AWARD CATEGORY:
NERI

AWARD NUMBER:
DE-FC07-06ID14734

AWARDEE NAME:
North Carolina State University

PROJECT TITLE:
Development and Utilization of Mathematical Optimization in Advanced Fuel Cycle Systems Analysis

Report category: Management

Report name: Final Report

PRINCIPAL INVESTIGATOR:
Paul J. Turinsky

RESEARCH ASSISTANT:
Ross Hays

Department of Nuclear Engineering
North Carolina State University
Raleigh, NC 27695-7909

September 2, 2011

Executive Summary

Over the past sixty years, a wide variety of nuclear power technologies have been theorized, investigated and tested to various degrees. These technologies, if properly applied, could provide a stable, long-term, economical source of CO2-free electric power. However, the recycling of nuclear fuel introduces a degree of coupling between reactor systems which must be accounted for when making long term strategic plans. This work investigates the use of a simulated annealing optimization algorithm coupled together with the VISION fuel cycle simulation model in order to identify attractive strategies from economic, evironmental, non-proliferation and waste-disposal perspectives, which each have associated an objective function.

The simulated annealing optimization algorithm works by perturbing the fraction of new reactor capacity allocated to each available reactor type (using a set of heuristic rules) then evaluating the resulting deployment scenario outcomes using the VISION model and the chosen objective functions. These new scenarios, which are either accepted or rejected according the the Metropolis Criterion, are then used as the basis for further perturbations. By repeating this process several thousand times, a family of near-optimal solutions are obtained.

Preliminary results from this work using a two-step, Once-through LWR to Full-recycle/FR-burner deployment scenario with exponentially increasing electric demand indicate that the algorithm is capable of finding reactor deployment profiles that reduce the long-term-heat waste disposal burden relative to an initial reference scenario. Further work is under way to refine the current results and to extend them to include the other objective functions and to examine the optimization trade-offs that exist between these different objectives.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The purpose of this work is to examine the use and application of mathematical optimization by 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. Presented below is 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 2011, 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 [7]. 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 the code user 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, 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.

[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.

The NFCSIM simulation program, developed at Los Alamos National Laboratory [11], 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 necessitate 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 [14, 15]. 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 The VISION Model

The Verfiable Fuel Cycle Simulaton model was developed through a collaboration between Idaho National Laboratory and numerous academic and laboratory contributors [9]. 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 location, mass and isotopic composition of all nuclear fuel, along with the number, types, and ages of the reactors, fuel fabrication facilities, separations 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 [18], 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.

1.3.1 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 plant. 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, leads 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 reactor-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 [9]. 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

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

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 [9]; 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 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.

1.3.2 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 [8].

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 and explore management alternatives.

Chapter 2

Optimization 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. To get an accurate accounting of the costs and benefits associated with a discrete segment of the larger energy-marketplace, 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 Socially-Conscious decision making methodology. 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 three simple metrics chosen for this work were selected because they relate to three key factors that will ultimately determine the acceptability of any proposed nuclear fuel cycle, e.g. waste disposal, proliferation resistance 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 it's own set of design

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

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 [17]. 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

$$H = \sum_{i \in Sources} H_i \tag{2.2}$$

The energy production term, E, is the deployed reactor capacity, 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.

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

Table 2.2: VISION Variables in the Weapons Usable Objective Function

Term	VISION Variable	Sample Time
W_t	`WU total`	All
E	See equation 2.3	

energy produced during the simulation (also in Joules) and multiplying by a factor of 1,000,000 (to give an aesthetically pleasing mantissa). Typical values are in the range of 150 to 3000 part-per-million.

2.1.2 Nonproliferation

From the beginning, the peaceful use of nuclear power has had a shadow cast on it by the military and destructive potential that it embodies. Indeed, most of the early research into nuclear power was for military purposes, with only incidental peacetime benefits. One concern that must be addressed in the nuclear fuel cycle is the potential ability of bad actors to divert fissile material from a facility for illicit purposes. Although a full analysis of the system would require detailed facility-level protection and monitoring plans, it is a rather simple matter to assume that, all else being equal, the more fissile material there is on hand in the system, the greater the potential diversion threat may be. Not all fissile material is created equal from a weapons standpoint; a common way to compare different isotopes is through the use of Pu-239 equivalence factors [5]. The VISION model uses the factors listed in table 2.3 to compute a 'Weapons Usable' value for fuel in all areas of the fuel cycle. These factors are inversely proportional to the bare-sphere critical mass of the isotope, and do not account for other usability factors such as gamma dose or heat production, or fuel form and diversion difficulty, e.g. where in the fuel cycle facilities material resides. This value, measured in equivalent tons of Pu-239, is integrated over the simulation time and normalized by the total electric energy production to determine the nonproliferation objective function value.

$$F_W = \frac{\Delta t \sum_{t=t_{start}}^{t_{stop}} W_t}{E} \tag{2.4}$$

Future improvements are planned for this objective function to account for the differences in vulnerability of the material at each stage in the fuel cycle.

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 [4]. 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

Table 2.3: Weapons Usable ^{239}Pu Equivalence Factors from VISION

Isotope	Factor	Isotope	Factor
^{231}Pa	0.06	^{241}Pu	0.79
^{233}U	0.62	^{241}Am	0.18
^{235}U	0.21	^{242m}Am	0.34
^{237}Np	0.17	^{243}Am	0.07
^{238}Pu	1.02	^{244}Cm	0.36
^{239}Pu	1.00	^{245}Cm	0.78
^{240}Pu	0.28	^{246}Cm	0.12

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	

produce the greatest amount of electrical energy. Mathematically, it is expressed simply by the equation below.

$$F_U = \frac{U_{initial} - U_{final}}{E} \tag{2.5}$$

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.

2.1.4 Economics

Ultimately, the key factor that will make or break any nuclear fuel cycle deployment scenario is cost. Traditionally, a large reason nuclear power has occupied the position it has in the national power grid is that is has a much lower marginal fuel cost than any other technology (that is also scalable and dispatch-able). 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 will be the most complex of all of the objective functions; it will track all major facilities, both through construction, operation and shutdown. Those facilities that run past the end of the planning horizon will have their costs and output adjusted to reflect their value within the planning horizon. Costs will be normalized by the net electric production to yield a net-present cost of electric power production that can be compared against other predictions of future energy prices. Finally, the economic objective function will be used in multi-objective optimizations to provide an estimate of the economic trade-off costs associated with minimizing the other objective functions.

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 these hard constraints 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, though 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

Table 2.5: 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

reactor type runs short of fuel. The constraint violation value is given by:

$$\Theta_{U.C.} = \sum_{i \in rx} C_i^{idle} \tag{2.6}$$

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

2.2.2 VISION Heuristic Invocation

As previously mentioned, some of the heuristic rules built into the VISION simulation act to constrain reactor build rates to levels that the available fuel supply can support. While this provides a check against allowing many grossly infeasible scenarios from going forward, it decouples the changes in input settings from any meaningful change in the output values. Thus, by penalizing the objective function for causing the invocation of these built-in heuristic rules, we constrain the optimization search to those areas where the output is sensitive to the input, and meaningful results may be obtained.

