University of Illinois Urbana-Champaign

Final Scientific/Technical Report

MEITNER

Award:	DE-AR0000983
Sponsoring Agency	USDOE, Advanced Research Project Agency – Energy (ARPA-E)
Lead Recipient:	University of Illinois Urbana Champaign
Project Team Members	Jiaqi Chen, Caleb S. Brooks Andrei Rykhlevskii, Mehmet Türkmen, Kathryn D. Huff Alvin Lee, Tomasz Kozlowski Zhen Li, Brent J Heuser, James F Stubbins
Project Title:	Enabling Load Following Capability in the Transatomic Power MSR
Program Director:	Dr. Robert Ledoux
Principal Investigators:	Dr. Caleb Brooks, Dr. Kathryn Huff, Dr. Tomasz Kozlowski, Dr. Brent Heuser, Dr. James Stubbins
Contract Administrator:	Advanced Research Projects Agency - Energy
Date of Report:	12/31/2021
Reporting Period:	06/27/2019 – 09/30/2021

The information, data, or work presented herein was funded in part by the Advanced Research Projects Agency-Energy (ARPA-E), U.S. Department of Energy, under Award Number DE-AR0000983. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

Please check the appropriate box:

☒This Report contains no Protected Data.

☐This Report contains Protected Data and the award allows data to be marked as protected.

☐This Report contains SBIR/STTR Data and the award allows data to be marked as SBIR data.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

PUBLIC EXECETIVE SUMMARY

This project is dedicated towards designing a fuel processing system that enables liquid-fueled molten salt reactors (MSR) to load follow by removing the dissolved xenon in the fuel salt. As one of the Gen-IV nuclear reactor concepts, the molten salt reactor receives increasing development interests in the recent years. One distinguishing feature of the liquid-fueled molten salt reactor is its improved ability to operate in a load-following mode by including the unique online fission product removal system. Load-following means that the reactor changes its power output based on the demand on the grid. Most of the current operating nuclear reactors have limited load-following ability and operate as the base load on the grid. Due to the rapid increase of solar energy, the requirement on load-following capacity is significantly increased because of the varying power output of the solar panels, yet the traditional load-following capacity is expected to decrease as the decarbonization of the grid continues. Therefore, the ability to perform load-following operation for the nuclear reactors will greatly enhance the resilience of the grid and make nuclear energy more economically competitive. This load-following feature is included in many commercial molten salt reactor designs, such as the designs by Transatomic Power, Terrestrial Energy, and Flibe Energy. Unfortunately, detailed analysis of the fuel processing system for commercial scale MSRs is still lacking, as well as how the fuel processing quantitively impacts the load-following operation. Moreover, experimental data for many of the underlying physics of fuel processing is limited. This project aims to pave the way for the fuel processing technology to advance to the commercial stage by performing combined experimental and simulation research.

During the project period, four interconnected aspects of the development of the fuel processing system in liquid-fueled molten salt reactors are investigated. These aspects are the simulation and analysis of the fission product removal system, the fuel cycle simulation, the coupled reactor neutronics and thermal hydraulics transient simulation, and the gaseous fission product removal experiment. Multiphase CFD simulations are performed for components of the processing systems, and simplified air-water experiments are carried out to provide validation data. It is concluded that the CFD simulation can satisfactorily predict the system level performance of the components, and engineering models are constructed based on this success. Fuel cycle analysis is performed for two representative MSR design, the MSBR and the Transatomic Power MSR. Open-source code SaltProc is developed to incorporate the unique fuel processing system of the MSRs. It is concluded that the removal of xenon is essential for load-following operation in thermal spectrum MSR and Molten Salt Breeder Reactor. For the Transatomic Power MSR, the xenon poisoning effect is negligible due to its relatively fast neutron spectrum, though the overall fuel cycle economics still benefits from the removal of xenon. Coupled reactor neutronics and thermal hydraulics transient simulation is performed specifically for the Transatomic Power MSR. It is concluded that the reactor core design could perform power ramping fast enough to satisfy load-following operation. Combining the findings from each aspect, it is concluded that the load-following operation of a thermal neutron MSR is dependent upon the removal of xenon, which could be achieved for a commercial sized reactor using continuous inert gas sparging in a separate system with reasonable dimensions.

ACKNOWLEDGEMENTS

The information, data, or work presented herein was funded in part by the Advanced Research Projects Agency-Energy (ARPA-E), U.S. Department of Energy, under Award Number DE-AR0000983. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

This project is supported by resource team at Argonne National Laboratory. The contributors are M.A. Rose, E. Wu, T. Lichtenstein, M. A. Williamson, and J. Krueger. This work also made use Idaho National Laboratory computing resources which are supported by the Office of Nuclear Energy of the U.S. Department of Energy and the Nuclear Science User Facilities under Contract No. DE-AC07-05ID14517.

RESEARCHER CREDIT SUMMARY

The simulation and analysis of the xenon removal system is carried out by Mr. Jiaqi Chen and Dr. Caleb Brooks, as detailed in Chapter 2. The fuel cycle simulation and analysis is performed by Dr. Andrei Rykhlevskii, Dr. Mehmet Türkmen and Dr. Kathryn Huff, and presented in Chapter 3. The coupled reactor core simulation is performed by Mr. Alvin Lee and Dr. Tomasz Kozlowski is provided in Chapter 4. The experimental study of the inert gas sparging in molten salt is performed by Dr. Zhen Li, Dr. Brent Heuser, and Dr. James Stubbins is summarized in Appendix A. Special thanks are given to Mr. Logan Crevelt for helping with the challenging experiment with molten salt.

ACCOMPLISHMENTS AND OBJECTIVES

This award allowed the project team at University of Illinois Urbana Champaign to demonstrate a number of key objectives. The focus of the project was on designing a fuel processing system that enables liquid-fueled MSRs to load follow.

A number of tasks and milestones were laid out in Attachment 3 of the original contract, the Technical Milestones and Deliverables, at the beginning of the project. The actual performance against the stated milestones is summarized here:

Table 1: Key milestones and deliverables.

Task/Milestone Title	Milestones and Deliverables
Task 1: Simulation and Analysis of the Xenon Removal System M1.1: CFD Model Release M1.2: CFD Sensitivity Analysis	M1.1: CFD model using existing data and new data from experiments shows sparging system can remove Xe at rate that is fast enough to allow TAP reactor to execute ≥ ±10% capacity/min power changes. **Actual Performance:** (12/28/2019) The CFD models are developed and validated against additional experiments using the same apparatus geometry for Task 4. The CFD model alone is not enough to determine whether the TAP rector could do power ramping at ±10% capacity/min.
	M1.2: CFD model used to create point design for MSR sparging system with allowable design margins sufficient to support the target reactor ramp rate of ≥ ±10% capacity/min. **Actual Performance:** (06/28/2021) The CFD models are used to simulate a prototypical xenon removal system. Regression models are created to predict the performance of the removal system. The TAP rector design does not benefit from removing xenon in terms of load following as suggested by results from Task 2. Though for thermal neutron reactor as MSBR, the efficiency of the xenon removal system will essentially eliminate the reactivity feedback preventing the load-following operation.
Task 2: Fuel Cycle Simulation and Analysis M2.1: Demonstrate SaltProc M2.2: Software and Database Release M2.3: SaltProc Sensitivity Analysis	M2.1: Initial demonstration of fuel cycle simulation package working together with Monte Carlo to complete full core TAP reactor depletion calculation. SaltProc will use separations efficiencies and dynamics based on work in Task 1 and will be coupled with Serpent 2 where Monte Carlo results will be done to <10% relative error accuracy. **Actual Performance:** (06/28/2019) Full core, full lifetime TAP reactor depletion calculations have been conducted. The input from Task 1 is used to set the realistic parameters in the fuel cycle simulation with SaltProc. Uncertainty analysis is performed, and the error is within the stated 10% accuracy.
	M2.2: Release a new SaltProc python package version incorporating user-parameterized components in the fuel salt

	processing design. Release a compatible .json database for use with the SaltProc python package. This citable data object will incorporate historical data regarding fuel salt processing component performance as well as parameters from experiments in Task 4. **Actual Performance:** (03/28/2020) The new SaltProc package is developed to incorporate the reduced order model developed by Peebles. Some inputs are taken from Task 1 to facilitate the removal system model. This version of code uses physical model instead of fixed removal efficiency compared with previous SaltProc and many other existing codes. However, input from Task 4 is not included due to experimental issues.
	M2.3: Recommend fuel processing system design (and feasible design space) that can achieve Xe removal for ≥ ±10% capacity/min reactor power ramping; approved by ARPA-E. **Actual Performance:** (03/28/2021) Sensitivity and design study are performed using the xenon migration model by Peebles and bubble separator model developed in Task 1. The criterion for sufficient removal efficiency is determined based on fuel cycle analysis using SaltProc and Serpent. However, large uncertainties for the recommended design exist due to the uncertainties of physical and chemical properties. The analysis also assumes batch-mode operation instead of continuous xenon removal.
Task 3: Coupled Reactor Core Simulation M3.0: Define operational conditions to be investigated M 3.1: Operation envelope for load-following operation (target ≥10% capacity/min) established M 3.2: Thermal hydraulics analysis of load-following operation	M3.0: Define expected operating conditions, including power levels, control rod positions, and operating transient and postulated accidents, to be investigated. Define the constraints that the core must meet, including maintaining criticality and sufficient shutdown margin while minimizing local power peaking. **Actual Performance:** (09/28/2018) The expected operating conditions, including power levels, control rod positions, and operating transient and postulated accidents, to be investigated are defined. Additionally, the constraints that the core must meet, including maintaining criticality and sufficient shutdown margin while minimizing local power peaking.
	M3.1: Establish operational envelope for load-following operation (target ≥ ±10% capacity/min) to meet constraints in M 3.0.**Actual Performance:** (09/28/2019) From the simulations, the TAP MSR was shown to be able to perform load-following operations at a fairly high rate of about ±220% of the rated power per minute without exceeding the limits identified in M3.0. The main constraint for further increase in reactor power up ramp was the exceeding of material temperature limits during the brief overpower expected during such power up ramp. On the other hand, the asymptotic nature

	of the decrease in reactor power during power down ramp appears to be the natural limit for the power down ramp rate.
	M3.2: Demonstrate that, with the designed sparging system, the core maintains sufficient flow while minimizing local temperature for all expected operating conditions defined in M 3.0. Do this using state-of-the-art thermal-hydraulics tools, a MOOSE-based code Moltres, and RELAP5-3D. **Actual Performance:** (06/28/2020) Temperature safety of the TAP MSR is studied by simulating its thermal-hydraulics behavior using the 1/8th core model as well as the unit cell model. An improvement regarding the control rods deployment is also made, leading to greater uniformity of the neutron flux and heat generation profile. The effects of gas entrainment are also studied. It is found that gas entrainment in the fuel salt introduces an overall positive reactivity to the reactor core. It is concluded that the TAP MSR core should be able to adequately perform load-following operations while doing so without exceeding its thermal safety constraints. The effect of sparging system is implicitly considered through the fuel composition and the entrainment ratio. RELAP5-3D simulation is not performed as it is not suited to model a liquid-fueled molten salt reactor.
Task 4: Experimental Investigation of Fission Gas Removal M4.1: Demonstrate ion implantation or gas sparging M 4.2: Single effect experiments M4.3: Multiple-effect experiments	M4.1: Demonstrate that ion implantation and/or gas sparging can achieve needed Xe concentration (>0.1 ppm) and that MS loop functions properly. **Actual Performance:** (09/28/2020) The finishing date is delayed compared with the original end date due to salt producing and analysis as well as technical challenges. Gas sparging is used to implant inert gas into FLiNaK. Krypton is used instead of xenon. The implantation is confirmed by flowing helium gas into the container and analyzing the Krypton signal in the residual gas analyzer.
	M4.2: Complete first single effect experiments to characterize Xe removal efficiency; feed this data into modeling work. **Actual Performance:** (12/28/2021) Nominal sparging cases are performed. Though it is later determined that a leakage exist in the sparging tube, and the results are actually for flowing helium cover gas. A flowing cover gas slightly increase the rate of dissolved inert gas extraction. The sparging effect is not successfully studied.
	M4.3: Determine optimal conditions for Xe and Kr removal with gas sparging and cover gas through multiple effect experiments; feed this data into modeling work. **Actual Performance:** (12/28/2021) Due to challenges with the experiment this milestone is not completed.

This project involves extensive modeling and simulation efforts, the model description and validation can be found in the corresponding sections of this report.

PROJECT ACTIVITIES

During the project period, four interconnected aspects of the development of the fission product removal system in liquid-fueled molten salt reactors are investigated. These aspects are basic properties of dissolved fission product and fission product removal experiment, the simulation and analysis of the fission product removal system, the fuel cycle simulation, and the coupled reactor neutronics and thermal hydraulics transient simulation. From the findings of each aspect, it is shown that removing the fission product continuously from the liquid-fueled thermal MSR is important to enable load-following operation. Based on the fuel cycle analysis and the coupled reactor core simulation, the safety limits are generally not exceeded during the operation.

PROJECT OUTPUTS

A. Journal Articles

- Chen, J., Brooks, C.S., 2021. Experiments and CFD Simulation of Mass Transfer and Hydrodynamics in a Cylindrical Bubble Column, Chemical Engineering Science, 234, 116435.
- Ashraf, O., et al., Preliminary design of control rods in the single-fluid double-zone thorium molten salt reactor (SD-TMSR). Annals of Nuclear Energy, 2021, 152: p. 108035.
- Ashraf, O., et al., Strategies for thorium fuel cycle transition in the SD-TMSR. Annals of Nuclear Energy, 2020. 148: p. 107656-107656.
- Turkmen, M., G.J.Y. Chee, and K.D. Huff, Machine learning application to single channel design of molten salt reactor. Annals of Nuclear Energy, 2021. 161: p. 108409.
- Rykhlevskii, A., J.W. Bae, and K.D. Huff, Modeling and simulation of online reprocessing in the thorium-fueled molten salt breeder reactor. Annals of Nuclear Energy, 2019. 128: p. 366-379.
- Ashraf, O., et al., Whole core analysis of the single-fluid double-zone thorium molten salt reactor (SD-TMSR). Annals of Nuclear Energy, 2020. 137: p. 107–115.

B. Papers (including conference papers and theses)

- Chen, J., Brooks, C.S., Simulation and analysis of a prototypical xenon removal system, ANS Winter Meeting, Chicago, IL, November 16-19, 2020.
- Chen, J., Brooks, C.S., Experimental study of fluid motion and mass transfer in a cylindrical bubble column, International Topical Meeting on Advances in Thermal Hydraulics, Paris, France, March 31-April 3, 2020.
- Chen, J., Brooks, C.S., CFD Simulation of Xenon removal by Helium sparging in molten salt, 18th International Topical Meeting on Nuclear Reactor Thermal Hydraulics, Portland, Oregon, USA, August 18-23, 2019.

- Chen, J., Brooks, C.S., Simulation and Analysis of a Prototypical Bubble Separator for Molten Salt Reactor, ANS Winter Meeting, Washington D.C., Nov. 31-Dec. 3, 2021.
- Chen, J., Validation of a multiphase CFD model with mass transfer for xenon removal in molten salt reactor., University of Illinois at Urbana-Champaign, 2020.
- Rykhlevskii, A., et al. The Impact of Xenon-135 on Load Following Transatomic Power Molten Salt Reactor. in Transactions of the American Nuclear Society. 2019. American Nuclear Society.
- Rykhlevskii, A., et al. Fuel Cycle Performance of Fast Spectrum Molten Salt Reactor Designs. in Proceedings of Mathematics and Computation 2019. 2019. American Nuclear Society.
- Rykhlevskii, A., Fuel Processing Simulation Tool for Liquid-fueled Nuclear Reactors. 2020, University of Illinois at Urbana-Champaign
- Lee, Alvin J.H., Neutronics and thermal-hydraulics analysis of Transatomic Power molten salt reactor (TAP MSR) core under load following operations, 2020, University of Illinois at Urbana-Champaign

C. Websites Featuring Project Work Results

- TAP MSR model for Serpent 2, https://zenodo.org/record/1450733#.W7vA8BNKi3A
- SaltProc output for TAP MSR and MSBR online reprocessing depletion simulations, https://databank.illinois.edu/datasets/IDB-7364919

FOLLOW-ON FUNDING

No follow-on funding has been received at the time of submission.

