ABSTRACT

HAYS, ROSS DANIEL. Stochastic Optimization for Nuclear Facility Deployment Scenarios. (Under the direction of Paul Turinsky.)

Single-use, low-enriched uranium oxide fuel, consumed through several cycles in a light-water reactor (LWR) before being disposed, has become the dominant source of commercial-scale nuclear electric generation in the United States and throughout the world. However, it is not without its drawbacks and is not the only potential nuclear fuel cycle available. Numerous alternative fuel cycles have been proposed at various times which, through the use of different reactor and recycling technologies, offer to counteract many of the perceived shortcomings with regards to waste management, resource utilization, and proliferation resistance. However, due to the varying maturity levels of these technologies, the complicated material flow feedback interactions their use would require, and the large capital investments in the current technology, one should not deploy these advanced designs without first investigating the potential costs and benefits of so doing. As the interactions among these systems can be complicated, and the ways in which they may be deployed are many, the application of automated numerical optimization to the simulation of the fuel cycle could potentially be of great benefit to researchers and interested policy planners.

To investigate the potential of these methods, a computational program has been developed that applies a parallel, multi-objective simulated annealing algorithm to a computational optimization problem defined by a library of relevant objective functions applied to the Verifiable Fuel Cycle Simulation Model (`VISION`, developed at the Idaho National Laboratory). The `VISION` model, when given a specified fuel cycle deployment scenario, computes the numbers and types of, and construction, operation, and utilization schedules for, the nuclear facilities required to meet a predetermined electric power demand function. Additionally, it calculates the location and composition of the nuclear fuels within the fuel cycle, from initial mining through to eventual disposal. By varying the specifications of the deployment scenario, the simulated annealing algorithm will seek to either minimize the value of a single objective function, or enumerate the trade-off surface between multiple competing objective functions. The available objective functions represent key stakeholder values, minimizing such important factors as high-level waste disposal burden, required uranium ore supply, relative proliferation potential, and economic cost and uncertainty. The optimization program itself is designed to be modular, allowing for continued expansion and exploration as research needs and curiosity indicate.

The utility and functionality of this optimization program are demonstrated through its application to one potential fuel cycle scenario of interest. In this scenario, an existing legacy LWR fleet is assumed at the year 2000. The electric power demand grows exponentially at a

rate of 1.8% per year through the year 2100. Initially, new demand is met by the construction of 1-GW(e) LWRs. However, beginning in the year 2040, 600-MW(e) sodium-cooled, fast-spectrum reactors operating in a transuranic burning regime with full recycling of spent fuel become available to meet demand. By varying the fraction of new capacity allocated to each reactor type, the optimization program is able to explicitly show the relationships that exist between uranium utilization, long-term heat for geologic disposal, and cost-of-electricity objective functions. The trends associated with these trade-off surfaces tend to confirm many common expectations about the use of nuclear power, namely that while overall it is quite insensitive to variations in the cost of uranium ore, it is quite sensitive to changes in the capital costs of facilities. The optimization algorithm has shown itself to be robust and extensible, with possible extensions to many further fuel cycle optimization problems of interest.

Stochastic Optimization for Nuclear Facility Deployment Scenarios

by
Ross Daniel Hays

A dissertation submitted to the Graduate Faculty of
North Carolina State University
in partial fulfillment of the
requirements for the Degree of
Doctor of Philosophy

Nuclear Engineering

Raleigh, North Carolina

2012

APPROVED BY:

James Wilson	Matthias Stallmann
Robert White	John Mattingly

Paul Turinsky
Chair of Advisory Committee

BIOGRAPHY

Ross Hays hails originally from the small town of Burlington, in southeastern Iowa. He is the older of two sons of Dr. Carl Hays and Billie Hays. After receiving degrees in Chemical Engineering and Applied Mathematics from Northwestern University in 2004, he moved to Raleigh, North Carolina to pursue graduate degrees in Nuclear Engineering.

ACKNOWLEDGEMENTS

This work would not have been possible without the continued assistance of a great many people. In particular, I would like to thank my advisor, Dr. Paul Turinsky for his expertise and guidance through all aspects of this project. Additionally, I would like to thank the team at the Idaho National Laboratory who developed the `VISION` fuel cycle model and patiently explained to me its inner workings; among this team particular gratitude is owed to Jake Jacobson, Steve Piet, Gretchen Matthern, and David Shropshire. Furthermore, Tyler Schweitzer was instrumental in creating, implementing and testing the initial forecasting techniques in the `VISION` model.

I must also express my continued thanks to John and Betsy Krsul for generously opening up their home to me while I worked at the Idaho National Lab for several summers. Similarly, I must express my grateful appreciation to my parents for their continued support during the many long years of graduate school.

I would like to thank the U.S. Department of Energy and the North Carolina State University Department of Nuclear Engineering which generously provided the funding which enabled this work to take place.

Last, but certainly not least, I wish to thank my many friends who have offered encouragement, motivation, and welcome and necessary distraction and relief from the rigors of work.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Chapter 1

Introduction

The purpose of this work is to develop the capability to determine strategies for the use and application of mathematical optimization employing simulated annealing to the future nuclear power enterprise in which a closed nuclear fuel cycle is used to more fully recover the energy content of uranium while reducing the nuclear waste disposal burden. This is accompanied by an overview of the nuclear fuel cycle and a selection of existing models and simulations thereof, the definitions, objectives and decision-variable space of the optimization, and the optimization algorithm itself. The discussion is capped off by the presentation of results obtained to date and of the future work planned.

1.1 The Nuclear Fuel Cycle

In the United States as of 2012, the commercial nuclear power industry consists of 69 Pressurized Water Reactors (PWRs) and 35 Boiling Water Reactors (BWRs). These reactors are fueled with uranium oxide pellets enriched to 3-5% by weight in the fissile ^{235}U isotope. After 18-24 months of operation, the fissile content of the fuel becomes too depleted to support the chain reaction, so some of the fuel in the reactor is discharged to the spent fuel pool and replaced with fresh fuel. This spent fuel is cooled under water until the spent fuel pool reaches capacity, at which point the oldest, most cooled fuel is removed and placed into dry steel and concrete casks. Under current law, the federal government is obligated to take possession of the spent fuel and safely dispose of it in a central repository. However, the siting and design of this repository has proven to be difficult for technical and political reasons; it remains unclear if, when and where the repository will ever open. Through chemical separation, it is possible to recover and reuse the vast majority of the spent fuel mass. The various actinide and fission product isotopes may then be dealt with in ways more suited to their particular chemical and radiological natures.

This process is quite expensive, and given current uranium and spent fuel disposal costs, there is no financial incentive for this undertaking.

Despite this expense, some countries deem the benefits worth the extra cost. The nation of France, for example, has sparse domestic fossil fuel reserves, so recycling affords the ability to extend nuclear fuel resources, ensuring the security of the energy supply. Of 58 operational PWRs, 20 currently use recovered plutonium in fuel[1] and one uses recovered uranium. In total, recovered spent fuel provides between 17% and 20% of the nuclear power output, reduces uranium usage by 30% and high-level waste material by 97% [2]. Depleted uranium tails and spent recycled fuel is currently stored for later use in Generation-IV reactors.

Nearly 70 years have elapsed since the first man-made nuclear reactor achieved criticality[2]; in that time, many alternative, advanced nuclear fuel cycling schemes have been imagined or investigated. A key feature of many of the so-called Generation-IV reactors is the ability to use a wide range of transuranics as the fuel source in a burn or breed manner. Complementing these advanced reactors are numerous advanced fuel recycling technologies, including aqueous and pyro-processing methods which minimize waste volumes and limit the separation of any material that would be attractive for fabricating nuclear weapons. Several of these schemes have been analyzed as integrated systems by the U.S. Department of Energy as part of the Global Nuclear Energy Partnership (GNEP) program and its predecessors. [3] Of particular interest are the one- and two-tier system deployment evolutions studied in the Dynamic Systems Analysis Report for Nuclear Fuel Recycle [8]. This report examined numerous costs and benefits arising out of a phased transition from the current once-through fuel cycle to a closed fuel cycle, where fast-spectrum reactors are used to recycle spent nuclear fuel. In the one-tier simulations, the transition is made directly to fast reactor recycling, whereas in the two-tier cases, an intermediate, LWR plutonium recycle step is included. Importantly, this study included estimates of the effects on atmospheric carbon dioxide levels and energy cost of various conversion ratios, carbon taxes, and spent fuel cooling times. This study relied heavily on the `VISION` model for its calculations.

Since a closed fuel cycle involves the continuous recycle of material with shortfalls of fuel material in fast reactors made up using reprocessed thermal reactors' fuel, a complex system dynamics occurs to assure proper material balances throughout the fuel cycle. Current fuel cycle codes are dependent upon proper user input to achieve the proper material balances. The work reported here automates this by the employment of mathematical optimization utilizing a fuel cycle code as the predictive engine. Of specific interest is determining the deployment schedule for nuclear fuel cycle facilities, e.g. thermal reactors, fast reactors, thermal reactor fuel fabrication and separation facilities, and fast reactor fuel fabrication and separation facilities,

[1] Four additional units are undergoing licensing for MOX fuel.

[2] Chicago Pile-1 first achieved a chain reaction on Dec. 2, 1942

[3] In 2005, GNEP replaced the Advanced Fuel Cycle Initiative, which in 2003 subsumed both the Generation-IV road-mapping activity and the Advanced Accelerator Applications program.

in such a manner to meet a specific objective with decisions made so as to avoid constraint violations. Chapter 2 of this report provides a detailed discussion of the optimization problem being addressed.

1.2 System Modeling Efforts

Over the years, numerous simulation codes and packages have been developed to investigate various aspects of the nuclear fuel cycle.

The NFCSIM simulation program, developed at Los Alamos National Laboratory [25], utilizes an object-oriented programming methodology to combine external criticality (LACE) and burn-up (ORIGEN) calculations with an internal time-stepping and facility control system to model fuel movements and facility operations through a variety of deployment scenarios. While the inclusion of these proven code packages into the NFCSIM system provides for a great deal of flexibility, power, and reliability, its use in this optimization project necessitates a good deal more computational power than currently available for this project.

The Argonne National Laboratory has developed two system dynamics based fuel cycle simulation codes – DYMOND and DANESS [31, 32]. DYMOND was developed to support the previously mentioned Generation-IV Road-mapping activities. It is specifically targeted to modeling fuel movements associated with the AFCI options within the U.S. DANESS expands upon DYMOND to include economic, environmental and socio-political aspects of multi-region fuel cycle systems utilizing many different reactor types and supporting technologies.

1.3 System Optimization Studies

Several studies may be found in the literature examining the process and outcomes of various nuclear deployment scenario optimization schemes. In [18] a linear programming algorithm is used to study the sizing and placement (domestic or abroad) of proposed fuel recycling facility and its effects upon the security of uranium supply, skilled domestic employment levels, and international balance of commerce. A recent paper by Lee et al. [19] examines many of the same issues addressed in this work, including the impact of uranium pricing and the sensitivity of fast reactor capital costs. However, several differences exist in the optimization approach (linear programming vs. simulated annealing) and the degree of sophistication in the underlying fuel cycle model – e.g., the `VISION` model explicitly models the construction of supporting fuel cycle facilities and tracks material at an isotopic level. Their work does examine several features not yet incorporated into this work, including the use of fast breeder reactors and the incremental ramp-up of fast reactor construction capacity. Another dynamic analysis of LWR/FBR deployment options for the Korean market was performed in [11], with an economic

analysis in [17]; these papers evaluated a series of reference deployment scenarios wherein a growing electric demand is met by a time-changing mix of reactor technologies. As in this work, LWR reactors of $1GW(e)$ nominal size are used until a $600MW(e)$ sodium-cooled fast reactor design reaches technological readiness in the year 2040, whereafter they are deployed in a ratio determined by previous equilibrium calculations [23]. However, these papers only examine a few, fixed deployment scenarios, and while they utilize the same AFCI Cost Basis Report as a source of economic data, the evaluation methodology differs significantly.

1.4 Thesis Organization

This document is organized to go from generality to specificity and back. This chapter examines the state and practice of nuclear fuel cycle modeling as well as the motivation for and general methods of numerical optimization. Chapter 2 delves into the specific algorithms, equations, methods and assumptions used to frame the optimization problem. The chapter begins by setting out the objective functions to be utilized for this work. This is followed by a discussion of the selected set of decision variables over which the optimization will take place. Finally, the parallel, multi-objective optimization algorithm itself is described. In Chapter 3, the results of a series of benchmarking tests are presented; these tests demonstrate the functionality and robustness of the optimization code and establish a set of reference results for later comparisons. In Chapter 4, the optimization code is applied to the examination of several interesting fuel cycle parameters, including fast reactor conversion ratio, the forecast price of uranium, and the estimated capital cost of building new fast reactor capacity. Finally, in Chapter 5, the previously described results are discussed in terms of their potential impact on the process of fuel cycle modeling itself, including various avenues for future exploration and certain pitfalls that may need to be addressed.

1.5 The `VISION` Model

The Verfiable Fuel Cycle Simulaton model was developed through a collaboration between Idaho National Laboratory and numerous academic and laboratory contributors [15]. It is based upon a System Dynamics methodology implemented in the commercial Powersim modeling environment. The Powersim software has several features which make it ideally suited for this work, most notably the ability to read input data from separate Excel spreadsheets and the ability to be operated programatically through an available Application Programming Interface (API). This API facilitates the usage of an optimization wrapper to run the thousands of necessary simulations in an automated manner without user input.

The `VISION` model tracks the construction, operation and decommissioning of the various

nuclear fuel cycle facilities and the movement of fuel therein in response to a user specified energy demand profile and constraints on predicted technological availability. Thus, for each time-step in the simulation, one may ascertain the mass, isotopic composition, and current usage of all nuclear fuel, along with the number, types, and ages of the reactors, fuel fabrication facilities, separation plants, etc. The specific settings used to determine the energy growth rate and the combinations of technologies available at different times are stored in individual *Base Cases* in the attached spreadsheets. There are 64 pre-configured base cases, with 5 additional user-configurable options.

In the `VISION` simulation, each reactor type is assumed to run identical fuel cycling schemes with fixed fresh and spent fuel recipes. For each reactor, only the total fuel mass and fuel type are tracked, not the individual isotopes. Furthermore, the recipes for the fuel to be recycled into the reactors will not match up precisely with the previously discharged spent fuel recipes. Because of this, recycling causes a buildup of certain isotopes, while creating shortages in others. If the fuel recipes were completely fixed and inflexible, these shortages would become a limiting factor, pacing the maximum growth rate of the nuclear enterprise. However, these isotopes are generally minor constituents of the fissile fuel, and growth will be paced by the major fissile isotopes (i.e. plutonium isotopes and the transuranic elements). The inclusion of variable fuel recipes into the `VISION` model would require at the least a way to model and compensate for the relative reactivity contributions of each isotope. Work to test this concept has been performed at INL [36], however this feature is not available in the `VISION` version used for this optimization work. Instead, the mass of recycled fuel that can be produced from a given batch of spent fuel is determined by the so-called *Pu Control Switch*, whereby the total mass of a specified set of fissile isotopes (minus minor recycling process loss) is conserved in the transition from spent fuel into fresh fuel. If there is a mismatch in isotopes between recipes, it either accumulates in the system, or is introduced as needed (without a physical source).

As the overall fuel mass available for recycling in the `VISION` model is determined by the active isotope specified in the Pu Control Switch, it alone will be used in the optimization calculations for determining the combination of Fast-burner and Thermal reactors.

1.5.1 Fuel Recipes and Limiting Isotopes

The process of designing and qualifying a fuel for use in a nuclear reactor is one of the most time consuming and difficult aspects of creating a nuclear power system. The fuel is subject to higher temperatures, greater heat flows and more intense radiation fields than any other part of the fuel cycle. Additionally, it provides the first line of containment for radioactive isotopes, helping to keep them in place both in normal and accident scenarios. The tight constraints placed on fuel performance lead to similarly tight tolerances on fuel chemistry; these, in turn,

lead to very demanding requirements for the reprocessing of spent fuel. The ability to control the fuel going into the reactor is thereby limited by the ability to control the chemistry in the fabrication and separations processes.

Knowing the fuel composition going into a reactor is only half of the simulation challenge. When nuclear fuel is burned in a reactor, an enormous array of nuclear reactions occur – some isotopes will fission into a set of lighter isotopes, while other isotopes will successively absorb neutrons to create heavier and heavier isotopes. All the while, these isotopes may undergo alpha, beta or other decay mechanisms, creating yet more isotopes. While the individual nuclear reactions may be modeled and understood at high levels of accuracy, it would be exceptionally difficult to attain such fidelity in a fuel cycle-wide model. Furthermore, it would be computationally prohibitive to repeat these calculations again and again for each scenario in a fuel cycle optimization calculation. In order to remain usable, the `VISION` model simplifies the calculations of fuel composition down to a set of table interpolations. For light-water reactor fuel, detailed calculations were performed for fresh and spent fuel designs supporting 33 and 100 GWD/MT-IHM[4] discharge burnups [15]. For fast burner reactors, similar calculations were performed for fuel with conversion ratios of 0.0 and 1.0, where conversion ratio is defined here as the ratio of production of transuranic material to total destruction of same. Fresh and spent fuel compositions are thus interpolated from the externally calculated recipes to produce the specified discharge burnups, which in turn are determined considering core average cycle burnup and core fresh fuel loading fraction.

When spent fuel is reprocessed, some chemical isotopes are retained, while some must be discarded. Because isotopic separation is extremely expensive even when dealing with nonradioactive fresh uranium fuel, it is considered unfeasible to attempt to isotopically separate any of the components of spent fuel. Instead, a variety of chemical separation techniques are utilized, none of which can affect the ratio of isotopes of a given element. This raises the possibility of encountering a mismatch between the relative amounts of each isotope in a spent fuel recipe and those in the feed fuel recipe that is to be fabricated. Because `VISION` lacks the capability to recompute fuel recipes on the fly, certain assumptions must be made to allow the fuel cycle to continue functioning. In order to understand these assumptions, one must first examine the way in which the `VISION` simulation tracks and monitors fuel usage and isotopics.

One major design criteria for the `VISION` model was that it be simple enough to run on a standard desktop PC in less than five minutes [15]; this ruled out isotope-level fuel tracking for most parts of the model. Instead, aggregate fuel masses of defined composition are tracked. This method works well for situations where fuel residence time is significantly shorter than decay half-lives and where fuel recipes are fixed. In areas where decay is a concern, such as wet and dry storage, isotopic decay is accommodated through the use of coupled chains of

[4]Gigawatt-days of thermal power per metric ton initial heavy metal content

recipes. For example, prior to irradiation, fresh fuel is measured by total mass per reactor type and recycling pass. After irradiation, when it is moved into wet storage, it is further divided into an `Age Wet Storage` and `Isotope` dimension. Each age recipe corresponds to the fuel composition after a given time; after the required storage time has elapsed, the fuel is moved into the next step in the fuel cycle. The reprocessing chain draws its fuel supply from dry storage, which has dimensions of `Reactors`, `Pass`, `Isotopes` and `Age Dry Storage`, and feeds it back into Fuel Fabrication, which only has dimensions `Reactors` and `Pass`. There will be at least one isotope that limits the total mass of fuel that may be extracted for reprocessing without altering the fixed fuel composition. However, as mentioned above, isotopic separation of individual chemical elements is impractical. Furthermore, in some cases the separation of certain chemical elements (e.g. plutonium), although feasible, is considered undesirable from a proliferation standpoint. For these reasons, when `VISION` calculates the mass of a given fuel-type that it can fabricate by reprocessing a given amount of spent fuel, it uses not a single limiting isotope, but a summation over a family of controlling isotopes. As an example consider a simple fuel recipe requiring 1-kg of Pu-239, 0.6-kg of Pu-240, and 0.5-kg of Cm-242. Further, assume that the spent fuel stockpile consists of 3-kg of each isotope. Strictly speaking, using this stockpile and fuel recipe, one may fabricate 6.3-kg of fuel, and will have 1.2-kg of Pu-240 and 1.5-kg of Cm-242 leftover. If, however, total Pu mass is said to be controlling, then 7.875-kg of fuel may be fabricated, leaving 1.125-kg Cm-242 and 0.75-kg Pu-240 leftover, while incurring an artificial deficit of 0.75-kg Pu-239. Thus, total mass is conserved, but the individual isotope balances are not.

The Socially Conscious Algorithm (described below) requires an estimation of both the mass of spent fuel both consumed and produced by each reactor type. In order to best match the `VISION` model, only the controlling isotopes are considered in these calculations. The fuel mass definitions above and the derivations below both reflect only the total mass of these controlling (or active-) isotopes.

1.5.2 System Dynamics Modeling

There are many approaches to modeling complex, interacting systems, be they chemical, mechanical, economic, political or social. System Dynamics, initially developed in 1956 by Jay Forrester at the MIT Sloan School of Management is a general approach to modeling large interacting systems; it was conceived as a method of applying engineering principles to the types of management problems where the complexity of feedback systems cause intuition and experience to fail to generate acceptable results [10].

An underlying premise of System Dynamics modeling is that though the individual actors in the simulation may exert relatively little influence on the overall behavior of the system, it is through their interactions that the observed, unexpected behaviors arise. The individual

interactions tend to occur rather predictably and according to simple rules, while the outcomes can vary wildly and counterintuitively. Thus, the first step in the development of a system dynamics model is to examine the organizational structure of the system under consideration, determine who or what all of the action points in the simulation are, then examine them to determine how they go about their decision-making process. Rather than rely on compiled data, historical trends or theories about what ought to be, the model maker focuses on what factors control decision making in the moment. The reasoning is that it better reflects the true nature of the decision making process, thereby recreating the actual behavior more closely.

The models of the individual actors in a system are interconnected through influence diagrams, which provide visual indication of the mathematical dependencies of each variable on each other. The system of equations is integrated forward in time using explicit numerical methods to generate an estimate of the overall system behavior. By design, these models tend to be simple, so that they run quickly, allowing the policies to be changed and adjusted in order to explore management alternatives.

1.6 Simulated Annealing Optimization

The simulated annealing algorithm has its roots in the stochastic methods used to estimate material properties for nuclear weapons design [21]. In this application, a series of random perturbations were used to generate a spatial distribution of atoms with the desired thermodynamic conditions. The connection between configuration, energy, the annealing process, and general optimization was later made by [16]. Since then, simulated annealing has been applied in various forms to a wide range of computational optimization problems.

In its most basic form, the simulated annealing (SA) algorithm seeks to find the minimum value of a given real-valued objective function, $f(X_i)$, of a set of configuration values $\{X_i\}$ through a series of random perturbations to the configuration values. What distinguishes SA from a truly random search or an incremental-improvement algorithm is the way in which certain perturbations are accepted while others are discarded. On sampling a new configuration, $\tilde{X}_i$, the new objective function value is compared against the prior value. If $f(\tilde{X}_i)$ is less than $f(X_i)$, then $\tilde{X}_i$ is accepted as the new base configuration, and the process repeats. If, however, $f(\tilde{X}_i)$ is greater than $f(X_i)$, then it is conditionally accepted with a probability given by the *Metropolis Criterion*:

$$p = \exp\left(-\frac{f(\tilde{X}_i) - f(X_i)}{T}\right) \tag{1.1}$$

where the parameter T, (the so-called *annealing temperature*), will be discussed later. By accepting new configurations according to these rules, the distribution of objective function values of the set of accepted solutions will approach the Boltzmann distribution for the given

annealing temperature. As can be seen in Figure 1.1, lower values of T correspond to a higher probability that the objective function value (energy) will be at a minimum. The high initial temperature allows the algorithm to explore the domain of possible solutions, ensuring that there is a finite probability that any given configuration is sampled. As the annealing temperature drops, the probability distribution becomes increasingly focused around those states which minimize the objective function values. The key distinction between the SA algorithm and the similar Iterative-Improvement algorithm is that in Iterative-Improvement only solutions which reduce the objective function value are accepted, whereas SA allows for the acceptance of inferior solutions. This prevents the SA algorithm from becoming trapped in locally optimum solutions and missing global optimized configurations.

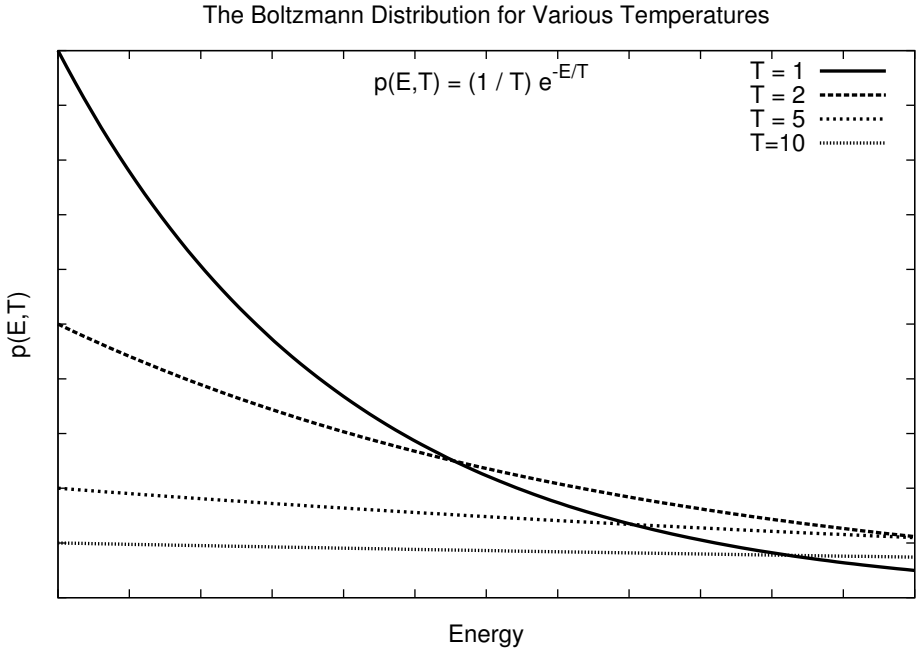

Figure 1.1: Boltzmann's Distribution

Throughout the optimization process, the annealing temperature, T, is adjusted according to the *cooling schedule*. Throughout the various cooling schedules proposed in the literature, several common features emerge. The first is that the initial temperature is high enough such that all, or nearly all of the proposed configurations will be accepted. This may entail evaluating a set of randomly sampled solutions, and setting the initial temperature value to be a multiple of the sample standard deviation of the observed objective function values. Alternatively,

one may select a threshold acceptance ratio for samples, and repeatedly double the initial temperature until this threshold is exceeded. Regardless of the method chosen, there is a trade-off between computational robustness and efficiency which must be addressed through testing on the optimization problem at hand. After the initial temperature has been selected, it is gradually reduced as further samples are taken. Each time when the temperature is reduced, the equilibrium sample distribution shifts, so additional configurations must be sampled until the objective function value distribution approaches the new equilibrium. The various cooling schedules differ in the way in which the initial and subsequent temperatures are calculated and the intervals at which the changes occur. Simple methods will have the temperature reduced by fixed fractions, with updates occurring after a predetermined number of samples or acceptances. More complicated methods will take into account such factors as the observed objective function distribution, the average accepted objective function value increase, or various other statistical tests to determine when and by how much to reduce the annealing temperature.

1.6.1 Parallelism

While the simulated annealing algorithm has the advantage of being quite simple and robust, its great disadvantage is the slow, asymptotic convergence to the set of optimum solutions. One obvious way to reduce the computational turnaround time for large problems is to parallelize the search process.

When metals are annealed, there is a potential that many atoms within the crystalline matrix will be in motion at any given time. Similarly, with simulated annealing, it possible to split the optimization search across several independent parallel processes. The parallelization may be accomplished at one (or more) of several levels in the computational process. In increasing order of computational granularity, one could 1) run several completely independent optimizations and correlate the final set of results, 2) utilize independent parallel optimization searches between annealing temperature updates, with limited coordination between processes at each update, 3) evaluate multiple proposed moves at one time, using some rules to select the surviving solution at each step, or 4) split the evaluation of the objective function itself across multiple processes. Of these options, the first offers only limited performance benefits while requiring potentially significant post-processing effort, the third is primarily effective at low annealing temperatures (but no more effective than the second option), and the fourth is only of very limited applicability for our application. Thus, the second option is the preferred approach for this work.

The parallel simulated annealing algorithm used for this work is quite similar to the synchronous Multiple-Markov-Chain Parallel Simulated Annealing (MMC-PSA) algorithm described in [20], and also used in previous work [12]. In this algorithm the traditional simulated annealing search is divided into P concurrent search segments, each following an independent path. When

the annealing temperature update criteria is satisfied (more than 87% of processes have sampled at least L_{chain} configurations or accepted at least L_{tran} configurations), all search processes come to a stop and the relevant optimization statistics are aggregated on the head process. These aggregated statistics are then checked against the termination criteria. If the termination criteria is not met, then the cooling parameters are updated and each search process begins again from a new starting point. For single-objective optimizations, the best archived solution is used as the starting point for all processes, while in multi-objective optimizations, each process starts with a different configuration, randomly selected from the non-dominated archive. This process is shown below in the flowchart below in Figure 1.2.

The parallel efficiency of this algorithm tends to be determined by two factors. First, because the evaluation time for each configuration will vary from process to process, there will be some unavoidable idle time in all but the last process to respond to an update command. In the limiting case, this idle time will be equal to the longest individual simulation time, thus the fraction of time spent waiting in any cooling step is proportional to $1/L_{tran}$. The second factor affecting parallel efficiency is the computation and communication time associated with the temperature update process. This factor is not of great consequence in this work, as it generally requires less than 30 seconds, which is negligible compared to the 5 to 8 hours of computation between updates.

Like the efficiency, the scalability of the algorithm is somewhat limited, though those limits are beyond the currently available computational resources. In order for the simulated annealing optimization algorithm itself to function properly, two conditions must be met at each temperature update. First, a sufficiently large number of samples must have been evaluated to ensure that the distribution of objective function values approaches thermal equilibrium. Secondly, the combined actions of the various accepted configuration perturbations must be sufficient to ensure that all feasible configurations within the search space have a reasonable probability of being sampled. Though splitting the annealing search amongst a greater number of processes can reduce the number of samples required on each individual segment, a certain (problem dependent) minimum segment length is still required. This effect can manifest itself through either poor convergence or poor parallel efficiency if the optimization is attempted on too many processes. It should also be noted that the maximum parallel extent will also be a function of the cooling temperature. During the later stages of the optimization very few of the candidate solutions will be accepted, and any additional search of the sampling space will be useful (up to the point of redundant search coverage).

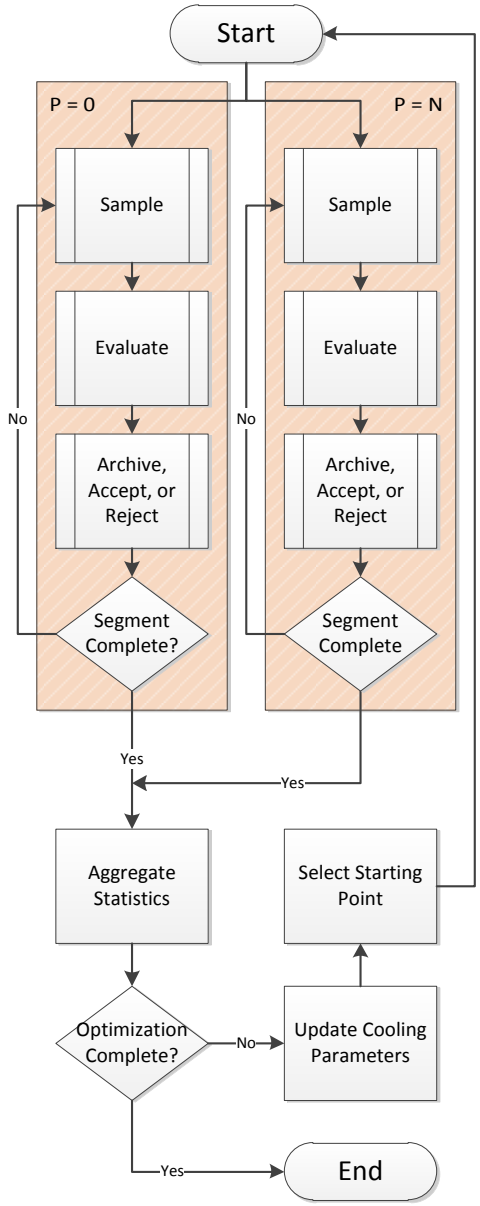

Figure 1.2: Parallel Simulated Annealing Algorithm Flowchart

1.6.2 Augmented Single Objective Optimization

In many optimization scenarios, a single objective function may be insufficient to fully represent various competing factors that are important in the quality of the final result. In these situations it may be possible to incorporate these subordinate factors into the objective function formulation through the use of penalty constraints. These penalty constraints, Θ_j, are incorporated into the objective function as shown in the equation below through the use of penalty multipliers, λ_j.

$$\bar{f}(X_i) = f(X_i) + \sum_j \lambda_j \Theta_j(X_i) \tag{1.2}$$

The penalty multipliers start with a low initial value, and are increased at each annealing temperature update with the goal of driving the average total constraint violation value (given by the product $\lambda_j \Theta_j(X_i)$) to zero before end of the cooling cycle so that none of the final configurations have violations. The precise algorithm used for incrementing the penalty multipliers is problem specific and considered part of the annealing schedule, which is discussed in the next chapter.

1.6.3 Multi-Objective Optimization

While single-objective optimization is quite useful for automatically finding optimum solutions when the objectives and search space are well known, this is often not what is desired. Oftentimes in engineering situations, there are multiple competing goals, and it is desired to know what exactly the trade-off surface between these goals looks like, both in terms of configurations and the objective values they give. For this, the multi-objective simulated annealing (MOSA) optimization schemes put forth by [30] and [9] are useful. In this method, rather than archiving the best N solutions (as ranked by a single augmented objective function), an archive of *non-dominated* solutions (defined below) is created. A solution, X, is said to dominate another solution, Y, if $f_i(X) < f_i(Y)$ for every objective function i. If neither X dominates Y, nor Y dominates X, then the solutions are said to be non-dominated. This is represented pictorially for two objectives in Figure 1.3 below.

Algorithmically, MOSA proceeds in a manner similar to SOSA, with two important exceptions. First, rather than having a single objective function and a single annealing temperature, there are multiple independent objective functions, each with their own annealing temperature. Second, any candidate solution that is accepted into the non-dominated archive is automatically accepted by the SA search algorithm. When a candidate solution is tested for acceptance into the non-dominated archive, it is compared against each of the archived solutions. Any previously archived solution dominated by the candidate solution is removed from the archive. Thus, as the optimization proceeds, the set of non-dominated solutions approaches the true *pareto-optimal* set of solutions that defines the objective trade-off surface. If K is the number of optimization

objectives, then the trade-off surface will have $K-1$ dimensions – therefore the total number of solutions which must be sampled in order to clearly describe the trade-off surface is a very strong function of K.

If the sampled candidate solution is not added to the non-dominated archive, then it is accepted as the current working solution with a probability given by

$$p = \prod_{i=1}^{K} e^{-\left(f_i^{n+1} - f_i^n\right)/T_i}$$

This equation is manipulated into the following equivalent form which has the benefit of being more robust when calculated with floating-point arithmetic.

$$p = \exp\left(-\sum_i \frac{\left(f_i^{n+1} - f_i^n\right)}{T_i}\right) \tag{1.3}$$

It should be noted that for certain combinations of objective values, values of $p > 0$ may occur. In these situations, the proposed solution will, of course, be accepted.

As with single-objective optimization, the annealing temperature and current working solutions must be periodically updated. As before, an adaptive Huang's temperature decrement factor is used, but with the addition of a factor to retard the cooling in proportion to the dimensionality of the search space.

$$\begin{aligned} T_i^{(k+1)} &= \max\left(\bar{\alpha}_{h,i}, \bar{\alpha}_{min}\right) \times T_i^{(k)} \\ \bar{\alpha}_{h,i} &= \alpha_{h,i}^{1/K} \\ &= \exp\left(\frac{-\lambda_h T_i^{(k)}}{s_i^{(k)} K}\right) \\ \bar{\alpha}_{min} &= \alpha_{min}^{1/K} \end{aligned} \tag{1.4}$$

where K is the number of active objectives, k is the cooling step number, i is an individual objective, and $s_i^{(k)}$ is the sample standard deviation of objective function value i across all processes over cooling step k. Once a new set of temperature values have been computed, a new starting position for each parallel segment is selected at random from the non-dominated archive.

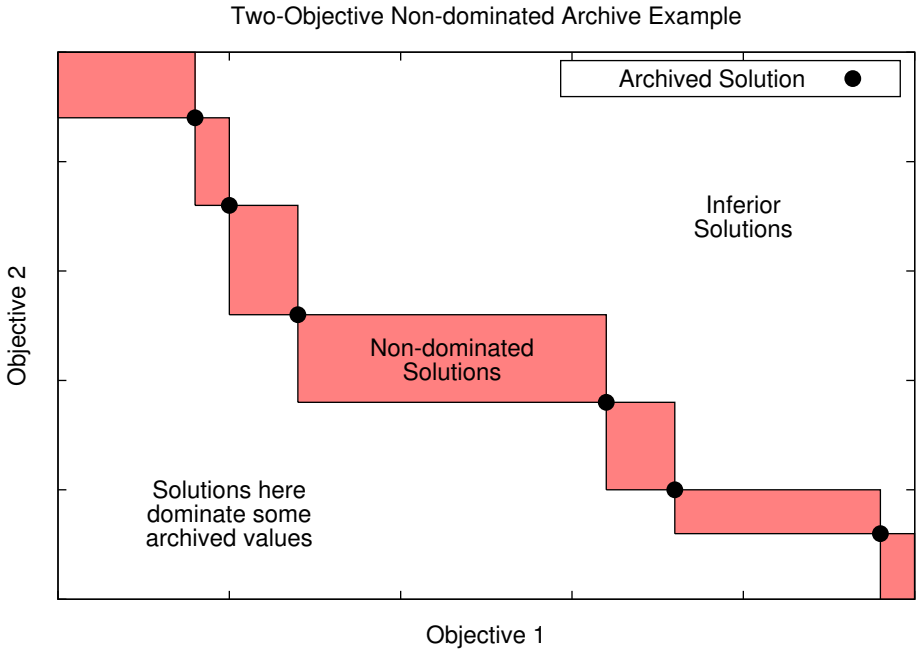

Figure 1.3: A two-objective example of a non-dominated archive

Chapter 2

Methodology

The simulated annealing optimization methodology can be divided into several sections. First is the definition of the decision variables (these establish the available parameter search space); this then ties into the selection of the objective functions, which the optimization algorithm will seek to minimize. For time dependent problems, such as are being addressed here, one needs to specify a time period over which the optimization is applied, referred to as the planning horizon. The difficulty often encountered is that decisions made within the planning horizon influence what occurs beyond the planning horizon. A method must be established for cleanly separating the costs and benefits of decisions made within the planning horizon from those assumed to be made after the planning horizon–this forms the basis for our maximal-FBR decision making methodology. Constraints need to be identified which will restrict the space of acceptable decision variables values. Finally, the optimization problem having been fully specified, the optimization algorithm itself may be discussed.

2.1 Objective Functions

The set of objective functions used in the optimization algorithm are a mathematical representation of what the system designers deem to be some of the important qualities of the fuel cycle strategy. The metrics chosen for this work were selected to relate to key factors that will ultimately determine the acceptability of any proposed nuclear fuel cycle, e.g. waste disposal, proliferation resistance, uranium utilization and, of course, cost.

2.1.1 Long-term Heat

Spent nuclear fuel and high-level reprocessing wastes present a significant, long-lasting disposal challenge. Unlike chemical hazards, radioactive waste cannot be mitigated through incineration

or dilution. The only methods devised so far involve either deep burial in solid, stable geologic formations, or further nuclear transmutation into short-lived isotopes. As transmutation would be exceptionally expensive, burial is currently the preferred option (it can always be disinterred and transmuted at a later time, should the need arise). It is hoped that once buried, the surrounding mass of rock will keep the radioisotopes safely shielded and contained for millenia. As in a reactor, a high-level waste repository relies on many protective layers, each of which has its own set of design limitations. Until recently, Yucca Mountain was the designated repository site for the United States. The design called for waste to be emplaced in a series of parallel tunnels, called drifts. A key limiting factor on the amount of waste that could be contained within the repository site was that the maximum temperature of the rock walls at the center point between the drifts could not exceed the boiling point of water [34,35]. This, in turn, limits the total amount of decay heat that the emplaced wastes may emit. Thus, the amount of heat released by the wastes present at the end of the simulation lifetime are used as a proxy for the overall waste disposal burden of the cycle.

The Long-term Heat Objective Function value is given by the equation:

$$f_{LTH} = \frac{H}{E}, \tag{2.1}$$

where H is the total heat load going into the geologic repository, which itself is given by the equation

$$H = \sum_{i=1}^{6} H_i, \tag{2.2}$$

where the individual H_i terms are as listed in Table 2.1. The energy production term, E, is the deployed reactor capacity[1], less any reactors that are not operating due to fuel shortages.

$$E = \sum_{i \in rx} \left[\left(\Delta t \sum_{t=t_{start}}^{t_{end}} C_{t,i} \right) - C_i^{idle} \right]. \tag{2.3}$$

The `VISION` variables corresponding to the terms above are given in Table 2.1. It should be noted that the heat produced by accumulated separated nuclear material remaining at the end of simulation is not included in this calculation, as this is classified as a resource instead of a waste stream.

The objective function value itself is calculated by combining the long-term heat terms for all waste streams at the end of the simulation (in Joules) and normalizing this value by the net electric energy produced during the simulation (also in Joules) and multiplying by a factor of

[1] The term "deployed reactor capacity", as used here and throughout this work, refers to the amount of energy actually produced during a specified operational period, incorporating any capacity or availability factors. This differs from the peak or nameplate capacity/rating that each plant may have.

Table 2.1: VISION Variables in the Long-term Heat Objective Function

Term	VISION Variable	Sample Time
H_1	LTH non-repository waste - total[100to1500]	2200
H_2	total LTH in repository[100to1500]	2200
H_3	LTH SF interim storage - total[100to1500]	2200
H_4	LTH SF mid aged storage - total[100to1500]	2200
H_5	LTH SF aged storage - total[100to1500]	2200
H_6	LTH legacy SF - total[100to1500]	2200
C^{idle}_{FBR}	Net Idle Reactor Capacity[FBR]	2200
C^{idle}_{LWR}	Net Idle Reactor Capacity[LWR]	2200
C^{idle}_{LWRmf}	Net Idle Reactor Capacity[LWRmf]	2200
$C_{t,FBR}$	deployed reactor capacity[FBR]	All
$C_{t,LWR}$	deployed reactor capacity[LWR]	All
$C_{t,LWRmf}$	deployed reactor capacity[LWRmf]	All

1,000,000 (to give an aesthetically pleasing mantissa). Typical values are in the range of 150 to 4000 part-per-million.

2.1.2 Nonproliferation

There is a great deal of overlap between those materials, equipment and skills required to operate a nuclear energy system (NES) and to fabricate a nuclear explosive device (NED). Although nuclear energy is capable of producing great benefits in terms of low-emissions electricity production, great care must be taken in the planning and operations stages to ensure that no material is surreptitiously diverted and applied toward dastardly ends. The science of analyzing the weaponization potential of a NES is, of course, closely tied to the science of weapons development. Though the more detailed risk factors and analyses are classified, several metrics are available in the open literature.

Weaponization risks of civilian nuclear energy systems are broken down into several categories based upon the level of technical sophistication and ability to act overtly of the would-be proliferator. If the proliferator is a nation already in possession of a weapon, then further overt diversion of material for weapons use is considered a matter of international arms-control negotiation. If the proliferator is instead a subnational group within a nation-state, then the covert diversion of material is considered a matter of physical protection and material controls and accountability. The term *proliferation resistance* generally refers to those characteristics of a NES that prevent the covert diversion of materials by a non-weapons state. In this work,

however, the term is more broadly applied to represent both resistance to state-sponsored diversion (which is generally not an issue in the US) and resistance to theft by individuals or subnational groups.

The design of a nuclear facility to prevent and detect the unauthorized diversion of nuclear materials goes hand-in-hand with the need to prevent and detect uncontrolled releases due to intentional or inadvertent operator or equipment failure. As safety is a design requirement for all facilities, the risks of material diversion are somewhat self-limiting (they become part of the design basis, which translate into increased construction and operations costs). Nonetheless, as a first assumption, a rough proliferation risk metric may be computed by considering the nuclear materials that must be handled and manipulated through the course of the nuclear fuel cycle.

The `VISION` model tracks nuclear materials through time as they move and transform through the various processes that make up the nuclear fuel cycle; let us refer to the fuel in each individual stage throughout the simulation as a *stream* of fuel. The physical, chemical and radiological properties of each stream largely determine their utility to a potential proliferator. However, the design of the facility for handling each stream plays a large role in either inviting or deterring the motivated proliferator; for example, a permanent waste repository will necessarily be designed to inhibit all physical access, while a fuel reprocessing facility will require some degree of human interaction to enable operations.

The nonproliferation objective function value is computed by integrating over the simulation time the estimated contributions to the proliferation risk of each stream, then summing across all of the streams. Two things must therefore be specified: 1) which of the streams in the `VISION` model are vulnerable or important, and 2) how to assess the relative vulnerability of the stream given its known properties. The first factor is determined by considering the physical accessibility of the individual streams. Facilities such as short-term storage and reprocessing facilities inherently require a greater level of interaction than long-term disposal facilities. The fuel actively powering an LWR or FBR reactor is likewise quite inaccessible, both due to the physical encumbrances of the reactor vessel assembly and the radiation fields generated during use. Currently the value for this factor is quite simplistic; long-term storage, reactor, and geologic repository streams are given a value of zero, while all others have a value of one. Those streams marked in Table 2.2 are considered to be accessible to the would-be proliferator.

Two approaches are used to assess the relative vulnerability of each stream; the weighted-weapons-usable method is based on a set of our own assumptions, and the second is based on a well known Figure-Of-Merit formula developed by Bathke [3]. It should be noted that though these metrics may be based on other established criteria, many assumptions and simplifications have been made, and the resulting values cannot be usefully compared to values generated outside of this work. The primary utility of these metrics is in making quantitative comparisons between competing configurations within a systems optimization framework.

Table 2.2: Dimensions of weapons usable streams within VISION

	Size of Stream Partition Dimensions†			
Location	Reactor	Pass	Age	Element
Fuel Fabrication	3	6		
Ready Fuel	3	6		
Wet Storage	3	6		
Dry Storage	3	6	5	
Separations	3*	6		6‡
Separated Material	2	6		6‡
Recycled Uranium		6		
Legacy Fuel		1	6	

† Not all streams will have fuel in each deployment scenario.
* Uses closely related FUELTYPE enumeration.
‡ Conservatively assumes complete separation of all weapons-usable elements.

Weighted Weapons-Usable

The Weighted Weapons-Usable approach combines three relevant stream properties to estimate a total attractiveness factor for the material contained therein.

Total Stream Concentration It is expected that regardless of who the potential proliferator at a given facility might be, they will only be able to divert material in the form in which it exists in the facility. Thus any secondary processes required to render the material into a weapons-usable form must be performed elsewhere, after the material has been successfully diverted. Therefore, the attractiveness of a stream to a proliferator is directly proportional to the concentration of the weapons usable isotopes within the stream itself. This factor is denoted mathematically as $\chi_{j,i}^{str}$, where i and j refer to the isotope and stream, respectively.

Isotopic Enrichment Assuming a proliferator has been able to divert a quantity of material from a given stream and chemically purify it, they are then left with a mixture with a fixed isotopic vector, which may be more or less favorable depending on the isotopic balance therein. (Although the chemical separation of each element is a relatively simple process, the isotopic enrichment of the desirable isotopes is exceedingly difficult.) Therefore, a factor is added to reflect the relative abundance of each isotope of a given element within each process stream. This factor is denoted as $\chi_{j,i}^{elem}$ where, as before, i and j denote individual isotopes and streams.

Critical Mass Even in an isotopically pure state, not all fissionable isotopes are equally usable; for example, the bare-sphere critical mass (BSCM) of Uranium-235 is 47.9-kg, compared to only 16.4-kg for Uranium-233 (at room temperature and pressure [6]). Because ^{239}Pu has a long history of weapons use, its BSCM is used as the normalizing factor for other isotopes in the `VISION` model. Therefore, the weapons-usability factor for each isotope is calculated by normalizing the inverse of its bare-sphere critical mass by the similar inverse value for ^{239}Pu. These factors are denoted w_i.