The VISION Heuristic penalty function is computed using equation 2.7 and the variables listed in table 2.5. 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.

$$\Theta_H = \frac{\sum_{t=t_{start}}^{t_{stop}} B_t P_{LWRmf}\, T_{LWRmf}}{\sum_{i \in rx} P_i T_i \sum_{t=t_{start}}^{t_{stop}} O_{t,i}} \tag{2.7}$$

2.2.3 Solution Dissimilarity

The solution dissimilarity metric can be used as a simple representation of the difference between two specified configurations. This is computed by taking the root-mean-square difference between

the fast reactor fraction of the deployed capacity between two evaluated configurations:

$$\Gamma_{dis} = \frac{1}{N}\sqrt{\sum_{j=1}^{N}\left(\tilde{f}(t_j) - f(t_j)\right)^2} \tag{2.8}$$

Where:

$$f(t_j) = C^{FBR}(t_j)/C^{total}(t_j)$$

This metric is not used by the optimization algorithm itself, but rather provides some indication to the user as to the overall difference between the optimization starting configuration and the current candidate configuration.

2.2.4 Build Discrepancy

The build discrepancy metric gives an indication of how closely fast reactor fraction of new capacity built in a given year by the VISION model matches the fraction requested in the input configuration value. The two primary reasons the values may not match for a given year are that 1) while the requested fraction can vary continuously, the fraction built is constrained to an integral number of reactors and must meet the power demand, and 2) the built-in heuristics in the VISION code may override the requested distribution (if heuristics are enabled). Mathematically this metric is given by the root-mean-square difference between the requested fast reactor fraction for year i, $\tilde{\chi}_i$ and the as-built fast reactor fraction for year i, χ_i.

$$\Gamma_{des} = \frac{1}{N}\sqrt{\sum_{i=1}^{N}\left(\tilde{\chi}_i - \chi_i\right)^2} \tag{2.9}$$

Where:

$$\tilde{\chi}_i = \tilde{B}_i^{FBR}/\tilde{B}_i^{total}$$

As with Γ_{dis}, this metric is not used by the optimization algorithm, but merely provides extra information to the user.

2.3 Closing the Planning Horizon: The Socially Conscious Approach

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 and have sufficient fuel to operate the reactors over their planned lifetimes. 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).

First, some terms and definitions are presented, followed by the derivation of the controlling equations and their implementation. Finally, there is a brief discussion of assumptions and potential limitations of this method.

2.3.1 Definitions

The derivation of the Socially Conscious Approach to closing the planning horizon involves a rather sizable list of terms and variables. For convenience, they are listed below in table 2.6. 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 2.6: Definitions used in the Socially Conscious 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.
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.

Continued on Next Page

Table 2.6: (continued)

Term	Description
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
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 Socially Conscious approach is to give a clean separation between those effects that occur after 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.

2.3.3 Forecasting Spent Fuel Inventories

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 [12, 13]. 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:

$$g_i = d_i^{met} - \sum_{rx' \in \{rx\}} \sum_{j=1}^{T_{rx'}^{life}} N_{rx',i-j} P_{rx'} \tag{2.10}$$

$$N_{LWRmf,i} = \begin{cases} \tilde{N}_{LWRmf,i} - 1 & \text{if } \tilde{N}_{LWRmf,i} P_{LWRmf} + \tilde{N}_{FBR,i} P_{FBR} - g_i > P_{LWRmf} \\ \tilde{N}_{LWRmf,i} & \text{otherwise} \end{cases} \tag{2.11}$$

$$N_{FBR,i} = \begin{cases} \tilde{N}_{FBR,i} - 1 & \text{if } P_{LWRmf} \geq \tilde{N}_{LWRmf,i} P_{LWRmf} + \tilde{N}_{FBR,i} P_{FBR} - g_i > P_{FBR} \\ \tilde{N}_{FBR,i} & \text{otherwise} \end{cases} \tag{2.12}$$

$$\tilde{N}_{rx,i} = \left\lceil \chi_{rx,i} \frac{g_i}{q_{rx}} \right\rceil \tag{2.13}$$

$$\tilde{M}_k^{SF} = \tilde{M}_{k-1}^{SF} + \sum_{rx' \in \{rx\}} \sum_{j=0}^{\bar{T}_{rx'}^{life}} N_{rx',k-j-T_{rx'}^{prep}} \dot{m}_{rx',j} + \dot{M}_k^{LSF} \tag{2.14}$$

where:

$rx \in \{LWRmf, FBR\}$

$i \in [2000, 2100]$

$k \in \left[2000, 2100 + T_{rx'}^{life}\right]$

$\dot{m}_{rx',j} \equiv$ Net SF production during year j of operation

$T_{rx'}^{prep} \equiv$ Fuel Preparation Time

$\bar{T}_{rx'}^{life} = \min\left(T_{rx'}^{life}, k - T_{rx'}^{prep}\right)$

$P_{rx} =$ Reactor Annual Energy Production

$\tilde{M}_{2000}^{SF} =$ Initial Legacy Fuel Inventory

$\dot{M}_k^{LSF} =$ legacy reactor fuel production at year k

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.

Because the goal of the socially conscious approach is to minimize the amount of usable spent fuel remaining at the end of 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.

2.3.4 Per-Reactor SF 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}. (Values are listed in the appendix, table A.3).

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

Table 2.7: 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 2.8: 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	$T_{pipeline}$	$\left(m^f_{b,P_1}\right) \frac{N_{batch}}{T_{cycle}}$
$T_{pipeline}+1$	$T_{pipeline}+T_{around}$	$\left(m^f_{b,P_2} - m^S_{b,P_1}\right) \frac{N_{batch}}{T_{cycle}}$
$T_{pipeline}+T_{around}+1$	$T_{pipeline}+2T_{around}$	$\left(m^f_{b,P_3} - m^S_{b,P_2}\right) \frac{N_{batch}}{T_{cycle}}$
$T_{pipeline}+2T_{around}+1$	$T_{pipeline}+3T_{around}$	$\left(m^f_{b,P_4} - m^S_{b,P_3}\right) \frac{N_{batch}}{T_{cycle}}$
$T_{pipeline}+3T_{around}+1$	$T_{pipeline}+4T_{around}$	$\left(m^f_{b,P_5} - m^S_{b,P_4}\right) \frac{N_{batch}}{T_{cycle}}$
$T_{pipeline}+4T_{around}+1$	$T_{lifetime}$	$\left(m^f_{b,P_5} - m^S_{b,P_5}\right) \frac{N_{batch}}{T_{cycle}}$

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 [12]). 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 2.7 and 2.8 below. 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.

2.3.5 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 2.9. 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, $\dot{M}^{LSF}_k$.

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 appro-

Table 2.9: 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.

priate 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 Reactos` [sic] category.

2.4 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 function (described below).

2.4.1 Perturbation Algorithm

The simulated annealing algorithm explores the configuration space by repeatedly perturbing the inputs to the simulation. The perturbation algorithm follows a simple, two step approach. First, suitable, randomly sampled LWRmf reactor within the planning horizon is replaced by the appropriate number of FBRs. Then, if this substitution causes a shortfall in the forecast of available LWR spent fuel, an LWRmf reactor is reinserted at different randomly sampled year.

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.

$$\Delta = \max\left(\left\lceil \frac{P_{LWR} - P_{surplus}}{P_{FBR}} \right\rceil, N_{FBR}\right) \tag{2.15}$$

$$\tilde{\chi} = \frac{1}{1 + \left(\frac{N_{LWR} - 1}{N_{FBR} + \Delta}\right)\frac{P_{LWR}}{P_{FBR}}} \tag{2.16}$$

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.

$$\Delta = \max\left(\left\lfloor \frac{P_{LWR} + P_{surplus}}{P_{FBR}} \right\rfloor, N_{FBR}\right) \tag{2.17}$$

$$\tilde{\chi} = \frac{1}{1 + \left(\frac{N_{LWR}+1}{N_{FBR}-\Delta}\right)\frac{P_{LWR}}{P_{FBR}}} \tag{2.18}$$

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.