CHAPTER 1: INTRODUCTION

1.1. Introduction to Molten Salt Reactor

Molten salt reactor (MSR) refers to the kind of nuclear reactors whose primary loops are filled with molten salt. Considering the fuel form of existing design, MSR could be subdivided into two different types as the liquid-fueled MSR and the solid-fueled MSR. The liquid-fueled MSR, unlike most reactor designs, uses a liquid fuel form, where fissile material is dissolved in the molten salt [1, 2]. This unique design introduces several distinguishing features compared with conventional light water reactors (LWR). The solid-fueled MSR design is introduced in recent years, with the ambition to overcome some technical difficulties with the liquid-fueled design while retaining favorable features of molten salt as a heat transfer fluid. The main characteristics of molten salt are their high boiling point, low vapor pressure, low chemical reactivity, and similar thermal conductivity, viscosity, and volumetric heat capacity to water. Low vapor pressure allows the reactor to operate at a much higher temperature than LWR with its loop pressure only slightly higher than atmosphere [3], which increases the thermal efficiency of the reactor without the risk introduced in a high-pressure system. The comparable heat capacity and thermal conductivity retain the exceptional safety margin of LWRs during accidents. Moreover, because water is not present in the primary loop, steam explosion and hydrogen explosion encountered in commercial light water reactors will not happen for a molten salt reactor.

The liquid-fueled MSR not only shares the features brought by utilizing molten salt, as discussed above, it also has several distinct virtues which make it favorable as an advanced reactor [4]. One important feature of a liquid-fueled MSR is the intrinsic safety against core melting. In case of power surge, the fuel salt will be drained from the primary loop into a storage tank where it is subcritical. Moreover, the liquid fuel possesses a larger thermal expansion coefficient compared with solid fuel, which leads to larger negative temperature coefficient for reactivity. Another inspiring design feature of a liquid-fueled MSR is online processing of the fuel salt. Since fissile material is dissolved in the molten salt, it is transported out of the reactor core region during circulation. By diverting a part of the fuel salt from the primary loop, unfavorable fission product could be removed, and the processed salt would reenter the primary loop. This feature, with online refueling, enables many interesting fuel-cycle designs, one of which is the Thorium-Uranium cycle. With different fuel compositions, the specific reactor could be a once-through convertor reactor, a self-sustainable reactor or a breeder reactor based on the conversion ratio [5, 6]. The neutrons could either be in the fast spectrum or the thermal spectrum based on specific reactor design [7]. However, the fact that the fission product is transported out of the pressure vessel into the whole primary loop may not be favorable in terms of radiation protection and plant maintenance. The radioactivity of the extracted fission product should be properly handled.

The ideas of using molten salt and dissolved fissile material were first proposed by Oak Ridge National Lab (ORNL) as part of its Aircraft Nuclear Propulsion project (ANP) in 1949, which aimed at creating a nuclear-powered aircraft for military use [8]. As the project went on, ORNL proposed the Aircraft Reactor Experiment (ARE) to build a low power test reactor, which was successfully operated with a maximum power of 2.5MW for 9 days in 1954. The design of a liquid-fueled reactor was successfully demonstrated in this project. However, due to the cost and the development of intercontinental ballistic missile, an aircraft with unlimited range lost its

military importance. As the interest into a nuclear-powered aircraft began to diminish, ORNL decided to adapt the ARE design into a civilian nuclear reactor [9]. These efforts led to the conceptional design of a Molten Salt Breeder Reactor (MSBR) operating on a thorium-uranium fuel cycle. The MSBR design is a two-region, two-fluid system. The fuel salt with uranium fluoride is mostly for power generation, which is surrounded by the blanket salt with thorium fluoride for breeding. On site fuel recycling is considered, combining the possibility of online processing and refueling. To further develop this idea, ORNL began its Molten Salt Reactor Experiment (MSRE) in 1960. The MSRE was designed to use essentially the same material as in the MSBR, with a different fuel salt component as $\mathrm{LiF} - \mathrm{BeF_2} - \mathrm{ZrF_4} - \mathrm{UF_4}$ $(65.0 - 29.1 - 5.0 - 0.9\ \mathrm{mole\%})$. For economy and simplicity, breeding was not part of the experiment. Thus, the only fissile material dissolved in the fuel salt was uranium, rather than a mixture of thorium and uranium. The blanket salt loop is also not included, since breeding is not a part of the experiment. The MSRE reactor first reached criticality on June 1, 1965, with addition of highly enriched U^{235} and was finally shutdown in December 1969. Various experiments were done during this period on reactor dynamics, material corrosion, fuel inspection, operation with U^{233}, xenon stripping, fission product deposition, tritium behavior, plutonium additions and others [2].

Despite the successful operation of MSRE, subsequent proposal for MSBR was rejected as the United States decided to focus on the research of liquid metal fast breeder reactor. The development status of MSBR is summarized in a later report [10]. However, over the years, the development of light water reactor reaches the point where major improvement is unlikely, when nuclear power needs to be more economically competitive while meeting the safety regulation. Thus, numerous existing nuclear reactor concepts were revisited and examined by the generation-IV international forum initiated by the Department of Energy in 2000. As a result, molten salt reactor was selected as one of the six generation-IV reactor concepts, and the interest has grown since. The research interest in the Europe has been focused on using molten salt reactor for burning the transuranic elements in spent fuel and molten salt breeder reactor. In their study, it is pointed out that a fast neutron reactor without graphite moderation possess a higher breeding ratio and could reduce the requirement for fuel processing rate. They proposed the MOSART design as a burner of spent fuel and the MSFR design for a fast spectrum breeder reactor. In the United States, the liquid-fueled MSR received many research and commercial interests in the recent years, especially on the topic of fuel processing. The Transatomic Power, Terrestrial Energy, ThorCon, and Flibe Energy are some of the organizations that have brought up conceptual designs of the liquid-fueled MSR. The solid fueled fluoride salt-cooled high-temperature reactor has been also studied by many universities and national labs. Generally, this type of reactor design utilizes solid fuel, and the fluoride salt is only used as a heat transfer liquid. The fluoride salt allows for high temperature operation with low pressure system, which is very favorable for reactor safety. The KR-FHR design by Kairos Power is in the early stage for licensing with the U.S. Nuclear Regulatory Commission (NRC). In China, the design of FHR and liquid-fueled MSR has been evaluated by the Chinese Academy of Science. Forced convection test loop and natural circulation loop were built for thermal hydraulic study. Recently, a 2 MWt liquid-fueled experimental reactor, TMSR-LF1, has been approved for construction by the regulation committee. For a more detailed review of the recent development of MSR, one could refer to the reviews by Serp et al. [4] and Locatelli et al. [11].

1.2. Load Following Operation through the Removal of Xenon

The liquid-fueled MSR, unlike most reactor designs, uses a liquid fuel form, where fissile materials (typically $^{233}U-^{232}Th$ for breeding reactors) [1, 2] are dissolved in the molten salt. This feature brings several benefits over the solid-fueled design. Continuous removal of fission products and refueling, for example, are two of them [12]. Among all the nuclides produced during fission, ^{135}Xe is of special importance because it is the major neutron poison produced in reactor operation, with a thermal neutron absorption cross section of 2.6×10^6 barns. It is mostly produced by decay of ^{135}I as [13]

$$^{135}I \xrightarrow{6.6h} {}^{135}Xe + \beta^-. \tag{1-1}$$

During normal operation of the nuclear reactor, the concentration of xenon is at equilibrium. Shortly after reactor shutdown, the concentration of xenon will increase significantly by ^{135}I decay for a short period, introducing extra negative reactivity. This phenomenon is known as iodine pit. A similar jump of xenon concentration also appears when the reactor is subjected to power change, which can dramatically slow down this process to raise the power back. By removing xenon continuously from the fuel salt, the power level of the reactor could vary more flexibly. This feature makes it possible for MSR to operate in a load-following manner. Load-following operations of a nuclear reactor are important maneuvers used to adjust the reactor power level in order to meet a changing external power demand or in general, to reach a desired reactor power state. The load-following capability of a nuclear reactor is an important measure in determining its competitiveness against other power sources for the provision of non-baseload power. While nuclear power has traditionally been associated with the provision of steady baseload power, the potential for the next generation nuclear reactors in taking on load-following roles is a significant consideration for the use of these reactors in the development of a resilient energy grid.

The MSR designs from Transatomic Power, Terrestrial Energy, and the Flibe Energy all include the ability of load following as one of their trademarks. Among them, the Transatomic Power Molten Salt Reactor (TAP) and the LFTR reactor from Flibe Energy include additional system to remove xenon from the fuel salt. The IMSR reactor from Terrestrial Energy relies on the cover gas to extract the gaseous fission product, which might result from its small modular design. In the effort to evaluate the load-following technology, the TAP and the MSBR are chosen in this project as the reference design for their similarity in the scale, salt composition and the amount of publicly available information.

1.3. Transatomic Power MSR Design Description

The TAP concept is a 1250 MW_t MSR with a LiF-based uranium fuel salt [14]. This concept uses configurable zirconium hydride ($ZrH_{1.66}$) rods as the moderator while most MSR designs usually propose high-density reactor graphite. Zirconium hydride can achieve the same degree of thermalization as graphite with a much smaller volume. Compared to graphite, which shrinks and swells over time under irradiation, the cladded zirconium hydride has a much longer lifespan in extreme operational conditions - high temperature, large neutron flux, chemically aggressive salt. Finally, zirconium hydride is a nonporous material that absorbs much fewer neutron poisons (e.g., krypton, xenon) compared to high-density reactor graphite [14-16].

Figure 1.1 renders of the primary and secondary loop of the TAP MSR seated inside a concrete nuclear island. Figure 1.2 shows the schematic design of a 520 MWe, 2-loop nuclear reactor system with an intermediate salt loop.

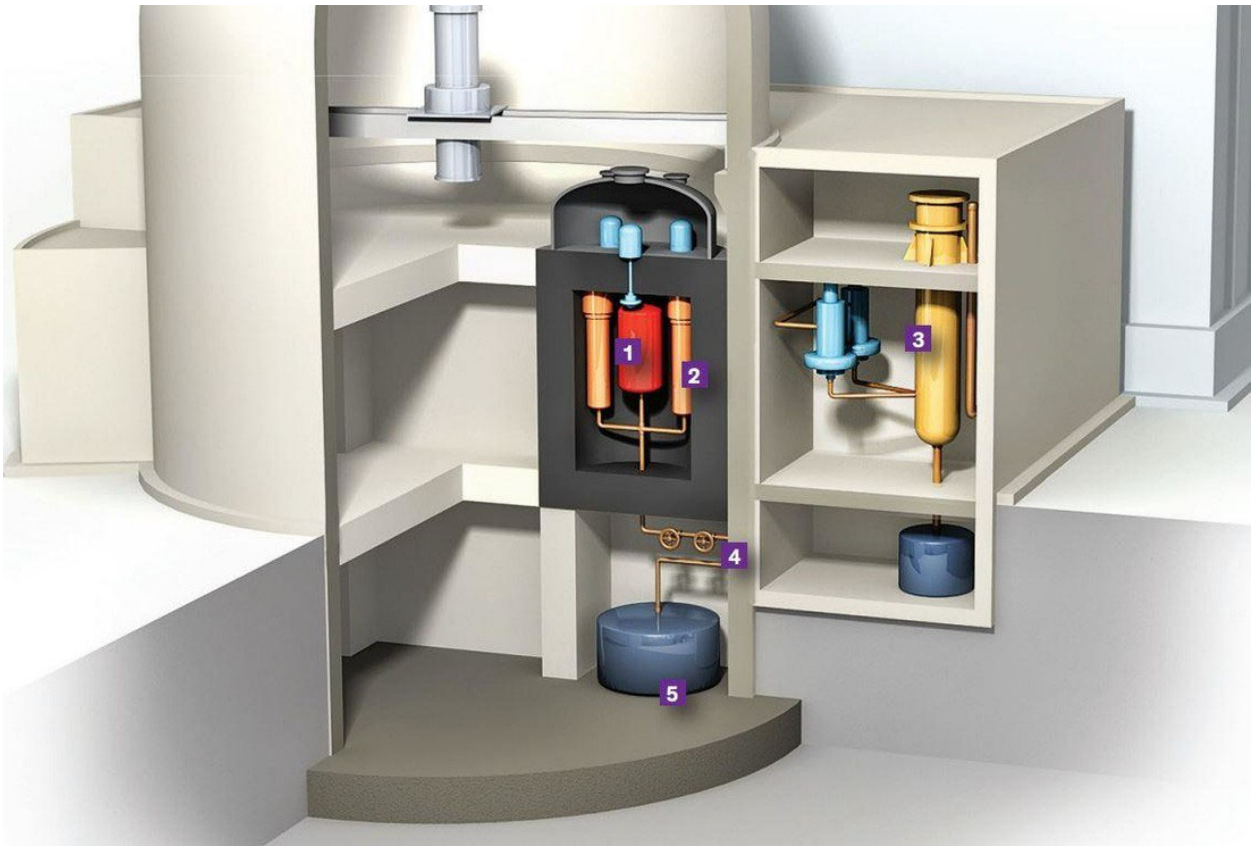

Figure 1.1: Rendering of the TAP MSR. The fission happens in the fuel salt inside the reactor vessel (1). The heat generated by a self-sustaining nuclear fission reaction would be transferred to the secondary salt by heat exchangers (2), which would boil water in the steam generator (3). Valves made of salt with a higher melting point (4) would melt in case of emergency, allowing the salt to drain into a drain tank (5), which can passively dissipate decay heat (reproduced from [17], illustration by Emily Cooper).

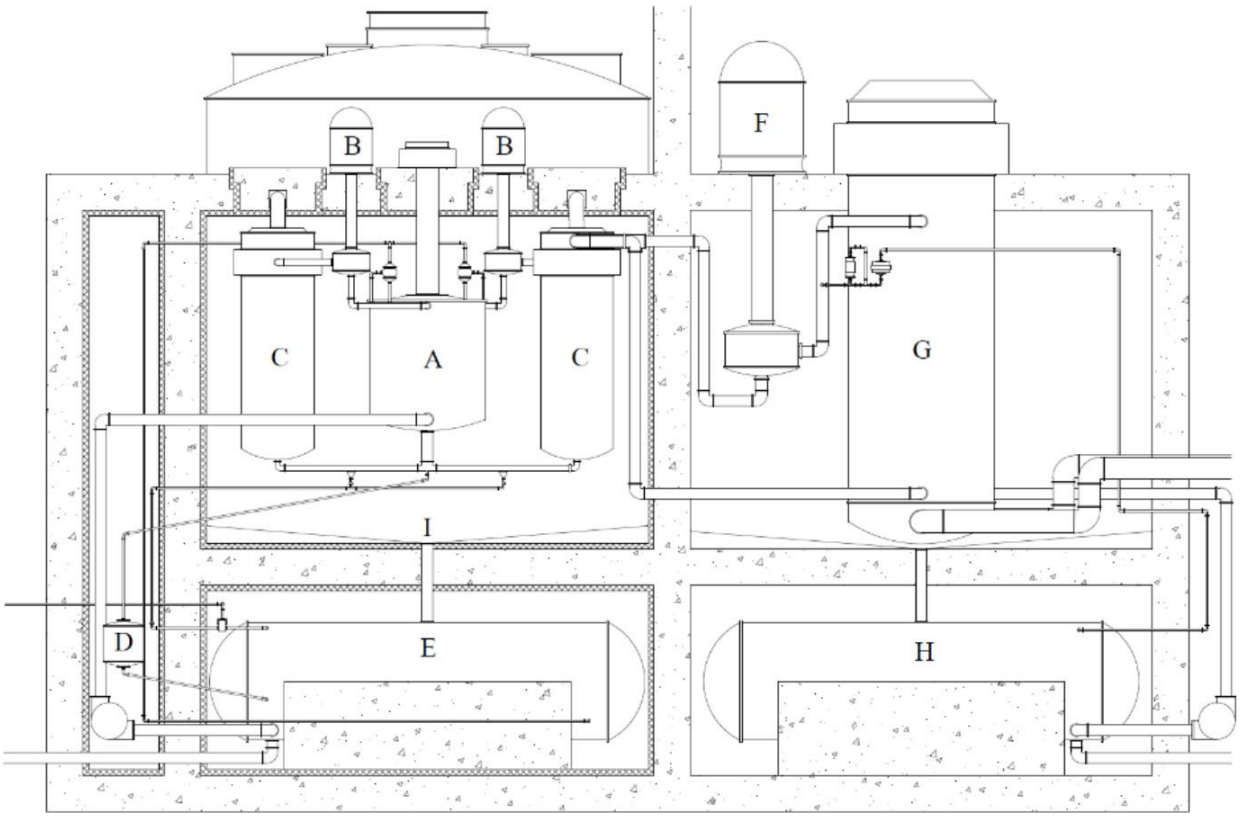

Figure 1.2: Simplified schematic of the TAP MSR primary and secondary loops (reproduced from the Transatomic Power Technical White Paper [14]) Figure legend: A) reactor vessel, B) fuel salt pumps, C) primary heat exchangers, D) freeze plug, E) primary loop drain tank, F) secondary loop salt pump, G) steam generator, H) secondary loop drain tank, I) fuel catch basin.