Mathematically, the Weighted Weapons Usable (WWU) objective function value is expressed as follows:

$$f_{WWU} = \int_{t_{start}}^{t_{end}} \Theta_{WWU}(t)\mathrm{d}t \tag{2.4}$$

$$\Theta_{WWU}(t) = \sum_{j=str} \sum_{i=iso} \chi_{j,i}^{str}(t)\chi_{j,i}^{elem}(t)w_i m_{j,i}(t) \tag{2.5}$$

$$\chi_{j,i}^{str}(t) = \frac{m_{j,i}(t)}{\sum_{i\in str} m_{j,i}(t)} \tag{2.6}$$

$$\chi_{j,i}^{elem} = \frac{m_{j,i}(t)}{\sum_{k\in E} m_{j,k}(t)} \tag{2.7}$$

Where:

$m_{j,i}(t) \equiv$ The mass of isotope i in stream j at time t.

$E_i \equiv$ The set of all tracked isotopes of the same element as isotope i

The units of the objective function correspond to the default mass and time units in the `VISION` code, namely Kiloton-yr/GW(e)-yr, however it is sometimes reported in tonne/GW(e).

The Bathke Approach

In collaboration with weapons experts, Bathke has developed a Figure of Merit equation that utilizes commonly available data such as the bare-sphere critical-mass and heat generation rate to rate the attractiveness of a material to proliferators of various degrees of sophistication. The equation below was developed with "sub-national groups," "less advanced proliferant nations,"

and "technically advanced proliferant states" [3] in mind.

$$FOM = 1 - \log_{10}\left(\frac{M}{800} + \frac{Mh}{4500} + \frac{M}{50}\left[\frac{D}{500}\right]^{\frac{1}{\log_{10} 2}}\right) \tag{2.8}$$

Where:

$M \equiv$ Bare-sphere critical mass [kg]

$h \equiv$ Heat generation rate [W/kg]

$D \equiv$ Dose rate of 0.2M at 1 meter [rad/h]

Materials possessing a FOM value greater than 2 are considered to be highly attractive for weapons use, while values between 1 and 2 are moderately attractive. Values below 1 are of only low attractiveness. FOM values for several possible weapons isotopes are plotted below in Figure 2.1.

For this work, we consider only the effects of the critical mass and the heat generation rate, taking no credit for the radiation dose. In the equations that follow, m_i, h_i, and w_i denote the mass in the stream, the specific heat output and the weapons-usable factor, respectively, of the constituent isotope i within a stream. The term M_i is the bare-sphere critical mass of isotope i. The heat generation term for the stream mixture is calculated by taking a mass-fraction weighted average of the constituent isotope heat generation rates of a stream (Constituent values sourced from Table 1 in [6], reproduced below in Table 2.3).

$$h(t) = \frac{\sum_i h_i m_i(t)}{\sum_i m_i(t)} \tag{2.9}$$

Similarly, the bare-sphere critical mass for the stream composition is estimated using the formula below:

$$M(t) = \left(\frac{\sum_i m_i(t)}{\sum_i \frac{1}{M_i} m_i(t)}\right) \tag{2.10}$$

This equation is then reformulated to utilize the weapons-usable factors mentioned in the previous section:

$$M(t) = M_{\text{Pu-239}} \frac{\sum_i m_i(t)}{\sum_i w_i m_i(t)} \tag{2.11}$$

If the calculated Figure-of-Merit is greater than zero, then it is multiplied by the total mass in the stream, integrated over the simulation time, and summed with the other streams to yield

Table 2.3: Weapons Usable Factors from [6]

Isotope	h_i[W/kg]	M_i[kg]	w_i
^{231}Pa	1.3	162.	0.06
^{233}U	0.281	16.4	0.62
^{235}U	0.000006	47.9	0.21
^{237}Np	0.021	59.	0.17
^{238}Pu	560.	10.	1.02
^{239}Pu	2.0	10.2	1.00
^{240}Pu	7.0	36.8	0.28
^{241}Pu	6.4	12.9	0.79
^{242}Pu	0.12	89.	0.11
^{241}Am	115.	57.	0.18
^{242m}Am*	–	–	0.34
^{243}Am	6.4	155.	0.07
^{244}Cm	2800.	28.	0.36
^{245}Cm	5.7	13.	0.78
^{246}Cm	10.	84.	0.12
^{247}Bk†	36.	10.	1.02
^{251}Cf‡	56.	9.	1.02

* Value from `VISION`, isotope not listed in source table.
† Isotope not tracked in `VISION`.
‡ Value not initially included in `VISION` input tables.

the overall proliferation resistance score.

$$f_{Ba} = \int_{t_{start}}^{t_{end}} \Theta_{Ba}(t)\mathrm{d}t \tag{2.12}$$

$$\Theta_{Ba}(t) = \sum_{j=str} \mathcal{F}_j(t) m_j(t) \tag{2.13}$$

$$\mathcal{F}_j(t) = \max\left(FOM_j(t), 0\right) \tag{2.14}$$

A comparison of the values of the two nonproliferation metrics may be found in Section 3.2.3.

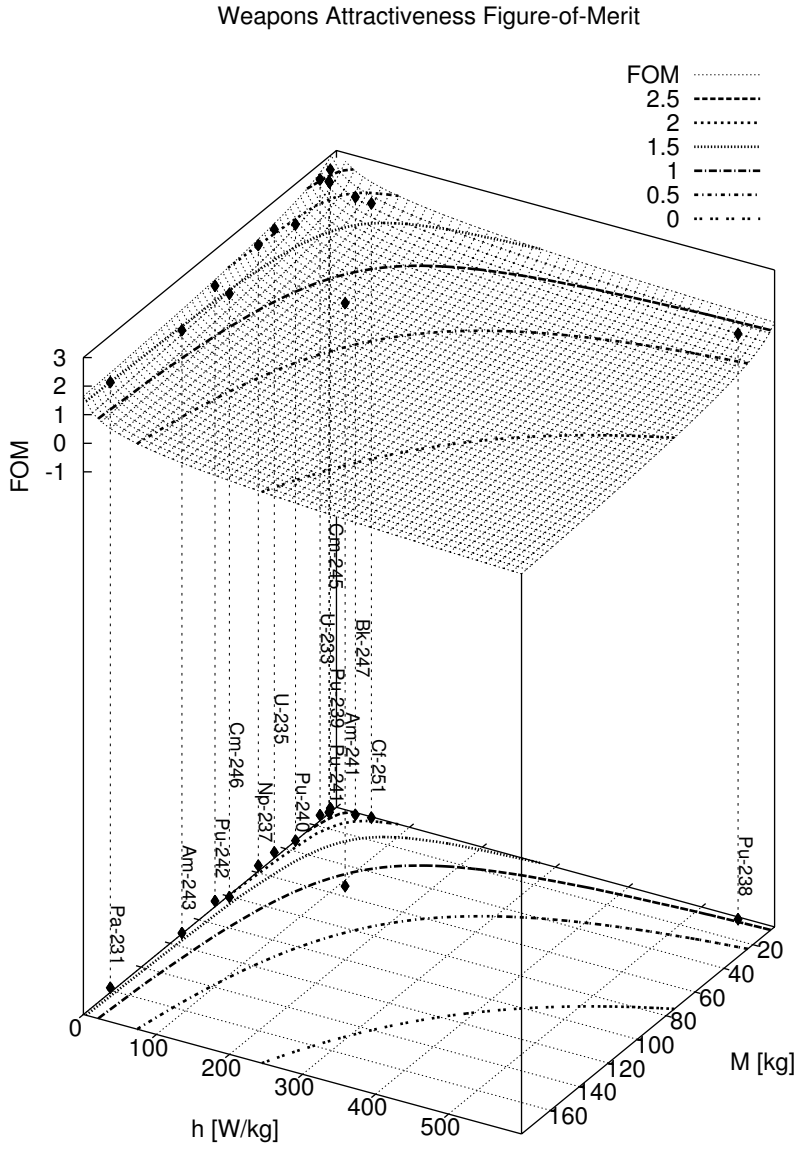

Figure 2.1: FOM values for weapons-usable isotopes (^{244}Cm off scale).

2.1.3 Uranium Utilization

Though the energy density of nuclear fuel is exceptionally high, the concentrations of uranium in ore are quite low (less than 0.3%) and the extraction methods can be quite disruptive [5]. Thus, where the minimization of waste heat serves as a metric for the environmental impact of the back end of the nuclear fuel cycle, the Uranium Utilization objective function serves as a metric for the environmental impact of the front-end. The goal, of course, is to use the least mass of ore to produce the greatest amount of electrical energy. Mathematically, it is expressed simply by the equation below.

$$f_U = \frac{U_{initial} - U_{final}}{E} \tag{2.15}$$

where E is the same net electrical energy production term employed previously. The variables $U_{initial}$ and U_{final} are the initial and final masses of unmined uranium ore as tracked within the `VISION` model (the variable names are listed in Table 2.4). This measure makes sense when deploying fast reactors since all isotopes of uranium are capable of directly fissioning or through neutron capture creating fissile isotopes.

Table 2.4: `VISION` Variables in the Uranium Utilization Objective Function

Term	**`VISION` Variable**	**Sample Time**
$U_{initial}$	`initial u resources`	2200
U_{final}	`U Ore`	2200
E	See Equation (2.3)	

Using the `VISION` default units, this objective has the units of Kt-ore/GW(e)-yr, however it is generally converted in this work to the more accessible units of g-ore/MWH.

2.1.4 Economics

Ultimately, the key factor that will make or break any nuclear fuel cycle deployment scenario is cost. Traditionally, the key strength of nuclear electric generation is its ability to provide stable base-load power with an exceptionally low marginal fuel cost. This is doubly true for the existing fleet of aging reactors, as the bonds used to initially finance their construction have been retired, and years of experience have led to gains in capacity factors and decreases in operating costs. The fuel savings are significantly offset, however, by the enormous capital and associated financing costs associated with the reactor system. Proposed methods of closing the fuel cycle, thereby recovering greater energy value from the fuel and minimizing waste, all

require the construction of several new types of facilities, which will further drive up system costs. If costs become too high, the market will naturally move to any cheaper, available options unless government policies promote a different behavior. Because the `VISION` model considers only the nuclear portion of energy production, the objective function we wish to minimize is the total cost of the nuclear enterprise under the constraint that the energy demands are met.

In analyzing the economics of a proposed fuel cycle strategy, two key issues must be addressed: 1) How to establish reasonable cost factors and their uncertainty bounds for all facilities and services, and 2) How to account for the time value of money, both in terms of general inflationary pressures and the need to ensure a reasonable return on investments. The economic objective function is the most complicated of the available objective functions; it tracks most major facilities through construction, operation and shutdown and accounts for costs on all major processing stages. The methodology used to calculate the net costs is adapted from that used in the Advanced Fuel Cycle Cost Basis [29]. Costs are normalized by the net electric production to yield a net unit cost of electric energy that may be compared against other predictions of future energy prices. The economic objective function is used in multi-objective optimizations to provide an estimate of the economic trade-off costs associated with minimizing the other objective functions.

Time-value of Money

The purchasing power of money tends to decrease over time through the well known process of inflation. To compensate, the prices and costs of materials and services tend to escalate. However, inflation and escalation rates do not necessarily coincide, and can vary due to a host of complicated factors including advances in technology, national monetary policy and the overall state of the economy. Among those who study it, there is some disagreement as to the importance and causality of the various factors that contribute. Thus the ability to make accurate predictions, even over relatively short time-frames, is quite limited. Over the 100+ year timespan considered in this work, it is unlikely that any prediction will remain valid for the entire span. Despite the certainty of discrepancy, an assumption must be made in order generate cost estimates; the simplest assumption is that cost escalation and monetary inflation are matched throughout the time period of the simulation. Thus, when levelized to a given year, the costs remain constant. The benefits of this assumption include its computational simplicity and the availability of comparative benchmark studies. The primary drawback is that any financial advantages to advancing or postponing large capital expenditures are lost.

The costs associated with the nuclear enterprise can be divided into three broad categories: Capital costs for the facilities, Operations and Maintenance costs for each facility, and service

and resource costs. Net costs are computed according to the following equation

$$C = \sum_{i \in rx} C_i^{CAP} N_i \tag{2.16}$$

where:

$$N_i = \text{Total number of facility type } i \text{ built.} \tag{2.17}$$

The accounting of these costs and their values are described in the following three sections.

Capital Costs

The largest single category of costs in almost any vision of the nuclear power enterprise are those costs associated with the construction and financing of the facilities themselves. The overall capital cost for facility type i is broken down into the following six component factors:

$$C_i^{CAP} = C_i^{ON} P_i T_{rec} F_D F_{DD} F_C F_B \tag{2.18}$$

where:

$$\begin{aligned} C_i^{ON} &= \text{Overnight Capital Cost per Unit Capacity} \\ P_i &= \text{Facility Capacity} \\ T_{rec} &= \text{Economic Recovery Time} \\ F_D &= \frac{r_D}{1-(1+r_D)^{-T_{rec}}} \\ F_{DD} &= (1+r_{DD}) \\ F_C &= (1+r_C) \end{aligned} \tag{2.19}$$

$$\begin{aligned} F_B &= \frac{e^{r_B T_B}-1}{r_B T_B} \\ r_D &= \text{Discount Rate of Return} \\ r_{DD} &= \text{Decontamination and Decommissioning Fund} \\ r_C &= \text{Contingency Fund} \\ r_B &= \text{Interest Rate on Construction Bonds} \\ T_B &= \text{Construction Time} \end{aligned} \tag{2.20}$$

The product of the specific overnight cost (C^{ON}) and the facility capacity (P) are the net cost of all land, equipment, construction and administration and other work required to license and build the plant from its component parts (less financing during construction). The factor accounting for the costs of financing during construction (F_B) assumes that the overnight cost is

expended at a uniform rate during a construction time (T_B), and that interest on the outstanding balance is compounded continuously at a rate of r_B. Contingency and Decontamination and Decommissioning (D&D) costs are estimated for some facilities to be a fixed fraction of the capital cost. The values for the contingency factor are related to the technological maturity of each facility. Unproven fast burner reactors naturally require a higher contingency than established LWR reactors. Likewise, the D&D factor is estimated based upon the amount of potentially activated hard components and structural material that must be disposed of at the end of the facility lifetime.

The largest of the capital cost factors is the discount rate of return on investment. For any investment of funds, the investors (be they individual, corporate, or government) expect to recover their original investment with interest. This money is paid out over the economic lifetime of the facility (taken here to be equal to the functional lifetime). For example, given a sixty year lifetime and a 10% rate of return, one finds F_D to be a substantial 6.02. The individual facility cost factors are listed in Table 2.1.4. As with all cost estimates, there is a degree of uncertainty in where the actual value will fall. As a first estimate, many of the capital cost factor values are assumed to be given by a triangular probability distribution, wherein the probability density of a given value occurring increases linearly from zero moving from the low to the nominal value, and decreases linearly to zero moving from the nominal to the high value (this is in keeping with the methodology used in the AFC Cost Basis Report). In most cases, these values are estimated by expert judgement, based upon a survey of relevant documentation and experience base. The interested reader is referred to the body of the Cost Basis Report itself, where the methodology and development process is fully explained.

Table 2.5: Capital Cost Components

Facility	**Component**	**Low**	**Nominal**	**High**	**Units**
LWR	C^{ON}	2,300	3,500	5,000	\$/kw(e)*
	r_D	5	10	15	%/yr
	r_C	10	11	12	%
	r_B	8	10	12	%/yr
	r_{DD}		0		%
	T_B		5		yr
	T_{rec}		60		yr
	P		1.05		GW(e)
FBR	C^{ON}	3,000	4,200	7,000	\$/kw(e)*
	r_D	5	10	15	%/yr
	r_C	20	25	30	%
	r_B	8	10	12	%/yr
	r_{DD}		0		%
	T_B		5		yr
	T_{rec}		60		yr
	P		0.6		GW(e)
UREX-1a Separations	C^{ON}	460	502	829	\$/kg†
	r_D	5	10	15	%/yr
	r_C	10	15	20	%
	r_B	8	10	12	%/yr
	r_{DD}		33		%
	T_B		5		yr
	T_{rec}		60		yr
	P		0.1		Kt/yr

Continued on Next Page

Table 2.5: continued

Facility	Component	Low	Nominal	High	Units
FBR Separations	C^{ON}	2,500	5,000	7,500	\$/kg[†]
	r_D	5	10	15	%/yr
	r_C	20	25	30	%
	r_B	8	10	12	%/yr
	r_{DD}		33		%
	T_B		5		yr
	T_{rec}		60		yr
	P		0.05[‡]		Kt/yr
FBR Fab	C^{ON}	1,200	1,800	2,700	\$/Kt[†]
	r_D	5	10	15	%/yr
	r_C	10	15	20	%
	r_B	8	10	12	%/yr
	r_{DD}		10		%
	T_B		5		yr
	T_{rec}		60		yr
	P				Kt/yr[§]

* AFCI Cost Basis Report (2009)

† VISION.ECON v2.1.9

‡ Scaled from 0.1 Kt/yr to reflect observed utilization rates.

§ Capital costs of FBR fuel fabrication are based solely on throughput, not facility size.

Operations and Maintenance Costs

The capital costs associated with building and financing a given facility do not cover the expenses associated with the continued operations of the facility. The ongoing operations and maintenance costs are accrued due to a complex and facility dependent mixture of factors, including replacement of depreciable components, workforce development and staffing costs, and ongoing security and regulatory compliance costs. To simplify the modeling of these costs,

they are divided into two categories: 1) those that occur at a fixed rate for a given facility size, regardless of its utilization, and 2) those that are proportional to the throughput of the facility (so-called variable costs). The values used for each of these factors are listed below in Table 2.6. As with the capital cost factors, the O&M cost factors are assumed to occur within a range given by a triangular probability distribution characterized by the listed low, nominal and high values. The costs for the actual nuclear materials required for each facility are accounted separately from Operations, and are listed in the next section.

Table 2.6: Facility Operations and Maintenance Cost Components

Facility	**Component**	**Low**	**Nominal**	**High**	**Units**
Light Water Reactor	Variable	0.80	1.80	2.50	mills/kWh*
	Fixed	55	66	80	$/kw(e)/yr*
Fast Burner Reactor	Variable	1.00	2.00	2.70	mills/kWh*
	Fixed	60	70	85	$/kw(e)/yr*
UREX-1a Separations	Variable	0.80	1.80	2.50	mills/Kt†
	Fixed	80	100	120	$MM/Kt/yr†
FBR Separations	Variable	1.00	2.00	2.70	mills/Kt†
	Fixed	80	100	120	$MM/Kt/yr†
FBR Fuel Fabrication	Variable	1.00	2.00	2.70	mills/Kt†
	Fixed	80	100	120	$/Kt/yr†

* AFCI Cost Basis Report (2009)
† VISION.ECON v2.1.9

Service and Materiel Costs

Costs for certain fuel cycle operations are computed directly based upon their usage, without calculating separate capital and operations costs. This is done for various reasons, depending on the operation in question. For example, uranium mining, conversion, enrichment, and LWR fuel fabrication are all mature technologies with an established industrial base; thus prices for each are determined by market forces. Other operations are likely to be co-located with other major fuel cycle facilities, and thus share most of the fixed capital expenses; examples of this primarily include waste conditioning costs (done at the reprocessing or reactor facility prior to shipment). As before, these costs are estimated to fall within a given range according to either a triangular

or uniform probability distribution as noted in Table 2.7. The total cost for each individual factor is computed by multiplying the unit cost of each materiel or service by the net amount required by a given deployment scenario (as indicated by the VISION model).

Table 2.7: Costs of Fuel Cycle Services and Materials

Service Component	**Low**	**Nominal**	**High**	**Units**	**Distribution**
Uranium Ore	30	75	260	\$/kg-U	Triangular*
Conversion	5	10	15	\$/kg-UF_6	Uniform*
Enrichment	85	110	135	\$/SWU	Uniform*
LWR Fuel Fabrication	200	250	300	\$/kg-U	Triangular*
Monitored Retrievable Storage	94	96	116	\$/kg-HM	Triangular*
Recycled Product Storage	3,300	4,400	6,000	\$/kg-Actinides	Triangular*
Dry Storage	100	120	300	\$/kg-SNF	Triangular*
Cladding Conditioning	200	540	1,800	\$/kg	Triangular†
Spent Fuel Packaging	50	100	130	\$/kg-HM	Triangular*
Spent Fuel Disposal	400	650	1000	\$/kg-HM	Triangular*
HLW Conditioning	2,200	5,000	6,600	\$/kg-HM	Triangular*
HLW Disposal	152	211	360	\$/kg-HM	Triangular‡
Aqueous RU Disposal	7	9	30	\$/kg	Triangular*
Pyroprocessing RU Disposal	75	93	150	\$/kg	Triangular*
DU Disposal	6	11	50	\$/kg-U	Triangular*
LLW Conditioning	1,000	1,500	4,200	\$/m^3	Triangular*
Surface Disposal	450	1,250	2,500	\$/m^3	Triangular*

* AFCI Cost Basis Report (2009)
† VISION.ECON v2.1.9
‡ AFC Economic Sensitivity Analysis (2006)

Figure 2.2 gives an idea of the relative contributions of the various fuel cycle components to the overall cost of the nuclear enterprise. In this figure, the total cost of two reference configurations (All LWR once-through, and Rear-loaded FBR) are divided according to three separate categories:

- Fuel cycle role (i.e. Front-end, Back-end, or Reactor).
- Technology Type (i.e. Once-through or Closed cycle).

- Expense category (i.e. Capital, Operations and Maintenance, or Service).

As the figure makes clear, the greatest contribution to overall cost is the capital cost of the reactor itself. This factor is driven primarily by the combination of the high system cost and the need to provide a healthy return on the investment.

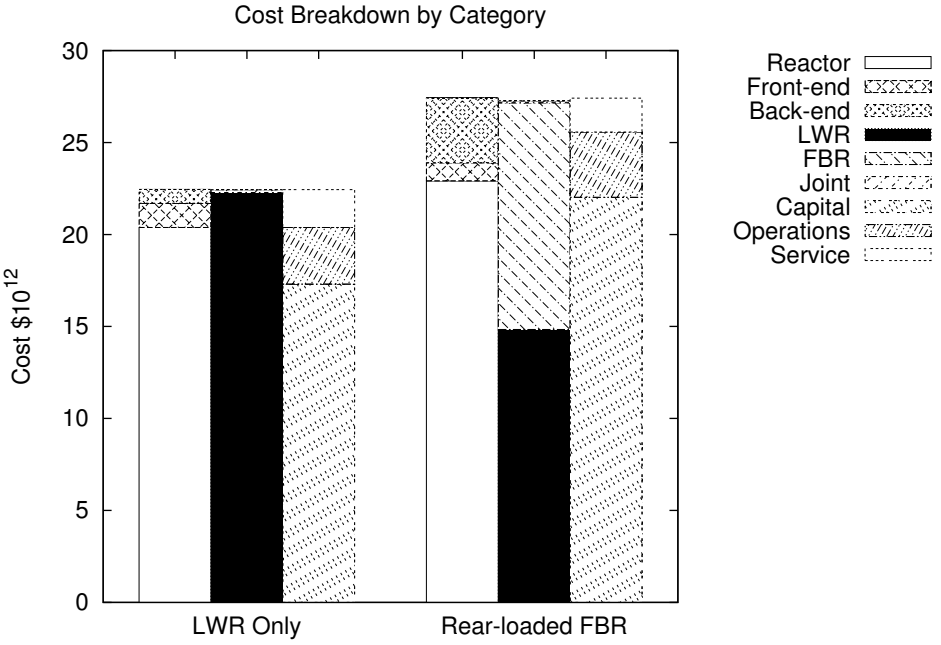

Figure 2.2: The total costs of two configurations broken down into three sets of cost categories.

Economic Uncertainties

As noted in previous sections, most of the cost factors exhibit some degree of uncertainty. This uncertainty carries through into the calculation of a net cost of electricity. There are numerous methods available for estimating the overall cost probability distribution; for this work, a simple and robust Monte-Carlo approach is used. For each optimization cycle, a representative distribution of cost factor values is generated by repeatedly sampling each individual cost factor using their respective probability distributions. This distribution set of cost factors is then applied to each evaluated deployment scenario to yield a total-cost distribution set, which is then sorted into ascending order and normalized by the net electric energy production. Thus, an estimate of the expected cost is given by the value in the midpoint of the sorted sample set. Likewise, the value nine tenths of the way through the list gives an estimate of the cost that has

a 90% chance of not being exceeded. A drawback of this simple analog sampling method is that the repeatability and accuracy of the cost predictions at the edges of the probability distribution can be quite large relative to those near the center. This effect is mitigated by restricting the electric cost objective function from utilizing these bounding values.

A simple test was performed to establish a reasonable value for the Monte-Carlo sampling size: for each of a set of Monte-Carlo sample sizes, five separate cost-factor distributions were generated using distinct seed values for the random number generator. These distributions were then applied to two representative deployment scenarios. The sample relative standard deviation observed between net electric costs of the five distributions was plotted at 5% cumulative-cost-probability increments for each sample size and reference configuration. These results are shown below in Figures 2.3 and 2.4. Due to the small number of distributions sampled in this test, considerable variation from one Monte-Carlo sample set to another, however the variation from one distribution sample to another tends to decrease as the Monte-Carlo sample size increases. For 800 samples, the variation between distributions is generally less than 1-2% for cumulative cost probabilities between 10% and 90%. It should be noted that in increasing the Monte-Carlo sample size from 100 to 800, no appreciable increase in computation time was observed. Additionally, as the initial sampling of cost factor values only occurs once per optimization cycle, the overall run-time penalty of choosing a large sample size is minimal. Therefore, for the optimization trials to follow a conservative Monte-Carlo sample size of 800 is utilized for results produced prior to 1-May-2012, and 2500 for those after. Finally, to ensure that the economics results would be comparable across the various optimization runs, a separate random-number generator with a fixed seed value was implemented within the objective function. This ensures that the same set of sample values are used each time (provided, of course, that the initialization routines are not modified between runs).

The unit electric cost relative probability for the two reference configurations is plotted below in Figure 2.5.

Electric Cost Objective Function Value

The numerical value for this objective function is the all-in unit cost of that which will not be exceeded for 90% of the cost distribution cases (given in units of mills per kWh or dollars per MWh). This was chosen to represent a conservative, risk-averse approach to estimating costs of the nuclear enterprise.

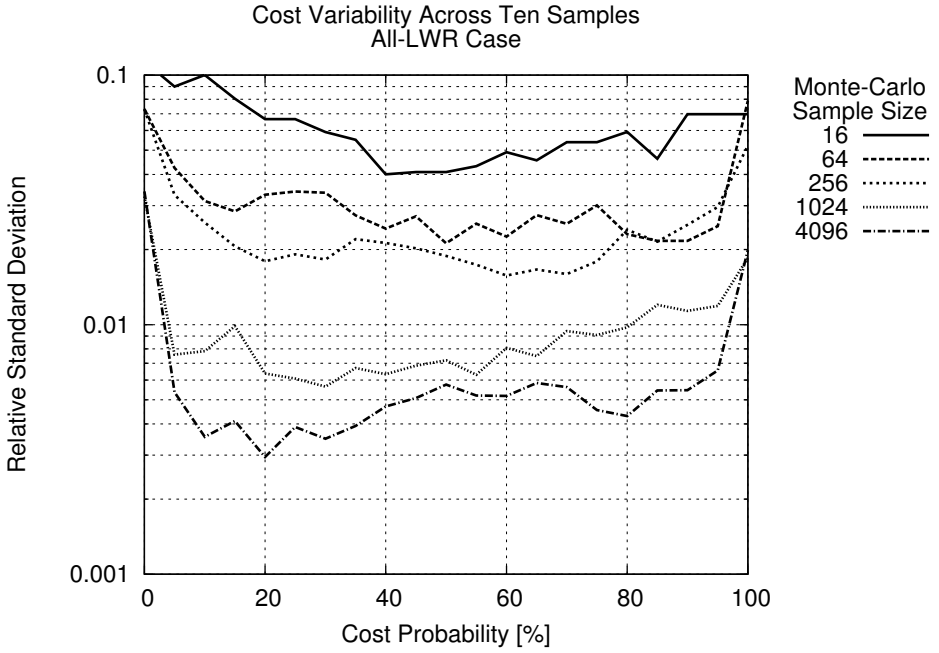

Figure 2.3: Cost variability by number of Monte-Carlo samples and cost-distribution level for five samplings with the all-LWR configuration.

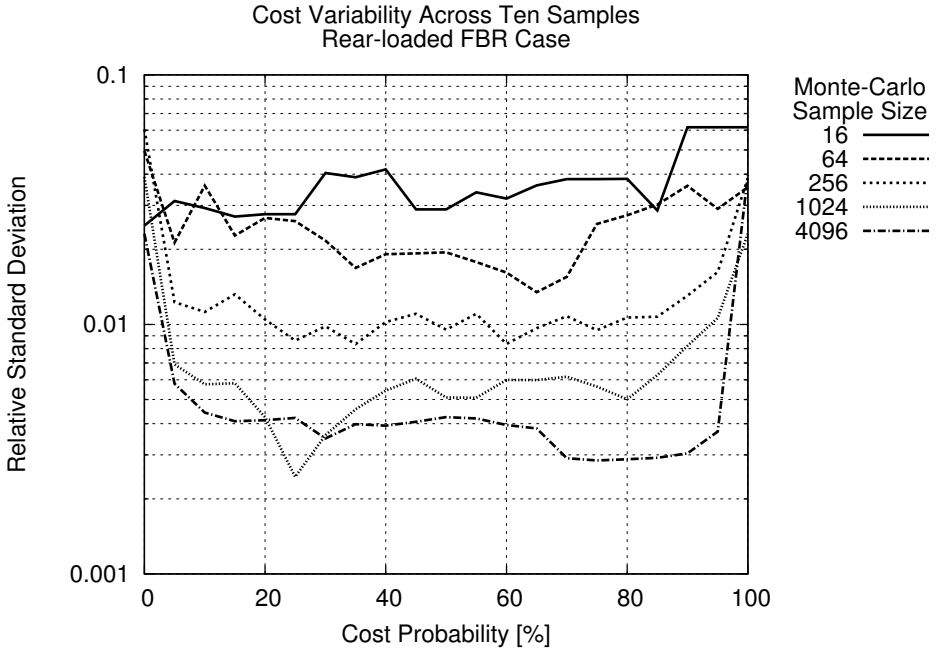

Figure 2.4: Cost variability by number of Monte-Carlo samples and cost-distribution level for five samplings with the rear-loaded FBR configuration.

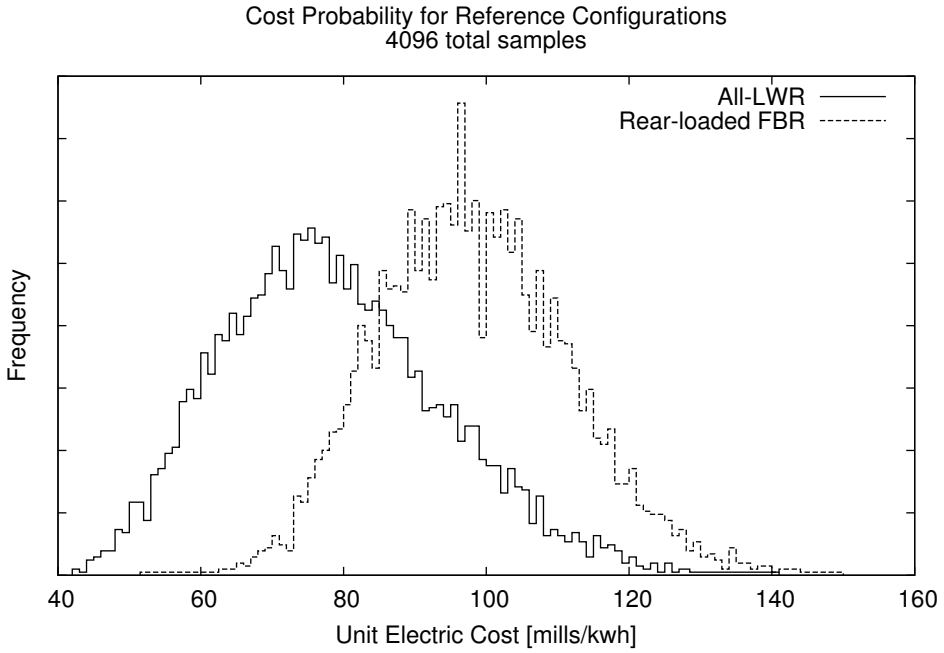

Figure 2.5: The probability of a given electric cost is plotted for two reference configurations.

2.2 Penalty Constraints and Other Functions

In addition to the primary optimization objectives there are a number of other constraining factors that either reflect the acceptability of a given candidate scenario or are useful for monitoring the progress of the optimization. The penalty constraints are delineated into two types: hard and soft. Hard constraints represent factors that are unacceptable in both final and intermediate scenarios. Any candidate scenario violating any hard constraint is automatically rejected. Currently, these constraints are used to detect and remove scenarios for which the `VISION` model does not produce valid results. Soft constraints represent those factors that are deemed undesirable in the final solution, but which are not ultimately fatal to its viability. Examples of this would include not meeting the prescribed electric power demand or deploying reactors for which sufficient fuel is not available. Soft constraint violations do not result in automatic disqualification of a candidate scenario; instead, the objective function value is penalized by an amount proportional to the extent of the violation and a corrective penalty multiplier. The violation value is dependent upon the formulation of the penalty constraint, which remains fixed throughout the optimization. The penalty multipliers, however, change adaptively throughout the optimization, gradually increasing in response to continued violations. Thus they steer the optimization toward acceptable solutions. There follows a description of the available soft penalty constraints and their formulations.

2.2.1 Unused Capacity

The `VISION` optimization strategy, through careful selection of input perturbation and use of built in heuristic rules attempts to constrain the simulation settings to those situations where fuel will be available for all of the reactors built. If, however, a pathological case manages to get past these checks, it will still be accounted for by explicitly tracking the amount and time for which each reactor type runs short of fuel. The constraint violation value is given by:

$$\Theta_{U.C.} = \sum_{i \in rx} C_i^{idle}, \tag{2.21}$$

where the terms and corresponding `VISION` variables are defined as above in Table 2.1.

2.2.2 `VISION` Heuristic Invocation

Within the `VISION` model there exist a set of heuristic rules which estimate and compare the mass of LWR spent fuel that will be available at a point in the future with the mass of the same required to operate a given number of fast burner reactors [27, 28]. These rules are available as an aid to a person using the `VISION` model in a standalone manner to manually explore

Table 2.8: Variables in the VISION Heuristic Penalty Function

Term	VISION Variable	Sample Time
B_i	Future Predictions Model.additional LWRmf reactors to be built[i]	All
$O_{i,rx}$	Future Predictions Model.reactor order rate for future years[rx,i]	All
P_{rx}	reactor power per CY[rx]	2200
T_{rx}	adjusted reactor and fuel fab lifetime[rx]	2200

fuel cycle deployment strategies. When active, these rules will prevent the construction of new fast burner reactors for which the heuristic spent-fuel-mortgaging model indicates there will be insufficient supply. In doing so, the heuristics decouple the action of the input variable selections (reactor fraction requested by year) from the output of the model (reactor fraction constructed by year). This decoupling is undesirable when using the VISION model within an automated optimization context, as the decoupling will truncate the expected increases in objective function value which are relied upon to guide the solution back toward a more feasible range. Additionally, the mortgaging scheme utilized by the VISION heuristics has lower fidelity and greater conservatism than the spent fuel forecasting scheme implemented directly in the optimization decision variable (see Section 2.3.3). Two remedies are available to counteract this occurrence: 1) disable the mortgaging heuristics within VISION, or 2) apply a penalty constraint function against the action of the heuristic rules.

The VISION heuristic penalty function is computed using Equation (2.22) and the variables listed in Table 2.8. The VISION heuristic rules act by requesting supplementary LWR reactors to meet energy demand when the forecast LWR SF inventories will not support the requested FBR build rate. The numeric value of this penalty function is the total fraction of energy produced by these so-called Bonus LWR reactors,

$$\Theta_H = \frac{\sum_{t=t_{start}}^{t_{stop}} B_t P_{LWRmf}\ T_{LWRmf}}{\sum_{i \in rx} P_i T_i \sum_{t=t_{start}}^{t_{stop}} O_{t,i}} \tag{2.22}$$

.

2.2.3 Smoothing Constraint

The infrastructure and expertise required to construct new nuclear power plants is neither inexpensive nor quickly developed. This is of particular concern for specialized components and competencies unique to nuclear construction; for example, there is currently only one foundry

in the world capable of forging a Generation III+ LWR pressure vessel. In the absence of any continued demand, it is difficult to justify the investment required to maintain – let along expand – these unique facilities. Furthermore, from a modeling perspective, it is difficult to properly and accurately estimate the build-out costs and carrying costs for these related industries. For these reasons and others, it is desirable that the reactor build rates produced by the optimization algorithm should exhibit a certain degree of "smoothness" in their variation from year to year.

$$\begin{aligned}
\Theta_s &= \max\left(\tilde{\Theta}_s^{N_s}, 0\right) \\
\tilde{\Theta}_s &= \frac{1}{N_{FBR}} \sum_{i=1}^{100} \left| O_{i,FBR} - \langle O_{i,FBR} \rangle_M \right| \\
&\quad + \frac{1}{N_{LWRmf}} \sum_{i=1}^{100} \left| O_{i,LWRmf} - \langle O_{i,LWRmf} \rangle_M \right| - \gamma_s
\end{aligned} \tag{2.23}$$

where:

$$i_{min} = \max\left(i - \frac{M-1}{2}, 0\right) \tag{2.24}$$

$$i_{max} = \min\left(i + \frac{M-1}{2}, 100\right) \tag{2.25}$$

$$\langle O_{i,rx} \rangle_M = \frac{\sum_{i'=i_{min}}^{i_{max}} O_{i',rx}}{1 + i_{max} - i_{min}} \tag{2.26}$$

$$N_{rx} = \sum_{i=1}^{100} O_{i,rx}$$

M is a constant, odd integer.

γ_s is a predetermined threshold value

N_s is a positive exponent value

The annual reactor order rate term, $O_{i,rx}$, is calculated by summing up the per-timestep reactor-startup rate reported by the `VISION reactors to start` variable. Equations (2.24) through (2.26) convert this into a moving-average reactor build rate with an M-year sliding window centered about the year i (ignoring those years that fall outside of the range $1 \leq i \leq 100$).

Through testing, parameter values $\gamma_s = 0.65$, $M = 5$, and $N_s = 3$ have been selected.

2.3 Decision Variables

The `VISION` model provides a wealth of parameters and options that may be adjusted to specify a wide range of fuel cycle options. Also provided are a large set of base-cases, which completely specify many commonly examined options (one-tier, two-tier, various recycling schemes, etc).

One common feature to most base-cases is the ability to specify the percentage of new reactor capacity to be built in any year that is filled by each of three main reactor types (LWR, LWRmf, and FBR). These reactor split profiles are used as the primary decision variable in simulations. Thus, the optimization algorithm is able to perturb the rate and time at which new reactor types are introduced or phased out in order to find a pattern that minimizes the selected objective functions. Although it is possible to generate candidate configurations through purely stochastic means, the various material feedback mechanisms in the fuel cycle make it unlikely that such configurations would be feasible, therefore greatly slowing the optimization search. Instead, by applying a simplified spent fuel forecast, stochastic perturbations may be applied in a way that gives a high confidence that the resulting configuration will not have any grossly infeasible features.

2.3.1 Definitions

The derivation of the spent fuel forecasting algorithm and the decision variables that follow require the introduction of a number of new terms and notation.. For convenience, they are listed below in Table 2.9. In all of these definitions, x can be any one of the reactor types used within the `VISION` model (LWR, LWRmf, FBR). In certain instances where the values for LWR and LWRmf reactors are identical they are further abbreviated to L.

Table 2.9: Definitions used in the Maximal FBR Approach

Term	**Description**
m^{f}_{b,P_i}	The mass of the limiting active isotope present in the $f$eed in a single $b$atch of Pass i fuel. (For simplicity of notation, the reactor type is not explicitly denoted in this variable. It may be determined from the context of the equations in which it is used.)
m^{S}_{b,P_i}	Likewise, the mass of the limiting active isotope present in a single batch of Pass i Spent discharged fuel.
N_{batch}	The number of fuel batches in a reactor at a given time.

Continued on Next Page

Table 2.9: (continued)

Term	Description
P, PH, BPH, BoPH, EoPH	Simulation time can be divided into two time periods (generically denoted P). Variables that pertain to the years within the Planning Horizon are denoted with a superscript PH. Those pertaining to the years Beyond the Planning Horizon are denoted with a superscript BPH. The beginning of the planning horizon (BoPH) is set at the year 2000, while the end of the planning horizon (EoPH) occurs at 2100.
X, LWR, FBR	The `VISION` simulation models used in this work contain two primary types of reactors: metal-oxide fueled reactors with a thermal neutron spectrum, and metallic-fueled reactors with a fast neutron spectrum. The thermal reactors may use either Uranium Oxide (UOX) or Mixed plutonium oxide (MOX) fuel. In equations, these reactor types are denoted with the LWR subscript. The fast reactors may operate with a variety of fuel recipes in either a breeding or a burning regime. In equations, they are denoted with the FBR subscript. In equations that apply to both reactor types, a generic x may be substituted for clarity and brevity.
ΔQ_x	The lifetime energy production of a reactor of type x
$\tilde{M}_k^{SF}$	The active isotope (TRU) mass of spent fuel existing at year k due to reactors built within the planning horizon.
$\dot{M}_k^{LSF}$	The active isotope (TRU) mass of spent fuel generated at year k due to legacy reactors
$\Delta \hat{M}_x$	The mass of spent fuel discharged by a single reactor of type x over its entire lifetime.
$\delta \hat{M}_{FBR}$	The mass of spent fuel discharged from the last operating FBR at shutdown.
$\dot{m}_x$	The mass of spent fuel discharged from a reactor of type x during a single year of operation.
$N_{x,k}$	The number of reactor type x to be built in year k.
d_k	The total power demand at year k
d_k^{net}	The net power demand (total less legacy contributions) at year k

Continued on Next Page

Table 2.9: (continued)

Term	Description
g_k	The gap between the electric power produced by reactors operating at year $k-1$ and the electric demand at year k.

2.3.2 Defining the Planning Horizon

Before deriving any equations or drawing any results, we must first fully define what is meant by the Planning Horizon. Simply put, the Planning Horizon is that period of time over which control is exerted over the fuel cycle, its facilities and its operations. For the present purposes, the Planning Horizon encompasses the years 2000 through 2100 in the `VISION` model. The goal of the Maximal FBR approach is to give a clean separation between those effects that occur after the end of the planning horizon (EoPH) that are due to actions occurring during the planning horizon from those due to actions occurring after the EoPH. Thus, when considering the planning horizon, we include any and all reactors constructed between the years 2000 and 2100, but also their fuel usage and energy production from 2100 through their end of life. In the calculations that follow, fresh, unirradiated nuclear fuel is NOT included, as it poses no long-term storage or disposal burden. Legacy spent fuel leftover from prior to the start of the planning horizon, however, is included. Two different approaches are used in classifying the spent fuel that is generated during the planning horizon itself.

The socially conscious approach to achieving closure of the planning horizon requires that the type-distribution of the fleet of reactors existing at the end of the active planning horizon will consume, to the fullest extent possible, all of the spent fuel available for recycling while maintaining sufficient fuel to operate each reactor over its planned lifetime. In other words, it aims to minimize the spent fuel burden carried forward beyond the planning horizon. The state of the system at the end of the planning horizon depends upon the decisions made at each previous time-step and upon the initial conditions. Therefore, this approach to closing the planning horizon takes the form of an integral, heuristic constraint on the simulation input parameters (specifically, the percentage of each reactor type requested). This approach also has the advantage of narrowing the configuration space for optimization, thus giving faster convergence and reasonable run times.

The dynamic grown approach to closing the planning horizon instead treats spent fuel as it is currently treated – as a resource for future use. Thus, instead of constraining the mixture of reactors to a relatively high percentage of fast reactors, a more conservative LWR-heavy portfolio may be utilized. As in the socially conscious approach, fast reactor deployment is

limited such that sufficient fuel stocks exist for all reactors.

2.3.3 Simplified Spent Fuel Forecasting

The current implementation of `VISION` uses an LWR spent-fuel mortgaging scheme to predict the amount of recyclable fuel available to supply future fast reactor needs [27, 28]. This mortgaging scheme forms the basis for the built-in heuristic rule set. In order to generate feasible input configurations for the optimization work, it is desired to implement a similar capability within the optimization framework. The two key differences between the built-in heuristic mortgaging scheme and the external forecasting scheme are that 1) the external forecast accounts for LWR SF production and consumption on a year-by-year basis for each reactor, where the mortgaging scheme allocates the entire reactor-lifetime fuel production/consumption in a single step, and 2) where the mortgaging scheme is used to limit the number of fast reactors such that no fuel shortages occur, the forecasting scheme has no such feedback and will happily compute infeasible scenarios.

The primary output of the forecasting method is, of course, the year-by-year estimate of the LWR spent-fuel inventory. In order to produce this a number of ancillary values must first be calculated. For computational reasons, these values are split between those that depend upon fixed base-case parameters and those that depend upon input variable parameters. The first fixed variable to be calculated is the electric demand profile for the planning horizon, d_i. The `VISION` model assumes that electric demand grows with a specified rate, r_j, for each year j in the planning horizon. Thus $d_j = (1 + r_j)d_{j-1}$ where $2001 < j < 2100$ and d_{2000} is the initial demand at year 2000. Next, the number, capacity, and remaining lifetime of the pre-existing legacy reactors are pulled from the `VISION` model; these are combined with the electric demand to calculate the net demand, d_j^{net}, for which new reactors must be built. Next, the per-operating-year spent fuel consumption or production profiles for each reactor type are computed and cached for later use. Finally, the spent fuel availability profile due to previously existing fuel and operating legacy reactors is computed and stored.

With the required scenario constants pre-calculated, it is possible to generate the forecast for a given reactor distribution profile. Given a reactor request profile, $\chi_{rx,i}$, which specifies what percentage of newly constructed capacity in year i will be of type rx, the actual number of reactors to build and spent fuel inventory can be calculated iteratively using the following

equations:

$$g_i = d_i^{net} - \sum_{rx} \sum_{j=1}^{T_{rx}^{life}} N_{rx,i-j} P_{rx} \tag{2.27}$$

$$N_{LWRmf,i} = \begin{cases} \tilde{N}_{LWRmf,i} - 1 & \text{if } \tilde{N}_{LWRmf,i} P_{LWRmf} + \tilde{N}_{FBR,i} P_{FBR} - g_i > P_{LWRmf} \\ \tilde{N}_{LWRmf,i} & \text{otherwise} \end{cases}$$

$$N_{FBR,i} = \begin{cases} \tilde{N}_{FBR,i} - 1 & \text{if } P_{LWRmf} \geq \tilde{N}_{LWRmf,i} P_{LWRmf} + \tilde{N}_{FBR,i} P_{FBR} - g_i > P_{FBR} \\ \tilde{N}_{FBR,i} & \text{otherwise} \end{cases}$$

$$\tilde{N}_{rx,i} = \left\lceil \chi_{rx,i} \frac{g_i}{q_{rx}} \right\rceil \tag{2.28}$$

$$\tilde{M}_k^{SF} = (1 - \lambda_{TRU}) \left\{ \tilde{M}_{k-1}^{SF} + \sum_{rx' \in rx} \sum_{j=0}^{\bar{T}_{rx',k}^{life}} N_{rx',k-j-T_{rx'}^{prep}} \dot{m}_{rx',j} + \dot{M}_k^{LSF} \right\} \tag{2.29}$$

where

$$\begin{aligned} rx &\equiv \{LWRmf, FBR\} \\ i &\in [2000, 2100] \\ k &\in \left[2000, 2100 + T_{rx'}^{life}\right] \\ \dot{m}_{rx',j} &\equiv \text{Net SF production during year } j \text{ of operation} \\ T_{rx'}^{prep} &\equiv \text{Fuel Preparation Time} \\ \bar{T}_{rx',k}^{life} &= \min\left(T_{rx'}^{life}, k - T_{rx'}^{prep}\right) \\ g_i &\equiv \text{Electric Power Gap at year } i \\ P_{rx} &\equiv \text{Reactor Annual Energy Production} \\ \tilde{M}_{2000}^{SF} &\equiv \text{Initial Legacy Fuel Inventory} \\ \dot{M}_k^{LSF} &\equiv \text{legacy reactor fuel production at year } k \\ \lambda_{TRU} &\equiv \text{TRU annual decay fraction} \end{aligned}$$

In these equations, the "Fuel Preparation Time" is the summation of the LWR spent fuel cooling time (5 years), the separation time (1 year), and the FR fuel fabrication time (1/2 year) rounded up to the nearest integer (7 years). The spent fuel forecast adds spent fuel to the inventory at the moment it is discharged from the thermal reactor, and removes it from the inventory 7 years before it is to be loaded into fast reactor. Although this does not reflect the actual fuel storage locations in the `VISION` model, it is functionally equivalent and allows for a degree of

computational simplification. Also note that while the forecast for electric power gap and reactor order rate applies only for the years 2000 through 2100, the LWR SF forecast is applicable through the lifetime of the last reactors ordered in the year 2100. A pictorial representation of the categorization of spent fuel in the forecasting method is shown in Figure 2.6.

