Method and Implementation of Geometric Optimization and Salt Selection of Molten Salt Reactors.

Method and Implementation of Geometric Optimization and Salt Selection of Molten Salt Reactors.

Elliott Berg, B.Eng.

A Thesis Submitted to the School of Graduate Studies in Partial Fulfilment of the Requirements for the Degree Doctor of Philosophy.

McMaster University DOCTOR OF PHILOSOPHY (2022) Hamilton, Ontario (Engineering Physics)

TITLE: Method and Implementation of Geometric Optimization and Salt Selection of Molten Salt Reactors

AUTHOR: Elliott Berg, B.Eng. (Carleton University)

SUPERVISORS: Professors John C. Luxat, Ph.D. and Adriaan Buijs, Ph.D.

NUMBER OF PAGES: xxviii, 281

Lay Abstract

Molten Salt Reactors (MSRs) often refer to design concepts that use liquid salt at high temperature and low pressure to remove heat from the core. Within this definition of MSRs, a large range of design concepts with many fundamental differences are viable and have been explored. This thesis aims to assess, on a cost basis, the feasibility of a large range of potential configurations. Two reactor design concepts, one with fuel encased in graphite pebbles and salt coolant and another consisting of salt with dissolved uranium fuel that circulates through a graphite-moderated core are analysed. Through a shared framework, the data generated from the analyses are used to relate reactor design/construction conditions as well as specific configuration parameters to cost. Probabilistic distributions are applied to many parameters consequential to determining cost in order to perform uncertainty and sensitivity analysis.

Abstract

A concerted effort is being coordinated among many developed countries to advance new nuclear power systems for commercial deployment. This project is called the Generation 4 International Forum. Of the design categories receiving significant attention is Molten Salt Reactors (MSRs). A shared feature is the use of molten salts at high temperature and low pressure to remove heat from the core. There are many viable MSR design configurations and current designs vary considerably, not only in the general design concept but within design concepts as well.

The primary objective of this study is to optimize, on a cost basis, the salt composition and key geometric parameters of MSR design concepts. A novel framework is developed to relate reactor design/construction conditions as well as specific configuration parameters to cost. The evaluation is broad in scope and is therefore divided into several metrics of performance, direct cost, waste, safety, proliferation, modularity and feasibility (technical difficulty). Two reactor classes are examined, the salt-cooled, pebble-bed Fluoride High-temperature Reactor (FHR) and the graphite-moderated, circulating-fuel reactor. Both design evaluations involve numerous analyses to generate data used in the evaluation, such as the calculation of temperature reactivity feedback coefficients and heat exchanger size optimization. The FHR analysis includes maximum fuel temperature calculations and depletion simulations that aim to model on-line refuelling. The circulating-fuel reactor analysis includes the use of novel methods to optimize the xenon removal rate and optimize the salt channel pitch. The cost evaluation was executed with both fixed values and probabilistic distributions applied to many consequential inputs. The results obtained using fixed values provide useful insights into the effect design parameters have on cost, while uncertainty analysis provides estimates of cost uncertainty, both overall and between configurations. Sensitivity analysis breaks down the cost uncertainty into component parts. This work aims to deepen the understanding of the costs and trade-offs associated with numerous design characteristics.

Acknowledgements

I am grateful for the opportunities and experiences that have come from undertaking this project and thus following my passion. I would like to thank my supervisors, Dr. John Luxat and Dr. Adriaan Buijs for supporting me in this project, their patience and for all else that has led to this moment. My Ph.D. supervisory committee members, Dr. Shinya Nagasaki and Dr. David Novog also deserve thanks for their valuable, constructive feedback and encouragement.

My fellow graduate students deserve acknowledgement. They fostered a collegial environment where assistance was offered without hesitancy and stimulating discussion was the norm. Both this thesis and the pleasure derived from my studies have benefited from our interactions.

Financial support provided by McMaster University, the Natural Sciences and Engineering Research Council of Canada and the University Network of Excellence in Nuclear Engineering is appreciated.

I would like to thank my parents and sister for their steadfast support and patience on this journey. And finally, a special thank you to my better half, Julia for her unwavering love, encouragement and for lending me her superb editing skills.

Contents

List of Figures

[1] The interpolations are provided to highlight tends and add clarity, the shape of the curves do not have a physical meaning.

[2]The interpolations are provided to highlight trends and add clarity, the shape of the curves do not have a physical meaning.
[3]The interpolations are provided to highlight trends and add clarity, the shape of the curves do not have a physical meaning.
[4]The interpolations are provided to highlight tends and add clarity, the shape of the curves do not have a physical meaning.
[5]The interpolations are provided to highlight tends and add clarity, the shape of the curves do not have a physical meaning.

List of Tables

List of Symbols

The descriptions presented in this thesis cross over multiple disciplines which each have their own established nomenclature. This text uses signs and symbols consistent with literature, but as a consequence some may possess different meanings depending on the context in which they are presented.

Symbols from Section 6.2.1 are not included in this list. That section uses many symbols specific to that particular analysis.

A	Activity coefficient
A	Area [cm^2], [m^2]
B or Bu	Burnup
C_p	Heat capacity [kJ/kgK]
D	Diffusion coefficient [cm]
D	Diameter [cm], [m]
ff	Friction Factor
G	Gibbs free energy
Gz	Graetz number
h	Heat transfer coefficient [W/m^2]
H	Height [cm], [m]
k	Thermal conductivity [W/m-K]
k	Multiplication factor
K	Equilibrium/solubility (Henry's) constant
l	Lifetime, incl. of prompt neutrons [s]
L	Lifetime
L	Length
L	Neutron leakage
L	Neutron diffusion length [cm]
$\dot{m}$	Salt flow rate [cm^3/s]
M	Molar mass [g/mole]
Nu	Nusselt number
P	Pitch [cm]
P	Non-leakage probability

P	Power [MW], [kW]
P	Pressure [Pa]
Pr	Prandtl number
Q	Heat transfer [W]
Q	Volumetric flow rate [cm^3], [m^3]
r	Spatial location
R	Radius [cm], [m]
R	Gas constant
R	Relative cost
R	Relaxation
Re	Reynolds number
S	Distribution of source neutrons
S	Power density [MW/m^3], [kW/l]
S	Solubility
t	Time [s], [min], [hr], [d], [yr]
T	Temperature [^{o}K], [^{o}C], [^{o}F]
T	Time [s]
u	Velocity [cm/s], [m/s]
V	Velocity [cm/s], [m/s]
V	Standard volume [cm^3], [m^3]
V	Molar volume [cm^3], [m^3]
y	Fraction of salt in core
z	Axial position [cm], [m]
Subscripts	
eff	Effective
inf	Infinite
t	Total
Greek Symbols	
α	Power value property of Pareto distribution
β	Delayed neutron fraction
β_{eff}	Effective delayed neutron fraction
β_{eff}/β	Correction factor
Δ	Change in
λ	Decay constant [s$^-$1]
θ	Probability of DNs emitted inside the core
π	Neutron or Fermi age
ϕ	Flux [n/cm^{2}s]
μ	Viscosity [cP]
μ	Minimum value (for Pareto distribution)

ρ	Density
pr or p	Prompt
$1/2$	half, typically for half-life
Units and Other	
eV	Electron Volts
MeV	Mega-electron Volts
g	Grams
kW	Kilowatt
l	Litre
mk	Milli-k [$k \cdot 10^{-3}$]
$mill$	1/10 of a cent
mol or $mole$	10^{23} atoms/molecules
MWe	Mega-Watts electric
$MWth$	Mega-Watts thermal
pcm	$\Delta k \cdot 10^{-5}$
ppm	Parts per million
μ	Micro (10^{-6})
USD	U.S. dollar (2021 if unspecified)
CAD	Canadian dollar (2021 if unspecified)
\$	USD if unspecified
$[x]$	Concentration [x 10^{24}/cm^3]
$\approx$	Approximately

List of Abbreviations

ABWR	Advanced Boiling Water Reactor
AECL	Atomic Energy of Canada Limited
AHTR	Advanced High Temperature Reactor
ANP	Aircraft Nuclear Propulsion program
*AP*1000	Advanced Passive 1000 (Westinghouse/Toshiba)
APR	Advanced Power Reactor (S. Korea)
ARE	Aircraft Reactor Experiment
BoP	Balance of Plant
CANDU	CANada Deuterium Uranium
CDF	Core Damage Frequency
COLEX	COLumn EXchange
CNRS	National Centre for Scientific Research (Grenoble, France)
CNSC	Canadian Nuclear Safety Commission
DFR	Duel Fluid Reactor
DGR	Deep Geological Repository
DOE	(U.S.) Department of Energy
DN	Delayed Neutron
DNF	Delayed Neutron Fraction
DNL	Delayed Neutron Lifetime
DPA	Displacements Per Atom
EIA	Energy Information Agency
EPR	European Pressurized water Reactor
EVOL	Evaluation and Viability of Liquid Fuel Fast Reactor System
FBR	Fast Breeder Reactor
FCC	Face Centered Cubic
FF	Fuel Fraction
FFPF	Fast Flux Peaking Factor
FHR	Fluoride-salt-cooled High-temperature Reactor
FNF	Fast Neutron Flux
FOAK	First Of A Kind

FOM	Figure Of Merit
FP	Fission Product
FPD	Full Power Days
FTC	Fuel Temperature Coefficient
GFR	Gas-cooled Fast Reactor
GHG	Green House Gas
GIF	Generation IV International Forum
GL	Graphite Lifetime
HCP	Hexagonal-close-packed
HL	Head Loss
HM	Heavy Metal
HTGR	High Temperature Gas Reactor
HTTR	High Temperature engineering Test Reactor
IMSR	Integral Molten Salt Reactor
IMSBR	Indian Molten Salt Breeder Reactor
IPK	Institute for Solid-State Nuclear Physics (Germany)
INES	International Nuclear Event Scale
JAEA	Japan Atomic Energy Agency
LCOE	Levelized cost of energy
LD	Lethal Dose
LEU	Low Enriched Uranium
LFR	Lead-cooled Fast Reactor
LFTR	Liquid Fluoride Thorium Reactor
LHS	Latin Hypercube Sampling
LPSC	Laboratory of Subatomic Physics and Cosmology (Grenoble, France)
LMTD	Log-Mean Temperature Difference
LWR	Light Water Reactor
MARR	Minimum acceptable rate of return
MCFR	Molten Chloride Fast Reactor (Terrapower)
MCSFR	Molten Chloride Salt Fast Reactor (Elysium Industries)
MOSART	Molten Salt Actinide Recycler & Transmuter
MSR	Molten Salt Reactor
MSRE	Molten Salt Reactor Experiment
MSBR	Molten Salt Breeder Reactor
MSTW	Molten Salt Thermal Wasteburner
MSFR	Molten Salt Fast Reactor
MSRP	Molten Salt Reactor Program
NOAK	Nth of a kind
NPP	Nuclear Power Plant

NRC	Nuclear Regulatory Commission (USA)
$O\&M$	Operation and Maintenance
$ORNL$	Oak Ridge National Laboratory
$PB-FHR$	Pebble-Bed Fluoride High-temperature Reactor
PD	Pressure Drop
PF	Packing Fraction
PNL	Prompt Neutron Lifetime
PP	Pumping Power
PPF	Power Peaking Factor
PWR	Pressurized Water Reactor
$(R)CC$	(Rank) Correlation Coefficient
$RETRU$	Rare Earth and TRansUranic
RR	Refuelling Rate
$SAMOFAR$	Safety Assessment of the MOlten salt FAst Reactor
$SCWR$	Supercritical Water-cooled Reactor
SFR	Sodium-cooled Fast Reactor
SNF	Spent Nuclear Fuel
SS	Stainless Steel
SSR	Stable Salt Reactor
$SSR-U$	SSR-Uranium (graphite moderated)
$SSR-W$	SSR-Wasteburner (fast spectrum)
SWU	Separative Work Unit
TAP	TransAtomic Power
TC or $TFRC$	Temperature Reactivity Feedback Coefficient
TMI	Three Mile Island
$TMSR-LF$	Thorium Molten Salt Reactor - Liquid Fuel (China)
$TMSR-SF$	Thorium Molten Salt Reactor - Solid Fuel (China)
$TRISO$	TRi-structural ISOtropic
TSL	Thermal Scattering Library
VHC	Volumetric Heat Capacity
$VHTR$	Very High Temperature Reactor

Declaration of Academic Achievement

This study was conducted by the author over the years 2017-2022. The work contained herein is the author's own except where stated otherwise in the text. This work has not been submitted for any other degree or professional qualification. The supervisory committee of Professors John Luxat, Adriaan Buijs, Shinya Nagasaki and David Novog provided advice and suggestions.

The data presented in this thesis was calculated or generated in simulations carried out by the author. Material properties and other such input data are explicitly referenced in the text. Missing property data, several cost-related input parameters and probability distributions, reactor model design choices and other facets of the analysis for which there are no established values were made though the synthesis of literature and/or relying on the author's personal judgement. The data analysis and interpretation are the author's own work with feedback from the supervisory committee members. The general structure and means of estimating costs for MSRs over a wide range of design parameters through the *Cost Estimation Framework* are (with guidance) of the author's own conception. All generated data was produced from models constructed by the author using the Monte Carlo neutron transport code, SERPENT 2 or Matlab.

Chapter 1

Introduction

Molten Salt Reactors (MSRs) often refer to a category of nuclear reactor that has liquid salt with dissolved fuel at high temperature and low pressure circulating through the core. The definition of a MSR can be broader than that to include all nuclear reactors that uses liquid salt at high temperature and low pressure to remove heat from the core. This more general definition, used by leading research organizations including Oak Ridge National Laboratory (ORNL) [1], allows for a large range of reactor design concepts with many fundamental differences.

The MSR was conceived of in the 1940's and was the leading candidate design for use in aircraft propulsion due to the high temperature and low weight. Though the project was cancelled prematurely, the MSR concept had proven its viability. Following the cancellation of the Aircraft Reactor Experiment (ARE), the focus shifted to civilian use, through the development of the Molten Salt Reactor Experiment (MSRE). In the 1960's a circulating fuel, graphite moderated MSR was operated for 4 years, to serve as an experiment for the development of a Molten Salt Breeder Reactor. The program eventually was cancelled and MSR development stagnated.

Interest has since renewed and there are many conceptual designs. The assortment of current designs concepts are diverse, some are based on the MSRE, some use TRISO fuel particles with a salt coolant, some have stationary molten salt fuel and yet other designs have circulating-fuel with a fast neutron spectrum.

1.1 Motivation

Limited resources, the threat of climate change and the lack of other scalable, economic, non-emitting energy sources, places ever-growing pressure on society to increase nuclear power production. The last three decades have been successively warmer than any preceding decade since 1850 and the data show a warming trend of 0.85 ^{o}C since 1880. The largest contribution to warming is by the increase in the atmospheric concentration of CO_2. Moving toward low carbon electricity production can have a major impact as the sector accounts for approximately 25% of green house gas emissions [2]. Potential solutions include renewable energies, like dispatchable hydro as well as intermittent wind and solar. Yet hydro faces the challenge of

progressively less suitable locations being available and the intermittent sources contain hidden costs that lead to difficulty with scale [3]. Though the world contains well-distributed and plentiful coal resources for the reasonably foreseeable future, burning coal emits CO_2 and other pollutants at a high intensity. Cleaner-burning natural gas with a reserves:production ratio of 52.6 [4] as well as an uneven distribution of reserves still emits CO_2 and has serious supply and price risks. Nuclear power is a solution that can address large energy requirements without jeopardizing health, the climate or grid reliability.

Unfortunately for the industry, nuclear power has generally struggled financially since the wave of rapid construction in the 1970's. The result has been an absence of reactor construction starts since the 1970's in the USA and the 1980's in France, Canada and (West) Germany. The approximately total increase in cost (inflation-adjusted) in these nations is; 300% from 1967-1978 in the USA, 75% from 1971-1991 in France, 200% from 1973-1983 in West Germany and 60% from 1971-1986 in Canada [5]. In non-Western European, non-North American nations, only one new nuclear power plant has delivered power since 1991. The Olkiluoto 3 reactor in Finland has recently come on-line [6] and five power reactors are under construction at current sites: Flamanville 3, Hinkley Point C1 & C2, and Vogtle 3 & 4. The reasons for the small number of reactor builds, construction delays and cost overruns are undoubtedly complex and unclear, but one particular technical feature appears to be an important underlying cause of cost increases and the resulting stagnation of the industry: The high-pressure core. This claim stems from cost increases that are very likely attributable to regulatory changes. For instance, reactors that were under construction during Three Mile Island (TMI) and eventually completed afterwards experienced median costs 2.8x higher than reactors that received operating licenses before TMI [5]. This phenomenon of costs being highly susceptible to the regulatory environment demonstrates the value of having a passively safe reactor operating at atmospheric pressure - generally Molten Salt Reactor designs inherently possess both of these features and therefore have the potential to lower costs and accelerate nuclear power deployment.

The motivation for evaluating MSR design concepts is two-fold: Education and the forging of consensus. A broad, objective investigation of MSRs that does not strictly focus on an aspect of technical performance, but relates performance to key practical cost considerations is uncommon is the literature. This thesis is intended to help develop a broad conceptual understanding of MSRs, through the optimization of several key design parameters of two MSR design types. As the understanding of various MSR design types expands, particularity relating to the advantages and challenges of design configurations, it is fair to suspect that support will coalesce around a smaller number of designs – thereby increasing the likelihood of MSR construction.

1.2 Research Overview

The objective of this dissertation is to provide a framework for the evaluation of, on a cost basis, MSR design types and to implement the framework upon two design types. The implementation involves the generation of wide-ranging reactor performance input data and a program that performs the evaluation. Through the

assignment of probability distributions to the set of cost related input parameters defined in the framework, the sensitivity of total cost to each of the input parameters is determined. This is a key result as it provides insight into the sources of cost uncertainty and therefore insight into which aspects of reactor design, construction and operation to direct both cost estimation and reduction efforts. The total cost is broken down into numerous components to provide an understanding of the configuration-dependent sources of cost. The design can be optimized through a comparison of total costs. The two design types evaluated are the Fluoride High-temperature Reactor (FHR) and circulating-fuel, graphite-moderated MSR.

The evaluation framework provides the basis to appropriately weigh and quantify the effect of individual design characteristics, in order to arrive at cost values over a range of design configurations. The costs are divided into six categories; direct cost, waste, safety, proliferation, modularity and feasibility. The design characteristics evaluated include geometric parameters and the salt composition.

Numerous Monte Carlo neutronics simulations are run and thermohyraulic calculations conducted to determine key performance measures such as the required refuelling rate and temperature reactivity feedback for all configurations in the evaluated design space. A novel depletion methodology is devised that is capable of closely simulating the continuous refuelling of fuel pebbles in the pebble-bed FHR. The equilibrium fuel composition of the depletion simulations is used in the corresponding static simulations. Similar depletion and static neutronics simulations are conducted for the MSR as well. 1-D thermohyraulics calculations based on Nusselt number relations are performed primarily to optimize the heat exchanger size for each salt given a cost of pumping and cost of heat exchanger volume.

The results from the neutronics simulations and thermohydraulic calculations go into a comprehensive evaluation program that calculates the expected costs of examined design configurations. The cost calculation requires numerous assumptions that have been estimated in the evaluation framework. Each cost assumption is given a probability distribution and uncertainty/sensitivity analysis is performed on the input cost assumptions. Uncertainty/sensitivity analysis is conducted on the cost types (categories), waste, safety, proliferation and feasibility costs as well.

Frameworks, guidelines, and/or methodologies for the estimation of nuclear power plant costs, such as the *Cost Estimating Guidelines for Generation IV Nuclear Energy Systems* [7], have been made but none have been found that are well suited for design optimization or that include key aspects of the evaluation framework, particularly the 'cost setting' methodology. Furthermore, no comparable comprehensive design optimizations have been found that seek to account for such a large range of direct and indirect costs.

1.3 Molten Salt Reactor Advantages and Disadvantages

This research focuses on MSRs because they have advantageous qualities over conventional designs. The hopeful scenario is that the advantageous qualities ultimately result in the construction of safe and cost ef-

fective MSRs that help usher in a new wave of nuclear power deployment.

The principal advantages of all MSRs that fit the definition used in this dissertation include:

- The retention of fission products (FPs): Most FPs are soluble in molten salt, so in a solid-fuel reactor FPs are likely to be contained in the event of a fuel failure and in a dissolved-fuel reactor the radiotoxicity of the fuel salt vapour and the FP load falling out of solution is manageable.
- High boiling point / low vapour pressure: The vapour pressure of most salts under consideration have a vapour pressure at typical operating temperatures ($\approx$600-700^oC) less than 100x atmospheric and a boiling point in excess of 1300^oC.
- High efficiency: High operating temperatures result in high thermal efficiencies.

Disadvantages:

- Corrosion: The deleterious effects of corrosion depend on the salt composition and the extent of impurities. What can be said is that corrosion effects are less understood and more of a concern than in water or gas cooled reactors.
- Lack of technological development: Only two demonstration reactors, the ARE and MSRE, have been operated and molten salts are not in widespread use even in non-nuclear engineering applications, so more technological development is required in comparison to other reactor types.

Pebble-bed reactor, including the FHR, specific advantages:

- Selective refuelling: Assuming that it is possible to continuously remove the fuel with the highest burnup, the average burnup is able to be doubled.
- On-line refuelling: Removes the need for burnable absorbers, enhancing the neutron economy and results in a high capacity factor.
- Simple structure: Very little or no core structural material which allows for greater tolerance in dimensional changes of the graphite and reduces/eliminates internal structure failure risks.

Pebble-bed reactor, including the FHR, specific disadvantages:

- Relatively high enrichment required (HALEU).
- Minimal fuel-moderator separation, which results in reduced neutronic efficiency.
- The coolant temperature reactivity feedback is positive in many configurations.
- Expense of particle/pebble manufacture.

Graphite-moderated, circulating-fuel MSR specific advantages:

- On-line refuelling: Removes the need for burnable absorbers, enhancing the neutron economy and results in a high capacity factor.

- On-line FP removal: Though it is difficult and likely uneconomic for most FPs it is particularly useful for the removal of ^{135}Xe.
- Negative salt temperature reactivity feedback.
- Salt drain: In a loss-of-flow incident the fuel coolant can be made to automatically drain into a strongly sub-critical arrangement.

Graphite-moderated, circulating-fuel MSR specific disadvantages:

- Graphite lifetime: The core operational lifetime is limited by the maximum fast neutron fluence through the graphite. As per present understanding, little can be done to effect the maximum energy generated per unit of core volume.
- Circulating-fuel: When in circulation, a reduction in the effective delayed neutron fraction.
- Corrosion: The presence of FPs in the salt (notably Te) increase corrosion rates compared to 'clean' salts.

1.4 Generation IV International Forum

The history of nuclear power reactor development is commonly categorized as a series of generations, from I in the 1950s to IV currently under development [8]. Generation I reactors were mainly for demonstration purposes and had comparatively low power output. A majority of nuclear power produced today comes from Generation II reactors, that had begun construction between the late 1960s and late 1970s [9]. Generation III reactors are a development on Generation II, incorporating evolutionary improvements acquired during the operating lifetime of the Generation II designs. The first Generation III reactor to begin operation was an Advanced Boiling Water Reactor (ABWR), on the Kashiwazaki-Kariwa site in Japan in 1996 [10]. Recent improvements in economics and safety have led contemporary designs to be termed III+. Developments have occurred in different areas of the globe and the first Generation III+ design in the US is considered to be the AP1000 [11].

Future conceptual nuclear reactor designs, termed Generation IV reactors, are in various design stages. These advanced reactors are intended for widespread deployment beyond 2030 and are being designed to meet wide ranging criteria involving enhanced safety, economics, sustainability, and proliferation resistance [12]. The Generation IV International Forum (GIF) is a cooperative endeavour intended to develop the research necessary to test and make available advanced nuclear designs [13]. In 2001, experts from the countries involved and international organisations began work developing a technology roadmap for Generation IV systems. This involved defining the goals, identifying promising concepts, evaluating them, and defining their R & D requirements. By the end of 2002, the work resulted in the determination of the six most promising systems. As of the updated 2014 technology roadmap the six reactor design are: [14]

1. Gas-cooled Fast Reactor (GFR)
   The GFR is a high-temperature helium-cooled fast-spectrum reactor with a closed fuel cycle. The reference design for the GFR is currently 2400 MWt with break-even breeding.

2. Lead-cooled Fast Reactor (LFR)
   LFRs are Pb (or Pb-Bi-alloy)-cooled reactors that operate at atmospheric pressure and at high temperature, similar to MSR's. Recent development has mostly come out of Russia, through the SVBR-100 and BREST-300 designs

3. Molten Salt Reactor (MSR)
   MSR development as it relates to the GIF includes both thermal and fast spectrum circulating fuel reactors and the salt-cooled Fluoride High Temperature Reactor (FHR) concept.

4. Sodium-cooled Fast Reactor (SFR)
   The SFR uses liquid sodium as a low-pressure reactor coolant. It has advantageous thermo-physical properties, however sodium does react chemically with air and water. Plant size options under consideration by the GIF range from small (50 MWe) to larger plants (1500 MWe). SFRs have been built before including the Phenix (France) and the BN-600 (Russia).

5. SuperCritical Water-cooled Reactor (SCWR)
   SCWRs are high-temperature, high-pressure, light water reactors that operate above the thermodynamic critical point of water (374 ^{o}C, 22.1 MPa). The design may be pressure vessel or pressure tube based and either light or heavy water moderated.

6. Very High Temperature Reactor (VHTR)
   The VHTR is a graphite-moderated, helium-cooled reactor with a thermal neutron spectrum that can supply heat up to 950 ^{o}C. It is the next development of gas-cooled reactors and has either a prismatic-block type or a pebble-bed core.

1.5 Thesis Outline

Including the introduction and conclusion, this thesis is comprised of seven chapters. Chapters 2 through 4 consist of discussion and results that apply to all MSRs, while Chapters 5 and 6 apply the content from Chapters 2 through 4 to two different design classes. Chapter 7 is the conclusion, which discusses results, compares the reactor classes, draws conclusions and discusses future work.

Chapter 2 introduces the evaluation program and describes the cost estimation framework. The evaluation program has a shared general structure among reactor classes and uses the cost estimation framework to determine costs. The exact setup, and assumed input values are specific to the reactor design class.

Chapter 3 explores the factors that are involved in the selection of alternative (non-LiF-BeF_2) salt compositions. It begins with a summary of historical literature, particularly the MSRP and lists salt compositions of current designs. Important considerations for the salt selection in any design class include, the interaction with graphite and the metal containment, (corrosion) heat exchanger requirements, cost of constituent materials, and neutron activation. Neutronics related effects, particularly the moderating effect and parasitic absorptions, are not discussed since the effects are best understood through interpreting the results of

design-specific full-core simulations. A discussion and list of property data is provided in Appendix C.

Chapter 4 discusses uncertainty and sensitivity analysis. It consists of a background literature discussion as well as an explanation and justification of the setup. Examples from the FHR and MSR are provided for clarity. A Monte Carlo approach is taken that assumes a 'black box model.' Each of the input variables is assigned one of two possible probability distributions. Select input variables are correlated and Latin Hypercube Sampling is used.

Chapter 5 consists of all content involved in the application of the evaluation program to the FHR. It begins with a description of the reactor and calculation of key parameters such as fuel temperature. The setup of depletion simulations that model the refuelling scheme as well as static simulations at (near) equilibrium fuel composition are described. The results of the simulations are then presented. The final section discusses aspects of the evaluation specific to the FHR and presents results of the evaluation with the expected input variable values and uncertainty/sensitivity analysis results.

Chapter 6 has the same general structure as Chapter 5. It begins with a description of the reactor and a discussion of the reference design parameter variables. The setup of the independent optimizations is described and results presented. Then the setup of the neutronics simulations are discussed and results provided. As in Chapter 5, the application of the evaluation to this reactor class is discussed and presents the results presented with expected input variable values and with uncertainty/sensitivity analysis.

1.6 Literature Review

This section is comparatively short relative to the volume of background research necessary to evaluate the design parameters of MSRs. It has been deemed more appropriate to include research on specific design parameters in the relevant chapter rather the literature review section. Instead this section focuses on the current MSR design concepts. It covers the technical concept, history and current status of the three leading design classes: both thermal and fast spectrum circulating-fuel reactors, as well as the FHR. Section 1.6.7 '*Current Designs*' briefly states key design parameters and discusses recent progress of other MSR designs under development.

1.6.1 History

Investigation of MSR technology began in the late 1940's as part of an American effort to develop nuclear powered aircraft. The success of the Aircraft Reactor Experient (ARE) but failure of the overall Aircraft Nuclear Propulsion (ANP) program, see Section 1.6.4, lead to a shift in focus toward civilian applications of MSRs. The first major question to address was, at what neutron spectrum the design should operate. Initially, non-moderated, intermediate spectrum (epi-thermal) reactors were rejected because they did not appear to have high enough breeding ratios. The fast spectrum was also rejected since studies indicated that very high power densities and fissile inventories were required. Necessarily high power densities appeared

difficult to achieve without the use of 'novel and untested heat removal methods' [15].

Research into MSRs following the experiments at ORNL in the 1960s greatly diminished. The United States had shifted focus to sodium cooled breeders and progress in Europe was minimal. Japan directed appreciable efforts toward the FUJI reactor and the Soviet Union researched molten salts but without a strong, dedicated focus on a particular reactor design. The FUJI reactor began development in the 1980's and came to focus on a 200 MWe design. The proposed design drew heavily on MSRP research. The design was a single salt, (no blanket salt) graphite moderated breeder [16]. MSR related research had been carried out in the Soviet Union/Russia since the 1970s. Research was primarily directed toward: [17]

- Physical and chemical properties of MSR materials.
- Subcritical MSR systems for minor actinide incineration.
- Thermal MSRs with a Th-U cycle.

Until recently, fast-spectrum MSRs were not considered because it was found that the PuF_3 solubility in the fluoride salts under examination did not exceed ≈ 3 mole% [18]. Up until approximately the year 2000, the year in which Generation Four International Forum (GIF) meetings began, fast spectrum molten salt reactors have received little attention and therefore have an under-developed research base compared to thermal spectrum MSRs.

1.6.2 Current Focus

The GIF has selected MSRs as one of six reactor technologies for further research and development. Originally, the focus was on thermal-spectrum, graphite-moderated designs. Since the early 2000's R & D has expanded to include a focus on both fast-spectrum and molten salt cooled, TRISO particle fuelled designs [14]. Governments, university researchers and private start-ups have invested significantly in all three MSR variants.

Start-ups have demonstrated a tendency to focus on designs similar to the Molten Salt Reactor Experient (MSRE) at ORNL. A likely reason for this is that companies can reduce development costs by relying on ORNLs research and testing. The Integral Molten Salt Reactor (ISMR) by Terrestrial Energy has a well-developed conceptual design and has undertaken pre-licencing. Seaborg Technologies out of Denmark, ThorCon, FLiBe Energy and Transatomic Power also have thermal-spectrum designs in various stages of early development. Fluoride High-temperature Reactor (FHR) R & D has been conducted at MIT, U.C. Berkeley, U. of Wisconsin, U. of New Mexico and ORNL with US Department of Energy (DOE) funding. The development of molten salt, pebble-bed reactor technology originally centered on Advanced High Temperature Reactor (AHTR) designs; the AHTR, PB-AHTR and a small, modular version the SmAHTR. The focus then shifted to the similar Mark 1 PB-FHR design [19]. Upon completion of the Mark 1 design in 2014, the rate of publications declined but the research helped lead to the founding of Kairos Power and the development of their FHR [20].

The French-lead European Molten Salt Fast Reactor (MSFR) can be seen as the leading candidate fast reactor. The project is being developed at the LPSC-CNRS in Grenoble France and is supported by the EVOL and SAMOFAR EURATOMS projects.[1] The reference 3000 MWt design configuration can function with various fuel loadings, in a breeder configuration and have a multi-decade lifetime, see Section 1.6.5 Other fast-spectrum MSR concepts currently being explored include the fast version of the Stable Salt Reactor (SSR-W) by Moltex Energy based in the UK, the Molten Chloride Fast Reactor by Terrapower and the Minor Actinide Recycling in molten Salts (MARS) project carried out in Russia, see Section 1.6.7.

1.6.3 Graphite and Neutron Moderation

A defining design parameter of a nuclear reactor is the neutron energy spectrum. Neutrons born from fission have a distribution of initial kinetic energy that peaks at approximately 1 MeV (fast). Elastic collisions with materials in the reactor (e.g., structural, coolant, moderator) slow down neutrons. Neutrons will continue to slow down toward thermal equilibrium (≈ 0.075 eV at 700 ^oC) as long as they are not absorbed.

The absorption cross-section of nuclear fuel materials varies with neutron speed. A portion of the energy spectrum is characterized by resonance absorptions. These are spikes in the absorption cross-section that occur between the thermal and fast speeds. In this range the absorption cross-section of fertile material is comparatively high. To avoid 'parasitic', or non-fission absorptions in this range, reactors are typically designed so that most absorptions are of either fast or thermal neutrons and are thus more likely to cause fissions.

Thermal spectrum reactors require a moderator to slow down fast neutrons created during fission through a series of collisions. An effective moderator has a low absorption cross-section and low atomic mass. The common moderator materials are carbon, hydrogen (protium) and deuterium. Beryllium has been used as well [21] but is generally impractical. Hydrogen/deuterium has been proposed for use as a component of zirconium hydride but has only been used as a component of water and carbon has been used in the form of graphite.

All MSR designs operate at a high temperature, as they must in order for the salt to be in a liquid state. The most common baseline temperature selected for MSR designs and related studies is 700 ^{o}C. At this temperature level a reasonable balance of structural material stability and both thermal efficiency and high solidification margins can generally be achieved. Water near this temperature is in the supercritical regime and must be held at extremely high pressure to maintain a reasonably high density. This requirement effectively rules out water as a moderator.

The FHR and circulating-fuel MSR design concepts analysed are both graphite moderated. The sublimation point of graphite is in excess of 3000 ^{o}C [22] and therefore is in no realistic danger of changing state in an MSR. Additionally, graphite does not react chemically with molten $LiF - BeF_2$ salt (FLiBe) [23] and no

[1] The organization involved are discussed in Section 1.6.5.

evidence has be found to suggest that graphite reacts to an appreciable extent with other fluoride salts either. Despite the lack of corrosion, graphite is affected by neutron irradiation. It first shrinks then swells. At a reasonably high core power density, graphite moderated MSR designs can only operate for approximately 3 to 10 years before the graphite must be replaced. The MSBR core graphite had a design lifetime of only 4 years, with a core power density of 22.2 $\frac{kW}{l}$ [24]. The modern IMSR and MSTW designs both have 7-year core graphite lifetimes, the MSTW has an explicitly stated power density of 12 $\frac{kW}{l}$ and the IMSR is about the same.[16]

The first MSR, the Aircraft Reactor Experiment (ARE) employed a Beryllium Oxide (BeO) moderator. This moderator has the advantage of Beryllium's low atomic mass and hence high slowing down power, while remaining a solid at high temperature and low pressure. However, beryllium is toxic, expensive and must be separated from the salt (see Section 1.6.4). Zirconium hydride (ZrH_x) has some of the same basic properties; high slowing down power and is solid at high temperature and low pressure [25]. It is also non-toxic and less costly. Yet, ZrH_x has a relatively high absorption cross section and must also be separated from the salt. Upon considering the issues associated with alternative moderators it is understandable why thermal spectrum MSR designs have near-universally selected graphite as the moderator material. Subsequent analysis of both the FHR and MSR only considers a graphite moderator.

Graphite quality

Nuclear grade is artifical or synthetic graphite made from petroleum coke with a coal-based binder. The manufacturing process involves first heat treating, crushing and grinding the coke to produce a specific particle size distribution. The coke is then combined with the coal-based binder, placed under pressure finally heated to a temperature of about 2800 ^{o}C [26]. The final product is designed to produce isotropic graphite with uniform material properties. To qualify as nuclear grade, graphite must have a density in excess of 1.5 g/cm^3 and less than 5 ppm Boron [27].

Prompt Neutron Lifetime (PNL, l)

Prompt neutrons are those born directly from fission. The PNL is time between the emission of a fission neutron and its absorption - if the $k_{eff} = 1$ the PNL is equal to the mean generation time; the time between emission and capture that results in fission. In fast reactors, prompt neutrons have a much shorter lifetime because the neutrons do not scatter about a moderator before being absorbed. However, the fast and thermal spectrum are not strict confines, reactors can have any range of neutron spectra and in thermal reactors the PNL depends in large part on the moderator material and level of moderation. The PNL is an important parameter because it, along with delayed neutrons determines the reactor period[2]during normal operation, which in turn determines the sensitivity to changes in reactivity. If the reactor reaches prompt criticality, the period is inversely proportional to only the PNL. The PNLs of reactors are approximately:

[2]The reactor period (or e-folding time) is the time required for the neutron density to change by a factor of e. It is inversely proportional to reactivity and therefore is effectively a measure of how sensitive a reactor is to changes in reactivity

- CANDU: $1 \cdot 10^{-3}$ s [28]
- LWR: $1 \cdot 10^{-4}$ s [28]
- Fast reactor generally: $1 \cdot 10^{-7}$ s [28] [29]

1.6.4 Oak Ridge National Laboratory Molten Salt Reactor Program

Beginnings: ARE

The realization of nuclear technologies in the 1940's led to the development of atomic weapons and military vehicle propulsion. During this time, the first MSR was conceptualized, and design began at ORNL in 1946 to use MSRs to propel a plane. The Aircraft Reactor Experiment's (ARE) specific objective was to 'build and operate a high-temperature low-power circulating-fuel reactor of materials which would be suitable for a high-power reactor' [30]. Fuel was a mixture of the fluorides of sodium, zirconium and 93.5% enriched [31] uranium. The reactor achieved criticality November 3, 1954.

The core consisted of 66 fuel tubes surrounded by a Beryllium Oxide moderator. Sodium coolant was pumped through spaces between the moderator and fuel tubes as well as used to cool the reflector and container. The fuel and sodium coolant were circulated through separate fin-tubed heat exchangers. The core shell and other structural material was composed of Inconel (Hastelloy N), a Nickel based alloy [31]. The reactor was designed to operate at a maximum temperature of 1500 ^{o}F (1090 K) and reached a peak power in excess of 2.5 MW [31]. The fuel composition by molar percent was: $53.09NaF - 40.73ZrF_4 - 6.18UF_4$, with a melting point of 1000 ^{o}F (810 K). The liquid fuel resulted in a strongly negative overall reactor temperature reactivity feedback coefficient of $-6 \cdot 10^{-5} \frac{\Delta k/k}{^oF}$ [30].

The ARE was a success, operation was very stable due to the high negative reactivity coefficient and it demonstrated load-following capability without control rods [32]. Yet it faced technical difficulties, high costs and competition from the advent of alternative forms of nuclear weapon delivery, particularly ICBMs. The Aircraft Nuclear Propulsion program (ANP) was cancelled in early 1961 by the new administration of President Kennedy. The total cost of the program was about $1 billion [33].

Molten Salt Reactor Experiment (MSRE)

The development of the ANP program and success of the Aircraft Reactor Experiment led to interest in using molten salts for civilian purposes [15]. The civilian molten salt power reactor program was initiated in 1958 [32]. The original design leaned heavily on ANP and was developed alongside Aqueous and Liquid Metal Fuel reactors. Subsequently, the thermal-spectrum molten salt reactor was selected as the primary fluid-fueled reactor concept and the MSRE design began in 1960 [32].

From the outset, potential breeding capabilities of a MSR were recognized [34]. The MSRE was intended not to simulate a final design for a commercial power plant but to conduct research for a future breeder reactor. Many technical and economic factors of the time, including concerns about the supply of uranium,

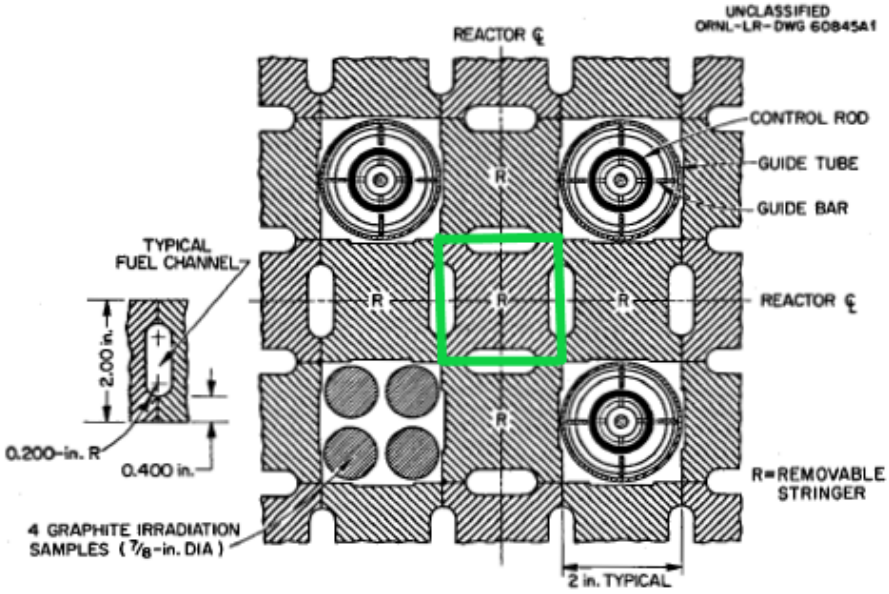

Figure 1.1: Fuel channels at a central location in the core. The green box indicates the outline of a single graphite 'cell.' Three control rods and the irradiation channel are shown [23].

combined to favour thermal spectrum breeders. The central problem of the thermal breeder was the need to prove the suitability of the in-core graphite [23].

The MSRE was designed for operation at 10 MWt. The reactor consisted of a 5 ft diameter, 8 ft high tank, containing a 55 in diameter by 64 in high graphite core structure [23]. Although fluorine, lithium and beryllium in the salt do have considerable moderating capability it is insufficient for a thermal-spectrum reactor, and so a moderator, graphite is required. The reactor was fuelled at first with 93.5% enriched ^{235}U [23], followed by ^{233}U only toward the end of operation [24]. The circulating fuel salt transferred heat to a secondary non-radioactive salt which dumped heat to a salt-to-air heat exchanger. The reactor first went critical in June 1965 and was shut-down in December 1969 [32].

The core was made up of 513 2.0x2.0 'cells.' Figure 1.1 shows several of these cells, the green box indicates the outline of one cell and the shaded area the graphite. The pill-shaped channels between the cells is where the fuel salt would flow in direct contact with the graphite. The fuel volume fraction of the core was 0.225 [23].

A thermal-spectrum MSR would function most effectively using a fuel salt composed of light, low-absorbing isotopes. Generally, non-metal polyatomic anions can be eliminated due to low stability (thermal, chemical or irradiation) [23] and elements with oxidation states greater than one can be eliminated because of corrosion issues and/or reactivity with air/water. That leaves chloride and fluoride salts, fluorine being the lighter, lower absorbing element, was selected. As for the cations, ^{7}Li and Be were selected in large part due to

Table 1.1: Physical properties of the MSRE fuel salt

Physical Property	Value
Liquidus temperature	450 ^{o}C
Density	2.34 g/cm^3
Heat capacity	1.9 kJ/kgK (water = 4.2)
Viscosity	7.5x water at 20 ^{o}C
Vapour pressure	Negligible

Table 1.2: Temperatures and processing in the MSBR

Salt	Inlet (^{o}C)	Outlet (^{o}C)	Processing Method	Processing Cycle Time (days)
Fuel	635	663	Fluorination and Distillation	47
Blanket	621	677	Fluorination	23
Secondary/Coolant	454	600	N/A	N/A

their low atomic mass and low-absorption to create a mix that results in an effective carrier salt. The fuel salt selected was therefore FLiBe or ^{7}LiF–BeF_2 with some ZrF_4 and UF_4 fuel. The mole percent of each constituent was as follows: 65 LiF, 29.1 BeF_2, 5 ZrF_4, 0.9 UF_4. The lithium was enriched to 99.99% ^{7}Li to prevent absorptions by ^{6}Li. FLiBe has favourable thermophysical properties including a relatively low liquidus temperature and low visocity. The physical properties of the fuel salt at the approximate operating temperature of 650 C are shown below in Table 1.1 [23]. The composition of the secondary, or coolant salt, was similar, by mole percentage: 66 LiF, 34 BeF_2 [23].

Molten Salt Breeder Reactor (MSBR) and program shutdown

The original design the Molten Salt Breeder Reactor (MSBR) was a 2220 MWt, 1000 MWe, ^{232}Th – ^{233}U cycle Breeder [35]. The reactor core would consist of 534 hexagonal, graphite fuel cells with a 4.8 in pitch resulting in a diameter of 10 ft and a height of 12.5 ft. The core composition would be 75% graphite, 18% fuel salt, 7% blanket salt. The core was to be surrounded by a 1.5 to 2.0 ft thick blanket composed of 27 mole % thorium molten salt. The reactor was to have an average core power density of 80 kw/l (80 MW/m^3) resulting in an estimated core graphite lifetime of 5 years [36]. Though it was realized in a later design iterations that the graphite could not withstand such a high fluence. As a result, the single-fluid design as of 1969 had an average power density of 22.2 kW/l and graphite lifetime of 4 years [24].

Table 1.2 provides a summary of the temperature of each salt and the processing which it undergoes [35]. Table 1.2 does not include Protactinium processing due to it not being confirmed at the time. Protactinium is produced by the irradiation of ^{232}Th and decays to fissile ^{233}U. The breeding ratio with Pa removal was found to be 1.071 and 1.049 without [36].

Due to a lack of government financial support the original 1000 MWe reactor was scaled down first to 300 MW and then to 100 MW. The MSBR was in competition with the Liquid Metal Fast Breeder Reactor

(LMFBR) and effectively lost out. In the end the whole project was terminated due to 'budgetary reasons.' It was reasoned that being concentrated in Oak Ridge, political support and technical expertise was spread too thin geographically. The director of the project Alvin Weinberg remarked, 'Our problem is not that our idea is a poor one – rather it is different from the main line...' [32].

Following the termination of MSBR development efforts, a study was conducted of alternative MSRs in 1976. With an emphasis on proliferation resistance, studies were then undertaken to develop a denatured Molten Salt Reactor (DMSR). Attention concentrated most on a conceptual design with a 30-year fuel cycle and no special chemical processing or fission product removal. The flux was sufficiently reduced to allow for a 30 year lifetime of the graphite. The design called for 1000 MWe and 2250 MWt and a core volume of 59.4 m^3. The salt composition would be similar to the MSBR with an intial fuelling of 20% enriched uranium and reach a maximum temperature of 704 ^oC. Gas sparging would still be employed to remove xenon [37].

Molten Salt Reactor Program key results and technological developments

The ORNL reports appeared confident that graphite was the clear choice for use as moderating material. Graphite does not react chemically with the fluoride fuel salt. The graphite lifetime is a major concern limited mostly by the total irradiation. High-grade graphite used in the MSRE is composed 10% by volume of voids. In part due to the high surface tension of the fuel salt, void impregnation by the fuel salt was not found to be a concern. Fast-neutron irradiation causes graphite to first shrink then expand. This phenomenon is the reason for the limited lifetime. Carbon is an effective moderator with a moderating ratio approximately equal to Be and 2.7x higher than light water, only approximately 1% of neutrons are absorbed in the graphite moderator. The fuel cells (Figure 1.1) are made small to reduce the fast-flux gradient across the graphite wall and therefore minimize the effect of radiation damage [38].

The primary structural material used in the MSRE was INOR-8 or now known as Hastelloy N. It is a nickel-based alloy with a chemical composition shown in Table 1.3. Fluoride salts must be at high temperatures and form no protective oxide layer with metals. Therefore, corrosion is governed by thermodynamic stability alone. The oxidation and diffusion into the salt of the Chromium alloying element was found to be the predominant corrosion mechanism. If the salt has an insufficient amount of cations and a high redox potential, then the container metal must supply the deficiency. Hastelloy N is easy to form, machine and weld to the required specifications, however, it difficult to connect to graphite. Graphite cannot be welded and has a much lower thermal expansion coefficient. The materials had to be brazed together with an intermediate [23]. The primary effect of irradiation is embrittlement by the production of helium. This embrittlement can be reduced by alloying with Titanium [24].

Considerable research was devoted to processing the salts. This includes but is not limited to the removal or recovery of oxygen, fission products, Uranium and Protactinium as well as the management of the redox potential. On average fission fragments have a combined valence of about +3.4 and the parent uranium

Table 1.3: Required chemical composition of Hastelloy N (INOR-8) [23].

Element	Percent[a]
Nickel	Remainder
Molybdenum	15.0-18.0
Chromium	6.0-8.0
Iron	5.00
Carbon	0.04-0.08
Manganese	1.0
Silicon	1.0
Tungsten	0.5
Aluminium + titanium	0.5
Copper	0.35
Cobalt	0.20
Phosphorus	0.015
Sulphur	0.02
Boron	0.01
Others, total	0.5

[a]Single values are maximum percentages unless otherwise specified.

cation is +4 (UF_4) [23]. This means that the redox potential must continuously be reduced through active external means. Gaseous fission products, notably Krypton and Xenon were bubbled out and collected. Other salt processing methods that were tested include:

- Fluorination is the means of adding excess fluorine to the salt mixture, thereby creating volatile compounds that are bubbled out. Uranium hexafluoride (UF_6) is more volatile than tetrafluoride, the main uranium-containing compound in the salt mixture. Upon fluorination the UF_6 would bubble out, be collected and put (back) into the fuel salt. The processing was 90-96% effective for UF_6. The same process was used to remove certain noble fission products as well. The most important of which are iodine, tellurium, riobium, ruthenium and antimony. The main disadvantage of fluorination is that the salt contains free fluorine ions and is therefore extremely corrosive. To protect the container, the process was carried out only 50 – 100 ^{o}F (28 – 56 ^{o}C) above the melting point and the container was lined with frozen salt [23].
- Vacuum Distillation, the removal of fission products through the vaporization of the carrier salt, 2LiF–BeF_2, at approximately 1000 ^{o}C. It was used remove rare earths, barium, strontium and yttrium [38].
- It was not an absolutely necessity that protactinium (Pa) be continuously removed in the MSBR, however the prospect did receive considerable attention. A method was developed during operation of the MSRE, liquid-metal bismuth extraction. However it was developed too late to be tested on the MSRE. The process involves placing fuel salt in contact with bismuth. Pa is extracted according to the following reaction:

$$PaF_4(salt) + 4Li(Bi) \rightarrow 4LiF(salt) + Pa(Bi)$$

Uranium, Thorium and Zirconium are partially removed as well and the mix is isolated until Pa decays to ^{233}U. Further processes then occur to return Th, Zr and U back to the fuel salt, with the net result being the removal of Pa [39].

The use of circulating fuel has a strong influence on the effect of delayed neutrons. In the MSRE where the fuel salt is in the core 1/3 of the time, the total effective delayed neutron fraction is about half that of a static fuel reactor. The Doppler and density reactivity feedback effects combine to generate a strong negative temperature coefficient of 6.4 to 9.9 pcm/F (11.5 to 17.8 pcm/C), depending on the fuel used [23]. During early design it was hoped that the negative temperature coefficient would sufficiently supply all control when power is above 1 MW. Analog simulation showed that this could not be possible and external control (i.e. control rods) would be required. The reactor was found to have a 'very sluggish' temperature response between 12 and 20 minutes due to the graphite. Ultimately components generally performed well including salt pumps that had run 'flawlessly' [35]. The Xenon poison fraction was found to be approximately 0.005 [38]. Dump tanks with freeze plugs that take 15 to 20 minutes to thaw were the safety barrier against overheating. It is safe to conclude that the MSRE was a success that proved the technical feasibility of thermal-spectrum FLiBe salt burner reactors and the many aspects of a breeding design.

1.6.5 Molten Salt Fast Reactor (MSFR)

The MSFR is a well developed MSR concept. Since 2004 the LPSC-CNRS in Grenoble France has focused R & D efforts on this design [40]. This endeavour is supported by the EVOL and SAMOFAR Euratom projects [41]. The research is part of the GIF collaboration in Molten Salt Reactor Technology. Given the complex structure of project support and the numerous parties involved, a list of organizations mentioned is provided below:

- LPSC-CNRS: Laboratory of Subatomic Physics and Cosmology (LPSC) is a division of the National Centre for Scientific Research (CNRS).
- EVOL Evaluation and Viability of Liquid Fuel Fast Reactor System. The objective of this project is to propose a pre-conceptual design of MSFR. It is coupled with the MARS (Minor Actinides Recycling in Molten Salt) project of Rosatom - the Russian state nuclear energy corporation [41].
- SAMOFAR Safety Assessment of the Molten Salt Fast Reactor, is a project in the Horizons 2020 research program, with the goal being to demonstrate key safety features of the MSFR.
- Euratom European Atomic Energy Community, which aims to pursue nuclear research and training activities with an emphasis on continually improving nuclear safety, security and radiation protection, notably to contribute to the long-term decarbonisation of the energy system in a safe, efficient and secure way.
- GIF Generation IV International Forum, see Section 1.4 for more information.

The current design is that of fuel salt circulating between the core and heat exchangers surrounded by a fertile blanket salt. The reactor operates on the thorium cycle but with an initial fuel salt composition consisting of thorium, enriched uranium, plutonium and minor actinides. The fuel and fertile blanket salts use

Table 1.4: Design parameters of the MSFR reference configuration [41].

Parameter	Value
Power	3000 MWth, 1300 MWe
Fuel salt melting point	585 C
Mean fuel salt temperature	725 C
Fuel salt density	$4.1 g/cm^3$
Breeding ratio	1.1
Core radius	1.06 to 1.41 m
Core height	1.60 to 2.26 m
Fuel salt cycle time	3.9 s
Total fuel salt volume	18 m^3

LiF carrier salt. Key reference design parameters are given Table 1.4.

The research being performed on the MSFR is extensive. Key design parameters and core geometry have been studied and optimized [42]. Neutronics and thermohydraucis have been coupled using the transient fission matrix for the purpose of transient modelling of velocity, power and temperature distribution. This method is used to model safety scenarios such as overcooling and reactivity insertion as well as load following [43]. A starting procedure has been modelled, including the analysis of abnormal transients [44]. Further, safety studies have been performed including studies that involve the draining tanks [45]. Most published research and development appears to be theoretical; no test reactor has been built and little tangible experimentation work has been conducted compared to the volume of theoretical research.

1.6.6 Fluoride-salt-cooled High-temperature Reactor (FHR)

Fluoride-salt cooled High-temperature Reactor (FHR) is a term used to describe the nuclear research concerning molten salt cooled pebble-bed reactors. Work under the GIF roadmap lead to recommendations in 2002 that molten salts be used as coolants and research had come to include pebble-beds as a means of fuelling [46]. The first pre-conceptual design to emerge from this program came in 2008 and in 2012 the DOE initiated a new Integrated Research Project (IRP) with the Massachusetts Institute of Technology, University of California Berkeley, and University of Wisconsin at Madison, to further develop the technical basis to design, develop, and license commercially attractive FHRs [47]. Three more molten salt cooled, pebble-bed designs followed:

- PB-AHTR, 2008: 900 MWt
- SmAHTR 2010: 125 MWt
- ORNL-AHTR 2012: 3600 MWt
- Mark 1 PB-FHR 2014: 236 MWt

As of the time of writing the Mk1 PB-FHR (Mark 1 Pebble-Bed FHR) is the design project to have generated the most publicly-available research. Conveniently, Mk1 PB-FHR conceptual design and technology is well

Table 1.5: Mk1 PB-FHR key design parameters [48].

Parameter	Value
Power	236 MWth, 100 MWe
Moderator	Graphite
Fuel	19.9% ^{235}U
Coolant Salt	$2^7 LiF - BeF_2$
Core average outlet temperature	700 ^{o}C
Main reactivity control	Negative temperature coefficient
Other control	8 buoyant control rods
Reactor vessel diameter	3.5 m
Pebbles used	TRISO coated particle fuel
Pebble diameter	3.0 cm
Uranium/Pebble	1.5 g

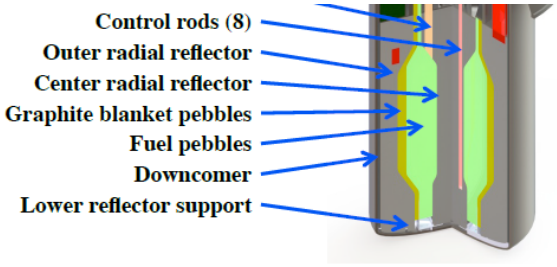

Figure 1.2: Mk1 PB-FHR reactor vessel geometry [48].

documented in the Mark-1 PB-FHR Technical Description (2014) [48]. The Mk1 PB-FHR design uses a nuclear air-Brayton combined cycle (NACC). This alternative design, with its co-firing capability, is to provide a new value proposition for nuclear power. However the core function is of primary interest to this thesis and its key parameters are listed in Table 1.5 and a diagram of the reactor vessel is shown in Figure 1.2.

1.6.7 Current Designs

This section covers current MSR design concepts under development that have not thus far been included in the literature review. Past designs that have achieved significant development are discussed as well. The intent of this section is to provide a broad, comprehensive picture of MSR development. A summary is provided in Tables 1.6 to 1.10.

[2]Moderator lifetime unclear and moderator may be continuously replaceable

Table 1.6: Summary of design parameters of current circulating-fuel, thermal spectrum MSR design concepts (part 1)

Parameter	IMSR	MSTW	ThorCon
Developer	Terrestrial Energy	Seaborg Technologies	Martingale
Fuel	5% ^{235}U	93 Th, 3.5 Pu, 3.5 U	80 Th, 20 U
	Unspecified	(1.1% ^{235}U)	(19.7% ^{235}U)
Power, MWe	190	100 or 115	250
Power, MWt	400	270	557
Moderator	Graphite	Graphite	Graphite
Moderator/core lifetime (yrs)	7	7	4
Reprocessing non-volatile FP	No	No	No
Carrier salt	FLiBe	Na-actinide fluoride	NaF–BeF_2
Primary reactivity control	FTC[a]	FTC[a]	FTC[a]
Peak salt temperature (^{o}C)	600 (secondary)	700 (core)	704 (core)

[a]Fuel Temperature Coefficient

Table 1.7: Summary of design parameters of current circulating-fuel, thermal spectrum MSR design concepts (part 2)

Parameter	TransAtomic	Fuji	TMSR-LF2
Developer	TransAtomic Power	ITMSF	CAS, China
Fuel	5% ^{235}U	98 Th, 2 ^{233}U	Th, U (19.75% ^{235}U)
Power, MWe	520	1500	168
Power, MWt	1250	3000	373
Moderator	Zirconium Hydride	Graphite	Graphite
Moderator/core lifetime (yrs)	N/A[a]	7	6-8
Reprocessing non-volatile FP	No	Yes	Yes
Carrier salt	LiF	FLiBe	FLiBe
Primary reactivity control	FTC[a]	FTC[a]	FTC[a]
Peak salt temperature (^{o}C)	650	704	700

[a]Moderator lifetime unclear and moderator may be continuously replaceable

Table 1.8: Summary of design parameters of current circulating-fuel, fast spectrum MSR design concepts

Parameter	MCSFR	MCFR	MOSART
Developer	Elysium Industries	TerraPower	MARS by ROSATOM
Fuel	many incl LEU & SNF	'many'	Th, TRU
Power, MWe	1000	—[a]	—
Power, MWt	—	—	2400
Reprocessing non-volatile FP	Yes	—	Yes
Carrier salt	Chloride	Chloride	15LiF-58NaF–27BeF_2
Primary reactivity control	FTC	FTC	FTC
Peak salt temperature (^{o}C)	>600	—	750

[a]Unspecified

Table 1.9: Summary of design parameters of non-circulating fuel current MSR design concepts

Parameter	KP-FHR	SSR-W	TMSR-SF2
Developer	Kairos	Moltex Energy	CAS, China
Fuel	TRISO (19.75% ^{235}U)	U, Pu, other TRU trichlorides	TRISO (19.75% ^{233}U)
Power, MWe	140	300/unit	168
Power, MWt	310	375	384
Moderator	Graphite	None	Graphite
Reprocessing non-volatile FP	No	No	Yes
Carrier salt	N/A	NaCl	N/A
Coolant salt	$2^7LiF - BeF_2$	Na, K, Zr, fluoride	FLiBe
Primary reactivity control	FTC	FTC	FTC
Peak salt temperature (^{o}C)	650	650 (coolant)	700

Table 1.10: Summary of design parameters of other current MSR design concepts

Parameter	DFR	IMSBR	LFTR
Developer	IPK	India	FLiBe Energy
Fuel	Actinide chlorides	Th/^{233}U cycle	Th/^{233}U cycle
Power, MWe	1000	250	
Power, MWt	—	—	
Moderator	None	None	Graphite
Moderator/core lifetime (years)	—	—	Undetermined
Reprocessing non-volatile FP	Yes	—	
Carrier salt	chloride	Chloride	FLiBe
Primary reactivity control	FTC	FTC	FTC
Peak salt temperature (^{o}C)	>600	—	750

Chapter 2

Evaluation and Cost Estimation Framework for Molten Salt Power Reactor Systems

This chapter discusses an evaluation program and cost estimation framework for the purposes of understanding the costs involved in MSR design configurations. The cost estimation framework is an approach to quantify the magnitude of a wide range of costs of differing configurations. The evaluation makes use of the framework, considering reactor-specific design characteristics, to arrive at a total expected cost over the range of geometric configurations and salt compositions under study.

2.0.1 Cost Estimating Guidelines for Generation IV Nuclear Energy Systems

The Cost Estimation Framework was initially developed without knowledge of the Cost Estimating Guidelines for Generation IV Nuclear Energy Systems [7] (Gen. IV Cost Estimation Guidelines, or just Guidelines). It is likely the study that is most closely aligned with the objectives of the Cost Estimation Framework. The purpose of the Gen. IV Cost Estimation Guidelines is to achieve, for Gen. IV nuclear energy systems, a life cycle (LCOE) cost advantage and a level of financial risk comparable to other energy projects. It does this through the development of standardized cost estimating protocols to assess, compare and eventually select future nuclear energy systems.

The key difference is that the Cost Estimation Framework is focused on cost estimation for the purposes of the design optimization rather than the selection among nuclear energy systems. Much of the Gen. IV Cost Estimation Guidelines content is focused on listing all the expected costs at comparably high level of detail. Many of list of costs are included at different levels of detail; cost categories such as 'Main heat transport system' or 'Safety systems' and those are further divided into sub-categories such as initial reactor coolant load, steam generators (if applicable), fluid drive circulation system, heat exchangers, etc. The Guidelines attempt to cover all costs, including direct costs like turbine and electrical equipment, indirect costs such as management and training as well as supplementary costs like shipping and insurance. Thesis is also considerable coverage of several issues, including site-specific considerations and comparing FOAK and NOAK costs.

The Guidelines discuss two approaches to calculating costs: bottom-up and top-down. The bottom up is a more conventional approach that requires an adequately detailed design to account for all construction commodities, plant equipment and labour hours. There is insufficient data for MSRs and this method is beyond scope. The top-down approach often calculates costs with standardized factors or formulas. The costs of bulk commodities and craft job hours are expressed as a fraction of process equipment cost. However even this cost estimation method likely requires a conceptual design with more detail than the reference MSR models include.

The majority of the individual costs listed, termed Code of Accounts (COA), and other issues discussed do not meaningfully contribute to the objective of determining costs as a function of configuration. There are however some numerical values that are useful in verifying in Cost Estimation Framework such as:

- The chosen discount rate of the Guidelines' authors is 5% real rate for plants operating under the more traditional 'regulated utility' model and 10% for a riskier 'deregulated' environment, where the plant must compete with other generation sources for revenues. The 10% discount rate is dismissed because, for the objective of configuration comparison, it would unreasonably favor configurations with a comparably high capital cost. In the Cost Estimation Framework the discount rate is effectively applied to both the costs and the revenues.
- Table 8.1 of the Guidelines lists many expected fuel cycles costs that can be directly compared to the costs used in the Evaluation. e.g. \$38.5/kgU_3O_8 (2021 USD: \$50/kgU_3O_8) compared to effectively \$63/kgU in the Evaluation. SWU is notable is that the value from the Guidelines, \$115/SWU (2021 USD: \$148) is considerably higher than the \$45 in the Evaluation. This is because the price has fallen considerably.

In summary, the Gen. IV Cost Estimation Guidelines would have been a useful document to take inspiration from and cite a limited number of specific values. However, it's focus and general structure differ considerably from the Cost Estimation Framework and Evaluation in this study.

2.1 Evaluation

There are two versions; one for evaluating the FHR and one for the Thermal-Spectrum, circulating-fuel (TC) design. The designs are evaluated over a range of geometric configurations and salt compositions. The evaluation is intended to appropriately weigh and quantify the effect of individual design characteristics, in order to arrive at a cost value for each design configuration under study. The fundamental objective is to optimize the reactor design class in question taking into account as many factors as reasonably possible. It is a useful tool for evaluating salts across many attributes relevant to reactor construction/operation.

The evaluation aims to assess the FHR design concept as a function of four design parameters listed below and discussed in Section 5.1.3: [itemsep = 0pt]

- Coolant salt composition.
- Particles per pebble (Packing Fraction, PF).
- Share of graphite only pebbles (blanks).
- Inner reflector thickness.

It aims to assess the TC design concept as a function of three design parameters listed below and discussed in Section 6.1.1:

- The non-uranium composition of the salt.
- Fraction of core volume taken up by the salt channels (Fuel Fraction or FF).
- Molar fraction of salt cations that are uranium (Mole % U).

For each design parameter in question, the reactor is evaluated over a range of realistic values. The evaluations are pieced together in order to gain insight into the impact various design choices have on reactor performance and ultimately to determine the most appropriate, lowest cost configuration.

A cost estimation requires assigning values to many inputs that do not have stated market prices, are a function of technological difficulty, are associated with regulatory challenges, or for some other reason can not be estimated accurately and/or precisely. The results obtained when using this framework are highly dependent on these assumed inputs, termed *input parameters* going forward. Many of these input parameters are selected intuitively and require at least some degree of user discretion. The objective of this study is to provide both the template and the implementation of a cost-optimization for MSRs. Doing so requires cost estimates of numerous input parameters, estimates that could not always be justified with supporting evidence. The authors do not claim that the estimates are all necessarily justified, however the optimization could not be conducted without them. Estimates are justified where possible and best efforts are taken to make reasonable postulations where they cannot be.

As a results of the considerable epistemic and stochastic uncertainty of the input parameters, uncertainty analysis is performed. As well, the effect of each input parameter on the total cost cannot be determined easily and/or precisely so sensitively analysis is conducted. A stochastic approach is taken for both the uncertainty and sensitivity analysis; parameters are assigned a probability distribution and the evaluation is run numerous times at pseudo-random points across each distribution (see Section 4.2).

The evaluation does not lend to design optimization (cost function minimization) through either a conventional brute force method or an exact approach using differential calculus. Rather each optimization is specific to the reactor type under analysis and is largely a high-level summation of a set of distinct but interrelated optimizations. The process to arrive at the various cost components varies considerably. The *Cost Estimation Framework* discusses how most of the costs are determined, reactor-specific cost calculations are discussed in their respective chapters.

The program iterates through the design space of the reactor type. For each configuration evaluated, it adds up all costs associated with each cost category and calculates the summation of all cost categories. Once that is complete, the total cost of each configuration is compared. The input data does not need to match the evaluated points in the possible design space. What is meant by input data are physical properties or data points extracted directly from neutronics simulations and thermohydraulics calculations. The evaluated data points are those selected by the user of the evaluation program for analysis. Thus interpolations must be made throughout to calculate costs as accurately as possible relative to the input data. Care must be taken to ensure that any evaluated data points that are not sufficiently close to the input data points are rejected.

The range of cost components included in the evaluation is extensive. The evaluation framework is meant to be able to include any type of cost that is a function of core design. An effective way to gain an understanding of the costs that can be included is to see the list of costs in FHR and MSR implementations, Table 5.12 and Table 6.14.

The cost outlay per value produced (electricity) can be expressed in different ways. The chosen scheme in this analysis is to move all costs and value produced to the moment of reactor start-up; all costs prior to that time have an interest cost and all costs following are discounted. To compare the cost of all scenarios equitably, the total value produced must be equal. Hence the value produced (and total cost) is adjusted relative to the average interest rate, lifetime and capacity factor. The scheme could instead have estimated the capital depreciation and O & M costs per year (adjusted for value produced) or capital depreciation and O & M costs per kWh. So long as implemented correctly, the relative cost differences between all the evaluated scenarios should be the same regardless of costing scheme chosen.

Several variables raise/lower the apparent cost but increase/reduce the value of or amount of produced energy. Without an adjustment, the sensitivity of the relevant input parameters has a sign opposite to what is should be (i.e. lower lifetime results in lower total cost). The input parameters adjusted are interest rate, lifetime and capacity factor.

- Interest rate: The term can be used interchangeably with *discount rate*.[1] The total cost is the cost at the time of reactor start-up. A higher interest rate has three effects. For one, it increases capital costs because capital costs are incurred prior to start-up and two, it decreases O & M costs since the time value of future cash flows decreases. The final effect is on the value of produced energy rather than costs; the value of future cash flows from energy sales decreases. The capital and O & M costs are accounted over throughout the evaluation while the value of future cash flows are adjusted at the end, in conjunction with lifetime, according to Equation 2.1.
- Lifetime: The term can be used interchangeably with *depreciation period*.[2] An increase in the reactor

[1] Interest rate is better suited for describing the cost of capital while discount rate is better suited for describing the amount the future cash flows are discounted. This analysis includes both capital and O & M costs; one term, *interest* rate was selected.

[2] Economic analysis of assets that generate cash typically includes a specification of the depreciation period - the useful life of the asset. Though depreciation period may be a more precise term, lifetime is simply a more convenient term.

lifetime has the effect of increasing the present value of all O & M costs. A longer lifetime makes for a higher return on investment and thus should result in a lower cost. Like with interest rates the difference in cost as a function of lifetime is taken into account throughout, while the value of produced energy is accounted for at the end of the evaluation, as per Equation 2.1.

- Capacity factor (CF): The CF does not effect the total amount of energy produced only the timing of energy production. A higher CF increases the present value (at start-up) of O & M costs and produced energy by an amount that is a function of the interest rate. There is therefore no net effect to the O & M costs[3] and capital costs are adjusted by a factor of $(1+i)^{((CF_\mu - CF)/2)}$ to account for the difference in the value of produced energy between the average CF, CF_μ and the CF of the individual iteration, CF. Thus if $CF > CF_\mu$, the adjustment due to ΔCF decreases total cost.

$$C_a = C \cdot \frac{(i_\mu^l - 1_\mu)/(i_\mu \cdot (1+i_\mu)_\mu^l))}{(i^l - 1)/(i \cdot (1+i)^l)} \tag{2.1}$$

Where C_a is the adjusted cost, i_μ and i the average and individual iteration interest rates, and l_μ and l the average and individual iteration lifetimes.

The comprehensive evaluation of a reactor design class, such as the FHR or TC (MSR) is a useful aid in understanding the contributing factors to total cost and optimizing the design. The fundamental objective is to minimize the cost of this reactor design taking into account as many design parameters as reasonably possible. In doing so the relative viability of the salts examined and the optimal geometric configuration of the reactor with each of the salts can be determined. It is a useful tool for evaluating salts across many attributes relevant to reactor construction/operation. Not only are the sum total outputs useful but the program is built in such a way that it is made clear how salt attributes as well as geometric parameters contribute to the different cost types.

2.1.1 Corrosion in the Evaluation Program

The level of corrosion inflicted by the fuel salts is primarily dictated by the Gibbs energy difference between the salt components and the metal container. The relative cost of corrosion for evaluation purposes is a function of the Gibbs energy difference and experimental results. It is impractically difficult to explicitly quantify the effects of impurities or irradiation, so the evaluation in this study treats those effects as equal for all salt compositions. A background on corrosion mechanics and corrosion findings is discussed in Section D.2.

The method used to quantify the relative magnitude of corrosion is a combination of experimental results and theory. The associated cost is proportional to the expected magnitude of corrosion. The experimental results are generally old, in short supply and conducted under inconsistent conditions (duration, temperature) such that the experimental results as a whole are insufficiently reliable on their own. And the theory on molten salt corrosion with metallic alloys has not been developed fully and precisely enough such that

[3] Expect to insurance, which is not a function of CF.

explicit estimates of can be made from first principals. Without either sufficient empirical data or theoretical models, corrosion is quantified as a combination of the two. The primary objective is to compare relative corrosion (mitigation) costs between salt compositions rather than accurately estimate the cost. The cost of corrosion is likely the least reliable aspect of the cost evaluation but given the importance of corrosion to MSRs it was included. The study is not proposing any corrosion cost estimation method but rather is explaining the method used to develop relative corrosion costs required for the evaluation.

The straightforward approach taken to quantify the cost requires the program user to set a total cost of $LiF-BeF_2$[4] corrosion and split that cost between direct cost and feasibility; done so through the evaluation input variables *Corrosion cost* and *Corrosion split*. Both the theoretical and experimental component of the cost of corrosion for other salts are equal to the cost of $LiF-BeF_2$ multiplied by a factor relative to $LiF-BeF_2$. The theoretical factor is calculated as follows, with $NaF-ZrF_4$ (59.5-40.5), salt 3 (s3) in the FHR analysis, as an example:

1. Calculate the weighted-by-constituent average Gibbs energy of each salt (G_{s3}) and the average Gibbs energy of relevant metal constituents (G_m). Though the structural material, Ni-based Hastalloy N, is composed of by weight percent Mo 15-18, Cr 6-8 and Fe 5 (max.) [23], only Ni and Cr are included in the average because of the high Ni concentration and high negative G_{CrF_2}.

$$G_{s3} = 0.595 \cdot G_{NaF} + 0.405 \cdot G_{ZrF_4} \tag{2.2}$$

$$G_m = \frac{G_{CrF_2} + G_{NiF_2}}{2} \tag{2.3}$$

2. Calculate the minimum Gibbs energy difference between a salt constituent fluoride and metal constituent fluoride (i.e. the difference between the lowest energy salt constituent fluoride and the highest energy metal fluoride, which is CrF_2.)

$$\Delta G_{min} = G_{ZrF_4} - G_{CrF_2} \tag{2.4}$$

3. Divide the weighted-average Gibbs energy difference of each salt by the average of all weighted-average differences ($R_{s3,w}$). Do the same for each minimum difference and it's average ($R_{s3,\mu}$).

$$R_{s3,w} = \frac{G_{s3} - G_m}{\sum_{i=1}^{5}(G_{si} - G_m)/5} \tag{2.5}$$

$$R_{s3,\mu} = \frac{\Delta G_{min}}{\sum_{i=1}^{5}(G_{xF_y} - G_{CrF_2})/5} \tag{2.6}$$

4. Take the reciprocal of the sum of the weighted-average and minimum difference to obtain the relative

[4]Doing so is inherently arbitrary; consideration for corrosion by impurities, radiation and interaction with graphite must be made here.

Table 2.1: Gibbs free energies of fluorides relevant to corrosion [50], data is provided per fluoride ion, at 700^{o}C and in units of kJ/mol.

Fluoride	ΔG	Fluoride	ΔG
LiF	-1046	NiF_2	-253
NaF	-948	CrF_2	-324
KF	-939		
BeF_2	-439		
ZrF_4	-398		
UF_4	-409		

cost:

$$C_{T,s3} = \frac{1}{(0.5 \cdot (R_{S3,w} + R_{s3,\mu})} \tag{2.7}$$

5. Determine the theoretical factor relative to FLiBe. This factor is indicative of the expensive of corrosion mitigation and consequences compared to the expense for FLiBe salt.

$$F_T = C_{T,s3}/C_{T,s1} \tag{2.8}$$

F_T is calculated based on linear functions of Gibbs free energy differences rather than the exponential functions. This is because exponential functions would only hold in an unrealistic, ideal environment that does not include the management of fluorine potential, the effect of oxidizing impurities or irradiation. It can be seen from experimental results, such as those listed in Table D.2 that a relationship between corrosion and Gibbs free energy exists. However the exact relation cannot be reliably determined and so the most simple function, a linear function, is assumed to allow for significant differences and yet prevent extreme differences in relative corrosion costs. The experimental data listed in Table D.2 is used to come up with the expected corrosion of the salts. The experimental factor (F_E) is simply the expected $[Cr]_{s3}/[Cr]_{s1}$ at 700^{o}C using the experimental setup from the respective data source ([49]). The theoretical and experimental components of the cost of corrosion are weighed equally[5], the relative cost of $NaF - ZrF_4$ corrosion is thus:

$$C_{s3} = C_{s1} \cdot (F_T + F_E)/2 \tag{2.9}$$

The Gibbs free energies of the fluorides involved are listed in Table 2.1. Different salt compositions and reactor core geometries would impact impurity and irradiation driven corrosion, however these effects cannot be quantified precisely enough to be considered.

2.2 Framework Introduction

In order to compare and optimize Molten Salt Reactor (MSR) designs a comprehensive framework on which to quantify costs is necessary. From this, cost minimization studies can be made that have the benefit of in-

[5]The theoretical and experiment results are weighted equally because without knowledge of which data is more accurate the average minimizes the expected deviation from the true value.

forming decision makers on how to efficiently direct research and development resources. The costs of any large research-design-construction project cannot be known with precision - especially with new technology and considerable regulatory constraints. Even if the estimated total cost differs significantly from the true value, there is still utility in conducting analysis to find an optimal design configuration (geometric arrangement and salt composition) so long as the inaccuracy of the cost assumptions are adequately consistent among design configurations. This approach to quantify costs over a range of configurations is the cost estimation framework, herein often referred to as 'the framework.' The objective of this section is to describe the cost estimation framework so that it can be applied to estimating the difference in cost between design configurations.

The cost optimization requires cost estimates of numerous input parameters, estimates that could not always be justified with supporting evidence. Despite the lack of reliability the estimates can still be useful as long as they are reasonable postulations. The contribution of this study is the *Cost Estimation Framework* and the implementation methodology, not the input parameter values. The unreliable estimates typically are associated with the determination of the reference costs - that is, costs associated with a select MSR configuration with typical/average characteristics. If the cost of an input parameter is lower than the true cost, the inaccurately low cost estimate lowers the cost of all configurations in the evaluation (though by varying amounts).

The framework does not account for all nuclear island or nuclear power plant costs, only those costs that depend on select geometric variables and the salt composition are accounted for. Although measured in dollars and representing dollar values, the configuration-specific 'total cost' is effectively the dependent variable in a cost function – also known as 'loss function.' The target of the 'total cost' value is not the most accurate estimate of an all-inclusive, sum-total cost but rather the sum-total of configuration-dependent costs. Some specific methodologies and the specific input variables required to arrive at the costs are discussed and listed in the respective FHR and MSR chapters.

The cost estimation framework is intended to be applicable to any nuclear reactor design that uses molten salt to remove heat from the core, including pebble-bed thermal-spectrum salt-cooled, thermal-spectrum circulating fuel, fast-spectrum circulating-fuel and molten-salt-fuelled molten-salt-cooled (known as stable-salt reactor) - a modified framework can likely be applied to non-MSR types as well. All reactor types that use molten salt to remove heat from the core, with or without dissolved fuel, will be referred to as MSRs throughout the remainder of this chapter.

Some costs can be calculated in a straightforward manner; a component may have specific cost at a specific time. That cost can be discounted through conventional economic analysis to a dollar value at a point in time that is comparable to the other costs. However, the cost of other consequences of nuclear power production, safety, proliferation mitigation etc. cannot be calculated in such a straightforward manner. A more effective means to estimate these non-straightforward consequences has been devised that involves separating them

into categories; the six categories, including the comparatively straightforward costs, are: direct, waste, safety, proliferation, modularity and feasibility.

The safety, proliferation and feasibility costs are calculated using broadly the same methodology. First the costs for a currently operational Light Water Reactor (LWR) are estimated, then the costs for a reference MSR are estimated based off of that information. This discussion constitutes the majority of the framework discussion. After the configuration specific cost is estimated relative to the reference configuration. The novel method of analyses used to compare the configuration specific to the reference cost is termed *Cost Setting*. The reference MSR is meant to represent a reactor with 'average' or 'typical' parameters - e.g., burnup, melting temperature, temperature reactivity feedback coefficients etc.

2.3 Direct Cost

Direct cost attempts to cover all straightforward costs associated with the basic construction and operation of the power plant. Most variables that contribute to direct cost include industrial goods or fabricated products, ideally with stated market prices. This generally covers all costs that are not strongly subject to the regulatory regime, management practices or operational procedures. The most clear contribution is material costs, which includes the various components of the salt. Included in this section is the cost of major nuclear island components such as the heat exchangers, the lost electric power revenue due to pumping requirements, continuously added fuel, some operating/maintenance costs and interest on capital. Direct cost does not include any balance of plant (BoP) construction costs unless they are a function of reactor configuration - given that it is assumed an optimization would have the same outlet temperature of the secondary salt in each scenario, the BoP costs are assumed to be equal for every configuration.

2.4 Waste

This analysis assumes that the fuelling cycle of all LWR and MSR design types are once through. The economics of reprocessing is likely more favourable with liquid fuel however reprocessing would have to be nearly an order of magnitude less costly than it is with current LWR systems to be cost competitive.[6] Bunn et .al [51] contains information and estimates, on the waste disposal costs of the current fleet of LWR reactors, that is useful in determining the waste disposal costs of potential MSR fuel cycles:

- The current US fuel repository program 'is financed by charging utilities a fee of 1 mill/kWh, which is equivalent to about $370/kg HM ($550, 2021) at time of discharge.'
- In the case of the U.S. Yucca Mountain Deep Geological Repository (DGR) total program costs are divided as follows:

[6] A 2005 study from Bunn et al. [51], assumes a 'central estimate' cost of repressed fuel for a LWR of $1000/kg HM ($1400/kg HM, 2021) though a 2009 DOE report assumes a reference cost of $3200/kgHM ($4250/kgHM, 2021) based on European experience [52] (pg D1-32). At the central estimated repressing, fuel fabrication, conversion and enrichment costs the study calculates a breakeven uranium price (where the once through and reprocessing fuel cycle costs are equal) of $368/kgU ($552/kgU in 2021).

  - 19%: Heat-related costs (repository construction and drip shield)
  - 53%: Those related to volume, mass, or number of items (repository emplacement operations and monitoring, waste package fabrication, and transportation)
  - 28% Other costs (siting, licensing, design, and engineering)

- 'Cost estimates produced by other countries for the disposal of spent fuel are roughly comparable.'

Using the categorical breakdown of LWR waste costs from Bunn et al., the total cost of waste disposal for MSRs can be derived for each cost category:

- Heat-related costs: The radiotoxicity/heat generation of fission products and actinides in LWR waste is approximately equal 100 years after reactor discharge, likely about the same amount of time before reactor waste is placed in a DGR. Heat generation is highest immediately following placement in the DGR, therefore the cost is determined by the heat generation rate upon initial placement in the DGR. As a result, the cost due to fission products and from actinides is assumed equal. The amount of fission products is nearly proportional to the thermal energy generated so the fission product contribution is proportional to thermal energy. For non-breeding reactors it is suggested that heat cost from actinide waste is treated as equal to fission products and also proportional to thermal energy. [7]

- Volume and mass or number of items: For LWRs it is assumed that the volume and mass contribution to this waste disposal cost is equal and the number of items is ignored. MSR volume and mass costs will be evaluated from the LWR reference; A burnup of 45.0 MWd/kgHM is assumed. Uranium dioxide fuel is 88% by mass uranium and has a density of approximately 11 g/cm^3. So measures of burnup by unit fuel are 39.5 MWd/kg fuel or 0.435 MWd/cm^3 fuel.

- Other costs: It is assumed that MSR waste would not need to be stored in a separate or unique facility to LWRs so the 'other costs' will be proportional to the heat-related and volume/mass costs.

The moderator of a thermal-spectrum, circulating-fuel MSR must be disposed of and may need be treated as high level waste. The cost of disposal is not a, or is a difficult to determine, function of the reactor configuration, so it is not included. In other MSR designs there is generally either no moderator (fast-spectrum) or the graphite moderator cannot easily be separated from the fuel (pebble-bed).

MSRs may produce non actinide, non fission product waste. The most serious radioactive waste is tritium and beryllium is a hazardous chemical waste. In lithium containing salts, tritium is produced primarily from neutron activation of ^6Li. With a half-life of 12.3 years, as a component of thermonuclear weapons and having a tendency to leak containment elements, it is treated as a safety and proliferation rather than a long-term waste concern (for more on detritation see Section 2.9.1). Beryllium fluoride is treated both as a safety and waste concern. It is quite toxic, the lethal dose 50% (LD_{50}) was found to be 98 mg/kg in rats [54], approximately equal to many mercury and arscenic compounds [55]. The handling costs of beryllium are captured by the safety category. The waste management costs are captured by the waste category and

[7] The heat generated from actinide waste a function of average burnup and the distribution of burnup.

Table 2.2: Summary of back-end fuel costs of currently operating U.S. NPPs derived using data from Bunn et. al. [51] and Supko [53].

	Cost (2021)	Unit
Dry cask facility	14	$
Dry cask at ToC[a]	8.8	$
Purchase & load casks	110	$/kg HM
Purchase & load	1,200,000	$/yr
Purchase & load ToC	20,500,000	$
DGR	1.07	$/MWh
DGR	12,900,000	$/yr
DGR	221,900,000	$
Total	234,800,000	$

[a]Time of commissioning

must be estimated imprecisely. For reference, the disposal fee in the United States for the similarily toxic mercury is $37,000/tonne [56]. Since used salt must be disposed of in a controlled fashion regardless due to it's radio-toxicity, the additional cost due to it's chemical toxicity likely would be less than chemical disposal cost alone. Thus an additional disposal cost of approximately 40% of estimated chemical disposal cost, $15,000/tonne or $15.00/kg Be is suggested.

2.5 Safety and Proliferation

Reactor safety and proliferation resistance costs are either part of the same calculation or calculated in a similar fashion thus it is convenient to discuss the costing of each together. For instance, worker monitoring and fuel tracking systems are in place to enhance safety and mitigate proliferation.

2.5.1 Capital Costs of Safety and Proliferation Mitigation

Capital costs of safety and proliferation mitigation: Generation III reactors

Addressing safety and proliferation concerns have an impact on reactor capital costs that is generally difficult to reliably attribute to specific components/systems. This may be a result of the complexity involved and the unique circumstance of each build. Table 2.3 shows a breakdown of reactor construction costs in the USA from the Energy Information Agency (EIA). It is useful to summarize the capital costs of recent Generation III/III+ reactor builds, see Table 2.4.[8]

It is well established that reactor construction costs, and in particular the overrun costs, have been on the rise since the 1970s [68]. The construction of physical structures does not explain the increase in costs. For instance, the change in the volume of structural materials is insufficiently small between Gen II and Gen III

[8]It is difficult to find accurate and reliable cost estimates for the builds in S. Korea, Japan and China, so the focus will be on Vogtle 3 & 4, Olkiluoto 3 and Flamanville 3.

Table 2.3: Base cost estimate for an advanced nuclear facility with a nominal capacity of 2,234 MWe [57].

Cost category	$ in 1,000,000s
Civil structural material & installation	1,927
Mech. equipment supply & installation	3,783
Elec. / I & C[a] supply & installation	701
Project Indirects[b]	3,029
Fee and contingency[c]	1,446
Owner costs[d]	2,395
Total project cost (excluding finance)	10,886 (2016)
Total project cost (excluding finance) /kW	5,945 (2016)
Total project cost (excluding finance)	12,347 (2021)
Total project cost (excluding finance) /kW	6,743 (2021)

[a]Instrumentation and Control

[b]Engineering, craft labour overtime & incentives, construction management, commissioning

[c]Contractor overhead costs, fees, profit, construction

[d]Development costs, preliminary feasibility & engineering, environmental studies & permitting, legal fees, insurance, taxes, tie-in to the electrical transmission system

Table 2.4: Capital costs and construction times of relevant recent nuclear power plant builds.

Reactor(s)	Country	Design	Capacity (GW, net)	Total project cost (B's)	Cost B's USD/GW, 2021	Construction[a] (start - criticality)
Vogtle 3 & 4	USA	AP1000	2.24	$25 + (2017) [58]	13.4 +	2009 - 2022 [59]
Olkiluoto 3	Finland	EPR	1.6	€8.5[b] (2012) [60]	8.0 +	2005 - 2022 [61]
Flamanville 3	France	EPR	1.6	€11 (2019) [62]	8.2	2007 - 2023 [63]
Hamaoka 5	Japan	ABWR	1.35	$2.73 (2007) [64]	2.88 +	2000 - 2004
Shika 2	Japan	ABWR	1.35	$2.73 (2007) [64]	2.88 +	2001 - 2005
Shin Hanul 1 & 2	S. Korea	APR	1.4	$6.7[c] (2012) [65]	6.1 +	2012 - 2019
Shin Kori 3 & 4	S. Korea	APR	1.4	$6.3[d] (2008) [65]	6.1 +	2008 - 2015
Taishan 1 & 2	China	EPR	3.2	$8[e] (2012) [67]	3.2 +	2009 - 2018

[a]construction start to (expected) first criticality. If multiple reactors, criticality of the first reactor.

[b]Aveva has not updated its 8.5 billion euro cost projection since 2012

[c]Assumed to be the same cost in 2021 USD as the Shin-Kori 3 & 4

[d]Currently there is no official publicly available construction cost data [66]

[e]Initial cost estimate

reactors [69] to explain the increase. Experience suggests that nuclear scale-up often undergoes 'negative learning' in which costs increase with accumulated experience [70]. Though construction cost overruns of effectively all types of large infrastructure projects, such as road and rail links, have been more common and greater in magnitude in the decades following the 1950's [71], it alone cannot explain the cost increases of Nuclear Power Plants (NPPs).

According to Cooper, [72] the drivers of cost increases are 'construction delays and cost overruns, as well as regulatory changes.' Cohen [73] finds that nearly all of the reasons for the increase in reactor construction costs through the 1970s and 1980s in the US were closely linked to more stringent and changing regulations. Davis [74], found that most studies attribute the increase in construction time to a 'rapidly evolving regulatory process.' Construction time is strongly correlated with capital cost; Harris et al. [68] found that for every year a NPP is delayed the Levelized-Cost-Of-Energy (LCOE) increases by approximately 8–10%. It can be presumed that the regulatory regime is responsible for a large portion of the increased costs of NPPs, either directly or due to intangible and/or unforeseen consequences such as prolonged construction time and re-manufactured parts.

A Pittsburgh Post-Gazette article investigating issues with the AP1000 build-out [75] contains insight into the causes of cost overruns. It finds construction before completion of design to be a major cause. As well it mentions cost overruns and delays that were directly or indirectly a result of addressing safety and/or regulations; for example the NRC mandated that the shield building must withstand an airplane crash - thus delaying final design certification. For certain components, if documentation were lost or their product tags weathered beyond recognition the component would be rendered useless. Though a major MIT study can be seen to be in contradiction [76], a large portion of costs are nonetheless assumed to be the result of a more stringent and evolving regulatory environment. The study instead finds the root cause of cost increases be to project and construction management practices. These two major explanations are not mutually exclusive and likely even reinforce one another, a flexible regulatory environment wouldn't be so punishing to relatively minor management mishaps. More financially successful builds in Asia (Korea and China) demonstrate that major cost overruns aren't an inevitability, the regulatory regime and/or management practices must have a significant effect.

Modern reactors being built in western countries, at Vogtle, Olkiluoto and Flammenville are or nearly are First-Of-A-Kind (FOAK) reactors with insufficient publicly available information necessary to build a ledger listing all costs. Thus the proportion of capital costs due to safety and proliferation mitigation cannot be determined directly, instead it must be estimated by examining expected construction costs and assuming a safety cost share for each cost category, an expected cost overrun and a safety cost share for the overrun. Initially safety will be considered in this analysis as it has a much larger and more predictable contribution to cost than proliferation mitigation. Table 2.5 shows the results of the analysis with proposed safety cost shares. The non safety share represents the cost if nuclear regulations have little effect and construction requires the only the precision, attention and documentation typical to other construction projects. A 100%

Table 2.5: Safety and proliferation mitigation share of costs for an advanced nuclear power plant (dual AP1000 reactors) Cost projections form EIA [57]. See Table 2.3 for cost category description. Proposed safety cost shares are by necessity selected in an intuitive fashion.

Cost category	Safety share, %	Total cost (2021 $ MM)	Safety cost, $
Civil, mech, elec material/ equipment and installation	15	7,340	1,101
Project indirects	25	3,468	867
Fee and contingency	15	1,655	248
Owner cost	30	2,742	823
Total no overrun	20	15,206	3,041
Cost of 50% overrun	60	7,602	4,561
Total incl. overrun	33.3	22,808	7,602
Cost of 100% overrun	60	15,206	9,123
Total incl. overrun	41.5	30,410	12,164
Cost of 200% overrun	60	30,410	18,246
Total incl. overrun	48	45,616	21,287

cost overrun will be taken as the expected or reference result, corresponding to a 1000 MWe reactor, this results in a cost approximately equal to that of Vogtle, Olkiluoto and Flammenville. A relation for overrun as a function of power (MWe) is provided in Section 2.6 on modularity.

Capital costs of safety and proliferation mitigation: MSR

Both the total capital costs of a MSR and the share of capital costs dedicated to safety and proliferation must be determined. Despite being inherently more safe, a MSR will have similar nuclear-specific and safety related costs such as licensing, craft labour and commissioning. Therefore, the cost categories of project indirects, fee & contingency and owner costs are expected to have the same shares dedicated to safety as a LWR. As well, the reactor will require security systems, containment, a drain tank and other safety and proliferation mitigation systems, so the safety shares of civil, mechanical and electrical materials/equipment is expected to also remain approximately the same (20%). A MSR would require craft labour and specialized components/materials that are subject to stringent quality assurance, same as a Generation III pressurized water reactor (PWR) - therefore, MSRs are likely subject to the same cost overrun risk. However, (thermal-spectrum) MSR designs generally have a lower power output than Generation III water-cooled reactors and thus are more modular - reducing the cost overrun risk and consequentially the safety costs. The proliferation mitigation as a percentage of capital cost is estimated using the same methodology, non-overrun and overrun costs of 0.05 and 0.10 are suggested.

The literature on unbiased, reliable MSR capital cost estimates is sparse. Extrapolated results from L. Samalova, O. Chvala and G. Maldonado [77], in a cost comparison of the AP1000 to IMSR600, IMSR300 and IMSR80, would suggest that a NPP with 500 MWth MSR units is expected to be equal in cost per unit energy to an AP1000 NPP. The Mark-1 PB-FHR design (salt-cooled) has more detailed publicly avail-

able cost estimates than any circulating-fuel MSR. The estimated capital cost for a single 100 MWe unit is $5,880/kW and $4,210/kW (2021 $) [78] for a 12 unit power plant. A paper by R. Moir [79] analysing cost comparisons involving a 1970s 1000 MWe MSR design from the ORNL Molten Salt Reactor Program (MSRP) finds that the cost of the MSR and PWRs of the era are nearly the same. The MSR designs from the MSRP were generally more complex as breeding/high-burnup was a priority - though that focus began to shift as time went on. As a whole, much literature appears to be in general agreement that non-cost-overrun capital cost (per unit energy) of an individual (modular) MSR reactor would be approximately equal to large Gen III LWRs and a power plant with many units would be somewhat ($\approx 10\%$) lower than a Gen III LWR NPP of the same capacity. Because MSRs are expected to have a lower capacity and therefore be more modular it is suggested that the cost overrun risk be lower - see Section 2.6.

The proposed method to compare safety mitigation costs of different configurations follows Equation 2.10, where R is the relative safety cost, Sc_{config} the configuration specific safety cost, Sc_{ref} the reference MSR safety cost and E_{sc} the safety cost exponent, which has a value less than 1.0.

$$R = \frac{Sc_{config}}{Sc_{ref}}^{E_{sc}} \tag{2.10}$$

2.5.2 Operation & Maintenance (O & M) Costs of Safety and Proliferation Mitigation:

A procedure is devised to arrive at the safety and proliferation O & M costs based on those costs for the current US nuclear fleet and the Core Damage Frequency (CDF) of each:

1. Determine the cost of safety and proliferation mitigation in the current nuclear fleet.
2. Determine the cost of safety as a function of reactor CDF.
3. Assuming a CDF of a reference MSR, calculate the O & M cost of a reference MSR.
4. Apply the same procedure and assume reference proliferation mitigation costs follow a similar risk-to-cost function.

O & M costs of safety and proliferation mitigation: cost of safety and proliferation mitigation in the current nuclear fleet.

The safety of nuclear reactors exceeds the economic optimum, so instead of applying any conventional economic risk optimization analysis an alternative technique has been devised: separate safety and proliferation mitigation costs from other O & M costs. To create this separation actual O & M costs are compared to a hypothetical NPP that is perfectly safe and proliferation resistant ('ideal'). It is inherently completely safe and proliferation resistant and so only requires only the security, monitoring, operational procedures etc., typical to non-nuclear industries. The only costs are those required for the reactor to physically operate; managing fuelling, control and necessary maintenance.

Table 2.6: Nuclear power plant categories of operation and maintenance costs.

Categories	Share of total cost (%)	Cost ($/MWh)	Share Safety	Share Proliferation	Cost Safety	Cost Proliferation
Work management	24	$4.06	0.525	0.125	$2.13	$0.51
Support services	32	$5.41	0.35	0.10	$1.89	$0.54
Operation	15	$2.54	0.5	0.0	$1.27	$0.00
Loss prevention	11	$1.86	0.55	0.45	$1.03	$0.83
Engineering	8	$1.35	0.5	0.0	$0.68	$0.00
Other	10	$1.69	0.2	0.0	$0.34	$0.00
Total	100	$16.90	0.43	0.11	$7.33	$1.88

To estimate the O & M costs of a perfectly safe and proliferation resistant plant, the share of cost in each category dedicated to safety and/or proliferation resistance must be estimated. Three variations are employed to estimate the overall share of NPP costs dedicated to safety and proliferation resistance in this way. The safety and proliferation portion of most cost categories must be estimated in a somewhat unsystematic, intuitive manner. Proposed safety and proliferation mitigation costs are provided.

1. Total O & M costs broken down in six categories.
2. Labour costs of 'Nuclear Electric Power Generation' broken down into approximately 100 distinct categories.
3. Breaking down O & M costs by Balance of Plant (BoP) and nuclear island.

Variation 1: Cost categories, listed in Table 2.6, are from the Energy Information Agency (EIA) [80]. The total cost of each category is multiplied by an estimate of the the share of safety and proliferation resistance costs in each category (columns 4 & 5) to come up with the safety and proliferation resistance cost of each cost category.

