization study. The first step is to generate cross sections for the materials that are to be irradiated, and the second step is to perform depletion and decay calculations to generate time-

dependent nuclide concentrations. For the purposes of this study, it is assumed that cross sections remain constant throughout the irradiation period, cross section generation is only performed once, while depletion calculations are completed during each iteration of the optimization routine. The details of each of these calculations are discussed below.

2.4.1 Nuclear Data Generation

In order to perform the depletion calculations that are described below, one first has to generate the cross sections data that will be used in these subsequent calculations. Cross section generation is performed using the SCALE code package system developed at Oak Ridge National Laboratory. Because cross section generation was completed before the release of SCALE 6.0 in February 2009, the following descriptions refer to the then-current SCALE 5.1 [8]. This code is used to perform an infinite, homogenous medium calculation on the UNF vector described in TABLE 2.1 to generate energy self-shielded, multigroup cross sections for use in depletion calculations.

As mentioned above, this study assumes that cross sections do not change with burnup. This is not entirely true in reality because the changes in isotopic concentrations that occur during irradiation alter the flux spectrum, which affects the effective cross sections. However, if the total effective cross section changes very little over time, this approximation may still be valid. This will be investigated in future work on this project by recalculating cross sections every few burnup steps and comparing the time-dependent total cross section to the initial total cross section. It should also be noted that while energy self-shielding corrections are performed, the effects of spatial self-shielding are neglected during cross section generation in the following methodology. These effects could alter resulting reaction rates, especially in the thermal energy range.

Within SCALE, the TRITON control module is used to call, in order, the functional modules BONAMI, WORKER, CENTRM, PMC, and NEWT. BONAMI performs Bondarenko calculations on master library cross sections to account for energy self-shielding effects; WORKER simply formats and passes data between other modules; CENTRM uses both pointwise and multigroup nuclear data to compute a continuous energy neutron flux by solving the Boltzmann transport equation using discrete ordinates; PMC takes the continuous energy neutron flux from CENTRM and calculates group-averaged cross sections; and NEWT collapses the cross sections from the master library energy group structure to a user-defined structure. After completing the TRITON sequence, an additional module, PALEALE, is called to convert cross section data for (n,fission), (n,gamma), and (n,2n) reactions from the binary working library file into text that can be imported into the depletion code.

The ENDF/B-VI, 238 energy group master library is used in the above calculation sequence but is collapsed to a user-defined 54 group structure. A smaller (~60) energy group structure is more suitable for the optimization method for reasons described in section 2.5. Initially, a 60 energy group structure was to be created by dividing the energy range covered in the 238 group structure (20 MeV – $1x10^{-5}$ eV) into 60 energy groups, each with equal lethargy. However, SCALE requires the energy boundaries of user-defined structures to be a subset of the boundaries of the initial fine group structure, so several of the very low energy groups were combined to meet this requirement, thus resulting in a 54 group structure [3].

The depletion calculations require additional nuclear data in the form of radioactive decay constants, decay yields, and branching ratios. The decay constants were collected from the Chart of Nuclides, while the rest of this information was acquired from a data sheet compiled from the ORIGEN 2.2 code, a predecessor of the ORIGEN-S module that is included in the SCALE code system [8, 9].

2.4.2 Depletion Calculations

A basic depletion code was written in the Fortran programming language. This code uses the matrix exponential expansion method of solving the rate equations that describe depletion and decay processes. Although this is the same method used in the SCALE depletion module, ORIGEN-S, a custom code was written because of speed considerations. Since this code needs to be called at each iteration of the optimization routine, the speed of the depletion calculation determines the overall speed of the optimization program. So by limiting the isotopes tracked to the actinides and simplifying decay chains by skipping short-lived nuclides, the custom code becomes drastically faster than a similar ORIGEN calculation. The list of nuclides tracked in this code is shown in TABLE 2.2 below.

TABLE 2.2 - Isotopes tracked in depletion calculation

U234
U235
U236
U238
Np237
Pu236
Pu238
Pu239
Pu240
Pu241
Pu242
Am241
Am242
Am243
Cm242
Cm243
Cm244
Cm245
Cm246
Cm247
Cm248
Bk249
Cf249
Cf250
Cf251
Cf252

As mentioned above, this code uses the matrix exponential expansion method to solve the system of rate equations that describe depletion and decay phenomena. These rate equations are shown as the time-dependent concentrations of isotope i (N_i) in EQUATION 2.1,

EQUATION 2.1

$$\frac{dN_i}{dt} = \sum_j \left[\gamma_{ji}\sigma_{f,j}(E)N_j(t)\phi(E) + \lambda_{ji}N_j(t)\right] + \sigma_{c,i-1}(E)N_{i-1}(t)\phi(E)$$

$$-\sigma_{f,i}(E)N_i(t)\phi(E) - \sigma_{c,i}(E)N_i(t)\phi(E) - \lambda_i N_i(t)$$

where (i=1,...,I) and

$\gamma_{ji}\sigma_{f,j}(E)N_j(t)\phi(E)$ is the yield rate of N_i due to the fission of all nuclide N_j;

$\lambda_{ji}N_j(t)$ is the rate of formation of N_i due to the decay of all nuclides N_j;

$\sigma_{c,i-1}(E)N_{i-1}(t)\phi(E)$ is the rate of transmutation into N_i due to neutron capture by N_{i-1};

$\sigma_{f,i}(E)N_i(t)\phi(E)$ is the destruction rate of N_i due to fission;

$\sigma_{c,i}(E)N_i(t)\phi(E)$ is the destruction rate of N_i due to neutron capture;

$\lambda_i N_i(t)$ is the destruction rate of N_i due to radioactive decay.

Spatial dependence is removed from concentrations by assuming a homogenous mixture of UNF and neutrons, and energy dependence is discretized so that neutron flux $\phi(E)$ is averaged into energy groups ϕ_k where k=1 is the highest energy group and k=K is the lowest energy group. If it is assumed that flux and cross sections σ remain constant over a time step Δt, then EQUATION 2.1 can be treated as a system of first-order, linear differential equations.

This system can be solved using the matrix exponential method. When written in matrix notation where A is the transition matrix containing the first order rate constants for the creation of nuclides i from nuclides j, EQUATION 2.1 has the known solution

EQUATION 2.2

$$\bar{N}(t) = \exp(\bar{\bar{A}}t)\,\bar{N}(0)$$

The matrix exponential function is defined as

EQUATION 2.3

$$\exp(\bar{\bar{A}}t) \equiv \bar{\bar{I}} + \bar{\bar{A}}t + \frac{(\bar{\bar{A}}t)^2}{2!} + \cdots = \sum_{m=0}^{\infty} \frac{(\bar{\bar{A}}t)^2}{m!}$$

The time dependent solution can be written in terms of the rate constants a_{ij} from the transition matrix A as

EQUATION 2.4

$$N_i(t) = N_i(0) + t\sum_j a_{ij} N_j(0) + \frac{t}{2}\sum_k \left[a_{ik} t \sum_j a_{kj} N_j(0) \right] + \frac{t}{3}\sum_m \left\{ a_{im} \frac{t}{2} \sum_k \left[a_{mk} t \sum_j a_{kj} N_j(0) \right] \right\} + \cdots$$

A recursion relation can be developed from EQUATION 2.4 that yields the terms

EQUATION 2.5

$$C_i^0 \equiv N_i(0)$$

EQUATION 2.6

$$C_i^{n+1} \equiv \frac{t}{n+1} \sum_j a_{ij} C_j^n$$

which requires storing only the C^n and C^{n+1} vectors, thus greatly reducing memory usage. The solution for the time-dependent concentration of nuclide i is then given by

EQUATION 2.7

$$N_i = \sum_{n=0}^{\infty} C_i^n$$

While EQUATION 2.7 yields the exact solution, an approximate solution can be calculated by truncating the infinite series once the inclusion of an additional term does not alter the nuclide

concentration [3]. For the precision necessary for this simplistic depletion study, EQUATION 2.4 was truncated after the third term.

2.5 Global Optimization Methodology

In order to find the best flux spectrum for a given application, an optimization scheme needs to be identified that can match the mathematical peculiarities presented by this problem. Many optimization methods cannot be used here because of the need to find global, not local, minima combined with the high non-linearity of the energy dependence of transmutation rates. Since the search space is so large (discussed further in section 2.5.2), the brute force method of evaluating every spectrum's effectiveness is unreasonable because of the massive computational time required. These considerations necessitate the use of a heuristic optimization algorithm. The simulated annealing method was selected from this class of routines, which includes genetic algorithms, graduated optimization, and hill climbing optimization, because of its simplicity and ability to find global minima. The following subsections will discuss the general simulated annealing routine and its application to spectral optimization.

2.5.1 Simulated Annealing

The simulated annealing optimization method gets its name from its mimicry of the physical process of metallurgical annealing, which involves heating up a solid so that its constituent parts have freedom to rearrange themselves and then slowing cooling them so they can achieve an arrangement that minimizes the energy of the system. In this analogy, the atoms of a metal represent segments of the solution and the energy of the system represents the cost functions that evaluate the performance of a given solution. In the physical process, the metallic solid, is

heated above its recrystalization temperature so that the atoms are no longer bound to each other in a rigid formation and are instead allowed to rearrange themselves freely. In simulated annealing, this represents the ability of the solution to change freely, even if it results in a drastic increase in the cost function. After reaching this high temperature state, the metal is slowly cooled down and as the temperature decreases, the constituent atoms rearrange themselves to minimize the energy of the system. As the solid approaches the recrystalization temperature, the atoms' movements become increasingly restricted and eventually settle upon a new crystalline structure. Analogously, as the optimization routine progresses, the probability of accepting solutions that increase the cost function continually decreases, eventually converging on a minimum. This analogy is clearly shown in FIGURE 2.2 below.

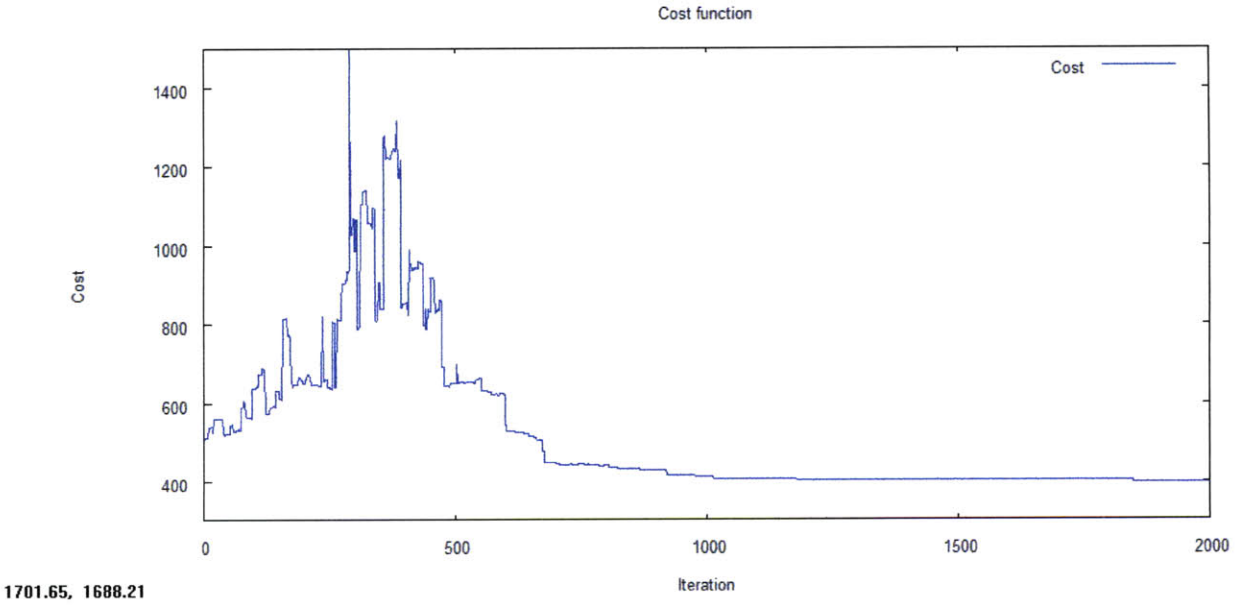

FIGURE 2.2 - Convergence using simulated annealing

Notice that the previous paragraph does not specify whether the minimum is local or global. This is because if the metal is cooled too quickly, its atoms do not have enough time to rearrange themselves in the structure that minimizes the energy of the system just as the optimization

routine has not sampled enough of the search space to converge upon the global minimum and instead converges upon a local minimum. Conversely, if the metal is cooled infinitely slowly, its atoms are guaranteed to find the arrangement that perfectly minimizes the system, just as the optimization routine is guaranteed to converge upon the global minimum given an infinite number of iterations. In the application of the simulated annealing method, this results in a tradeoff between accuracy and computational time that has to be balanced by the user.

This analogy is easily translated into a general optimization routine. First, one makes an initial guess at the solution and evaluates the subsequent cost value. Next, a neighbor function makes a new solution by altering the initial solution and the new solutions cost is evaluated. If the new cost is lower than the previous lowest cost, then both the solution and cost are saved. If the new cost is higher than the previous lowest cost, a random number is generated and compared to an acceptance probability function based on the temperature of the system; if the random number is higher than the acceptance probability, then even though it has a higher cost, both the solution and cost are saved and vice versa. After checking the acceptance of the new solution, the temperature is decreased and convergence is checked. If the solution has converged, then the optimization stops; otherwise, iteration continues by generating a new solution from the previous saved solution [10]. This process is presented in a much clearer manner in FIGURE 2.3 below. The details of these steps are presented in the following subsection.

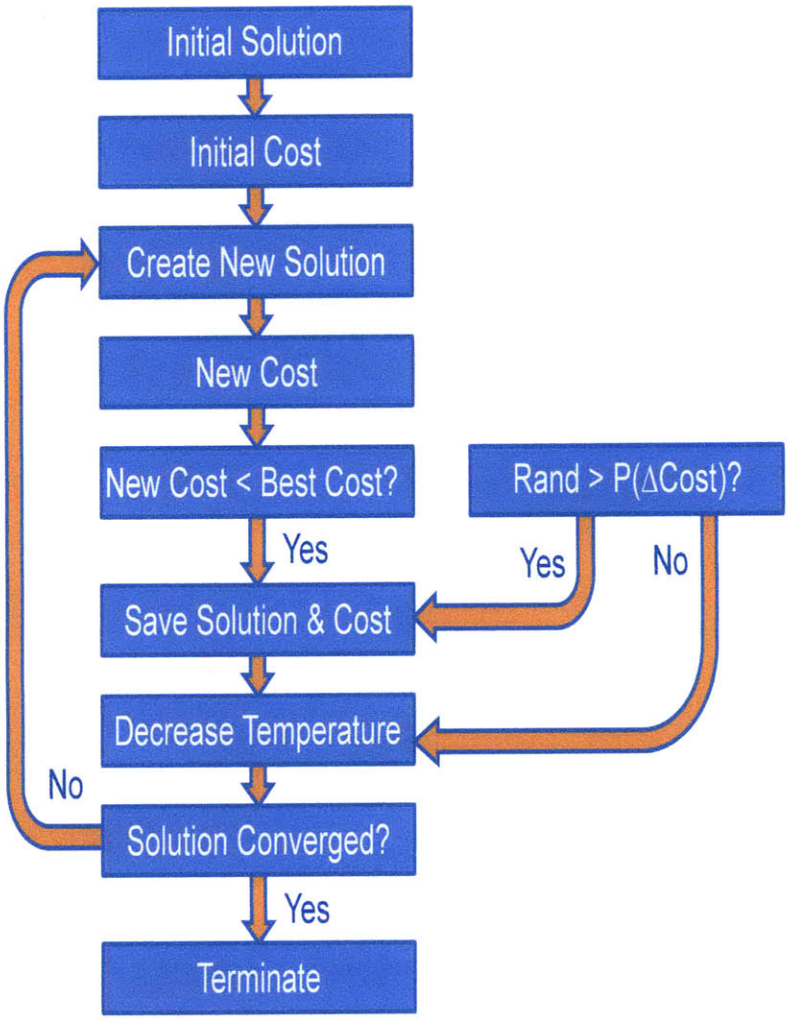

FIGURE 2.3 – General simulated annealing routine

2.5.2 Application to Spectral Optimization

In the application of the general simulated annealing optimization methodology, the mathematical solution mentioned in the above discussion is the neutron flux spectrum and the energy of the system is the cost function that is to be minimized. Cost functions will be further described in section 2.6, but for current discussion, it should be noted that these cost functions can be anything that can be calculated from the time-dependent nuclide concentrations output from the depletion calculations (decay heat, plutonium content, etc.).

The flux spectrum, which is the mathematical solution in this problem, is discretized into 54 energy groups for reasons discussed below and the flux values within each energy group are also discretized logarithmically between the upper and lower flux bounds of 10^{8}-10^{14} n/cm^{2}-s. This flux discretization allows the use of a very simple neighbor function while still maintaining meaningful spacing between flux levels.

As mentioned above, the energy spectrum is discretized into 54 energy groups. This structure was chosen because when finer structures are used (like the standard 238 group structure in SCALE), the optimization routine becomes very inefficient as it tries to account for cross section peaks and valleys in the resolved resonance range. By reducing the number of energy groups, the effects of resonance are still accounted for because they are flux-averaged over the energy group and only the general shape of neutron is optimized. This is all that is required in this spectral optimization study because it would be impossible to create a flux spectrum with many narrowly spaced peaks, while a very general flux shape may be possible to replicate by combining target materials with different cross section features.

The description of the application of simulated annealing to spectral optimization will follow the flow chart in FIGURE 2.3. Therefore, the first step is to define the initial guess solution. This, like many of the other optimization parameters, cannot be determined *a priori* and must be chosen based on a trial and error approach. This study uses flat spectra with all energy group fluxes equal to 10^{10} n/cm^{2}-s for initial solutions. It should be noted that almost all initial guesses will allow convergence to global minima except those that include initial group fluxes equal to the upper or lower bounds. This initial solution is used as an input to the depletion calculations described in section 2.4.2 and this outputs time-dependent nuclide concentrations, which are then used to calculate the value of the cost function. The initial solution and cost are set as the saved solution and cost and then a neighbor function is called to generate a new solution from the saved solution.

The neighbor function is another problem-specific aspect of simulated annealing implementation that must be determined through trial and error. This led to the implementation of a neighbor function that randomly alters the flux value of 1, 2, or 3 energy groups by a factor of 10 or 100. The flux of a selected energy group can be increased or decreased as long as the new value is still within the bounds of 10^8-10^{14} n/cm2-s. Each energy group's flux level is independent and is free of physical and geometrical constraints. This is not necessarily realistic because energy group fluxes are not entirely independent because of self-shielding effects and some flux spectra may not be able to be reproduced in physical systems. An initial guess solution and the new solution that is generated from it through the neighbor function are shown in FIGURE 2.4 and FIGURE 2.5 below.

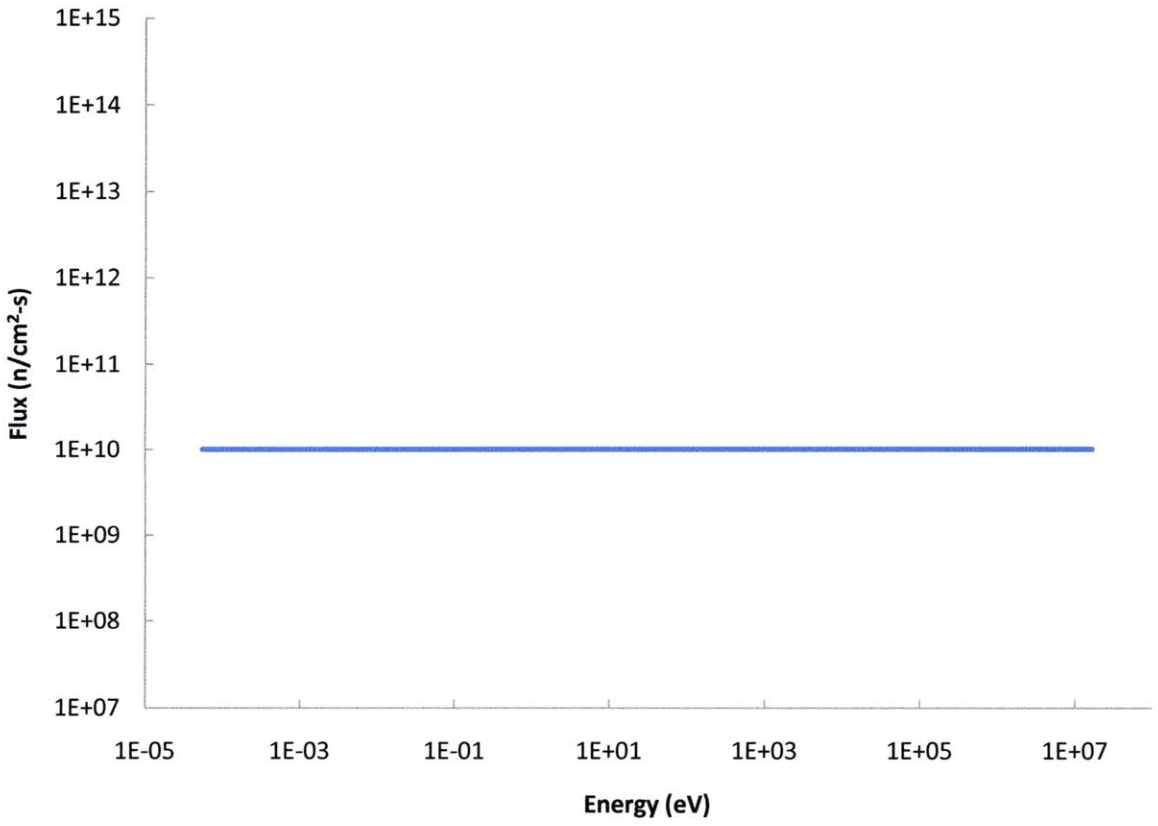

FIGURE 2.4 - Initial flux spectrum solution

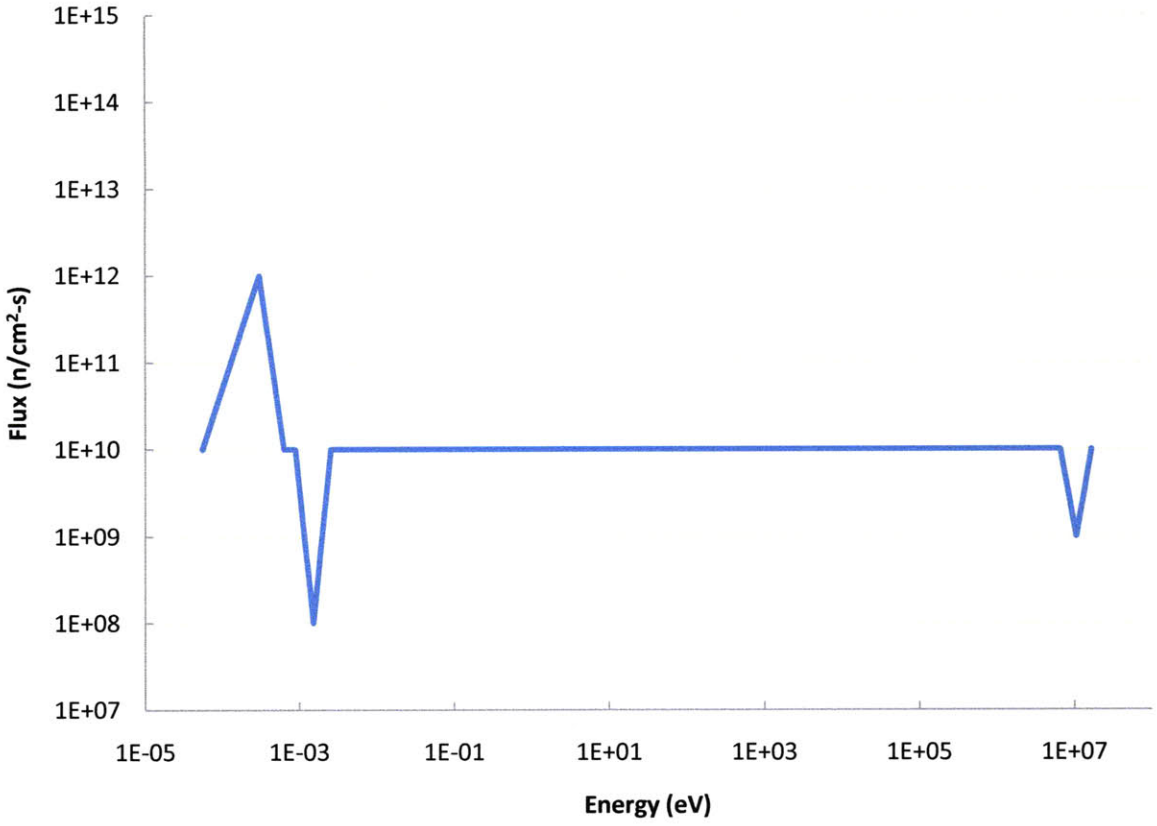

FIGURE 2.5 - New solution created by invoking the neighbor function

The decision to save the new solution or keep the previous solution is determined by an acceptance function, which determines the probability of saving a new solution even if it has a higher cost than the previous solution. This probability is decided by the increase in cost from the previous saved value, ΔC, and the "temperature" of the system, T. The equation for the acceptance probability is

EQUATION 2.8 - Acceptance probability [10]

$$P(\Delta C) = e^{-\Delta C/T}$$

which is analogous to the probability of an increase in system energy in physics,

EQUATION 2.9 - Probability of physical energy increase [10]

$$P(\Delta E) = e^{-\Delta E/kT}$$

where E is the energy of the system and k is the Boltzmann constant [10]. This is applied by calling a random number in the Fortran optimization code and if the random number is greater than $P(\Delta C)$ then the solution is saved even though it increases cost and if the random number is less than $P(\Delta C)$, then the solution and its associated cost are discarded.