All of the transuranic isotopes to a greater or lesser extent are subject to radioactive decay during storage. While the majority of these transuranic isotopes decay into other transuranic isotopes, a certain fraction, λ_{TRU}, each year will decay into lighter elements. In order to estimate this decay factor for the stored TRU, it is necessary to assume a basis composition vector to which the decay matrix may be applied – for this the `VISION` LWR spent fuel recipe with a burnup of 51GWD/MTU was chosen, yielding a value of $\lambda_{TRU} = 0.00022$.

Because the goal of the socially conscious approach is to minimize the amount of usable spent fuel remaining at the end of the planning horizon, the ability to quickly forecast this value allows the optimization wrapper to both quickly verify the validity of a proposed configuration and iteratively approach certain defined configurations without invoking the heuristic rules embedded in VISION. One such defined configuration is the so-called rear-loaded initial configuration, which is used for many of the optimization tests. In this configuration, the optimization wrapper starts with an initial reactor request distribution of 100% LWR reactors. Then, working backward from the year 2100 in one year increments, it substitutes FBRs for LWRs until the forecast LWR spent fuel inventory at the end of the planning horizon reaches the smallest non-negative value possible. (In practice, it iterates until a negative value is obtained, then steps back). By defining the initial configuration algorithmically, it is not necessary to recalculate the required distribution manually every time a parameter value changes.

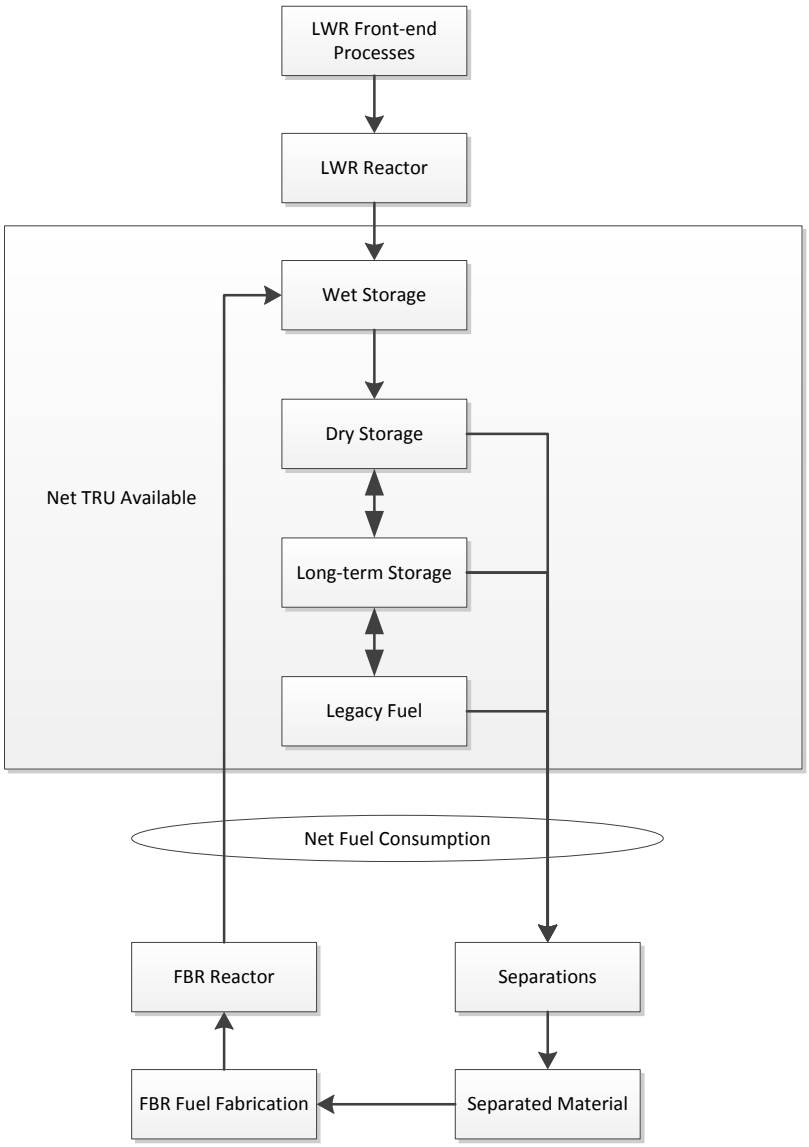

Figure 2.6: The LWR SF forecasting method aggregates and predicts spent fuel masses according to three primary categories in order to simplify calculations.

Spent Fuel Production and Consumption

In order to forecast the LWR SF levels through the simulations, we must first calculate the production or consumption profiles for each reactor type. This, in turn, requires specifying certain aspects of the fuel cycling scheme. It is standard practice in most power reactors to shut down for refueling every 12-24 months (denoted T_{cycle} herein). At each outage, only a certain fraction of the fuel is replaced; by having each *batch* of fuel in the core for multiple cycles, a greater amount of energy may be extracted, and the uranium economy improves. (These improvements are of course balanced against the down time required for each outage). Thus, at any one time, the reactor contains N_{batch} batches of fuel, each with an initial feed mass of m^f_{b,P_i}, where P_i indicates how many times (passes) the fuel has been recycled (P_0 is made from virgin ore, P_1 is made from recycled P_0, and so on until pass 5, where it is considered to have reached equilibrium). On discharge, the total mass and composition will have changed slightly, so it is designated separately by m^S_{b,P_i}.

In the `VISION` base case used for this work, two types of reactors are used, each with their own fuel cycling strategy. The Light Water Reactor (LWR) and Light Water Reactor mixed-fuel (LWRmf) varieties are typical of the current US fleet of Pressurized Water Reactors. At startup, they are loaded with N_{batch} batches of fuel made from virgin ore, and they cycle through an additional batch every T_{cycle} years. At each cycle until the end of life, they discharge one batch of spent fuel. When they are shut down, the remaining N_{batch} batches are discharged. LWR and LWRmf operating parameters are listed in Table A.2.

The Fast Burner Reactor (FBR) type is typical of a liquid-metal cooled, fast-spectrum reactor. Though the thermal power levels are assumed in `VISION` to be much lower than the LWR reactors, the temperature and power density are considerably higher. Furthermore, whereas the LWR fleet is fueled by virgin, Pass 0 fuel, each FBR recycles its spent fuel, using LWR spent fuel to make up the difference between the spent and feed fuel compositions. Indeed during the early cycles of a new fast reactor the fuel material is coming completely from spent LWR fuel since it takes a number of years to recycle its own discharged fuel back into the reactor. Newly discharged spent fuel must cool in wet storage for a given time, T_{wet}, before it can be sent through separations and fuel fabrication (which require time T_{sep} and T_{fab}, respectively). The sum of these three times is called the *pipeline time*, denoted $T_{pipeline}$ for simplicity. This number is then added to the time that the fuel spends in the reactor to give the so-called *around time*, T_{around} (i.e. the time for fuel to make one pass around a cycle of irradiation and recycling). (See also Equations 2.11 through 2.14 in [27]). The net effect is that each reactor type has a defined LWR spent-fuel production or consumption profile, the formulas for which are listed in Tables 2.10 and 2.11. These values are computed on a per-time-step basis within the `VISION` simulation and subsequently converted to an annual basis within the optimization wrapper.

Table 2.10: Spent Fuel Production Profile for Light Water Reactors

From Time	**To Time**	**SF Mass**
0	$T_{lifetime}$	$m^S_{b,P_0} \frac{N_{batch}}{T_{cycle}}$
$T_{lifetime}$	–	$m^S_{b,P_0} \frac{N_{batch}}{\Delta t}$

Table 2.11: Spent Fuel Consumption Profile for Fast Burner Reactors

From Time	**To Time**	**SF Mass**
0	–	$m^f_{b,P_1} \frac{N_{batch}}{\Delta t}$
0	$T_{pipeline}$	$\left(m^f_{b,P_1}\right) \frac{N_{batch}}{T_{cycle}}$
$T_{pipeline}+1$	$T_{pipeline}+T_{around}$	$\left(m^f_{b,P_2} - m^S_{b,P_1}\right) \frac{N_{batch}}{T_{cycle}}$
$T_{pipeline}+T_{around}+1$	$T_{pipeline}+2T_{around}$	$\left(m^f_{b,P_3} - m^S_{b,P_2}\right) \frac{N_{batch}}{T_{cycle}}$
$T_{pipeline}+2T_{around}+1$	$T_{pipeline}+3T_{around}$	$\left(m^f_{b,P_4} - m^S_{b,P_3}\right) \frac{N_{batch}}{T_{cycle}}$
$T_{pipeline}+3T_{around}+1$	$T_{pipeline}+4T_{around}$	$\left(m^f_{b,P_5} - m^S_{b,P_4}\right) \frac{N_{batch}}{T_{cycle}}$
$T_{pipeline}+4T_{around}+1$	$T_{lifetime}$	$\left(m^f_{b,P_5} - m^S_{b,P_5}\right) \frac{N_{batch}}{T_{cycle}}$

Initial Spent Fuel Stocks

The Legacy spent fuel category encompasses both spent fuel existing prior to the start of the simulation and spent fuel produced by the continued operation of those reactors that exist at the start of the simulation. The initial spent fuel inventories are spread amongst several stages of the fuel cycle, and are specified through base-case settings for the variables listed in Table 2.12. By combining the year-by-year spent fuel production profile with the legacy reactor retirement profile it is a straightforward matter to calculate the legacy spent fuel mass produced at each year of the simulation, $\dot{M}_k^{LSF}$.

The legacy reactors (and all other reactors in the `VISION` model) retire through a three stage process that accounts for the time required to build replacement capacity and to produce fuel for the reactors. When the time reaches the `adjusted legacy reactor retirement date`, the appropriate number of retiring legacy reactors are transferred from the `Legacy Reactors` category, to the `Legacy Reactors Near Retirement` category. At this time, construction of replacement power begins. After a number of years (4 for all Light Water reactors), the reactors move to the `Legacy Reactors Near Shutdown` category. From this point on, no further fuel is ordered. Finally, after this last delay (3.25 years), the legacy reactors are moved into the `Retired Legacy Reactors` category.

Reference Configurations

The spent fuel forecasting algorithm is utilized to determine several of the initial configurations used for the optimization search. In the Rear-loaded FBR configuration, a date for changing over from all LWR reactor construction to all FBR construction is determined by iteratively forecasting the spent fuel inventory and moving the changeover date back from 2100 until a fuel shortage occurs. When the fuel shortage is encountered, the fraction of FBRs requested in the changeover year is then reduced until the shortage is alleviated. In the Constant-request FBR configuration, the requested fraction of FBRs for new construction in the years 2040 through 2100 is increased in 1% increments from zero until a fuel shortfall is predicted. It is then decreased again in 0.1% increments until the shortage is removed.

Limitations

The `VISION` model determines fuel movement based on the active isotope and total fuel mass, while enforcing a fixed set of fuel recipes. As a result, some of the uncontrolled isotopes are not conserved through the simulation (mass may be created or destroyed).

This forecast technique also neglects the decay of transuranic elements. Calculations of total transuranic decay rates using the `VISION` decay matrix and LWR spent fuel recipe show that

Table 2.12: VISION Initial Fuel Mass Terms

Description	VISION Variable	Sample Time
The mass of spent fuel initially sitting dry storage.	INITIAL SF DRY STORAGE	2200
The mass of spent fuel initially stored in a Monitored Retrievable Storage site.	INITIAL SF MRS	2200
The mass of spent fuel initially in the reactor spent fuel storage pools.	INITIAL SF WET STORAGE	2200
The mass of spent fuel initially existing due to legacy reactors, but not assigned to a storage location.	adjusted initial legacy SF	2200
Total initial legacy TRU inventory[1]	total initial SF active isotope	2200
The year at which the legacy reactors begin the retirement process	adjusted legacy reactor retirement date	2200
The number of years legacy reactors run after ordering (but before fabricating) the last batch of fuel	legacy reactors retirement delay[LWR]	2200
The number of years legacy reactors run after fabricating the last batch of fuel	legacy reactors shutdown delay[LWR]	2200
The rate at which legacy reactors retire	Adjusted Legacy Reactor Retirement Rate	2200
The initial number of legacy reactors	adjusted initial legacy reactors	2200

[1] This variable was added to VISION specifically for the optimization work, and is not part of the standard distribution.

only 0.0224% of TRU mass is lost to decay per year. Thus five years of storage will result in losses similar to the assumed 0.1% separations process loss.

At startup, the first $N_{batch} - 1$ batches of fuel to be discharged from a fast reactor will not have reached isotopic equilibrium; therefore the discharge isotopics will be somewhat different than the current model calculations. Similarly, at shutdown all N_{batch} batches of fuel are assumed to have full discharge burn-up, when in reality only one batch would.

2.3.4 Maximal FBR Perturbation

The simulated annealing algorithm explores the configuration space by repeatedly perturbing the inputs to the simulation.

The Maximal FBR perturbation algorithm is named such because it treats leftover LWR spent fuel as a waste material which is better disposed of through recycling than by leaving it as-is for subsequent generations to grapple with. This algorithm follows a simple, two step approach to maximize the number of FBRs built while maintaining sufficient LWR spent fuel supplies to ensure continuous operations. First, a randomly sampled LWRmf reactor within the planning horizon is replaced by the appropriate number of FBRs to meet electric demand. Then, if this substitution causes a shortfall in the forecast of available LWR spent fuel, an LWRmf reactor is reinserted at randomly sampled year prior to the forecast shortage.

The initial donor year, T_a, is selected by first compiling, and then sampling from, a list of all years between the fast reactor phase-in year (2040) and the end of the active planning horizon (2100) in which the forecast indicates an LWRmf reactor will be built. Let N_{LWR} , N_{FBR} , and $P_{surplus}$ be the currently forecast values for the numbers of LWRmf reactors and FBR reactors and the electrical power capacity to be built beyond demand (due to integer reactor ordering), respectively at year T_a. We wish to estimate the fraction of new power that must come from FBR reactors, $\tilde{\chi}$, if one of the LWRmf reactors is canceled. Therefore, additional FBR capacity must be substituted to meet the demand. However, given that there is not a one-to-one correspondence between LWRmf power and FBR power, and that reactors must be built in integral units, the following equations are used to compute the new fast reactor fraction.

$$\Delta = \max\left(\left\lceil \frac{P_{LWR} - P_{surplus}}{P_{FBR}} \right\rceil, N_{FBR}\right) \tag{2.30}$$

$$\tilde{\chi} = \frac{1}{1 + \left(\frac{N_{LWR} - 1}{N_{FBR} + \Delta}\right)\frac{P_{LWR}}{P_{FBR}}} \tag{2.31}$$

This new value is substituted into the requested distribution profile, and the deployment forecast is regenerated. The substitution of FBRs in place of LWRs will increase the overall TRU consumption, and can be expected to cause a fuel shortage at some point in the forecast – the

second step of the algorithm addresses this concern.

If a shortfall in LWR spent fuel inventory is predicted, then a list is compiled of all of the years preceding the shortfall in which a fast reactor is built. From this list, a recipient year, T_b is sampled, and, in a similar manner to the first step, an LWRmf reactor is built instead. The new requested FBR fraction for T_b is again denoted by $\tilde{\chi}$, given below.

$$\Delta = \max\left(\left\lfloor \frac{P_{LWR} + P_{surplus}}{P_{FBR}} \right\rfloor, N_{FBR}\right) \tag{2.32}$$

$$\tilde{\chi} = \frac{1}{1 + \left(\frac{N_{LWR}+1}{N_{FBR}-\Delta}\right)\frac{P_{LWR}}{P_{FBR}}} \tag{2.33}$$

Once again, the deployment forecast is recomputed with the new distribution, and step two is repeated until no shortfalls occur.

This entire perturbation sequence is repeated multiple times so that the configuration space may be more quickly spanned at high annealing temperatures. The exact number of repetitions is uniformly and randomly sampled to be between one and fifteen. In order to track the acceptance probabilities of a given number of perturbations relative another, an overall *perturbation distance* value is reported for each new sample. This value is simply the sum total of all LWRmf capacity removed from or added to the forecast reactor build profile.

2.3.5 Dynamic Growth Perturbation

The ability of spent nuclear fuel to be recycled proves that it is not necessary to unilaterally treat it all as waste, requiring prompt and permanent disposal. Instead it may be regarded as a valuable (though expensive and difficult to store) energy source. With this in mind, the perturbation constraint maximizing the requested number of FBRs may be relaxed, thus increasing the available optimization search space. In addition to covering a greater search range, this will allow for better determination of trade-off costs for reactor construction, uranium utilization and repository loading objectives, which depend more on total reactor number than on the actual deployment timing.

Like the socially conscious perturbation algorithm, the dynamic growth perturbation algorithm utilizes the simplified spent fuel forecast to ensure continuous operation of all reactors built. It differs from the previous algorithm by first randomly selecting (with equal probability) whether it will remove an LWR or an FBR from the request. This ensures that the total FBR count will increase or decrease with equal likelihood. As before, these single-reactor perturbations are repeatedly applied a (randomly sampled) number of times. Then a final SF forecast check is performed, and any fuel shortfalls are addressed by randomly removing FBRs from the request prior to the forecast shortfall time.

2.4 Parallel Simulated Annealing Framework

Much effort has gone into the design and development of a scalable and extensible computational framework for implementing a parallel simulated annealing (PSA) algorithm. The basic design parameters were largely specified by the requirements of the Powersim software package in which `VISION` was implemented. While the Powersim Studio program offers a rich graphic interface that is useful for the development and testing of systems dynamics models, it is poorly suited for the kind of repetitive batch operation that the simulated annealing algorithm requires. Fortunately, the underlying Powersim Simulation Engine exposes an Application Programming Interface (API) that allows the programmer to directly manipulate the model without invoking any interactive or graphical elements. Unfortunately, the Powersim Simulation Engine is currently only available for use on the Microsoft Windows operating system; this prevents its use on most high performance computing resources, which are predominantly Linux based. This interface is accessible through any of the .NET compatible languages; C# was chosen for this project. One advantage afforded by the use of the .NET platform was the ability to make use of a great deal of pre-existing code in the Framework Class Library [7]. These provide many of the basic building blocks required to support efficient inter-machine communication, process management and data handling.

In the PSA framework, one master process manages all program I/O, controls the annealing algorithm, and manages communication with other machines. This master process connects with separate server processes (possibly running on other machines) from which it receives so-called SimulationDrone instances. These SimulationDrone instances, which reside on the remote machines, perform the heavy lifting of repeatedly running the `VISION` simulation and report the results back to the master process. Currently the scaling of this method is limited only by the number of computers available to be dedicated to this project and by the archive memory requirements on the head process.

2.4.1 Cooling Schedule

In simulated annealing, the choice of the annealing temperature parameter and the Markov-chain update conditions comprise the so-called *cooling schedule*, which plays a critical role in determining the success and efficiency of the optimization. To guarantee convergence on the global optima, the distribution of accepted configurations must remain at equilibrium as the annealing temperature goes to zero. In practice, it would take far too long to achieve a true equilibrium state; instead, one tries to achieve a *quasi-equilibrium* state [33] by the end of each cooling step. This can be achieved either by making large reductions in annealing temperature followed by long cooling steps or by taking small, incremental temperature decrements and correspondingly shorter cooling steps. The latter approach is generally preferred.

In order to maximize the efficiency of the synchronous parallel annealing algorithm (and to maintain simplicity overall) the individual cooling chain segments are fixed to a constant length. After sampling a total of L_{chain} samples, or accepting a total of L_{tran} samples each process signals its readiness to update. (Currently, values of 75 and 30, respectively, are being used.) During the initial cooling step, where all proposed transitions are accepted, a separate segment length, L_{surv}, is used (currently set to 50). When more than 87% of the individual processes indicate readiness, an update is initiated.

The initial value of the annealing temperature is chosen by examining the outcome of the first cooling step. Because it is desired that the majority of samples be accepted early in the cooling cycle, the initial temperature is chosen to be a fixed multiple, B_T, of the standard deviation of the sampled objective values.

$$T_0 = B_T \sigma_F \tag{2.34}$$

Currently, B_T has been assigned a value of 2.0.

As it is desired to keep the sampled objective function distribution near to equilibrium after each temperature update, it is necessary to update the temperature in small increments such that the change to the equilibrium distribution is small. This is accomplished by using Huang's algorithm [1], which uses Equation (2.35) below to ensure that the expectation value of the new distribution is within one standard deviation of the previous distribution.

$$T_{k+1} = T_k e^{-\lambda_H T_k / \sigma_F} \tag{2.35}$$

The value of λ_H, defined by the user, is always less than 1.0, and is currently set at 0.7. Finally, to prevent quenching in the event that only a narrow distribution of samples has been accepted, the overall temperature reduction factor is limited to a minimum value of λ_c, generally greater than 0.5.

In order to develop the annealing algorithm and to measure the convergence of the results, a simple set of search termination criteria are used; the search is terminated after a total of L_{lngth} samples have been evaluated across all processes. The ability to restart a saved optimization run with different cooling parameters further enhances the ability to test end-of-cooling convergence properties and cooling parameter values.

2.4.2 Constraint Multipliers

During a single-objective optimization cycle, the penalty constraints are incorporated into the objective function through the use of penalty multiplier values, resulting in the *Augmented*

Table 2.13: Definitions of Optimization Terms

Term	Description
f	Unaugmented Objective Function
$\overline{f}$	Augmented Objective Function
Θ_i	Penalty Constraint Function
$\lambda_{i,k}$	Penalty Constraint Multiplier i at cooling step k
$\Lambda_{i,k}$	Mean value of Θ_i over cooling step k
$\mathcal{F}_{\tilde{N}}$	Final mean constraint violation reduction factor
χ_p	Foreshortening factor for constraint violation reduction
L_{lngth}	The specified maximum number of configurations to evaluate during a single cooling cycle.
L_{chain}	The specified maximum number of configurations to evaluate on a single parallel thread during a single cooling step.
L_{tran}	The specified maximum number of configurations to accept on a single parallel thread during a single cooling step.
L_{surv}	The specified maximum number of configurations to evaluate on a single parallel thread during the temperature initialization step.
N_{thread}	The number of parallel threads used in the optimization.
$N_{sampled}$	The total number of configurations sampled across all processes.

Objective Function, $\overline{f}_{obj}$.

$$\overline{f}_{obj} = f_{obj} + \sum_i \lambda_i \Theta_i \tag{2.36}$$

Thus penalized, the augmented objective function tends to constrain the solution chain to those solutions that minimize the constraint violations. The multipliers are increased in response to continued constraint violations, thus increasing the pressure on the optimization to reject offending solutions.

The algorithm used to increment the penalty multipliers is similar to that used in the `FORMOSA-B` code [4, 22]. The method aims to achieve a specified total reduction factor in the mean constraint violation value in even logarithmic increments. To do this, first define the average

constraint violation reduction factor for cooling step $k+1$,

$$\mathcal{F}_{k+1} = \Lambda_{k+1}/\Lambda_k \tag{2.37}$$

Next, define the average penalty reduction factor for cooling step $k+1$,

$$\mathcal{C}_{k+1} = \frac{\lambda_{k+1}\Lambda_{k+1}}{\lambda_k\Lambda_k} \tag{2.38}$$

Substitute (2.37) into (2.38) and solve for λ_{k+1} to get

$$\lambda_{k+1} = \frac{\mathcal{C}_{k+1}}{\mathcal{F}_{k+1}}\lambda_k \tag{2.39}$$

Assume that $\mathcal{C}_k$ and $\mathcal{F}_k$ remain constant throughout the cooling cycle, and require that by cooling step $\tilde{N}$, (defined below) that an overall specified reduction factor of $\mathcal{F}_{\tilde{N}}$ will have been achieved. It follows that if the mean constraint violation is to be reduced by the required factor by the requested cooling step, that

$$\mathcal{F}_{k+1} = 10^{\left(\frac{\log\left(\mathcal{F}_{\tilde{N}}\Lambda_0/\Lambda_k\right)}{\tilde{N}-k}\right)} \tag{2.40}$$

It is desired that the constraint violations be effectively removed before the end of the cooling cycle (after N steps), so the foreshortened value $\tilde{N}$ is used.

$$\tilde{N} = k + \chi_p\left[\frac{L_{lngth} - N_{sampled}}{L_{chain}N_{thread}}\right] \tag{2.41}$$

The constant parameter $\chi_p < 1$ is specified by the user (currently set to 0.75).

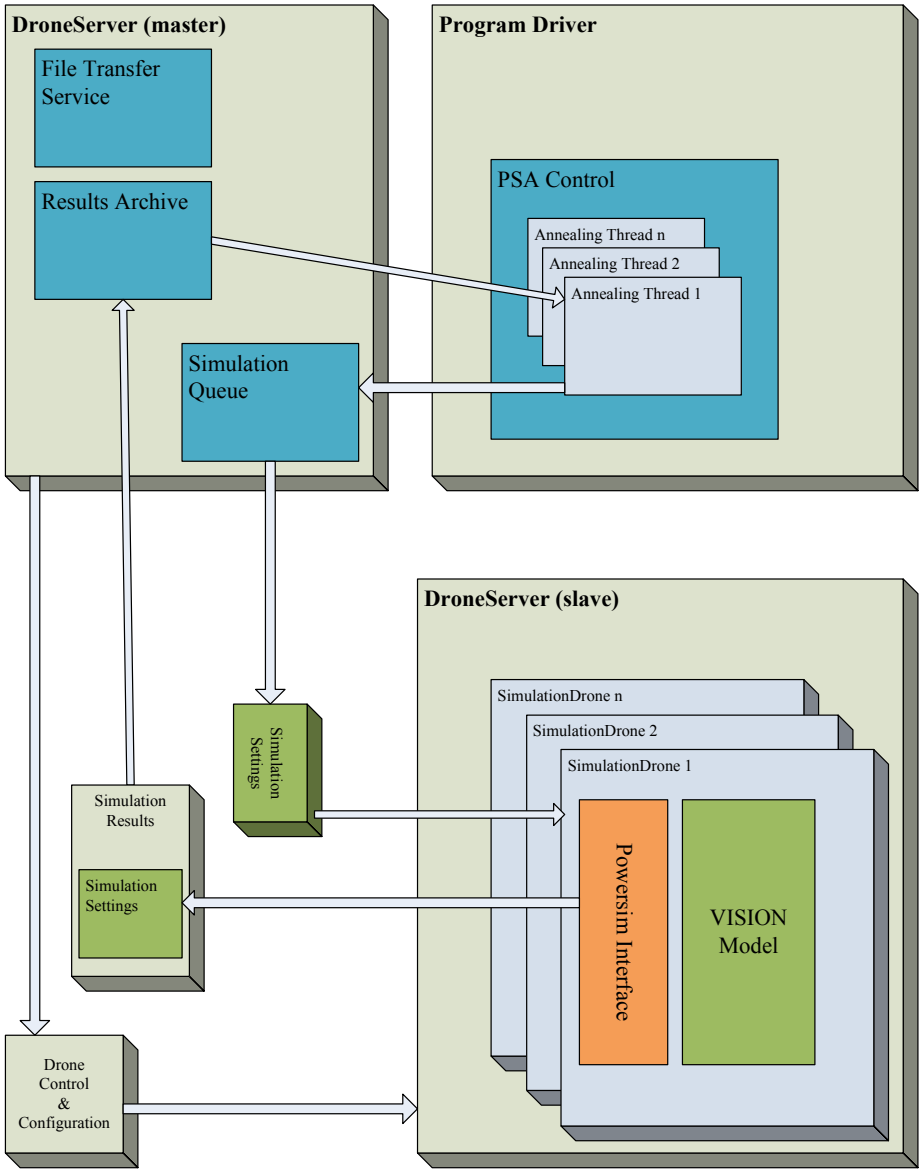

Figure 2.7: Parallel Simulated Annealing Framework Design Diagram

Chapter 3

Test Cases and Benchmarks

Prior to this work, most usage of the VISION model involved carefully planned and specified deployment scenarios, each designed to function within the domain over which the VISION model is valid. Although the spent fuel forecasting algorithm is specifically designed to prevent the inadvertent specification of the main class of infeasible scenarios, it is prudent to deliberately test the VISION model on these same situations in order to strengthen and verify the underlying computational logic.

Extensive modifications have been made to the VISION model to ensure that it functions according to the assumptions and requirements of the optimization methodology and to reduce its run-time. Among these modifications are the removal of all output-oriented spreadsheet connections, the streamlining of many computationally expensive calculations, the ability to override the built-in heuristic rules limiting the maximum number of fast reactors to order in a given year, and the ability to run simulations over an extended 200 year time horizon.

The goal of this chapter is to demonstrate that the upgraded VISION model and the optimization wrapper function properly and robustly. To this end, four sets of tests are documented. In the first set, the requested fraction of fast burner reactors is ramped up from zero to well beyond the level that can be supported by available LWRsf stocks. This demonstrates not only the ability of the forecasting algorithm to predict shortage conditions but also the ability of the VISION code to gracefully handle difficult deployment scenarios. The second set of tests defines a set of reference configurations which are used to provide benchmarks for demonstrating fuel and reactor aging within the model and for comparison with later optimization results. The final two sets of tests demonstrate the function and performance of the optimization algorithm itself, starting with a single objective, then moving on to multiple objectives.

3.1 Spent Fuel Forecasting

The spent fuel forecasting algorithm was developed to provide an inexpensive *a priori* estimate of recyclable fuel availability in order to detect and correct infeasible configurations without undergoing the computational expense of running the whole `VISION` model. The forecasting model is based on a greatly simplified view of the fuel cycle, it therefore should be tested to show that it is capable of performing its designated task. One simple way to do this is to incrementally increase the requested fraction of fast reactors from a known configuration until fuel shortages occur. The results of this test are shown in Figure 3.1. The top half of the plot shows the final forecast LWRsf inventory level, while the bottom half shows the net electric energy production. As can be seen, while the forecast inventory values are positive, the net energy production remains nearly constant (the variability is due to mismatch in output between the fast and thermal reactors). As the forecast LWRsf level initially drops below zero, some fast reactors run out of fuel and are idled (as indicated by the red band). Finally, as the forecast becomes increasingly negative, idled capacity remains nearly constant, while total energy output drops. This is due to requested fast reactors lacking in startup fuel, which prevents them from ever reaching operational status. That the LWRsf forecast line crosses below zero coincidentally with the first appearance of idled fast reactor capacity gives strong indication that the LWRsf forecasting algorithm is sufficiently accurate for the present needs.

This test also provides the opportunity to examine the way in which the various objective function values change in response to an increase in fast reactor construction. The values of each objective, normalized by their respective All-LWR configuration values are plotted below in Figure 3.2. As is expected, costs and nonproliferation metrics increase with increasing FBR construction, while long-term heat and uranium consumption decrease. It is interesting to note that the weighted-weapons usable value shows a much greater increase than the Bathke nonproliferation metric; as is shown in the next section, this is due to the much greater emphasis in the former placed on separated material.

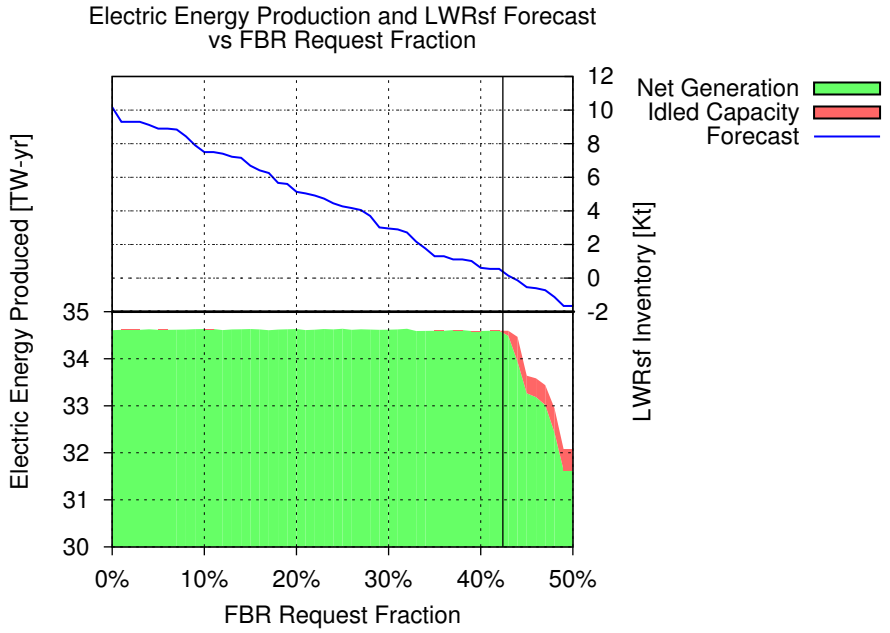

Figure 3.1: The LWRsf forecast is used to provide *a priori* estimates of whether a given configuration can be supported. This plot shows the accuracy of the method.

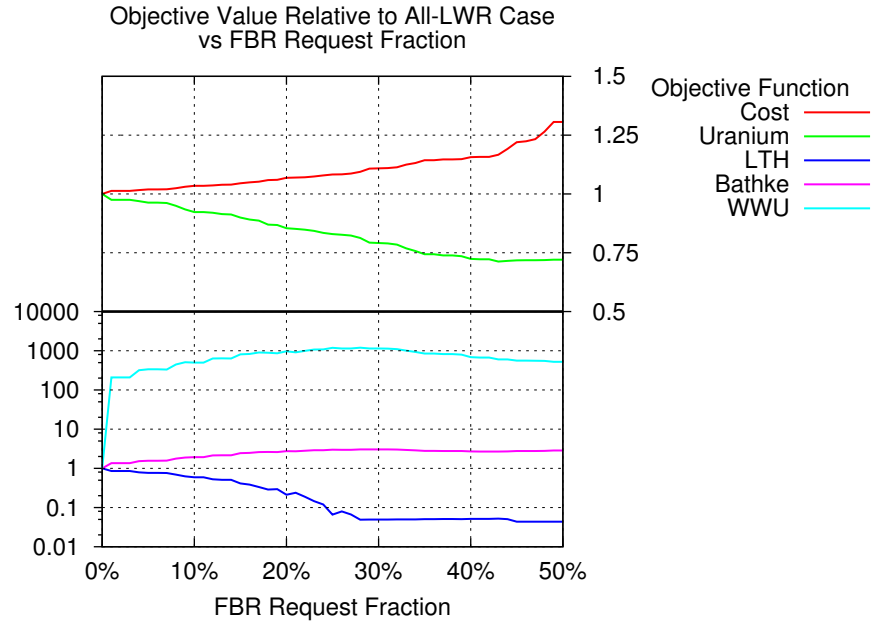

Figure 3.2: Increasing the fast reactor fraction beyond that which can be supported leads to energy shortfalls which adversely affect most objective function values.

3.2 Reference Configurations

Although the base-case specifications set out the fraction (or number) of each type of reactor to build on a year-by-year basis, it is useful to define an additional set of reference configurations for testing and initialization purposes. Each configuration is specified by means of a simple set of portable rules, rather than through a more cumbersome list of values. Furthermore, each configuration, as described below, serves a specific purpose in the testing of the `VISION` model and the optimization routines.

All LWR This case shows what might occur if nuclear power continues to meet the same fraction of total electric demand while utilizing only the current once-through fuel cycle. The inherent simplicity of the fuel cycle provides a good first step for verifying the output from `VISION` and the LWR spent fuel production terms in the forecasting methodology.

FBR Impulse The isotopics and makeup fuel ratios for fast reactors evolve as the reactors age and continue to recycle their own wastes. The FBR Impulse case substitutes LWRs for a single year (2095) with FBRs. All the fast burner reactors thus progress through the reactor aging chain together, so the process of isolating any problems that arise in the fuel recycling process is greatly simplified. Additionally, as the fuel usage in each fast reactor will be identical, this provides a straightforward benchmark for the FBR fuel usage terms of the forecasting methodology.

Rear Loaded FBR Although the simulated annealing algorithm is generally considered to be robust and capable of being started from any valid configuration, it is expected that it will converge more quickly if it is started in what we expect is a near optimum configuration. To this end, the Rear-Loaded FBR configuration is specified by progressively substituting FBRs in place of LWRs, starting at the end of the planning horizon, until just before the spent fuel forecast indicates a shortage will occur. This minimizes LWR spent fuel carryover and is consistent with the socially-conscious decision variable perturbation approach. With the current base case settings, this works out to be all LWRs before 2084 and all FBRs from 2086 to 2100 (with a mix of reactors in 2085).

Constant Request The Constant-Request configuration is similar to the Rear-Loaded FBR configuration in that it estimates a near-optimum configuration by utilizing the LWR spent-fuel forecast. It differs in that for the years 2040 through 2100 a fixed fraction of new capacity is met by fast reactors. This fraction is incrementally increased from zero until just before a shortfall is forecast (42.4%, currently).

For comparison, Figures 3.4 and 3.5 show the rate at which LWRs and FBRs (respectively) are built for each of the four base configurations. Figure 3.6 depicts the contributions of each

reactor type to the overall electric power output. Figure 3.3 then compares the forecast LWRsf TRU inventories for each of the configurations. It should be noted that the power output of each LWR, at 0.945 GW(e), is nearly twice that of each FBR, which generates 0.51 GW(e).

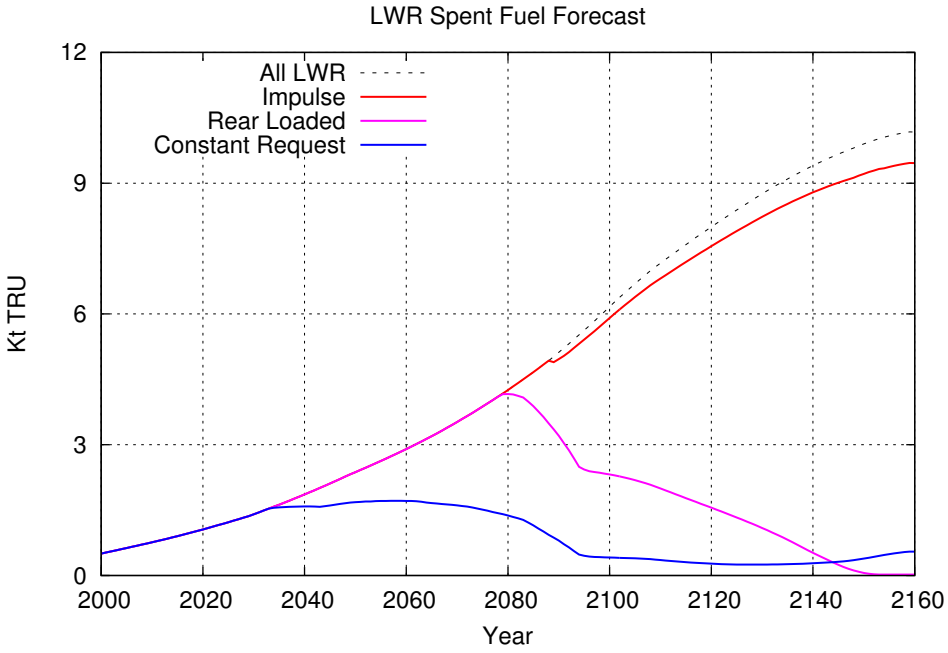

Figure 3.3: The forecast mass of transuranic material available in LWR spent fuel for reuse in fast burner reactors.

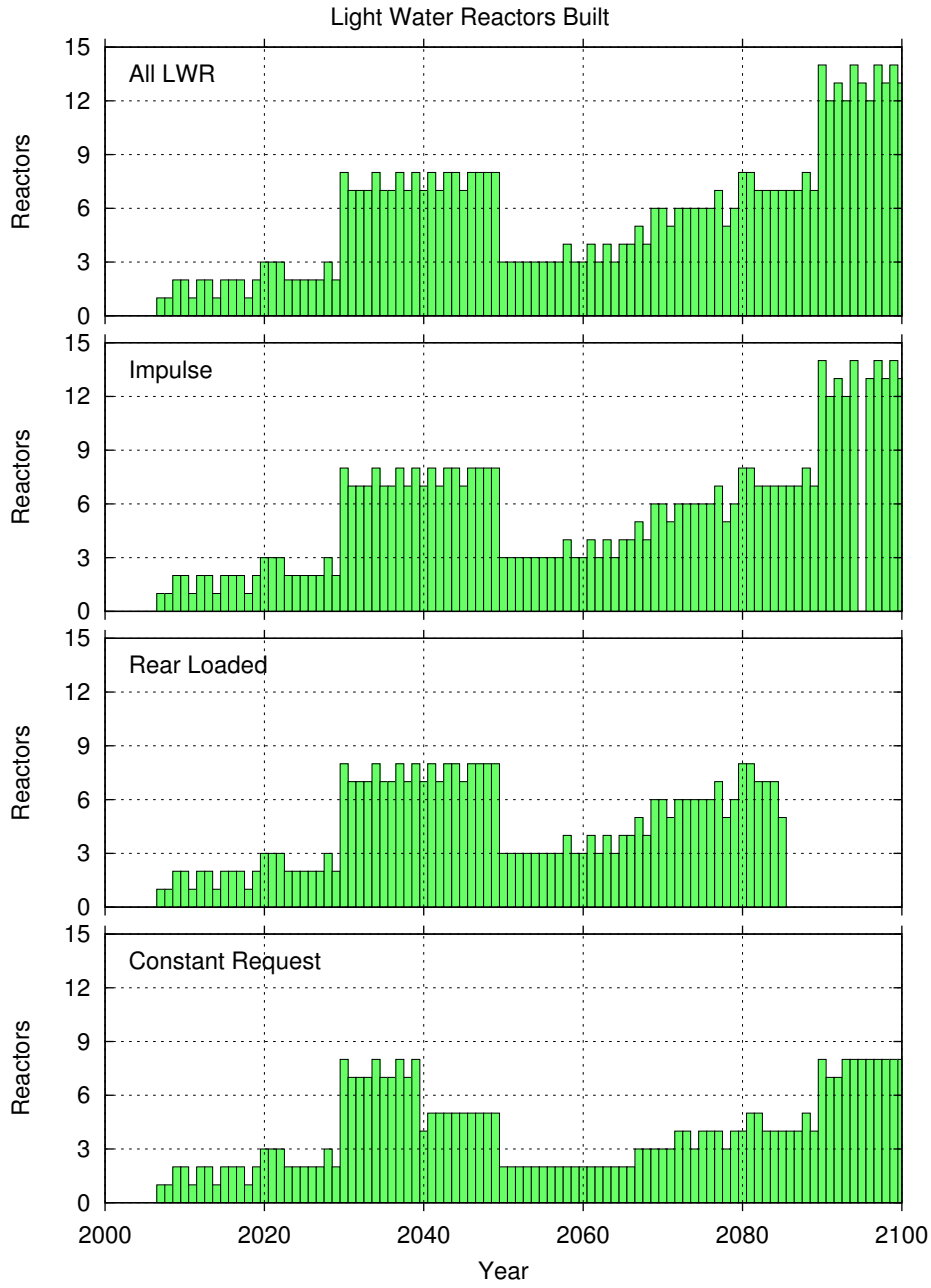

Figure 3.4: These plots show the number of light water reactors built in a given year for each of the four reference configurations. The increase in construction between the years 2030 and 2050 is due to the retirement of the legacy reactor fleet (see Figure 3.6 for details.)

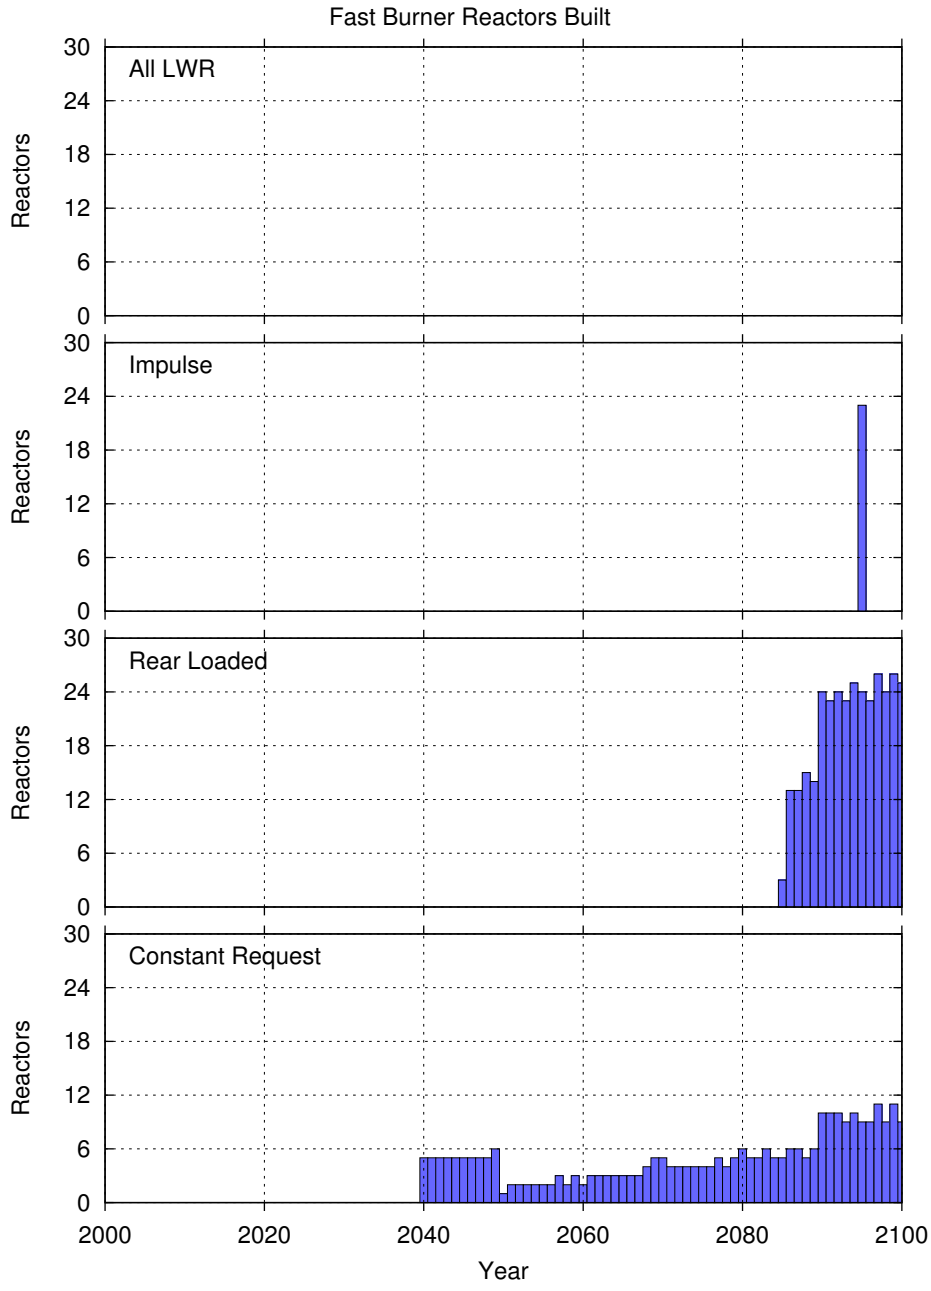

Figure 3.5: These plots show the number of fast burner reactors built in a given year for each of the four reference configurations.