Variation 2: A sample of the safety and proliferation resistance calculations of a select number of labour categories is shown in Table 2.7. The data is from the U.S. Bureau of Labor Statistics [81] and the total labour cost at all nuclear reactors is $3,973.9 million. For each category the share of costs that are a result of addressing safety and proliferation concerns are estimated. Overall, 43.6% of labour costs were estimated to result from addressing safety and 14.2% proliferation concerns. The remaining 42.2% of all other labour costs, represent hypothetical labour costs of an 'ideal' reactor.

Variation 3: The total cost estimate for a 'ideal' NPP compared to a real NPP was determined by dividing the total cost into components and estimating the 'ideal' cost of each. Operations and maintaince were each divided into BoP and nuclear island to arrive at a more reliable 'ideal' NPP cost. Using the proposed costs, the total O & M cost of ideal reactors was found to be $8.71/MWh compared to $16.90/MWh for real NPPs - see Table 2.8. Thus $8.19/MWh or 48.5% of O & M costs are the result of addressing to safety and proliferation concerns. Several different methods were employed to estimate the cost of each category in an

Table 2.7: Nuclear power plant sample categories of labour costs. The safety and/or proliferation resistance costs of labour categories without a specified safety and/or proliferation resistance share are proportional to the average safety and proliferation share of the labour categories with a specified safety and proliferation share (i.e. the number of extra human resource workers due to safety concerns is dependent on the number of extra security guards due to safety concerns.)

Occupation[a]	Mean annual salary ($ 1000's)	% of total employment	Proportionality	Safety	Proliferation
Management	$90.4	7.4	1.0	—	—
Compliance officers	$91.3	0.6	—	0.8	0.2
Human resource workers	$83.6	0.5	1.0	—	—
Information security analysts	$74.1	0.8	—	0.4	0.6
Drafter, eng. technicians	$79.7	2.0	0.7[b]	—	—
Security guards	$48.3	10.1	—	0.2	0.8
Industrial machinery mechanic	$72.8	2.5	—	0.2	0

[a]The full dataset is with all occupations can be provided upon request.

[b]In cases where the proportionality is less than 1, the share of costs that is proportional to safety and proliferation resistance is set by the proportionality variable and the rest of the costs are not dependant on safety and/or proliferation resistance.

ideal reactor including U.S. EIA data on the O & M costs of nuclear, coal and gas power plants, see Table 2.9 [80]. The rationale and assumptions made to determine the values in Table 2.8 are:

- Balance of Plant (BoP) maintenance: $2.48/MWh = average(0.6·(coal maintenance) + gas maintenance) = average(0.6·$4.54 + $2.24).[9]
- The BoP operations and maintenance costs are the same in the ideal NPP and the real NPP.
- The nuclear island maintenance costs in an ideal reactor are $\frac{1}{2}$ the maintenance costs in a real reactor.
- A NPP is similar in complexity to a coal power plant. Complexity referring to the size and number of systems - a NPP may be less intuitive or straightforward but does not have large/convoluted on-going fuelling systems or emissions mitigation systems. Thus the nuclear island operating costs are assumed to be 80% of the coal operating costs ($5.01 for BoP and coal island) for an ideal NPP.
- The BoP operating costs are $2.00/MWh; this assumption is by necessity largely discretionary.

The reasonable similarity of the results using each method suggests reliability and some degree of validity of both the methodolgy and proposed safety and proliferation shares. The results are summarized in Table 2.10. Variations 1 and 2 are assumed to be more comprehensive and reliable so more weight is given to those costs share. Going forward the safety cost share used is 43% and the proliferation mitigation share is 12%. Further a study from the American Action Forum found that regulatory compliance costs at NPPs average $60 million [82], or 54% of the total estimated safety and proliferation mitigation costs of Variation 1.

[9]It has been assumed that approximately 60% of maintenance costs at a coal power plant are directed toward the turbine, pumps and condenser while the other half is for the boilers and scrubbers.

Table 2.8: Cost breakdown by balance of plant and nuclear island for a hypothetical perfectly safe and proliferation resistant (ideal) NPP compared to a real NPP.

Category	Cost Ideal ($/MWh)	Cost Total ($/MWh)
BoP[a] operations	2.00	2.00
BoP maintenance	2.48	2.48
Nuc. island operations	2.01[b]	8.27
Nuc. island maintenance	1.93	3.86
Distribution/transmission	0.28	0.28
Total	$8.71	$16.90

[a] Balance of Plant
[b] (0.8)·5.01 - 2.00

Table 2.9: Operations and maintenance costs of thermal electricity production [80]. Note that 'nuclear operations' + 'nuclear maintenance' = 'total' in Table 2.8.

Category	Cost ($/MWh) Operations	Cost ($/MWh) Maintenance
Nuclear	$10.27	$6.63
Coal	$5.01	$4.54
Gas	$2.45	$2.24

Table 2.10: Cost of safety and proliferation resistance in U.S. nuclear reactor.

	Cost type	Share of costs (%)	Mean annual cost, $1000's/NPP	$/MWh
Variation 1	total, safety	43	$87,960	$7.33
	total, proliferation	11	$22,560	$1.88
Variation 2	labour, safety	44	$26,250	$2.19
	labour, proliferation	14	$8,550	$0.71
Variation 3	total, safety and proliferation	48	$98,296	$8.19

O & M costs of safety and proliferation mitigation: The cost of safety mitigation as a function of risk (CDF)

To meet the objective of this section, a function that relates the total cost of safety to a quantitative measure of safety is required. CDF is selected as the measure of safety. The function cannot be determined directly since insufficient data is available on both the cost of individual safety measures and the associated quantitative increase in safety. However, the cost of safety mitigation for currently operational reactors has been estimated in the previous section: 43% of operating costs, or $7.27/MWh. That can serve as a starting point for a safety cost - CDF function. A goal of safety systems is to be independent so that each system reduces risk by a multiplicative factor less than one. This results in a clear exponential relation between risk and cost. J. K. Vaurio [83], and A. Mancuso et. al [84] each discuss an example that demonstrates this relationship for a particular failure mode. This exponential relationship appears to hold true for both capital and operating costs. The vast majority of power reactors are a similar type (water-cooled), built at approximately the same time and so have a similar CDF. Therefore there is not enough variation in CDF to reliably determine a relation to cost - so an alternative approach is taken:

For the last two decades the rate of nuclear power accidents or incidences (events) with a cost in excess of $20 MM has been constant at roughly 0.003 events per reactor per year [85]. With the assumption that $20 MM + events in a hypothetical reactor without safety systems/protocols would likely lead to core damage, the CDF of a reactor without safety systems/protocols would be approximately equal to the frequency of $20 MM + events. This assumption can be used to establish two data points and from the data points a function relating reactor safety cost to CDF can be made. The first data point is at an O & M safety cost of $0.00/MWh with a CDF of $3 \cdot 10^{-3}$/yr and the second is at a cost of $7.27/MWh with a CDF of $1.0 \cdot 10^{-5}$/yr. (CDF of $1.0 \cdot 10^{-5}$/yr is the estimated average value for reactors in operation, based on [86]) The resulting relation for the O & M cost of safety, using an exponential fit, is shown in Equation 2.11:

$$Cost_{safety}[\$/MWh] = m \cdot (\log(CDF_x)) + b \tag{2.11}$$

Where m is the slope, equal to -2.93 and b is the y-intercept equal to -7.39. Equation 2.11 demonstrates that for a given reactor, the higher safety (smaller CDF value) obtained, the higher the excess cost. Consequently, an inherently more safe design with an equal non-safety cost will cost less to achieve an equal level of safety.

O & M costs of safety and proliferation mitigation: O & M cost of a reference MSR

With a conservative CDF estimate for a MSR of $1 \cdot 10^{-7}$, (see Section 2.5.4), at an equal safety cost to LWRs, a data point is established: CDF = $1 \cdot 10^{-7}$, cost = $7.27/MWh. With the assumption that the $\frac{d(CDF)}{d(Cost)}$ (i.e. the slope) is equal for currently operating reactors and future MSRs, the equivalent relation to Equation 2.11 for a MSR can be made:

$$Cost_{safety}[\$/MWh] = -2.93 \cdot (\log(CDF_x)) - 13.2 \tag{2.12}$$

Modern designs generally have CDFs much higher than regulator targets. As well, all nuclear reactor accidents that registered 4 or higher on the INES scale occurred at reactors constructed before 1983. (1978 if the uniquely unsafe conditions of Soviet nuclear power is ignored) Thus an inherently more safe reactor is likely to be operated at a higher safety level (CDF) than a conventional reactor - how much safer cannot be determined precisely. A proposed level of safety is the geometric average of the MSR CDF and the CDF of reactors currently in operation. A MSR would therefore operate at a CDF of 10^{-6}.

$$Cost_{safety}[\$/MWh] = -2.93 \cdot (\log(10^{-6})) - 13.2 = 4.33 \tag{2.13}$$

Equation 2.13 demonstrates that the safety cost of an MSR would be \$4.33 at a CDF of $1.0 \cdot 10^{-6}$. Assuming an equal non-safety cost, that is 31% of total operating costs. This compares to typical currently operational LWRs as follows:

- Typical currently operational LWR: $\frac{\$7.27/MWh}{\$16.90} = 43\%$
- Reference MSR: $\frac{\$4.33}{\$16.90 \cdot (1-0.43)+\$4.33} = 31\%$

It is proposed that the safety mitigation cost of different configurations use the same relation as in capital costs, see Equation 2.10.

2.5.3 Insurance costs: The Cost of Previous (Generation II/III) Accidents/Incidents.

The final separate component of safety related costs is that of insurance. The cost of insurance is derived from the cost of events (accidents or incidents) of past reactor operation. The risk of MSR events is assumed to be a function of the reference CDF and the safety of each particular configuration (geometry and salt composition) relative to the reference MSR safety value.

A statistical analysis of nuclear power incidents and accidents by S. Wheatley, B. Sovacool and D. Sornette [87] provides insight into the cost of these events. The study considers only events with a cost above \$20 million USD, the dataset includes 175 events. It is mentioned that the current rate of events is expected to be 0.003/reactor/year. Post-1979 (post-TMI), the frequency and severity of events \$20-1000 million USD follows a Pareto distribution with $\alpha = 0.5 - 0.6$. Using this data the expected statistical average cost of \$20-1000 million accidents is \$0.40 million/reactor/year, see Equations 2.14 through 2.18. Extrapolating the distribution to include all accidents \$1-1000 million results in a cost of \$0.46 million/reactor/year.

$$Cost = \frac{real\ events}{Pareto\ events} \cdot ((Pareto\ events)(\frac{cost}{event})) \tag{2.14}$$

$$Cost = \frac{0.003}{\int_{20}^{1000} \frac{\alpha\mu^{\alpha}}{x^{\alpha+1}}} \cdot \int_{20}^{1000} \frac{\alpha\mu^{\alpha}}{x^{\alpha+1}} \cdot x \tag{2.15}$$

Where α is the power value property of the Pareto distribution - equal to 0.55, μ is the minimum x value

and x is the cost in millions of USD.

$$Cost = \frac{0.003}{\frac{\alpha\mu^{\alpha}}{-\alpha} \cdot (x^{-\alpha})\big|_{20}^{1000}} \cdot \frac{\alpha\mu^{\alpha}}{1-\alpha} \cdot (x^{1-\alpha})\big|_{20}^{1000} \tag{2.16}$$

$$Cost = \frac{0.003}{-20^{0.55}(1000^{-0.55} - 20^{-0.55})} \cdot (\frac{0.55 \cdot 20^{0.55}}{0.45})(1000^{0.45} - 20^{0.45}) \tag{2.17}$$

$$Cost = \frac{0.003}{0.884} \cdot 118 = \$0.40\, million \tag{2.18}$$

This technique can predict the cost of relatively small events with sufficient statistical accuracy to be relevant but provides little or questionable utility in predicting extremely costly events. Two events are considered extreme: Fukushima and Chernobyl. Wheatley, Sovacool and Sornette found that the cost of Fukushima and Chernobyl is equal to nearly 5 times the sum of the other 173 events in the dataset [87]. It is unlikely that the risk of a Fukushima and Chernobyl is the same as it has been over the last 4 decades but it is unclear how much lower the risk is. If the risk of an extreme event is 5 times the risk of \$20-1000 million events, the risk of an extreme event is $5 \cdot (\$0.4 \text{ million}) = \$2.0$ million. Thus a conservative estimate of total risk of safety incidents/accidents is $\$2.0 + \$0.46 \approx \$2.5$ million/reactor/year.

2.5.4 Insurance costs: The Cost of MSR Accidents/Incidents and the Cost of Insurance

Studies of probability and consequence (risk) of potential MSR accidents/incidences were not found in the current literature. Neither were probabilistic risk/safety assessment data found for any other low-pressure, higher-temperature, reactors with unreactive molten coolant (lead-cooled). No Gen IV reactor has issued nor submitted a design certification application to the NRC which would include this information. However the water cooled NuScale design has and it is largely comparable to Gen IV reactors. It is a modern passively safe Small Modular Reactor (SMR). According to an NRC document, *'Probabilistic risk assessment and severe accident evaluation for new reactors,'* [88] the CDF, total all hazards, for the NuScale design is $3.78 \cdot 10^{-9}$.[10]

The CDF of medium-sized MSR designs are also likely to be low. A conservative estimate is on the order of $1 \cdot 10^{-7}$, the basis for the estimate is:

- Successive NPP designs tend to increase in safety.
- The NuScale design shares many objectives/features with MSRs and has a very low CDF probability.
- MSRs are often regarded as more safe than other designs, those that are notably water-cooled [90] [91], though insufficient evidence exists to verify these claims.

The CDFs of reactors currently in operation, consisting largely of Gen II designs, are mostly between $1 \cdot 10^{-5}$ and $1 \cdot 1^{-4}$ [92] - the lower value, $1 \cdot 10^{-5}$ is recommended as an input to Equation 2.19 due to enhanced safety practices since initial licensing. The real CDF differences between reactor designs is likely less than

[10] The AP1000 for comparison has a CDF of $2.97 \cdot 10^{-7}$ [89]

estimated/nameplate CDF values and even if the CDF did accurately reflect safety differences it would unlikely be viewed this way by insurers. Thus a function to smooth CDF differences is suggested, Equation 2.19, it also accounts for the safety of each particular configuration relative to the reference MSR safety value through the $\frac{Sc_{config.}}{Sc_{ref.}}$ fraction:

$$Risk[\frac{1}{GW-yr}] = (\frac{\$2.5\,MM}{[GW-yr]})\frac{(\frac{Sc_{config.}}{Sc_{ref.}})\log(CDF_{ref.\,MSR})}{\log(CDF_{Gen\,II})} \tag{2.19}$$

Sc_{config} is the configuration specific safety cost, $Sc_{ref.}$ the reference MSR safety cost, $CDF_{Gen\,II}$ the CDF of a typical Gen. II reactor and CDF_{MSR} the CDF the reference MSR. It is assumed that the probability of an accident/incident is independent of reactor size and that the average damage of an event is proportional to the thermal power - thus the overall cost of insurance is directly proportional to thermal power. The total insurance cost is the product of the expected risk and the insurance premium, the proposed insurance premium increases cost by a factor of two.

2.5.5 Proliferation mitigation costs of a reference MSR

It is challenging to estimate proliferation mitigation costs of MSRs because of the unique fuel structures and fuelling systems. All fuel used in current operating reactors is solid and apart from one pebble-bed reactor in China, the HTR-10, is fabricated into bundles that are large enough to individually have a serial number and be tracked. Many MSR designs are different in that fissionable material is liquid during operation or in small freely-moving pebbles. MSR designs have various fuelling schemes most of which involve continuous fuel addition. This makes it difficult to track the fuel. Nonetheless, the fuel of a thermal-spectrum reactor is of low fissile content, likely less than 5% for a circulating fuel reactor and less than 20% for a pebble-bed reactor. Fuel in the reactor or spent fuel is also difficult to recover. In a circulating-fuel reactor the fuel is dissolved in salt along with most fission products. In a pebble-bed reactor the fuel is encased in TRISO particles dispersed throughout each pebble. This fuel is specifically designed with multiple layers of protection and is difficult to purposely extract. Table 2.11 examines some cost categories associated with proliferation mitigation. It is likely that proliferation mitigation costs of a reference MSR design would be somewhat higher than a currently operational Gen II reactor. The estimated costs for a Gen II reactor as stated in Table 2.10 is 12% of the $16.90/MWh total operating costs, $2.03/MWh[11] and for a reference MSR the estimated cost is $2.50 - $3.00 for solid fuel and dissolved fuel respectively.

The proposed method to compare both the capital and O & M proliferation mitigation costs of different configurations is similar to that of safety costs, it follows Equation 2.20, where R is the relative proliferation cost, Pc_{config} the configuration specific proliferation cost, Pc_{ref} the reference MSR proliferation cost and

[11] This cost estimate is for a NOAK reactor, additional costs associated with a FOAK reactor are covered in the section addressing feasibility.

Table 2.11: MSR proliferation mitigation costs relative to Gen II reactors.

Labour category	Circulating fuel	Pebble bed	Notes
Compliance	same	same	
IT security	same	same	
Physical security	same	same	
Record keeping	slightly more	slightly more	likely more stringent requirements
Materials handling	slightly more	slightly more	more difficult to handle
Inspectors/testers/sorters	more	slightly more	more inspections/testing required to verify no fuel diversion

E_{pc} the proliferation cost exponent, which has a value less than 1.0.

$$R = \frac{Pc_{config}}{Pc_{ref}}^{E_{pc}} \tag{2.20}$$

^{239}Pu cost quantification

The cost of proliferation is calculated relative to the reference proliferation cost ($2.30-$3.00, see above) and is a function of the ^{239}Pu isotope concentration and total amount of ^{239}Pu produced. This section discusses how the cost is calculated. The relative Pu cost is one component of multiple relative costs (tritium production, uranium enrichment - if applicable) that determine the relative and ultimate proliferation cost. The proliferations costs are to be calculated using the novel 'Cost Setting' method as described in the upcoming Section 2.8.

The ^{239}Pu concentration can be considered to have reached an asymptote at approximately 30 MWd/kgHM which is approximately equal to 0.008 kgPu/kgHM. Up until 30 MWd/kgHM the ^{239}Pu is assumed to follow a square root function, Equation 2.21. Before including the effect of the proliferation cost exponent, the Pu proliferation cost is proportional to both the amount of net ^{239}Pu produced by the reactor per unit of energy generated and the concentration of ^{239}Pu as a share of all Pu isotopes. This calculation requires an input cost for a reference configuration, the suggested value is $\$5{\cdot}10^6$ for a reactor that has an equilibrium average burnup of 40 MWd/kgHM. Equation 2.22 shows how the Pu proliferation cost is calculated.

$$\begin{aligned} &x <= 30\,[MWd/kgHM]\text{: } Pu[kgPu/kgHM] = 0.00146 \cdot \sqrt{x} \\ &x > 30\,[MWd/kgHM]\text{: } Pu[kgPu/kgHM] = 0.00146 \cdot \sqrt{30} = 0.008 \end{aligned} \tag{2.21}$$

$$C_{Pu,i} = C_{Pu,ref} \frac{B_{ref}}{B_i} \frac{\left[^{239}Pu_i\right] \cdot \frac{\left[^{239}Pu_i\right]}{[Pu_i]-\left[^{239}Pu_i\right]}}{\left[^{239}Pu_i\right] \cdot \frac{\left[^{239}Pu_{ref}\right]}{[Pu_{ref}]-\left[^{239}Pu_{ref}\right]}} \tag{2.22}$$

The subscript i is assigned to values associated with the particular configuration analysed and ref to values associated with reference configuration. C is the cost, B the average burnup of removed fuel, $[^{239}Pu]$ the

concentration of the ^{239}Pu isotope in the fuel and $[Pu]$ the concentration of all Pu isotopes in the fuel. The burnup is required to make the cost proportional to the amount of fuel removed from the reactor.

2.6 Modularity

Because the framework is built on cost additions rather than cost reductions, the cost of modularity is the additional cost as a result of a lack of modularity. It is not a specific cost that is directly attributable to any particular process or component rather it reflects the difference in the cost of factory and on-site construction. A large, less modular reactor requires more on-site construction and less factory construction as a proportion of total construction costs. Studies show that on-site construction costs are more expensive, have risen faster in recent decades and more likely to go over budget and/or be delayed than factory work [93]. The extra on-site construction cost is the summation of a share of the overrun cost and the share of the base (non-overrun) cost. Both the overrun cost and the share of base cost are a function of reactor size (power (P), MWe). The two suggested functions are shown in Equations 2.23 and 2.24. The share of the cost overrun that is due to the modularity cost (ie. lack of modularity) is 1.0 minus the cost overrun share due to safety and proliferation (using proposed values that is equal to 0.30). The proposed reference share of base costs is 0.15 for a reference 500 MWe reactor. Like many other function constants it cannot be known precisely and an alternative value may be used as per user discretion.

$$Percent\ overrun = \sqrt{100 + 10 \cdot P(MWe)} \tag{2.23}$$

$$Share\ of\ base = (\text{ref. share}) \frac{(P_i(MWe))^{E_{mod}}}{(P_{ref.}(MWe))^{E_{mod}}} \tag{2.24}$$

Where, $P_{ref.}$ the MWe of a reference reactor, P_i the MWe of the reactor under evaluation and E_{mod} the modularity share exponent, it's proposed value is 0.5. It is assumed that it is only technologically and/or economically feasible to construct fast-spectrum MSRs at a large scale, ($\approx$ 1.0 GW) so in this framework modularity is only relevent in thermal-spectrum designs.

2.7 Feasibility

Feasibility refers to technological challenge, required preparedness and complexity; the amount of fundamental research that must be conducted before detailed design and construction, as well as operational difficultly and expected maintenance. It is perhaps best understood by the input parameters that contribute to feasibility, see Table 2.12.

Because the feasibility cost is composed of largely unrelated input parameters, the framework requires estimates of the feasibility cost of each individual input parameter. Where possible, empirical sources or analytical processes should be used to estimate cost, however they may largely have to be determined intuitively. The feasibility cost of each parameter is divided into two categories listed below and the division

Table 2.12: Reference technological development cost and reactor-specific share of cost to feasibility

Feasibility parameters	Development cost share	Each-build cost share
T_{melt}	0.2	0.8
P_{vap}	0.2	0.8
Corrosion	0.7	0.3
Xe removal[a]	0.5	0.5
Salt processing[a]	0.5	0.5

[a]If applicable

between the two categories is shown in Table 2.12.

- Basic technological (tech) development costs: Required R & D costs before detailed design begins.
- Costs specific to each reactor build (each-build): Cost that contribute to the construction of each new reactor whether through higher construction or O & M costs.

The reference basic tech development cost is the product of a basic research cost constant and a value representing the share of basic research costs borne by an individual reacted slated for construction. The configuration specific tech development cost ($Td_{config.}$) is to be determined relative the reference tech development cost:

$$Td_{config.} = Td_{ref.} \cdot \frac{C_{config.}}{C_{ref.}} \tag{2.25}$$

Where $Td_{ref.}$ is the reference basic tech development cost, $C_{config.}$ the configuration specific total cost and $C_{ref.}$ the reference total cost. $C_{config.}$ and $C_{ref.}$ should be calculated using the 'cost setting' methodology, discussed in Section 2.8 and should make use of the 'development cost share' or the 'each-build cost share.' The relevant cost share should be multiplied by associated relative cost.

The basic technological development cost can be estimated using values in the Gen IV roadmap and by extrapolating the experience of other reactor development programs:

- The 2002 Generation IV Roadmap [94] provides an R & D schedule and cost for each of the 6 reactor types during the 'performance phase' (resolve issues and determine viability), see Table 2.13. It also mentions '... at least six years and several US $ billion will be required for detailed design and construction of a demonstration system.'
- The Clinch River breeder reactor was expected to cost twice as much as a LWR - $3.4 vs $1.7 billion in 1981 ($8.4 vs $4.2 billion, 2021 $) though the original cost estimate was significantly lower [95]. The 350 MWe reactor was never completed and the total sunk cost was 'about' $1.7 billion ($4.3 billion, 2021 $) according to the project's final report [96].
- An estimate of the total CANDU reactor development cost could not be found however it can be roughly estimated from Atomic Energy of Canada Limited (AECL) funding. In fiscal years ending 1953 to 1962, approximately the height of CANDU development, AECL received a total of $262

Table 2.13: Technological development cost of Generation IV reactors according the R & D schedule of the 2002 Generation IV Roadmap [94].

Reactor system	R & D, 2002 $ MM	R & D, 2021 $ MM
Gas-cooled fast reactor	940	1438
Lead-cooled fast reactor	990	1515
Molten salt reactor[a]	1000	1530
Sodium-cooled fast reactor	160	245
Supercritical-water-cooled reactor	500	762
Very-high-temperature reactor	670	1025

[a]Fast or thermal spectrum unspecified

million CAD in federal funding [97] ($2,490 million in 2021 CAD). Other parties were involved and contributed to the financing of CANDU development including Ontario Hydro and Canadian General Electric [98].

- The development cost of the NuScale reactor is somewhat unclear, it was expected to cost $1.4 billion [99] (2018) and the cost of a 12-module plant $4.2-6.1 billion [100] (2019).

Pre-construction costs for a revolutionary design such as the CANDU or the sodium-cooled fast breeder were in excess of $2.0 billion. It would be naive to expect much less for a reactor that differs so much from any commercial reactors in operation. Even the development of the NuScale design, a small PWR, is expected to cost $1.4 billion. The roadmap estimates basic R & D for a MSR to cost approximately $1.5 billion - and in this type of situation it would be prudent to assume some level of cost overrun. Detailed design can be expected to cost in excess of $1.0 billion. Thus an optimistic MSR technological development cost is estimated to be $3.0 billion; $1.5 billion for R & D and $1.0 for detailed design and $0.5 for licensing, contingency, cost overruns etc. The R & D is dependent on the type of MSR design. The R & D estimated from the Generation IV Roadmap is quite extensive compared to the requirements of some proposed designs. As well it is expected that the organization developing a new reactor will not bare the entire technological development cost themselves. Thus the a technological development cost of $1.0 is proposed instead.

The other contributor to feasibility cost is the each-build portion. It is suggested to estimate it's cost by adjusting the non-configuration direct cost (non-config. or C_N) as shown in Equation 2.26. The non-config. cost is intended to account for nuclear island capital costs that are not captured by the other specific, configuration dependant costs. These non-config. costs are assumed not to fit clearly into any of the other cost categories and to be more effected by technological challenge and complexity. There are two underlying assumptions behind the feasibility each-build cost calculation. One, that the non-config. cost represents minimum cost for a theoretical reactor design without a technological challenge. And another that the cost would increase by an amount both proportional to the non-config. cost and a function of the configuration-specific to reference feasibility difference. The target each-build cost is therefore the additional expected cost beyond a (theoretical) design with no serious technology challenge. Equation 2.26 suggests how the

ΔC_N (ie. each build cost) should be calculated.

$$\Delta C_N = C_N \cdot ((\frac{C_{config.}}{C_{ref.}})^{E_{E,B}}) \tag{2.26}$$

C_N is the non-configuration, $C_{config.}$ the config. specific each-build feasibility and $C_{ref.}$ the reference each-build feasibility cost. $E_{E,B}$ is the each-build exponent. Proposed values for C_N, $E_{E,B}$ and F_{exp} (Equation 2.27) are those used in the FHR and MSR evaluations, Table 5.10/5.11 and 6.12/6.13 respectively.

The feasibility cost (tech development + each-build) should be adjusted to the reactor size through the function shown in Equation 2.27. P is the reactor power, P_{ref} is the reference reactor power (FHR: 100, MSR: 200 MWe) and F_{exp} the feasibility size exponent.

$$feasibility\ cost = \frac{(P(MWe))^{F_{exp}}}{(P_{ref}(MWe))^{F_{exp}}} \tag{2.27}$$

2.8 Cost Setting

It is very difficult to estimate the cost components of the safety, proliferation and feasibility costs. No adequately precise, valid and justifiable means to directly assign a cost to issues such as 'the safety mitigation costs resulting from the salt vapour pressure' (safety P_{vap}) or 'the proliferation mitigation costs resulting from the production of tritium' could be though of. As a result, a completely different method was devised that involves a more accurate estimation of a smaller number of dollar-value costs while making use of relative costs to determine the cost differences between configurations.

The costs of a reference MSR can be estimated to a sufficiently valid and justifiable degree based off of LWR costs, previous nuclear R & D projects and the limited MSRP experience. The reference MSR is meant to represent a reactor with 'average' or 'typical' parameters - eg. burnup, melting temperature, temperature reactivity feedback coefficients etc. What then remains to estimate configuration-dependent costs are the relative weights of each cost component and the input parameter cost effects (ex. how much lower cost a $T_{melt} = 500\ ^oC$ is compared to $T_{melt} = 550\ ^oC$). Though the relative weights and cost effects are also difficult to estimate, the magnitude of inaccuracy and the consequences of inaccuracy are much more limited.

To illustrate the greater reliability and accuracy of the cost setting method consider the following hypothetical: The true O & M safety P_{vap} cost should be \$1.1 million/year. Though one could easily select a value of \$0.2 or \$5.0 million. However, with cost setting, the relative cost (configuration divided by reference) is not extremely far off from equalling 1.0, since core geometries and salt compositions generally do not differ in extreme ways. It is also easier to estimate relative costs than absolute costs. It is fair to suggest that an estimate of the safety P_{vap} cost relative to other safety costs such as temperature coefficients and the ^{135}Xe poison fraction would be more accurate than an estimate of the absolute cost.

The specific cost setting process is as follows:

1. Input data: (If applicable) An attribute is added to the salt (s3.Tmelt = T_{melt}).[12]
2. Cost effect: Sometimes the input data is equal to the cost effect, otherwise it is calculated from input data through an attribute-specific function (T_{melt} effect = $600/(600 - T_{melt})$).
3. Cost setting: An attribute value equal to the expected value for the reference reactor is matched with a corresponding relative cost (eg. $T_{melt} = 500^{o}C$ corresponds to a T_{melt} relative cost of 5). Most cost setting involves at least some degree of arbitrary selection.
4. Create ratio: The ratio between the cost effect of the configuration under evaluation and the cost effect using the expected value for the reference reactor is determined.
5. Cost calculation: The cost calculation process described is conducted for the safety and proliferation cost categories as well as the 'each build' and 'technology development' costs of the feasibility category. Each cost effect ratio is multiplied by its corresponding relative cost to arrive at a value that represents the relative configuration-specific cost. All of the relative costs are simply summed together as well, to produce a value that represents the relative reference configuration cost. The configuration-specific cost is thus the ratio of these two values subject to an exponent, 'E', as per Equation 2.29. The cost of each category of the configuration under evaluation (Config. Category Cost) is thus the ratio multiplied by the category cost for the reference configuration.

$$\begin{aligned} \text{Rel. Config. Cost} &= \frac{(T_m \text{ effect})_i}{(T_m \text{ effect})_{\text{ref.}}} \cdot (T_m \text{ rel. cost}) + \frac{(P_{vap} \text{ effect})_i}{(P_{vap} \text{ effect})_{\text{ref.}}} \cdot (P_{vap} \text{ rel. cost}) + \ldots + \text{etc.} \\ \text{Rel. Ref. Cost} &= (T_m \text{ rel. cost}) + (P_{vap} \text{ rel. cost}) + \ldots + \text{etc.} \end{aligned} \tag{2.28}$$

$$\text{Config. Category Cost} = \left(\frac{\text{Rel. Config. Cost}}{\text{Rel. Ref. Cost}}\right)^{E} \cdot \text{Ref. Category Cost} \tag{2.29}$$

2.8.1 Example

This example involves 5 of the cost components that make up the safety cost category, T_{melt}, P_{vap}, ^{6}Li, Xe Poison Fraction (PF) and TC, and aims to calculate the O & M safety cost of a hypothetical configuration. The input data associated with the configuration under evaluation is 500 ^{o}C, 2 mmHg, 99.998% Li with [Li] = $0.010 \cdot 10^{24}$, a PF of 0.02 and a TC of -2.1 pcm/^{o}C. The reference reactor input data is 550 ^{o}C, 10 mmHg, 99.995% Li with [Li] = $0.013 \cdot 10^{24}$, a PF of 0.05 and a TC of -1.5 pcm/^{o}C. The configuration-specific ratios

[12] In the evaluation program each salt is treated as an object (eg. s3) and attributes, such as T_{melt} representing the melting temperature (T_{melt}) are properties.

are as follows[13]:

$$\begin{aligned} R_{T_{melt}} &= \frac{1/(600-500)}{1/(600-550)} \\ R_{P_{vap}} &= \frac{2}{10} \\ R_{^6Li} &= \frac{(1-0.99998)(0.01)}{(1-0.99995)(0.013)} \\ R_{PF_{Xe}} &= \frac{0.02}{0.05} \\ R_{TC} &= \frac{-1/-2.1}{-1/-1.5} \end{aligned} \tag{2.30}$$

The relative costs that relate to the safety calculation ($Crel_x$) are 1.6, 1.5, 3.7, 4 and 9 for each of the cost components respectively. These values are used in calculating the total relative configuration-specific cost and the summation of all values is equal to the relative reference cost, $Crel_{ref.}$ ('Rel. Ref. Cost' in Equation 2.29). The safety O & M cost of the reference reactor design, $C_{ref.}$, is \$4.33/MWh and no exponents are used in this example. There is now enough information to calculate the configuration specific cost ($C_{config.}$):

$$\begin{aligned} C_{config.} &= C_{ref.}\frac{R_{T_{melt}}Crel_{T_{melt}} + R_{P_{vap}}Crel_{P_{vap}} + R_{^6Li}Crel_{^6Li} + R_{PF_{Xe}}Crel_{PF_{Xe}} + R_{TC}Crel_{TC}}{Crel_{ref}} \\ C_{config.} &= (4.33)\frac{(2.0\cdot 1.6 + 0.2\cdot 1.5 + 0.3\cdot 3.7 + 0.4\cdot 4 + 0.7\cdot 9)}{1.6+1.5+3.7+4+9} \\ C_{config.} &= \$2.74 \end{aligned} \tag{2.31}$$

2.9 Further Discussion

2.9.1 Cost of Detritation; A Safety and Proliferation Cost

No study that specifically estimates the cost of MSR tritium removal was found and the detritiation cost cannot be derived from that of LWRs or CANDUs because those designs must only detritiate water not molten (fuel) salt. The Molten Salt Reactor Program (MSRP) was found to have provided a categorical breakdown of a reactor salt processing system in a report from 1965 on the costs of a 1000 MWe MSR [101]. The planned processing system was far broader in scope than just tritium removal and the literature suggests that tritium removal/management was less of a concern at the time. The report estimates the capital cost of a processing plant to be \$25 million (\$218 million, 2021) and the cost of the 'reactor plant' (nuclear island minus some civil work) to be \$51 million (\$445 million, 2021). Since a tritium removal system would be much less complex then a full processing plant, it is assumed that the cost of a tritium removal system would be $\approx$ 1/10th the cost of a processing plant, or \$20 million for a reference size of 1000 MWe. And the

[13] the numerator is configuration-specific, the denominator is for the reference configuration.

proposed function that determines the capital cost for a reactor of any size is shown in Equation 2.32.

$$C_{detrit} = 10^6 \cdot \frac{(2 + 18 \cdot (R \cdot P)^{0.75}}{P_{ref.}^{0.75}} \tag{2.32}$$

C_{detrit} is the detritiation capital cost, R is the tritium generation rate relative to the reference thermal flux and reference ^{6}Li concentration and P is the electric power in MW. The selection of the proposed constants (values '2', '18' and the exponent '0.75') in Equation 2.32 is guided by the MSRP cost report [101], but they were chosen somewhat intuitively. The reference tritium generation is from a reactor with a Li concentration of 0.025 $10^{24}/cm^3$ ($[Li_{ref}]$) with Li enriched to 99.995%. ($[^6Li_{ref}]$) These reference values along with the reference flux, ($\phi_{ref.}$) as determined from a selected reference configuration, are used to solve for the relative ^6Li production rate:

$$= \phi_{ref.} \cdot \frac{[Li_i][^6Li_{optimum}]}{[Li_{ref}] \cdot [^6Li_{ref}]} \tag{2.33}$$

The algorithm tests a range of $[^6Li]$ in an effort to determine the lowest cost; $[^6Li_{optimum}]$ represents the enrichment level that results in the lowest overall cost considering the costs of enrichment, tritium removal and the neutron poisoning effect. The O & M cost of a detritiation system is determined by assuming a detritiation cost per MWh for the reference ^6Li production ($1.00/MWh is suggested) and adjusting it to the actual ^6Li production.

2.9.2 Input Parameter Summary

The cost estimation framework has been developed to serve as a basis for estimating the cost over a range of configurations for MSR design types. The value of the cost functions at each configuration is highly dependent on the selection of the unknown input constants (inputs or input parameters). It is recognized that many of the framework inputs have considerable uncertainty and/or require intuitive selection, yet the inputs are necessary to implement the framework. Input uncertainties are significant and therefore uncertainty analysis is implemented; it is extended to sensitivity analysis to determine the effect of individual inputs.

Inputs that are common among all MSR types are listed in Table 2.14 and those specific to a particular MSR type are not included. Some inputs that are not discussed in this framework and involve relatively straightforward cost calculations are also listed. The input parameters listed are either elaborated upon in footnotes, self-explanatory or discussed in a cost category section. These input parameters attempt to capture all major costs that are a function of configuration and thus necessary for a design optimization.

2.10 Cost Estimation Framework Conclusion

A MSR power plant has never been constructed and only two relatively low-power test reactors have, the ARE and MSRE, more than 50 years ago. Thus there is little relevant and applicable operational cost data, so cost estimates must largely be extrapolated from data on construction and operation of LWRs. A comprehensive framework for comparing MSR design configurations on a cost basis has been provided in this

Table 2.14: Cost estimation framework input parameters. These values cannot be determined precisely; where possible, empirical sources or analytical processes have and should be employed. Nonetheless, they largely have been/must be determined intuitively. Proposed values are provided for many of the inputs throughout the cost category sections or can be found in the reactor-specific chapter (Tables 5.10/5.11 and 6.12/6.13).

Associated category	Input parameter	Notes
Direct	pumping multiplier [a]	
	cost of the heat exchanger [b]	
	salt component prices (^{7}Li, Be, Zr)	
	separative work price	
	U_3O_8 price	
Safety	cost overrun due to safety	as in Table 2.5
	share capital costs due to safety	as in Table 2.5
	reference OM safety cost	see Section 2.5.2
	reference insurance cost per MWth [c]	see Sections 2.5.2 through 2.5.4.
	insurance premium	see Section 2.5.4
	CDF	see Section 2.5.4
Proliferation	cost overrun due to proliferation	see Section 2.5.5
	share capital cost due to proliferation	see Section 2.5.5
	reference OM proliferation cost	see Section 2.5.2
	enrichment cost at a reference concentration [d]	
	Pu cost at reference concentration [e]	
Modularity	reference share regular cost modularity	see Section 2.6
Feasibility	reference technological development cost	see Section 2.7
	share tech. development cost born each reactor	see Section 2.7
	feasibility splits: T_{melt}, P_{vap} etc.	see Section 2.7
Safety and proliferation	detritation constants 1, 2 and 3	see Section 2.9.1
Safety	safety cost exponent	see Section 2.5.1
Proliferation	proliferation cost exponent	see Section 2.5.5
Modularity	regular cost moderation exponent	see Section 2.6
Feasibility	feasibility size exponent	see Section 2.7

[a] The factor by which the cost of pumping coolant salt through the core and heat exchanger is multiplied by. It is greater than 1.0 to account for the capital cost of the pump.

[b] per-unit of primary tube length

[c] expected risk at reference CDF

[d] The cost of other levels of enrichment is estimated relative to this input.

[e] The cost of other Pu concentrations is estimated relative to this input.

study. The objective is to develop useful metrics on which to evaluate designs across a range of potential geometries and salt compositions.

The cost estimation framework is structured around six cost categories: direct cost, waste, safety, proliferation, modularity and feasibility. Direct costs are comparatively straightforward but creative methods are required for quantifying the other costs. The waste cost is derived from total and categorically-broken-down costs of LWR high-level waste storage. The safety cost takes a similar, but more complex approach. Capital and Operation and Maintenance (O & M) costs are treated separately, though, for both costs, the cost setting method is employed to determine the cost of each configuration relative to the reference. The cost of insurance is treated separately and is a function of CDF. The proliferation resistance costs are calculated in the same way as safety costs. Modularity, of thermal-spectrum MSRs, is just a function of size since the power densities are similar among designs. Finally, feasibility, or technology difficultly/complexity, is split into pre-detailed design R & D costs and costs specific to each build. A total cost for the required R & D of a reference MSR must be assumed and the feasibility cost is relative to it.

Though numerous examinations of MSR designs have been conducted, these examinations have a small number of variables, do not account for many relevant costs and/or fail to account for interrelated effects. The cost estimation framework provides a structure to quantify the effect that a wide range of design parameters have on all types of costs, not just those with a straightforward calculation. It is centred around comparing specific configurations to a reference MSR design rather than assigning each cost component a cost directly. The key development that allows for all these costs to be determined and compared is the 'cost setting' method. The application of this framework to the evaluation of MSR design types would contribute to the general understanding of MSRs and assist informing decision makers on how to efficiently direct research resources. It may prove particularly useful in the conceptual design phase when the salt composition and general geometric layout are in flux. This dissertation applies the framework to both a salt-cooled (Chapter 5) and graphite-moderated, circulating-fuel (Chapter 6) design type.

Chapter 3

Salt Considerations

3.1 Introduction

Molten salt literature has a strong focus on FLiBe (the $^7LiF - BeF_2$ system), since the MSRE used a FLiBe salt and FHR designs have selected FLiBe coolant. Despite the favourable neutronics properties, enriched lithium and beryllium are expensive, beryllium is toxic, and both generate non-insignificant amounts of tritium.[1] These issues provide a motivation to explore alternative salts. This chapter explores several considerations, not captured by neutronics simulations that are relevant to the selection of reactor carrier/coolant salts. Corrosion is an important consideration, however a theoretical discussion is covered in Appendix D.2 and the generation of cost estimates for the evaluation program is discussed in Section 2.1.1.

This chapter begins with a discussion of the early history of salt selection through the ORNL MSRP and lists select coolant/carrier salt compositions of past and present MSR design concepts. Following that is a discussion of the process used to select salts to investigate in the FHR and MSR evaluations. The next section discusses the optimization that is used to determine the heat exchanger size and pumping power required for each evaluated salt composition. Then the risks, potential removal mechanisms, etc. of Tritium are discussed along with the [^{7}Li] enrichment optimization scheme employed. Other salt composition related considerations that are independent of reactor physics are also discussed.

3.1.1 MSRP

The ARE and especially the subsequent ORNL Molten Salt Reactor Program provide the most comprehensive source of molten salt research for reactor applications. The rationale behind the selection of the ARE carrier salt, $53.09NaF$–$40.73ZrF_4$–$6.18UF_4$ is unclear as it is difficult to find documentation from the design phase in the late 1940's and early 1950's; the reactor when critical November 3, 1954. It is known that this salt has quite favourable nuclear properties, though it has a relatively high melting point of 538 C [103].

[1] Tritium: maximum concentration at ground level directly downwind was estimated to be a factor of ten below the AEC's maximum permissible concentration for uncontrolled areas [102]

Table 3.1: Salt composition of various past and present MSR design concepts covered in Section 1.6.7.

Reactor	Developer	Reactor type	Carrier Salt
MSRE	ORNL MSRP	Circulating-fuel	FLiBe
MSBR	ORNL MSRP	Circulating-fuel	FLiBe
MSTW	Seaborg Technologies	Circulating-fuel	Na-actinide Fluoride
ThorCon	Martingale	Circulating-fuel	$NaF - BeF_2$
TransAtomic	TransAtomic Power	Circulating-fuel	LiF
Fuji	ITMSF	Circulating-fuel	FLiBe
TMSR-LF2	CAS, China	Circulating-fuel	FLiBe
Mark-1 FHR	Universities, US DOE Funded	FHR	$2^7LiF - BeF_2$
TMSR-SF2	CAS, China	FHR	FLiBe
SSR-U	Moltex Energy	Stable-salt	$33NaF - 30RbF - 37UF_4$ (Carrier) $39ZrF_4 - 1ZrF_2 - 60NaF$ (Coolant)
SSR-W	Moltex Energy	Stable-salt	NaCl (Carrier) Na, K, Zr, Fluoride (Coolant)

Soon after the operation of the ARE, research was conducted on salts other than the $NaF - ZrF_4$ system. The 1957 paper, *Preliminary Study of Molten Salt Power Reactors* [104] indicates that by that point sufficient research had been conducted on the $^7LiF - BeF_2$ system to deem it 'attractive as a fuel carrier.' This is in contrast to the $NaF - ZrF_4$ system which was found to have several serious disadvantages, including an 'intolerably high' epithermal/intermediate neutron energy capture cross section and a tendancy of ZrF_4 to evaporate and crystallize on exposed surfaces. Subsequent ORNL MSRE reports study the $^7LiF - BeF_2$ system, thus suggesting that it had been effectively selected as the carrier salt at this point.

The same 1957 paper [104] finds that the $^7LiF - BeF_2$ system has 'satisfactorily low' capture cross sections, attractive physical properties (heat transfer properties, melting temperature and viscosity), adequate solubility and a 'low' corrosion rate with Inconel. Subsequent research generally agrees with these conclusions particularly for FLiBe. Ultimately the cost of ^{7}Li, Be toxicity and tritium production issues were not enough to reconsider using the $^7LiF - BeF_2$ system over the life of the MSRP.

3.1.2 State of Salt Selection, Table 3.1

Table 3.1 lists the selected coolant/carrier salt of several past and present MSR designs concepts. There is considerable diversity in the compositions, indicating that it is not immediately clear which composition(s) are most suitable. Thus the salt composition is expected to be an important, vast area for potential research and optimization.

3.2 Salt Composition Selection

3.2.1 Salt Component Selection

The salt composition, $LiF - BeF_2 - (UF_4)$, or FLiBe, is well studied and has been proposed in many designs so it is included in both the FHR and MSR analyses. Other viable salts to analyse are not as obvious. Process of elimination, was largely employed to arrive at alternative salt compositions. A similar process was used by Grimes [105] in the selection of salt compositions for the MSRE and for subsequent MSR designs. Upon inspection only alkali, alkali earth, transition metals and Al may contain viable cation elements and only reactive non-metals are viable anions. The key selection considerations include:

- (Thermal) neutron capture cross section: most elements are eliminated this way. The radiative capture cross section (N,G) of candidate salt constituent elements is shown in Table 3.2. Light elements tend to have a sufficiently low enough cross section. The alkali metals Na and Rb and the alkali earth metals Be through Ca have adequately low thermal cross sections. Among transition metals only Zirconium is a viable option. K, has an absorption cross section than makes the element a marginally viable candidate.
- Melting point: Most fluoride salt mixes have a sufficiently low melting temperature, (T_{melt}) but the melting points of binary salts (LiF, NaF etc.) are too high. Notably the T_{melt} of Al, Mg and Ca fluoride mixtures are too high as well.
- Irradiation, chemical and/or thermal stability: most polyatomic ions are eliminated as viable salt components.
- Corrosion: many remaining stable polyatomic ions are eliminated, including, S_2^{2-}, OH^-, Si^{4-}.
- Reactivity with air/water: many remaining stable, single-element anions are eliminated, including, O_2^{2-}, N^{3-}, Si^{4-}, P^{3-}.
- Heat transport qualities (density, viscosity, heat capacity, thermal conductivity)

No common polyatomic anions were found to have adequate chemical stability at high temperature, have a low enough melting point, be non-reactive with air/water and be sufficiently non-corrosive. Natural Cl^- has a prohibitively high absorption cross section and enriched $^{37}Cl^-$ is more expensive with little benefit over F^-, so only flouride salts are considered.

Since enriched Li has been produced and has otherwise very favourable properties, it is included. In the FHR the only other alkali metal examined is Na, K has too high of a capture cross section. In the MSR both Na and K are examined. The MSR design allows for components with a higher capture cross section because the share of core volume taken up by the carrier salt can be considerably lower (up to ≈10x). Though Rubidium is otherwise viable it is left out of both the FHR and MSR examinations because it does not have any clear advantages over Na, is expensive (≈$100/g) and the total market is small, consumption in the United States is approximately 2000 kg/year [106]. Among alkali earth metals only Be is viable; Mg

Table 3.2: Radiative capture cross section (N,G) of candidate salt composition constituent elements. Data from [107] and [108]

Element/isotope	0.0705 eV [barns]
^{6}Li	562[a]
^{7}Li	0.027
^{23}Na	0.314
K	1.26
^{85}Rb (72.2%)	0.296
^{87}Rb (27.8%)	0.720
^{9}Be	0.0060
Mg	0.379
Ca	0.258
^{27}Al	0.140
Zr	0.111
S	0.312
^{31}P	0.101
Si	0.0963

[a] Production of a triton (n,t)

containing salts that are otherwise viable, have high T_{melt}s that preclude further examination.

In general, 'lighter' (low-Z) salts tend to exhibit better heat transfer and nuclear performance metrics [49]. Section 3.3 provides a complete pumping power and heat exchanger sizing calculation for each salt.

3.2.2 Molar Composition Selection

Molten fluoride salts mixes tend to easily form homogeneous mixes where at least one mixing ratio is a eutectic point. The mixing generally results in a considerable melting point depression from the constituent simple salts; often $> 150^{o}C$ in a binary mix and $\approx 40^{o}C$ from a binary to ternary mix.

A eutectic mixture is a homogeneous mixture of two or more substances that melts/solidifies at a single temperature; neither of the substances solidify and crystallize out of solution before the other. The temperature at which crystallization occurs is called the eutectic temperature.

The formation of eutectics is driven by Gibbs Free Energy (G), the maximum thermodynamic potential work at constant temperature and pressure. Systems tend toward the most thermodynamically stable arrangement which occurs where G is at a minimum. The Gibbs free energy of formation is negative for all candidate salt mixes in the liquid phase at the eutectic point(s). Therefore, there is a thermodynamic tendency favouring homogeneous mixing of the salt in the liquid phase. Since the liquid phase is made more stable (by a lower G) than the liquid phase of either simple salt there must exist a temperature range lower than the T_{melt} of either simple salt where a liquid mixture exists.

Binary mixes generally have few or even one eutectic point while ternary mixes have several. The aim is to select salt mixes for analysis at eutectic points because of the low melting temperatures and because property data is generally produced at these points. For ternary salts, the eutectic composition is selected largely on the basis of choices from previous authors as well as capital cost, neutronics, safety, among other considerations. Compositions with high mole fractions of Li and/or Be tended to be avoided in an effort to highlight the difference with $LiF - BeF_2$.

3.3 Heat Exchanger Pumping Power and Sizing Requirements

3.3.1 Introduction

It is not immediately clear what the Pressure Drop (PD) of salt through the heat exchanger is nor the required size of the heat exchanger. These costs are non-trivial and differ with salt composition. At operating temperature ($\approx$700 ^oC for molten salts and 300 ^oC for water) the viscosity of molten salts is generally between 1 and 10 cP and the heat capacity generally between 1.0 and 2.4 kJ/kgK, while, for reference, light water is only 0.086 cP, and 4.2 kg/kgK, see Appendix C. Thus the required pumping power through the heat exchanger is a cost worthy of consideration. The cost associated with the size of the heat exchanger can be substantial because potentially expensive corrosion-resistant metal is required and the heat exchanger adds to the volume of salt in the primary loop. Ergo, it is imperative to calculate the pumping power required and corresponding heat exchanger size.

A PD calculation through the reactor core in not covered in this analysis. The PD through a FHR core is taken from literature and adjusted according the relative flow velocity and salt composition of the particular configuration, more on this in Section 5.2.1. The thermal-spectrum circulating-fuel reactor core PD analysis is more complex and is in the circulating-fuel reactor chapter, specifically Section 6.2.3.

The purpose of this analysis is to determine the pumping power and heat exchanger size required for candidate primary circuit salts across a range of flow velocities. The heat exchanger size refers specifically to the amount of primary fluid piping; for a given tube diameter this can be measured in pipe length, area (circumference $\cdot$ length) or volume. Going forward this analysis opts to use area. In following with the rest of this thesis, the evaluation of the reactor has been limited to the core and primary salt selection; not the balance of plant. Therefore the secondary salt composition and heat exchanger layout are not optimized. The objective is to determine the cost differences among the salt compositions, not optimize the heat exchanger design generally.

3.3.2 Setup

All MSRs require a heat exchanger to transfer heat from the primary circuit salt to the secondary fluid. The design of a primary heat exchanger necessarily depends only on the power generated, temperature differ-

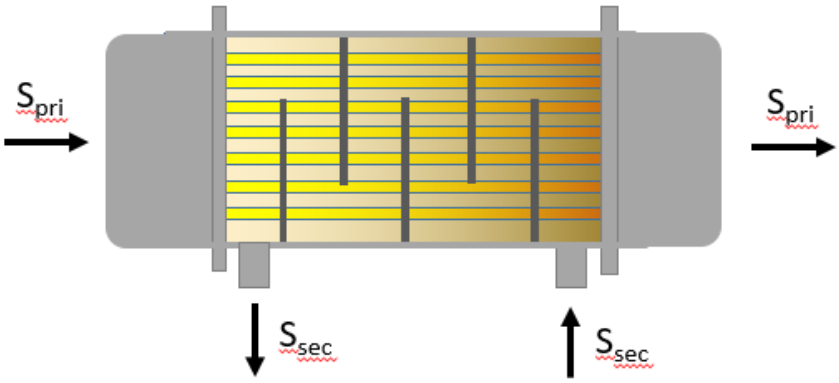

Figure 3.1: Schematic of the assumed heat exchanger design used for the calculation of pressure drop and heat exchanger size across different salt compositions and primary salt velocities. Though only five baffles are shown, the calculations are performed as if there were infinitely many baffles.

ences and properties of the salt - not the geometry of the core, the neutron-spectrum nor any other core design parameter. The heat exchanger design is assumed to be of a baffled, cross-flow, shell-and-tube geometry. A schematic of the design is shown in Figure 3.1; notably the temperature difference between the primary and secondary salt is assumed to be smooth along the direction of the primary salt flow - as if the heat exchanger has infinitely many baffles.[2] The secondary salt is considered beyond the scope of optimization thus one salt, $LiF-NaF-KF$ (46.5-11.5-42), has been used for all calculations.[3] The salt is circulated through the heat exchanger once, so the temperature difference of the primary salt is equal to that of the core.

The results are strongly dependant on property data (k, Cp, ρ and μ) and the Nusselt number (Nu) correlations selected. The salts examined are those from the two MSR design classes; the salts without uranium from the FHR and with uranium from the circulating fuel MSR. The uranium molar concentration($[UF_4]$) is a variable and is generally analysed at a $[UF_4]$ of 0% and 23%, the salt mixes are listed below. The concentrations analysed are limited by available data and cover the range of plausible $[UF_4]$. Interpolation is used to determine the pumping power and heat exchanger area required of intermediate $[UF_4]$.

- $^7LiF-BeF_2$ (67-33-0)
- $NaF-BeF_2$ (57-43)
- $NaF-ZrF_4$ (FHR: 59.5-40.5, MSR: 51-49)
- $LiF-NaF-BeF_2$ (15-58-27)
- $LiF-NaF-ZrF_4$ (26-37-37)
- $NaF-KF$ (40-60)

[2] This is design was selected over other designs that are easier to analyse, such as a counter-flow, because this design is more common and realistic.

[3] $LiF-NaF-KF$ is selected as the secondary salt because it has been 'extensive studied' [109], has desirable thermal and chemical properties and is not composed of any expensive or poisonous cation elements.

- $LiF - BeF_2 - UF_4$ (74-2-23)
- $LiF - UF_4$ (77-33)
- $LiF - NaF - BeF_2 - UF_4$ (23-39-15-23)
- $NaF - BeF_2 - UF_4$ (57-20-23)
- $NaF - ZrF_4 - UF_4$ (51-26-23)
- $NaF - KF - UF_4$ (45-32-23)

Assumptions:

- Conductive heat transfer resistance within metal piping is insignificant.
- Fully developed flow throughout.
- Shell and Tube heat exchanger with a staggered, hexagonal arrangement of primary fluid tubes.
- The secondary fluid contacts the tubes at a right angle.
- Incoming primary fluid is of uniform temperature.
- Uniform Reynolds and Prandlt numbers are applied across the length of the heat exchanger. The density and viscosity used in calculating the Reynolds and Prandlt numbers are determined using the average temperature of the salt.
- Log mean temperature difference approach is taken to solve for the temperature difference between the primary and secondary salts.
- Since insufficient data was found, the thermal conductivity and heat capacity at 700 ^oC was applied throughout.
- The inlet temperature of the secondary salt is equal to $T_{sec,out} - 0.75 \cdot \Delta T_{core}$.
- The cross-sectional area through which the secondary salt flows is on average 3x the cross-sectional area of the primary salt tubes.[4]

Key structural dimensions:

- Tube Diameter: 0.025 m
- Pitch distance: 3·(Tube Radius)

Procedure:

The setup parameters are either input variables that are generated elsewhere or must be assumed. The average salt temperature is 650 ^{o}C and the primary salt inlet and outlet temperatures are determined from this assumption and the core temperature difference. The secondary salt outlet temperature is set at 920 K in all instances in order to prevent for the need to adjust for thermal efficiency. And the inlet temperature is 920 K - ($T_{sec,out} - 0.75 \cdot \Delta T_{core}$). The heat transfer from one salt to the other must balance, as in Equation 3.1

[4]The cross-sectional area (XA) is 2.50x the primary salt XA area in the 3·R arrangement, it is however assumed that a portion of the salt will flow along the edge in the absence of primary tubes.

so the velocity (u) of the secondary salt can be solved for. The ρ, C_p and ΔT of the primary and secondary salts are similar (and salts with a relatively high density tend to have a relatively low heat capacity). The approximately value of u_{sec}/u_{pri} is 0.44, see Equation 3.2. A_{pri} is the cross sectional area of the primary salt tubes and A_{sec} the cross sectional area through which the secondary salt flows.

$$\begin{aligned} \dot{m}_{pri} C_{p,pri} \Delta T_{pri} &= \dot{m}_{sec} C_{p,sec} \Delta T_{sec} \\ (\rho u A C_p \Delta T)_{pri} &= (\rho u A C_p \Delta T)_{sec} \end{aligned} \tag{3.1}$$

$$. \frac{u_{sec}}{u_{pri}} \approx \frac{A_{pri} \Delta T_{pri}}{A_{sec} \Delta T_{sec}} \approx \frac{1}{3} \frac{1}{0.75} \approx 0.44 \tag{3.2}$$

The calculation of the primary salt pressure drop and heat exchanger area is repeated across a range of average core temperature differences (ΔT) and primary fluid flow velocities. For instance the FHR uses data from ΔT's of 62, 71, 80 and 89 ^{o}C. The velocities tested are 0.8, 1.0, 1.3, 1.6, 2.0, 2.5, 3.0, 3.5, 4.0, 4.5, and 5 m/s. Results for intermediate temperatures/velocities are interpolated. The procedure to arrive at the PD and heat exchanger area is as follows:

1. Use the log-mean-temperature-difference (LMTD) approach to arrive at the log-mean ΔT between the primary and secondary salt, as per Equation 3.3.
2. Find the primary salt heat transfer coefficient (h_{pri}) using the Nusselt number, $\overline{Nu} = \bar{h}L/k$. Equation 3.4 shows the Nu for the primary fluid with a Reynolds number (Re) > 12000 and Equation 3.5 for the primary fluid with a Re $3000 - 12000$.
3. Find the h_{sec} using the Nusselt number correlation shown in Equation 3.6.
4. Calculate the heat transfer surface area, $A_s = \frac{P}{h_{eff} \Delta T}$, where P is the power, ΔT is the LMTD and h_{eff} is the combined resistance, $h_{eff} = \frac{1}{\frac{1}{h_{pri}} + \frac{1}{h_{sec}}}$.
5. From A_s and u_{pri} calculate the number of tubes (n_t), tube length (L) and total material area. n_t is calculated as per Equation 3.7, where $A_{xs,tot}$ is the cross-sectional area of all tubes and $A_{xs,tube}$ the cross-sectional area of one tube. Tube length is, $L = \frac{A_s}{n_t D_t \pi}$ and total material area is just, $M_A = n_t \cdot L \cdot C$, where C is the circumference. The heat exchanger geometry is implicitly adjusted to accommodate n_t and L; a low n_t, high L heat exchanger would be relatively long with a small circumference.
6. Calculate the Head Loss (HL) as per Equation 3.8. The HL is converted to PD as required in the optimization.

$$\Delta T = \frac{(T_{pri,in} - T_{sec,out}) - (T_{pri,out} - T_{sec,in})}{log(\frac{(T_{pri,in} - T_{sec,out})}{T_{pri,out} - T_{sec,in}})} \tag{3.3}$$

$$Nu_{pri} = 0.023 \cdot Re_{pri}^{0.8} \cdot Pr_{pri}^{0.3} \tag{3.4}$$

$$f = (0.79 \cdot ln(Re_{pri}) - 1.64)^{-2},\ into$$
$$Nu_{pri} = \frac{(f/8) \cdot (Re_{pri} - 1000) \cdot Pr_{pri}}{(1 + 12.7 \cdot (f/8)^{0.5}) \cdot ((Pr_{pri}^{2/3}) - 1))} \tag{3.5}$$

$$Nu_{sec} = 0.35 \cdot (pitch_t \cdot pitch_l)^{0.2} \cdot Re_{sec}^{0.6} \cdot Pr_{sec}^{0.36} \tag{3.6}$$

$$n_t = \frac{A_{xs,tot}}{A_{xs,tube}} = \frac{\frac{\dot{m}}{u_{pri}\rho_{pri}}}{A_{xs,tube}} \tag{3.7}$$

$$ff(Re) : 0 = \frac{-1}{\sqrt{ff}} + 1.93 \cdot log(Re\sqrt{ff}) - 0.537$$
$$HL = ff(Re)(\frac{1}{2 \cdot 9.81})u^2 \frac{L}{D} \tag{3.8}$$

Property data

Density (ρ) and especially viscosity (μ) are strongly dependant on temperature and so each property was also calculated as a function of temperature. Thermal conductivity (k) and heat capacity (Cp) were taken as constant with a temperature at 700 C. If density values were unavailable, density was calculated using the additive principle with all data from [110]. The data sources for μ, k and Cp, came from various sources, see Table 3.3. Property data for other relevant salts can be found in Appendix C. The units of μ are cP, of k are W/mK and of Cp are kJ/kgK.

Optimization

The material area and pumping power requirements over a range of (average) velocities is useful as it provides the necessary data to optimize the salt velocity. This data set is generated for candidate salts at each of the tested core temperature differences.

In order to cost optimize the circulation velocity a value must be assigned to \$/m^2 of heat transfer surface area and \$/kW of pumping power required. These values are reactor evaluation specific and thus assigned in the respective evaluation. Although the material area cost does not include the cost of salt inside the primary salt tubing, the evaluation program accounts for that cost as well. The pumping cost is comprised of two parts: the direct cost of electrical power that could otherwise be sold and capital costs. The pumping power required, as shown in Figure 3.2, is the necessary mechanical power, or conceptually the power requirement of a 100% efficient pump. To account for pump inefficiency, secondary salt pumping power and the pump capital costs, the total pumping related costs are estimated to be approximately 2.0 times the cost of electricity of a 100% efficient pump. For example, assuming an electricity cost of \$0.10/kWh, a 30 year lifetime and a MARR of 7%, one kW of mechanical pumping requirements costs \$10,950/kW, and

Table 3.3: Properties of the salts analysed. Rao-Turnbull and Dulong-Petit are predicted values, not empirical measurements. In some cases the predicted values have a pattern of deviating from measured values, so the author's judgement is used to select an estimate instead. $LiF - BeF_2$ is only analysed at 0% UF_4 because the UF_4 is expected to displace only Be, thus intermediate compositions would be interpolated between $LiF - BeF_2$ and $LiF - UF_4$. The binary salt LiF has an infeasibly high melting point so it is only analysed at 20% UF_4.

Salt composition	Property	700^oC, 0% UF_4	Source	700^oC, 20% UF_4	Source
$LiF - BeF_2$	μ (cP)	5.5	[111], pg. 571	–	
(67-33)	k (W/m-K)	1.1	[112]	–	
	Cp (kJ/kGK)	2.24	[112]	–	
$NaF - BeF_2$	μ	7.0	[111], pg. 571	14.0	est.
(57-43)	k	2.4	[113]	1.5	est.
	Cp	2.2	[113], Pri: [114]	1.8	Dulong-Petit
$NaF - ZrF_4$	μ	4.57	[49] [115]	8.0	Based on [116]
(59.5-40.5)	k	0.49	[115]	0.35	est.
	Cp	1.2	[115]	1.0	est.
$LiF - NaF - BeF_2$	μ	1.8	[111]	3.7	est.
(15-58-27)	k	0.72	Rao-Turnbull	0.9	est.
	Cp	1.86	Dulong-Petit	1.05	est.
$LiF - NaF - ZrF_4$	μ		–		
(26-37-37)	k	1.0		–	
	Cp	1.27		–	
LiF	μ	–		10.1	[117]
(100)	k	–		1.0	[117]
	Cp	–		1.59	[117]
$NaF - KF$	μ	3.5	est.[a]	9.8	[116]
(40-60)	k	0.9	est.	1.5	Based on [118]
	Cp	1.3	Dulong-Petit	0.96	[119]

[a] same as $LiF - NaF - KF$, however it is nearly irrelevant because low $[UF_4]$ are infeasible with this salt.

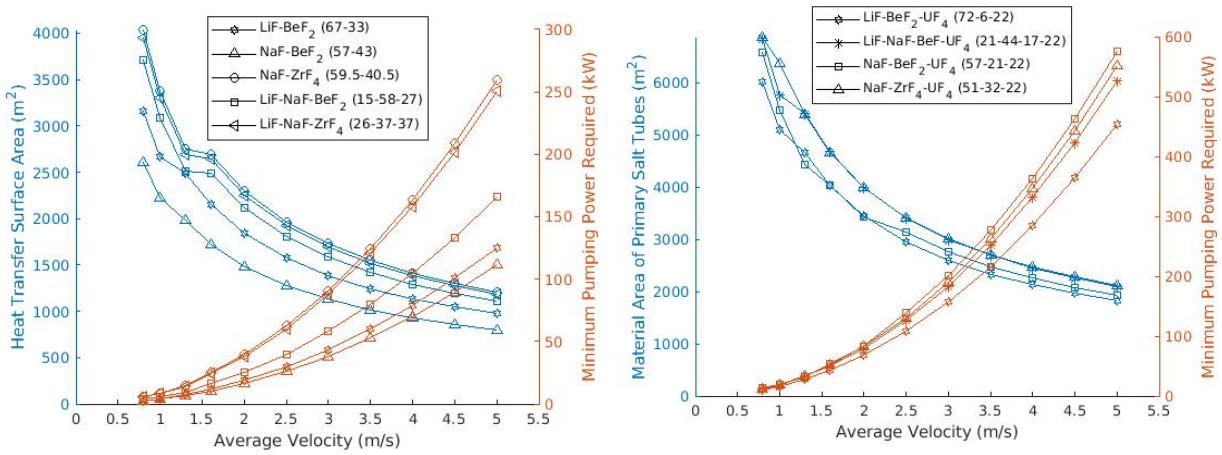

(a) FHR, 200 MWth, 80 ^oC average ΔT

(b) MSR, 400 MWth, 80 ^oC average ΔT

Figure 3.2: Heat transfer surface area and pumping power required in a molten salt heat exchanger for select salts of the two reactor classes. Several surface area curves are not smooth because the Reynold's number switches from below to above 12000, and hence the empirical correlation switches, at $\approx 1.5 m/s$.

$21,900/kW accounting for inefficiency and capital costs. The results of the velocity/cost optimization are shown for the FHR in Section 5.2.1 and for the MSR in Section 6.2.3.

3.4 Tritium

The production of tritium is an issue primarily because it poises an ongoing safety risk and it is caused by a parasitic neutron absorption. In the MSRE the salt contained less than 0.01% ^6Li, yet because of the large reaction cross section (522 barns at operating temperature) it was responsible for over 80% of ^6Li production [102].

The tritium generation rate is not only a function of the salt composition but also specific design features [120] and reactor geometry. Tritium production depends on the salt fraction in the core, the uranium enrichment level and the fuel concentration. Though the primary determinent of tritium production is fundamentally the neutron spectrum, as evidenced from MSRE experience: 'In operations with ^{233}U the fissile material concentration was much lower, the thermal neutron flux much higher and the fast neutron flux about the same as in the ^{235}U operation. These differences account for the changes in tritium production rates from ^{235}U to ^{233}U operation.' [102].

There are a variety of potential means to remove tritium, for instance, a carbon absorber bed, gas sparging or permeators. Tritium can also migrate through metals and thus escape through container walls or the heat exchanger [120]. However removal/migration is riddled with uncertainty. For instance, there are 'order-of-magnitude' uncertainties for maximum solubility in high temperature graphite [120] [121] and MSRE finding for tritium in radiator cooling air, an indirect measurement of the amount of tritium diffusing

through the walls of the heat exchanger are unreliable. It was expected that a 'substantial fraction' of tritium leaves the salt through heat exchanger [102]. The two means of reducing uncontrolled tritium release are to increase the efficiency of removal from the salt or create barriers (coatings) to slow tritium transport through the heat exchangers [120].

The two major deleterious effects of tritium generation are the decrease in k_{eff} and the safety concerns associated with the production of a mobile radioactive substance. A simple optimization of ^6Li (for applicable salts) is performed in Section 3.4.1 that balances the costs of ^7Li enrichment against the costs associated with k_{eff} and safety concerns.

3.4.1 ^{7}Li Enrichment Optimization

^{6}Li in the reactor is deleterious for two reasons:

1. Neutron poison: The ^{6}Li poisoning is small yet has a non-negligible impact on fuel requirements.
2. Tritium production: ^{6}Li atoms that absorb neutrons decay into a helium and a tritium atom.

Neutron poison: The ^{6}Li concentration ($[^6\text{Li}]$) decreases exponentially; depending on the reactor configuration, with a decay constant of approximately 90 days. The $[^6\text{Li}]$ can therefore, for evaluation purposes, be assumed to have an impact equal to that of the initial $[^6\text{Li}]$ for a period of 90 days; since $\int_0^{\inf} e^{-\lambda x} = 1/\lambda$. This assumption is used to calculate the $[^6\text{Li}]$ poison effect on the minimum required initial enrichment and the increased refuelling rate.

Though the reactor requires a higher enrichment than it otherwise would without ^{6}Li, the higher level of enrichment would also require less refuelling once the initial ^{6}Li is effectively burned out. Rather than calculate the cost of higher initial U enrichment and at the same time adjust for the reduced amount of fuel added while the ^{6}Li is being burned out, it is less error-prone and more straightforward to assume the same initial enrichment and calculate the cost of additional refuelling. A lower initial k_{eff} will require more refuelling during the burn out period.

To relate the ^{6}Li concentration of the salt initially in the loop to the Refuelling Rate (RR), the $\frac{dRR}{dk_{eff}}$ and $\frac{dk_{eff}}{d[^6Li]}$ must be determined. Multiple depletion simulations are run in an effort to determine $\frac{dRR}{dk_{eff}}$ for each salt across a range geometric configurations. $\frac{dk_{eff}}{d[^6Li]}$ is determined for a reactor with a reference configuration (e.g., salt = 1, mole % U = 0.1, FF = 0.094). Equation 3.9 is used to relate the ΔRR to the corresponding $\Delta[^6Li]$.

$$\Delta RR = \frac{dRR}{dk_{eff}} \cdot \frac{dk_{eff}}{d[^6Li]} \cdot ([^6Li]_i - [^6Li]_{ref}) \tag{3.9}$$

For the MSR, $\frac{dRR}{dk_{eff}}$ is explicitly solved for through neutronics simulations while for the FHR it is equal to $1/(n \cdot \Delta k_{eff}$ per step) where n is the number of fuel groups/zones. To find the cost of the ΔRR, multiply ΔRR by the cost of the added fuel and adjust for the salt loop volume (salt channels and heat exchanger) relative

to the reference configuration.

To calculate the impact of the added ^{6}Li to the total refuelling, use the same procedure but instead of adjusting with the salt loop volume adjust to the total volume of refuelled salt discounted to the present (again relative to the reference configuration).

Tritium production: The user must input a [Li] (cm^3) and enrichment associated with a particular cost. The [Li], enrichment, cost combination is for a reference thermal flux, heat exchanger size and FF. The cost resulting from the generation of tritium is equal to the input reference cost multiplied by [^{6}Li]/[^{6}Li$_{ref}$] and adjusted for thermal flux, heat ex. volume and FF. ([^{6}Li] = [Li][1-enrichment]). [5]

^{7}Li enrichment cost and enrichment optimization: A Li price at 99.995% Li is a program input. The assumed function relating [^{7}Li] to price is a square root. This function is selected because it simple and similar to the relationship between the enrichment and the number of stages in a Crown Ether Separation Design from Table 3.3.3 in a 2012 report, *Lithium Isotope Enrichment: Feasible Domestic Enrichment Alternatives* [122]; it is assumed that cost is roughly proportional to the number of stages:

$$\$_i = \$_{99.995} \cdot \frac{\sqrt{1-0.99995}}{\sqrt{1-[^7Li]_i}} \tag{3.10}$$

The program loops through a set of Li enrichment values and calculates the cost of neutron poisoning, tritium cost and cost to enrich Li. The three values are summed together for each level of enrichment and the Li enrichment value associated with the lowest cost is saved.

3.5 Additional Considerations

3.5.1 Costs and Supply of MSR Salt Component Materials

Evaluating the cost of candidate salts is made difficult by the cost variance with purity. It is generally unclear what level of purity is necessary and there is often no straightforward, quantitative means to cost-optimize the purity. The costs are specified at expected purity levels and though the costs are used to inform the selection of evaluation input parameter values they are not necessarily the estimated costs specified in this section.

Lithium-7: A small market currently exists for lithium-7, as lithium hydroxide is used in PWRs to stabilize pH levels and thereby reduce corrosion. In the USA this market is 400 kg of ^{7}Li per year [122].

It is effectively a requirement for lithium to be enriched ^7Li as the thermal cross section is only 0.045 barns while for ^6Li it's effectively[6] 940 barns at MSR operating temperatures (700 ^{o}C) [123]. The thermal cross-

[5] Only the initial amount of ^{6}Li is of consideration because the tritium removal/handling/monitoring systems must be designed for maximum tritium production. Relatively minor differences in added ^{6}Li through refuelling are effectively inconsequential.

[6] The thermal cross section is 522 b, but approximately 940 when integrating over the Maxwellian temperature distribution.

section is dominated by the reaction that produces tritium and a residual (an α particle).

^7Li has a natural isotopic abundance of approximately 92.5% atomic percent. The ^7Li concentration in PWRs must be 99.92% [122]. The latest FHR design assumes a 99.995% ^7Li concentration [120] and ^7Li used in the MSRE was enriched to at least 99.99% [124].

U.S. PWRs rely exclusively on Chinese markets for ^7Li. The Chinese produce ^7Li through an environmentally questionable method, the column exchange (COLEX) separation process [122]. A Mercury-lithium mix flows counter to aqueous lithium hydroxide in packed columns. ^6Li moves preferentially in the mercury phase and ^7Li to the aqueous solution [125]. The U.S. employed the COLEX process from 1955 to 1963. Because a significant amount of mercury is required and there exists many opportunities for leaks into the environment the EPA stopped and prohibited further use of the COLEX process in the U.S [122].

The cost of ^7LiF acquired for the MSRE was \$16.50/lb [124], the reactor went critical in 1965. The cost of the ^7Li alone (minus F and salt preparation cost) would be approximately \$54.00/lb^7Li ($\approx$ \$464.00/lb^7Li or 1022.00/kg^7Li 2021 USD). In 2010, the estimated price for enriched ^7Li from China was \$15,000/kg^7Li at a 99.99% [122].

It has been estimated that FHRs would require approximately 21-56 metric tonnes of ^7Li per GWe of capacity [122]. At current prices this is not economically feasible but would be at the \$1022.00/kg cost from the MSRE: it would cost \$21.5 mil to \$57.2 mil/GWe. A recent study [122] identifies two 'fairly well established' alternative enrichment technologies that have the potential to replace the COLEX process: atomic vapour laser isotope separation (AVLIS) and crown ether enrichment. It goes on to conclude that a crown ether facility could profitably produce ^7LiCl at \$83.33/kg, (\$546/kg^7Li) with a 20 ton/year throughput.

Sodium: Sodium is likely to be purchased as NaF. Standard high purity or industrial grade NaF may not be sufficiently pure but is only on the order of \$1100/ton.

Beryllium: The cost of BeF_2 acquired for the MSRE was \$5.70/lb [124] ($\approx$ \$52.08/lb or 114.70/kg 2021 USD). In the US from 2013 to 2017 the price of Beryllium metal has averaged \$512/kg and 200 tons was consumed in 2017 [126].

Zirconium: Normal commercial grades of zirconium compounds may contain from 1 to 3% hafnium as an impurity and would invoke a severe penalty in neutron economy if used in the MSRE [124]. Zirconium is widely used in the nuclear industry, it is the fuel cladding of choice for water cooled reactor and the principle constituent of CANDU pressure tubes. Reactor-grade zirconium contains $< 0.01\%$ hafnium [127] and ranges between \$40 and \$80 per kilogram [128] (\$48 and \$96, 2021 USD).

Chlorine: Chlorine would likely only be used in fast reactor and it is likely that natural chlorine be separated

Table 3.4: Estimated minimum cost of MSR salt components, $ USD 2021

Component	Estimated Minimum Cost	Grade/Purity
^{7}Li	\$550/kg	99.995% ^{7}Li
NaF	\$1100/ton	Industrial grade
Be	\$550/kg	unspecified
Zr	\$70/kg	Reactor grade
^{37}Cl	\$750/kg	$< 99\%$
F	\$820/ton	100% ^{19}F

to produce ^{37}Cl [129]. ^{35}Cl has an absorption cross-section of 44 barns at 0.025 eV while ^{37}Cl has cross-section of 0.43 at the same temperature [130]. The absorption cross sections are more similar and much smaller toward fast-spectrum neutrons but there is inevitably some thermalization in a fast reactor. There is no current, commercial chlorine enrichment operation, however based on costs of enriching D_2O and ^{7}Li, a rough estimate is \$750/kg (\$28.0/mole).

Fluorine: It impractical to purchase pure, gaseous fluorine, instead it is likely to be obtained as a salt compound bonded to the cations of the salt mix, as was done in the MSRP [124]. Thus it tricky to isolate the cost of fluorine alone. An estimate can be made from the section on sodium; industrial grade NaF is \$1100/ton.

3.5.2 Neutron Activation of Salt Constituents

All data is from the ENDF/B database [130] unless otherwise mentioned. Only the thermal and resonance-integrated capture cross-sections are provided because the magnitude of the cross-section is generally far greater than it is for fast neutrons. The only activation precursor of concern is residual ^{6}Li in enriched ^{7}Li (99.99% +) and a summary of the activation of the constituent elements is provided in Table 3.5.

Lithium: The majority of reactions that consume a neutron occur with ^{6}Li as the cross-section is far greater. Most often it absorbs a neutron to produce tritium and a residual α particle. Captures in ^{7}Li lead to ^{8}Li which quickly decays, first through $\beta-$ emission ($T_{1/2} = 0.84s$) to ^{8}Be then by α ($T_{1/2} = 8.2 \cdot 10^{-17}s$) to ^{4}He.

Sodium: The only naturally occurring, stable isotope is ^{23}Na, it has a cross section of 0.29 barns at 0.0839 eV and 0.31 barns integrated over the resonance region. Upon capture ^{24}Na decays through $\beta-$ emission ($T_{1/2} = 15hr$) to stable ^{24}Mg. Thus neutron activation is not a concern.

Potassium: Potassium has two naturally occurring, stable isotopes, ^{39}K (93.26%) and ^{41}K (6.73%) with respective thermal capture cross sections of 1.2 and 2.5 barns at 0.0839 eV. ^{40}K ($T_{1/2} = 1.25 \cdot 10^{9}$) makes up the remaining 0.01% of naturally occurring K. There is a 89.28% probability that it undergoes a $\beta-$ decay to stable ^{40}Ca and a 10.72% probability of $\beta+$ emission to stable ^{40}Ar. ^{40}K is the ground state daughter

Table 3.5: Activation products of salt components.

Component	Long-term activation product	$T_{1/2}$	σ^a activation (b)	σ^a total (b)
Li	Tritium	13.3 yr	0.05	0.075
Na	N/A	–	–	0.3
K	^{40}K	$1.25 \cdot 10^9$ yr	1.1	1.3
Be	^{10}Be	$1.5 \cdot 10^6$ yr	0.005	0.005
Zr	^{93}Zr	$1.6 \cdot 10^6$ yr	0.02	0.1
F	N/A	–	–	0.013

[a] Approximate cross-section under likely neutron temperature and isotopic composition.

product of a capture in ^{39}K. ^{42}K decays through $\beta-$ emission ($T_{1/2} = 12.4hr$) to stable ^{42}Ca.

Beryllium: The only naturally occurring, stable isotope is ^{9}Be, it has a capture cross section of 0.0055 barns at 0.0839 eV and 0.0038 barns integrated over the resonance region. Upon capture ^{10}Be decays through $\beta-$ emission ($T_{1/2} = 1.5 \cdot 10^6 yr$) to stable ^{10}B. Thus neutron activation is of minor concern.

Zirconium: There are five naturally occurring isotopes: ^{90}Zr at 51.45%, ^{91}Zr at 11.22%, ^{92}Zr at 17.15%, ^{94}Zr at 17.38% and ^{96}Zr at 2.80%. The weighted-average cross-section at 0.0839 eV is 0.104 barns. Neutrons captures in ^{92}Zr, ^{94}Zr and ^{96}Zr produce ^{93}Zr, ^{95}Zr and ^{97}Zr that decay by $\beta-$ emission with respective half-lives of $T_{1/2} = 1.6 \cdot 10^6 yr$, $T_{1/2} = 64$day and $T_{1/2} = 17$hr. The decay product of ^{93}Zr, ^{93}Nb, is stable, though ^{95}Nb and ^{97}Nb further decay by $\beta-$ emission, with respective half lives of $T_{1/2} = 35$day and $T_{1/2} = 72$min to stable ^{95}Mo and ^{97}Mo. Only the production of ^{93}Zr is of concern in regard to the production of long-term radioactive waste. A ^{93}Zr isotope is produced in 20.5% of captures at 0.0839 eV.

Fluorine: The capture cross-section of ^{19}F, the only stable isotope, is 0.0096 barns at 0.025 eV, 0.0053 barns at 0.084 eV (700 ^oC) and 0.02 barns over the integrated resonance region. ^{20}F decays with a half-life of 11 seconds into stable ^{20}Ne.

3.5.3 Thermal Scattering

When neutrons are moderated to thermal energy in thermal reactors, the chemical binding of the scattering nuclei in the moderator material affects the scatting cross section by way of the energy and angular distribution of secondary neutrons.

Thermal scattering libraries (TSL) are widely available for H_2O, D_2O and graphite. However the TSL for key light element constituents, Li, Be and F are generally not readily available and if they were it would be as a pure substance. For instance the only materials which include either Li, Be or F in the ENDF/B-VIII library are Be-metal and BeO [107]. The factors contributing to the difference in TSL from the free-gas model, such as rotational and vibrational modes are specific to the chemical structure of the material rather

than just the material components. Molten fuel salts vary considerably in composition and chemical structure and thus cannot be treated as having a single thermal scatting library.

Though not available in distributed libraries (ENDF, JENDL etc.) a small number of published studies have investigated the TSL of molten salts or molten salt constiuent materials. Mei et al., generated the TSL for FLiBe ($2LiF - BeF_2$) [131] and LiF and BeF_2 crystals using NJOY [132]. The TSL for FLiBe was applied to the MSRE (fuel fraction 22.5%) and the resulting k_{eff} was found to be approximately 0.5% lower than the free-gas model. The k_{eff} difference is strongly dependent on temperature and the size of multi-ion clusters.

The available TSLs and published results in literature are insufficient to determine the k_{eff} difference due to thermal scattering in molten fuel salt. Further it is far beyond scope and necessity of this study to construct the TSL for each candidate salt.

Chapter 4

Uncertainty and Sensitivity Analysis

The analysis of uncertainty of outcomes is often an important part of modelling, especially of complex systems [133], and the evaluation of an FHR/MSR involves a complex model with inputs that have considerable uncertainty. Further, understanding of the sensitivity of individual input variables to those outcomes can provide beneficial insight into the effects of altering the input variables. Sensitivity analysis is useful in understanding sources of uncertainty and determining the inputs that are most consequential to the total cost.

Models are the primary means available to analyse both physical and economic risk of nuclear power plants. Models predict the values of various system performance indicators. The model used to estimate the cost of MSRs, based on the framework described in Chapter 2, involves a considerable amount of uncertainty in the total cost as well as uncertainty in the cost sensitivity to individual input variables. It involves two types of variables: input (independent or explanatory) variables and the output (dependent or response) variable. The output variable represents cost.

This chapter focuses on the identification and quantification of uncertainty in model predictions. The uncertainty analysis portion attempts to describe the probabilities of all possible outcomes. The other component of analysis, sensitivity analysis, aims to assess how much model output values are affected by changes in model input values.

Multiple techniques have been developed to conduct uncertainty analysis, most situations make use of either a Monte Carlo, differential analysis or response surface method [134] [135] [136] [137]. Monte Carlo is seen to be the most common and widely applicable. The Monte Carlo approach that is taken assumes a 'black box' model - a model where the theoretical interactions between inputs and outputs not available or unduly complex; the variable relationships can only be inferred from the model outputs. Each input is assigned a probability distribution from which each iteration is independently drawn.

The generation of a random sample depends on being able to generate uniformly distributed random numbers from the interval [0,1]. These numbers may be referred to as pseudorandom number because they are generated by algorithmic processes rather than by a truly random process. This work takes the capability to

generate random numbers for granted and does not discuss the issue any further.

4.1 Uncertainty Analysis Overview

Uncertainty analysis is an inquiry of the confidence of results. Through the quantification of uncertainties it aims to determine the reliability of model predictions, accounting for various sources of uncertainty in the input variables. Mathematically it can be represented as:

$$Y = f(X) \tag{4.1}$$

Where Y represents the output variables $(y_1, y_2, ..., y_n)$, X the inputs variables $(x_1, x_2, ..., x_n)$ and f is the model. All or some of X are assigned a probability distribution. The objective is to determine the probability distributions of all Y.

An effective Monte Carlo analysis would accurately represent the probability distributions, sample broadly, efficiently and representatively. Repeated sets of uncertainty analysis should result in adequately similar means and variances of the overall cost. A set of uncertainty analysis runs contains many samples; each sample has a unique value assigned to each variable according to a probability distribution. Each set will produce a distribution of total cost values (for example, see Figures 5.41 and 6.33 from Chapters 5 and 6), that distribution will have a mean and a variance.

4.2 Variance Reduction

Variance reduction refers to techiques used to increase the precision of the estimates of Monte Carlo methods that can be obtained for a given simulation or computational effort. Popular methods employed include common random numbers, antithetic variates, control variates, importance sampling, moment matching, quasi random variables and stratified sampling [138] [139] [140] [141]. A common stratified sampling technique, Latin Hypercube Sampling, (LHS) is the primary means of variance reduction applied to the MSR evaluations. Common random number was employed as well. The others were found to incompatible with LHS, the model distributions and/or added a layer of complexity deemed unnecessary.

Importance sampling is where low probability sections of the distribution are sampled more often but are given a relatively low weight and high probability sections less often, with a high weight. Importance sampling was tested but deemed on balance to be unnecessary or even problematic for the following reasons:

- The use of LHS and a high number of samples (2000) resulted in an adequately low variance such that importance sampling is unnecessary.
- It adds complexity
- Curse of dimensionality: The high number of variables involved resulted in extreme weights, a substantial share were orders of magnitude from unity.

- The weight is proportional to $R_w^{\sqrt{n}}$, where n is the number of variables and R_w is the geometric average of weights and inverse weights (1/w) greater than or equal to 1 (if $w_1 = 0.5$, $w_2 = 2$, $R_w = \sqrt[2]{(1/0.5) \cdot 2} = 2$).
- For example, with 47 (about the number of variables with a Pearson distribution) variables, the greatest weight ratio (1:R_w) that can be reasonably be employed is 1:1.5. And with 23 variables, the most is 1:2.

The sampling algorithm makes use of LHS, a method for generating near-random parameter values from a multidimensional distribution [142]. The range of each variable is divided into n equally probable intervals. A sample point is placed in a random position in a randomly selected unused interval for each variable. This process is repeated until a position has been selected for all intervals. The advantage of this process over simple random sampling is that the entire range of each variable is sampled. The samples selected from the input variables are independent. The algorithm does not attempt to optimize the Latin Hypercube design (add orthogonality) because:

1. Optimizing the Latin Hypercube design becomes a combinatorial problem, and can be very expensive computationally [143]. Both the FHR and MSR models make use of approximately 70 variables thus a space filling optimization would likely have little effect and use considerable computer resources. Resources that are likely better spent running more iterations of the model.
2. Each iteration runs relatively quick so many iterations of each configuration can be run. The more iterations, the more completely the design space is sampled. It is generally argued that lower correlation is desired with orthogonality [144] [143], so a straightforward LHS sampling would avoid a limitation of maximum correlation.
3. The procedure to input specific correlations and achieve orthogonality is not well established and would be difficult to implement.

4.3 Probability Distributions

There is a wide variety of uncertainty of the mean value, bounds and shape of the distribution of the input parameters. The objective was to select distributions that resemble the expected shape, that allow for considerable shape control by the defining parameters and are adequately simple. Ideally and where possible, distributions are developed through an expert review process [145] or are derived from historical data. Several input costs have a price distribution resembling a Gaussian normal with a 'significant' positive skewness and kurtosis [146].

Two types of probability distributions are used, Pearson and double trapezoid ('double-trap'). The Pearson is generally selected when the input variable does not have a clear upper bound and likely is continuous with a central tendency. The double-trap is used when the expected variance is high relative to the mean and very little is known of the shape of the distribution.

4.3.1 Pearson Distribution

The Pearson distribution is a family of continuous probability distribution functions that satisfy Equation 4.2 [147]. Many forms of the Pearson distribution can be thought of as an expansion on the standard normal distribution; the most simple form of the Pearson distribution functions. The Pearson distribution considers two additional variables to the standard normal: skewness and kurtosis. Skewness is the third standardized moment and intuitively represents which side of the distribution has the larger tail. Kurtosis is the fourth standardized moment and intuitively represents how thick the tails are.

It is impossible to account for all interdependencies, bounds, discontinuities, non-linearities etc. associated with each distributed variable. Therefore no probability distribution can be known/proven. A standard Gaussian normal distribution requires that each variable be independent and random. Though a common assumption, it is naive and empirically unsound on metrics relating to economics, complex systems and large human projects.[1] Accounting for skewness and kurtosis better model the 'tails' thereby resulting in a more accurate distribution.

The family of Pearson distributions is selected as the best practically tenable probability distribution. Though not a general assumption, independent and random processes are involved so distributions similar to the standard normal are considered. Pearson distributions constitute the 'vast majority' of continuous probability density functions [147]. As well, there exists a lack of clear, viable alternative distributions. Few alternatives to the Gaussian have a central tendency and can be applied broadly to a range of many variables.

The Pearson distribution allows for the individual adjustment of skew and kurtosis of each variable. The skew and kurtosis applied to each individual variable is guided by empirical observation but still requires the use of intuition applied in an inescapably arbitrary fashion. Commodity prices were found to have 'significant' positive skewness and high kurtosis according to a 1993 OECD paper [146].

$$\frac{df(x)}{dx} = \frac{(x-a)f(x)}{dx^2+cx+b} \tag{4.2}$$

4.3.2 Double Trapezoid Distribution

The other probability distribution (PDF function) used, is assigned the name *'double-trapezoid'* it is composed of two side-by-side trapezoids. Each PDF of the two trapezoids is defined by the mean, minimum, maximum and peak:average ratio. The peak:average ratio is equal for each trapezoid and is the ratio of the maximum probability density (where the trapezoid meets the other) to the average probability density. In PDFs where the distribution mean is not halfway between the minimum and maximum bound, the average probability density of one trapezoid relative to the other is increased to achieve the target mean. Figures 4.4 and 4.5 are examples of this PDF, that demonstrate the effects of the parameters.

[1] Stock market returns [148], commoditiy prices [146], schedule overruns in construction projects [149] etc.

The double-trap PDF is generally used when one or more of the following apply:

- The expected variance is high relative to the mean. Inputs must be either strictly positive or negative for this condition to apply - in most cases in does.
- Little is known of the shape of the distribution.
- Inputs have (a) clear bound(s). Even if the exactly position of the bound(s) is not clear, an input is likely better represented by a bounded PDF.

An example is the exponent involved in estimating the expected cost overrun (in %), E, in Equation 4.3. Predicting a cost overrun is very unreliable, the shape of the distribution can only be roughly estimated from historical data. Bounds are necessary because a 99%+ percentile data point in a tailed distribution would likely result in an unreasonable expected overrun value.

$$overrun = (100 + 10 \cdot P_{MWe})^E \tag{4.3}$$

The objective of the construction of this PDF is to be roughly representative, be straightforward, sample broadly and cover the entire range of realistic values. Variables that are not bounded at zero tend to have a central tendency, so this was set as a requirement. The benefit of complexity depends on the knowledge of shape of the distribution. The Pearson distribution is already used when there is relatively high knowledge of the shape, thus there is relatively low knowledge of the shape of the remaining distributions and little benefit in complexity. As well, the whole range should be sampled relatively evenly when the level of knowledge of the shape is low.

Many distributions were quickly eliminated from consideration. All non-linear or tailed distributions tend to be rather complex and/or do not have an adequately high probability density across the entire range, so log-normal, Chi-square, Weibull etc. were ruled out. Two common linear distributions considered were:

- Flat distribution, which has two issues: no central tendency and a fixed mean for given bounds. However, it samples well across the entire range.
- Triangular distribution, which also has two issues: low probability near the tails and a large difference in the mean and mode in non-symmetric distributions. It however has a central tendency and does not have a fixed mean for given bounds.

The double-trap is an attempt to rectify the issues with the flat and triangular distributions while maintaining the advantages of each. The peak:average ratio used for all ‘double-trap’ input parameters in both the FHR and MSR analysis is 1.33.

4.4 Correlation

Two widely used possibilities exist for defining correlations between variables: the Pearson correlation coefficient (CC) and the Spearman rank correlation coefficient (RCC) [145]. The CC is the familiar least

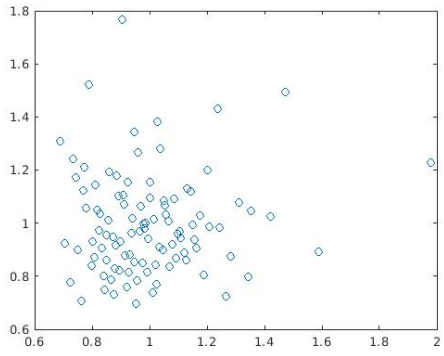

Figure 4.1: Example of a two dimensional LHS design space. 100 data points of each variable are generated from a Pearson distribution - see Section 4.3. The mean is 1.0, variance: 0.2, skew: 1, kurtosis: 7.

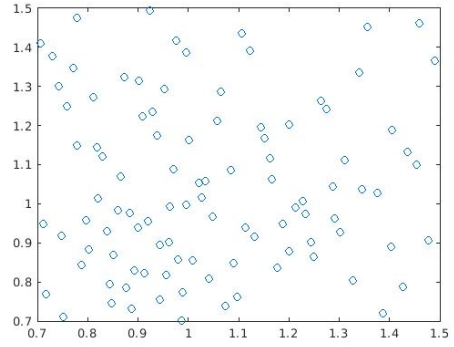

Figure 4.2: Example of a two dimensional LHS design space. 100 data points of each variable are generated from a 'double trapezoid' distribution - see Section 4.3. The mean is 1.0, lower bound: 0.7 and upper bound: 1.5.

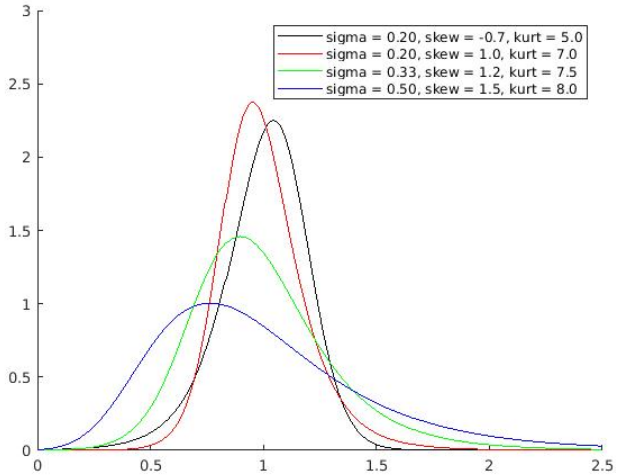

Figure 4.3: Examples of the Pearson probability density function, each with a different variance, skew and kurtosis. All examples have a mean (μ) equal to one.

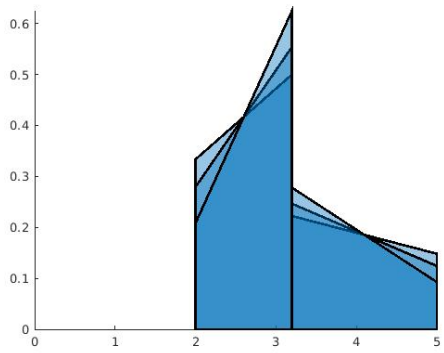

Figure 4.4: Examples of asymmetric double-trapezoid probability density functions. Shown with three peak:average ratios: 1.2, 1.33 and 1.5. Mean: 3.2, Min: 2, Max: 5.