Up to this point, the temperature of the optimization routine has remained relatively unclear. This is because besides being an indication of the number of completed iterations, the temperature is an almost entirely arbitrary variable. Like many of the other parameters in simulated annealing, this one can only be determined through trial and error but while the other somewhat arbitrary values can be set using a level of intuition, this parameter has no physical relation to the problem being solved.

The temperature of the system is dictated by a cooling schedule, which sets the initial temperature and the amount the temperature is decreased after each iteration. In this study, the initial temperature is set to 1,000 and it is decreased by 2.5% after each iteration. This cooling schedule results in a convenient balance between accuracy of solution and speed of convergence for this particular problem.

The final detail in the implementation of simulated annealing to spectral optimization is the convergence criteria. This, like temperature, is completely user-determined and in this study,

the solution is considered converged when the cost has not changed over the last 500 iterations and the current cost is equal to the best saved cost.

2.6 Cost Functions

One of the overarching goals of this project is to apply the general optimization methodology for isotope management developed here to aid the Advanced Fuel Cycle Initiative in advancing reprocessing knowledge and technology to allow the United States to transition to a closed fuel cycle. Therefore, although this optimization technique could be applied to many different areas in isotope management like medical isotope production, space-nuclear isotope production, fission product destruction, etc., the major optimization goals of this study will focus on the AFCI's objectives of reducing the long-term radiological impact of waste, developing a cheaper used fuel repository, and reducing proliferation risks.

The first cost function to be discussed is the heat generation rate of used nuclear fuel at various times after irradiation. Since decay heat production rate is, on a very simplistic level, an indicator of overall radioactivity, its evaluation at different post-irradiation times allows fundamental insight into the relationship between flux spectra and effectiveness of transmuting short, medium, and long-lived isotopes. This cost function is also important to reprocessing applications, because it can be used to minimize heat and thus ease handling of the fuel for immediate reprocessing after irradiation, after 10 years in a spent fuel pool, after 100 years in dry cask storage, etc. This cost function is evaluated by multiplying the mass of each nuclide at a certain time step by its specific heat generation rate.

The second cost function is the integrated heat production over a specified time interval. This is calculated by approximating the integral of decay heat generation rate over the specific time. More specifically, the decay heat rate for all isotopes is calculated at each time step and then the

trapezoidal rule is used to approximate the integral of these values. The main purpose of this cost function is to minimize total heat production over very long time periods, which would allow increased storage capacity at heat-limited repositories and thus reduce the cost per unit mass of storing used fuel. This cost function is also applicable to reducing the heat over a specific time period to allow easier reprocessing and, conversely, maximizing this cost could be used to increase proliferation resistance since the used fuel would be more difficult to handle.

The third major cost function to be analyzed is the isotopic composition of the plutonium vector after irradiation. Since the isotope of interest in terms of nonproliferation is Pu-239 because its small critical mass and low spontaneous fission rate make it an ideal material for nuclear weapons. Since these attributes are significantly less ideal for Pu-238 and Pu-240, increasing their concentrations relative to Pu-239 makes used fuel less attractive as weapons material. Therefore, two cost functions are used: one that minimizes the concentration of Pu-239 and one that maximizes Pu-238 and Pu-240 concentrations.

2.7 Results

2.7.1 Decay Heat Generation Rate

The optimization methodology developed in this study was first applied to minimizing decay heat generation rate immediately after irradiation (0 years of cooling), after 100 years of cooling, and after 1,000 years of cooling. The results of the optimized solutions are compared to the cost function generated by the natural decay of UNF. The decay-only reference case was created by running the same irradiation-then-decay calculations described in section 2.4 but with all energy group flux values set to zero. This ensures that the cost functions are evaluated at the same time for both the reference and optimized cases since the decay-only case is "irradiated"

with zero flux for 1,000 days before the cooling period begins. The five largest contributors to decay heat generation rate in unirradiated UNF are shown in FIGURE 2.6 below.

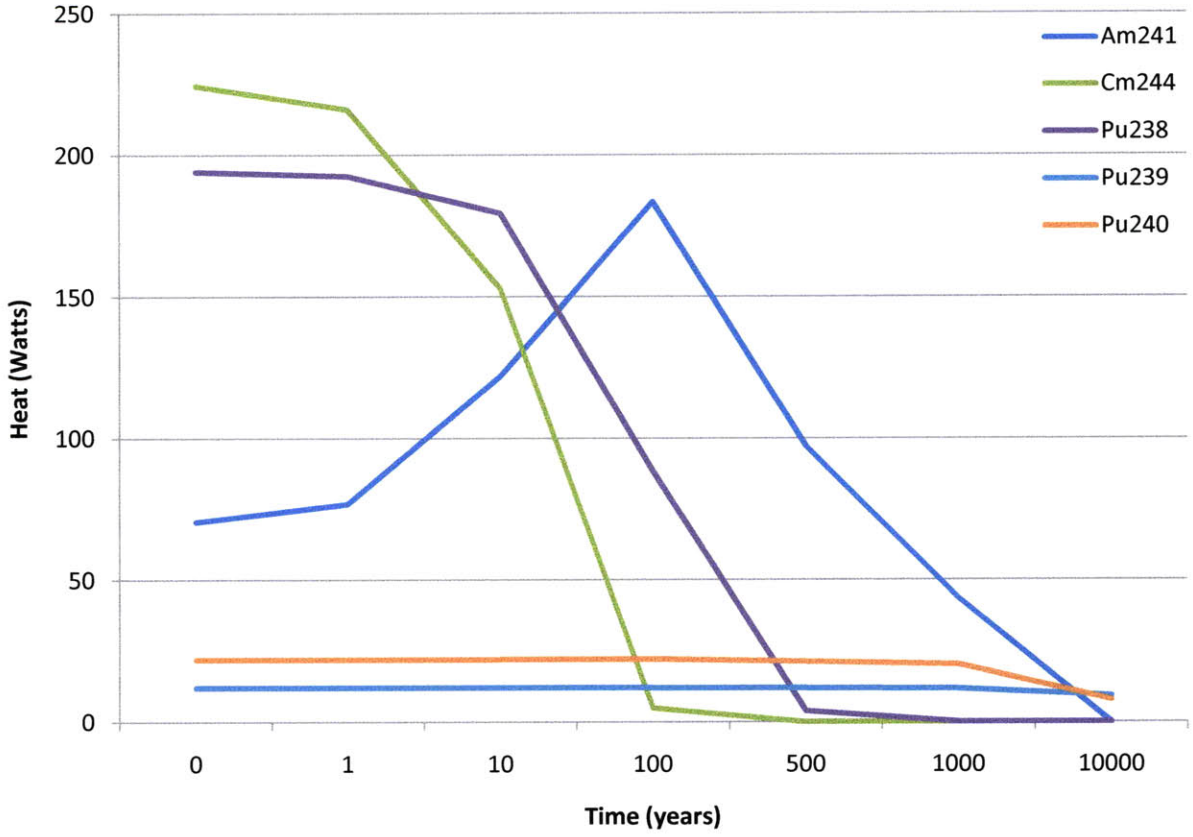

FIGURE 2.6 - Decay heat generation in unirradiated UNF

The first of the optimization cases to be examined is minimizing decay heat generation rate immediately after irradiation. As is shown in TABLE 2.3 and FIGURE 2.7, irradiating UNF with the optimal flux spectrum results in a 9.3% reduction in heat rate immediately after irradiation. The flux spectrum that causes this 9.3% reduction is shown in FIGURE 2.8 and its effects on the major contributors to decay heat generation rate are shown in FIGURE 2.9.

TABLE 2.3 - Decay heat rate after zero years cooling results

	UNF Decay	Optimized
Watts	530.63	481.31
% Change	---	-9.29

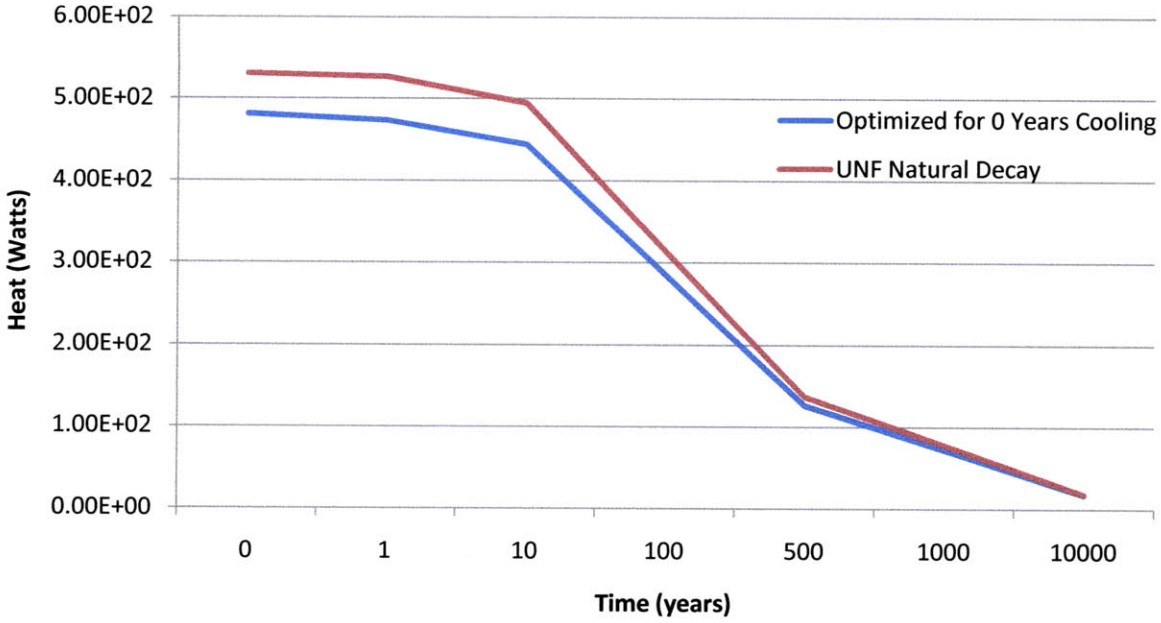

FIGURE 2.7 - Decay heat rate optimized for zero years cooling

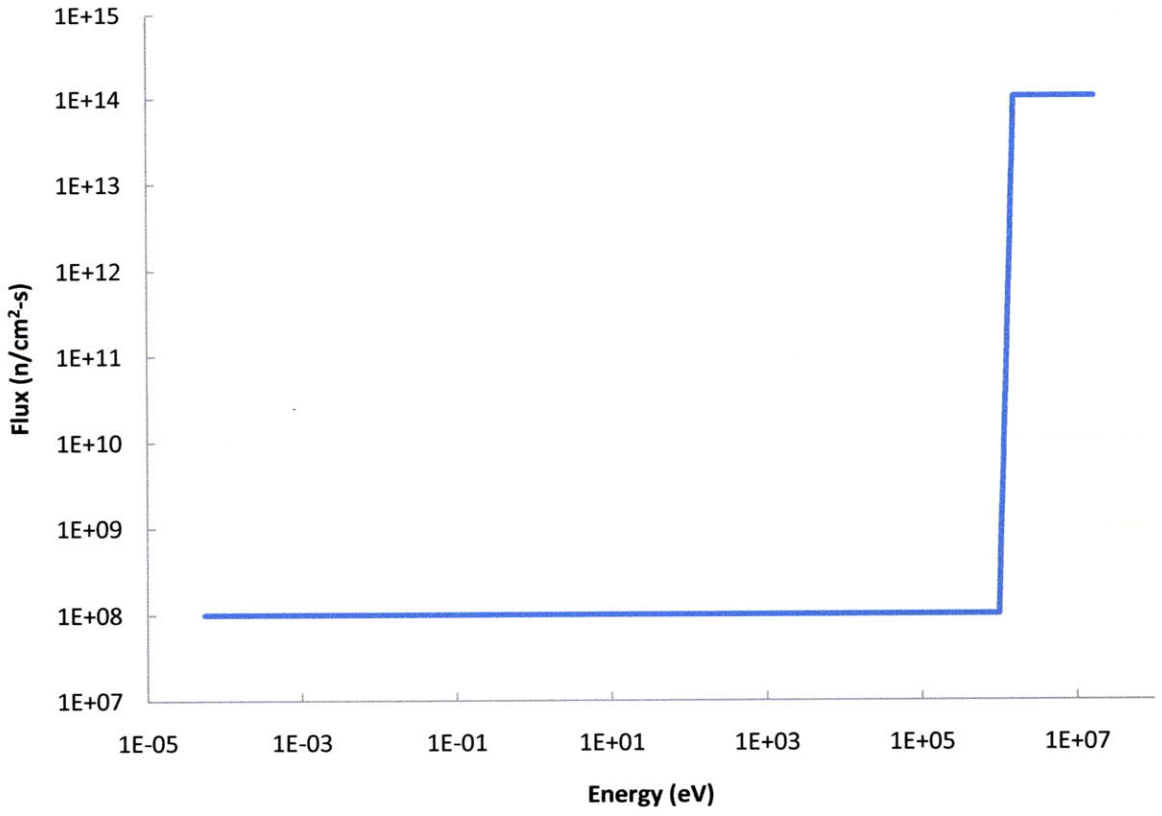

FIGURE 2.8 - Optimal flux spectrum for decay heat after zero years cooling

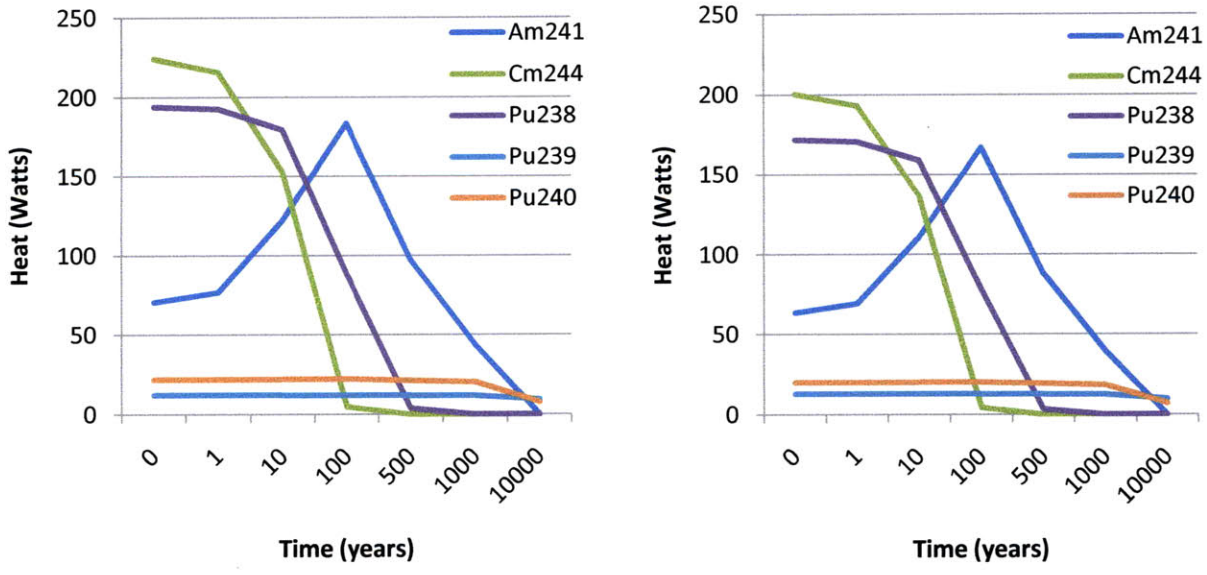

FIGURE 2.9 - Heat rate for natural decay (L) and optimized for zero years cooling (R)

The optimized flux spectrum shows that fast neutrons effectively reduce short term heat and that all neutrons below 1.5 MeV should be avoided. FIGURE 2.9 shows that the optimal flux noticeably decreases the concentration of all of the major heat producing isotopes except for Pu-239, which is slightly increased, relative to the reference case. The information illustrated in these two figures implies that the fast spectrum reduces concentrations of almost all isotopes through fast fission events and that slower neutrons are not useful because they are less likely to cause fission. This hypothesis is confirmed by examining the capture and fission cross sections of the important isotopes presented in the following figures.