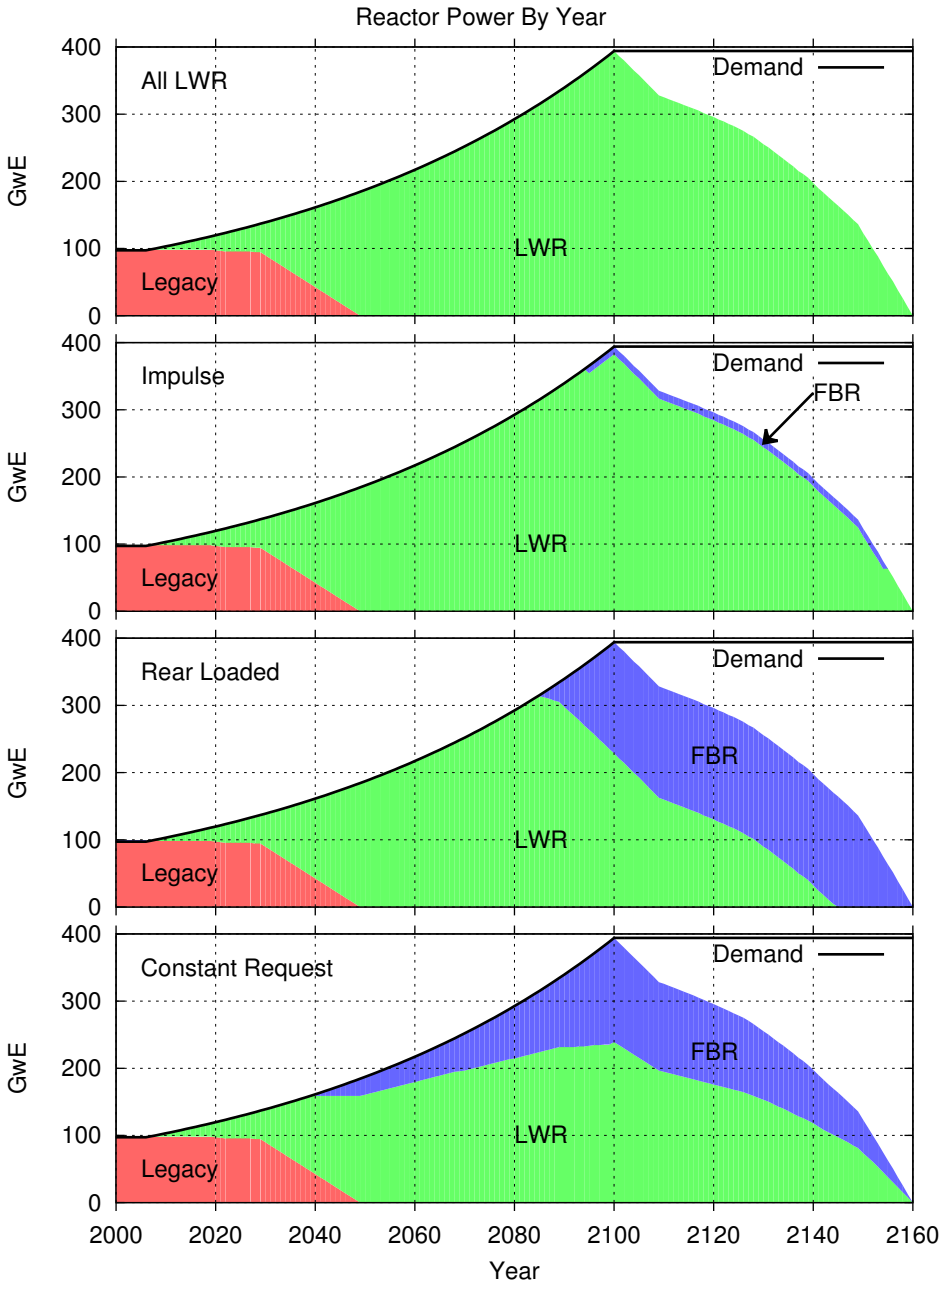

Figure 3.6: These plots show the power produced annually by each reactor type for each of the four reference configurations.

3.2.1 Long-term Heat

The long-term heat objective function, like the other objective functions, takes an extensive fuel cycle quantity (in this case, decay heat) and normalizes it by the net electric energy production. Because all cases considered here have the same demand growth rate this normalization factor plays only a minor role; however, it will become active if fuel shortages cause reactors to shut down or if the demand growth parameters are varied. The long-term heat objective values for the four reference configurations are shown below in Figure 3.7. The `VISION` model currently uses two very simple algorithms for building and operating separations facilities; namely, it builds sufficient capacity in fixed increments to meet the peak demand rate required by the reactors, and then continues operating the separations facilities at full capacity as long as separable spent fuel remains available. In situations which have few FBRs (such as the impulse case), this leads to a considerable accumulation of unused separated material, and a corresponding diminution in the amount of long-term heat going into the repository. Given the considerable expense associated with both the operation of the separations facility and the secure separated material storage facility, this represents an unlikely plan of action. The revision of these algorithms within the `VISION` model would be a useful task for future work.

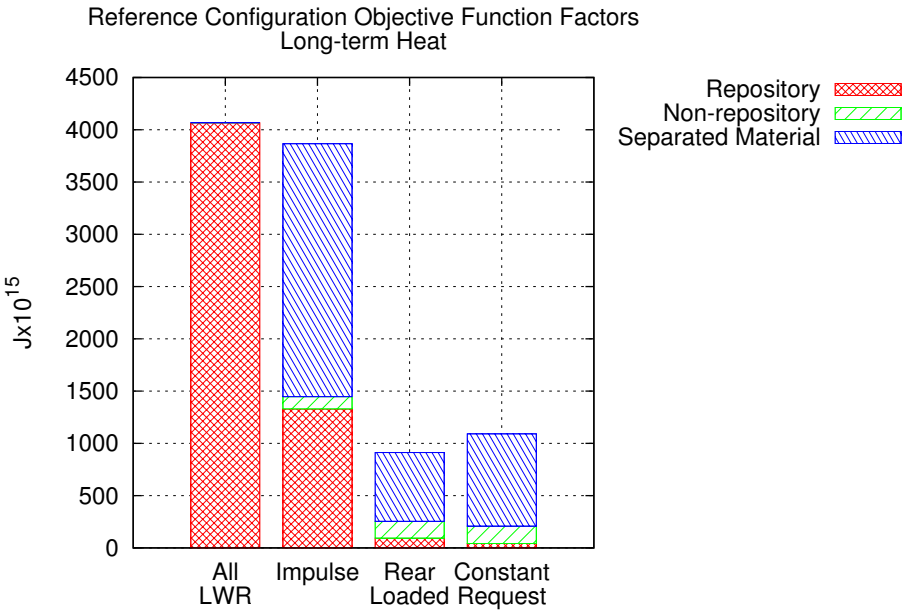

Figure 3.7: Reference Configuration Long-term Heat Loading. Note that separated material is not considered to be a waste product, and therefore does not contribute to the LTH objective function value.

3.2.2 Uranium Utilization

In contrast to the long-term heat objective, the uranium utilization objective, plotted below in Figure 3.8, shows relatively little sensitivity to the presence of fast reactors. This is largely due to the number of LWRs built early in the planning horizon (before FBR technology is available), and built later as required to support the fast reactor fleet. These values depend, of course, on the choice of transuranic conversion ratio for the FBR fleet; an examination of these effects is undertaken in Section 4.1.

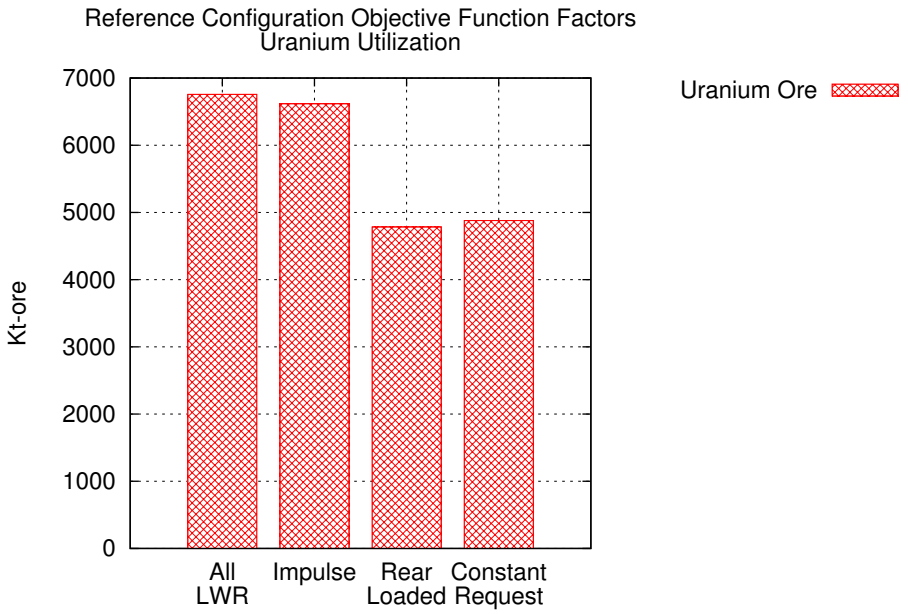

Figure 3.8: Reference Configuration Uranium Utilization

3.2.3 Nonproliferation Metrics

Figure 3.9 shows the relative non-proliferation risks presented by the four reference configurations according to both the weighted-weapons usable and the Bathke metrics. As expected, the WWU metric depends predominantly on the amount of material separated for re-fabrication into fast reactor fuel. Somewhat counterintuitively, the impulse case shows a greater risk than either of the full-FBR configurations. This is likely due to the previously noted separations facility ordering and operating algorithms within the `VISION` model. This separated material then persists through the end of the simulation, whereas in the full-FBR configurations it is

eventually consumed. Similarly, the constant request configuration shows a slightly higher proliferation risk than the rear loaded configuration. By starting fast reactor operations earlier in the planning horizon, the constant request configuration has a longer exposure of the relatively risky separations and recycled-uranium streams than does the rear loaded configuration.

While the categorical breakdown of the two nonproliferation metrics shown in Figure 3.9 indicates that the two metrics emphasize the risks of different factors, the relative proportions of these factors remains somewhat stable across a wide range of fuel cycle scenarios. Thus, there is a clear correlation between the overall values of the two metrics. Figure 3.10 below shows a visual comparison of the objective values calculated for both methods for over 10,000 sampled configurations.

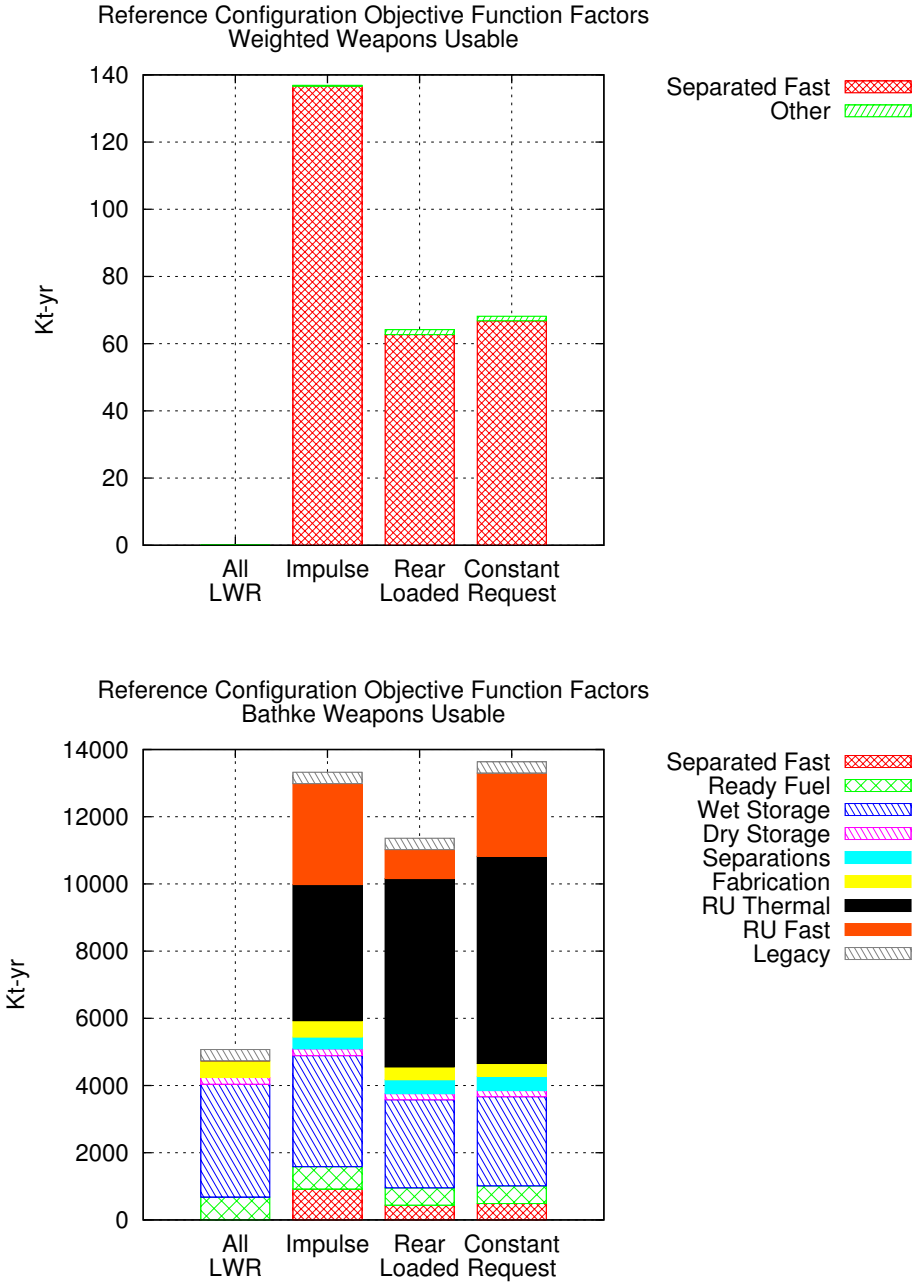

Figure 3.9: Reference Configuration Nonproliferation Metrics

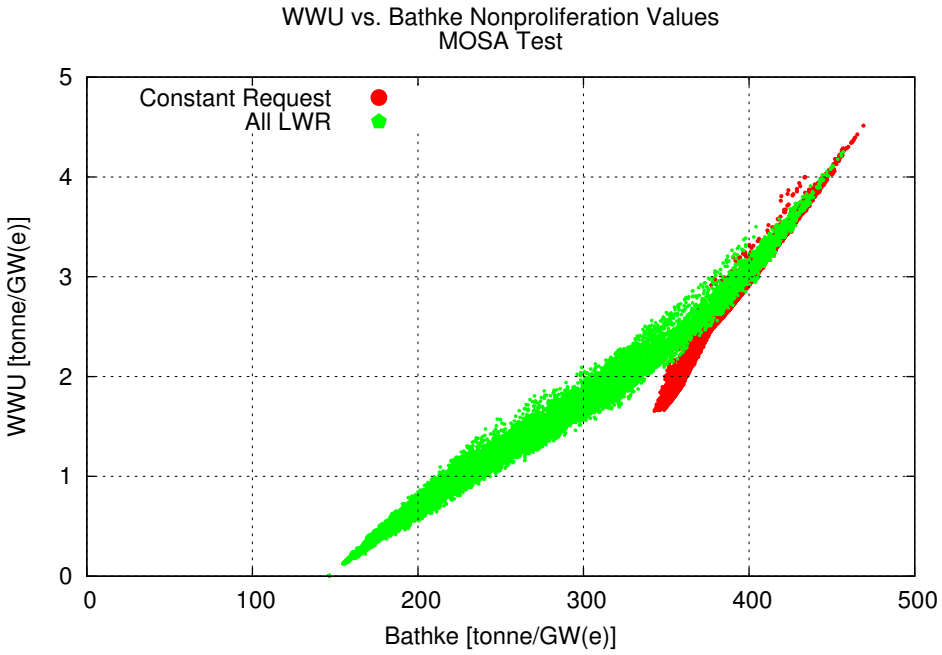

Figure 3.10: The overall value of Bathke and WWU nonproliferation metrics are highly correlated for the sampled configurations.

3.2.4 Economics Metrics

It was discussed in the previous chapter that a great deal of uncertainty exists around the estimation of all manner of future costs. This carries forward into the cost estimates for the reference configurations. Figure 3.11 shows that the cost probability range for any single configuration encompasses nearly a factor of three difference. As expected, the greatest difference in cost is between those configurations with many FBRs and those with few.

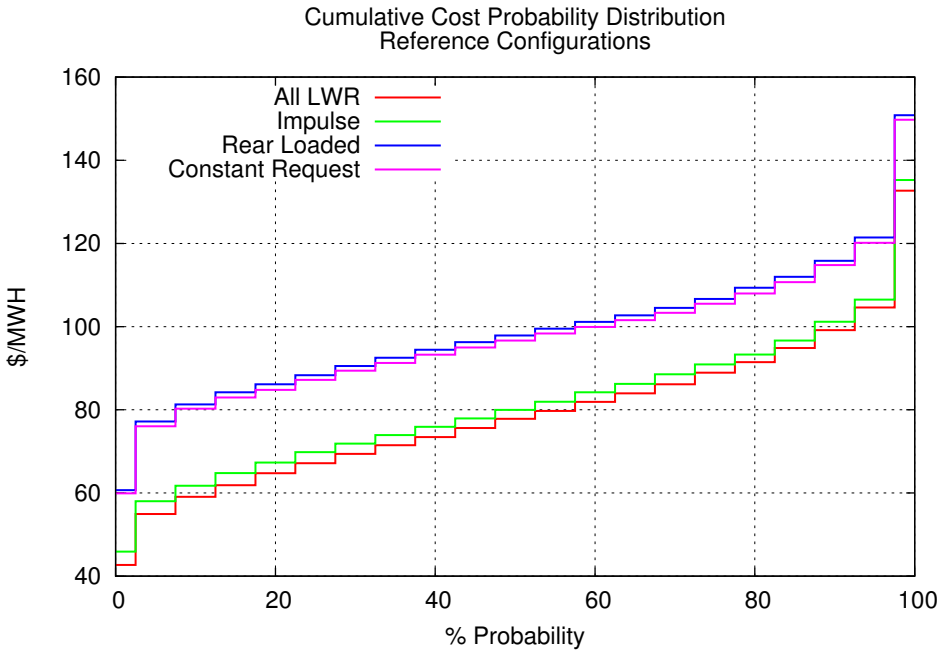

Figure 3.11: Reference Configuration Cost Probability Distribution

Drilling down category by category, it is found that, as expected, reactor capital is the dominant cost factor. Because of the multitude of cost factors and differences in scale involved, the cost factors are divided among Figures 3.12 through 3.14.

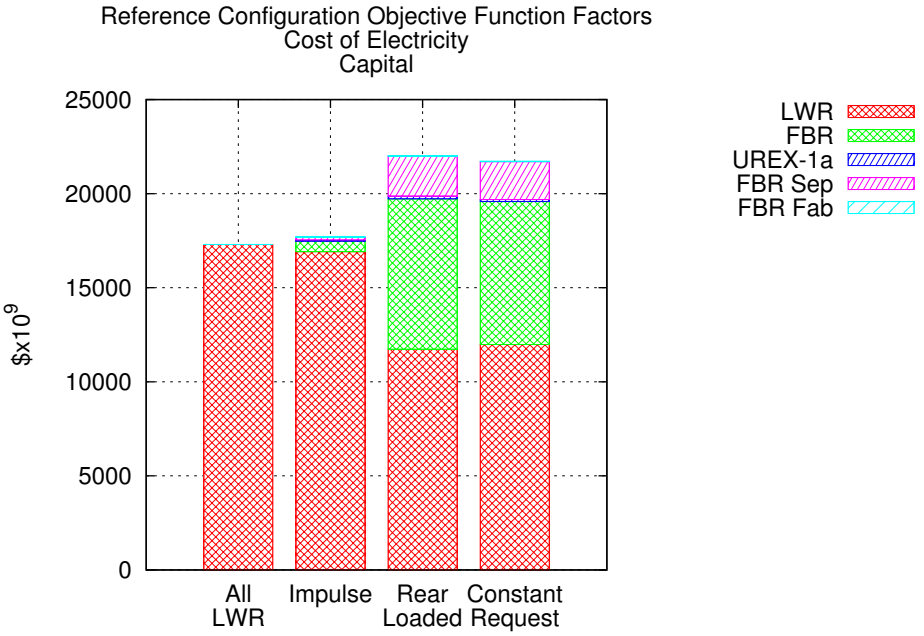

Figure 3.12: Reference Configuration Capital Cost-of-Electricity

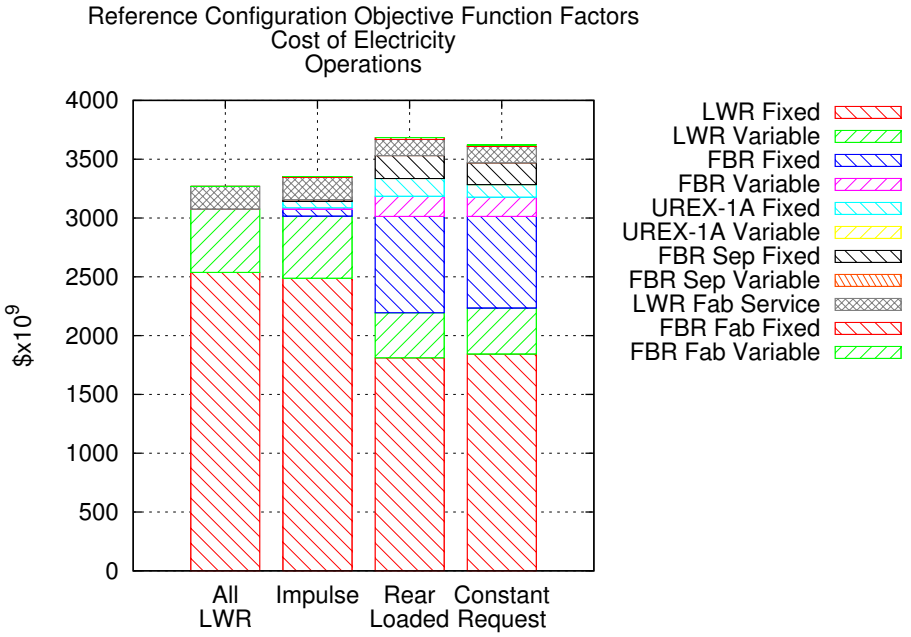

Figure 3.13: Reference Configuration Cost-of-Electricity Operations & Maintenance

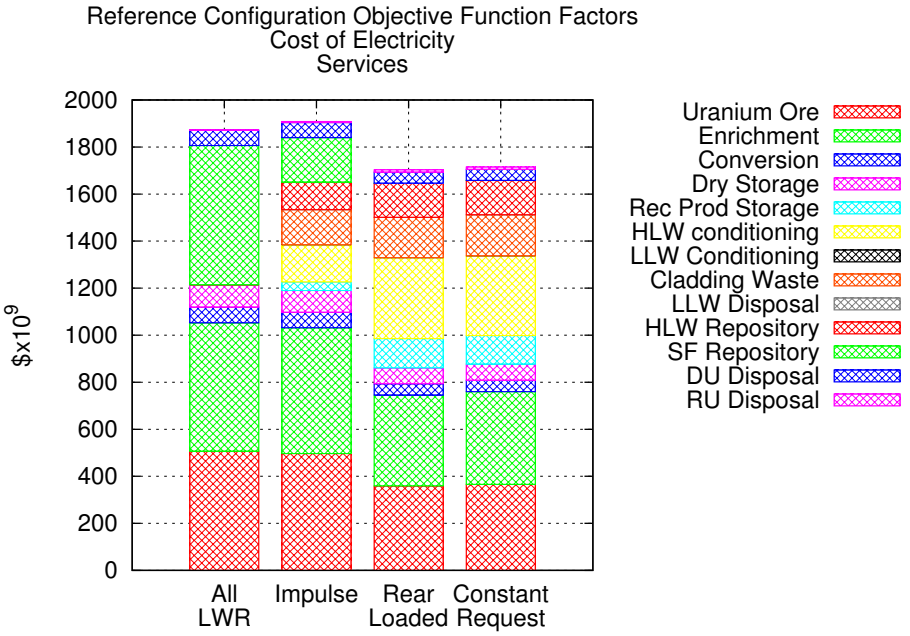

Figure 3.14: Reference Configuration Cost-of-Electricity Services

3.3 Optimization Benchmarks

Two sample optimization tests are shown here to demonstrate the capability of the code. In the first case, the socially conscious perturbation approach is used to optimize on a single long-term heat objective function. In the second case, the dynamic-growth perturbation approach and multi-objective optimization algorithm are used to explore the trade-offs between the uranium utilization, electric cost, and Bathke nonproliferation objectives. A key strength of the simulated annealing algorithm is the ability to converge to a similar set of solutions with little regard to the specified initial solution. To demonstrate this capability, each test case was performed twice, each time with a different initial configuration. As is shown, though the rate of convergence may be affected, the final set of solutions is quite similar.

Finally it must be noted that shortly after these calculations were performed, an error was discovered in the `VISION` input data, whereby the specified fast reactor discharge burnup exceeded the correct value by 34%. This erroneously high burnup value, when processed through the `VISION` model, causes a corresponding decrease in the FBR fuel consumption rate, thereby increasing the maximum number of fast reactors that may be built. Although this error means that the optimized objective values, scenario configurations and trade-off surfaces computed in these two test cases are not correct in an absolute sense, they still have value as a demonstration of the optimization methodology and implementation. This caveat applies to the results in this section only.

3.3.1 Single Objective

The two tests below test the robustness of the simulated annealing algorithm. Each applies an identical cooling schedule to a different initial configuration in order to see whether the optimization tends toward a similar result. Though the final optimized configurations found by each run are slightly different, both test runs were kept somewhat short in order to minimize the overall computational turnaround time. As a result, acceptance ratios at the end of each cooling cycle were still quite high, so the final configurations may not have completely converged on a solution. It is expected that by extending the optimization the final optimized configurations will approach each other more closely. The optimization settings utilized for these two runs are shown below in Table 3.1, while pertinent optimization results are included in Table 3.2.

The annealing temperatures for each cooling step are shown in Figure 3.15. For both tests, the annealing temperature decrement ratios were constrained to the minimum value of 0.75 at each update step. The higher initial temperature for the constant request run was due to a greater variance in the objective function value during the initialization phase of the

Table 3.1: Test Settings

Objective	Long-term Heat
Constraints	none
$N_{archive}$	35
N_{thread}	16
L_{lngth}	8000
L_{chain}	75
L_{tran}	40
L_{surv}	50
χ_p	0.75
λ_H	0.70
λ_c	0.70

Table 3.2: Test Results

Initial Configuration	Rear Loaded	Constant Request
Annealing Segments	12	12
Total Samples	10578	10184
Accepted Samples	8238	8103
Rejected Samples	2340	2081

optimization algorithm. Comparing the sizable decrease in the annealing temperature to the more modest decrease in the sample acceptance ratio (Figure 3.16) indicates that the choice of initial temperature and minimum temperature decrement may be unnecessarily high. While this would provide some margin against trapping in local minima and allow for greater coverage of the configuration space, it also prolongs the optimization cycle and increases run times.

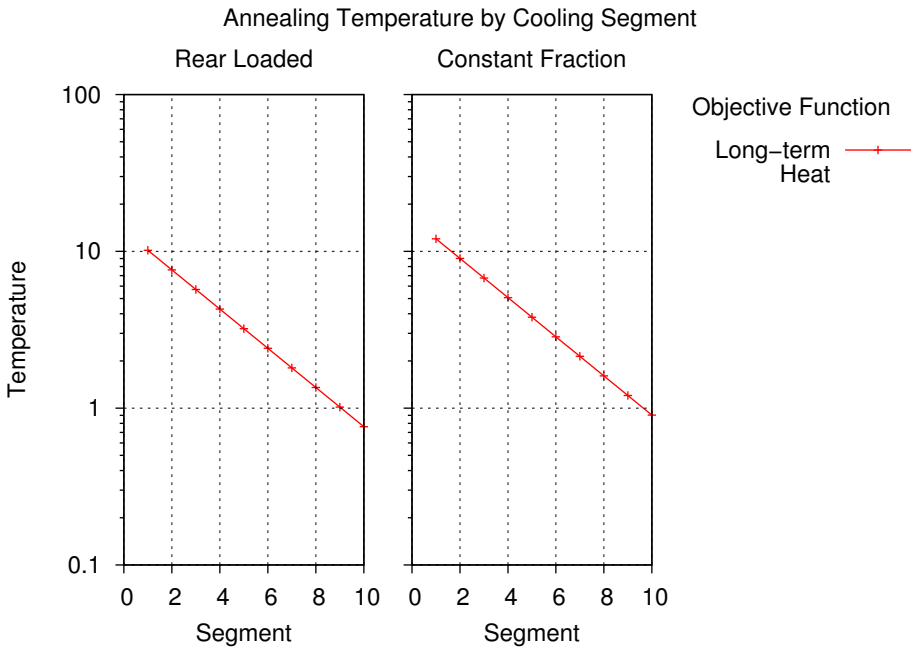

Figure 3.15: Single objective test annealing temperature by cooling segment

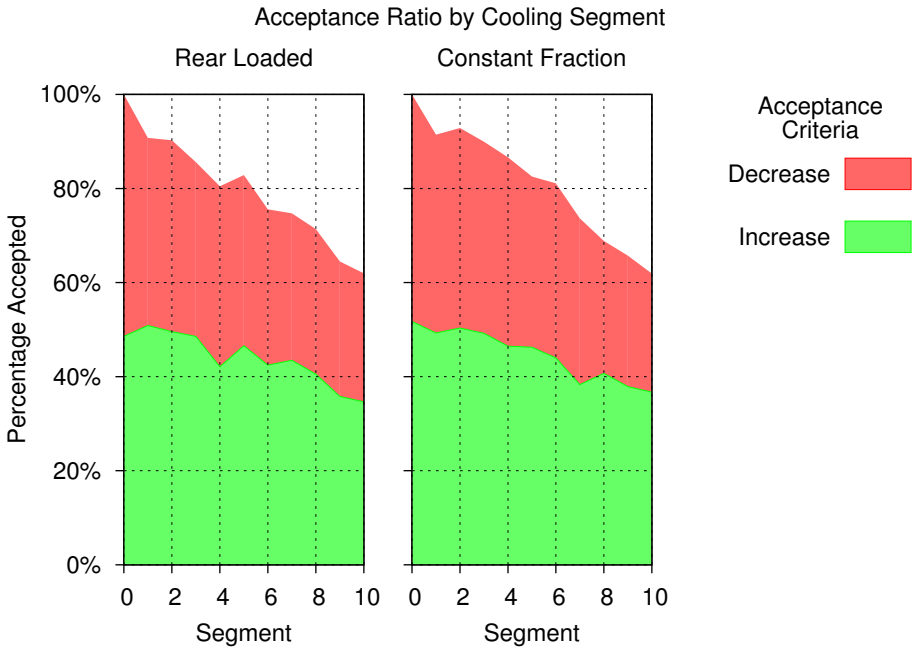

Figure 3.16: Single objective test acceptance ratio by cooling segment

Please note that in many of the plots that follow, the values are plotted against the so-called *adjusted sample number*. The adjusted sample number is computed by adding the maximum adjusted sample number (across all processes) from the previous cooling step to the number of the sample relative to the current cooling step. This ensures that when multiple parallel processes are plotted on the same axis the cooling step increments are all aligned together.

The reactor capacity displaced by the perturbation algorithm is sampled uniformly within a given range (in this case, 1 to 15 individual perturbations). Thus, it is possible to transit the configuration space in large steps while still retaining the ability to make small movements within a close neighborhood of a minimum. Therefore, it is expected that as the cooling cycle progresses, the amount of reactor capacity swapped from one reactor type to another (the so-called *Perturbation Distance*) by each accepted perturbation will decrease. However, as is evident in Figure 3.17, this was not observed in the rear loaded run (similar results for the constant request run are omitted). This is most likely due to the short length of the runs and the relatively high annealing temperatures.

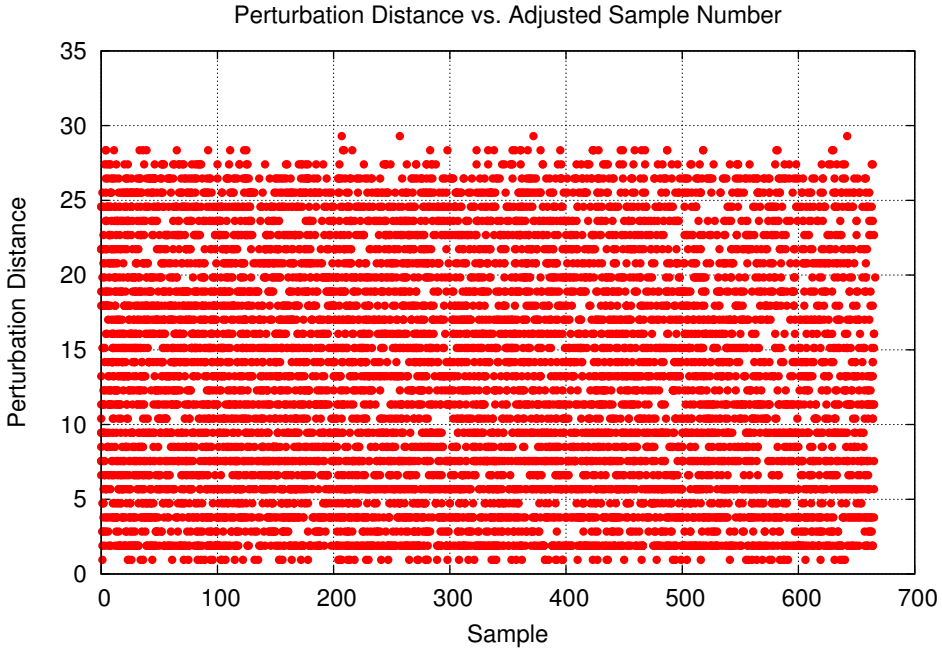

Figure 3.17: Single objective test rear loaded configuration perturbation size by sample number

The socially conscious perturbation method works to maximize the total number of fast reactors built, thereby minimizing the amount of LWR spent fuel remaining at the end of the

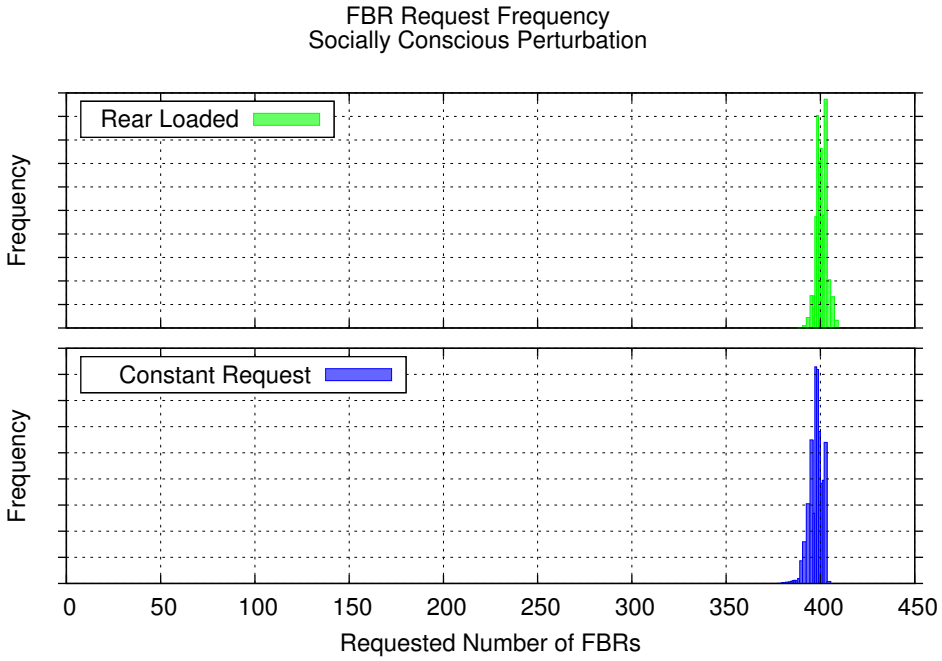

Figure 3.18: The socially conscious perturbation algorithm attempts to maximize the number of FBRs built in order to minimize LWRsf carryover.

simulation. Thus, as expected, there is very little variation to be seen in the total number of fast reactors built in each of the sampled configurations (Figure 3.18).

The minimization of long-term heat was the objective for these optimization runs, and as expected, the results show a general downward and convergent trend (see Figures 3.19, 3.20, and 3.21). Although during the temperature initialization segment the rear loaded run exhibited higher long-term heat values than the constant request run, those values were more closely distributed. This leads to the lower initial temperature value seen for the rear loaded run. The relatively small difference between the objective values of the best and worst archived configurations indicates that the objective function surface is relatively flat near the vicinity of the archived solutions.

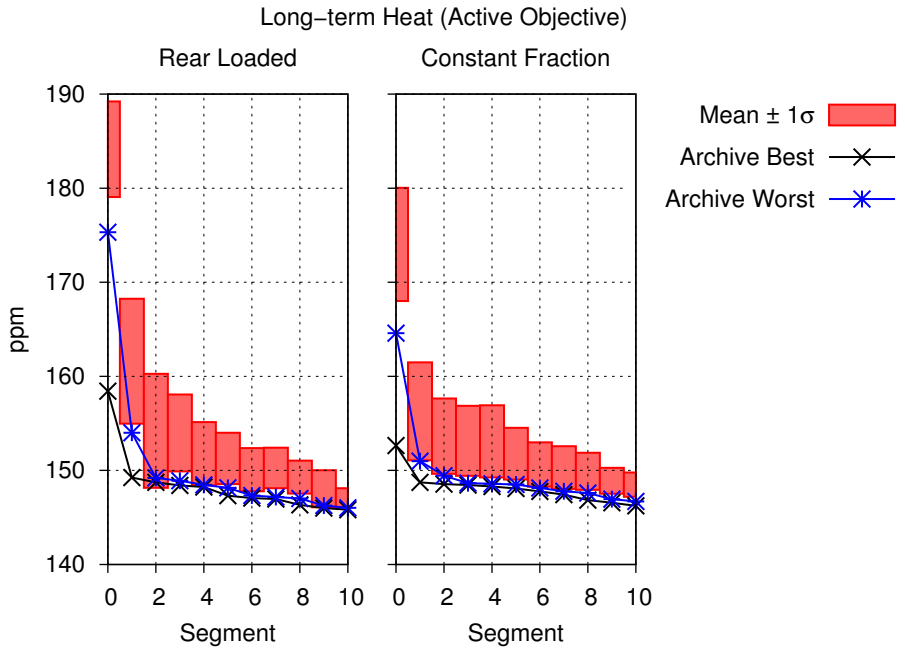

Figure 3.19: These plots show, for each annealing segment, the best and worst archived solution long-term heat objective values, and the accepted-sample mean LTH objective value distribution.

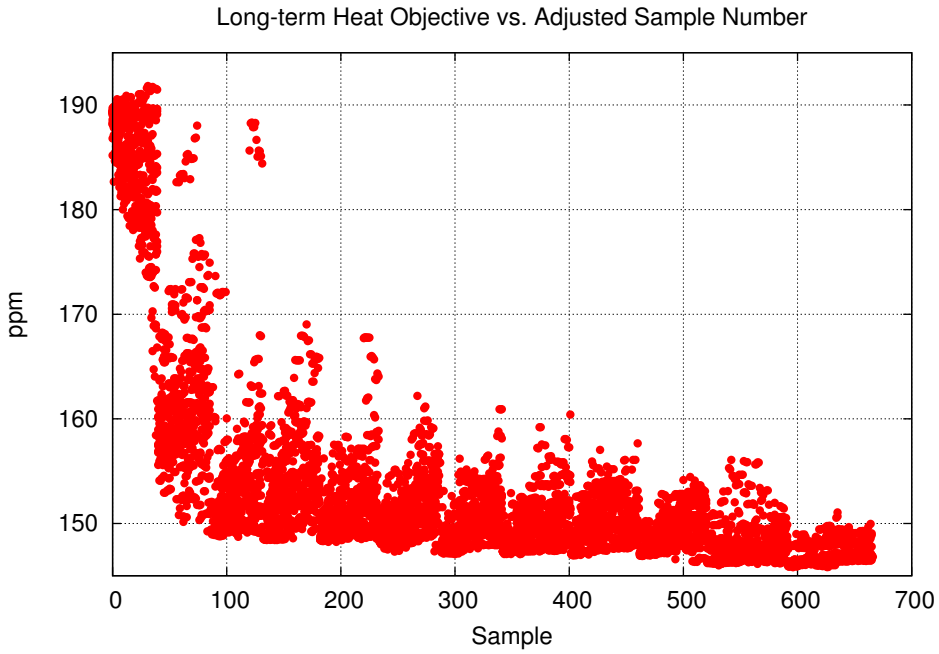

Figure 3.20: Single objective test rear loaded initial configuration accepted solution long-term heat objective values.

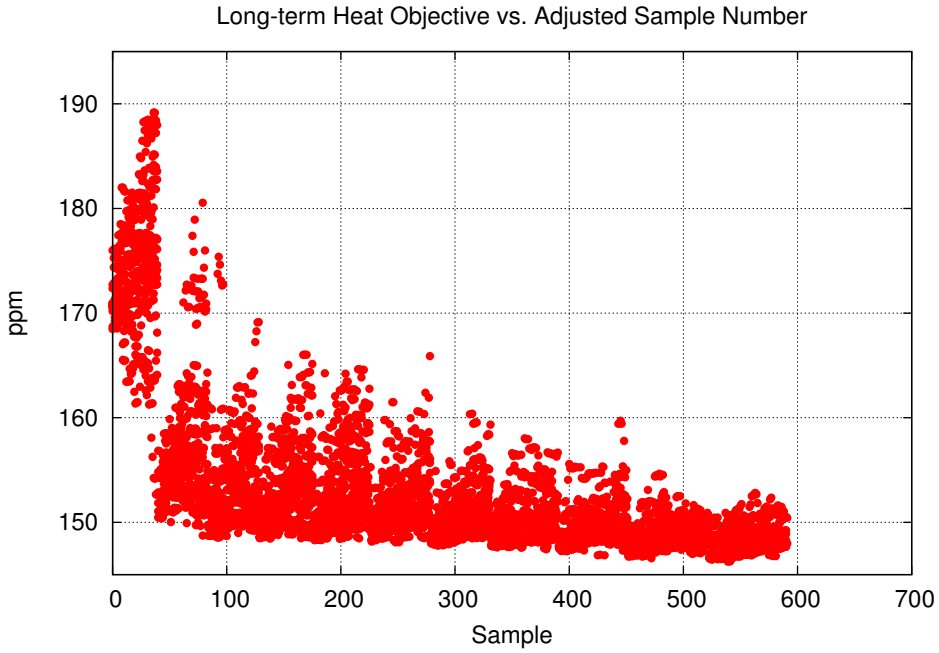

Figure 3.21: Single objective test constant request initial configuration accepted solution long-term heat objective values.

Fast burner reactors rely on recycled light water reactor fuel to supply their needs. A significant fraction of the lifetime LWR spent fuel production occurs at end-of-life, when the entire reactor core is discharged according to the VISION model. An alternative scenario would be for the partially burnt fuel from the shutdown LWR to be shipped and loaded into an LWR that is continuing operation. Likewise, the demand for fuel in an FBR is greatest at the beginning-of-life, when the first core is fabricated entirely from LWR spent fuel. Factor in the cooling, reprocessing and fabrication steps required in between, and it becomes apparent that in order to maximize the number of fast reactors that can be built, one must build them as late as possible. In other words, if the FBRs are to fully utilize the LWRsf that will be produced, they must continue to operate after all of the LWRs have shut down. In the rear loaded initial configuration run, all of the fast reactors are built at the end of the planning horizon, whereas in the constant request initial configuration run, they are spaced equally throughout; this gives rise to the results below in Figures 3.22, 3.23, and 3.24. It is interesting to compare Figures 3.19 and 3.22 and note that in several instances, counter to expectation, the worst archived solution (greater long-term heat) will exhibit a better (lower) uranium utilization value. These values typically move together, both decreasing as greater numbers of fast reactors are built. These seemingly reversed situations are likely attributable to changes in the overall energy production, which can vary slightly from scenario to scenario depending upon how much the deployed power output exceeds the demand requirement (which is itself due to reactor sizing). Additionally, the difference between the best and worst archived objective function values is quite small.

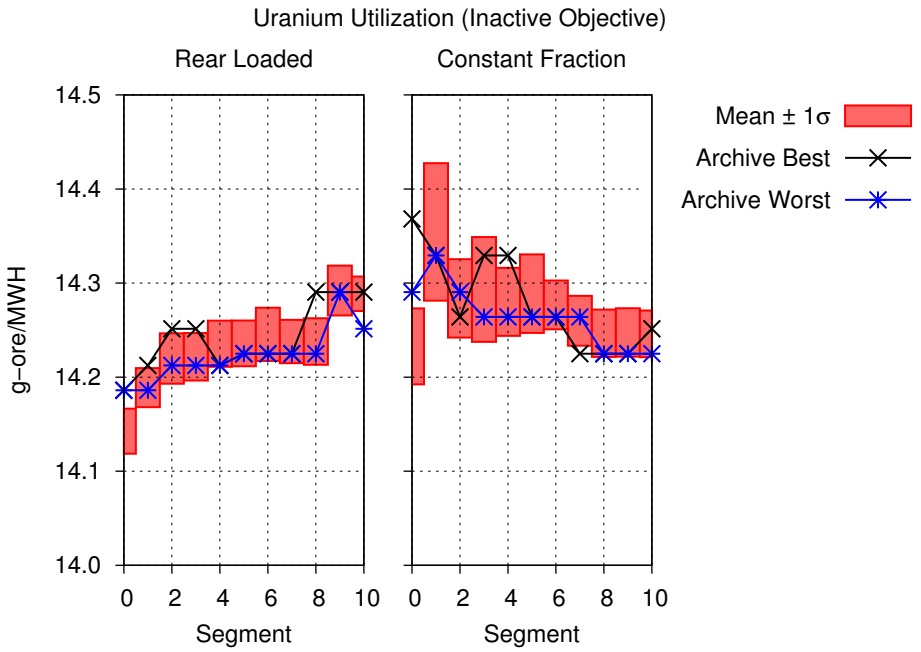

Figure 3.22: These plots show, for each annealing segment, the best and worst archived solution uranium utilization objective values, and the accepted-sample mean objective value distribution.

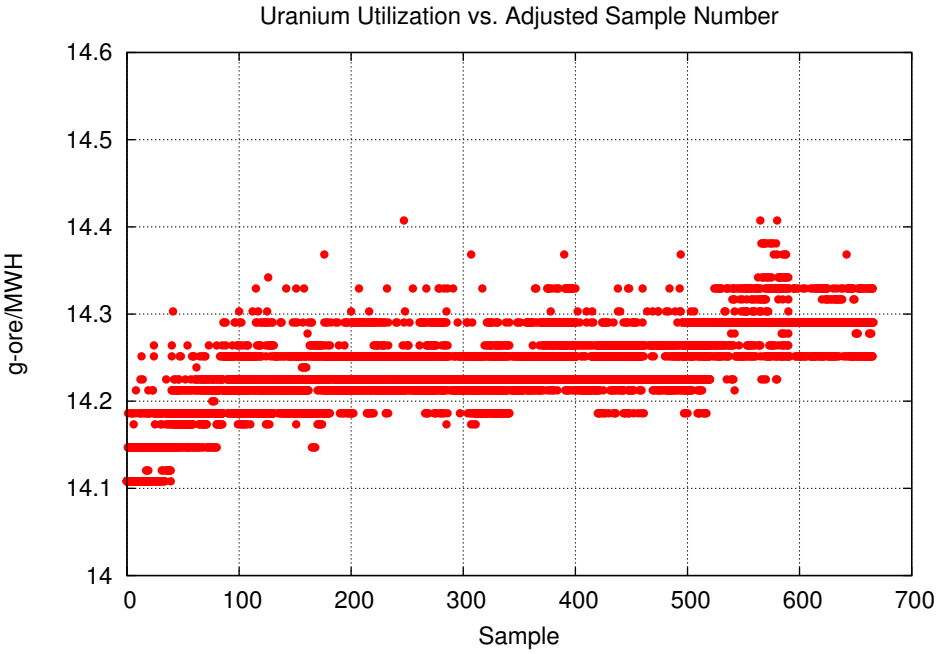

Figure 3.23: Single objective test rear loaded initial configuration accepted solution uranium utilization values. The observed banding structure is due to variations in the total number of fast and thermal reactors ordered.

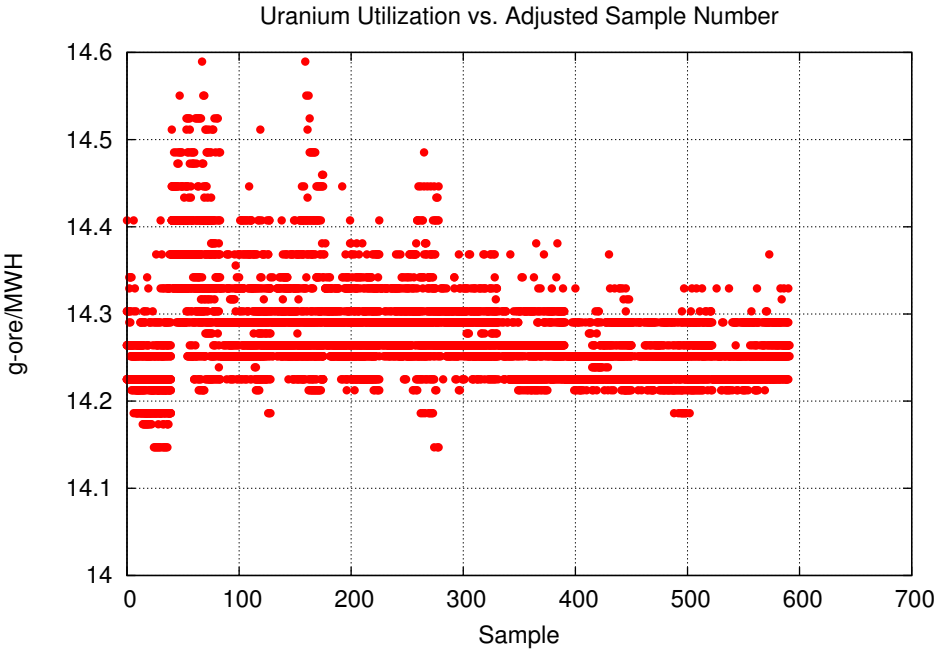

Figure 3.24: Single objective test constant request initial configuration accepted solution uranium utilization values.