The TAP design (Figure 1.3) is very similar to the original Molten Salt Reactor Experiment (MSRE) design developed by Oak Ridge National Laboratory (ORNL) [2] but with two major innovations: the fuel salt composition and the moderator. The LiF-BeF_2-ZrF_4-UF_4 salt used in MSRE has been substituted with a LiF-UF_4 salt allowing the uranium concentration within the fuel salt to be increased from 0.9 to 27.5% while maintaining a relatively low melting point (490°C compared with 434°C for the original MSRE's salt) [16]. The graphite has a very high thermal scattering cross section which would make it an excellent moderator, but it has a few major drawbacks. First, due to the low lethargy gain per collision, the core requires a large volume of graphite to reach criticality, leading to a larger core and obstructing the core power density. Second, even special reactor-grade graphite has relatively high porosity, meaning, it holds gaseous Fission Products (FPs) (e.g., tritium, xenon) in its pores. Third, the reactor graphite lifespan in a commercial reactor is only 10 years [1]. To resolve these issues, the TAP concept uses an alternative moderator, zirconium hydride, allowing for a more compact core and a significant increase in power density. These two innovative design choices, together with a configurable moderator (the moderator-to-fuel ratio can be changed during regular maintenance shutdown), facilitate the commercial deployment of this conceptual design viable in the commercially available 5% LEU fuel cycle.

The primary loop of the TAP MSR consists of the reactor core volume moderated by the silicon carbide (SiC) cladded zirconium hydride rods, pumps, and primary heat exchanger. The pumps circulate the LiF-(Act)F_4 fuel salt through the primary loop. The pumps, vessels, tanks, and piping are made of a corrosion resistant nickel-based alloy (similar to Hastelloy-N[1]) in various molten salt environments. Inside the reactor vessel, near to the zirconium hydride moderator rods, the fuel salt is in a critical configuration and generates heat. Table 1.1 contains details of the TAP system design taken from the technical white paper [14], the neutronics overview [15], and the ORNL analysis of the TAP design [16, 18].

Table 1.1: Summary of principal data for the TAP MSR (reproduced from [14, 18]).

Parameter	Value
Thermal power	1250 MW_{th}
Electric power	520 MW_e
Gross thermal efficiency	44%
Outlet temperature	620°C
Fuel salt components	LiF-UF_4
Fuel salt composition	72.5-27.5 mole%
Uranium enrichment	5% ^{235}U
Moderator	Zirconium Hydride ($ZrH_{1.66}$) rods (with silicon carbide cladding)
Neutron spectrum	Thermal/Epithermal

[1] Hastelloy-N is very common in reactors now but have been studied and developed at ORNL in a program that started in 1950s.

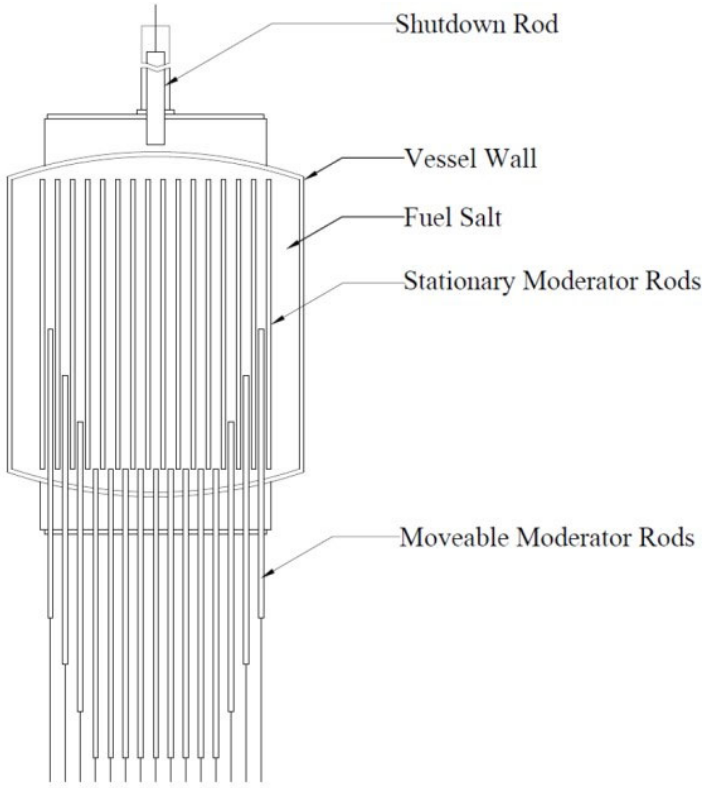

Figure 1.3: Schematic view of the TAP MSR showing the movable moderator rod bundles and the shutdown rod (figure reproduced from Transatomic Power White Paper [14]).

In the TAP core (Figure 1.4), lattices of SiC clad moderator rods form the moderator assemblies around which the fuel salt flows (Figure 1.3). The TAP reactor pressure vessel is a cylinder with an inner radius of 150 cm, a height of 350 cm, and a wall thickness of 5 cm. The moderator-to-fuel ratio, or salt volume fraction (SVF), in the core can be varied during operation to shift the spectrum from intermediate to thermal energies to maximize fuel burnup. Intermediate energies are used at Beginning of Life (BOL) and are shifted to thermal at End of Life (EOL). During operation the SVF can be varied by inserting fixed-sized moderator rods from the bottom of the reactor vessel, similar to moving the control rods in a Boiling Water Reactor (BWR), as shown in Figure 1.3. For the TAP reactor, EOL occurs when the maximum number of moderator rods are inserted into the core and further injection of fresh fuel salt does not change criticality. Unmoderated salt flowing in the annulus between the core and the vessel wall provides potential reduction of fast neutron flux at the vessel structural material [15].

The TAP nuclear island contains a FP removal system. Gaseous FPs are continuously removed using an off-gas system while liquid and solid FPs are extracted via a chemical processing system. A small quantity of fresh fuel salt is regularly added to the primary loop as byproducts are gradually removed. This process maintains a constant fuel salt mass and keeps the reactor critical. In contrast with the MSBR reprocessing system, the TAP does not require a protactinium separation and isolation system because it operates in a single-stage uranium-based fuel cycle. The authors of the TAP concept detailed three distinct fission product removal methods [15].

Off-Gas System: Gaseous fission products such as krypton and xenon are removed, compressed, and stored temporarily until they have decayed to background radiation levels. Trace amounts of tritium are also removed and bottled in a liquid form via the same process. The off-gas system also removes a small fraction of the noble metals.

Metal Plate-Out/Filtration: Removes solid noble and semi-noble metal fission products as they plate out onto a nickel mesh filter located in a side stream of the primary loop.

Liquid Metal Extraction: Lanthanides and other non-noble metals stay dissolved in the fuel salt. These elements generally have a lower capture cross section and thus absorb fewer neutrons than ^{135}Xe but their extraction is essential to ensuring normal operation. In the TAP reactor, lanthanide removal is accomplished via a liquid-metal/molten salt extraction process similar to that developed for MSBR by ORNL [1]. The process converts the dissolved lanthanides into a well-understood oxide waste form, similar to that for Light Water Reactor (LWR) spent nuclear fuel (SNF). This oxide waste comes out of the TAP reprocessing plant in ceramic granules, which can be sintered into another convenient form for storage.

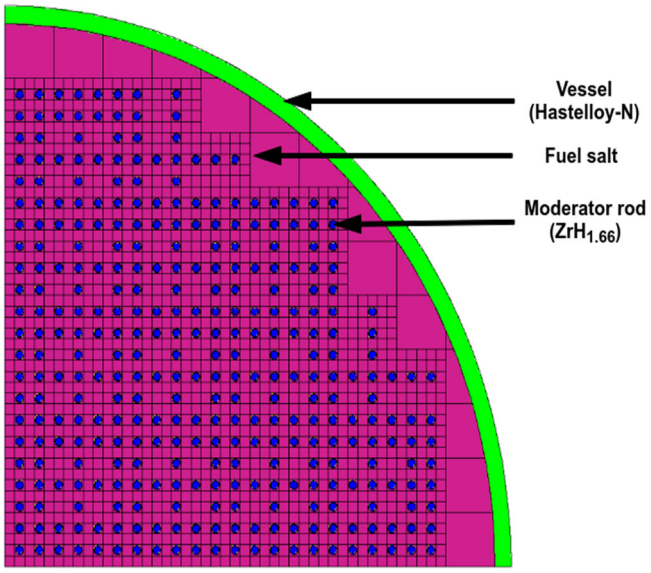

Figure 1.4: The TAP MSR schematic core view showing moderator rods in the BOL position defined by their point design (figure reproduced from ORNL/TM-2017/475 [18]).

Figure 1.5 shows a principal design of the TAP primary loop including an off-gas system, nickel mesh filter, and lanthanide chemical extraction facility. Similar to MSBR, the off-gas system is based on a simple process of helium sparging through the fuel salt with consequent gas bubbles removed before returning the fuel salt back to the core. One very notable difference is the MSBR gas separation system helium injection and subsequent transport of the voids run throughout the primary loop, including the core, for at least 10 full loops [1]. This system presents a significant concern to the safety and stability of operation due to the increase of void fraction in the fuel salt when it enters back to the core, causing unpredictable changes in reactivity. This

drawback can be overcome by using an effective gas separator to strip helium/xenon bubbles before returning the salt back to a primary loop (Figure 1.5, blue block).

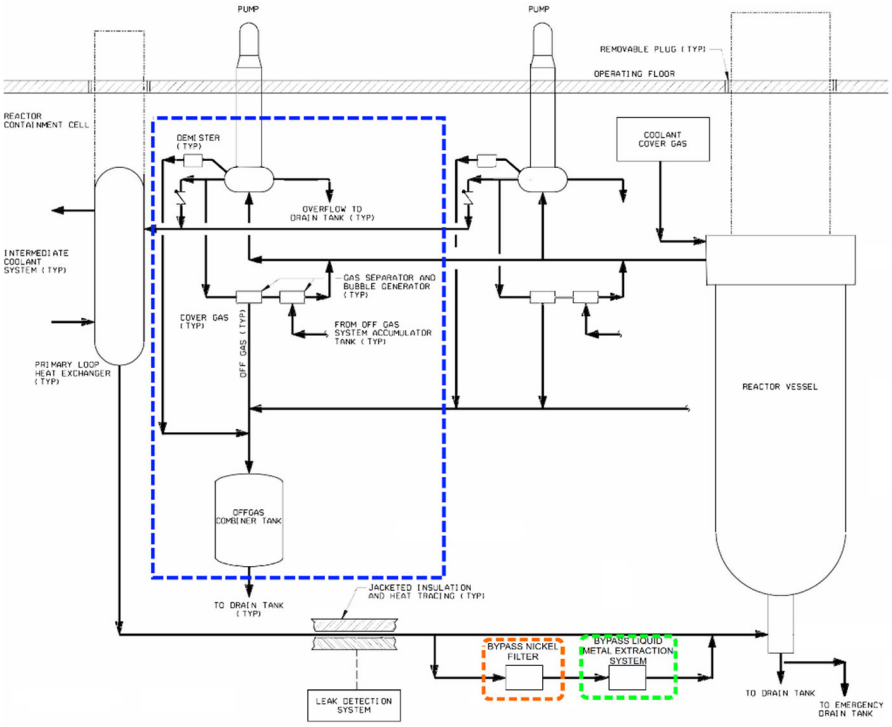

Figure 1.5: Simplified TAP primary loop design including off-gas system (blue), nickel filter (orange) and liquid metal extraction system (green) (reproduced from [19]).

Solid noble and semi-noble metal fission products tend to plate out onto the metal surfaces including piping, heat exchanger tubes, reactor vessel inner surface, etc. Previous research by ORNL [1] concluded that about 50% of noble and semi-noble metals would plate out inside MSBR systems without any special treatment. To improve the extraction efficiency of these fission products, the TAP concept employs a nickel mesh filter located in a bypass stream in the primary loop (Figure 1.5, orange block). The main idea of this filter is to create a maze with a large metal (nickel) surface area. The fuel salt flows throughout the filter and the noble metals plate-out on the filter internal surface.

This Liquid Metal Extraction process for the TAP concept has been adopted from the MSRE. The MSRE demonstrated a liquid-liquid extraction process for removing rare earths and lanthanides from fuel salt and estimated efficiency of this process. In fact, due to similarities in reprocessing schemes, the TAP project reported almost the same set of elements for removal and similar effective cycle times[2] as suggested for the MSBR (Table 1.2) The TAP neutronics white paper specifies additional low-probability fission products and gases that should be removed during operation [15]. These elements are categorized into the previously defined processing groups, but the removal rates of most of these elements (except hydrogen) are meager.

[2]The MSBR program defined “cycle time” as the time required to remove 100% of atoms of a target nuclide from a fuel salt [1].

Table 1.2: The effective cycle times for fission products removal from the TAP MSR (reproduced from [20] and [15]).

Processing group	Nuclides	Removal Rate (s^{-1})	Cycle time (at full power)
Elements removed in MSBR concept and adopted for the TAP [1]			
Volatile gases	Xe, Kr	5.00E-2	20 sec
Noble metals	Se, Nb, Mo, Tc, Ru, Rh, Pd, Ag, Sb, Te	5.00E-2	20 sec
Seminoble metals	Zr, Cd, In, Sn	5.79E-8	200 days
Volatile fluorides	Br, I	1.93E-7	60 days
Rare earths	Y, La, Ce, Pr, Nd, Pm, Sm, Gd	2.31E-7	50 days
	Eu	2.32E-8	500 days
Discard	Rb, Sr, Cs, Ba	3.37E-9	3435 days
	Additional elements removed [15, 20]		
Volatile gases	H	5.00E-2	20 sec
Noble metals	Ti, V, Cr, Cu	3.37E-9	3435 days
Seminoble metals	Mn, Fe, Co, Ni, Zn, Ga, Ge, As	3.37E-9	3435 days
Rare earths	Sc	3.37E-9	3435 days
Discard	Ca	3.37E-9	3435 days

1.4. Molten Salt Breeder Reactor Design Description

The thorium-fueled MSBR was developed in the early 1970s by ORNL, specifically to explore the promise of the thorium fuel cycle, which uses natural fertile thorium feed material instead of enriched uranium fissile fuel. With continuous fuel reprocessing, the MSBR realizes the advantages of the thorium fuel cycle because the ^{233}U bred from ^{232}Th is almost instantly[3] recycled back into the core [14]. The chosen fuel salt, LiF-BeF_2-ThF_4-UF_4, has a melting point of 499°C, low vapor pressure at operating temperatures, and beneficial flow and heat transfer properties [83].

Figure 1.6 shows the MSBR vessel which has a diameter of 680 cm and a height of 610 cm. It contains a molten fluoride fuel-salt mixture that generates heat in the active core region and transports that heat to the primary heat exchanger by way of the primary salt pump. In the active core region, the fuel salt flows through channels in moderating and reflecting graphite blocks. Fuel salt at 565°C enters the central manifold at the bottom via four 40.64-cm-diameter nozzles and flows upward through channels in the lower plenum graphite. The fuel salt exits at the top at about 704°C through four equally spaced nozzles, which connect to the salt-suction pipes leading to primary circulation pumps. The fuel salt drain lines connect to the bottom of the reactor vessel inlet manifold.

[3] The fertile ^{232}Th is transmuted into the ^{233}Th after capturing a neutron. Next, this isotope decays to the ^{233}Pa ($\tau_{1/2}$=21.83m), which finally decays to the ^{233}U ($\tau_{1/2}$=26.967d).