2.5 Limitations of this Calculation

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 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 been fully burned; 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.

2.6 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 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 ideal for developing and testing systems dynamics models, it is poorly suited for the kind of repetitive, large-batch operation that the simulated annealing algorithm requires. Fortunately, the underlying Powersim Simulation Engine exposes a series of properties and methods that allow the programmer to interface directly with the simulation engine without invoking any unnecessary elements. Unfortunately, the Powersim Simulation Engine is currently only available for use on the Microsoft Windows operating system; this precludes the use of on campus High Performance Computing resources, which are 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 [6]. 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 runs 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.

2.6.1 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 [16] 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 the process each process signals its readiness to update. (Currently, values of 75 and 30, respectively, are being

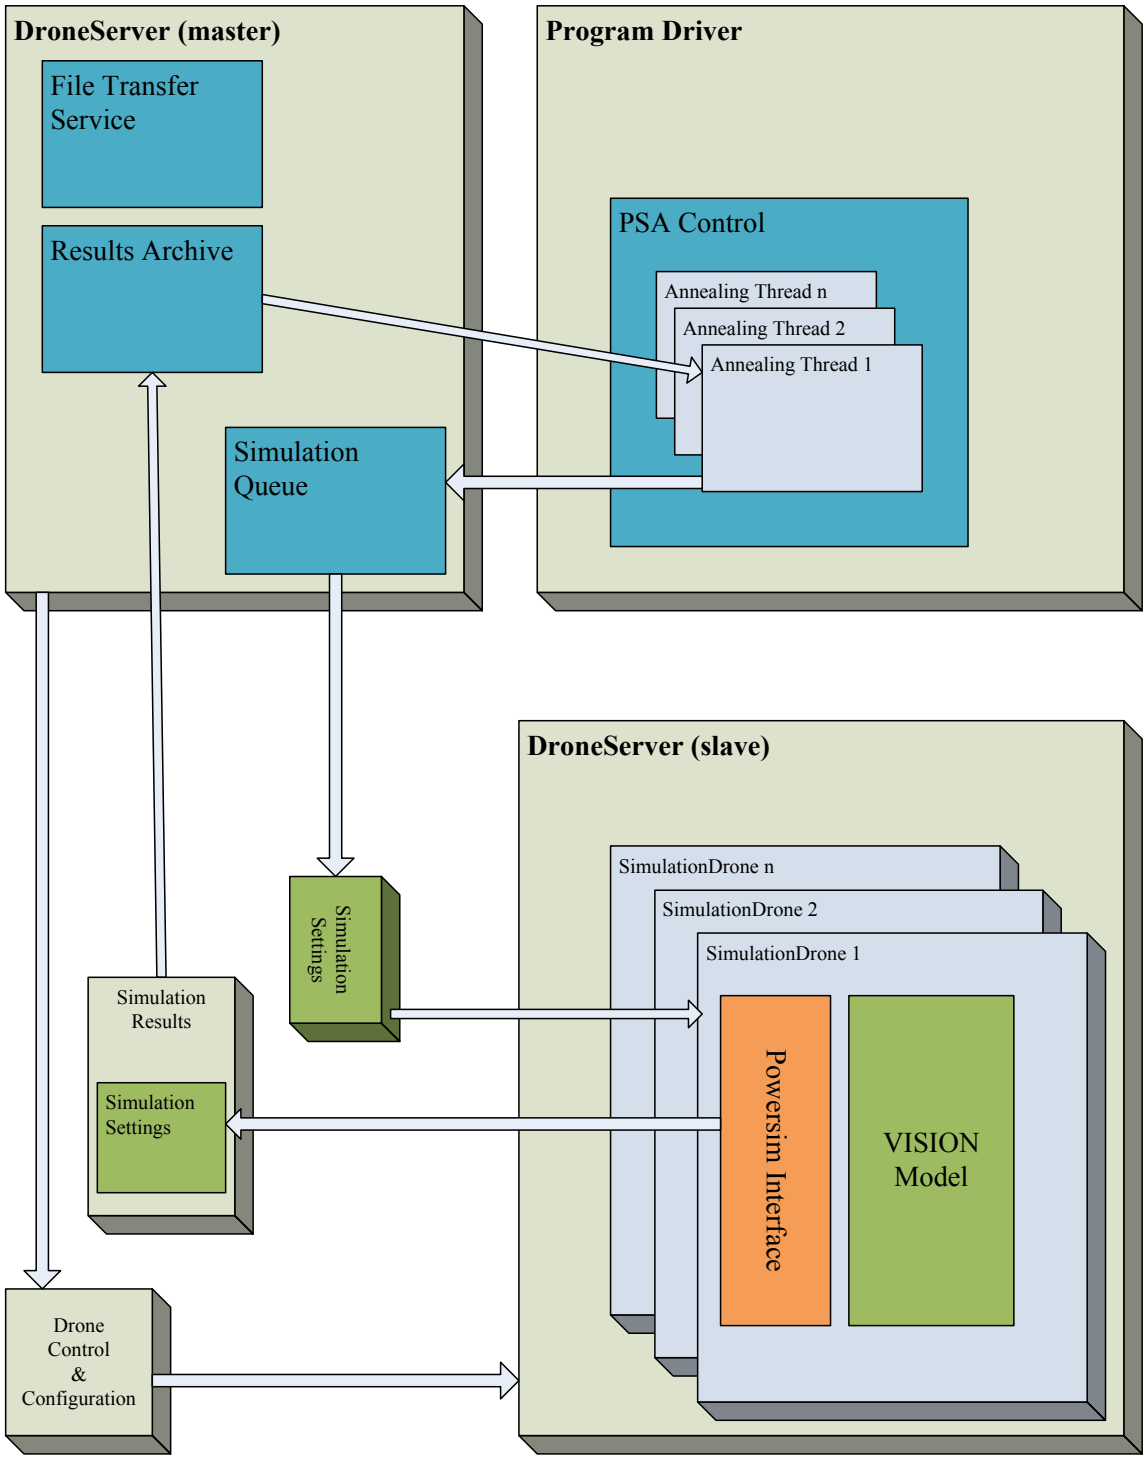

Figure 2.1: Parallel Simulated Annealing Framework Design Diagram

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.

$$T_0 = B_T \sigma_F \tag{2.19}$$

Currently, B_T has been assigned a value of 2.0.

As is it 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 2.20 below to ensure that the expectation value of the new distribution is within one standard deviation of the previous distribution.

$$T_{k+1} = T_k e^{-\lambda_H T_k / \sigma_F} \tag{2.20}$$

The value of λ_H, defined by the user, is always less than 1.0, and is currently set at 0.7. 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 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.