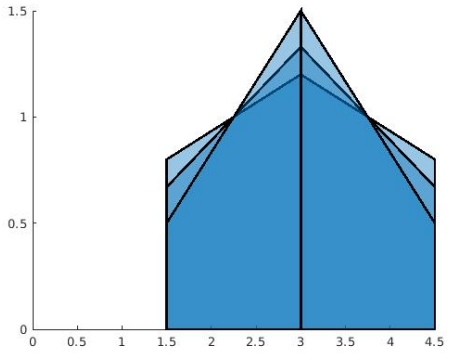

Figure 4.5: Examples of symmetric double-trapezoid probability density functions. Shown with three peak:average ratios: 1.2, 1.33 and 1.5. Mean: 3, Min: 1.5, Max: 4.5.

squared regression. The RCC is calculated similarly to the CC but with rank-transformed data, according to Equation 4.4

$$R = \rho_{R(X),R(Y)} = 1 - \frac{6\Sigma(R(X) - R(Y))}{n(n^2 - 1)} \tag{4.4}$$

Where $\rho_{R(X),R(Y)}$ is the CC but with rank variables, R(X) the rank of X, R(Y) the rank of Y and n the number of data points.

The expected correlations between select individual variables were added to the sample creation process. A straightforward method to add CCs could not be found, so instead correlation was imposed through RCCs. A common [145] [150] [151] [151] and effective method for implementing rank correlations on generated random data has been proposed by Iman and Conover [152]. It has several advantages including the preservation of the stratification from LHS and distribution independence (any sampling can be used). The technique involves inputting the correlations into a symmetric matrix, performing a Cholsky decomposition on the matrix and then multiplying the rank position of data by Cholsky decomposition. A detailed explanation is not provided in this document.

A matrix representation of all correlations must be positive definite in order for the set of correlations to be feasible - meeting this criterion becomes more difficult as the number of variables grow. All non-zero symmetric entries into a identity matrix have the effect of making the determinant more negative. Thus large correlation matrices cannot have many high correlations. Most variables combinations are uncorrelated and those that are were not able to be correlated more than approximately r = 0.30.

4.5 Sensitivity Analysis

Sensitivity analysis is conducted to gain insight into the magnitude of the effect that each input variable has on the total cost as well as the reliability of the estimate of the effect. It is based on the same setup

as the uncertainty analysis; the same probability distributions, Monte Carlo LHS approach and input correlations. The sensitivity analysis is based on multivariable linear regression and is based on the following assumptions: [153] [154]

- A linear relationship between the dependent variable and each of the independent variables. In a black-box model this assumption cannot be confirmed with certainty. The model is however largely additive (a structure of f(X) = x1 + x2 + ... + xn) and an additive model is linear. Some variables are inherently non-linear, the interest rate is a good example, as it discounts future costs by $(1-i)^n$. However the few variables that are non-linearity, only deviate from linearity by a small amount (see Figures 5.49 to 5.53 and Figures 6.41 to 6.46.
- The input variables are not highly correlated - they do not exhibit multicollinearity. It can be quantified by the Variance Inflation Factor (VIF):

$$VIF = \frac{1}{1-R^2_{X_k}} \tag{4.5}$$

  Where R_{X_k} is obtained by regressing X_k on all X except X_k. If the VIF exceeds 5, the multicollinearity is considered high [155]. For a VIF > 5, $R^2_{X_k} > 0.8$, this is never the case since the maximum correlation coefficient between any two variables in both the FHR and MSR analysis is set at only 0.3.
- Homoscedasticity, which is where the variance of the residuals is consistent over the range of the output. In both the FHR and MSR models this assumption does not hold true. It does not prevent the use of multiple linear regression, this is discussed alongside the multivariate normality.
- Independence of observations.
- Multivariate normality: in all evaluations, the residuals are not normally distributed. Thus both the homoscedasticity and multivariate normality assumptions are clearly false, see for example Figure 4.8. The multiple linear regression model can still be sufficiently valid as long as the variance of the residuals is low.[2] A non-normal, heteroscedastic distribution of residuals is non-problematic if the variance of the residuals is insignificant. In the case of the FHR evaluations, the variance of the residuals is 44.1 to 100.1 (corresponding to a standard deviation of ≈4% of the average cost) for the lowest cost configurations of each salt. This is low relative to the variance of the predicted values, 552 to 912. The variance is similarly low for the MSR configurations evaluated. Thus if non-predictable, normally-distributed, randomly generated variance were added to the model the residuals would likely appear sufficiently homoscedastic and normally distributed, though multiple linear regression would, if anything, be less valid. It should be noted that a frequently mentioned consequence of heteroscedastic residuals is an increase in the variance of the predicted coefficients. Histograms of several regression coefficients show that the actual variance differs little from the expected variance, for instance Figure 4.7, thus this issue is only a minor concern. Figures 5.54 and 6.47 show histograms from the FHR and

[2] This is a mathematical model where the output is an exact (complex) function of many inputs. This is no random variation added apart from the values of the inputs and there are no environmental or other non-controllable factors contributing non-predictable variation. As the result it is unsurprising that the variance of the residuals is very low. It is presumed that nearly all of the variance arises from deviations in linearity.

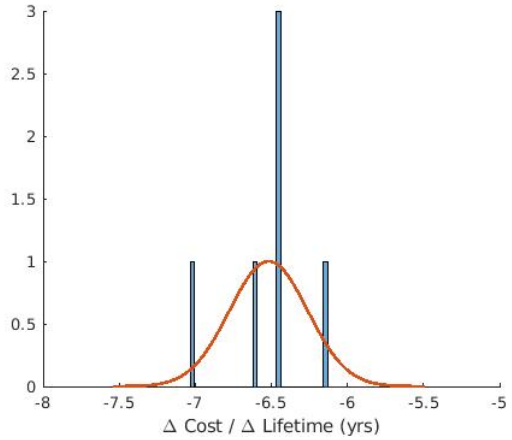

Figure 4.6: LiNaZrF

Figure 4.7: Histogram of the regression coefficient of lifetime for all LiNaZrF configurations with 2000 samples overlaid with a normal curve that has a variance equal to the average expected variance of the coefficient.

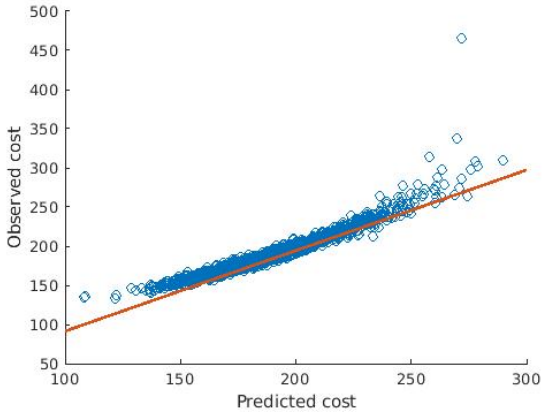

Figure 4.8: Predicted compared to observed cost for the lowest cost FHR configuration with NaBeF coolant salt. The difference between the two values (deviation from the line) represents the residuals.

MSR chapters respectively.

Notably multi-variable linear regression does not require normally distributed input variables. The shape of the input distribution is only a concern if the residual normality is insufficient.

Chapter 5

Fluoride High Temperature (FHR) Reactor Analysis

5.1 Introduction

The Fluoride-salt-cooled High-temperature Reactor (FHR) is the principal molten salt cooled, pebble-bed reactor conceptual design. The purpose of this chapter is to evaluate potential FHR geometries and coolant salt compositions and thereby optimize the overall design.

The analysis requires a reference design from which changes in design parameters are evaluated. The first pre-conceptual FHR design, the AHTR, emerged in the early 2000's [156] and the latest design to have undergone the considerable development with publically-available documentation is the Mark 1 Pebble-Bed (PB) FHR. In 2012 the DOE through its Nuclear Energy University Program, awarded a 3-year Integrated Research Project to a network of U.S. universities [47]. The conceptual design and technology is well documented in the Mark-1 PB-FHR Technical Description (2014) [48].

The analysis conducted in this chapter uses, as a reference design, a simplified version of the Mark-1 PB FHR. The design concept is examined across the entire design space of realistic configurations. The performance data generated from neutronics simulations and thermohyraulic analysis is used by the evaluation program (see Chapter 2) to come up with cost estimates across the entire design space. The evaluations are pieced together and trends are discerned in order to gain insight into the impact various design features have on reactor performance and ultimately to determine the lowest cost design configuration.

The Mark-1 PB-FHR core contains 3.0 cm diameter pebbles containing 0.081 cm diameter TRISO fuel particles with 19.9% enriched $UC_{1.5}O_{0.5}$ fuel. All subsequence analysis in this paper makes use of this same general pebble/particle geometry since it is well-developed and is a necessary limitation to restrict the bounds of the 'optimization' to a level which is feasible. The Chinese Academy of Science (CAS) is persuing a similar TRISO particle fuelled, FLiBe cooled, design with plans for a demonstration reactor by

2025 and a commercial reactor by 2030 [157] [158]. [1]

Limiting the general structure of the pebble/particle design considerably narrows the scope of potential salt-cooled designs. Other assumptions made to further narrow the scope, simplify the analysis and remove infeasible options include limiting the design to:

- Cylindrical shape with a central graphite core.
- 21.65 MWth/m^3 power density.
- 200 MWth.
- Continuous recirculation refuelling scheme in which the burnup level of an outgoing pebble determines whether it is recirculated or discarded.

The remaining variable design parameters that constitute the possible design space are discussed below. They can be grouped into two categories; the coolant salt composition and geometry. The main impact of the geometric parameters is on the level of moderation.

5.1.1 Selection of Salt Compositions

Both the Mark 1 PB-FHR and the TMSR-SF2 specify the use of FLiBe salt, $2LiF - BeF_2$ with enriched lithium ($\approx$ 99.995% ^{7}Li). FLiBe has desirable physical, chemical and neutronics properties however beryllium is toxic and ^{7}Li is expensive; in 2010, the estimated price for enriched ^{7}Li from China was $15,000/kg at a 99.99% [122]. Alternative salts have a different sets of trade-offs that could make them advantageous in certain situations.

Various coolant salt compositions are potentially feasible in this design class. The scope of this study limited the evaluation to 5 salt systems. $LiF - BeF_2$, with a 67-33 molar composition or FLiBe is a well studied salt that has been proposed in many MSR designs, so it was included. Other viable salts are not as obvious. Process of elimination, as discussed in Section 3.2, was largely employed to arrive at the other salt compositions selected for analysis.

Property data is more widely available at, and the lowest T_{melt} is typically a mixture at a eutectic point - thus the salts should be eutectic fluoride mixtures of the viable cations Li(LiF), Na(NaF), Be(BeF_2) and Zr(ZrF_4). Among binary mixtures, $LiF - NaF$ and $BeF_2 - ZrF_4$ do not have sufficiently high melting points [111] [159]. $LiF - ZrF_4$ though viable, was rejected because it has only minor advantages over either $NaF - ZrF_4$ and $LiF - BeF_2$. Therefore, $LiF - BeF_2$, $NaF - BeF_2$ and $NaF - ZrF_4$ remain. Among ternary mixtures $LiF - NaF - BeF_2$ and $LiF - NaF - ZrF_4$ are the two most advantages mixes. The particular eutectic points

[1] In 2011 the Chinese Academy of Science (CAS) effectively restarted their Thorium Molten Salt Reactor (TMSR) program. The TMSR project pursues both circulating liquid fuel and TRISO particle fuelled, graphite moderated concepts. The pebble-bed (solid fuel) demonstration reactor, the TMSR-SF2, will be 168 MWe, 384 MWt and FLiBe cooled. It is to contain TRISO fuel in 3 cm diameter pebbles with 19.75% ^{235}U held in a 2.5 m diameter, 4.2 m height core. The temperature reactivity feedback coefficient is -5.90 pcm/K, the inlet temperature is 600 C and outlet 700 C [157] [158].

were selected on the basis of available property data, conclusions in literature and authors judgement. The five salts analysed are:

1. $LiF - BeF_2$ (67-33), or FLiBe/Flibe
2. $NaF - BeF_2$ (57-43), or NaBeF
3. $NaF - ZrF_4$ (59.5-40.5) or NaZrF
4. $LiF - NaF - BeF_2$ (15-58-27) or LiNaBeF
5. $LiF - NaF - ZrF_4$ (26-37-37) or LiNaZrF

Corrosion is an important consideration for FHRs; the relative corrosion of candidate salt compositions was taken into consideration. However, it is a complex subject, as discussed in Section D.2, and no compositions were explicitly excluded due to metal container compatibility.

In any reactor salt system, a purged cover gas will likely be required. The transport of salt vapor in this cover gas system can be problematic. As a result, salts with low vapor pressures are preferred. This is acutely a concern with Zr-containing salts. Williams et al. [49] in the assessment of the $LiF - NaF - ZrF_4$ system found 26-37-37 and 33-24-43 mixes to be the most promising $LiF - NaF - ZrF_4$ coolants because '(a) they maintain $ZrF_4 < 40$ mole % (necessary for low vapor pressure) and (b) they provide a significant melting point depression (40–64^{o}C) compared to the $NaF - ZrF_4$ binary system.' The 26-37-37 mix has been selected for analysis because it has lower Li and Zr concentrations. For more on the constituent mole fraction selection see Section 3.2.2.

5.1.2 Reference Reactor Key Design Specifications

The reference design or reference model is the design from which all tested parameters are evaluated relative to. The geometry is simple; it does not include control rods, structural material or other components that have a relatively minor impact on reactor neutronics. The level of burnup of the fuel pebbles in the reactor is spread across a continuum. Fresh fuel is continually added, replacing fuel with the highest burnup. In an attempt to capture the continuum, the fuel divided into eight equal-sized groups, each with a unique burnup. The approach often taken to model the range of temperatures throughout the core is to divide the coolant salt into 15 'zones' with unique temperature and corresponding density. More on the fuel groups and temperature zones in Section 5.4.2. The basic specifications of the reference reactor are listed in Table 5.1.

Salt temperature

The salt temperature is inherently somewhat arbitrary and selected to be approximately that of other FHR designs. The Mark 1 PB-FHR has inlet and average-outlet temperatures of 600 and 700 ^oC [48] respectively. The TMSR-SF2 also has inlet and outlet temperatures of 600 and 700 ^oC [157] [158]. In this analysis the maximum salt temperature in each configuration is in the inner and top most zone, with a temperature applied to the zone of 691 - 713 ^oC - the temperature varies with the inner radius. The minimum temperature of 615

Table 5.1: FHR Reference Reactor Key Specifications.

Parameter	Value
Coolant Salt	$2^7 LiF - BeF_2$
Avg. Coolant Salt Outlet Temp	690 ^oC (963 K)
Avg. Coolant Salt Inlet Temp	610 ^oC (883 K)
Reactor Vessel H	3.0 m
Reactor Vessel Outer D	2.1 m
Reactor Vessel Inner D	0.7 m
Radial Reflector Thickness	0.45 m
Axial Reflector Thickness	0.50 m
Thermal power	200 MW
Thermal power density	21.65 MW/m^3

^oC in each configuration is in the outer and bottom zone. It is assumed that a portion of the salt mixes and/or heat is conducted radially. The amount of heat transferred between the inner and middle and between the middle and outer radial zones is proportional to the temperature difference. An amount is transferred such that the temperature difference between the inner zone and the lowest temperature zone at the outlet is 50% of what it would be without mixing/heat transfer. Volume differences between the radially zones are taken into account. Figures 5.1 and 5.2 show the positioning of the temperature zones for an inner radius of 35 cm. The axial zone boundaries are equally spaced and the radial zone boundaries are listed below.

- 25 cm: 25 50 80 105 cm
- 35 cm: 35 60 80 105 cm
- 45 cm: 45 65 85 105 cm

Fuel temperature

An important variable necessary for accurately modelling FHR designs is the fuel temperature. The higher the temperature the higher the risk of fuel melt and all the associated issues. It is quite possibly the limiting factor to the maximum power for most core configurations. For the purposes of the evaluation framework, the safety cost of the fuel temperature is assumed to follow:

$$C_T = \frac{T_{max}}{T_{max} - T} - 1 \tag{5.1}$$

Where C_T is the cost, T_{max} is the maximum fuel temperature the reactor could conceivably operate at and T is the highest fuel temperature in the proposed design.

No published documentation, including the comprehensive Mark 1 PB-FHR technical report [48], precisely specifies the fuel temperature, so it must be calculated from first principles where possible or deduced from other reactor designs with TRISO fuel particles. Moreover, no other salt-cooled designs have public documentation that includes the expected fuel temperature, so gas-cooled reactors with TRISO fuel particles

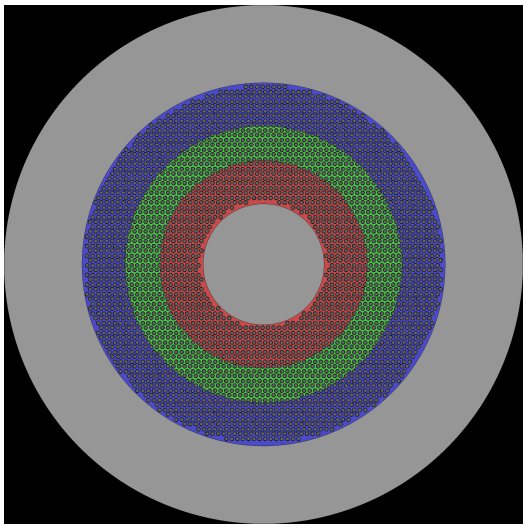

Figure 5.1: Radial view of the reference FHR geometry with pebbles in a hexagonal-closed-packed arrangement and highlighting the radial zones.

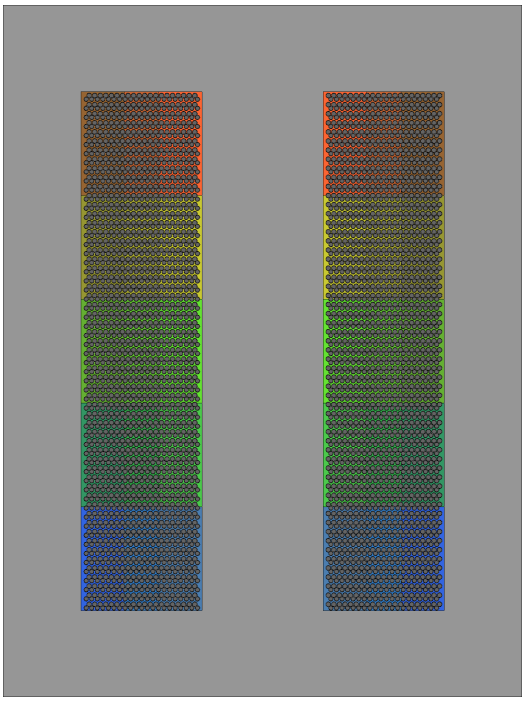

Figure 5.2: Axial view of the reference FHR geometry with pebbles in a hexagonal-closed-packed arrangement and highlighting the axial zones.

must be examined instead. Simplifying assumptions, particularly that property data is constant with temperature, are applied so that the temperature differences (ΔT) are proportional to the power in each zone. Subsequent calculations are of the average not the peak ΔT.

The ΔT between the fuel and salt is calculated for the reference pebble and particle geometry, then adjusted for the configuration under consideration according to Equation 5.2:

$$\Delta T = \frac{100}{100 - B} \cdot \left(\frac{40}{PF} \cdot \Delta T_{KM} + \Delta T_{MS}\right) \tag{5.2}$$

Where B is the percentage share of pure graphite pebbles (blanks), ΔT_{KM} the ΔT from the fuel kernel to matrix, ΔT_{MS} the ΔT from the matrix to coolant salt and PF the packing fraction of particles.

The thermal resistances are grouped into three categories:

1. Fuel kernel to graphite matrix (ΔT_{KM})
2. Graphite matrix to pebble surface
3. Pebble surface to coolant salt (ΔT_{MS})

Fuel kernel to graphite matrix: The fuel kernel is surrounded by 4 layers of different matrial and is at a small scale such that difficult-to-predict contact resistances are assumed relatively high. Thus the thermal

conductivity cannot be reliably calculated from first principles. The High Temperature engineering Test Reactor (HTTR) by the Japan Atomic Energy Agency (JAEA) [160] is sufficiently documented such that through an extrapolation of its geometry and operating temperatures the ΔT between the TRISO particles and the graphite matrix in the FHR can be estimated. Assumptions made include:

- The heat transfer coefficient between the fuel compact and the particles in the HTTR and between the graphite matrix in the pebbles and the particles is the same. Both cases are graphite-to-graphite contact.
- All heat transfer coefficients are temperature independent.
- The HTTR reactor has a power distribution with a PPF ($\frac{P_{peak}}{P_{avg}}$) of 1.5. [2]

Table 5.2 lists design specifications of the HTTR and FHR relevant to the calculation of the fuel temperature in the FHR as well as intermediate calculated values. The key factor is the *'relative FHR TRISO heat transfer,'* which attempts to captures the relative difference in thermal resistance between the HTTR and the FHR, the higher the number the lower the relative resistance of the FHR particles. It is calculated according to Equation 5.3.

$$1.38 = \frac{6}{4}x + \frac{160}{205}(1-x) \tag{5.3}$$

Where x, equal to 0.83, is the contact resistance share, $\frac{6}{4}$ the relative contact resistances and $\frac{160}{205}$ the relative conductive resistance. The contact resistance ratio, $\frac{6}{4}$, is the relative difference in fuel kernal surface area per unit fuel kernal volume. The conductive resistance ratio, $\frac{160}{205}$, is a ratio of the thickness of the surrounding layers.

Next the Δ T between the graphite matrix and the pebble surface is calculated. The can be done using Fourier's law, shown rearranged for solving ΔT in Equation 5.4.

$$\Delta T = \frac{qd}{k} \tag{5.4}$$

Where q is the thermal flux (W/m^2), k the thermal conductivity and d the distance. Inputs to Equation 5.4 necessary to solve for ΔT:

- The reactor has a thermal power of 200 MW and contains 375,000 pebbles in core, thus power per pebble is $200 \cdot 10^6 W/375,000 = 533 W/pebble$.
- The surface area of a pebble is:

$$A = 4\pi r^2 = 4\pi(0.015m)^2 = 0.00283m^2 \tag{5.5}$$

- The thermal conductivity of graphite used in this calculation is 77 W/mk [161].
- The average distance between matrix graphite and the pebble surface is 0.0013 m.

[2] The radial PPF is given in Table 2.4 and the axial power distribution is shown in Figure 2.16. The product of the two is is approximately 1.5, [160] however the PPF of the HTTR is dependant on burnup.

Table 5.2: Comparison of select design parameters of the HTTR and FHR relevant to the calculation of the average fuel temperature in FHR designs

Parameter	HTTR	FHR
Volume per fuel compact	10.782 cm^3	—
Fuel compacts per fuel rod	14	—
No. of fuel rods	280	—
Particle packing fraction	0.30	0.40
Volume fuel/volume fuel compact	0.1276	—
TRISO particle outer diameter	920	810 μm
Fuel kernel diameter	600	400 μm
Total fuel volume	6933.9	59,723 cm^3
Max. temp of fuel compact	1270 ^oC	—
Max fuel temp	1420 ^oC	—
Compact matrix to fuel ΔT [a]	150 ^oC	—
Relative FHR TRISO heat transfer	1.38	—
Max matrix to fuel ΔT, FHR particles, PPF = 1.5	109 ^oC[b]	84 ^oC[c]
Thermal power	30	200 MW
Thermal power/cm^3 fuel	4.33	3.35 kW/cm^3 fuel
Avg. matrix to fuel ΔT, FHR particles, PPF = 1.5	—	**56** ^oC[d]

[a] 1420-1270 ^oC

[b] The matrix to fuel ΔT if the HTTR used FHR TRISO particles

[c] $84 = 109 \cdot 3.35/4.33$

[d] ΔT_{max}/PPF = 84/1.5

The thermal flux is the quotient of power per pebble (533 W) and the surface area (0.00283 m^2). Substituting all the values into Fourier's law:

$$\Delta T = \frac{533/0.00283 \cdot 0.0015}{77} = 3.3^o C \tag{5.6}$$

Finally the Δ T between the pebble surface and the molten salt coolant must be calculated. This is calculated using Newton's law of cooling as shown in Equation 5.15.

$$Q = hA\Delta T \tag{5.7}$$

Where Q is the total heat transfer (W), h the heat transfer coefficient (W/m^2K) and A the pebble surface area. As before, the power per pebble is 533 W and the surface area of a pebble is 0.00283 m^2. The heat transfer coefficient h is solved for using the following Nusselt number correlation: [162]

$$Nu = \frac{hD}{k} = 2 + [0.4Re^{1/2} + 0.06Re^{2/3}]Pr^{0.4}(\frac{\mu_\infty}{\mu_s})^{1/4} \tag{5.8}$$

The Reynolds and Prandlt numbers are a function of salt properties. The temperature used is an average of the inlet and outlet temperature of a core with an inner radius of 35 cm. An example of it is solved for using FliBe salt at 650 ^{o}C. First the velocity must be calculated:

$$\begin{aligned} Q[m^3/s] &= \frac{P}{Cp \cdot \Delta T \cdot \rho} = \frac{200 \cdot 10^6}{2420 \cdot 80 \cdot 1931} = 0.535 \\ u[m/s] &= \frac{Q}{0.9 \cdot (1 - PF)X} = \frac{0.41}{0.9 \cdot 0.4 \cdot \pi(1.05^2 - 0.35^2)} = 0.39 \end{aligned} \tag{5.9}$$

Where Q[m^3/s] is the volumetric flow of salt through the core, P[W] reactor power, Cp [J/kgK] the heat capacity, ρ [kg/m^3] the salt density, 0.9 a non-dimensional factor to account for flow components perpendicular to the net velocity, PF the packing fraction of the pebbles and X the radial cross-section of the core. Now the Reynold's and Prandlt number:

$$Re_D = \frac{\rho u D}{\mu} = \frac{1931[kg/m^3]0.39[m/s]0.030[m]}{0.0044[Pa \cdot s]} = 5135 \tag{5.10}$$

$$Pr = \frac{C_p \mu}{k} = \frac{2420[J/kgK]0.0044}{1.1[W/m^2 - K]} = 9.68 \tag{5.11}$$

Substituting in the values in the Nusselt number correlation

$$Nu = \frac{hD}{k} = 2 + [0.4(5135)^{1/2} + 0.06(5135)^{2/3}](9.68)^{0.4}(\frac{0.0044}{\approx 0.0035})^{1/4} \tag{5.12}$$

$$Nu = \frac{hD}{k} = 2 + [28.66 + 17.86]2.626 = 124.1 \tag{5.13}$$

Table 5.3: Summation of temperature increases used in the calculation average fuel temperature.

Thermal resistance	ΔT
Fuel kernel to matrix	56
Matrix to pebble surface	3.3
Pebble surface to salt coolant	41.4
Total average Δ T	100.7

Table 5.4: FHR fuel temperature groupings. ox indicates 'outer,' mx 'middle' and ix 'inner' zone.

$NaF-BeF_2$, PF 40, blanks 250, rad 35cm														
Low					Mix						High			
o5	m5	i5	o4	m4	(m1	o1	m3	o3	m2	o2)	i4	i1	i3	i4
975	978	1021	1025	1026	(1036	1038	1055	1057	1058	1060)	1091	1102	1135	1137
975	+3	+43	+4	+1	(+10	+2	+17	+2	+1	+2)	+31	+11	+33	+2

Solving for h:

$$h = \frac{Nu \cdot k}{D} = \frac{124.1 \cdot 1.1}{0.03} = 4552[W/m^2K] \tag{5.14}$$

Substituting into Newton's law of cooling, Equation 5.15, and rearranged to solve for ΔT:

$$\Delta T = \frac{Q}{hA} = \frac{533W}{4552[W/m^2K] \cdot 0.00283[m^2]} = 41.4K \tag{5.15}$$

Fuel temperature in each zone:

Working with the assumption that thermal conductivities are temperature independent, the fuel to coolant ΔT in each zone is calculated. The fuel ΔT in each zone is proportional to the relative power density (individual zone power density divided by average power density) and is a function of the number of 'blank' pebbles and the PF (see Equation 5.2). The relative power density of each zone is also a function of the inner radius.

As mentioned in the section introduction, it is not practical to have a unique fuel temperature in each zone, so the fuel temperature in each zone is rounded to either a 'high,' 'low,' or a mix of both temperature(s). Multiple fuel temperature calculations were made with the temperature arranged in ascending order and three zone groupings corresponding to high, low or mix. The objective of the high/low/mix assignments was to minimize the ΔT between the high/low/mix temperature and the unique calculated temperature of each zone. Ideally the number of high and low temperature zones would be equal or nearly so. Despite varying fuel temperatures among the geometric configurations, a similar grouping emerged in most cases. The high and low temperatures used are the fission-reaction-weighted-average of the zones included in the respective high and low groupings. An example, along with the chosen groupings in provided in Table 5.4 and Figure 5.3 shows what fuel mix is used in each zone, blue is exclusively low, purple a mix and red exclusively high temperature.

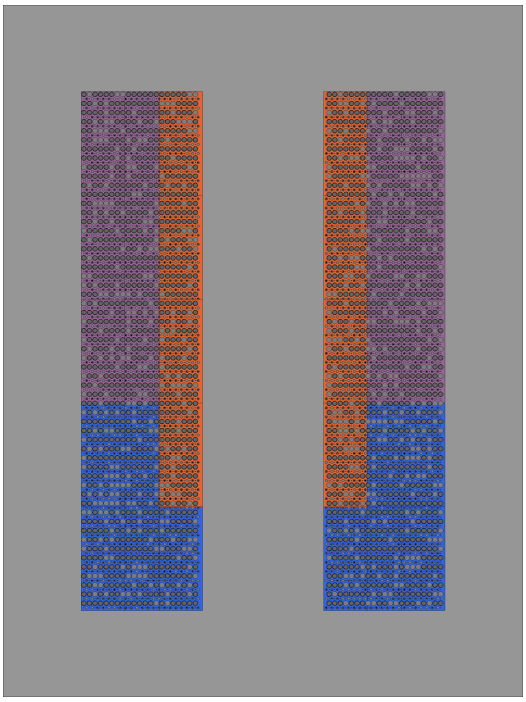

Figure 5.3: Fuel mix is used in each zone; blue color indicates an exclusively low, purple a mix and red exclusively high temperature zone.

5.1.3 Selection of Design Parameters to be Evalated

The primary objective of the FHR analysis is to determine the configurations that result in a low total cost within a design space specified by impactful design parameters. The entire design space is evaluated rather than evaluating parameters independently in order to avoid running the risk of neglecting any key variable dependencies and instead taking into account any interdependence. The trade-off is that it necessitates a large number of neutronics simulations that increase exponentially with increasing parameters - the so called *'curse of dimensionality.'* Available computational resources limited the number of geometric design parameters to three.

All neutronics calculations are performed using the Monte Carlo reactor physics code SERPENT 2 [163]. Four pebble-level design parameter optimizations were considered (1-4) and two reactor-level (5-6):

1. TRISO fuel particle design.
2. The addition of pure graphite pebbles.
3. The diameter of the pebbles.
4. PF of the particles inside the pebbles.
5. The diameter to height ratio.
6. The radius of the inner graphite reflector.

TRISO fuel particles are spherical particles of nuclear fuel surrounded by multiple layers of different types of graphite and a layer of SiC. The fuel temperature and the net concentration of fuel in the core are likely the two key performance metrics effected by the TRISO design. The smaller the TRISO particles the higher the heat conductivity and lower the fuel temperature. Though a key metric, the net concentration of fuel can vary by particle PF inside the pebbles instead. Ultimately, the TRISO design optimization was avoided primarily because the manufacturing costs of various geometries would be very difficult to predict; any results would have been prohibitively unreliable.

The addition of pure graphite pebbles among the active fuel pebbles is a feature that had not been considered in the Mark 1 design [48] and to the extent of the author's knowledge could not be found in other FHR designs either. As a potentially novel contribution, it was prioritized for analysis. The fundamental trade-off in this analysis is between a higher fuel temperature and higher neutronic efficiency. It is assumed that graphite pebbles would be dispersed throughout the core stochastically.

The PF of the particles is an important parameter to optimize because is has a proportional effect on the key metric of core fuel concentration. Without this design parameter, the entire design space examined risks only evaluating sub-optimal fuel concentrations. Given the importance of fuel concentration, failing to optimize the PF risks omitting the lowest total cost configuration from evaluation. The share of graphite pebbles and pebble diameter parameters alone risk failing to generate a large enough range of fuel concentrations.

Many of the trade-offs associated with the reactor diameter to height ratio are difficult to quantify. A smaller diameter would have a higher salt temperature difference through the core but would be likely be easier to transport and perhaps construct. The optimal neutronics efficiency would be approximately where diameter = height. Apart from the neutronics it is difficult to meaningfully quantify the other key trade-offs.

For a given fuel concentration, a larger pebble diameter is expected to increase fuel temperature, lower refuelling costs and slightly enhance neutronic efficiency through fuel-moderator separation. However the magnitude of the fuel temperature and neutronic efficiency effects would be less than with graphite pebbles at the same fuel concentration.[3] All three of these trade-offs can, to a large degree, be replicated by and are better highlighted by a combination of PF and graphite pebble optimization, therefore including pebble diameter as well is somewhat redundant.

The diameter of the central graphite reflector has an effect on the flux shape (Power Peaking Fraction, PPF), salt velocity and the level of moderation.[4] The impact of the central reflector size is dependent on the level of moderation. Like other design parameters, the important trade-off is between fuel temperature and neu-

[3] In this reactor design the fuel-moderator separation is much less than the neutron mean free path. As a result, the stochastic arrangement of blank pebbles enhances neutronic efficiency compared to simple patterned arrangements.

[4] In the interest of maintaining a consistent total core volume, the outer diameter and/or height would have to be adjusted to compensate. The outer diameter was elected to remain constant since it is a more consequential parameter to manufacturing and transportation. Moreover a constant outer diameter would result in a greater variance in moderation.

tronic efficiency.

In summary the three variable parameters that constitute the design space are:

1. Share of graphite pebbles (blanks).
2. Packing Fraction (PF) of the particles inside the pebbles.
3. Radius of the inner graphite reflector (inner radius).

5.2 Pressure Drop and Heat Exchanger Optimization

5.2.1 Core Pressure Drop

It is beyond the scope of this work to calculate the Pressure Drop (PD) of a pebble bed FHR from first principles or with simulation software. Instead PD data is taken from the Mark 1 design. Though it does not explicitly provide the PD of either the entire primary loop or the core, it does mention that the primary loop is designed to keep head losses under 2 m and suggests a HL across the core of 0.28 m (5.4 kPa) [48]. This piece of information is used as a reference from which to estimate the PD through the core for other salt compositions and core configurations as well.

First, the relative PD for an alternative salt composition is calculated. To do so the heat exchanger pumping power for the reference salt composition ($LiF-BeF_2$) and the salt under examination (salt X) are extracted from the heat exchanger analysis. The ratio of pumping powers of salt X and the reference is then calculated (R_{XR}); a velocity must be selected to obtain the pumping powers, any velocity can be chosen since the ratio is not expected to be a substantial function of velocity. The temperature difference is equal regardless of reactor length, so the flow rate must differ among the inner radii. PD is a function of velocity squared (V^2) and the reference configuration is 5.4 kPa at a 35cm inner radius. The pump power draw of the other inner radii configurations are adjusted according to the relative flow velocity, and the velocity is relative to the radial cross sectional area; it is adjusted as per Equation 5.16. Next the power required to pump $LiF-BeF_2$ through the core is calculated, according to Equation 5.17, where P_{th} is reactor thermal power. Finally, the core pumping power for the salt under examination is calculated as the product of Equations 5.16 and 5.17. The implicit simplifying assumption is that the power ratio R_{XR} between two salts through a pipe and through a pebble-bed core is equal and independent of velocity.

$$R_{XR} = R_{XR} \cdot \frac{105^2 - 35^2}{105^2 - r_x^2} \tag{5.16}$$

$$P = (PD) \cdot \frac{\dot{m}}{\rho} = \frac{(PD)\frac{P_{th}}{C_p \Delta T}}{\rho} \tag{5.17}$$

$$P_x = P \cdot R_{XR} \tag{5.18}$$

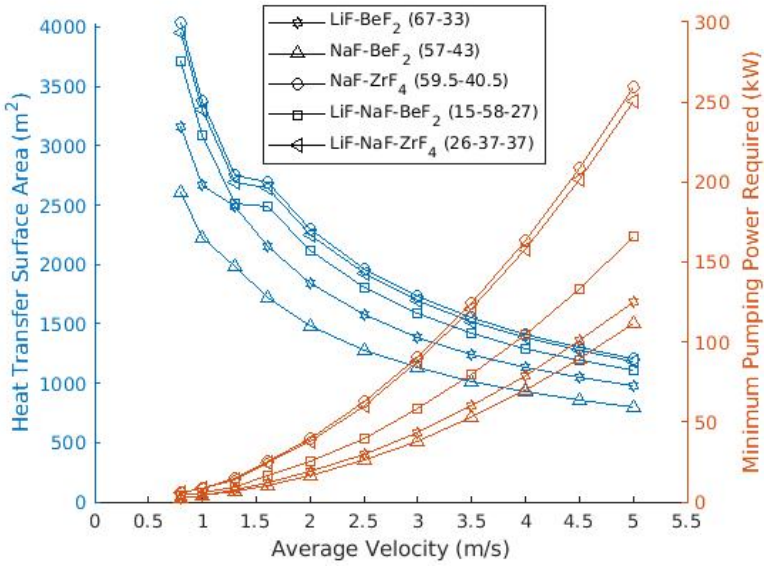

Figure 5.4: Heat transfer surface area and pumping power required in a molten salt heat exchanger for a 200 MWth FHR. Several surface area curves are not smooth because the Reynold's number switches from below to above 12000, and hence the empirical correlation switches, at $\approx 1.5 m/s$.

5.2.2 Pumping Power and Material Requirements Optimization in the Molten Salt Heat Exchanger

There is a trade off between required pumping power and tubing material as shown in Figure 5.4. The purpose of this analysis is to optimize the pumping power requirements and the heat exchanger pipe area (size) for each of the candidate coolant salts. The setup, assumptions and procedure to arrive at a required pumping power and pipe area are discussed in Section 3.3. This section only discusses and presents the results.

In order to cost optimize the circulation velocity, a value must be assigned to $/$m^2$ of heat transfer surface area and $/kW of pumping power required. The expected salt heat exchanger cost is $2500/$m^2$. To account for pump inefficiency, secondary salt pump power and pump capital costs, the expected total pumping related costs are estimated to be 2.0 times the cost of electricity of a 100% efficient pump. Both the heat transfer surface area and pumping power prices are evaluation program variables that are given a probability distribution, see Section 5.6.3. The optimum circulation velocity as a function of pumping power cost and heat exchanger surface area is presented in Figure 5.5 and the corresponding total cost in Figure 5.6.

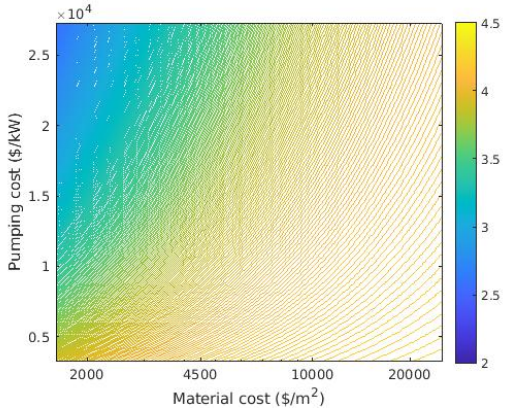

Figure 5.5: Optimum average velocity of $2^7LiF - BeF_2$ salt as a function of heat transfer surface area and pumping power cost.

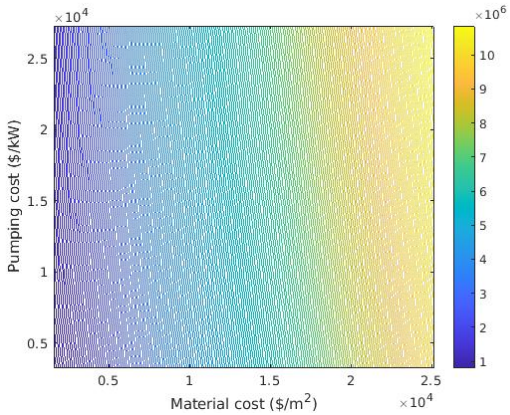

Figure 5.6: Minimum total cost ($) using $2^7LiF - BeF_2$ salt as a function of heat transfer surface area and pumping power cost.

5.3 Geometry Considerations

5.3.1 Packing Arrangement of the Particles and Pebbles

The proposed design of the PB-FHR Mark 1 [48] calls for a pebble packing factor - the share of space occupied by the component constituent (the remainder is composed of matrix) - of 0.6. Fortunately SERPENT 2 has an algorithm that can output the location of stochastically dispersed spheres in a matrix, given inputs of the matrix geometry and radius of the dispersed sphere. In theory, through this algorithm, the geometry of reactors such as pebble-bed reactors with double heterogeneity can be accurately modelled. However it is difficult - this study failed - to achieve the 0.6 packing fraction (PF) for models with many pebbles; the 200 MWth reference design has 375,000 fuel pebbles in the core. Therefore a regular distribution of pebbles was used; the common crystal structures hexagonal-close-packed (HCP) and face-centred-cubic (FCC) were tested. It is unclear which arrangement is more similar to a fully accurate disperse distribution and the k_{eff} difference between HCP and FCC was within the margin of error. A HCP arrangement was arbitrarily selected for subsequent analysis.

The k_{eff} difference between disperse and HCP packing was tested at PFs lower than 0.6 that allowed for disperse packing. The objective was to use these results to determine if the k_{eff} difference at a PF of 0.6 can be predicted or if it is small enough that it can be assumed negligible. It was found that the difference is very small, within the margin of error, at all 4 PF values tested. The results are shown in Figure 5.7. The disperse and HCP packing resulted in a slightly different total number of pebbles - this difference is negligible but has been noted in the Figure.

There are only 2976 to 6548 fuel particles inside each pebble with a PF of only 0.25 - 0.55 so they were able to be modelled stocastically. The same SERPENT 2 algorithm was used to set the location of the particles inside the pebbles.

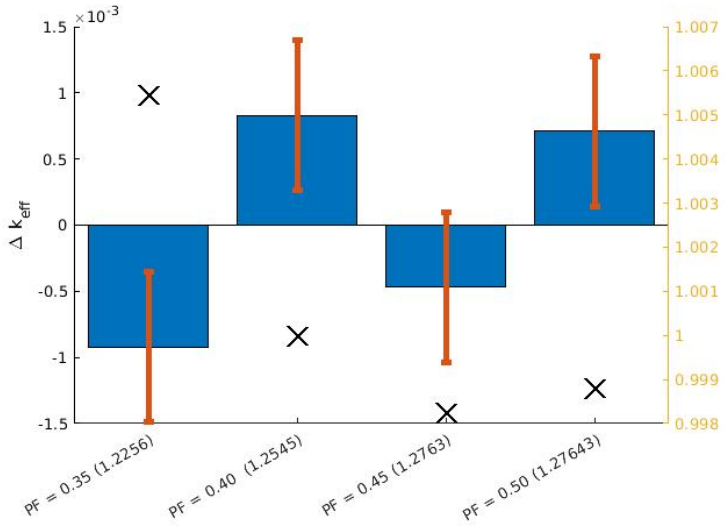

Figure 5.7: Disperse k_{eff} - HCP k_{eff} difference at various PFs. The thick blue bars indicate the k_{eff} difference, the thin orange bars the statistical error and the black X's: (num. of disperse pebbles)/(num. of HCP pebbles) (scale on right-hand-side of plot)

Table 5.5: TRISO fuel particle geometry. The fuel kernal is solid sphere. The surrounding layers are listed in order of radial distance.

Composition	Radius/thickness (μm)	Atomic density ($10^{24}/cm^3$)
Fuel kernel ($UC_{1.5}O_{0.5}$)	200	0.071653
Buffer	100	0.052645
Pyrolytic carbon	35	0.095262
Silicon carbide (SiC)	35	0.095448
Pyrolytic carbon	35	0.095262

5.3.2 Particle and Pebble Geometry

The spherical particle geometry consists of a fuel kernel surrounded by several layers. The temperature used for materials surrounding the fuel is approximately equal to the average fuel temperature and the densities are from [164] and [165]. The outer radii of each concentric layer of TRISO fuel used throughout the analysis is taken from the *Technical Description of the "Mark 1" PB-FHR* [48], see Table 5.5.

The pebble geometry is provided in Table 5.6. The pebble coating is the same high density as that of the graphite matrix. The Mark-1 FHR specifies a density of 1.54 g/cm^3, (0.07728 $10^{-24}/cm^3$) [166] for the inner graphite of the pebbles. The choice of a comparatively low density is to offset the high density of the TRISO particles and make the pebbles neutrally buoyant in the salt. The temperatures of the pebble graphite, coating, matrix and inner fill were all made to equal approximately the average fuel temperature, while the temperature of the reflector graphite is made to approximately equal the average salt temperature.

An aside on buoyancy considerations: Literature on the issues poised by differences in density between

Table 5.6: Fuel pebble geometry; layers are listed in order of radial distance.

Composition	radius/thickness (cm)	atomic density ($10^{-24}/cm^3$)
Inner graphite	1.25	–
Fuel layer graphite	1.4	0.08774
Outer graphite	1.5	0.08774

pebble fuel and the coolant fluid could not be found. Clearly buoyancy was a concern in the FHR design but it seems quite possible that buoyancy was achieved because it was very convenient to do so ($LiF-BeF_2$ and nuclear graphite have nearly the same density) rather than out of necessity. Coolant salts other than $LiF-BeF_2$ generally have higher densities, so neutral buoyancy may be difficult to achieve within the confines of this evaluation. One or both of the following strategies could be employed to increase the density of fuel (or graphite) pebbles:

1. Restrict the design space to high fuel concentrations. That would require high particle PFs, a thicker fuel layer and/or a larger fuel kernal.
   - Lower enrichment as required so long as sufficiently high burnup can be maintained.
2. Addition of a zirconium metal (density 6.5 g/cm^3) layer to pebbles

A Matlab program was built to stochastically assign each pebble as fuel or pure graphite such that the share of each pebble type is equal to exactly the predetermined level.[5] Radial plots of the reference reactor geometry with pure graphite pebbles comprising 37.5% of pebbles are shown in Figures 5.8a and 5.8b. An Axial plot is shown in Figure 5.9.

5.4 Neutronics Analysis: Methodology

The analysis consists of both depletion and static neutronics simulations. The fuel groupings and temperature differ in each. The objective is to produce accurate data sets of reactor performance metrics over a range of configurations such that the range covers the realistic potential design space.

Simulations were run at all stated geometric configurations for all five salts. Key performance parameters extracted were the k_{eff}, flux and pnl. At each of these configurations & salt compositions the k_{eff} at a higher salt temperature and higher fuel temperature were determined as well, in order to calculate the temperature reactivity feedback coefficients.

5.4.1 Depletion Simulations

The dual purpose of the depletion simulations is to determine the refuelling rate and fuel composition to be used in subsequent static simulations. To reliably determine these values, the reactor must be modelled close to its equilibrium state - a depletion simulation with uniform burnup producing the result shown in Figure

[5]The reactor core contains 375,000 pebbles, enough to ensure that differences between any pair of particular stochastic pebble distributions have a negligible impact on the results.

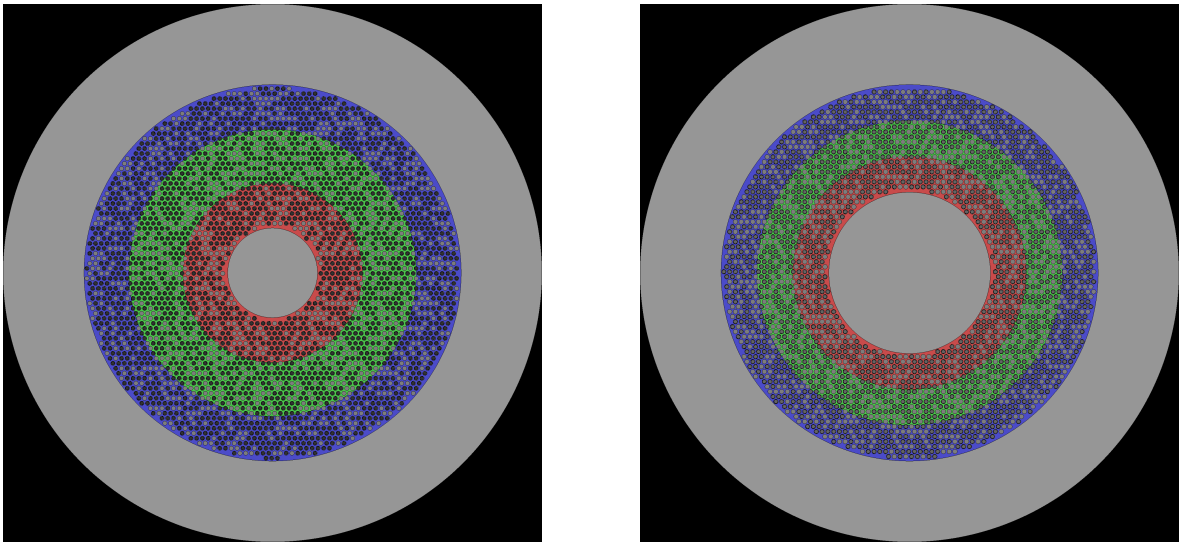

(a) Inner radius: 25 cm

(b) Inner radius: 45 cm

Figure 5.8: Radial view of the reference FHR geometry with pure graphite pebbles comprising 37.5% pebbles.

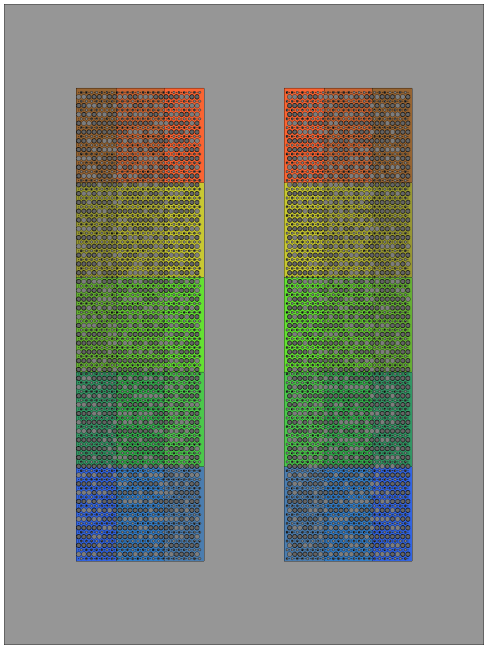

Figure 5.9: Axial view of the reference FHR geometry with pure graphite pebbles comprising 37.5% pebbles; inner radius of 25 cm.

5.10 would be untenable. There is a continuum of burnup levels among the fuel pebbles and the pebbles are assumed well mixed such that burnup does not vary as a function of macro-level position. This presents a modelling challenge; infinitely many depletion zones would be required to model the fuel composition with maximum precision. And a high number of depletion zones comes with challenges, higher simulation time, higher memory usage and more cumbersome file management. Further, the coolant salt temperature is a continuum that would also require infinitely many zones to achieve maximum precision - modelling that aspect faces the same challenges.

The accuracy of the key target values, refuelling rate and fuel composition, is much more dependant on the number of fuelling zones used rather than the salt temperature. Furthermore, discrepancies in k_{eff} due to a uniform temperature can be corrected in subsequent static simulations. Thus only one salt temperature was applied throughout the core.

The reactor fuel is divided into eight equally sized depletion zones (fuel groups). There are an equal number of pebbles in each group spread randomly throughout the core. The algorithm that sets the location of pebbles in the core (see Section 5.3.2) also stochastically assigns a fuel group to each pebble. The probability of a fuel pebble being assigned to each group is continually adjusted so that the final count is equal among all 8 groups. The initial loading is of entirely fresh fuel with no burnable absorbers and in the startup phase of the burnup calculation no fuel is added. Once a k_{eff} of 1.0 is reached, fresh fuel replaces the fuel in one of the eight groups. This is repeated 12 times until the fuel groups have reached a near equilibrium burnup distribution. A (stepwise) spectrum of fuel burnup is obtained that is randomly dispersed throughout the core.

The selection of eight depletion zones was largely driven by the trade-off between achieving the maximum fuel burnup and minimizing the number of times pebbles are cycled through the core and either replaced with fresh fuel or recirculated. The system must replace pebbles within a burnup range; the pebble with the highest burnup cannot always be the next one replaced. For this refuelling model, assuming complete, stochastic mixing of the pebbles and assuming an 'ideal' burnup distribution (see Figure 5.12), the number of fuel groups is equal to the average number of times a pebble exits the core and is evaluated for disposal/recirculation.[6] It is recognized that this assumption cannot be precisely true for a finite number of fuel groups/evaluations; the process of recirculation is expected to drive most of the mixing (no stirring or other mixing mechanism is analysed nor suggested).

A realistic maximum recirculation rate is difficult to determine, it was not found to be well established in literature or vendor designs - an insight in the recirculation rate is from the Xe-100 pebble-bed, gas-cooled reactor, pebbles are 'circulated through the core up to six times.' [167]. No serious impediments to a higher recirculation rate were conceived off and as the technology matures the recirculation rate should increase so a somewhat higher rate was selected. The benefits of an increased recirculation rate can be substantial,

[6] Assuming the burnup is effectively and consistently evaluated correctly from the radioactive signature

the expected average relative burnup of a removed pebble (R_b) as a percentage of the infinite zones case is provided in Equation 5.19. n is the number of fuel zones.

$$R_b = 100 \cdot (1 - \frac{1}{2 \cdot n}) \tag{5.19}$$

The computational resources available was also a consideration - more zones would have required trade-offs in other aspects of the simulation, such as the number of burnup steps. Refuelling could have been modelled with many more depletion zones and have pebbles from multiple high burnup zones refuelled, however this additional precision is not expected to yield substantially different results on any of the important metrics.

To obtain an accurate k_{eff} of a continuous fuel distribution using a model of discrete fuel compositions, the burnup to k_{eff} relation should be linear, more specifically $\frac{d^2k_{eff}}{dBu^2}$ should be small relative to $\frac{dk_{eff}}{dBu}$. As well the deviation from $\frac{keff_{fuel1}+keff_{fuel2}}{2} = keff_{fuel1\&fuel2}$, what can be termed as superposition, introduces error in a piecewise fuel distribution model. Fortunately, as shown in Figure 5.10, $\frac{d^2k_{eff}}{dBu^2}$ is very small relative to $\frac{dk_{eff}}{dBu}$ beyond the first few days. Superposition of k_{eff} does however introduce some error. A piecewise fuel distribution model of continuous burnup is mathematically analogous to numerical integration using the trapezoidal rule. The method quickly approachs the true value as the number of samples taken increases ($O(x^2)$), so for this situation with nearly linear functions (low $\frac{d^2k_{eff}}{dBu^2}$) and tolerable deviation from superposition, 8 groups is sufficient from a determination of k_{eff} perspective as well.

To simulate continuous refuelling, the depletion zones are refuelled one at a time. SERPENT 2 has a useful refuelling routine for applying changes during the depletion simulation; the 'rfr' and 'rfw' cards, that read and write the fuel compositions. A simplified example of the refuelling scheme is provided in Figure 5.11 and the results from SERPENT simulations are in Section 5.5.1. The refuelling procedure is as follows:

1. An initial depletion simulation is run with entirely fresh fuel and a series of increasing time steps - the goal is for the k_{eff} of the final time step to be close to 1 without going below 1. The fuel composition is saved.
2. The output file is then read and k_{eff} extracted.
3. A depletion simulation (depletion step) is run with a low cumulative burnup, the total burnup depends on the salt composition and increases with fuel concentration. The day steps used depend on whether refuelling takes place.
    - If $k_{eff} > 1$, the current depletion run uses the fuel composition of the last depletion step of the previous run in each zone. Two relatively long time steps are used.
    - If $k_{eff} < 1$, the current depletion run copies the fuel composition for all but the one of the zones. The highest burnup fuel is replaced with fresh fuel. Shorter, increasing time steps are used because of the added fresh fuel.[7]

[7] Saturating FPs in the fresh fuel have a significant effect but generally approach saturation within hours/days. As a result depletion curves rapidly approach a linear relation, as shown in Figure 5.10.

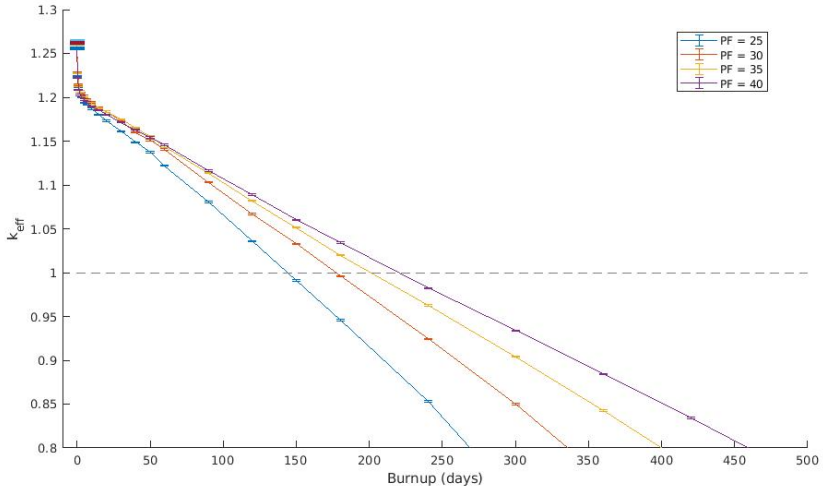

Figure 5.10: Depletion simulation of an FHR with FLiBe salt at multiple Packing Fractions. Contains only one depletion zone and begins with fresh fuel.

- If it is the 12th refuelling, where the final zone is refuelled, the program exits the loop. At the conclusion of this step the reactor can be considered nearly at equilibrium.

4. Repeat from step 2.

The depletion simulations are too computationally expensive to run for each configuration, so they were run for each salt and only as a function of Packing Fraction (PF). The reactor fuel composition and performance metrics are similar for two reactors with the same fuel concentration. Fuel concentration can be varied within the confines of this study either through the PF or the blank %. The PF was selected with the blank % held at 0%. The inner radius is expected to have little impact and the median dimension of 35 cm was used.

The target metrics, refuelling time and fuel composition, in the ideal case both require equilibrium fuel composition. However the fuel composition only approaches equilibrium, there is no clear point where equilibrium is reached. This scheme refuels when k_{eff} falls below 1 and that does not necessarily happen at regular intervals. An equilibrium fuel composition could be made with a more complex scheme involving burnable absorbers but that comes with trade-offs as well.

Nonetheless equilibrium is approached quite quickly, 12 total refuellings produces a burnup distribution that closely matches the ideal case, see Figure 5.12. No appreciable (potential) concerns were found or conceived of that result from a small deviation from equilibrium for the same k_{eff}. The fuel composition used in subsequent static simulations is the composition from the end of the depletion simulation of the 12^{th} and final refuelling. The refuelling rate can much more easily be adjusted to account for the deviation from equilibrium. The Δk_{eff} with refuelling is dependent on the burnup of the fuel it is replacing, see Figure 5.13. Therefore using the highest burnup fuel of the final refuelling step and the highest burnup fuel of the ideal

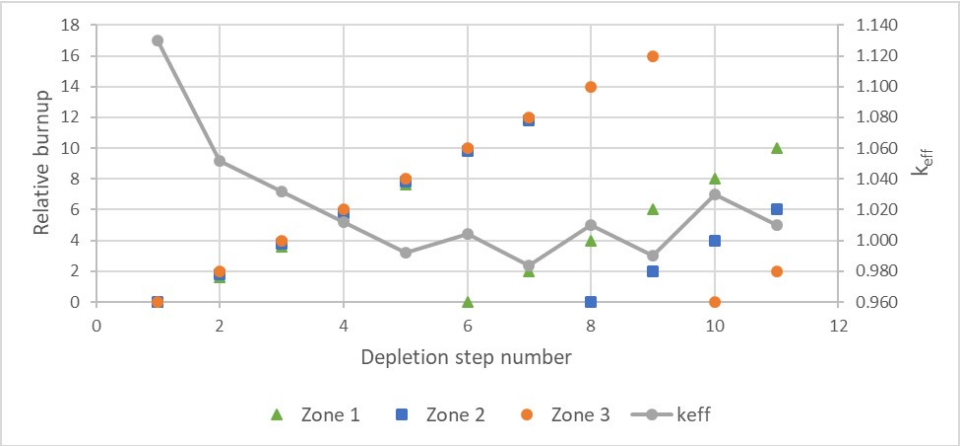

Figure 5.11: Simplified example refuelling scheme with 3 depletion zones instead of 8. Arbitrary values only, for demonstration purposes. Notice how the k_{eff} increases when the burnup of one of the zones goes back to zero (is refuelled). Initially the zones have the exact same burnup but the displayed burnup is slightly shifted to highlight the overlapping data points.

case an adjustment to Δk_{eff} is made.

$$\Delta k_{eff,r,eq} = \Delta k_{eff,r} \cdot \frac{m \cdot Bu_{eq} + b}{m \cdot Bu + b} \tag{5.20}$$

Where m and b are the coefficients for the equation of the line of ideal burnup. From $\Delta k_{eff,r,eq}$ the refuelling time can be calculated using Equation 5.21.

$$t_r = n \cdot t_s \cdot \frac{|\Delta k_{eff,no}| + \Delta k_{eff,r}}{|\Delta k_{eff,no}|} \tag{5.21}$$

Where t_r is the refuelling time of a fuel refuelling cycle, t_s is the length of each simulated depletion step and n represents the number of fuel groups. $\Delta k_{eff,r}$ is the adjusted k_{eff} at the end of the final refuelling step minus k_{eff} at the end of the step before the final refuelling step (i.e. $k_{eff,r,eq}$). And $\Delta k_{eff,no}$ is the Δk_{eff} over a time period of length t_s with no refuelling.

Additionally, the predictor-corrector was used with the following number of histories:

- Regular: 500 active, 40 inactive cycles of 5000 neutrons
- Final two steps: 2000 active, 100 inactive cycles of 5000 neutrons

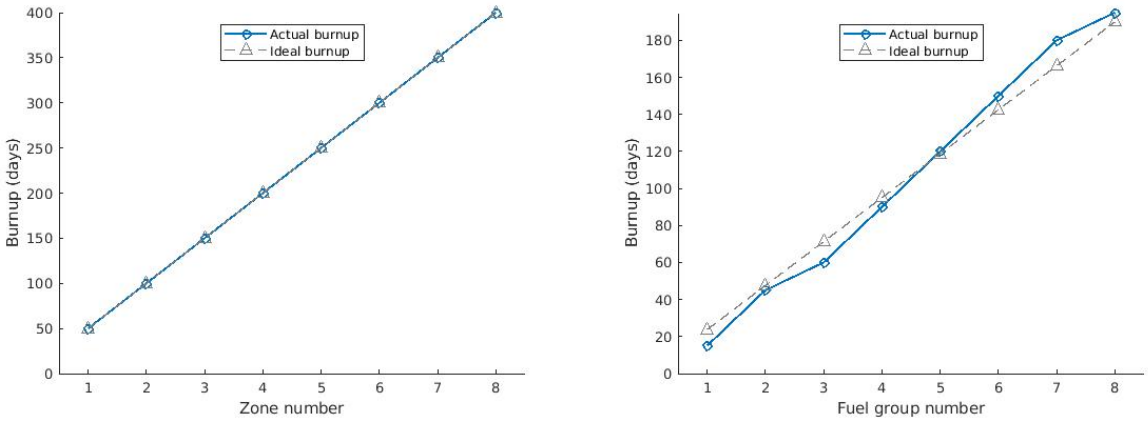

(a) FLiBe PF 35

(b) NaBeF PF 40

Figure 5.12: Burnup distribution among the 8 zones in the refuelling model. The simulated (actual) is compared to the ideal burnup at equilibrium. The average burnup is equal in each case. See Appendix A.1 for the results of the other PFs.

Depletion step time selection

The required accuracy of the intermediate burnup steps is quite low because the refuelling rate is calculated by the change in k_{eff} between two steps once equilibrium is reached, see Equation 5.21, rather than the time required to complete 8 refuellings. Moreover, the refuelling time for each configuration is adjusted as per the results of the corresponding static simulations. Thus short burnup steps are not required. An attempt is made to select burnup steps that are reasonably close to the minimum time between refuelling steps, the steps used are listed in Table 5.7.

The selection of the day steps before the first refuelling is inherently arbitrary. The steps should be progressively increasing in length since the effect of neutron poisons (such as Xe) and fissile isotopes (such as Pu) initially changes quickly but over time the concentrations move toward saturation/equilibrium.

5.4.2 Static Simulations

In the static simulations where an accurate determination of k_{eff} is the primary goal, the reactor geometry and materials should be modelled more precisely. The active portion of the core is split into zones of uniform salt temperature; three in the radial and five in the axial direction for a total of 15 zones. The geometry with the temperature zones is shown in Figures 5.1 and 5.2. Ideally each of the 15 zones would have a unique fuel temperature as well, however available computational resources did not allow for 15$\cdot$ (8 fuel groups per zone) $=$ 120 unique fuel compositions. 8 fuel compositions, taken from the depletion calculations are at minimum required, but 16 fuel groups were used, 8 at a higher and 8 a lower temperature. More on this in the discussion on fuel temperature in Section 5.1.2, specifically note Figure 5.3. Simulations are performed across the entire evaluated design space:

1. Packing Fraction (PF) of the particles inside the pebbles [%]:

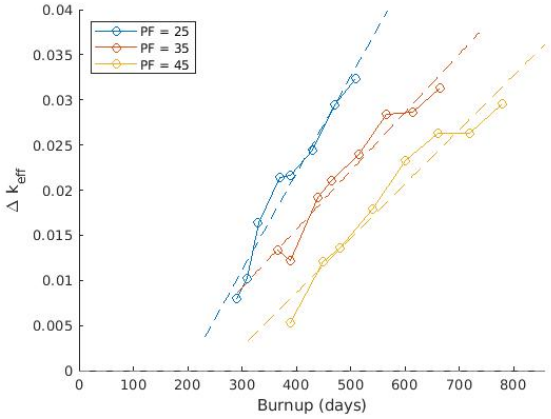

(a) FLiBe PF 35

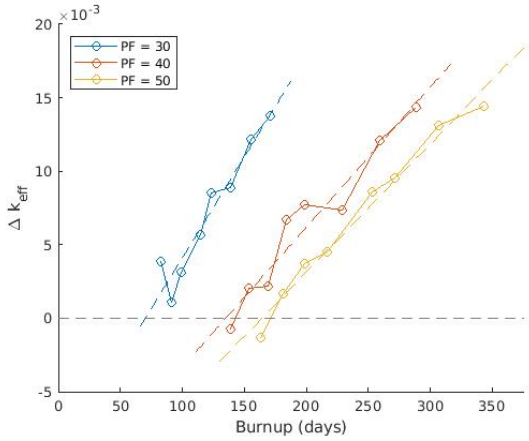

(b) NaBeF PF 40

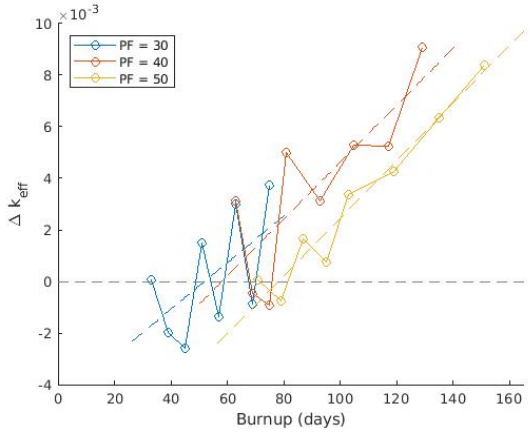

(c) NaZrF PF 40

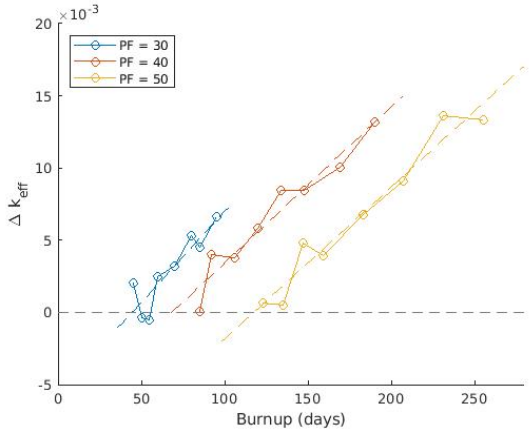

(d) LiNaBeF PF 40

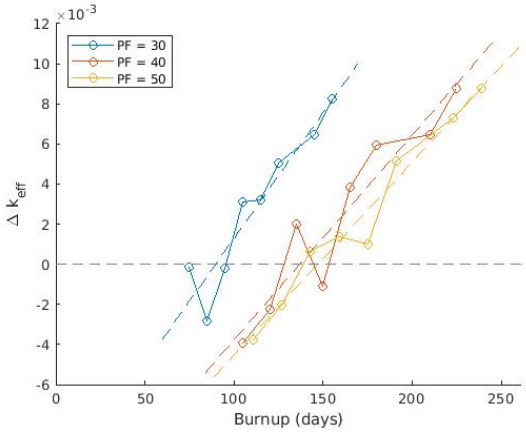

(e) LiNaZrF PF 40

Figure 5.13: Δk_{eff} as a function of the burnup of the replaced fuel.

Table 5.7: Total depletion per step (days).

Salt	PF	Initial burnup	Burnup/step
$2LiF-BeF_2$	25	70	20
	35	70	25
	45	70	30
$NaF-BeF_2$	30	19	8
	40	19	15
	50	19	18
$NaF-ZrF_4$	30	15	6
	40	15	8
	50	15	10
$LiF-NaF-BeF_2$	30	12	5
	40	17	7
	50	17	12
$LiF-NaF-ZrF_4$	30	15	10
	40	15	15
	50	15	18

- $LiF-BeF_2$: 25, 30, 35, 40, 45
- Other salts: 30, 35, 40, 45, 50

2. Share of graphite pebbles [%]: 0, 12.5, 25.0, 37.5

3. Diameter of the inner radius [cm]: 25, 35, 45

The static simulations involve looping through all configurations and each time concatenating an input file comprised of numerous sections, including:

1. The 15 salt compositions, one for each zone; temperature is a function of zone location and inner radius, and salt density is a function of temperature.

2. The location of the particles and pebbles.

3. Fuel composition:
   - Determine the depletion calculation with the core fuel concentration that is closest to the configuration under analysis.
   - Import the fuel composition from the end of the last refuelling step of the associated depletion simulation.
   - Use the imported the fuel composition and replace the temperature with the high or low (or a combination of both) fuel temperature value(s) depending on the zone.

Additionally a so called 'initial' run of each configuration is conducted to find the power in each zone. It assumes a uniform salt to fuel temperature difference and uses a low neutron population of 200 active cycles of 5000 neutrons. The input file is then rebuilt, except the zone powers from the 'initial' run are used to calculate the temperatures in each zone. 2000 active and 100 inactive cycles of 5000 neutrons are used for

the final run.

The fuel compositions used in the static simulations are taken from the associated depletion simulation with the closest matching fuel density.[8] Figure 5.14 shows the fuel composition used in each configuration.

The calculation of the time required to refuel the entire core (8 refuellings) use a combination of depletion and static simulation data. It is calculated according to Equation 5.22.

$$t_r = n \cdot t_i \frac{\Delta k_{eff,r} + \Delta k_{eff,sd}}{\Delta k_{eff,no}} \tag{5.22}$$

Where n representing the number of fuel groups is equal to 8, t_i is the burnup time per step, $\Delta k_{eff,r}$ is the Δk_{eff} due to a refuelling, $\Delta k_{eff,sd}$ is the Δk_{eff} between the static simulation and the corresponding depletion simulation (the values in Figure 5.14) and $\Delta k_{eff,no}$ is the reduction in Δk_{eff} over one time step.

This calculation is performed with t_i, $\Delta k_{eff,r}$, $\Delta k_{eff,sd}$ and $\Delta k_{eff,no}$ data from the depletion simulation with the nearest core fuel concentration (the colour groups in Figure 5.14). The same calculation is performed at the second closest fuel concentration, and the final refuelling time, t_r, is actually a linear interpolation of the two.

5.5 Neutronics Simulations Results

The objective of the neutronics simulations is to provide accurate reactor performance data necessary for the evaluation program. The reactor start-up procedure is not considered and the data generated is intended to represent the reactor in the equilibrium state. Depletion simulations are conducted primarily to determine the steady-state fuel composition and refuelling rate. The fuel composition from the depletion simulations is used in static simulations. The static simulations are run at the expected temperatures and at higher salt/fuel temperatures in order to calculate temperature reactivity feedback coefficients (TRFC).

5.5.1 Depletion Simulations: Results and Discussion

The effective multiplication factor (k_{eff}) over the duration of the depletion simulation, is shown in Figures 5.15 to 5.19 for all configurations. There are 12 piecewise refuellings that occur at burnup levels indicated by the vertical dashed lines. Notice that refuellings takes place if the k_{eff} at the end of the previous step is < 1. Each data point represents the beginning/end of a burnup step.

Because a maximum of one fuel group is refuelled per step, by step 8 each fuel group has a unique burnup level. Beyond step 8, the burnup among the fuel groups quickly approaches the equilibrium (ideal) distribution, Figure 5.12 demonstrates this by comparing the ideal distribution to that after 12 refuelling cycles. The change in the fuel group burnup distribution over time is shown in Figure 5.20. The burnups shown

[8] If the difference between two depletion simulations is equal, the higher fuel density composition is selected.

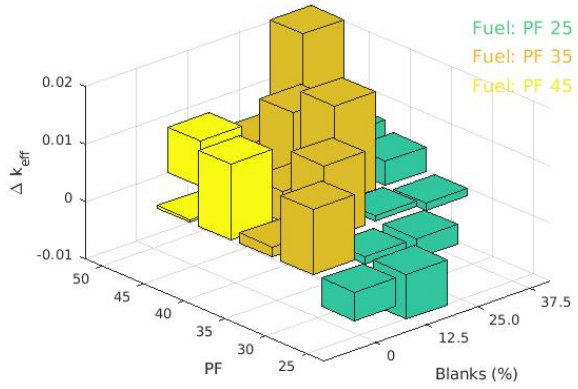

(a) FLiBe PF 35

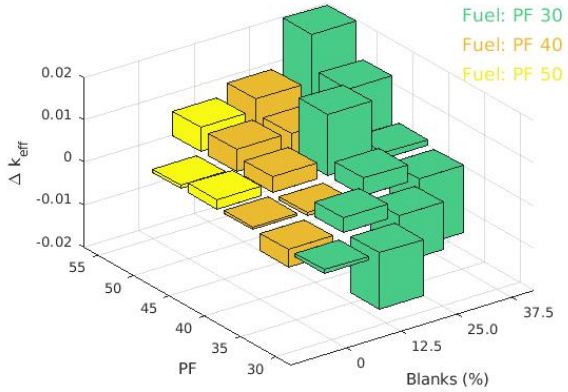

(b) NaBeF PF 40

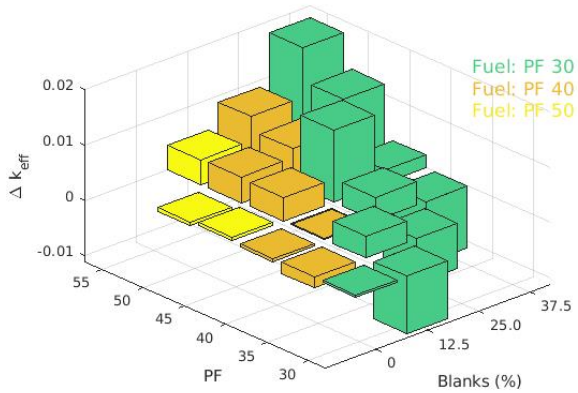

(c) NaZrF PF 40

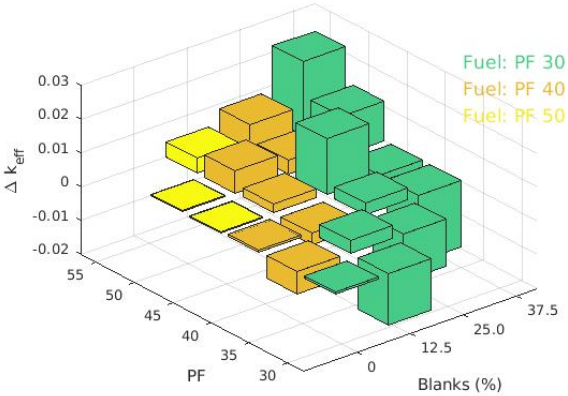

(d) LiNaBeF PF 40

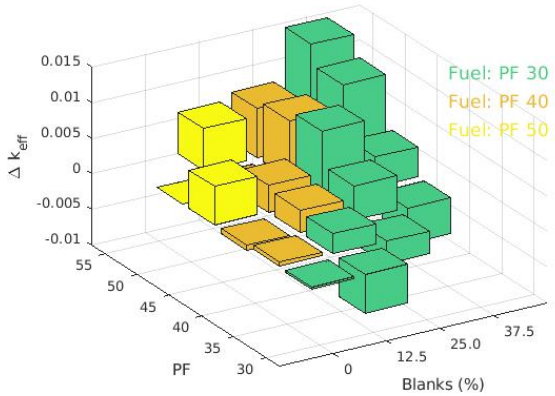

(e) LiNaZrF PF 40

Figure 5.14: Static simulations compared to the corresponding depletion simulation step from which the fuel composition is taken. The bar height shows the Δk_{eff} value, as a function of blanks and PF. The color shows which PF was used to generate the fuel composition. Data is from an inner radius of 35 cm.

Table 5.8: Key depletion simulation results.

Salt	PF	4^{st} refuel	12^{th} refuel	Refuelling period
$2LiF-BeF_2$	25	370	650	280
	35	515	865	350
	45	540	1020	480
$NaF-BeF_2$	30	115	235	120
	40	184	379	195
	50	217	469	252
$NaF-ZrF_4$	30	51	105	54
	40	81	177	96
	50	95	191	96
$LiF-NaF-BeF_2$	30	60	135	75
	40	120	260	140
	50	159	351	192
$LiF-NaF-ZrF_4$	30	105	215	110
	40	150	315	165
	50	159	335	176

Table 5.9: Daysteps used in each depletion run.

Salt	PF	Burnup steps	
$2LiF-BeF_2$	all	Initial: [0.5 1 1.5 2 2 3 5 5 10 10 15 15 20]	
$NaF-BeF_2$	all	Initial: [0.5 1 1.5 2 2 3 4 5]	
$NaF-ZrF_4$	all	Initial: [0.5 1 1.5 2 2 3 5]	
$LiF-NaF-BeF_2$	all	Initial: [0.5 1 1.5 2 2 3 5]	
$LiF-NaF-ZrF_4$	all	Initial: [0.5 1 1.5 2 2 3 5]	
Salt	PF	Refuel	No refuel
$2LiF-BeF_2$	25	[0.5 1 2.5 6 10]	[8 12]
	35	[0.5 1 2.5 5 7 9]	[10 15]
	45	[0.5 1 2.5 5 8 13]	[13 17]
$NaF-BeF_2$	30	[0.5 1 1.5 2 3]	[3 5]
	40	[0.5 1 2 3 4 4.5]	[6 9]
	50	[0.5 1 2 3.5 5 6]	[8 10]
$NaF-ZrF_4$	30	[0.5 1 1.5 3]	[2 4]
	40	[0.5 1 1.5 3]	[2 4]
	50	[0.5 1 2.5 4]	[3 5]
$LiF-NaF-BeF_2$	30	[0.5 1 1.5 2]	[2 3]
	40	[0.5 1 2 3.5]	[3 4]
	50	[0.5 1 2 3.5 5]	[5 7]
$LiF-NaF-ZrF_4$	30	[0.5 1 2 2.5 4]	[4 6]
	40	[0.5 1 2.5 5 6]	[6 9]
	50	[0.5 1 2.5 5 7]	[7 9]

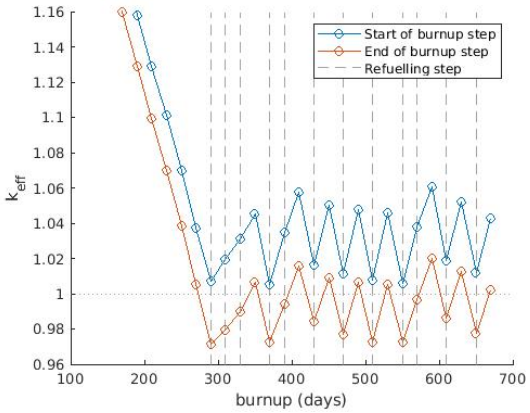

(a) FLiBe PF 25

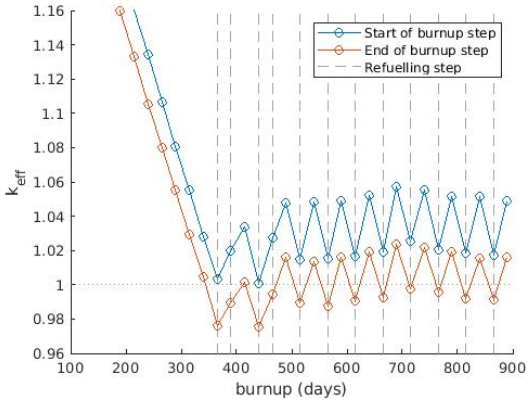

(b) FLiBe PF 35

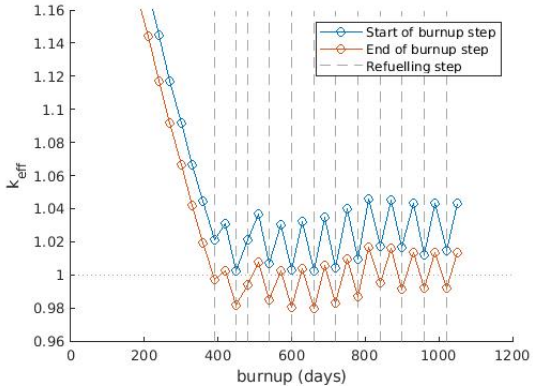

(c) FLiBe PF 45

Figure 5.15: k_{eff} of FLiBe salt configurations over the duration of the depletion simulation with 12 piecewise refuellings.

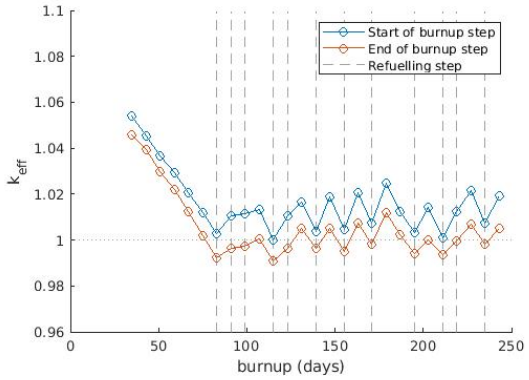

(a) NaBeF PF 30

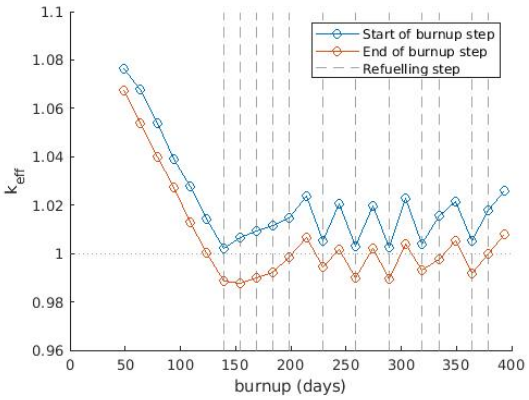

(b) NaBeF PF 40

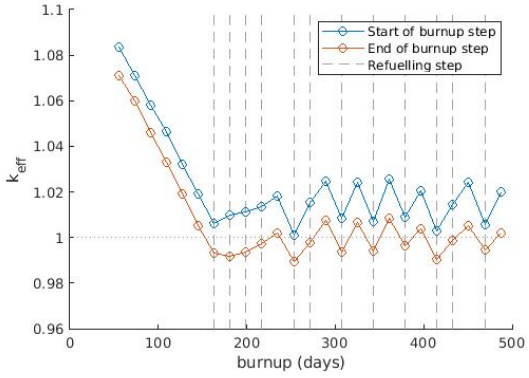

(c) NaBeF PF 50

Figure 5.16: k_{eff} of NaBeF salt configurations over the duration of the depletion simulation with 12 piecewise refuellings.

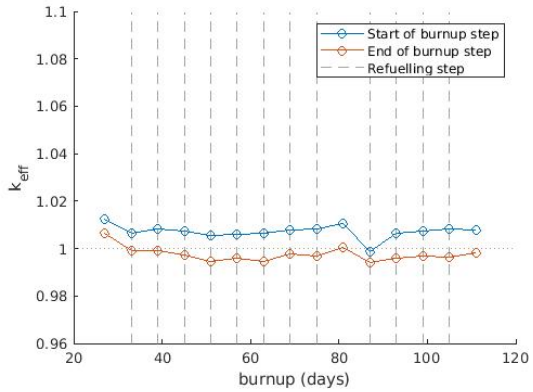

(a) NaZrF PF 30

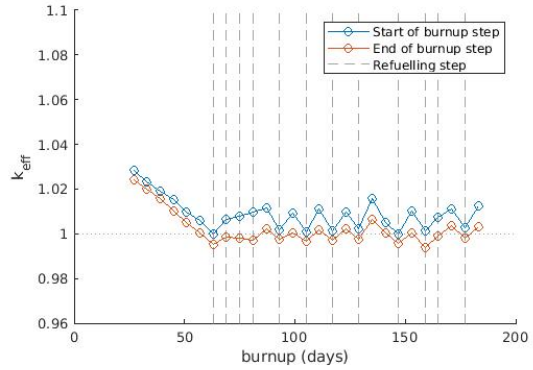

(b) NaZrF PF 40

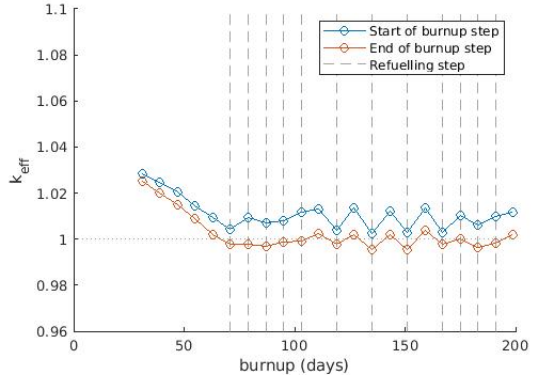

(c) NaZrF PF 50

Figure 5.17: k_{eff} of NaZrF salt configurations over the duration of the depletion simulation with 12 piecewise refuellings.

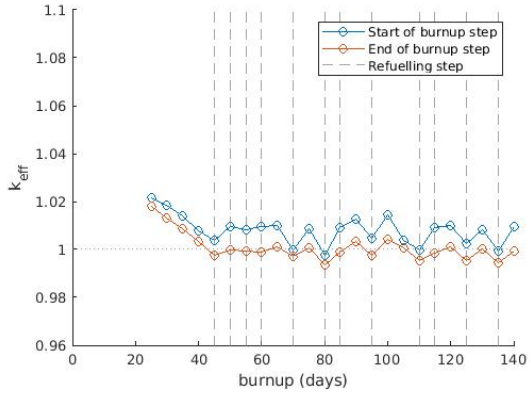

(a) LiNaBeF PF 30

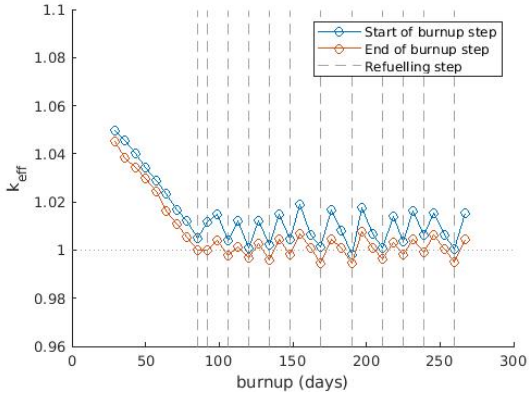

(b) LiNaBeF PF 40

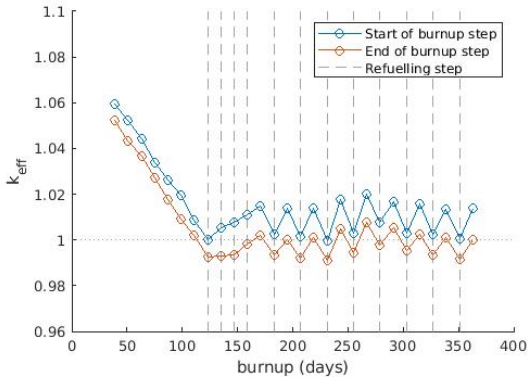

(c) LiNaBeF PF 50

Figure 5.18: k_{eff} of LiNaBeF salt configurations over the duration of the depletion simulation with 12 piecewise refuellings.

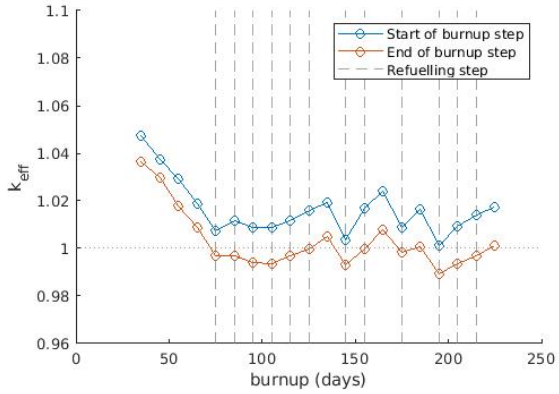

(a) LiNaZrF PF 30

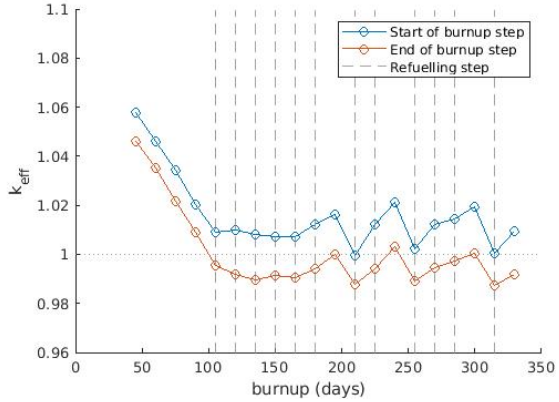

(b) LiNaZrF PF 40

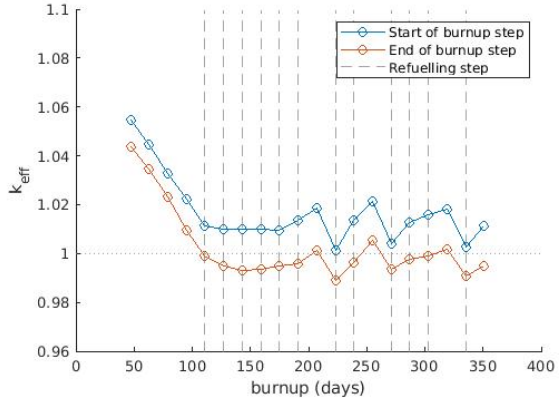

(c) LiNaZrF PF 50

Figure 5.19: k_{eff} of LiNaZrF salt configurations over the duration of the depletion simulation with 12 piecewise refuellings.

are those immediately preceding refuelling. The results are only provided for each salt at the median PF, the remaining results are in Appendix A.1. The change in the average burnup is largely a function of the maximum burnup of the previous refuelling. The average burnup demonstrates a pattern of decreasing a small amount from the initial onset of refuelling as it moves toward equilibrium.

5.5.2 Static Simulations: Results and Discussion

So far, the only output extracted from the simulations is the k_{eff}, it determines the excess reactivity, maximum burnup and temperature reactivity feedback coefficients. Additionally, the prompt neutron lifetime (PNL) and thermal flux were found for all configurations at equilibrium. The thermal flux is relevant to the evaluation in order to estimate Xe poisoning and ^{6}Li parasitic absorption (and consequentially tritium build up) of Li bearings salts. The PNL is also factored into the evaluation, the results are shown in Figures 5.21 and 5.22. The inner radius size has the same affect in all salts as it does in FLiBe. It is assumed that the PNL difference between adjusting the level of moderation with the share of graphite pebbles and by way of the particle PF is negligible.

The impact that the inner radius has on k_{eff} is shown in Figure 5.23. Each data point is the average of blank shares 0, 12.5% and 25.0%. The error shown is ± 1.0 SD of the individual run. The impact is fairly consistent among the salt compositions. The k_{eff} gain that results from a smaller inner radius has been found to be a function of PF. The higher the PF, the lower the Δk_{eff}. It is suspected that this pattern is due to the differential moderation; at a low PF and high overall moderation, the relatively low moderation gain of the large inner radius provides little benefit, while at low overall moderation the moderation gain provides more of a neutronic efficiency gain.

The impact that the number of graphite pebbles (blanks) has on k_{eff} is shown in Figure 5.24. The data presented is at an inner radius of 35 cm and the k_{eff} values shown are relative to the k_{eff} of the depletion run from which the fuel composition is taken. The data presented is only for configurations with a fuel composition taken from the corresponding depletion run with a PF of 30 (PF of 25 for $LiF - BeF_2$), notice how the data points shown match the PF 30 data points in Figure 5.14. Additional results of all salts for a PF of 40 are shown in Appendix A.1. There is a clear trend, in both PF 30 and PF 40, of increasing k_{eff} with the number of blanks, holding the overall core-average fuel density constant.

5.5.3 Temperature Reactivity Feedback Coefficients

There are two important temperature reactivity feedback coefficients (TRFCs or TCs) to analyse in a salt-cooled design; the coefficient of the salt and of the fuel. The TC of the salt includes the change in k_{eff} due to the temperature of the salt but more consequently the density. The temperature coefficient of fuel is governed by the change in Doppler broadening with temperature.

The temperature of the fuel and the coolant are different and the difference can vary. Thus it is preferable that

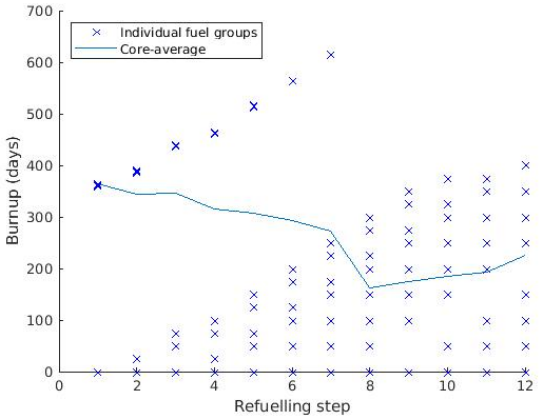

(a) FLiBe PF 35

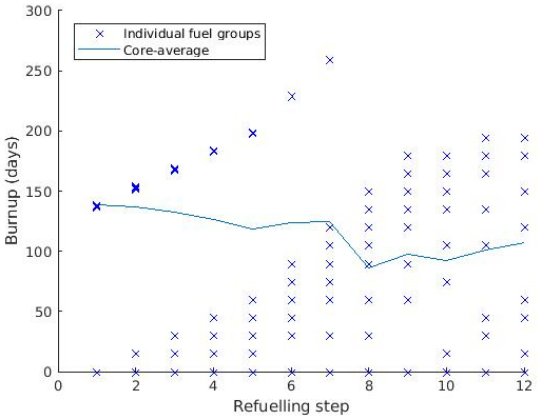

(b) NaBeF PF 40

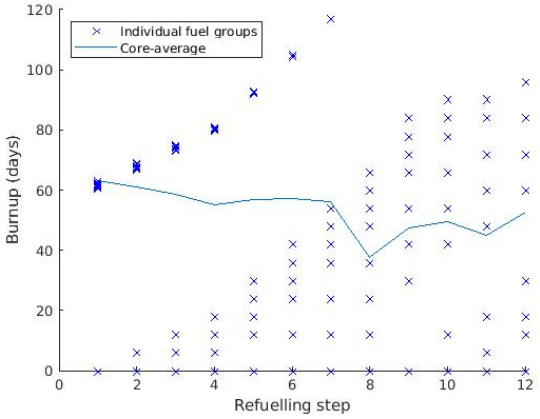

(c) NaZrF PF 40

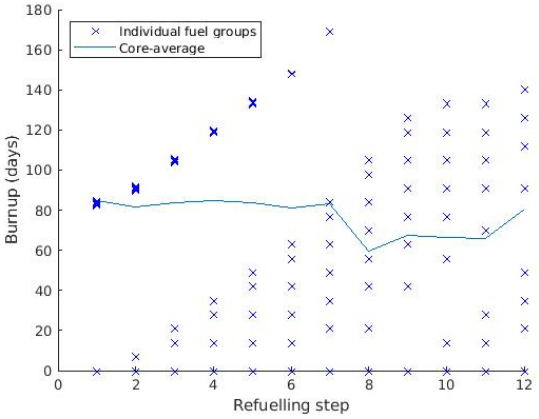

(d) LiNaBeF PF 40

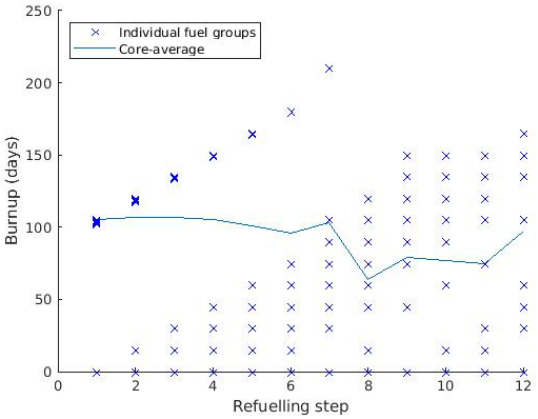

(e) LiNaZrF PF 40

Figure 5.20: Burnup of each group as the simulation progresses. Additional results for the other PFs can be found in Appendix A.1. Fuel groups with the same burnup are slightly offset to indicate that multiple groups occupy the data point.

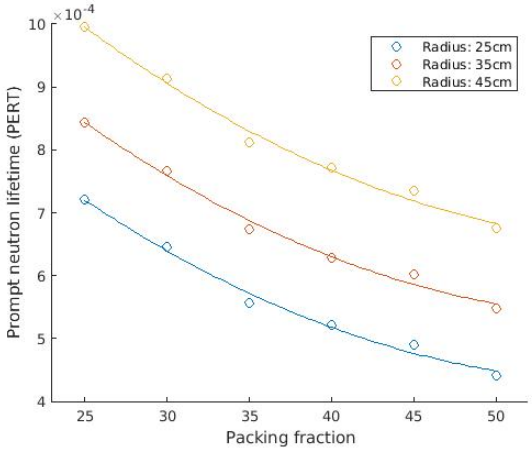

Figure 5.21: PNL as a function of the particle PF and inner radius diameter for FLiBe salt with 12.5% blanks, calculated using the PERT method.