Figure 1.7 shows the configuration of the MSBR vessel, including the “fission” (zone I) and “breeding” (zone II) regions inside the vessel. The core has two radial zones bounded by a solid cylindrical graphite reflector and the vessel wall. The central zone, zone I, in which 13% of the volume is fuel salt and 87% is graphite, is composed of 1,320 graphite cells, 2 graphite control rods, and 2 emergency shutdown rods. The under-moderated zone, zone II, in which 37% of the volume is fuel salt and 63% is graphite, and radial reflector, surrounds the zone I core region and serves to diminish neutron leakage. Zones I and II are surrounded radially and axially by fuel salt (Figure 3.3); this space for fuel is necessary for the injection and flow of molten salt.

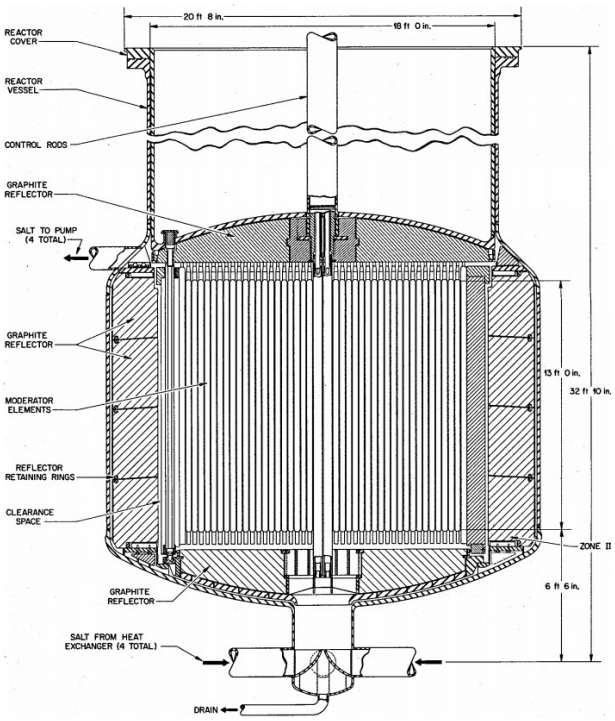

Figure 1.6: XZ view of the MSBR (reproduced from Robertson et al. [83]).

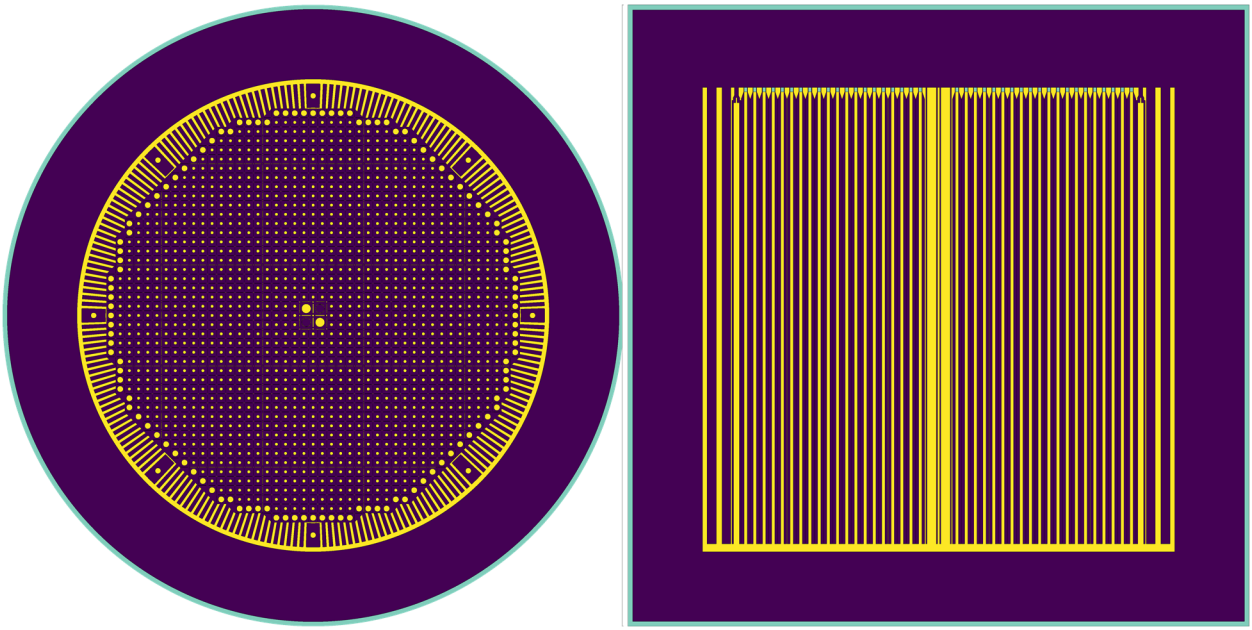

Figure 1.7: XY (left) and XZ (right) views of a Serpent MSBR model (reproduced from Rykhlevskii et al. [88]).

Since reactor graphite experiences significant dimensional changes due to neutron irradiation, the reactor core was designed for periodic replacement. Based on the experimental irradiation data from the MSRE, the core graphite lifetime is about 4 years, and the reflector graphite lifetime is 30 years [83].

The core design also has eight symmetric graphite slabs with a width of 15.24 cm in zone II, one of which is illustrated in Figure 1.8. The holes in the centers are for the core lifting rods used during the core replacement operations. These holes also allow a portion of the fuel salt to flow to the top of the vessel for cooling the top head and axial reflector. Figure 1.8 also shows the 5.08-cm-wide annular space between the removable core graphite in zone II-B and the permanently mounted reflector graphite. This annulus consists entirely of fuel salt, provides space for moving the core assembly, helps compensate for the elliptical dimensions of the reactor vessel, and serves to reduce the damaging flux at the surface of the graphite reflector blocks.

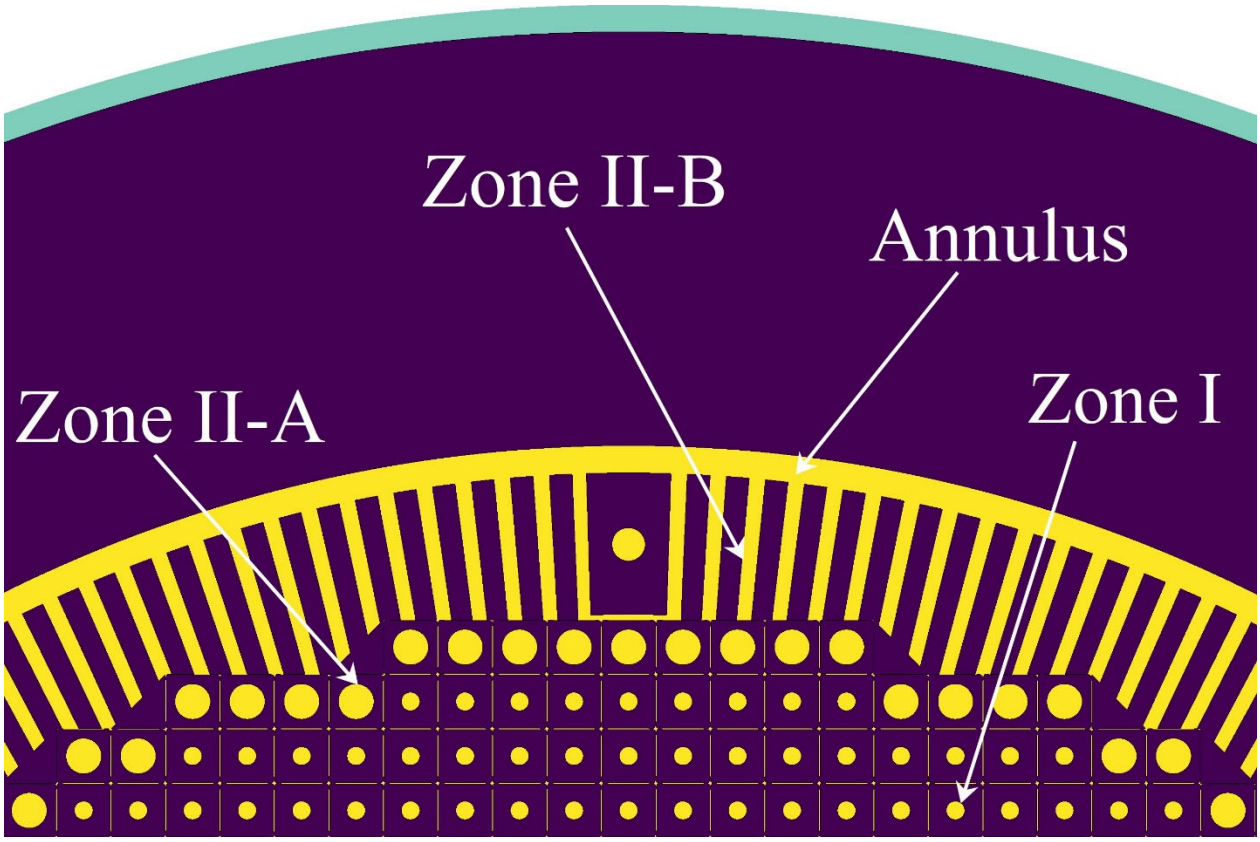

Figure 1.8: Detailed view of the MSBR two-zone model. Yellow represents fuel salt, purple represents graphite, and aqua represents the reactor vessel (reproduced from Rykhlevskii et al. [88]).

^{135}Xe is a strong neutron poison, and some fraction of this gas is absorbed by graphite during MSBR operation. ORNL calculations showed that for unsealed commercial graphite with a helium permeability of 10^{-5} cm^2/s, the calculated ^{135}Xe poison fraction[4] is less than 2% [1]. This parameter can be improved by using experimental graphite types or by applying sealing technology. The effect of the gradual poisoning of the core graphite with xenon is outside of the scope of this work.

The central region of the core, called zone I, is made up of graphite elements, each 10.16cm 10.16cm×396.24cm and has 13% fuel salt by volume. Zone I has 4 channels for control rods: two for graphite rods, which both regulate and shim during normal operation, and two for backup safety rods consisting of boron carbide clad to assure sufficient negative reactivity for accidents.

Zone I graphite elements have a mostly rectangular shape with lengthwise ridges at each corner that leave space for salt flow around the elements. Figure 1.9 shows the elevation and plan views of graphite elements of zone I [83] and their Serpent model [92].

[4] The original ORNL report by Robertson et al. [1] defined ^{135}Xe poison fraction as the number of neutrons absorbed by ^{135}Xe compared with the total number of neutrons (both fast and thermal) absorbed by ^{233}U.

Zone II, which is undermoderated, surrounds zone I. Combined with the bounding radial reflector, zone II serves to diminish neutron leakage. Two kinds of elements form this zone: large-diameter fuel channels (zone II-A) and radial graphite slats (zone II-B).

Zone II has 37% fuel salt by volume, and each element has a fuel channel diameter of 6.604 cm. The graphite elements for zone II-A are prismatic, with elliptical dowels running axially between the prisms. These dowels isolate the fuel salt flow in zone I from that in zone II. Figure 1.10 shows the shapes and dimensions of these graphite elements and their Serpent model. Zone II-B elements are rectangular slats spaced far enough apart to provide the 0.37 fuel salt volume fraction. The reactor zone II-B graphite 5.08cm-thick slats vary in the radial dimension (average width is 26.67cm) as shown in Figure 1.8. Zone II serves as a blanket to achieve the best performance: a high breeding ratio and a low fissile inventory. The harder neutron energy spectrum in zone II enhances the rate of thorium resonance capture relative to the fission rate, thus limiting the neutron flux in the outer core zone and reducing the neutron leakage [83].

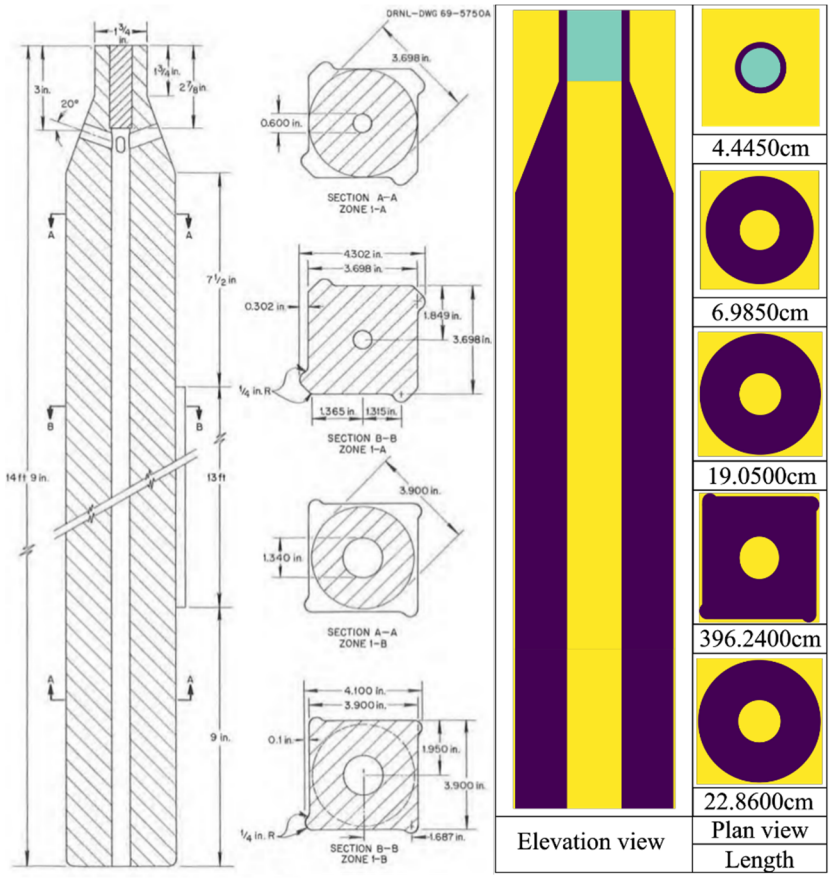

Figure 1.9: Graphite moderator elements for zone I: reference design (left) [83] and Serpent model (right) [92]. Yellow represents fuel salt, purple represents graphite, and aqua represents the reactor vessel (reproduced from Rykhlevskii et al. [88]).

The sophisticated, irregular shapes of the fuel elements challenge an accurate representation of zone II-B. The suggested design [83] of zone II-B has eight irregularly-shaped graphite elements as well as dozens of salt channels. These graphite elements were simplified into right-circular cylindrical shapes with central channels. Figure 1.8 illustrates this core region in the Serpent model.

The volume of fuel salt in zone II was kept exactly at 37% so this simplification did not impact the core neutronics. Simplifying the eight edge channels was the only simplification made to the MSBR geometry in the work presented in Chapter 3.