2.6.2 Constraint Multipliers

During an optimization cycle, the penalty constraints are incorporated into the objective function through the use of penalty multiplier values, resulting in the so-called *Augmented Objective Function*, $\overline{F}_{obj}$.

$$\overline{F}_{obj} = F_{obj} + \sum_i \lambda_i \Theta_i \tag{2.21}$$

The augmented objective function, thus penalized, 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 borrowed from the `FORMOSA-B` code [3, 10], which aims to achieve, in even logarithmic increments, a specified total reduction factor in the mean constraint violation value. First, define the average constraint violation reduction factor

Table 2.10: Optimization Terms Defined

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
$f_{\tilde{N}}$	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.

for cooling step $k+1$,

$$f_{k+1} = \Lambda_{k+1}/\Lambda_k \tag{2.22}$$

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

$$C_{k+1} = \frac{\lambda_{k+1}\Lambda_{k+1}}{\lambda_k\Lambda_k} \tag{2.23}$$

Substitute 2.22 into 2.23 and solve for λ_{k+1} to get

$$\lambda_{k+1} = \frac{C_{k+1}}{f_{k+1}}\lambda_k \tag{2.24}$$

Assume that C_k and f_k remain roughly constant throughout the cooling cycle, and require that by cooling step $\tilde{N}$, (defined below) that an overall specified reduction factor of $f_{\tilde{N}}$ 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

$$f_{k+1} = 10^{\left(\frac{\log\left(f_{\tilde{N}}\Lambda_0/\Lambda_k\right)}{\tilde{N}-k}\right)} \tag{2.25}$$

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.

$$\tilde{N} = k + \chi_p\left[\frac{L_{lngth} - N_{sampled}}{L_{chain}N_{thread}}\right] \tag{2.26}$$

The constant parameter $\chi_p < 1$ is specified by the user.

Chapter 3

Progress to Date

To date, the simulated annealing framework implementation is largely complete and tested. Some additional work will be required to fine-tune the cooling parameters and to implement the multi-objective optimization capability.

Extensive modifications have been made to the VISION model to reduce its run-time and to ensure that it functions according to the assumptions and requirements of the optimization methodology. Among these modifications are the removal of all output-oriented spreadsheet connections, the streamlining of most of the embedded `VBFUNCTION`-defined variables, the ability to override the built-in heuristic rules controlling 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 degree of confidence that one can have in the results of an automatic optimization scheme are directly related to one's confidence in the underlying model. Therefore, a great deal of effort has gone into testing and verifying many aspects of the VISION model. Because the number and variability of the random perturbations attempted by the annealing algorithm far exceed the bounding cases previously tested in the model, a number of modifications and clarifications were required to ensure valid results could be obtained.

The single objective optimization capability is being fine-tuned and tested using the Long-term Heat objective function; optimization testing using other objective functions is expected to begin shortly.

The Socially Conscious Approach to closing the planning horizon is also being fine-tuned and tested to ensure that it functions as anticipated and meets the requirements.

3.1 Test Cases and Results

Throughout the development work to date, numerous test cases have been performed a number of times with varying, but improving, degrees of success. In order to keep the voluminous output data in order, each run is named after the date on which it was started or the run which it is continuing. Thus, as the run described below was the first run started on 27 July, it is named *7-27a*.

Overall, the results of test 7-27a indicate that the optimization algorithm performs as intended.

Although subsequent testing has revealed that there were several inconsistencies in the forecasting model at the time of the test, it is not believed that this invalidates the optimization method itself, but it does indicate the potential for greater improvement in the results obtained. This work is currently ongoing. As seen in figure 3.1, the fraction of samples accepted at each cooling step tends to decrease as the optimization proceeds. The relatively high final acceptance ratio of approximately 40% indicates that either the algorithm had not yet converged, or that there exist a number of nearby optimal solutions. However, the perturbation distances of accepted histories, shown in figure 3.2 remain rather uniformly distributed, indicating that the previously noted high final acceptance ratio is likely due to insufficient cooling.

Table 3.1: Test 7-27a Settings

Parameter	**Value**
Objective	Long-term Heat
N_{thread}	16
L_{lngth}	8000
L_{chain}	75
L_{tran}	40
L_{surv}	50
χ_p	0.75
λ_H	0.70

Table 3.2: Test 7-27a Results

Parameter	**Value**
Total Samples	9251
Accepted Samples	6130
Rejected Samples	3121