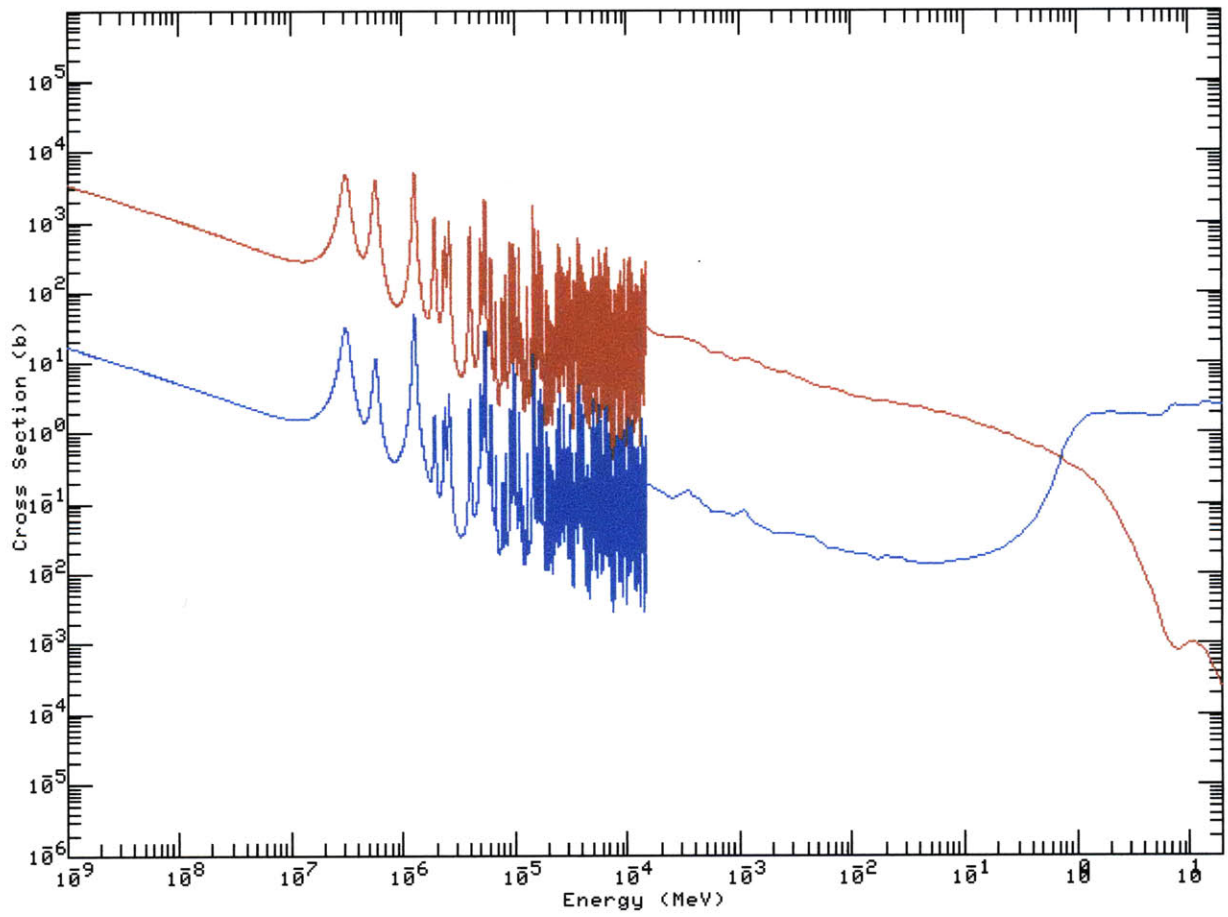

FIGURE 2.10 - Capture (red) and fission (blue) cross sections for Am-241

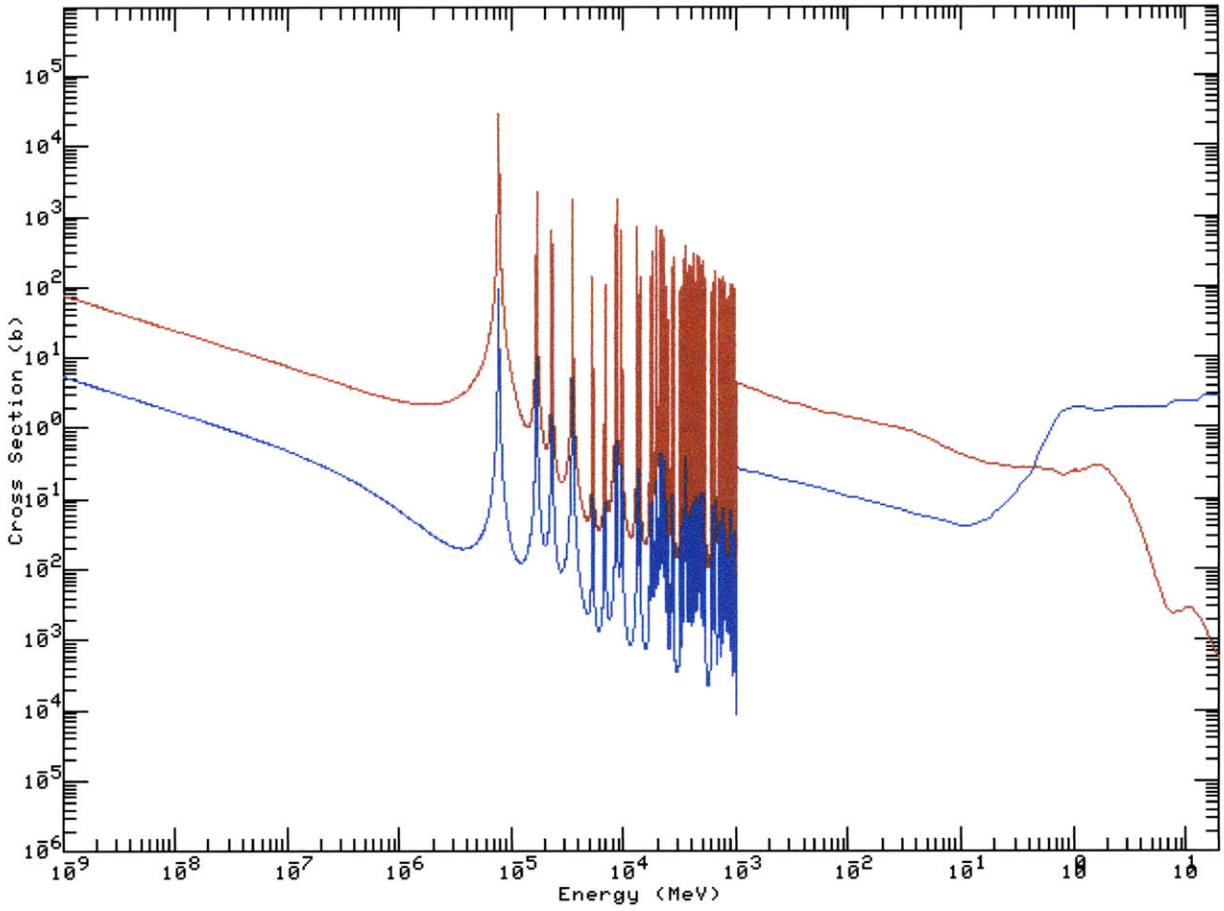

FIGURE 2.11 - Capture (red) and fission (blue) cross sections for Cm-244

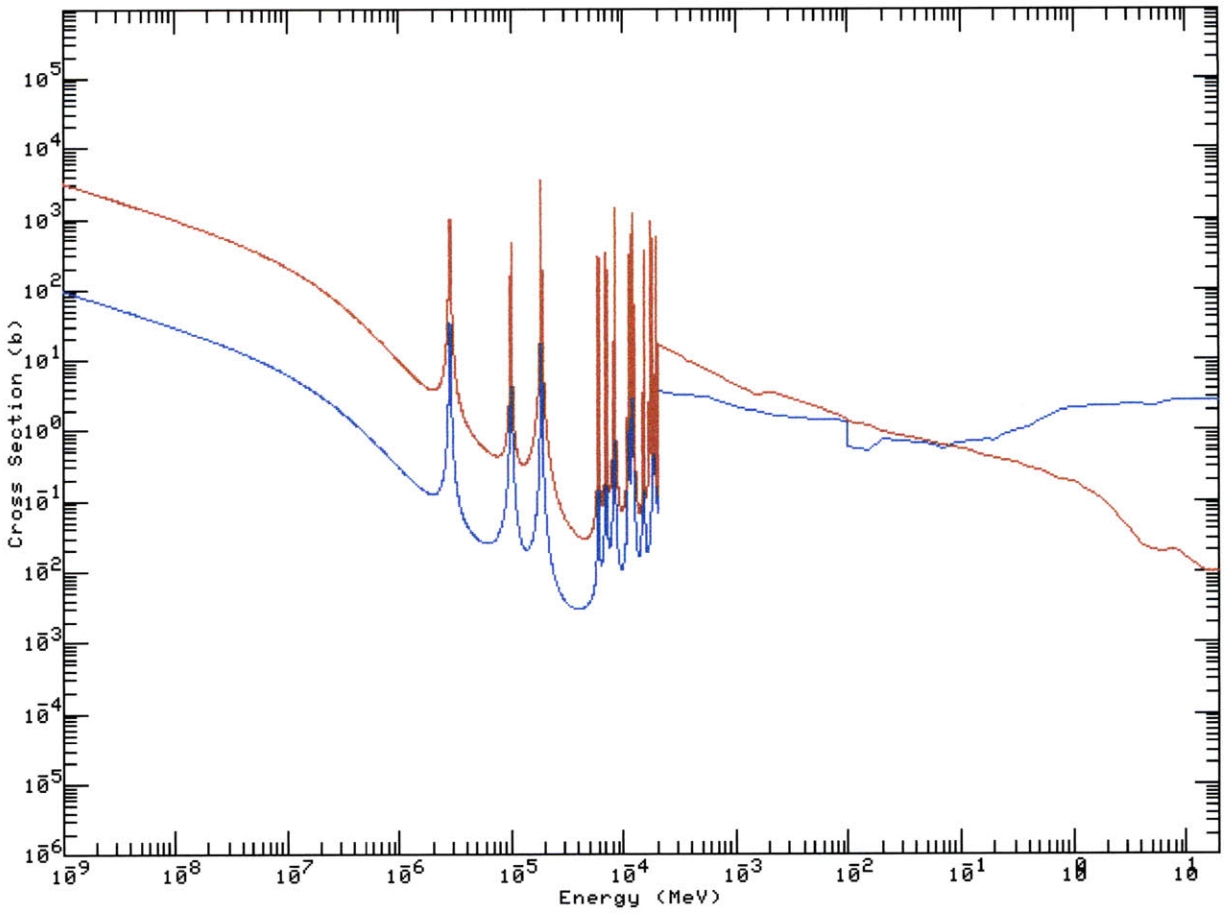

FIGURE 2.12 - Capture (red) and fission (blue) cross sections for Pu-238

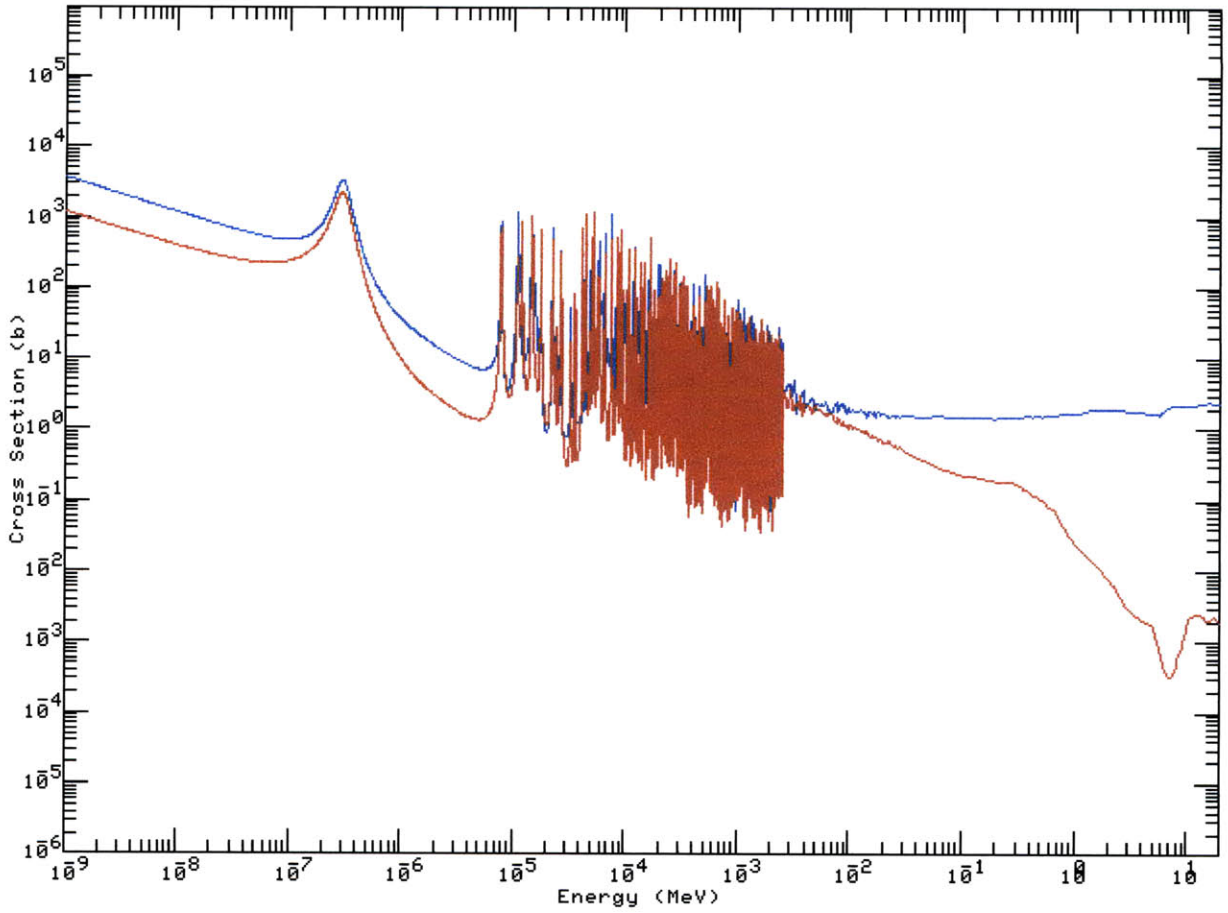

FIGURE 2.13 - Capture (red) and fission (blue) cross sections for Pu-239

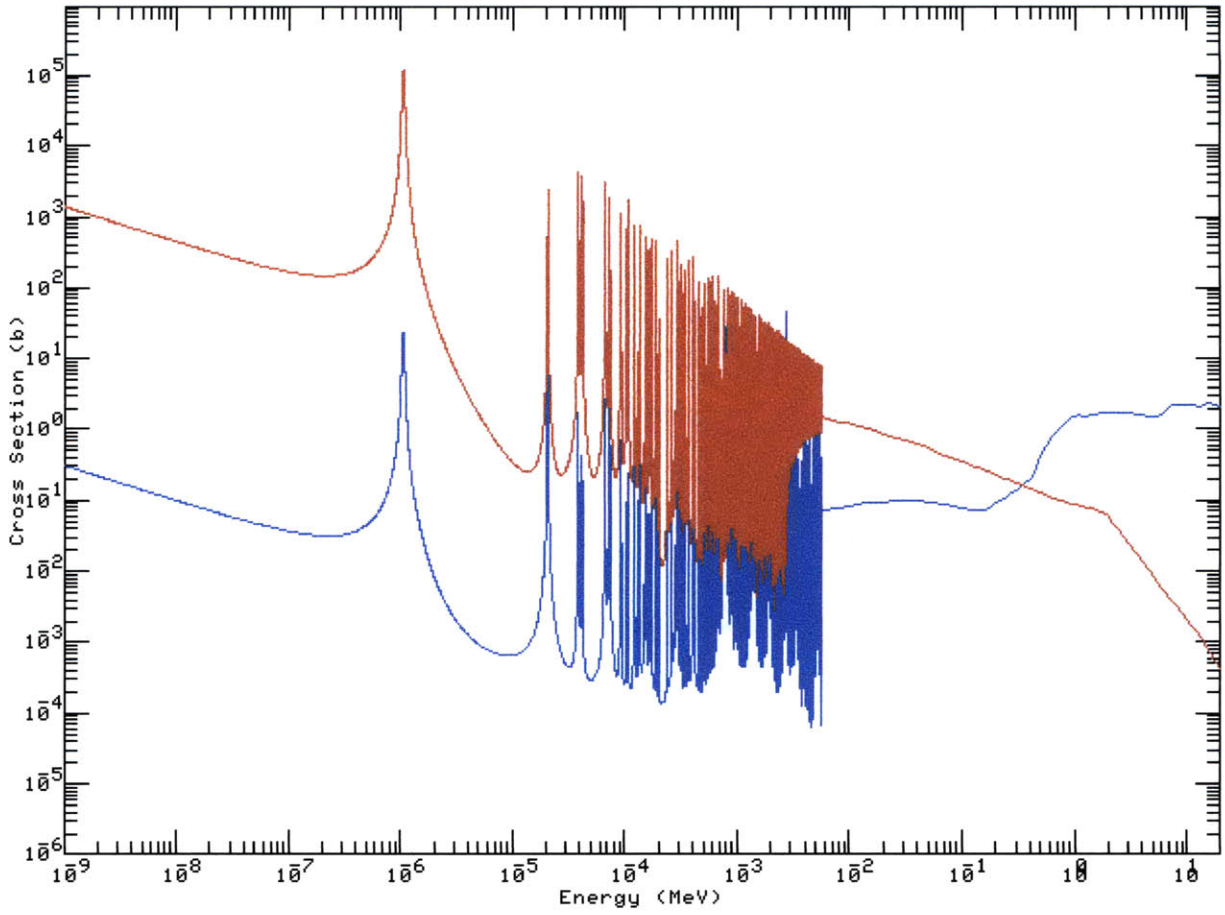

FIGURE 2.14 - Capture (red) and fission (blue) cross sections for Pu-240

With the exception of Pu-239, all of cross sections exhibit a transition between capture-dominant and fission-dominant regimes between 0.5-1 MeV, with the fission/capture ratio increasing with higher energy. This explains the need for an exclusively fast flux for immediately reducing high heat producing isotope concentrations since fission destroys the nuclide (in its current form) immediately. Conversely, capture only transforms the initial isotope to another heat producing actinide and unless the daughter isotope has a half-life on the order of days or less, this transmutation will have no effect on the immediate heat production rate. Also, the slight increase in Pu-239 can be explained by the smaller fission/capture ratio for Pu-238 relative to the other cross section figures above.

Next, the decay heat rate cost function was evaluated after 100 years of cooling. As TABLE 2.4 - Decay heat rate after 100 years cooling results shows, irradiation by an optimal flux spectrum results in a reduction in decay heat rate of almost 62%, which is much more significant than the zero years cooling case. FIGURE 2.15 shows that the optimal flux minimizes heat at 100 years by considerably increasing the heat production during the first 15 years of cooling.

TABLE 2.4 - Decay heat rate after 100 years cooling results

	UNF Decay	Optimized
Watts	313.55	119.56
% Change	---	-61.87

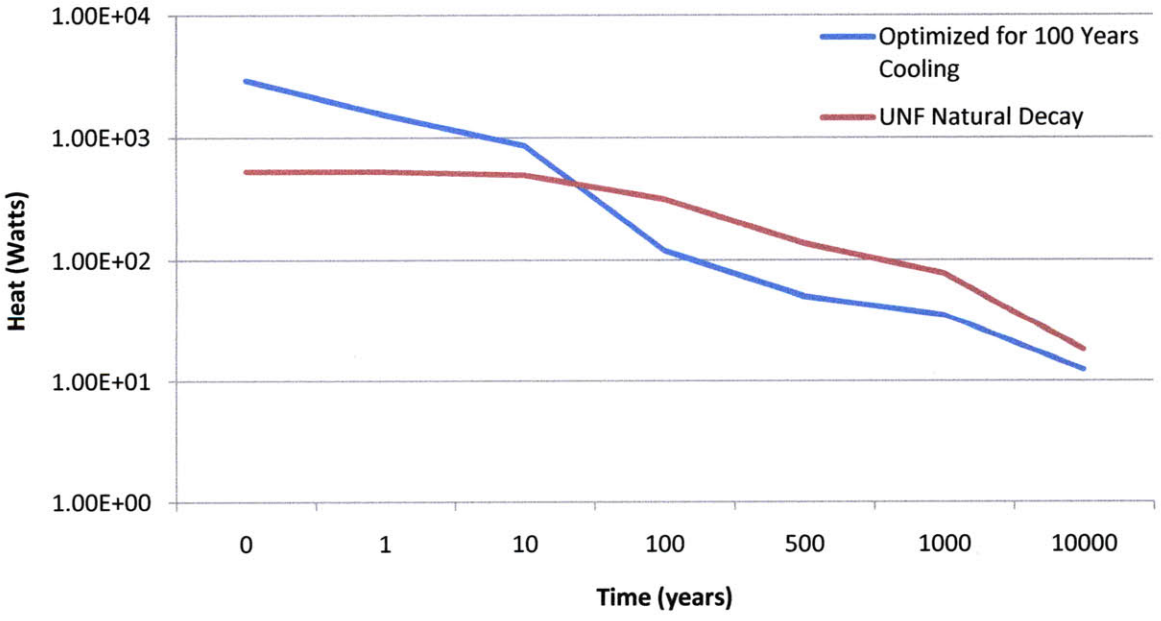

FIGURE 2.15 - Decay heat rate optimized for 100 years cooling

The flux spectrum that minimizes heat after 100 years of cooling is illustrated in FIGURE 2.16 and it shows the same fast flux peak as in the optimal spectrum for zero years cooling with an additional peak at 0.675 eV and a thermal peak. This spectrum's effects on isotopic concentrations are shown in FIGURE 2.17.

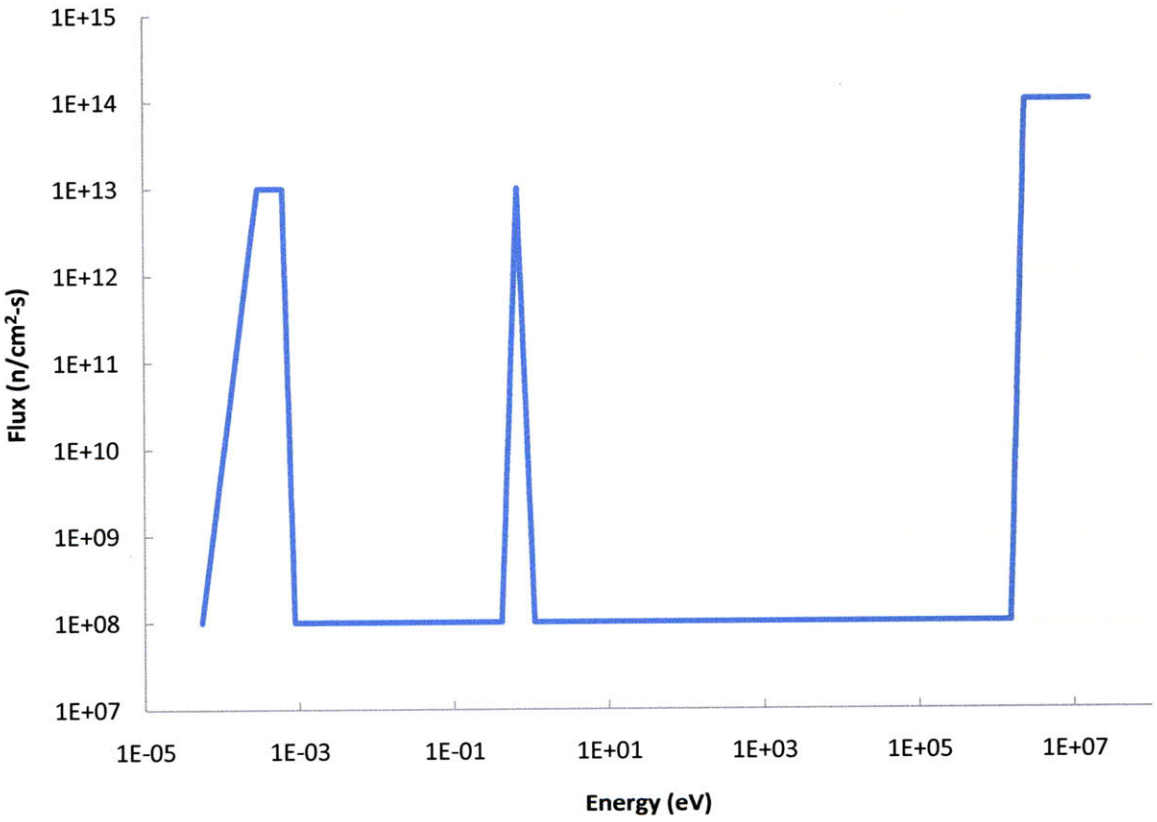

FIGURE 2.16 - Optimal flux spectrum for decay heat after zero years cooling

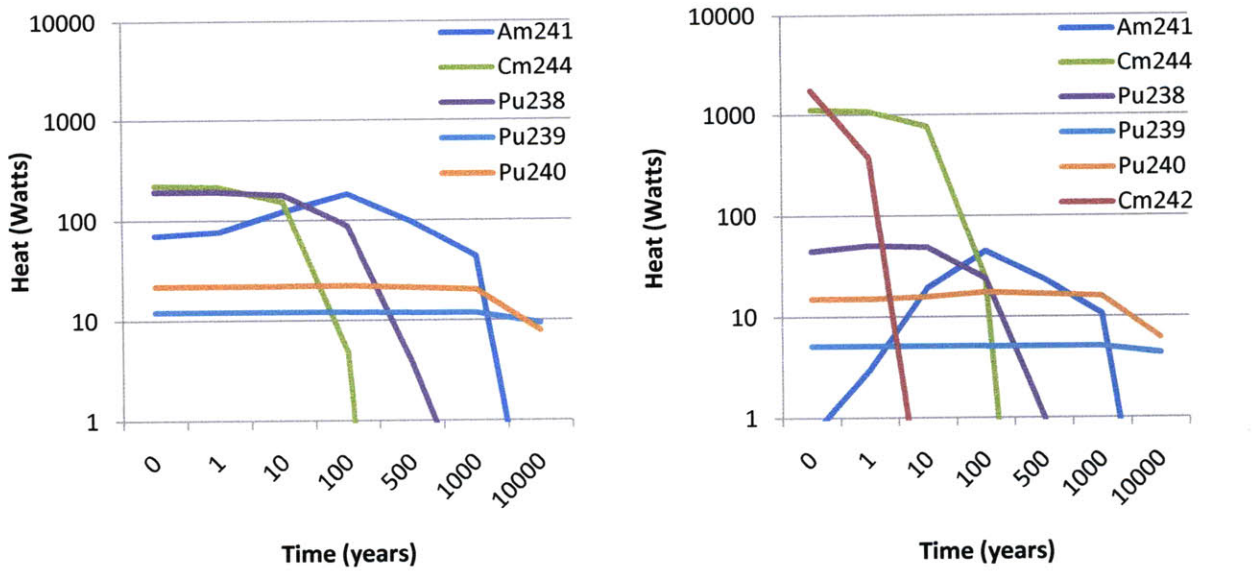

FIGURE 2.17 - Heat rate for natural decay (L) and optimized for 100 years cooling (R)

Once again, the fast peak is used to reduce the concentrations of all isotopes through fast fission. The flux spectrum peak around 0.675 eV corresponds to the approximately 5000 b thermal resonances in both Pu-239's capture and fission cross sections. This peak is almost entirely responsible for the approximately 60% reduction of heat from Pu-239 over the entire cooling period. The largest contribution to the reduction in decay heat, however, comes from the thermal peak which is responsible for the almost entire transmutation of Am-241 into Cm-242 through neutron capture and beta decay of Am-242 (half-life of 16 hours [9]). This peak also is responsible for pushing the other isotopes towards higher actinides through neutron capture and decay, thus resulting in the significant increase in both Cm-242 and Cm-244. This transmutation scheme minimizes decay heat rate after 100 years of cooling by essentially transmuting as many of the longer-lived isotopes into shorter-lived higher actinides, especially Curium.