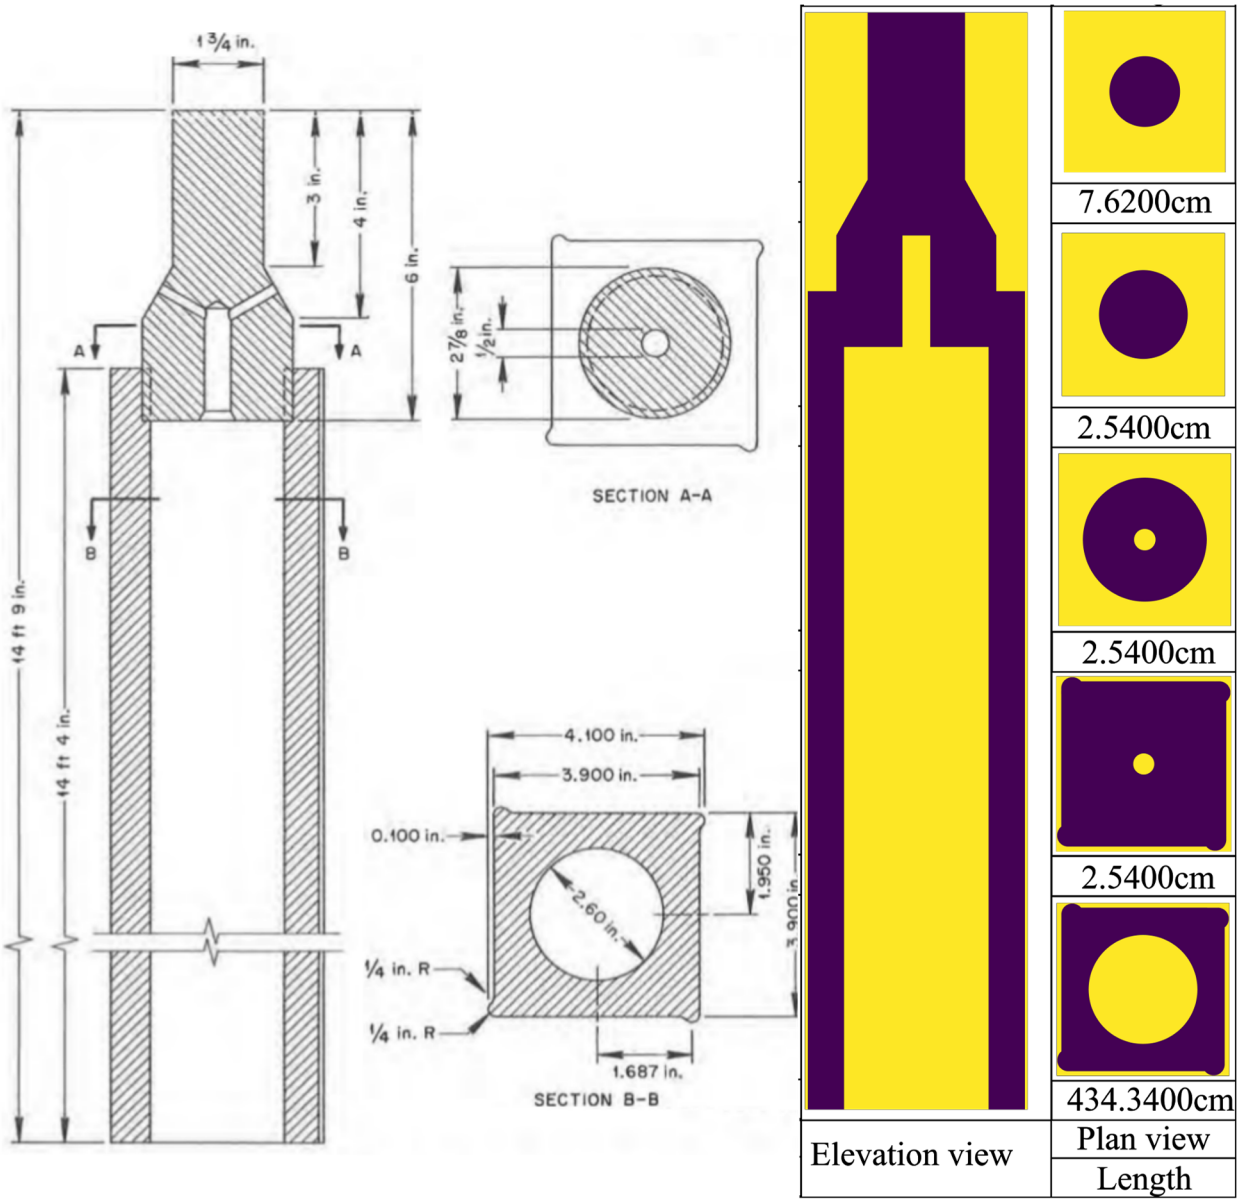

Figure 1.10: Graphite moderator elements for zone II-A: reference design (left) [83] and Serpent model (right) [92]. Yellow represents fuel salt, purple represents graphite, and aqua represents the reactor vessel (reproduced from Rykhlevskii et al. [88]) The fuel salt, reactor graphite, and modified Hastelloy-N are all materials invented at ORNL specifically for the MSBR. The fuel salt selected for use in the MSBR is LiF-BeF_2-ThF_4-$^{233}UF_4$ (71.75-16-12-0.25 mole% which has density of 3.35 g/cm^3 [83]. The lithium in the molten salt fuel is fully enriched to 99.995% ^{7}Li because ^{6}Li is an extremely strong neutron poison and becomes tritium upon neutron capture.

The specific temperature was fixed for each material and stays constant during reactor operation. The isotopic composition of each material at the initial state was described in detail in the MSBR conceptual design study [83] and has been applied to the Serpent model without any modification. Table 1.3 is a summary of the major MSBR parameters used to inform the Serpent model [83].

Table 1.3: Summary of principal data for the MSBR (reproduced from Robertson et al. [83]).

Parameters	Value
Thermal power	2250 MW_{th}
Electric power	1000 MW_e
Gross thermal efficiency	44.4%
Salt volume fraction (Zone I)	0.13
Salt volume fraction (Zone II)	0.37
Fuel salt inventory (Zone I)	8.2 m^3
Fuel salt inventory (Zone II)	10.8 m^3
Fuel salt inventory (annulus)	3.8 m^3
Total fuel salt inventory	48.7 m^3
Fissile mass in fuel salt	1303.7 kg
Fuel salt components	LiF-BeF_2-ThF_4-$^{233}UF_4$
Fuel salt composition	71.75-16-12-0.25 mole%
Fuel salt density	3.35 g/cm^3

As mentioned in section 2.1, the MSBR design requires online reprocessing to completely remove neutron gaseous FPs (Xe, Kr) and noble metals (e.g., Se, Nb, and Mo) every 20 seconds. The ^{232}Th in the fuel absorbs thermal neutrons and produces ^{233}Pa, which then decays into the fissile ^{233}U. Protactinium presents a challenge since it has a large absorption cross section in the thermal energy spectrum. Moreover, ^{233}Pa left in the core produces ^{234}Pa and ^{234}U, neither of which are useful as fuel. Accordingly, ^{233}Pa is continuously removed from the fuel salt into a temporary storage tank to allow ^{233}Pa to decay to ^{233}U without the corresponding negative neutronic impact. The reactor chemical processing system must separate ^{233}Pa from the molten salt fuel over 3 days, hold it while ^{233}Pa decays into ^{233}U, and return it to the primary loop. This feature allows the reactor to avoid neutron losses to protactinium, lowers in-core fission product inventory, and increases the efficiency of ^{233}U breeding.

Table 1.4 summarizes a full list of nuclides and their cycle time used for modeling salt treatment and separations [83]. The removal rates vary among chemical elements in this reactor concept and dictate the necessary resolution of depletion calculations. If the depletion time intervals are short, an enormous number of depletion steps are required to obtain the equilibrium composition. On the other hand, if the depletion calculation time interval is too long, effective multiplication factor k_{eff} would be lower than expected in reality due to higher equilibrium concentration of strong poisons (e.g., ^{135}Xe) in fuel salt. To compromise, a 3-day time interval was selected for depletion calculations to correlate with the removal interval of ^{233}Pa as suggested by Powers et al. [77]. Finally, ^{232}Th was continuously added every 3 days to maintain the initial mass fraction of ^{232}Th in the fuel salt.

Table 1.4: The cycle times for protactinium and fission product removal from the MSB (reproduced from Robertson et al. [83]).

Processing group	Nuclides	Cycle time (at full power)
Rare earths	Y, La, Ce, Pr, Nd, Pm, Sm, Gd	50 days
	Eu	500 days
Noble metals	Se, Nb, Mo, Tc, Ru, Rh, Pd, Ag, Sb, Te	20 sec
Semi-noble metals	Zr, Cd, In, Sn	200 days
Gases	Kr, Xe	20 sec
Volatile fluorides	Br, I	60 days
Discard	Rb, Sr, Cs, Ba	3435 days
Protactinium	^{233}Pa	3 days
Higher nuclides	^{237}Np, ^{242}Pu	16 years

1.5. Structure of the Project Report

The project aims at the design of a fuel processing system that enables the liquid-fueled MSRs to load follow. The prototypical system is similar to the inert gas sparging system considered in the MSBR design [10]. Initially, the reference reactor design is the TAP reactor, but later the MSBR design is also included to extend the application to thermal neutron reactors. The development of the system is broken into four aspects from top to the bottom of this technology, which are covered in different chapters of this report. These aspects are basic properties of dissolved fission product and fission product removal experiment, the simulation and analysis of the fission product removal system, the fuel cycle simulation, and the coupled reactor neutronics and thermal hydraulics transient simulation.

The most basic aspect of the processing system is the properties of molten salt and dissolved fission product. The knowledge of solubility and diffusivity of the gaseous fission product in different molten salt is very limited. Though indirect measurement of the xenon removal is performed in the MSRE through the reactivity calculation, direct measurement of the dissolved xenon removal in molten salt has not been done in the past. A molten salt sparging experiment is designed and carried out during the project period using residual gas analyzer in a vacuum system. The experimental design and the results are included in Appendix A. The salt used in the experiment is obtained from the Argonne National Lab and the Oak Ridge National Lab. Property measurements are made at the Argonne National Lab and the results are summarized in Appendix B [21, 22].

Chapter 2 summarizes the second level of the processing system, the specific design and evaluation of the removal system. The physical process is simulated with computational fluid dynamics. Before the simulation tools are used on the removal system, a series of validation experiments and simulations are performed in a geometry similar to the molten salt sparging experiment. With confidence in the constitutive models, simulations are performed in the prototypical removal system under various conditions. Comparisons are also made with existing literature data relevant to the removal system to confirm the model applicability under the different flow conditions, though these comparisons could not be considered as a rigorous validation. The performance of the removal system is analyzed using global quantities as the removal efficiency

and the separation efficiency. Regression models are constructed to predict the system performance under different conditions based on simulation results. These simulation results could be feed into the fuel cycle analysis as the underlining model for the fission product removal system. On the other hand, the general requirement for fission product removal system under normal load following conditions could be obtained from the reactor transient simulation and the fuel cycle simulation. Based on the input from these high-level simulations, the specifications of the xenon removal system could be determined.

Chapter 3 describes the third aspect of the development, the fuel cycle simulation and analysis. In liquid-fueled MSRs, the fuel is not confined in the reactor core, which presents a significant challenge to the traditional neutronics code. The flowing of the fuel salt and the necessity to model the full primary loop require new simulation capacity. To address these problems, the SaltProc Python toolkit is developed and coupled to the neutronics code Serpent. The reactor core is modeled in Serpent for neutronics calculation, tracking the evolution of isotopes in the core. The SaltProc uses the output from Serpent, simulates the fuel processing and refueling operations and provides the modified input to Serpent. The fuel processing is modeled using realistic engineering model from the literature and the output from the second level of development, rather than a prescribed constant which could not reflect the change of system performance under different operation conditions. Using the coupled strategy, fuel depletion and load-following simulations are carried out for TAP MSR and MSBR. The fuel depletion simulations are compared with literature results for validation. The load following simulations track the change of reactivity with power maneuvering, where a reactivity constantly larger than unity is required to follow the load. Moreover, reactivity swing after power maneuvering should be limited or short-term load-follow operations are not possible. Suggested removal efficiency is found through the simulations for the processing system design.

Chapter 4 describes the fourth aspect of the system development, the transient coupled rector core simulation. In the fuel depletion calculation, the load-following capacity is evaluated with excessive reactivity. However, excessive reactivity is not the only parameter that determines whether the reactor could follow the load. Changing the reactor power is related to the control of the reactor system. Safety criteria must be met during the process, which requires detailed knowledge of the temperature, flow, neutron flux, control rod movement, etc. This dictates a transient coupled simulation of the reactor core, which is carried out using the Moltres, a MOOSE-based Finite Element (FE) Partial Differential Equation (PDE) solver that is built specifically to simulate a neutron diffusion (and multiplication) system with advective fuel. The fuel composition is provided by the fuel depletion calculation presented in Chapter 3, and the material cross section used in the neutronics calculation is obtained using Serpent. Based on the simulation results, the safety criteria during the power ramping process are assessed such as the maximum temperature of different components. In addition, the effects of gas entrainment and xenon concentration are evaluated which are closely related to the fuel processing system.

After the descriptions of the four aspects of the system development are given, the conclusions and findings from these tasks are summarized in Chapter 5. Based on the findings, an outlook of the design of the processing system is given. The limitations of the analysis are discussed, and possible future improvements are suggested.

CHAPTER 2: SIMULATION AND ANALYSIS OF THE XENON REMOVAL SYSTEM

In this chapter, the efforts towards the design analysis of a xenon removal system are presented. In Section 2.1, a literature review on gaseous fission product removal and multiphase mass transfer is carried out. The motivation of xenon removal is briefly discussed, and the difficulty of designing a xenon removal system is reviewed, which leads to the usage of CFD simulation. In Section 2.2, the basics of Eulerian two-fluid model and species mass transfer are discussed briefly. A series of validation experiments are presented in Section 2.3. The void fraction, velocity field and mass transfer data from the validation experiment serve as a guidance for model selection and a benchmark of the final CFD model as shown in Section 2.4. In Section 2.5 and Section 2.6, the validated model is applied to the removal system. Simulation results are compared with experimental data for bubble separator and engineering mass transfer model in pipe flow. Sensitivity study is carried out for design parameters, based on which regression models are derived.

2.1. History of Fission Product Removal

As a reactor type that is still being actively studied, many aspects of the molten salt reactor are still under investigation. The idea of fission product removal is among one of these aspects. In this section, the development of fission product removal, along with the associated modeling efforts are reviewed.

2.1.1. Experience with Gaseous Fission Product Removal

As discussed in Chapter 1, the idea of MSR originates from the ANP project at ORNL. During the operation of Aircraft Reactor Experiment, a steady state xenon poisoning experiment was carried out. The reactor was maintained at constant power for 25 hours, and the xenon poisoning was measured by the change of control rod value to maintain the reactor in steady state. The measured xenon poisoning is only 5% of the theoretical value if all xenon is contained in the fuel salt [8]. This deviation was attributed to the swirling of the fuel at the fuel pump, where the fission-product gases could have escaped to the cover gas. This argument was reinforced by the fission-product gases detected in a leak of the gas systems. In ORNL's subsequent proposal for Aircraft Reactor Test [23], this feature was utilized, and a xenon removal system by helium scrubbing at the mixing chamber was considered.

As a transition from the military-based project to a civilian power reactor, the MSBR design and the MSRE project were proposed as the interest in a nuclear-powered aircraft diminished. Many design considerations in the MSRE were inherited from the ANP project, including the xenon removal system at the fuel pump [24]. During the operation of MSRE, a series of experiments and studies were carried out on xenon poisoning. It is concluded that most of the xenon poisoning result from neutron absorption by ^{135}Xe within the core graphite [25]. One solution is to use graphite with lower xenon diffusivity and void fraction. Additionally, this xenon poisoning could be reduced by processing of the fuel salt. With careful analysis of the system, an empirical model for steady state xenon poison calculation was established [26]. In this model, xenon poisoning is estimated based on balance between source and sink terms. The source terms are identified as decay of ^{135}I and small portion of direct generation from fission, which are functions of reactor power level. The sink terms, on the other hand, are more complicated. Five

primary sink terms are considered, including burnup, decay, migration to core graphite, transfer to helium bubble and removal via the xenon stripper. The xenon stripper, or the xenon removal system, refers to the special design of MSRE's fuel pump, where a portion of the discharged fuel salt is sucked from the pump and sprayed into the cover gas in the fuel pump bowl through a spray ring. Large interfacial area is formed for xenon stripping in this process, whose performance is approximated by a stripping efficiency coefficient, defined as the percentage of xenon transferred to the cover gas in the diverted fuel salt. Another important sink term is the transfer to helium bubbles. Helium bubbles are formed at the fuel pump bowl by helium bubbler and salt spray. Some of the bubbles are further carried into the loop where xenon could escape from the salt to the bubbles. These circulating bubbles could escape to the cover gas as passing through the fuel pump again, and the xenon absorbed in these bubbles could level the system. The helium bubble renewal efficiency, which is defined as the percentage of circulating bubbles that are renewed in the fuel pump, is used to describe this process. The model was first evaluated using a gas containing ^{85}Kr tracer in a preliminary experiment. Some of the model constants were determined in these tests. Good agreement was found between the model and experiment. By applying the model to evaluate xenon poisoning, it is concluded that the transfer to helium bubbles term has a huge impact on the final derived xenon poisoning. With 1% of helium void fraction, xenon poisoning would decrease by around 80%, if the stripping efficiency and helium bubble renewal efficiency are both as assumed to be 10%. During their derivation, the bubble size was assumed to be 0.254 mm in diameter and the mass transfer coefficient was taken as 1.7×10^{-4} m/s. However, certain disagreement was shown when compared with the actual ^{135}Xe poisoning later in the MSRE operation, which could partly come from the uncertainty in the model parameters.