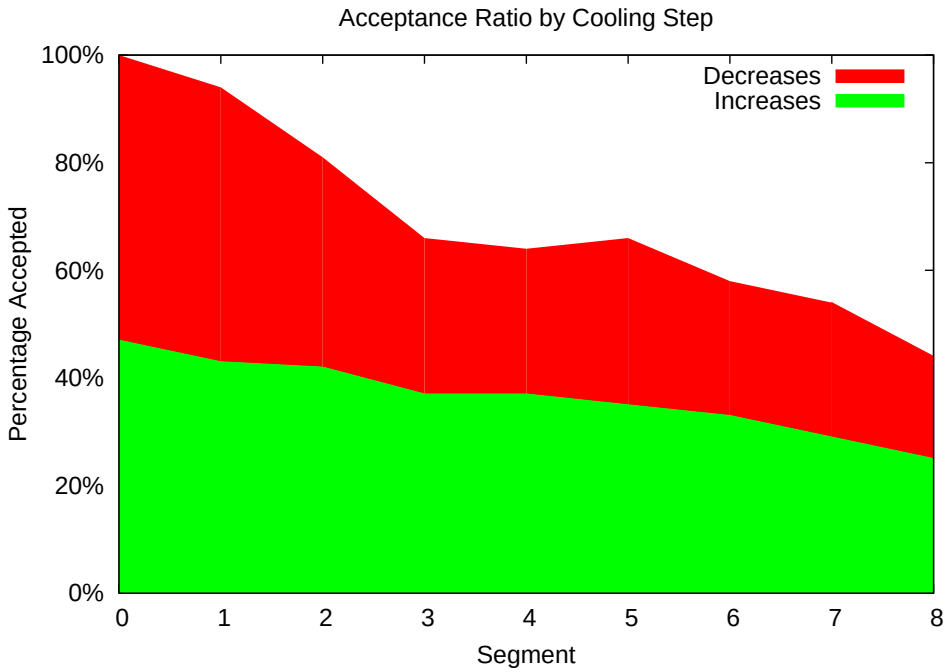

Figure 3.1: Test 7-27a Percentage Accepted by Cooling Step

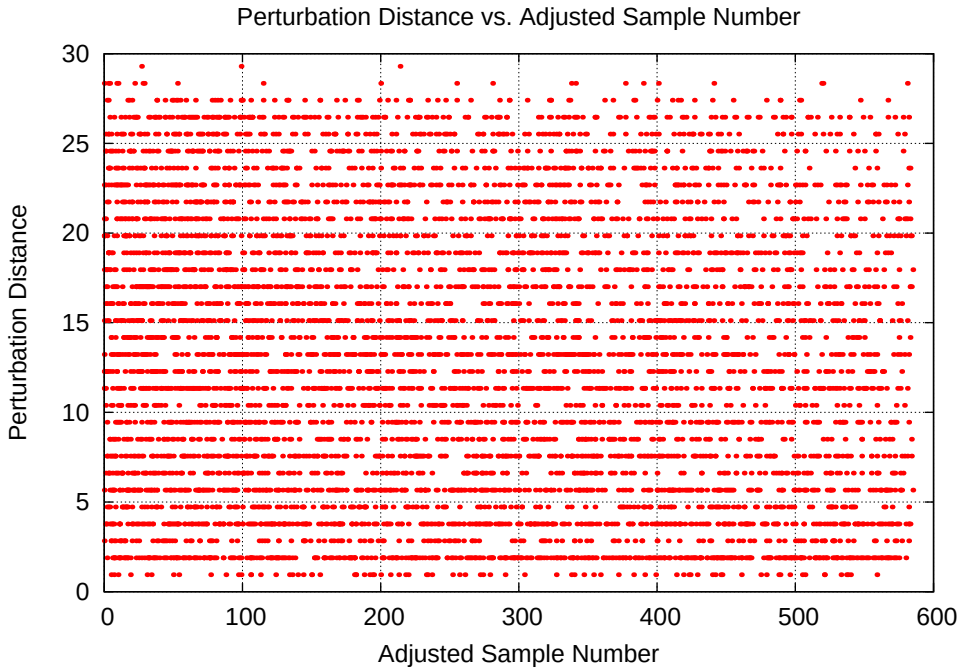

Figure 3.2: Test 7-27a Configuration Perturbation Size by Sample Number

Many of the following plots are paired so that the upper plot shows the mean and deviation of all accepted solutions along with the best and worst solutions from the archive, listed by annealing segment. The lower plot shows the saved values for each accepted sample according to the adjusted sample number. The adjusted sample number is computed simply by adding the maximum adjusted sample number of the previous cooling step to the index of the given sample on a parallel process relative to the start of the last annealing step.

Figures 3.3 and 3.4 show that the optimization algorithm is indeed working to minimize the Long-term Heat objective function value with each cooling step. In particular, it can be noted in figure 3.3 that for the final cooling steps the best and worst archived solutions have nearly identical objective function values, indicating that there exists a family of nearby optimum solutions. In addition, figure 3.4 shows that although after the fifth cooling step the minimum accepted objective function value does not markedly decrease, the average value of the accepted solutions does continue to drop.

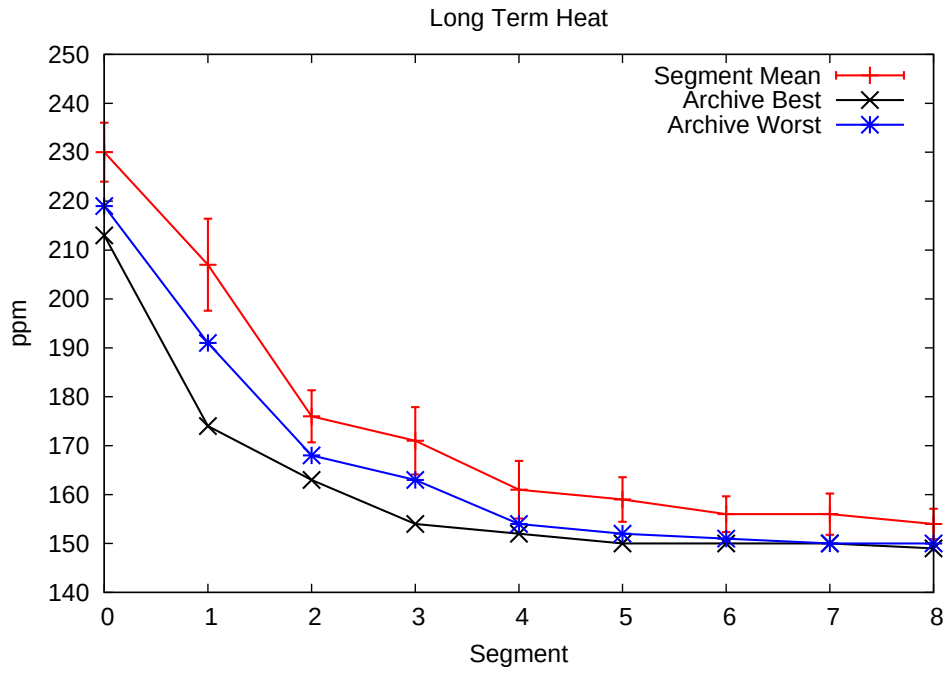

Figure 3.3: Test 7-27a Accepted Sample Average Long Term Heat Objective by Cooling Step

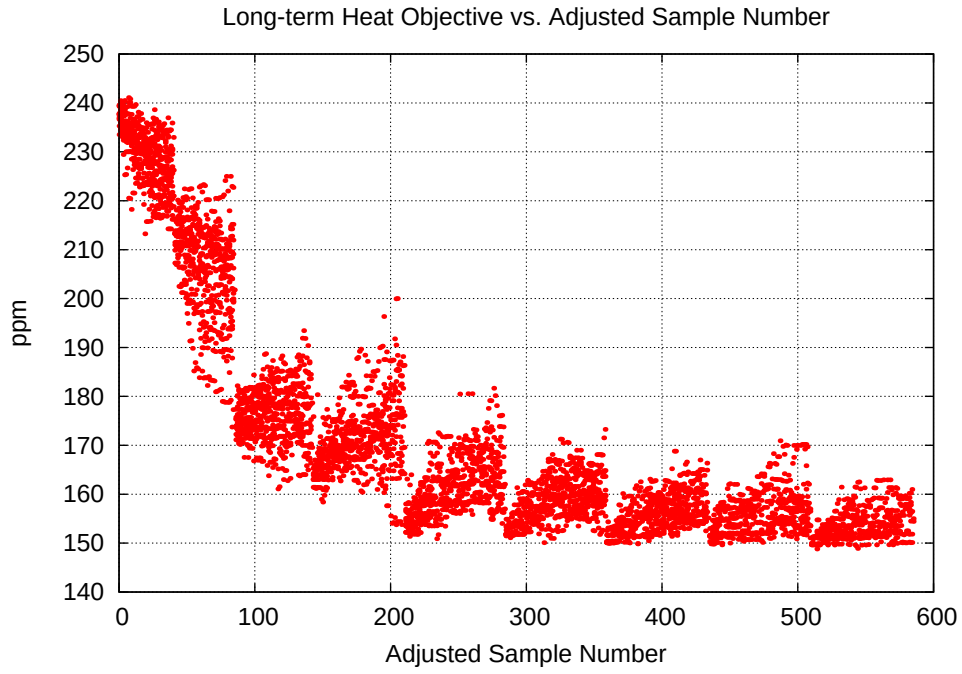

Figure 3.4: Test 7-27a Long Term Heat Objective Value by Sample Number

One feature of the perturbation algorithm is that while trying to minimize the leftover LWR SF inventory, it might select a configuration that will cause some FBRs to run out of fuel. While such configurations are obviously non-optimal, their acceptance as part of the annealing chain may be useful to promote the full coverage of the configuration space. Thus, the idled (or unused) capacity penalty function is used to progressively guide the optimization away from these configurations as the cooling cycle evolves. Figures 3.5 and 3.6 indicate that while there were a small number of accepted configurations with constraint violations, none were accepted into the solution archive, and the segment mean constraint violation value was successfully reduced during the cooling cycle. Figure 3.7 shows how the penalty multiplier value was adaptively increased throughout the cooling cycle in order to further decrease the number of configurations accepted in spite of constraint violations.