The final analysis of decay heat production rate minimization was evaluated for irradiated UNF after 1,000 years of cooling. This analysis results in an optimal flux spectrum, illustrated in FIGURE 2.18, that reduces the decay heat production rate after 1,000 years of cooling by 67.7% by increasing the heat released over the first 50 years. FIGURE 2.18 and the similar plots for zero and 100 years of cooling show the trend that as one optimizes for minimal decay heat after longer cooling times, both the initial magnitude of heat production and the length of the increased heat production period increase.

TABLE 2.5 - Decay heat rate after 1,000 years cooling results

	UNF Decay	Optimized
Watts	77.73	25.11
% Change	---	-67.69

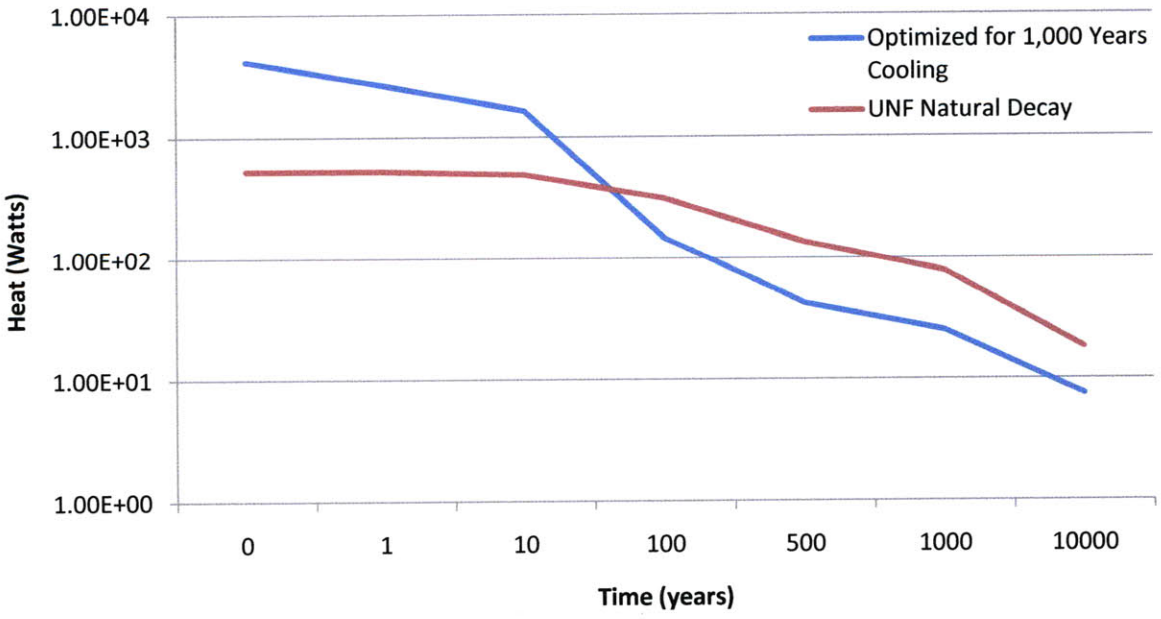

FIGURE 2.18 - Decay heat rate optimized for 1,000 years cooling

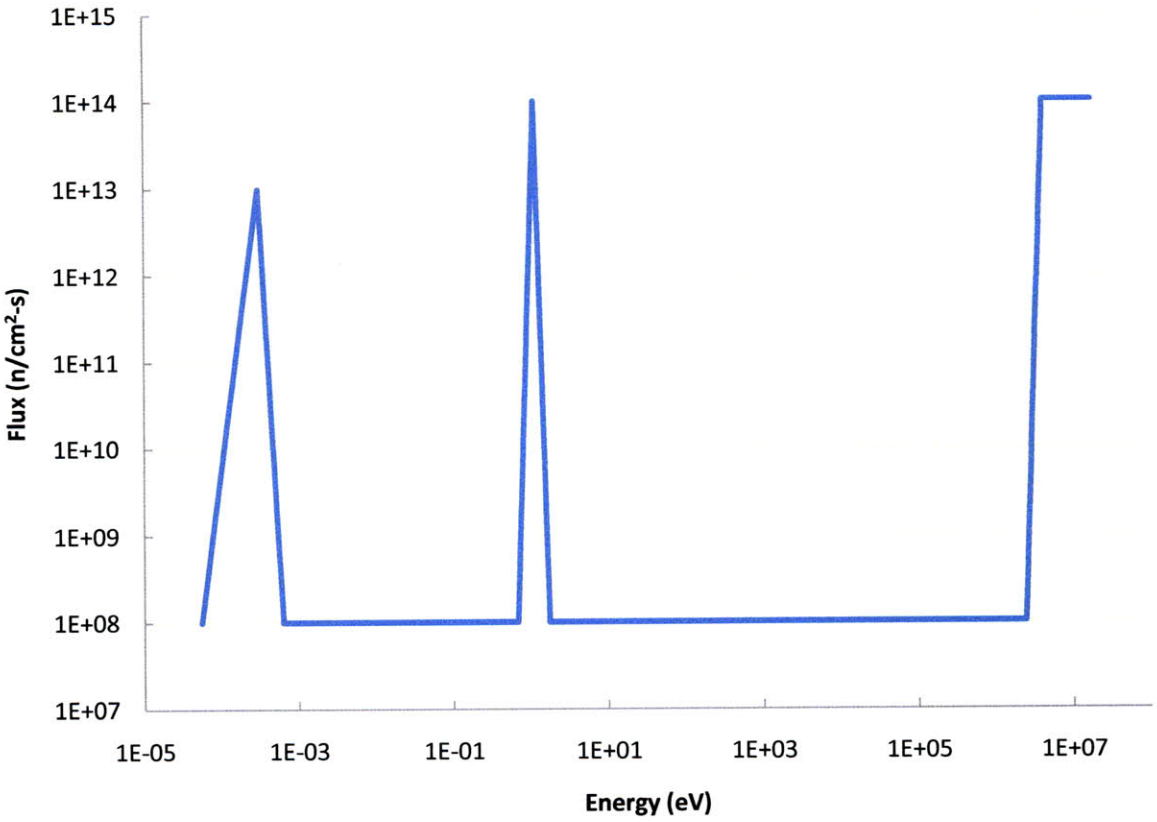
1E+15
1E+14
1E+13
1E+12
1E+11
1E+10
1E+09
1E+08
1E+07
Flux (n/cm²-s)
1E-05
1E-03
1E-01
1E+01
1E+03
1E+05
1E+07
Energy (eV)

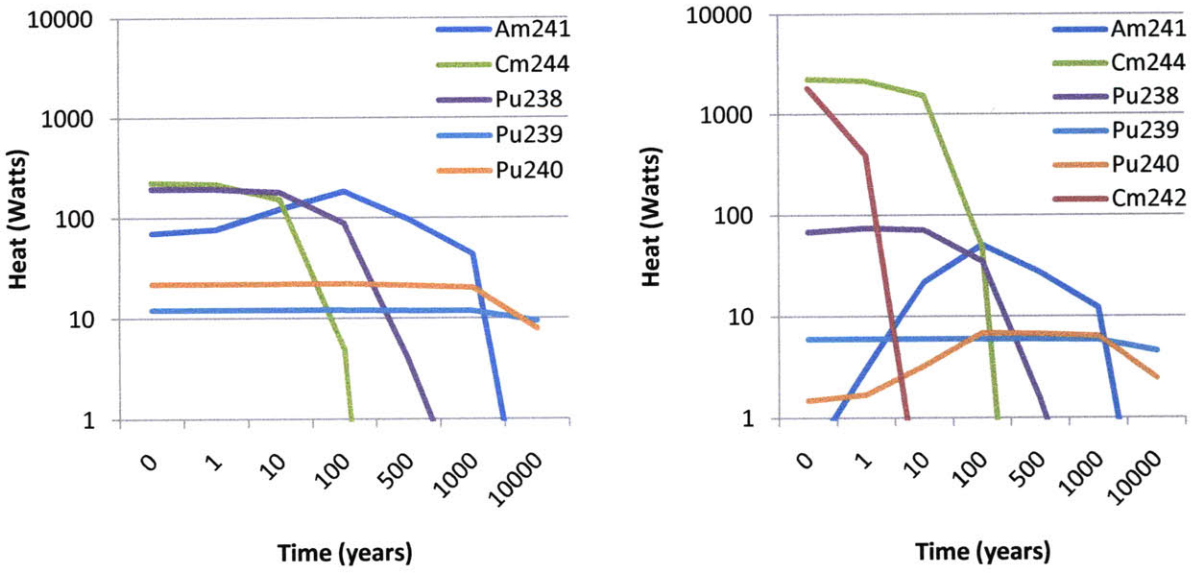

FIGURE 2.20 - Heat rate for natural decay (L) and optimized for 1,000 years cooling (R)

The optimal spectrum for 1,000 years of cooling is very similar to the one for 100 years of cooling except that the fast peak includes one less energy group, the middle peak is at 1.1 eV instead of 0.675 eV, and the thermal peak is slightly narrower. It follows that this spectrum produces similar effects on nuclide concentrations as those optimized for 100 years of cooling. The differences are that this spectrum transmutes even more Am-241 into Cm-242 and that Pu-240's initial concentration is significantly lower than when the spectrum is optimized for 100 years of cooling. The Pu-240 reduction is caused by the flux peak at 1.1 eV because of this isotope's enormous and only partially shielded neutron capture resonance (10,000 b) at this energy, as can be seen in FIGURE 2.14. This leads to large increases in Pu-241 concentration, which decays to Am-241 through beta decay with a half-life equal to 14.3 years [9].

2.7.2 Integrated Decay Heat Production

Minimizing Decay Heat Production

The next logical step in this optimization study is, instead of minimizing decay heat at a specific time, minimizing the total amount of decay heat produced over a period of time. By optimizing for one specific instance, the cost function will be as low as possible at that particular time but this may come at the expense of performance over long periods of time elsewhere in the depletion simulation. This may be acceptable for some applications, but for long-term UNF disposal, the integrated heat production over the entire storage period is a more appropriate cost function. By setting the cost function equal to an approximation of the integral of decay heat produce from zero to 10,000 years of cooling, it was shown that by applying an optimal incident flux spectrum to UNF for 1,000 days the total amount of energy released during storage can be reduced by 46.7%. These results are shown in TABLE 2.6 and FIGURE 2.21 below.

TABLE 2.6 – Integrated heat production over 10,000 years results

	UNF Decay	Optimized
kW-years	185.07	98.68
% Change	---	-46.68

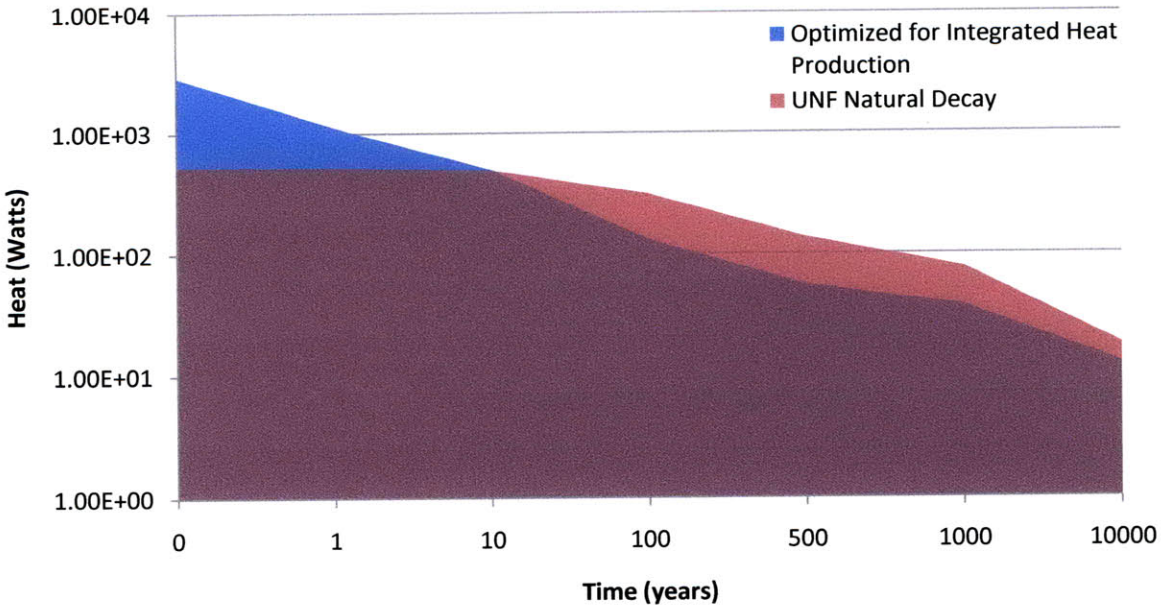

FIGURE 2.21 – Integrated heat production over 10,000 years of cooling

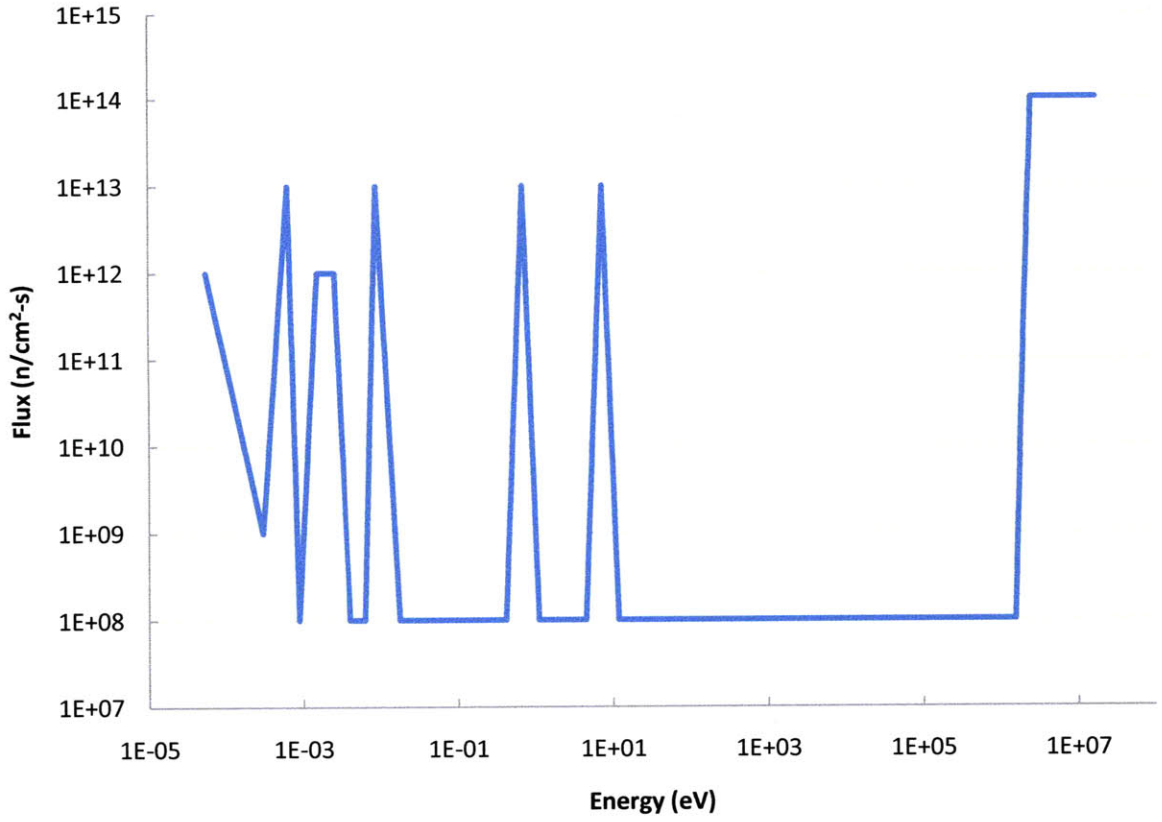

FIGURE 2.22 - Optimal flux spectrum for integrated heat production over 10,000 years

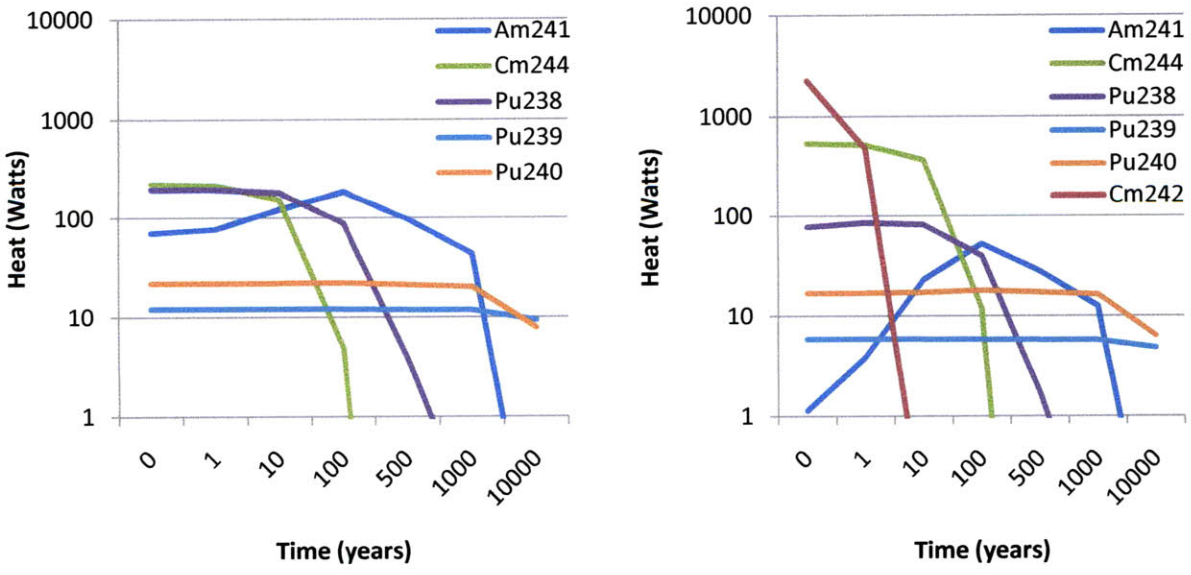

FIGURE 2.23 - Heat rate for natural decay (L) and optimized for 1,000 years cooling (R)

The optimal spectrum for this case is illustrated in FIGURE 2.22 and its effects on nuclide concentrations are shown in FIGURE 2.23. Although this spectrum is significantly different in appearance from those in the previous subsection, its effects on isotope concentrations are similar to those used to minimize decay heat after 100 years of cooling. The main difference between these two cases is that nuclide concentrations are not pushed as far up the actinide chain as they are when optimizing just for minimal heat at 100 years of cooling.

Since this spectrum is considerably more complex than any of those previously illustrated, a detailed analysis of each individual spectral peak's contribution to minimizing the cost function was conducted. The procedure involved evaluating the cost function from a flat flux at 10^8 n/cm^2-sec and then reevaluating the cost function after adding one of the peaks to the flat flux spectrum. The results are presented below.

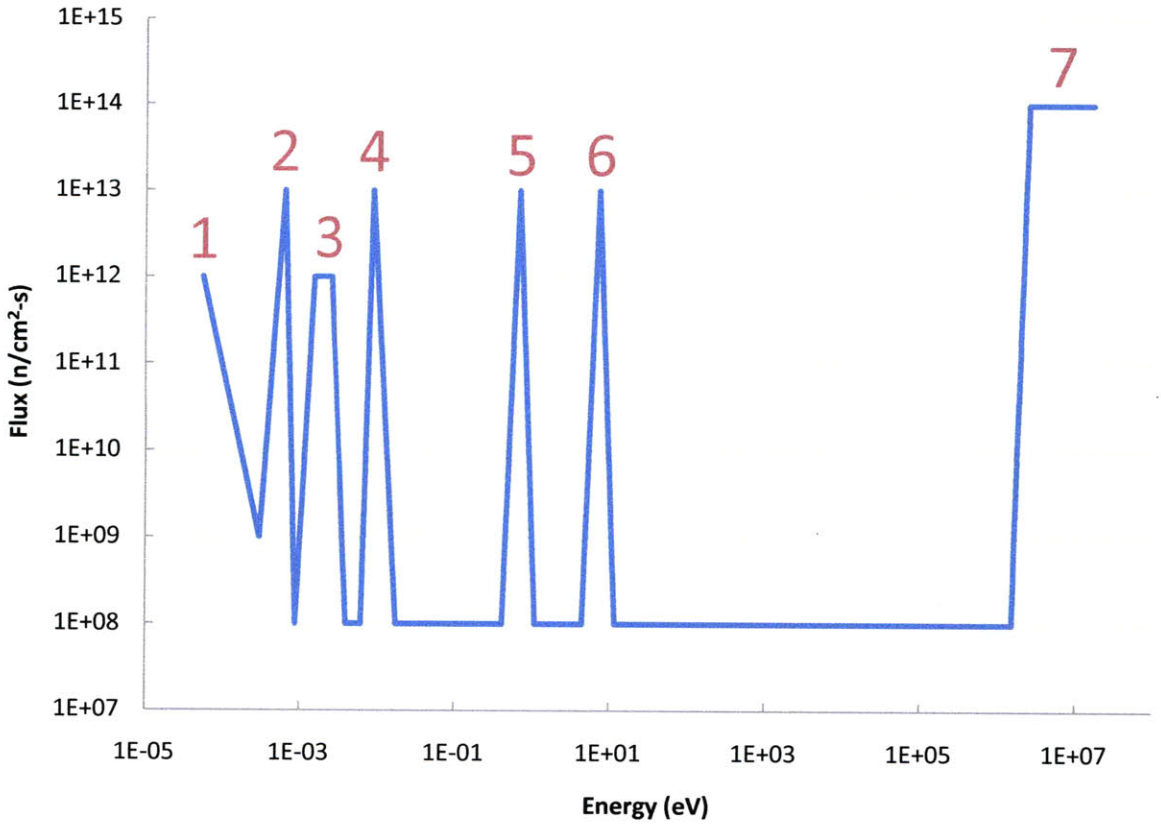

FIGURE 2.24 - Integrated heat production over 10,000 years spectrum with peaks labeled

TABLE 2.7 - Contribution of individual peaks to cost

	Flat Flux	Peak 1	Peak 2	Peak 3	Peak 4	Peak 5	Peak 6	Peak 7
kW-Years	185.06	142.50	113.56	163.27	149.87	263.99	182.06	170.67
% Change	---	-23.00	-38.63	-11.77	-19.02	42.65	-1.62	-7.77

The information presented in TABLE 2.7 clearly exemplifies the nonlinearity of this optimization problem; while most of the peaks effectively reduce the cost function, peak 5 when included on its own actually increases the cost by over 40%, but when coupled with the 6 other peaks, it contributes to an overall decrease in cost.

A reduction in integral decay heat production by almost 47% is a very exciting result especially when considering the impact this would have on the cost per unit mass of heavy metal of heat-limited geologic repositories; however, the spectrum that permits this reduction is too detailed to be replicable in an actual physical system. While the individual peaks in the optimal spectrum could not be reproduced physically, a simpler version of this spectrum, shown in FIGURE 2.25, could be approximated in reality by introducing efficient moderating materials that have large scattering cross sections and low absorption that together cover the epithermal energy range, leaving only the fast and thermal flux peaks shown below. This spectrum's performance is also compared to that of natural decay and the optimal flux spectrum in TABLE 2.8.