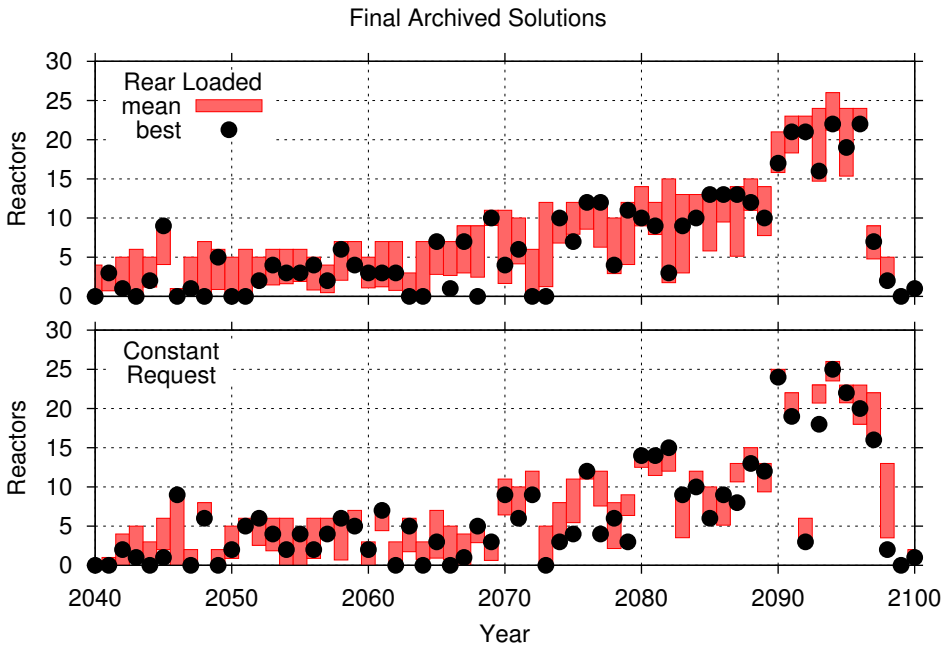

Figure 3.25: Single objective test case best solution archive.

Finally, a comparison of the configurations contained in the best-solution archive at the end of the two optimization tests shows that despite the difference in starting configuration between the two tests, they are converging toward a similar set of optimum configurations. Figure 3.25 shows the year-by-year number of fast reactors requested by the best archived solution as solid points, and the average number of reactors requested by year which are greater and lesser than the overall archive mean for that year (with the range depicted by the solid bar). The year-to-year variation makes it difficult to directly compare the similarity of the archived solutions; Figure 3.26 instead plots a five-year moving-average fast reactor construction rate for the best archived solution from each optimization test. By averaging the fast reactor build rates, the apparent similarity between the results becomes much more evident.

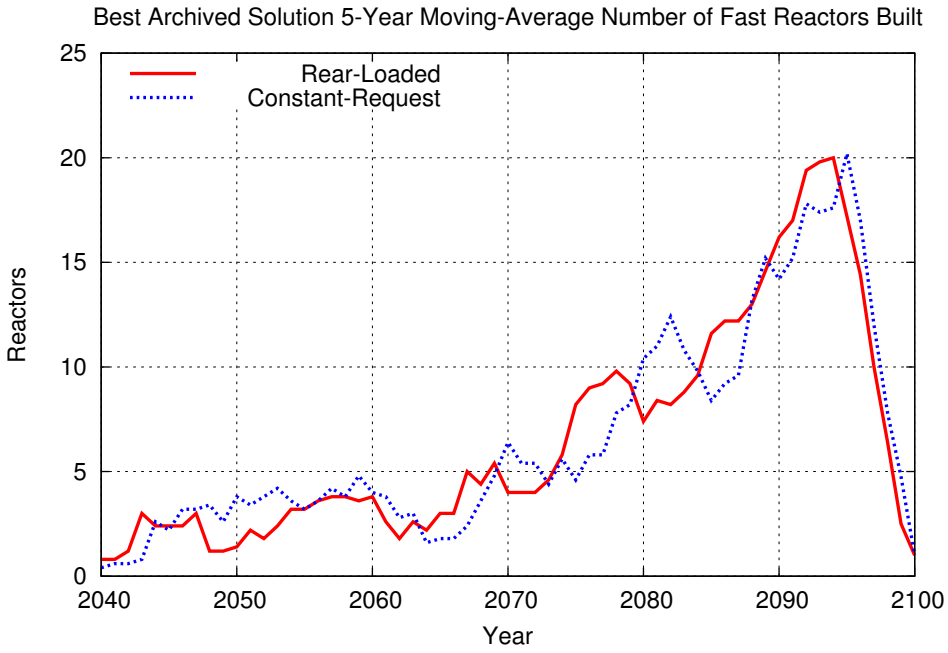

Figure 3.26: Five-year averaged FBR build rate for the single best solution found during each run of the single objective test case.

3.3.2 Multiple Objective

To test the full capabilities of the optimization code a set of multi-objective optimizations was performed utilizing the dynamic growth perturbation method. These tests develop the two-dimensional trade-off surface between the uranium utilization, electric cost, and Bathke non-proliferation objectives. The first test case starts with the constant request configuration, while the second starts with the vastly different All-LWR configuration. The test case settings are listed below in Table 3.3, and the resulting values in Table 3.4.

Table 3.3: Multiple objective test settings

Objectives	Uranium Utilization
	Cost of Electricity
	Bathke Non-proliferation
Constraints	Smoothness
$N_{archive}$	N/A
N_{thread}	16
L_{lngth}	15,000
L_{chain}	100
L_{tran}	60
L_{surv}	60
λ_H	0.70
λ_c	0.70

Table 3.4: Multiple objective test results

Initial Configuration	Constant Request	All-LWR
Annealing Segments	11	11
Total Samples	18,304	16,772
Accepted Samples	11,363	8,964
Rejected Samples	6,901	7,808
Net Archived Samples	2,858	3,128

Though the individual objective temperatures plotted below in Figure 3.27 show a good

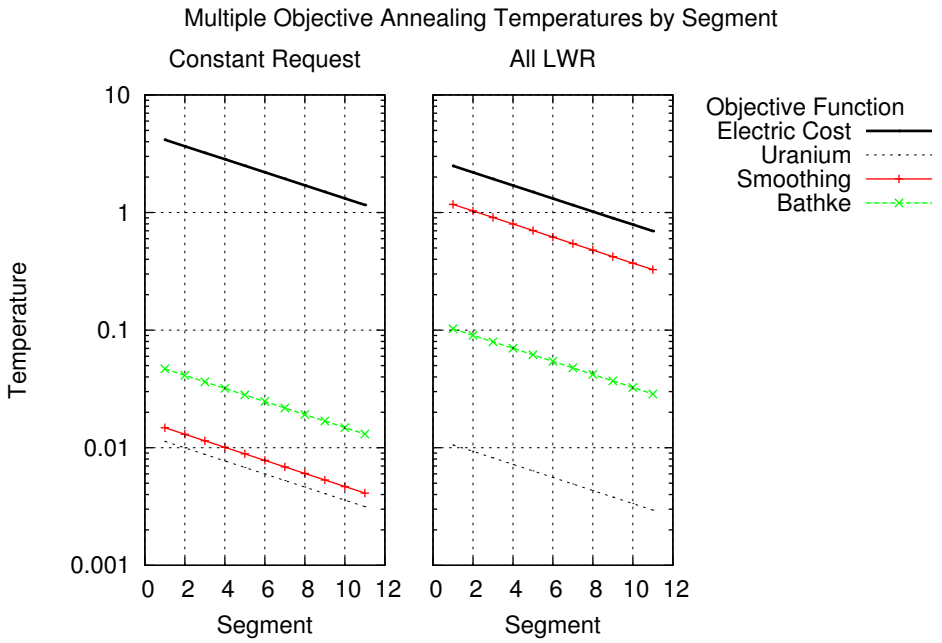

Figure 3.27: MOSA Objective Temperatures

deal of variation in initial value between the two runs, they all exhibit the same reduction rate throughout the annealing process. This indicates two things: 1) the temperature initialization search did not fully cover the configuration space, leading to notably different initial objective samples, and 2) the adaptively computed temperature decrement for each objective at each cooling step was lower than the minimum allowed temperature decrement. This could either indicate that an insufficient number of samples were taken at each cooling step, or –more likely– that the minimum allowed temperature decrement is more conservative than necessary. It is also interesting to note in Figure 3.28 that the acceptance ratio does not monotonically decrease. This is likely due to the changes in the slope of the objective function trade-off surface at different points in the configuration space. At each temperature update, the random selection of a segment starting point from the non-dominating archive may lead certain segments to have markedly more acceptances relative to other segments. (This is particularly true when the vicinity of the starting point is poorly represented in the non-dominating archive, leading to a large number of solutions being accepted due to inclusion in the archive rather than due to relative objective function values.)

For multi-objective optimization, the primary output is the trade-off surface between objectives. As a 3D surface is very difficult to render adequately on paper, the trade-off surfaces created by these test cases is presented as a set of three two-dimensional projections. The

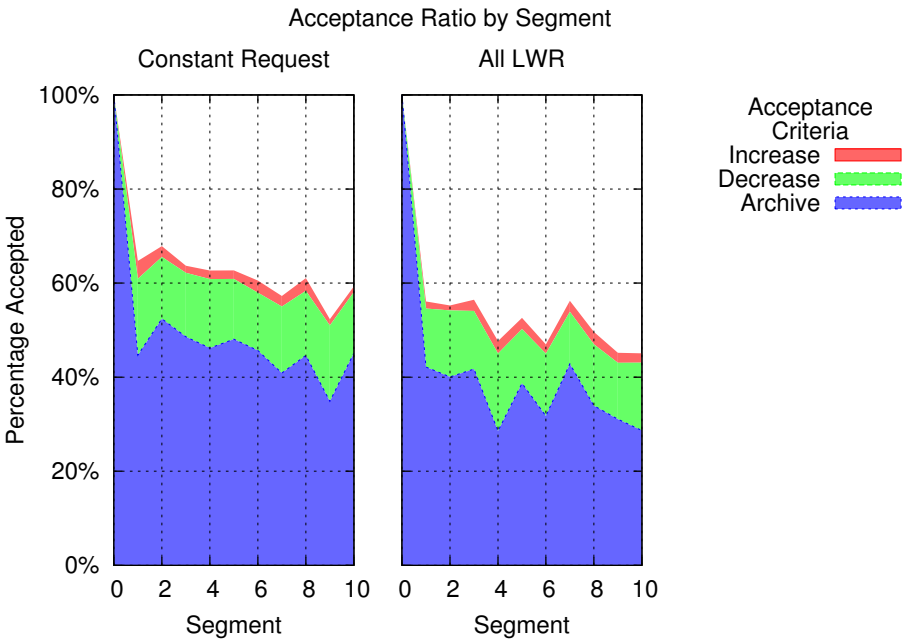

Figure 3.28: MOSA Acceptance Ratio

archived points plotted in red were generated by the optimization with the constant request initial configuration, while those in green were from the All-LWR initial configuration test. Figure 3.29 shows a simple, monotonic trade-off between the uranium utilization and the cost of electricity in the fuel cycle. This is as expected; fast reactors are more fuel efficient, yet more costly. Figures 3.30 and 3.31 show the relationship between the nonproliferation metric and uranium utilization and cost (respectively). These figures show that while the current once-through, All-LWR fuel cycle is the most proliferation resistant scenario, full spent fuel recycling presents less proliferation risk than intermediate fuel cycles. This is due to the trade-off between the large mass of relatively proliferation resistant LWR spent fuel in the first case, and smaller mass of less resistant separated material in the full recycling scenarios. It results because the `VISION` algorithm reprocesses LWR spent fuel after a cooling time whether needed or not for FBRs, producing an inventory of separated material. If the `VISION` algorithm only reprocessed LWR spent fuel when needed to fuel FBRs, this behavior would not occur.

Figure 3.32 shows that while each individual test case had somewhat poor coverage of the search space, the combination of the two has complete (though not even) coverage. This issue may be addressed through either increasing the number of samples in each cooling segment, or through increasing the number of individual perturbations performed in a single sampling.

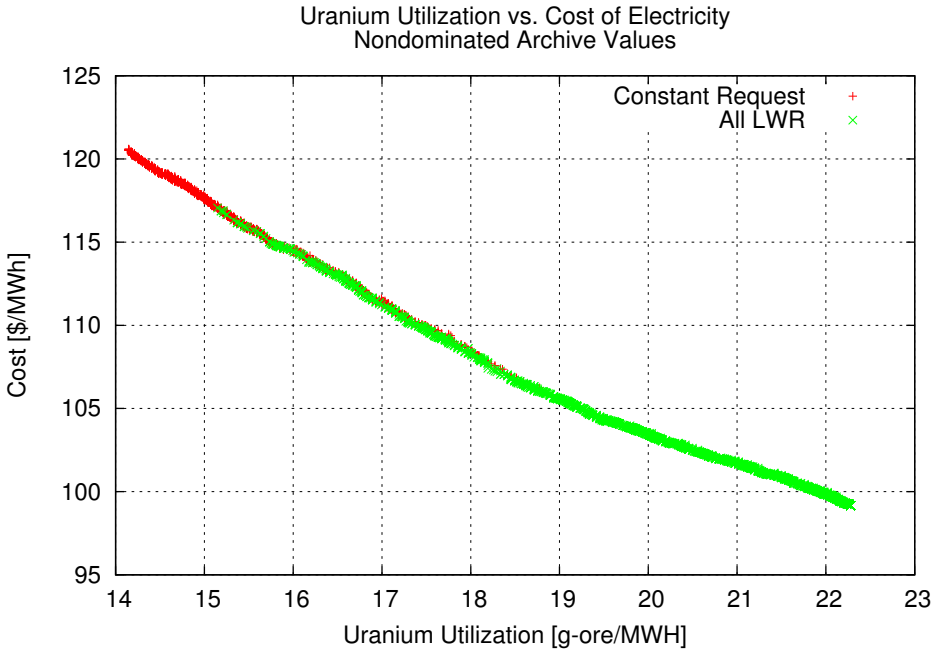

Figure 3.29: MOSA Cost and Uranium Utilization Trade-off Surface

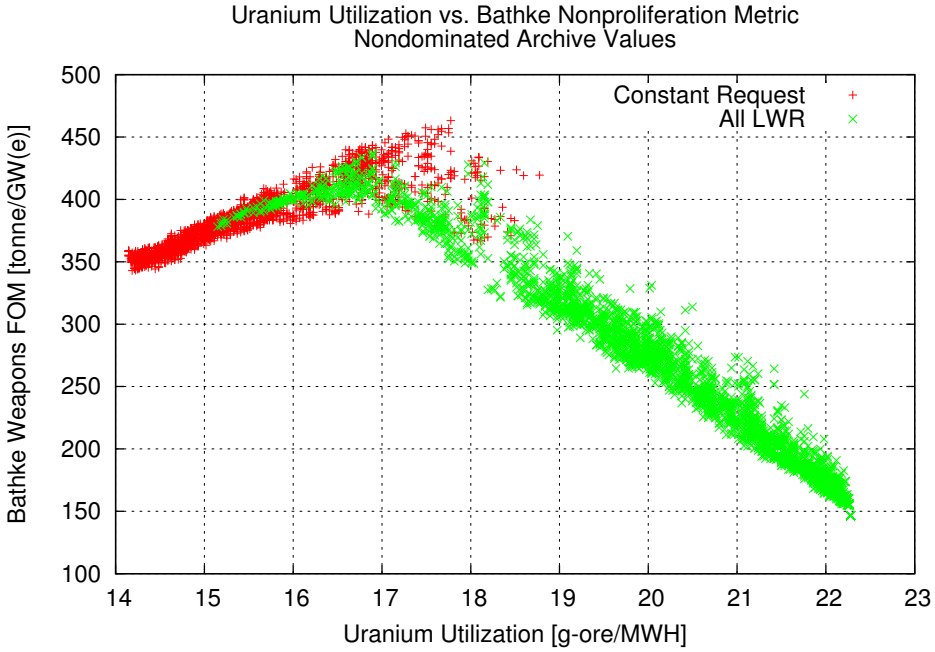

Figure 3.30: MOSA Non-proliferation and Uranium Utilization Trade-off Surface

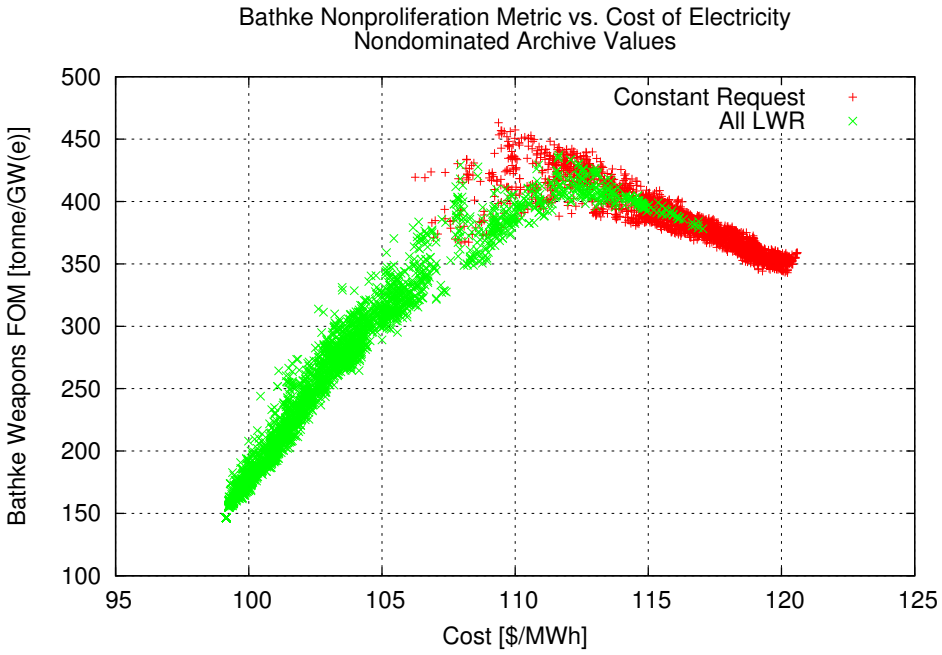

Figure 3.31: MOSA Non-proliferation and Cost Trade-off Surface

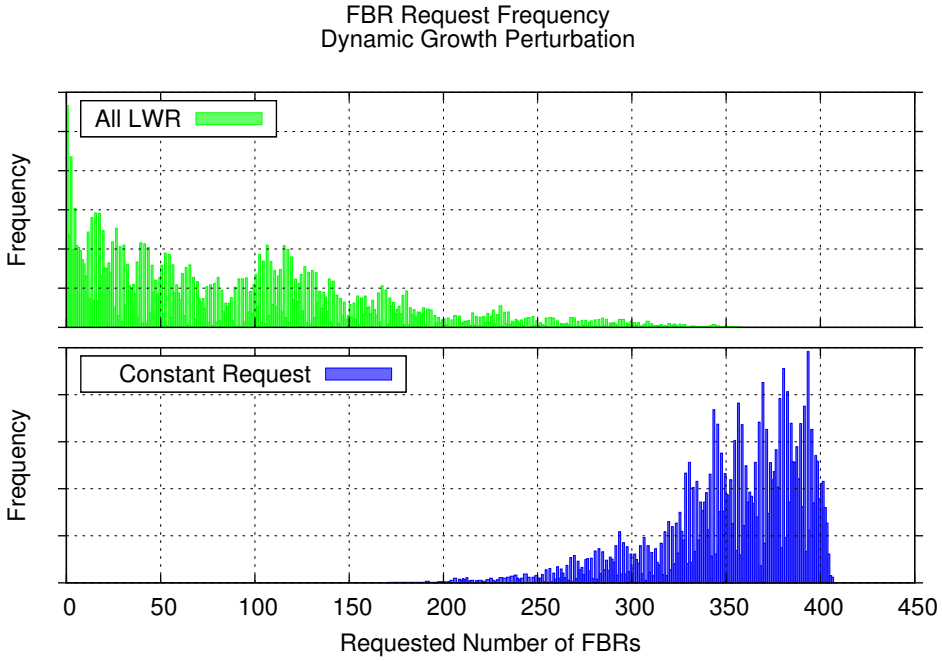

Figure 3.32: Tracking the number of fast reactors requested in each sampled configuration gives an indication of the search space coverage obtained.

3.4 Scalability

While the size and connectedness of the optimization search space are the key factors in determining the rate of convergence of the simulated annealing algorithm, the number of parallel search processes and active objective functions play a large role as well. By employing a larger number of parallel processes in the optimization, fewer samples are required on each individual process in order to reach a suitable distribution of solutions for the next annealing temperature update – thereby shortening the computational turnaround time. The extent of these improvements is limited by the requirement that each process must be able to traverse the entire search space during a given segment in order to avoid trapping in local minima. However, this restriction is somewhat offset by the fact that, at each annealing temperature update, every parallel process starts anew from a randomly selected configuration in the non-dominated archive; this increases the diversity of solutions being sampled at a given time relative to when fewer processes are utilized.

In multi-objective optimization, new non-dominated results are always accepted by the search algorithm. As the number of active objective functions increases, the dimensionality of the trade-off surface also increases. If the objective function values are continuously or finely distributed, then the majority of acceptances will be due to acceptance in the non-dominated archive. While this mechanism does improve the resolution of the trade-off surface, it exhibits poor convergence properties (bordering on a brute-force approach) while providing only marginal benefits to the user. In this section, a selection of optimization tests are compared to examine the effects of these scalability concerns on the optimization problem at hand.

The convergence properties for the five test cases listed in Table 3.5 are compared based on the change in overall acceptance rates and acceptance mechanisms as the optimization progresses, as shown in Figures 3.33, 3.34, and 3.35. Two observations are made from the plots in Figure 3.33: first, the ratio of those configurations accepted with dominant, non-dominant, or un-dominant remain rather consistent in the two-objective test cases (A through D), whereas a much smaller fraction of acceptances in the three-objective case (Test E) is due to un-dominant solutions. Secondly, in Tests A through D, it can be seen that the overall acceptance ratio eventually dropped to the level at which temperature updates were triggered by all threads reaching the pre-set sample limit (rather than the pre-set acceptance limit, as evidenced by the flat peak on the total number of samples per cooling step). Test E, however, does not exhibit this behavior, indicating that the overall acceptance ratio remained somewhat higher throughout the optimization.

It is expected that as the annealing temperatures decrease throughout the optimization cycle, the accessible space of configurations will shrink, and the configurations will trend with increasing probability toward areas where the objective function values are minimized. In

Table 3.5: Relevant optimization parameters for a selection of MOSA trials used to examine the effects of process and objective count on algorithm performance.

Test*	A	B	C	D	E
Settings:					
Objectives†	2	2	2	2	3
Constraints‡	1	1	1	1	1
N_{proc}	8	8	12	16	16
L_{chain}	150	150	100	75	100
L_{tran}	60	60	40	30	40
Samples:					
Requested	7,000	7,000	7,000	7,000	15,000
Total	8,521	7,797	7,671	7,672	15,913
Accepted	4,218	4,285	3,715	2,983	11,481
Rejected	4,287	3,492	3,940	4,679	4,420
Cooling Steps	7	7	6	6	12
Runtime:					
Total [h:m:s]	79:31:58	72:55:25	48:09:07	36:10:21	75:32:13
Per Sample [min]	4.48	4.49	4.52	4.53	4.56

* These optimization tests are culled from runs presented in later sections examining the effects of changes to various aspects of the fuel cycle analysis. Tests A, B and C are used in Section 4.2 examining the price of uranium. Test D is used in Section 4.3, examining the effect of FBR cost reductions. Test E is used in Section 4.1, where various TRU conversion ratios are investigated.

† All tests optimized on Cost and Uranium Utilization, with Test E also optimizing on Long-term Heat.

‡ Smoothing constraint.

addition, it is expected that randomly applied smaller perturbations will generally have a smaller effect on objective values than larger randomly applied perturbations. Combining these two observations, one would therefore expect to see the distribution of perturbation size amongst accepted perturbations to progressively gravitate toward smaller perturbations as the optimization progresses. As with the previous acceptance ratio plots, the plots in Figures 3.36, 3.37, and 3.38 show that this trend is primarily evident in the two-objective optimization cases. The three-objective case again appears to converge much more slowly.

The somewhat limited time and computational resources available for this project, coupled to the wide range of desired test cases restricts the ability to fully examine the parallel scalability and long-run asymptotic convergence properties of this optimization algorithm. Furthermore, these scalability tests were set up with the primary intention of gathering useful data about fuel cycle scenario trade-off surfaces, not to explore the useful limits of the algorithm itself. Thus, unfortunately the data available at this juncture do not provide any comprehensive insights into the ultimate scalability or parallel efficiency. It can be noted, however, that the computational time required per configuration sample increases slightly as the number of processes increases. This represents a modest 1.8% increase in computation time in moving from 8 to 16 processes (based on a comparison of Test A and Test E run-times in Table 3.5). This increase is mostly likely attributable to the increased overhead involved in coordinating the temperature updates among the higher number of processes.)

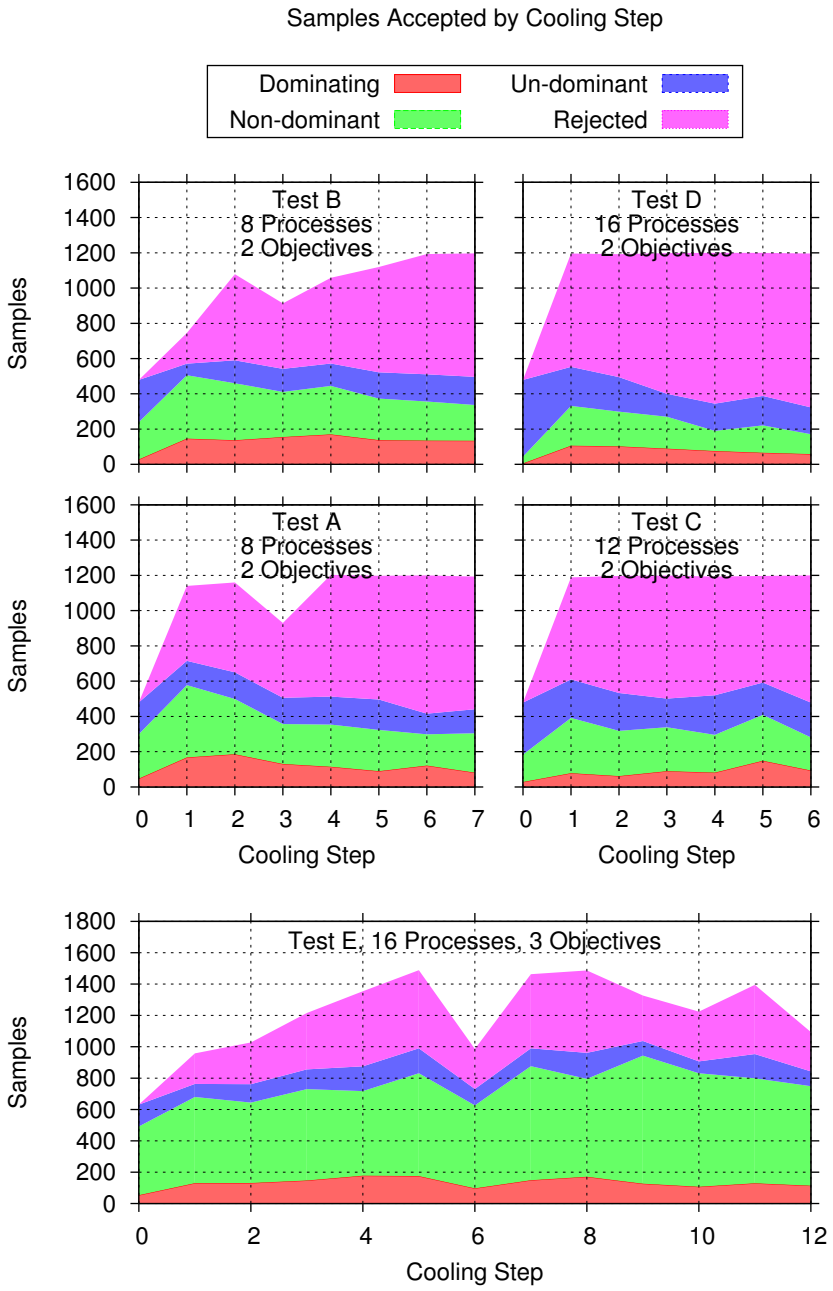

Figure 3.33: While the sample acceptance mechanism does not exhibit a strong dependency on the number processes used, the fraction accepted due to inclusion in the non-dominating archive increases greatly with the addition of an additional optimization objective.

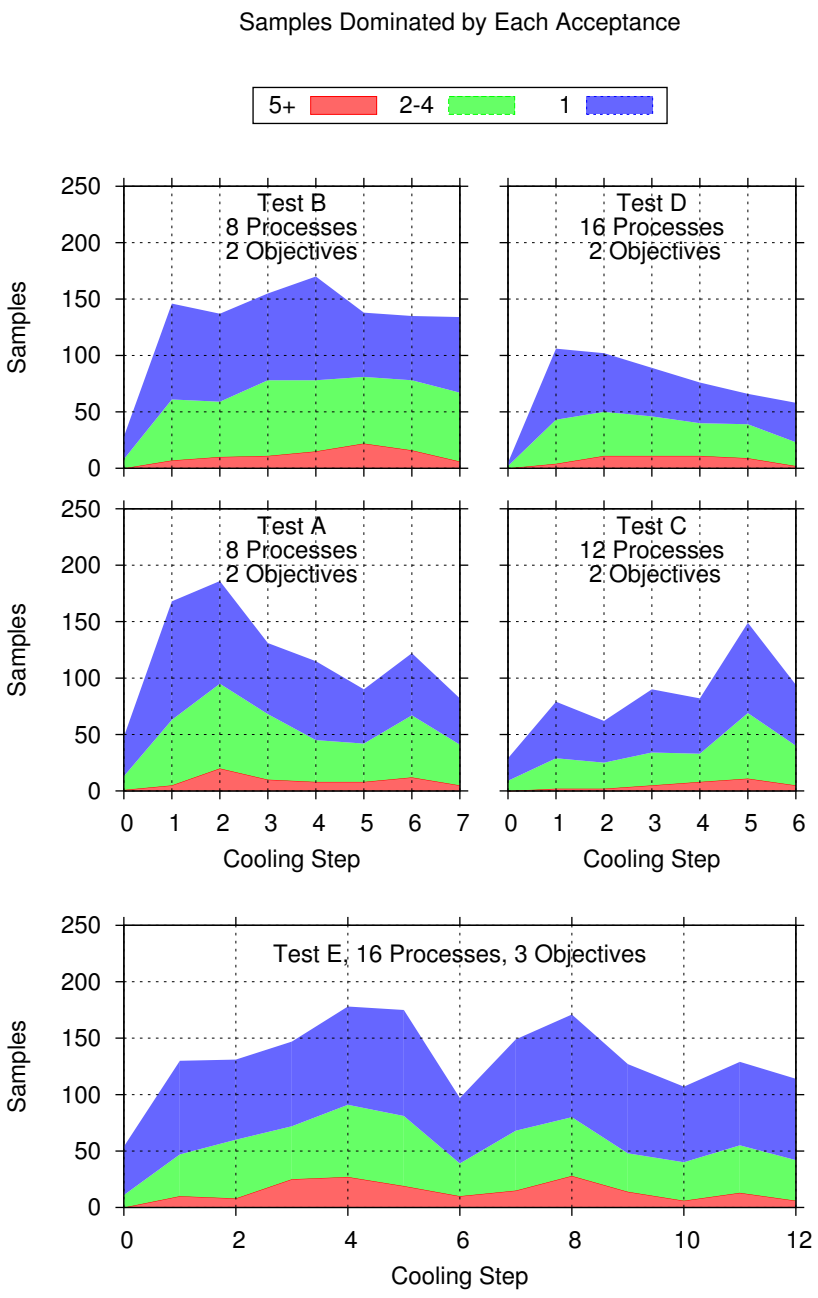

Figure 3.34: Any configurations existing in the non-dominating archive that are dominated by a new addition are themselves removed. The number of configurations removed for each archive addition gives an indication as to whether the optimization is advancing the trade-off surface, or merely refining it.

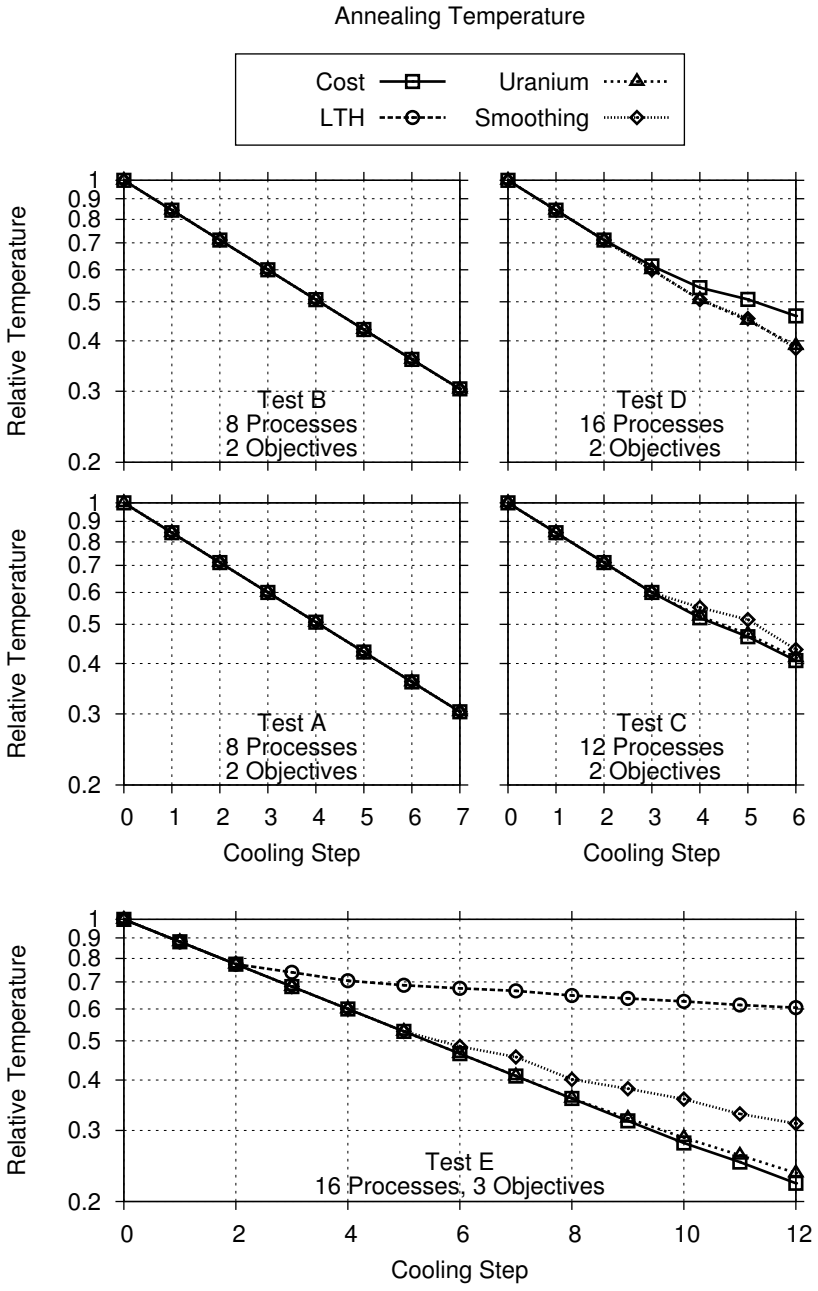

Figure 3.35: The annealing temperatures employed for each objective vary independently based upon the distribution of accepted objective function values. Temperatures which decrease at an exponentially constant rate throughout the optimization are limited by the maximum temperature decrement, whereas those which decrease at a varying rate do so due to the action of the adaptive cooling schedule algorithm. Note that the maximum decrement rate differs between the two- and three-objective tests due to the $1/K$ exponent factor of equation 1.4.

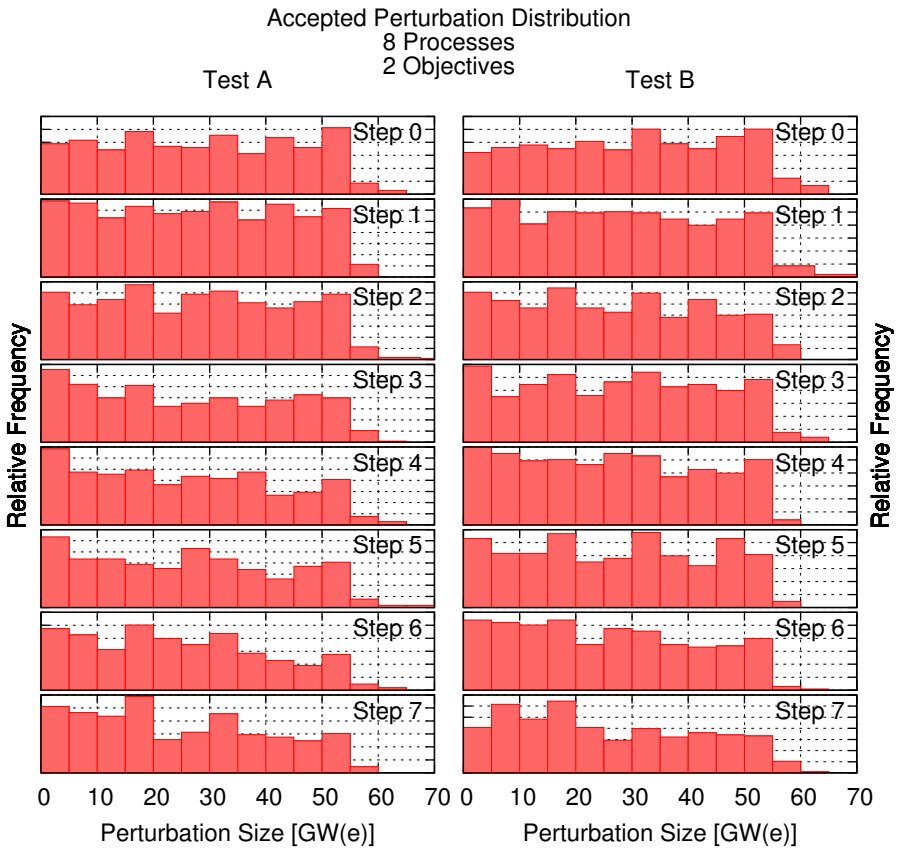

Figure 3.36: The distribution of accepted perturbation distances in optimization tests A through D show a slight preference for smaller perturbations in the later cooling steps than in the initial steps.

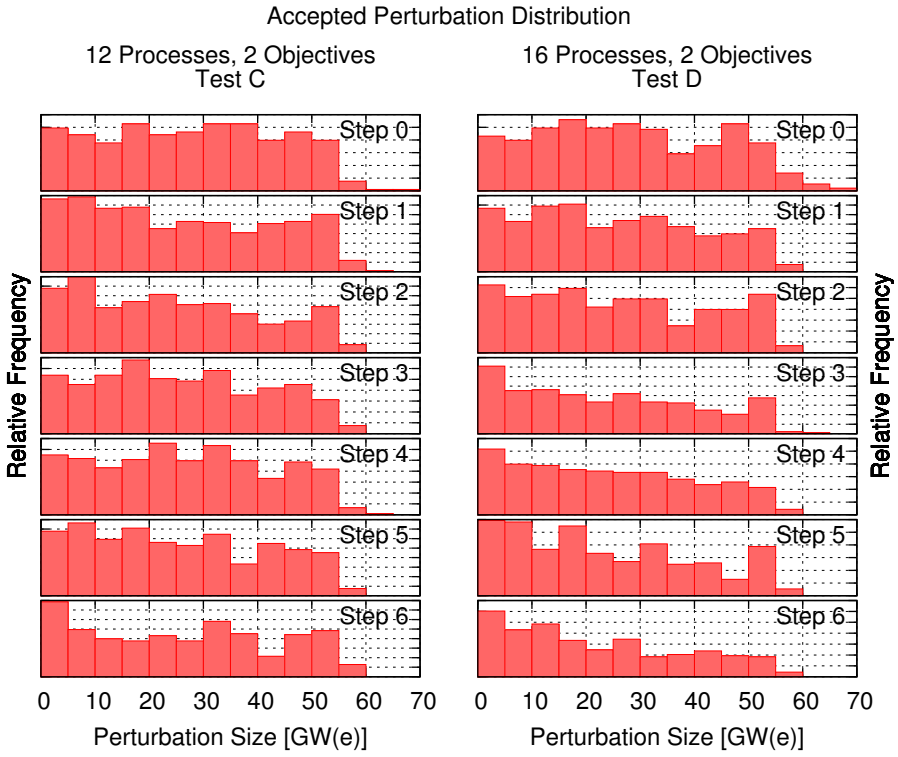

Figure 3.37: The distribution of accepted perturbation distances in optimization tests A through D show a slight preference for smaller perturbations in the later cooling steps than in the initial steps. (Continued)

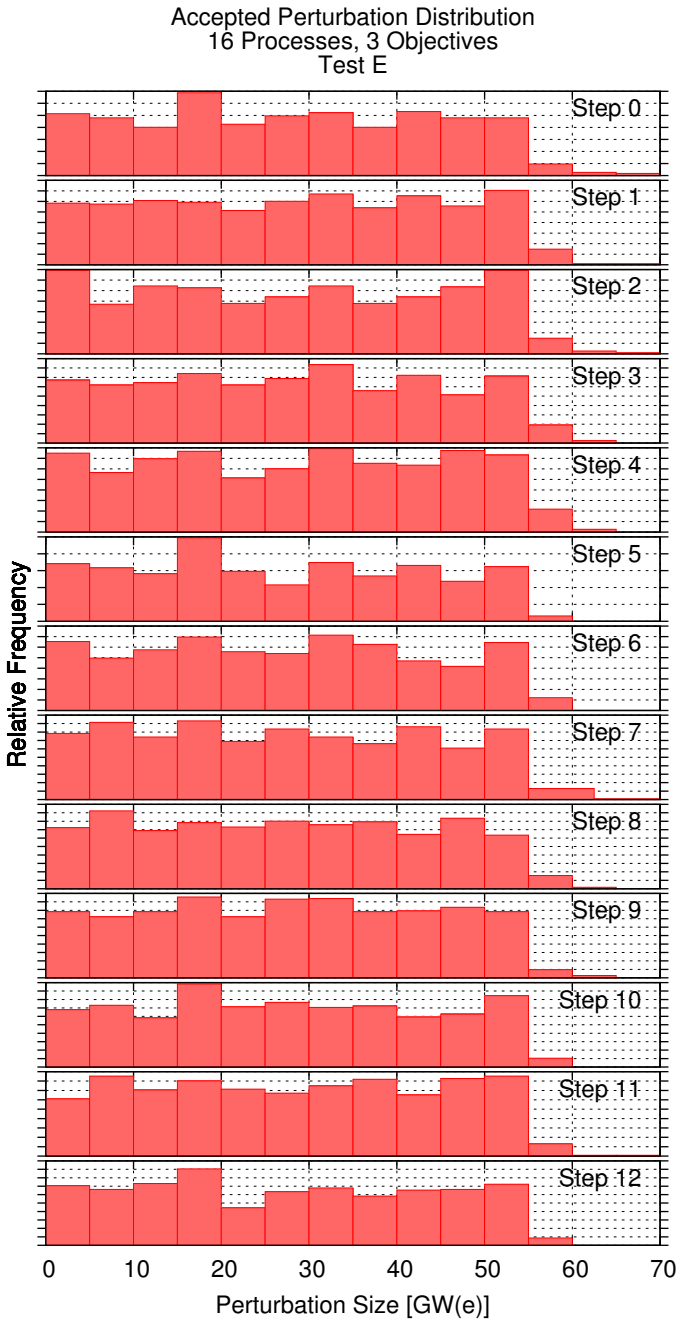

Figure 3.38: Test E does not exhibit as marked a preference for smaller perturbations in the later cooling stages as Tests A through D. This is likely due to the addition of the third objective function.

Chapter 4

Optimization Tests

Having established in the previous chapter the functionality of both the modified VISION model and the optimization framework, we may now apply the optimization in an effort to gain insights into the potential future direction of the nuclear fuel cycle as it relates to the chosen objectives. The test cases selected below were chosen to provide both a demonstration of the optimization capabilities developed to date and to examine several fuel cycle parameters of interest. A full examination of all available combinations and perturbations, however, remains beyond the scope of this project.

4.1 Fast Reactor Conversion Ratio

While the base-case settings for the VISION "Once-through then Fast Burner Reactor" scenario specify that the FBR reactors have a transuranic conversion ratio of 0.5, additional specification sets are available within the model for CRs of 0.0, 0.25, 0.75, and 1.00 (please see Appendix B for details). The change in fuel usage factors that accompanies a change in conversion ratio will have an effect on both the range of feasible deployment scenarios and the objective function trade-off curves. For example, a lower conversion ratio will increase the TRU consumption rate for an FBR, which will in turn restrict the maximum number of FBRs that may be built. It will also, for a fixed number of reactors (and therefore, cost) decrease the long-term repository heat burden relative to the default case. However, because each reactor will consume more TRU, the potential improvements in uranium utilization will be more limited than in higher CR scenarios. In this section, we examine and attempt to quantify the change in objective trade-offs due to a change in fast reactor transuranic conversion ratio.

Three test cases were examined using identical optimization parameters, but with different conversion ratios of 0.25, 0.50, and 0.75. Each test utilized the Long-term Heat, Uranium Uti-

lization, and Economics objectives, as well as the smoothing constraint function. Computations were performed on all 16 available processes, with a minimum of 15,000 samples being requested. Annealing temperature updates occurred after at least 14 of 16 processes had sampled 100 configurations or accepted 40. The objective trade-off surfaces contained in the non-dominating archive at the end of the optimization cycle are plotted below in Figures 4.1 and 4.2.

As is expected, an increase in the number of deployed fast reactors results in an increase in the total cost of electricity, while the uranium utilization and long-term heat loading objectives show improvement. An initially unexpected feature of the observed trade-off surface, however, is the manner in which the long-term heat objective rapidly decreases for the first few fast reactors before settling at a nearly constant low value. A close examination of the `VISION` model reveals that this effect is due to the simple method currently used to predict the sizing, construction, and operation of the aqueous reprocessing facilities used to treat the LWR spent fuel. In this method, reprocessing facilities are built as needed such that their annual capacity will meet the maximum annual fuel demand of the FBR fleet. This fuel demand rate has a very sharp peak when the first reactor core is produced, but then quickly drops away. Because the ordering method does not predict this demand shape ahead of the time required to build the reprocessing facility, it cannot pro-actively accumulate a stock of separated material in order to spread the production rate out over a more reasonable time frame. Furthermore, once the reprocessing facility has been built, it continues to operate at full capacity until it exhausts the available supply of spent fuel. This transfers the majority of the heat-emitting isotopes from the spent fuel (waste) stock, to the separated material (resource) stock, even though there is no foreseeable use for it. In configurations approaching the upper limit of FBR deployment levels, this approach has little effect on the overall fuel cycle (apart from low utilization of reprocessing facilities). However, for configurations in which the overall FBR deployment level is well short of the maximum, this leads to significant excess storage costs and the accumulation of politically problematic separated plutonium stocks. Improvements to the `VISION` algorithm are not addressed in this work, as they are expected to have been incorporated into more general ongoing improvements to the `VISION` model.

It is also interesting to note that though the extent of the trade-off curve between cost of electricity and uranium utilization differ between the three cases tested, the curves align very closely over those portions where their extents do coincide. Thus, if only a few fast reactors are to be built, the conversion ratio at which they operate will have little effect on the total mass of uranium required by the system as a whole, allowing those who plan the system to consider other factors as well.