Though in the MSRE project, the xenon removal system is integrated in the fuel pump. For a large scale MSR at higher power, separate xenon removal device is required to achieve the same level of xenon removal. Based on the study mentioned before and several other research [8, 27], physical removal of fission product by inertial gas sparging raised special interest because of its simplicity and compatibility with the fuel salt system [28]. In these reports, the lack of experimental xenon diffusivity and equilibrium ratio data were identified as a significant limitation, and further experiments were suggested. A prototypic gaseous fission product removal system was designed [10, 29] based on these studies, which includes a gas generator to inject helium bubbles and a gas separator to remove them. The separator adopted a rotational flow pattern to push the bubbles into the center of the pipe, and subsequently exit from the outlet located at the center. The liquid entrainment and pressure drop are the limiting factors for this design. In view of xenon removal and neutronics, the separation efficiency of the separator should be as large as possible. In order to design a gaseous fission product removal system, accurate prediction of mass transfer coefficient, void fraction and interfacial area concentration is necessary. However, these quantities were assumed to be uniform in the preliminary design and were calculated based on experiments with simplified geometry and flow conditions. Thus, considerable uncertainty could be expected for these results.

After the research of MSR was ceased in the US, few studies have been carried out on fission product removal, until recently. With the effort to develop more advanced and safer Gen-IV nuclear reactors, the MSR once again raises global research interest, and the possibility of fission product removal with inert gas sparging is being inspected again. Rubio et al. [30] carried out a scaled experiment to study tritium extraction by inert gas sparging and ultrasonic enhancement, where the Schmidt number of the scaling fluid is adjusted to the Schmidt number of FliBe by

changing its composition. In this experiment, oxygen was dissolved in the scaling fluid and was removed by inert gas sparging. The mass transfer coefficient is calculated based on the oxygen concentration change with time. The experiment is carried out in a co-current pipe flow. The bubbles are injected upstream to the ultrasonic affected region and the oxygen concentration is measured downstream. After the measurement point, the bubbles are separated in a vertical cylinder. The upper region of the cylinder is filled with gas and connected to the outlet. The mixture enters the cylinder in the middle and flows downwards. The bubbles are driven by buoyancy and exit the cylinder at the top. From the result, it is concluded that by applying ultrasonic wave onto the two-phase mixture, the mass transfer could be enhanced considerably, which is attributed to the increased interfacial area because of breakup. Kanai et al. [31] studied two-phase mass transfer in molten salt for bubble column application. The salt composition used in the experiment was a combination of Li_2CO_3, Na_2CO_3 and K_2CO_3. The concentration of CO_2 in the off gas was measured with an infrared detector and the concentration of CO_2 in the liquid was calculated based on mass balance and time integration. The sensitivity of temperature and gas injection rate on the mass transfer rate is reported. A decrease of mass transfer coefficient is found with increasing temperature, which is different from the observation in common air-water bubble columns. Chen et al. studied the effect of salt properties on the volumetric mass transfer coefficient with computational fluid dynamics [32]. Apart from the studies of mass transfer, the design of a pipeline gas separator is also studied. Contrary to the pipeline separator design, Funahashi et al. proposed a separator with pick-off rings and a vertical cylinder [33]. Zheng et al. carried out a series of experiments with an inclined tube [34]. Moreover, experiments and analysis were carried out regarding the general thermal hydraulics of molten salt. Bardet et al. performed an analysis for possible scaling fluid in MSR development [35]. Salt loops were built at University of Wisconsin to study the corrosion, heat transfer characteristics and natural circulation of molten salt [36, 37]. Forced convection loops with FLiNaK and nitrates, and a natural circulation loop with nitrate were constructed and tested at Shanghai Institute of Applied Physics as part of their TMSR program [38]. A natural circulation experiment with FLiNaK and heat pipes was conducted by Liu et al. [39]at Xi'an Jiaotong University.

2.1.2. Modeling of Mass Transfer in Multiphase Flow

In order to design a xenon removal device, information on local distribution of species concentration, relative velocity and void fraction are needed to determine the overall transport rate. Due to the challenge and cost of carrying out experiment with molten salt and xenon, designing and testing a gaseous fission product removal device completed based on experiment is continually becoming less practical. Therefore, simulation approach becomes a natural choice for verifying the feasibility of this technology. With the help of computational fluid dynamics (CFD), a design analysis of fission product removal system could be carried out in a cost-effective way. Typically, in dealing with large scale two-phase flow phenomenon, Eulerian-Eulerian two fluid model is utilized.

However, the modeling of two-phase flow mass transfer is far from mature. Despite extensive simulation and experimental research, a consensus on how to model the interfacial forces, turbulence, bubble diameter and even interpretation of the governing equations is hardly reached. For a detailed discussion on this topic, one could refer to some existing literature reviews [40-42]. As for mass transfer, this issue becomes more subtle, since the rate of transfer not only relies on the mass transfer coefficient, but also relies on liquid velocity, void fraction distribution and

interfacial area concentration. Without proper modeling of the quantities mentioned before, validation of mass transfer models is hardly convincing. Nevertheless, various mass transfer models have been used for CFD simulation in literature. In an early attempt to simulate mass transfer in bubble column, Krishna et al. [43] adopted a constant mass transfer coefficient of 4×10^{-4} m^2/s and reported good agreement with experimental data. Wang et al. [44] examined mass transfer coefficient based on penetration theory and surface renewal model. Wiemann et al. [45] chose an empirical correlation of Sherwood number to calculate the coefficient. In general, the mass transfer models used in CFD simulation are mostly from the studies for single bubble mass transfer. The major improvement compared with the 1D or 0D mass balance analysis with theoretical or experimental correlation is that the local distribution of velocity, void fraction and turbulent quantities are considered. The transient behavior of the system could be properly captured. With physical constitutive models, the CFD simulation could be applied to more complicated flow conditions and geometries. Additional description of mass transfer modeling is included in work of Kulkarni [46] and Rzehak et al. [47].

In the field of nuclear engineering, ORNL carried out some studies on mass transfer in pipe flow. Preliminary experiments and analysis were carried out to study this two-phase mass transfer mechanism after the MSRE experiment [28, 48], the focus of which was to obtain an experimental correlation for mass transfer in bubbly pipe flow that can be further used in the design of a xenon removal system. Peebles [25] carried out a theoretical study for the mass transfer between liquid and spherical bubble. The author first analyzed the simple case of a spherical bubble moving in a stagnant liquid, considering mobile and rigid bubble interface. Assumptions are made for the velocity field in order to arrive at a solution. However, these assumptions would not be correct in a turbulent liquid. Moreover, the relative velocity between phases used in the derivation is only applicable to laminar flow. In the large-scale xenon removal system, it is expected that the flow would be turbulent. In order to make an estimation, the author related the relative velocity between phases in turbulent cases to the fluctuation of liquid velocity. This assumption is questionable since the derived relative velocity is not related to the orientation of the flow. The derived turbulent relative velocity is plugged back into the laminar flow solution, which is expected to underestimate the mass transfer since the vortex near the bubble interface is not considered. After a review of the literature, the author stated that the knowledge at that time was not enough to arrive at a firm analysis of the mass transfer process, and further experiments were suggested. Subsequently, a more thorough study was carried out by Kress [48]. In this study, a comprehensive review of the single bubble mass transfer analysis and pipe flow experiment is first conducted. The scaling of pipe Sherwood number to Schmidt number is reported in the range of $1/3$ to $1/2$. The power dependence on Reynold number, on the other hand, is reported to be 0.9 to 1.1 for conduits, and 0.6 to 0.8 for stirred vessels. Dissolved oxygen experiments were carried out to measure the mass transfer coefficient and to separate the coefficient from the interfacial area concentration. The measurement of interfacial area was obtained from images taken at the region of interest. By assuming a fixed bubble distribution and estimating the gas void fraction from experimental correlation, the interfacial area is derived from the bubble number density. This is simply done by counting the number of bubbles in the images, which is much easier than measuring the size directly from the image. From the experimental result, a correlation is proposed for the Sherwood number as

$$Sh_L = 0.34\, Re_L^{0.94} Sc^{0.5} \left(\frac{d}{L_c}\right), \tag{2-1}$$

where the characteristic length L_c used here is the hydraulic diameter of the pipe. Additionally, with the usage of the surface renewal model, a similar relation is obtained as

$$Sh_L \sim Sc^{0.5} Re_L^{0.92} \left(\frac{d}{L_c}\right)^{\frac{1}{3}}. \quad (2\text{-}2)$$

In arriving at the experimental and theoretical results, many assumptions were made along with the usage of experimental correlation. The local distribution of flow variables is also ignored, which could be a source of uncertainty in the results.

Based on the discussion above, in order to gain confidence in the simulation, model validation with experiments is preferable and necessary. The validation should not only include the mass transfer rate, but also other flow quantities. Regarding this issue, experiments on bubble columns are helpful. Bubble columns refer to the kind of apparatus where gas is sparged into a liquid-filled vessel in form of bubbles. Various multiphase phenomenon could occur during the operation of a bubble column, including bubble deformation, bubble induced turbulence, bubble coalescence, bubble breakup, and mass transfer across the interface. The large interfacial area between bubbles and the liquid phase makes bubble column a great multiphase reactor and is therefore widely used in chemical, metallurgical and pharmaceutical industries [49]. With its simplicity, different shapes of bubble columns are also widely used in research on two-phase mass transfer and general two-phase flow dynamics. Early experiments of bubble columns generally targeted at global characteristic like total gas hold up and volumetric mass transfer coefficient [50, 51]. With the development of flow measurement technique, localized data of the flow field and phase distribution became available. Computer automated radioactive particle tracking (CARPT) was used by Devanathan et al. [52] and Degaleesan et al. [53] to obtain liquid phase velocity distribution. In the meantime, long existing methods of X-ray tomography and conductivity probe was adopted by Kumar et al. [54] and Buwa et al. [55] to measure void fraction distribution. With the help of particle image velocimetry, simultaneous measurement of liquid and gas phase velocity was carried out by Hassan et al. [56]. Later improvement of the laser technique and high-speed camera allows for high time resolution measurement [57]. More sophisticated measurement arrangement with two cameras greatly improves the accuracy of phase discrimination, and more accurate measurement of the flow field could be obtained [58]. Another widely used method is flow visualization and image processing, sometimes coupled with PIV measurement. Two-phase flow visualization could supply information on phase distribution, bubble size distribution and information on gas phase velocity [47, 55, 59, 60]. As for mass transfer measurement, by virtue of the development of dissolved gas sensors, local concentrations can be measured with fair accuracy and time resolution, such as widely done for dissolved oxygen [61-63].

2.2. Physical Model Description

The physical model used in this study is the Eulerian two-fluid model coupled with species transport to account for the xenon removal. A detailed description of the constitutive models and corresponding sensitivity analysis is given in a previous report [64]. The model selected in the validation studies are summarized in Table 2.1. For the bubble separator, modifications are made to incorporate the coalescence and breakup of the bubbles. These changes are described in detail in this section. The turbulent dispersion force and the mass transfer source term are implemented with field functions in STAR CCM+.

Table 2.1. Constitutive models used in the pipeline removal simulation.

Mechanism	Constitutive Model
Turbulence	Realizable $k-\epsilon$ + Issa's Turbulence Response [65]
Bubble Induced Turbulence	Sato's Bubble Induced Turbulence [66]
Drag	Tomiyama with Richardson-Zaki Drag Correction [67, 68]
Lift	Tomiyama [69]
Turbulent Dispersion	Lopez de Bertodano [70]
Virtual Mass	Not Included
Wall Lubrication	Not Included
Mass Transfer Coefficient	Small Eddy Cell Model [71]
Henry's Law Coefficient $\boldsymbol{K_c}$	Blander et al. [72]

2.2.1. Governing Equations

The general conservation equation of the two-fluid model is,

$$\frac{\partial \alpha_k \rho_k \, \phi_k}{\partial t} + \nabla \cdot (\alpha_k \rho_k \, \phi_k \boldsymbol{u_k}) = -\nabla \cdot \left(\alpha_k \left(\boldsymbol{J_k} + \boldsymbol{J_k^T}\right)\right) + I_{k,\phi}. \tag{2-3}$$

$\boldsymbol{J_k}$ is the flux for ϕ_k, $I_{k,\phi}$ is the source term, and $\boldsymbol{J_k^T}$ is the turbulence flux from the derivation. For a more rigorous derivation, the source term could be split into the interaction between phases and pure source from the environment. The corresponding equations for mass conservation, momentum and species transport are included in the simulation. The mass conservation equation is,

$$\begin{gathered} \frac{\partial \alpha_l \rho_l}{\partial t} + \nabla \cdot (\alpha_l \rho_l \boldsymbol{u_l}) = \Gamma_l, \\ \frac{\partial \alpha_g \rho_g}{\partial t} + \nabla \cdot \left(\alpha_g \rho_g \boldsymbol{u_g}\right) = \Gamma_g, \\ \Gamma_g + \Gamma_l = 0. \end{gathered} \tag{2-4}$$

The source term Γ is related to the mass transfer at the interface, which could either be related to phase change or chemical equilibrium. The third equation reflects the conservation of mass at the interface where it is assumed that the interface does not store mass.

The momentum conservation equation is,

$$\begin{gathered} \frac{\partial \alpha_k \rho_k \boldsymbol{u_k}}{\partial t} + \nabla \cdot (\alpha_k \rho_k \boldsymbol{u_k} \boldsymbol{u_k}) = -\nabla(\alpha_k p_k) + \nabla \cdot \left(\alpha_k (\mathfrak{G}_k + \mathfrak{G}_k^T)\right) + \alpha_k \rho_k \boldsymbol{g} + \boldsymbol{M_k}, \\ \boldsymbol{M_g} + \boldsymbol{M_l} = \boldsymbol{M_m}. \end{gathered} \tag{2-5}$$

The term $\mathfrak{G}_k$ and $\mathfrak{G}_k^T$ are the viscous and turbulent stress tensor. The second equation describes the relationship between the source terms. The summation of the source terms is not equal to zero,

indicating that the interface can store momentum. The term $\boldsymbol{M_m}$ represents the contribution from surface tension, when external momentum source such as magnetic force does not exist.

The transport equation for the mass fraction of species i in phase k is,

$$\frac{\partial \alpha_k \rho_k Y_{k,i}}{\partial t} + \nabla \cdot \left(\alpha_k \rho_k Y_{k,i} \boldsymbol{u_k}\right) = \nabla \cdot \left(\alpha_k \rho_k \left(D_{k,t} + D_{k,i}\right) \nabla Y_{k,i}\right) + \Gamma_{k,i},$$

$$D_{l,t} = \frac{\mu_t}{\rho_l Sc_t}, \qquad Sc_t = 0.7. \tag{2-6}$$

$Y_{k,i}$ is the mass fraction of species in phase k, $D_{k,t}$ is turbulent diffusivity, $D_{k,i}$ is the molecular diffusivity, and $\Gamma_{k,i}$ is the mass flux of species i transferred to phase k. The source term also appears in the continuity equation for phase k. The source term in the species transport equation is modeled according to the surface renewal stretch model [73] and the small eddy cell model [71],

$$\Gamma_{l,i} = k a_i \left(K_\rho^i \rho_g Y_g^i - \rho_l Y_l^i\right),$$

$$k = C \sqrt{D_m \sqrt{\frac{\epsilon}{\nu}}}, \qquad C_{SRS} = \sqrt{\frac{4}{\pi}}, \qquad C_{SEC} = 0.4. \tag{2-7}$$

Temperature variation during the normal operation of MSR certainly exists, but the time scale for major changes is expected to be much longer than the flow residence time in the removal system. Therefore, in the present study, the phases are assumed to be isothermal, and the energy equation is not included in the calculation.

2.2.2. Interfacial Area Concentration and Bubble Diameter

Transport between phases is controlled by interfacial area concentration a_i. Thus, precise prediction of a_i is crucial in two-phase flow modeling. There are multiple ways of computing a_i in existing literature. One method is establishing a transport equation for the interfacial area concentration, as described by Ishii et al. [74]. This method which is still under development requires a proper boundary condition and initial condition for a_i. The other method to calculate a_i is an algebraic equation including a bubble diameter d. Using spherical bubble assumption, the interfacial area is calculated as,

$$a_i = \frac{6\alpha}{d}, \tag{2-8}$$

where α is the dispersed phase void fraction. This equation is appropriate when the void fraction is low. Additionally, in the pipeline removal process, the bubble interaction is limited, meaning d could be treated as a constant design parameter related to the bubble generator.