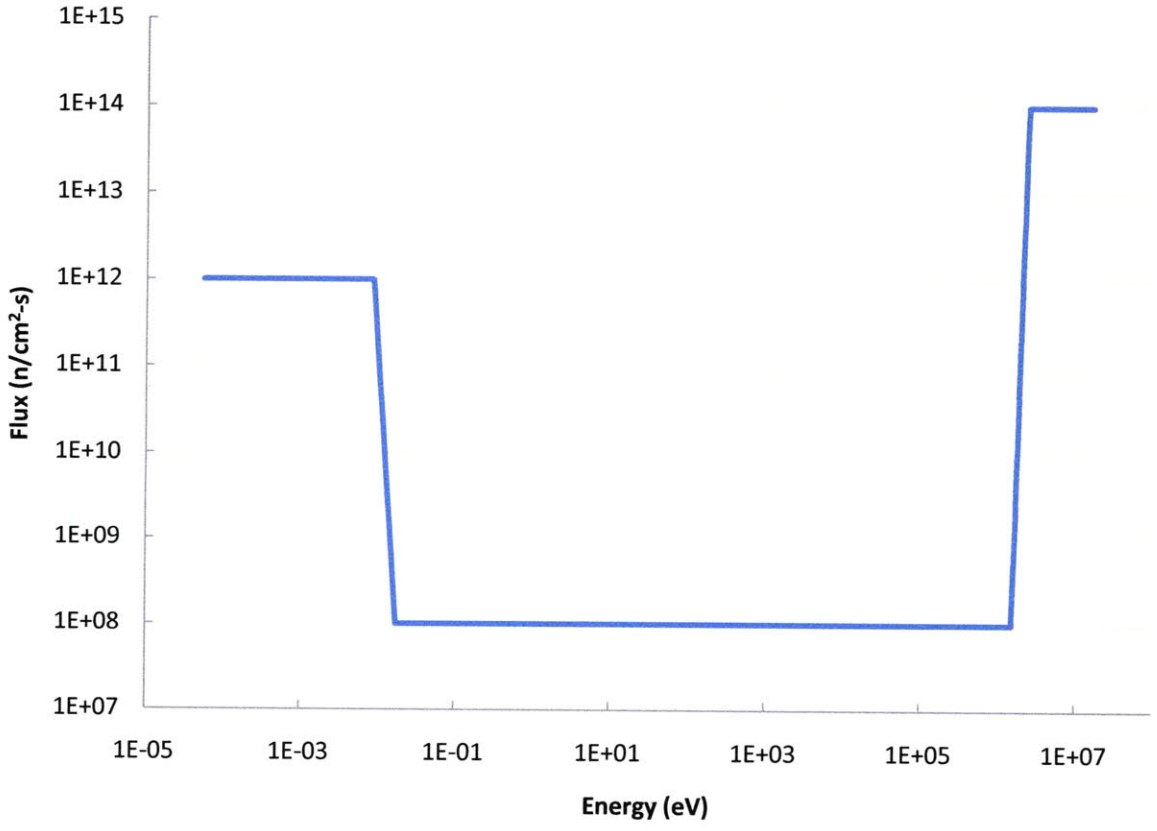

FIGURE 2.25 - Simplified flux spectrum for minimizing integral heat production

TABLE 2.8 - Performance of fast, thermal, simplified, and optimal spectra relative to natural decay

	UNF Decay	Fast Peak	Thermal Peak	Simplified Spectrum	Optimal Spectrum
kW-years	185.07	170.67	117.29	109.35	98.68
% Change	---	-7.78	-36.62	-40.91	-46.68

While the simplified spectrum is not as tailored to the actinides' individual resonances as the optimal spectrum, it is still able to reduce integral heat production over 10,000 years by almost 41%. What is perhaps even more surprising is the large discrepancy between the fast and thermal peaks' abilities to decrease the concentration of long-term heat producing isotopes, with the thermal spectrum reducing the cost function by over five times as much as the fast peak.

Maximizing Decay Heat Production

An additional case related to those above was performed next, but instead of using the integral of heat produced as the cost function, the inverse of the integrated heat production over 10,000 years was taken as the cost function. By minimizing this function, the optimization routine is essentially finding the maximum integral heat production. The solution in this case indicates what flux spectrum should be avoided when trying to decrease heat production. The results of this case show that integrated heat production can be increased by almost 800% by irradiating UNF with the indicated flux spectrum. The numerical results are shown in TABLE 2.9 and the associated flux spectrum is shown in FIGURE 2.26. This spectrum's general shape is very close to the opposite of the spectrum that minimizes the same cost function, which verifies that the optimization methodology used here provides consistent results. It should also be noted that the procedure used here to maximize heat production could also be used to optimize flux spectra for the production of certain isotopes.

TABLE 2.9 – Maximized integrated heat production over 10,000 years results

	UNF Decay	Maximized
kW-years	185.07	1662.65
% Change	---	798.39

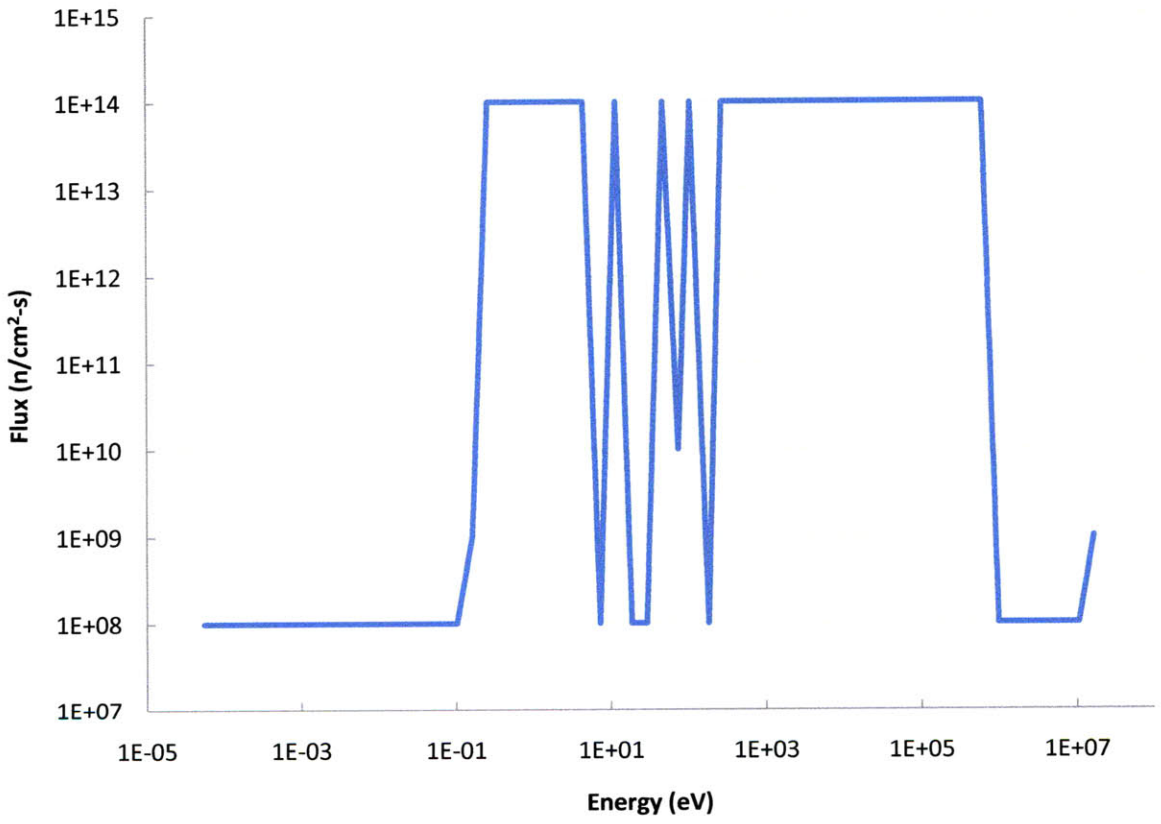

FIGURE 2.26 - Spectrum to maximize integrated heat production

Effects of Actinide Separation Schemes

The final cases analyze the effects of typical actinide separation schemes on integrated heat production over 10,000 years of cooling. In these cases, selected isotopes are removed from the

UNF vector and allowed to decay while the remaining isotopes are irradiated; all isotopes are then recombined and allowed to decay for 10,000 years. The optimizations performed are for the transmutation of all transuranic elements (TRU); neptunium, plutonium, and americium; neptunium and plutonium; and only plutonium.

TABLE 2.10 – Effects of actinide separation schemes on minimizing integrated heat production

	UNF Decay	TRU	Np Pu Am	Np Pu	Pu
kW-years	185.07	98.68	103.70	111.99	112.17
% Change	---	-46.68	-43.97	-39.49	-39.39

The numerical results from these cases are shown in TABLE 2.10 and their respective spectra are shown in FIGURE 2.27, FIGURE 2.28, and FIGURE 2.29; the spectrum for TRU transmutation is not included below because it is identical to the spectrum for optimizing integrated heat production for the entire UNF vector in FIGURE 2.22. This information indicates that separating uranium from TRU has essentially no effect on long-term heat production and that plutonium dominates the other actinides in terms of total heat production over 10,000 years. These conclusions, however, may not hold in the presence of more realistic flux spectra.

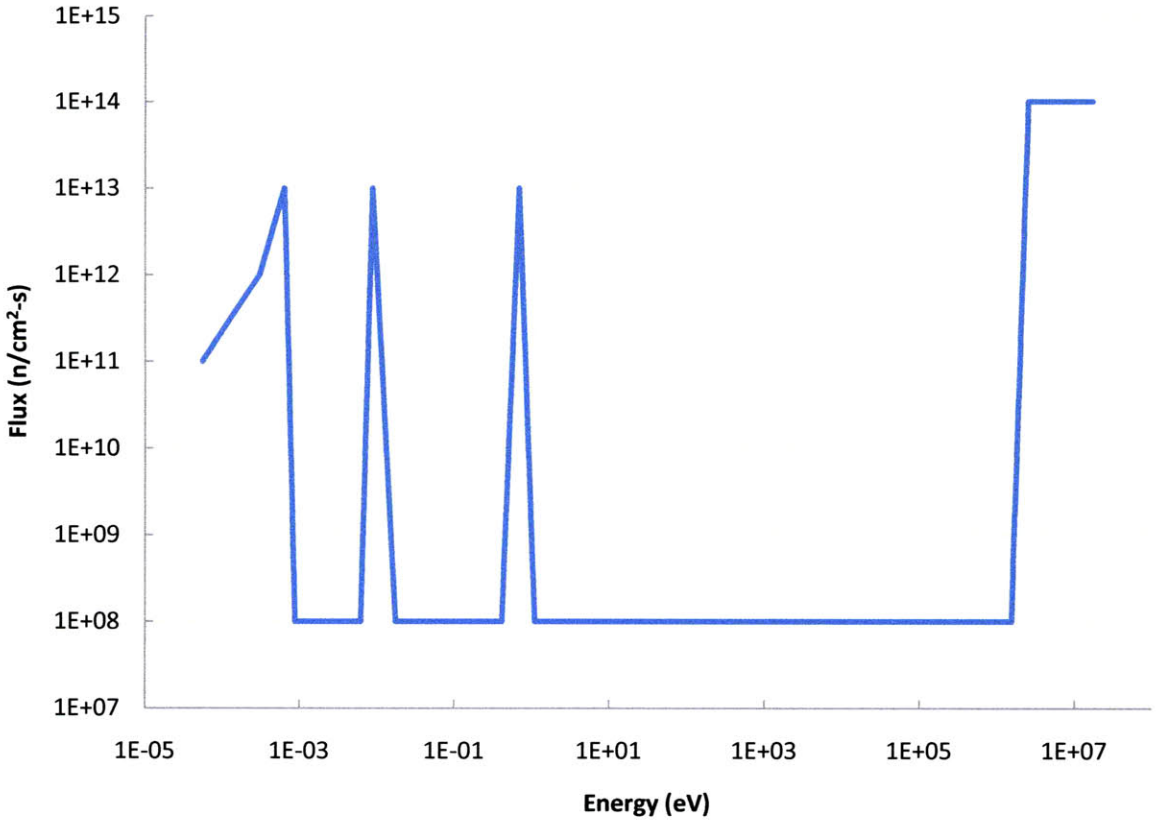

FIGURE 2.27 - Optimal spectrum for Np Pu Am transmutation

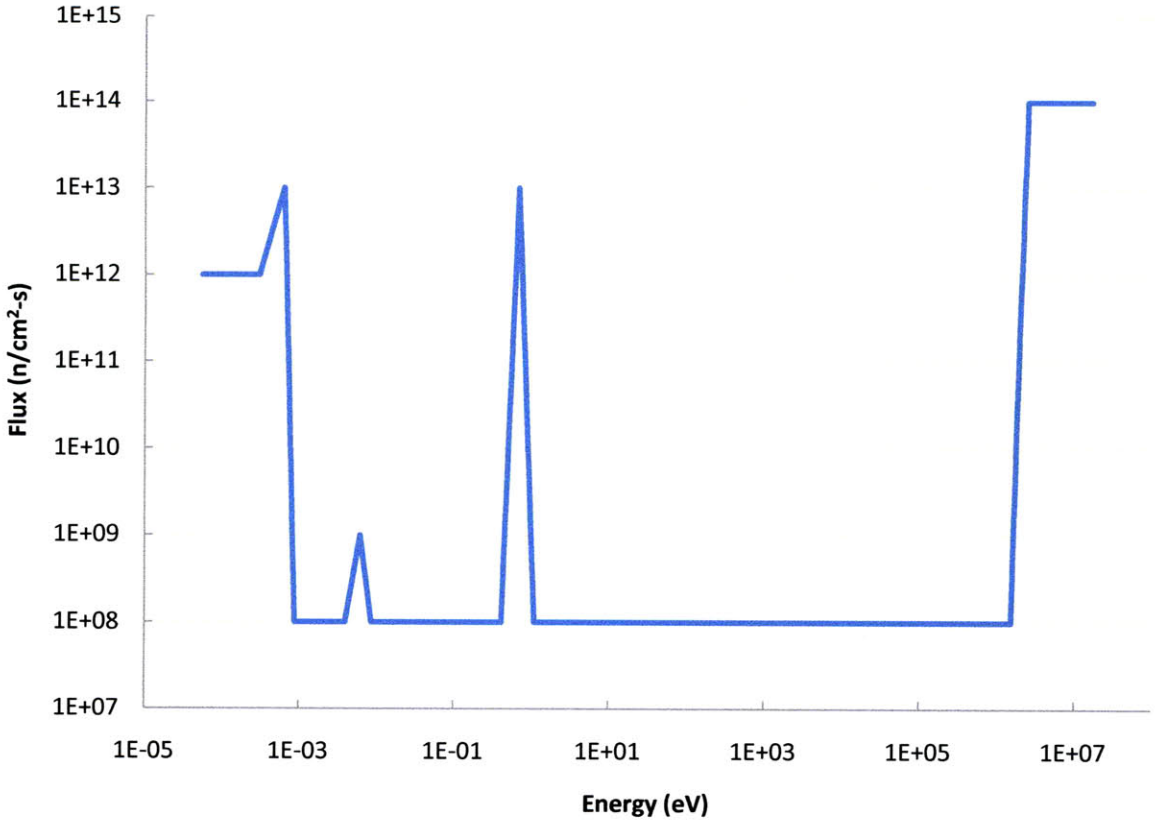

FIGURE 2.28 - Optimal spectrum for Np Pu transmutation

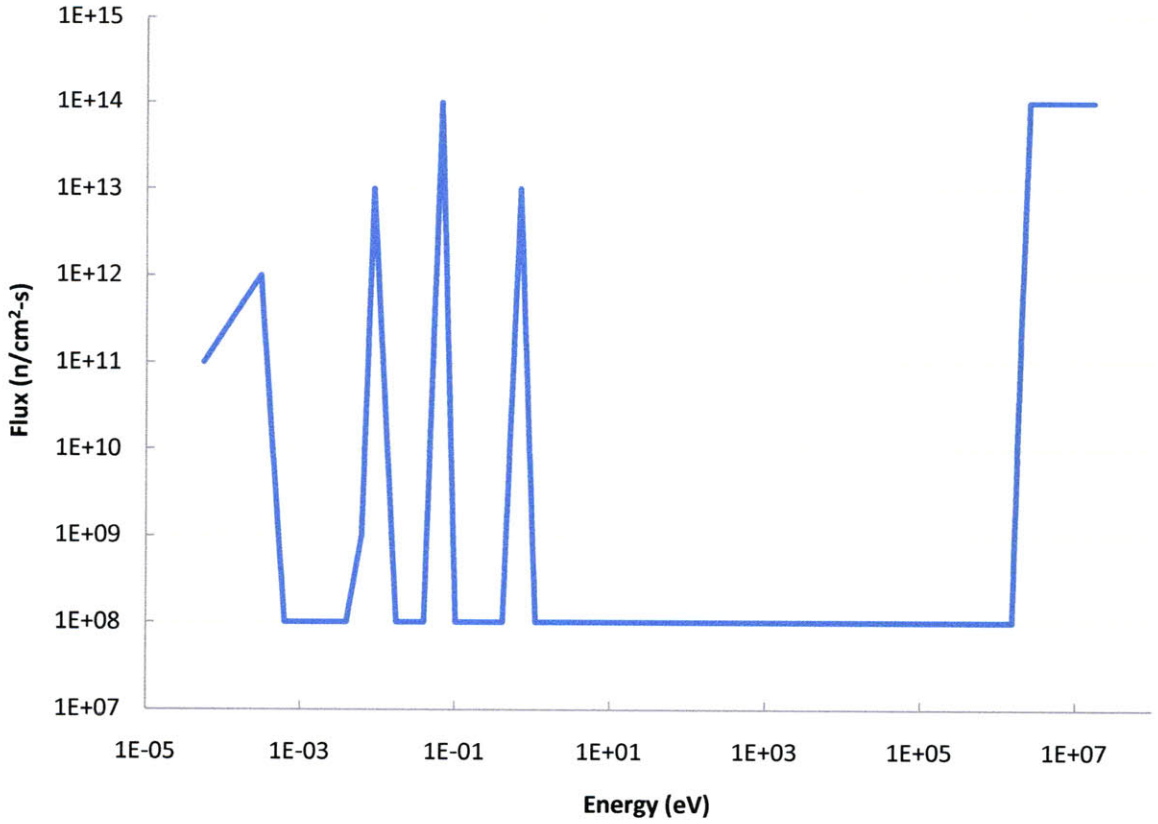

FIGURE 2.29 - Optimal spectrum for Pu transmutation

2.7.3 Plutonium Content

The same optimization routine used in the above studies was also applied to managing the isotopics of the plutonium vector in UNF. The major application for this study is in reducing proliferation risk by making the plutonium in UNF less attractive for weapons material. This can be accomplished by reducing the concentration of Pu-239, the primary isotope used in nuclear weapons, or by increasing the concentrations of Pu-238 and/or Pu-240, which act as contaminants that are difficult to separate from Pu-239. Many other metrics have been proposed for proliferation resistance of plutonium that could also be used as our cost function in future studies [20]. Several separate cases are presented here to analyze plutonium's spectral

dependences and how these can be used to increase the proliferation resistance of UNF. As in the previous studies, UNF is irradiated for 1,000 days but the plutonium content is evaluated immediately after irradiation; there is no decay case because, with the exclusion of Pu-241, the major isotopes of plutonium have half-lives on the order of thousands of years [9].

The first case studied here is the most obvious: finding an optimal spectrum to minimize the Pu-239 content in UNF. The results of this case indicate that exposing UNF to the flux spectrum illustrated in FIGURE 2.30 for 1,000 days can transmute almost 96% of the initial Pu-239 mass, while reducing the weight percent of Pu-239 from 52.1% to 3.7%. The numerical data are presented in TABLE 2.11 and TABLE 2.12.

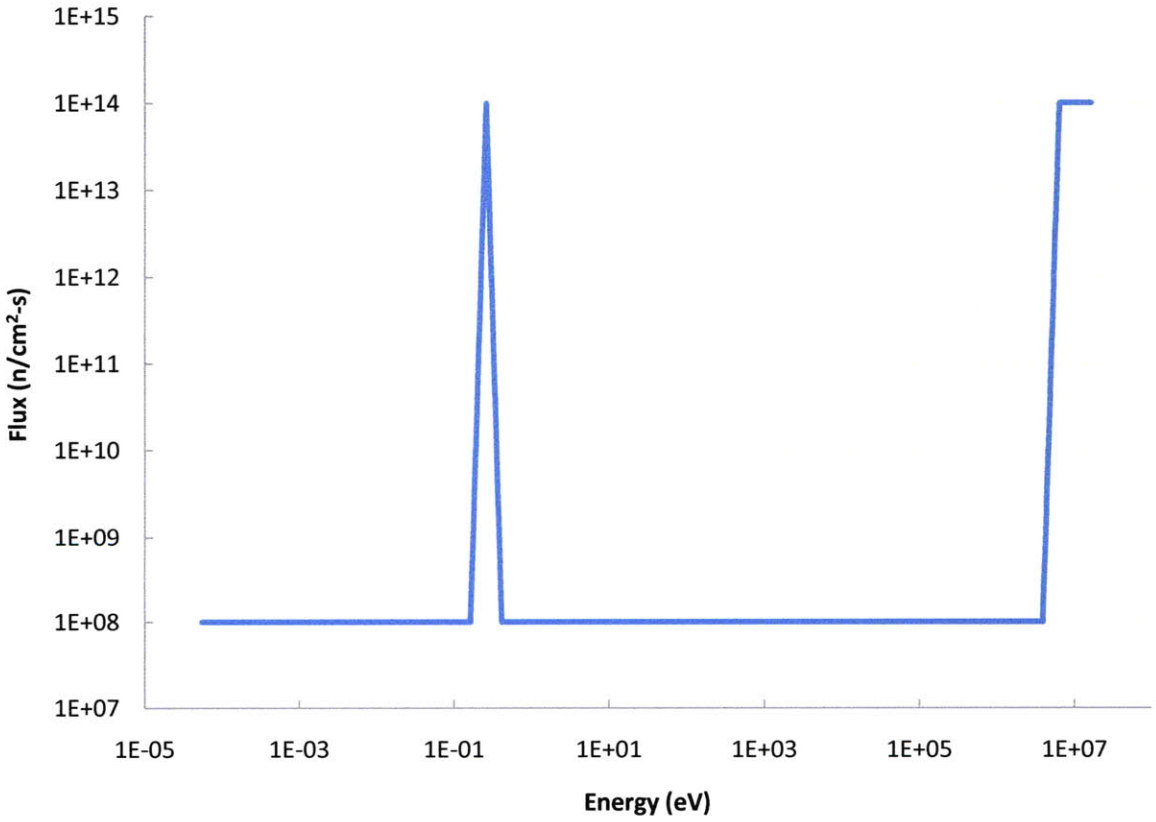

FIGURE 2.30 - Optimal spectrum for transmutation of Pu-239

TABLE 2.11 - Results of Pu-239 transmutation optimization

	Transmuted	Natural Decay
Initial Pu-239 (kg)	6.3309	6.3309
Final Pu-239 (kg)	0.2612	6.3305
% Change	-95.87	-0.01

TABLE 2.12 - Composition of plutonium vector

	Initial (wt%)	Natural Decay (wt%)	Transmuted (wt%)
Pu-238	2.88	2.86	8.09
Pu-239	52.11	52.85	3.71
Pu-240	25.50	25.92	45.37
Pu-241	11.61	10.35	3.89
Pu-242	7.91	8.02	38.94

The spectrum is easily explained by examining Pu-239's cross sections. The peak at 0.2625 eV corresponds to the large thermal resonance in both the capture and fission cross sections and the fast peak is only present at energies of approximately 6 MeV and above because that is where the depression in the capture cross section maximizes the fission/capture ratio. The spectrum is very stratified because higher fluxes at any other energies would increase production of Pu-239 through neutron capture in Pu-238. This is in agreement with the spectrum produced when optimizing for maximum Pu-239 mass that is shown in FIGURE 2.31. This spectrum increases Pu-239 mass by 874% while also increasing Pu-238 mass by 841% and Pu-240 mass by 460%

and creates a plutonium vector with weight percents of 73.7% Pu-239, 20.7% Pu-240, and 3.9% Pu-238.