The selection of FBR conversion ratio plays a large role in determining the total number of FBR reactors that may be built during the planning horizon. Lower conversion ratios naturally favor the goals of quickly minimizing the LWR spent fuel burden, while higher ratios can

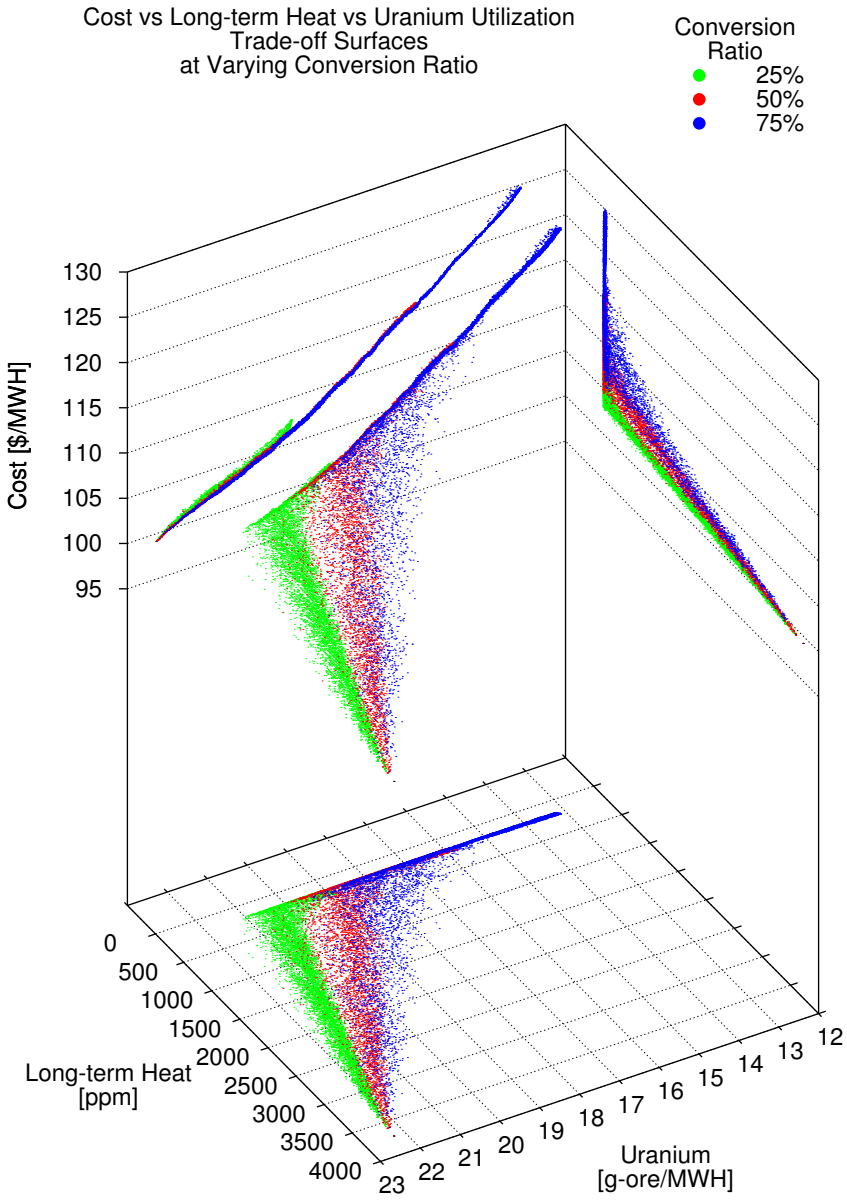

Figure 4.1: The shape and extent of the trade-off surface varies with transuranic conversion ratio.

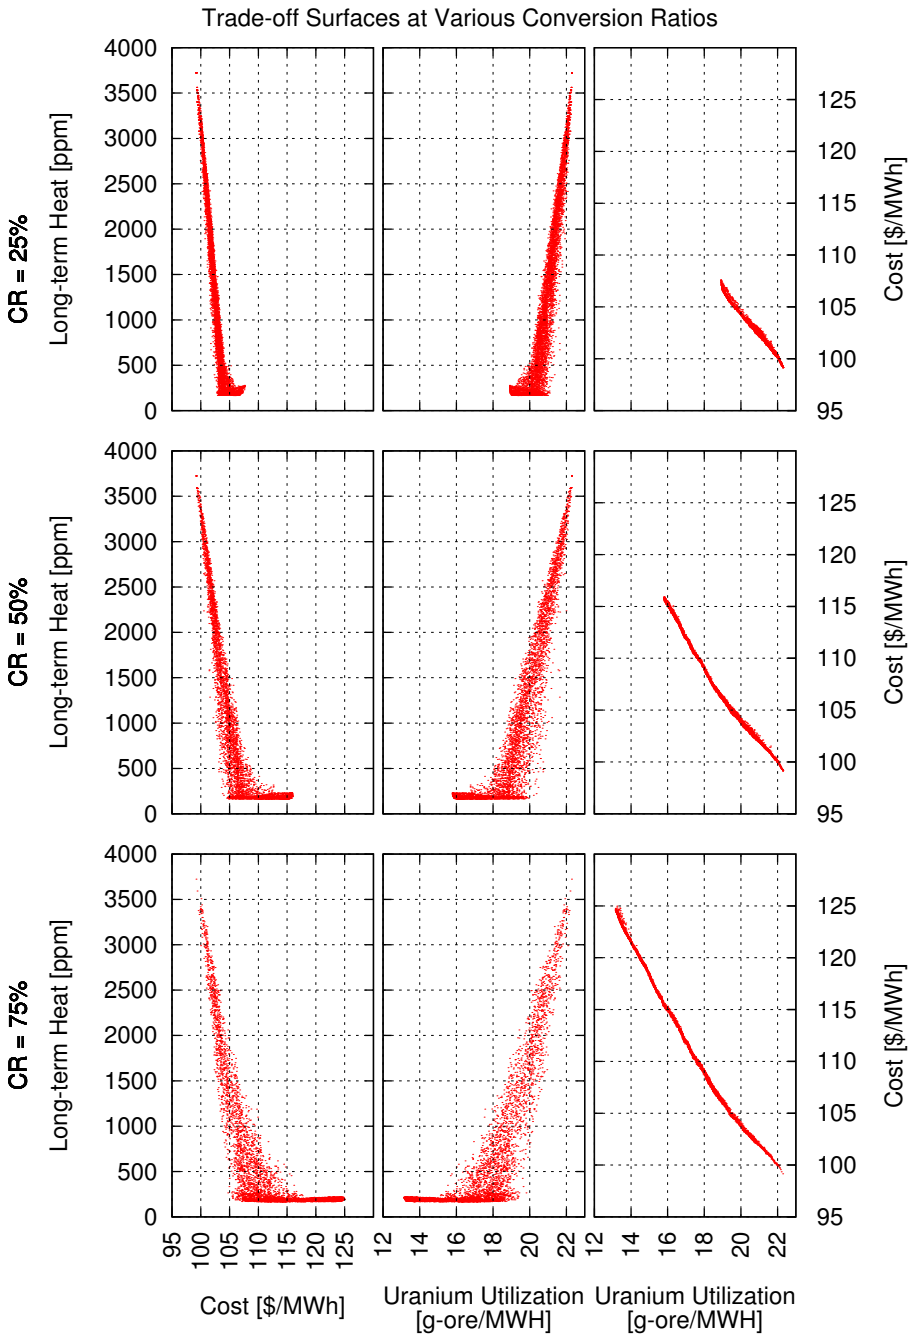

Figure 4.2: The shape and extent of the trade-off surface varies with transuranic conversion ratio.

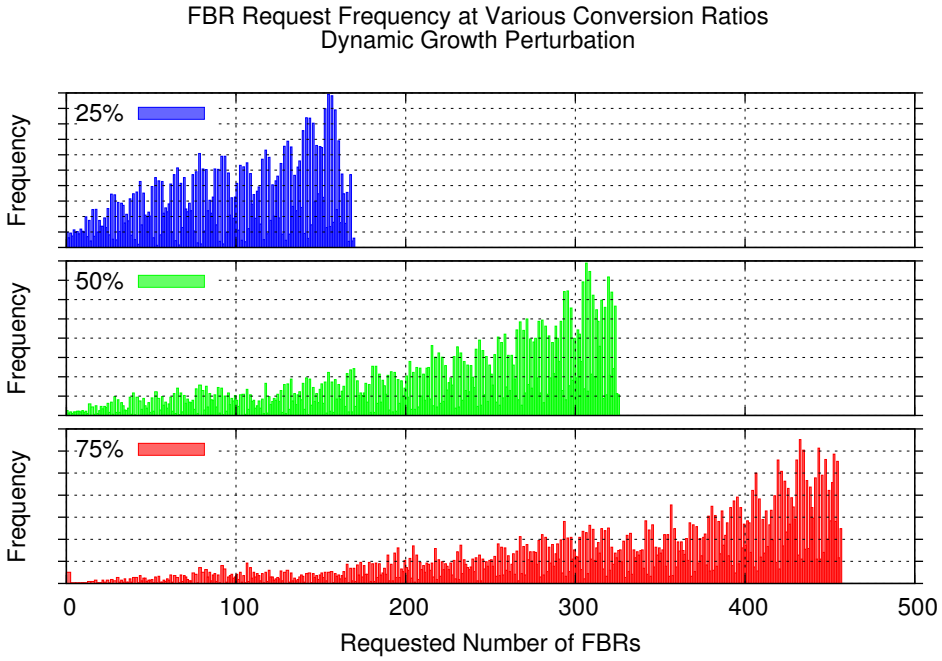

Figure 4.3: FBRs with higher conversion ratios consume less TRU, therefore allowing a greater number to be built. The finely structured variation in these most likely due to the combination of the fact that the ratio of power output between fast and thermal reactors is slightly greater than 1:2, and the electric demand gap is a function of both continuous demand growth and reactor construction history.

ultimately extract greater energy from said fuel (albeit, at a greater cost). Figure 4.3 shows just how great is the variation between feasible FBR build-out levels at the tested conversion ratios.

Finally, it is instructive to note the slope of the trade-off surface between uranium utilization and cost. The cost of electricity increases by approximately \$25/MWH, while uranium utilization improves by 9 grams of ore per MWH. Thus, the marginal cost of uranium saving is about \$2,780 per kilogram (assuming all other costs remain fixed.) As this is more than an order-of-magnitude beyond current uranium prices, and well outside of the forecast future price range, one may conclude that fast burner reactor usage will never be justified on the grounds of fuel economy alone.

4.2 Uranium Supply Models

The estimations of uranium ore price in the AFCI Cost Basis Report [29] are given by a constant-value triangular probability distribution with minimum, nominal, and maximum values of \$30, \$70, and \$260/kg respectively. Historically, however, mineral prices are known to vary in response to a number of factors such as the quality and availability of ore reserves, the level of demand, and the rate of change of demand. A comparative study of historic price trends of various minerals as they relate to the future price of uranium was performed by Schneider, et. al. [26], leading to a best-estimate forecast of

$$P(t) = P_0 e^{Mt} \tag{4.1}$$

where P_0 is the price of uranium in the year 2005 (equal to \$100 year 2005 valued dollars), t is the number of years from 2005, and M is a fitting parameter with a value of $-0.0118 \pm 0.01358 year^{-1}$. It should be noted that though the cited paper concludes that such factors as the availability of substitutes, the diversity of end-users, and the knowledge of the quantity and quality of ore reserves do influence the historic price trends of many minerals, the magnitude of these effects is not quantified, and is therefore not explicitly considered here.

Figure 4.4 shows how the uranium price trends change in time for the two models considered. Figure 4.5 charts the probability that the uranium price will be less than a given value at given time points for the two pricing models.

Four multi-objective optimization tests are used here to compare the effects of uranium pricing on the cost trade-off surface. The 0.50 conversion ratio test from the previous section provides a comparative baseline for the next three tests in which the fixed uranium price distribution is replaced by a Schneider-like exponential growth distribution. In the first test, the nominal values of $M = -0.0118 \pm 0.01358 year^{-1}$ are used. In the remaining two tests, the value of M is fixed to plus or minus one standard deviation from the nominal value. In order to decrease the computational time required for these tests and to gauge the performance under varying conditions, these tests used only two objectives, sampled a minimum of 7,000 configurations, and utilized either 8 or 12 parallel processes. The threshold number of samples and acceptances per segment required before triggering a temperature update was adjusted so that the total number of samples across all parallel processes would remain nearly constant (please see Section 3.4 for details). Figure 4.6 shows the resulting trade-off surface from the first two tests, while Figure 4.7 contains a comparison of the three Schneider-like tests.

As can be seen in Figure 4.5, the expectation value for uranium price – the price that has a cumulative probability of 0.5 – is lower for the Schneider price forecast than for the AFC prediction. This relationship holds for most probabilities and times apart from the highest end of the price distribution at the later years of the planning horizon. This observation explains why

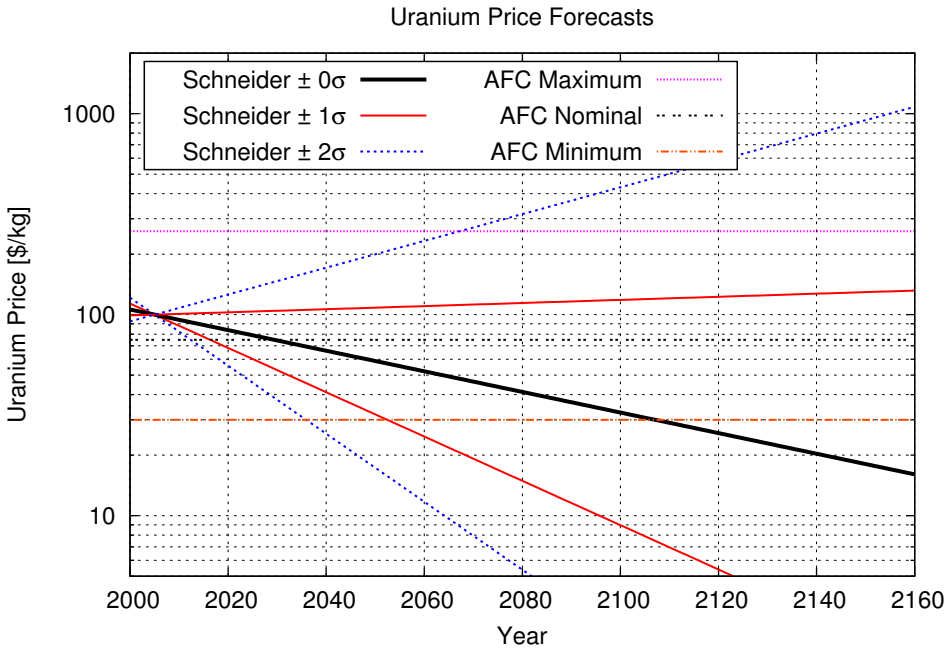

Figure 4.4: The Schneider model forecasts that uranium prices will decline on average, however there is considerable uncertainty and variation in this estimate. The AFCI model uses a single, average cost value throughout the planning horizon.

in Figure 4.6 the [90% downside] cost of electricity calculated using the Schneider price forecast is lower than that calculated using the AFC reference price. Furthermore, those deployment scenarios which rely heavily on FBRs (those with lower Uranium Utilization objective values) will have a greater fraction of their uranium usage occuring early in the planning horizon, when the Schneider forecast upper-bound price most greatly under-cuts the AFC predicted prices. This explains why the difference between the trade-off surfaces widens as the uranium utilization objective value decreases.

To test the sensitivity of the total electric cost to the Schneider uranium price escalation exponent, M, two further optimization tests were performed with its value fixed to plus or minus one standard deviation. The effect of this on the trade-off surface can be seen in Figure 4.7; the overall variation in electric cost for a fixed level of uranium utilization was quite low, between 1 and 2%, considerably smaller than the anticipated precision of the overal cost estimation. Thus, variations in the modeling and pricing of uranium have little effect on overall power costs (as would be expected given the capital intensive, low-fuel-cost nature of nuclear power in general).

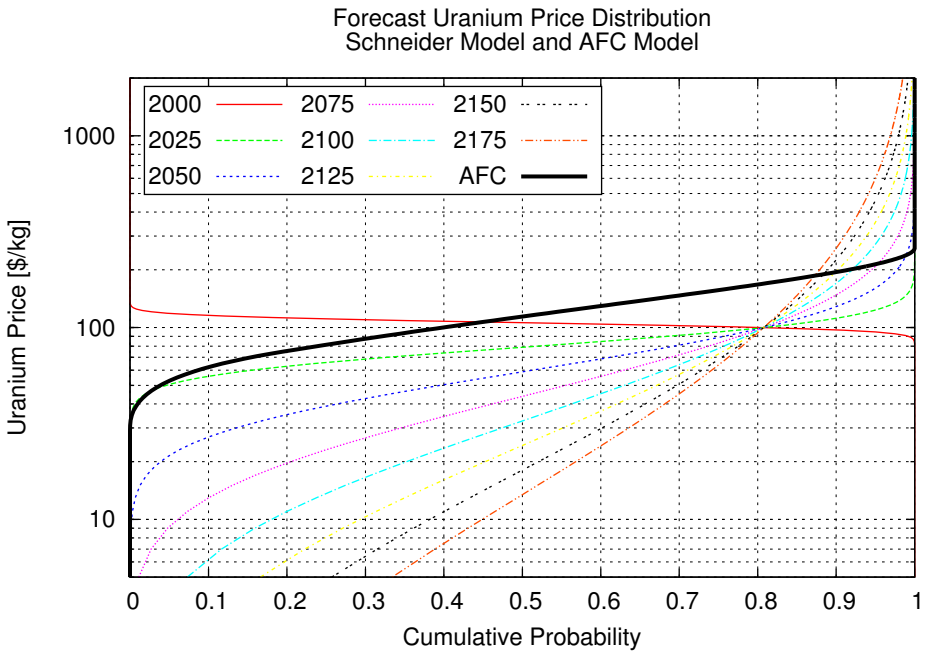

Figure 4.5: This plot shows the probability (along the x-axis) with which the price of uranium will be less than a given level (on the y-axis) for several times within the simulation, and for the AFCI reference triangular-probability price distribution.

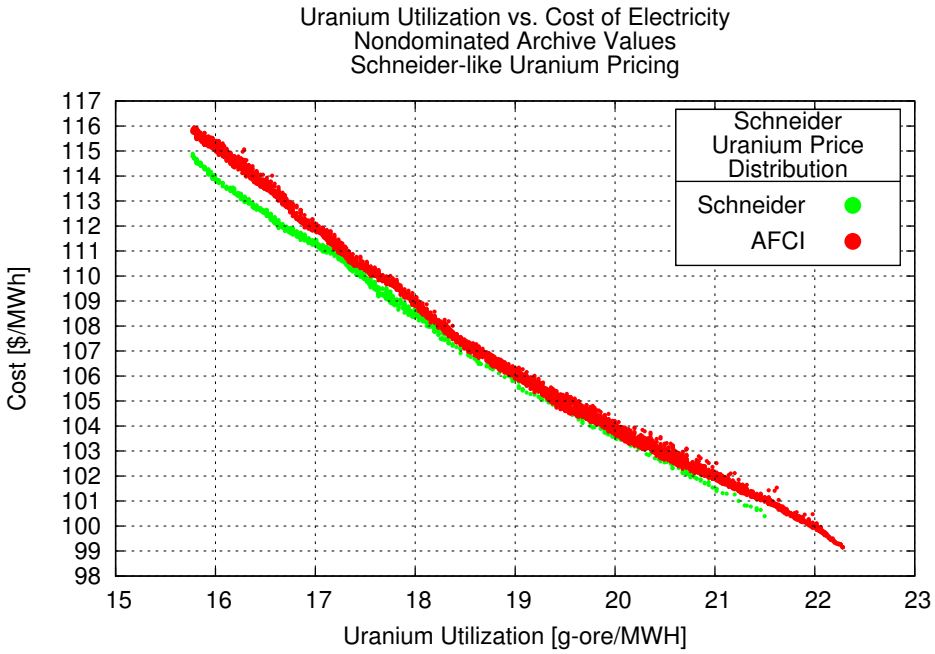

Figure 4.6: There is a great deal of overlap between the predicted downside electric costs using both the AFCI and the Schneider uranium cost models in those configurations that are heavily dependent on LWRs. In highly FBR dependent scenarios, the Schneider uranium cost model exhibits slightly lower overall costs.

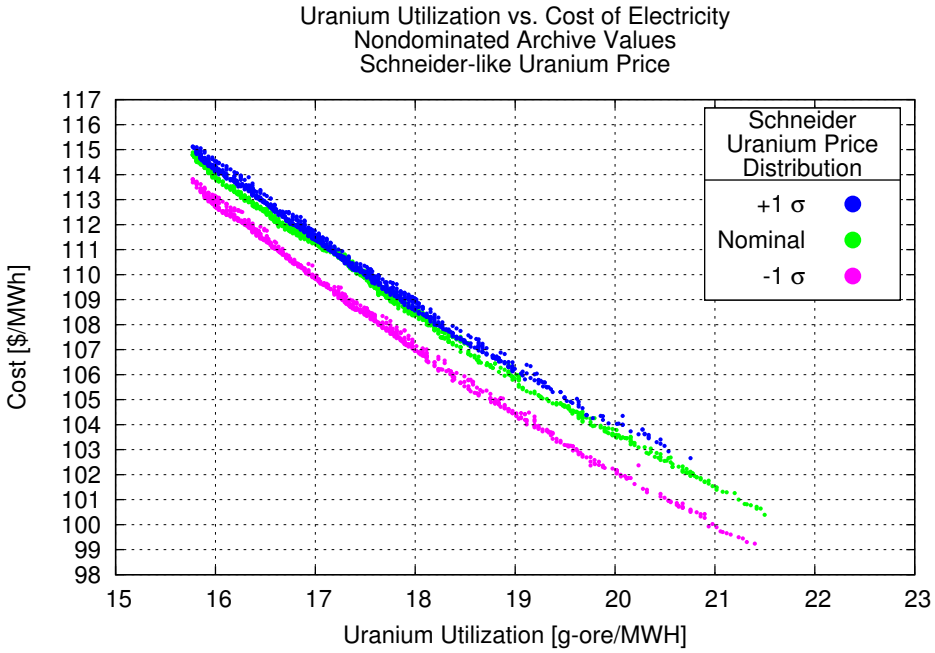

Figure 4.7: The overall cost of electricity has relatively little variation with respect to differing uranium price predictions; about 1-2%, with lower sensitivity in configurations with greater levels of spent fuel recycling. The cost-of-electricity objective value series for the nominally distributed Schneider uranium price forecast does not fall directly in between the series with the fixed plus and minus one sigma uranium price distribution due to the fact that the cost-of-electricity objective itself takes values at the 90% cost probability distribution, thereby penalizing cost uncertainties in addition to cost increases.

4.3 Fast Reactor Learning Curve

In all endeavors, certainty comes through knowledge, and knowledge comes through experience. Thus the relatively small experience base for fast burner reactors, their fuels, and their associated processes translates into a high degree of uncertainty around their costs and operations. This uncertainty leads to increased financial risks, which in turn drives up the cost projections, both through higher contingency factors (F_C) and higher projected return rates for investors (r_D). The current economic cost treatment assumes that the cost factors for each type of facility remain constant throughout the simulation. However, it would be reasonable to predict that as the nuclear experience base grows during the course of the simulation, many of these cost factors that relate to design and operational uncertainty would therefore decrease.

In order to estimate the impact of decreasing capital costs on objective trade-offs and optimized configurations, one must first assign some assumed functional form for the rate of reduction. For simplicity, assume that the cost of each subsequent unit decreases exponentially according to the total number of units built.

$$C(i) = C_0 e^{-\lambda i} \tag{4.2}$$

where λ is a chosen parameter such that a specified overall reduction, C_0/C_N, will occur after a given number of units, N.

$$\lambda = \frac{1}{N} \ln \left(\frac{C_0}{C_N} \right) \tag{4.3}$$

It should be noted that if a reduction factor of $C_0/C_N = 1.0$ is defined, then $\lambda = 0.0$, and no cost reduction will occur, thus preserving previous calculations.

As with the investigations of uranium pricing, the effects of FBR price reductions are examined by running two additional optimization test, each with a different FBR Learning Factor value. Once again, the $CR = 0.5$ nominal test case is used as the basis for comparison. In the first test, the value of λ is chosen such that the price of a fast reactor drops by a factor of two at the 300^{th} reactor; the second test increases the reduction factor to four. On first glance it may appear that such a large reduction in overall per-unit capital costs would be quite difficult to attain. However, the combined action and variability of the various cost compounding factors – in particular the Discount Rate of Return, r_D, the contingency factor, r_C, and the decontamination and decommission fund rate, r_{DD} – at the very least makes this effect possible. The likelihood of this reduction happening, of course, depends upon many societal, governmental, and economic factors, and will remain unknown and uncertain until ultimately tested by experience.

As can be seen in the earlier results, and again in Figure 4.8, a decision to change from a fully LWR-powered deployment scenario to a maximally FBR-powered scenario results in an

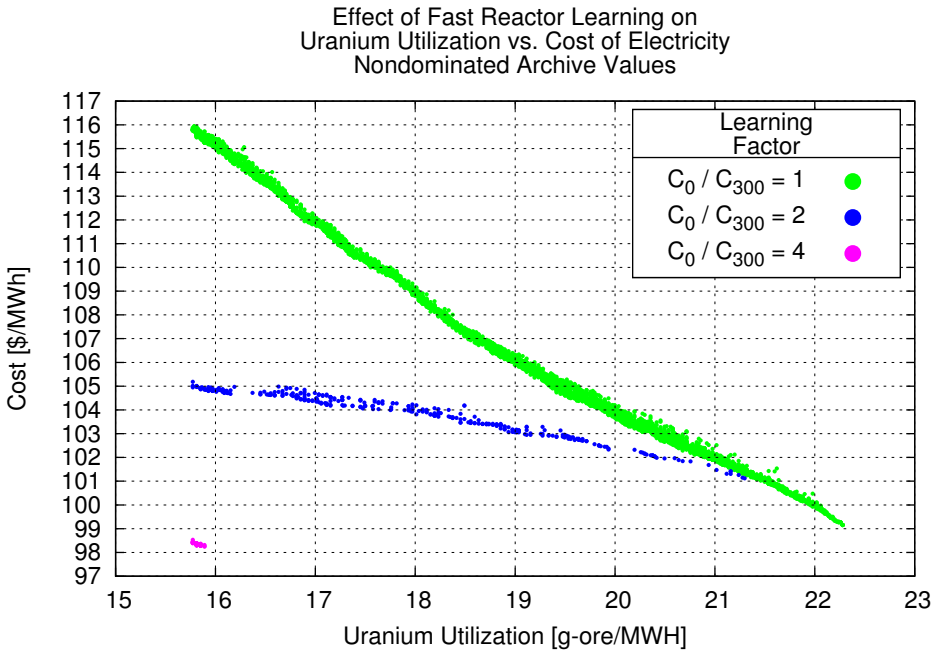

Figure 4.8: Bringing down the capital cost of fast reactors greatly reduces the cost of transitioning to a predominantly fast-reactor-powered nuclear fleet. If capital costs are reduced by a factor of four, then FBRs – even with their associated facilities – become cheaper than LWRs, thereby collapsing the cost vs. uranium utilization trade-off surface.

increase in electric cost of approximately 17%. If, however, fast reactor costs can be brought down by a factor of two over the course of the first 300 reactors, the overall cost increase due to fast reactor use is relatively modest 6%. Further increasing the fast reactor cost reduction factor to four causes the marginal cost of fast reactor power to drop below that of light water reactors, at a cost savings of about 1% overall. This also causes uranium utilization and cost to cease to be competing objectives – they increase or decrease concurrently. This causes the trade-off surface archived by the non-dominated archive to collapse to a much smaller set of solutions which dominate those scenarios utilizing more LWRs. The effect of this collapse of the trade-off surface is apparent in the distribution of the number of fast reactors requested by each random sample (see Figure 4.9).

An important caveat regarding these results would be to note that in these tests the capital cost reductions were applied only to the fast reactors. In reality, many of the drivers that one may postulate that would cause such a reduction would apply equally to light water reactors (and any other competing energy sources or investment opportunities).

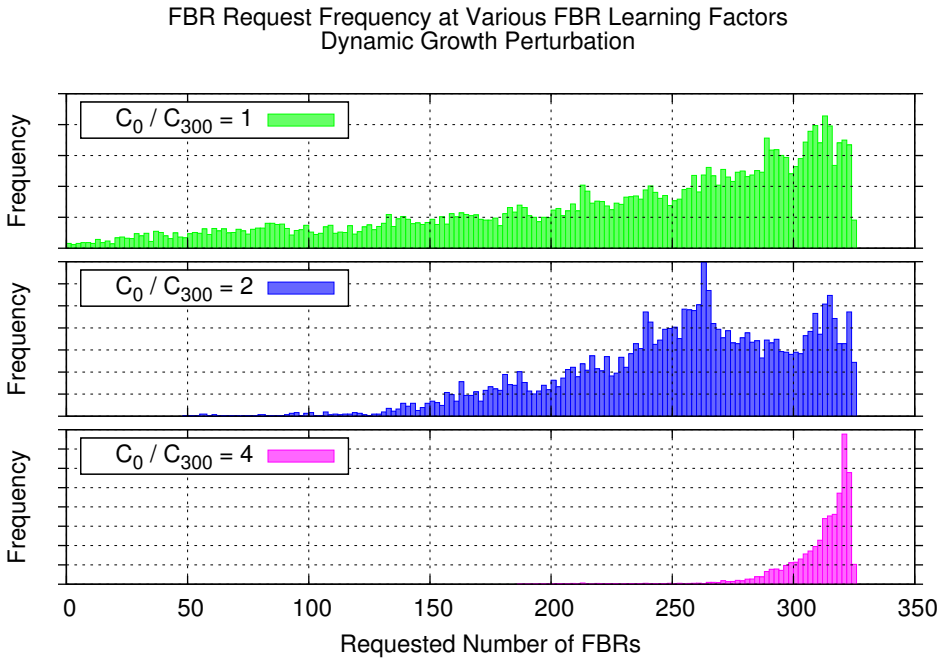

Figure 4.9: The lower costs of fast reactors in the $C_0/C_{300} = 4$ case cause fast reactor heavy configurations to dominate all others, thus confining the sampling range to only those configurations.

Chapter 5

Conclusions

The methods and results presented thus far represent a first attempt at developing an extensible and automated fuel-cycle deployment scenario optimization capability. These results serve to both re-affirm expected outcomes and known shortcomings in available data as well as illuminate and enumerate with greater detail the specific forms and values that objective trade-offs may take in a given scenario. Additionally, the process of developing these early results has yielded a number of insights into the modeling and optimization process itself. Certain of these insights were readily incorporated into the current body of work, while the implementation of others remains speculative, pending further investigation and analysis. This chapter begins with a brief overview of what new contributions have been made through this work, followed by a brief summary of the state of each aspect of this project, including lessons and key insights as well as indicated directions for further exploration.

5.1 New Contributions

This work provides several new contributions to the general field of nuclear facility deployment optimization, and extends several others beyond their previous bounds. Primary among these is the application of multi-objective simulated annealing optimization to the problem of selecting between various facility types and technologies. Although many studies have examined a carefully selected set or range of fuel cycle deployment scenarios – and have made comparative judgments between them on the basis of similar objective functions – none seem to have gone so far as to incorporate an automated configuration generation and optimization mechanism. In addition, the incorporation of a standalone fuel-cycle simulation (`VISION`), the use of modular frameworks for the objective functions and decision variables, and the parallel capability provide for a level of flexibility, scalability and customizability not seen in any other fuel cycle optimization work.

Furthermore, the method used to establish closure in the accounting of costs associated with the planning horizon – by simulating all existing facilities through shutdown (with material disposition) – was developed here specifically for optimization; however, it could also be useful more generally for analysis of system or process alternatives. Finally, while the specific isotopic weighting scheme used in the weighted-weapons usable metric was created specifically for this work, it does not have a wide ranging or rigorous basis in the available non-proliferation literature, and thus may be of somewhat limited utility.

5.2 Fuel Cycle Modeling

The decision to utilize the `VISION` fuel cycle model (version 2.1.9, in particular) has driven a majority of the design decisions in this work. These range from the most basic choices, such as the programming language and computational platform, to more subtle decisions, such as the LWR spent-fuel forecasting algorithm and the approximation that the composition of TRU within FBR fuel is fungible. Despite this, many of the basic methods used here are expected to be indirectly portable to other fuel cycle models, provided that the requisite fuel cycle inputs and outputs may be made available.

The application of a stochastic optimization algorithm to a given fuel cycle model places very strong limitations on both the robustness and performance of the model. The ability of the model to run quickly and to fail elegantly and unambiguously will play a large role in determining the overall performance of the optimization. In this regard, the version of the `VISION` model utilized here was somewhat wanting, partially owing to the limitations of the underlying Powersim framework, and partially owing to its original scope of development. A great deal of investigation and trial-and-error testing was required in order to verify that the model produced the expected results, both in nominal and upset conditions (shortage of fuel, reactors unable to meet electric demand, reactor end-of-life shutdown, etc.) These modifications resulted in a much more robust, though possibly slightly slower model. Subsequent refactoring of select methods was performed for the purpose of enhancing the computational speed. These changes were primarily focused on those methods which utilized the `VBFUNCTION` feature of Powersim, which has very poor performance when dealing with nested iteration loops. It should be noted that recent releases of the `VISION` model (version 3 and higher) have entirely eliminated the use of `VBFUNCTION` calls, thus reducing the runtime to nearly half that of prior versions. Furthermore, the new `VISION` releases have introduced a great deal more flexibility in the specification of fuel cycle options, including the ability to specify the use of up to ten reactor types in a single simulation and the ability to specify a variety of fall-back fuel supply strategies. Finally, the adoption of the latest `VISION` release should address a number of the quirks encountered in the earlier versions, particularly the previously mentioned troublesome aqueous separations fuel loading algorithm.

It is difficult to quantify the sensitivity of the objective function values to the fidelity of the fuel cycle model without the use of comparative tests. However, given the aforementioned approximations used in the accounting of transuranic material in the fast reactor fuel cycle, it may be reasonably expected that objective functions directly tied to TRU inventory will be greatly affected. This would be of particular concern for any future objective functions that take into account spontaneous neutron generation rates, as these depend almost entirely on only the heaviest of the actinide elements. Objective functions which do not require TRU inventories, such as Uranium Utilization and Cost of Electricity, will be relatively insensitive to these modeling approximations.

If one were to design a fuel cycle simulation model from scratch for use with an external optimization engine, one key feature to ensure scalability would be the ability to run on multiple different operating systems. For small scale use, an ad-hoc cluster of workstations running Microsoft Windows is adequate, but for large scale computations, Unix/Linux clusters are the predominant choice.

5.3 Objective Function Formulation

In the work presented here, the computational time required to evaluate the objective functions themselves was only a tiny fraction of the time required to evaluate the fuel cycle model. As such, the sophistication or lack thereof of the objective functions does not figure greatly in the overall optimization turnaround time. However, these objective functions are merely mathematical surrogates for a considerably more nuanced, poorly defined value structure. They necessarily contain a host of assumptions and approximations that belie their ability to accurately reflect those features of the fuel cycle that one would consider important. Thus, the user of the optimization code must be vigilant that the results generated truly represent a suitably optimized solution, and not merely an exploitation of some mathematical idiosyncrasy.

The objective functions used in this work to date represent first approximations of a selected subset of all the possible options; there are several obvious improvements which could be useful in the future. For simplicity, the Cost of Electricity objective function assumes that monetary inflation and cost escalation occur at the same rate; the ability to examine these counteracting forces independently – particularly their potential variability in time – could possibly lend a great deal of insight into the interplay between capital and operating costs in determining overall costs. This objective functions could also benefit from a more sophisticated treatment of the various financing methods available for building new reactors, either through private or government investment. Finally, it would be useful to compare the results obtained using the current 90% cost probability metric to those from using a 50% cost probability value; this would give some insight into the effects of cost uncertainty on the optimization outcomes.

The long-term heat objective function is currently the only waste-minimization metric in use. This choice was made early in the project on the basis that the heat loading in the drift tunnels is the limiting factor in loading the proposed Yucca Mountain high-level waste repository. In the interim, however, it has become less certain as to if and when this site will ever become available. Therefore this objective might benefit from the inclusion of other factors more relevant to other repository sites and designs (and continued dry-cask storage).

The non-proliferation objective functions primarily consider the form, makeup, and total mass of fuel in each facility type in the fuel cycle. As the `VISION` fuel cycle model gives no detailed information regarding the design on the individual facilities, there is no basis for determining the relative proliferation resistance that might be ascribed to the physical and administrative layout of each facility, let alone a way to compare between different facility options. If non-proliferation concerns are to play a key role in the deployment of future fuel cycles, it would be useful to incorporate a more advanced safeguards analysis capability, perhaps through the addition of a variety of facility safeguards design options (and their associated economic cost impacts).

5.4 Optimization Algorithms

Previous experience with the simulated annealing algorithm indicated that its application to this optimization problem had a high likelihood of success. However, this is not to say that it is the only, or even necessarily the best, algorithm suitable for this problem. The desire to retain the flexibility to mix and match continuous and combinatorial input decision variables quickly rules out the direct application of most direct-solution methods, such as conjugate gradient, steepest descent, and others. (Note that it might be possible to apply such methods to certain, well defined subsections of the optimization problem as a whole, but in these situations their functionality would be such that it is an acceleration method for the overriding optimization algorithm). Furthermore, the potentially large size of the combinatorial search space greatly restricts the usability of any brute-force enumeration or exhaustive search algorithms.

Among the available stochastic methods that are applicable to this problem, genetic algorithms (GA) seem to offer the greatest potential for improving the scalability and performance of the program. In particular, the ability to evaluate an arbitrarily large populations of candidate solutions – as computational resources permit – allows for a much more efficient foreshortening of runtime than the iterative search segments of parallel simulated annealing. Additionally, the mutation operator required by the GA method is a direct extension of the existing SA perturbation operation. The primary disadvantage of GA compared to the SA would appear to be the relatively higher level of complexity, which could potentially slow and hinder the development and verification process.

5.5 Maximizing Computational Performance

A key advantage of the simulated annealing optimization algorithm is its relative simplicity – while there are many options, schemes, and methods for improving its performance, very few of them are necessary to produce a working product. The primary focus of this work was to develop and prove the functionality and applicability of this algorithm to the problem at hand; as such, the cooling parameter values are mostly conservatively chosen best estimates. Robustness was given priority over runtime minimization. Therefore considerable performance improvements may likely be gained through a thorough investigation on the cooling schedule details.

Another challenge inherent to the development of the simulated annealing algorithm was in finding ways to manage the voluminous output data generated by the repeated simulation runs. Each simulation entails the generation and storage of three sets of data: 1) the set of input values that define the scenario, 2) the intermediate output values required by the objective functions, and 3) the objective function values themselves. In addition to the required intermediate output values, any number of optional simulation variables may be specified by the user in order to track or compare any value of interest. Currently, all of the variables and values mentioned above are optionally output to a flat text file (for later analysis), while only the inputs and objective function values are retained directly in memory. The non-dominated archive – required in the multi-objective optimization trials – does not have a predetermined size, and may grow to such a size as to outgrow the available computer memory. At the moment this growth is accommodated by caching the input and intermediate values of all archived simulations into a series of binary file volumes. These volumes are updated and synchronized at every cooling parameter update step. Considerable performance and reliability gains would likely be obtained by separating those portions of the code-base responsible for the optimization itself, the archive management, and the input/output management – in particular, the use of a standalone database for simulation results offers much promise.

5.6 Nuclear Fuel Cycle Deployment Scenarios

The optimization tests performed here, though somewhat limited in scope, nevertheless provide a good deal of insight into the trade-offs and limitations involved in planning for continued expansion of nuclear power. The overriding observation with regard to the implementation and expansion of fast reactor technology is that because (in this scenario) they are entirely fueled by LWR spent fuel – not naturally available – their available fuel supply is vastly more limited than that of traditional LWRs. These limitations in fuel supply therefore restrict the number of units which may be successfully operated, even in combination with additional LWRs. Even the change from a burning to a breeding regime would have a relatively minor impact over the

short-term given the relatively low demand growth rate, the long doubling times for breeder reactors, and the fraction of the planning horizon in which fast reactors are available to be built.

The theoretical limitations on the maximum build-out of fast reactors becomes less of a concern when one examines the estimated costs associated with the construction and operation of these reactors and their associated facilities. This occurs both due to the relatively greater costs of the facilities themselves and due to the much greater uncertainty in those cost values. These uncertainties are certain to remain prohibitive to large scale deployment barring the successful construction and operation of a demonstration unit (which itself would require a narrowing of the technological and design uncertainties). In order to maximize the immediacy of the benefits of fast reactor deployment, it might be prudent to focus initially on the design of low conversion-ratio designs, then move to a higher conversion ratio. As the predicted demand growth rate continues to increase throughout the planning horizon, this would maximize the reduction of the existing waste burden without unduly restricting future resource utilization goals.

There are many interesting avenues for further exploration within the current "Once-through then Fast-burner Reactor" base case. These include examining the effects of varying the fast reactor conversion ratios in time, and exploring the effects of different, possibly varying, energy demand growth rates. As stated previously, a key limitation to the deployment of fast reactors using only recycled LWR spent fuel as a source of fissile material is the considerable inventory required to generate the initial reactor loading.

In addition to the current base case considered here, there are a host of alternative technologies and potential growth scenarios which would be interesting to investigate. In particular, those scenarios that allow for either a transition from one regime to another during the simulation (e.g. LWR to burner to breeder) merit further analysis, as it is unlikely that only a single specific technology will be adopted over the long term. An even more interesting study would be to model some form of market competition between the various technologies (including fossil, renewable, and alternative energy sources), in order to estimate trade-offs for the energy sector at large.

Finally, through GNEP and other programs, there is considerable interest in expanding access to nuclear power while limiting weapons proliferation risk through a demarcation between so-called Supplier and User States. The Supplier States – which generally are already in possession of nuclear weapons technology – would retain the sensitive technologies, while selling their services to the User States at a competitive rate. Thus User States avoid the expense and political repercussions of developing a native nuclear enterprise, while Supplier States potentially gain additional market access and help maintain geopolitical stability. The application of multi-objective optimization to these multi-region scenarios could provide useful insight into a number of key concerns, such as 1) the allocation of costs between states given overall national goals, 2)

ways to assure sufficient energy security for the user states so as to prevent their developing their own technology (while also assuring efficient utilization of supplier state resources and an ability to detect and respond to potential diversion efforts), and 3) a globally secure, yet equitable, disposition of waste products.

In summary, the multi-objective simulated annealing algorithm applied here to the `VISION` deployment scenario model is a simple and robust, if computationally intensive, way of probing and examining issues relevant to the continued and expanded utilization of nuclear power both within the U.S. and worldwide. The trial results presented here reinforce many prior results and expectations regarding advanced technology deployment options, while serving to comprehensively fill in many intermediate deployment options and schedules which may not have been explicitly examined previously. Through continued development of `VISION` and other fuel cycle simulations, refinements in the optimization algorithm and program structure, and further development of topical and relevant objective functions and decision variables, there is a great deal of benefit to be realized in the continuation of this line of research.

REFERENCES

[1] Emile Aarts and Jan Korst. *Simulated Annealing and Boltzmann Machines.* Wiley - Interscience Series in Discrete Mathematics and Optimization. John Wiley and Sons, 1989.

[2] World Nuclear Association. Nuclear Power in France, January 2011. Accessed 25-February-2011.

[3] C.G. Bathke, R.K. Wallace, J.R. Ireland, MW Johnson, K.R. Hase, G.D. Jarvinen, B.B. Ebbinghaus, B.W. Sleaford, B.A. Collins, M. Robel, et al. An assessment of the attractiveness of material associated with a mox fuel cycle from a safeguards perspective. Technical report, Los Alamos National Laboratory (LANL), 2009.

[4] Electric Power Research Center. *FORMOSA-B Code Methodology and Usage Manual - Version 3.2.* North Carolina State University, 2004.

[5] R.G. Cochran and N. Tsoulfanidis. *The Nuclear Fuel Cycle: Analysis and Management.* American Nuclear Society, 1999.

[6] Nuclear Energy Research Advisory Committee. Annex: Attributes of proliferation resistance for civilian nuclear power systems. Technical report, U.S. Department of Energy, `http://www.ne.doe.gov/neac/neacPDFs/FinalTOPSRptAnnex.pdf`, 2000.

[7] K. Cwalina and B. Abrams. *Framework design guidelines.* Addison-Wesley, 2005.

[8] B. Dixon, S. Kim, D. Shropshire, S. Piet, G. Matthern, and B. Halsey. Dynamic systems analysis report for nuclear fuel recycle. Technical report, Idaho National Laboratory (INL), 2008.

[9] Philippe Engrand. A multi-objective optimization approach based on simulated annealing and its application to nuclear fuel management. *ASME Conference Proceedings*, 1997:416–423, 1997.

[10] J.W. Forrester. System dynamics and the lessons of 35 years. *The Systemic Basis of Policy Making in the 1990s*, 29:4224–4, 1991.

[11] Fanxing Gao and Won Il Ko. Dynamic analysis of a pyroprocessing coupled sfr fuel recycling. *Science and Technology of Nuclear Installations*, 2012, 2012.

[12] R. Hays and P. Turinsky. Bwr in-core fuel management optimization using parallel simulated annealing in formosa-b. *Progress in Nuclear Energy*, 53(6):600–606, 2011.

[13] EA Hoffman. Updated design studies for the advanced burner reactor over a wide range of conversion ratios. *Argonne National Laboratory, ANL-AFCI-189*, 2007.

[14] EA Hoffman, WS Yang, RN Hill, et al. Preliminary core design studies for the advanced burner reactor over a wide range of conversion ratios. Technical report, Argonne National Laboratory (ANL), 2006.

[15] J.J. Jacobson, A.M. Yacout, G.E. Matthern, S.J. Piet, and D.E. Shropshire. VISION: Verifiable Fuel Cycle Simulation Model. Technical report, INL/CON-08-15051, Idaho National Laboratory (INL), 2009.

[16] S. Kirkpatrick, C. D. Gelatt Jr., and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, May 1983.

[17] Won Il Ko and Fanxing Gao. Economic analysis of different nuclear fuel cycle options. *Science and Technology of Nuclear Installations*, 2012, 2012.

[18] P.L. Kunsch and J. Teghem Jr. Nuclear fuel cycle optimization using multi-objective stochastic linear programming. *European Journal of Operational Research*, 31(2):240 – 249, 1987. ¡ce:title¿First EURO VII special issue: Location and transportation problems production systems¡/ce:title¿.

[19] J. Lee, Y.H. Jeong, Y.I. Chang, and S.H.Chang. Linear Programming Optimization of Nuclear Energy Strategy with Sodium-cooled Fast Reactors. *Nuclear Engineering and Technology*, 43(4):383–390, 2011.

[20] Soo-Young Lee and Kyung Geun Lee. Synchronous and asynchronous parallel simulated annealing with multiple markov chains. *IEEE Transactions on Parallel and Distributed Systems*, 7(10):993–1008, October 1996.

[21] Nicholas Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller. Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21(6):1087–1092, June 1953.

[22] Brian R. Moore, Paul J. Turinsky, and Atul A. Karve. Formosa-b: A boiling water reactor in-core fuel management optimization package. *Nuclear Technology*, 126(1):153–169, March 1998.

[23] Byung Heung Park, Fanxing Gao, Eun ha Kwon, and Won Il Ko. Comparative study of different nuclear fuel cycle options: Quantitative analysis on material flow. *Energy Policy*, 39(11):6916 – 6924, 2011. ¡ce:title¿Asian Energy Security¡/ce:title¿.

[24] S.J. Piet, S.E. Bays, and E.A. Hoffman. Description of transmutation library for fuel cycle system analyses. Technical report, FCRD-SYSA-2010-000116, INL/EXT-10-19545, 2010.

[25] E.A. Schneider, C.G. Bathke, and M.R. James. NFCSim: A dynamic fuel burnup and fuel cycle simulation tool. *Nuclear technology*, 151(1):35–50, 2005.

[26] Erich A. Schneider and Neil Shah. Near term deployment, long term impact: Uranium price over the lifetime of new capacity. *ASME Conference Proceedings*, 2008(48175):891–899, 2008.

[27] T.M. Schweitzer. Improved Building Methodology and Analysis of Delay Scenarios of Advanced Nuclear Fuel Cycles with the Verifiable Fuel Cycle Simulation Model (VISION). Master's thesis, North Carolina State University, 2008.

[28] T.M. Schweitzer, P.J. Turinsky, and J.J. Jacobson. Improved Forecasting Method for Building Reactors and Nuclear Fuel Cycle Facilities with the Verifiable Fuel Cycle Simulation Model (VISION). *Transactions of the American Nuclear Society*, 99:189–190, 2008.

[29] DE Shropshire, JD Williams, KA Smith, BW Dixon, M. Dunzik-Gougar, RD Adams, D. Gombert, JT Carter, E Schneider, and D Hebditch. Advanced Fuel Cycle Cost Basis. Technical report, INL/EXT-07-12107 Rev. 2, Idaho National Laboratory (INL), 2009.