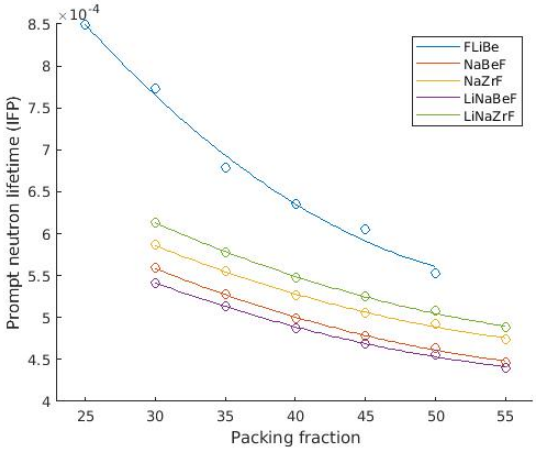

Figure 5.22: PNL of each salt as a function of the particle PF at an inner radius of 35 cm with 12.5% blanks, calculated using the IFP method.

both coefficients are negative rather than just the sum of the coefficients as the fuel and coolant temperature do not necessarily vary in tandem. The results for all five salt analysed are shown in Figures 5.25 to 5.32. A comprehensive breakdown of the results is provided for salts $LiF-BeF_2$ and $NaF-BeF_2$. $LiF-BeF_2$, the reference, well-studied composition, is examined at unique data points and has TCs that differ significantly from the alternatives. The alternative salts generally have very similar TC behavior; only the comprehensive results of one composition, $NaF-BeF_2$, is provided to demonstrate key trends.

The salt density effects k_{eff} in two ways; the salt is a parasitic absorber thus reducing k_{eff} but it also, at least somewhat, has a moderating effect. The moderating effect tends to, all else being equal, increase k_{eff}. And the lower the moderation, the greater the change in k_{eff} due to additional moderation. Therefore, if the coolant salt has low parasitic absorption and the level of moderation is sufficiently low, the coolant salt TC can be negative. A negative density TC only occurred in configurations with $LiF-BeF_2$ salt. It is unsurprising that $LiF-BeF_2$ has the most negative density TC, as it has the lowest capture cross section and lowest average atomic weight. The density TCs were found to decrease with increasing core fuel concentration (particle PF), this matches expectations.[9]

The results were examined to determine whether the share of blank pebbles effects the TCs independent of the core fuel concentration. The results are shown both for the density and Doppler TCs at a inner radius of 35cm; $LiF-BeF_2$ salt in Figure 5.27 and $NaF-BeF_2$ in Figure 5.30. Both the density and Doppler TCs are similar enough, when controlling for core fuel concentration, to allow for TC determination as a function of fuel concentration rather than both particle PF and blank percentage separately.

[9]Configurations with low fuel concentration have high thermal flux and therefore relatively high parasitic absorptions in the salt, so changes in density have a large (positive) impact on the parasitic absorption portion of the density TC. At low fuel concentrations (high moderation) changes in density have only a small (or even negative) impact on the salt moderation portion of the density TC. Therefore the total density TC becomes less positive as the fuel concentration increases.

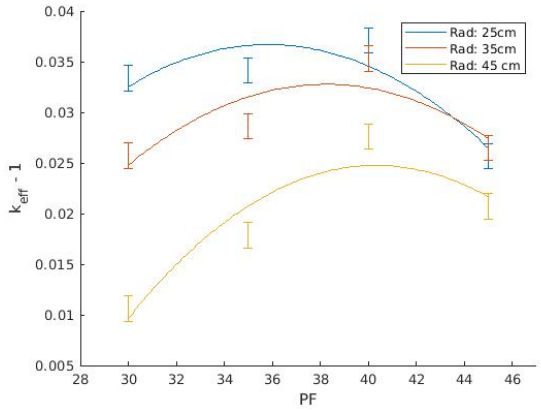

(a) FLiBe

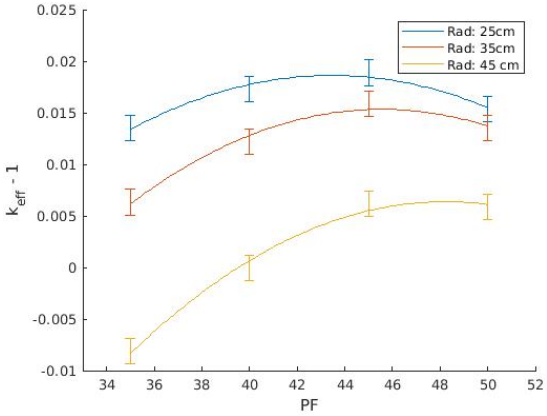

(b) NaBeF

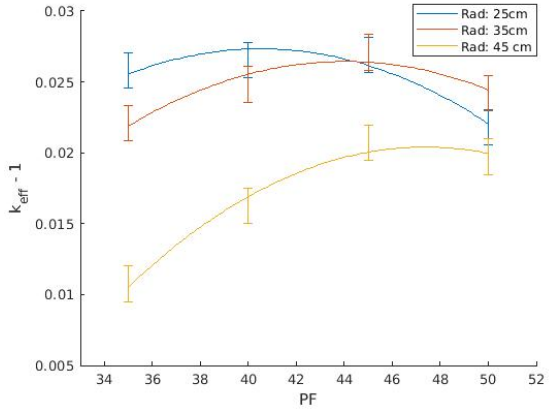

(c) NaZrF

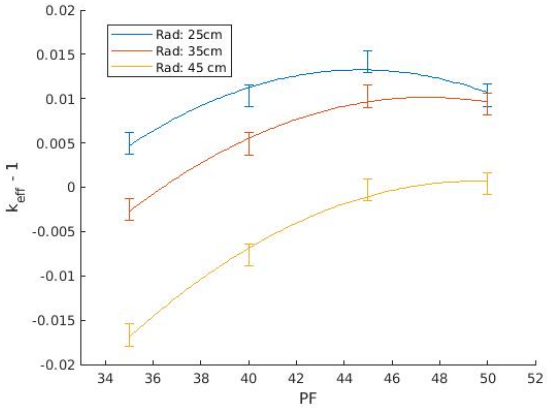

(d) LiNaBeF

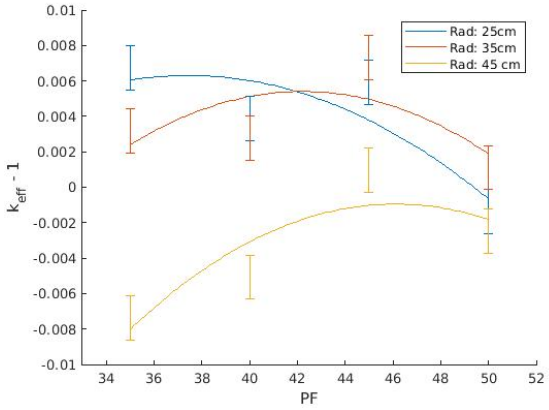

(e) LiNaZrF

Figure 5.23: k_{eff} as a function of particle PF for three inner radii dimensions.[a] Each data point is the average of blank shares 0, 12.5% and 25.0%.

[a] The interpolations are provided to highlight tends and add clarity, the shape of the curves do not have a physical meaning.

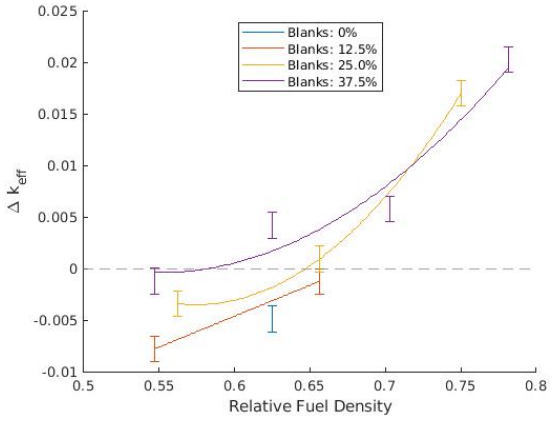

(a) FLiBe

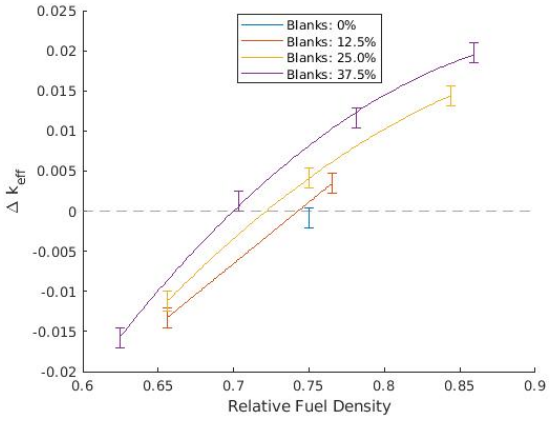

(b) NaBeF

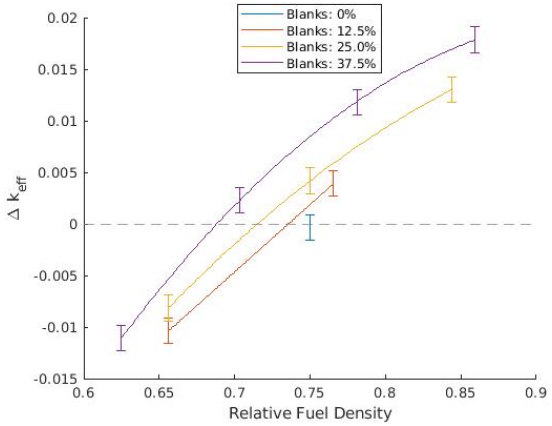

(c) NaZrF

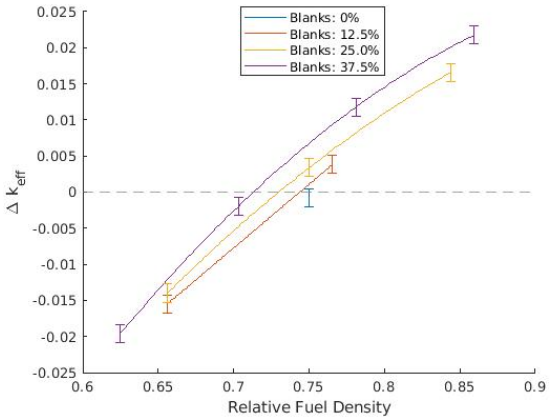

(d) LiNaBeF

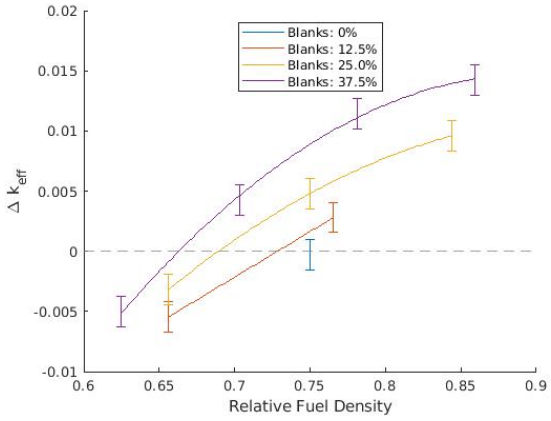

(e) LiNaZrF

Figure 5.24: k_{eff} as a function of fuel density with varying shares of blank pebbles, relative to 0% blanks and a PF of 40.[a] All data points shown for each salt only include only those generated with an alike fuel composition. The blank share used to generate all fuel compositions is 0%. FLiBe uses a fuel composition from a PF of 25, the other salts from a PF of 30.

[a]The interpolations are provided to highlight tends and add clarity, the shape of the curves do not have a physical meaning.

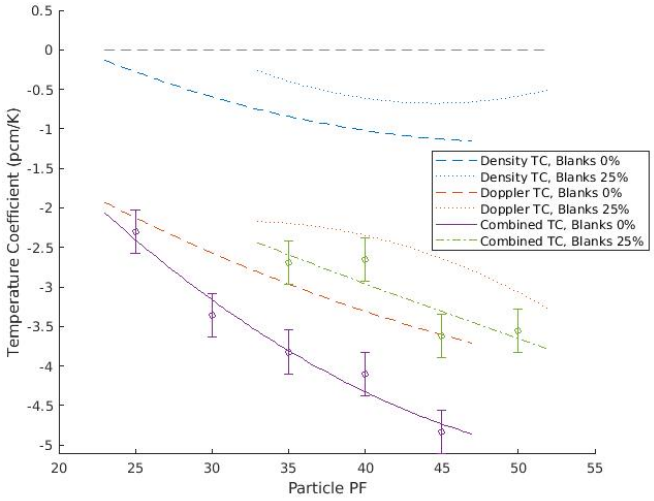

Figure 5.25: Temperature coefficients of $LiF - BeF_2$ coolant salt at 0% blanks and 25% blanks at a inner radius of 35cm.

The salt density and fuel Doppler TCs were calculated at each inner radius. For $LiF - BeF_2$ salt the density TCs showed little variation with inner radius, see Figure 5.26, any trend is difficult to separate from the statistical error. The fuel Doppler TCs did exhibit a small increase in the TC with increasing inner radius. It is suspected that is due to differences in moderation. For salt $NaF - BeF_2$ both the density and Doppler TCs showed a small increase with increasing inner radius, see Figure 5.29. The other alternative salts exhibited a very similar trend.

The TCs are calculated at equilibrium conditions. The reactor models are the same as those used in the static simulations, Section 5.4.2, with a fuel 8-group fuel composition extracted from the depletion simulations. The only difference is the salt/fuel temperature. Both the fuel and salt temperature change are uniform, the same ΔT is applied throughout the salt/fuel.

The salt and fuel TCs of each configuration are calculated with two sets of temperatures, the reference temperatures modelling normal operation and salt/fuel temperatures 400 ^{o}C higher. The k_{eff} at the reference temperatures is used because the k_{eff} values have already been calculated, Section 5.5.2. A 400 ^{o}C ΔT was selected to balance Monte Carlo statistic error that decreases and non-linearity that increases with rising ΔT. It was realized that the TCs are largely a function of salt composition and core fuel concentration. Therefore all salts and a range of fuel concentrations are analysed but only two blank shares were analysed, 0 and 25%. The same number of active (2000), inactive (100) and neutrons/cycle (5000) are simulated as in the static simulations. The k_{eff} error is $\approx$0.38 mk, resulting in Δk_{eff} error of 0.54 mk ($0.38 \cdot \sqrt{2}$) and a TC error of $\pm$0.13 pcm/K.

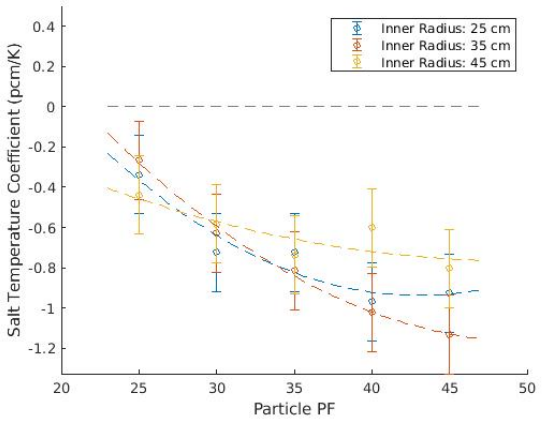

(a) Salt density TC, Blanks: 0%

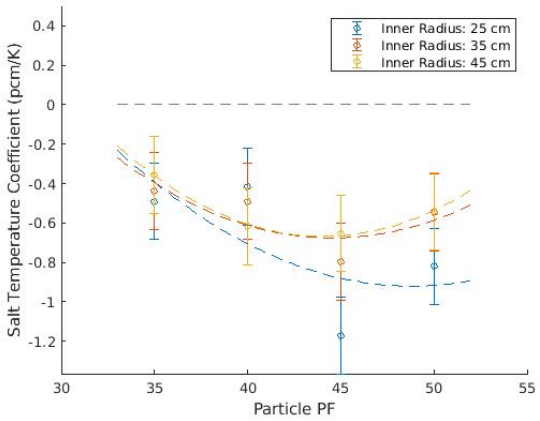

(b) Salt density TC, Blanks: 25%

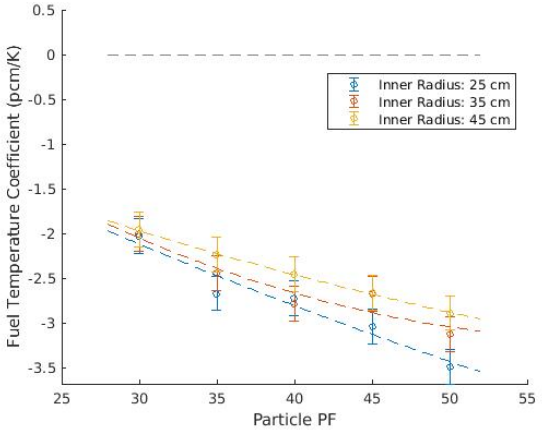

(c) Fuel Doppler TC, Blanks: 0%

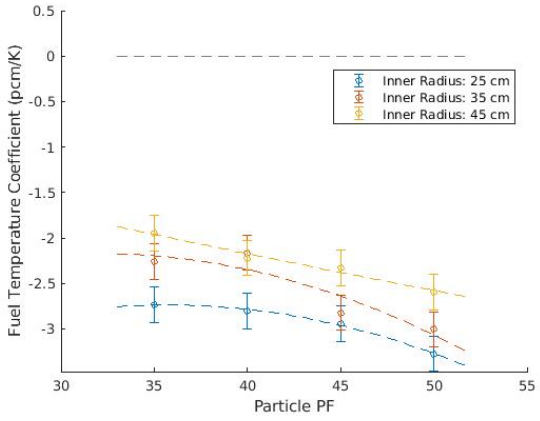

(d) Fuel Doppler TC, Blanks: 25%

Figure 5.26: Temperature coefficients of $LiF - BeF_2$ coolant salt; trend lines of the particle PF fitted for each of the inner radii.

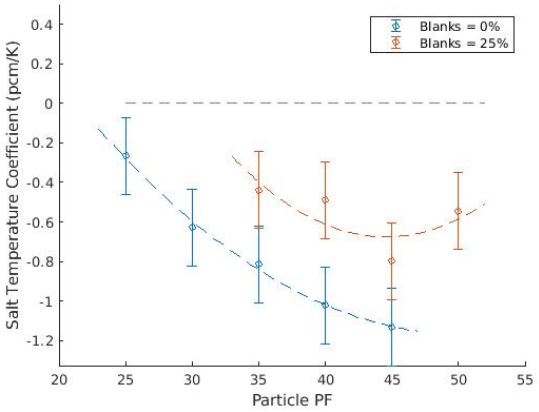

(a) Salt density TC, actual particle PF

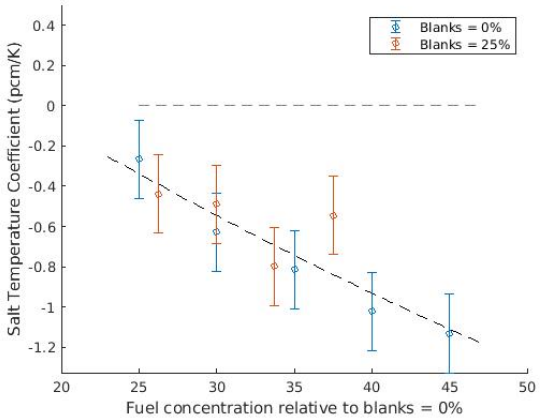

(b) Salt density TC, particle PF relative to 0% blanks

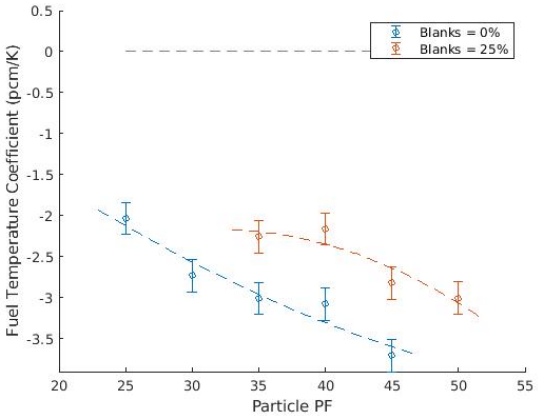

(c) Fuel Doppler TC, actual particle PF

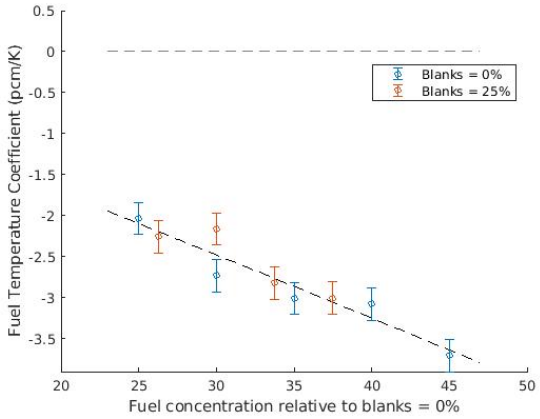

(d) Fuel Doppler TC, particle PF relative to 0% blanks

Figure 5.27: Blanks % comparison of the salt density and fuel Doppler temperature coefficients using $LiF - BeF_2$ coolant salt. Trends lines are fitted of the particle PF/fuel concentration.

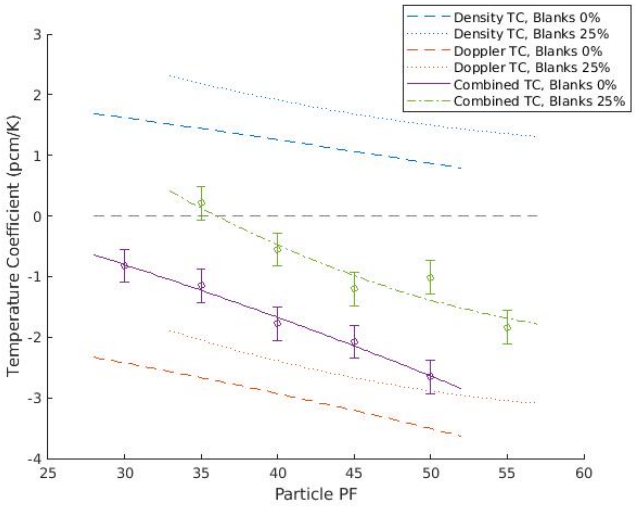

Figure 5.28: Temperature coefficients of $NaF - BeF_2$ coolant salt at 0% blanks and 25% blanks at a inner radius of 35cm.

Graphite TCs are not analysed. The amount of (fuel and blank) pebble graphite relative to fissile fuel is low in this reactor type, so the pebble graphite TC is not expected to be significant. The graphite of the outer and central reflector is also expected to have low magnitude TCs. Due to thermal mass and non-zero heat-transfer resistance, changes in temperature of the graphite occur with a significant time lag. Therefore graphite TCs can be considered to have only a minor impact in the initial stages of unplanned reactivity insertion transients. As it relates to TCs, reactor behaviour during unplanned reactivity insertion transients is the primary performance concern. The impact of graphite TCs in accident scenarios and also normal operation is expected to be relatively minor; moreover it is beyond scope and difficult to quantify the associated costs and risks.

5.6 Evaluation

5.6.1 Summary

This section discusses an evaluation of a range of geometric configurations and salt compositions of the salt-cooled, pebble-bed Fluoride High-temperature Reactor (FHR) design. The program is built upon the Framework discussed in Section 2. The evaluation is intended to appropriately weigh and quantify the effect of individual design characteristics in order to arrive at a cost value for each design configuration under study. The fundamental objective is to optimize this reactor design taking into account as many factors as reasonably possible. It is a useful tool for evaluating salts across many attributes relevant to reactor construction/operation. All costs are measured in units of 2021 USD at the time of reactor start-up. Capital costs have an interest expense while both the value of electricity produced and O & M costs are discount to the time of reactor start-up.

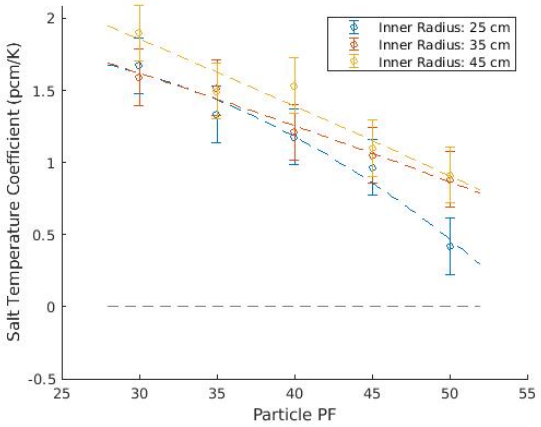

(a) Salt density TC, Blanks: 0%

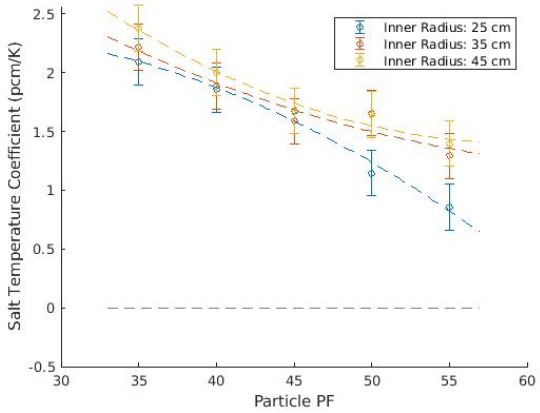

(b) Salt density TC, Blanks: 25%

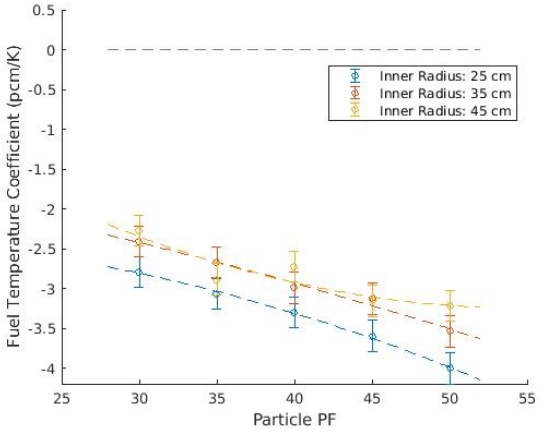

(c) Fuel Doppler TC, Blanks: 0%

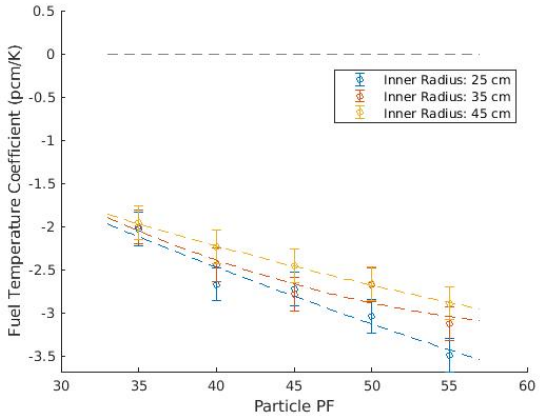

(d) Fuel Doppler TC, Blanks: 25%

Figure 5.29: Temperature coefficients of $NaF - BeF_2$ coolant salt; trend lines of the particle PF are fitted for each of the inner radii.

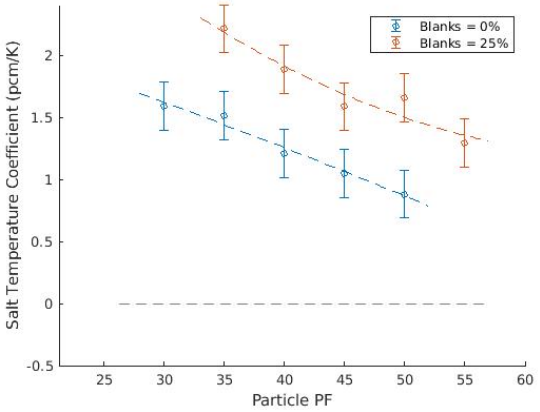

(a) Salt density TC, actual particle PF

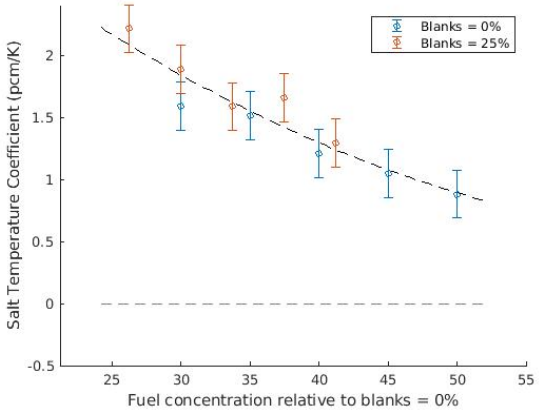

(b) Salt density TC, particle PF relative to 0% blanks

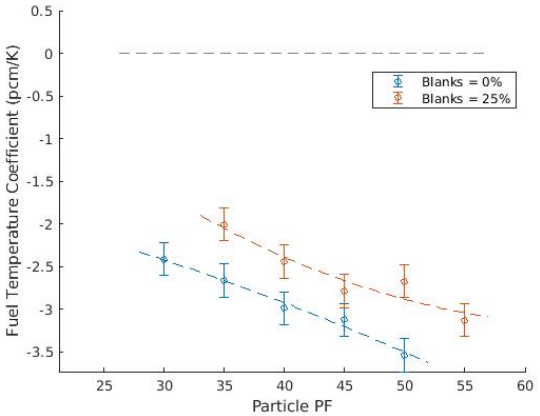

(c) Fuel Doppler TC, actual particle PF

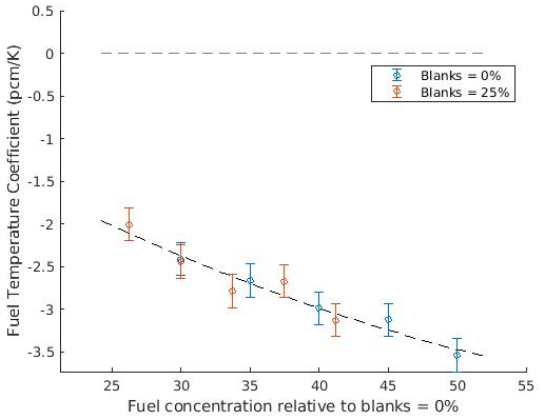

(d) Fuel Doppler TC, particle PF relative to 0% blanks

Figure 5.30: Blanks % comparison of the salt density and fuel Doppler temperature coefficients using $NaF - BeF_2$ coolant salt. Trends lines are fitted of the particle PF/fuel concentration.

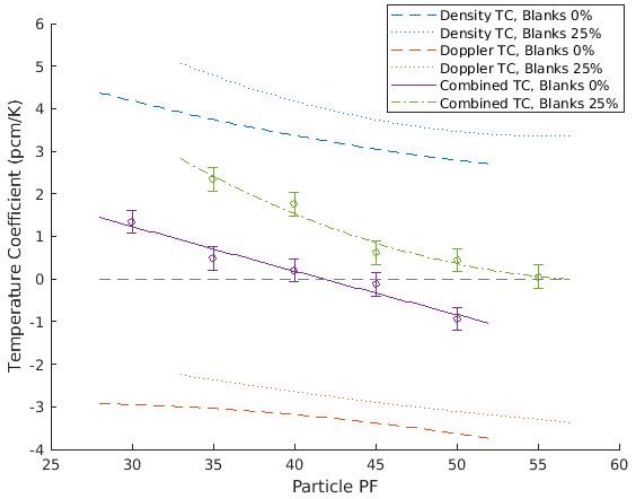

(a) $NaF - ZrF_2$

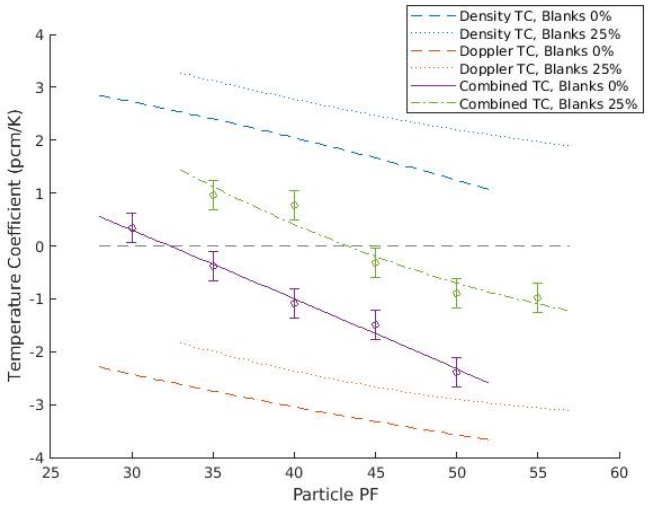

(b) $LiF - NaF - BeF_2$

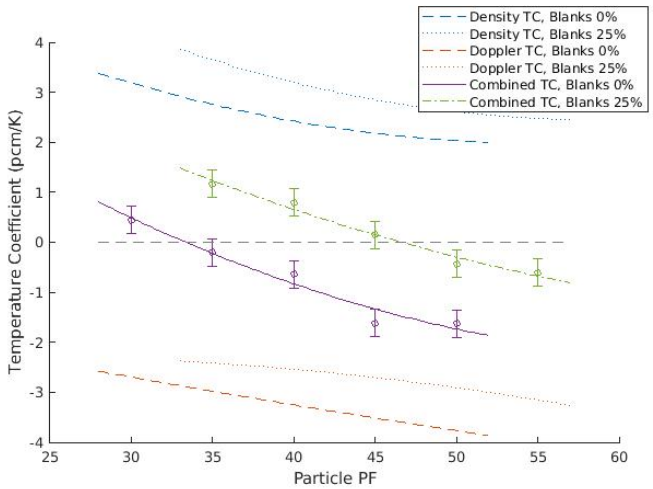

(c) $LiF - NaF - ZrF_4$

Figure 5.31: TCs of $NaF - ZrF_2$, $LiF - NaF - BeF_2$ and $LiF - NaF - ZrF_4$ at Blanks of 0% and 25% with an inner radius of 35cm.

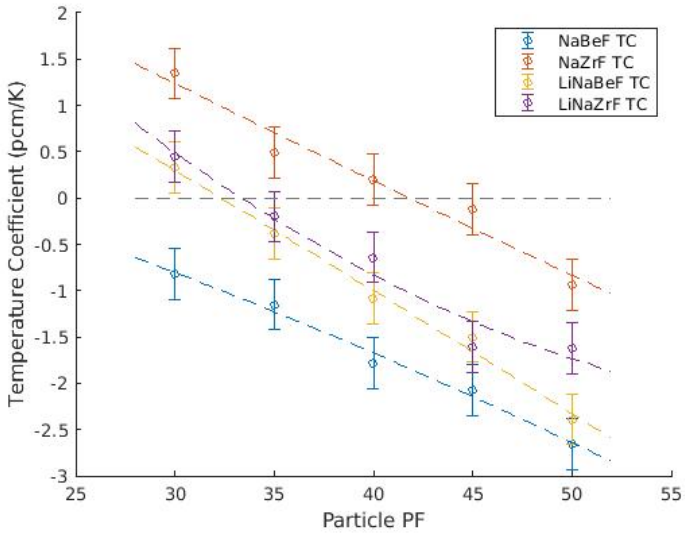

Figure 5.32: Combined (salt & fuel) Temperature coefficients of alternative (non-$LiF - BeF_2$) coolant salts at 0% blanks and an inner radius of 35cm.

The evaluation aims to assess the FHR design concept as a function of three design parameters listed below and discussed in Section 5.1.3.

- Particles per pebble. (Packing Fraction, PF)
- Share of graphite only pebbles (blanks)
- Inner reflector thickness

The results obtained when using this framework are highly dependent on the assumed input constants. Many of these input values are selected intuitively and require at least some degree of user discretion. Consequently uncertainty analysis is performed. The effect of each input variable on the total cost is initially unclear so sensitivity analysis is conducted as well. The uncertainty/sensitivity analysis assigns variables a probability distribution and the evaluation is run numerous times at pseudo-random points across each distribution (see Section 4.2).

5.6.2 Objective and Capabilities

The two fundamental objectives are to find low-cost reactor designs and to quantify the effect numerous input parameters have on the total cost. In doing so the relative viability of the salts examined and the optimal geometric configuration of the reactor using each of these salts can be determined. It is a useful tool for evaluating salts across many attributes relevant to reactor construction/operation. Not only are the sum total outputs useful but the program is built in such a way that it can be clearly seen which salt attributes as well as geometric parameters are responsible for which cost components.

5.6.3 FHR Evaluation Program Specifics

There are a total of 69 input parameters - or as per sensitivity analysis, explanatory variables. The variables defining the distribution of each input parameter used are listed in Tables 5.10 and 5.11. All of the input parameters have a configuration-dependent effect on cost. The costs are divided into six categories: direct cost, waste, safety, proliferation, modularity and feasibility. Costs of the same category are generally calculated in the same or similar fashion. Direct costs include industrial goods or fabricated products, often with a known market prices, consequently the cost calculations are comparatively straightforward. The other categories require more indirect costing, methods employed are discussed in Chapter 2. Each cost is made up of a set of cost components, see Table 5.12. The input parameter values used in the analysis with fixed inputs are the expected/mean values (μ) from the uncertainty analysis.

A separate uncertainty/sensitivity analysis is also performed on the cost categories. All of the input variables are set at constant values while the cost categories are given a double trapezoid probability distribution between 0.7 and 1.3. The *Direct cost* category is not included because the uncertainty stems from the uncertainty in the individual cost components. Whereas with the other cost categories there exists uncertainty in converting a *issue* (ex. safety) into a cost value.

In FHR designs with a central graphite reflector, the central reflector recieves a high fast neutron dose and therefore is expected to be the limiting attribute of operational lifetime, however it is assumed to be replaceable as a modular unit as in the Mark-1 design [48]. The maximum fast neutron flux of the central graphite reflector is configuration dependent, it is particularly sensitive to the reflector radius. The maximum fast flux decreases (and lifetime increases) with radius. Yet the size and therefore cost of the reflector increases with radius. As well the replacement is not expected to disrupt operation for a long period of time and may even be replaceable during operation. Considering that trade-off appears to balance well and the cost/time uncertainties are large, the evaluation program does not take it's replacement into consideration.

Without a central reflector or with a replaceable one, the operational lifetime is assumed to be limited by the irradiation damage sustained by the outer graphite reflector. The maximum fast neutron fluence of solid blocks of graphite can be inferred from MSBR research (Section 1.6.4), where the core was expected to sustain four years of operation with an average core power density of 22.2 MW/m^3, for a total of 88.8 $MW \cdot yrs/m^3$. However the reflector would likely have a geometry that can sustain greater dimensional change than a solid block. In the reference FHR design the maximum outer reflector fast flux is ≈0.35x the maximum fast flux in the core and because of the central graphite reflector the peak:avg fast flux is lower. A factor of 1.33 is selected to account for the peak:avg differences and the additional fluence tolerance. With a 21.65 MW/m^3 power density, the lifetime is expected to be:

$$15\ yrs = \frac{1.3 \cdot (88.8\ MW \cdot yrs/m^3)}{(0.35 \cdot 21.65\ MW/m^3)} \tag{5.23}$$

The average core inlet and outlet temperatures used in the evaluation program analysis are 615 and 685 ^oC.

Table 5.10: The defining distribution parameters of all input parameters using the Pearson probability distribution. Costs are in units of 2021 USD.

Input variable	Pearson probability distribution			
	mean (μ)	S.D. (σ)	skew	kurtosis
Construction time (yr)	6	2	1.7	8
Interest rate	0.05	0.015	1	6
Lifetime (yr)	15	1.5	1	6
MWh cost (\$/MWh)	100	7	1	6
Pumping multiplier	1.5	0.25	1	6
Material cost (\$/m^2)	2500	700	1.2	7.5
^{7}Li price (\$/kg)	1500	450	1.5	8
Beryllium price (\$/kg)	500	100	1.5	8
Zirconium price (\$/kg)	100	20	1.5	8
Separative work[a]	400	80	1.5	8
$U_3O_8^a$	600	120	1.5	8
Fuel pebble manufacture (\$/peb.)	4	2	1.5	8
Graphite pebble manufacture (\$/peb.)	3	1.5	1.5	8
Ref. refuelling OM (\$/MWh)	1	0.5	1.5	8
Ref. refuelling cap ($\$ \cdot 10^6$)	5	2.5	1.5	8
Ref. O & M safety (\$/MWh)	3.03	1	1	6
Ref. O & M prolif (\$/MWh)	1.66	0.2	1	6
Be handling	15	4	1	7
Detritation O & M (\$/MWh)	1	0.4	1	7
Ref. tech. develop ($\$ \cdot 10^6$)	1000	100	1.2	7.5
Share tech.	0.025	0.007	1.5	8
Rel. corrosion 1, 2, 3, 4 and 5	1	0.15	1	7
Non-configuration ($\$ \cdot 10^6$)	11.5	1.5	1.5	7
Ref. cost MWth (\$/MWth)	2500	650	1.2	7.5
Insurance premium ($\$10^6$)	2	0.5	1.2	7.5
CDF	$1.0 \cdot 10^{-6}$	$0.25 \cdot 10^{-6}$	1.2	7
Detritation C1	2	0.7	1.2	7
Detritation C2	18	3	1	7
Ref. Pu rel. cost	5	0.5	1	7
T_{melt} rel. cost	6	1.5	1	7
P_{vap} rel. cost	3	1	1	7
Corrosion rel. cost	5	0.5	1	7
Be rel. cost	6	1.5	1	7
^{6}Li effect rel. cost	15	3	1	7
PNL rel. cost	1	0.2	1	7
Xe PF safety rel. cost	4	1	1	7
Temp. feedback rel. cost	10	2	1	7
Fuel temp. rel. cost	10	2	1	7

[a] \$ to enrich 1 kg to 5% with 0.2% in tails.

Table 5.11: The defining distribution parameters of all input parameters using the double trapezoid probability distribution.

	Double Trapezoid probability distribution		
Input variable	mean	min	max
Reference PD (Pa)	5400	3500	7500
Capacity factor	0.87	0.72	0.98
Cost overrun safety	0.6	0.45	0.75
Share cap safety	0.2	0.05	0.4
Cost overrun prolif	0.10	0.05	0.15
Share cap prolif	0.05	0.02	0.10
Ref. share modularity	0.15	0.05	0.30
T_{melt} split	0.8	0.6	0.95
P_{vap} split	0.8	0.6	0.95
Corrosion split	0.8	0.6	0.95
T_{melt} max (^oC)	585	550	615
Detritation C3	0.75	0.55	0.95
Non-config. exp.	0.7	0.5	0.9
Rel. safety cost exp.	0.5	0.3	0.7
Rel. prolif cost exp.	0.5	0.3	0.7
Regular cost modularity exp.	0.5	0.3	0.7
Overrun exp.	0.5	0.4	0.6
Feasib each exp.	0.35	0.25	0.45
Refuel cap exp.	0.5	0.3	0.7
Flux to Xe PF exp.	0.6	0.4	0.8

Table 5.12: Cost components associated with each cost category.

Cost	Input parameter
Direct cost	Pumping power in heat exchanger
	Heat exchanger
	Pumping power in core
	^{7}Li, in core
	Beryllium, in core
	Zirconium, in core
	Initial fuel pebbles
	Initial graphite pebbles
	Salt in the heat exchanger
	Refuelling capital cost
	Refuelling O & M[a]
	Detritation capital cost
	Detritation O & M
	Corrosion
	Structural materials
	Interest on capital cost
Waste	Deep geological repository
	Be toxicity
Safety	Vapour pressure
	Melting temperature
	^{6}Li/Tritium
	Be toxicity
	PN generation time
	Xe safety
	Fuel temperature
	Doppler coefficient
	Density coefficient
	Insurance
Proliferation	^{6}Li/Tritium
	Plutonium production
Modularity	Modularity
Feasibility	Vapour pressure
	Melting temperature
	Corrosion

[a]Includes the U_3O_8, separative work and fuel manufacturing costs.

Table 5.13: Reference values used in cost setting. A description of cost setting provided in Section 2.8.

Parameter	Value	Associated input variable
T_{melt} ref.	550 ^oC	T_{melt} cost
P_{vap} ref.	10 mmHg	P_{vap} cost
Corrosion ref.	1	Corrosion cost
Xe PF	0.05	Xe PF safety cost
[Be]	0.015 $10^{24}/cm^3$	Be cost
[Li]	0.013 $10^{24}/cm^3$	^{6}Li effect cost
^7Li enrichment	0.99995	^{6}Li effect cost
PNL	0.001 s	PNL cost
TC	-1.5 pcm/K	Temp. feedback cost
Max. fuel ΔT	500 ^oC	Fuel temp. cost

The reference power used in the calculation of the cost-share due to (a lack of) modularity is 500 MWe and the reference CDF used in the insurance cost calculation is $1.0 \cdot 10^{-5}$.

Iteration

Many calculations in the evaluation program use data points generated in other sections. The program is ordered in such a way to minimize situations where input data has not been explicitly calculated in an earlier section, however it cannot be completely eliminated. One example is the ^{7}Li enrichment optimization and pump/heat exchanger size optimization: The optimal ^{7}Li enrichment is a function of the salt volume - and by extension the heat exchanger size, and the heat exchanger size is a function of the salt cost per unit volume - and by extension the ^{7}Li enrichment level. Thus it is necessary for this section of the program to iterate until convergence; no more than 3 iterations are required. The costs involved and the iteration methodology are discussed in Section 3.4.1.

Cost of corrosion in the evaluation program

The cost of corrosion inflicted by the coolant salts (and corrosion mitigation) in the evaluation program is a function of the relative expected corrosion depth. The corrosion depth is relative to uranium-free FLiBe salt. The method used to quantify the relative magnitude of corrosion is a function of the Gibbs energy differences and experimental results. This method is explained in Section 2.1.1.

5.7 Results

5.7.1 Results: with Input Variables at Mean/Expected Values

Figures 5.33 to 5.40 show results derived from setting all the input parameters (listed in Section 5.6.3) to fixed values. An advantage to analysing fixed input data is that the trends are more clear than with uncertainty analysis. As well the lower computing requirements allow for more configurations to be analysed.

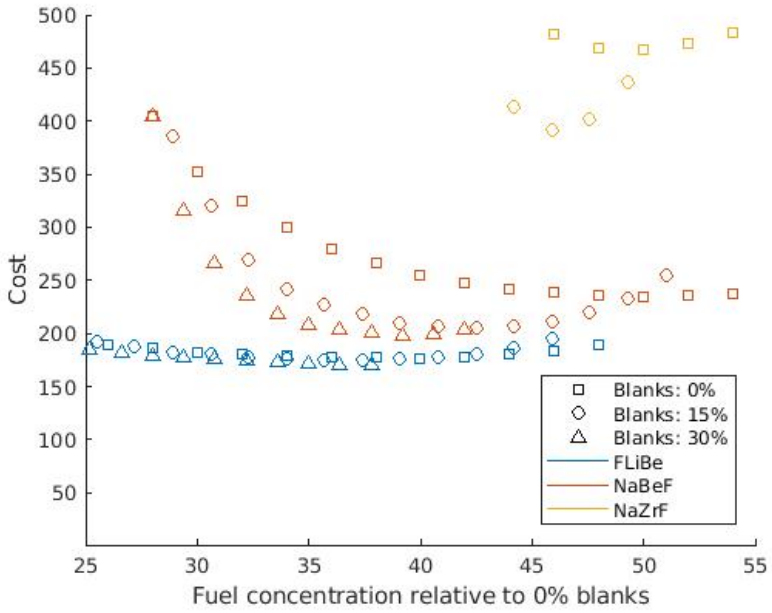

Figure 5.33: Total cost over a range of PFs and several blank shares at an inner radius of 35cm. Results shown comparing FLiBe, NaBeF and NaZrF salts where the x-axis is the PF that corresponds to the PF with 0% blanks of an equal core fuel concentration.

The selected input parameter values are the expected mean values; the expected mean from the expected probability distributions used in uncertainty/sensitivity analysis. Figures 5.33 and 5.34 show the calculated cost with the inner radius fixed at 35 cm. The independent variable is the core fuel concentration, scaled to PFs with 0% blank pebbles. Figure 5.35 shows the same data but with the actual PF, not scaled to the core fuel concentration. Figures 5.36 and 5.37 are similar style plots but with share of blanks fixed at 15%. Presenting results of all five salts is generally too cluttered, so the data presented are split into two, each including FLiBe salt so that comparisons to the reference salt composition can easily be made.

At the configuration that results in the lowest total cost, the cost is broken down into the six categories, Figure 5.38. The contributing input parameters to the 'Direct cost' and 'Safety' cost for FLiBe, NaBeF and NaZrF salts are shown respectively Figures 5.39 and A.13. The breakdown for the other salts and of the other cost categories into the cost of their inputs is provided in Appendix A.1.

5.7.2 Results: Uncertainty Analysis

Uncertainty analysis results are only provided for FLiBe, NaBeF, LiNaBeF and LiNaZrF salts going forward. NaZrF is left out because it is feasible over less of the analysed design space and is significantly higher cost (at most configurations) compared with the other salts. The uncertainty/sensitivity analysis was originally conducted with a sample size of 400 for all feasible configurations across the following design parameter values: (a total of 81 potential configurations per salt)

- Blanks: 0, 15, 30%

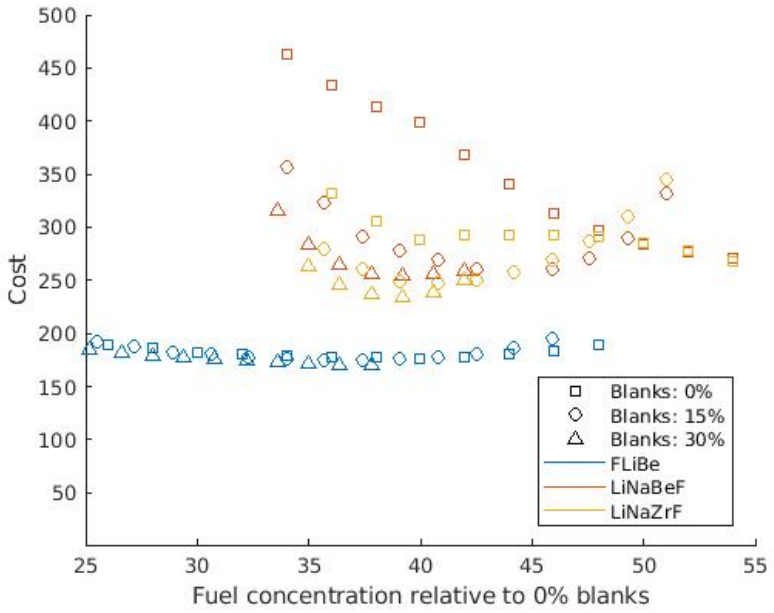

Figure 5.34: Total cost over a range of PFs and several blank shares at an inner radius of 35cm. Results shown comparing FLiBe, LiNaBeF and LiNaZrF salts where the x-axis is the PF that corresponds to the PF with 0% blanks of an equal core fuel concentration.

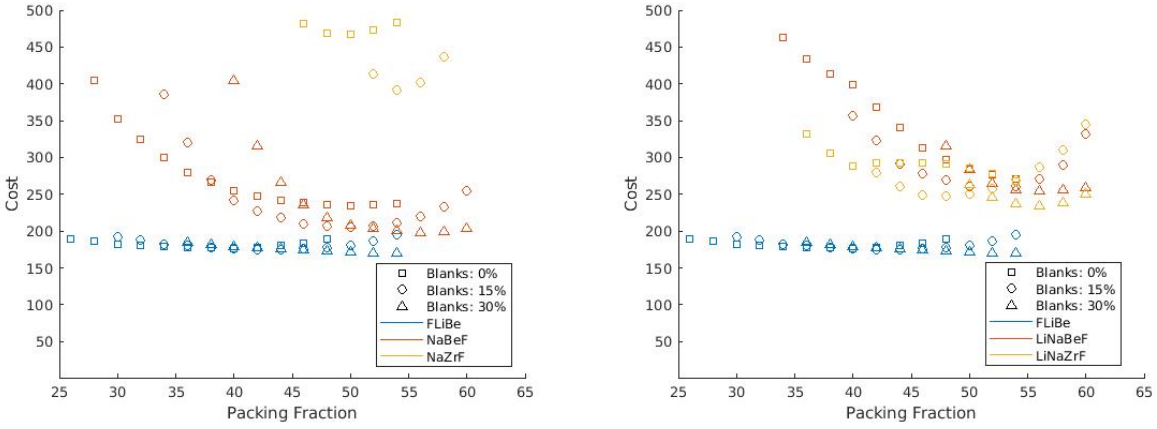

(a) Salts: Flibe, NaBeF and NaZrF

(b) Salts: Flibe, LiNaBeF and LiNaZrF

Figure 5.35: Total cost over a range of PFs and several blank shares at an inner radius of 35 cm.

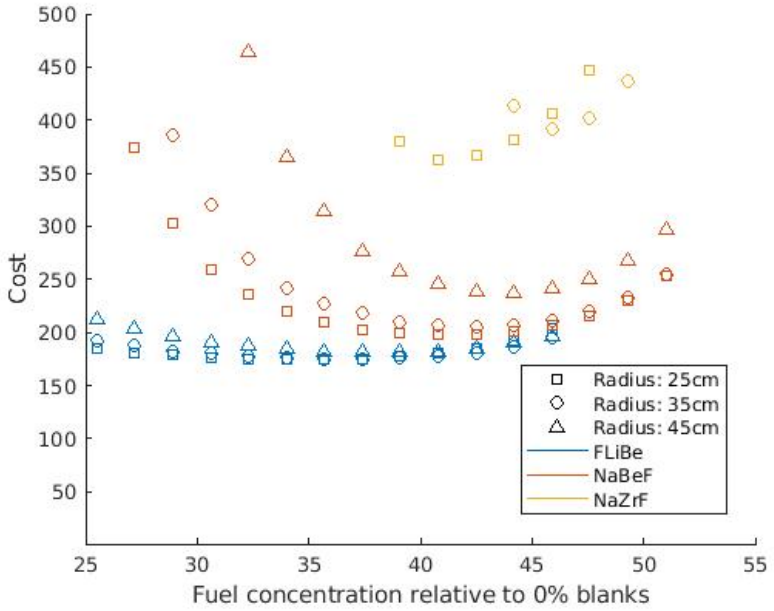

Figure 5.36: Total cost of each PF and inner radius at a blank share of 15%. Results shown comparing FLiBe, NaBeF and NaZrF salts where the x-axis is the PF that corresponds to the PF with 0% blanks of an equal core fuel concentration.

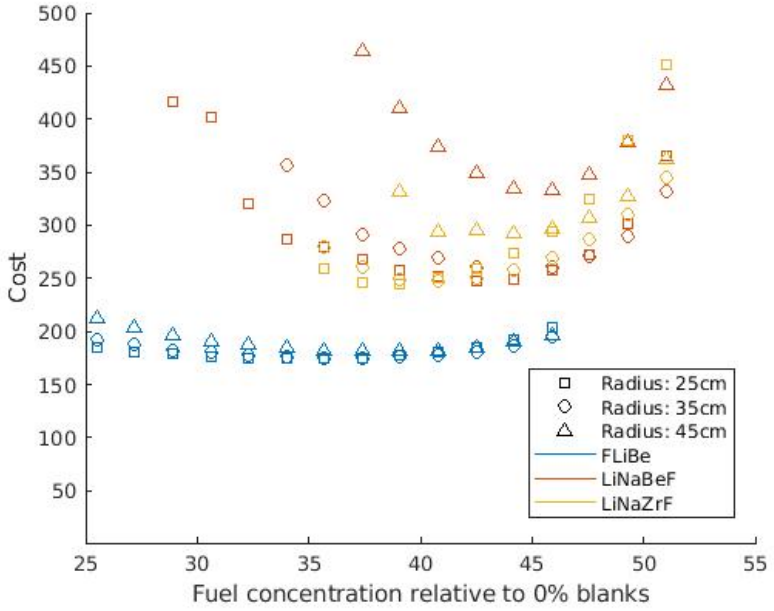

Figure 5.37: Total cost of each PF and inner radius at a blank share of 15%. Results shown comparing FLiBe, LiNaBeF and LiNaZrF salts where the x-axis is the PF that corresponds to the PF with 0% blanks of an equal core fuel concentration.

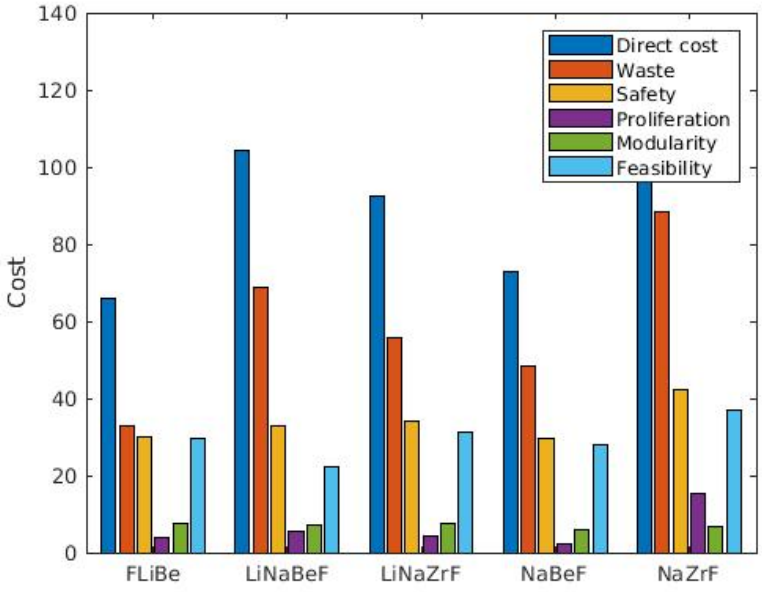

Figure 5.38: The cost of each of the six categories is shown for each salt in a FHR. The cost values are given at the graphite pebble, PF and inner radius combination that results in the lowest feasible total cost.

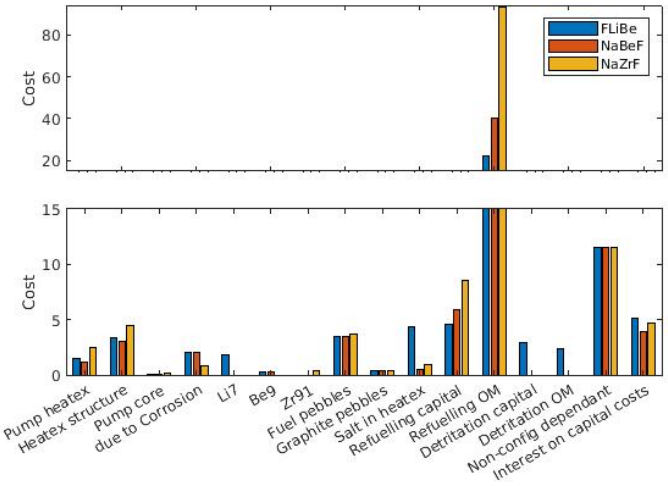

Figure 5.39: Breakdown of the 'Direct cost' category of cost into components. The cost values are given at graphite pebble, PF and inner radius combination that results in the lowest feasible total cost.

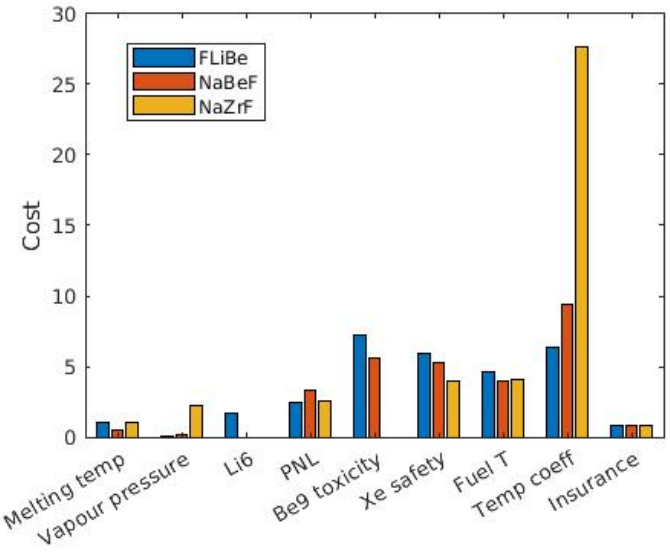

Figure 5.40: Breakdown of the 'Safety' category of cost into components. The cost values are given at graphite pebble, PF and inner radius combination that results in the lowest feasible total cost. Only one cost resulting from TCs is calculated, it is a function of the summation of the density and doppler TCs.

- PFs: 26, 30, 34, 38, 42, 46, 50, 54, 58
- Inner radii: 25, 35, 45 cm

This sample size is sufficient to calculate the mean and standard deviation (SD). With cost SDs of approximately 30, the deviations from the means are only $\approx$1.75 ($35/\sqrt{400}$) - or only slightly higher due to non-normality. Therefore these results are used to produce Figures 5.43 and 5.45. The cost distribution cannot be assumed normal in the right side tail so a large number of samples is required to determine important percentiles of the distribution, such as the 95 and 97.5, with sufficient precision. All other uncertainty (and sensitivity) analysis results shown use 20 repetitions of 100 sample bins for a total of 2000 samples. Using 2000 samples and making no assumptions about the distribution, it was found that the SD of the 95^{th} percentile for the lowest cost configuration is between 3.3 and 4.5 (1.1% and 2.0%[10]) depending on the salt composition and the SD of the 97.5 percentile is between 4.7 and 7.9 (1.8% and 2.8%). This level of precision was deemed sufficient considering the substantial uncertainty of many input parameters.

It become clear that much of the range includes many sub-optimal configurations so 2000 samples were applied only to a smaller range that includes the optimal configurations: (a total of 6 potential configurations per salt)

- Blanks: 30%
- PFs: 50, 54, 58
- Inner radii: 25, 35 cm

[10] $100 * \frac{deviation}{95^{th} \text{ percentile cost}}$

Table 5.14: Lowest cost configuration of the data analysed in the uncertainty analysis.

Salt	Blanks %	PF	Inner radius
FLiBe	30	54	35cm
NaBeF	30	54	25cm
LiNaBeF	30	58	25cm
LiNaZrF	30	54	25cm

The distributed variables used in the uncertainty analysis have a mean equal to that used for the fixed values. Figure 5.41 shows the uncertainty analysis results for the lowest cost configuration of four of the salts. Vertical lines representing the average cost of the other salts, at their respective lowest cost configuration, are shown for comparison. The lowest-cost configuration of each salt includes 1-3 clearly outlying data point(s). These 'outliers' are only included in Figure 5.41, all other results do not display these points. All points, including the outliers, are however included in calculating averages, variances, sensitivities, etc. If results are only presented for 1 configuration or 1 configuration per salt, the results are for the configuration that results in the lowest expected total cost. The lowest cost configuration of each salt is listed in Table 5.14.

A double histogram plot of the lowest cost *FLiBe* configuration compared to the lowest cost alternative salt configuration is provided in Figure 5.42 (a), Figure 5.42 (b) shows the lowest cost configuration of the other two salts. Uncertainty analysis is compared to set values over a range of PFs in Figure 5.43. FLiBe is compared against NaBeF because the two salts compositions have the lowest expected cost among all salts. The other sub-figure compares the other two salt compositions, LiNaBeF and LiNaZrF.

Results presented so far have shown the overall uncertainty. A likely more important presentation of data than the distributions of the total costs are the distributions of the relative cost. If the input parameters are treated as independent, as in Figure 5.42, there is often significant overlap in the cost distributions; overlap proportional to the probability that the salt with the higher expected cost has a lower cost than the other salt. However, in situations where the combination of input parameters results in high costs (high CDF in the histogram of total cost) for one configuration, it tends to result in a high cost for all configurations. Figure 5.44 shows the difference in cost between salts at their lowest cost configurations. The same salts are shown in Figures 5.42 and 5.44, take note of the dramatically lower probability of the salt with the higher expected cost ((a) NaBeF, (b) LiNaBeF) having a lower cost. This perspective suggests that, for the input parameters values used, FLiBe results in the lowest total cost under almost all scenarios.

Correlation:

Input parameter correlations are assigned relative to a maximum. The maximum used in the analysis is a correlation coefficient of 0.30. The relative correlations are either, 0.5, 0.7 or 1.0, each of which occurs with approximately equal frequency. A higher correlation coefficient is preferred but the positive definite requirement of the correlation matrix, see Section 4.4, limited the maximum correlation to little more than 0.30. The correlations were found to have little impact; approximately 220 correlations are assigned out of

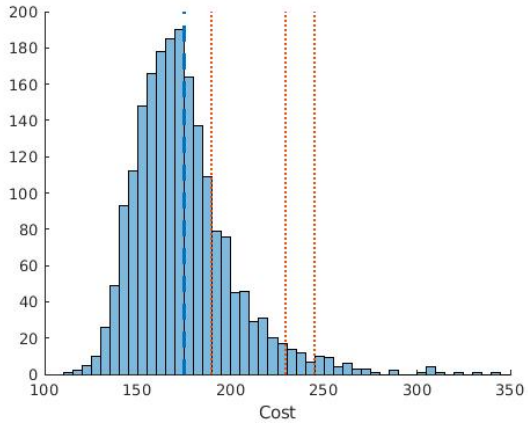

(a) Flibe, blanks 30%, PF 54, radius 35cm.

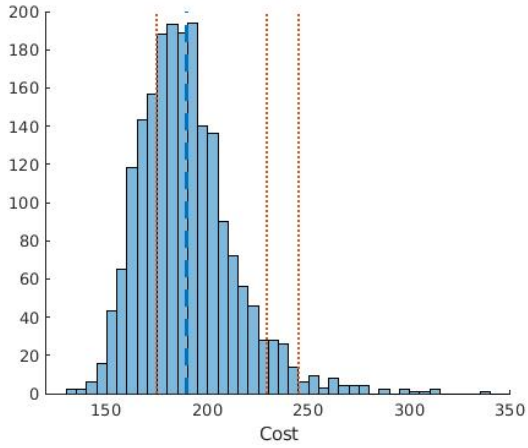

(b) NaBeF, blanks 30%, PF 54, radius 25cm.

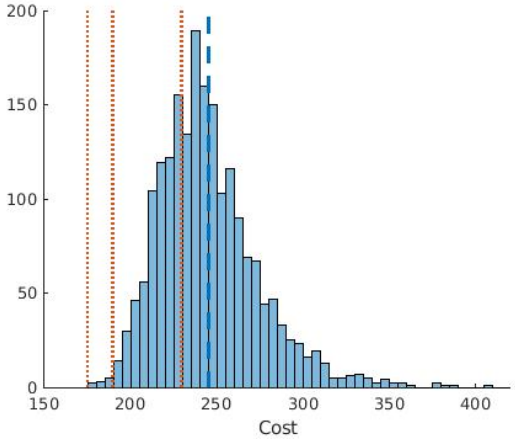

(c) LiNaBeF, blanks 30%, PF 58, radius 25cm.

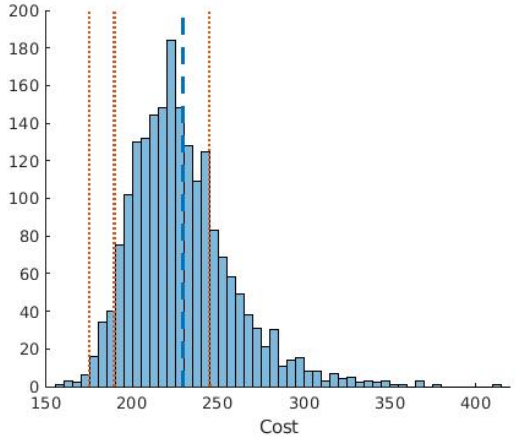

(d) LiNaZrF, blanks 30%, PF 54, radius 25cm.

Figure 5.41: Results of the uncertainty analysis of the geometric configuration with the lowest cost. The corresponding blanks, PF and inner radius is noted.

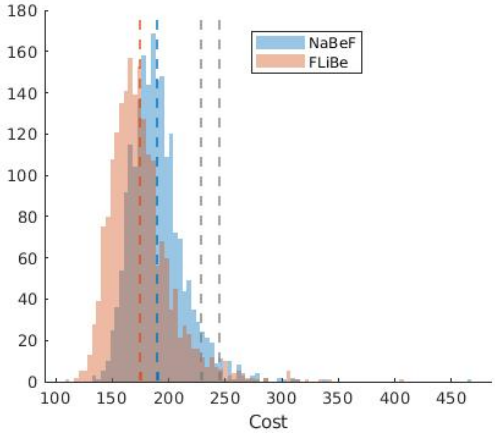

(a) NaBeF and FLiBe

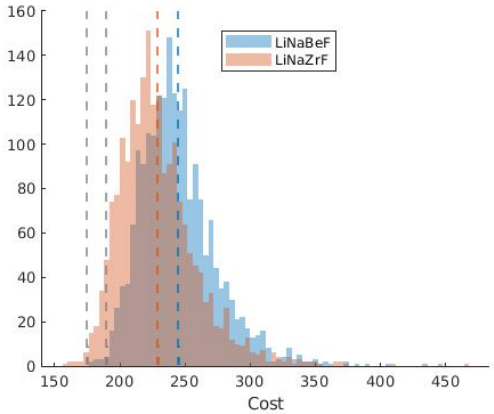

(b) LiNaBeF and LiNaZrF

Figure 5.42: Overlapping histogram plots of the total cost. Data presented is at the lowest cost configuration.

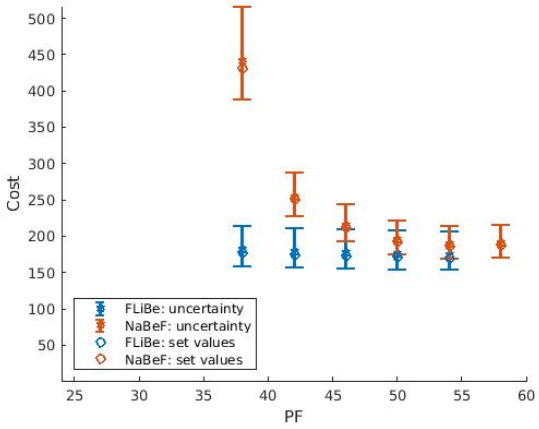

(a) FLiBe and NaBeF

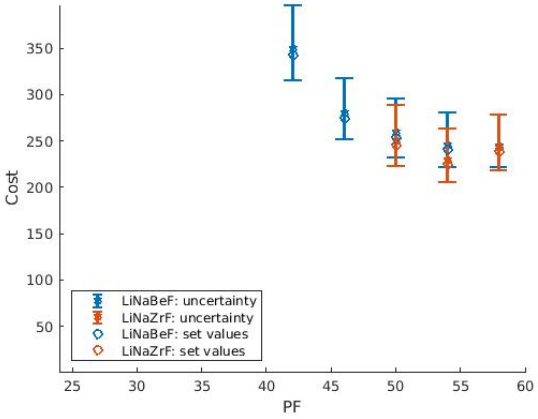

(b) LiNaBeF and LiNaZrF

Figure 5.43: Comparison of uncertainty analysis and fixed values. Results shown with 15% blanks and at an inner diameter of 35 cm across all evaluated PFs. Uncertainty analysis shown with an error range of one standard deviation.

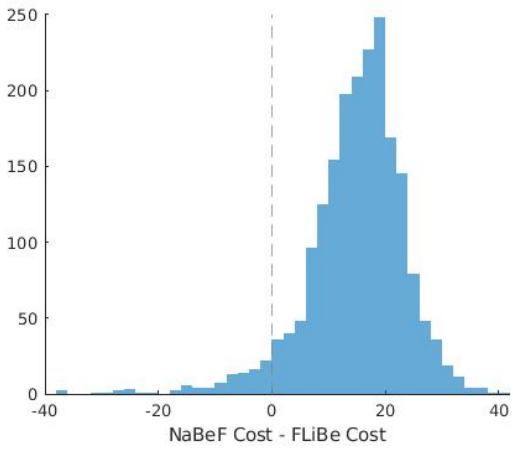

(a) FLiBe and NaBeF

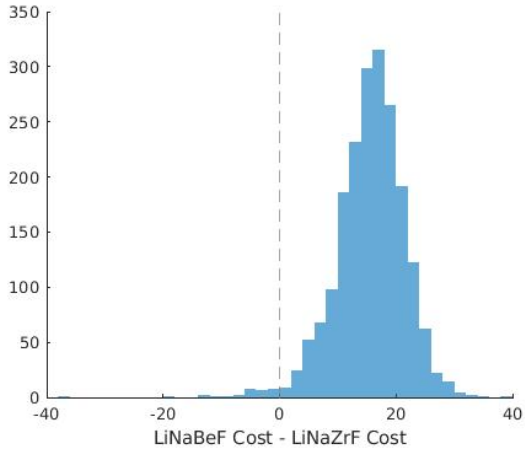

(b) LiNaBeF and LiNaZrF

Figure 5.44: Histogram of the difference in cost between two salts at each combination of input parameter values.

Table 5.15: Sensitivity of total cost to each cost category at the lowest cost configurations. The values shown are scaled relative to the standard deviation, thus the sensitivities are listed as $\Delta Cost$, where $\Delta Cost = b_i\sigma_i$.

Cost category	FLiBe	NaBeF	LiNaBeF	LiNaZrF
Waste	5.07	7.49	10.6	8.59
Safety	5.03	5.11	5.69	5.66
Proliferation	0.634	0.259	0.684	0.569
Feasibility	2.89	2.39	2.13	4.92

a possible 2380 $((69 \cdot 69)/2)$ with an average coefficient of only 0.22. $(0.30(0.5+0.7+1.0)/3)$ However, the correlations do tend to be assigned to input parameters that have high impact on total cost. In general it is difficult to draw many useful conclusions. It was found that the level of correlation does not change the optimal configuration, it mainly has the effect of mildly increasing the variance of the total cost. Comprehensive documentation is not deemed worthwhile.

Uncertainty at the cost category level:

Uncertainty analysis is also performed at the cost category level; with categories direct cost, waste, safety, proliferation, modularity and feasibility. This must be a separate analysis because the addition of cost category uncertainty on top of individual input parameter uncertainty would account for the same underlying uncertainty twice. The cost uncertainty provides less insights than the comprehensive input parameter uncertainty so only limited results are shown: Figures 5.46, 5.47 and 5.48 as well as Table 5.15. A comparatively small number of samples are used (400) because no outliers are generated and a normal distribution can be assumed.[11] The uncertainty in the direct costs category is fundamentally uncertainty at the level of the input parameters, not the cost category level. Modularity is closely tied to the capital cost portion of direct costs. Direct and modularity costs are fixed while the remaining cost categories are treated as uncertain - therefore the cost category analysis does not account for all uncertainty. In each run the total value of each distributed cost category is multiplied by a factor selected from a probability distribution; a double-trapezoid probability distribution with a mean of 1.0, minimum of 0.7 and maximum of 1.3.

Uncertainty analysis compared to fixed values at the lowest cost configuration for each of the four salts is shown in Figure 5.45. In three of four lowest-cost configurations, the uncertainty analysis average is lower than the corresponding fixed-values; it is unclear why. When assigning all input parameters as probabilistic variables and creating various non-linearities in the process, the average instead becomes somewhat higher ($\approx$5%) than the fixed values calculation for all lowest-cost configurations. The variance when using probabilistic input parameters is considerably higher than with probabilistic factors assigned to several cost categories.

[11] The normal distribution is assumed because the analysis cost calculation is additive and the input distributions have fixed boundaries.

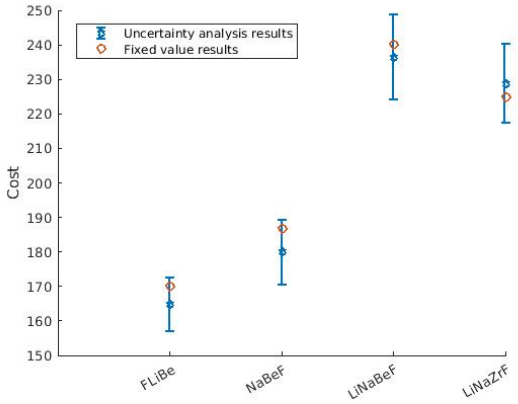

(a) Cost category uncertainty

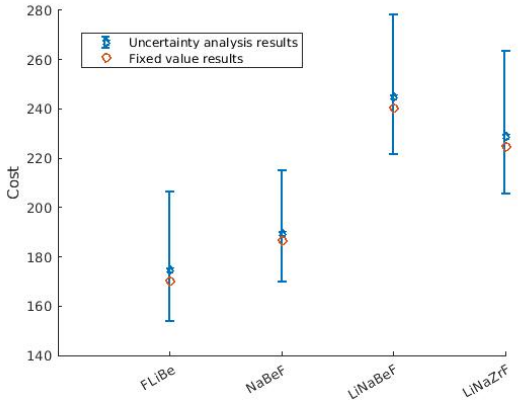

(b) Input parameter uncertainty

Figure 5.45: Comparison of uncertainty analysis and fixed values at the configurations with the lowest cost. Uncertainty analysis shown with an error range of one standard deviation.

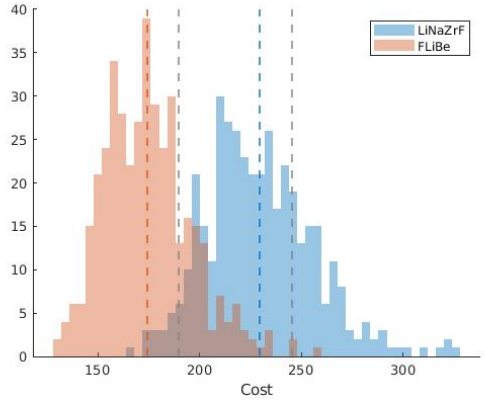

(a) LiNaZrF and FLiBe

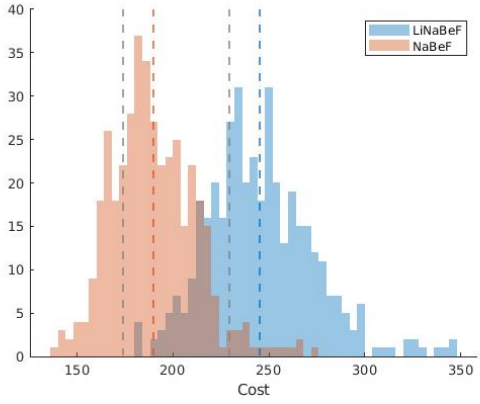

(b) LiNaBeF and NaBeF

Figure 5.46: Overlapping histogram plots of the total cost with uncertainty at the cost category level. Data presented is at the lowest cost configuration.

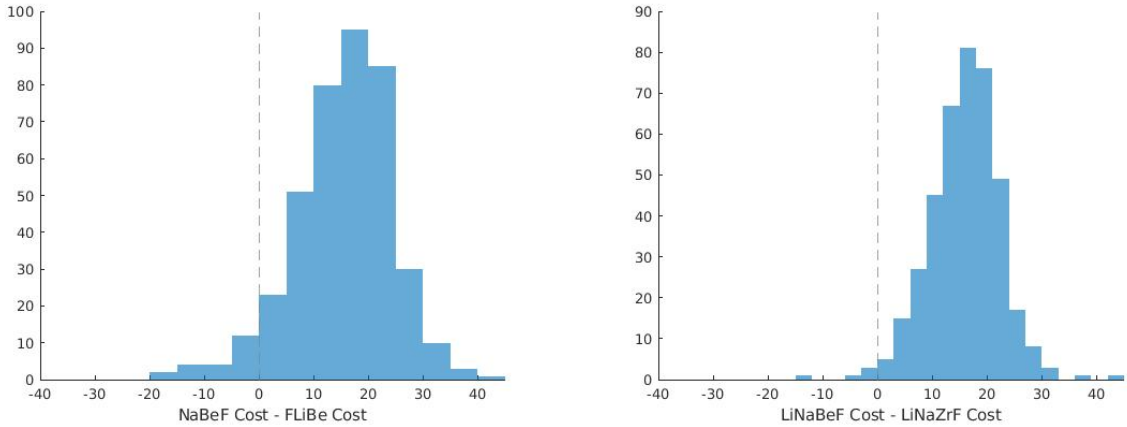

(a) FLiBe and NaBeF

(b) LiNaBeF and LiNaZrF

Figure 5.47: Histogram of the difference in cost between two salts at each combination of cost category values. Data presented is at the lowest cost configuration.

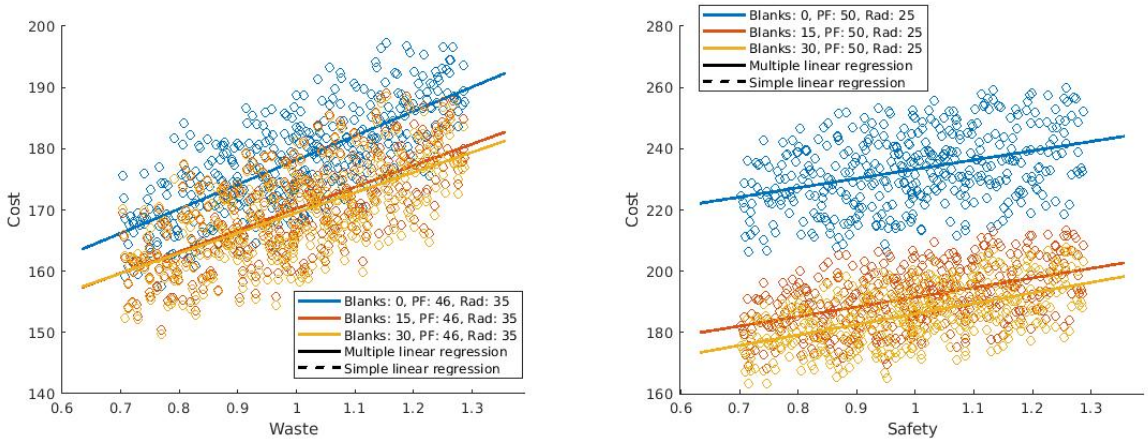

(a) FLiBe: Waste

(b) NaBeF: Safety

Figure 5.48: Total cost at three configurations as a function of select cost categories. The data points shown indicate the cost category factor (x-axis) and total cost (y-axis) of each run. The resulting coefficient from performing multiple linear regression is shown as the slope of the solid line. Simple linear regression is provided for comparison.

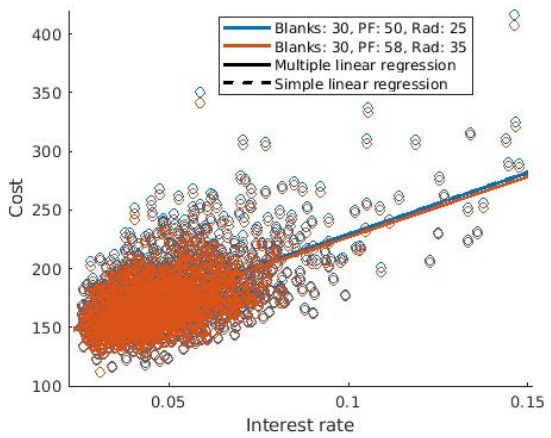

(a) FLiBe

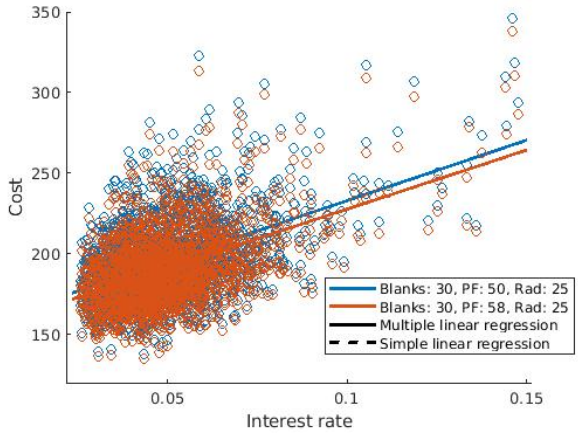

(b) NaBeF

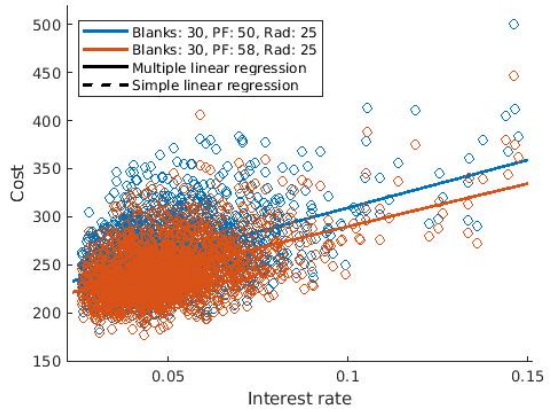

(c) LiNaBeF

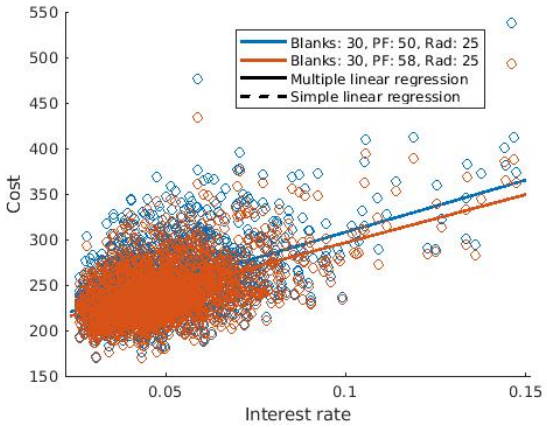

(d) LiNaZrF

Figure 5.49: Total cost of four salts at two low cost configurations as a function of interest rate. The data points shown indicate the input parameter value (x-axis) and total cost (y-axis) of each run. The resulting coefficient from performing multiple linear regression is shown as the slope of the solid line. Simple linear regression is provided for comparison.

5.7.3 Results: Sensitivity Analysis

Using the same 2000 data points as in uncertainty, sensitivity analysis was performed to gain insight into the effect each of the input variables have on the total cost. The analysis is based upon linear multiple regression.

A list of regression coefficients of select input parameters at the lowest cost configuration for the four salts are provided in Table 5.16. Many input parameters have little impact on total cost; parameters with a large impact, that differ significantly between salt compositions and/or have unexpected/interesting results are listed. A complete list is provided in Appendix A.1. Economic variables have high-magnitude coefficients, though they tend to be similar among the different configurations, the interest rate is a notable exception. Most other non-economic, high-magnitude coefficients are similar among configurations as well. Variables that are similar among configurations are relatively inconsequential to salt composition and core geometry

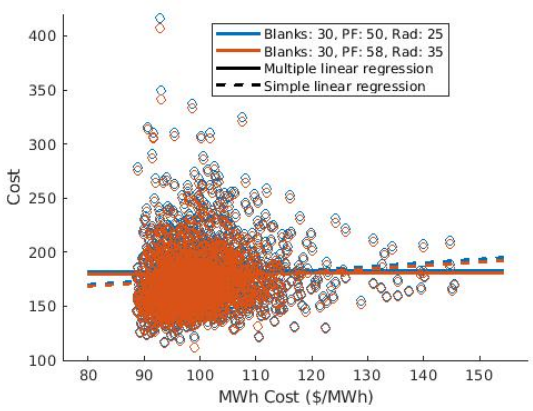

(a) MWh cost (for pump)

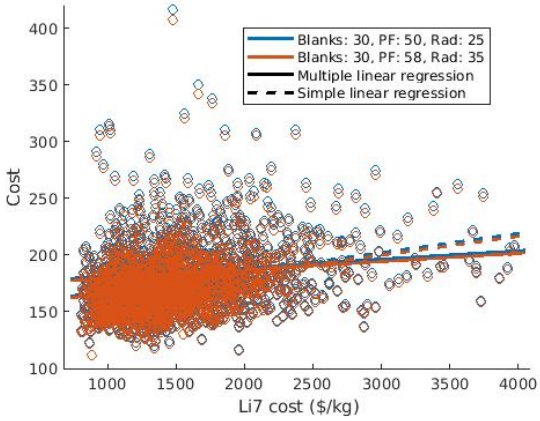

(b) ^{7}Li price

Figure 5.50: Total cost of FLiBe salt at two low cost configurations as a function of two input parameters. The data points shown indicate the input parameter value (x-axis) and total cost (y-axis) of each run. The resulting coefficient from performing multiple linear regression is shown as the slope of the solid line. Simple linear regression is provided for comparison.

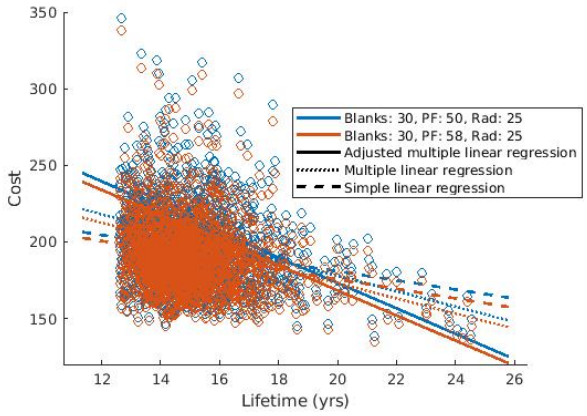

(a) Lifetime[a]

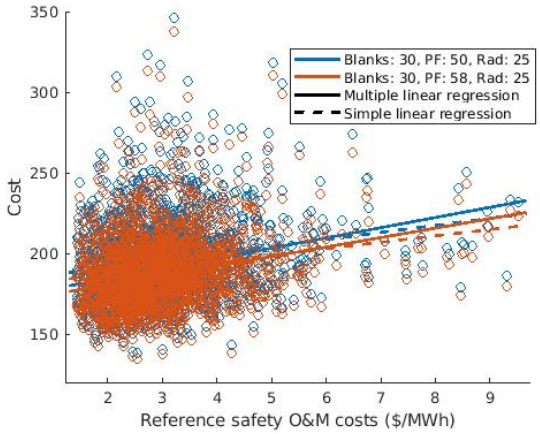

(b) Reference safety O & M cost

[a]Unreasonable correlations with the *Lifetime* variable were used. The *Lifetime* sensitivities have been adjusted to obtain the expected sensitivity in the absence of any correlations.

Figure 5.51: Total cost of NaBeF salt at two low cost configurations as a function of two input parameters. The data points shown indicate the input parameter value (x-axis) and total cost (y-axis) of each run. The resulting coefficient from performing multiple linear regression is shown as the slope of the solid line. Simple linear regression is provided for comparison.

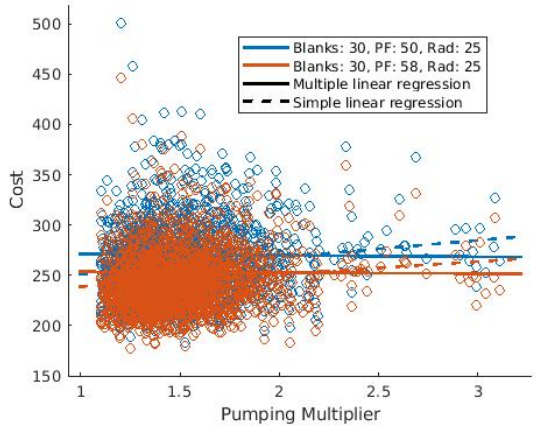

(a) Pumping multiplier

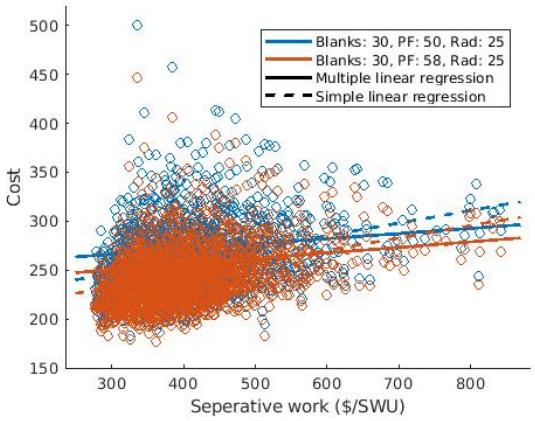

(b) Separative work

Figure 5.52: Total cost of LiNaBeF salt at two low cost configurations as a function of two input parameters. The data points shown indicate the input parameter value (x-axis) and total cost (y-axis) of each run. The resulting coefficient from performing multiple linear regression is shown as the slope of the solid line. Simple linear regression is provided for comparison.

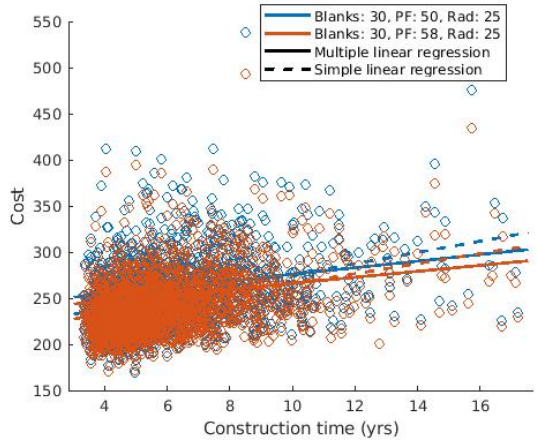

(a) Construction time

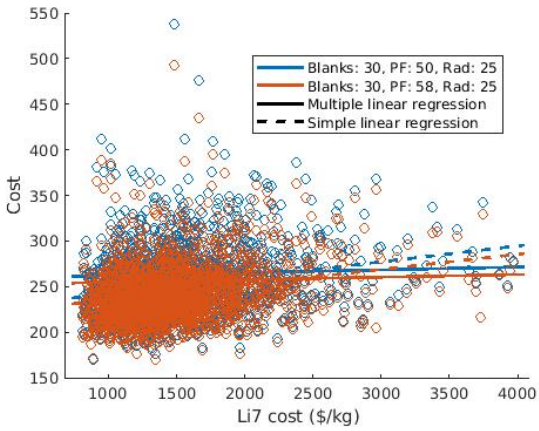

(b) ^{7}Li price

Figure 5.53: Total cost of LiNaZrF salt at two low cost configurations as a function of two input parameters. The data points shown indicate the input parameter value (x-axis) and total cost (y-axis) of each run. The resulting coefficient from performing multiple linear regression is shown as the slope of the solid line. Simple linear regression is provided for comparison.

selection. Interest rate and fuel pebble costs (separative work cost, U_3O_8 price and fuel pebble manufacturing cost) appear to have the greatest impact on configuration selection, particularly the salt composition. This suggests that much of the uncertainty between configurations, especially if one is FLiBe salt, has to do with the difference in timing of the costs; FLiBe configurations have costs that, on average, occur earlier in time.

The input parameters involved in setting the cost for a reference reactor, such as 'Share tech. develop,' 'Ref. tech. develop' and 'Ref. safety O & M costs,' generally have a large impact on the total cost. While the relative costs, 'T_{melt} rel. cost,' 'P_{vap} rel. cost' etc., have a much lesser cost effect. Often the associated sensitivity is even negative. For the relative costs this may initially seem counter-intuitive because one would surmise that a higher cost of effect/system X should only result in a higher total cost. The relative costs however collectively indicate only the importance of each cost component to calculating the ratio of the configuration-specific to reference design cost. For instance a higher relative cost of T_{melt} means a somewhat lower relative cost of all other relevant components. This outcome is the result of the novel 'cost setting' method discussed in detail in Section 2.8.

Figures 5.49 to 5.53 show the data points and regression results of select configurations. Figures 5.54a and 5.54b show histograms of the regression coefficient of all configurations with 2000 samples, each for a select input parameter, overlaid with a normal curve that has a mean equal to the average coefficient value and a variance equal to the average expected variance of the coefficient. If the analysed configurations are all the same, the histogram of coefficients should match the expected normal distribution of coefficients. The deviation between actual and expected coefficient variance is relatively low, as evidenced by the data in Figure 5.54 and the similar slopes in Figures 5.50 to 5.53. The regression coefficient variance was also tested for several other input parameters and the with the 400 sample dataset. Thus it is expected that the sensitivity analysis results presented for the lowest cost configurations would be similar to results for other configurations as well.

Multi-variable regression analysis validity:

The sensitivity analysis is conducted using multi-variable linear regression that requires several assumptions to hold true, see Section 4.5. Notable assumptions are the linear relationship between the dependent and independent variables, homoscedasticity and a normal distribution of the residuals. There are too many variables to always be confident in the underlying relationship but Figures 5.49 to 5.53 consistently appear to be close to linear. As well, the inherent structure of the evaluation is largely additive and therefore largely linear. The validity of the homoscedasticity and multi-variable normality can be inferred from Figure 5.55, showing the predicted (from multiple regression) and observed costs for the lowest cost configuration of each FHR coolant salt. These two assumptions are discussed Section 4.5 since the residuals in the FHR and MSR evaluations are similar. The author's basic conclusion is that the low variance allows for substantial tolerance of heteroscedasticity and non-normality, thus making the model, specifically the calculated sensitivity coefficients, sufficiently valid.

Table 5.16: Sensitivity of total cost to key input explanatory variables at the lowest cost configurations. The values shown are scaled relative to the standard deviation, thus the sensitivities are listed as $\Delta Cost$, where $\Delta Cost = b_i \sigma_i$. The sensitivities are ordered in decreasing average magnitude.

Variable	Salt 1	Salt 2	Salt 4	Salt 5
Interest rate	16.5	11.8	14.2	16.3
Lifetime (yrs)[a]	-14.7	-12.1	-14.6	-14.1
Fuel pebble manufacturing cost	4.2	7.07	10.5	8.45
Ref. safety OM costs	6.02	6.43	6.69	6.72
Share cap safety costs	6.25	4.98	6.36	6.87
Ref. refuel cap	4.08	5.64	7.12	6.35
Construction time (yrs)	6.78	4.1	5.66	6.51
Non-configuration	4.84	5.14	4.69	5.27
Overrun exp.	4.32	3.49	4.22	4.58
U_3O_8	2.15	3.67	5.78	4.36
Heat exchanger material cost ($\$/m^2$)	3.67	2.92	4.01	4.28
Separative work	1.77	2.95	4.67	3.49
Share tech. develop	3.25	3.29	1.95	3.12
Ref. share modularity	2.76	1.95	2.57	2.83
T_{melt} max	2.66	1.58	1.89	2.11
Safety cap. exp.	-2.5	-1.84	-1.82	-1.87
^{7}Li cost ($\$/kg$)	3.36	-0.155	0.47	1.27
Ref. tech. develop	1.12	1.26	0.673	1.03
Share cap prolif	0.626	0.513	0.866	0.717
Temp. feedback rel. cost	0.0644	0.373	0.668	0.793
Ref. OM prolif	0.354	0.0449	0.55	0.379

[a]Unreasonable correlations with the *Lifetime* variable were used. The *Lifetime* sensitivities listed have been adjusted to obtain the expected sensitivity in the absence of any correlations. See Table A.1 for unadjusted values.