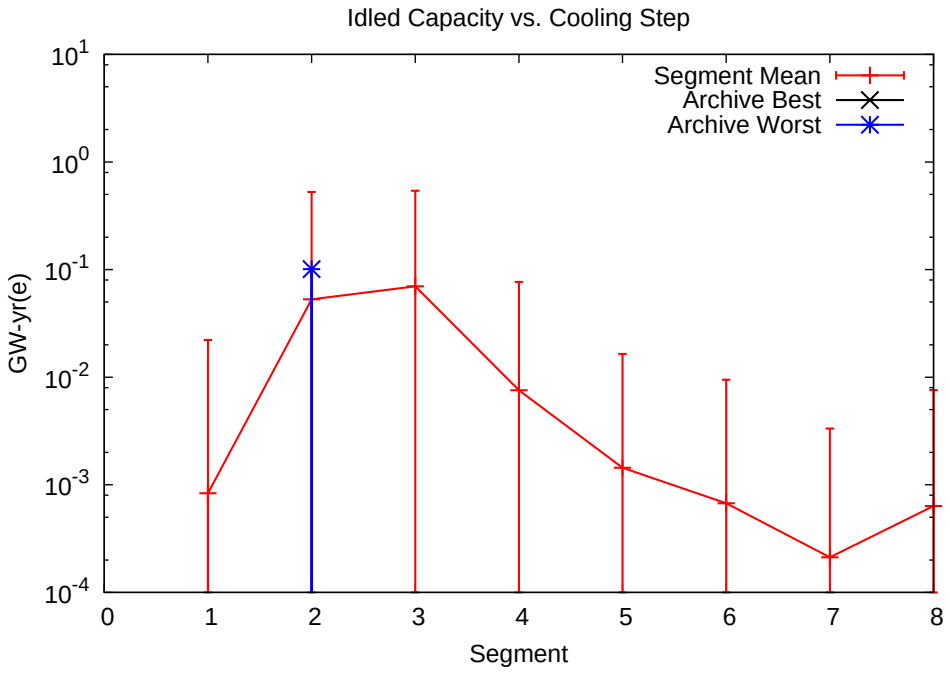

Figure 3.5: Test 7-27a Accepted Sample Average Unused Capacity Constraint by Cooling Step

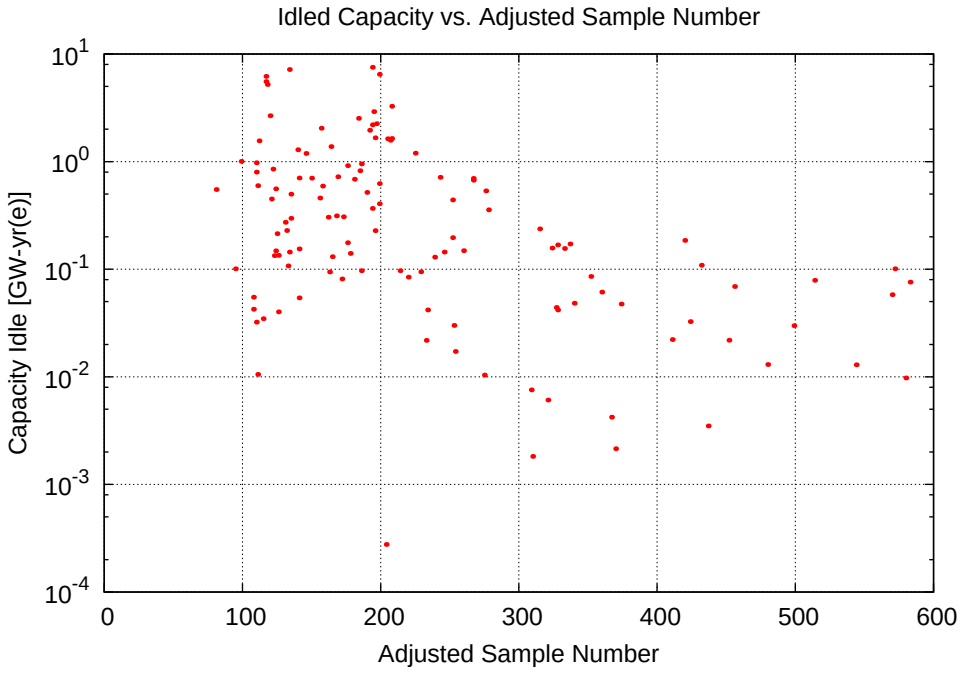

Figure 3.6: Test 7-27a Unused Capacity Constraint by Sample Number

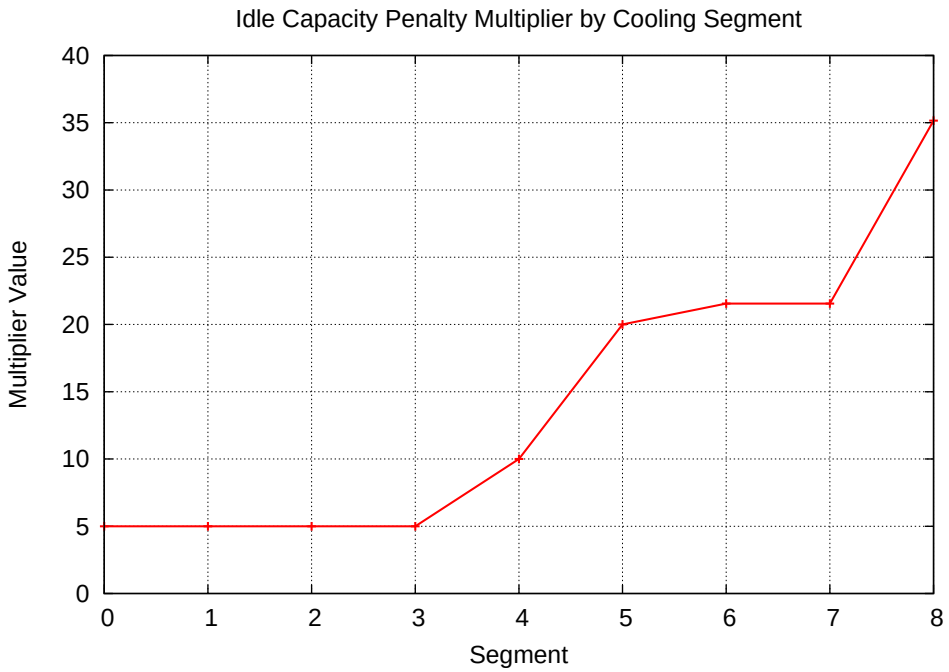

Figure 3.7: Test 7-27a Idle Capacity Penalty Multiplier by Cooling Step

Figures 3.8 and 3.9 indicate while the overall uranium utilization for all accepted solutions shows a slight variation from one configuration to another, both the best and worst archived solutions had identical results for most cooling steps.