[30] A. Suppapitnarm, K. A. Seffen, G. T. Parks, and P. J. CLARKSON. A simulated annealing algorithm for multiobjective optimization. *Engineering Optimization*, 33(1):59–85, 2000.

[31] L. Van Den Durpel, D.C. Wade, and A. Yacout. Daness: a system dynamics code for the holistic assessment of nuclear energy system strategies. In *24th International Conference of the System Dynamics Society 2006*, 2006.

[32] LGG Van Den Durpel, AM Yacout, and DC Wade. Status on developments and applications of the integrated nuclear energy system code DANESS. *Transactions of the American Nuclear Society*, 96:212–214, 2007.

[33] P.J.M. van Laarhoven and E.H.L. Aarts. *Simulated Annealing: Theory and Applications*. Mathematics and Its Applications. D. Reidel Publishing Company, Dordrecht, Holland, 1987.

[34] RA Wigeland, TH Bauer, TH Fanning, and EE Morris. Spent nuclear fuel separations and transmutation criteria for benefit to a geologic repository. *Waste Management*, 4, 2004.

[35] R.A. Wigeland, EE Morris, and TH Bauer. Criteria derived for geologic disposal concepts. In *Ninth Information Exchange Meeting on Actinide and Fission Product Partitioning & Transmutation (9IEMPT), Nîmes, France*, 2006.

[36] S.K. Yee. Nuclear fuel cycle modeling approaches for recycling and transmutation of spent nuclear fuel. Master's thesis, The Ohio State University, 2008.

APPENDICES

Appendix A

VISION Base Case Settings

To simplify the process of specifying an individual fuel cycle deployment scenario, the VISION code incorporates a large number of predefined base-cases. These base-cases cover most of the commonly requested scenarios, ranging from a simple phase-out of the current reactor fleet, to a two-tier evolution from the current state to a system with full spent fuel recycling and breeding reactors. For the current optimization studies a one-tier base-case with fast reactor recycle is used. Initially, only thermal reactors are available and their fuel is not recycled. After the year 2040, fast burner reactors (conversion ratio of 0.5) may be built. These FBRs are initially fueled entirely with material separated from LWR spent fuel (plus depleted uranium, as necessary). Having been sufficiently cooled after irradiation, the FBR spent fuel is recycled back into the reactor; after repeated recycle iterations, fuel composition reaches an equilibrium.

It should be noted that some of the fuel cycle parameters specifed below are specified directly while others are calculated using the direct values. For example, reactor electric power and thermal efficiency are combined to calculate the reactor thermal power. Similarly, reactor thermal power is then combined with capacity factor and average fuel discharge burnup to compute the annual fuel consumption rate. The annual fuel consumption can then be multiplied by the cycle time and the number of batches in the core to give the total reactor fuel load.

The separations technology modeled in this base-case is based on the UREX-1a process for thermal fuel and a pyroprocessing method for fast fuel. The fuel is partitioned between various output streams with the efficiencies listed below in table A.1

Several reactor parameters pertinent to the optimization methodology are listed below in table A.2.

Table A.1: Separations Efficiencies

Thermal Oxide Fuel

Stream	He	FP[1]	U	TRU[2]	^{3}H	C	Kr	Sr	Tc	I	Cs
PNAC[3]	—	—	—	0.999	—	—	—	—	—	—	—
RU[4]	—	—	0.999	—	—	—	—	—	—	—	—
Iodine	—	—	—	—	—	—	—	—	—	0.999	—
Gas	0.999	—	—	—	0.999	—	0.999	—	—	—	—
Tc	—	—	—	—	—	—	—	—	0.749	—	—
Cs, Sr	—	—	—	—	—	—	—	0.999	—	—	0.999
Lanth FP	0.001	1.000	0.001	0.001	0.001	1.000	0.001	0.001	0.001	0.001	0.001
UDS[5]	—	—	—	—	—	—	—	—	0.250	—	—

Fast Metallic Fuel

Stream	He	FP[1]	U	TRU[2]	^{3}H	C	Kr	Sr	Tc	I	Cs
RU PNAC	—	—	0.034	0.999	—	—	—	—	—	—	—
RU	—	—	0.934	—	—	—	—	—	—	—	—
Gas	0.999	—	—	—	0.999	—	0.999	—	—	—	—
Cs, Sr	—	—	—	—	—	—	—	0.999	—	0.999	0.999
Lanth FP	0.001	1.000	—	0.001	0.001	1.000	0.001	0.001	0.001	0.001	0.001
Zr, SS[6]	—	—	0.034	—	—	—	—	—	0.999	—	—

[1] Fission products isotopes include: ^{226}Ra, ^{228}Ra, ^{206}Pb, ^{207}Pb, ^{208}Pb, ^{210}Pb, ^{228}Th, ^{229}Th, ^{230}Th, ^{232}Th, ^{209}Bi, ^{227}Ac, ^{231}Pa

[2] Transuranic isotopes include: ^{237}Np, ^{238}Pu, ^{239}Pu, ^{240}Pu, ^{241}Pu, ^{242}Pu, ^{244}Pu, ^{241}Am, ^{242m}Am, ^{243}Am, ^{242}Cm, ^{243}Cm, ^{244}Cm, ^{245}Cm, ^{246}Cm, ^{247}Cm, ^{248}Cm, ^{250}Cm, ^{249}Cf, ^{250}Cf, ^{251}Cf, ^{252}Cf

[3] Recycled Uranium

[4] Recycled Plutonium, Neptunium, Americium and Curium

[5] Undissolved Solids

[6] Zirconium and stainless steel cladding and structural material

Table A.2: Basic Reactor Data

	LWR	**FBR**
Thermal Power	3088 MWth	1667 MWth
Electric Power	1050 MWe	600 MWe
Efficiency	34%	36%
Capacity Factor	90%	85%
$T_{lifetime}$	60 yr	60 yr
T_{cycle}	1 yr	0.605 yr
N_{batch}	5	6.2
Fuel Form	oxide	metallic
Conversion Ratio	—	0.5
Total Core Mass	100 tonnes	9.3 tonnes
Discharge Burn-up	51 GWD/MT-iHM[2]	131.9 GWD/MT-iHM[3]

[1] Conversion Ratio is defined as the ratio of mass of transuranic isotopes created over those destroyed during fuel irradiation [8]
[2] Legacy fuel from Gen-II reactors has 33 GW-day/tonne burn-up
[3] MT-IHM stands for Metric Ton of Initial Heavy-Metal content.

Appendix B

VISION Fuel Recipes

In order to meet the run-time constraints the VISION model foregos expensive reactor physics calculations in favor of a simpler set of pre-calculated values for reactor parameters such as fuel composition, cycle length, and discharge burnup. For the Fast Burner Reactor, many of these values are calculated in [14], and [13], and subsequently archived in Transmutation Library for Fuel Cycle Systems Assessment [24].

The Fast Burner Reactor model in VISION is based on calculations made for the Advanced Burner Reactor (ABR) as part of the Global Nuclear Energy Partnership (GNEP) program. The ABR was in turn based on the General Electric S-PRISM pool-type sodium-cooled fast breeder reactor. The S-PRISM reactor core is composed of 138 driver fuel assemblies, 97 fertile blanket assemblies, and 126 reflector assemblies, with a total effective core diameter of 2.71m. Each assembly is approximately 2m tall, with 1m of active fuel. Thermal power is nominally 1,000MW (linear scaling to 1,667MWth is assumed for the VISION model). As the ABR is intended to function as a net consumer of transuranic material, its fuel design does away with the fertile blanket assemblies, replacing them instead with active driver fuel.

The TRU conversion ratio (that is, the ratio of TRU material created to consumed) is varied by adjusting the design and composition of the fuel assemblies. In order to produce a lower conversion ratio, one must reduce the fraction of fertile material (^{238}U) within the fuel. This reduction then causes several important changes to the thermal and neutronic properties of the core. First of all, because the ^{238}U makes a smaller contribution to the heat generation rate than the remaining TRU elements. Thus, in order to keep within thermal cladding and fuel melt limits, the fuel pins themselves must be made smaller in diameter with decreasing conversion ratio. Secondly, the displacement of ^{238}U leads to a reduction in in-situ breeding of fissile Pu, which in turn increases the cycle reactivity swing. The increased reactivity swing then places limits on the maximum allowed cycle length, and therefore the number of fuel

batches within the core. Thirdly, by removing uranium from the core, the specific burnup of the fuel is correspondingly increased; however, in these calculations, fast neutron fluence to the cladding, rather than burnup, was the factor that most limited the permissible fuel irradiation levels. Finally, although the displacement of non-fissile uranium from the fuel reduces the overall fuel mass, it greatly increases the enrichment, which heightens nonproliferation and security concerns.

For further details regarding the core design and computational processes involved, please see [14].

Table B.1: `VISION` Fuel Cycle Parameters

Reference	TRU Conversion	Burnup	N_{batch} †	T_{cycle}
	–	GWD/MTiHM	–	EFPD
`VISION` 2.1.9	0.50	176.6*	7.3	157
ANL-AFCI-189	0.00	239.9	9.972	132
	0.25	171.7	8.07	154
	0.50	131.9	6.237	221
	0.75	99.6	6.246	232
	1.00	73.0	3.518	370

* It was discovered after-the-fact that this value was a typographical error. Results produced using this incorrect value are used in Section 3.3 solely to demonstrate the functionality of the optimization algorithm itself.

† N_{batch} values calculated from "Fuel residence time" values listed in `VISION` fuel recipe spreadsheet version 5.

Figure B.1, below, shows the U and TRU fractions of the fresh and spent Advanced LWR fuel that is primarily used within the `VISION` model. Also shown is the composition of the existing legacy LWR fuel (with a nominal 33GWD/MTU burnup).

Figures B.2 through B.5 show the U and TRU portions of the fresh and spent FBR fuel for both startup and equilibrium recycle cores for various conversion ratios. Note that the overall mass of fuel loaded into the reactors will also vary with conversion ratio.

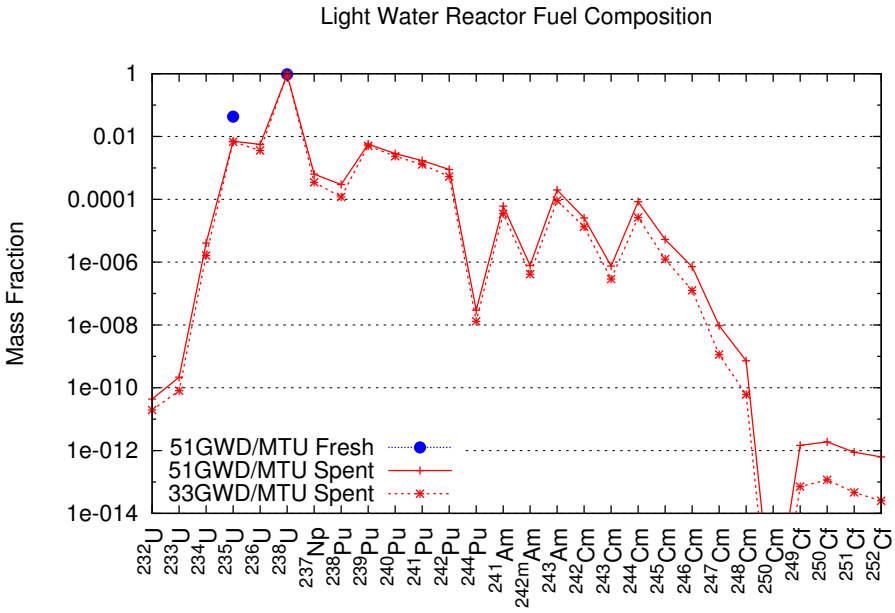

Figure B.1: Partial LWR Fuel Compositions

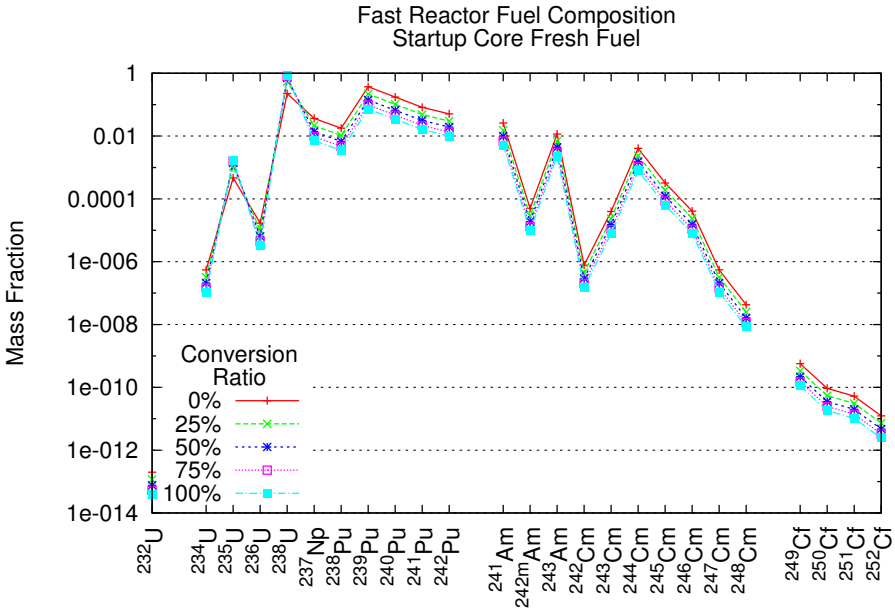

Figure B.2: Partial FBR startup core fresh fuel compositions

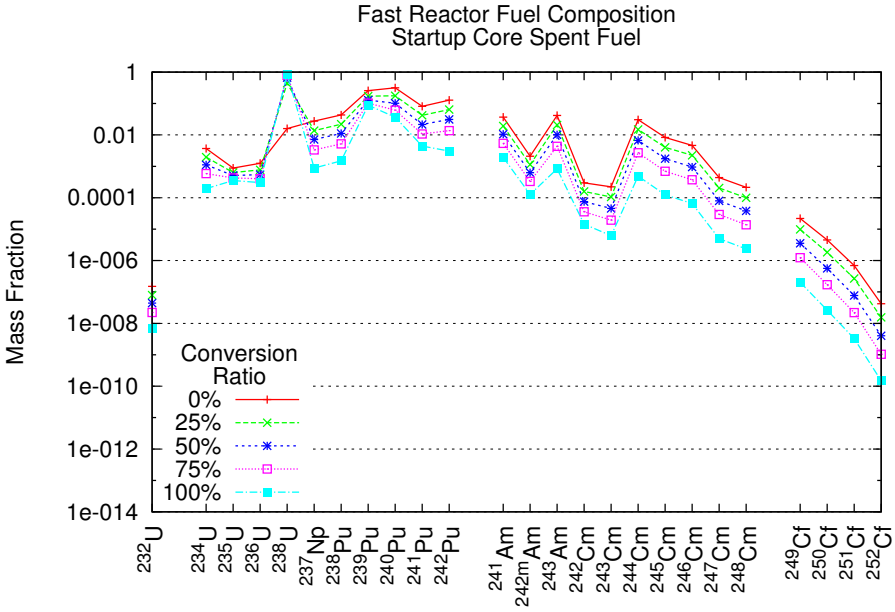

Figure B.3: Partial FBR startup core spent fuel composition

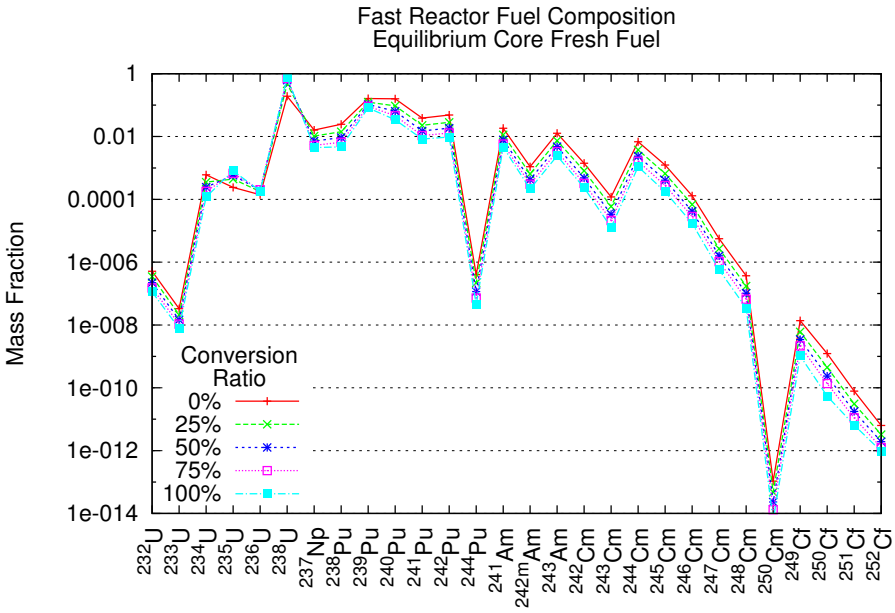

Figure B.4: Partial FBR equilibrium core fresh fuel composition

For completeness, the tables below list in their entirety the fuel recipes used by the `VISION` model.

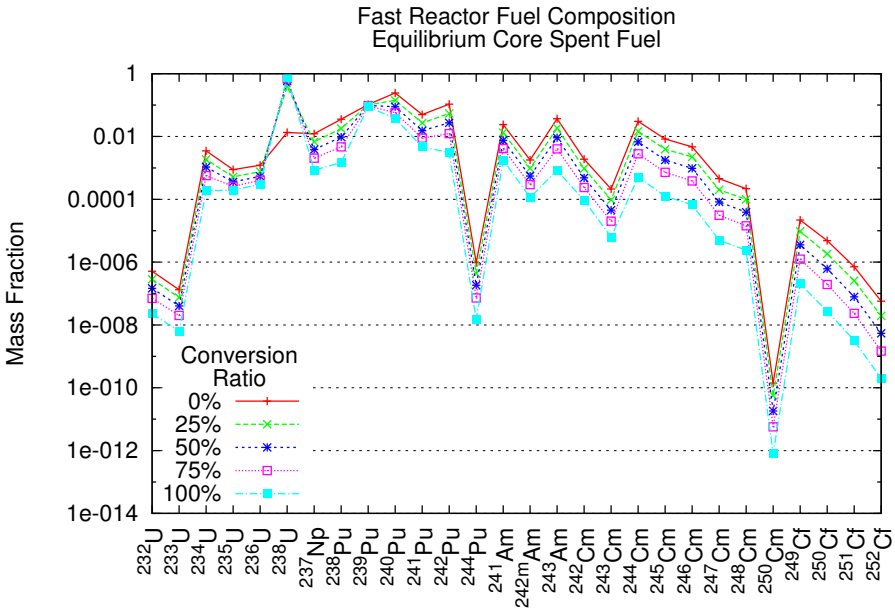

Figure B.5: Partial FBR equilibrium core spent fuel composition

Table B.2: VISION LWR Fuel Recipes by Mass Fraction

Isotope	Fresh 51MWD/MTU	Spent 51MWD/MTU	Legacy 33MWD/MTU
^{232}U	–	4.36×10^{-11}	1.96×10^{-11}
^{233}U	–	2.17×10^{-10}	8.09×10^{-11}
^{234}U	–	4.06×10^{-6}	1.62×10^{-6}
^{235}U	4.28×10^{-2}	7.07×10^{-3}	6.76×10^{-3}
^{236}U	–	5.60×10^{-3}	3.60×10^{-3}
^{238}U	9.57×10^{-1}	9.22×10^{-1}	9.46×10^{-1}
^{237}Np	–	6.33×10^{-4}	3.47×10^{-4}
^{238}Pu	–	2.95×10^{-4}	1.18×10^{-4}
^{239}Pu	–	5.68×10^{-3}	5.04×10^{-3}
^{240}Pu	–	2.87×10^{-3}	2.36×10^{-3}
^{241}Pu	–	1.70×10^{-3}	1.29×10^{-3}
^{242}Pu	–	8.88×10^{-4}	5.32×10^{-4}
^{244}Pu	–	2.96×10^{-8}	1.31×10^{-8}
^{241}Am	–	6.08×10^{-5}	3.48×10^{-5}
^{242m}Am	–	7.82×10^{-7}	4.15×10^{-7}
^{243}Am	–	1.98×10^{-4}	9.07×10^{-5}
^{242}Cm	–	2.55×10^{-5}	1.34×10^{-5}
^{243}Cm	–	7.49×10^{-7}	2.89×10^{-7}
^{244}Cm	–	8.42×10^{-5}	2.65×10^{-5}
^{245}Cm	–	5.27×10^{-6}	1.25×10^{-6}
^{246}Cm	–	7.15×10^{-7}	1.26×10^{-7}
^{247}Cm	–	9.44×10^{-9}	1.14×10^{-9}
^{248}Cm	–	7.23×10^{-10}	6.10×10^{-11}
^{250}Cm	–	4.19×10^{-18}	2.64×10^{-19}
^{249}Cf	–	1.47×10^{-12}	7.11×10^{-14}
^{250}Cf	–	1.89×10^{-12}	1.17×10^{-13}
^{251}Cf	–	8.89×10^{-13}	4.64×10^{-14}
^{252}Cf	–	6.30×10^{-13}	2.52×10^{-14}

Continued on Next Page

Table B.2: (continued)

Isotope	Fresh 51MWD/MTU	Spent 51MWD/MTU	Legacy 33MWD/MTU
^{226}Ra	–	1.94×10^{-16}	4.11×10^{-17}
^{228}Ra	–	2.72×10^{-20}	7.46×10^{-21}
^{206}Pb	–	1.11×10^{-19}	1.47×10^{-20}
^{207}Pb	–	8.20×10^{-17}	1.86×10^{-17}
^{208}Pb	–	2.11×10^{-13}	4.02×10^{-14}
^{210}Pb	–	4.83×10^{-18}	1.26×10^{-18}
^{228}Th	–	4.20×10^{-13}	1.38×10^{-13}
^{229}Th	–	1.40×10^{-13}	3.23×10^{-14}
^{230}Th	–	1.58×10^{-11}	4.94×10^{-12}
^{232}Th	–	4.13×10^{-10}	1.70×10^{-10}
^{209}Bi	–	3.78×10^{-17}	7.13×10^{-18}
^{227}Ac	–	1.21×10^{-15}	5.20×10^{-16}
^{231}Pa	–	2.48×10^{-11}	1.66×10^{-11}
^{3}H	–	8.59×10^{-8}	5.81×10^{-8}
^{4}He	–	9.29×10^{-7}	2.80×10^{-7}
^{14}C	–	4.07×10^{-11}	2.64×10^{-11}
C*	–	–	–
^{81}Kr	–	4.24×10^{-11}	2.35×10^{-11}
^{85}Kr	–	3.46×10^{-5}	2.35×10^{-5}
Kr*	–	5.63×10^{-4}	4.18×10^{-4}
^{90}Sr	–	7.92×10^{-4}	5.23×10^{-4}
Sr*	–	5.76×10^{-4}	3.64×10^{-4}
^{99}Tc	–	1.15×10^{-3}	7.80×10^{-4}
Tc*	–	7.87×10^{-4}	1.04×10^{-3}
^{129}I	–	2.72×10^{-4}	1.82×10^{-4}
I*	–	2.44×10^{-5}	4.52×10^{-7}
^{134}Cs	–	2.28×10^{-4}	1.24×10^{-4}
^{135}Cs	–	6.14×10^{-4}	3.40×10^{-4}
^{137}Cs	–	1.82×10^{-3}	1.21×10^{-3}
Cs*	–	1.62×10^{-3}	1.12×10^{-3}
FP†	–	4.04×10^{-2}	3.07×10^{-2}
Act‡	–	8.73×10^{-5}	1.18×10^{-4}

* Non-enumerated isotopes
† Non-enumerated Fission Products
‡ Non-enumerated Actinides

Table B.3: VISION 0% Conversion Ratio FBR Fuel Recipes by Mass Fraction

	Fresh					Spent				
Isotope	Pass 1	Pass 2	Pass 3	Pass 4	Pass 5	Pass 1	Pass 2	Pass 3	Pass 4	Pass 5
^{232}U	1.99×10^{-13}	3.77×10^{-8}	7.53×10^{-8}	1.13×10^{-7}	1.51×10^{-7}	5.14×10^{-7}	5.14×10^{-7}	5.13×10^{-7}	5.13×10^{-7}	5.13×10^{-7}
^{233}U	–	–	–	–	–	3.31×10^{-8}	5.79×10^{-8}	8.27×10^{-8}	1.08×10^{-7}	1.32×10^{-7}
^{234}U	5.51×10^{-7}	9.29×10^{-4}	1.86×10^{-3}	2.79×10^{-3}	3.71×10^{-3}	6.06×10^{-4}	1.33×10^{-3}	2.05×10^{-3}	2.77×10^{-3}	3.49×10^{-3}
^{235}U	4.64×10^{-4}	5.69×10^{-4}	6.73×10^{-4}	7.78×10^{-4}	8.82×10^{-4}	2.40×10^{-4}	4.00×10^{-4}	5.60×10^{-4}	7.20×10^{-4}	8.81×10^{-4}
^{236}U	1.69×10^{-5}	3.26×10^{-4}	6.34×10^{-4}	9.43×10^{-4}	1.25×10^{-3}	1.40×10^{-4}	4.11×10^{-4}	6.82×10^{-4}	9.54×10^{-4}	1.22×10^{-3}
^{238}U	2.26×10^{-1}	1.73×10^{-1}	1.21×10^{-1}	6.83×10^{-2}	1.59×10^{-2}	1.92×10^{-1}	1.48×10^{-1}	1.03×10^{-1}	5.82×10^{-2}	1.35×10^{-2}
^{237}Np	3.66×10^{-2}	3.43×10^{-2}	3.21×10^{-2}	2.98×10^{-2}	2.76×10^{-2}	1.61×10^{-2}	1.52×10^{-2}	1.42×10^{-2}	1.32×10^{-2}	1.22×10^{-2}
^{238}Pu	1.78×10^{-2}	2.41×10^{-2}	3.05×10^{-2}	3.68×10^{-2}	4.32×10^{-2}	2.47×10^{-2}	2.74×10^{-2}	3.00×10^{-2}	3.27×10^{-2}	3.54×10^{-2}
^{239}Pu	3.70×10^{-1}	3.42×10^{-1}	3.14×10^{-1}	2.86×10^{-1}	2.58×10^{-1}	1.62×10^{-1}	1.48×10^{-1}	1.33×10^{-1}	1.19×10^{-1}	1.04×10^{-1}
^{240}Pu	1.74×10^{-1}	2.10×10^{-1}	2.45×10^{-1}	2.81×10^{-1}	3.16×10^{-1}	1.59×10^{-1}	1.80×10^{-1}	2.01×10^{-1}	2.22×10^{-1}	2.44×10^{-1}
^{241}Pu	8.23×10^{-2}	8.20×10^{-2}	8.17×10^{-2}	8.14×10^{-2}	8.11×10^{-2}	3.88×10^{-2}	4.16×10^{-2}	4.45×10^{-2}	4.73×10^{-2}	5.01×10^{-2}
^{242}Pu	5.09×10^{-2}	7.01×10^{-2}	8.92×10^{-2}	1.08×10^{-1}	1.28×10^{-1}	4.86×10^{-2}	6.32×10^{-2}	7.78×10^{-2}	9.24×10^{-2}	1.07×10^{-1}
^{244}Pu	–	–	–	–	–	4.03×10^{-7}	5.43×10^{-7}	6.84×10^{-7}	8.24×10^{-7}	9.65×10^{-7}
^{241}Am	2.61×10^{-2}	2.88×10^{-2}	3.15×10^{-2}	3.42×10^{-2}	3.70×10^{-2}	1.84×10^{-2}	1.98×10^{-2}	2.11×10^{-2}	2.25×10^{-2}	2.39×10^{-2}
^{242m}Am	4.95×10^{-5}	5.63×10^{-4}	1.08×10^{-3}	1.59×10^{-3}	2.10×10^{-3}	1.10×10^{-3}	1.27×10^{-3}	1.44×10^{-3}	1.61×10^{-3}	1.79×10^{-3}
^{243}Am	1.16×10^{-2}	1.90×10^{-2}	2.65×10^{-2}	3.40×10^{-2}	4.15×10^{-2}	1.27×10^{-2}	1.87×10^{-2}	2.47×10^{-2}	3.08×10^{-2}	3.68×10^{-2}
^{242}Cm	7.75×10^{-7}	7.52×10^{-5}	1.50×10^{-4}	2.24×10^{-4}	2.98×10^{-4}	1.41×10^{-3}	1.53×10^{-3}	1.65×10^{-3}	1.77×10^{-3}	1.89×10^{-3}
^{243}Cm	3.96×10^{-5}	8.52×10^{-5}	1.31×10^{-4}	1.76×10^{-4}	2.22×10^{-4}	1.20×10^{-4}	1.43×10^{-4}	1.66×10^{-4}	1.90×10^{-4}	2.13×10^{-4}
^{244}Cm	4.06×10^{-3}	1.07×10^{-2}	1.73×10^{-2}	2.39×10^{-2}	3.05×10^{-2}	6.85×10^{-3}	1.27×10^{-2}	1.86×10^{-2}	2.45×10^{-2}	3.04×10^{-2}
^{245}Cm	3.23×10^{-4}	2.32×10^{-3}	4.32×10^{-3}	6.32×10^{-3}	8.32×10^{-3}	1.24×10^{-3}	3.02×10^{-3}	4.79×10^{-3}	6.57×10^{-3}	8.35×10^{-3}
^{246}Cm	4.07×10^{-5}	1.19×10^{-3}	2.34×10^{-3}	3.48×10^{-3}	4.63×10^{-3}	1.30×10^{-4}	1.27×10^{-3}	2.40×10^{-3}	3.54×10^{-3}	4.67×10^{-3}
^{247}Cm	5.50×10^{-7}	1.10×10^{-4}	2.20×10^{-4}	3.30×10^{-4}	4.40×10^{-4}	5.60×10^{-6}	1.18×10^{-4}	2.30×10^{-4}	3.43×10^{-4}	4.55×10^{-4}
^{248}Cm	4.28×10^{-8}	5.39×10^{-5}	1.08×10^{-4}	1.61×10^{-4}	2.15×10^{-4}	3.70×10^{-7}	5.50×10^{-5}	1.10×10^{-4}	1.64×10^{-4}	2.19×10^{-4}
^{250}Cm	–	–	–	–	–	1.07×10^{-13}	3.52×10^{-11}	7.03×10^{-11}	1.05×10^{-10}	1.41×10^{-10}
^{249}Cf	5.67×10^{-10}	5.47×10^{-6}	1.09×10^{-5}	1.64×10^{-5}	2.19×10^{-5}	1.37×10^{-8}	5.45×10^{-6}	1.09×10^{-5}	1.63×10^{-5}	2.18×10^{-5}
^{250}Cf	9.25×10^{-11}	1.13×10^{-6}	2.26×10^{-6}	3.39×10^{-6}	4.52×10^{-6}	1.23×10^{-9}	1.21×10^{-6}	2.42×10^{-6}	3.63×10^{-6}	4.84×10^{-6}
^{251}Cf	5.24×10^{-11}	1.74×10^{-7}	3.48×10^{-7}	5.22×10^{-7}	6.96×10^{-7}	7.81×10^{-11}	1.78×10^{-7}	3.55×10^{-7}	5.33×10^{-7}	7.11×10^{-7}
^{252}Cf	1.24×10^{-11}	1.05×10^{-8}	2.10×10^{-8}	3.16×10^{-8}	4.21×10^{-8}	6.35×10^{-12}	1.42×10^{-8}	2.84×10^{-8}	4.26×10^{-8}	5.68×10^{-8}

Continued on Next Page

Table B.3: (continued)

	Fresh					Spent				
Isotope	Pass 1	Pass 2	Pass 3	Pass 4	Pass 5	Pass 1	Pass 2	Pass 3	Pass 4	Pass 5
^{226}Ra	–	–	–	–	–	3.98×10^{-14}	1.99×10^{-13}	3.59×10^{-13}	5.18×10^{-13}	6.78×10^{-13}
^{228}Ra	–	–	–	–	–	6.86×10^{-22}	3.59×10^{-21}	6.50×10^{-21}	9.41×10^{-21}	1.23×10^{-20}
^{206}Pb	–	–	–	–	–	4.95×10^{-19}	3.88×10^{-18}	7.27×10^{-18}	1.07×10^{-17}	1.40×10^{-17}
^{207}Pb	–	–	–	–	–	1.11×10^{-16}	8.69×10^{-16}	1.63×10^{-15}	2.39×10^{-15}	3.14×10^{-15}
^{208}Pb	–	–	–	–	–	9.90×10^{-10}	1.08×10^{-9}	1.17×10^{-9}	1.25×10^{-9}	1.34×10^{-9}
^{210}Pb	–	–	–	–	–	1.72×10^{-17}	1.10×10^{-16}	2.03×10^{-16}	2.96×10^{-16}	3.89×10^{-16}
^{228}Th	–	–	–	–	–	2.12×10^{-9}	2.22×10^{-9}	2.32×10^{-9}	2.42×10^{-9}	2.52×10^{-9}
^{229}Th	–	–	–	–	–	8.91×10^{-11}	9.66×10^{-11}	1.04×10^{-10}	1.12×10^{-10}	1.19×10^{-10}
^{230}Th	–	–	–	–	–	3.20×10^{-9}	1.12×10^{-8}	1.92×10^{-8}	2.72×10^{-8}	3.52×10^{-8}
^{232}Th	–	–	–	–	–	9.73×10^{-12}	4.12×10^{-11}	7.28×10^{-11}	1.04×10^{-10}	1.36×10^{-10}
^{209}Bi	–	–	–	–	–	6.44×10^{-13}	7.06×10^{-13}	7.67×10^{-13}	8.29×10^{-13}	8.90×10^{-13}
^{227}Ac	–	–	–	–	–	3.77×10^{-15}	2.43×10^{-14}	4.49×10^{-14}	6.54×10^{-14}	8.59×10^{-14}
^{231}Pa	–	–	–	–	–	1.67×10^{-10}	8.29×10^{-10}	1.49×10^{-9}	2.15×10^{-9}	2.81×10^{-9}
^{3}H	–	–	–	–	–	7.25×10^{-7}	7.03×10^{-7}	6.82×10^{-7}	6.60×10^{-7}	6.39×10^{-7}
^{4}He	–	–	–	–	–	1.83×10^{-4}	2.17×10^{-4}	2.51×10^{-4}	2.85×10^{-4}	3.19×10^{-4}
^{14}C	–	–	–	–	–	–	–	–	–	–
C*	–	–	–	–	–	–	–	–	–	–
^{81}Kr	–	–	–	–	–	1.43×10^{-10}	1.39×10^{-10}	1.36×10^{-10}	1.32×10^{-10}	1.28×10^{-10}
^{85}Kr	–	–	–	–	–	1.23×10^{-4}	1.19×10^{-4}	1.16×10^{-4}	1.12×10^{-4}	1.08×10^{-4}
Kr*	–	–	–	–	–	1.78×10^{-3}	1.73×10^{-3}	1.67×10^{-3}	1.62×10^{-3}	1.56×10^{-3}
^{90}Sr	–	–	–	–	–	2.17×10^{-3}	2.10×10^{-3}	2.03×10^{-3}	1.97×10^{-3}	1.90×10^{-3}
Sr*	–	–	–	–	–	1.56×10^{-3}	1.51×10^{-3}	1.46×10^{-3}	1.41×10^{-3}	1.37×10^{-3}
^{99}Tc	–	–	–	–	–	6.66×10^{-3}	6.46×10^{-3}	6.26×10^{-3}	6.06×10^{-3}	5.86×10^{-3}
Tc*	–	–	–	–	–	1.16×10^{-6}	1.13×10^{-6}	1.11×10^{-6}	1.08×10^{-6}	1.06×10^{-6}
^{129}I	–	–	–	–	–	2.17×10^{-3}	2.11×10^{-3}	2.04×10^{-3}	1.98×10^{-3}	1.92×10^{-3}
I*	–	–	–	–	–	7.88×10^{-4}	7.66×10^{-4}	7.44×10^{-4}	7.22×10^{-4}	7.00×10^{-4}
^{134}Cs	–	–	–	–	–	3.08×10^{-4}	2.99×10^{-4}	2.91×10^{-4}	2.83×10^{-4}	2.74×10^{-4}
^{135}Cs	–	–	–	–	–	1.17×10^{-2}	1.13×10^{-2}	1.10×10^{-2}	1.06×10^{-2}	1.03×10^{-2}
^{137}Cs	–	–	–	–	–	1.02×10^{-2}	9.94×10^{-3}	9.64×10^{-3}	9.33×10^{-3}	9.03×10^{-3}
Cs*	–	–	–	–	–	1.02×10^{-2}	9.89×10^{-3}	9.58×10^{-3}	9.27×10^{-3}	8.97×10^{-3}
FP†	–	–	–	–	–	2.68×10^{-1}	2.70×10^{-1}	2.73×10^{-1}	2.75×10^{-1}	2.77×10^{-1}
Act‡	–	–	–	–	–	–	–	–	–	–

* Non-enumerated isotopes
† Non-enumerated Fission Products
‡ Non-enumerated Actinides

Table B.4: VISION 25% Conversion Ratio FBR Fuel Recipes by Mass Fraction

	Fresh					Spent				
Isotope	Pass 1	Pass 2	Pass 3	Pass 4	Pass 5	Pass 1	Pass 2	Pass 3	Pass 4	Pass 5
^{232}U	1.17×10^{-13}	1.98×10^{-8}	3.96×10^{-8}	5.94×10^{-8}	7.92×10^{-8}	3.48×10^{-7}	3.32×10^{-7}	3.17×10^{-7}	3.01×10^{-7}	2.86×10^{-7}
^{233}U	–	–	–	–	–	2.10×10^{-8}	3.44×10^{-8}	4.78×10^{-8}	6.12×10^{-8}	7.47×10^{-8}
^{234}U	3.24×10^{-7}	4.93×10^{-4}	9.87×10^{-4}	1.48×10^{-3}	1.97×10^{-3}	3.52×10^{-4}	7.29×10^{-4}	1.11×10^{-3}	1.48×10^{-3}	1.86×10^{-3}
^{235}U	1.09×10^{-3}	9.72×10^{-4}	8.59×10^{-4}	7.45×10^{-4}	6.32×10^{-4}	4.43×10^{-4}	4.62×10^{-4}	4.81×10^{-4}	5.00×10^{-4}	5.20×10^{-4}
^{236}U	9.95×10^{-6}	2.01×10^{-4}	3.93×10^{-4}	5.84×10^{-4}	7.76×10^{-4}	1.78×10^{-4}	3.23×10^{-4}	4.69×10^{-4}	6.15×10^{-4}	7.61×10^{-4}
^{238}U	5.44×10^{-1}	5.19×10^{-1}	4.94×10^{-1}	4.69×10^{-1}	4.44×10^{-1}	4.75×10^{-1}	4.53×10^{-1}	4.31×10^{-1}	4.09×10^{-1}	3.87×10^{-1}
^{237}Np	2.15×10^{-2}	1.96×10^{-2}	1.76×10^{-2}	1.57×10^{-2}	1.37×10^{-2}	1.03×10^{-2}	9.40×10^{-3}	8.51×10^{-3}	7.62×10^{-3}	6.74×10^{-3}
^{238}Pu	1.05×10^{-2}	1.33×10^{-2}	1.62×10^{-2}	1.91×10^{-2}	2.19×10^{-2}	1.42×10^{-2}	1.51×10^{-2}	1.61×10^{-2}	1.71×10^{-2}	1.80×10^{-2}
^{239}Pu	2.17×10^{-1}	2.06×10^{-1}	1.95×10^{-1}	1.83×10^{-1}	1.72×10^{-1}	1.25×10^{-1}	1.19×10^{-1}	1.13×10^{-1}	1.07×10^{-1}	1.01×10^{-1}
^{240}Pu	1.03×10^{-1}	1.21×10^{-1}	1.39×10^{-1}	1.58×10^{-1}	1.76×10^{-1}	9.46×10^{-2}	1.06×10^{-1}	1.18×10^{-1}	1.30×10^{-1}	1.42×10^{-1}
^{241}Pu	4.84×10^{-2}	4.68×10^{-2}	4.51×10^{-2}	4.35×10^{-2}	4.19×10^{-2}	2.28×10^{-2}	2.39×10^{-2}	2.49×10^{-2}	2.60×10^{-2}	2.71×10^{-2}
^{242}Pu	3.00×10^{-2}	3.85×10^{-2}	4.70×10^{-2}	5.56×10^{-2}	6.41×10^{-2}	2.84×10^{-2}	3.49×10^{-2}	4.14×10^{-2}	4.79×10^{-2}	5.44×10^{-2}
^{244}Pu	–	–	–	–	–	2.19×10^{-7}	2.73×10^{-7}	3.28×10^{-7}	3.82×10^{-7}	4.37×10^{-7}
^{241}Am	1.54×10^{-2}	1.63×10^{-2}	1.73×10^{-2}	1.83×10^{-2}	1.93×10^{-2}	1.14×10^{-2}	1.18×10^{-2}	1.23×10^{-2}	1.28×10^{-2}	1.33×10^{-2}
^{242m}Am	2.91×10^{-5}	3.06×10^{-4}	5.82×10^{-4}	8.59×10^{-4}	1.14×10^{-3}	6.46×10^{-4}	7.27×10^{-4}	8.08×10^{-4}	8.89×10^{-4}	9.71×10^{-4}
^{243}Am	6.81×10^{-3}	1.03×10^{-2}	1.38×10^{-2}	1.73×10^{-2}	2.08×10^{-2}	7.33×10^{-3}	1.01×10^{-2}	1.30×10^{-2}	1.58×10^{-2}	1.86×10^{-2}
^{242}Cm	4.56×10^{-7}	4.01×10^{-5}	7.98×10^{-5}	1.19×10^{-4}	1.59×10^{-4}	8.11×10^{-4}	8.47×10^{-4}	8.83×10^{-4}	9.19×10^{-4}	9.55×10^{-4}
^{243}Cm	2.33×10^{-5}	4.39×10^{-5}	6.46×10^{-5}	8.52×10^{-5}	1.06×10^{-4}	5.92×10^{-5}	6.84×10^{-5}	7.75×10^{-5}	8.67×10^{-5}	9.59×10^{-5}
^{244}Cm	2.39×10^{-3}	5.53×10^{-3}	8.67×10^{-3}	1.18×10^{-2}	1.49×10^{-2}	3.77×10^{-3}	6.47×10^{-3}	9.17×10^{-3}	1.19×10^{-2}	1.46×10^{-2}
^{245}Cm	1.90×10^{-4}	1.15×10^{-3}	2.11×10^{-3}	3.06×10^{-3}	4.02×10^{-3}	6.56×10^{-4}	1.47×10^{-3}	2.29×10^{-3}	3.10×10^{-3}	3.92×10^{-3}
^{246}Cm	2.40×10^{-5}	5.80×10^{-4}	1.14×10^{-3}	1.69×10^{-3}	2.25×10^{-3}	6.84×10^{-5}	6.13×10^{-4}	1.16×10^{-3}	1.70×10^{-3}	2.25×10^{-3}
^{247}Cm	3.24×10^{-7}	5.12×10^{-5}	1.02×10^{-4}	1.53×10^{-4}	2.04×10^{-4}	2.68×10^{-6}	5.22×10^{-5}	1.02×10^{-4}	1.51×10^{-4}	2.01×10^{-4}
^{248}Cm	2.52×10^{-8}	2.51×10^{-5}	5.01×10^{-5}	7.52×10^{-5}	1.00×10^{-4}	1.74×10^{-7}	2.53×10^{-5}	5.03×10^{-5}	7.54×10^{-5}	1.00×10^{-4}
^{250}Cm	–	–	–	–	–	4.72×10^{-14}	1.41×10^{-11}	2.81×10^{-11}	4.22×10^{-11}	5.62×10^{-11}
^{249}Cf	3.34×10^{-10}	2.46×10^{-6}	4.92×10^{-6}	7.38×10^{-6}	9.84×10^{-6}	6.07×10^{-9}	2.39×10^{-6}	4.77×10^{-6}	7.14×10^{-6}	9.52×10^{-6}
^{250}Cf	5.44×10^{-11}	4.60×10^{-7}	9.19×10^{-7}	1.38×10^{-6}	1.84×10^{-6}	4.50×10^{-10}	4.61×10^{-7}	9.22×10^{-7}	1.38×10^{-6}	1.84×10^{-6}
^{251}Cf	3.08×10^{-11}	6.69×10^{-8}	1.34×10^{-7}	2.01×10^{-7}	2.67×10^{-7}	3.07×10^{-11}	6.38×10^{-8}	1.28×10^{-7}	1.91×10^{-7}	2.55×10^{-7}
^{252}Cf	7.30×10^{-12}	3.91×10^{-9}	7.82×10^{-9}	1.17×10^{-8}	1.56×10^{-8}	3.23×10^{-12}	4.89×10^{-9}	9.77×10^{-9}	1.47×10^{-8}	1.95×10^{-8}

Continued on Next Page

Table B.4: (continued)

	Fresh					Spent				
Isotope	Pass 1	Pass 2	Pass 3	Pass 4	Pass 5	Pass 1	Pass 2	Pass 3	Pass 4	Pass 5
^{226}Ra	–	–	–	–	–	2.50×10^{-14}	1.03×10^{-13}	1.81×10^{-13}	2.59×10^{-13}	3.37×10^{-13}
^{228}Ra	–	–	–	–	–	8.00×10^{-22}	2.37×10^{-21}	3.93×10^{-21}	5.50×10^{-21}	7.06×10^{-21}
^{206}Pb	–	–	–	–	–	3.66×10^{-19}	2.03×10^{-18}	3.70×10^{-18}	5.37×10^{-18}	7.04×10^{-18}
^{207}Pb	–	–	–	–	–	7.50×10^{-17}	4.39×10^{-16}	8.03×10^{-16}	1.17×10^{-15}	1.53×10^{-15}
^{208}Pb	–	–	–	–	–	6.71×10^{-10}	6.94×10^{-10}	7.17×10^{-10}	7.40×10^{-10}	7.63×10^{-10}
^{210}Pb	–	–	–	–	–	1.16×10^{-17}	5.67×10^{-17}	1.02×10^{-16}	1.47×10^{-16}	1.92×10^{-16}
^{228}Th	–	–	–	–	–	1.40×10^{-9}	1.39×10^{-9}	1.39×10^{-9}	1.39×10^{-9}	1.38×10^{-9}
^{229}Th	–	–	–	–	–	5.82×10^{-11}	5.96×10^{-11}	6.11×10^{-11}	6.25×10^{-11}	6.40×10^{-11}
^{230}Th	–	–	–	–	–	1.91×10^{-9}	5.92×10^{-9}	9.93×10^{-9}	1.39×10^{-8}	1.80×10^{-8}
^{232}Th	–	–	–	–	–	1.18×10^{-11}	2.89×10^{-11}	4.60×10^{-11}	6.31×10^{-11}	8.02×10^{-11}
^{209}Bi	–	–	–	–	–	4.16×10^{-13}	4.29×10^{-13}	4.43×10^{-13}	4.56×10^{-13}	4.69×10^{-13}
^{227}Ac	–	–	–	–	–	2.40×10^{-15}	1.25×10^{-14}	2.27×10^{-14}	3.28×10^{-14}	4.30×10^{-14}
^{231}Pa	–	–	–	–	–	1.00×10^{-10}	4.38×10^{-10}	7.75×10^{-10}	1.11×10^{-9}	1.45×10^{-9}
^{3}H	–	–	–	–	–	4.77×10^{-7}	4.68×10^{-7}	4.58×10^{-7}	4.48×10^{-7}	4.38×10^{-7}
^{4}He	–	–	–	–	–	1.01×10^{-4}	1.15×10^{-4}	1.29×10^{-4}	1.42×10^{-4}	1.56×10^{-4}
^{14}C	–	–	–	–	–	–	–	–	–	–
C*	–	–	–	–	–	–	–	–	–	–
^{81}Kr	–	–	–	–	–	8.81×10^{-11}	8.62×10^{-11}	8.43×10^{-11}	8.24×10^{-11}	8.05×10^{-11}
^{85}Kr	–	–	–	–	–	8.18×10^{-5}	8.01×10^{-5}	7.84×10^{-5}	7.67×10^{-5}	7.50×10^{-5}
Kr*	–	–	–	–	–	1.18×10^{-3}	1.16×10^{-3}	1.13×10^{-3}	1.11×10^{-3}	1.08×10^{-3}
^{90}Sr	–	–	–	–	–	1.46×10^{-3}	1.43×10^{-3}	1.40×10^{-3}	1.37×10^{-3}	1.33×10^{-3}
Sr*	–	–	–	–	–	1.06×10^{-3}	1.03×10^{-3}	1.01×10^{-3}	9.89×10^{-4}	9.66×10^{-4}
^{99}Tc	–	–	–	–	–	4.38×10^{-3}	4.29×10^{-3}	4.20×10^{-3}	4.11×10^{-3}	4.01×10^{-3}
Tc*	–	–	–	–	–	8.84×10^{-7}	8.74×10^{-7}	8.64×10^{-7}	8.54×10^{-7}	8.43×10^{-7}
^{129}I	–	–	–	–	–	1.41×10^{-3}	1.38×10^{-3}	1.35×10^{-3}	1.32×10^{-3}	1.29×10^{-3}
I*	–	–	–	–	–	5.06×10^{-4}	4.96×10^{-4}	4.86×10^{-4}	4.76×10^{-4}	4.66×10^{-4}
^{134}Cs	–	–	–	–	–	1.99×10^{-4}	1.94×10^{-4}	1.90×10^{-4}	1.85×10^{-4}	1.81×10^{-4}
^{135}Cs	–	–	–	–	–	7.63×10^{-3}	7.47×10^{-3}	7.31×10^{-3}	7.15×10^{-3}	6.99×10^{-3}
^{137}Cs	–	–	–	–	–	6.72×10^{-3}	6.58×10^{-3}	6.43×10^{-3}	6.29×10^{-3}	6.15×10^{-3}
Cs*	–	–	–	–	–	6.67×10^{-3}	6.53×10^{-3}	6.39×10^{-3}	6.25×10^{-3}	6.12×10^{-3}
FP†	–	–	–	–	–	1.73×10^{-1}	1.74×10^{-1}	1.75×10^{-1}	1.76×10^{-1}	1.76×10^{-1}
Act‡	–	–	–	–	–	–	–	–	–	–