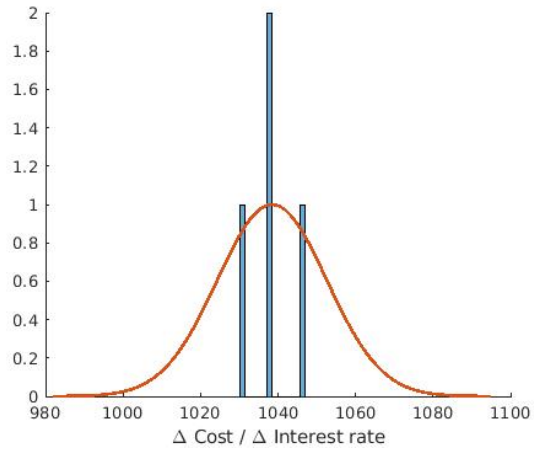

(a) LiNaBeF: interest rate

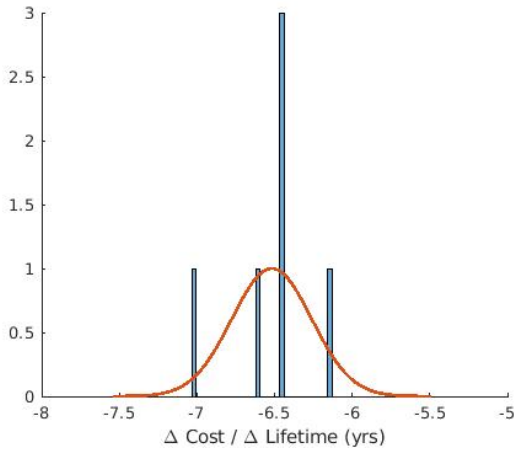

(b) LiNaZrF: lifetime

Figure 5.54: Histogram of the regression coefficient of the input in question for all configurations overlaid with a normal curve with a mean equal to the average coefficient value and a variance equal to the average expected variance of the coefficient.

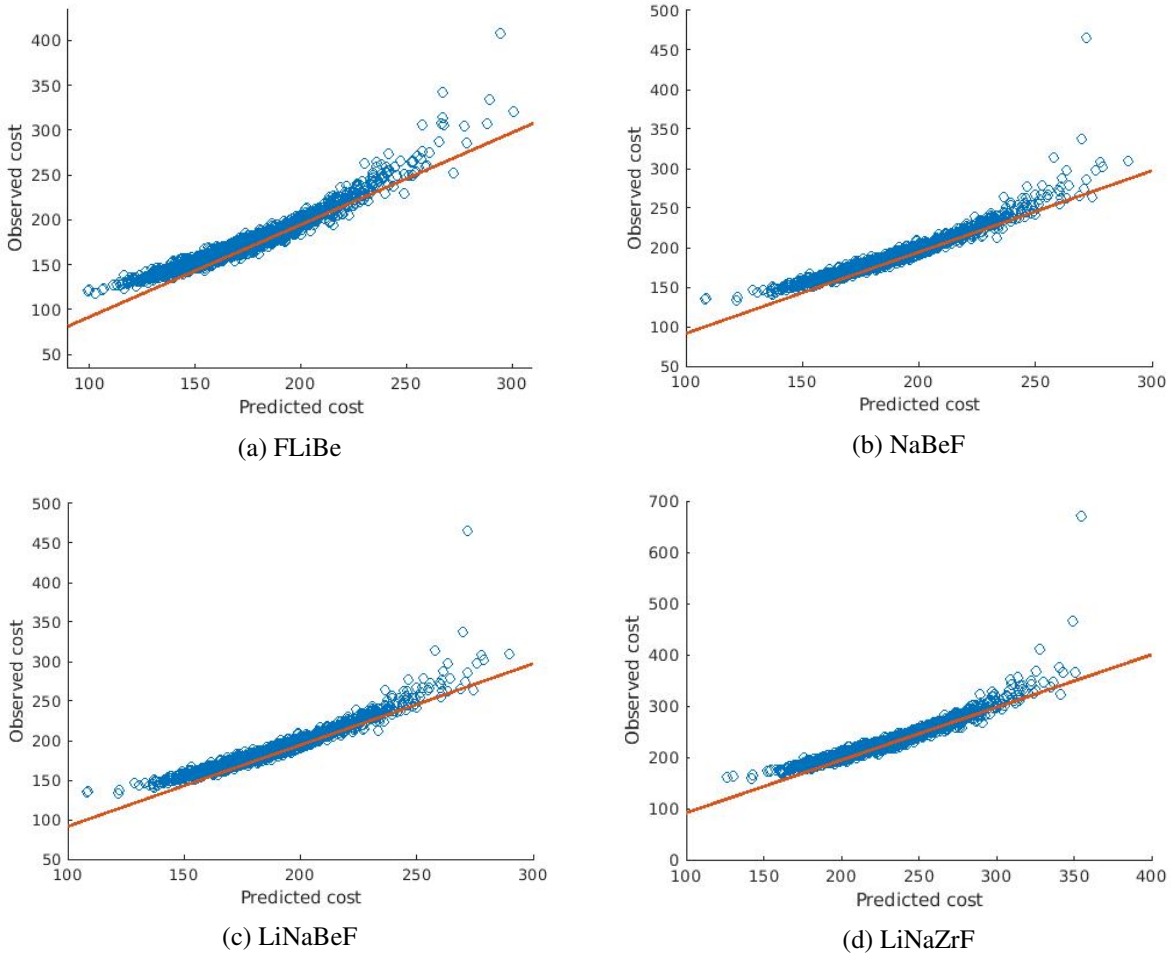

(a) FLiBe

(b) NaBeF

(c) LiNaBeF

(d) LiNaZrF

Figure 5.55: Predicted compared to observed cost for the lowest cost configuration of each FHR coolant salt. The difference between the two values (deviation from the line) represents the residuals.

5.8 Conclusions

A large potential design space of a pebble bed reactor cooled with molten salt (FHR) was analysed. Specifically, five salt compositions and a range of graphite pebble (blank) shares, TRISO fuel particle packing fractions (PF) and inner graphite reflector radii. A large dataset covering salt properties, heat exchanger requirements, expected burnup, fuel temperature, temperature coefficients etc. was developed to supply a comprehensive evaluation program that aims to determine the relative total cost of configurations across the design space.

The double heterogeneity of the fuel particles in the pebbles, and the pebbles in the core, was modelled successfully with no indication of significant error that must be addressed. The fuel (Doppler) temperature coefficient (TC) is consistently negative. While the salt (density) TC is positive for all salt compositions except $2LiF - BeF_2$ salt, commonly referred to as *FLiBe*. Both TCs become more negative with increasing core fuel concentration, either through increasing the PF or decreasing the % of blanks. Most compositions have a net negative TC (fuel + salt), the ones that do not were eliminated from further analysis.

The objective of this analysis was to optimize the geometric configuration and coolant salt composition of a FHR. The results of the evaluation program suggest that the configuration that would result in the most feasible, lowest cost design would use FLiBe. The most attractive alternative coolant salt is NaBeF. It is only a worthwhile alternative if the price of enriched ^{7}Li is high and fuel related costs are low; input parameters related to these costs have regression coefficients that differ most among configurations, see Table 6.18. According to this analysis, such a situation has only approximately a 6% probability of occurring. The order of lowest cost configurations for the remaining salts analysed is quite clear: *LiNaZrF*, *LiNaBeF*, *NaZrF*.

It should be repeated that this analysis only considers costs that are a function of salt composition and analysed geometric design parameters (configuration). It is recognized that the analysis likely even missed some significant, configuration-dependent costs. The total costs are in units of millions of USD, 2021. The total cost, ≈200 million is only $\approx 1/4$ of the total expected value of electricity generated by this reactor.[12] Considering the non-configuration dependent costs is important as it puts the cost differences in the optimization in context; since the relative differences in cost are only $\approx 1/4$ as large, other considerations such as availability of material become more of a concern.

The substitution of some fuel pebbles for graphite pebbles, while adjusting the particle PF to maintain a constant fuel concentration resulted in a consistently higher k_{eff}. For alternative salts (non-FLiBe), the evaluation consistently found that the higher k_{eff} and corresponding lower fuel costs from the addition of graphite pebbles outweighed other effects and lowered the overall total cost. This analysis assumes that pebbles in the core are well mixed, both in the level of burnup and the distribution of graphite pebbles. Though nothing was conceived of that precludes mixing, it is unclear how realistic it is to construct a reac-

[12]The value of 100 MW of electricity produced with a CF of 0.92, by this hypothetical FHR, over the 15 year lifetime, at the time of commissioning is $837 million.

tor adhering to this assumption. If the pebbles deviate significantly from well-mixed or if the management of non-fuel pebbles adds (an) unforeseen complication(s), graphite pebble containing configurations would have a higher cost.

The accuracy of many sections of the study is limited by the accuracy of the property data available. Few salt compositions have sufficiently comprehensive, reliable and accessible data; this study is limited to compositions that do. Neutronics calculations involving salts with lithium would benefit from the availability of thermal scattering data. Many cost components as well as the respective probability distributions are difficult to estimate. Other cost input data would likely result in much different total costs. However the more important relative costs would remain more similar and as a result lead to similar conclusions.

The uncertainty and sensitivity analysis using distributed value input parameters provided some useful insights.

- The variance in total cost of the lowest cost configurations is high ($\sigma \approx$30 - 45) compared to the variance in the difference in cost between the configurations ($\sigma \approx$10). This suggests that the total cost is more susceptible to the circumstances of a FHR build than to the configuration selected.
- The probability that the lower average cost salt has the highest expected cost, if the input parameter set is independent for each configuration (overlap in total cost distribution), is much higher than the probability accounting for the input parameter set (< 0 in average cost difference distribution).
- Total cost distributions have similar relative standard deviation (σ/μ) and similar shape - notably positive skew.
- The expected cost as calculated by uncertainty analysis (the average) is approximately 3-5% higher, depending on the configuration compared to the corresponding fixed values calculation.
- The variance in the cost difference tends to be highest when one of the salts being compared is *FLiBe*, therefore suggesting that the costs involved are the most unique.
- Few of the sensitivity analysis coefficients are particularly notable. The total cost is generally most sensitive to the interest (or discount) rate, lifetime and the fuel pebbles manufacturing cost. The cost is sensitive to the fuel pebble manufacturing cost largely because of the high relative standard deviation ($\sigma/\mu = 0.5$) rather than the high proportion of costs. For the many input variables with low sensitivity coefficients, the lack of a definite basis for the estimated cost distribution is of little consequence; an incorrect estimated cost has little effect on the total cost.

Chapter 6

Graphite-moderated, Circulating-fuel Molten Salt Reactor

6.1 Introduction

The purpose of this chapter is to evaluate the graphite-moderated, circulating fuel Molten Salt Reactor (MSR) design type. The design concept is examined across the entire design space of realistic reactor geometries and fuel salt compositions. Reactor performance data generated from neutronics simulations and several other analyses is used by the evaluation program (see Chapter 2) to come up with cost estimates across the entire design space. The evaluations are pieced together and trends are discerned in order to gain insight into the impact various design features have on reactor performance, draw conclusions on cost sensitivity and determine the lowest cost design configuration.

6.1.1 Selection of Design Parameters Evaluated

The primary objective of this analysis is to determine the configurations that result in low total costs within a design space specified by impactful design parameters. The entire design space is evaluated rather than evaluating parameters independently in order to take into account interdependence. The trade-off associated with evaluating the entire design space is that is necessitates a large number of neutronics runs that increases exponentially with increasing parameters. Available computation resources limited the number of design parameters to three. The variable design parameters selected are:

1. The non-uranium composition of the salt, also referred to as the 'carrier salt.' The carrier salt nuclides can have a considerable parasitic absorption effect and light elements can have some effect on moderation. This is discussed in Section 6.1.2.
2. The molar fraction of salt cations that are uranium (mole % U or molU). The core fuel concentration is a parameter of major importance that should be captured by the variable parameters, mole % U allows for the examination of a wide range of fuel concentrations.

- The maximum range over which molU is evaluated is 0.05 to 0.27. Infeasible salt compositions (eg. prohibitively high melting temperature) are not included.

3. The fraction of core volume taken up by the salt channels (Fuel Fraction or FF). The value of FF is proportional to the amount of both carrier salt and the fuel in the core. Though either molU or FF is able to capture the effect of core fuel concentration, fixing the value of either one fails to allow for the independent analysis of both the core fuel concentration and the concentration of carrier salt isotopes.

    - The FF is generally evaluated over a range of 0.06 to 0.15. The FF of the MSRE was 0.225 [23], the average for the MSBR 0.206 [168] and the average for the FUJI-U3 0.36 [169]. Modern designs face a different set of objectives compared to the reactors of the MSRP or the FUJI, principally a high priority on minimizing capital costs (reduced FF reduces the salt inventory), avoiding high enrichment levels and do not require breeding capabilities.

For each configuration (a particular carrier salt, molU and FF), channel pitch, ^{7}Li concentration (if applicable), heat exchanger size and xenon removal rate are optimized.

Three design specifications were considered for inclusion as variable design parameters in the evaluation but were ultimately left out:

1. Channel pitch: It is a somewhat consequential parameter that can vary significantly and the optimal value for each configuration does not have a intuitively obvious range. Naturally, it would be a candidate variable parameter. It was left out because the other three design parameters are more impactful and is not strongly dependent on the values of the other design parameters.

2. Diameter to height ratio (D:H): The exact optimum (highest k_{eff}) D:H is close to 1.0 for all configurations and is assumed to have a relatively weak dependence on the values of the other variable parameters.

3. Flux flattening through core zoning: Two means of flux flattening were investigated: 1. Variation in salt channel radius. 2. Absorbing adjuster rods. Adjuster rods reduce k_{eff} by magnitudes that may not be tolerable. Channel radius variation demonstrated a relatively small reduction in the peak fast flux and a very small reduction in k_{eff} so it is the preferred means. However, flux flattening was ultimately not included, see Section 6.3.

6.1.2 Salt Selection

A wide range of fuel salts can be used in this type of MSR. The most commonly proposed carrier salt in graphite-moderated, circulating-fuel MSR designs is $^7LiF - BeF_2$ (TMSR-LF2, Fuji etc.), however several other salt compositions, $NaF - BeF_2$ (ThorCon), $NaF - AnF_x$[1] (MSTW) etc. have been proposed as well. The range of proposed salts suggests that trade-offs exist, and a thorough examination of several salts could

[1] An represents an actinide mix.

be useful.

The carrier salt composition must be selected across the range of examined UF_4 molar percentages. The clean (uranium and FP free) carrier salt is typically a mixture at a eutectic point. Ideally the carrier salt is at a molar ratio that results in the lowest melting temperature across the range of molU values examined. Data is not available for all salt compositions to precisely calculate the molar ratios and preforming the necessary thermochemical Gibbs free energy analysis is beyond the scope of this work. Instead tenary phase diagrams are examined to come up with the carrier salt molar ratios. At a molU of 23% the tenary diagram is read to determine the carrier salt composition that results in the lowest melting temperature. Intermediate carrier salt compositions are interpolated between the clean salt and salt with 23% molU. Often the mole percentage of one component is unchanged while another is decreased to allow for the addition of uranium, see Table 6.1

The salts selected for subsequent analysis were the ones deemed most feasible. If the melting point of a composition with a particular mole % U is below the minimum specified value (585 ^{o}C by default) it is deemed infeasible and analysis is not performed on it.

The solubility of Fission Products (FPs) and higher actinides is not included in the evaluation because there is little difference in solubility between compositions, data is unreliable/lacking and it is a complex issue with costs that are difficult to quantify. No fluoride salts were found to have a solubilities low enough to warrant exclusion.

The selection of the carrier salts is largely derived from a process of elemination of salt composition elements and it is influenced by the proposed salts selected in other designs/research projects. The process of elimination used, based off of a process by Grimes [105] for MSRP designs, is discussed in Section 3.2.

A relatively high parasitic absorption cross section threshold and the addition of Uranium allows for more viable salt mixtures than can be evaluated. The scope of this study limits the number of salts to 6. The carrier salt compositions without Uranium are generally a eutectic fluoride mixture of the cations remaining after the process of elimination, Li^+, Na^+, K^+, Be^{2+} and Zr^{2+}. Quaternary salts are not considered because they possess the disadvantages associated with each of the constituent elements and their complexity makes it more difficult to draw conclusions from. Several tenary salts are viable options, $LiF-NaF-KF$ (46.5-11.5-42), $LiF-NaF-ZrF_4$ (26-37-37), $LiF-KF-BeF_2$ (32.5-32.5-35) [170], etc., but due to scope constraints only $LiF-NaF-BeF_2$ (15-58-27) is evaluated. Some mixes are only viable at high mole % U, including LiF, and $NaF-KF$.

Each salt has been proposed as a carrier salt for a MSR design concept or has at least been considered and studied. The salts in general sufficiently differ in composition so as to cover a broad range of nuclear, physical and chemical properties. The salts selected for analysis are listed below. Unless otherwise specified, salts 1, 3, 4 and 5 are analysed across a molU range of up to 5.0 - 22.0% while maximum melting temperatures

limit the analysis of salts 2 and 6 to between 22.0 - 27.0% molU.

1. $2^7LiF - BeF_2$ or FLiBe, 99.995% ^7Li: Many current or past thermal-spectrum, circulating-fuel MSR design concepts select a Li – Be fluoride carrier salt with natural Li enriched to ^{7}Li [16][158]. FLiBe is the chosen coolant of the FHR and has been studied for use in fusion reactors [48][171][172]. The most studied composition, a Li to Be ratio of 2:1, is at a eutectic point with a comparatively low melting temperature [23]. If the ratio differs from 2:1, typically it is a greater share of Li, as in the MSBR design (originally 3.4:1) [24].

2. LiF: The primary motivation for analysing this salt is that it is the selected carrier salt of both the fast-spectrum MSFR and the thermal-spectrum TransAtomic Power reactor designs, see Section 1.6.7. It can be considered that LiF is $^7LiF - BeF_2$ with 0% Be and so the properties of FLiBe of any molar ratio can be roughly estimated by interpolating the results of LiF and 2:1 FLiBe.

3. LiF–NaF–BeF_2, mole % 15-58-27: This salt is analysed because it is proposed for use in the Molten Salt Actinide Recycler & Transmuter (MOSART) design concept, see Section 1.6.7 and the $LiF - NaF - BeF_2$ system has in numerous instances been studied for nuclear applications, incl. [173] [174] [175] [49] [176].

4. NaF–BeF_2, mole % 57-43: This salt has been suggested as a potential carrier/coolant salt by multiple sources [118] [177] and is the chosen salt of the ThorCon reactor, see Section 1.6.7.

5. $NaF - ZrF_4$, mole % 51-49: This salt also has been suggested as a potential carrier salt by multiple sources [118] [177] and was the carrier salt of the ARE, see Section 1.6.4.

6. $NaF - KF$, mole % 40-60: The expense and unreliable supply of ^{7}Li and the chemical toxicity of Be, are the motivation to evaluate salts that do not contain Li and Be. Of the salts listed so far, only $NaF - ZrF_4$ does not contain either isotope/element. K is included because it is not a component of any of the other salts. The remaining available systems are $NaF - KF$, $KF - ZrF_4$ and $NaF - KF - ZrF_4$. $KF - ZrF_4$ is not selected because the chemical similarity of Na and K may result in a lack of differentiation between $KF - ZrF_4$ and $Na - ZrF_4$. $NaF - KF - ZrF_4$ is not selected because of a lack of property data and it may fail to highlight the imapact and general suitability of K. The $NaF - KF - UF_4$/$NaF - KF - ThF_4$ system was researched in the MSRP [174] [23] [111].

6.1.3 The Fuelling Scheme

The reactor models are continuously refuelled. The scheme makes use of the SERPENT 2 [163] mflow card. Using mflow, nuclides can be continuously added or removed at a rate specified by a value mathematically equivalent to a decay constant. The fuelling rate is adjusted between each burn up step in an effort to maintain the target k_{eff}. The target k_{eff}, 1.02, is greater than 1.00 to make available excess reactivity. That excess reactivity would be taken up in a detailed design by the flux flattening mechanism, detectors and other features that are not accounted for in the model. Depletion simulations begin with fresh fuel and the

Table 6.1: Composition of reactor salts with and without uranium. For intermediate amounts of uranium ($0 < U < 23\%$) the concentration of the non-uranium components is linearly interpolated between the two values. The sources generally refer to a phase diagram of the respective salt; not all phase diagrams were found, the referenced data was not verified and the linear assumption does not necessarily result in optimal compositions at each molU.

Composition clean	Composition max molU	Source
$LiF - BeF_2$ (67-33)	$LiF - BeF_2 - UF_4$ (74-2-23)	[175]
LiF —	$LiF - UF_4$ (77-23)	[178]
$LiF - NaF - BeF_2$ (15-58-27)	$LiF - NaF - BeF_2 - UF_4$ (23-39-15-23)	[175] & [179] [a]
$NaF - BeF_2$ (57-43)	$NaF - BeF_2 - UF_4$ (57-20-23)	[111] (pg. 578)
$NaF - ZrF_4$ (51-49)	$NaF - ZrF_4 - UF_4$ (51-26-23)	[111] (pg. 577)
$NaF - KF$ (40-60)	$NaF - KF - UF_4$ (45-32-23)	[119] & [111] (pg. 576) [b]

[a] [179] for clean salt; [175] for an approximate eutectic point with UF_4
[b] [119] for clean salt; [111] for an approximate eutectic point with UF_4

enrichment level that results in a k_{eff} of 1.02. The fuel added is the same composition as fresh fuel but with a somewhat higher enrichment. Fuel salt is removed from the loop at the same volumetric rate as the added fuel. The enrichment of the added fuel salt is a function of the initial enrichment, see Section 6.3.5.

6.1.4 Reference Reactor Key Design Specifications

The analysis requires a reference design, or reference model, from which changes in design parameters are evaluated. The reference geometry is simple; it does not include control rods, detectors or any other components that have a minor impact on reactor neutronics. The basic parameters are listed in Table 6.3 and the geometry is shown in Figures 6.1 and 6.2. Fortunately, proposed graphite moderated MSR designs generally share many of the same characteristics; fluoride salt, a size of roughly 400 MWth, cylindrical core with circular channels in a hexagonal pitch with either LEU fuel or a ^{232}Th - ^{233}U fuel cycle. Key specifications of some of these designs are listed in Table 6.2 and a list of key parameters of the reference design are provided below:

- Power density: The magnitude of the power density was selected in order to achieve the desired core lifetime. The core lifetime is limited by the maximum neutron fluence [180] (more on this in Section 6.1.4). Since multiple proposed reactor designs, including the IMSR and MSTW have a 7 year lifetime (see Section 1.6.7), it was determined that a 7 year lifetime and 12.5 MW/m^3 average power density should be a reasonable, realistic trade-off between core lifetime and power density.

- Share of salt volume in core: The share of the total salt volume in the core (core salt share) has an

Table 6.2: Summary of Design Parameters of Current Circulating-Fuel, Thermal Spectrum MSR Design Concepts.

Parameter	IMSR	TMSR-LF2	ThorCon
Developer	Terrestrial Energy	CAS, China	Martingale
Fuel	5% ^{235}U	Th, U (19.75% ^{235}U)	80 Th, 20 U (19.7% ^{235}U)
Power, MWe	190	168	250
Power, MWt	400	373	557
Moderator	Graphite	Graphite	Graphite
Moderator/core lifetime (yrs)	7	6-8	4
Reprocessing non-volatile FP	No	Yes	No
Carrier salt	—	FLiBe	NaF-BeF_2
Primary reactivity control	FTC	FTC	FTC[a]
Peak salt temperature (C)	600 (secondary)	700 (core)	704 (core)

[a]Fuel Temperature Coefficient

impact on the results of depletion simulations; initially the concentration of Fission Products (FP) is inversely proportional the to core salt share, however, as the salt mix approachs equilibrium the concentration of FP is independent of core salt share. The MSRE had 11260 lbs (5112 kg) of fuel salt, a 0.225 fuel fraction and a core volume 2.88 m^3 [23], so $1/3.36$ of the salt was in core. Since fuel salt is expensive and modern designs often have a more compact loop, the configurations analysed generally have a lower core salt share that largely depends on the core fuel fraction. The volume of salt outside the core is assumed to be 12% of the core volume in many neutronics simulations.

- Reflector thickness: The reflector thickness was selected to be 50 cm with the intent of striking a balance between material requirements and reflector effectiveness. The diffusion length of thermal neutrons in graphite is 53.5 cm and so beyond about 50 cm the reactivity gain of increasing thickness is considerably decreased. One graphite moderated MSR that provides a clear specification of reflector thickness is the Mk-1 PB-FHR [48], the effective reflector thickness (including reflector pebbles) is about 50 cm.

Key assumptions and bounds of the reference design are listed below, they are selected with the aim of resulting in reactor design concepts that are similar to others proposed. As well they narrow the scope, simplify the analysis and remove infeasible options from wasteful analysis:

- Limiting the design to a cylindrical shape where diameter = height.
- Uniform of 22 cm pitch throughout.
- No flux flattening.
- 400 MWth and 12.5 MWth/m^3 power density.

Table 6.3: MSR Reference Reactor Key Specifications.

Parameter	Value
Avg. salt outlet temp	968 K
Avg. salt inlet temp	888 K
Active region D & H	3.44 m
Reflector thickness	0.50 m
Thermal power	400 MWth
Thermal power density	12.5 MW/m^3

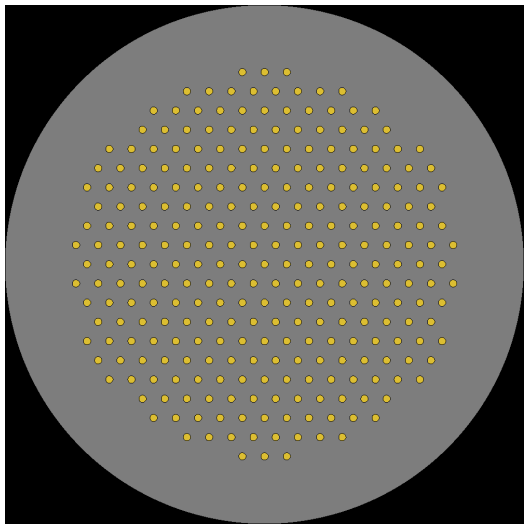

Figure 6.1: Radial view of reactor geometry with a pitch of 19 cm and a diameter & height of 344 cm corresponding to 400 MWth at 12.5 MW/m^3.

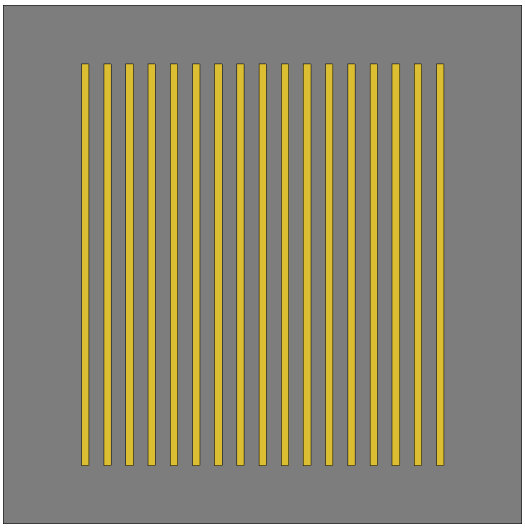

Figure 6.2: Axial view of reactor geometry with a pitch of 19 cm and a diameter & height of 344 cm corresponding to 400 MWth at 12.5 MW/m^3.

Salt temperature

The selected salt temperature is inherently somewhat arbitrary though MSR designs tend to specify an average core salt outlet temperature of approximately 700 ^{o}C (973 K). Since nuclear heat is generated in the circulating salt, there is no heat transfer to a coolant, so the temperature difference through the core is not required to be high, Δ80 ^{o}C is selected. The chosen average core outlet temperature is 715 ^{o}C and so the flux-weighted average temperature is 675 ^{o}C.

Two types of analyses involve the salt temperature: pumping power & heat exchanger size optimization (pump-heatex) and SERPENT 2 neutronics simulations. The pressure drop through each channel is assumed equal, thus the resulting channel velocities are nearly equal. The temperature at the heat exchanger primary salt inlet is the mass-weighted average of the core outlet, and the temperature of the outlet is equal to the core inlet. A temperature rise in the maximum power channel of 50, 65, 80 and 100 ^{o} C is analysed, though only Δ80 ^{o}C is ultimately selected for the reference design and evaluated.

SERPENT 2 simulations were run with a uniform temperature; the difference in important parameters such

as k_{eff} and flux with and without temperature zones was found to be very low compared to other design parameters and reasonably consistent among configurations. Therefore, the uniform, flux-weighted average temperature is applied to the fuel salt throughout the core and assumed to have a negligible impact on relative cost evaluations. For depletion simulations especially, it would be a considerable challenge to accurately model multiple zones due to the continuous mixing.

Graphite lifetime and temperature

The Graphite Lifetime (GL) is the limiting factor in the reactor operational lifetime. Graphite undergoes dimensional changes with fast neutron fluence; the graphite first starts to shrink and then begins to expand quite rapidly. The rate of dimensional change is a function of temperature. The issue is further complicated by the fact that the minimum dimensional size reached is also a function of temperature. However, using experimental data from Molten Salt Breeder Reactor (MSBR) research [181], a relation between the maximum GL and temperature, for core temperatures in the operating range, can be estimated: $\frac{d(GL)}{dT} = 0.00235$, with GL the fractional change and T the temperature in ^{o}C/K. The real relation is non-linear but the difference in temperature between different core configurations is small enough such that the non-linear effects are of minor importance.

The differences in fast flux, and consequentially graphite lifetime, among configurations is not a part of the analysis because fast flux is nearly proportional to power density in a thermal reactors and the power density is equal for all configurations. The channel pitch however can vary over a large range and cause significant temperature differences (ΔT) between the graphite surface and the graphite interior. The ΔT between the graphite surface and graphite interior of the highest power channel of the MSBR design is 65 ^{o}C. This design had a power density of 40 MW/m^3 and a pitch of 12 cm, most current designs and the reference design selected in this chapter, specify a lifetime of approximately 7 years, corresponding to a power density of approximately 12.5 MW/m^3. MSBR findings are used as the basis for calculations of the change in GL due to ΔTs in the graphite. The average ΔT between the graphite interior and surface can be solved according to Equation 6.1:

$$T_{interior} - T_{salt} = \frac{65^{o}\,C}{(40\,MW/m^3)(12\,cm)} \cdot S(MW/m^3) \cdot P(cm) \tag{6.1}$$

Where S is the power density and P is the pitch. Using the $\frac{d(GL)}{dT}$ value, the relative (to a reference pitch) GL can be calculated, see Equation 6.2. In addition to the linear $\frac{d(GL)}{dT}$ assumption, several others are made in the analysis:

- The average temperature of the graphite is equal to average of the surface and interior temperatures.
- The within-hex stresses do no limit lifetime, only the maximum fluence.
- The conductive resistance through the graphite is much higher than the convective resistance with the salt - at pitch of 20 cm the conductive resistance $\approx$ 10x convective.
- $\frac{d^2(ch.rad.)}{d(pitch/2)^2} = 0$, thus the fast flux is constant with pitch.

$$\Delta(GL) = \frac{d(GL)}{dT}(T_{interior} - T_{salt}) \tag{6.2}$$

P. R. Kasten et al. [181] in discussing stresses generated in graphite through irradiation, concluded that 'other factors, such as the influence of dimensional changes on graphite permeability' limit graphite exposure. Thus it is assumed that dimensional changes are the limiting factor of graphite lifetime as opposed to differential distortions and internal stresses.

Since the within-hex stresses are assumed not to limit lifetime, the peak and/or average fast flux within the maximum flux (center) hex are linearly proportional to GL. It is unclear whether the peak flux/fluence or the average is more consequential to the GL so the effect is taken to be the average of the two. The peak/average fast flux are extracted from simulations that vary in molU, FF and pitch.

6.2 Independent Non-Neutronics Optimizations

6.2.1 Xenon Poisoning

Introduction and xenon in the Molten Salt Reactor Experiment (MSRE)

A key advantage MSRs have over other reactor types is the on-line removal of gaseous fission products to achieve any desired, yet feasible, residence time. In this study the xenon poising fraction is optimized for each reactor configuration. The analysis allows for 5 variable inputs:

1. Gas stripping/sparging efficacy
2. Salt penetration into the graphite
3. Mass transfer coefficient to the graphite
4. Share of the salt outside the core
5. (Relative) thermal flux

The xenon transport in these reactors is a complex issue and not well understood. Considerable uncertainty surrounds the movement of xenon between the salt channels and the graphite. The amount that salt penetrates into the graphite and the amount which xenon soaks into the graphite moderator (regulated by the mass transfer coefficient) are consequential unknowns governed by the chemical/structural nature of the graphite and salt; further each are time-dependent as the graphite undergoes structural changes with irradiation. The cumulative ^{135}Xe yield at thermal incident neutron energies for all two main fissile isotopes [182] is shown in Table 6.4. This analysis assumes that 2/3 fissions are from ^{235}U and the remaining 1/3 from ^{239}Pu.

In the MSRE it was found that large changes to the diffusivity of xenon in graphite did not affect the value of the xenon poison fraction. This is because the transfer of xenon to the graphite is governed by the mass transfer coefficient to the graphite surface rather than xenon diffusivity. The contribution of xenon in graphite to the poison fraction is not explicitly specified. However it can be inferred from the relation between the mass transfer coefficient and the poison fraction. The poison fraction at the expected/measured

Table 6.4: Cumulative Xe-135 yield (%) of the three main fissile isotopes [182].

Isotope	Yield % (n Energy 0.0253 eV)
^{235}U	6.539
^{239}Pu	7.608
(2/3) ^{235}U + (1/3) ^{239}Pu	6.895

mass transfer coefficient of approximately 0.0005 cm/s is 1.7% and experimental results can be extrapolated to indicate that the poison fraction would be 1.0% if the mass transfer coefficient were zero (ie. no xenon soaked in the graphite). Therefore the poison fraction of xenon in the graphite of the MSRE is 0.7% or 70% higher than if no xenon had soaked into the graphite [38].

Analytical xenon poison fraction calculation

The xenon poison fraction is solved using the point kinetics simplification. Several key additional assumption are made in the analysis:

- The reactor is at steady state; key variables are constant: reactor power, the penetration of xenon into the graphite (i.e. the volume of salt in the graphite is constant), salt channel xenon concentration and salt-in-graphite xenon concentration.
- The xenon concentration in the bulk graphite is small relative to the infiltrated and channel salt. As a result, the mass transfer depends only on xenon concentration in the fuel channels.
- Xenon production in the graphite occurs from fission of the infiltrated fuel salt in the graphite.
- The salt composition, apart from xenon concentration, is the equal in the graphite and the salt channels.[2]
- The shape of the neutron flux is ignored; it is averaged cross the core, largely because mixing results in a uniform xenon distribution in the salt - thus allowing the use of point kinetics equations. Given the imprecision of the inputs and exploratory nature of this analysis a precise flux distribution requiring computer simulation would be of relatively little benefit.
- Xenon diffuses from the salt channels into the bulk graphite but does not diffuse from the infiltrated salt to either the bulk graphite or the salt channels. This assumption generates little error because the salt is expected to penetrate less than 1% of the graphite volume.

The listed assumptions implicitly assume that the xenon perfectly diffuses once it is in the graphite. The justification for the assumption is that the MSRE program found that the xenon concentration ([Xe]) in the bulk graphite is governed by the mass transfer coefficient to the graphite surface rather than xenon diffusivity [38]. Therefore the amount of Xe entering the graphite is not a function of Xe distribution but rather the

[2]It is recognized that the salt would move back and forth between the graphite and the channels. However this is expected to occur over a much longer time period than the lifetime of xenon in the core. The xenon $T_{1/2} \approx$ 1-5 hours whereas salt lifetime in the graphite $T_{1/2} \approx$ weeks. The salt lifetime in the graphite is however expected to be considerably shorter than consequential changes to the fuel make-up, $T_{1/2} \approx$ months - thus supporting the validity of the uniform salt composition assumption.

Table 6.5: List of constant and variable symbols.

Symbol	Value/units	Meaning
$[N_g]$	cm^{-3}	Concentration of Xe atoms in graphite
$[N_c]$	cm^{-3}	Concentration of Xe atoms in salt
$[N_b]$	cm^{-3}	Concentration of Xe atoms in bulk graphite
λ_X	$2.09 \cdot 10^{-5} s^{-1}$	^{135}Xe decay constant
λ_{strip}		Xe stripping/sparging removal rate
σ_a^X	1.0 Mb ($10^{18}\, cm^2$)	Micro absorption cross section
$N_{X,total}$		Total no. of Xe atoms in core
$P_{X,total}$		Total production of Xe
V_c	cm^3	Salt channel volume
V_l	cm^3	Total salt loop volume
V_g	cm^3	Volume of salt impregnated in graphite
ϕ	cm^{-2}	Flux, averaged over core volume
k	cm/s	Mass transfer coefficient
A	cm^2	Surface area of the salt-channel to graphite boundary

[Xe] of the channel and infiltrated salt.

Equation 6.3 shows the standard xenon concentration balance and Equation 6.4 shows that same concentration balance with stripping (or sparging) of noble gases from the salt mixture. The calculation of the xenon concentration in graphite $[N_g]$ and in the salt channels $[X_c]$ is built upon these equations.

$$[N_X] = \frac{(\gamma_X + \gamma_I) \cdot \Sigma_f \cdot \phi}{[N_X] \cdot \sigma_a^X \cdot \phi + \lambda_X} \tag{6.3}$$

$$[N_X] = \frac{(\gamma_X + \gamma_I) \cdot \Sigma_f \cdot \phi \cdot \frac{V_c}{V_l}}{[N_X] \cdot \sigma_a^X \cdot \phi \cdot \frac{V_c}{V_l} + \lambda_X + \lambda_{X,strip}} \tag{6.4}$$

Table 6.5 catalogues the meaning of each symbol used in the equations of this section.

This model also takes into account xenon migration from the salt to the graphite. This is represented by Equation 6.5. $\dot{m}$ is the mass flow rate to the surface, k the convective mass transfer coefficient, A the surface area, $[N_c]$ the xenon concentration in the salt channels and $[N_b]$ the xenon concentration in the bulk graphite. Units for each variable are shown in the the equation. As discussed in the list of assumptions it is assumed that $[N_b]$ is insignificantly low therefore subsequent calculations will use just $[N_c]$ instead of $([N_c] - [N_b])$.

$$\dot{m}[s^{-1}] = k[cm/s]A[cm^2]([N_c][cm^{-3}] - [N_b][cm^{-3}]) \tag{6.5}$$

The xenon poison fraction is equal to $N_x \cdot \sigma_a^X \cdot \phi$, where σ_a^X is the Xe micro absorption cross section, N_x the number of xenon atoms in the core and ϕ the thermal flux . The σ_a^X is the collapsed cross-section, calculated by integrating the xenon cross-section at an average of 700 ^{o}C over a thermal flux range of 0 to 0.62 eV.[3] N_x

[3]Cross-section data is from the JEFF-3.2 library.

is the sum of N_g (Xe atoms in salt which has infiltrated the graphite), $N_c \cdot \frac{V_c}{V_l}$ (Xe atoms in the salt channels) and N_b (Xe atoms in the bulk graphite). Each is an unknown and so 3 equations are required to solve for them:

$$\Delta N_X = P_{X,in} - P_{X,out}$$
$$\Delta N_X = (\gamma_X + \gamma_I)\Sigma_f \phi V_c - \lambda_X [N_g] V_g - \sigma_a^X \phi [N_g] V_g - (\lambda_X + \lambda_{X,strip})[N_c] V_l - \sigma_a^X \phi [N_c] V_c - \lambda_X N_b - \sigma_a^X \phi N_b \tag{6.6}$$

$$\frac{V_g}{V_c} = \frac{\text{fissions in graphite}}{\text{fissions in channels}}$$
$$\frac{V_g}{V_c} = \frac{\lambda_X [N_g] V_g + \sigma_a^X \phi [N_g] V_g}{(\lambda_X + \lambda_{X,strip})[N_c] V_l + \sigma_a^X \phi [N_c] V_c + kA[N_c]} \tag{6.7}$$

$$\Delta N_b = kA[N_c] - \lambda_X N_b - \sigma_a^X \phi N_b \tag{6.8}$$

ΔN_X is the difference in $P_{X,in}$, the production of and $P_{X,out}$, the removal, decay and burn-out of ^{135}Xe atoms. $P_{X,in} = (\gamma_X + \gamma_I)\Sigma_f \phi$, but more precisely $P_{X,in}$ can be made equal to the input fission rate multiplied by $(\gamma_X + \gamma_I) = 0.06895$. γ_X and γ_I are precise empirical values, V_c is determined by the reactor geometry and V_g is one of the independent variables. Now Equations 6.6, 6.7 and 6.8 can be solved for $[N_g]$, $[N_c]$ and $[N_b]$. First substitute Equation 6.8 into Equation 6.6 and P_X into $(\gamma_X + \gamma_I)\Sigma_f \phi$, to get Equation 6.9. Solve 6.6 and 6.7 for $[N_c]$ then substitute $[N_c]$ into Equation 6.10 to find $[N_{X,g}]$.

$$P_{X,total} = \lambda [N_g] V_g + \sigma_a^X \phi [N_g] V_g + (\lambda_X + \lambda_{strip})[N_c] V_l + \sigma_a^X \phi [N_c] V_c + kA[N_c] \tag{6.9}$$

Simplify,

$$P_{X,total} = [N_g] V_g (\lambda_X + \sigma_a^X \phi) + [N_c](V_l(\lambda_X + \lambda_{X,strip}) + V_c \sigma_a^X \phi) + kA[N_c] \tag{6.10}$$

Rearrange Equation 6.10 to obtain Equation 6.11 and sub into a simplified version of Equation 6.7, Equation 6.12, to produce Equation 6.13.

$$[N_g] V_g = \frac{P_{X,total} - [N_c](V_l(\lambda_X + \lambda_{X,strip}) + V_c \sigma_a^X \phi) - kA[N_c]}{\lambda_X + \sigma_a^X \phi} \tag{6.11}$$

$$\frac{V_g}{V_c} = \frac{[N_g] V_g (\lambda_X + \sigma_a^X \phi)}{[N_c](V_l(\lambda_X + \lambda_{X,strip}) + V_c \sigma_a^X \phi) + kA[N_c]} \tag{6.12}$$

$$\frac{V_g}{V_c} = \frac{\frac{P_{X,total} - [N_c](V_l(\lambda_X + \lambda_{X,strip}) + V_c \sigma_a^X \phi) - kA[N_c])}{\lambda_X + \sigma_a^X \phi}(\lambda_X + \sigma_a^X \phi)}{[N_c](V_l(\lambda_X + \lambda_{X,strip}) + V_c \sigma_a^X \phi) + kA[N_c]} \tag{6.13}$$

Table 6.6: Delayed Neutron Fraction (DNF) and Yield Data.

Isotope and Energy	Fraction (DNF)	Prompt Yield	Delayed Yield
^{235}U Thermal	0.00680	2.366	0.0162
^{239}Pu Thermal	0.00238	2.876	0.00686

Isolate $[N_c]$,

$$P_{X,total} = [N_c](\frac{V_g}{V_c}((V_l(\lambda_X + \lambda_{X,strip}) + V_c\sigma_a^X\phi) + kA) + (V_l(\lambda_X + \lambda_{X,strip}) + V_c\sigma_a^X\phi) + kA) \tag{6.14}$$

$$P_{X,total} = [N_c](\frac{V_g}{V_c} + 1)((V_l(\lambda_X + \lambda_{X,strip}) + V_c\sigma_a^X\phi) + kA) \tag{6.15}$$

$$[N_c] = \frac{P_{X,total}}{(\frac{V_g}{V_c} + 1)((V_l(\lambda_X + \lambda_{X,strip}) + V_c\sigma_a^X\phi) + kA)} \tag{6.16}$$

And for $[N_g]$ sub in $[N_c]$ from Equation 6.16 into Equation 6.9.

$$P_{X,total} = (\lambda_X + \sigma_a^X\phi)[N_g]V_g + \frac{P_{X,total}(V_l(\lambda_X + \lambda_{X,strip}) + V_c\sigma_a^X\phi + kA)}{(\frac{V_g}{V_c} + 1)(V_l(\lambda_X + \lambda_{X,strip}) + V_c\sigma_a^X\phi + kA)} \tag{6.17}$$

And simplify,

$$(\lambda_X + \sigma_a^X\phi)[N_g]V_g = \frac{P_{X,total}((\frac{V_g}{V_c} + 1)(V_l(\lambda_X + \lambda_{X,strip}) + V_c\sigma_a^X\phi + kA) - P_{X,total}(V_l(\lambda_X + \lambda_{X,strip}) + V_c\sigma_a^X\phi + kA)}{(\frac{V_g}{V_c} + 1)(V_l(\lambda_X + \lambda_{X,strip}) + V_c\sigma_a^X\phi + kA)} \tag{6.18}$$

$$[N_g] = \frac{P_{X,total}((\frac{V_g}{V_c})V_l(\lambda_X + \lambda_{X,strip}) + V_c\sigma_a^X\phi + kA)}{(\frac{V_g}{V_c} + 1)(V_l(\lambda_X + \lambda_{X,strip}) + V_c\sigma_a^X\phi + kA)(\lambda_X + \sigma_a^X\phi)V_g} \tag{6.19}$$

$$[N_b] = \frac{kA[N_c]}{(\lambda_X + \sigma_a^X\phi)} \tag{6.20}$$

6.2.2 The Delayed Neutron Fraction (DNF or β) and the Effective DNF (β_{eff})

The delayed neutron fraction is the share of neutrons born from precursor fission product nuclides. All else being equal, a higher DNF results in a greater margin to prompt criticality and a longer reactor period. The DNFs of MSR applicable fissile isotopes are listed in Table 6.6 (sources: [183] verified with [184]).

Fuel circulation complicates the issue of Delayed Neutrons (DN). Fissions occur over a spatial distribution inside the core and DNs are released a significant time period after the fission reaction has taken place. This means that the spatial distribution of the delayed neutron emission is shifted depending upon the decay constant value, λ, circulation velocity and the geometry of the reactor. This effect can be conveniently

Table 6.7: Results of 6-group Delayed Neutron Calculations for the MSRE [185].

	1	2	3	4	5	6
$t_{1/2}$	55.7	22.7	6.22	2.30	0.61	0.23
θ	0.364	0.371	0.458	0.709	0.960	0.994
P_i/P_{pr}	0.676	0.718	0.868	0.906	1.010	1.031
$B_{eff} \cdot 10^4$	0.52	3.73	4.99	16.98	7.18	2.77
$B \cdot 10^4$	2.11	14.02	12.54	25.28	7.40	2.70

quantified by adjusting the DNF (β), to create a smaller value, β_{eff}. The ratio β_{eff}/β is equal to the number of next-generation neutrons that a DN will on average produce divided the neutrons that it would produce if the fuel were static [185]. A function is developed to calculate β_{eff}/β over a range of thermal-spectrum circulating fuel reactor geometries and the results are verified against analysis conducted for the MSRE and Molten Salt Fast Reactor (MSFR).

β_{eff} in the MSRE: The calculation of the β_{eff} at ORNL for the MSRE takes an analytical approach. This summary of the methodology and finding is based on work by Haubenreich [185]. In circulating-fuel reactors three effects should be taken into consideration:

1. DNs emitted outside the core.
2. The spatial distribution of DN emission within the core. (ie. greater share being emitted close to the edge of the core and a smaller share emitted from the center)
3. DNs have a softer energy spectrum. The reduced energy of DNs has the effect of reducing leakage.

Each of these effects were incorporated into the calculation of β_{eff}. Delayed neutrons were divided into the usual 6 groups ($\sum_{i=1}^{6} \beta_i$), β_{eff} was solved individually for each group and all groups summed at the end.

$$\beta_{i,eff} = \frac{\beta_i \theta_i P_i}{P_{pr}} \tag{6.21}$$

The calculation of β_{eff} begins with Equation 6.21. Where θ_i is the probability of DNs being emitted inside the core, P_i is the DN non-leakage probability and P_{pr} is the prompt neutron non-leakage probability. P_{pr} is a function of the flux shape. The final β_{eff} value works out to be:

$$\sum \beta_{i,eff} = \sum \frac{\beta_i \theta_i P_i}{P_{pr}} = 0.003617 \ (361.7 pcm) \tag{6.22}$$

The physical delayed neutron fraction of ^{235}U is 0.006405 (640.5 pcm), thus β_{eff} of the MSRE is 56.5% that of physical ^{235}U. Table 6.7 lists the key values calculated for each of the six groups.

Effective delayed neutron fraction calculation

A program modelled on analytical calculations of the MSFR [186] has been built to calculate the β_{eff} over any range of geometries required in the evaluation program. This model is only able to calculate the correction factor (β_{eff}/β) for cylindrical geometries; MSR design concepts are (near) universally cylindrical and

the evaluation program assumes all configurations are cylindrical as well.

Aufiero, 2014 [186] uses three methods to calculate the correction factor: analytical, deterministic and Monte-Carlo:

1. Analytical approach: An integral transform operator calculates the decay probability of a precursor in position r_o , given that it was produced in the position r. The geometry is discretized and calculations are looped over all positions.
2. Deterministic: Precursor transport is modelled using CFD techniques.
3. Monte Carlo: SERPENT-2 is extended to account for the presence of circulating fuel.

The deterministic and Monte Carlo methods are more accurate and precise but are beyond scope and require considerably more computing power than the analytical approach, for analyzing a large range of geometries this is impractical. There is considerable deviation between the uniform and non-uniform velocity case using the analytical approach on the MSFR see *Aufiero, 2014* [186]. The MSFR has an open core, with significant velocity differences; the evaluation program however concerns a thermal spectrum core with equal-size channels and an equal pressure drop through each channel. Though the higher temperature central channels have a higher average temperature and thus lower density and higher velocity, the effect is small enough that the velocity between channels and within each channel can be assumed uniform for the purposes of this analysis.

The replicated analytical model was run over the MSFR geometry with both 25 bins in the axial/radial directions and 16 bins. Despite being a fast reactor the MSFR in the uniform velocity case is a valid comparison because salt moves unidirectionally through a cylindrical core. These runs are shown in Figure 6.3 over a range of circulation times and overlaid with results from *Aufiero, 2014*.

The analytical model was also tested against the calculated results [185] of the MSRE. The MSRE calculations accounted for the difference in neutron energy of delayed and prompt neutrons, it was found to effect the correction factor by 5.5% in the most effected neutron group. In Figure 6.4 the analytical model is compared against MSRE results that both do account for and do not account for the different emission spectra. Relevant design parameters of the MSRE used in the calculation are provided in Table 6.8. The validation against the MSFR and MSRE calculations suggest the model is sufficiently accurate to determine the effect geometric and operational designs parameters have on the correction factor.

The relationship of the correction factor to relevant design parameters is considered. For the analytical model with set fissile material data the correction factor is a function of 3 variables: extrapolation distance relative to the geometric diameter/height (assuming diameter = height and a cylindrical core), circulation time (T) and the fraction of total salt in the core (y). The correction factor over a feasible range of T and y values is provided in Figure 6.5 for ^{235}U fuel. The correction factor calculation assumes the same 2/3:1/3, ^{235}U:^{239}Pu fission ratio as in the Xe poison fraction analysis (Section 6.2.1).

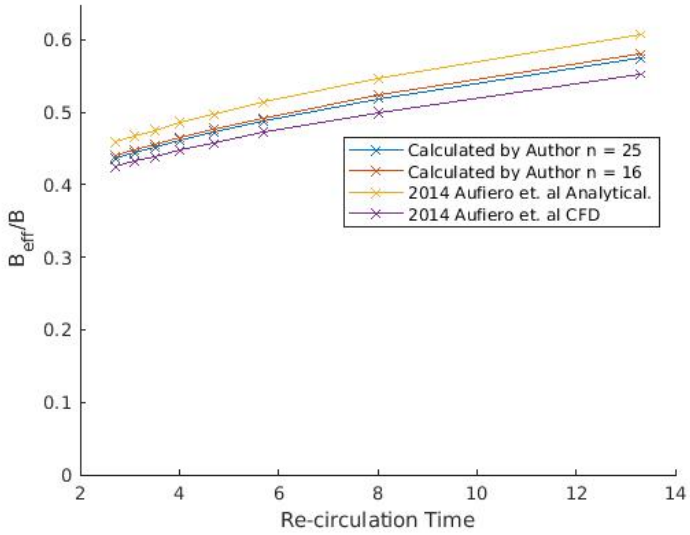

Figure 6.3: Calculation of β_{eff}/β of the MSFR calculated by the author compared to the calculation by Aufiero et. al., 2014 [186]. The number of bins (n) used in both the radial and axial direction are indicated. The author's calculation uses the same analytical methodology as Aufiero et. al.

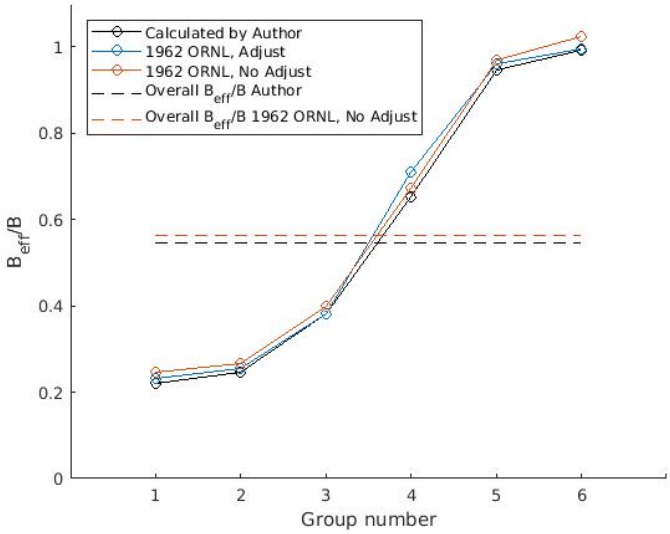

Figure 6.4: Comparison of calculated β_{eff}/β to the β_{eff}/β calculated by ORNL for the MSRE in 1962. ORNL results adjusted to account for the softer delayed neutron spectrum (Adjust) and not adjusted (No Adjust) are both shown.

Table 6.8: MSRE Geometry Relevant to the β_{eff}/β Verification Calculation.

MSRE Geometry	Measurement used in [185]	SI
Radius	27.75 in	0.705 m
Height	68.9 in	1.75 m
T (circ. time)	25.82 s	25.82 s
y (frac. salt in core)	0.362	0.362

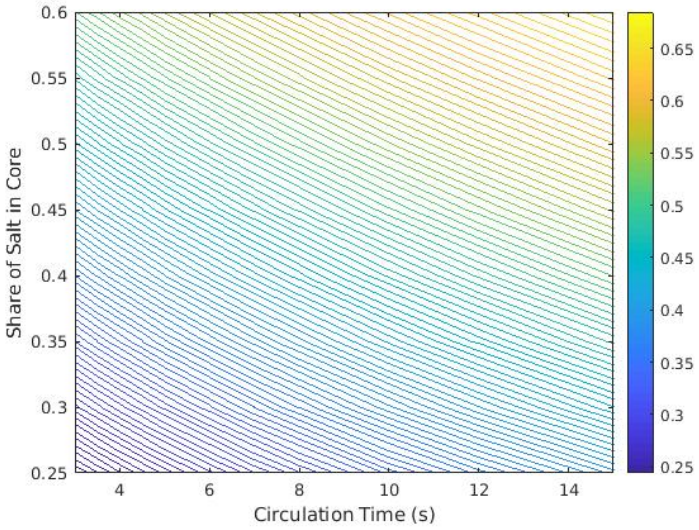

Figure 6.5: β_{eff}/β as a function of circulation time and share of salt in core. Diameter and height of 2.24 m; extrapolation distance beyond core boundary of 18 cm.

6.2.3 Pumping Power and Heat Exchanger Sizing

Pressure drop in the MSR core

A model has been built to calculate the Pressure Drop (PD) and required pumping power though a thermal-spectrum, circulating-fuel reactor core. The Head Loss (HL) is calculated initially, then converted to PD and ultimately pumping power. The HL through the core is calculated analytically through first principles as a function of pitch, fuel fraction, thermal power and average inlet to outlet salt temperature difference (average ΔT). The reactor geometry is discretized into individual salt channels and each channel is discretized into a specific number (16) of axial cells of equal length. To model a realistic scenario where the inlet and outlet salt are at the same pressure, the program iterates toward a solution where the HL is the same in each channel and the maximum temperature change is equal to the input value. Physical properties of the salt, density, viscosity and velocity are tracked individually for each cell to account for variability with temperature.

The HL in a MSR core is a function of many variables, the salt properties, the temperature difference(s) through the core, pitch, fuel salt fraction, reactor power, power density and fission distribution. The power density and fission distribution are constant across all configurations for which the HL is calculated. Although the HL is calculated across a range of average salt ΔT values and thermal power, the evaluation program only uses $\Delta 80\ ^{o}$C and 400 MW.

Assumptions in the HL calculation model:

- Conduction between salt and graphite is insignificant
- Fully developed flow

- The PD is equal through all salt channels
- Heat capacities used are the heat capacities at 700 ^{o}C
- Reynolds and Prandlt number are calculated using the average temperature of the salt and applied across the length of the heat exchanger

Set of variable input values over which calculations are conducted:

- Pitch: 10, 12, 15, 19, 22, 26, 30, 34, 40 cm
- FF: 0.05, 0.06, 0.075, 0.094, 0.12, 0.15, 0.19
- Salt composition: All 6 salts at a 0 % U and 23% U

Head Loss (HL) model calculation process:

1. Divide the core by each fuel salt channel and each of those channels into 16 equally sized length-wise (axial) bins. Each of the separations is referred to as a 'zone.'
2. Run SERPENT 2 simulations that measure the fission response function in the 16 length-wise bins in each lattice position (fuel salt channel). Size the data so it can be used to calculate the heat generated in each zone.
3. Assign an initial guess of an average velocity in each zone.
4. For all zones, take the inlet temperature and add the temperature rise in the zone, see Equation 6.23.
5. Using the new temperatures, calculate the density and viscosity in each zone.
6. Adjust the velocities so that (velocity)·(density) is constant along each channel length.
7. Calculate the HL in each zone using Equation 6.24 where friction(Re) represents the friction factor as calculated in Equation 6.25.
8. Identify the channel with the highest HL/PD, adjust the velocity (u) of all the other channels to match, then reduce u in all channels by a common factor in order to maintain an equal average u.
9. Find the channel with the highest dT, calculate the difference compared to the target dT and adjust the u in all zones by a common factor such that dT = dT.
10. Test if the process has converged or else repeat.

$$dT(c,z) = \frac{\dot{g}(c,z)}{A \cdot u(c,z) \cdot \rho(c,z) \cdot C_p} \tag{6.23}$$

Where c is the channel reference, z is the axial position, dT the change in temperature across the zone, $\dot{g}$ the zone power, u the average velocity, ρ the density (a function of temperature).

$$HL(c,z) = friction(Re) \cdot \frac{1}{2 \cdot 9.81} \cdot u(c,z)^2 \cdot \frac{\frac{L}{z}}{D} \tag{6.24}$$

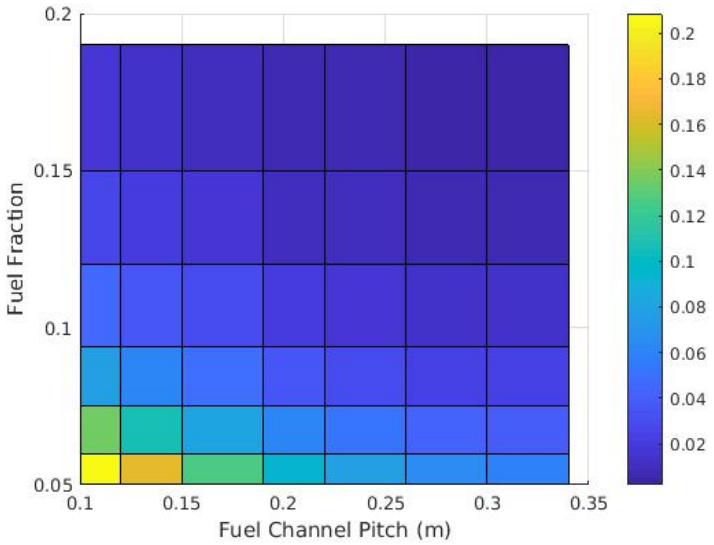

Figure 6.6: Head Loss (m) across the core as a function of pitch and fuel fraction with $LiF - UF_4$ (77.5-22.5) salt at an average core ΔT of 50^{o}C and a thermal power of 400 MW. Where lines meet in the heatmap indicate the data points.

Where friction is the friction-factor function, Re is the Reynolds number, L is the channel length and z the number of axial divisions. In Equation 6.25, ff is the friction factor (thus $ff = friction(Re)$).

$$0 = \frac{-1}{\sqrt{ff}} + 1.93log(Re\sqrt{ff}) - 0.537 \tag{6.25}$$

Results

The program was run over a range of pitch and fuel fraction values at an average power density (12.5 MW/m^3), a sample of the results are shown in Figure 6.6. The power density is proportional to a graphite lifetime of approximately 7 years and is the same as used throughout this chapter. The results shown are with an average temperature increase of 50 ^oC though the core and a total power of 400 MWth. The evaluation program converts HL to PD and ultimately to required pumping power. The results confirm that PD is proportional to $1/(FF)^2$ and $1/(pitch)^2$ as expected.

Pumping power and molten salt heat exchanger size optimization

There is a trade off between required pumping power and the size of the heat exchanger, as shown in Figure 6.7. The purpose of the analysis is to optimize the pumping power requirements and the heat exchanger pipe area (size) for each of the candidate salts. The setup, assumptions and procedure to arrive at a required pumping power and pipe area is discussed in Section 3.3. This section only discusses and presents the results.

In order to cost optimize the circulation velocity a value must be assigned to \$/m^2 of heat transfer surface area and \$/kW of pumping power required. The expected salt heat exchanger cost is \$3500/m^2. To account

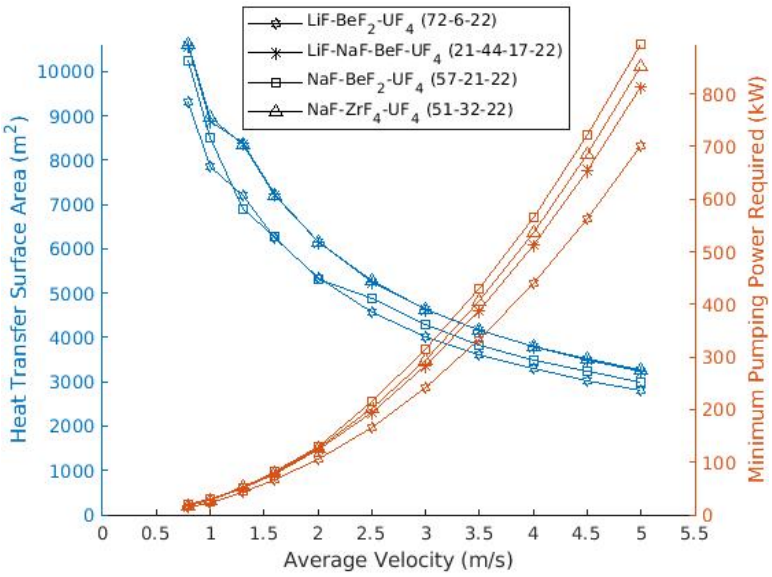

Figure 6.7: Heat transfer surface area and pumping power required in a MSR heat exchanger with an 50 ^oC average ΔT for select salts at mole U of 22%. Several surface area curves are not smooth because the Reynold's number switches from below to above 12000, and hence the empirical correlation switches, at $\approx 1.5 m/s$.

for pump inefficiency, secondary salt pump power and pump capital costs, the expected total pumping related costs are estimated to be 2.0 times the cost of electricity of a 100% efficient pump. Both the heat transfer surface area and pumping power prices are evaluation program variables that are given a probability distribution, see Section 5.6.3. The optimum circulation velocity as a function of pumping power cost and heat exchanger surface area is presented in Figure 6.8 and the corresponding total cost in Figure 6.9.

6.3 Neutronics Simulations

The core graphite structure can be arranged in many different ways, core designs of MSRs generally come in one of two forms: 2-D rectangular or 2-D hexagonal prism with vertical fuel channels. The hexagonal prism is chosen because it is the more common design of the two and the offsetting rows/columns result in a more even distribution of neutrons. The cross-sectional area of the prisms, a value proportional to the square of the pitch, is a key parameter that must be determined. The MSRE had a semi-rectangular structure with pill-shaped fuel channels between graphite rods, see Section 1.6.4, with an effective cross-section of 3.125 in^2 (20.16 cm^2) per fuel channel. The rational behind the small fuel channels was to minimize the fast flux gradient, thereby maximizing core lifetime. The primary trade-off associated with fuel channels of small pitch is a reduction in k_{inf}, see Moser, Wheeler, Chvala, 2017 [187]. Other more modern studies and reactors designs don't employ such a small cross-section. For instance, A study from 2012 on the effects of core-zoning, K. Nagy et al [188], had a reference design with a 140 cm^2 cross-section and the MSTW burner has a cross-section of approximately 120 cm^2 [16]. The channel pitch can be a more convenient metric so it is the preferred metric in subsequent analysis.

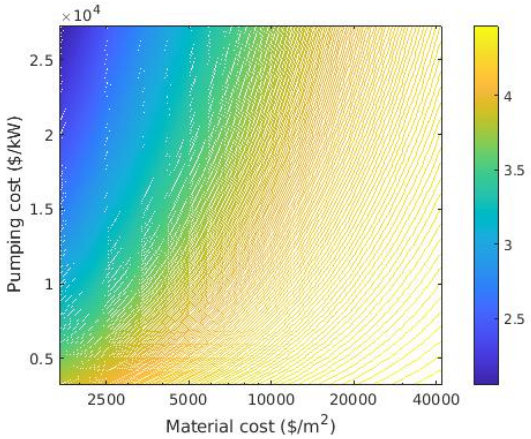

Figure 6.8: Optimum average velocity [m/s] of $LiF - BeF_2$, 22% mole U salt in the primary loop over as a function of heat transfer surface area and pumping power cost.

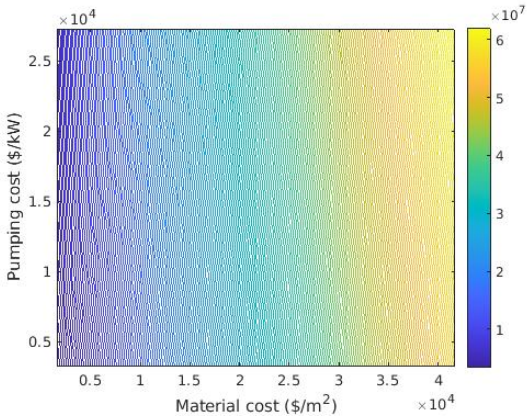

Figure 6.9: Minimum total cost ($) using $LiF - BeF_2$, 22% mole U salt as a function of heat transfer surface area and pumping power cost at the reference power level.

Multiple sets of neutronics simulations were conducted over the range of evaluated salt compositions and geometries. For each set of simulations the range of configurations, objective and key additional details are discussed along with a presentation of results. Two simulation types were run: static and depletion. All simulations have a target k_{eff} value of 1.02; 20 mk is added to account primarily to account for flux flattening which is not included, and also for $\Delta\beta_{eff}$ (delayed neutron precursor movement), sensors, etc. All neutronics simulations are conducted using SERPENT 2.1.30 [163] with JEFF 3.1.1 nuclear data [108].

6.3.1 Static, Enrichment

- Objective: Primarily to determine minimum required enrichment for the range of configurations evaluated.
- Data points: Variation of salt (all 6 salts), molU (0.05, 0.065, 0.08, 0.10, 0.13, 0.17, 0.22, 0.27) and FF (0.05, 0.06, 0.075, 0.094, 0.12, 0.15, 0.19)
- Details: The target k_{eff} is achieved through iteration with varying uranium enrichment (en). An initial guess of en is made and it is adjusted in subsequent iterations according to Equation 6.26. The relaxation variable is continuously corrected based on the accuracy of previous iterations and the precision (number of simulated neutrons) increases with increasing accuracy of the previous run. Once the k_{eff} is adequately accurate in relation to the target and precise the enrichment value is saved and the process moves on to the next configuration.

$$en = en \cdot (1 + relaxation(\Delta k_{eff})) \tag{6.26}$$

- Results: see Figures 6.10, 6.11 and 6.12.

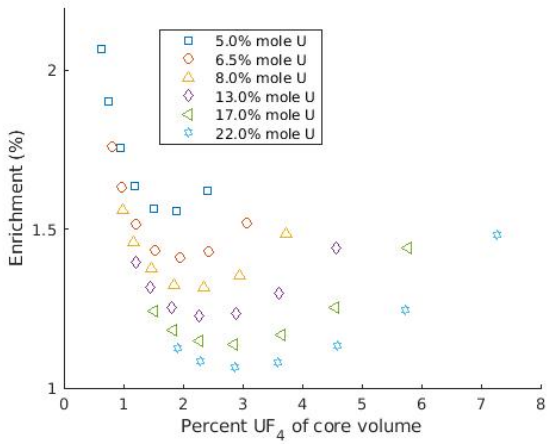

Figure 6.10: Minimum required uranium enrichment for configurations with $LiF - BeF_2$ salt as a function fuel concentration.

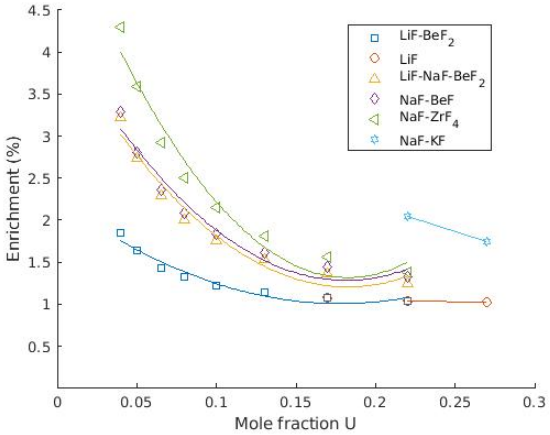

Figure 6.11: Minimum required uranium enrichment as a function of mole % U at a fuel fraction of 0.094.[a]

[a]The interpolations are provided to highlight trends and add clarity, the shape of the curves do not have a physical meaning.

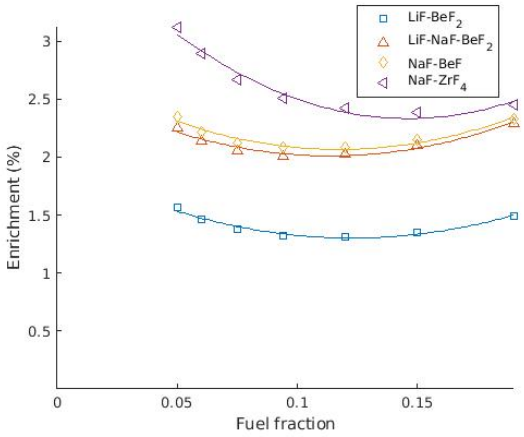

(a) 8.0% mole U

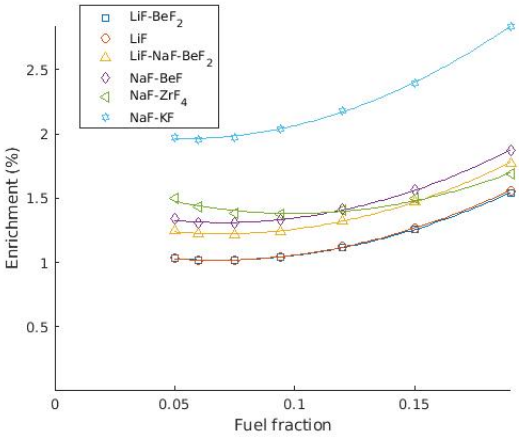

(b) 22.0% mole U

Figure 6.12: Minimum required uranium enrichment as a function of fuel fraction.[a]

[a]The interpolations are provided to highlight trends and add clarity, the shape of the curves do not have a physical meaning.

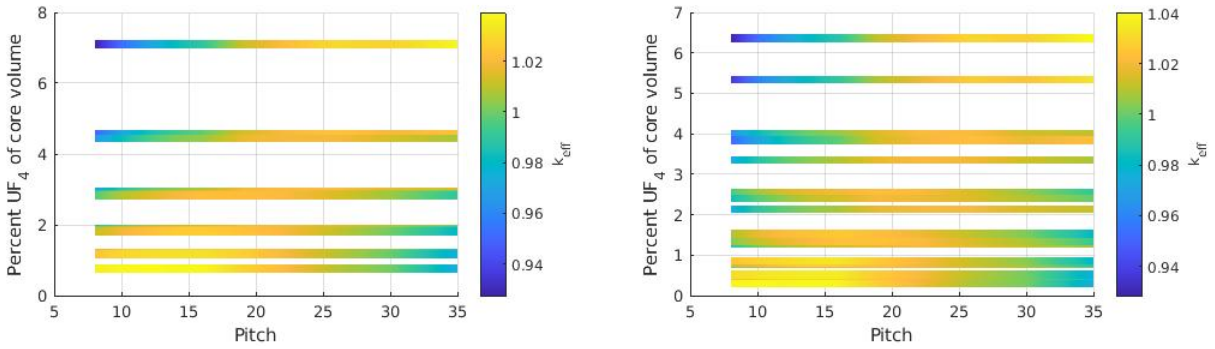

(a) $LiF - BeF_2$ (b) $NaF - ZrF_4$

Figure 6.13: k_{eff} vs pitch at several core fuel concentrations.

6.3.2 Static, Pitch

- Objective: Determine the k_{eff} as a function of pitch.
- Data points: Variation of salt (all 6 salts), molU (0.05, 0.08, 0.13, 0.22, 0.27), FF (0.06, 0.094, 0.15), pitch (8, 10, 12, 15, 19, 22, 26, 30, 35)
- Details: All configurations use the enrichment value from the *Static, enrichment* simulations. k_{eff} as a function of pitch depends on the salt, molU and FF so simulations were run at multiple salt/molU/FF configurations.
- Results: A sample of k_{eff} vs pitch for the core fuel concentrations analysed are shown in Figures 6.13 with $LiF - BeF_2$ and $NaF - ZrF_4$ carrier salts. Notice that the fuel concentration is shifted slightly higher in $LiF - BeF_2$ compared to $NaF - ZrF_4$ even though the identical mole fractions and FFs were analysed; this is due to the higher volume fraction of UF_4 in $LiF - BeF_2$ carrier salt. Some molU and FF combinations result in nearly equal fuel concentrations so the horizontal gradient bars overlap. The pitch with the highest k_{eff} is only a function of molU and FF insofar as it effects core fuel concentration. The remaining results for the other salts are provided in Appendix A.2.

6.3.3 Other Simulations

Static, peak fast flux

- Objective: The fast flux within a single hex and throughout the core broadly is examined.
- Data points: Variation of molU (0.08, 0.17), FF (0.06, 0.15) and pitch (10, 14, 19, 26, 35)
- Details: The neutronics simulations are 2D, periodic boundary conditions, covering the area of one hex. Flux measurements used an energy grid with boundaries of $1{\cdot}10^{-11}$, $6.25{\cdot}10^{-7}$, $1{\cdot}10^{-5}$, 10 MeV and a 40 by 40 equally-spaced dimensional grid.
- Results: Full plotted/tabulated results are not provided. Trends found include:

  - An increasing pitch results in a decreasing avg. fast flux (-8% over the pitch range) and increasing peak fast flux (+38% over the pitch range)
  - An increasing fuel fraction results in a decreasing peak fast flux (-17% over the fuel fraction range) and less than a 5% impact on the average fast flux
  - Mole % U has less than a 5% impact on average and peak fast flux over the range evaluated.

Static, flux flattening

- Objective: Test potential means of flux flattening to obtain an estimate of the k_{eff} loss and relative peak fast flux reduction as well as to determine if a comprehensive flux flattening optimization is worthwhile.

- Details: Two means of flux flattening were investigated: 1. Variation in the salt channel radius, 2. Absorbing adjuster rods. The fast flux is measured as the Fast Flux Peaking Factor (FFPF) rather than the absolute peak fast flux because it is easier to understand and the fast flux relative to the unflattened configuration is the key measure of interest. Adjuster rods cause a large and predictable reduction in peak fast flux, however they reduce k_{eff} by magnitudes that may not be tolerable. In the configurations analysed, the channel radius variation generally causes only a small reduction in the FFPF. It generally resulted in a more moderate, even beneficial, change in k_{eff}. Given the inconsistency and strong configuration-dependence of the channel diameter results, it is difficult to draw conclusions, however it is suspected to be the preferred flux flattening mechanism. Flux flattening was not included primarily because it is expected to have little impact on the relative cost of various reactor configurations and requires considerable computational resource.

- Results: See Figures A.21 and A.22 in Appendix A.2. The Fast Flux Peaking Factor (FFPF) and the $\frac{k_{eff}}{dFFPF}$ are provided for several configurations.

6.3.4 Depletion, Excess k_{eff}:

- Objective: Relate differences in k_{eff} to the difference in the flow rate of fresh fuel salt into the core. (find $\frac{d(\dot{m})}{dk_{eff}}$)
- Data points: Variation in molU (0.05, 0.08, 0.13, 0.22), FF (0.06, 0.15) boron concentration in the graphite (0, $1{\cdot}10^{-6}$, $2{\cdot}10^{-6}$). 39 burnup steps are used to simulate a 7 year period.
- Details: Several calculations in the evaluation program involve determining the increase/reduction in cost associated with a reduction/increase in k_{eff}. Several variable factors increase/decrease k_{eff}; the movement of delayed neutron precursors, ^{135}Xe poisoning, FP removal, etc. The cost of a reduction in reactivity cannot be determined directly. It is the sum of cost to adjust the initial enrichment back to criticality plus the cost to adjust the rate of fresh fuel salt addition into the core. For instance, a lower β_{eff}/β increases the required initial enrichment and refuelling requirements. This cost varies among the different salt configurations. To calculate $\frac{d(\dot{m})}{dk_{eff}}$, a means of altering k_{eff} that is constant with burnup must be introduced. Parasitic absorption, in the form of boron, is introduced to the graphite

at 3 different concentrations to simulate an addition of negative reactivity. The graphite material is not set as burnable and so has a constant composition over the course of the burnup run. Depletion calculations are run at the molU, FF and boron concentrations listed. First the total amount of salt added over the depletion simulation is calculated as per Equation 6.27. That result is then used to calculate $\frac{d(\dot{m})}{dk_{eff}}$ in Equation 6.28.

$$M = \sum \dot{m}_i V_i B_i F \tag{6.27}$$

Where M is the total added salt [cm^3], $\dot{m}_i$ the fractional feed rate, V the total salt volume and B the burnup [days] of step i. F is factor to convert B into [s].

$$\frac{d(\dot{m})}{dk_{eff}} = \frac{M_B - M_0}{keff_B - keff_0} \tag{6.28}$$

Where M_B and $keff_B$ represent the values of the borated run, while M_0 and $keff_0$ are the values without boron.

This analysis uses $LiF - BeF_2$ salt and a channel pitch of 19 cm. The salt feed is composed of 5% enriched uranium. The initial enrichment varies with each run to achieve a target k_{eff} of 1.02, the enrichment values used are those determined in the analysis discussed in Section 6.3.1. Computational expense limited the analysis to one carrier salt and one feed enrichment level.[4] The salt composition is expected to make little difference since differences in moderation and parasitic absorption are largely captured by the core fuel concentration and initial enrichment. There is little difference in cost across a wide range of feed-rates/feed-enrichments because of the cost-per-unit-volume and total-salt-volume trade-off. Moreover the $\frac{d(\dot{m})}{dk_{eff}}$ is not nearly as consequential to the overall feed rate cost as the base feed rate (Section 6.3.7).

- Results: See Table 6.9 for $\frac{d(\dot{m})}{dk_{eff}}$ values and Equation 6.29 for the corresponding regression equation; $[UF_4]$ is the core fuel concentration. Long depletion runs are required so the dataset is limited. The core fuel concentration and minimum enrichment explain nearly all of the variance ($\approx$25% for $[UF_4]$ and 70% for minimum enrichment) so the results are adequately useful for the evaluation.

$$\frac{d(\dot{m})}{dk_{eff}} = -27.25 + 1.462 \cdot [UF_4] + 23.19 \cdot [^{235}U] \tag{6.29}$$

6.3.5 Depletion, Enrichment of Added U:

- Objective: Approximate the optimal enrichment of the added fuel salt.
- Data points: Variation in molU (0.065, 0.22), FF (0.06, 0.15), enrichment (3, 4, 5, 7%)

[4]The evaluation calculates the cost adjustment associated with changes in the feed rate that arise from ^{135}Xe poisoning, FP removal, etc., with a U feed enrichment of 5.0% even if the base U feed rate differs.

Table 6.9: Change in mflow ($\dot{m}$) per year as a function of k_{eff}.

molU	FF	$[UF_4]$ (%)	Enrichment (%)	$\frac{d(\dot{m})}{dk_{eff}}$ $[10^3 \cdot \frac{cm^3}{yr}]$
0.05	0.075	0.95	1.754	14.83
0.08	0.075	1.47	1.378	6.18
0.13	0.075	2.28	1.147	2.93
0.22	0.075	3.55	1.015	1.63
0.05	0.120	1.52	1.566	12.06
0.08	0.120	2.36	1.318	5.67
0.13	0.120	3.64	1.169	6.20
0.22	0.120	5.67	1.120	6.59

Table 6.10: Enrichment of Uranium in continuously added fuel salt.

Ini. Enrichment (%)	Added Enrichment (%)
3.5 < En	Too high, N/A
1.6 < En < 3.5	5.0
1.2 < En < 1.6	$3.0 + (5.0 - 3.0) \cdot (En - 1.2)/(1.6 - 1.2)$
En < 1.2	3.0

- Details: There exists a trade-off regarding the enrichment of added salt: low enrichment is less costly per unit volume and removes more FPs (through the shorter average salt loop residence time) but requires more salt addition and produces a greater volume of waste. Only fresh UF_4 is added and UF_4, noble gas and noble metal FPs are removed. The carrier salt and other FPs are not removed. The exchange of only UF_4 is to isolate the UF_4 effect, exchanging entire salt compositions would hide the UF_4 exchange rate and produce results valid for only one salt.
- Results: The feed-rate (mflow) of added UF_4 is shown in Figure 6.14 for several configurations and the UF_4 feed-rates over the entire depletion calculations are provided for one sample configuration in Figure 6.15. There appears to be insufficient benefit to enrichment beyond 5% to justify exceeding the maximum 5% enrichment of most current commercial reactors. Thus the enrichment is constrained to a maximum of 5%. As well, the added enrichment should be greater than or nearly double the initial enrichment. The added enrichment level was selected according to Table 6.10.

6.3.6 Depletion, on-line FP Removal:

- Objective: Evaluate the sensitivity of k_{eff} to the removal rate of multiple sets of similar removed FPs.
- Data points: Variation in the sets of FP's removed (2 sets), the removal rate (4 speeds) and molU (0.065, 0.13, 0.22)
- Details: Simulations used $NaF - BeF - UF_4$ salt, a pitch of 19 cm and a FF of 0.094. For the salt volume outside of the core, 12% of the core volume was selected. It was decided not to include the salt-outside-core percentage in the evaluation of processing as it has little impact. The removal rate - the inverse of cycle time - refers to the average rate at which an element/isotope is removed from the

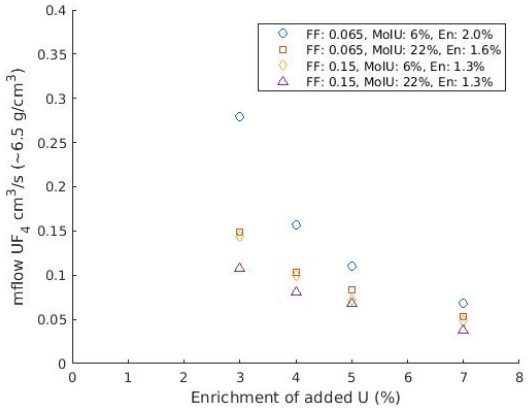

Figure 6.14: UF_4 feed-rate (mflow UF_4) after 1080 days of full-power operation for several configuration with $LiF - BeF_2$ carrier salt. The En % in the legend indicates the initial enrichment.

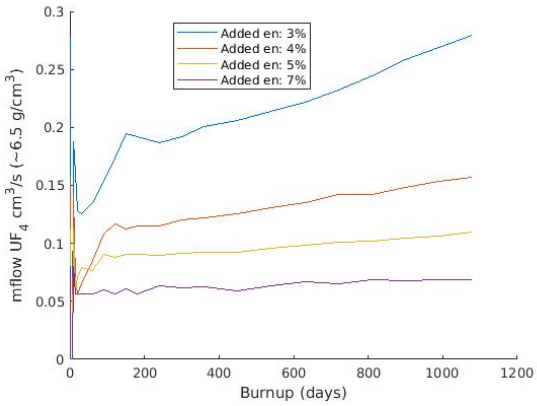

Figure 6.15: UF_4 feed-rate (mflow UF_4) over the course of a 1080 day depletion simulation of a $LiF - BeF_2$ salt, molU 6.5%, FF 0.06 configuration with an initial enrichment of 2.0%.

reactor, it is measured as a fraction of element/isotope removed per second (s^{-1}). The elements are grouped into two sets according to similar chemical characteristics such that they can be processed using the same technique(s). A 'very high fraction of isotopes' including ^{140}Ba, ^{141}Ce, and ^{91}Sr [35] that form stable soluble fluorides are infeasibly difficult to process and thus will not be examined. Noble gases are also not examined because the removal of ^{135}Xe, the only noble gas of consequence to reactor operation, is covered in a different analysis, see Section 6.2. The other two sets along with the removal rates are listed below. Additional discussion of fuel-salt processing generally can be found in Appendix D.3. The evaluation program has a reference processing removal rate of $3{\cdot}10^{-5}$ of set 1 and does not process set 2, processing in the evaluation program is discussed in Section 6.5.6.

1. Set 1: Noble metals: Mo, Ru, Rh, Tc, Te, Pd, Ag, Nb
   $1.60{\cdot}10^{-4}$, $8.02{\cdot}10^{-5}$, $4.01{\cdot}10^{-5}$, $1.925{\cdot}10^{-5}$
2. Set 2: Rare earth:[5] Sc, Y, La, Ce, Pr, Nd, Pm, Sm, Eu, Gd, Tb, Dy, Ho, Er, Tm, Yb, Lu
   $2.41{\cdot}10^{-6}$, $8.02{\cdot}10^{-7}$, $3.85{\cdot}10^{-7}$, $1.92{\cdot}10^{-7}$

- Results: see Figure 6.16.

6.3.7 Central Depletion Analysis

The central depletion analysis is a large set of simulations that provide key data to the evaluation program. The main objective is to find the salt flow rate in/out of the reactor over time for all configurations. Available computing resources only allow for a limited number of depletion simulations, especially when the cumulative time is high. The 6 salts were analysed at the molU fractions listed in Table 6.11 and at each of three

[5] Includes some non-rare earth elements with similar chemical characteristics

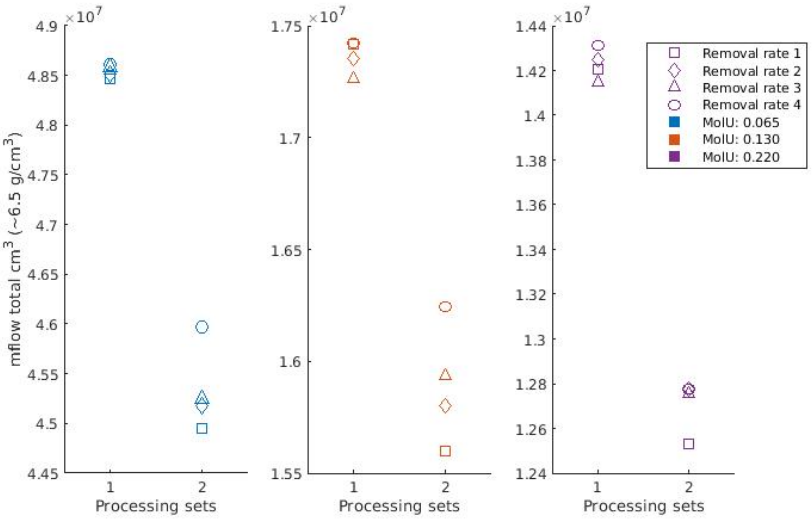

Figure 6.16: Total added fuel salt (mflow total) over the lifetime of the depletion simulation for different FP processing schemes. The depletion simulations run cover FP set 1 and set 2, the 4 corresponding removal rates and 3 mole % U values.

Table 6.11: Mole U fraction in central depletion calculation.

$LiF-BeF_2$	0.05	0.065	0.08	0.10	0.13	0.17	0.22	–
LiF	–	–	–	–	–	–	0.22	0.27
$LiF-NaF-BeF_2$	0.05	0.065	0.08	0.10	0.13	0.17	0.22	–
$NaF-BeF_2$	0.05	0.065	0.08	0.10	0.13	0.17	0.22	–
$NaF-ZrF_4$	0.05	0.065	0.08	0.10	0.13	0.17	0.22	–
$NaF-KF$	–	–	–	–	–	–	0.22	0.27

FFs, 0.06, 0.094 and 0.15. Some simulation parameters and common reactor design specifications are listed below:

- Salt volume outside the core: 12% of core volume
- Pitch: 19 cm
- Salt temperature 665 ^{o}C
- Monte Carlo simulation software: SERPENT 2
  - 100 active and 40 inactive cycles of 10000 neutrons

The flow rate of fresh salt into the loop and depleted salt leaving cannot be known a priori, so an initial guess is made and based on the results of previous depletion steps is updated along the way. The material flow rate (mflow or $\dot{m}$) step-to-step is adjusted as a function of both the deviation from the target k_{eff} and change in k_{eff}. An algorithm that merely considers the deviation from the target k_{eff} would often cause the k_{eff} to continue trending beyond the target value for several steps and thus result in oscillations. The k_{eff} deviation (from the target) and k_{eff} change are treated separately until the end and have unique relaxation factors. The depletion simulation process for each configuration is as follows:

1. Find the initial enrichment level of the configuration from the static runs (Section 6.3.1).
2. Calculate the total salt volume.
3. Determine the enrichment of the added salt as per Table 6.10.
4. Make the salt composition
   - If first step: Write fresh fuel
   - Else: Copy the fuel composition from the end of the previous run
5. If past the first few steps, update the relaxations: Use Equation 6.30 to calculate the relative undershoot/overshoot (S) of the previous mflow rate. Substitute the undershoot/overshoot result into Equation 6.31, to update the 'deviation' relaxation factor.

$$S = \frac{keff_t - keff_p}{keff_t - keff_{2p}} \tag{6.30}$$

   Where $keff_t$ is the target $keff$, $keff_{2p}$ is the $keff$ at the end of the burnup step from two steps previous and $keff_p$ is the $keff$ at the end of the previous burnup step. The minimum magnitude of the denominator is 0.015 to avoid division by zero.

$$R_d = R_d(1 + 0.25 \cdot S/{C_b}^{0.4}) \tag{6.31}$$

   Where R_d is the deviation relaxation, S is the relative under/overshoot and C_b is the burnup count. (C_b is equal to the larger of 10 or the current burnup step count)

   The 'slope' relaxation factor is calculated in a similar fashion:

$$R_s = R_s(1 + 0.25 \cdot (\Delta keff/\Delta keff_p)/{C_b}^{0.5}) \tag{6.32}$$

   Where $\Delta keff$ is the change in $keff$ over the most recent two runs and $\Delta keff_p$ is the $\Delta keff$ of the previous step. The ratio intends to capture the relative change in slope. The minimum magnitude of $\Delta keff_p$ is 0.01, again to avoid division by zero.
6. Calculate the mflow rate:
   - If first step: Make an initial guess of $\dot{m}$[6]
   - Else: Adjust the $\dot{m}$ using a combination of adjustments from the deviation from the target $keff$ and $\Delta keff$ (slope). The relative change in $\dot{m}$ from deviation ($\Delta\dot{m}_d$) is calculated according to Equation 6.33 and from slope ($\Delta\dot{m}_s$) according to Equation 6.34. The two $\Delta\dot{m}$ values are then added together to determine the new $\dot{m}$ in Equation 6.35.

$$\Delta\dot{m}_d = R_d \cdot 250(keff_t - keff_p)/t_c \tag{6.33}$$

[6]Initial guess: $\dot{m} = \frac{0.065 \cdot 0.05(V_c 0.06+1.7)}{molU[^{235}U_{in}]V_s}$, where V_c is the core volume, $[^{235}U_{in}]$ the enrichment of added U and V_s the salt loop volume.

$$\Delta\dot{m}_s = R_s \cdot 250 \cdot \Delta keff \cdot \frac{t_p/t_c}{\sqrt{t_p \cdot t_c}} \tag{6.34}$$

Where t_p and t_c are the length of the previous and current burnup steps (days) respectively. To avoid unreasonably large changes in $\dot{m}$ the maximum change from either adjustment is 0.35 (a 35% change).

$$\dot{m} = (1 + \Delta\dot{m}_d + \Delta\dot{m}_s) \tag{6.35}$$

7. Create a 1-step depletion input file, by combining the geometry, salt composition, mflow rate and other components of the calculation.
8. Repeat from 4 or until 2520 ($7 \cdot 360$) cumulative day steps have passed.[7]

Results

The full salt flow dataset generated is too large for all of it to be provided. Fortunately it was found that many significant nuclides approach equilibrium (for example see Figure 6.17) and the mflow stabilizes relatively early in the 7 year depletion analysis. As a result, the average mflow deviates little from the long-run equilibrium value.

Though the mflow data is noisy and volatile on occasion, it averages out effectively as evidenced by the continuous, logical patterns in the average mflow data (Figure 6.18). The noise in mflow has little effect on the average mflow rate because inaccurately low or high values must be balanced out in subsequent depletion steps. The first 15-20 depletion steps are especially noisy as the algorithm is trying to find the correct mflow rate and the correct mflow rate changing with burn-up. Sample plots of k_{eff} and the mflow rate for two configurations are shown in Figure 6.19. Additional mflow samples for several other configurations are provided in Appendix A.2. The algorithm tracked the target k_{eff} (1.02) quite well and the noisy initial depletion steps are of little utility, so plots are provided of all FFs tested for several salt and mole % U configurations in Figure 6.20.

6.3.8 Additional Details

- Core boundary: For each pitch and FF, each channel was made to be fully within the core boundary. The fuel channel diameter is calculated not a fraction of a graphite hex volume but by dividing (core radial area)(FF) by the number of channels within the reactor boundary. This way the intended FF is accurate and no fuel channels are cut-off by the reactor boundary.
- Graphite temperature: graphite temperature differences were found to have an inconsequential effect on k_{eff}.
- Adjustments in the composition of the fuel necessary to regulate the redox potential are ignored.

[7] Day steps: 0.5, 0.5, 1, 1, 2, 2, 3, 5, 7, 8, 10, 15, 15, 20, 20, 25, 25, 30, 40, 50, 60, 80, 90, 90, 90, 90, 90, 90, 90, 90, 90, 120, 120, 120, 120, 120, 180, 180, 180, 180

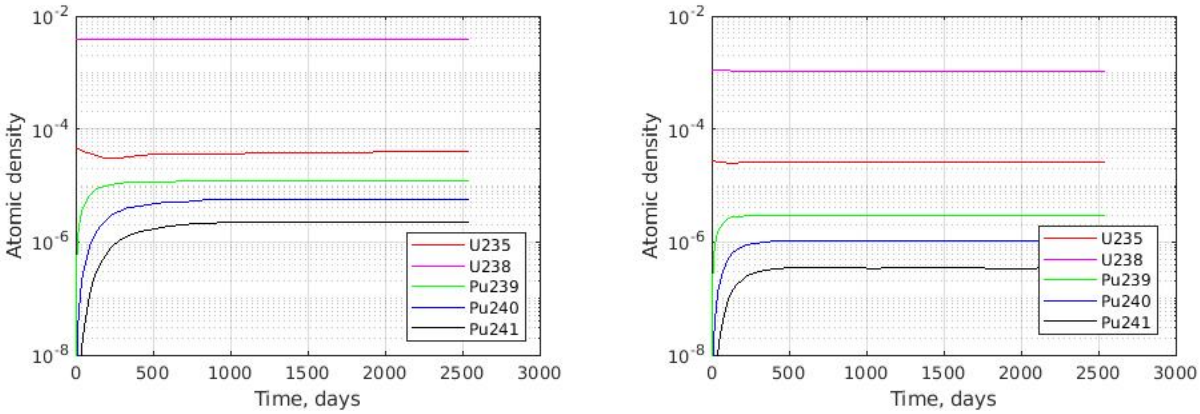

(a) $LiF - BeF_2 - UF_4$, mole U = 13%, FF = 0.094 (b) $NaF - ZrF_4 - UF_4$, mole U = 8%, FF = 0.094

Figure 6.17: Atomic density of several major nuclides over the course of the depletion simulation.

6.3.9 Temperature Reactivity Feedback Coefficients (TRFCs or TCs)

For the majority of an MSR's lifetime it would operate at a near-equilibrium level of burnup - ie. a level of burnup where the salt composition is effectively constant. The TCs are therefore calculated at a near-equilibrium level; the fuel salt composition used in all TC calculations is taken from 4 years into the corresponding depletion calculation. Two main effects combine to determine the TCs; Doppler broadening and density change. The Doppler broadening effect includes changes in temperature of the carrier salt components and FPs but the effect is almost exclusively from the change in fuel temperature.

The Doppler broadening effect on fuel salt TCs is always negative and reasonably consistent across reactor geometry and fuel salt composition, while the TCs due to density changes vary, and are largely a function of the level of over/under moderation - they can be either negative or positive. The TCs are calculated by adjusting the temperature of the salt or the temperature used to calculate density by both -200 and +200 ^oC.

The graphite TC also has a relatively small impact on neutronics, namely k_{eff} and due to the thermal mass and non-zero heat-transfer resistance, changes in temperature of the graphite occur with a significant time-lag. Therefore the graphite TC can be considered to have a negligible impact in the initial stages of unplanned reactivity insertion transients. As it relates to TCs, reactor behaviour during unplanned reactivity insertion transients is the primary performance concern. The graphite TC is nonetheless expected to have a (relatively minor) impact on both normal operation and accident scenarios but it is beyond scope and difficult to quantify the associated costs and risks.