However, in the bubble separator, the bubbles are collected at the center of the pipe where significant breakup and coalescence exist. Therefore, the constant bubble diameter and the spherical bubble assumption are no longer valid. To account for the reduction of interfacial area in the bubble separator, the single-speed adaptive multi-size group (AMUSIG) model and symmetric interfacial area equation are used. The interfacial area in the bubble separator is calculated as,

$$a_i = \frac{6\alpha(1-\alpha)}{d_{sm}}. \tag{2-9}$$

The additional $(1-\alpha)$ factor is included to ensure the interfacial area decreases to zero when the void fraction reaches one. The bubble diameter is defined as the Sauter mean diameter, calculated from the additional single-speed AMUSIG model,

$$\frac{\partial n_i}{\partial t} + \nabla \cdot \left[n_i \left(\boldsymbol{u_i} + D_T \nabla \ln \frac{\alpha_i}{n_i} \right) \right] = S_i,$$

$$\alpha_i = \frac{\alpha}{N}, \qquad \boldsymbol{u_i} = \boldsymbol{u_g}.$$

$$d_{sm} = \frac{\sum n_i d_i^3}{\sum n_i d_i^2}. \tag{2-10}$$

In this equation, n_i is the number density of ith bubble size group, D_T is the number density diffusion coefficient, and S_i is the source term from the coalescence and breakup model. The additional diffusive term arises from the difference between time averaged and phase averaged quantities, which shares a similar origin to the turbulent dispersion force in the momentum equation. The coalescence model used in the simulation is based on the study of Luo [75], and the breakup model is based on the study of Martínez Bazán et al. [76]

The physical modeling of the void core discussed above is based on modifications to the dispersed bubbly flow model. The combination of the symmetrical interfacial area concentration model and multi group model is phenomenological, rather than based on first principle. Its applicability is examined by comparison with literature data in the following sections.

2.2.3. Properties of Molten Salt

To simulate the xenon removal process in the molten salt reactor, the properties of molten salt and solute xenon are necessary. A short literature review is included here on the properties of FLiNaK, which is the candidate salt for Transatomic Power MSR, which is also the salt used in the ongoing xenon removal experiment. FLiNaK has a composition of 46.5% LiF, 11.5% NaF and 42% KF. This composition is chosen around the eutectic point of the mixture to minimize its melting point.

The density of FLiNaK is shown in Figure 2.1, where the dashed line indicates unknown effective range. The red bold line is used in the simulation, which is the correlation given by Janz [77]. The error bar is plotted to show the uncertainty of the experiment, if reported. The style for other figures is the same. The uncertainty about density measurement is small compared with other properties, and the difference between different measurements could be attributed to the impurities in the salt. Also, different measurements basically fall in the uncertainty region of other correlation. As can be seen from the figure, the density of FLiNaK drops with increasing temperature, as one would expect. In the simulation, the density at 923K is used, which is the operation temperature in the MSRE. The corresponding equations in the figure are given in Table 2.2.

Table 2.2: Density correlations of FLiNaK.

Correlation (kg/m^3)	Effective Range (K)	Uncertainty	Source
$2555 - 0.6T$	Unknown	Unknown	Hoffman et al. [78]
$2729.3 - 0.73T$	Unknown	$\pm 5\%$	Powers et al. [3, 79]
$2579.3 - 0.624T$	$940 - 1170$	$\pm 2\%$	Janz [77]
$2655.64 - 0.68T$	Unknown	Unknown	Cohen et al. [3, 80]
$2603 - 0.669T$	$750 - 1100$	Unknown	Salanne et al. [81]

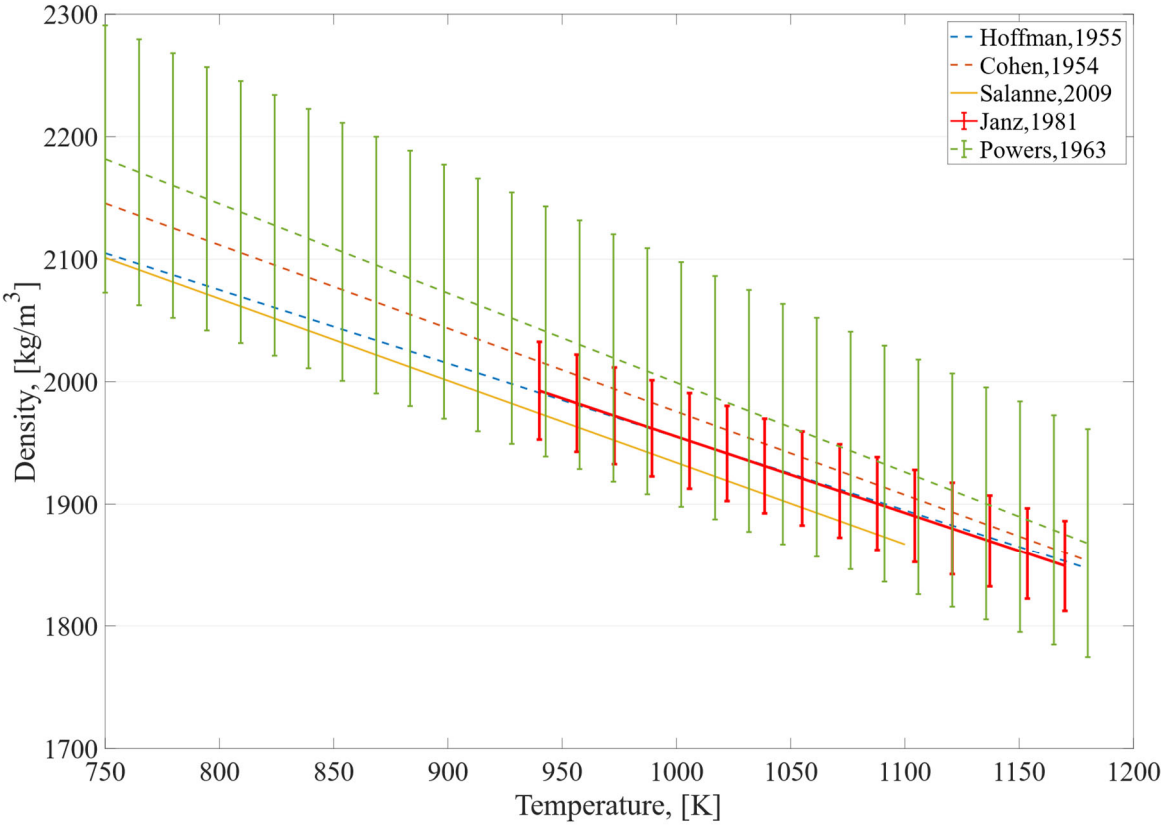

Figure 2.1: Correlations for density of FLiNaK. (Dashed lines represent correlations not based on experiments. The length of solid corresponds to the applicable range of experimental correlations. The red bold line is used in Chapter 2 when calculating the properties at different temperature.)

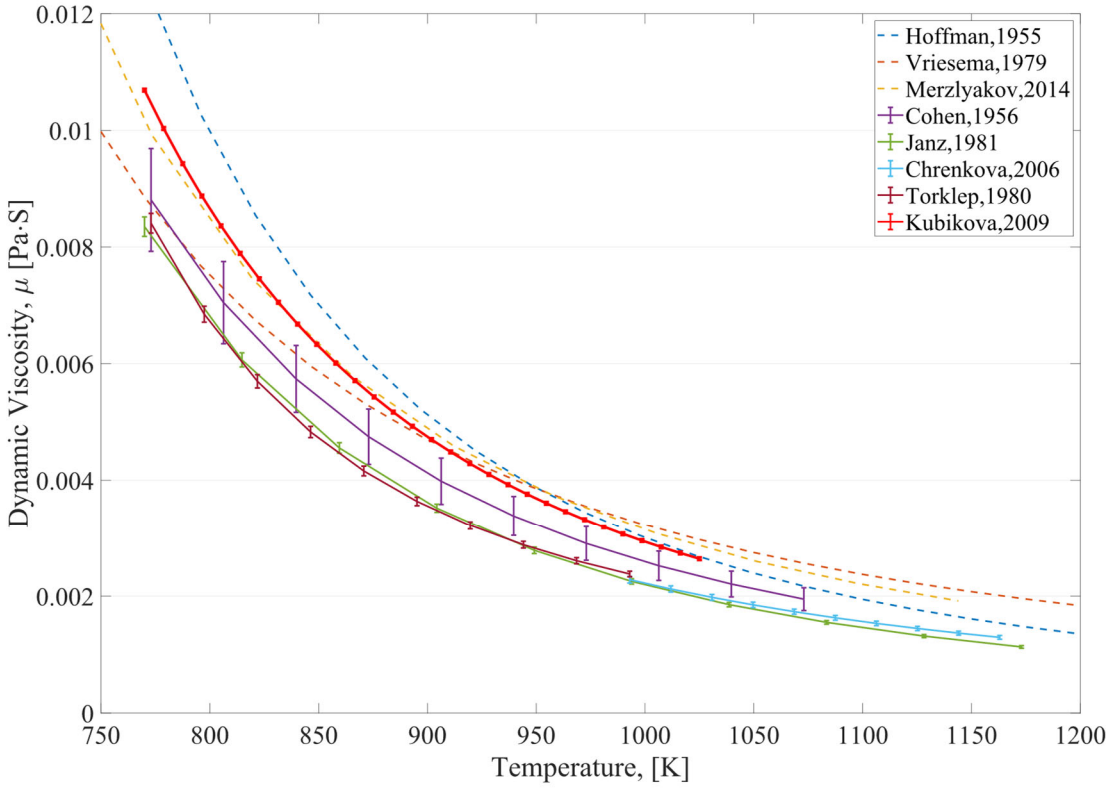

Figure 2.2: Correlations for viscosity of FLiNaK. (Dashed lines represent correlations not based on experiments. The length of solid corresponds to the applicable range of experimental correlations. The red bold line is used in Chapter 2 when calculating the properties at different temperature.)

The viscosity of FLiNaK is quite sensitive to the temperature of the salt as shown in Figure 2.2. Especially, the viscosity changes significantly near the melting point. This sensitivity is not present in water-based heat transfer system, which might require special care when performing a safety analysis. As for the temperature used in the simulation, at 923K, the viscosity is around five times of that of water at room temperature. Considering the density of the salt, the kinematic viscosity is around 2.5 times of water at room temperature. For a scaling experiment, 30% $CaCl_2$ solution has been used in literature to match the kinematic viscosity of the liquid. The correlation given by Kubíková et al. [82] is selected in this document. The difference between measurements is larger than that for density, which is likely from the impurities in the salt. The well-known Stokes–Einstein equation indicates that the viscosity of a mixture will increase with impurities. Alternatively, the different could be caused by the slight difference in the salt composition used in different experiments.

Table 2.3: Viscosity correlations of FLiNaK.

Correlation ($kg/m \cdot s$)	Effective Range	Uncertainty	Source
$2.5 \times 10^{-5} e^{\frac{4790}{T}}$	Unknown	Unknown	Hoffman et al. [78]

Correlation ($kg/m \cdot s$)	Effective Range	Uncertainty	Source
$4 \times 10^{-5} e^{\frac{4170}{T}}$	$773 - 1073$	10%	Cohen et al. [83]
$3.877 \times 10^{-5} e^{\frac{4327}{T}}$	$770 - 1025$	2.4×10^{-5}	Kubíková et al. [82]
$1.1 \times 10^{-4} e^{\frac{3379}{T}}$	Unknown	Unknown	Vriesema [84]
$2.487 \times 10^{-5} e^{\frac{4478.62}{T}}$	$770 - 1173$	$\pm 2\%$	Janz [77]
$6.23 \times 10^{-5} e^{\frac{3921.4}{T}}$	$727 - 1144$	Unknown	Merzlyakov et al. [85]
$e^{-3.049} \times 10^{-3} e^{\frac{3847}{T}}$	$993 - 1163$	$\pm 2.5\%$	Cibulková et al. [86]
$1.633 \times 10^{-3} e^{\left(-\frac{2762.9}{T} + \frac{3.115E6}{T^2}\right)}$	$773 - 993$	$\pm 2\%$	Tørklep et al. [87]

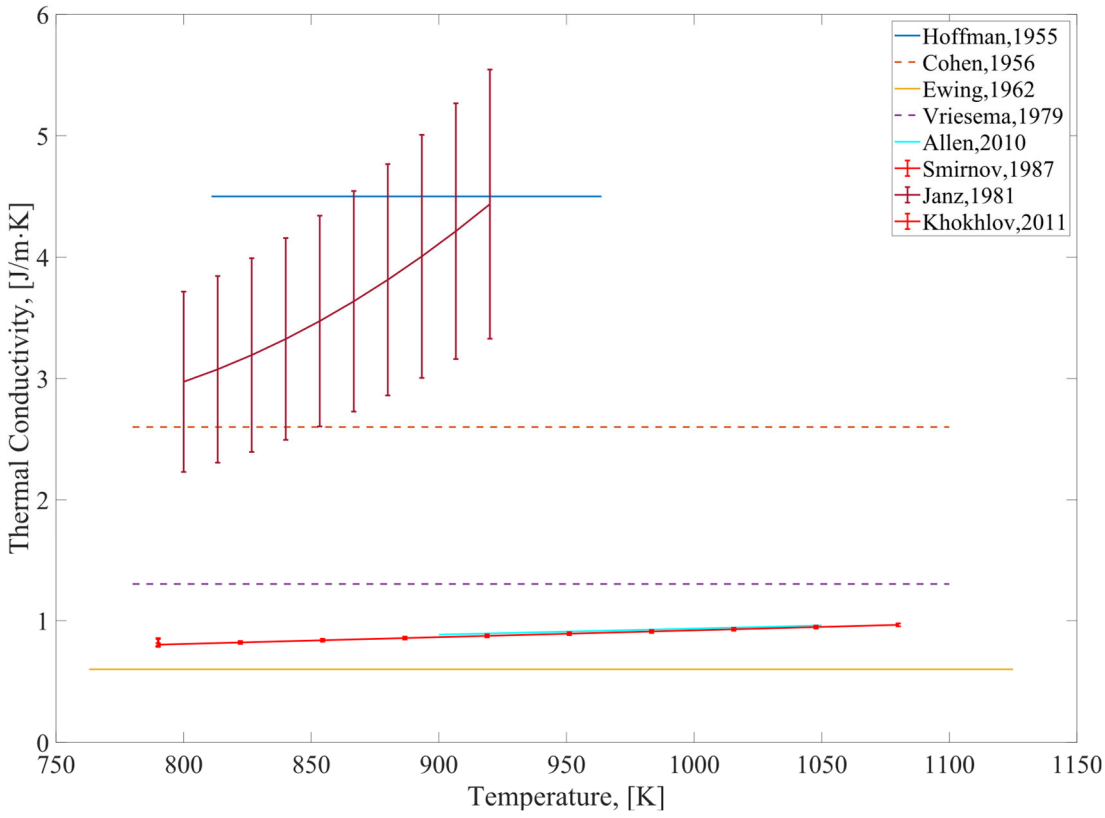

Figure 2.3: Correlations for thermal conductivity of FLiNaK. (Dashed lines represent correlations not based on experiments. The length of solid corresponds to the applicable range of experimental correlations. The red bold line is used in Chapter 2 when calculating the properties at different temperature.)