* Non-enumerated isotopes
† Non-enumerated Fission Products
‡ Non-enumerated Actinides

Table B.5: VISION 50% Conversion Ratio FBR Fuel Recipes by Mass Fraction

	Fresh					Spent				
Isotope	Pass 1	Pass 2	Pass 3	Pass 4	Pass 5	Pass 1	Pass 2	Pass 3	Pass 4	Pass 5
^{232}U	7.56×10^{-14}	1.11×10^{-8}	2.21×10^{-8}	3.32×10^{-8}	4.42×10^{-8}	2.21×10^{-7}	2.03×10^{-7}	1.84×10^{-7}	1.66×10^{-7}	1.47×10^{-7}
^{233}U	–	–	–	–	–	1.47×10^{-8}	2.11×10^{-8}	2.75×10^{-8}	3.39×10^{-8}	4.02×10^{-8}
^{234}U	2.10×10^{-7}	2.79×10^{-4}	5.58×10^{-4}	8.38×10^{-4}	1.12×10^{-3}	2.45×10^{-4}	4.49×10^{-4}	6.54×10^{-4}	8.59×10^{-4}	1.06×10^{-3}
^{235}U	1.40×10^{-3}	1.18×10^{-3}	9.54×10^{-4}	7.31×10^{-4}	5.08×10^{-4}	5.82×10^{-4}	5.26×10^{-4}	4.70×10^{-4}	4.14×10^{-4}	3.58×10^{-4}
^{236}U	6.43×10^{-6}	1.41×10^{-4}	2.77×10^{-4}	4.12×10^{-4}	5.47×10^{-4}	1.95×10^{-4}	2.81×10^{-4}	3.68×10^{-4}	4.54×10^{-4}	5.40×10^{-4}
^{238}U	7.04×10^{-1}	6.95×10^{-1}	6.85×10^{-1}	6.75×10^{-1}	6.65×10^{-1}	6.22×10^{-1}	6.13×10^{-1}	6.04×10^{-1}	5.95×10^{-1}	5.86×10^{-1}
^{237}Np	1.39×10^{-2}	1.22×10^{-2}	1.05×10^{-2}	8.81×10^{-3}	7.12×10^{-3}	7.16×10^{-3}	6.33×10^{-3}	5.49×10^{-3}	4.65×10^{-3}	3.82×10^{-3}
^{238}Pu	6.76×10^{-3}	7.84×10^{-3}	8.91×10^{-3}	9.99×10^{-3}	1.11×10^{-2}	9.32×10^{-3}	9.38×10^{-3}	9.44×10^{-3}	9.51×10^{-3}	9.57×10^{-3}
^{239}Pu	1.41×10^{-1}	1.38×10^{-1}	1.35×10^{-1}	1.31×10^{-1}	1.28×10^{-1}	1.05×10^{-1}	1.03×10^{-1}	1.01×10^{-1}	9.95×10^{-2}	9.76×10^{-2}
^{240}Pu	6.63×10^{-2}	7.53×10^{-2}	8.43×10^{-2}	9.33×10^{-2}	1.02×10^{-1}	6.37×10^{-2}	6.99×10^{-2}	7.61×10^{-2}	8.23×10^{-2}	8.86×10^{-2}
^{241}Pu	3.13×10^{-2}	2.89×10^{-2}	2.64×10^{-2}	2.40×10^{-2}	2.15×10^{-2}	1.51×10^{-2}	1.52×10^{-2}	1.53×10^{-2}	1.53×10^{-2}	1.54×10^{-2}
^{242}Pu	1.94×10^{-2}	2.23×10^{-2}	2.53×10^{-2}	2.83×10^{-2}	3.12×10^{-2}	1.86×10^{-2}	2.07×10^{-2}	2.29×10^{-2}	2.50×10^{-2}	2.72×10^{-2}
^{244}Pu	–	–	–	–	–	1.16×10^{-7}	1.33×10^{-7}	1.49×10^{-7}	1.66×10^{-7}	1.83×10^{-7}
^{241}Am	9.93×10^{-3}	1.01×10^{-2}	1.02×10^{-2}	1.04×10^{-2}	1.05×10^{-2}	7.89×10^{-3}	7.81×10^{-3}	7.74×10^{-3}	7.66×10^{-3}	7.59×10^{-3}
^{242m}Am	1.88×10^{-5}	1.71×10^{-4}	3.24×10^{-4}	4.76×10^{-4}	6.29×10^{-4}	4.27×10^{-4}	4.57×10^{-4}	4.88×10^{-4}	5.18×10^{-4}	5.48×10^{-4}
^{243}Am	4.40×10^{-3}	5.81×10^{-3}	7.23×10^{-3}	8.64×10^{-3}	1.01×10^{-2}	4.75×10^{-3}	5.85×10^{-3}	6.94×10^{-3}	8.04×10^{-3}	9.13×10^{-3}
^{242}Cm	2.94×10^{-7}	1.90×10^{-5}	3.77×10^{-5}	5.65×10^{-5}	7.52×10^{-5}	4.83×10^{-4}	4.83×10^{-4}	4.83×10^{-4}	4.83×10^{-4}	4.83×10^{-4}
^{243}Cm	1.51×10^{-5}	2.27×10^{-5}	3.04×10^{-5}	3.80×10^{-5}	4.57×10^{-5}	3.36×10^{-5}	3.64×10^{-5}	3.93×10^{-5}	4.22×10^{-5}	4.51×10^{-5}
^{244}Cm	1.54×10^{-3}	2.84×10^{-3}	4.13×10^{-3}	5.42×10^{-3}	6.72×10^{-3}	2.37×10^{-3}	3.47×10^{-3}	4.58×10^{-3}	5.68×10^{-3}	6.78×10^{-3}
^{245}Cm	1.23×10^{-4}	5.30×10^{-4}	9.36×10^{-4}	1.34×10^{-3}	1.75×10^{-3}	4.06×10^{-4}	7.48×10^{-4}	1.09×10^{-3}	1.43×10^{-3}	1.77×10^{-3}
^{246}Cm	1.55×10^{-5}	2.50×10^{-4}	4.84×10^{-4}	7.19×10^{-4}	9.53×10^{-4}	4.20×10^{-5}	2.74×10^{-4}	5.05×10^{-4}	7.37×10^{-4}	9.68×10^{-4}
^{247}Cm	2.09×10^{-7}	2.00×10^{-5}	3.98×10^{-5}	5.96×10^{-5}	7.94×10^{-5}	1.59×10^{-6}	2.18×10^{-5}	4.20×10^{-5}	6.22×10^{-5}	8.24×10^{-5}
^{248}Cm	1.63×10^{-8}	9.52×10^{-6}	1.90×10^{-5}	2.85×10^{-5}	3.80×10^{-5}	1.02×10^{-7}	9.80×10^{-6}	1.95×10^{-5}	2.92×10^{-5}	3.89×10^{-5}
^{250}Cm	–	–	–	–	–	2.31×10^{-14}	4.55×10^{-12}	9.07×10^{-12}	1.36×10^{-11}	1.81×10^{-11}
^{249}Cf	2.16×10^{-10}	8.88×10^{-7}	1.78×10^{-6}	2.66×10^{-6}	3.55×10^{-6}	3.45×10^{-9}	8.89×10^{-7}	1.77×10^{-6}	2.66×10^{-6}	3.55×10^{-6}
^{250}Cf	3.52×10^{-11}	1.42×10^{-7}	2.83×10^{-7}	4.25×10^{-7}	5.66×10^{-7}	2.36×10^{-10}	1.54×10^{-7}	3.07×10^{-7}	4.60×10^{-7}	6.13×10^{-7}
^{251}Cf	1.99×10^{-11}	1.93×10^{-8}	3.85×10^{-8}	5.77×10^{-8}	7.70×10^{-8}	1.76×10^{-11}	1.97×10^{-8}	3.94×10^{-8}	5.91×10^{-8}	7.88×10^{-8}
^{252}Cf	4.72×10^{-12}	1.01×10^{-9}	2.02×10^{-9}	3.02×10^{-9}	4.03×10^{-9}	1.92×10^{-12}	1.35×10^{-9}	2.71×10^{-9}	4.06×10^{-9}	5.41×10^{-9}

Continued on Next Page

Table B.5: (continued)

	Fresh					Spent				
Isotope	Pass 1	Pass 2	Pass 3	Pass 4	Pass 5	Pass 1	Pass 2	Pass 3	Pass 4	Pass 5
^{226}Ra	–	–	–	–	–	1.78×10^{-14}	6.97×10^{-14}	1.22×10^{-13}	1.73×10^{-13}	2.25×10^{-13}
^{228}Ra	–	–	–	–	–	9.96×10^{-22}	2.23×10^{-21}	3.46×10^{-21}	4.69×10^{-21}	5.93×10^{-21}
^{206}Pb	–	–	–	–	–	2.45×10^{-19}	1.44×10^{-18}	2.64×10^{-18}	3.84×10^{-18}	5.03×10^{-18}
^{207}Pb	–	–	–	–	–	5.72×10^{-17}	3.08×10^{-16}	5.60×10^{-16}	8.11×10^{-16}	1.06×10^{-15}
^{208}Pb	–	–	–	–	–	4.68×10^{-10}	4.60×10^{-10}	4.52×10^{-10}	4.43×10^{-10}	4.35×10^{-10}
^{210}Pb	–	–	–	–	–	8.07×10^{-18}	3.96×10^{-17}	7.12×10^{-17}	1.03×10^{-16}	1.34×10^{-16}
^{228}Th	–	–	–	–	–	9.57×10^{-10}	9.11×10^{-10}	8.64×10^{-10}	8.18×10^{-10}	7.71×10^{-10}
^{229}Th	–	–	–	–	–	3.87×10^{-11}	3.79×10^{-11}	3.70×10^{-11}	3.62×10^{-11}	3.54×10^{-11}
^{230}Th	–	–	–	–	–	1.36×10^{-9}	3.83×10^{-9}	6.29×10^{-9}	8.75×10^{-9}	1.12×10^{-8}
^{232}Th	–	–	–	–	–	1.39×10^{-11}	2.61×10^{-11}	3.83×10^{-11}	5.05×10^{-11}	6.27×10^{-11}
^{209}Bi	–	–	–	–	–	2.68×10^{-13}	2.64×10^{-13}	2.60×10^{-13}	2.56×10^{-13}	2.52×10^{-13}
^{227}Ac	–	–	–	–	–	1.77×10^{-15}	8.36×10^{-15}	1.50×10^{-14}	2.15×10^{-14}	2.81×10^{-14}
^{231}Pa	–	–	–	–	–	7.10×10^{-11}	2.73×10^{-10}	4.75×10^{-10}	6.77×10^{-10}	8.79×10^{-10}
^{3}H	–	–	–	–	–	3.32×10^{-7}	3.30×10^{-7}	3.27×10^{-7}	3.25×10^{-7}	3.22×10^{-7}
^{4}He	–	–	–	–	–	6.54×10^{-5}	6.96×10^{-5}	7.39×10^{-5}	7.82×10^{-5}	8.25×10^{-5}
^{14}C	–	–	–	–	–	–	–	–	–	–
C*	–	–	–	–	–	–	–	–	–	–
^{81}Kr	–	–	–	–	–	5.62×10^{-11}	5.58×10^{-11}	5.55×10^{-11}	5.51×10^{-11}	5.48×10^{-11}
^{85}Kr	–	–	–	–	–	5.72×10^{-5}	5.68×10^{-5}	5.64×10^{-5}	5.59×10^{-5}	5.55×10^{-5}
Kr*	–	–	–	–	–	8.39×10^{-4}	8.32×10^{-4}	8.24×10^{-4}	8.17×10^{-4}	8.10×10^{-4}
^{90}Sr	–	–	–	–	–	1.04×10^{-3}	1.03×10^{-3}	1.03×10^{-3}	1.02×10^{-3}	1.01×10^{-3}
Sr*	–	–	–	–	–	7.57×10^{-4}	7.50×10^{-4}	7.44×10^{-4}	7.37×10^{-4}	7.31×10^{-4}
^{99}Tc	–	–	–	–	–	3.08×10^{-3}	3.05×10^{-3}	3.03×10^{-3}	3.01×10^{-3}	2.98×10^{-3}
Tc*	–	–	–	–	–	6.21×10^{-7}	6.22×10^{-7}	6.23×10^{-7}	6.23×10^{-7}	6.24×10^{-7}
^{129}I	–	–	–	–	–	9.77×10^{-4}	9.69×10^{-4}	9.62×10^{-4}	9.55×10^{-4}	9.47×10^{-4}
I*	–	–	–	–	–	3.45×10^{-4}	3.43×10^{-4}	3.41×10^{-4}	3.38×10^{-4}	3.36×10^{-4}
^{134}Cs	–	–	–	–	–	1.26×10^{-4}	1.26×10^{-4}	1.25×10^{-4}	1.24×10^{-4}	1.23×10^{-4}
^{135}Cs	–	–	–	–	–	5.31×10^{-3}	5.27×10^{-3}	5.23×10^{-3}	5.19×10^{-3}	5.15×10^{-3}
^{137}Cs	–	–	–	–	–	4.67×10^{-3}	4.63×10^{-3}	4.60×10^{-3}	4.56×10^{-3}	4.53×10^{-3}
Cs*	–	–	–	–	–	4.67×10^{-3}	4.63×10^{-3}	4.59×10^{-3}	4.56×10^{-3}	4.52×10^{-3}
FP†	–	–	–	–	–	1.19×10^{-1}	1.20×10^{-1}	1.20×10^{-1}	1.21×10^{-1}	1.21×10^{-1}
Act‡	–	–	–	–	–	–	–	–	–	–

* Non-enumerated isotopes
† Non-enumerated Fission Products
‡ Non-enumerated Actinides

Table B.6: VISION 75% Conversion Ratio FBR Fuel Recipes by Mass Fraction

	Fresh					Spent				
Isotope	Pass 1	Pass 2	Pass 3	Pass 4	Pass 5	Pass 1	Pass 2	Pass 3	Pass 4	Pass 5
^{232}U	5.27×10^{-14}	5.58×10^{-9}	1.12×10^{-8}	1.67×10^{-8}	2.23×10^{-8}	1.50×10^{-7}	1.30×10^{-7}	1.10×10^{-7}	9.00×10^{-8}	7.00×10^{-8}
^{233}U	–	–	–	–	–	1.07×10^{-8}	1.31×10^{-8}	1.54×10^{-8}	1.77×10^{-8}	2.00×10^{-8}
^{234}U	1.46×10^{-7}	1.47×10^{-4}	2.94×10^{-4}	4.40×10^{-4}	5.87×10^{-4}	1.80×10^{-4}	2.76×10^{-4}	3.72×10^{-4}	4.68×10^{-4}	5.64×10^{-4}
^{235}U	1.57×10^{-3}	1.29×10^{-3}	1.00×10^{-3}	7.14×10^{-4}	4.27×10^{-4}	6.68×10^{-4}	5.65×10^{-4}	4.63×10^{-4}	3.61×10^{-4}	2.59×10^{-4}
^{236}U	4.48×10^{-6}	1.05×10^{-4}	2.05×10^{-4}	3.05×10^{-4}	4.05×10^{-4}	2.02×10^{-4}	2.53×10^{-4}	3.03×10^{-4}	3.53×10^{-4}	4.04×10^{-4}
^{238}U	7.93×10^{-1}	7.92×10^{-1}	7.90×10^{-1}	7.88×10^{-1}	7.87×10^{-1}	7.06×10^{-1}	7.04×10^{-1}	7.03×10^{-1}	7.01×10^{-1}	6.99×10^{-1}
^{237}Np	9.69×10^{-3}	8.11×10^{-3}	6.52×10^{-3}	4.93×10^{-3}	3.34×10^{-3}	5.26×10^{-3}	4.45×10^{-3}	3.65×10^{-3}	2.85×10^{-3}	2.04×10^{-3}
^{238}Pu	4.71×10^{-3}	4.83×10^{-3}	4.95×10^{-3}	5.06×10^{-3}	5.18×10^{-3}	6.56×10^{-3}	6.10×10^{-3}	5.63×10^{-3}	5.17×10^{-3}	4.70×10^{-3}
^{239}Pu	9.80×10^{-2}	9.94×10^{-2}	1.01×10^{-1}	1.02×10^{-1}	1.03×10^{-1}	9.20×10^{-2}	9.25×10^{-2}	9.30×10^{-2}	9.34×10^{-2}	9.39×10^{-2}
^{240}Pu	4.63×10^{-2}	5.01×10^{-2}	5.39×10^{-2}	5.77×10^{-2}	6.15×10^{-2}	4.60×10^{-2}	4.89×10^{-2}	5.18×10^{-2}	5.48×10^{-2}	5.77×10^{-2}
^{241}Pu	2.18×10^{-2}	1.91×10^{-2}	1.63×10^{-2}	1.35×10^{-2}	1.07×10^{-2}	1.07×10^{-2}	1.02×10^{-2}	9.77×10^{-3}	9.32×10^{-3}	8.88×10^{-3}
^{242}Pu	1.35×10^{-2}	1.36×10^{-2}	1.36×10^{-2}	1.37×10^{-2}	1.37×10^{-2}	1.30×10^{-2}	1.28×10^{-2}	1.27×10^{-2}	1.25×10^{-2}	1.23×10^{-2}
^{244}Pu	–	–	–	–	–	7.21×10^{-8}	7.21×10^{-8}	7.21×10^{-8}	7.21×10^{-8}	7.21×10^{-8}
^{241}Am	6.93×10^{-3}	6.54×10^{-3}	6.15×10^{-3}	5.76×10^{-3}	5.37×10^{-3}	5.74×10^{-3}	5.34×10^{-3}	4.94×10^{-3}	4.54×10^{-3}	4.14×10^{-3}
^{242m}Am	1.31×10^{-5}	9.27×10^{-5}	1.72×10^{-4}	2.52×10^{-4}	3.32×10^{-4}	3.04×10^{-4}	3.01×10^{-4}	2.98×10^{-4}	2.95×10^{-4}	2.92×10^{-4}
^{243}Am	3.07×10^{-3}	3.39×10^{-3}	3.70×10^{-3}	4.02×10^{-3}	4.34×10^{-3}	3.32×10^{-3}	3.50×10^{-3}	3.68×10^{-3}	3.85×10^{-3}	4.03×10^{-3}
^{242}Cm	2.05×10^{-7}	9.10×10^{-6}	1.80×10^{-5}	2.69×10^{-5}	3.58×10^{-5}	3.23×10^{-4}	3.01×10^{-4}	2.80×10^{-4}	2.59×10^{-4}	2.37×10^{-4}
^{243}Cm	1.05×10^{-5}	1.27×10^{-5}	1.49×10^{-5}	1.71×10^{-5}	1.93×10^{-5}	2.13×10^{-5}	2.10×10^{-5}	2.08×10^{-5}	2.05×10^{-5}	2.02×10^{-5}
^{244}Cm	1.08×10^{-3}	1.49×10^{-3}	1.90×10^{-3}	2.31×10^{-3}	2.72×10^{-3}	1.62×10^{-3}	1.92×10^{-3}	2.22×10^{-3}	2.52×10^{-3}	2.82×10^{-3}
^{245}Cm	8.56×10^{-5}	2.38×10^{-4}	3.90×10^{-4}	5.42×10^{-4}	6.95×10^{-4}	2.76×10^{-4}	3.87×10^{-4}	4.98×10^{-4}	6.10×10^{-4}	7.21×10^{-4}
^{246}Cm	1.08×10^{-5}	1.01×10^{-4}	1.92×10^{-4}	2.82×10^{-4}	3.73×10^{-4}	2.84×10^{-5}	1.17×10^{-4}	2.05×10^{-4}	2.94×10^{-4}	3.83×10^{-4}
^{247}Cm	1.46×10^{-7}	7.41×10^{-6}	1.47×10^{-5}	2.19×10^{-5}	2.92×10^{-5}	1.04×10^{-6}	8.54×10^{-6}	1.60×10^{-5}	2.35×10^{-5}	3.10×10^{-5}
^{248}Cm	1.14×10^{-8}	3.46×10^{-6}	6.91×10^{-6}	1.04×10^{-5}	1.38×10^{-5}	6.63×10^{-8}	3.62×10^{-6}	7.17×10^{-6}	1.07×10^{-5}	1.43×10^{-5}
^{250}Cm	–	–	–	–	–	1.31×10^{-14}	1.44×10^{-12}	2.87×10^{-12}	4.29×10^{-12}	5.72×10^{-12}
^{249}Cf	1.50×10^{-10}	3.07×10^{-7}	6.14×10^{-7}	9.22×10^{-7}	1.23×10^{-6}	2.17×10^{-9}	3.15×10^{-7}	6.28×10^{-7}	9.41×10^{-7}	1.25×10^{-6}
^{250}Cf	2.45×10^{-11}	4.28×10^{-8}	8.57×10^{-8}	1.28×10^{-7}	1.71×10^{-7}	1.37×10^{-10}	4.87×10^{-8}	9.73×10^{-8}	1.46×10^{-7}	1.94×10^{-7}
^{251}Cf	1.39×10^{-11}	5.51×10^{-9}	1.10×10^{-8}	1.65×10^{-8}	2.20×10^{-8}	1.10×10^{-11}	5.86×10^{-9}	1.17×10^{-8}	1.76×10^{-8}	2.34×10^{-8}
^{252}Cf	3.29×10^{-12}	2.64×10^{-10}	5.24×10^{-10}	7.84×10^{-10}	1.04×10^{-9}	1.25×10^{-12}	3.69×10^{-10}	7.37×10^{-10}	1.11×10^{-9}	1.47×10^{-9}

Continued on Next Page

Table B.6: (continued)

	Fresh					Spent				
Isotope	Pass 1	Pass 2	Pass 3	Pass 4	Pass 5	Pass 1	Pass 2	Pass 3	Pass 4	Pass 5
^{226}Ra	–	–	–	–	–	1.41×10^{-14}	4.37×10^{-14}	7.32×10^{-14}	1.03×10^{-13}	1.32×10^{-13}
^{228}Ra	–	–	–	–	–	1.12×10^{-21}	2.06×10^{-21}	3.00×10^{-21}	3.94×10^{-21}	4.89×10^{-21}
^{206}Pb	–	–	–	–	–	2.05×10^{-19}	9.64×10^{-19}	1.72×10^{-18}	2.48×10^{-18}	3.24×10^{-18}
^{207}Pb	–	–	–	–	–	5.16×10^{-17}	2.00×10^{-16}	3.48×10^{-16}	4.96×10^{-16}	6.44×10^{-16}
^{208}Pb	–	–	–	–	–	3.49×10^{-10}	3.20×10^{-10}	2.91×10^{-10}	2.62×10^{-10}	2.33×10^{-10}
^{210}Pb	–	–	–	–	–	6.57×10^{-18}	2.56×10^{-17}	4.46×10^{-17}	6.36×10^{-17}	8.26×10^{-17}
^{228}Th	–	–	–	–	–	6.86×10^{-10}	6.12×10^{-10}	5.38×10^{-10}	4.64×10^{-10}	3.90×10^{-10}
^{229}Th	–	–	–	–	–	2.74×10^{-11}	2.50×10^{-11}	2.26×10^{-11}	2.02×10^{-11}	1.78×10^{-11}
^{230}Th	–	–	–	–	–	1.04×10^{-9}	2.35×10^{-9}	3.66×10^{-9}	4.96×10^{-9}	6.27×10^{-9}
^{232}Th	–	–	–	–	–	1.51×10^{-11}	2.37×10^{-11}	3.22×10^{-11}	4.08×10^{-11}	4.94×10^{-11}
^{209}Bi	–	–	–	–	–	1.88×10^{-13}	1.73×10^{-13}	1.57×10^{-13}	1.42×10^{-13}	1.26×10^{-13}
^{227}Ac	–	–	–	–	–	1.47×10^{-15}	5.13×10^{-15}	8.78×10^{-15}	1.24×10^{-14}	1.61×10^{-14}
^{231}Pa	–	–	–	–	–	5.50×10^{-11}	1.61×10^{-10}	2.67×10^{-10}	3.73×10^{-10}	4.79×10^{-10}
^{3}H	–	–	–	–	–	2.55×10^{-7}	2.55×10^{-7}	2.55×10^{-7}	2.55×10^{-7}	2.55×10^{-7}
^{4}He	–	–	–	–	–	4.58×10^{-5}	4.45×10^{-5}	4.32×10^{-5}	4.19×10^{-5}	4.06×10^{-5}
^{14}C	–	–	–	–	–	–	–	–	–	–
C*	–	–	–	–	–	–	–	–	–	–
^{81}Kr	–	–	–	–	–	4.03×10^{-11}	4.04×10^{-11}	4.04×10^{-11}	4.05×10^{-11}	4.06×10^{-11}
^{85}Kr	–	–	–	–	–	4.42×10^{-5}	4.42×10^{-5}	4.41×10^{-5}	4.41×10^{-5}	4.41×10^{-5}
Kr*	–	–	–	–	–	6.53×10^{-4}	6.52×10^{-4}	6.51×10^{-4}	6.51×10^{-4}	6.50×10^{-4}
^{90}Sr	–	–	–	–	–	8.17×10^{-4}	8.16×10^{-4}	8.14×10^{-4}	8.13×10^{-4}	8.12×10^{-4}
Sr*	–	–	–	–	–	5.94×10^{-4}	5.94×10^{-4}	5.93×10^{-4}	5.92×10^{-4}	5.91×10^{-4}
^{99}Tc	–	–	–	–	–	2.38×10^{-3}	2.38×10^{-3}	2.38×10^{-3}	2.38×10^{-3}	2.38×10^{-3}
Tc*	–	–	–	–	–	4.91×10^{-7}	4.94×10^{-7}	4.97×10^{-7}	5.00×10^{-7}	5.02×10^{-7}
^{129}I	–	–	–	–	–	7.46×10^{-4}	7.47×10^{-4}	7.47×10^{-4}	7.48×10^{-4}	7.48×10^{-4}
I*	–	–	–	–	–	2.61×10^{-4}	2.61×10^{-4}	2.62×10^{-4}	2.62×10^{-4}	2.62×10^{-4}
^{134}Cs	–	–	–	–	–	9.12×10^{-5}	9.13×10^{-5}	9.14×10^{-5}	9.15×10^{-5}	9.16×10^{-5}
^{135}Cs	–	–	–	–	–	4.07×10^{-3}	4.08×10^{-3}	4.08×10^{-3}	4.08×10^{-3}	4.08×10^{-3}
^{137}Cs	–	–	–	–	–	3.58×10^{-3}	3.58×10^{-3}	3.58×10^{-3}	3.58×10^{-3}	3.58×10^{-3}
Cs*	–	–	–	–	–	3.59×10^{-3}	3.59×10^{-3}	3.60×10^{-3}	3.60×10^{-3}	3.60×10^{-3}
FP†	–	–	–	–	–	9.07×10^{-2}	9.07×10^{-2}	9.06×10^{-2}	9.06×10^{-2}	9.06×10^{-2}
Act‡	–	–	–	–	–	–	–	–	–	–

* Non-enumerated isotopes
† Non-enumerated Fission Products
‡ Non-enumerated Actinides

Table B.7: VISION 100% Conversion Ratio FBR Fuel Recipes by Mass Fraction

	Fresh					Spent				
Isotope	Pass 1	Pass 2	Pass 3	Pass 4	Pass 5	Pass 1	Pass 2	Pass 3	Pass 4	Pass 5
^{232}U	3.91×10^{-14}	1.77×10^{-9}	3.54×10^{-9}	5.30×10^{-9}	7.07×10^{-9}	1.13×10^{-7}	9.03×10^{-8}	6.77×10^{-8}	4.51×10^{-8}	2.25×10^{-8}
^{233}U	–	–	–	–	–	7.79×10^{-9}	7.43×10^{-9}	7.08×10^{-9}	6.72×10^{-9}	6.37×10^{-9}
^{234}U	1.08×10^{-7}	4.89×10^{-5}	9.76×10^{-5}	1.46×10^{-4}	1.95×10^{-4}	1.20×10^{-4}	1.37×10^{-4}	1.54×10^{-4}	1.70×10^{-4}	1.87×10^{-4}
^{235}U	1.68×10^{-3}	1.35×10^{-3}	1.02×10^{-3}	6.85×10^{-4}	3.54×10^{-4}	8.12×10^{-4}	6.57×10^{-4}	5.03×10^{-4}	3.49×10^{-4}	1.95×10^{-4}
^{236}U	3.32×10^{-6}	7.87×10^{-5}	1.54×10^{-4}	2.29×10^{-4}	3.05×10^{-4}	1.84×10^{-4}	2.14×10^{-4}	2.44×10^{-4}	2.75×10^{-4}	3.05×10^{-4}
^{238}U	8.46×10^{-1}	8.50×10^{-1}	8.53×10^{-1}	8.57×10^{-1}	8.61×10^{-1}	7.69×10^{-1}	7.72×10^{-1}	7.75×10^{-1}	7.78×10^{-1}	7.81×10^{-1}
^{237}Np	7.18×10^{-3}	5.61×10^{-3}	4.03×10^{-3}	2.46×10^{-3}	8.80×10^{-4}	4.46×10^{-3}	3.56×10^{-3}	2.66×10^{-3}	1.75×10^{-3}	8.51×10^{-4}
^{238}Pu	3.49×10^{-3}	2.99×10^{-3}	2.50×10^{-3}	2.00×10^{-3}	1.50×10^{-3}	4.67×10^{-3}	3.87×10^{-3}	3.08×10^{-3}	2.28×10^{-3}	1.48×10^{-3}
^{239}Pu	7.26×10^{-2}	7.66×10^{-2}	8.06×10^{-2}	8.46×10^{-2}	8.86×10^{-2}	8.15×10^{-2}	8.36×10^{-2}	8.57×10^{-2}	8.78×10^{-2}	8.99×10^{-2}
^{240}Pu	3.43×10^{-2}	3.49×10^{-2}	3.54×10^{-2}	3.60×10^{-2}	3.66×10^{-2}	3.42×10^{-2}	3.49×10^{-2}	3.56×10^{-2}	3.64×10^{-2}	3.71×10^{-2}
^{241}Pu	1.62×10^{-2}	1.32×10^{-2}	1.03×10^{-2}	7.37×10^{-3}	4.43×10^{-3}	8.32×10^{-3}	7.41×10^{-3}	6.50×10^{-3}	5.59×10^{-3}	4.68×10^{-3}
^{242}Pu	1.00×10^{-2}	8.27×10^{-3}	6.53×10^{-3}	4.79×10^{-3}	3.05×10^{-3}	9.68×10^{-3}	8.03×10^{-3}	6.39×10^{-3}	4.74×10^{-3}	3.10×10^{-3}
^{244}Pu	–	–	–	–	–	4.45×10^{-8}	3.70×10^{-8}	2.95×10^{-8}	2.20×10^{-8}	1.45×10^{-8}
^{241}Am	5.13×10^{-3}	4.33×10^{-3}	3.53×10^{-3}	2.73×10^{-3}	1.92×10^{-3}	4.59×10^{-3}	3.89×10^{-3}	3.18×10^{-3}	2.47×10^{-3}	1.76×10^{-3}
^{242m}Am	9.73×10^{-6}	3.94×10^{-5}	6.91×10^{-5}	9.88×10^{-5}	1.28×10^{-4}	2.14×10^{-4}	1.90×10^{-4}	1.65×10^{-4}	1.40×10^{-4}	1.15×10^{-4}
^{243}Am	2.28×10^{-3}	1.91×10^{-3}	1.55×10^{-3}	1.19×10^{-3}	8.33×10^{-4}	2.42×10^{-3}	2.02×10^{-3}	1.63×10^{-3}	1.24×10^{-3}	8.42×10^{-4}
^{242}Cm	1.52×10^{-7}	3.70×10^{-6}	7.25×10^{-6}	1.08×10^{-5}	1.44×10^{-5}	2.36×10^{-4}	2.00×10^{-4}	1.64×10^{-4}	1.27×10^{-4}	9.14×10^{-5}
^{243}Cm	7.78×10^{-6}	7.39×10^{-6}	6.99×10^{-6}	6.60×10^{-6}	6.20×10^{-6}	1.28×10^{-5}	1.12×10^{-5}	9.54×10^{-6}	7.91×10^{-6}	6.28×10^{-6}
^{244}Cm	7.98×10^{-4}	7.20×10^{-4}	6.42×10^{-4}	5.65×10^{-4}	4.87×10^{-4}	1.10×10^{-3}	9.51×10^{-4}	8.04×10^{-4}	6.58×10^{-4}	5.11×10^{-4}
^{245}Cm	6.35×10^{-5}	7.86×10^{-5}	9.37×10^{-5}	1.09×10^{-4}	1.24×10^{-4}	1.74×10^{-4}	1.62×10^{-4}	1.50×10^{-4}	1.37×10^{-4}	1.25×10^{-4}
^{246}Cm	8.00×10^{-6}	2.23×10^{-5}	3.67×10^{-5}	5.10×10^{-5}	6.53×10^{-5}	1.73×10^{-5}	2.96×10^{-5}	4.18×10^{-5}	5.41×10^{-5}	6.64×10^{-5}
^{247}Cm	1.08×10^{-7}	1.33×10^{-6}	2.55×10^{-6}	3.77×10^{-6}	4.99×10^{-6}	5.69×10^{-7}	1.68×10^{-6}	2.79×10^{-6}	3.91×10^{-6}	5.02×10^{-6}
^{248}Cm	8.41×10^{-9}	5.91×10^{-7}	1.17×10^{-6}	1.76×10^{-6}	2.34×10^{-6}	3.45×10^{-8}	6.21×10^{-7}	1.21×10^{-6}	1.79×10^{-6}	2.38×10^{-6}
^{250}Cm	–	–	–	–	–	6.09×10^{-15}	2.02×10^{-13}	3.98×10^{-13}	5.94×10^{-13}	7.90×10^{-13}
^{249}Cf	1.11×10^{-10}	5.08×10^{-8}	1.02×10^{-7}	1.52×10^{-7}	2.03×10^{-7}	1.04×10^{-9}	5.12×10^{-8}	1.01×10^{-7}	1.52×10^{-7}	2.02×10^{-7}
^{250}Cf	1.82×10^{-11}	6.58×10^{-9}	1.31×10^{-8}	1.97×10^{-8}	2.63×10^{-8}	5.45×10^{-11}	6.90×10^{-9}	1.37×10^{-8}	2.06×10^{-8}	2.74×10^{-8}
^{251}Cf	1.03×10^{-11}	8.30×10^{-10}	1.65×10^{-9}	2.47×10^{-9}	3.29×10^{-9}	6.49×10^{-12}	8.12×10^{-10}	1.62×10^{-9}	2.42×10^{-9}	3.23×10^{-9}
^{252}Cf	2.44×10^{-12}	3.91×10^{-11}	7.58×10^{-11}	1.13×10^{-10}	1.49×10^{-10}	9.51×10^{-13}	4.92×10^{-11}	9.74×10^{-11}	1.46×10^{-10}	1.94×10^{-10}

Continued on Next Page

Table B.7: (continued)

	Fresh					Spent				
Isotope	Pass 1	Pass 2	Pass 3	Pass 4	Pass 5	Pass 1	Pass 2	Pass 3	Pass 4	Pass 5
^{226}Ra	–	–	–	–	–	8.90×10^{-15}	1.61×10^{-14}	2.32×10^{-14}	3.04×10^{-14}	3.75×10^{-14}
^{228}Ra	–	–	–	–	–	$8.14\times10^{-2}2$	$1.39\times10^{-2}1$	$1.96\times10^{-2}1$	$2.54\times10^{-2}1$	$3.11\times10^{-2}1$
^{206}Pb	–	–	–	–	–	1.38×10^{-19}	3.26×10^{-19}	5.14×10^{-19}	7.02×10^{-19}	8.90×10^{-19}
^{207}Pb	–	–	–	–	–	3.20×10^{-17}	6.17×10^{-17}	9.15×10^{-17}	1.21×10^{-16}	1.51×10^{-16}
^{208}Pb	–	–	–	–	–	2.10×10^{-10}	1.72×10^{-10}	1.35×10^{-10}	9.68×10^{-11}	5.90×10^{-11}
^{210}Pb	–	–	–	–	–	4.23×10^{-18}	8.83×10^{-18}	1.34×10^{-17}	1.80×10^{-17}	2.27×10^{-17}
^{228}Th	–	–	–	–	–	4.40×10^{-10}	3.56×10^{-10}	2.73×10^{-10}	1.89×10^{-10}	1.05×10^{-10}
^{229}Th	–	–	–	–	–	1.56×10^{-11}	1.28×10^{-11}	9.93×10^{-12}	7.11×10^{-12}	4.28×10^{-12}
^{230}Th	–	–	–	–	–	6.67×10^{-10}	9.79×10^{-10}	1.29×10^{-9}	1.60×10^{-9}	1.91×10^{-9}
^{232}Th	–	–	–	–	–	1.20×10^{-11}	1.74×10^{-11}	2.28×10^{-11}	2.82×10^{-11}	3.37×10^{-11}
^{209}Bi	–	–	–	–	–	1.06×10^{-13}	8.74×10^{-14}	6.83×10^{-14}	4.93×10^{-14}	3.02×10^{-14}
^{227}Ac	–	–	–	–	–	8.94×10^{-16}	1.68×10^{-15}	2.47×10^{-15}	3.26×10^{-15}	4.05×10^{-15}
^{231}Pa	–	–	–	–	–	3.29×10^{-11}	5.73×10^{-11}	8.18×10^{-11}	1.06×10^{-10}	1.31×10^{-10}
^{3}H	–	–	–	–	–	1.89×10^{-7}	1.91×10^{-7}	1.93×10^{-7}	1.95×10^{-7}	1.97×10^{-7}
^{4}He	–	–	–	–	–	2.83×10^{-5}	2.41×10^{-5}	1.99×10^{-5}	1.58×10^{-5}	1.16×10^{-5}
^{14}C	–	–	–	–	–	–	–	–	–	–
C*	–	–	–	–	–	–	–	–	–	–
^{81}Kr	–	–	–	–	–	2.57×10^{-11}	2.61×10^{-11}	2.65×10^{-11}	2.69×10^{-11}	2.73×10^{-11}
^{85}Kr	–	–	–	–	–	3.31×10^{-5}	3.33×10^{-5}	3.36×10^{-5}	3.39×10^{-5}	3.42×10^{-5}
Kr*	–	–	–	–	–	4.83×10^{-4}	4.87×10^{-4}	4.91×10^{-4}	4.95×10^{-4}	4.98×10^{-4}
^{90}Sr	–	–	–	–	–	6.15×10^{-4}	6.19×10^{-4}	6.23×10^{-4}	6.27×10^{-4}	6.31×10^{-4}
Sr*	–	–	–	–	–	4.49×10^{-4}	4.52×10^{-4}	4.55×10^{-4}	4.58×10^{-4}	4.61×10^{-4}
^{99}Tc	–	–	–	–	–	1.75×10^{-3}	1.77×10^{-3}	1.78×10^{-3}	1.80×10^{-3}	1.82×10^{-3}
Tc*	–	–	–	–	–	4.41×10^{-7}	4.46×10^{-7}	4.52×10^{-7}	4.57×10^{-7}	4.63×10^{-7}
^{129}I	–	–	–	–	–	5.40×10^{-4}	5.46×10^{-4}	5.52×10^{-4}	5.58×10^{-4}	5.63×10^{-4}
I*	–	–	–	–	–	1.89×10^{-4}	1.91×10^{-4}	1.94×10^{-4}	1.96×10^{-4}	1.98×10^{-4}
^{134}Cs	–	–	–	–	–	6.15×10^{-5}	6.23×10^{-5}	6.31×10^{-5}	6.39×10^{-5}	6.47×10^{-5}
^{135}Cs	–	–	–	–	–	2.97×10^{-3}	3.00×10^{-3}	3.03×10^{-3}	3.06×10^{-3}	3.09×10^{-3}
^{137}Cs	–	–	–	–	–	2.62×10^{-3}	2.65×10^{-3}	2.67×10^{-3}	2.70×10^{-3}	2.73×10^{-3}
Cs*	–	–	–	–	–	2.63×10^{-3}	2.66×10^{-3}	2.68×10^{-3}	2.71×10^{-3}	2.74×10^{-3}
FP†	–	–	–	–	–	6.57×10^{-2}	6.56×10^{-2}	6.54×10^{-2}	6.53×10^{-2}	6.52×10^{-2}
Act‡	–	–	–	–	–	–	–	–	–	–

* Non-enumerated isotopes
† Non-enumerated Fission Products
‡ Non-enumerated Actinides

Appendix C

Interactive Content

The following three figures show the segment-by-segment development of the trade-off surfaces during the optimization run. Two similar three-objective MOSA runs are included in each plot. One run starts with the Rear-loaded FBR Configuration while the second starts with an All-LWR configuration. As the animations show, the objective trade-off surfaces contained within the non-dominating archives converge toward consistent values as the optimization progresses.

The functionality of these figures has been verified while using the Adobe Acrobat PDF reader (version 9.5), but as with all static documents, future compatibility cannot be guaranteed.

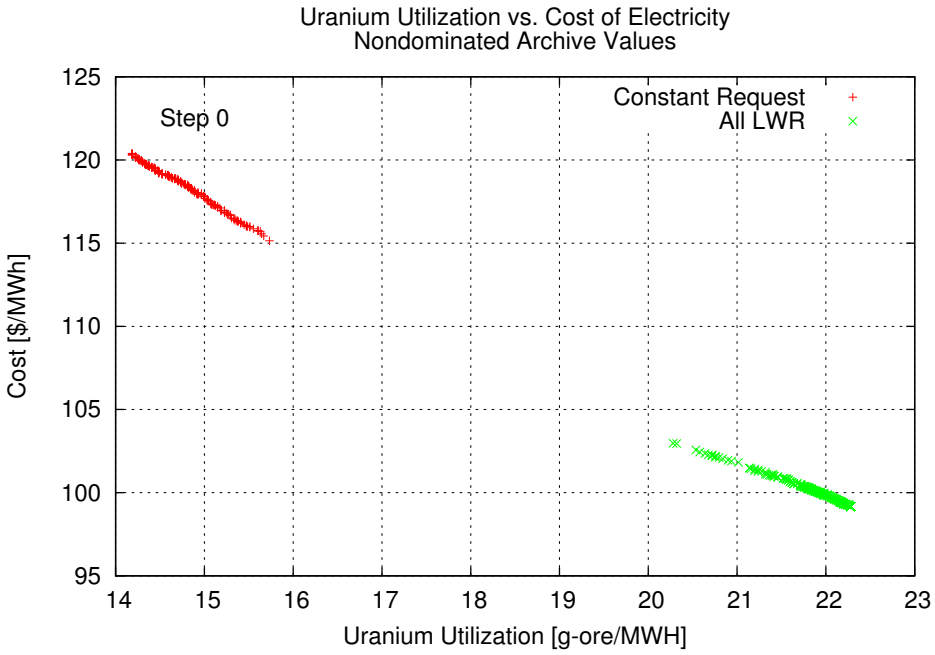

Figure C.1: Step-by-step animation of Figure 3.29

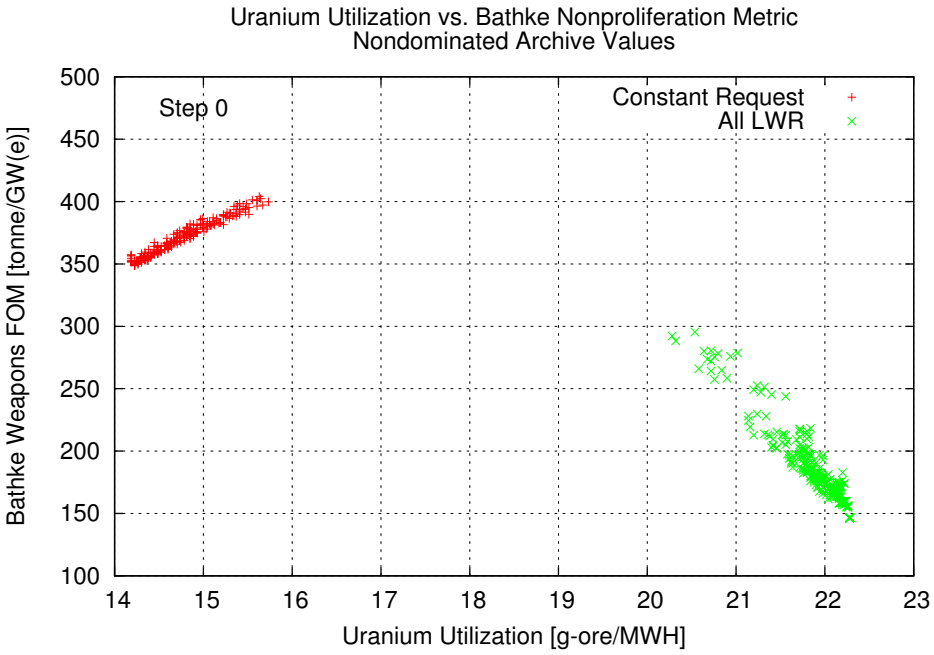

Figure C.2: Step-by-step animation of Figure 3.30

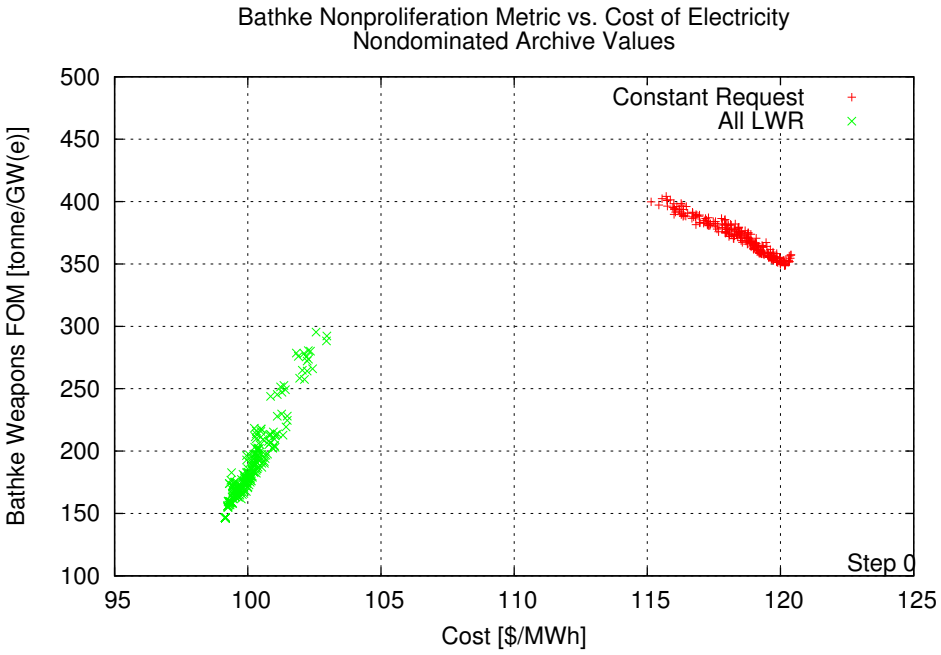

Figure C.3: Step-by-step animation of Figure 3.31