Each composition/geometry at which a TC is calculated must correspond to a depletion run so that a near equilibrium salt composition can be extracted. The TC calculations are no more computationally expensive than the depletion runs thus the TCs are calculated for all of the depletion run configurations in the *Central depletion analysis*, Section 6.3.7. That is: all six salts, molU of 0.05, 0.065, 0.08, 0.10, 0.13, 0.17, 0.22,

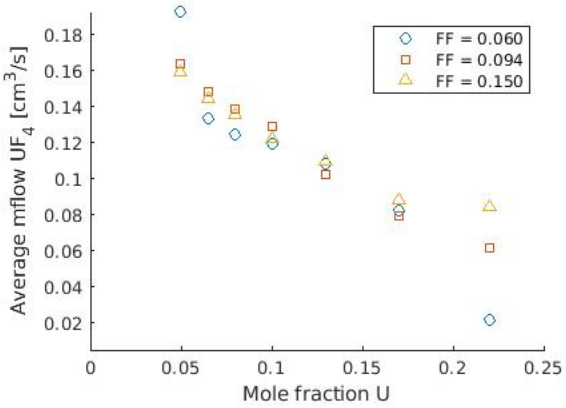

(a) $LiF - BeF_2 - UF_4$

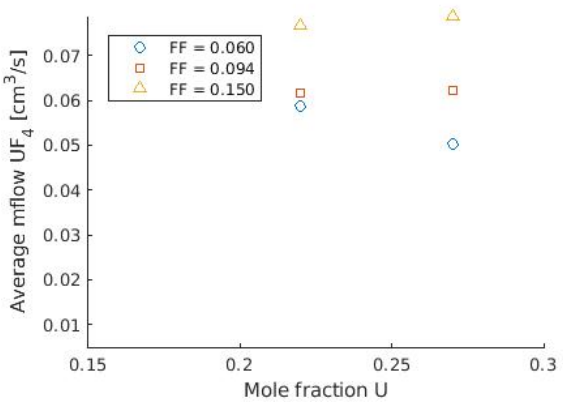

(b) $LiF - UF_4$

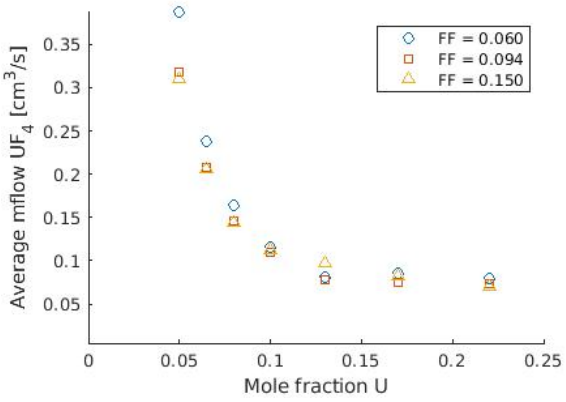

(c) $LiF - NaF - BeF_2 - UF_4$

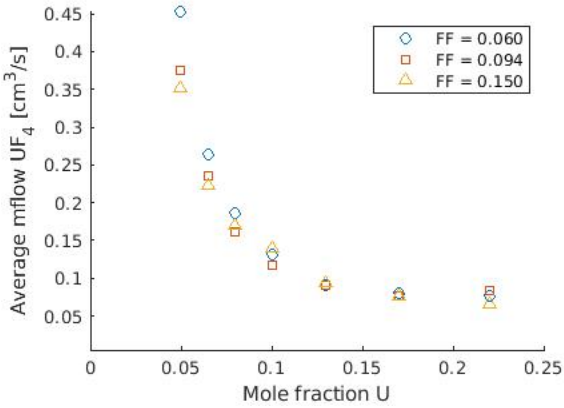

(d) $NaF - BeF_2 - UF_4$

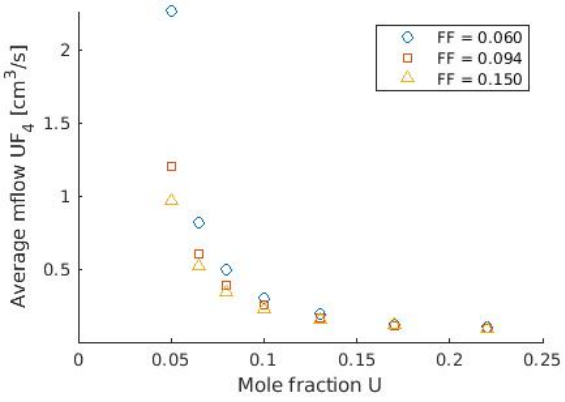

(e) $LiF - ZrF_4 - UF_4$

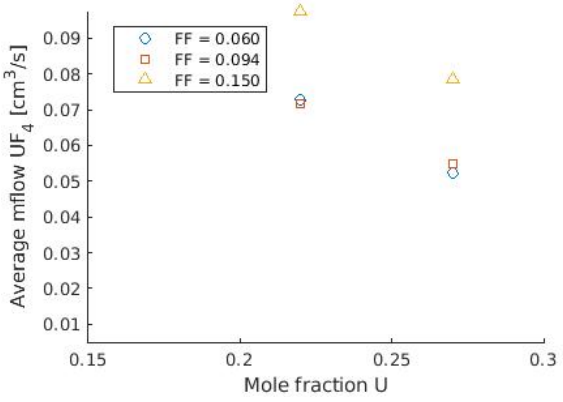

(f) $NaF - KF - UF_4$

Figure 6.18: Average rate of fuel salt addition (mflow) for the reactor configurations analysed over the 7 year depletion simulations.

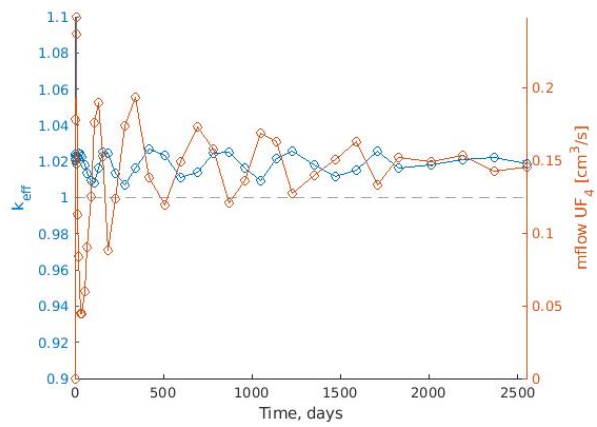

(a) $LiF - NaF - BeF_2 - UF_4$, mole U = 8%, FF = 0.094

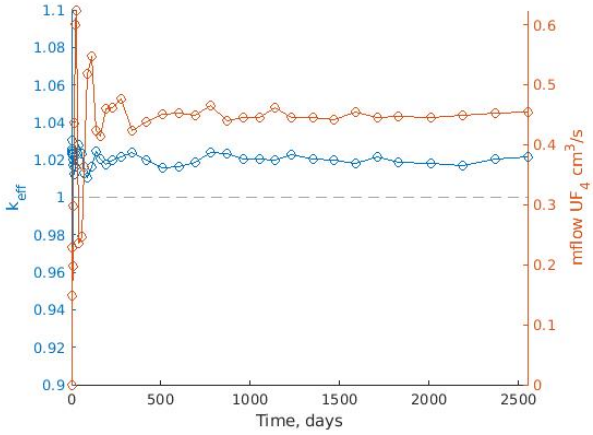

(b) $NaF - BeF_2 - UF_4$, mole U = 5%, FF = 0.094

Figure 6.19: Sample plots of k_{eff} and the rate of added fuel salt (mflow) over time.

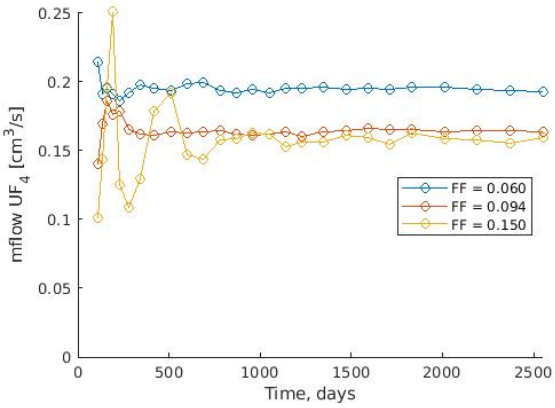

(a) $LiF - BeF_2 - UF_4$, mole U = 5%

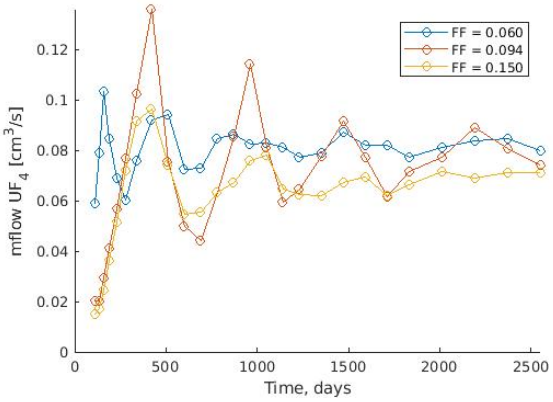

(b) $LiF - NaF - BeF_2 - UF_4$, mole U = 22%

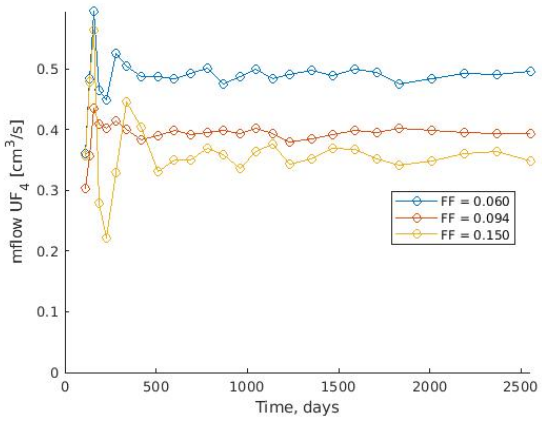

(c) $NaF - ZrF_4 - UF_4$, mole U = 8%

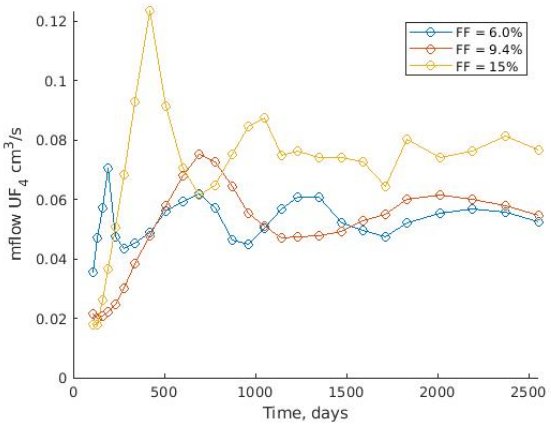

(d) $NaF - KF - UF_4$, mole U = 27%

Figure 6.20: Sample plots of the rate of added fuel salt (mflow) for each FF analysed.

0.27 and FF of 0.06, 0.094, 0.15.[8] All simulations are run with 5 million neutrons resulting in a k_{eff} error of $\approx \pm 0.00055$, a Δk_{eff} error of $\approx \pm 0.00078$ ($0.00055 \cdot \sqrt{2}$) and a TC error of ± 0.19 pcm/K.

TC results for the four salts feasible over a large molU range are shown in Figures 6.21. For all carrier salt compositions, increasing fuel density, either through the molU or FF, results in an increased TC through the density effect. This is in line with expectations, the TC is expected to be higher in under moderated conditions (high fuel density). The TCs of $LiF-BeF_2$, $LiF-NaF-BeF_2$ and $NaF-BeF_2$ were found to all be very similar. The density effect of $NaF-ZrF_4-UF_4$ salts is a stronger function of both molU and FF than in other salts. This is likely explained by the fact that ZrF_4 has a high molar volume relative to other carrier salt components and that ZrF_4 density has a relatively high temperature dependence ($1.06 \cdot 10^{-3} \frac{g}{cm^3}/^oC$). The Doppler effect has a fairly consistent impact of approximately -2 to -3 pcm/K across all salt compositions, molUs and FFs; though at low FFs it increases with higher molU but at high FFs it decreases.

The results of all salt compositions are compared in Figure 6.22, which shows results with a FF of 0.094 and in Figure 6.23, with mainly 13% mole U. Results of the same format, at the remaining FF (0.06 and 0.15) and two additional molU values (6.5 and 22.0%) are in Appendix A.2.

Impact of Fission Product (FP) concentrations on fuel salt temperature coefficients

The composition of FP's removed through on-line processing and the removal rate (decay constant, $[\lambda\ s^{-1}]$) effects the overall FP concentrations in the fuel salt. The FP concentration effect on fuel TCs was analysed. It is unclear what significant mechanisms may be responsible, for instance, a low FP removal rate requires a somewhat higher uranium enrichment to counteract parasitic FP absorptions. That higher enrichment may effect the fuel TC.

Due to computational expense, TC (both Doppler and density) calculations were initially limited to three reactor configurations: $NaF-BeF_2-UF_4$ salt, FF = 0.094 at molU = 6.5, 13.0 and 22.0%. The three configurations were run at each of the following processing scenarios, with element sets 1 & 2 and the respective removal rates equal to those in Section 6.3.6.

- No processing.
- Set 1: Processing of noble metals at four processing rates measured in decay constants: $1.60 \cdot 10^{-4}$, $8.02 \cdot 10^{-5}$, $4.01 \cdot 10^{-5}$, $1.925 \cdot 10^{-5}$
- Set 2: Processing of rare earth elements at four processing rates measured in decay constants: $2.41 \cdot 10^{-6}$, $8.02 \cdot 10^{-7}$, $3.85 \cdot 10^{-7}$, $1.92 \cdot 10^{-7}$

The impact of processing (ΔTC) is calculated according to Equation 6.36 and the results are shown in Figures 6.24 and 6.25.

$$\Delta TC_{pro} = TC_{pro} - TC_{no\,pro} \tag{6.36}$$

[8] If each configuration is feasible from a maximum enrichment and maximum melting point standpoint.

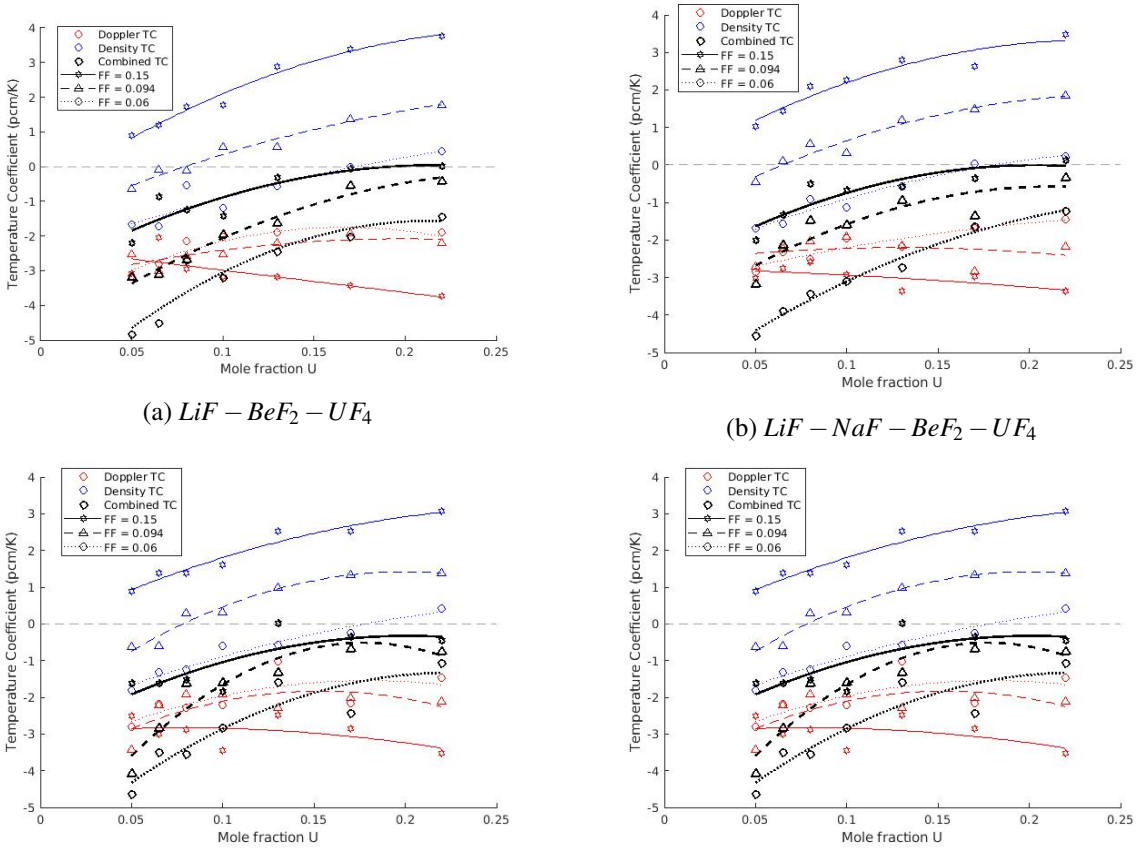

(a) $LiF - BeF_2 - UF_4$

(b) $LiF - NaF - BeF_2 - UF_4$

(c) $NaF - BeF_2 - UF_4$

(d) $NaF - ZrF_2 - UF_4$

Figure 6.21: Fuel Temperature Coefficients (TC) of salts with a large feasible mole % U range. The net (combined) TCs are shown alongside the component TCs of the Doppler (temperature) and density effects.[a]

[a] The interpolations are provided to highlight tends and add clarity, the shape of the curves do not have a physical meaning.

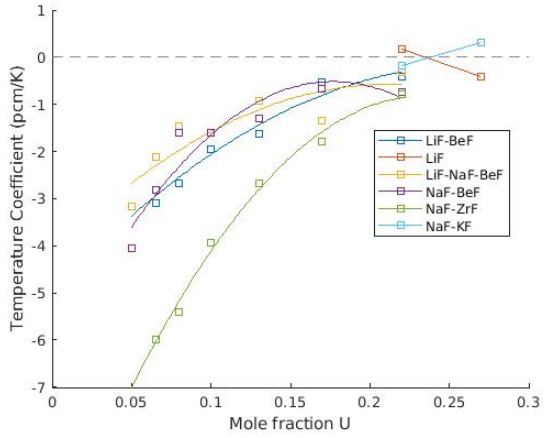

Figure 6.22: Temperature coefficients vs the evaluated molU fractions for all fuel salts at a FF of 0.094.[a]

[a]The interpolations are provided to highlight tends and add clarity, the shape of the curves do not have a physical meaning.

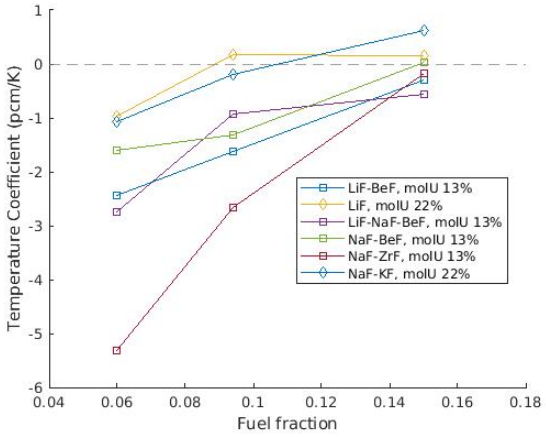

Figure 6.23: Temperature coefficients vs the evaluated FFs for all fuel salts. The mole % U of each salt is 13% if feasible and 22% otherwise (see legend).

Where TC_{pro} is the TC with processing and $TC_{no\,pro}$ is the TC without processing. The results of the ΔTC_{pro} calculations, as shown in Figures 6.24 and 6.25 appear to have too much noise to warrant further analysis or inclusion into the evaluation program.

6.4 Evaluation

6.4.1 Summary

This analysis takes generated neutronics and non-neutronics data as well as gathered salt property data of the graphite-moderated, circulated-fuel MSR and inputs it into an evaluation program. The program is built upon the cost estimation framework discussed in Section 2.1. The objective of the evaluation is to appropriately weigh and quantify the effect of individual design characteristics in order to arrive at a cost value for each design configuration under study. As many cost factors as reasonably possible are taken into account. All costs are measured in units of 2021 USD at the time of reactor start-up. Capital costs have an interest expense while both the value of electricity produced and O & M costs are discount to the time of reactor start-up.

The evaluation fundamentally aims to assess the thermal spectrum circulating-fuel MSR design concept as a function of the parameters discussed in Section 6.1.1:

- Carrier salt composition.
- Mole % U in salt.
- Fuel salt fraction.

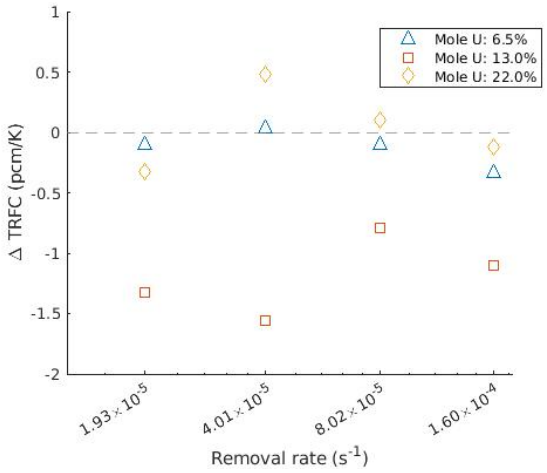

Figure 6.24: The difference in T(RF)C (combined density and Doppler effect) of the fuel salt with removal of noble metals (Set 1) and without removal. Carrier salt composition of $NaF - BeF_2$ and FF of 0.094.

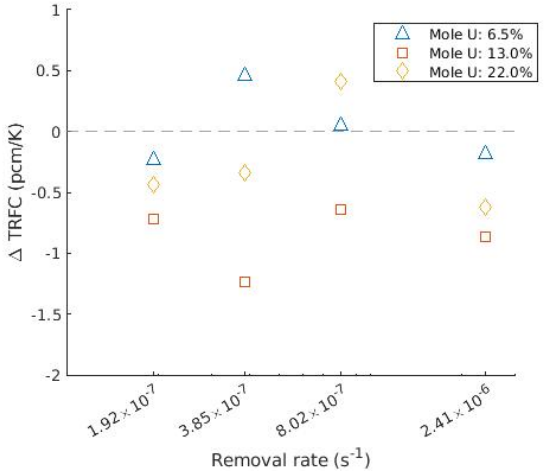

Figure 6.25: The difference in T(RF)C (combined density and Doppler effect) of the fuel salt with removal of rare earth elements (Set 2) and without removal. Carrier salt composition of $NaF - BeF_2$ and FF of 0.094.

The results obtained when using this framework are highly dependent on the assumed input values. Many of these input values are selected intuitively and require at least some degree of user discretion. Consequently uncertainty analysis is performed. The effect of each input variable on the total cost is initially unclear so sensitivity analysis is conducted as well. The uncertainty/sensitivity analysis assigns variables a probability distribution and the evaluation is run numerous times at pseudo-random points across each distribution (see Section 4.2).

6.4.2 Objective and Capabilities

The two fundamental objectives are to find low-cost reactor designs and to quantify the effect numerous input parameters have on the total cost. In doing so the relative viability of the salts examined and the optimal geometric configuration of the reactor using each of these salts can be determined. It is a useful tool for evaluating various salt compositions across many attributes relevant to reactor construction/operation. Not only are the sum total outputs useful but the program is built in such a way that it can be seen how much various salt attributes, construction conditions, component costs etc. contribute to both the total cost and the uncertainty of the total cost.

6.4.3 MSR Evaluation Specifics

There are a total of 71 input parameters - or as per sensitivity analysis, explanatory variables. The variables defining the distribution used are listed in Table 6.12 and 6.13. All of the input parameters have a configuration-dependent effect on cost. The costs are divided into six categories: direct cost, waste, safety, proliferation, modularity and feasibility, the cost categories as discussed further in Chapter 2. Each cost is

made up of a set of cost components, see Table 6.14. The values used in the analysis with fixed inputs are the expected/mean values (μ) from the uncertainty analysis. Many of the cost components are calculated according to the *cost setting* method (see Section 2.8) and reference parameters values are required to perform the calculations; these reference values are listed in Table 6.15.

A separate uncertainty/sensitivity analysis is also performed on the cost categories. All of the the input parameters are set at constant values while the cost categories are given a double trapezoid probability distribution between 0.7 and 1.3. The *Direct cost* category and it's corollary *Modularity* are not included because the uncertainty stems from the uncertainty in the individual cost components. Whereas with the other cost categories there exists uncertainty in converting a *issue* (ex. safety) into a cost value.

The graphite, in a graphite moderated MSR with an average power density of 12 MW/m^3 can only withstand neutron fluence from approximately 7 years of operation. The power plant will likely continue to operate beyond this point however. It is unclear in a commercial operation whether only the graphite, the entire core or entire core and salt loop, including heat exchanges would be replaced. Moreover it is unclear whether, following the core/graphite replacement, the reactor would restart with the same or with fresh fuel salt.

The costs involved are approximately evenly split between capital and continuous costs. The continuous, Operations and Maintenance (O & M) costs, are nearly constant per unit power delivered regardless of the number of graphite/core swaps are included the analysis. The per unit power capital costs however can differ significantly with the number of graphite/core swaps. Though, if the entire core and salt loop are swapped and all capital costs must be repeated every 7 years, it makes no difference and the relative cost of reactor configurations are the same regardless of the number of swaps included in the analysis. This evaluation makes exactly this assumption, because it simplifies the analysis and is likely to be implemented, at least in early commercial designs. This assumption is corroborated by the fact that the Terrestial Energy plans to implement full core replacement with their IMSR design [189].

6.5 Key Evaluation Program Optimizations/Calculations

6.5.1 Iteration

Many calculations in the evaluation program use data points generated in other sections. The program is ordered in such a way to minimize situations where input data has not been explicitly calculated in an earlier section, however it cannot be completely eliminated. One example is the ^{7}Li enrichment optimization and pump/heat exchanger size optimization: The optimal ^{7}Li enrichment is a function of the salt volume - and by extension the heat exchanger size, and the heat exchanger size is a function of the salt cost per unit volume - and by extension the ^{7}Li enrichment level. Thus it is necessary for this section of the program to iterate until convergence; no more than 3 iterations are necessary. The costs involved and the iteration methodology is discussed in Section 3.4.1.

Table 6.12: The defining distribution parameters of all input parameters using the Pearson probability distribution. Costs are in units of 2021 USD.

Input variable	Pearson probability distribution			
	mean (μ)	S.D. (σ)	skew	kurtosis
Construction time (yr)	5	2	1.7	8
Interest rate	0.05	0.015	1	6
Lifetime (yr)	7	0.5	1	6
MWh cost ($/MWh)	100	7	1	6
Pumping multiplier	1.5	0.25	1	6
Material cost ($/$m^2$)	3500	900	1.2	7.5
^{7}Li price ($/kg)	1500	450	1.5	8
Be price ($/kg)	500	100	1.5	8
Zr price ($/kg)	100	20	1.5	8
Separative work[a]	400	80	1.5	8
$U_3O_8^a$	600	120	1.5	8
Ref. OM safety ($/MWh)	3.03	1	1	6
Ref. OM prolif ($/MWh)	1.80	0.2	1	6
Be handling	15	4	1	7
Detritation OM ($/MWh)	1	0.4	1	7
Ref. tech. develop ($10^6)	1000	100	1.2	7.5
Share tech.	0.025	0.007	1.5	8
Rel. corrosion 1, 2, 3, 4, 5, 6	1	0.15	1	7
Non-configuration ($10^6)	20	2.5	1.5	7
Ref. cost MWth ($/MWth)	2500	650	1.2	7.5
Insurance premium ($10^6)	2	0.5	1.2	7.5
CDF	$1.0{\cdot}10^{-6}$	$0.25{\cdot}10^{-6}$	1.2	7
Min. enrich. for cost	3	0.7	1	7
Detritation C1	2	0.7	1.2	7
Detritation C2	18	3	1	7
Processing 1 direct cost ($10^6)	10	3	1.2	7
Enrich rel. cost at 7%	3	0.5	1	7
Ref. Pu rel. cost	4	1	1	7
T_{melt} rel. cost	6	1.5	1	7
P_{vap} rel. cost	3	1	1	7
Corrosion rel. cost	5	0.5	1	7
Xe removal rel. cost	1	0.3	1	7
Be rel. cost	4	1.5	1	7
^{6}Li effect rel. cost	15	3	1	7
PNL cost ($10^6)	1	0.2	1	7
Xe PF safety rel. cost	4	1	1	7
Temp. feedback rel. cost	10	2	1	7
β_{eff} rel. cost	5	1.2	1	7
Processing 1 rel. cost	5	1.5	1	7
V_g/V_c	0.01	0.002	1.2	1.7
Mass transfer coeff.	0.0005	0.0001	1.2	7
Time outside[b]	1	0.15	1	7

[a] $ to enrich 1 kg to 5% with 0.2% in tails.

[b] Time for delayed neutron precursors to travel from the core outlet to core inlet.

Table 6.13: The defining distribution parameters of all input parameters using the double trapezoid probability distribution.

Input variable	Double Trapezoid probability distribution mean	min	max
Capacity factor	0.87	0.72	0.98
Cost overrun safety	0.6	0.45	0.75
Share cap safety	0.2	0.05	0.4
Cost overrun prolif	0.1	0.05	0.15
Share cap prolif	0.05	0.02	0.10
Ref. share modularity	0.15	0.05	0.30
T_{melt} split	0.8	0.6	0.95
P_{vap} split	0.8	0.6	0.95
Corrosion split	0.3	0.15	0.50
Xe removal split	0.5	0.2	0.8
Processing 1 split	0.5	0.15	0.75
Processing 2 split	0.5	0.15	0.75
Detritation C3	0.75	0.55	0.95
Non-config. exp.	0.7	0.5	0.9
Rel. safety cost exp.	0.5	0.3	0.7
Rel. prolif cost exp.	0.5	0.3	0.7
Regular cost modularity exp.	0.5	0.3	0.7
Overrun exp.	0.5	0.4	0.6
Feasib each exp.	0.35	0.25	0.45
Processing exp.	0.5	0.3	0.7
Fresh contribution[a]	0.3	0.1	0.6
Flux to Xe PF exp.	0.6	0.4	0.8

[a]Share of the PN generation time cost from the fresh fuel calculation

Table 6.14: Cost components associated with each cost category.

Cost	Cost component
Direct cost	Pumping power in heat exchanger
	Heat exchanger structure
	Pumping power in core
	^{7}Li in core initially
	Be in core initially
	Zr in core initially
	U in core initially[a]
	Salt initially in heat exchanger
	Xe removal
	Δ$ due to GL
	Processing set 1
	Processing set 2
	^{7}Li from refuelling
	Be from refuelling
	Zr from refuelling
	U from refuelling
	Detritation capital cost
	Detritation O & M
	Corrosion
	Non-configuration
	Interest on capital cost
Waste	Deep geological repository
	Be toxicity
Safety	Vapour pressure
	Melting temperature
	^{6}Li/Tritium
	Be toxicity
	PN generation time
	Xe safety
	β_{eff}
	Doppler coefficient
	Density coefficient
	Insurance
Proliferation	^{6}Li/Tritium
	Uranium enrichment
	Plutonium production
Modularity	Modularity
Feasibility	Vapour pressure
	Melting temperature
	Xe removal
	Processing set 1
	Processing set 2
	Corrosion

[a] Includes the cost of U_3O_8 and separative work.

Table 6.15: Reference values used in cost setting. A description of cost setting is provided in Section 2.1.

Parameter	Value	Associated input variable
T_{melt} ref.	550 ^oC	T_{melt} cost
P_{vap} ref.	10 mmHg	P_{vap} cost
Corrosion ref.	1	Corrosion cost
Xe PF	0.05	Xe PF safety cost
Xe solubility	1^a	Xe removal cost
Xe stripping time	100s	Xe removal cost
[Be]	0.015 $10^{24}/cm^3$	Be cost
[Li]	0.025 $10^{24}/cm^3$	^{6}Li effect cost
^7Li enrichment	0.99995	^{6}Li effect cost
PNL	0.001 s	PNL cost
TC	-2 pcm/K	Temp. feedback cost
β_{eff}/β	0.25	β_{eff} cost
Processing set 1 solubilities	4^a	Processing 1 cost
Processing set 1 removal rate	$3 \cdot 10^{-5}$	Processing 1 cost
Processing set 2 solubilities	2^a	Processing 2 cost
Processing set 2 removal rate	$3 \cdot 10^{-7}$	Processing 2 cost

[a] Solubilities are in measured in $K_p \cdot 10^3$ $(mol/(cm^3 \cdot atm))$ at 700^{o}C

6.5.2 Pitch Optimization

The channel pitch is program optimized (not an arbitrary input). The impact channel pitch has on reactor performance parameters generally cannot be considered in isolation as it strongly depends on the fuel fraction and salt composition. There are four considerations in the optimization:

1. Graphite lifetime (GL)
2. Effect on k_{eff}
3. Core pumping power required
4. Xenon poisoning

Graphite lifetime (GL): The reactor lifetime is limited by the dimensional changes that take place in the graphite moderator. The dimensional changes are caused by fast neutrons, the effect is cumulative and is a function of temperature [181].[9] A higher pitch has two effects which can impact the GL: higher fast flux (FF) gradients and higher average graphite temperature. The average graphite temperature is higher because the graphite surface temperature is (assumed) independent of pitch while the temperature difference between the graphite surface and interior increases with pitch.

[9] It was also found that 'the fast flux decreases about 5% from the surface of the graphite to it's interior ... ', which originally was expected to impact GL. Instead the induced stress was calculated to be much less than the tensile strength of graphite thereby leading to the conclusion that volumetric changes were permissible from viewpoint of stresses and that other factors (total volume, shape, porosity, etc.) limit graphite exposure.

The interior graphite temperature is estimated using an equation derived from Molten Salt Breeder Reactor experimental results [181], Equation 6.37. The temperature difference between the graphite surface and graphite interior of the highest power channel the MSBR is 65^oC. This design had a power density of 40 MW/m^3 and a pitch of 12 cm, most current designs and analysed design configurations throughout this paper use a power density of approximately 12.5 MW/m^3. It is assumed that the average temperature of the graphite is equal to average of the surface and interior temperatures.

$$T_{interior} - T_{salt} = \frac{65^o\,C}{(40MW/m^3)(12cm)} \cdot S(MW/m^3) \cdot P(cm) \tag{6.37}$$

Where S is the power density and P is the pitch. Simplified to include a single constant:

$$T_{interior} - T_{salt} = 0.1354 \cdot S \cdot P \tag{6.38}$$

At an operating temperature near 700 ^{o}C, $\frac{dGL}{dT} = 0.00235$, as derived from results in *P. R. Kasten et al. 1969* [181]. The change in GL is calculated relative to the peak Fast Neutron Flux (FNF), average FNF and temperature difference of an arbitrary reference configuration of mole U of 10%, fuel salt fraction of 0.1 and pitch of 19 cm. $\frac{\Delta L}{L}$ is the relative change in lifetime. Since the induced stresses will likely counteract the volumetric changes to some degree a term is added, '$factor$' of 0.7, to dampen the change in lifetime that FNF calculations would suggest.

$$\frac{\Delta L}{L} = \frac{\frac{FNF_{peak,ref}}{FNF_{peak}} + \frac{FNF_{avg,ref}}{FNF_{avg}}}{2} \cdot \frac{dT_{peak}}{dT_{peak,ref}} \cdot \frac{dGL}{dT} \cdot factor \tag{6.39}$$

Effect on k_{eff}: SERPENT simulations were carried out over a range of pitchs, for each of a range of salt/molU/FF configurations and the analogue k_{eff} extracted and compared to a reference value. A reduction in k_{eff} increases cost through higher initial enrichment and a higher average refuelling rate - the calculation process is covered in Section 6.5.3.

Core pressure drop (PD): The PD is calculated for each salt with 0% U and 22% U and the PD is interpolated for values in between. If the pitch/molU/FF do not align precisely with directly evaluated data points each must be interpolated - thereby resulting in a triple interpolation.

Xenon poisoning: The methodology for calculating the xenon poisoning of a given configuration is covered in Section 6.2.1. For each pitch of each salt/molU/FF configuration the Xe poisoning is calculated for a range of half lives: 1, 2, 3, 5, ... 1000, 2000, 5000 seconds. Similar to ^{7}Li enrichment optimization (Section 3.4.1), xenon poisoning optimization must balance:

1. The effect on reactor performance and safety. Xe poisoning lowers k_{eff} increasing cost through higher initial enrichment and a higher average refuelling rate (see Section 6.5.3). The safety cost portion of the input/reference xenon poison cost is either estimated (first iteration) or calculated from saved data (subsequent iterations). The value is then multiplied by the PF of the particular iteration and divided by the reference PF of the safety calculations.

2. The cost of xenon stripping/sparging; both the direct cost and the increased feasibility cost. Using the 'cost setting' method outlined in Section 2.1 the *Direct* and *Feasibility* cost are determined for each xenon stripping half life. The direct/feasibility cost relative to the reference cost is:

$$Cost_{rel} = \frac{T_{in}}{T_i} \cdot \frac{c_i}{c_{in}} \tag{6.40}$$

Where T_{in} is the reference half life and T_i the half life of the particular iteration. c_{in} and c_i represent the reference and particular iteration solubility respectively.

For each iteration of a xenon stripping rate the cost of poisoning and removal are summed to obtain the total cost associated with xenon. For each pitch, the xenon stripping rate that results in the lowest cost is determined, with both the stripping rate and associated cost saved.

Final step The sum of each of the four pitch optimization considerations is generated for each of the pitchs in the range. The pitch with the lowest cost is then selected and used to calculate GL, the k_{eff} adjustment, core PD and the xenon stripping/sparging rate.

6.5.3 Cost of Reduced k_{eff}

A reduction in k_{eff} increases cost through higher initial enrichment and a higher average refuelling rate. The cost of an increased initial enrichment is calculated according to the following process:

$$\Delta en = \frac{d(en)}{dk_{eff}} \cdot \Delta k_{eff} \tag{6.41}$$

Where en is the U enrichment level. Using Δen the difference in the cost of U/cm^3 (ΔU_{cc}) is then calculated. From that the overall cost difference:

$$\Delta\$ = \Delta U_{cc} \cdot V_{salt} \tag{6.42}$$

And the cost of an increased refuelling rate due is calculated according to the following equation:

$$\Delta cost = \frac{d\dot{m}}{dk_{eff}} \cdot (A \rightarrow P) \cdot \Delta k_{eff} \tag{6.43}$$

Where $\frac{d\dot{m}[kg/yr]}{dk_{eff}}$ is a measure of the change in the rate of added fuel salt as a function of k_{eff} and $(A \rightarrow P)$ is the annual cost of each year of operation discounted to the present. $\frac{d\dot{m}}{dk_{eff}}$ is a function of geometry and salt composition.

6.5.4 Cost of β_{eff} in Evaluation

Since the proportion of fissions from ^{235}U and other nuclides is not tracked and is not expected to differ significantly among configurations only the β_{eff}/β, the correction factor, is considered in the evaluation. β_{eff}/β is a function of circulation time (T) and fraction of salt in the core. T is equal to the summation of the time in the heat exchanger, the core and the piping between the core and heat exchanger (input parameter

time between). The β_{eff}/β related cost is determined using the 'cost setting' method, with the cost effect determined by the relation $1/(\beta_{eff}/\beta)$.

6.5.5 Cost of Corrosion in Evaluation

The cost of corrosion inflicted by the fuel salts (and corrosion mitigation) in the evaluation program is a function of the relative expected corrosion depth. The corrosion depth is relative to uranium-free FLiBe salt. The method used to quantify the relative magnitude of corrosion is a function of the Gibbs energy differences and experimental results. This method is explained in Section 2.1.1.

6.5.6 Cost of Processing

The evaluation is set up such that the Fission Products (FP) processing contributes to both the direct and feasibility cost categories. If the costs and effects of FP processing were accurately and precisely known, the chemical groupings removed and the removal rate should be determined through optimization. FP processing metrics, especially costs, are very difficult to predict and the chemical grouping(s) (set(s)) and removal rates were simply selected; only set 1 isotopes (in addition to Xe) with a removal rate equal to the reference value undergo FP processing, see Table 6.15. Despite set 2 not included in the evaluation, suggested values are included in the input parameter probability distributions in Tables 6.12 and 6.13, the reference removal rate in Table 6.15. A cost of zero is shown in the results.

The relative cost of FP processing of the various salts was made a function of solubility; it was assumed that a higher solubility would make FP extraction more difficult and thus more expensive. Insufficient data was found to estimate the solubilities of sets 1 and 2 in each of the salt compositions directly, so instead they were estimated by the known solubility of other elements/chemical groups. The relative solubility (S) of each salt was calculated as the geometric mean of the Henry's constant/solubility of the average of noble gases, Xe and Rare-Earth & TRansUranic (RETRU) elements, as per Equation 6.44. The Henry's constant/solubility values used in the evaluation are shown in Table A.2 of Appendix A.2 and a precise processing definition and chemical grouping discussion is provided in Appendix D.3.

$$S = (S_N S_{Xe} S_{RETRU})^{1/3} \tag{6.44}$$

6.6 Results

6.6.1 Results: with Input Variables at Mean/Expected Values

Figures 6.26 to 6.32 show results derived from fixing all the input parameters at mean/expected values. An advantage to analysing fixed input data is that the trends are more clear than with uncertainty analysis. As well the lower computing requirements allow for more configurations to be analysed. The selected input parameter values are the expected mean values; the expected mean from the expected probability distributions used in uncertainty/sensitivity analysis. The data in Figures 6.26 to 6.29 are arranged to clearly show

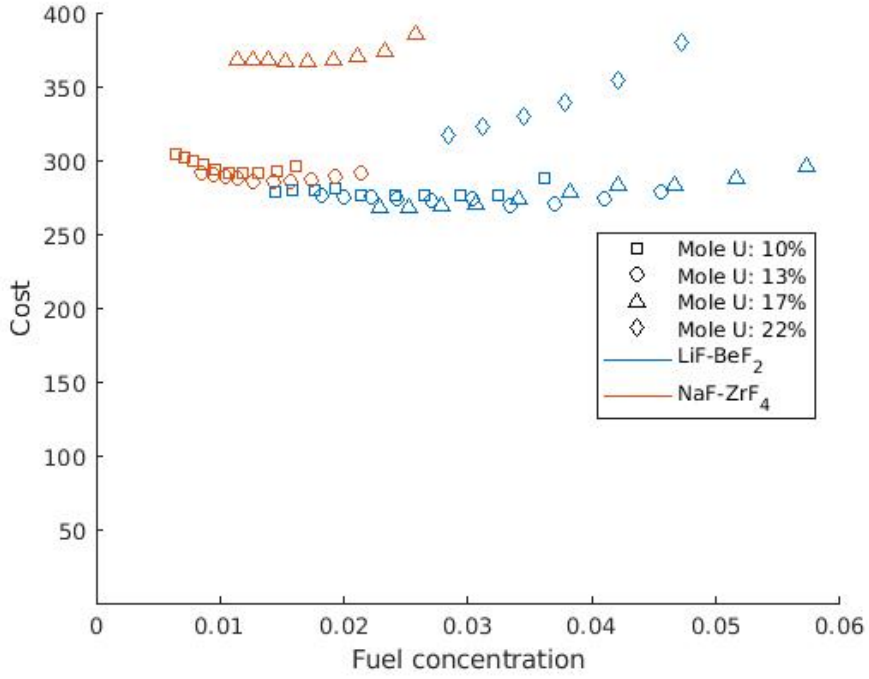

Figure 6.26: Total cost of all $LiF - BeF_2$ and $NaF - ZrF_4$ configurations evaluated using set design parameter values.

the relationship between cost and fuel concentration. The x-axis fuel concentration scale is *fraction of core volume*. The data presented covers most of the analysed design space; configurations with a low (< 10%) mole % U are left out of Figures 6.26 and 6.27 because they consistently result in a higher cost than at least one configuration with a different mole % U. Presenting all data of all 6 salts in one figure would be too cluttered, Figure 6.29 is a compromise that aims to still show key trends in one figure. Data points for some salts at 22% mole U do not appear in the respective figures because the salt has a(n) (expected) melting point that is infeasibly high.

At the configurations that result in the lowest total cost, the cost is broken down into the six components, see Figure 6.30. The contributing input parameters to the 'Direct' cost for three salts is shown in Figure 6.31 and for the other three salts the 'Safety' cost category is shown. The breakdown for the other cost categories for all salts is provided in Appendix A.2.

6.6.2 Results: Uncertainty Analysis

The uncertainty/sensitivity analysis is conducted for all 6 salts. The uncertainty/sensitivity analysis was originally conducted with a sample size of 500 for all feasible configurations across the the following design parameters values: (a total of 30 potential configurations per salt)

- Mole % U: 8.0, 10.0, 13.0, 0.17, 0.22, 0.27
- FF: 0.060, 0.075, 0.094, 0.120, 0.150

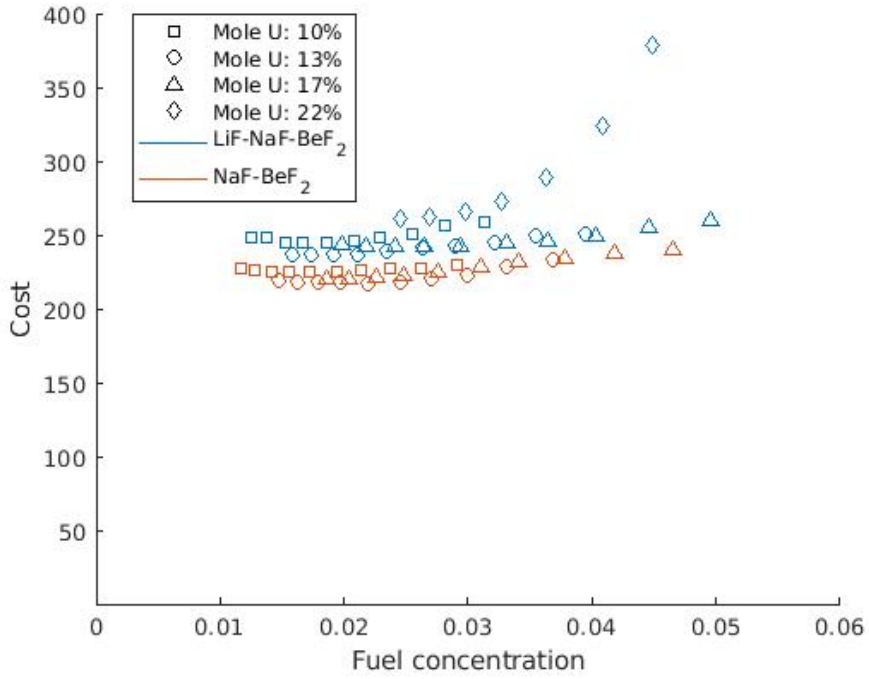

Figure 6.27: Total cost of all configurations evaluated using set design parameter values.

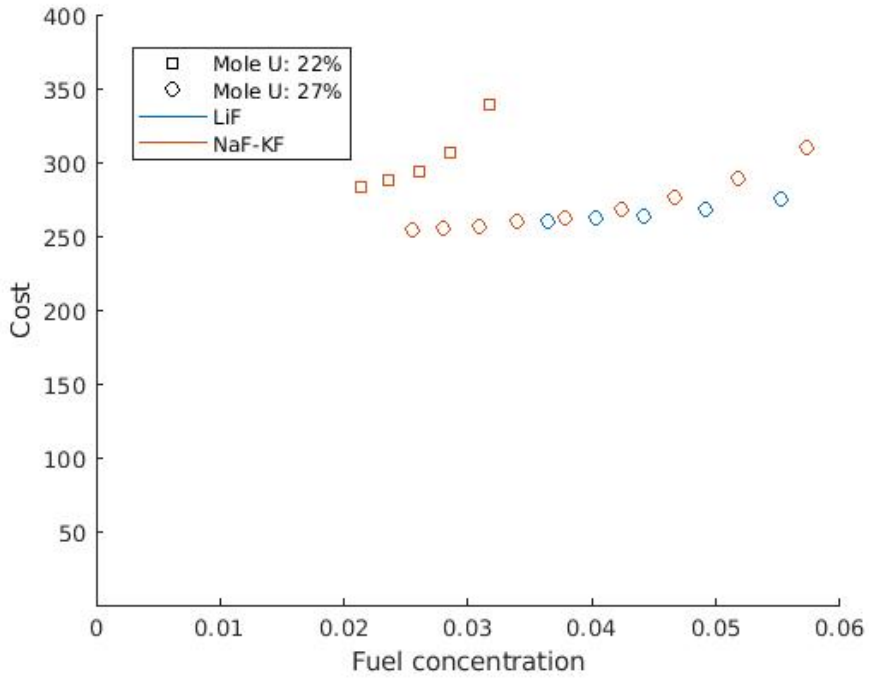

Figure 6.28: Total cost of all LiF and $NaF-KF$ configurations evaluated using set design parameter values.

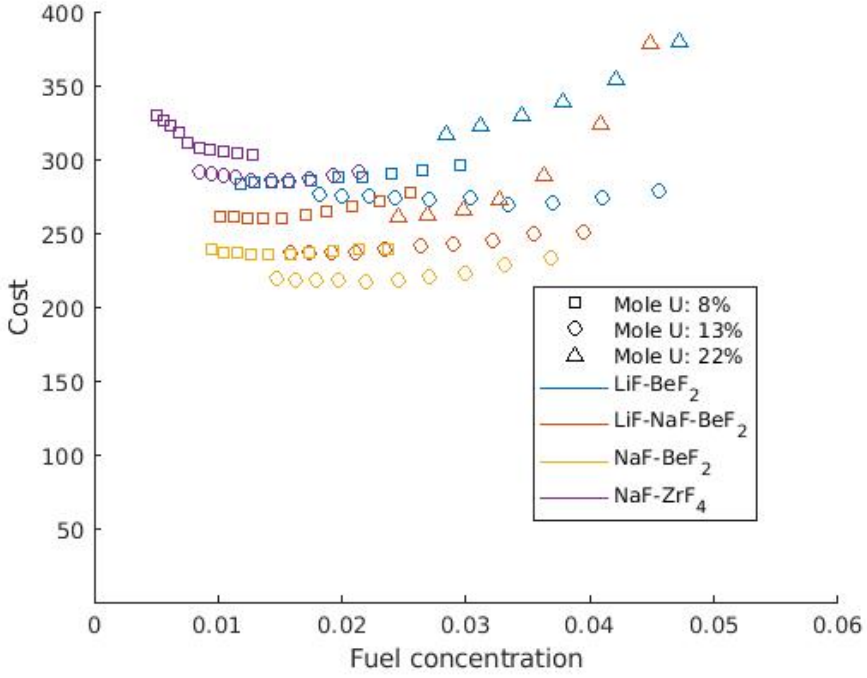

Figure 6.29: Total cost of $LiF - BeF_2$, $LiF - NaF - BeF_2$, $NaF - BeF_4$ and $NaF - ZrF_4$ configurations evaluated at three mole % U values using set design parameter values.

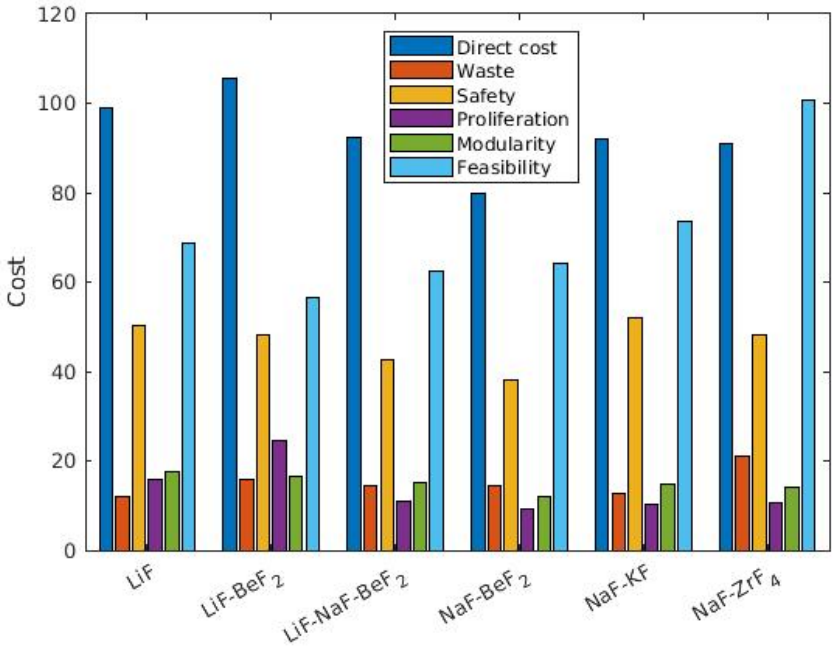

Figure 6.30: The cost of each of the six categories is shown for each salt. The cost values are given at the mole % U and FF combination that results in the lowest feasible total cost.

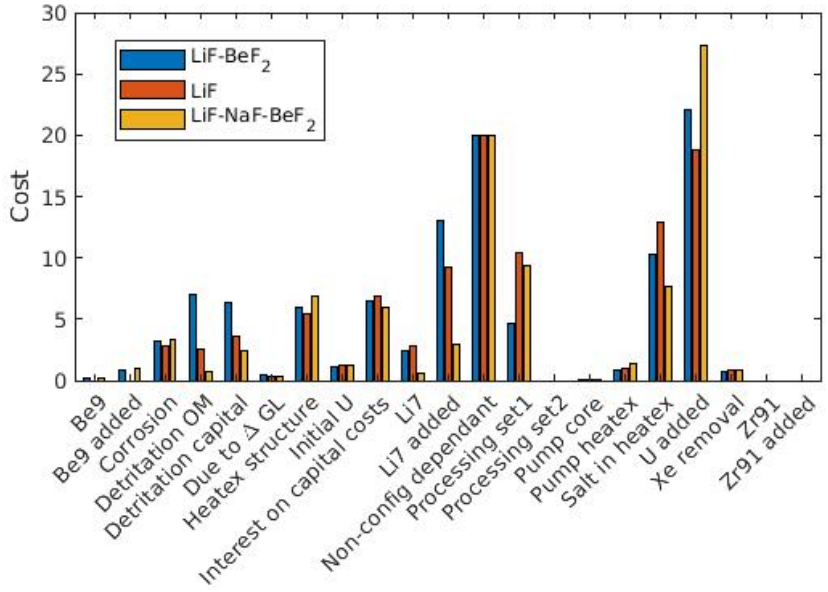

Figure 6.31: Breakdown of the 'Direct cost' category of cost into components for three salts. The cost values are given at the mole % U and FF combination that results in the lowest feasible total cost.

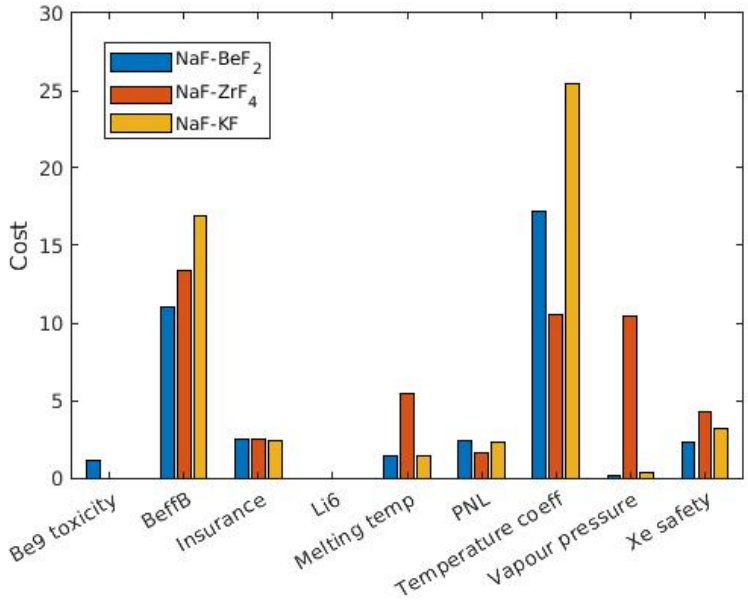

Figure 6.32: Breakdown of the 'Safety cost' category of cost into components for three salts. The cost values are given at the mole % U and FF combination that results in the lowest feasible total cost.

Table 6.16: Lowest cost configuration for each salt in the uncertainty analysis.

Salt	Mole % U	FF
(1) $LiF-BeF_2-UF_4$	13	0.094
(2) $LiF-UF_4$	27	0.075
(3) $LiF-NaF-BeF_2-UF_4$	13	0.075
(4) $NaF-BeF_2-UF_4$	13	0.094
(5) $NaF-ZrF_4-UF_4$	13	0.094
(6) $NaF-KF-UF_4$	27	0.060

A sample size of 500 is sufficient to calculate the mean and standard deviation (SD). With cost SDs of approximately 45, the deviations from the means are only ≈1.8 ($40/\sqrt{500}$) - or only slightly higher due to non-normality. Thus these results are used to produce Figures 6.35 and 6.37. The cost distribution cannot be assumed normal in the right side tail so a large number of samples is required to determine important percentiles of the distribution, such 95 and 97.5. All other uncertainty (and sensitivity) analysis results shown use 20 repetitions of 100 sample bins for a total of 2000 samples. Using 2000 samples and making no assumptions about the distribution, it was found that the SD of the 95^{th} percentile for the lowest cost configuration is between 2.4 and 5.2 (0.8% and 1.4%[10]) depending on the salt composition and the deviation from the expected 97.5 percentile is between 3.3 and 8.9 (0.9% and 2.3%). This level of precision was deemed sufficient considering the substantial uncertainty of many input parameters.

It become clear that much of the range includes many sub-optimal configurations so 2000 samples were applied only to a smaller range that includes the optimal configurations: (a total of 12 potential configurations per salt)

- Mole % U: 13.0, 0.17, 0.22, 0.27
- FF: 0.060, 0.075, 0.094

The distributed variables used in the uncertainty analysis have means equal to the respective fixed input values. Figure 6.33 shows the uncertainty analysis results for the lowest cost configuration of each salt. Vertical lines representing the average cost of the other salts are included for comparison. The lowest-cost configuration of each salt includes 1-4 clearly outlying data point(s). These 'outliers' are only included in Figure 6.34, all other results do not display these points. All points, including the outliers, are however included in calculating averages, variances, sensitivities, etc. If results are only presented for 1 configuration or 1 configuration per salt, the hresults are for the configuration that results in the lowest expected total cost. The lowest cost configuration of each salt is listed in Table 6.16.

Double histogram plots of two of the salt compositions at the lowest cost mole % U and FF configuration are shown in Figure 6.34. Figure 6.34 (a) compares the two salt compositions with the lowest cost configurations. Figure 6.34 (b) compares the two salt compositions that are only feasible at high mole % U (>≈ 20% U). The lowest-cost configurations of each salt have similar, variance and skew.

[10] $100*\frac{deviation}{95^{th}\text{ percentile cost}}$

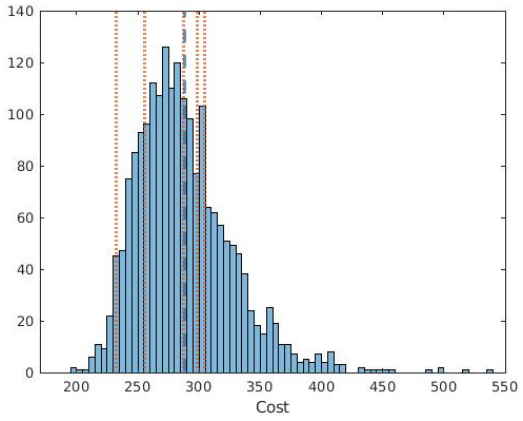

(a) $LiF - BeF_2$, mole U: 22%, FF: 0.060.

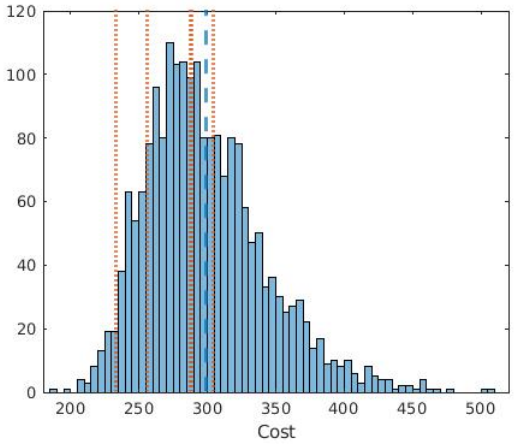

(b) LiF, mole U: 27%, FF: 0.075.

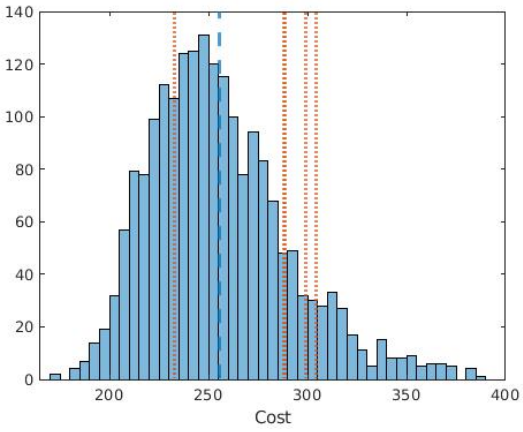

(c) $LiF - NaF - BeF_2$, mole U: 13%, FF: 0.075.

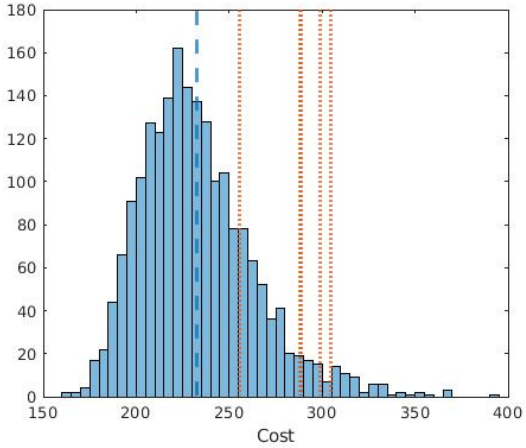

(d) $NaF - BeF_2$, mole U: 17%, FF: 0.060.

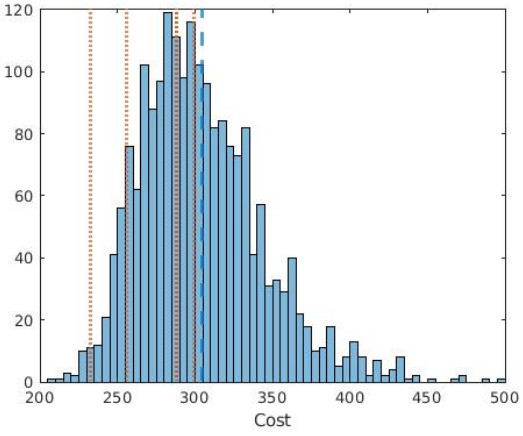

(e) $NaF - ZrF_4$, mole U: 13%, FF: 0.094.

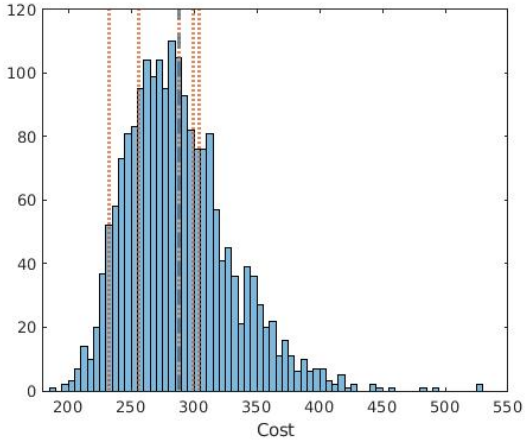

(f) $NaF - KF$, mole U: 27%, FF: 0.060.

Figure 6.33: Histogram of uncertainty analysis results of the geometric configurations with the lowest cost. The corresponding mole % U and FF is noted.

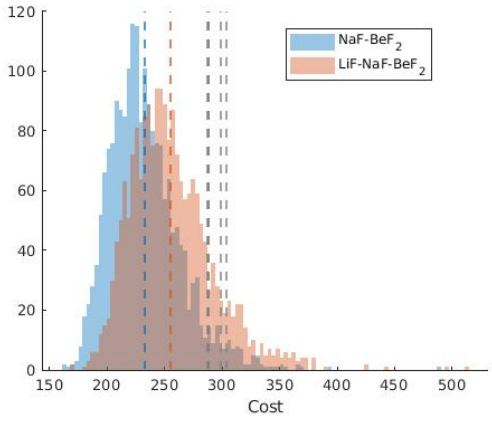

(a) $LiF - NaF - BeF_2$ and $NaF - BeF_2$

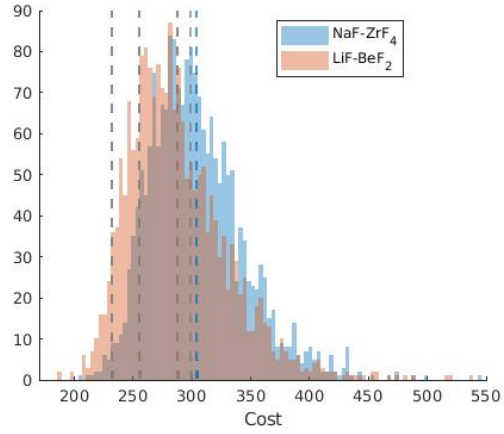

(b) $NaF - ZrF_4$ and $LiF - BeF_2$

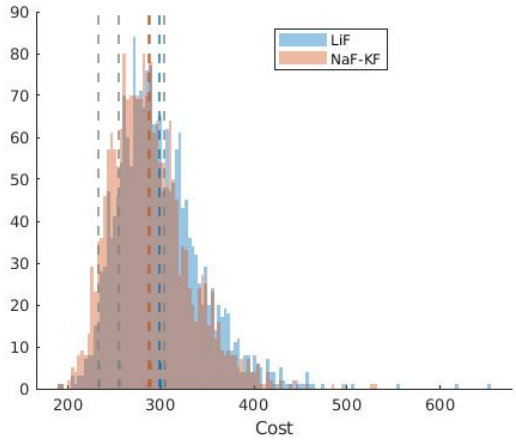

(c) LiF and $NaF - KF$

Figure 6.34: Overlapping histogram plots of the total cost. Data presented is at the lowest cost configuration.

Results presented so far have shown the overall uncertainty. A likely more important presentation of data than the distributions of the total costs are the distributions of the relative cost. If the input parameters are treated as independent, as in Figure 6.34, there is often significant overlap in the cost distributions; overlap proportional to the probability that the salt with the higher expected cost has a lower cost than the other salt. However, in situations where the combination of input parameters results in high costs (high CDF in the histogram of total cost) for one configuration, it tends to result in high costs for all configurations. Figure 6.36 shows the difference in cost between salts at their lowest cost configurations. The same configurations from Figures 6.36 (a) and (b) are shown in Figures 6.34 (a) and (b), take note of the much lower probability of the salt with the higher expected cost (ex. (a) $LiF - NaF - BeF_2$, (b) LiF) having a lower cost. This perspective suggests that, for the input parameters values used, $NaF - BeF_2$ results in the lowest total cost under almost all scenarios.

Correlation:

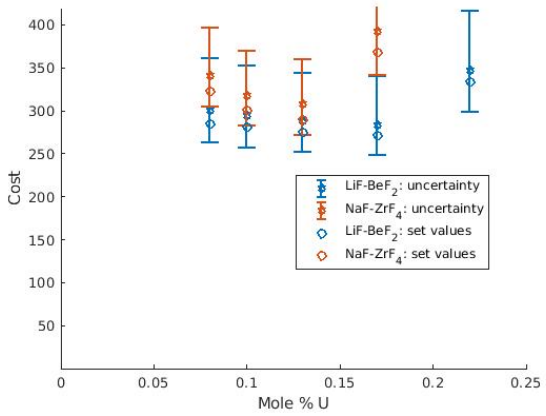

(a) $LiF - BeF_2$ and $NaF - ZrF_4$

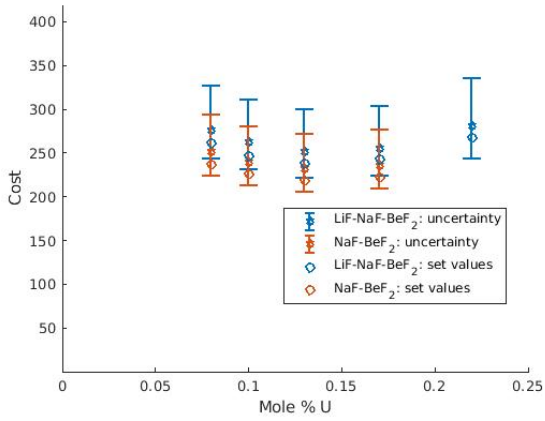

(b) $NaF - ZrF_4$ and $NaF - BeF_2$

Figure 6.35: Comparison of uncertainty analysis and fixed values. Results are shown at a FF of 0.075 at all feasible mole % U values. Uncertainty analysis is shown with an error range of one standard deviation.

Some input parameter values are correlated as in the FHR analysis. The same maximum correlation coefficient of 0.30, with relative correlations of either 0.5, 0.7 or 1.0, was assigned. The alike correlation methodology and parameters yielded similar results, again it is difficult to draw many useful conclusions. It was found that the level of correlation does not change the optimal configuration, it mainly has the effect of mildly increasing total cost variance. Once again, comprehensive documentation is not deemed worthwhile.

Uncertainty at the cost category level:

Uncertainty analysis is also performed at the cost category level; with categories direct cost, waste, safety, proliferation, modularity and feasibility. This must be a separate analysis because the addition of cost category uncertainty on top of individual input parameter uncertainty would account for the same underlying uncertainty twice. The cost uncertainty provides less insights than the comprehensive input parameter uncertainty so only limited results are shown: Figures 6.38, 6.39 and 6.40 as well as Table 6.17. Both the input parameter and cost category uncertainty cost differences, Figures 6.36 and 6.39, demonstrate that there is much more uncertainty between physically/chemically dissimilar compared to alike salt compositions. A comparatively small number of samples (500) are used because no outliers are generated and a normal distribution can be assumed.[11] The uncertainty in the *direct costs* category is fundamentally uncertainty at the level of the input parameters, not the cost category level. The *modularity* category is closely tied to the capital cost portion of direct costs. Direct and modularity costs are fixed while the remaining cost categories are treated as uncertain - therefore the cost category analysis does not account for all uncertainty. As in the FHR analysis each cost category is assigned a double-trapezoid probability distribution with a mean of 1.0, minimum of 0.7 and maximum of 1.3.

When using probabilistic input parameter values the average cost is in each case higher than the correspond-

[11] The normal distribution is assumed because the analysis cost calculation is additive and the input distributions have fixed boundaries.

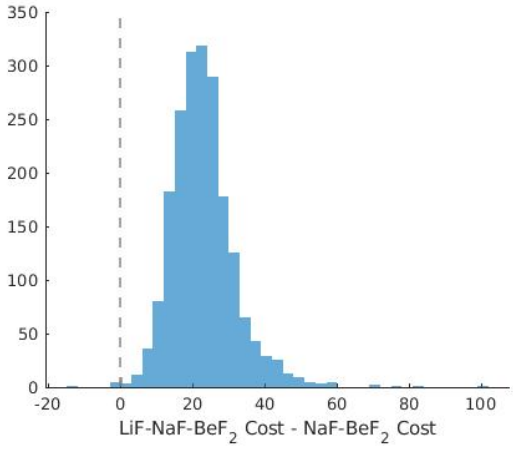

(a) $LiF - NaF - BeF_2$ and $NaF - BeF_2$

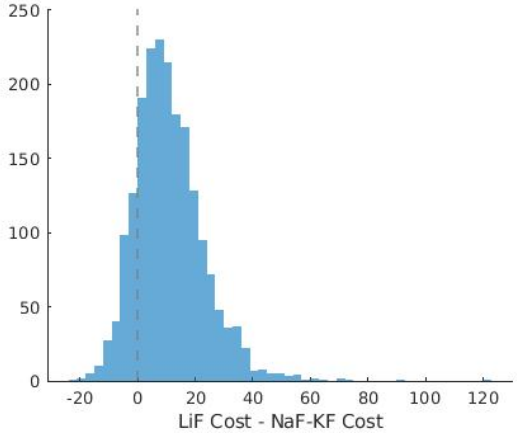

(b) LiF and $NaF - KF$

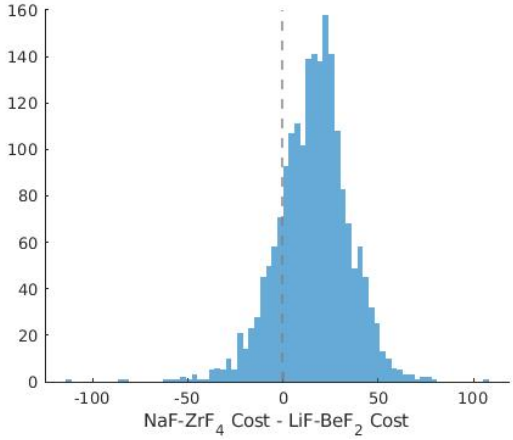

(c) $LiF - BeF_2$ and $NaF - ZrF_4$

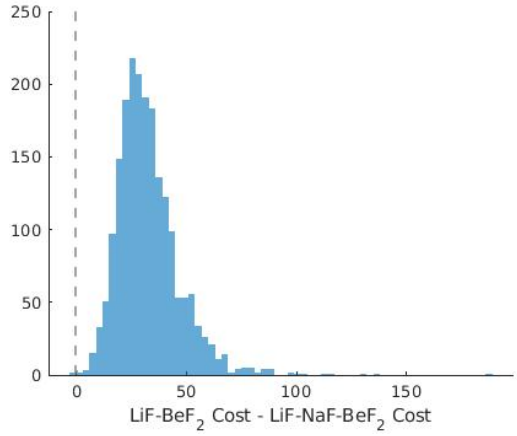

(d) $LiF - BeF_2$ and $LiF - NaF - BeF_2$

Figure 6.36: Histogram of the difference in cost between two salts of each uncertainty run, i.e. each cost difference is the result of total costs calculated using the same input parameters. The histograms are generated using data at the configuration that results in the lowest total cost.

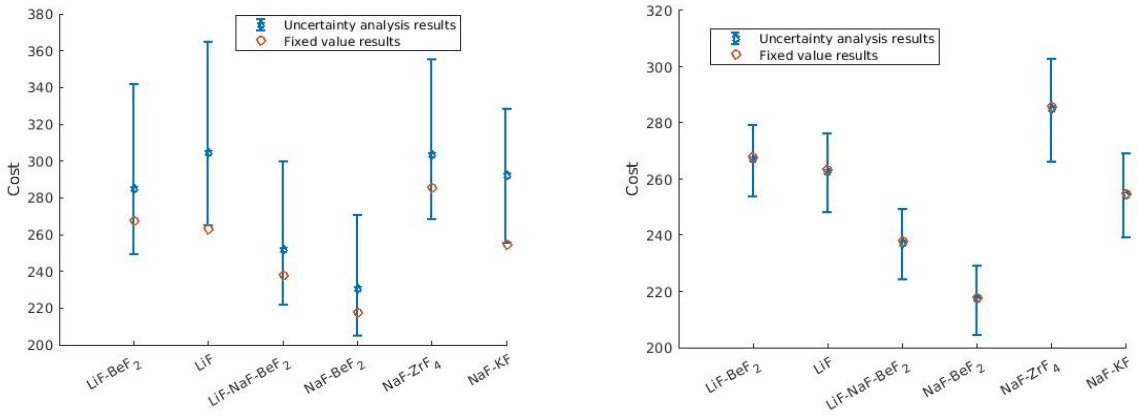

(a) Input parameter uncertainty

(b) Cost category uncertainty

Figure 6.37: Comparison of uncertainty analysis and fixed values at the configurations with the lowest cost. Uncertainty analysis shown with an error range of one standard deviation.

Table 6.17: Sensitivity of total cost to each cost category at the lowest cost configurations. The values shown are scaled relative to the standard deviation, thus the sensitivities are listed as $\Delta Cost$, where $\Delta Cost = b_i \sigma_i$

Cost category	FLiBe	LiF	LiNaBeF	NaBeF	NaZrF	NaKF
Waste	2.46	1.84	2.24	2.21	3.27	1.93
Safety	7.46	7.76	6.59	5.9	7.44	8.04
Proliferation	3.8	2.45	1.67	1.42	1.64	1.59
Feasibility	8.76	10.6	9.62	9.92	15.5	11.3

ing fixed costs, however it is not consistent among configurations, see Figure 6.37. For configurations with a salt that is viable over a large mole % U range, those in Figure 6.35, and performing uncertainty analysis on input parameters, the average and fixed cost differs less than in the case of the other two salts. It is unclear why uncertainty analysis increases the expected cost by a greater amount only for LiF and $NaF-KF$ (this is consistent for all mole % U and FF configurations, not just the lowest cost). This result is not explored further since the difference is no more than 1/2 of a standard deviation. The average-fixed difference is negligible with cost category uncertainty. The variance when using probabilistic input parameter results is considerably higher than with probabilistic factors assigned to several cost categories.

6.6.3 Results: Sensitivity Analysis

Using the same 2000 data points as in uncertainty, sensitivity analysis was performed to gain insight into the effect each of the input variables has on the total cost. The analysis is based upon linear multiple regression.

A list of regression coefficients of select input parameters at the lowest cost configuration of each salt is provided in Table 6.18. Many input parameters have little impact on total cost; parameters with a large impact, that differ significantly between salt composition and/or have unexpected/interesting results are listed. A complete list is provided in Appendix A.2. Many high-magnitude coefficients are similar among configu-

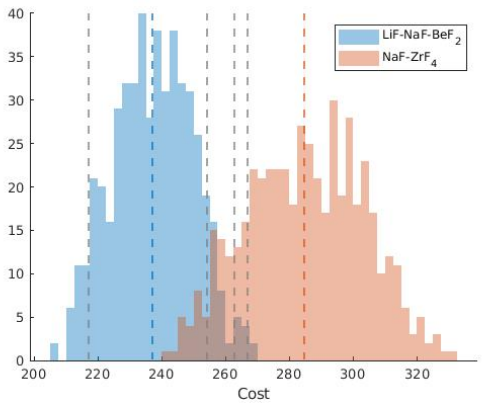

(a) $LiF - NaF - BeF_2$ and $NaF - ZrF_4$

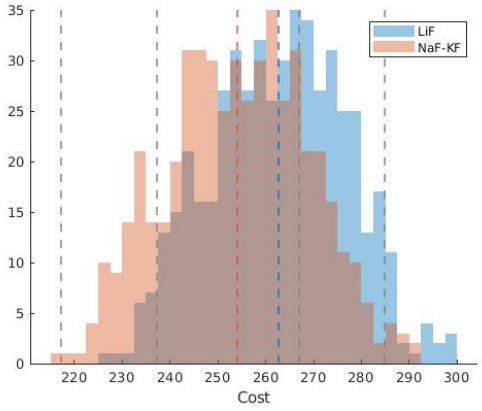

(b) LiF and $NaF - KF$

Figure 6.38: Overlapping histogram plots of the total cost with uncertainty at the cost category level. Data presented is at the lowest cost configuration.

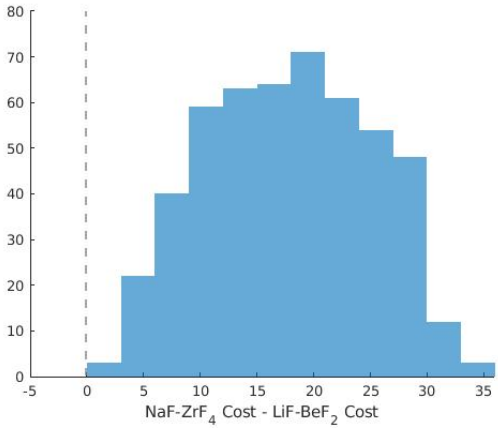

(a) $NaF - ZrF_4$ and $LiF - BeF_2$

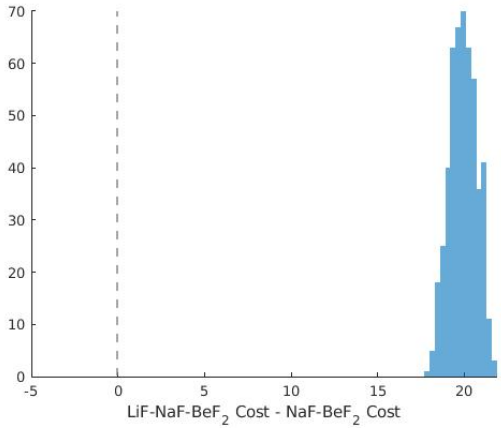

(b) $LiF - NaF - BeF_2$ and $NaF - BeF_2$

Figure 6.39: Histogram of the difference in cost between two salts at each combination of cost category values. Data presented is at the lowest cost configuration.

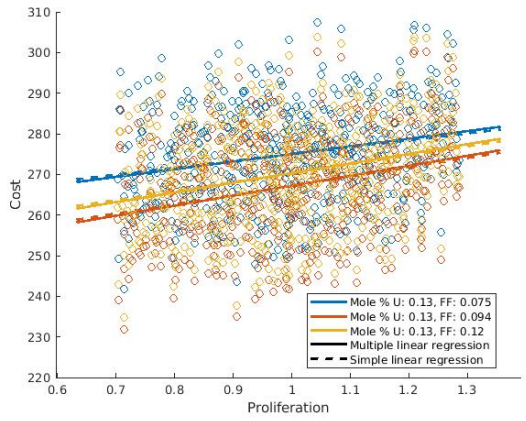

(a) $LiF - BeF_2$: Proliferation

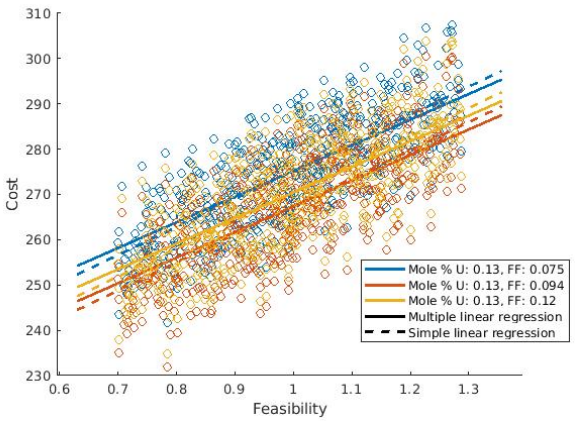

(b) $NaF - BeF_2$: Feasibility

Figure 6.40: Total cost at two low cost configurations as a function of select cost categories. The data points shown indicate the input parameter value (x-axis) and total cost (y-axis) of each run. The resulting coefficient from performing multiple linear regression is shown as the slope of the solid line. Simple linear regression is provided for comparison.

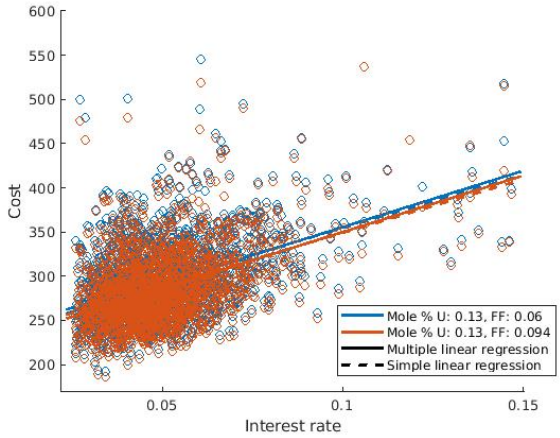

(a) Interest rate

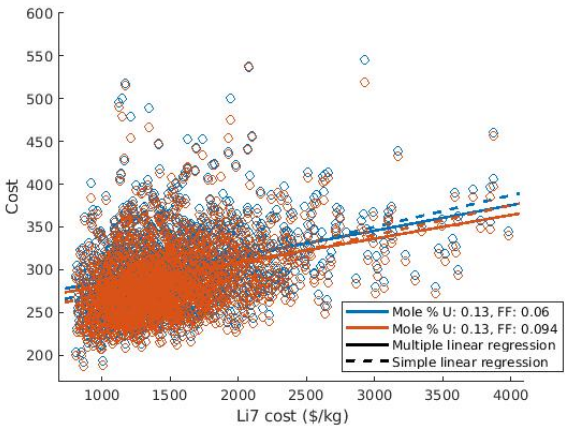

(b) ^{7}Li cost

Figure 6.41: Total cost of $LiF - BeF_2$ salt at two low cost configurations as a function of two input parameters. The data points shown indicate the input parameter value (x-axis) and total cost (y-axis) of each run. The resulting coefficient from performing multiple linear regression is shown as the slope of the solid line. Simple linear regression is provided for comparison.

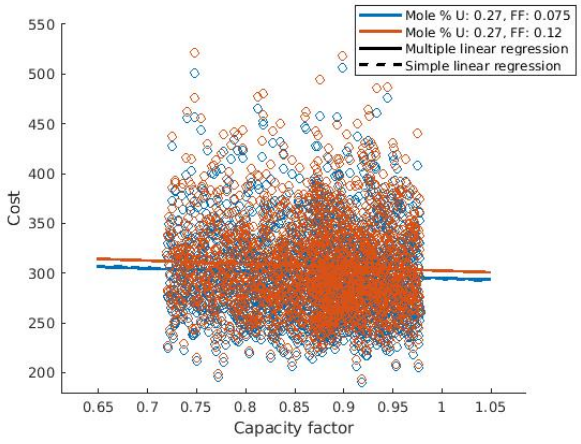

(a) Capacity factor

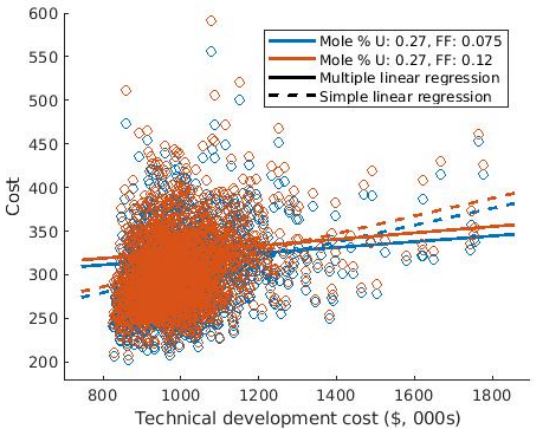

(b) Technical development cost ($, 000's)

Figure 6.42: Total cost of LiF salt at two low cost configurations as a function of two input parameters. The data points shown indicate the input parameter value (x-axis) and total cost (y-axis) of each run. The resulting coefficient from performing multiple linear regression is shown as the slope of the solid line. Simple linear regression is provided for comparison.

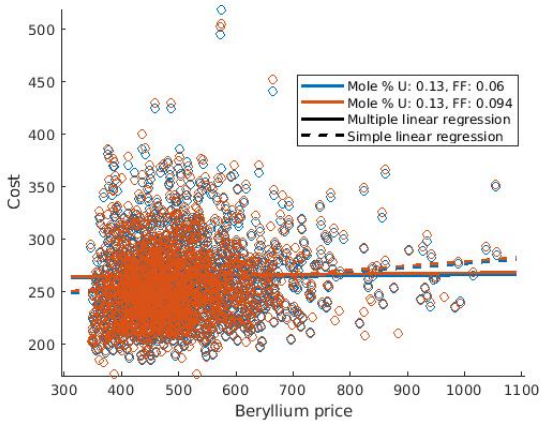

(a) Beryllium price ($/kg)

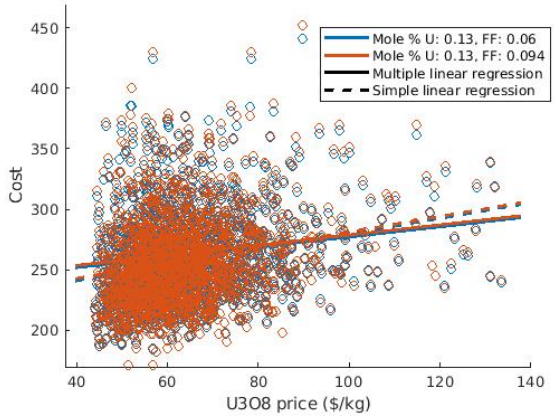

(b) U_3O_8 price ($/kg)

Figure 6.43: Total cost of $LiF - NaF - BeF_2$ salt at two low cost configurations as a function of two input parameters. The data points shown indicate the input parameter value (x-axis) and total cost (y-axis) of each run. The resulting coefficient from performing multiple linear regression is shown as the slope of the solid line. Simple linear regression is provided for comparison.

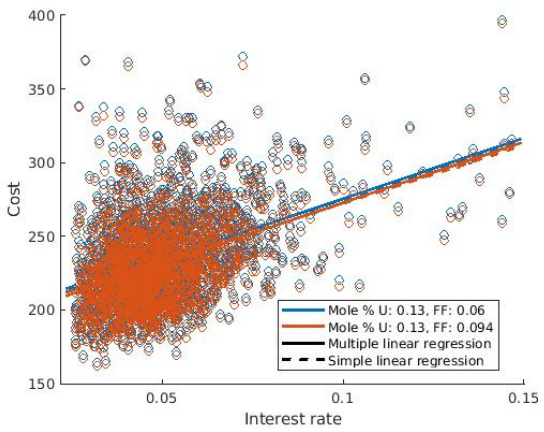

(a) Interest rate

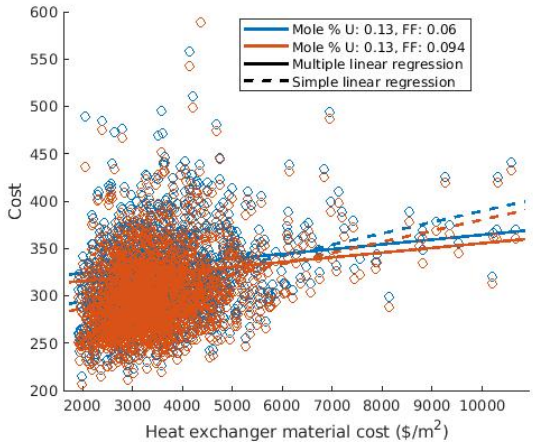

(b) U_3O_8 price ($/kg)

Figure 6.44: Total cost of $NaF - BeF_2$ salt at two low cost configurations as a function of two input parameters. The data points shown indicate the input parameter value (x-axis) and total cost (y-axis) of each run. The resulting coefficient from performing multiple linear regression is shown as the slope of the solid line. Simple linear regression is provided for comparison.

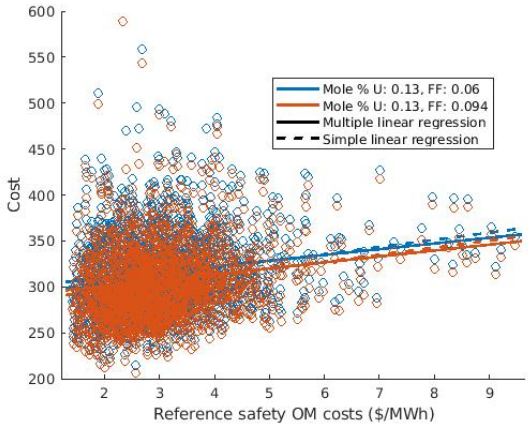

(a) Ref. safety O & M costs ($/MWh)

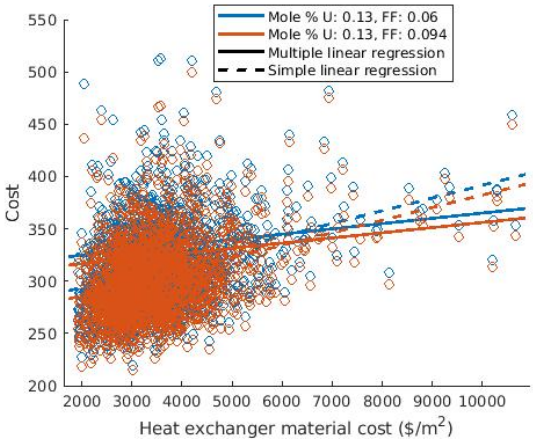

(b) Heat exchanger material cost ($/m^2)

Figure 6.45: Total cost of $NaF - ZrF_4$ salt at two low cost configurations as a function of two input parameters. The data points shown indicate the input parameter value (x-axis) and total cost (y-axis) of each run. The resulting coefficient from performing multiple linear regression is shown as the slope of the solid line. Simple linear regression is provided for comparison.

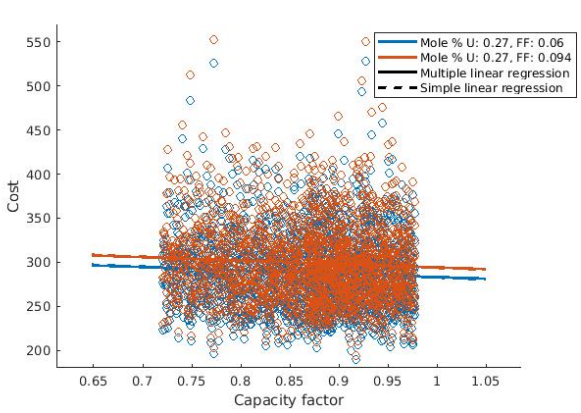

(a) Capacity factor

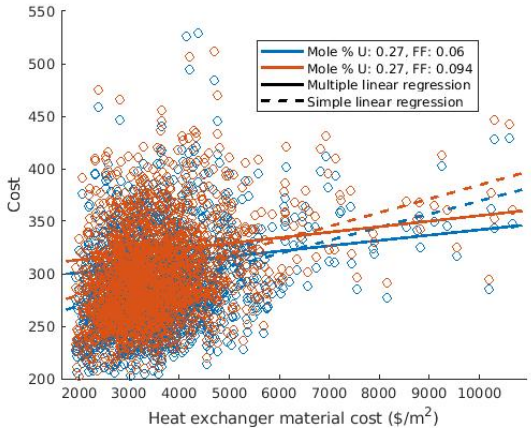

(b) Heat exchanger material cost ($/m^2)

Figure 6.46: Total cost of $NaF - KF$ salt at two low cost configurations as a function of two input parameters. The data points shown indicate the input parameter value (x-axis) and total cost (y-axis) of each run. The resulting coefficient from performing multiple linear regression is shown as the slope of the solid line. Simple linear regression is provided for comparison.

rations. These variables are thus relatively inconsequential to salt composition and core geometry selection. Interest rate, fuel costs (separative work cost, U_3O_8 price) and the ^{7}Li price are expected to have the greatest impact on configuration selection, particularly the salt composition. The low magnitude, even positive in one case, sensitivities of the *Lifetime* parameter may be surprising. A higher (longer) lifetime results in more electricity produced and only increases costs by the additional O & M required. In the FHR the sensitivity to the lifetime is substantially higher, this is partly due to the higher relative variance assigned to the distribution.

The input parameters involved in setting the cost for a reference reactor, such as 'Share tech. develop,' 'Ref. safety OM costs' and 'Non-configuration,' generally have a large impact on the total cost. While the relative costs, 'T_{melt} rel. cost,' 'P_{vap} rel. cost' etc. have a much lesser cost effect and often the associated sensitivity is negative. This may initially seem counter-intuitive because one would surmise that a higher cost of effect/system X should only result in a higher total cost. The relative costs however collectively indicate only the importance of each cost component to calculating the ratio of the configuration-specific to reference design cost. For instance a higher relative cost of T_{melt} means a somewhat lower relative cost of all other relevant components. This outcome is the result of the novel 'cost setting' method discussed in detail in Section 2.8.

Figures 6.41 to 6.46 show the data points and regression results of select configurations. Figure 6.47 shows histograms of the regression coefficient of the input in question for all configurations with the specified carrier salt overlaid with a normal curve that has a mean equal to the average coefficient value and a variance equal to the average expected variance of the coefficient. If the analysed configurations are all the same, the histogram of coefficients should match the expected normal distribution of coefficients. The deviation

Table 6.18: Sensitivity of total cost to key input explanatory variables at the lowest cost configurations. The values shown are scaled relative to the standard deviation, thus the sensitivities are listed as $\Delta Cost$, where $\Delta Cost = b_i\sigma_i$. The sensitivities are ordered in decreasing average magnitude.

Variable	Salt 1	Salt 2	Salt 3	Salt 4	Salt 5	Salt 6
Interest rate	20.5	21.6	17.3	13.3	17.3	18.8
Non-configuration	10.5	11.5	11	10.9	13.4	11.9
Overrun exp.	12.6	13	10.6	8.76	11.5	11.9
Share cap safety costs	12	12.1	9.56	7.34	10	11.4
Share tech. develop	5	9.63	6.25	6.66	12.6	10.5
^{7}Li cost (\$/kg)	12.8	9.28	2.7	0.511	0.621	0.405
Construction time (yrs)	6.81	6.48	5.42	3.54	4.43	5.07
Ref. safety OM costs	5.75	6	5.76	5.78	6.79	7.36
Processing 1 direct rel. cost	1.98	7.28	4.51	3.53	7.87	7.79
Ref. share modularity	5.88	6.18	5.15	4.03	4.87	5.56
U_3O_8	4.23	4.23	5.4	5.07	4.56	7.3
Heat exchanger material cost ($\$/m^2$)	4.36	3.88	5.11	3.71	5.07	5.18
Share cap prolif	4.85	4.37	2.99	2.41	3	3.11
Ref. tech. develop	2.03	3.66	2.39	2.5	4.87	4.07
Processing 1 rel. cost	0.288	5.89	2.63	1.43	3.54	5.71
Separative work	2.25	1.91	3.16	3.19	3.16	4.11
Ref. Pu rel. cost	3.43	3.23	2.35	2.27	2.65	2.41
Safety cap. exp.	-2.96	-2.15	-2.83	-2.52	-1.4	-1.02
Capacity factor	-1.32	-2.15	-1.58	-1.52	-0.898	-2.52
Corrosion rel. cost	2.13	0.143	1.77	3.13	-0.189	1.63
^{7}Li effect rel. cost	-0.712	-1.31	-1.55	-1.59	-1.91	-1.95
Lifetime (yrs)[a]	1.46	-1.44	-1.36	-0.154	-1.44	-3.46
Temp. feedback rel. cost	-0.116	0.645	0.11	0.45	-0.517	1.15
Ref. OM prolif	0.772	0.329	0.0159	-0.0387	-0.0594	-0.0281
β_{eff} rel. cost	-0.0171	0.0538	0.142	0.118	0.06	0.0245

[a]Unreasonable correlations with the *Lifetime* variable were used. The *Lifetime* sensitivities listed have been adjusted to obtain the expected sensitivity in the absence of any correlations. For unadjusted values see Table A.3.

between actual and expected coefficient variance is relatively low, as evidenced by the data in Figure 6.47 and the similar slopes of the different configurations in Figures 6.41 to 6.46. The regression coefficient variance was also tested for several other input parameters and the with the 400 sample dataset.