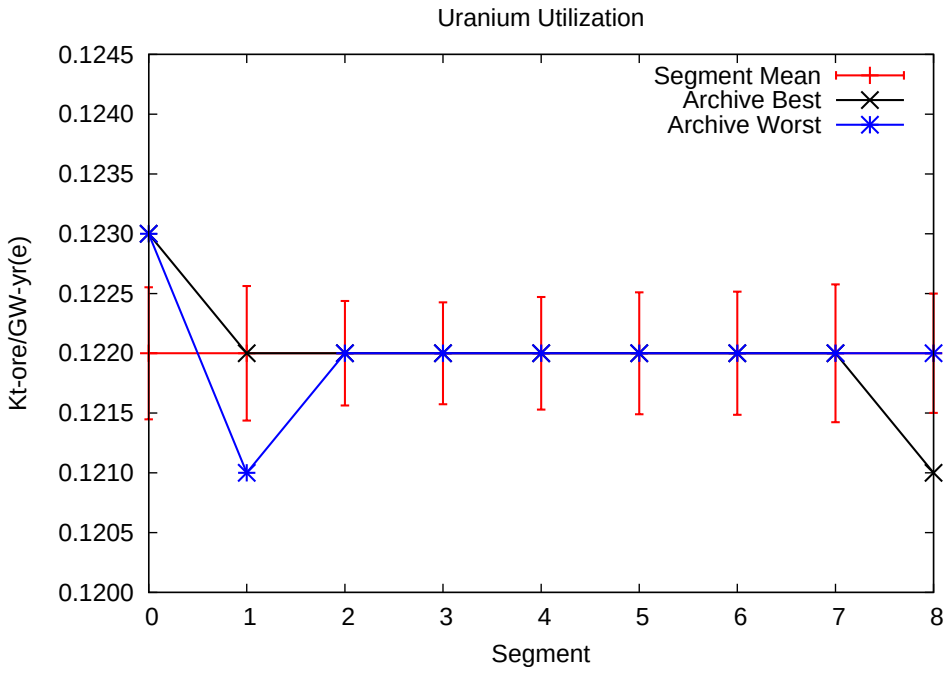

Figure 3.8: Test 7-27a Accepted Sample Average Uranium Utilization Objective by Cooling Step

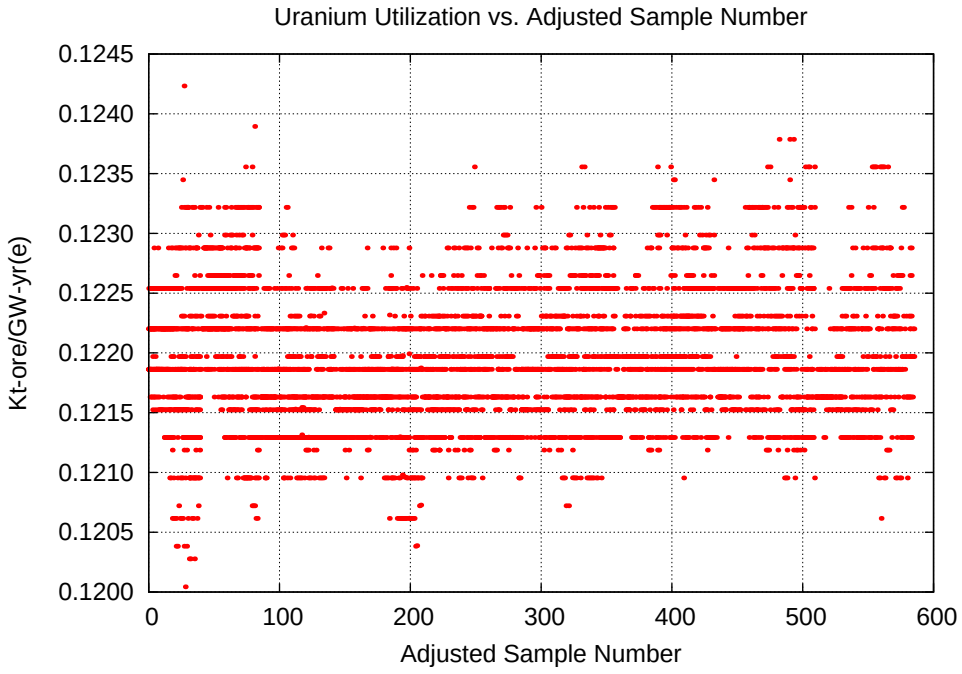

Figure 3.9: Test 7-27a Uranium Utilization Objective by Sample Number

Figures 3.10 and 3.11 show that, after an initial drop, the total weapons usability metric tends to hover near the same value as the optimization progresses, indicating that it is loosely coupled to the long-term heat objective value (as would be expected, because many of the actinides active in the weapons metric also emit decay heat and are produced and consumed through similar processes).

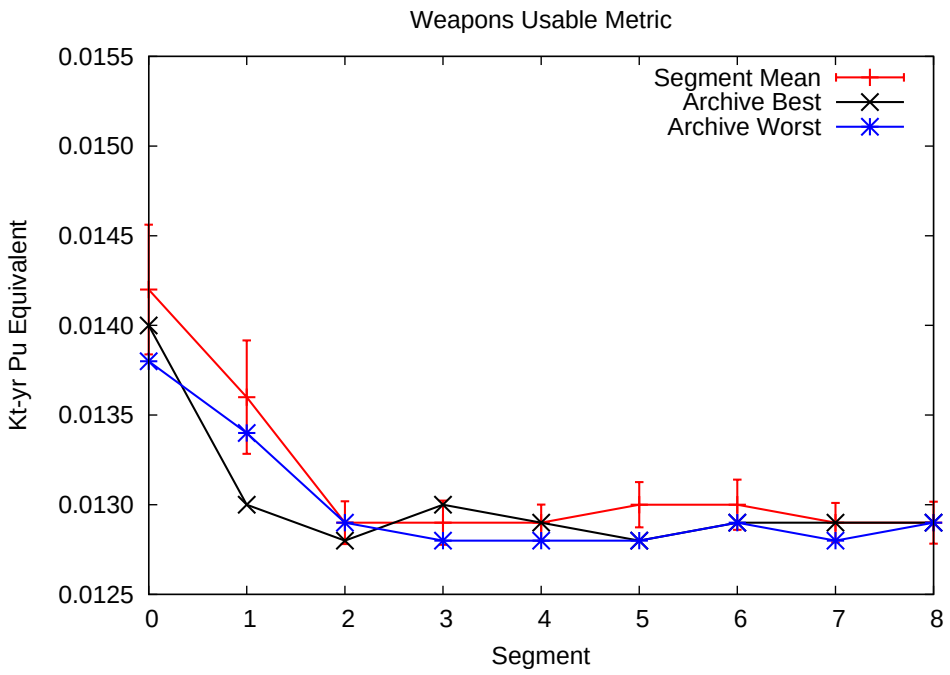

Figure 3.10: Test 7-27a Accepted Sample Average Weapons Usable Objective by Cooling Step