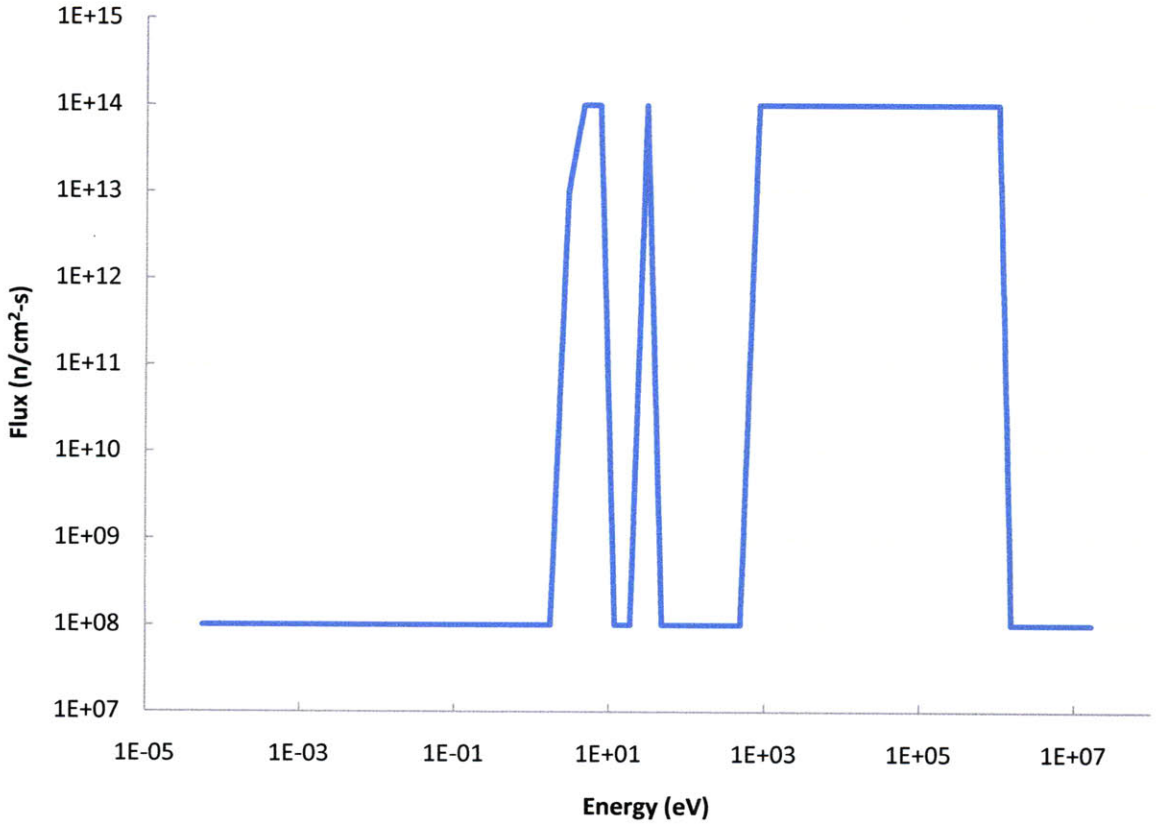

FIGURE 2.31 - Optimal spectrum for maximizing Pu-239 production

Optimization cases were also performed to maximize Pu-238. The only features in the optimal spectrum, FIGURE 2.32, that account for the 1300% increase in Pu-238 concentration are three flux peaks in the 1-100 eV range. By looking only at the Np-237 (the main source of Pu-238 production) cross sections in FIGURE 2.33, one would assume there would be a flux plateau and not three separate peaks in this energy range to take advantage of Np-237's gigantic low-energy resonances; the depressions, however, are quickly accounted for by examining the Pu-238 cross

sections in FIGURE 2.12, which shows three distinct resonances at the same energies of the flux depressions.

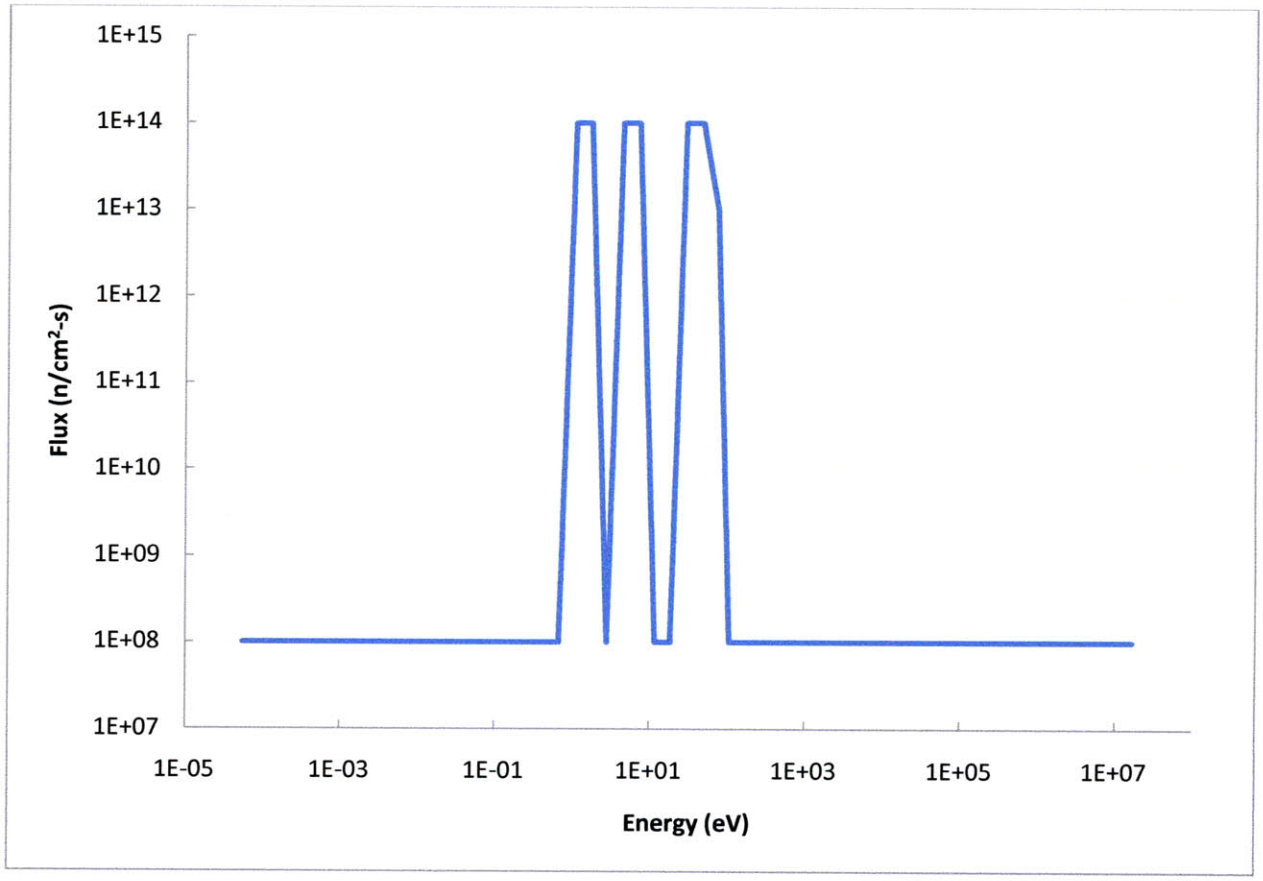

FIGURE 2.32 - Optimal spectrum for maximizing Pu-238 production

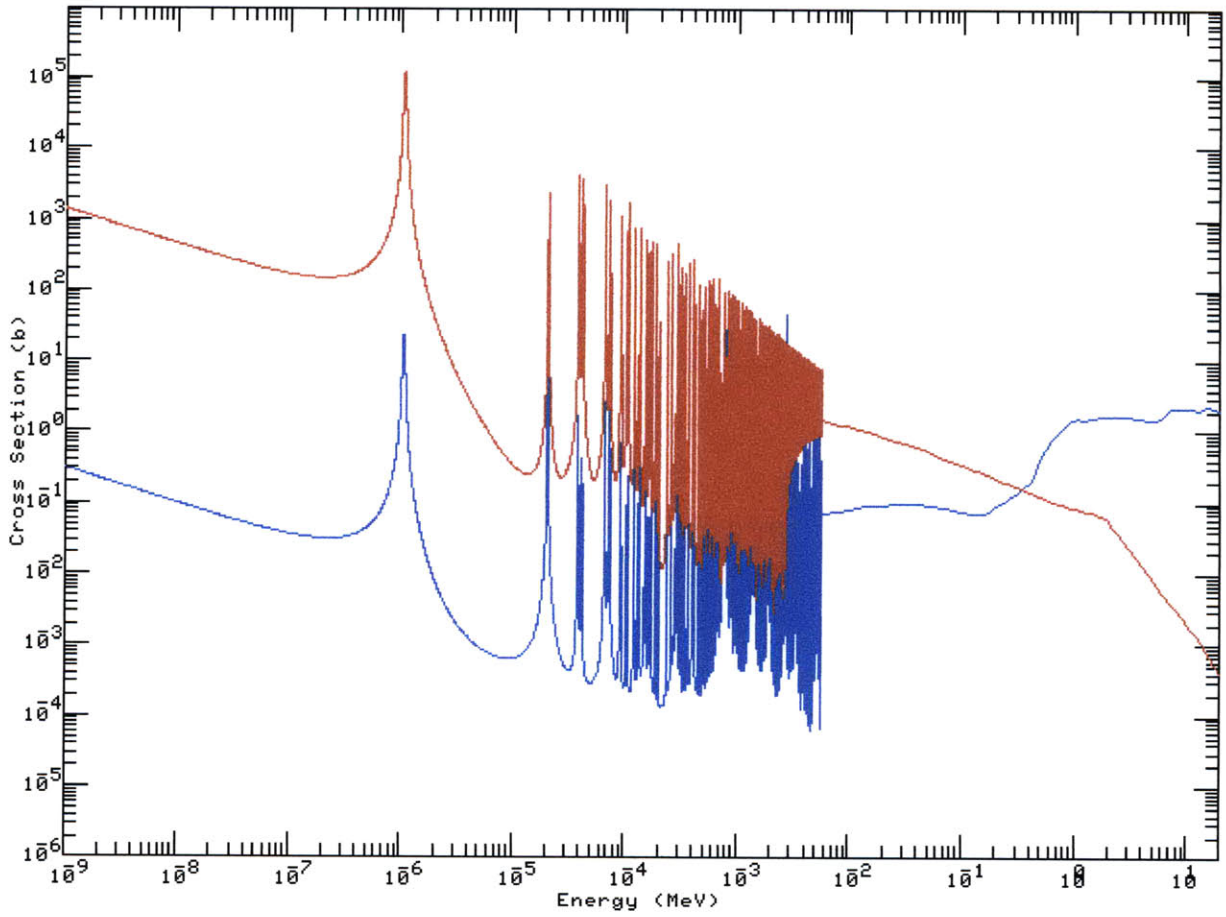

FIGURE 2.33 - Capture (red) and fission (blue) cross sections for Np-237

Another optimization case was performed for maximizing Pu-240. The resulting spectrum is shown in FIGURE 2.34 and its shape can be easily explained by looking at the cross sections of Pu-239 and Pu-240. The peak around 0.26 eV is from Pu-239's lowest energy capture resonance and the flux depressions around 1 eV and 20 eV are caused by Pu-240's very sharp resonances at those energies. Although the two isotopes have roughly similar capture cross sections in the 100 eV to 1 MeV, the corresponding flux plateau increases the amount of Pu-240 because the initial concentration of Pu-239 in UNF is twice that of Pu-240, thus leading to twice as many Pu-240 nuclei produced than lost through neutron capture in this energy region.

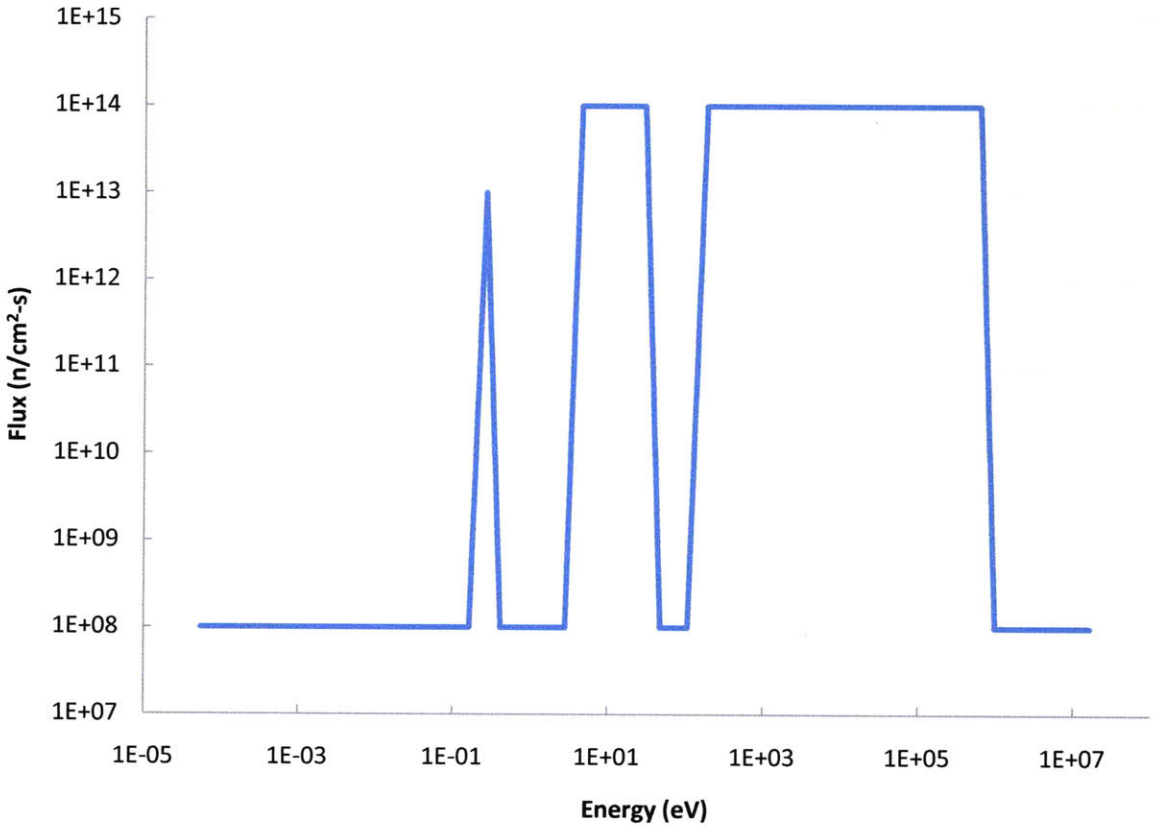

FIGURE 2.34 - Optimal spectrum for maximizing Pu-240 production

The main purpose of this study on plutonium isotopic optimization was to find the best way to reduce proliferation risks through the transmutation of UNF. To determine the relative effectiveness of each of the above plutonium optimization cases, their resulting plutonium vectors are presented in TABLE 2.13. From the data below, it is clear that by optimizing for minimal Pu-239 directly is much more effective than trying to increase Pu-238 or Pu-240. It also appears that while optimizing for Pu-238 production increases its relative concentration in the plutonium vector by more than 400%, it also significantly increases the concentration of Pu-

239; this essentially rules out maximizing Pu-238 as a way to reduce proliferation risks, but it could still be used to produce Pu-238 for space-nuclear applications.

TABLE 2.13 – Comparison of plutonium vectors from different isotopic optimizations

	Initial UNF (wt%)	Natural Decay (wt%)	Minimum Pu-239 (wt%)	Maximum Pu-238 (wt%)	Maximum Pu-240 (wt%)
Pu-238	2.88	2.86	8.09	12.88	3.44
Pu-239	52.11	52.85	3.71	69.18	30.28
Pu-240	25.50	25.92	45.37	2.03	58.44
Pu-241	11.61	10.35	3.89	8.41	2.68
Pu-242	7.91	8.02	38.94	7.49	5.16

3 MATERIAL OPTIMIZATION

3.1 Introduction

In the previous chapter, it was shown that transmutation of UNF can drastically alter its physical properties to make it more amenable to many fuel cycle-related applications. The efficiency of transmutation, however, is dictated by the incident flux spectrum and although optimal flux spectra were identified in the previous chapter, replicating these energy distributions in reality is not a trivial matter. In the previous study, it was assumed that the flux level in each energy group could be set independently of all other groups and with no dependence on physical properties such as geometry and material composition. These assumptions do not hold true when examining real physical systems; flux spectra are shaped by the material through which neutrons pass, group fluxes are strongly influenced by those of the adjacent groups, flux spectra do not remain constant in physical space, etc. So while an optimal flux spectrum may be able to drastically reduce UNF long-term heat production or make it much less attractive as weapons material, the spectrum's sharp peaks and depressions cannot be reproduced in reality. This does not however mean that similar results cannot be obtained; it just provides theoretical insights that can direct further study into optimal transmutation. It follows that the next step in this project is to apply transmutation optimization to a more realistic situation.

3.2 General Methodology

Because of the direct relationship between flux spectra and the media through which neutrons pass, it is logical that optimization of material composition follows the study of optimal flux

spectra. While the general optimization routine implemented here can be applied to almost any design variable, this project focuses on the effects of material composition on transmutation properties. By varying the assembly composition, one can quantitatively determine optimal material combinations and concentrations for specific fuel cycle goals.

On a fundamental level, the methodology used here is very similar to the steps in the spectral optimization study. First, the problem parameters such as geometric dimensions, neutron fluxes, residence time, etc. must be determined for the specific situation to be studied. As is discussed in section 3.3, this study focuses on the transmutation of actinides in fast reactor target assemblies. Therefore, problem parameters will be based on documentation on the specific reactor modeled and calculated values generated from ERANOS, a fast reactor analysis code [21].

Next, the optimization search space needs to be defined, which in this study means creating a list of useable reactor materials. This list is then used as the input for the optimization scheme. However, because of its more complex search space and increased calculation time required for each iteration, this study doesn't implement a direct optimization scheme like that used in the spectral optimization; instead JMP, a statistical analysis program, is used to generate a list of material combinations to analyze that ensures the search space is sufficiently covered. The JMP functionality used in this study implements a design of experiment technique which takes the experimenter's knowledge of a problem as input to determine which experiments need to be performed to obtain complete results, thereby avoiding additional, time-consuming experimentation that provides no new information [19].

The neutronics calculations required for each iteration are very similar to those used in the spectral optimization study. The two main differences are that cross section generation will be performed using a one-dimensional, instead of an infinite homogeneous, transport calculation to account for spatial and energy self-shielding effects. Depletion calculations will be performed

with ORIGEN-S, the depletion module in the SCALE code system, instead of the simplified code written for spectral optimization. The time-dependent nuclide concentrations output by the depletion calculation will be used to evaluate cost functions in the same way as in the previous study.

3.3 Case Study Description

The optimization methodology described in the previous section could be applied to essentially any design parameter in any type of reactor. However, many prior fuel cycle studies have concluded that fast reactors are best for transmutation because of their favorable neutron balances, so this study will examine actinide transmutation in fast reactor target assemblies. The Prototype Fast Breeder Reactor, discussed in more detail in section 3.3.1, was selected for analysis because its design characteristics are representative of typical fast reactors and its core incorporates radial blanket (or target) assemblies for actinide management.

The material optimization case study models the transmutation of UNF in these target assemblies. In each run, the target assembly will be modeled with different volume fractions of UNF and various combinations of other reactor materials. Then depletion calculations will compute the time-dependent nuclide concentrations specific to that mixture of materials needed to evaluate several cost functions.

3.3.1 Fast Reactor Description

The Prototype Fast Breeder Reactor (PFBR), which is currently under construction in India, was selected because its core characteristics are representative of typical fast reactors and its design includes the use of blanket assemblies for breeding and/or transmutation. The PFBR is a 1250

MWt (500 MWe) sodium cooled fast reactor. The core has two zones of mixed oxide (MOX) driver fuel surrounded by two rows of blanket assemblies, two rows of reflector, and one row of shielding [14]. The core layout is shown in FIGURE 3.1. For simplicity in neutronics calculations, the PFBR core layout is altered slightly by modeling only one row of blanket assemblies followed by one row of reflector.

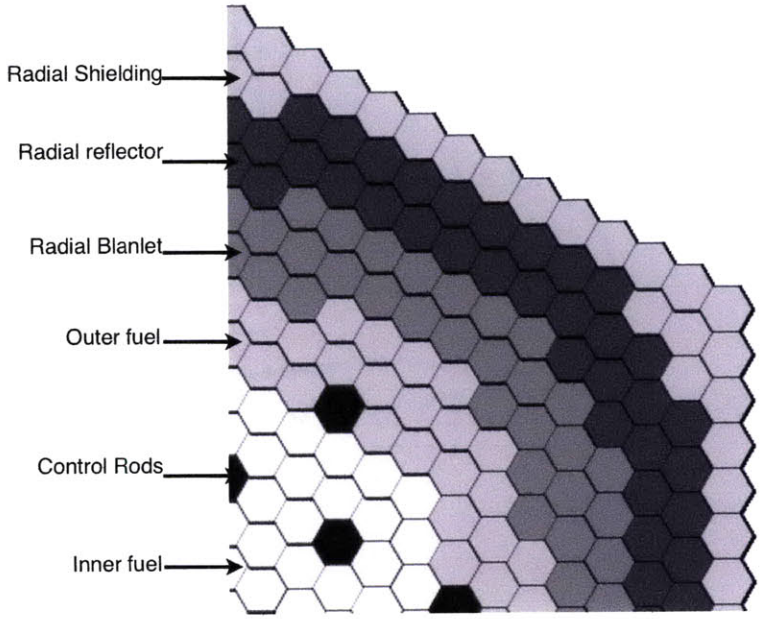

FIGURE 3.1 - Radial layout of PFBR core [14]

Although the blanket assemblies in the PFBR are loaded with depleted uranium (99.69% U238, 0.22% U235, and 0.09% U236) to breed plutonium, blanket assemblies in this study will be used as targets for transmutation of minor actinides to reduce the impact of UNF, not to breed additional fuel. It is noted, however, that this optimization method could be utilized in such a study. The structural material for both the fuel assemblies and the target assemblies is HT9

stainless steel, which is also the material used in the radial reflector assemblies [14]. This and other materials are discussed in more detail in section 3.3.3.