Multi-variable regression analysis validity:

The sensitivity analysis is conducted using multi-variable linear regression that requires several assumptions to hold true, see Section 4.5. Notable assumptions are the linear relationship between the dependent and independent variables, homoscedasticity and a normal distribution of the residuals. There are too many variables to always be confident in the underlying relationship but Figures 6.41 to 6.46 consistently appear to be close to linear. As well, the inherent structure of the evaluation is largely additive and therefore largely linear. The validity of the homoscedasticity and multi-variable normality can be inferred from Figure 6.48, showing the predicted (from multiple regression) and observed costs for the lowest cost configuration of

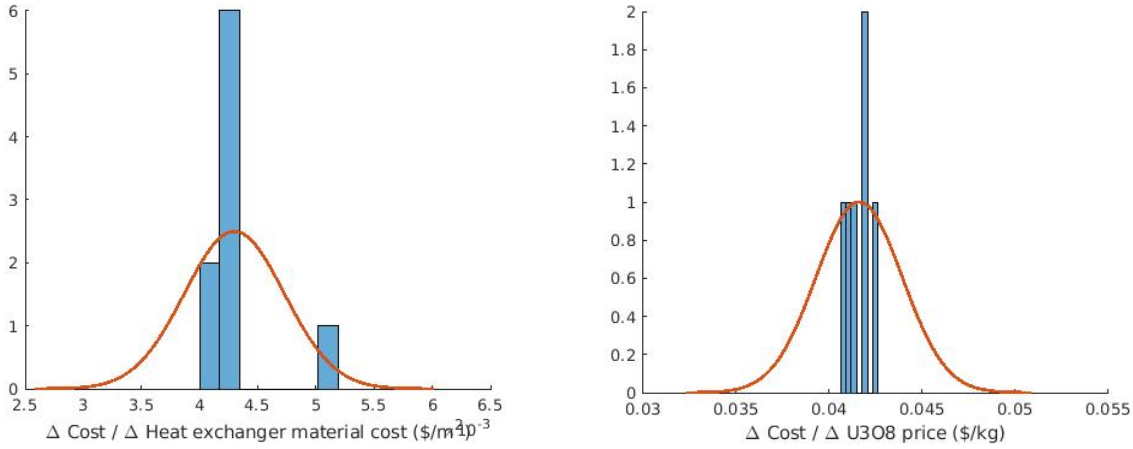

(a) $LiF - BeF_2$: interest rate

(b) $NaF - BeF_2$: U_3O_8 price ($/kg)

Figure 6.47: Histogram of the regression coefficient of the input indicated for all analysed configurations with 2000 samples, overlaid with a normal curve with a mean equal to the average coefficient value and a variance equal to the average expected variance of the coefficient.

each MSR coolant salt. These two assumptions are discussed Section 4.5 since the residuals in the FHR and MSR evaluations are similar. The author's basic conclusion is that the low variance allows for substantial tolerance of heteroscedasticity and non-normality, thus making the model, specifically the calculated sensitivity coefficients, sufficiently valid.

6.7 Conclusions

There are many viable graphite-moderated, circulating-fuel MSR configurations. This analysis covers a large range of those configurations. Specifically 6 salt compositions, 8 to 27% (depending on the salt) mole fraction of uranium cations (mole % U) and a core fuel fraction between 0.06 to 0.15. A large dataset covering salt properties, heat exchanger requirements, initial enrichment, on-line refuelling rate, temperature coefficients etc. was developed to supply a comprehensive evaluation program. The evaluation aims to determine the relative total cost across the range of analysed configurations.

Three major independent, non-neutronics analysis were covered: optimization of the rate of on-line removal of xenon-135 gas, calculation of the effective delayed neutron fraction correction factor (B_{eff}/B) and the optimization of the channel pitch. Additional calculations/neutronics simulations were conducted in an attempt to generate all other data that would be of significance to the evaluation program. The accuracy of the xenon-135 optimization is limited by the present knowledge of graphite behaviour under radiation and its interaction with both xenon and fuel salt. The accuracy of many other analysis are limited by the reliability of available property data.

A comprehensive evaluation of a large range of potential MSR designs has been conducted. The analy-

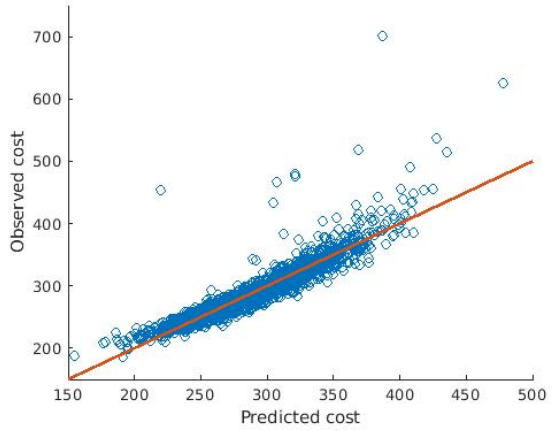

(a) $LiF - BeF_2$

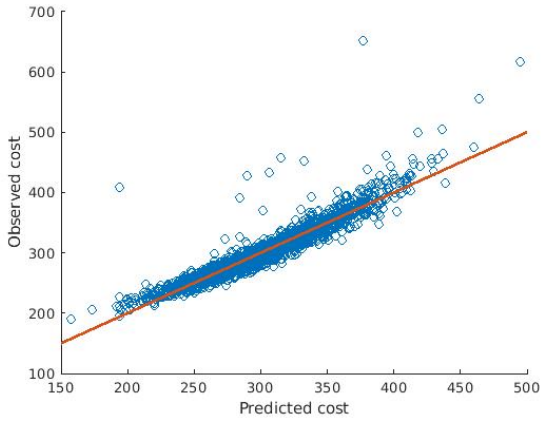

(b) LiF

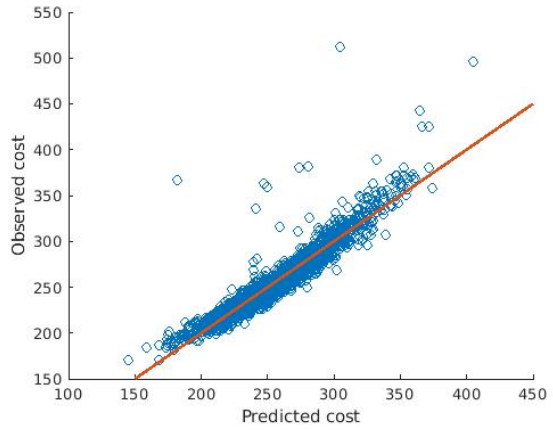

(c) $LiF - NaF - BeF_2$

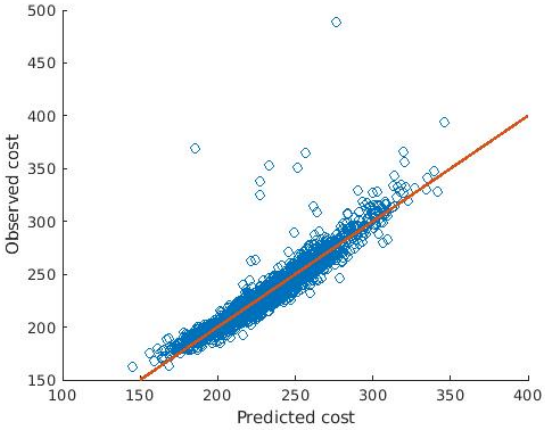

(d) $NaF - BeF_2$

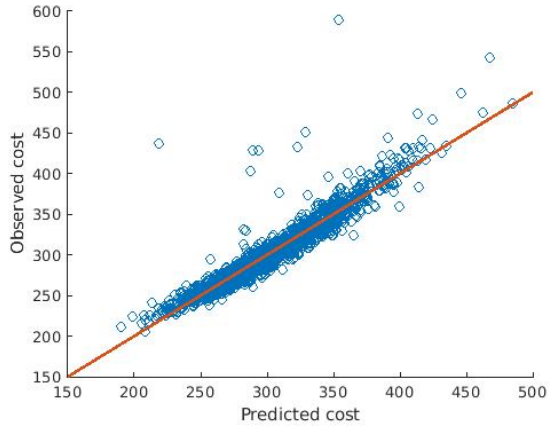

(e) $NaF - ZrF_4$

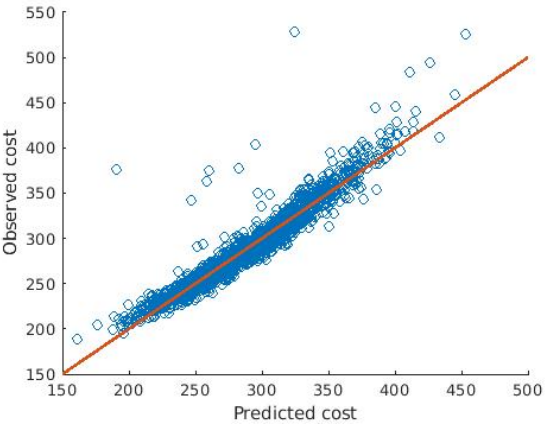

(f) $NaF - KF$

Figure 6.48: Predicted compared to observed cost for the lowest cost configuration of each MSR coolant salt. The difference between the two values (deviation from the line) represents the residuals.

sis succeeded in identifying trends that lead to a particular geometric configuration and salt composition that is likely to result in the lowest total cost. For nearly all configurations (mole % U and FF) where $NaF - BeF_2 - UF_4$ is a possible carrier salt, the evaluation suggests that it is the carrier salt that results in lowest total cost; at least several configurations of the other salts result in adequately low costs such that they can be considered viable alternatives. The benefits of the superior neutronics of ^{7}Li appear to only outweigh the cost of enriching Li and managing tritium production if the costs of doing so are significantly lower than the expected cost estimates employed. The optimum core fuel salt concentration is approximately 2.0% and so the optimum mole % U or FF value is largely dependent on the other. Among the configurations analysed, higher mole % U and lower FF combinations tended to have lower costs, this is expected to be largely the result of higher parasitic absorptions in the carrier salt.

It should be repeated that this analysis only considers costs that are a function of salt composition and analysed geometric design parameters (configuration). It is recognized that the analysis likely even missed some significant, configuration-dependent costs. The total costs are in units of millions of USD, 2021. The total cost, ≈240 million is only $\approx 1/4^{th}$ of the total value of electricity generated by this reactor.[12] Considering the non-configuration dependent costs is important as it puts the cost differences in the optimization in context; since the relative differences in cost are only $\approx 1/4$ as large, other considerations such as availability of material become more of a concern.

The uncertainty and sensitivity analysis using distributed value input parameters provided some useful insights.

- The variance in total cost of the lowest cost configurations is high (≈ 35 - 45) compared to the difference in cost between the configurations. This suggests that the total cost is more susceptible to the circumstances of a MSR build than to the configuration selected.
- The probability that the evaluation finds a lower average cost salt to have a higher cost, if the input parameter set is independent for each configuration (overlap in total cost distribution, Figure 6.34), is much higher than the probability accounting for the input parameter set (< 0 in lower average cost difference distribution, Figure 6.36).
- Total cost distributions have similar relative variance (σ/μ) and similar shape - notably positive skew and excess kurtosis.
- The expected cost of the lowest-cost configurations as calculated by uncertainty analysis (the average) are approximately $15-40 MM (6-16%) higher, depending on the salt composition, compared to the corresponding fixed values calculation.
- There is much more uncertainty in the cost difference between physically/chemically dissimilar, compared to alike, salt compositions.

[12]The value (at the time of commissioning) of 200 MW of electricity produced with a CF of 0.92, by this hypothetical MSR, over the 7 year lifetime, with a 5% interest rate is $933 billion.

- Few of the sensitivity analysis coefficients are particularly notable. The interest (or discount) rate, non-configuration, and technological development cost parameters, total development costs and share borne by one reactor, are unsurprisingly very consequential. For the many input variables with low sensitivity coefficients, the lack of a definite basis for the estimated cost distribution is of little consequence; an incorrect estimated cost has little effect on the total cost.

Chapter 7

Conclusion

This dissertation presents a means of evaluating, on a cost basis, reactors that use molten salt as the heat transport fluid. The evaluation is then applied to two reactors types, the Fluoride High-temperature Reactor (FHR) and the graphite-moderated circulating-fuel concept.

The evaluation of both reactor types make use of much of the same methodology, research and analysis. In general, areas in which the study of each overlap were documented together. The evaluation and cost estimation framework, corrosion, heat exchanger analysis, the cost and supply of salt constituents, ^{7}Li enrichment optimization as well as the methodology of the uncertainty and sensitivity analysis are discussed in Chapters separate from either reactor analysis.

A summary of the approach taken to arrive at the total costs is provided. In doing to contributions to knowledge generated over the course of this study are highlighted. Conclusions specific to the FHR and circulating-fuel type are generally covered in their respective chapters. Throughout this dissertation the results of the evaluation of the two reactors types have not been compared. Despite the challenges of doing so and difficulty devising valid conclusions, a limited comparison is made and discussed.

Given the relatively early stage of development that MSRs are in today, there is considerable opportunity to expand the knowledge base. This chapter includes recommendations for future work, both cost related and otherwise.

7.1 Summary, Conclusions and Contributions

The construction of new nuclear power plants has stagnated in many areas of the world. There is a renewed interest in alternative reactor design types that can potentially speed-up nuclear power deployment; among those alternative design types are MSRs. There exists an assortment of MSR designs that differ significantly and thus development is divided amongst them. This dissertation aims to shrink what is considered viable design space by presenting a broad, investigation of MSRs that does not strictly focus on one particular aspect of technical performance but relates performance to key practical cost considerations. Such an inves-

tigation is uncommon, and of this breadth (to the author's knowledge) non-existent in literature.

Typical design optimization studies generally only consider a small number of effects. Though such analysis is more straightforward, the results may provide little utility on their own. The 'Cost Estimation Framework' seeks to provide a novel means of incorporating many, often unrelated, cost components. It is built upon dividing costs among 6 categories, 'direct,' 'waste,' 'safety,' 'proliferation,' 'modularity,' and 'feasibility.' The reactor evaluation is fundamentally a means of comparing costs and is based upon both the 'Cost Estimation Framework' and the division into the 6 categories.

The evaluation requires a large amount of input data. Much of the input data generated and/or gathered is applicable to all MSR design types. The methodology and resulting data on corrosion, the heat exchanger analysis, the cost and supply of salt constituents, ^{7}Li enrichment optimization is documented in a chapter separate from reactor-specific analyses. Similarly, methodology and parameters of the uncertainty and sensitivity analysis to be conducted for each reactor type is discussed separately, including the novel 'double-trapezoid probability distribution.'

The FHR type is analysed across a wide design space. In addition to examining salt composition and TRISO particle packing fraction, the share of graphite pebbles is varied. This unique approach to fuel-moderator separation results in higher neutronic efficiency. It is assumed that the pebbles are well mixed and the reactor undergoes continuous refuelling, however this presents modelling challenges for the necessary depletion simulations. The solution devised was to divide the pebbles among dispersed fuel zones and to refuel one zone at a time as required. A key finding of the evaluation is that the addition of graphite pebbles for a given core fuel concentration generally results in a low cost. Another is that $2LiF - BeF_2$ (FLiBe) salt results in the lowest cost in $\approx 95\%$ of scenarios (scenarios or runs defined by the input parameters in the uncertainty analysis).

The graphite-moderated, circulating-fuel MSR is analysed across a wide design space as well. The salt composition, mole % uranium cation (mole % U) and fuel fraction are the design parameters examined. The reactor is continually refuelled in a setup where fuel is added at the same volumetric flow rate as fuel is removed. A series of depletion simulations where the salt addition/removal rate is updated between burnup steps was developed. The data generation involved novel optimizations of the xenon removal rate and channel pitch. The evaluation demonstrated that the costs involved for all salts are quite similar and thus the salt selection is not extremely consequential. It is found that $NaF - BeF_2$ salt results in the lowest cost in all but outlier scenarios.

Both evaluations show smooth results and distinct patterns, a promising indication that they were executed properly. The sensitivity analysis, particularly of input parameters, proved to be effective and informative. The uncertainty and sensitivity analysis with distributed value cost categories results in similar variance and sensitivity coefficients among all configurations. For many input parameters a reliable and valid method to

determine probability distribution parameters (μ, min, max) was not found. The cost category uncertainty analysis appears to provide little insight other than demonstrating by comparison the utility of the input parameter analysis.

7.2 Discussion of Reactor Class Comparison

The results of the evaluation of the FHR and MSR are compared here. This is a difficult comparison to make. The focus of each evaluation is principally on the more reliable and useful comparison among configurations of the same type.

A comparison of uncertainty between an FHR and MSR configuration would carry much of the configuration-independent uncertainty from each while comparison of configurations of the types can focus only on the uncertainty that is the result of configuration differences. It is recognized that the parameters of many input distributions are incorrect by significant amounts but the evaluation results can still be useful if the same incorrectly value(s) are applied to all configurations. Many input parameters are specific to or must be applied separately to either reactor. For instance, though each reactor type has a technological development cost, different systems must be developed, thus the uncertainty would be largely independent of the other reactor - the same development cost cannot be applied to each. The consequence of the independent uncertainties is that uncertainty in the cost difference between the two reactor types is much greater than the uncertainty between configurations of the same type.

Consideration has been given to the effect that reactor size/power has on total cost. Costs consider power where possible. O & M costs are generally a function of either fuel usage or a reference safety/proliferation cost per unit thermal power and so are roughly proportional to power. Capital costs are more difficult, some costs such as those related to heat exchanger and the initial cost of fuel/salt components are inherently proportional (or nearly so) to power. Others such as modularity, detritation and corrosion capital costs must be adjusted to the power at the discretion of the user of the evaluation program.

This analysis likely fails to compare reactors types with differing power outputs and lifetime effectively. A significant portion of capital costs must be adjusted to power at user discretion and all reactors compared must use the same reference power for adjustment. These cost adjustments relative to a universal reference power have not been made. As well, the cost is not adjusted relative to some universal reference lifetime. As a result, and all else being equal, reactors with a shorter lifetime will appear to have a lower total cost even though they are less economical.[1]

Since the expected electric power and lifetime of the FHR is 100 MWe and 15 years, while the MSR is 200 MWe and 7 years, each have a (similar at-reactor-start-up) present value of electricity sales. Thus a cost

[1] Note that differences in lifetime among configurations of the same reactor type are appropriately accounted for (longer lifetime results in lower cost), these comments pertain specifically to different reactor types.

comparison is not exceedingly unreasonable. The expected total cost, as determined using distributed input parameters, of the lowest cost configuration of each salt composition evaluated is listed below. The MSR costs are also adjusted to account for the difference in present value of electricity sales at a interest rate of 5.0%.

- FHR: 175, 190, 245, 229
- MSR: 288, 299, 256, 232, 304, 288
- MSR adjusted 258, 268, 230, 208, 273, 258

It can be concluded, according to the evaluation program setup and input parameter values selected, that at the lowest-cost configuration of each reactor type, the FHR results in a significantly lower cost. However there is more variance in cost among the lowest cost salts of the FHR. The total cost is more sensitive to salt composition likely because there is a much higher volume in the core and therefore higher parasitic absorptions and higher salt component costs. The difference in cost is likely not large enough to draw any definite conclusions because of the substantial differences between the two designs types and the difficulty in quantifying the timing and cost associated with replacing graphite.

A major point of uncertainty for both reactors is the lifetime. Given the need to swap out graphite, it is unclear what point should be considered the end of life for the FHR and how to best handle core/graphite swaps in the MSR. The choice of how to handle this issue (which has been discussed in both the FHR and MSR chapters) is likely to be the determining factor in which reactor type one considers lower cost.

7.3 Recommendations for Future Work

The cost estimation framework and associated evaluations consist of many aspects that would benefit from further development. Few studies exist that explore costs involved in alternative, Generation IV, reactor design concepts, so there is considerable opportunity to expand the knowledge base. Specific recommendations of future research, both cost related and of other phenomenon, include:

- The cost to address safety and proliferation concerns is undoubtedly significant. However the costs are difficult to quantify. A more thorough understanding of the underlying causes that give rise to these costs would be beneficial.
- Repeated assembly in a factory setting results in lower manufacturing costs than on-site construction. The difficulty lies in quantifying the expected cost differences. Several lines of potential research may stem from this idea. Research may be conducted to estimate the expected increase in cost that cannot be attributed to specific material, components or tasks. This (lack of) modularity research should be expanded to estimate the cost difference as a function of reactor size/power and other configuration parameters.
- The cost of preventing/addressing corrosion of metals by molten salts is another aspect that is difficult to quantify. It is unlikely that a reliable and valid means of accurately predicting that cost can be

developed, however comparisons between molten salt compositions would be more straightforward and useful. Experimental work aiming to replicate reactor conditions would be the most useful; under irradiation, in contact with graphite, with impurities, with fission products etc.

- Thermal scattering libraries for any molten salts are not included in standard nuclear libraries. MSR research would greatly benefit from easily accessible libraries of salts compositions. The development of thermal scattering libraries requires considerable molecular dynamic analysis that is well beyond the scope of this project.
- The available data on heat capacity and thermal conductivity is quite limited and often unreliable. More reliable predictive methods or an expansion of the experimental literature would be useful.
- Reliable theoretical predictions, or measurements of, the fuel temperature of molten salt cooled designs were not found. The fuel temperature is an important metric that may be the limiting factor of core power density. Studies that aim to determine determine the difference in temperature between the fuel and coolant as a function of particle/pebble geometry and power would benefit the field.
- The degree of homogeneity of pebble dispersion (ie. mixing) in an FHR is unclear. This study assumed that pebbles, by fuel composition, are randomly dispersed. Studies assessing the validity of that assumption, as well as other pebble distributions (partial mixing, longitudinal burnup gradient etc.) would make clear to researchers which potential pebble distribution and refuelling schemes to investigate. Moreover it would provide insight into the challenges associated with having randomly dispersed graphite pebbles throughout the core.
  - Even without knowing the amount of pebble mixing it may be useful to conduct non-fully-mixed simulations. However, this would poise a considerable modelling challenge. A conceivable method would be to assign each pebble to a burnup zone based on the expected position and burnup of the previous refuelling step as well as the increase in burnup at that position since the last refuelling. This would require tracking the position and burnup of each pebble between steps as well as re-making the position of the zones each step. It would be interesting to see the effect that mixing has on the expected refuelling rate.
- Consider the evaluation of different numbers of fuel zones and different pebble discard probabilities. It is most simple for the number of fuel zones to equal to 1/discard-probability but this is only the minimum number of fuel zones. Additional fuel zones are not intended to change the refuelling function but rather model it more precisely. Ultimately, this analysis should culminate in a plot showing refuelling rate as a function of discard-probability and the number of refuelling zones.
  - Determining the refuelling rate as a function of the number of refuelling steps before equilibrium could be done as well. However, the distribution of burnup steps can be tracked quite easily and is it quite an even distribution (see Figures A.1 to A.5) so this should be less of a priority.
  - A potential way to reduce computational requirements is first to conduct a limited number, or even just one, simulation with very high precision, i.e. many of burnup zones, many refuelling steps and short burnup steps. Then compare the results to (a) corresponding lower-precision

simulations of the same configuration(s). Finally, adjust the results of the simulations for all other configurations based of the high-low precision comparison.

- A further understanding of the costs of many processes involved in the operation of a MSR but not widely practised among the current fleet of reactors would be beneficial. These processes include but are not limited to, xenon removal, tritium removal/management, fission product processing and refuelling in a salt-cooled pebble-bed reactor. Not only costs, but the physical mechanics of xenon transport and fission product processing should be further studied as well.
- The MSR analysis did not have temperature zones, the fuel salt and moderator were all one uniform temperature. Though from experience in the analysis of the FHR, the effect of temperature zones was known to be small, future work that would see greater benefit from modelling precision should consider multiple temperature zones.
- Combining the analysis of the movement of delayed neutron precursors with neutronics simulations would lead to valuable transient safety analysis. This would be a considerable undertaking however; existing neutronics codes would have to undergo substantial modification. The resulting reactivity changes as a function of reactor power and circulation speed over time would provide useful insights into reactor behaviour in both normal operation and accident scenarios.
- Both the FHR and MSR analysis do not consider graphite temperature reactivity feedbacks, the reasoning is discussed in Section 5.5.3 and 6.3.9. The inclusion of graphite temperature feedbacks could prove valuable but a proper, thorough analysis would likely have to consider the effect of the feedbacks in transient conditions - thereby making the analysis complex and computationally expensive.
- Future reactor evaluations based off the framework/analysis presented in the work should consider one-at-a-time sensitivity analysis, where each variable is evaluated at several values while all other variables are held at their expected values. The uncertainty and sensitivity analysis was set-up aiming to capture, as comprehensively as reasonably possible, low-probability values (tails in the probability distributions) and non-linear effects. It can be safe to conclude that this was achieved. It did not cause issues in the uncertainty analysis but the heteroscedasticity and non-normality may be somewhat problematic in the sensitivity analysis. Further investigation could test one-at-a-time sensitivity analysis, if it deviates significantly from multi-variable regression than one-at-a-time should be used, whereas if the coefficients are very similar then multi-variable regression is likely valid and may be the preferred option because it captures the effects of multiple inputs having low-probability (tail) values.

Bibliography

[1] "Fluoride-Salt-Cooled High-Temperature Reactors," https://www.ornl.gov/content/fluoride-salt-cooled-high-temperature-reactors, Accessed: 2022-05-29.

[2] O. E. et al, "Climate Change 2014," Intergovernmental Panel on Climate Change, Tech. Rep., 2014.

[3] P. L. Joskow, "Comparing the costs of intermittent and dispatchable electricity generating technologies," American Economic Review, vol. 101, no. 3, pp. 238–41, May 2011. [Online]. Available: http://www.aeaweb.org/articles?id=10.1257/aer.101.3.238

[4] BP, "Statistical Review of World Energy," BP, Tech. Rep., 2017.

[5] J. R. Lovering, A. Yip, and T. Nordhaus, "Historical construction costs of global nuclear power reactors," Energy Policy, vol. 91, pp. 371–382, 2016.

[6] "Finland opens nuclear power plant amid concerns of Europe energy war," https://www.theguardian.com/environment/2022/mar/12/finland-opens-nuclear-power-plant-amid-concerns-of-europe-energy-war, Accessed: 2022-05-09.

[7] T. E. M. W. G. of the Generation IV International Forum, "Cost Estimating Guidelines for Generation IV Nuclear Energy Systems," Gen IV International Forum, Tech. Rep., 2007.

[8] "From Gen I to Gen III," http://www.iaea.org/inis/collection/NCLCollectionStore/_Public/44/078/44078368.pdf, International Atomic Energy Agency, Accessed: 2018-06-28.

[9] "Generation I reactors," http://www.radioactivity.eu.com/site/pages/Generation_I_Reactors.htm, Radioactivity.eu with the help of EDP Sciences and IN2P3, Accessed: 2018-06-28.

[10] "Development and Global Deployment of ABWR," https://www-pub.iaea.org/mtcd/meetings/PDFplus/2009/cn164/Sessions/08KS_1_Shimizu.pdf, International Atomic Energy Agency, Accessed: 2018-06-28.

[11] E. L. Lisbeth Gronlund, David Lochbaum, "Nuclear Power in a Warming World: Assessing the Risks, Addressing the Challenges," Union of Concerned Scientists USA, Tech. Rep., 2007.

[12] "Generation IV Goals," https://www.gen-4.org/gif/jcms/c_9502/generation-iv-goals, Accessed: 2018-06-24.

[13] "Generation IV Goals," https://www.gen-4.org/gif/jcms/c_9260/public, Accessed: 2018-06-24.

[14] T. R. U. for Generation IV Nuclear Energy Systems, "Technology Roadmap Update for Generation IV Nuclear Energy Systems," Gen IV International Forum, Tech. Rep., 2014.

[15] M. Rosenthal, P. Kasten, and R. Briggs, "Molten-salt reactors—history, status, and potential," Nuclear Applications and Technology, vol. 8, no. 2, pp. 107–117, 1970.

[16] International Atomic Energy Agency, Ed., Advances in Small Modular Reactor Technology Developments. International Atomic Energy Agency, 2016 edition.

[17] Developing the Next Generation of MSR Systems in Russian Feberation, ser. Technical Meeting on the Status of MSR Technology. International Atomic Energy Agency (IAEA), 2016.

[18] Thomas J. Dolan, Ed., Molten Salt Reactors and Thorium Energy. Woodhead Publishing, 2017 edition.

[19] Haneklaus, N and Kendrick, J and Huddar, L and Zweibaum, N and Peterson, PF and Hughes, J and Blandford, E and Lv, Q and Sun, X and Yoder, G, "Thermal Hydraulic Benchmarking Exercises to Support Fluoride-Salt-Cooled, High-Temperature Reactor (FHR) Licensing," in International Topical Meeting on Nuclear Reactor Thermal Hydraulics (NURETH-16), vol. 8, 2015.

[20] "Kairos Power and its Vision," https://kairospower.com/company/, Accessed: 2022-05-09.

[21] Bettis, ES and others, "The Aircraft Reactor Experiment—Design and construction, Nucl," Sci. Engng, pp. 804–25.

[22] J. Abrahamson, "Graphite sublimation temperatures, carbon arcs and crystallite erosion," Carbon, vol. 12, no. 2, pp. 111–141, 1974.

[23] R. Briggs, "Molten salt reactor program semiannual progress report for period ending july 31, 1964," ORNL, Tech. Rep., 1964.

[24] ——, "Molten salt reactor program semiannual progress report for period ending july 31," ORNL, Tech. Rep., 1969.

[25] Zirconium Hydride, chemicalland21.com, information sheet.

[26] B. Marsden, "Graphite technology course," NRC, Tech. Rep., 19xx.

[27] A. Bolewski Jr, M. Ciechanowski, A. Dydejczyk, and A. Kreft, "A practical method for measuring the boron equivalent of graphite impurity," Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms, vol. 237, no. 3-4, pp. 602–612, 2005.

[28] B. Rouben, "Kinetics with delayed neutrons," McMaster University, lecture notes, EP 4D03/6D03, 2017.

[29] "Prompt Neutron Lifetime," https://www.nuclear-power.net/nuclear-power/fission/prompt-neutrons/prompt-neutron-lifetime-pnl/, Accessed: 2018-06-24.

[30] E. Bettis, W. Cottrell, E. Mann, J. Meem, and G. Whitman, "The aircraft reactor experiment—operation," Nuclear Science and Engineering, vol. 2, no. 6, pp. 841–853, 1957.

[31] E. S. Bettis, R. W. Schroeder, G. A. Cristy, H. W. Savage, R. G. Affel, and L. F. Hemphill, "The aircraft reactor experiment—design and construction," Nuclear Science and Engineering, vol. 2, no. 6, pp. 804–825, 1957. [Online]. Available: https://doi.org/10.13182/NSE57-A35495

[32] J. C. Gehin, "History of the ORNL Molten Salt Program," Oak Ridge National Laboratory, Tech. Rep., 2015, workshop.

[33] Mark D. Bowles, Robert S. Arrighi, NASA's nuclear Frontier. NASA, 2014.

[34] "MSRE's 50th: The short-lived Molten Salt Reactor Experiment is far from forgotten," https://www.ornl.gov/news/msres-50th, Oak Ridge National Laboratory, Accessed: 2018-06-25.

[35] R. Briggs, "Molten salt reactor program semiannual progress report for period ending august 31, 1967," ORNL, Tech. Rep., 1967.

[36] P. R. Kasten, E. Bettis, and R. C. Robertson, "Design studies of 1000-mw (e) molten-salt breeder reactors," Oak Ridge National Lab., Tenn., Tech. Rep., 1966.

[37] J. Engel, H. Bauman, J. Dearing, W. Grimes, H. McCoy, and W. A. Rhoades, "Conceptual design characteristics of a denatured molten-salt reactor with once-through fueling," Oak Ridge National Lab., TN (USA), Tech. Rep., 1980.

[38] R. Briggs, "Molten salt reactor program semiannual progress report for period ending february 28, 1966," Report No. ORNL-3122,(Oak Ridge National Laboratory), 1966.

[39] J. Hightower Jr, "Process technology for the molten-salt reactor 233 u–th cycle," Oak Ridge National Lab., Tech. Rep., 1975.

[40] D. D. Gerardin, M. Allibert, D. Heuer, A. Laureau, E. Merle-Lucotte, and C. Seuvre, "Design evolutions of molten salt fast reactor," in nternational Conference on Fast Reactors and Related Fuel Cycles: Next Generation Nuclear Systemes for Sustainable Development (FR17), 2017.

[41] E.Merle, "Concept of European Molten Salt Fast Reactor (MSFR)." Generation IV International Forum, 2017.

[42] R. Li, S. Wang, A. Rineiski, D. Zhang, and E. Merle-Lucotte, "Transient analyses for a molten salt fast reactor with optimized core geometry," Nuclear Engineering and Design, vol. 292, pp. 164–176, 2015.

[43] A. Laureau, D. Heuer, E. Merle-Lucotte, P. Rubiolo, M. Allibert, and M. Aufiero, "Transient coupled calculations of the molten salt fast reactor using the transient fission matrix approach," Nuclear Engineering and Design, vol. 316, pp. 112–124, 2017.

[44] D. Heuer, A. Laureau, E. Merle-Lucotte, M. Allibert, and D. Gerardin, "A starting procedure for the msfr: Approach to criticality and incident analysis," in International Congress on Advances in Nuclear Power Plants (ICAPP2017), 2017, p. 17219.

[45] E. Merle-Lucotte, M. Allibert, D. Heuer, M. Brovchenko, A. Laureau, V. Ghetta, and P. Rubiolo, "Preliminary design studies of the draining tanks for the molten salt fast reactor," in European Nuclear Conference (ENC 2014), 2014, pp. 220–230.

[46] "History (FHR)," http://fhr.nuc.berkeley.edu/history/, Accessed: 2018-06-26.

[47] C. Forsberg, L.-w. Hu, P. Peterson, and K. Sridharan, "Fluoride-salt-cooled high-temperature reactor (fhr) for power and process heat," Massachusetts Inst. of Technology (MIT), Tech. Rep., 2015.

[48] A. T. Cisneros, J. K. Choi, A. Y. Chong, M. Fratoni, and S. Hong, "Technical description of the "mark 1" pebble-bed fluoride-salt-cooled high-temperature reactor (pb-fhr) power plant," UCBTH-14-002, Tech. Rep., 2014.

[49] D. Williams, L. Toth, K. Clarno et al., Assessment of Candidate Molten Salt Coolants for the Advanced High Temperature Reactor (AHTR). United States. Department of Energy, 2006.

[50] J. H. Yoo, B. G. Lee, Y. Kang et al., "Thermodynamic data-base for metal fluorides," 2001.

[51] M. Bunn, J. P. Holdren, S. Fetter, and B. Van Der Zwaan, "The economics of reprocessing versus direct disposal of spent nuclear fuel," Nuclear Technology, vol. 150, no. 3, pp. 209–230, 2005.

[52] D. Shropshire, K. Williams, W. Boore, J. Smith, B. Dixon, M. Dunzik-Gougar, R. Adams, D. Gombert, and E. Schneider, "Advanced fuel cycle cost basis," Idaho National Laboratory (INL), Tech. Rep., 2009.

[53] E. Supko, "Minimizing risks associated with post-shutdown spent fuel storage and llw disposal," in Infocast Nuclear Power in the Competitive Era postconference workshop, Developing Risk Strategies for Successful Decommissioning, Washington, DC, 1998.

[54] D. Naglav, M. R. Buchner, G. Bendt, F. Kraus, and S. Schulz, "Off the beaten track—a hitchhiker's guide to beryllium chemistry," Angewandte Chemie International Edition, vol. 55, no. 36, pp. 10 562–10 576, 2016.

[55] C. Cooksey, "Health concerns of heavy metals and metalloids," Science progress, vol. 95, no. 1, pp. 73–88, 2012.

[56] D. Haught, "Elemental mercury management and storage fees," USA, Department of Energy, Tech. Rep., 2019.

[57] U. EIA, "Capital cost estimates for utility scale electricity generating plants," US Department of Energy, Energy Information Administration. Parts on PHP and batteries http://www. eia. gov/analysis/studies/powerplants/capitalcost/pdf/capcost_assumption. pdf, 2016.

[58] R. Gold, "Tab swells to $25 billion for nuclear-power plant in Georgia," https://www.wsj.com/articles/tab-swells-to-25-billion-for-nuclear-power-plant-in-georgia-1501691212, Accessed: 2019-08-29.

[59] "Vogtle in-service dates and cost forecast revised," https://world-nuclear-news.org/Articles/In-service-dates-and-cost-forecast-revised-for-Vog, 2021, Accessed: 2021-09-24.

[60] "Areva's Finland reactor to start in 2019 after another delay," https://www.reuters.com/article/us-finland-nuclear-olkiluoto/arevas-finland-reactor-to-start-in-2019-after-another-delay-idUSKBN1CE1ND, 2017, Accessed: 2019-08-29.

[61] "Finland's Olkiluoto 3 nuclear reactor faces another delay," https://www.reuters.com/business/energy/finlands-olkiluoto-3-nuclear-reactor-faces-another-delay-2021-08-23/, 2021, Accessed: 2021-09-24.

[62] "Nuclear: the Flamanville EPR is experiencing new delays and will not be able to start before 2022," https://www.lemonde.fr/economie/article/2019/06/20/nucleaire-nouveau-retard-confirme-pour-l-epr-de-flamanville_5478853_3234.html, 2019, Accessed: 2019-08-29.

[63] "Edf's levy confirms flamanville epr start-up in 2023," https://www.montelnews.com/news/1194318/edfs-levy-confirms-flamanville-epr-start-up-in-2023, 2021, accessed: 2021-09-14.

[64] Y. Du and J. E. Parsons, "Update on the cost of nuclear power," Center for Energy and Environmental Policy Research (CEEPR) No, pp. 09–004, 2009.

[65] "KHNP granted consruction permit for Shin-Kori 3 and 4," http://www.world-nuclear-news.org/Articles/KHNP-granted-construction-permit-for-Shin-Kori-3-a, 2008, Accessed: 2019-08-29.

[66] L. T. Rush, "Critical assessment of techniques, markets and overall economics of generation iii+ and iv reactors," Ph.D. dissertation, Massachusetts Institute of Technology, 2018.

[67] "Taishan Nuclear Power Plant," https://www.power-technology.com/projects/taishan-nuclear-power-plant/, 2012, Accessed: 2019-08-29.

[68] G. Harris, P. Heptonstall, R. Gross, and D. Handley, "Cost estimates for nuclear power in the uk," Energy Policy, vol. 62, pp. 431–442, 2013.

[69] P. F. Peterson, H. Zhao, and R. Petroski, "Metal and concrete inputs for several nuclear power plants," University of california, berkeley, vol. 4153, pp. 94 720–1730, 2005.

[70] A. Grubler, "The costs of the french nuclear scale-up: A case of negative learning by doing," Energy Policy, vol. 38, no. 9, pp. 5174–5188, 2010.

[71] B. Flyvbjerg, M. K. Skamris Holm, and S. L. Buhl, "How common and how large are cost overruns in transport infrastructure projects?" Transport reviews, vol. 23, no. 1, pp. 71–88, 2003.

[72] M. Cooper, "The economics of nuclear power: Renaissance or relapse?" Nonuclear. se, p. 34, 2009.

[73] B. L. Cohen, "Costs of nuclear power plants-what went wrong," in The Nuclear Energy Option. Springer, 1990, pp. 145–157.

[74] L. W. Davis, "Prospects for nuclear power," Journal of Economic perspectives, vol. 26, no. 1, pp. 49–66, 2012.

[75] A. Litvak, "Westinghouse sold an unfinished product, then the problems snowballed," https://www.post-gazette.com/business/powersource/2017/10/23/westinghouse-sold-an-unfinished-product-then-the-problems-snowballed/stories/201710290008, 2017, Accessed: 2019-09-04.

[76] D. Petti, P. J. Buongiorno, and M. Corradini, "The future of nuclear energy in a carbon-constrained world," 2018.

[77] L. Samalova, O. Chvala, and G. I. Maldonado, "Comparative economic analysis of the integral molten salt reactor and an advanced pwr using the g4-econs methodology," Annals of Nuclear Energy, vol. 99, pp. 258–265, 2017.

[78] C. Andreades and P. Peterson, "Mk1 pebble-bed fluoride-salt-cooled high-temperature reactor capital cost estimation," Transactions of the American Nuclear Society, vol. 113, 2015.

[79] R. Moir, "Cost of electricity from molten salt reactors," Nuclear technology, vol. 138, no. 1, pp. 93–95, 2002.

[80] U. E. I. Agency, "Electric power annual 2017," U.S. Energy Information Agency, Tech. Rep., 2018.

[81] U. B. of Labor Statistics, "May 2018 national industry-specific occupational employment and wage estimates, naics 221113 - nuclear electric power generation," U.S. Bureau of Labor Statistics, Tech. Rep., 2018.

[82] D. G. S. Batkins, P Rossetti, Putting Nuclear Regulatory Costs in Context, https://www.americanactionforum.org/research/putting-nuclear-regulatory-costs-context/, 2017.

[83] J. K. Vaurio, "Safety-related decision making at a nuclear power plant," Nuclear engineering and design, vol. 185, no. 2-3, pp. 335–345, 1998.

[84] A. Mancuso, M. Compare, A. Salo, and E. Zio, "Portfolio optimization of safety measures for reducing risks in nuclear systems," Reliability Engineering & System Safety, vol. 167, pp. 20–29, 2017.

[85] S. Wheatley, B. K. Sovacool, and D. Sornette, "Reassessing the safety of nuclear power," Energy Research & Social Science, vol. 15, pp. 96–100, 2016.

[86] D. T. J. Gaertner, K. Canavan, "Safety and Operational Benefits of Risk-Informed Initiatives," Electric Power Research Institute, Tech. Rep., 2008.

[87] S. Wheatley, B. Sovacool, and D. Sornette, "Of disasters and dragon kings: a statistical analysis of nuclear power incidents and accidents," Risk analysis, vol. 37, no. 1, pp. 99–115, 2017.

[88] U. N. R. Commission, "Probabilistic risk assessment and severe accident evaluation for new reactors," U.S. Nuclear Regulatory Commission, Tech. Rep., 2019.

[89] ——, "Ap1000 probabilistic risk assessment," U.S. Nuclear Regulatory Commission, Tech. Rep., 2003.

[90] B. M. Elsheikh, "Safety assessment of molten salt reactors in comparison with light water reactors," journal of radiation research and applied sciences, vol. 6, no. 2, pp. 63–70, 2013.

[91] D. LeBlanc, "Molten salt reactors: A new beginning for an old idea," Nuclear Engineering and design, vol. 240, no. 6, pp. 1644–1656, 2010.

[92] S. Dingman, S. Camp, J. LaChance, and M. Drouin, "Core damage frequency prespectives for bwr 3/4 and westinghouse 4-loop plants based on ipe results," Sandia National Labs., Albuquerque, NM (United States), Tech. Rep., 1995.

[93] J. Buongiorno, M. Corradini, J. Parsons, and D. Petti, "Nuclear energy in a carbon-constrained world: Big challenges and big opportunities," IEEE Power and Energy Magazine, vol. 17, no. 2, pp. 69–77, 2019.

[94] U. D. N. E. R. A. Committe and the Generation IV International Forum, "A technology roadmap for generation iv nuclear energy systems," US DOE Nuclear Energy Research Advisory Committe and the Generation IV International Forum, Tech. Rep., 2002.

[95] J. Boudreau, "The american breeder reactor program gets a second chance," Los Alamos Science, vol. 2, no. 2, pp. 118–119, 1981.

[96] B. R. Corporation, "Final report, the clinch river breeder reactor plant project," Breeder Reactor Corporation, Tech. Rep., 1985.

[97] M. o. F. Government of Canada, "Budget speech 1955, 1957, 1961 and 1963," Government of Canada, Tech. Rep., 1955 to 1963.

[98] G. L. Brooks, "A short history of the candu nuclear power system," Atomic Energy of Canada Ltd., Tech. Rep., 1993.

[99] J. Reyes, "Overview of the nuscale design and licensing," in Generation IV and Small Modular Reactors. NuScale Power, 2018, p. slide 3.

[100] A. Cho, "Smaller, safer, cheaper: One company aims to reinvent the nuclear reactor and save a warming planet," https://www.science.org/content/article/smaller-safer-cheaper-one-company-aims-reinvent-nuclear-reactor-and-save-warming-planet, Accessed: 2022-08-14.

[101] L. Alexander, W. Carter, C. Craven, D. Janney, T. Kerlin, and R. Van Winkle, "Molten salt converter reactor. design study and power cost estimates for a 1000 mwe station," Oak Ridge National Lab., Tenn., Tech. Rep., 1965.

[102] P. Haubenreich, "Tritium in the msre: Calculated production rates and observed amounts," Oak Ridge National Lab.(ORNL), Oak Ridge, TN (United States), Tech. Rep., 1970.

[103] W. B. Cottrell, H. Hungerford, J. Leslie, and J. L. Meem, "Operation of the aircraft reactor experiment," Oak Ridge National Lab., Tenn., Tech. Rep., 1955.

[104] H. MacPherson, L. Alexander, D. Carrison, J. Estabrook, B. Kinyon, L. Mann, J. Roberts, F. Romie, and F. VonderLage, "A preliminary study of molten salt power reactors," Oak Ridge National Lab., Tenn., Tech. Rep., 1957.

[105] W. Grimes, "Molten-salt reactor chemistry," Nuclear Applications and Technology, vol. 8, no. 2, pp. 137–155, 1970.

[106] "RUBIDIUM," https://pubs.usgs.gov/periodicals/mcs2022/mcs2022-rubidium.pdf, Accessed: 2022-06-04.

[107] D. A. Brown, M. Chadwick, R. Capote, A. Kahler, A. Trkov, M. Herman, A. Sonzogni, Y. Danon, A. Carlson, M. Dunn et al., "Endf/b-viii. 0: The 8th major release of the nuclear reaction data library with cielo-project cross sections, new standards and thermal scattering data," Nuclear Data Sheets, vol. 148, pp. 1–142, 2018.

[108] A. Santamarina, D. Bernard, P. Blaise, M. Coste, A. Courcelle, T. Huynh, C. Jouanne, P. Leconte, O. Litaize, S. Mengelle et al., "The jeff-3.1.1 nuclear data library," JEFF report, vol. 22, no. 10.2, p. 2, 2009.

[109] S. Guo, J. Zhang, W. Wu, and W. Zhou, "Corrosion in the molten fluoride and chloride salts and materials development for nuclear applications," Progress in Materials Science, vol. 97, pp. 448–487, 2018.

[110] W. Haynes, "Density of molten elements and representative salts," CRC Handbook of Chemistry and Physics, 2012.

[111] J. A. Lane, Fluid fuel reactors. Addison-Wesley Pub. Co., 1958, page 571.

[112] M. Rosenthal, "Molten salt reactor program semiannual progress report for period ending august 31, 1968," ORNL, Tech. Rep., 1968.

[113] R. Briggs, "Molten salt reactor program semiannual progress report for period ending december 31, 1965," ORNL, Tech. Rep., 1965.

[114] W. Powers, "Aircraft nuclear propulsion project quarterly progress report for period ending march 10, 1956," Report No. ORNL-2061,(Oak Ridge National Laboratory), 1956.

[115] R. R. Romatoski and L.-W. Hu, "Fluoride salt coolant properties for nuclear reactor applications: A review," Annals of Nuclear Energy, vol. 109, pp. 635–647, 2017.

[116] S. Cohen and T. Jones, "Viscosity measurements on molten fluoride mixtures," Oak Ridge National Lab., Tenn., Tech. Rep., 1957.

[117] H. Rouch, O. Geoffroy, P. Rubiolo, A. Laureau, M. Brovchenko, D. Heuer, and E. Merle-Lucotte, "Preliminary thermal–hydraulic core design of the molten salt fast reactor (msfr)," Annals of Nuclear Energy, vol. 64, pp. 449–456, 2014.

[118] W. Grimes, D. Cuneo, F. Blankenship, G. Keilholtz, H. Poppendick, and M. Robinson, "Chemical aspects of molten fluoride salt reactor fuels," Fluid-Fueled Reactors, p. 569, 1958.

[119] Y. Hirose, "Fast-spectrum fluoride molten salt reactor (ffmsr) with ultimately reduced radiotoxicity of nuclear wastes," IntechOpen, 2020.

[120] C. Forsberg, L. Hu, P. Peterson, M. Fratoni, K. Sridharan, and E. Blandford, "Integrated fhr technology development: tritium management, materials testing, salt chemistry control, thermal-hydraulics and neutronics with associated benchmarking," FHR Annual Report, 2015.

[121] R. Causey, M. Baskes, and K. Wilson, "The retention of deuterium and tritium in poco axf-5q graphite," Journal of Vacuum Science & Technology A: Vacuum, Surfaces, and Films, vol. 4, no. 3, pp. 1189–1192, 1986.

[122] T. Ault, K. Brozek, L. Fan, M. Folsom, J. Kim, and J. Zeismer, "Lithium isotope enrichment: feasible domestic enrichment alternatives," Department of Nuclear Engineering, University of California, Berkeley, Tech. Rep. UCBTH-12-005, 2012.

[123] C. N. D. Center, "Cendl-2, the chinese evaluated nuclear data library for neutron reaction data," Report IAEA-NDS-61, Rev. 3, Tech. Rep., 1996.

[124] J. H. Shaffer, "Preparation and handling of salt mixtures for the molten salt reactor experiment," Oak Ridge National Lab., Tech. Rep., 1971.

[125] T. B. Rhinehammer and L. J. Wittenberg, "Evaluation of fuel resources and requirements for the magnetic fusion energy program," Mound Lab., Miamisburg, OH (USA), Tech. Rep., 1978.

[126] U. G. Survey, "Beryllium," https://minerals.usgs.gov/minerals/pubs/commodity/beryllium/mcs-2018-beryl.pdf, US Geological Survey, Tech. Rep., 2018.

[127] M. Smolik, A. Jakóbik-Kolon, and M. Porański, "Separation of zirconium and hafnium using diphonix® chelating ion-exchange resin," Hydrometallurgy, vol. 95, no. 3-4, pp. 350–353, 2009.

[128] E. Collins and G. Del Cul, "Recycle of zirconium from used nuclear fuel cladding: A major element of waste reduction-11336."

[129] D. Holcomb, G. Flanagan, B. Patton, J. Gehin, R. Howard, and T. Harrison, "Fast spectrum molten salt reactor options," ORNL/TM-2011/105, 2011.

[130] ENDF, "Endf/b-vii summary documentation," National Nuclear Data Center, Brookhaven National Laboratory, Tech. Rep., 2011.

[131] L. Mei, X. Cai, D. Jiang, J. Chen, Y. Zhu, Y. Liu, and X. Wang, "The investigation of thermal neutron scattering data for molten salt flibe," Journal of Nuclear Science and Technology, vol. 50, no. 7, pp. 682–688, 2013.

[132] L. Mei, X. Cai, D. Jiang, J. Chen, W. Guo, and W. Xiong, "Investigation of thermal neutron scattering data for bef2 and lif crystals," Journal of Nuclear Science and Technology, vol. 50, no. 4, pp. 419–424, 2013.

[133] J. C. Helton, J. D. Johnson, C. J. Sallaberry, and C. B. Storlie, "Survey of sampling-based methods for uncertainty and sensitivity analysis," Reliability Engineering & System Safety, vol. 91, no. 10-11, pp. 1175–1209, 2006.

[134] J. C. Helton, "Uncertainty and sensitivity analysis techniques for use in performance assessment for radioactive waste disposal," Reliability Engineering & System Safety, vol. 42, no. 2-3, pp. 327–367, 1993.

[135] M. Buyukada, "Probabilistic uncertainty analysis based on monte carlo simulations of co-combustion of hazelnut hull and coal blends: Data-driven modeling and response surface optimization," Bioresource technology, vol. 225, pp. 106–112, 2017.

[136] S. K. Seaholm, E. Ackerman, and S.-C. Wu, "Latin hypercube sampling and the sensitivity analysis of a monte carlo epidemic model," International journal of bio-medical computing, vol. 23, no. 1-2, pp. 97–112, 1988.

[137] R. L. Iman and J. C. Helton, "An investigation of uncertainty and sensitivity analysis techniques for computer models," Risk analysis, vol. 8, no. 1, pp. 71–90, 1988.

[138] J. P. Kleijnen, A. Ridder, and R. Rubinstein, "Variance reduction techniques in monte carlo methods," 2010.

[139] P. Glasserman, "Variance reduction techniques," in Monte Carlo Methods in Financial Engineering. Springer, 2004, pp. 185–279.

[140] H. Yang and S. An, "Variance reduction techniques for basket option," in 2010 International Conference on Management and Service Science. IEEE, 2010, pp. 1–4.

[141] E. Saliby and R. Paul, "A farewell to the use of antithetic variates in monte carlo simulation," Journal of the Operational Research Society, vol. 60, no. 7, pp. 1026–1035, 2009.

[142] A. Olsson, G. Sandberg, and O. Dahlblom, "On latin hypercube sampling for structural reliability analysis," Structural safety, vol. 25, no. 1, pp. 47–68, 2003.

[143] F. A. Viana, "Things you wanted to know about the latin hypercube design and were afraid to ask," in 10th World Congress on Structural and Multidisciplinary Optimization, vol. 19, no. 24.05. sn, 2013.

[144] D. M. Steinberg and D. K. Lin, "A construction method for orthogonal latin hypercube designs," Biometrika, vol. 93, no. 2, pp. 279–288, 2006.

[145] J. C. Helton and F. J. Davis, "Latin hypercube sampling and the propagation of uncertainty in analyses of complex systems," Reliability Engineering & System Safety, vol. 81, no. 1, pp. 23–69, 2003.

[146] Various, "Commodity price variability: It's nature and it's causes," Organization for Economic Co-operation and Development, Tech. Rep., 1993.

[147] B. Lahcene, "On pearson families of distributions and its applications," African Journal of Mathematics and Computer Science Research, vol. 6, no. 5, pp. 108–117, 2013.

[148] F. M. Longin, "The asymptotic distribution of extreme stock market returns," Journal of business, pp. 383–408, 1996.

[149] P. E. Love, C. Sing, X. Wang, D. J. Edwards, and H. Odeyinka, "Probability distribution fitting of schedule overruns in construction projects," Journal of the Operational Research Society, vol. 64, no. 8, pp. 1231–1247, 2013.

[150] C. J. Sallaberry, J. C. Helton, and S. C. Hora, "Extension of latin hypercube samples with correlated variables," Reliability Engineering & System Safety, vol. 93, no. 7, pp. 1047–1059, 2008.

[151] H. Yu, C. Chung, K. Wong, H. Lee, and J. Zhang, "Probabilistic load flow evaluation with hybrid latin hypercube sampling and cholesky decomposition," IEEE Transactions on Power Systems, vol. 24, no. 2, pp. 661–667, 2009.

[152] R. L. Iman and W.-J. Conover, "A distribution-free approach to inducing rank correlation among input variables," Communications in Statistics-Simulation and Computation, vol. 11, no. 3, pp. 311–334, 1982.

[153] M. N. Williams, C. A. G. Grajales, and D. Kurkiewicz, "Assumptions of multiple regression: Correcting two misconceptions," Practical Assessment, Research, and Evaluation, vol. 18, no. 1, p. 11, 2013.

[154] C. S. Parke, "Module 7: evaluating model assumptions for multiple regression analysis," Essential first steps to data analysis: Scenario-based examples using SPSS, pp. 147–178, 2013.

[155] S. Sheather, A modern approach to regression with R. Springer Science & Business Media, 2009.

[156] D. Ingersoll, C. Forsberg, L. Ott, D. Williams, J. Renier, D. Wilson, S. Ball, L. Reid, W. Corwin, G. Del Cul et al., Status of Preconceptual Design of the Advanced High-Temperature Reactor (AHTR). United States. Department of Energy, 2004.

[157] Thorium Molten Salt Reactors (TMSR) Development in China, Institut of Advanced Nuclear Energy. Technical Meeting on the Status of Molten Salt Reactor Technology, IAEA, 2016.

[158] Thorium Molten Salt Reactors (TMSR) Development in China, Shanghai Institute of Applied Physics (SINAP). PSI, Switzerland, 2017.

[159] O. Beneš and R. Konings, "Thermodynamic study of lif–bef2–zrf4–uf4 system," Journal of alloys and compounds, vol. 452, no. 1, pp. 110–115, 2008.

[160] S. Saito, T. Tanaka, and Y. Sudo, "Design of high temperature engineering test reactor (httr)," Japan Atomic Energy Research Inst., Tech. Rep., 1994.

[161] D. McEligot, W. D. Swank, D. L. Cottle, and F. I. Valentin, "Thermal properties of g-348 graphite," Idaho National Lab.(INL), Idaho Falls, ID (United States), Tech. Rep., 2016.

[162] T. L. Bergman, F. P. Incropera, D. P. DeWitt, and A. S. Lavine, Fundamentals of heat and mass transfer. John Wiley & Sons, 2011.

[163] J. Leppänen, M. Pusa, T. Viitanen, V. Valtavirta, and T. Kaltiaisenaho, "The serpent monte carlo code: Status, development and applications in 2013," Annals of Nuclear Energy, vol. 82, pp. 142–150, 2015.

[164] F. A. Strantz, "Density determination of tristructural-isotropic nuclear fuel using multiple projection x-ray radiography," 2011.

[165] P. A. Demkowicz, "Triso fuel: design, manufacturing, and performance," Idaho National Lab.(INL), Idaho Falls, ID (United States), Tech. Rep., 2019.

[166] A. T. Cisneros, R. O. Scarlat, M. R. Laufer, E. Greenspan, and P. F. Peterson, "Pebble fuel design for the pb-fhr," in Proceedings of the International Congress on Advances in Nuclear Power Plants (ICAPP'12), 2012.

[167] "X-energy is Developing a Pebble Bed Reactor That They Say Can't Melt Down," https://www.energy.gov/ne/articles/x-energy-developing-pebble-bed-reactor-they-say-cant-melt-down, Accessed: 2023-02-16.

[168] Y. Jeong, J. Park, H. C. Lee, and D. Lee, "Equilibrium core design methods for molten salt breeder reactor based on two-cell model," Journal of Nuclear Science and Technology, vol. 53, no. 4, pp. 529–536, 2016. [Online]. Available: https://doi.org/10.1080/00223131.2015.1062812

[169] K. Mitachi, T. Yamamoto, and R. Yoshioka, "Self-sustaining core design for 200 mwe molten-salt reactor with thorium-uranium fuel: Fuji-u3-(0)," Proceedings of TU2007, pp. 4–6, 2007.

[170] D. F. Williams, "Additional physical property measurements and assessment of salt compositions proposed for the intermediate heat transfer loop," ORNL/GEN4/LTR-06-033, 2006.

[171] R. Anderl, D. Petti, and G. Smolik, "Filbe molten salt research for tritium breeder applications," Tech. Rep., 2004.

[172] G. Smolik, M. Simpson, P. Pinhero, M. Hara, Y. Hatano, R. Anderl, J. Sharpe, T. Terai, S. Tanaka, D. Petti et al., "Beryllium interactions in molten flibe," Proc. ICFRM-12, 2006.

[173] P. Yvon, Structural materials for generation IV nuclear reactors. Woodhead publishing, 2016.

[174] R. Thoma, "phase diagrams of nuclear reactor materials," ORNL, Tech. Rep., 1959.

[175] E. Capelli, O. Beneš, and R. Konings, "Thermodynamic assessment of the lif–naf–bef2–thf4–uf4 system," Journal of Nuclear Materials, vol. 449, no. 1-3, pp. 111–121, 2014.

[176] O. Beneš and R. Konings, "Thermodynamic calculations of molten-salt reactor fuel systems," in Molten Salts Chemistry. Elsevier, 2013, pp. 49–78.

[177] I. Cool, "Molten salt coolants for high temperature reactors," 2009.

[178] J. Ocádiz-Flores, A. Gheribi, J. Vlieland, K. Dardenne, J. Rothe, R. Konings, and A. Smithr, "New insights and coupled modelling of the structural and thermodynamic properties of the lif-uf4 system," Journal of Molecular Liquids, p. 115820, 2021.

[179] R. Thoma, H. INSLEY, B. Landau, H. Friedman, and W. Grimes, "Phase equilibria in the alkali fluoride-uranium tetrafluoride fused salt systems: Iii, the system naf–lif–uf4," Journal of the American Ceramic Society, vol. 42, no. 1, pp. 21–26, 1959.

[180] J. McWherter, "Molten salt breeder experiment design bases," Oak Ridge National Lab., Tech. Rep., 1970.

[181] P. Kasten, E. Bettis, W. Cook, W. Eatherly, D. Holmes, R. Kedl, C. Kennedy, S. Kirslis, H. McCoy, A. Perry et al., "Graphite behavior and its effects on msbr performance," Nuclear Engineering and Design, vol. 9, no. 2, pp. 157–195, 1969.

[182] ENDF, "Endf/b-viii summary documentation," National Nuclear Data Center, Brookhaven National Laboratory, Tech. Rep., 2018.

[183] many, "Delayed neutron data for the major actinides," NEA Nuclear Science Committee, OECD, Tech. Rep. Volume 6, 2002.

[184] "Average number of neutrons emitted per fission," https://www-nds.iaea.org/sgnucdat/a6.htm, International Atomic Energy Agency, Accessed: 2018-06-24.

[185] P. Haubenreich, "Prediction of effective yields of delayed neutrons in msre," ORNL-TM-038, USA: ORNL, Tech. Rep., 1962.

[186] M. Aufiero, M. Brovchenko, A. Cammi, I. Clifford, O. Geoffroy, D. Heuer, A. Laureau, M. Losa, L. Luzzi, E. Merle-Lucotte et al., "Calculating the effective delayed neutron fraction in the molten salt fast reactor: analytical, deterministic and monte carlo approaches," Annals of Nuclear Energy, vol. 65, pp. 78–90, 2014.

[187] D. Moser, A. Wheeler, and O. Chvála, "Lattice optimization for graphite moderated molten salt reactors using low-enriched uranium fuel," Annals of Nuclear Energy, vol. 110, pp. 1–10, 2017.

[188] K. Nagy, J. Kloosterman, D. Lathouwers, and T. Van der Hagen, "The effects of core zoning on the graphite lifespan and breeding gain of a moderated molten salt reactor," Annals of Nuclear Energy, vol. 43, pp. 19–25, 2012.

[189] "IMSR Technology. Terrestrial Energy's IMSR: Safe, clean, low-cost and high-impact," https://www.terrestrialenergy.com/technology/, Terrestrial Energy, Accessed: 2018-06-29.

[190] G. J. Janz, Molten salts handbook. Elsevier, 2013.

[191] M. Blander, Molten salt chemistry. Interscience Publishers, 1964.

[192] V. Khokhlov, V. Ignatiev, and V. Afonichkin, "Evaluating physical properties of molten salt reactor fluoride mixtures," Journal of fluorine chemistry, vol. 130, no. 1, pp. 30–37, 2009.

[193] G. Watson, R. Evans III, W. Grimes, and N. Smith, "Solubility of noble gases in molten fluorides. in lif-bef2." Journal of Chemical and Engineering Data, vol. 7, no. 2, pp. 285–287, 1962.

[194] M. S. R. Tomkins, N Bansal, "Solubility data series: Gases in molten salts," 1979.

[195] E. G. Ward, Watson, "Report ornl 2931," Report ORNL 2931, pp. 29–31, 1960.

[196] W. Grimes, N. Smith, and G. Watson, "Solubility of noble gases in molten fluorides. i. in mixtures of naf–zrf4 (53–47 mole%) and naf–zrf4–uf4 (50–46–4 mole%)," The Journal of Physical Chemistry, vol. 62, no. 7, pp. 862–866, 1958.

[197] V. Ignat'ev, A. Merzlyakov, V. Subbotin, A. Panov, and Y. V. Golovatov, "Experimental investigation of the physical properties of salt melts containing sodium and lithium fluorides and beryllium difluoride," Atomic Energy, vol. 101, no. 5, pp. 822–829, 2006.

[198] Y. Hirose and Y. Takashima, "The concept of fuel cycle integrated molten salt reactor for transmuting pu+ ma from spent lwr fuels," 2001.

[199] C. N. A. C. Z. Bahri, W. M. Al-Areqi, M. F. M. Ruf, and A. A. Majid, "Characteristic of molten fluoride salt system lif-bef2 (flibe) and lif-naf-kf (flinak) as coolant and fuel carrier in molten salt reactor (msr)," in AIP Conference Proceedings, vol. 1799, no. 1. AIP Publishing LLC, 2017, p. 040008.

[200] V. Sobolev, "Database of thermophysical properties of liquid metal coolants for gen-iv," 2011.

[201] C. Fiorina, D. Lathouwers, M. Aufiero, A. Cammi, C. Guerrieri, J. L. Kloosterman, L. Luzzi, and M. E. Ricotti, "Modelling and analysis of the msfr transient behaviour," Annals of Nuclear Energy, vol. 64, pp. 485–498, 2014.

[202] E. Merle-Lucotte, D. Heuer, M. Allibert, M. Brovchenko, N. Capellan, and V. Ghetta, "Launching the thorium fuel cycle with the molten salt fast reactor," in Proceedings of ICAPP, 2011, pp. 2–5.

[203] G. J. Janz and R. Tomkins, "Physical properties data compilation relevant to energy storage. iv.–molten salts: Data on additional single and multicomponent salt systems," National Bureau of Standards, v+ 861, 26 x 20 cm(, 1981, 1981.

[204] G. J. Janz, R. Tomkins, C. Allen, J. Downey Jr, G. Garner, U. Krebs, and S. K. Singer, "Molten salts: Volume 4, part 2, chlorides and mixtures—electrical conductance, density, viscosity, and surface tension data," Journal of Physical and Chemical Reference Data, vol. 4, no. 4, pp. 871–1178, 1975.

[205] P. Nelson, D. Butler, M. Chasanov, and D. Meneghetti, "Fuel properties and nuclear performance of fast reactors fueled with molten chlorides," Nuclear Applications, vol. 3, no. 9, pp. 540–547, 1967.

[206] A. V. Grosse, "Viscosities of liquid sodium and potassium, from their melting points to their critical points," Science, vol. 147, no. 3664, pp. 1438–1441, 1965.

[207] F. Kanda and R. Colburn, "The absolute viscosity of some lead-tin alloys," Physics and Chemistry of Liquids, vol. 1, no. 2, pp. 159–170, 1968.

[208] D. Williams et al., "Assessment of candidate molten salt coolants for the ngnp/nhi heat-transfer loop," ORNL/TM-2006/69, Oak Ridge National Laboratory, Oak Ridge, Tennessee, 2006.

[209] D. F. Williams, G. D. D. Cul, L. M. Toth, and E. D. Collins, "The influence of lewis acid/base chemistry on the removal of gallium by volatility from weapons-grade plutonium dissolved in molten chlorides," Nuclear Technology, vol. 136, no. 3, pp. 367–370, 2001.

[210] M. Taube and J. Ligou, "Molten chlorides fast breeder reactor problems and possibilities," Eidgenoessisches Inst. fuer Reaktorforschung, Tech. Rep., 1972.

[211] M. Taube, "Fast reactors using molten chloride salts as fuel," INFCE (Switzerland), Tech. Rep., 1978.

[212] O. Beneš and R. Konings, "Thermodynamic properties and phase diagrams of fluoride salts for nuclear applications," Journal of Fluorine Chemistry, vol. 130, no. 1, pp. 22–29, 2009.

[213] K. A. Sense, R. W. Stone, and R. Filbert Jr, "Vapor pressure and equilibrium studies of the sodium fluoride-beryllium fluoride system," Battelle Memorial Inst., Columbus, Ohio, Tech. Rep., 1957.

[214] A. Tosolin, O. Beneš, J.-Y. Colle, P. Souček, L. Luzzi, and R. Konings, "Vaporization behaviour of the molten salt fast reactor fuel: The lif-thf 4-uf 4 system," Journal of Nuclear Materials, 2018.

[215] R. C. Robertson, "Conceptual design study of a single-fluid molten-salt breeder reactor." comp.; Oak Ridge National Lab., Tenn., Tech. Rep., 1971.

[216] J. Barton and E. Bloom, "A boiling point method for determination of vapor pressures of molten salts," The Journal of Physical Chemistry, vol. 60, no. 10, pp. 1413–1416, 1956.

[217] W. Grimes, "Chemical research and development for molten-salt breeder reactors." Oak Ridge National Lab., Tenn., Tech. Rep., 1967.

[218] U. Kose, U. Koç, L. ERBAY, E. Ogut, and H. Ayhan, "Heat exchanger design studies for molten salt fast reactor," EPJ Nuclear Sciences & Technologies, vol. 5, 2019.

[219] J. Fink and L. Leibowitz, "Thermodynamic and transport properties of sodium liquid and vapor," Argonne National Lab., Tech. Rep., 1995.

[220] O. Beneš and R. Konings, "Molten salt reactor fuel and coolant," 2012.

[221] D. J. Rogers, T. Yoko, and G. J. Janz, "Fusion properties and heat capacities of the eutectic lithium fluoride-sodium fluoride-potassium fluoride melt," Journal of Chemical and Engineering Data, vol. 27, no. 3, pp. 366–367, 1982.

[222] W. Powers, "Molten salt reactor quarterly progress report for period ending oct 31, 1958," Report No. ORNL-2626,(Oak Ridge National Laboratory), 1966.

[223] N. Araki and Y. Kato, "Measurements of thermophysical properties of molten fluoride salt containing thorium," in Research on thorium fuel, 1987.

[224] A. A. Redkin, Y. P. Zaikov, I. V. Korzun, O. G. Reznitskikh, T. V. Yaroslavtseva, and S. I. Kumkov, "Heat capacity of molten halides," The Journal of Physical Chemistry B, vol. 119, no. 2, pp. 509–512, 2015.

[225] V. Glushko, L. Gurvich, G. Bergman, I. Veitz, V. Medvedev, G. Khachkuruzov, and V. Yungman, "Thermodynamic properties of individual substances vol iv," High-Temperature Institute, State Institute of Applied Chemistry, National Academy of Sciences of the USSR, Moscow, 1982.

[226] Y. Nagasaka, N. Nakazawa, and A. Nagashima, "Experimental determination of the thermal diffusivity of molten alkali halides by the forced rayleigh scattering method. i. molten licl, nacl, kcl, rbcl, and cscl," International Journal of Thermophysics, vol. 13, no. 4, pp. 555–574, 1992.

[227] J. Bockris and N. Richards, "The compressibilities, free volumes and equation of state for molten electrolytes: some alkali halides and nitrates," Proc. R. Soc. Lond. A, vol. 241, no. 1224, pp. 44–66, 1957.

[228] I. Murgulescu and C. Telea, "Heat-capacities in molten alkaline halides," REVUE ROUMAINE DE CHIMIE, vol. 22, no. 5, pp. 683–689, 1977.

[229] N. D. G. W. D. Powers, S. I. Cohen, Reactor Handbook; Engineering. Atomic Energy Commission, 1963.

[230]

[231] Chembase, "Plutonium (3+) ion trifluoride, chem. id. 129029," 2012-2014, data retrieved from,http://en.chembase.cn/molecule-129029.html.

[232] R. D. Shannon, "Revised effective ionic radii and systematic studies of interatomic distances in halides and chalcogenides," Acta crystallographica section A: crystal physics, diffraction, theoretical and general crystallography, vol. 32, no. 5, pp. 751–767, 1976.

[233] R. D. Mariani and D. Vaden, "Modeled salt density for nuclear material estimation in the treatment of spent nuclear fuel," Journal of Nuclear Materials, vol. 404, no. 1, pp. 25–32, 2010.

[234] M. Rosenthal, "The development status of molten-salt breeder reactors," ORNL-4812, 1972.

[235] K. Cornwell, "The thermal conductivity of molten salts," Journal of Physics D: Applied Physics, vol. 4, no. 3, p. 441, 1971.

[236] M. Smirnov, V. Khokhlov, and E. Filatov, "Thermal conductivity of molten alkali halides and their mixtures," Electrochimica acta, vol. 32, no. 7, pp. 1019–1026, 1987.

[237] W. Gambill, "Fused salt thermal conductivity," Chem. Eng., vol. 66, 1959.

[238] C. Raseman, H. Susskind, G. Farber, W. McNulty, and F. Salzano, "Engineering experience at brookhaven national laboratory in handling fused chloride salts," Brookhaven National Lab., Upton, NY, Tech. Rep., 1960.

[239] D. Williams, L. Toth, and L. UT-Battelle, "Chemical considerations for the selection of the coolant for the advanced high-temperature reactor," ORNL/GEN4/LTR-05-011, Oak Ridge National Laboratory, Oak Ridge, TN, 2005.

[240] W. Manly, J. Coobs, J. DeVan, D. Douglas, H. Inouye, P. Patriarca, T. Roche, and J. Scott, "Metallurgical problems in molten fluoride systems," Oak Ridge National Lab., Tenn., Tech. Rep., 1958.

[241] T. L. Brown, H. E. LeMay Jr, B. E. Bursten, and C. J. Murphy, Chemistry: The central science. Southeast, 1978.

[242] S. Lower, "Free energy and equilibrium," https://chem.libretexts.org/Bookshelves/Physical_and_Theoretical_Chemistry_Textbook_Maps/Supplemental_Modules_(Physical_and_Theoretical_Chemistry)/Thermodynamics/Chemical_Energetics/Free_Energy_and_Equilibrium, Accessed: 2022-06-30.

[243] M. Anderson, K. Sridharan, D. Morgan, P. Peterson, P. Calderoni, R. Scheele, A. Casekka, and B. McNamara, "Heat transfer salts for nuclear reactor systems-chemistry control, corrosion mitigation, and modeling," Univ. of Wisconsin, Tech. Rep., 2015.

[244] D. Olander, "Redox condition in molten fluoride salts: Definition and control," Journal of Nuclear Materials, vol. 300, no. 2-3, pp. 270–272, 2002.

[245] K. Sridharan and T. Allen, "Corrosion in molten salts," in Molten Salts Chemistry. Elsevier, 2013, pp. 241–267.

[246] D. Wilson, "Corrosion issues in molten fluoride salts," in ANS Annual Conference, 2006.

[247] L. Olson, "Materials corrosion in molten lif-naf-kf eutectic salt, nuclear engineering," 2009.

[248] C. Baes Jr, "The chemistry and thermodynamics of molten salt reactor fuels," Journal of Nuclear Materials, vol. 51, no. 1, pp. 149–162, 1974.

[249] H. McCoy and B. McNabb, "Intergranular cracking of inor-8 in the msre." Oak Ridge National Lab., Tenn., Tech. Rep., 1972.

[250] G. Zheng, D. Carpenter, L.-W. Hu, and K. Sridharan, "High temperature corrosion of structural alloys in molten li2bef4 (flibe) salt," Advances in Materials Science for Environmental and Energy Technologies V: Ceramic Transactions, Volume 260, vol. 260, pp. 93–101, 2016.

[251] G. Zheng and K. Sridharan, "Corrosion of structural alloys in high-temperature molten fluoride salts for applications in molten salt reactors," JOM, vol. 70, no. 8, pp. 1535–1541, 2018.

[252] Victor Ignatiev, "Materials and metals in MSR." National Research Center "Kurchatov Institute", 2017, http://samofar.eu/wp-content/uploads/2017/07/Ignatiev_MSR-materials.pdf.

[253] W. Manley et al., "Construction materials for molten-salt reactors," Fluid-Fueled Reactors, pp. 595–603, 1958.

[254] D. Williams, D. Wilson, L. Toth, J. Caja, and J. Keiser, "Research on molten fluorides as high temperature heat transfer agents," in Proceedings Global 2003, Embedded Topical in 2003 American Nuclear Society Winter Meeting, 2003, pp. 16–20.

[255] G. Zheng, B. Kelleher, G. Cao, M. Anderson, T. Allen, and K. Sridharan, "Corrosion of 316 stainless steel in high temperature molten li 2 bef 4 (flibe) salt," Journal of Nuclear Materials, vol. 461, pp. 143–150, 2015.

[256] J. Keiser, "Compatibility studies of potential molten-salt breeder reactor materials in molten fluoride salts," Oak Ridge National Lab., Tech. Rep., 1977.

[257] J. Koger and A. Litman, "Compatibility of molybdenum-base alloy tzm with lif–bef2–thf4–uf4 (68–20–11.7–0.3 mole%) at 1100 c," ORNL-TM-2724, Oak Ridge National Laboratory, 1969.

[258] S. Delpech, E. Merle-Lucotte, D. Heuer, M. Allibert, V. Ghetta, C. Le-Brun, X. Doligez, and G. Picard, "Reactor physic and reprocessing scheme for innovative molten salt reactor system," Journal of fluorine chemistry, vol. 130, no. 1, pp. 11–17, 2009.

[259] "David LeBlanc on IMSR fuel and waste scenario's," http://www.daretothink.org/david-leblanc-on-imsr-fuel-and-waste-scenarios/, Accessed: 2018-06-24.

[260] J.-L. Do, D. Tan, and T. Friscic, "Oxidative mechanochemistry for direct, room-temperature, solvent-free conversion of palladium and gold metals into soluble salts and coordination complexes," Angewandte Chemie, 2018.

[261] M. Aufiero, A. Cammi, C. Fiorina, J. Leppänen, L. Luzzi, and M. Ricotti, "An extended version of the serpent-2 code to investigate fuel burn-up and core material evolution of the molten salt fast reactor," Journal of Nuclear Materials, vol. 441, no. 1-3, pp. 473–486, 2013.

[262] C. W. Forsberg, S. Lam, D. M. Carpenter, D. G. Whyte, R. Scarlat, C. Contescu, L. Wei, J. Stempien, and E. Blandford, "Tritium control and capture in salt-cooled fission and fusion reactors: status, challenges, and path forward," Nuclear Technology, vol. 197, no. 2, pp. 119–139, 2017.

Appendix A

Additional Results

A.1 Additional FHR results

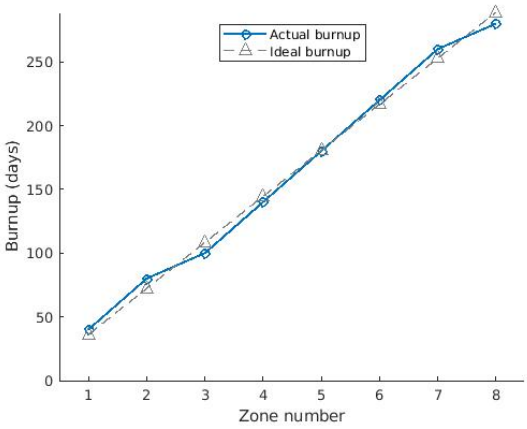

(a) FLiBe PF 25

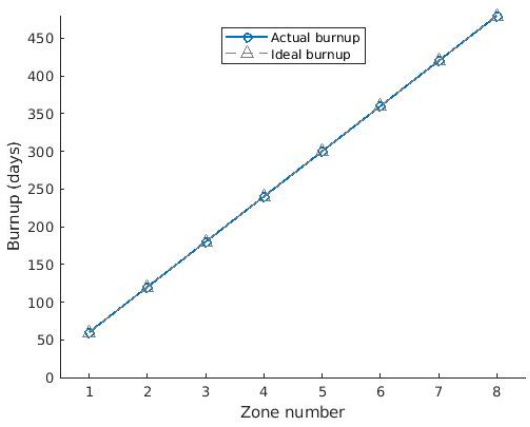

(b) FLiBe PF 45

Figure A.1: FLiBe salt: Burnup distribution among the 8 zones in the refuelling model. The simulated (actual) is compared to the ideal burnup at equilibrium. The average burnup is equal in each case.

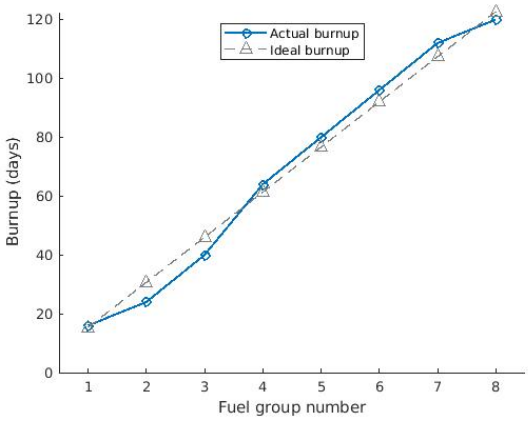

(a) NaBeF PF 30

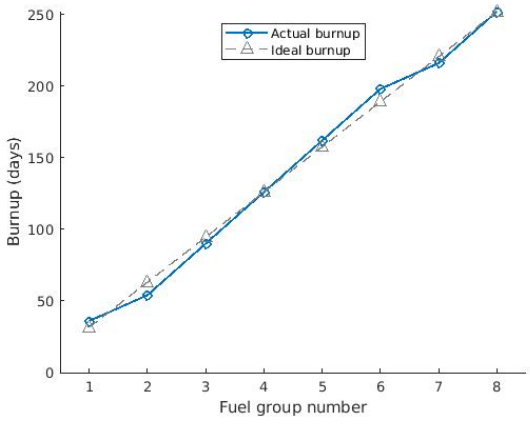

(b) NaBeF PF 50

Figure A.2: NaBeF salt: Burnup distribution among the 8 zones in the refuelling model. The simulated (actual) is compared to the ideal burnup at equilibrium. The average burnup is equal in each case.

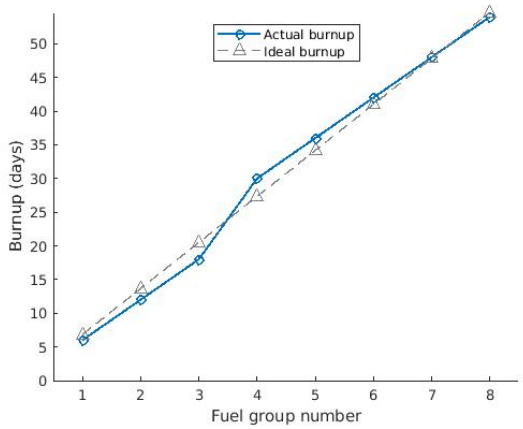

(a) NaZrF PF 30

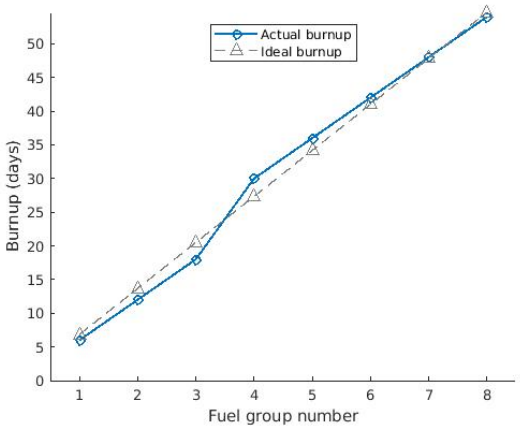

(b) NaZrF PF 50

Figure A.3: NaZrF salt: Burnup distribution among the 8 zones in the refuelling model. The simulated (actual) is compared to the ideal burnup at equilibrium. The average burnup is equal in each case.

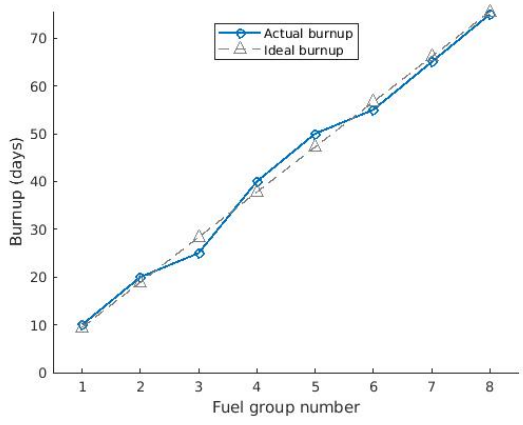

(a) LiNaBeF PF 30

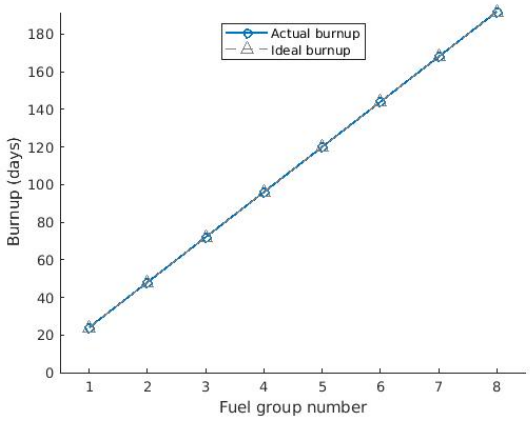

(b) LiNaBeF PF 50

Figure A.4: LiNaBeF salt: Burnup distribution among the 8 zones in the refuelling model. The simulated (actual) is compared to the ideal burnup at equilibrium. The average burnup is equal in each case.

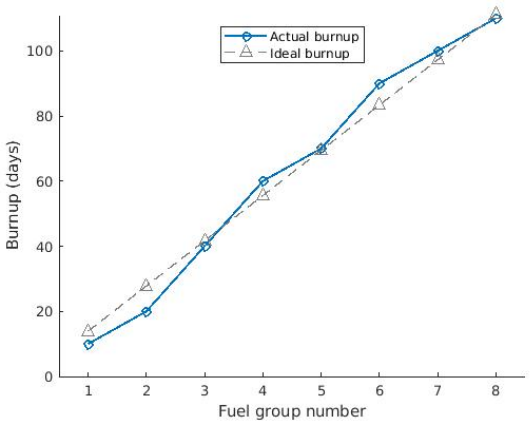

(a) LiNaZrF PF 30

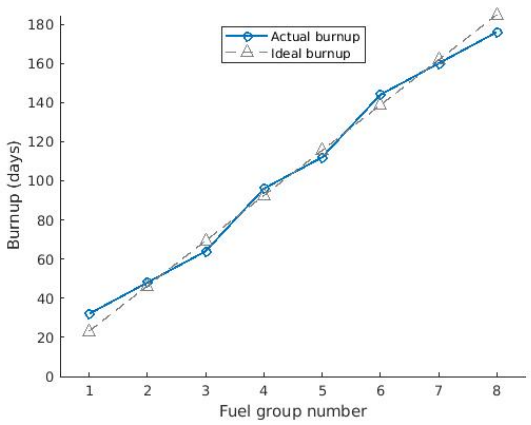

(b) LiNaZrF PF 50

Figure A.5: LiNaZrF salt: Burnup distribution among the 8 zones in the refuelling model. The simulated (actual) is compared to the ideal burnup at equilibrium. The average burnup is equal in each case.

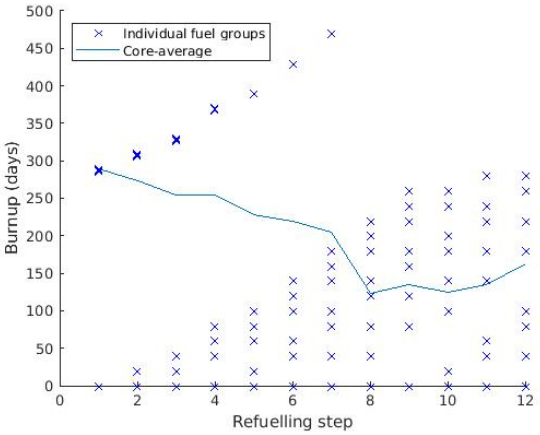

(a) FLiBe PF 25

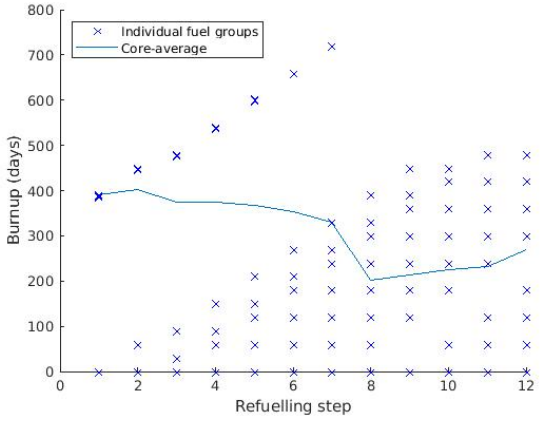

(b) FLiBe PF 45

Figure A.6: FLiBe: Burnup of each fuel group as the simulation progresses.

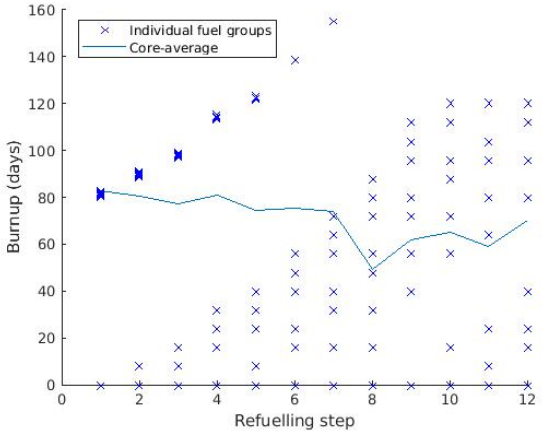

(a) NaBeF PF 30

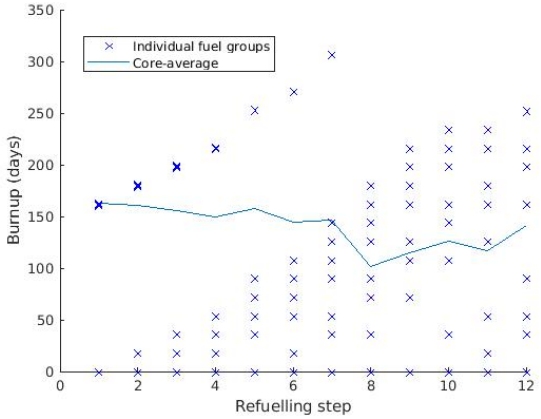

(b) NaBeF PF 50

Figure A.7: NaBeF: Burnup of each fuel group as the simulation progresses.

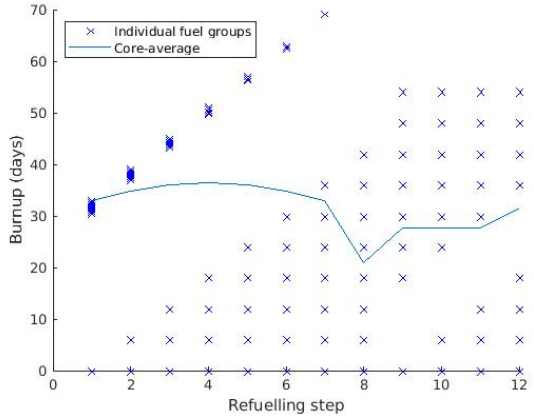

(a) NaZrF PF 30

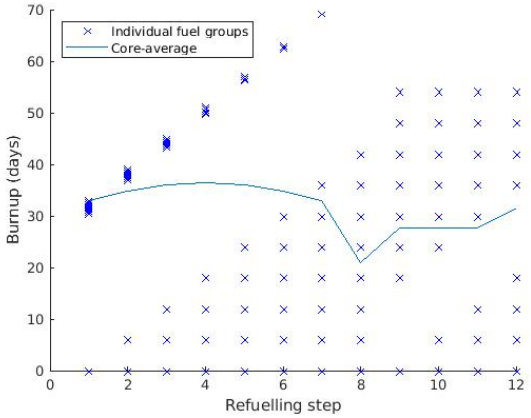

(b) NaZrF PF 50

Figure A.8: NaZrF: Burnup of each fuel group as the simulation progresses.

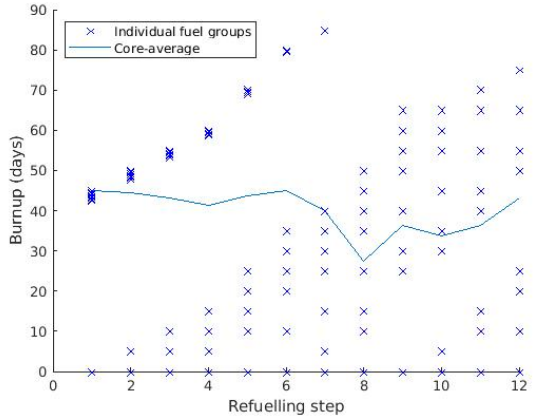

(a) LiNaBeF PF 30

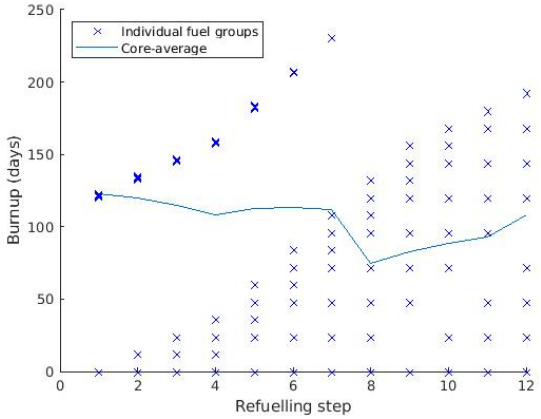

(b) LiNaBeF PF 50

Figure A.9: LiNaBeF: Burnup of each fuel group as the simulation progresses.

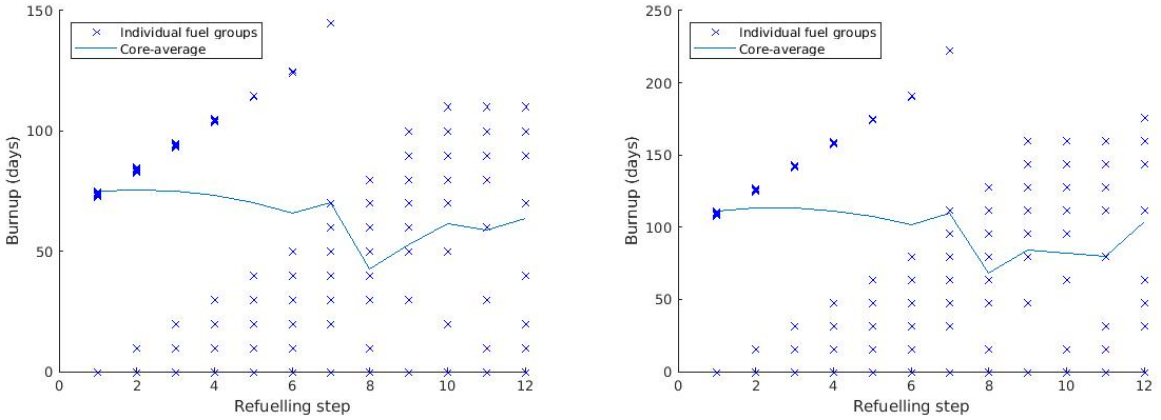

(a) LiNaZrF PF 30

(b) LiNaZrF PF 50

Figure A.10: LiNaZrF: Burnup of each fuel group as the simulation progresses.

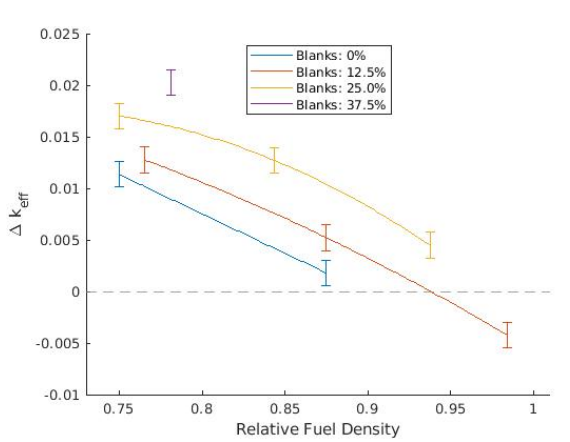

(a) FLiBe

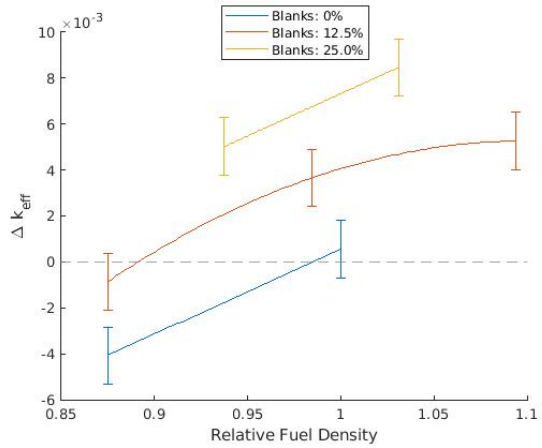

(b) NaBeF

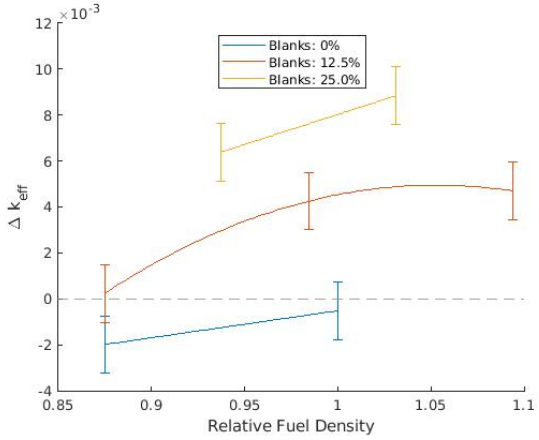

(c) NaZrF

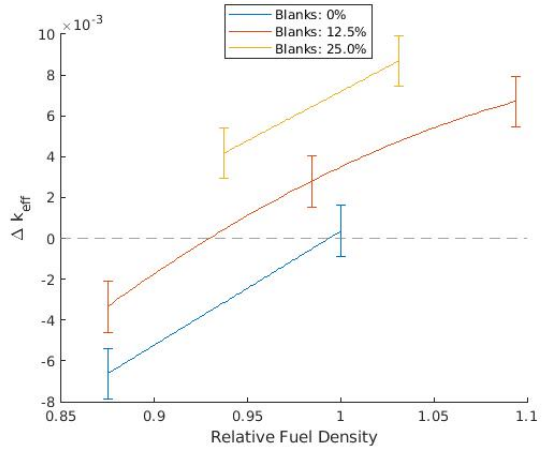

(d) LiNaBeF

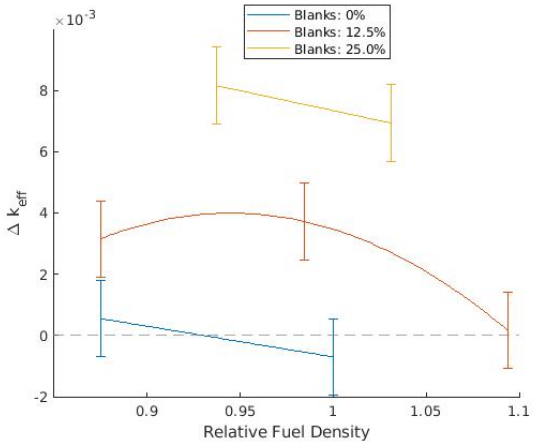

(e) LiNaZrF

Figure A.11: k_{eff} as a function of fuel density with varying shares of blank pebbles, fuel density is relative to 0% blanks and a PF of 40. All data points shown for each salt only include those generated with an alike fuel composition. The blank share used to generate all fuel compositions is 0%. FLiBe uses a fuel composition from a PF of 35, the other salts from a PF of 40.