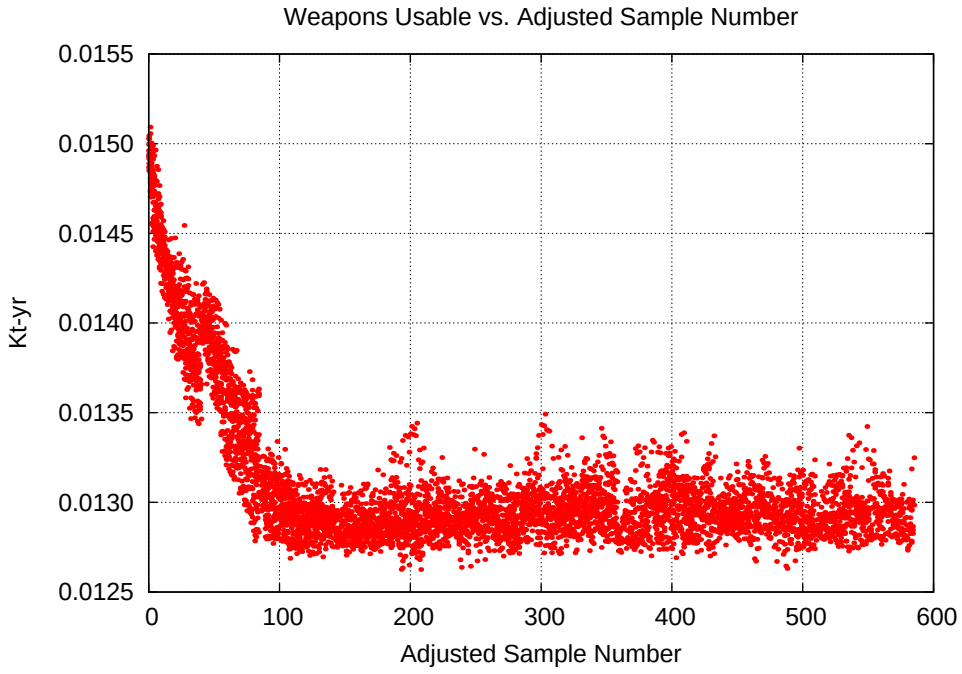

Figure 3.11: Test 7-27a Weapons Usable Objective by Sample Number

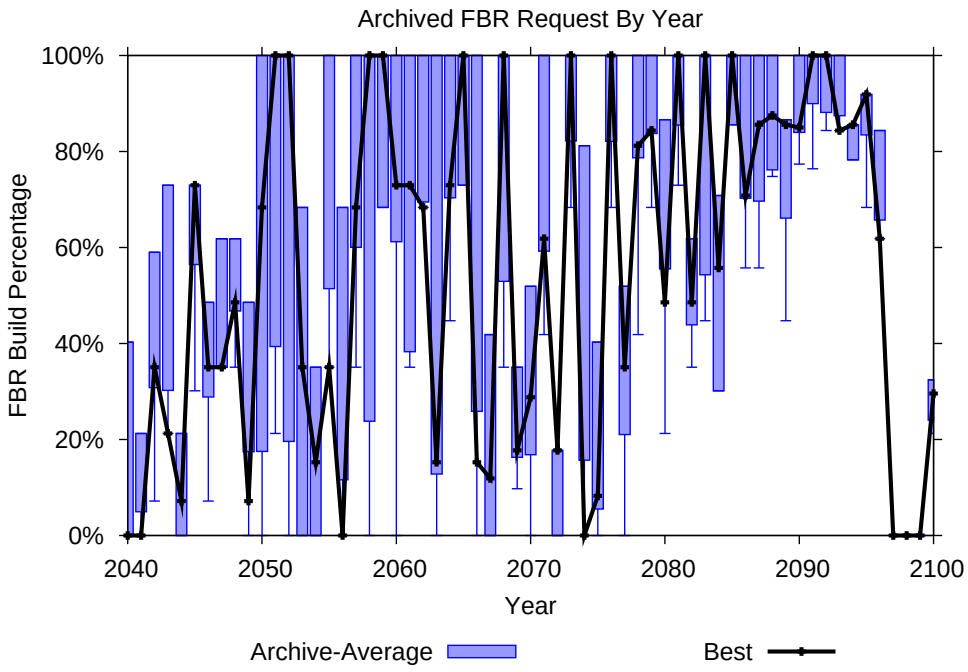

Figure 3.12: Test 7-27a Best-Solution Archive FBR Request Distribution

Finally, figure 3.12 depicts the distribution in the requested FBR build profile in the archive of configurations from the final cooling step. The error-bars indicate the extent of the accepted configurations, while the solid bars indicate the average distribution; the best configuration is plotted as a solid black line. Because of the increasing electrical demand, a given variation in the percentage of FBR capacity requested toward the end of the planning horizon will correspond to a greater variation in the actual number of FBRs requested than it would at the beginning of the planning horizon. One interesting feature to note is that all of the archived configurations show a generally increasing FBR request from 2070-2090, yet request zero FBRs for the years 2097-2099, followed by a slight uptick at 2100. The reason for this is not yet known. Clearly from a new build viewpoint, a smooth variation in the start of new builds per year with time is desired, which can be addressed by a penalty function.

Chapter 4

Plan of Future Work

There are several goals remaining in this work – first among them are the complete specification, implementation, and testing of the economics objective function. Furthermore, it is desired to refine the Weapons Usable objective function to account for the relative vulnerabilities of various types of fuel cycle facilities. As there are a number of objective functions available, it is also desired to include a multi-objective simulated annealing (MOSA) capability within the code; this will be used to estimate the trade-off surfaces that exist between select objectives. Finally, the Socially Conscious methodology restricts the decision space to an area that minimizes residual TRU, suggesting an eventual phase-out of the nuclear enterprise. We wish to explore a Dynamic Growth case, where all TRU existing prior to the end of the planning horizon has been consumed and a stationary reactor distribution at the end of the planning horizon will accommodate an assumed constant future growth rate without introducing any additional TRU disposal burden. The specifics of this scenario have not yet been finalized, and are subject to change.

Bibliography

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

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

[5] 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.

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

[7] 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.

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

[9] 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.

[10] 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.

[11] 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.

[12] 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.

[13] 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.

[14] 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.

[15] 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.

[16] 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.

[17] 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.

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

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 recycled LWR spent fuel, transitioning to a time-varying mixture of recycled FBR and LWR spent fuel as they age.

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.

In order for the Socially Conscious approach to forecast LWR SF inventories efficiently, it must compute the mass of TRU required for each batch of spent and fresh recycled fuel. These are listed below in table A.3.

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]	0.000	0.000	0.000	0.999	0.000	0.000	0.000	0.000	0.000	0.000	0.000
RU[4]	0.000	0.000	0.999	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Iodine	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.999	0.000
Gas	0.999	0.000	0.000	0.000	0.999	0.000	0.999	0.000	0.000	0.000	0.000
Tc	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.749	0.000	0.000
Cs, Sr	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.999	0.000	0.000	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]	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.250	0.000	0.000

Fast Metallic Fuel

Stream	He	FP[1]	U	TRU[2]	^{3}H	C	Kr	Sr	Tc	I	Cs
RU PNAC	0.000	0.000	0.034	0.999	0.000	0.000	0.000	0.000	0.000	0.000	0.000
RU	0.000	0.000	0.934	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Gas	0.999	0.000	0.000	0.000	0.999	0.000	0.999	0.000	0.000	0.000	0.000
Cs, Sr	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.999	0.000	0.999	0.999
Lanth FP	0.001	1.000	0.000	0.001	0.001	1.000	0.001	0.001	0.001	0.001	0.001
Zr, SS[6]	0.000	0.000	0.034	0.000	0.000	0.000	0.000	0.000	0.999	0.000	0.000

[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%
Load Factor	90%	85%
$T_{lifetime}$	60 yr	60 yr
T_{cycle}	1 yr	0.434 yr
N_{batch}	5	7.3
Fuel Form	oxide	metallic
Conversion Ratio	-	0.5
Total Core Mass	100 tonnes	9.3 tonnes
Discharge Burn-up	51 GW-day/MT-IHM[2]	176.6-176.9 GW-day/MT-IHM[3]

[1] Conversion Ratio is defined as "the ratio of mass of transuranic isotopes created over those destroyed during fuel irradiation" [7]
[2] Legacy fuel from Gen-II reactors has only 33 GW-day/tonne burn-up
[3] MT-IHM stands for Metric Ton of Initial Heavy-Metal content.

Table A.3: TRU Batch Loading

		Tonnes/batch	
	Pass	**Feed**	**Spent**
LWR	0	0.000	0.247
FBR	1	0.374	0.300
	2	0.386	0.310
	3	0.397	0.321
	4	0.410	0.331
	5	0.422	0.343