3.3.2 Used Nuclear Fuel Description

The same UNF vector described in section 2.3.1 is used as the basis for the material optimization study. However, the uranium, neptunium, and plutonium in UNF can be recycled as fuel for the driver core, so this study will focus on the transmutation of only the minor actinides. The UNF vector is, therefore, altered by removing uranium, neptunium, and plutonium through a partitioning scheme like TRUEX or pyro-processing [15, 16]. After removing these isotopes, the weight fractions of the remaining minor actinides (MA) were recalculated with the results shown in TABLE 3.1.

TABLE 3.1 - Initial minor actinide isotopic composition

Isotope	wt%
Am241	5.55E+01
Am242	1.71E-01
Am243	3.25E+01
Cm242	1.74E-03
Cm243	9.19E-02
Cm244	1.10E+01
Cm245	6.56E-01
Cm246	8.51E-02
Cm247	1.35E-03
Cm248	9.82E-05
Bk249	1.27E-07
Cf249	1.37E-06
Cf250	4.49E-07
Cf251	2.37E-07
Cf252	4.39E-10
TOTAL	100.00

3.3.3 Potential Assembly Materials

The list of materials included in this study was compiled by reviewing journal articles for descriptions of materials that have been successfully tested in existing reactors. Materials that would be chemically incompatible with sodium were not included. The list was also reduced by including only one example from groups of very physically similar materials (the various types of stainless steel, for example) unless they were thought to have significantly different neutronic properties. The results of this search are shown in TABLE 3.2 - Potential assembly materials. This list includes two materials, HT9 stainless steel and sodium, that must be included in every assembly design to conform to the PFBR design for cooling and structural requirements. Each

run must also include the minor actinide vector in the combination of materials. Additional information on each of these materials is included in the APPENDIX. Future studies will look at the transmutation of natural, depleted, and recycled uranium as well as thorium in fast reactor target assemblies.

TABLE 3.2 - Potential assembly materials

MINOR ACTINIDES	B_4C	YH_2	V	ZrH2
SODIUM	MgO	Al_2O_3	Mo	LEAD
HT9	ZrN	Gd_2O_3	Cr	GRAPHITE
CaH_2	SiC	Y_2O_5Zr	Be	CADMIUM
SS316	LiH	CeO_2	Nb	SPINEL
	WC	BeO	Zr	

3.3.4 Constraints

Since the objective of this study is to optimize transmutation in a more realistic situation, constraints must be set on the combinations of materials to ensure the target assembly design remains practical. This is done by placing limits on the amount of each material that can be placed in the assembly. Because of the format of ERANOS and SCALE input files, it is convenient to specify the amount of a material to be used by assigning it a fraction of the assembly volume. Therefore, maximum and minimum volume fractions are set for each material, with each run being required to include at least a minimum amount of sodium and HT9 because any assembly must be cooled and have a structural component and cladding. These limits are determined by the volume fractions of these two materials in the actual PFBR design, which are given in TABLE 3.3. Bounds were also set on the amount of minor actinides in

the assembly, but the upper limit is set more conservatively than the volume of fuel in the PFBR design because high concentrations of minor actinides harden the flux spectrum and diminish safety factors through decreasing rod worth, decreasing doppler coefficient, and increasing sodium void coefficient [17]. The volume fraction constraints used in this study are shown in TABLE 3.4 below.

TABLE 3.3 - PFBR homogenous assembly parameters [14]

	Volume Fractions (%)	
	Fuel Assembly	Blanket Assembly
Fuel	30.67	49.42
Void (He)	4.41	2.91
HT9	23.94	19.07
Sodium	40.98	28.60

TABLE 3.4 – Constraints on material volume fractions

	Volume Fractions (%)
	Target Assembly
Minor Actinides	10 - 20
HT9	19 - 62
Sodium	28 - 71
Other Materials	0 - 43

The second major constraint in this study is the irradiation time. Whereas in the spectral optimization study, the residence time in a reactor for UNF was set at 1,000 days to

approximate that of a typical LWR assembly, irradiation time is now determined by the fluence limit on pin cladding. The fluence limit for HT9 stainless steel, which is the cladding material in the PFBR, is approximately $4x10^{23}$ n/cm^2 [18]. Therefore, for every analysis of a combination of materials, the total scalar flux is calculated during the transport calculation for cross section processing and the irradiation time is calculated by dividing the fluence limit by the total scalar flux, assuming that the flux in the target assembly is constant over time. This is a valid assumption since the core is at constant power and the target assembly produces very little power, especially when actinide content is limited to below 20%.This ensures that each combination of materials is exposed to the same number of neutrons, thus avoiding biasing the results towards highly multiplying or low absorption materials simply because of their neutron balances, not necessarily their transmutation properties. In addition, the constraint on cooling time remains 10,000 years after the irradiation period to simulate permanent geologic storage, as was done in the spectral optimization study.

3.4 Design of Experiment

Instead of using a direct optimization routine like simulated annealing, this material composition optimization study uses the statistical analysis program, JMP, to perform experimental design. This method involves using JMP to determine, based on the volume fraction constraints, the number of runs and their associated material combinations that must be completed to cover the search space effectively and efficiently [19]. This method was chosen over simulated annealing because the longer computation time for neutronics calculations coupled with the trial and error approach required for determining optimization routine parameters would have been too inefficient. The mixture design of experiment method was used in JMP because it is tailored to design experiments in which various components are

combined so that the sum of the proportions of the mixture's constituents always equals one. The mixture experimental design method differs from similar techniques in that the factors cannot vary independently of one another; instead, changing the proportion of one constituent inherently means changing the proportion of another. In mixture problems, the search space is in the shape of a simplex and constraints like those placed on volume fractions in TABLE 3.4 have the effect of slicing edges off the search space. JMP uses the extreme vertices method to do experimental design on models such as these. The extreme vertices method selects cases at the each intersection of the search space and additional points between each vertex depending on the amount of *a priori* knowledge the experimenter has of the solution [19]. When applied to the 28 materials in this study and with the lack of information about the solution, JMP created a list of 11,408 separate cases to run.

3.5 Neutronic Methodology

In order to generate flux-weighted cross sections and perform depletion calculations, the flux spectrum within the target assembly must be determined. This spectrum is determined by the material composition of the target assembly and the flux entering it from the reactor. Determining this flux spectrum could be accomplished by performing a three-dimensional, full-core calculation for each combination of target materials but unfortunately these calculations are very computationally expensive (approximately one hour to complete). To circumvent the need for a full-core calculation for every combination of materials, an approximation is made: the flux entering the target assembly remains largely the same regardless of the material composition of the target. This is valid because the volume of the core is much larger than the volume of target assemblies (181 driver fuel assemblies versus 68 target assemblies), the target assemblies produce very little power relative to the driver assemblies, and the target assemblies

are located along the periphery of the core [14]. This means that the composition of the target assembly will have very little effect on the power distribution within the core and on the incident flux spectrum. Therefore, a full-core calculation is done to determine the incident flux on the target assembly from the core and this flux is used as a boundary condition for much faster one-dimensional transport calculations that model the target assembly and the adjacent reflector. The validity of the approximation is evaluated in section 3.5.3. As is explained in more detail in subsections 3.5.1 and 3.5.2, ERANOS is used to perform the full-core calculation and SCALE is used to generate flux-weighted cross sections and perform depletion calculations for each set of target assembly materials.

3.5.1 Fast Reactor Calculations

ERANOS is a deterministic fast reactor neutronics code developed by the CEA that performs core calculations using either a diffusion theory solver or a SN transport approximation solver. The diffusion solver performs calculations in two- and three-dimensional hexagonal and Cartesian coordinates while the SN solver uses R-Z cylindrical, two-dimensional X-Y, or R-θ coordinates. The latest version, ERANOS 2.1, was released in 2006 and has been extensively benchmarked against existing fast reactors [21].

To determine the flux incident on a target assembly, a full-core calculation of the PFBR was performed using the JEFF3.1 cross section library collapsed to 33 energy groups by the ECCO module and the SN solver in R-Z coordinates. Energy-dependent angular flux values, $\varphi(r,E,\Omega_n)$, at the core-target interface were extracted from the ERANOS output file as well as the discrete ordinates quadrature weighting factors, w_n. These were used to determine the scalar flux entering the target assembly using

EQUATION 3.1 [5]

$$\Phi(\mathrm{r}, E) = \sum_{n} w_n \varphi(\mathrm{r}, E, \Omega_n)$$

where the summation is only over angles, Ω_n, that enter the target assembly from the core region. The resulting incident flux spectrum is shown in FIGURE 3.2 – Flux spectrum incident upon target assembly.

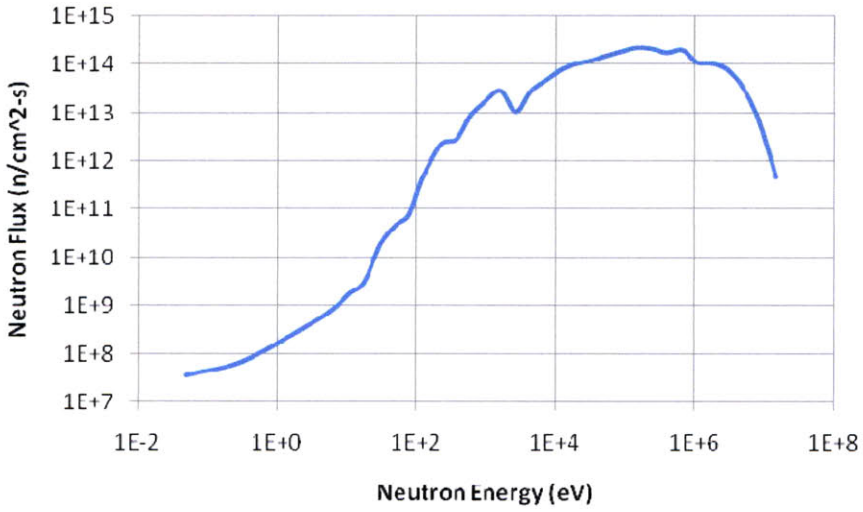

FIGURE 3.2 – Flux spectrum incident upon target assembly

3.5.2 Target Assembly Calculations

For each combination of target assembly materials, a transport calculation must be performed to generate flux-weighted cross sections that are then used in a depletion calculation. As in the spectral optimization study, the SCALE code package will be used for these calculations. However, instead of performing an infinite, homogeneous medium transport calculation, the system will be modeled as two, one-dimensional slabs with a flux boundary condition. The

system consists of two zones, one for the target assembly and one for the reflector assembly, which are illustrated below, with an incident neutron flux on the left and a vacuum boundary on the right. Each slab is 13.5 cm wide to match the assembly pitch in the PFBR [14].

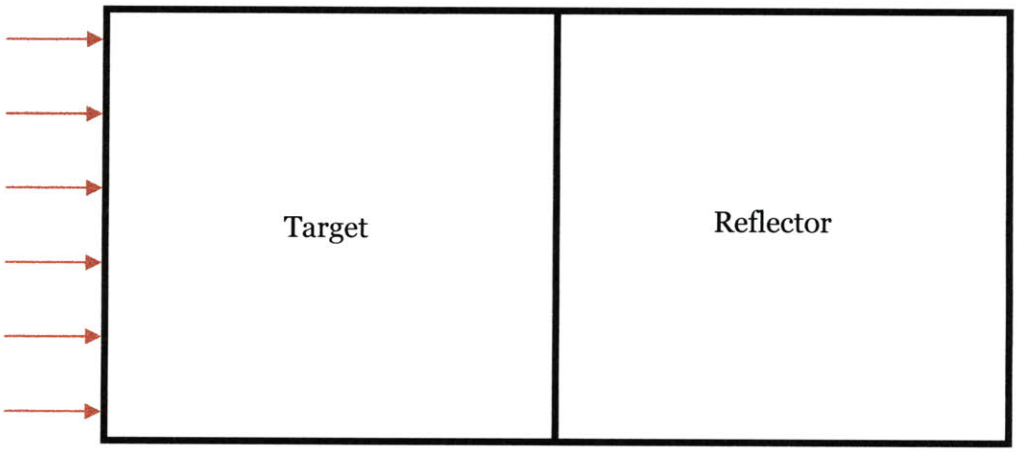

FIGURE 3.3 - One-dimensional model of target and reflector assemblies

To perform these calculations, first a combination of materials is selected from the list generated by JMP and formatted to match the SCALE material composition input format by a Fortran script that then generates an input file that calls several SCALE functional modules. The first of these is BONAMI, which performs Bondarenko calculations on pointwise continuous master library cross sections to account for resonance self-shielding effects in energy. The CENTRM module is then called to construct a pointwise continuous flux spectrum, which is then passed to PMC to collapse cross sections into a flux-weighted multigroup format. XSDRNPM solves the Boltzmann transport equation using the SN method for the one-dimensional model shown in FIGURE 3.3 to generate cell-weighted, multigroup cross sections that account for spatial self-shielding effects. WAX and COUPLE are used to format these cross sections and update the three-group master libraries for use in depletion calculations by ORIGEN-S [3].

Before ORIGEN-S performs the depletion calculations, execution is halted while a Fortran script extracts the total scalar flux within the target assembly calculated by XSDRNPM and inputs this value as the flux value used by ORIGEN-S during the irradiation period. This script also uses the scalar flux value to calculate how long until the HT9 cladding fluence limit is reach and passes this value to ORIGEN-S as the length of the irradiation period. Execution is then restarted and ORIGEN-S performs a depletion calculation with an irradiation period lasting until the cladding fluence limit is reached followed by a 10,000 year cooling period.

These calculations are automated using a Fortran script, so after each SCALE calculation is complete, another combination of materials is read from the list created by JMP and the neutronics calculations are performed again. This sequence is repeated until all 11,408 combinations of materials have been evaluated, resulting in approximately 250 GB of data.

3.5.3 Validation

To ensure the approximation that target assembly composition does not significantly influence the neutron current into the target, the assembly-averaged flux spectrum within the target was calculated using a full-core model and the one-dimensional model illustrated in FIGURE 3.3. Both models simulated the PFBR core surrounded by one row of target assemblies composed of HT9 stainless steel, sodium, and LWR UNF (TABLE 2.1) using the volume fractions of a PFBR blanket assembly (TABLE 3.3) and one row of reflector. The full-core calculation was performed with the ERANOS diffusion solver in three-dimensions using the 172 group JEFF3.1 cross section library. The one-dimensional calculation was performed using the same sequence of SCALE modules as described in section 3.5.2 but without the depletion step. The resulting flux spectra are shown in FIGURE 3.4.

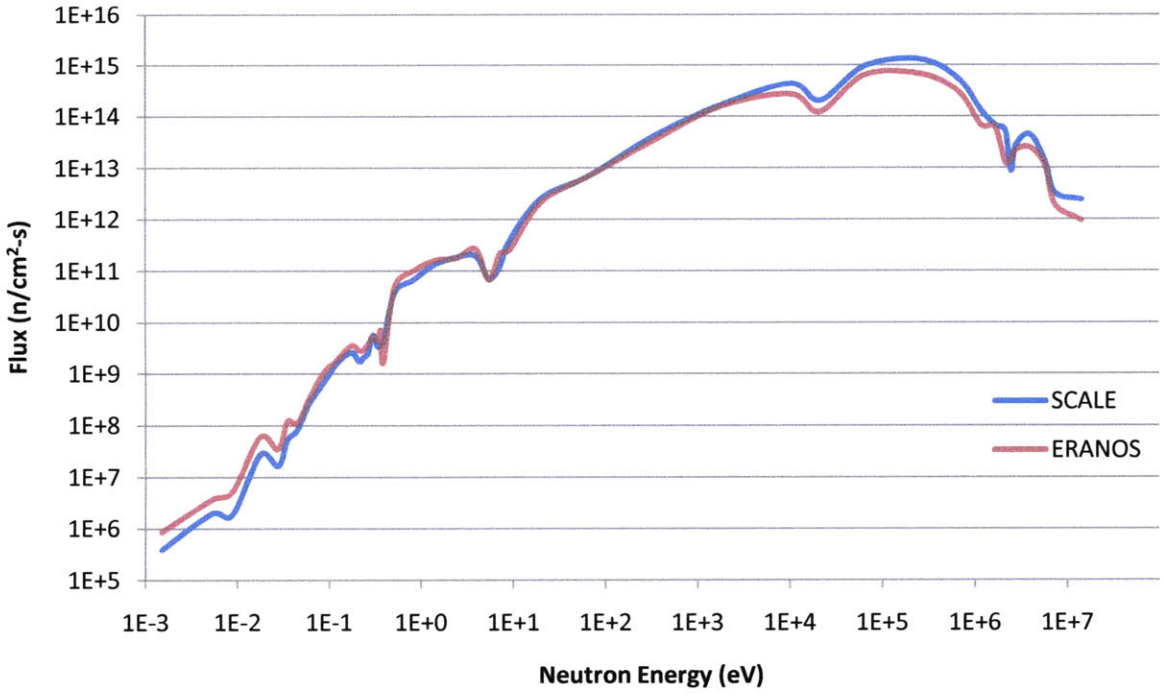

FIGURE 3.4 - Assembly-averaged flux spectra in target

Although the difference in complexity between the full-core and one-dimensional models is great, the above figure shows remarkable accuracy in the shape of the flux spectrum produced by the slab model. The spectra match particularly well in the epithermal region and the discrepancies in the low flux region may be due to less accurate cross section handling in the thermal region by ERANOS since it is a fast reactor neutronics code. The deviation in magnitude in the fast region, where neutrons are more likely to leak out of the system because of their large mean free paths, is likely caused by the differences in geometry; the full-core ERANOS model allows neutrons to leak axially whereas the one-dimensional SCALE model is infinite in the axial direction so no axial leakage occurs. Regardless, the agreement between the spectra suggests that the approximations allowing the use of the one-dimensional model with a flux boundary are valid for the purposes of this study.

3.5.4 Cost Functions

After all neutronics calculations are performed, another Fortran script is called to evaluate cost functions for each combination of target materials. This script extracts pertinent data and calculates integrated decay heat production values for each SCALE output file.

Integrated decay heat production is evaluated over the first 10 years, 100, 1,000, and 10,000 years of cooling. Because ORIGEN-S follows actinides as well as fission products, integrated decay heat production is evaluated three times, once for all isotopes, once for actinides only, and once for fission products only. These values are divided by the initial mass of minor actinides (kW-years/gram MA) in the target assembly so that results are not biased toward small minor actinides loadings.

The plutonium content will not be optimized in this study as it was for spectral optimization because of the removal of plutonium and its parent elements, uranium and neptunium, from the UNF vector. Although plutonium will still be present after irradiation because of alpha decay of curium, its concentration is miniscule and optimization of its isotopic concentration would provide little insight into transmutation properties of materials.

3.6 Results

Two reference cases are evaluated for benchmarking the transmutation effectiveness of various material combinations. The first of these is called the "natural decay" case and it models the natural decay of the minor actinide vector with no irradiation. The second is called the "typical target" case, which models a target assembly using only those materials present in the PFBR blanket assembly design. Its composition is 20% MA, 30% sodium, and 50% HT9 by volume.

These reference cases allow comparison among the decay heat released after transmutation in optimal target assembly designs, transmutation in standard assemblies like the blanket assemblies in the PFBR, and immediate storage of untransmuted UNF.

The first cost function to be evaluated in this study is the integrated heat production from actinides and fission products over the first 10 years of cooling. All of the best performing material combinations have minimal amounts of the required materials (10% MA, 28% sodium, and 19% HT9) and fill the remaining volume hydride materials: zirconium hydride, lithium hydride, yttrium hydride, and calcium hydride. The three best performing target assemblies are shown in TABLE 3.5. The optimal combination includes the maximum amount of zirconium hydride (43%) and results in an integrated heat production value of 4.2 W-years/gram MA, while the natural decay case produces a total of 3.9 W-years/gram MA over the first ten years. Thus, integrated decay heat production over 10 years actually increases by 7% through irradiation in an optimal target assembly. This is explained by noting that the integrated actinide heat production value for the ZrH2 target is 3.6 W-years/gram MA, which is slightly lower than that of natural decay. The remaining 0.6 W-years/gram MA is released by the fission products created during irradiation. When compared to the 7.6 W-years/g MA produced in the typical target case, the ZrH2 target vastly outperforms by almost 45%.

TABLE 3.5 - Top performing target assemblies for integrated heat production over 10 years

Case	ZrH_2 (%)	LiH (%)	YH_2 (%)	W-years/g MA
1	43	0	0	4.20
2	22	22	0	4.29
3	0	43	0	4.34
4	14	14	14	4.41
Natural Decay	0	0	0	3.90
Typical Target	0	0	0	7.57

The presence of a very efficient moderating material like zirconium hydride in a target assembly optimized to reduce short-term heat production is very surprising given the results of the spectral optimization study, which found that an exclusively fast flux is optimal for reducing short-term heat production. These two results can be reconciled by examining the depletion calculations used in each study; the spectral optimization study uses a depletion code that only tracks actinides, so when actinides fission, their fission products are not tracked and thus do not contribute to heat production cost function. In this study, however, fission products are also tracked and since many fission products produce large amounts of heat, the fissioning of actinides caused by a fast flux becomes less effective relative to the neutron capture that dominates in a thermal flux.

The same calculations were performed for integrated decay heat production over 100 years, 1,000 years, and 10,000 years of cooling but the target assembly with the maximum amount of ZrH_2 performed best in all of them. The rest of the hydrides followed closely behind, with LiH as the second best performer, YH_2 in third, and CaH_2 in fourth. As the cooling length increase, so did the savings in integrated heat production by ZrH_2 versus natural decay or a typical target assembly. The results of these two comparisons are shown numerically in TABLE 3.6 and TABLE 3.7 and graphically in FIGURE 3.5.

TABLE 3.6 - Integrated heat production of natural decay and ZrH_2 target assembly

Cooling time (years)	W-years/g MA		Change (%)
	Natural Decay	ZrH_2	
100	14.4	15.1	5.3
1,000	43.0	23.6	-45.2
10,000	69.7	31.5	-54.9

TABLE 3.7 - Integrated heat production of typical and ZrH2 target assemblies

Cooling time (years)	W-years/g MA		Change (%)
	Typical Target	ZrH_2	
100	20.9	15.1	-27.6
1,000	41.2	23.6	-42.7
10,000	61.2	31.5	-48.6

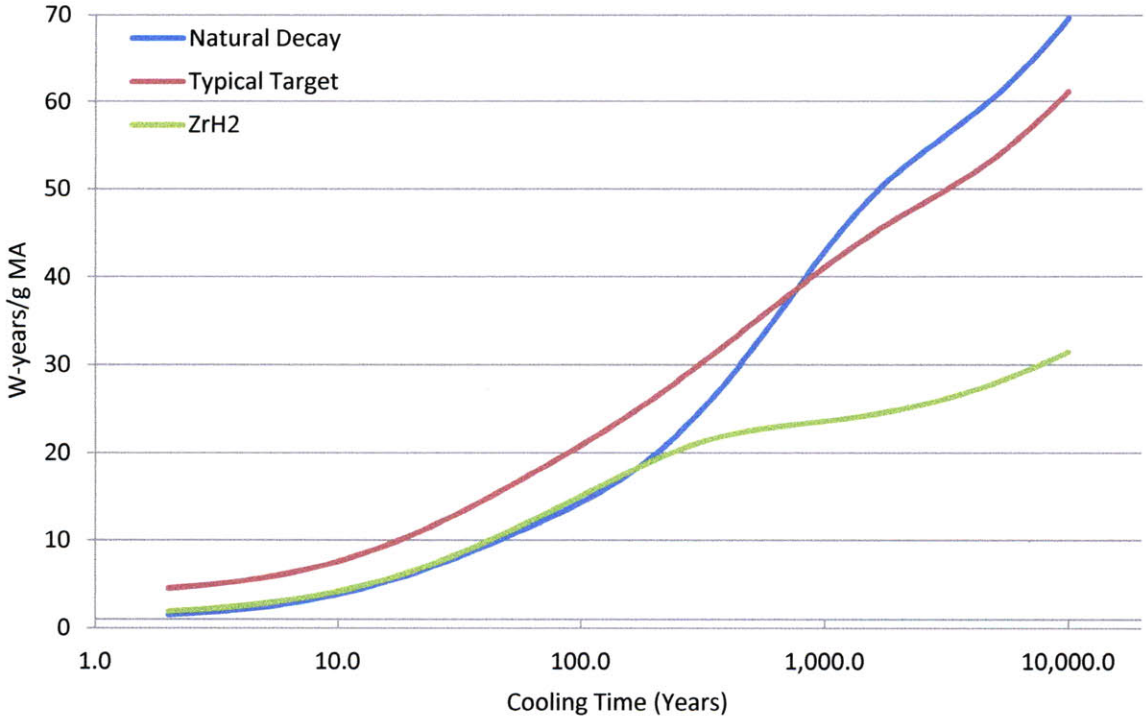

FIGURE 3.5 - Integrated decay heat production

Although this trend is rather apparent, the average integrated decay heat production versus ZrH_2 volume fraction is plotted in FIGURE 3.6. This shows that as ZrH_2 content increases, the total amount of decay heat released generally decreases, with intermediate peaks caused by cases that include effective materials to fill the remaining volume. This general trend is confirmed by the contour map in FIGURE 3.7, which shows that integrated decay heat is minimized by maximizing ZrH_2 content and minimizing MA content.