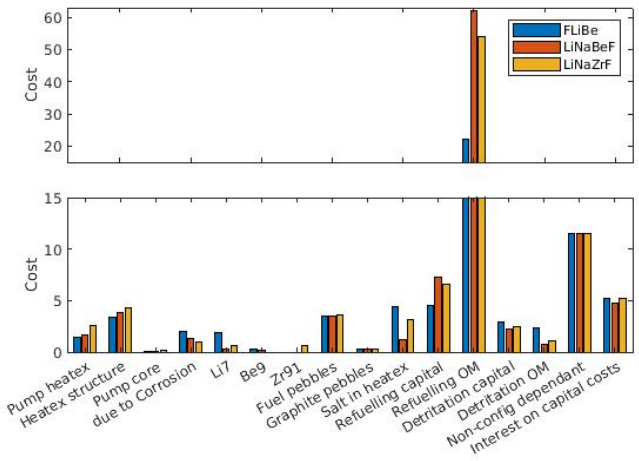

Figure A.12: Breakdown of the 'Direct cost' category of cost into components. The cost values are given at graphite pebble, PF and inner radius combination that results in the lowest feasible total cost.

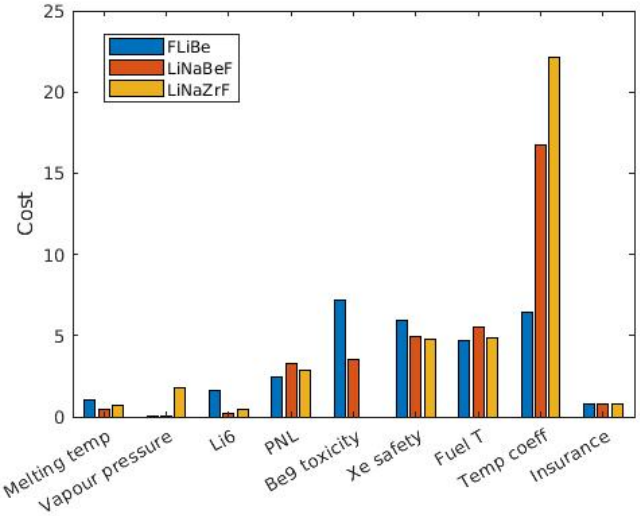

Figure A.13: Breakdown of the 'Safety' category of cost into components. The cost values are given at graphite pebble, PF and inner radius combination that results in the lowest feasible total cost. Only one cost resulting from TCs is calculated, it is a function of the summation of the density and doppler TCs.

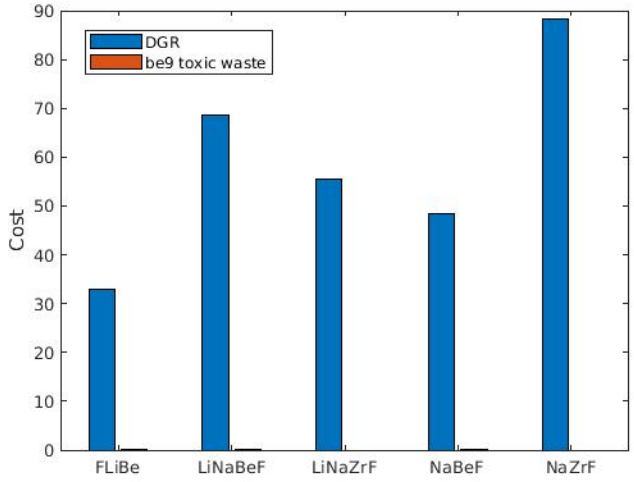

Figure A.14: Breakdown of the 'Waste' category of cost into components. The cost values are given at the configuration that results in the lowest feasible total cost.

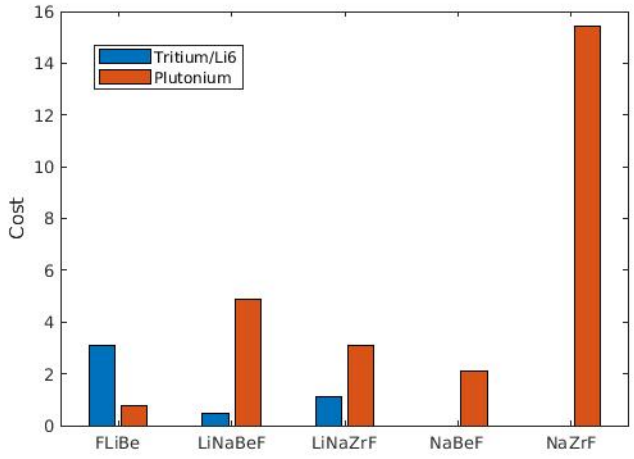

Figure A.15: Breakdown of the 'Proliferation' category of cost into components. The cost values are given at the configuration that results in the lowest feasible total cost.

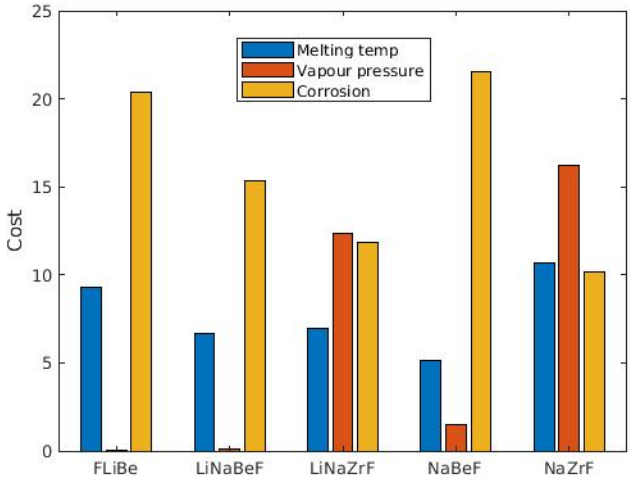

Figure A.16: Breakdown of the 'Feasibility' category of cost into components. The cost values are given at the configuration that results in the lowest feasible total cost.

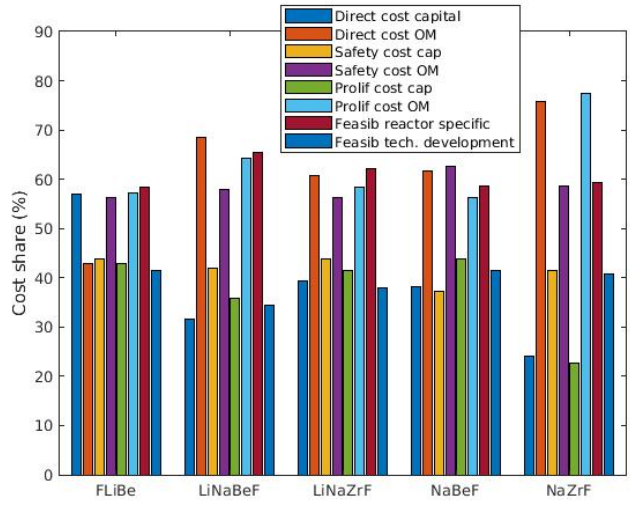

Figure A.17: Additional breakdown of several cost categories. The cost values are given the configuration that results in the lowest feasible total cost.

Table A.1: Sensitivity of total FHR cost to each input explanatory variable at the lowest cost configurations. The values shown are scaled relative to the standard deviation ($b_i \cdot \sigma_i$), thus the sensitivities listed are $\frac{d(Cost)}{d\sigma}$.

Variable	Salt 1	Salt 2	Salt 4	Salt 5
Construction time (yrs)	6.78	4.1	5.66	6.51
Interest rate	16.5	11.8	14.2	16.3
Lifetime (yrs)[a]	-8.34	-7.93	-10.1	-9.93
MWh Cost ($/MWh)	0.137	0.179	0.152	0.243
Pumping Multiplier	-0.324	-0.175	-0.232	-0.126
Heat exchanger material cost ($\$/m^2$)	3.67	2.92	4.01	4.28
Core pressure drop (kPa)	-0.326	-0.157	-0.265	-0.349
Capacity factor	-0.735	-0.781	-0.867	-0.952
^{7}Li cost ($\$/kg$)	3.36	-0.155	0.47	1.27
Beryllium price	0.718	0.68	0.711	0.361
Zirconium price	-0.18	-0.129	-0.175	0.708
Separative work	1.77	2.95	4.67	3.49
U_3O_8	2.15	3.67	5.78	4.36
Fuel pebble manufacturing cost	4.2	7.07	10.5	8.45
Graphite pebble manufacturing cost	0.107	0.166	0.107	0.0642
Ref. refuel OM	0.153	0.0143	0.116	0.102
Ref. refuel cap	4.08	5.64	7.12	6.35
Cost overrun due to safety	-0.324	-0.365	-0.256	-0.333
Share cap safety costs	6.25	4.98	6.36	6.87
Ref. safety OM costs	6.02	6.43	6.69	6.72
Cost overrun prolif	-0.626	-0.528	-0.499	-0.616
Share cap prolif	0.626	0.513	0.866	0.717
Ref. OM prolif	0.354	0.0449	0.55	0.379
Ref. share modularity	2.76	1.95	2.57	2.83
Be handling	-0.317	-0.207	-0.292	-0.34
Detritation OM	1.25	0.251	0.574	0.784
Ref. tech. develop	1.12	1.26	0.673	1.03
Share tech. develop	3.25	3.29	1.95	3.12
T_{melt} split	-0.0833	-0.222	-0.104	-0.239
P_{vap} split	-0.324	-0.152	-0.249	0.0198
Corrosion split	0.534	0.779	0.514	-0.117
Rel. corrosion[b]	1.87	1.86	1.44	1.16
Non-configuration	4.84	5.14	4.69	5.27
Ref. cost MWth	0.0104	0.14	0.122	0.112
Insurance premium	-0.166	-0.0665	-0.141	-0.151
CDF	-0.0106	-0.0638	-0.0292	0.00907
T_{melt} max	2.66	1.58	1.89	2.11
Min. consequence enrich	-0.0267	-0.0413	0.00234	-0.00165
Detritiation C1	1.21	-0.116	1.29	1.33
Detritiation C2	0.798	0.346	0.591	0.656
Detritiation C3	-0.804	-0.0191	-0.429	-0.567
Non-config. exp.	-0.447	-0.342	-0.432	-0.468
Relative prolif cost exp.	-0.553	-0.342	-0.428	-0.526
Regular cost modularity exp.	-1.17	-0.814	-1.01	-1.09
Overrun exp.	4.32	3.49	4.22	4.58
Feasib each exp.	-1.07	-1.1	-1.25	-0.772
Feasib size exp.	0.313	0.198	0.224	0.243
Safety cap. exp.	-2.5	-1.84	-1.82	-1.87
Safety OM exp.	-0.212	-0.14	-0.161	-0.222
Refuel cap exp.	0.0762	0.58	1.33	0.971
Flux to Xe PF exp.	-0.274	-0.136	-0.215	-0.222
Ref. Pu cost	-0.0446	0.234	0.617	0.395
T_{melt} cost	-0.781	-1.38	-0.765	-1.2
P_{vap} cost	-1.66	-1.23	-1.36	1.17
Corrosion cost	0.783	0.936	0.501	0.0458
Be cost	0.589	0.318	0.0801	-0.428
^{6}Li effect cost	-0.186	-0.626	-0.906	-0.653
PNL cost	-0.0678	0.0626	0.0448	-0.0166
Xe PF safety cost	0.437	0.349	0.32	0.383
Temp. feedback cost	0.0644	0.373	0.668	0.793
Fuel temp. cost	-0.324	-0.322	-0.264	-0.178
T_{melt} safety weight	0.851	0.591	0.542	-1
P_{vap} safety weight	0.696	1.07	0.431	0.993
Corrosion feasib weight	0.945	1.09	0.761	-0.307
^{6}Li effects safety weight	-0.701	-0.652	-0.434	-0.643

[a] Unadjusted

[b] Though the probability distribution is same for all salts, it must be split into 5 variables, one for each salt. Otherwise the total corrosion cost would be variable but not the relative corrosion difference between salts. The relative corrosions variables are combined into one line for convenience.

A.2 Additional MSR results

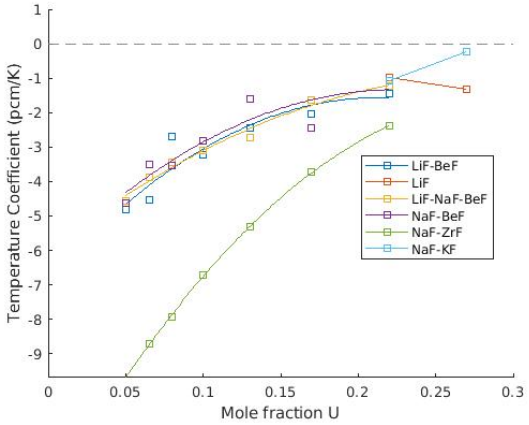

(a) FF: 0.06

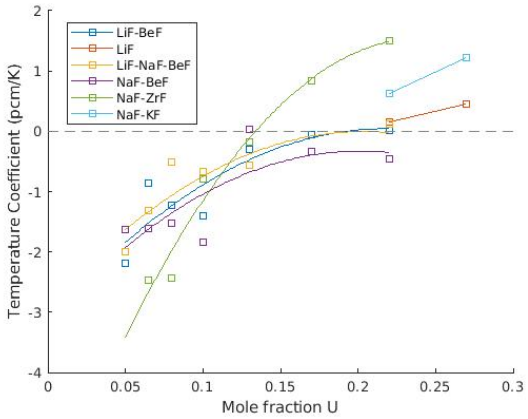

(b) FF: 0.15

Figure A.18: Temperature coefficients vs the evaluated mole % U fractions for all fuel salts at a constant FF.

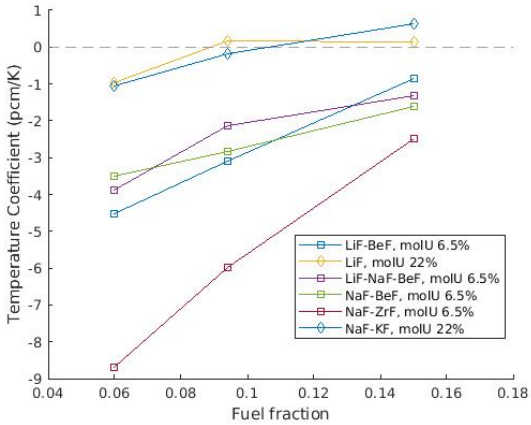

(a) Mole U: 6.5%

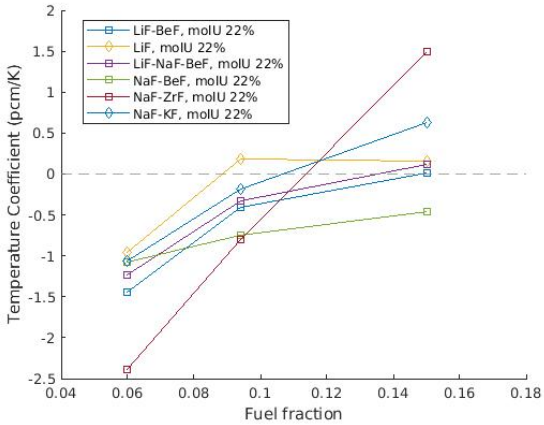

(b) Mole U: 22.0%

Figure A.19: Temperature coefficients vs the evaluated fuel fractions for all fuel salts a constant mole % U. The mole % U is at the stated value if feasible and 22% otherwise (see legend).

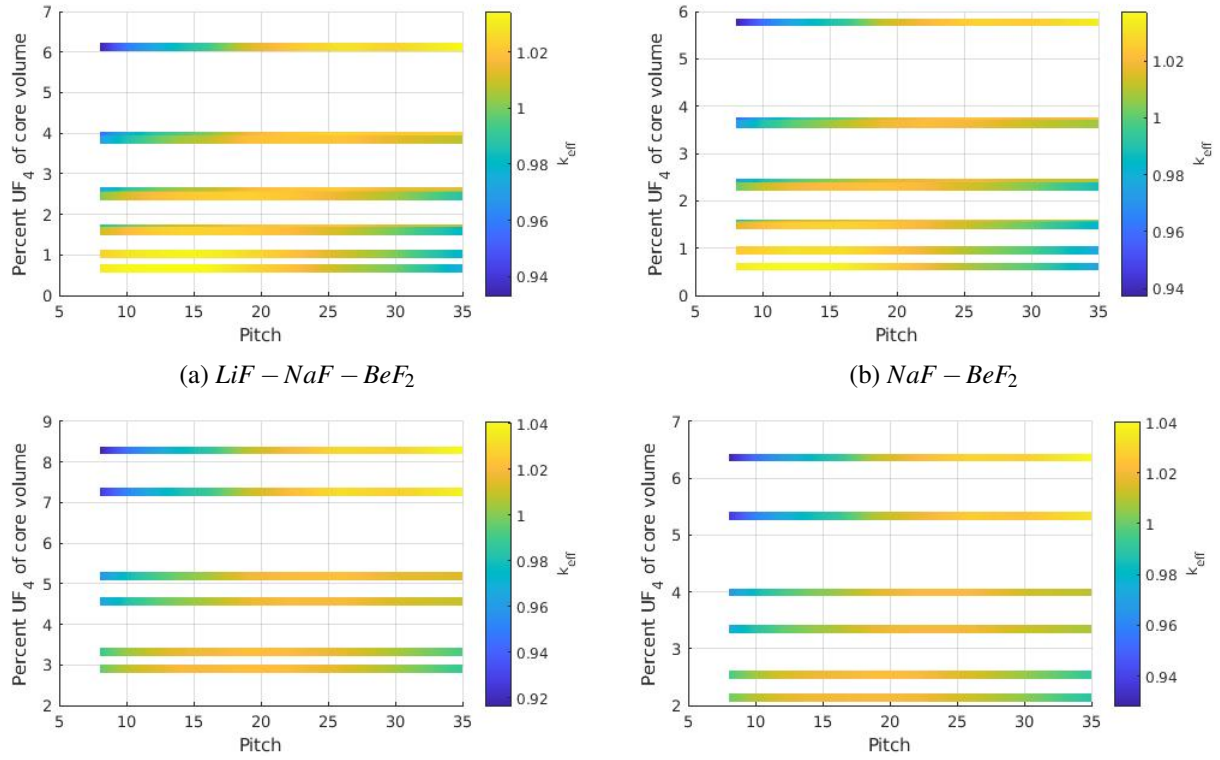

(a) $LiF - NaF - BeF_2$ (b) $NaF - BeF_2$

(c) LiF (d) $NaF - KF$

Figure A.20: k_{eff} vs pitch at several core fuel concentrations. LiF and $NaF - KF$ carrier salts are only feasible at high molU fractions.

Table A.2: Henry's Constant ($K_p \cdot 10^3$, $mole/(cm^3 \cdot atm)$) at 700 ^{o}C for He and Xe and solubility (mole %) at 700 ^{o}C of Rare-Earth & TRansUranic (RETRU) elements in candidate carrier salts. The sources do not necessarily specify the stated value; in some cases for solubility for the candidate salt mixture could not be found and so is estimated primarily from the value of a similar salt composition. The RETRU solubilities, particular at 22% U, do not have a specific source but rather are estimated from the synthesis of information.

	Salt 1	Salt 2[a]	Salt 3	Salt 4	Salt 5	Salt 6[a]
Noble	10.7 [193]	17.5	17.5[a]	16.7 [194] [195]	36.5 [196]	17.5
Xe	0.51 [193]	1.35	1.35[a]	0.86 [194] [195]	3.25 [196]	1.35
RETRU 0% U	0.6 [192]	16	1.5 [197]	2 (est.)	2.2 [198]	16
RETRU 22% U[b]	2	4	3	2	1.6	4

[a] Data could not be found, assumed the same as FLiNaK ([190] [191] [192]).
[b] All estimated, based on [192].

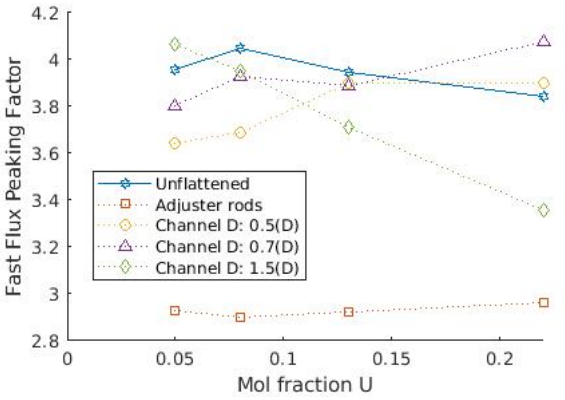

Figure A.21: Fast Flux Peaking Fraction as a function of mol fraction U, for the $LiF - BeF_2 - UF_4$ salt system at a FF of 0.094. The channel diameters (D) listed are the channel diameters of the inner zone relative to the outer zone.

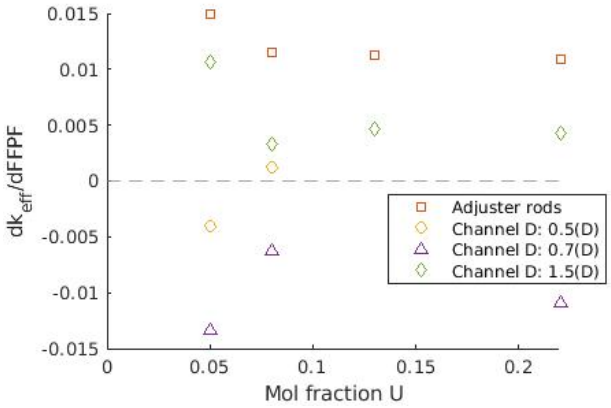

Figure A.22: Change in k_{eff} compared to the FFPF difference (FFPF unflattened - FFPF of flattening mechanism) of several flux flattening mechanisms and mol fraction U values. Analysis conducted on the $LiF - BeF_2 - UF_4$ salt system at a FF of 0.094. The channel diameters (D) listed are the channel diameters of the inner zone relative to the outer zone. Data points where the magnitude of the FFPF difference is small (< 0.07) are left out to avoid division by a small number.

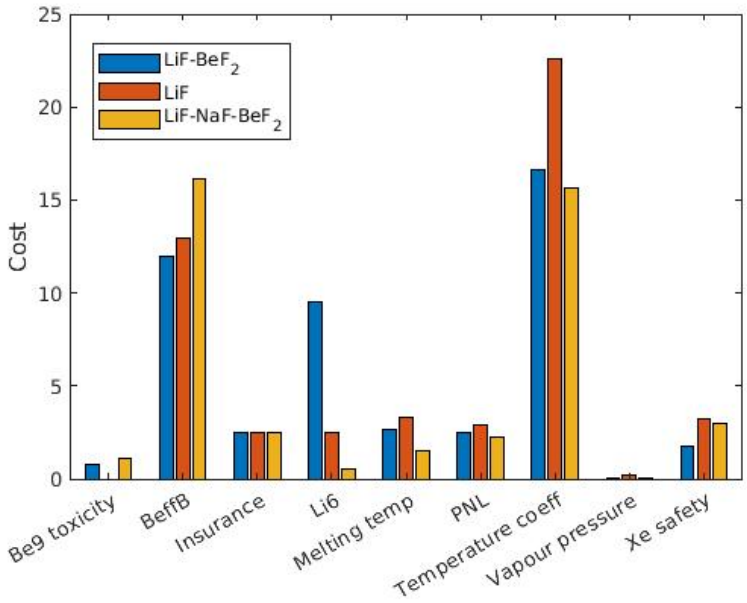

Figure A.23: Breakdown of the 'Safety cost' category of cost into components for three salts. The cost values are given at the mole % U and FF combination that results in the lowest feasible total cost.

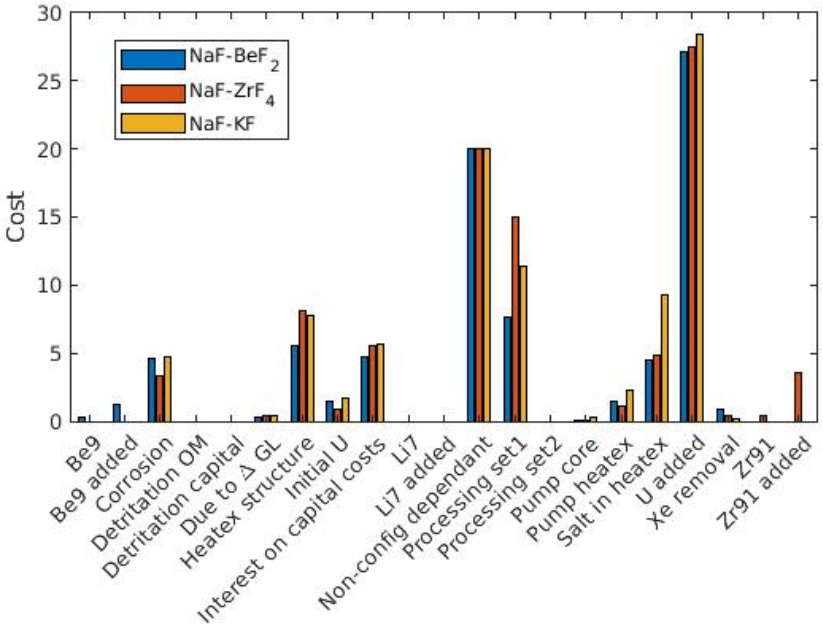

Figure A.24: Breakdown of the 'Direct cost' category of cost into components for three salts. The cost values are given at the mole % U and FF combination that results in the lowest feasible total cost.

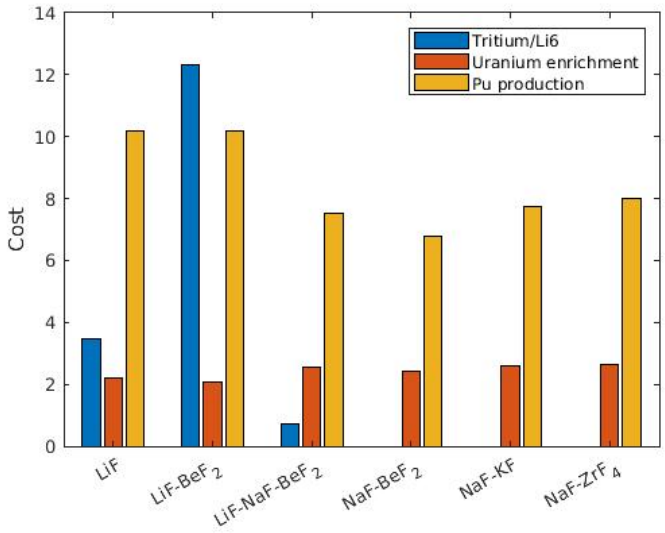

Figure A.25: Breakdown of the 'Proliferation' category of cost into components for all salts. The cost values are given at the mole % U and FF combination that results in the lowest feasible total cost.

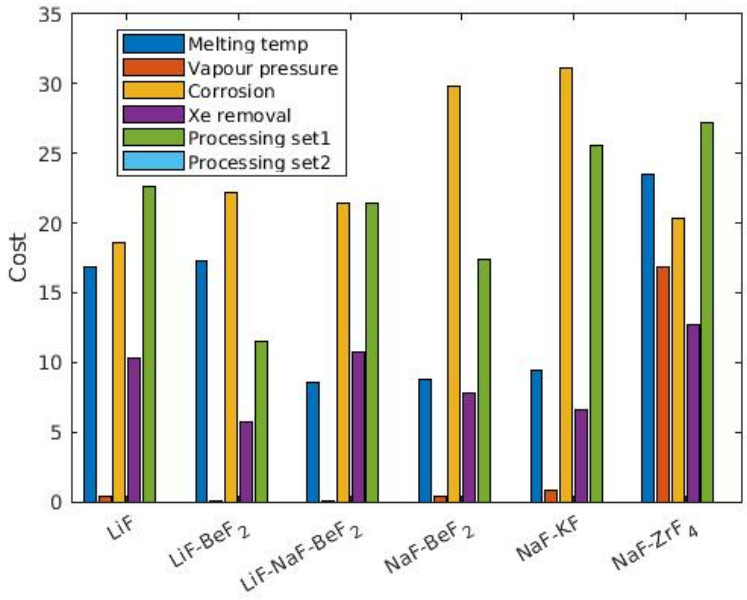

Figure A.26: Breakdown of the 'Feasibility' category of cost into components for all salts. The cost values are given at the mole % U and FF combination that results in the lowest feasible total cost.

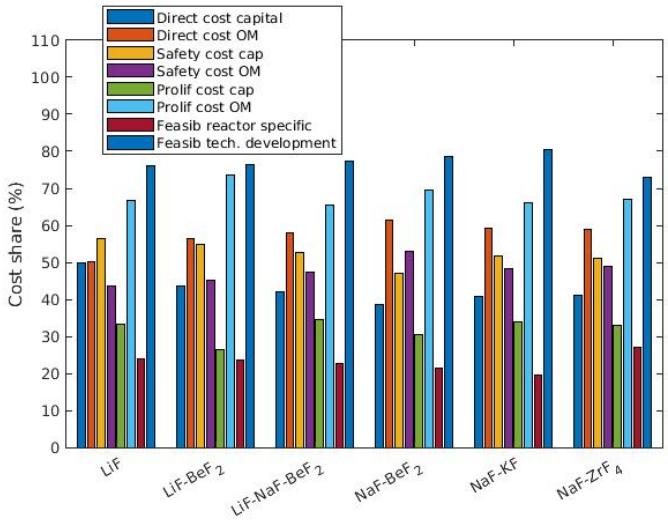

Figure A.27: Additional breakdown of several cost categories. The cost values are given the mole % U and FF combination that results in the lowest feasible total cost.

Table A.3: Sensitivity of total MSR cost to each input explanatory variable at the lowest cost configurations. The values shown are scaled relative to the standard deviation ($b_i \cdot \sigma_i$), thus the sensitivities listed are $\frac{d(Cost)}{d\sigma}$.

Variable	Salt 1	Salt 2	Salt 3	Salt 4	Salt 5	Salt 6
Construction time (yrs)	6.81	6.48	5.42	3.54	4.43	5.07
Interest rate	20.5	21.6	17.3	13.3	17.3	18.8
Lifetime (yrs)[a]	2.7	-0.467	-0.279	1.09	-0.196	-2.38
MWh Cost ($/MWh)	1.03	0.861	0.877	0.847	0.979	0.724
Pumping Multiplier	-0.268	-0.18	-0.19	-0.191	-0.121	-0.0539
Heat exchanger material cost ($\$/m^2$)	4.36	3.88	5.11	3.71	5.07	5.18
Capacity factor	-1.32	-2.15	-1.58	-1.52	-0.898	-2.52
^{7}Li cost ($/kg)	12.8	9.28	2.7	0.511	0.621	0.405
Beryllium price	0.213	-0.225	0.493	0.516	-0.224	-0.147
Zirconium price	0.493	0.442	0.264	0.0817	1.81	0.16
Separative work	2.25	1.91	3.16	3.19	3.16	4.11
U_3O_8	4.23	4.23	5.4	5.07	4.56	7.3
Cost overrun due to safety	0.0727	-0.267	-0.395	-0.375	-0.463	-0.222
Share cap safety costs	12	12.1	9.56	7.34	10	11.4
Ref. safety OM costs	5.75	6	5.76	5.78	6.79	7.36
Cost overrun prolif	-0.154	-0.589	-0.503	-0.408	-0.579	-0.601
Share cap prolif	4.85	4.37	2.99	2.41	3	3.11
Ref. OM prolif	0.772	0.329	0.0159	-0.0387	-0.0594	-0.0281
Ref. share modularity	5.88	6.18	5.15	4.03	4.87	5.56
Be handling	-0.125	-0.191	-0.0459	0.00403	-0.121	-0.109
Detritation OM	1.43	1.04	-0.0296	-0.593	-0.583	-0.523
Ref. tech. develop	2.03	3.66	2.39	2.5	4.87	4.07
Share tech. develop	5	9.63	6.25	6.66	12.6	10.5
T_{melt} split	-0.0478	0.118	-0.222	-0.31	-0.0647	-0.154
P_{vap} split	0.177	-0.0413	-0.146	-0.101	0.448	-0.139
Corrosion split	1.04	0.962	1.16	1.61	0.401	1.77
Xe remove split	0.522	0.731	0.579	0.349	0.185	0.397
Processing 1 split	-1.17	-2.28	-0.948	-0.934	-2.7	-2.21
Processing 2 split	0.792	1.72	1.12	0.918	2.21	1.27
Rel. corrosion[b]	1.89	0.234	1.78	3.53	-0.579	0.332
Non-configuration	10.5	11.5	11	10.9	13.4	11.9
Ref. cost MWth	-0.263	-0.231	-0.0239	-0.0465	-0.185	-0.0846
Insurance premium	-1.19	-0.963	-0.662	-0.702	-0.917	-0.731
CDF	0.142	0.151	0.125	0.102	0.117	0.209
Min. consequence enrich	-0.31	-0.349	-0.225	-0.0788	-0.235	-0.175
Detritiation C1	2.82	2.79	2.5	0.777	0.944	0.871
Detritiation C2	2.22	1.76	1.12	0.884	1.05	0.861
Detritiation C3	-1.63	-1.2	-0.548	0.225	0.296	0.276
Processing 1 direct cost	1.98	7.28	4.51	3.53	7.87	7.79
Relative prolif cost exp.	0.31	0.105	-0.149	-0.0933	-0.157	-0.207
Regular cost modularity exp.	-1.61	-1.75	-1.49	-1.16	-1.55	-1.55
Overrun exp.	12.6	13	10.6	8.76	11.5	11.9
Feasib each exp.	-1.19	-0.173	-0.925	-0.862	1.68	-0.271
Feasib size exp.	-0.622	-0.542	-0.521	-0.526	-0.598	-0.475
Safety cap. exp.	-2.96	-2.15	-2.83	-2.52	-1.4	-1.02
Safety OM exp.	0.627	0.634	0.543	0.484	0.594	0.527
Processing exp.	-0.482	-0.385	-0.327	-0.26	-0.268	-0.327
Fresh contrib	-0.335	-0.359	-0.272	-0.164	-0.219	-0.248
Flux to Xe PF exp.	-0.332	-0.41	-0.239	-0.188	-0.268	-0.242
Enrich at % rel. cost	-0.771	-0.74	-0.255	-0.0583	-0.255	-0.278
Ref. Pu rel. cost	3.43	3.23	2.35	2.27	2.65	2.41
T_{melt} rel. cost	0.0269	-0.77	-1.51	-1.62	-0.207	-2.43
P_{vap} rel. cost	-1.7	-2.08	-1.77	-1.62	3.15	-2.09
Corrosion rel. cost	2.13	0.143	1.77	3.13	-0.189	1.63
Xe removal rel. cost	0.423	0.276	0.125	0.183	0.157	0.201
Be rel. cost	-0.955	-0.993	-0.668	-0.564	-1.01	-1.04
^{7}Li effect rel. cost	-0.712	-1.31	-1.55	-1.59	-1.91	-1.95
PNL rel. cost	0.165	0.217	0.179	0.224	0.0708	0.152
Xe PF safety rel. cost	-0.479	-0.79	-0.26	-0.226	-0.0411	-0.765
Temp. feedback rel. cost	-0.116	0.645	0.11	0.45	-0.517	1.15
β_{eff} rel. cost	-0.0171	0.0538	0.142	0.118	0.06	0.0245
Processing 1 rel. cost	0.288	5.89	2.63	1.43	3.54	5.71
Processing 2 rel. cost	-0.153	-0.165	-0.0648	-0.145	-0.301	-0.226
Vg/Vc	0.145	0.181	0.13	0.0254	0.0946	0.0797
Mass transfer coeff.	0.402	0.169	0.301	0.289	0.229	0.136
Time outside	0.504	0.456	0.215	0.311	0.336	0.2

[a] Unadjusted

[b] Though the probability distribution is same for all salts, it must be split into 6 variables, one for each salt, to account for the relative corrosion difference between salts. The relative corrosion variables are combined into one line for convenience.

Appendix B

Salt Section Optimization Program

B.1 File Description

The program is divided into a number of scripts, functions and classes. In the main folder are the files:

- **costSettingInputs, script:** Cost setting values are input in this file. The values are meant to represent an approximate average property value or metric of performance and each is associated with a corresponding relative cost.
  - In uncertainty analysis, the relative costs are assigned here each run through. For fixed-variable runs, the average value of each relative cost is assigned here.
- **corrosion, script:** Determines the relative corrosion cost of each salt.
  - A comprehensive predictive model of corrosion is beyond the scope of this thesis. Instead corrosion is predicted by combining experimental results with differences in Gibbs free energy.
  - Assumptions/scope: corrosion with Ni-based Hastelloy N - the container material used in the MSRP. The cost due to Gibbs free energy is proportional to the average difference in Gibbs between the salt and container components averaged with the minimum Gibbs free energy difference between salt and container components.
  - MSR: insufficient experimental results were found to determine the relative corrosion impact as a function of uranium concentration; thus it is assumed corrosion is constant across all uranium concentrations.
- **deltaKeffLi6, function** Inputs the ^{6}Li atomic density and outputs Δk_{eff} due to the presence of ^{6}Li.
- **dgrCostFn, function:** A description of how the DGR cost is calculated can be found in Section 2.4.
  - FHR inputs: refuelling duration, adjusted refuelling time, blanks, particle PF, power in MWe, no. of pebbles, no. of particles per pebble, capacity factor.
  - MSR inputs: salt composition, mole % U, core salt volume, heat exchanger volume, power in MWe, capacity factor.
  - Output: expected cost of a DGR.

- **eachSaltClass, object:** It stores the properties/attributes of each salt (ex. T_{melt}) and has various methods most of which are used to determine the value of particular properties.
- **econDiff, function:** The construction time, interest rate and plant lifetime sample of each uncertainty/sensitivity analysis iteration as well as the mean/expected value of each of these parameters is assigned in this function. The function accepts two values, say varb1 and varb2. Varb1 specifies whether a capital cost (Equation B.1), continuous cost (Equation B.2) or end-of-life cost (Equation B.3) calculation is to be made or if the construction time/interest rate/plant lifetime is to be returned. If the calculation is to be made, varb2 specifies whether the result of Equation B.1, B.2 or B.3 is returned or the same Equation B.1, B.2 or B.3 result relative to the result with mean/expected values. If varb1 denotes a return of construction time, interest rate or plant lifetime, varb2 specifies which of those to return. For Equations B.1, B.2 and B.3, T_c is the construction time, i the interest rate and L the lifetime.

$$\text{Capital} = (1+i)^{T_c/2} \tag{B.1}$$

$$\text{Continuous} = ((1+i)^L - 1)/(i(1+i)^L) \tag{B.2}$$

$$\text{End-of-life} = (1+i)^{-L} \tag{B.3}$$

- **feasibCnsq, script:** Determines the cost components that contribute to the feasibility category. The costs are a function of the input and reference values used in cost setting that relate to feasibility (T_{melt}, P_{vap}, etc.). The costs are made up of both an 'each-build' and technological development component.
- **findTC, function:** Returns an interpolated fit of the density or Doppler (Temperature) reactivity feedback Coefficient (TC) for the particular blanks, particle PF and inner radius configuration input (or mole % and FF).
- **fuelMassFraction, function:** Returns the mass fraction of U in the fuel salt. The mass fraction is used in the conversion of volumetric flow rate of the salt mix to mass flow rate of U.
  - MSR only
- **inputCost, object:**
  - Stores the total cost of each of the 6 categories of each salt along with a vector of the cost components.
  - Store the cost-category weighted total cost and vector of cost-category weighted cost components.
- **inputs, script:** Most numeric inputs/assumptions are specified here. Ex. pumping and heat exchanger related cost values, core power (MWe and MWth), core power density, salt component costs, any reference costs.

  - In uncertainty analysis, the input parameter values are assigned here each run through. For fixed parameter runs, the average value of each input parameter is assigned here.

- **inputsSalt, script:** Properties of the salt are added to the salt object's (instance of 'eachSaltClass') properties here; melting temperature, vapour pressure, solubility, corrosion and salt component atomic densities.
  - Only MSR, in the FHR the relevant properties are added in the inputs script.

- **mainAll, script:** The core of the program, mainly consisting of sequential steps that add cost components - either in the file itself or it directs to other files.
  - Key metrics of each configuration are saved as attributes in the respective instance of the class 'eachSaltClass' or class 'inputCost'.

- **matrixSaltdT, function:** Calculates the (reactor average) graphite matrix to coolant salt temperature difference.
  - FHR only

- **mflowAdjstr, function:** Adjusts the fuel salt flow rate to account for the difference in lifetime from the expected value.
  - MSR only

- **pnlFn, function:** Performs a quadratic (double) interpolation to determine the prompt neutron generation time for the particular blanks, particle PF and inner radius (or enrichment and mole % uranium).

- **presentOMCostCalc, script:** Calculates the reference O & M safety and proliferation cost at the start-up time from $/MWh input values. Also calculates the total of all safety relative costs and all proliferation relative costs.

- **pumpHeatexScript, script:**
  - Determines the heat exchanger design that results in the minimum overall cost given the pumping cost ($/kW), the volumetric cost of the salt ($/$cm^3$) and cost of the heat exchanger structure ($/(heat transfer surface area)) as a function of salt composition.
  - Inputs four matrices: 'optRangeV' and 'optRangeC' that store the optimum average velocity and minimum cost respectively of each salt. The matrix 'power' stores the required pump size (kW) and 'allNoTubes' stores the number of primary side tubes in the heat exchanger.
  - Calculated values include: Cost of the heat exchanger structure, heat exchanger primary fluid volume and pumping power cost.

- **prolifCnsq, script:** Determines the cost components that contribute to the proliferation category and splits it between capital and O & M costs. Like 'safetyCnsq,' the capital cost portion of the proliferation cost is proportional to the total direct capital cost and overrun cost while the O & M cost is calculated relative to a user-input reference O & M proliferation cost. The cost of specific configurations relative to the reference is a function of the input and the reference values used in cost setting.

- **quadInterp3, function:** Inputs the configuration under evaluation, the configuration data points and the associated data. Then outputs the results of a quadratic interpolation if possible or else a linear interpolation.
  - FHR only (because there are more than two non-salt design parameters)
- **rhoNew, function:** Outputs the salt density given a carrier salt, temperature and mole % U.
  - the mole % U can be set to zero so it can and is used for the FHR as well.
- **runner, script:** Starting file that initiates the evaluation program. The design configurations evaluated and whether uncertainty analysis is being conducted are input here.
  - Configurations, FHR version: range of inner radius, particle PF and share blanks.
  - Configurations, MSR version: range of mole % U and fuel fraction.
  - Key results of the evaluation runs are saved in matrix form as .mat files.
- **safetyCnsq, script:** Determines the cost components that contribute to the safety category and splits it between capital and O & M costs. The capital cost portion of the safety cost is proportional to the total direct capital cost and the overrun cost while the O & M cost is calculated relative to a user-input reference O & M safety cost. The cost of specific configurations relative to the reference is a function of the input and the reference values used in cost setting.
- **surrDataNum, function:** Returns the number of surrounding data points for the input configuration.
  - Ex. if a configuration lies just outside the edge of a 2D grid of data points there are two surrounding points; within the grid there are four.
  - The minimum number of data points is set as two but can increased. If any calculation does not meet this requirement the configuration ceases to be evaluated.
- **wasteCnsq, script:** Determines the value of the cost components that contribute to the waste category. The major cost contributor is the DGR, the cost of which is determined in the 'dgrCostFn' function.
- **thermFluxCalc, function:** Outputs the thermal flux as a function of configuration.
  - FHR version: specifically a function of inner radius and relative fuel concentration.
  - MSR version: specifically a function of fuel concentration and enrichment.
- **tripleInterp, function:** Determines a value from a matrix of data through triple linear interpolation.
  - MSR only
- **UpriceFn, function:** Outputs the price of uranium in $/kg given the enrichment, separative work cost and U_3O_8 price.
- **weighting, script:** It assigns weights that do not effect any input parameters values directly but rather the way in which inputs produce cost values.
  - Relative weighting of input parameters to cost categories: Ex. If T_{melt} safety = 3, T_{melt} feasibility = 8, $\frac{3}{8}$ of the T_{melt} cost, as set in the costSettingInputs script, contributes to the safety cost.

- **xePoisonFrac, script:** Outputs the xenon poison fraction (at the reference thermal flux level) given the volume of salt in the graphite, volume of salt in the channels, mass transfer coefficient from the salt to the graphite, the power density and the Xe stripping half-life.
  - MSR only

B.1.1 Uncertainty analysis related:

- **correlMatrixBuilder, script:** Creates a symmetric matrix of the following variable names:
  1 AB AC ... AX
  AB 1 BC ... BX
  AC BC 1 ... CX
  ...
  AX BX CX ... XX
  Each variable represents a correlation coefficient between two variables, AC the correlation between the input variable associated with A and that with C. X represents a letter corresponding to the number of input variables. If greater than 26 it becomes AA(27), eventually BA(53) etc.

- **distCreator, script:** The probability distribution parameters of the input parameters are input here and it divides each distribution into bins of equal probability.
  - If uncertainty analysis: Inputs the correlation matrix as well.

- **distCreator2, script:** Creates a permutation of integers for each of the input parameters to store the order of bin number draws. Then it draws a value for each parameter from the bin associated with the run count. The location of the draw within the bin is random.

- **distCreatorCosts, script:** Has the same function as distCreator and distCreator2 but the cost categories of waste, safety, proliferation and feasibility are the only input parameters. There is also no correlation.

- **doubleTrap, function:** Inputs the mean, minimum and maximum as well as the number of bins and creates a vector of equal-probability bins between the minimum and maximum (see Section 4.3.2).

- **inputCorrels, script:** The file where correlation coefficients are input, so that correlations do not have to be entered directly into the correlation matrix. All non-specified correlations are automatically set to zero.

- **inputMatricies, script:** Matrices of data generated from SERPENT simulations or thermohyraulic calculations.
  - Only MSR

- **pearsonMakeBins, function:** Inputs the mean, standard deviation, skew and kurtosis as well as the number of bins and creates a vector of equal-probability bins (see Section 4.3.1).

B.2 Evaluation Program: Procedural Walk-through

The program is initiated by running the 'runner.m' script. All other subsequent script and function calls are made automatically. For each configuration the 'runner.m' file initiates the following sequence to calculate costs:

1. Call the main program ('mainAll') and pass the design parameters values (FHR: blanks, PF and radius, MSR: mole % U and FF).
2. Create an instance of the class 'eachSaltClass' and of 'inputsCost' for each salt.
3. Run inputs (MSR only: and 'inputsSalt'), this will assign attributes to many of the properties in each instance.
4. Run 'weightings' to assign in all weightings values.
5. Run 'costSettingInputs' to assign in the value of parameters involved in the 'cost setting' process.
6. Determine and assign to the properties of each salt instance, the relative thermal flux and Prompt Neutron Lifetime (PNL).
7. FHR only: perform necessary interpolations to determine the initial k_{eff}.
8. FHR only: calculate the (at reactor start-up) value of added fuel and graphite pebbles over the reactor lifetime.
9. MSR only: calculate the (at reactor start-up) value of added fuel over the reactor lifetime.
   - Perform necessary interpolations to assign in $\frac{d\dot{m}}{dk_{eff}}$.
10. Determine the optimum ^{7}Li enrichment level for the applicable salts. This is done by finding the optimum balance of the cost of enrichment and the effects of Tritium production.
11. Calculate the \$/cm^3 of all expensive salt components.
12. From the pumping cost and the cost of the heat transfer surface area of the heat exchanger along with the optRangeV (optimum avg. salt velocity) and optRangeC (optimum cost) matrices determine the optimum pumping power and heat exchanger size. This calculation involves the 'pumpHeatexScript.'
    - MSR only: Interpolate between the optimum at 0 % U and 22 % U to determine the optimum velocity/cost at the given mole % U.
13. Calculate the overall pumping power cost, costs of the heat exchanger and cost of the salt in the heat exchanger.
14. MSR only: Determine costs involved with the processing of fission products.
15. MSR only: Determine the optimum pitch and determine the corresponding lifetime, adjusted k_{eff}/salt-flow-rate, core pumping power required and optimum Xe removal rate.

16. MSR only: Calculate the effective delayed neutron fraction and the associated cost(s).
17. Determine the costs associated with detritation for the relevant salt compositions.
18. FHR only: Determine the maximum temperature difference between the fuel and salt coolant.
19. Calculate the relative salt melting temperature cost: $T_{max}/(T_{max} - T_{melt})$.
20. Determine the proliferation effect of plutonium relative to the reference. It is largely a function of burnup.
21. Determine the temperature coefficients and assign to the corresponding properties of each salt instance.
22. Divide 'direct costs' between capital and O & M and determine the interest costs during construction.
23. Determine the modularity costs.
24. Calculate the cost components of the cost categories 'waste,' 'safety,' 'proliferation,' and 'feasibility.' Most of the waste cost is associated with the DGR, the function dgrCostFn calculates that cost. The costs components of the other three categories are generally calculated in accordance with the 'cost setting' method.
25. Determine the cost of insurance.
26. For each cost category, assign into the instances of the 'inputsCosts' class a vector of the value of each cost component. As well store the sum of each vector.
27. Make adjustments to all costs for capacity factor, lifetime and interest rate (lifetime and interest rate can also be adjusted after the evaluation is complete).
28. Fixed values only: Save the instances of the 'eachSaltClass' and 'inputsCosts' classes.
29. Save the total cost of each configuration and other parameters of the calculation as desired.

The 'eachSaltClass' class has a property, 'possible', that signifies whether the salt at the configuration evaluated is a feasible option. It is initially set 'true' and is switched to 'false' the program comes across a test indicating that the configuration is infeasible.

Appendix C

Physical Properties

Physical property values of the molten salt compositions examined throughout this paper are required for thermohyraulic and various other analysis. Many physical properties are relevant to the feasibility of a particular salt composition, including viscosity, density, melting point, vapour pressure and the heat transport characteristics, heat capacity and thermal conductivity. These properties in all MSRs as well as the solubility of actinide fuels and fission products in circulating-fuel reactors is of major significance to safe and effective operation. This section lists property data at temperatures in the reactor operating range and discusses theory, trends, data sources and uncertainty.

Before listing properties it is important to point out that the knowledge of properties at one exact salt composition may not be useful as there exists a spectrum of potential properties across all different possible compositions in that system. The issue becomes more complex as more unitary salts are added to the mix.

Due to the large range of evaluated salt compositions, experimental data alone is insufficient to determine necessary property values, thus predictive methods and interpolation are used as required. The salt compositions without uranium ('clean') are generally at eutectic points, so experimental data is available for most physical properties.

Though the use of chloride salts are likely only feasible in fast-spectrum reactors, the property values of select chloride salts are provided give a sense of the thermohyraulic trade-offs that would be involved. As well, the property values of binary salts and other common liquid/molten heat transport fluids are provided for comparison.

It can be useful to classify the examined salts. All fluoride salt systems analysed can be grouped among the following three categories

- Alkali salts; The cation(s) is/are only alkali metal(s).
- BeF_2 salts; Are made up of (an) alkali salt(s) and BeF_2, ex. $LiF - BeF_2$ (FLiBe)
- ZrF_4 salts; Are made up of (an) alkali salt(s) and ZrF_4, ex. $NaF - ZrF_4$

C.1 Basic Physical Properties

The basic physical properties to be examined are, density, viscosity, vapour pressure and melting point. The magnitude of each property is vital to effective and safe operation.

C.1.1 Density

Adequately precise and accurate density data exists for nearly all relevent, simple salts with one cation and one anion such as LiF [110]. These simple salts can be combined into binary, ternary and higher-order mixes to make any MSR fuel or carrier salt. Fortunately the density of each mix can be estimated to a sufficient degree of accuracy [192] using the additive principle.

For the FHR reactor evaluation the density of only five salts must be known. These compositions are Uranium free, at eutectic points and have been considered for reactor applications so density predictions specific to the composition that do not rely on the additive principle are generally available. However, in the case of the MSR where a range of density values must be known across all UF_4 concentrations, the additive principle is generally the only option. Since the share of the salt in the core is substantially lower in circulating-fuel designs ($\approx$15%) than salt-cooled designs (40%), the density value has less of an effect on neutronics.

Density (ρ $[g/cm^3]$) as a function of temperature for molten salts has been found to be linear. Equation C.1 is used in density calculations of binary/ternary salt mixtures and other heat transport fluids in Table C.1, ρ_y is effectively the y-intercept, k is a linear factor and T the temperature; Kelvin or Celsius is specified in the Table C.1. For the unitary salt ZrF_4 Equation C.1 is used but for all other unitary salts, Equation C.2 is used. ρ_m represents the density at the melting point.

$$\rho = \rho_y - k \cdot T \tag{C.1}$$

$$\rho = \rho_m - k \cdot (T - T_m) \tag{C.2}$$

C.1.2 Viscosity

Molten salts are Newtonian fluids that exhibit the typical exponential decrease in viscosity, μ, with temperature. Williams, Toth, Clarno [49] found that the data on binary mixtures is 'fairly complete' and that candidate salts exhibit 'reasonably' low viscosities.

Exact measured viscosities of selected canadidate salts are provided in Table C.2 and notable trends of fluoride salts applicable to thermal reactors found by [49] using data from [190] and [116] are listed below:

- In BeF_2 and ZrF_4 systems, lighter alkali elements result in a higher viscosity. The difference between alkali elements is on average approximately 20%. This is due to the heavier, less electronegative

Table C.1: Density of Select Molten Salts. Composition specific density predictions for the five evaluated coolant salts and $FLiNaK$, the secondary salt, are shown along with the unitary salts used in additive principle density calculations. Other heat transport fluids are shown for comparison.

Fluoride Salt Mixtures				
Composition	ρ at 700 ^oC	ρ_y	k $(\frac{g}{cm^3}/^oC)\cdot 10^{-3}$)	Source
$LiF-NaF-KF$ (46.5-11.5-42)	2.14	2.58	0.624 (^oC)	[199]
$LiF-BeF_2$ (67-33)	1.906	2.214	0.44 (^oC)	[199]
$NaF-BeF_2$ (57-43)	2.01	2.27	0.37 (^oC)	[118]
$NaF-ZrF_4$ (59.5-43)	2.96	3.58	0.89 (^oC)	[103]
$LiF-NaF-BeF_2$ (15-58-27)	2.05	2.57	0.53 (K)	[192]
$LiF-NaF-ZrF_4$ (26-37-37)	2.92	3.53	0.87 (^oC)	[49]
Unitary Salts, all [110]				
Composition	ρ at 700 ^oC	ρ_m	k $(\frac{g}{cm^3}/^oC)\cdot 10^{-3}$)	
LiF	1.88	1.81	0.490	
NaF	2.14	1.95	0.636	
KF	2.01	1.91	0.651	
BeF_2	1.96	1.96	0.015	
ZrF_4	3.56	4.30	1.06 [a]	
UF_4	6.82	6.49	0.992	
$LiCl$	1.46	1.50	0.432	
$NaCl$	1.61	1.56	0.543	
$NaUl_3$	5.63	4.84	0.794	
Other Heat Transport Fluids				
Composition	ρ	ρ_y	k $(\frac{g}{cm^3}/^oC)\cdot 10^{-3}$)	Source
Light Water, 300 ^oC	0.710			
Na, 550 ^oC	0.821	1.014	0.235 (K)	[200]
Pb, 550 ^oC	10.4	11.44	1.28 (K)	[200]

[a] Derived from [103]

elements more readily giving up electrons to the Be^{2+} or Zr^{4+} ions and thus making the mix less like a glass. The phenomenon is explained further in Sections C.1.3 and D.2.1.

- Of the three binary fluoride salt system categories, alkali fluoride mixes have the lowest viscosity, $\approx 3.0cP$ at 700 ^oC, followed by ZrF_4 salts, $\approx 4.5cP$ and BeF_2 salts the highest at $\approx 7.0cP$.
- As the share of BeF_2 in the salt raises, the viscosity strongly increases. In the $LiF - BeF_2$ system at 700 ^oC: BeF_2, 31 mole % = 4.9 cP, BeF_2, 50 mole % = 8.0 cP and BeF_2, 75 mole % = 120 cP.
- Adjusting the concentration of ZrF_4 has a comparatively small impact; it is difficult to draw clear conclusions on the viscosity effect of a change in the ZrF_4 concentration.

Dynamic viscosity (μ) values of various molten salts are provided in Table C.2. As Newtonian fluids, molten salt viscosity can be predicted with Equation C.3. The fluoride salt mixes, most of which are included in the FHR/MSR evaluations, have the empirical A and B constants from Equation C.3. The viscosity was not found for all fuel (UF_4 bearing) salts so estimates at the author's discretion had to be made. The estimated fuel salt viscosities are listed in Section 3.3.2. Inspection of Table C.2, reveals that for all salt systems examined, increases in the concentration of actinide fuel resulted in increased viscosity. As well nowhere was it found that any fluoride salt viscosity values are prohibitively high.

$$\mu = Ae^{B/T} \tag{C.3}$$

C.1.3 Vapour Pressure

In any high-temperature salt system a purged cover gas is necessary to contain the salt mix. ARE and MSRE experience show that very low salt vapour pressures (< 1 mm Hg) simplify the off-gas system design and that certain vapour species can present problems [49]. The off-gas system refers to the system that processes gaseous fission products and tritium.

Vapour pressure is very sensitive to salt composition. Alkali fluorides have a low pressure compared to systems containing BeF_2 and especially ZrF_4. Systems with high ZrF_4 concentrations are problematic. The ARE with $NaF - ZrF_4 - UF_4$, 53.09-40.73-6.18 salt exhibited a vapour containing nearly pure ZrF_4 with a pressure of approximately 10 mm Hg at 900 ^oC[1] [208], over 100 times higher than the MSRE. Pure ZrF_4 sublimes into a 'snow' rather than condensing into a liquid and draining back into the salt reservoir [49].

The mechanisms responsible for the variance in vapour pressure can largely be explained with the acid-base theory; 'Acidic' constituents like Zr^{4+} and Be^{2+} are volatile and are suppressed by the 'basic' F^- constituent. Heavier cations, like Na^+, are less attached to F^- than Li^+, the free F^- ions then go on to suppress the Zr^{4+} and Be^{2+} ions. Thus salts with Na instead of Li support a higher concentration of BeF_2 or ZrF_4.

[1] One standard atmosphere of pressure is equal to 760 mm Hg.

Table C.2: Viscosity of molten salts applicable to reactor applications for which data is available.

Composition	Viscosity at 700 ^oC (cP)	A	B	Source
Fluoride Salt Mixtures				
$LiF-NaF-KF$ (46.5-11.5-42)	2.91	0.04	4170	[111] (pg. 571)
$LiF-BeF_2$ (67-33)	4.89	0.116	3755	[115] [111]
$NaF-BeF_2$ (57-43)	6.98	0.0346	5164	[111]
$NaF-ZrF_4$ (59.5-40.5)	4.57	0.0767	3977	[49] [115]
$LiF-NaF-BeF_2$ (16-56-28)	4.0	0.111	3486	[116]
$NaF-KF-UF_4$ (46.5-26-27)	9.8	0.0767	4731	[116]
$NaF-ZrF_4-UF_4$ (50-25-25)	8.5	–[a]	–	[116]
$NaF-UF_4$ (66.7-33.3)	10.25	0.181	3727	[116]
$LiF-UF_4-ThF_4$ (77.5-2.5-20)	10.1	0.225	3689	[201] / [202]
$NaF-ZrF_4-UF_4$ (50.0-46-4)	5.37	0.981	3805	[118]
$LiF-BeF_2-ThF_4-UF_4$ (71.7-16-12-0.3)	6.78	–	–	[112]
Unitary Salts				
Composition	Viscosity at 700 ^oC (cP)[b]			Source
LiF	3.31			[203]
NaF	6.21			[203]
$LiCl$	1.18			[203]
$NaCl$	1.35			[203]
Chloride Salts for Fast Reactor				
$NaCl-MgCl_2$ (58-42)	1.26			[204]
Actinide, Mg mix-UCl_3-PuCl_3				[205]
(82-12.6-5.4)	2.9			
(70-0-30)	3.3			
(50-27.5-22.5)	3.9			
Other Heat Transport Fluids				
Composition	Viscosity			Source
Light Water, 300 ^oC	0.086			
Na, 550 ^oC	0.22			[206]
Pb, 550 ^oC	1.3			[207]

[a] not provided

[b] If $T_{melt} > 700$, extrapolated value

Polyvalent Zr and Be are less volatile when their valance shell is satified by the necessary covalent (coordination) bonds [49]. The reasonable maximum ($<\approx$ 10 mm Hg) Zr concentration is about 25-45 mole % and 35-45 % for Be [49] and [209] depending on the alkali cations and temperature. The vapour pressures of PuF_3, UF_4 and ThF_4 are negligibly small at temperatures practical to reactor applications [111], thus the addition of fuel to a carrier salt will lower the vapour pressure.

Chloride salts have a considerably higher vapour pressure than comparable fluorides, see Table C.3. Pu and U with greater than 3 bonded chloride ions causes unacceptably high vapour pressures. And the vapour pressure of UCl_3 is approximately 25% greater than $PuCl_3$ [210]. There is little data on the vapour pressures of chloride salts with actinide fuel, but it seems safe to assume that the vapour pressure decreases with the addition of actinide fuel, as evidenced by [211] and the low vapour pressures of UCl_3 and $PuCl_3$.

The functions relating vapour pressure to temperature are not provided because insufficient data was found and the functions are not required for any thermohyraulic calculations. The relative vapour pressure is the data of consequence for evaluation purposes. It is assumed that the relative pressures (P_1/P_2) hold regardless of temperature.

Melting Point

The melting point of the salt is an important property because, for a given operating temperature, it determines the margin to solidification or freeze-temperature margin. A requirement was established for the ARE, MSRE and MSBR that the freeze-temperature margin be greater than 100 ^{o}C throughout the plant. Thus the melting point of the salt was limited to 525 ^{o}C [217]. However, some more current designs allow for a higher melting point, the MSFR design for instance has a melting point of 585 ^{o}C [44] and a core-inlet temperature of 620 ^{o}C [218] - resulting in a freeze temperature margin of just 35 ^{o}C.

A lower melting temperature is continually more desirable as a greater freeze-temperature margin allows for potentially reduced material demands as well as simplification of components and systems. However the relationship between the melting temperature and the cost of reactor operation is unclear; it effects the operational difficultly and safety, the cost of which cannot be quantified in a unambiguous manner. The melting temperature, like many other design aspects is quantified according the 'cost setting' method discussed in Section 2.1.

Extensive phase diagram data exists for the salt systems examined. Alkali fluorides, ZrF_4 and BeF_2 salts all have broadly similar melting points and many compositions have a melting point below 525 ^{o}C. Binary eutectic salts have considerably lower melting temperatures than their individual unitary constituents. Ternary eutectic salt mixtures cause comparatively modest melting point depressions in ZrF_4 and BeF_2 salts (approximately 40-60 ^{o}C) over binary eutectics [49].

An FHR is likely to use a binary or ternary eutectic coolant salt, while a thermal spectrum MSR is likely

Table C.3: Vapour pressures of molten salts applicable to reactor applications for which data is available.

Composition	Vapor Pressure at 900 ^oC (mm Hg)[a]	Source
Fluoride Salt Mixtures		
$LiF-NaF-KF$ (46.5-11.5-42)	0.5	[208]
$LiF-BeF_2$ (67-33)	0.05	[212]
$NaF-BeF_2$		[111] [213]
(41-59)	2.41	
(50-50)	1.05	
(60-40)	0.23	
$NaF-ZrF_4$ (50-50)	32	[111]
$NaF-ZrF_4$[b] (57-43)	16	[118]
$LiF-NaF-BeF_2$ (15-58-27)	0.05	est.[c]
$LiF-UF_4-ThF_4$ (77.5-2.5-20)	0.15	[214]
$NaF-ZrF_4-UF_4$ (53-43-4)	21	[118]
$LiF-BeF_2-ThF_4-UF_4$ (71.7-16-12-0.3)	<0.1	[215]
Unitary Salts		
LiF	0.1	[208] [d]
NaF	–[e]	
BeF_2	4.06	[190][f]
ZrF_4	780	[38]
UF_4[g]	0.9	[118]
$LiCl$	7	[208] [203]
$NaCl$	2.5	[208][h]
UCl_3	≈ 0.08	[210]
Chloride Salts for Fast Reactor		
$NaCl-MgCl_2$ (58-42)	<2.5	[208]
$NaCl-PuCl_3$ (85-15)	$\approx NaCl$ (2.5)	[211]

[a] 1 atm = 760 mmHg
[b] Interpolate to find the vapour pressure of $NaF-ZrF_4$, 59.5-40.5.
[c] estimate; low [BeF_2], lower P_{vap} than $NaF-BeF_2$ (60-40)
[d] Original source not listed; 0.40 according to [190]
[e] Melting point $>900^o$C
[f] 12.0 according to [38]
[g] 'Negligably small' for ThF_4, UF_4 and PuF_3
[h] Original source not listed; 1.58 according to [216]

to use a ternary (incl. UF_4) eutectic salt as it is assumed the modest melting point depression does not fully compensate for additional difficulties caused by adding in a fourth unitary salt. Table C.4 provides the melting point of some candidate salt mixes at atmospheric pressure. In many cases it is the minimum melting point for the given system. The boiling point is provided for some salts as well, though the boiling point is often well in excess of the $\approx 700\,^oC$ maximum salt temperature. The non-actinide melting point data in Table C.4 is from [208], actinide compositions from [211] and MSFR fuel ($LiF - ThF_4 - UF_4$) from [44].

C.2 Heat Transport Properties

C.2.1 Heat Capacity

Fluoride salts have relatively high Volumetric Heat Capacities (VHC), similar to that of liquid water. The product of density and heat capacity (C_p) is volumetric heat capacity. The VHC is an important value as it largely corresponds to a fluid's heat transport effectiveness. Since heat capacities are provided, by convention, per unit mass and per unit mass is required for heat transfer calculations, that is the unit of the empirical and predicted heat capacity data.

There exists no fundamental theory used to generally predict heat capacity, though it has been found that the heat capacity per mole of each atom in a mixture is approximately 8 calories/^{o}C or 33.5 J/^{o}C [49]. This empirical finding, credited to Dulong and Petit is [113] as shown below in Equation C.4:

$$C_p = 8 \cdot \sum X_i N_i / \sum X_i M_i \tag{C.4}$$

Where X_i is the mole fraction, N_i atoms per salts constituent i (ex. 3 for BeF_2) and M_i the formula weight of component i. This equation is in strong agreement with most measured heat capacity's of light salts; the Dulong-Petit finding demonstrates a pattern of increasingly under-predicting as the actinide concentration increases.

Similar to much of the other property data for actinide-containing salts, measured values are old. The experimental accuracy of these early measurements is no better than $\pm$ 10%. The measured heat capacity values found are listed in Table C.5, along with the source and Dulong-Petit estimation. The VHCs are listed in Table C.6.

C.2.2 Thermal Conductivity

Thermal conductivity is the most difficult fluid property to measure [49] as evidenced by early thermal conductivities measurements that were recorded as four times greater than values that are now believed to be true [116]. The wrong measurements lead to the postulation that a film resistance must exist and hence a film resistance coefficient was assumed. However it is found that 'the overall heat transfer coefficient of the MSRE heat exchanger did not change during 22,000 h of salt circulation and 13,000 equivalent full-power

Table C.4: Melting point of selected molten salts applicable to reactor applications [208], [211], [44], [103], [219], [175], [220]

Composition	Melting point (^{o}C)	Boiling point (^{o}C)
Ternary Fluoride Salts		
$LiF - NaF - KF$ (46.5-11.5-42)	454	1570
$LiF - NaF - ZrF_4$ (26-37-37)	436	
$LiF - NaF - BeF_2$ (15-58-27)	350	
Binary Fluoride Salts		
$LiF - BeF_2$ (67-33)	460	≈1400
$NaF - BeF_2$ (57-43)	340	≈1400
$NaF - ZrF_4$ (59.5-40.5)	500	≈1350
$NaF - KF$ (40-60)	718	
Actinide Fluoride Salts		
$LiF - ThF_4 - UF_4$ (77.5-20-2.5)	585	
$LiF - BeF_2 - UF_4$ (70-20-10)	525	
$NaF - UF_4$ (78-22)	720	
$NaF - KF - UF_4$ (55-25-20)	630	
Unitary Salts		
LiF	848	1681
NaF	995	1704
$LiCl$	610	1382
$NaCl$	808	1465
Chloride Salts for Fast Reactor		
$NaCl - MgCl$ (58-42)	445	$>$ 1465
$NaCl - PuCl_3$		
(64-36)	453	
(80-20)	640	
$NaCl - UCl_3$		
(65-35)	520	
(80-20)	660	
Other Heat Transport Fluids		
Na	98	883
Light Water	0.0	100

Table C.5: Heat capacities of molten salts applicable to reactor applications.

Composition	Measured Heat Capacity (kJ/KgK)	Dulong-Petit Prediction (kJ/KgK)	Source
Fluoride Salt Mixtures			
$LiF - NaF - KF$ (46.5-11.5-42)	1.9	1.62	[221]
$LiF - BeF_2$ (67-33)	2.42	2.37	[112]
$LiF - NaF - BeF_2$ (15-58-27)	–	1.86	
$LiF - NaF - ZrF_4$ (26-37-37)	–	1.27	
$NaF - BeF_2$ (57-43)	2.2	1.84	[113] [114]
$NaF - ZrF_4$ (50-50)	1.2	1.15	[113] [222]
$NaF - ZrF_4$ (59.5-40.5)	–	1.17	[115]
$NaF - ZrF_4 - UF_4$ (50.0-46-4)	1.1	1.57	[118] (1958)
$LiF - AnF_4$ (77.5-22.5)	1.59	0.99	[117][a]
$LiF - BeF_2 - ThF_4$ (64-18-18)	1.23	1.12	[223] [212]
$LiF - BeF_2 - ThF_4 - UF_4$ (71.7-16-12-0.3)	1.36	1.31	[112]
$NaF - KF - UF_4$ (48.2-26.8-25)	0.96	0.80	[119]
Unitary Salts			
LiF	2.46	2.58	[224] [225]
NaF	1.67	1.60	[224] [225]
$LiCl$	1.47	1.58	[226] [227]
$NaCl$	1.18	1.15	[226] [228]
Chloride Salts for Fast Reactor			
$NaCl - MgCl_2$ (58-42)	1.08	1.10[b]	[229]
Actinide, Mg mix-$UCl_3 - PuCl_3$			[205]
(82-12.6-5.4)	–	0.96	
(70-0-30)	–	0.84	
(50-27.5-22.5)	–	0.71	
Other Heat Transport Fluids			
Na	1.28	–	[219]
Light Water	4.19	–	[230]

[a] Extrapolated, recent, seemingly accurate
[b] Dulong-Petit

Table C.6: Volumetric heat capacities of molten salts applicable to reactor applications.

Composition	Density 700 ^{o}C $(g/cm^3)^{abcd}$	Heat Capacity (KJ/KgK)	Volumetric Heat Capacity (MJ/m^3K)
Fluoride Salt Mixtures			
$LiF-NaF-KF$ (46.5-11.5-42)	1.99	1.9	3.77
$LiF-BeF_2$ (67-33)	1.92	2.42	4.64
$LiF-NaF-BeF_2$ (15-58-27)	2.05	1.86	3.81
$LiF-NaF-ZrF_4$ (26-37-37)	2.98	1.27	3.78
$NaF-BeF_2$ (57-43)	2.05	2.2	4.51
$NaF-ZrF_4$ (50-50)	3.14	1.2	7.23
$NaF-ZrF_4$ (59.5-40.5)	2.96	1.2	3.55
$NaF-ZrF_4-UF_4$ (50.0-46-4)	2.65	1.57	4.17
$LiF-AnF_4$ (77.5-22.5)	4.38	1.59	6.97
$LiF-BeF_2-ThF_4$ (64-18-18)	3.68	1.23	4.53
$LiF-BeF_2-ThF_4-UF_4$ (71.7-16-12-0.3)	3.25	1.36	4.42
$NaF-KF-UF_4$ (48.2-26.8-25)	3.98	0.96	3.82
Unitary Salts			
LiF	1.88	2.46	4.62
NaF	2.14	1.67	3.57
$LiCl$	1.46	1.47	2.15
$NaCl$	1.61	1.18	1.90
Chloride Salts for Fast Reactor			
$NaCl-MgCl_2$ (58-42)	1.65	1.08	1.78
Actinide, Mg mix-UCl_3-PuCl_3			
(82-12.6-5.4)	2.17	0.96	2.08
(70-0-30)	2.57	0.84	2.16
(50-27.5-22.5)	3.31	0.71	2.35
Other Heat Transport Fluids			
Na	0.83[e]	1.28	1.06
Light Water	0.75	4.19	3.14

[a] If $T_{melt} > 700^oC$ density is extrapolated density from T_{melt}

[b] ZrF_4 density is calculated using the additive principle and ARE density data [103]

[c] PuF_3 density original source could not be found; verified with ionic radii calculation [231] [232]

[d] $PuCl_3$ same estimation technique used in [233]

[e] Source: [219]

hours of operation, thus indicating no buildup of scale and no evidence of gas filming' [234].

Thermal conductivity is primarily driven by vibration rather than a diffusion mechanism. The comparatively high thermal diffusivity compared to mass diffusivity and the high Lorenz number ($\propto$ thermal/electrical conductivity) indicate that the contribution of diffusion conduction is less than 5%. As a result molten salts have thermal conductivity values similar to solid dielectrics (insolators) [235].

The paper [235] contains what is considered the 'most successful' [49] model for predicting the thermal conductivity of molten salts, see Equation C.5.

$$k\,[W/m-K] = 0.119 \cdot T_M^{0.5} \cdot V_m^{0.667}/(M/n)^{1.167} \tag{C.5}$$

Where T_m is the melting point (K), V_m the molar volume ($cm^3/mole$), M the average formula weight of the salt ($\sum X_i M_i$) and n the number of discrete ions (2 for simple salts like $NaCl$).

The model, here on referred to as Rao-Turnbull, predicts that increasingly complex salt mixtures exhibit lower conductivity as a result of disruption to the vibrational modes. Compared to many other physical properties thermal conductivity is only weakly a function of temperature and uncertainty is high so recommended values, even for well-studied salts such as FLiBe, are not a function of temperature [115]. The thermal conductivities of some candidate salts as well as liquid sodium and water are provided in Table C.7. The data from Table C.7 suggests that the Rao-Turnbull predictive method is only effective for light element mixes and still produces considerable error.

Table C.7: Thermal conductivities of molten salts applicable to reactor applications.

Composition	Measured Conductivity (W/m-K)	Rao-Turnbull Prediction (W/m-K)	Temp (^{o}C)	Source
Fluoride Salt Mixtures				
$LiF-NaF-KF$ (46.5-11.5-42)	0.90	0.68	700	[236]
$LiF-BeF_2$ (67-33)	1.1	0.79	700	[112]
$LiF-NaF-BeF_2$ (15-58-27)	–	0.66	700	
$LiF-NaF-ZrF_4$ (26-37-37)	–	0.36	700	[115][a]
$NaF-BeF_2$ (57-43)	2.4	0.58	700	[113] (1965)
$NaF-ZrF_4$ (59.5-40.5)	–	0.36	700	[115][b]
$NaF-ZrF_4-UF_4$ (50-46-4)	2.3	0.60	–	[118] (1958)
$LiF-AnF_4$ (77.5-22.5)	1.01[c]	0.29	700	[117]
$LiF-BeF_2-ThF_4$ (64-18-18)	0.97	0.54	700	[223] [212]
$NaF-KF-UF_4$ (46.5-26-27.5)	1.73	0.31	700	[118]
Unitary Salts				
LiF	1.35	1.18	848	[236]
NaF	1.1	0.94	995	[236]
$LiCl$	0.78	0.92	610	[236]
$NaCl$	0.88	0.85	808	[112][d]
Chloride Salts for Fast Reactor				
$NaCl-MgCl_2$ (58-42)	–	0.43	700	[208]
Actinide, Mg	–	1.02	–	[205][e]
mix-UCl_3-PuCl_3				
(82-12.6-5.4)	–	1.02	–	
(70-0-30)	–	0.86	–	
(50-27.5-22.5)	–	0.66	–	
Other Heat Transport Fluids				
Na	69	–	500	[219]
Light Water	0.545	–	300	[162]

[a] Khokhlov value recommended instead, 0.53 W/m-K

[b] Khokhlov value recommended instead, 0.49 W/m-K

[c] Recent, seemingly accurate

[d] Could be 0.70, contradiction with [236]

[e] The thermal conductivity was estimated using an empirical correlation [237] that produced satisfactory agreement with $LiCl-KCl$ experimental data [238]

Appendix D

Additional Appendices

D.1 Thorium

Molten salt reactors are often associated with the use of Thorium fuel. Thorium is a fertile fuel, upon the absorption of a neutron it transmutes into ^{233}Pa which decays into fissile ^{233}U.

Protactinium-233 has a substantial half life of 27 days and a cross section of 21.9 barns at 0.0718 eV, approximately the mode neutron energy of a Maxwellian distribution at 650 ^{o}C. For comparison the cross section of ^{233}U is 28.7 barns at 0.0739 eV [108]. An absorption by ^{233}Pa is considered parasitic, so a reactor with thorium fuel would benefit from separating the ^{233}Pa from areas of high flux.

The thermal fission neutron yield of ^{233}U is 2.50 so a maximum of 20% (1-2/2.50) of fission neutrons could escape or be lost to parasitic absorption in order to achieve net breeding in a thermal reactor on the ^{233}Th-^{233}Pa-^{233}U cycle. The Molten Salt Breeder Reactor design was intended to do just that. It required a fertile blanket salt and extensive salt processing to achieve the necessary neutronic efficiency, see Section 1.6.4. At that time nuclear fuel prices and supply concerns drove the motivation to achieve net breeding [23]. Today capital costs are more and uranium supply less of a concern. This is reflected in modern designs such as the IMSR that do not intend to use a breeding cycle.

It would be even more difficult and perhaps prove impossible to construct a salt-cooled, thermal-spectrum, thorium-fuelled breeder. On-line reprocessing would likely prove more difficult, many designs would require core structural material and many designs would have a less than ideal fuel-moderator separation. The MSFR design has demonstrated that the ^{233}Th-^{233}Pa-^{233}U cycle is viable option in molten salt fast spectrum designs. The simplicity of the core, with no moderator or structural material, allows for thorium fuel with limited processing required.

The viability of thorium is heavily dependent on the MSR design type. The use of some thorium (breeding ratio < 1) is likely impractical for thermal-spectrum salt-cooled MSRs but an option worth exploring in thermal-spectrum circulating-fuel MSRs. In fast spectrum MSR designs a thorium breeder a is promising

fuel cycle, however partially thorium fuelled designs are likely worth exploring as well.

D.2 Corrosion

The scope of this work does not allow for a detailed, comprehensive explanation of all aspects of corrosion. Corrosion of metals by fluoride molten salts is the primary focus of this section; interactions with graphite are briefly explored in the upcoming section, D.2.4. It is assumed that the structural material interacting with the salt is Ni-based and contains Cr, particular Hastelloy N.

D.2.1 Metallic Corrosion Mechanisms

The purpose of this section is to outline metallic corrosion mechanisms that serve as a basis from which estimates of the differences in corrosion rates of various fluoride salt compositions can be made. The cost of managing corrosion in the evaluation program is a function of the expected corrosion rate. Corrosion of metals by molten salts occurs through fundamentally different mechanisms than corrosion by oxidation in conventional media like air or water [239]. Most standard metal alloys undergo an oxidation reaction on the surface that produces a stable, protective layer. This occurs in molten salt melts as well, however, unlike the more conventional oxidizing media, the resultant protective surface film tends to be completely soluble in the molten salt [240]. Without a protective surface coating, the corrosion depends on dissolution mechanics and thermodynamics.

Basic salt chemistry [239]

The interactions between component ions/atoms can differ considerably depending on the molten salt composition. In alkali fluorides the thermal energy of each constituent atom is high enough to overcome the coulomb forces holding atoms together above the melting point. As a result, the salt disassociates into ions. For instance, LiF becomes:

$$LiF \rightarrow Li^{+} + F^{-} \tag{D.1}$$

Other molten salts such as BeF_2 are by contrast non-ionic. Though the Be-F bond is considered ionic, molten BeF_2 forms glassy networks of chained Be and F atoms that results in high viscosity [241]. Mixtures of simple salts like either LiF or BeF_2 can result is yet more chemical structures/behaviours. The commonly proposed salt composition, $2^7Li - BeF_2$, results in the formation of BeF_4^{2-}:

$$BeF_2 + 2Li^{+} + 2F^{-} \rightarrow BeF_4^{2-} + 2Li^{+} \tag{D.2}$$

The glassy network of BeF_2 is thus broken and the viscosity significantly reduced.

Alkali fluorides, such as LiF-NaF-KF and simple LiF aren't able to capture free F- ions, see Equation D.2. The consequence of these free mobile fluorides is greater corrosion [239].

Reversiable reactions and redox potential [239]

Molten salt corrosion mechanics are governed by thermodynamics that constantly push the system toward equilibrium. The reactions involved in the fundamental corrosion mechanics of molten salts are reversible. A reversible reaction is a reaction which results in a mixture of products and reactants – it never reaches full completion, the share of both reactants and products is > 0. It can be understood through Gibbs free energies. Gibbs free energy is a measure of a substance's ability to do reversible work. The more negative a metal fluoride's free energy of formation ($\Delta G^0_{MF_x}$) the more prone the metal is to attack in fluoride salt.

ΔG is the total change in Gibbs energies of all substances in the reaction mixture. It varies continuously as the composition changes, reaching zero at equilibrium. Whereas the standard Gibbs energy change, ΔG^o, has a single value for a particular reaction at a given temperature and pressure. In the ideal case of gaseous chemical reactions, the reaction is represented as:

$$aA + bB \iff cC + dD \tag{D.3}$$

And,

$$\Delta G = G^o + RTInP_r \tag{D.4}$$

P_r is the ratio of partial pressures or concentrations. At equilibrium P_r is equal to the equilibrium constant, K_{eq}. Rearranging Equation D.4, subbing in K_{eq} for P_r and making $G = 0$:

$$\begin{aligned} K_{eq} &= \frac{[C]^c[D]^d}{[A]^a[B]^b} \\ K_{eq} &= exp(\frac{-G^o}{RT}) \end{aligned} \tag{D.5}$$

The bracketed constants (ex. $[A]$) are activities (pressures, in the ideal case, or concentrations). When K_{eq} is greater than 1, the reaction favours the products and when K_{eq} is less than one, the reactants. The value of the activities and consequentially K_{eq} is a function of both Gibbs free energies of formation, G^0, the fluorine gas constant, R, and Temperature, T.

However, using the ratio of mole fractions or partial pressures is not reliable in ionic fluids (molten salts) [242], instead an activity coefficient must be employed:

$$A_i = x_i\gamma_i \tag{D.6}$$

Where A_i is the activity of the substance, x_i is the mole fraction and γ_i is the activity coefficient. The activity coefficients are assumed to be equal to 1, this will be subsequently explained.

Combining the partial reactions of fluorine with the salt cations and fluorine with the metal container, as

Table D.1: Gibbs free energies of fluoride salts relevant to corrosion [50].

Fluoride	ΔG @ 700^{o}C	$\Delta G/X, MF_x$
LiF	-1046	-1046
NaF	-948	-948
KF	-939	-939
BeF_2	-878	-439
ZrF_4	-1591	-398
UF_4	-1636	-409
NiF_2	-506	-253
CrF_2	-648	-324

well as using the activity coefficient assumption, K_{eq} can be simplified to:

$$K_{eq} = exp(\frac{-(\Delta G_{MF_x} - \Delta G_{CF_x})}{RT}) \tag{D.7}$$

Where G_{CF} and G_{MF} are the Gibbs free energy of the cation-fluoride and structural metal-fluoride ionic bonds respectively. $\Delta G_{MF} - \Delta G_{CF}$ is > 100 KJ/mol so the K_{eq} is an extremely low value. Even if the activity coefficient caused an increase in K_{eq} by 2 or 3 orders of magnitude, the reaction will still strongly favour the reactants (salt cations). Thus a pure salt is effectively non-corrosive. Table D.1 lists the Gibbs free energies of molten salts relevant to corrosion analysis.

Fundamentally, the minimum rate of corrosion, as driven by thermodynamics and ignoring the effect of irradition and impurities, is dictated by the redox (reduction-oxidation) potential. It is the tendency of a substance to acquire or lose electrons when a new element is introduced [243]. In a molten salt-structural metal interaction the molten salt is composed of components that are strongly electronegative (flourine) and thus have a strong negative redox potential. Fluorine is the oxidizing agent while the metal is the reducing agent.

The fluorine potential of a salt is not necessarily constant; it can be controlled by the introduction/removal of redox agents. It is commonly controlled by either hydrofluorination or metal additions. The addition of a reducing agent, typically a metal, has the effect of reducing the magnitude of the fluorine potential [243] [244]. Hydrofluorination is typically achieved through the process of sparging/bubbing HF/H_2 gas through the salt. The greater the $[HF]$ and lower the $[H^2]$, the greater the fluorine potential and higher the rate of corrosion. Chromium, is the least thermodynamically stable structural material and will corrode most significantly by the reaction:

$$[Cr] + 2[HF] \iff [CrF_2] + [H^2] \tag{D.8}$$

Using Equation D.5 which ignores activity coefficients,

$$K_{eq} = \frac{[Cr][HF]^2}{[CrF_2][H^2]} = exp(-G^o/RT) \tag{D.9}$$

It can be rearranged to solve for $[CrF_2]$,

$$[CrF_2] = \frac{[Cr][HF]^2}{[H^2]} \cdot exp(\frac{-G^o}{RT}) \tag{D.10}$$

By controlling the amounts of $[HF]$ and $[H^2]$, the rate of Cr corrosion can be controlled. The objective is not simply to minimize $[CrF_2]$ production since there exists a trade-off between corrosion and solubility. A more negative redox potential increases the rate of corrosion but also increases solubility. Thus it is especially important to properly manage the redox potential in circulating fuel reactor with dissolved fission products.

D.2.2 Impurity Driven Corrosion

Over the long-term corrosion tends to be dominated by thermodynamics, though impurities can have a significant pernicious effect. Moisture, hydroxides and oxides along with metallic impurities can trigger corrosion reactions [245]. These impurities can arise from manufacture/storage, neutron transmutation of salt constituents or fission products. Impurity-driven corrosion brings about an initial rapid increase of the maximum depth of attack in Inconel 600, a Ni-based high temperature alloy. However, the depth of attack increases linearly at a relatively low rate after using up all impurities [245] [246].

Moisture is 'perhaps the most deleterious contaminant in molten fluoride salts' since it triggers reactions with metal fluorides that generate highly corrosive gaseous hydrofluoric acid [245]. Equations D.11 and D.12 describe the generation of hydrofluoric acid when the alkali metal fluoride (MF) reacts with moisture [247] [245].

$$MF(l) + H_2O(g) = MOH(l) + HF(g) \tag{D.11}$$

$$2MF(l) + H_2O(l) = M_2O(l) + 2HF(g) \tag{D.12}$$

The generated oxides such as BeO, Li_2O, ZrO_2 or UO_2 tend to precipitate out of solution due their high melting point, 2530^oC, 1700^oC, 2715^oC and 2865^oC [248].

The fission product that has attracted the most concern with respect to corrosion is Tellurium. According to experiments performed for the MSRP [249], the sulfur, selenium, and tellurium family were found to trigger intergranular cracks in Hastelloy N, and tellurium penetrated metal alloys to depths approximately equal to those of the cracks. The cracking was found to become more severe as the concentration of tellurium increased, demonstrating the detrimental effect tellurium has on the embrittlement of Hastelloy N [249].

D.2.3 Irradiation Effects on Corrosion

Corrosion of structural alloys in molten fluoride salt has been found to accelerate in a nuclear reactor environment. Zheng et al. find that in a reactor environment, 'the strong radiation fields can exacerbate alloy

Table D.2: Equilibrium level of dissolved Cr in various fuel salts [49].

Salt Mixture	Mole % ZrF_4 or BeF_2	[Cr] at 600 C, ppm	[Cr] at 800 C, ppm
$FLiNaK$	0	1100	2700
$LiF-ZrF_4$	48	2900	3900
$NaF-ZrF_4$	50	2300	2550
$NaF-ZrF_4$	41	975	1050
$NaF-LiF-ZrF_4$	23	550	750
$LiF-BeF_2$	48	1470	2260

corrosion in molten salt, but the mechanisms are not conclusively understood.' [250] [251]. The magnitude of the accelerated corrosion varies considerably depending on the structural alloy and the whether graphite is present; from $\approx 2x-10x$ greater weight change in a reactor environment.

D.2.4 Molten Salt Interaction With Graphite

Graphite does not react chemically to any appreciable extent, nor it is wetted by molten fluoride salts. However graphite can be somewhat permeated by fuel salt and Fission Products (FP) can diffuse into or deposit onto graphite. Nonetheless the overriding concern in using graphite remains the dimentional changes and stresses associated with irradiation; a comprehensive 1969 MSRP paper on Graphite Behavior [181] clearly focuses on that issue the most.

Once again, the MSRP provides the most revelant and comprehensive experimental data. MSRE tests indicate that fuel salt permeation is of little to no concern. Initially, the magnitude of permeation was unclear as evidenced by graphite-fuel temperature difference calculations that relied on knowing fuel permeation as a volume percentage of graphite [23]. The assumed permeation range was 0.0 - 2.0 %. Later, in-pile tests found 'no permeation of fuel salt into the graphite' [181]. Gaseous diffusion can have a considerable impact, it was assumed that the ^{135}Xe poison fraction of the MSBR would be 2.25% with uncoated graphite. It appears that it can be reduced to 0.5% with a pyrolytic carbon seal on the graphite surfaces. It was found that 99% of deposited noble-metal FP were within 5 mm of graphite surfaces, while 'relatively heavy' deposits of noble-metal FP were observed on the Hastelloy N specimens [181].

D.2.5 Experimental Data

The experimental results are generally old, in short supply and conducted under inconsistent conditions (duration, temperature). The experimental data source used in evaluations is from a comparison of the performance of various fluoride salts provided in an ORNL report on Molten Salt Coolants for the AHTR [49]. Useful performance results are shown in Table D.2, the UF_4 concentration of each mix is 1.5 - 4.1%. Note that very basic ($FLiNaK$) and very acidic ($LiF-ZrF_4$) salts, in Table D.2, showed the worst performance.

D.2.6 Fluoride Salt Corrosion With Various Metal Alloys

The objective of this section is to compare the corrosion of various structural metal alloys due to molten salt. The most widely studied and commonly proposed structural material for use in MSRs is Hastelloy N (INOR-8). It show excellent compatibility with common fluoride salt constituents, including LiF, BeF_2, ThF_4 and UF_4. These fluorides are much more stable towards Hastelloy N than structural metal fluorides such as NiF_2, BeF_2 and CrF_2 [234], pg. 198.

Corrosion is dominated by the oxidation and removal of Chromium, the least noble constituent [38]. A major motivation for the addition of Chromium is to increase corrosion resistance in air [177]. Stainless steels are more susceptible to corrosion than Ni-base alloys in large part because they have a greater chromium concentration [252].

Inconel vs INOR-8 (Hastelloy N)

The two structural materials to have gone extensive corrosion testing over the development of the MSRP are Inconol and INOR-8, now called Hastelloy-N. The alloy with better corrosion resistance clearly proved to be Hastelloy-N, thus the MSRE was constructed with that material [23]. The major constituents of the Ni-based INOR-8 (Hastelloy N) by weight percent are Mo 15-18, Cr 6-8, Fe 5 (max) [23] and of Inconel, Cr 14-17, Fe 6-10, Mo 0 [253].

Stainless Steel

As mentioned earlier, Stainless Steels (SS) have a higher chromium content then commonly-proposed nickel-base alloys such as Hastelloy N. The corrosion resistance of 316 SS, common to the nuclear industry, has been tested with the fluoride salt $FLiNaK$; at 815 C and over a 500 hr duration the corrosion depth was 4 mm [49]. A similar test on INOR-8 (Hastelloy N) but with a 3048 hr duration resulted in only 0.1 mm of corrosion [254].

A study on the corrosion of 316 SS in 700 ^oC $2LiF - BeF_2$ salt found a corrosion rate of 17.1 $\mu m/yr$ and 31.2 $\mu m/yr$ in the presence of graphite [255]. In comparison, Hastelloy N was found to have a corrosion rate of approximately 25 $\mu m/yr$ at 704 ^oC [256].

The lack of experimental corrosion test results that can provide a useful comparison of SSs to nickel-base alloys prevents clear conclusions from being made. Hastelloy N performed significantly better than 316 SS with $FLiNaK$ salt at 815 C but the performance was approximately equal with $2LiF - BeF_2$ near 700 C. Taking into account these test results and that SSs are considered more susceptible to corrosion because of the higher chromium content, a safe conclusion would be: Stainless steels are generally equal or more susceptible to corrosion than nickel-base alloys but results vary considerably depending on the temperature, salt and purity in question [252].

Molybdenum Alloys

The molybdenum base alloy, TZM (0.5 Ti, 0.08 Zr, 0.02% C) was studied at ORNL the late 1960's and early 1970's for use in vacuum distillation. The TZM alloy showed a very small amount of attack by the fluoride salt $LiF - BeF_2 - ThF_4 - UF_4$, 68-20-11.7-0.3 at 1100 ^oC for 1011 hours. Corrosion manifested itself as leaching of Ti and possibly Zr. The report concluded, that 'the magnitude and mechanism of corrosion indicate no serious problems for long-term use of TZM in the vacuum distillation processing scheme for the MSBR.' The TZM alloy was selected over pure molybdenum because it is stronger and usually more fabricable [257].

As TZM corrodes and becomes more similar in composition to pure Mo, the strength decreases but the corrosion resistance increases. TZM is easier to work than unalloyed Mo and is strengthened by cold working and precipitation hardening. The ultimate tensile strength of TZM is 'double or triple' that of unalloyed molybdenum. The greatest concern associated with the use of TZM appears to be the degradation of strength properties due to leaching of Ti and Zr over time. The results from this test suggest that TZM has greater corrosion resistance at very high temperatures than any other metal alloy tested for compatibility with molten salts [257].

D.3 Processing

D.3.1 Processing Introduction

Processing refers to purposeful processes undertaken to alter the composition of the fuel salt. The major objectives of processing typically are to remove fission products, regulate the redox potential and add/remove fuel. As fissions occur, fuel is consumed, fission products are created and there becomes a discrepancy between the fluoride ions required and fluoride ions available to oxidize fuel, fission products and carrier cations. Fluoride ion levels, or redox potential, is maintained at a level that ensures adequate solubility and that corrosion is necessarily mitigated. The primary purpose of FP removal is to minimize parasitic absorptions. All else being equal, breeder reactors require more fuel salt processing and may require a separate fertile blanket that itself requires processing (particularly the removal of protactinium) as well. In a burner reactor fuel management generally just involves the addition (continual or batch) of fuel over time.

The processing rate - the inverse of cycle time - refers to the average rate at which an element is removed from the reactor, it is measured as a fraction of element removed per second (s^{-1}).

D.3.2 Fission Product Processing

The deleterious effects and the potential methods of processing vary widely among the FPs. The parasitic absorption and corrosion effects of the FP should be considered when selecting a processing scheme, yet it is beyond scope to detail the effects for all FPs of consequence. Many elements have similar chemical behaviour and so can be processed (removed from the salt) using the same or similar methods. The groupings

listed below are sufficiently chemically similar such that it is safe to assume that they can be processed though the same mechanism(s).

- Noble Gases: It can be assumed that all past and current MSR designs employ some mechanism to remove noble gas fission products (typically bubbling of an inert gas). – ex. MSRE [23], MSFR [258] and IMSR [259]. The reactivity effect (poison fraction) of noble gases is dominated by a single isotope, ^{135}Xe.

- Noble Metals: The elements included in this category are the noble metals of appreciable quantities: Mo, Ru, Tc, Te, Nb [35]. The noble metals are a separate category in this analysis because they are much less soluble and more volatile than other FPs. The two properties are linked, the primary reason noble metals are volatile seems to be because they are weakly soluble [260]. This weak attachment to the salt (i.e. solubility) allows noble metal fluorides to escape solution through plating out on surfaces and/or diffusing into the sparging gas.

- Rare-earths: Rare-earth metals (or lanthanides + Sc + Y) can be thought of as moderately soluble. In the original MSBR design the first processing step, fluorination, would remove volatile FP but leave most other FP, including rare-earths, in solution. Rare earths were to be removed through precipitation and filtration in a HF 90% solution [23]. In the meantime a new and more effective method of FP removal was developed using liquid bismuth. This method involved a chemical reduction of materials from the fuel salt to the bismuth [39]. The quickest processing cycle time of rare-earth removal in the MSBR as of 1967 was a 10-day cycle - though the design was in flux and cycle times were generally not confirmed [35]. The MSFR design removes rare-earths on a 450-day cycle [261].

- Stable, Soluble Fluorides: A 'very high fraction of isotopes' including ^{140}Ba, ^{141}Ce, and ^{91}Sr that form stable soluble fluorides are present in circulating fuel [35]. These isotopes are the most difficult to extract from the salt. They can either be removed though vacuum distillation or by discarding the salt. The MSBR planned to use vacuum distillation but the design is unique in the tough demands placed on ensuring a high neutron economy. Other reactor designs such as the MSFR and IMSR do not propose vacuum distillation.

Tritium

The appropriate level of attention to tritium removal is unclear. Studies at ORNL on the MSRE were conducted on tritium solubility, diffusion through metal and methods of tritium removal. Tritium solubility is approximately $1.5 \cdot 10^{-4}$ mole solute/solvent at operating temperature and pressure. Tritium can be removed through the addition of HF and sparging, as well there is high uptake in graphite [262] [113]. The primary source of tritium production is the irradiation of ^{6}Li and it is therefore only a concern in MSRs that use lithium in the carrier/coolant salt.