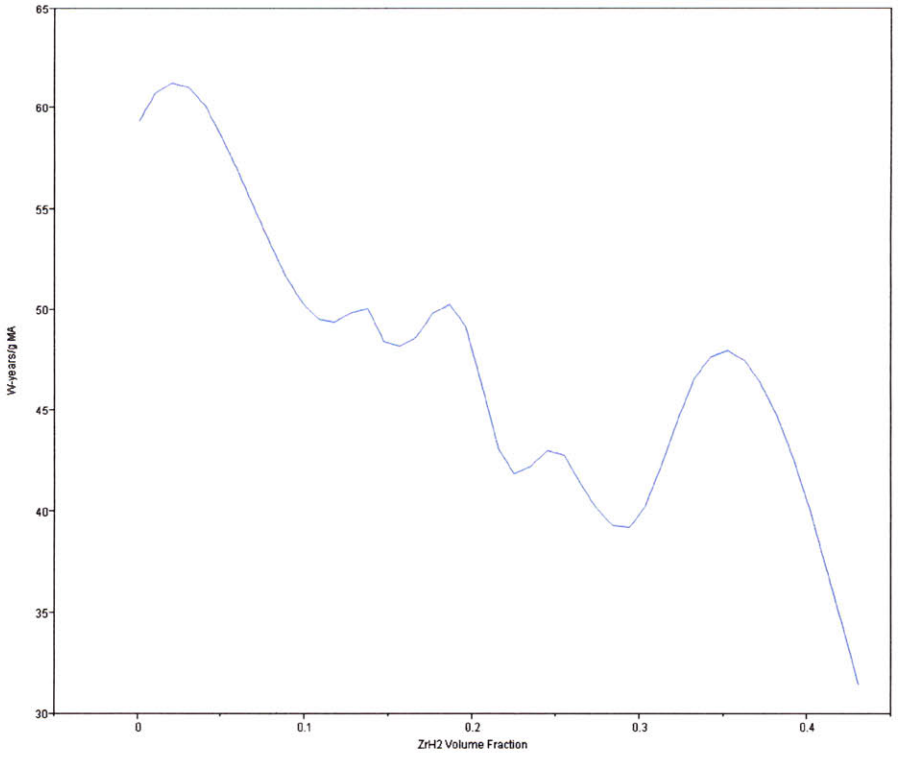

FIGURE 3.6 - Average effect of ZrH_2 content on integrated heat production over 10,000 years

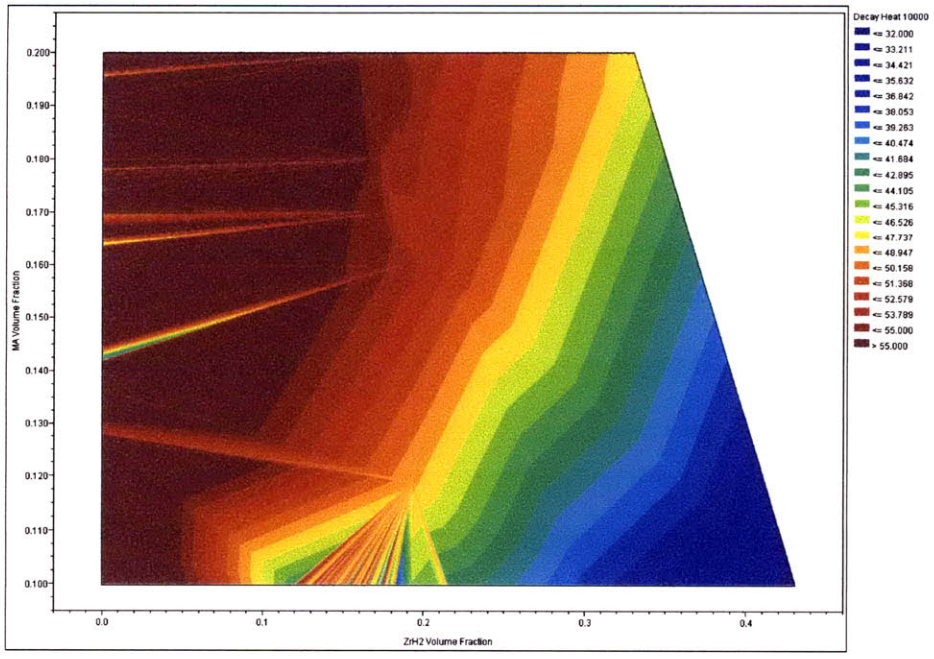

FIGURE 3.7 – ZrH_2 and MA content versus integrated heat production over 10,000 years

It was concluded by the results of the spectral optimization study that a thermal spectrum is more effective at transmuting actinides to reduce long-term decay heat production. This conclusion is supported by the results presented in this section since the materials that were found to optimize transmutation were also the best moderators. The moderating ability of the hydrides comes from their high density of hydrogen and zirconium hydride is particularly effective because of its very low probability of parasitic neutron absorption. The flux spectrum in FIGURE 3.8 is a testament to zirconium hydride's moderating ability as it increases the thermal flux by several orders of magnitude relative to a target assembly comprised of minor actinides, HT9 stainless steel, and sodium coolant.

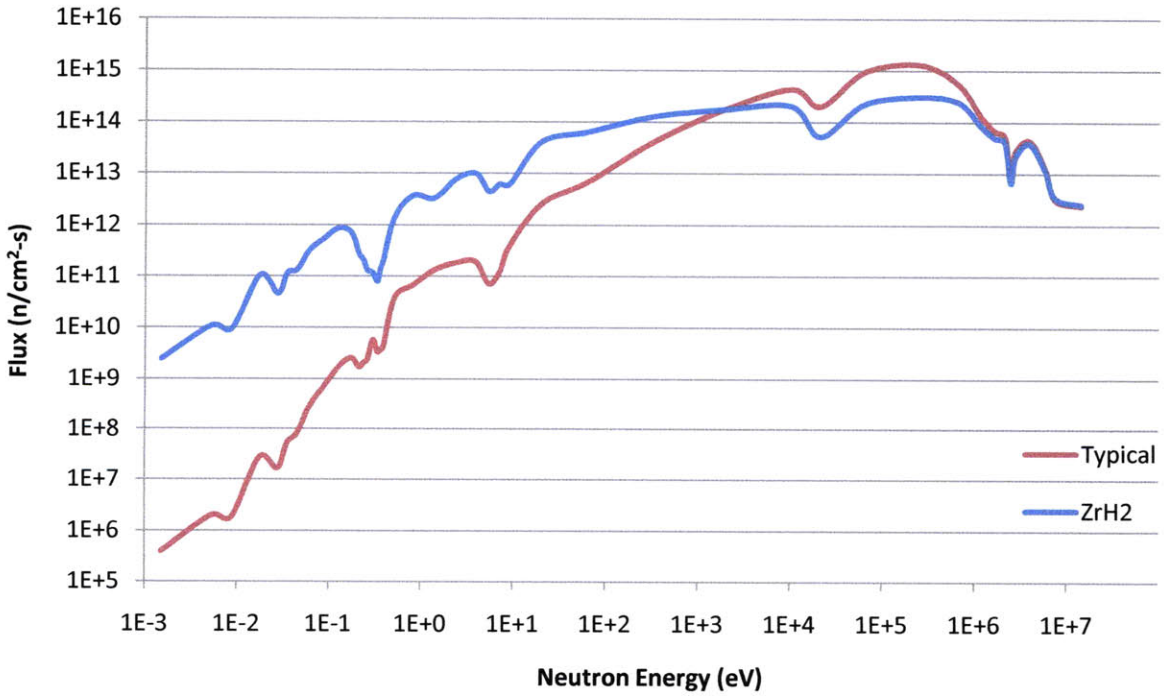

FIGURE 3.8 - Flux spectra in a typical target assembly and a ZrH_2 target assemby

4 CONCLUSIONS AND FUTURE WORK

4.1 Conclusions

The spectral optimization study was performed with the understanding that all of the assumptions necessary to perform such a study ensured that it would be a largely theoretical exercise. This, however, did not dampen the enthusiasm behind the project because hypothetical exercises can occasionally breathe insight into otherwise mystifying phenomena. Although the energy dependence of transmutation properties is not completely unknown, it appears that it has been largely unexplored, at least quantitatively.

The resulting information from this highly approximate optimization study indicated that while the details of optimal flux spectra may be difficult to parse, their general shapes can be understood by examining only a few key cross sections. Thus, for immediate destruction of an isotope, a fast flux is needed, whereas a thermal spectrum is more important for long-term alteration of isotopic concentrations.

Although the optimization of plutonium isotopics was an interesting theoretical exercise, the unusual flux shapes required for optimal transmutation would be all but impossible to replicate in a reactor. The results for the integrated decay heat production, however, were both theoretically interesting as well as potentially reproducible in physical situations. These showed that by keeping only the fast and thermal flux regions of an incident spectrum, the total decay heat produced by UNF over 10,000 years can be reduced by almost 50%. It was also found that this long-term transmutation is almost exclusively due to the thermal part of the spectrum.

Because of the approximations made in the spectral optimization study, a more realistic situation needed to be modeled to prove the applicability of the methodology developed here.

This led to the study of optimal target assembly material compositions for transmutation of UNF in a fast reactor. By evaluating thousands of material combinations, it was found that the best materials for transmutation are those that are the most efficient moderators. By loading the target assembly with zirconium hydride, which provides both a high density of hydrogen and a small absorption cross section, the long-term decay heat production can be reduced by over 50% compared to the natural decay of minor actinides. This result encourages future work on this topic as it reinforces the dominance of a thermal flux for long-term transmutation as well as proving the applicability of the optimization methodology for isotope management outlined in this work to realistic situations.

4.2 Future Work

Because the research area of optimal transmutation is largely unexplored, there are a number of interesting topics that remain unstudied. Several of the topics that can be viewed as an extension of this work include:

- Combined target and reflector material optimizations
- Multivariate optimization
- Target material optimization with natural, depleted, and recycled uranium
- Target material optimization with thorium
- Transmutation optimization in light water reactors
- Core geometry optimization
- Economic cost function evaluation for transmutation optimization

5 REFERENCES

[1] T. Todd, "Spent Nuclear Fuel Reprocessing," (2008) [Online]. http://www.ne.doe.gov/pdfFiles/NRCseminarreprocessing_Terry_Todd.pdf

[2] Argonne National Laboratory, "Advanced Fuel Cycle Initiative," (2009) [Online]. http://www.ne.anl.gov/research/afc/

[3] Oak Ridge National Laboratory, "SCALE: A Modular Code System for Performing Standardized Computer Analyses for Licensing Evaluations," (2009).

[4] U.S. Department of Energy, "DOE Fundamentals Handbook: Nuclear Physics and Reactor Theory," (1993).

[5] J.J. Duderstadt and L. J. Hamilton, *Nuclear Reactor Analysis,* John Wiley & Sons, (1976).

[6] U.S. Nuclear Regulatory Commission, "Licensing Yucca Mountain: Fact Sheet." (2009).

[7] U.S. Nuclear Regulatory Commission, "Disposal of High-Level Radioactive Wastes in a Geologic Repository at Yucca Mountain, Nevada," 10 Code of Federal Regulations Part 63, (2007)

[8] Oak Ridge National Laboratory, "SCALE Newsletter, January" (2009, January).

[9] *Chart of the Nuclides*, 16th ed., Knolls Atomic Power Laboratory (2002).

[10] D.T. Phamand and D. Karaboga, *Intelligent Optimisation Techniques*, 2nd ed., Springer-Verlag Limited, London (2000).

[11] H.W. Wiese,"Actinide transmutation properties of thermal and fast fission reactors including multiple recycling," *Journal of Alloys and Compounds*, **271-273**, 522 (1998).

[12] M. Salvatores et al., "The Physics of TRU Transmutation – A Systematic Approach to the Intercomparison of Systems," *The Physics of Fuel Cycles and Advanced Nuclear Systems: Global Developments (PHYSOR 2004),* Chicago, Illinois, April 25-29, 2004, American Nuclear Society (2004).

[13] T. Toshikazuand et al., "Interpretation of actinide transmutation in thermal and fast reactors," *Progress in Nuclear Energy*, **40**, 449 (2002).

[14] N. Stauff, B. Forget, and M. Driscoll, "Resolution of proliferation issues for a SFR blanket with a specific application," *GLOBAL 2009 – The Nuclear Fuel Cycle*, Paris, September 6-11, 2009, French Nuclear Energy Society (2009).

[15] M. Ozawa et al., "Separation of actinides and fission products in high-level liquid wastes by the improved TRUEX process," *Journal of Alloys and Compounds*, **271-273**, 538 (1998).

[16] T. Inoue, "Actinide recycling by pyro-process with metal fuel for future nuclear fuel cycle system," *Progress in Nuclear Energy*, **40**, 3-4, 547 (2002).

[17] K. Tucek et al., "comparison of sodium and lead-cooled fast reactors regarding severe safety and economical issues," *13th International Conference on Nuclear Engineering*, Beijing, May 16-20, 2005, Chinese Nuclear Society (2005).

[18] D.C. Crawford et al., "Requirements and long-term development plan for fast-spectrum transmutation fuels in the U.S.," *Seventh Information Exchange Meeting on Actinide and Fission Product Partitioning and Transmutation*, October 14-16, 2002, Jeju, Korea (2002).

[19] *JMP 8 Design of Experiments Guide*, SAS Institute Inc., Cary, NC (2008).

[20] N. Stauff, "Resolution of proliferation issues for a Sodium Fast Reactor blanket," Massachusetts Institute of Technology, Cambridge, MA, M.S. Thesis (2008).

[21] J.M. Ruggieri et al., "ERANOS 2.1 : International Code System for GEN IV Fast Reactor Analysis," *Proc. Int. Congress on Advances in Nuclear Power Plants (ICAPP '06)*, Reno, Nevada, June 4–8, 2006, 233, American Nuclear Society (2006).

[22] S. Yamanaka, et al., "Thermal and mechanical properties of zirconium hydride," *Journal of Alloys and Compounds*, **293-295**, 23 (1999).

[23] J. G. Reynolds, "Spinel structure and liquidus temperature relationships in nuclear waste glass," *Journal of Material Science*, **40**, 3987 (2005).

[24] Y. Kato, et al., "Kinetic study of the hydration of magnesium oxide for a chemical heat pump," *Applied Thermal Engineering*, **16**, 853 (1996).

[25] S. J. Zinkle, et al., "Overview of materials technologies for space nuclear power and propulsion," *AIP Conf. Proc.*, **608**, 1063 (2002).

[26] A. Lofberg, et al., "Mechanism of WO3 Reduction and Carburization in CH4/H2 Mixtures Leading to Bulk Tungsten Carbide Powder Catalysts," *Journal of Catalysis*, **189**, 170 (2000).

[27] R. Van Houten, "Selected engineering and fabrication aspects of nuclear metal hydrides (Li, Ti, Zr, and Y)," *Nuclear Engineering and Design*, **31**, 434 (1974).

[28] R. J. M. Konings, et al., "The influence of neutron irradiation on the microstructure of Al2O3, MgAl2O4, Y3Al5O12 and CeO2", *Journal of Nuclear Materials*, **254**, 135 (1998).

[29] K. Une and S. Kashibe, "Fission gas release during postirradiation annealing of UO2-2 wt% Gd2O3 fuels," *Journal of Nuclear Materials*, **189**, 210 (1992).

[30] M. Burghartz, et al., "Some aspects of the use of ZrN as an inert matrix for actinide fuels," *Journal of Nuclear Materials*, **288**, 233 (2001).

[31] K. Sasaki, et al., "Helium release and microstructure of neutron-irradiated SiC ceramics," *Journal of Nuclear Materials*, **179-181**, 407 (1991).

[32] B.A. Loomis, et al., "Evaluation of low-activation vanadium alloys for use as structural material in fusion reactors," *Journal of Nuclear Materials*, **179-181**, 148 (1991).

[33] C. Degueldre, et al., "Behaviour of implanted xenon in yttria-stabilised zirconia as inert matrix of a nuclear fuel," *Journal of Nuclear Materials*, **289**, 115 (2001).

[34] T.A. Taylor, et al., "Development of several new nickel aluminide and chromium carbide coatings for use in high temperature nuclear reactors," *Thin Solid Films*, **107**, 427 (1983).

[35] M. K. Meyer, et al., "Low-temperature irradiation behavior of uranium-molybdenum alloy dispersion fuel," *Journal of Nuclear Materials*, **304**, 221 (2002).

[36] H. Kleykamp, "The chemical state of the fission products in oxide fuels," *Journal of Nuclear Materials*, **131**, 221 (1985).

[37] L. Green, "Energy storage via calcium hydride production," *Intersociety Energy Conversion Engineering Conference Proceedings*, **1**, 949 (1976).

[38] International Nuclear Safety Center, "INSC Material Property Database," (2004) [Online]. http://www.insc.anl.gov/matprop/

APPENDIX

This appendix contains the information characterizing each material used in the target assembly material composition optimization study in chapter 3. The list of materials that have been tested in reactors was compiled from journal articles and references for each material are listed in the following table. The table also lists the density and isotopic composition of each material. For simplicity, the hydrides were assumed to take the form of $X\text{-}H_2$, except for lithium hydride (LiH).

Material	Reference	Density	Element	wt %	Atomic Mass
MINORACT	3	13.65090	Am241	5.5513E+01	241.0568
			Am242	1.7100E-01	242.0600
			Am243	3.2500E+01	243.0610
			Cm242	1.7397E-03	242.0584
			Cm243	9.1943E-02	243.0610
			Cm244	1.0980E+01	244.0626
			Cm245	6.5574E-01	245.0653
			Cm246	8.5100E-02	246.0669
			Cm247	1.3510E-03	247.0725
			Cm248	9.8242E-05	248.0721
			Bk249	1.2741E-07	249.0797
			Cf249	1.3735E-06	249.0797
			Cf250	4.4938E-07	250.0763
			Cf251	2.3747E-07	251.0799
			Cf252	4.3945E-10	252.0815
ZRH2	22	5.61000	Zr90	50.33709	89.9047
			Zr94	17.00405	93.9063
			Zr92	16.77903	91.9050
			Zr91	10.97730	90.9056
			H	2.01631	1.0078
SS316	25	7.94000	Fe	68.37500	55.8447
			Cr	19.00000	51.9957

			Ni	9.50000	58.6868
			Mn	2.00000	54.9381
			Si	1.00000	28.0853
			C	0.08000	12.0111
			P	0.04500	30.9741
SODIUM	14	0.97000	Na	100.00	22.9895
SPINEL	23	3.64500	Mg	17.08420	24.3051
			Al	42.14870	26.9818
			O	40.76710	15.9954
CADMIUM	38	8.64200	Cd	100.00	112.4258
B4C	25	2.52000	B10	14.43150	10.0129
			B11	63.86850	11.0093
			C	21.70000	12.0111
MGO	24	3.58000	Mg	60.30360	24.3051
			O	39.69640	15.9954
GRAPHITE	38	2.30000	C12	100.00	12.0000
HT9	14	7.61100	Fe	85.51306	55.8447
			Cr	11.56950	51.9957
			Ni	0.50300	58.6868
			Mo92	0.22394	91.9068
			Mo94	0.13959	93.9051
			Mo95	0.24024	94.9058
			Mo96	0.25171	95.9047
			Mo97	0.14411	96.9060
			Mo98	0.36414	97.9054
			Mo100	0.14532	99.9075
			Mn	0.90543	54.9381
LEAD	38	11.34400	Pb	100.00	207.2000
LIH	25	0.82000	Li7	87.30000	6.9410
			H	12.70000	1.0078
WC	26	8.64000	W	93.87000	183.8400

			C	6.13000	12.0111
YH2	27	4.29000	Y89	97.78250	88.9059
			H	2.21750	1.0078
AL2O3	28	3.91000	Al27	52.93000	26.9818
			O16	47.07000	15.9954
GD2O3	29	7.40700	Gd152	0.16752	151.9198
			Gd154	1.85231	153.9208656
			Gd155	12.65631	154.922622
			Gd156	17.61780	155.9221227
			Gd157	13.55556	156.9240
			Gd158	21.65313	157.9241039
			Gd160	19.29738	159.9270541
			O16	13.20000	15.9954
ZIRCN	30	7.09000	Zr90	44.60715	89.9047
			Zr94	15.06846	93.9063
			Zr92	14.86905	91.9050
			Zr91	9.72774	90.9056
			N14	13.30000	14.0070
SIC	31	3.21000	Si	70.00000	28.0855
			C	30.00000	12.0111
NB	38	8.57000	Nb93	100.00	92.9064
V	32	5.96000	V	100.00	50.9415
Y2O5ZR	33	5.80000	Y89	50.90000	88.9059
			O16	22.90000	15.9954
			Zr90	13.47990	89.9047
			Zr94	4.55356	93.9063
			Zr92	4.49330	91.9050
			Zr91	2.93964	90.9056
CR	34	7.20000	Cr	100.00	51.9957
BE	25	1.85000	Be9	100.00	9.0122

BEO	25	3.02000	Be9	36.00000	9.0122
			O16	64.00000	15.9954
ZR	38	6.49000	Zr90	51.45000	89.9047
			Zr94	17.38000	93.9063
			Zr92	17.15000	91.9050
			Zr91	11.22000	90.9056
MO	35	10.20000	Mo92	14.84000	91.9068
			Mo94	9.25000	93.9051
			Mo95	15.92000	94.9058
			Mo96	16.68000	95.9047
			Mo97	9.55000	96.9060
			Mo98	24.13010	97.9054
			Mo100	9.63000	99.9075
CEO2	36	7.13000	Ce140	72.18490	139.9054
			Ce142	9.21510	141.9092
			O16	18.60000	15.9954
CAH2	37	1.70000	Ca	95.21000	40.078
			H	4.79000	1.0078