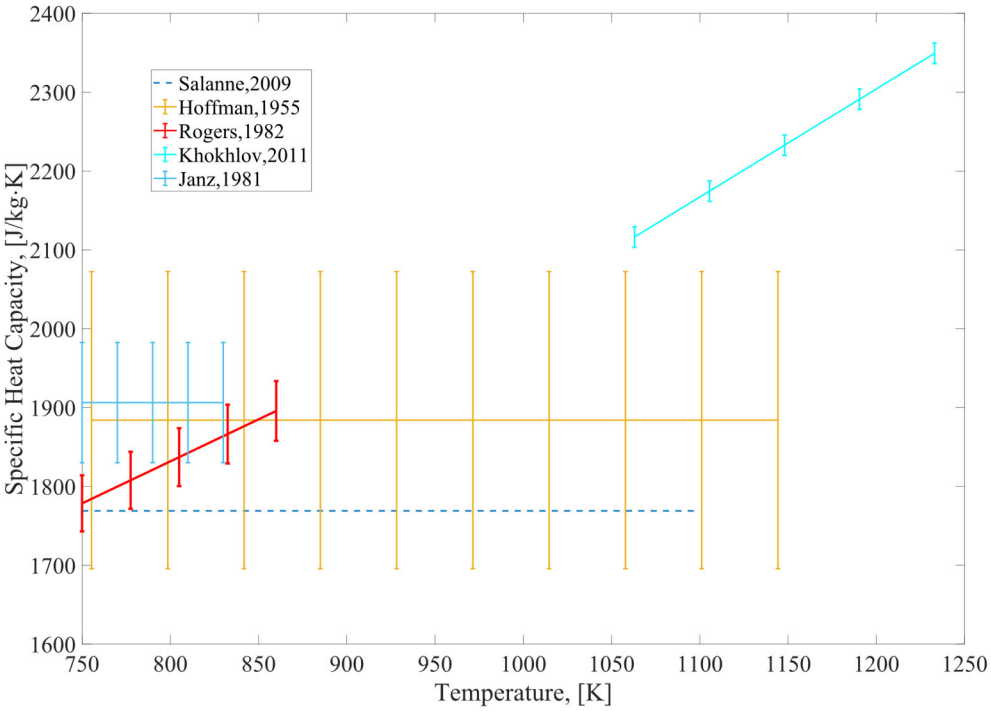

Figure 2.4: Correlations for specific heat capacity of FLiNaK. (The length of solid corresponds to the applicable range of experimental correlations. The red bold line is used in Chapter 2 when calculating the properties at different temperature.)

The correlations for thermal conductivity and specific heat capacity are also included here, though in the current CFD model the liquid is assumed to be isothermal. From Figure 2.3 and Figure 2.4, the dispute in literature is obvious. Ewing et al. [88] argued that the difference with the measured thermal conductivity is likely from the radiation and natural convection, which is not considered in some of the early experiments [89]. Without the radiation, the thermal conductivity would be overestimated. This could explain the relatively large value given in the figure. In this document, the recent correlation given by Khokhlov et al. [90] is recommended, which basically predicts the same thermal conductivity given by Smirnov et al. [91].

Table 2.4: Thermal conductivity correlations of FLiNaK.

Correlation, $[J/kg \cdot K]$	Effective Range	Uncertainty	Source
4.5	$810.93 - 963.71$	Unknown	Hoffman et al. [78, 92]
2.6	Unknown	Unknown	Cohen et al. [83]
0.6	$763 - 1125$	Unknown	Ewing et al. [88]
1.3	Unknown	Unknown	Vriesema [84]
$0.36 + 5.6 \times 10^{-4}T$	$790 - 1080$	± 0.012	Smirnov et al. [91]
$24.288 - 6.042 \times 10^{-2}T + 4.222 \times 10^{-5}T^2$	$800 - 920$	$\pm 25\%$	Janz et al. [77]
$-0.34 + 0.5 \times 10^{-3} + 32/M$	$790 - 1080$	$\pm 4\%$	Khokhlov et al. [90]

The specific heat capacity of FLiNaK is shown in Figure 2.4. In the review by Williams [93], it is argued that the dependency of heat capacity with temperature is small. Salanne et al. [81] reported a constant heat capacity of $1769 J/kg \cdot K$, and similarly Hoffman et al. [78, 92] reported a constant value of $1884.06 J/kg \cdot K$. On the other hand, the data taken by Rogers et al. [94] suggests that the heat capacity increases with temperature in the operation range of MSR. Nevertheless, the existing measurements roughly fall within each other's uncertainty range. In this document, the correlation by Rogers et al. [94] is suggested.

Table 2.5: Specific heat capacity correlations of FLiNaK.

Correlation $(J/kg \cdot K)$	Effective Range	Uncertainty	Source
1884.06	$755.37 - 1144.26$	$\pm 10\%$	Hoffman et al. [78, 92]
$979.22 + 1.0657T$	$750 - 860$	$\pm 2\%$	Rogers et al. [94]
1906.3	$750 - 830$	$\pm 4\%$	Janz et al. [77]
$660 + 1.37T$	$1063 - 1233$	± 13	Khokhlov et al. [90]
1769	$750 - 1100$	Unknown	Salanne et al. [81]

Table 2.6: Surface tension correlation of FLiNaK.

Correlation	Effective Range	Uncertainty	Source
$0.2726 - 1.014 \times 10^{-4} T$	$770 - 1040$	$\pm 2\%$	Janz et al. [77]

The surface tension of the salt is important in the calculation of drag for distorted bubbles and the bubble diameter. Only limited data could be found for this quantity. The correlation reported by Janz et al. [77] is used in this document.

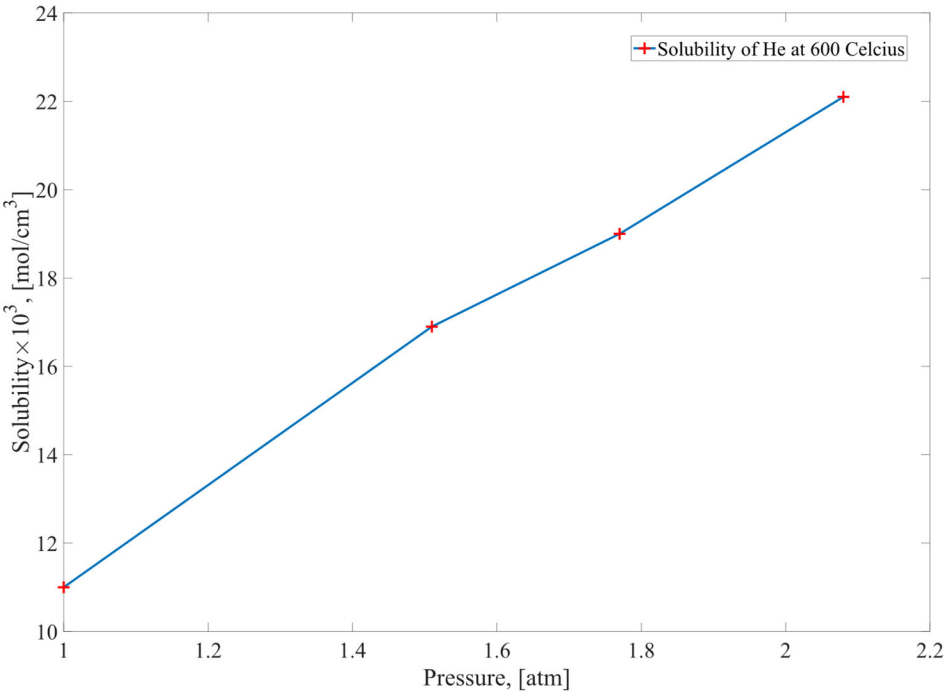

Figure 2.5: Solubility of helium in FLiNaK with partial pressure based on the experiment of Blander et al. [72] at fixed temperature.

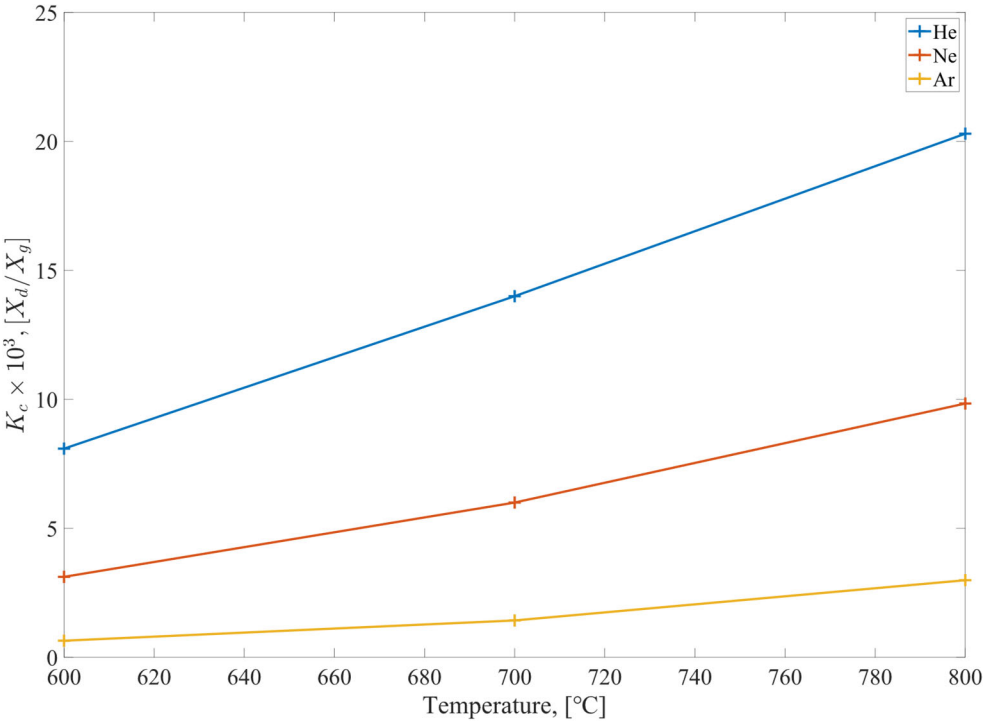

Figure 2.6: Solubility of noble gases in FLiNaK at different temperature based on experiment of Blander et al. [72].

The solubility data of noble gas in molten salt is quite limited. In terms of solubility of noble gas in FLiNaK, only one experiment could be found in the literature, which is conducted by Blander et al. [72] at ORNL. It is found in the experiment that the addition of 4% UF4 to the salt has basically no impact on the solubility. The solubility of gases in liquid could generally be expressed by Henry's law, that the equilibrium molar fraction of the dissolved gas is proportional to the gas' partial pressure in the gas phase as

$$P_g^i = HX_{l,i}^e. \tag{2-11}$$

This relation is usually valid in low pressure and dilute mixture cases. Blander et al. [72] measured the solubility of Helium in FLiNaK at different partial pressure, as shown in Figure 2.5. The linear relationship between the solubility and partial pressure is a confirmation of Henry's law for the inert gas-FLiNaK system.

Then, the solubility of He, Ne, Ar in the fluoride at different temperatures are experimentally measured, as shown in Figure 2.6. The solubility data are expressed in terms of $K_{k,i}^c$, which is the equilibrium ratio in Henry's law in different formats, as

$$K_{k,i}^{\rho} = \frac{\rho_k Y_{k,i}^e}{\rho_j Y_{j,i}^e} = K_{k,i}^c. \tag{2-12}$$

The equilibrium ratio will increase with temperature and decrease with molecule mass or diameter, as shown in the figure.

Xenon solubility data for FLiNaK are not available, but the theoretical values are listed, as shown in Table 2.7. For none reacting gas-liquid system, the solubility of gases could be calculated based on classic chemical equilibrium theory as

$$RTlnK^{\rho} = T\Delta S - \Delta H, \tag{2-13}$$

$$lnK^{\rho} = -\frac{\Delta\mu}{k_B T}.$$

According to the second equation, the solubility could be calculated with the chemical potential $\Delta\mu$ of the dissolution process. An early theory assumes the dissolution is governed by the "surface energy" to form cavities in the salt, which is used by Blander et al. [72, 95] A more sophisticated model would consider the effect of volume exclusion, dispersion interaction and polarization of gas and salt [96]. Nevertheless, the theoretic value given here is likely within the right magnitude according to the comparison of theoretic value and measurement for other gases. Therefore, the theoretical solubility of Xenon would be used in our simulation, until experimental becomes available.

Table 2.7: Solubility of noble gases in FLiNaK [72].

Gas	Temp. (℃)	Measured $K_{k,i}^{\rho} \times \mathbf{10^3}$	Calculated $K_{k,i}^{\rho} \times \mathbf{10^3}$
He	600	8.09	28.3
	700	14.0	46.8
	800	20.3	70.7
Ne	600	3.12	3.94
	700	6.00	8.63
	800	9.84	16.4
Ar	600	0.645	0.146
	700	1.43	0.509
	800	2.99	1.41
Xe	600	-	0.011
	700	-	0.057
	800	-	0.212

2.3. Validation Experiments

Limited data for validation of CFD is available in literature, and what is available only includes measurement of one or two of the three important quantities related to liquid-gas mass transfer. For experiments focused on hydrodynamics, void fraction and liquid velocity profiles might be available. For mass transfer experiments, void fraction and volumetric mass transfer coefficient are usually measured. Though some simulation works validated their model in these three aspects [44], the fact that different geometries and conditions were used in those experiments would inevitably reduce the fidelity of the validation. In the attempt to fill this gap and provide confidence for our CFD model, a series of experiments on the same apparatus have been carried out with flow visualization, PIV and dissolved oxygen measurement. The apparatus used in these experiments is a simplified version of an ongoing two-phase mass transfer experiment in molten salt, as a first step towards the design of a gaseous fission product removal system. The experimental data obtained here will be used to validate our CFD model, which will be further

applied to the molten salt experiment for comparison. The validated model with information supplied by the molten salt experiment is used in the design analysis of the fission product removal system later in this chapter.

The experiments are carried out in a small-scale cylindrical bubble column made of acrylic, as shown in Figure 2.7. The bubble column is 2.5 inch in diameter and 6.55 inch in height. This aspect ratio is slightly smaller compared with industrial cylindrical bubble columns but is adopted anyway in accordance with the ongoing molten salt experiment. Different pieces of acrylic are cut and glued together. A 0.12-inch hole at the center of the bottom serves as the gas inlet. A corresponding acrylic flange could be attached at the top of the bubble column, where sparging tube and other instruments could be installed. This design is to replicate the other ongoing molten salt experiment with air-water, as shown in Figure 2.8. For the model validation experiments, this setup introduces unnecessary complexity to the geometry modeling. Therefore, in most of the experiments, the gas is injected from the bottom. In the experiment, the bubble column is filled with deionized water. The gas used in this experiment is air, from the laboratory compressed air supply. The flowrate is measured with Dwyer RMB-49-SSV panel mounted gas flow meter, which has a measurement range from 0.5SCFH to 15SCFH and 5% accuracy. The cylinder is surrounded by a square water box filled with water to limit the refraction caused by the column curvature. With the water box and thin wall thickness, the influence of refraction becomes negligible. This is confirmed by pictures of an immersed standard ruler at different location. The spacing of the ticks is practically uniform. The actual physical length for pixels could be easily derived based on the known dimensions of the bubble column. However, for the part that is not covered by the water box, the refraction is not negligible, and the corresponding area in the obtained videos is cropped out. In the following section, procedures and results of each experiment will be discussed in detail.

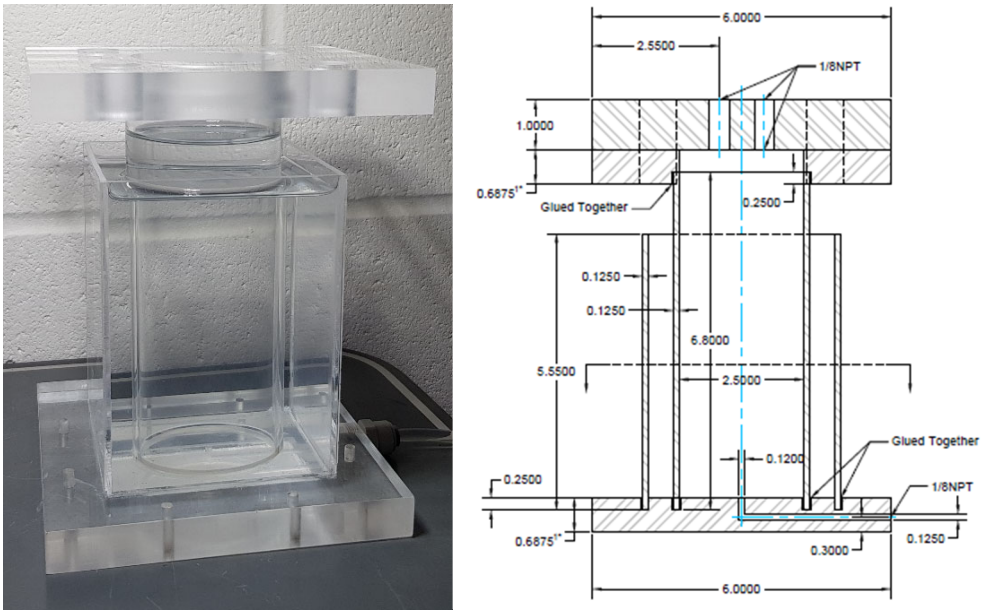

Figure 2.7: Bubble column with water box used in our experiment, dimensions based on the corresponding molten salt experiment shown in Figure 2.8. Sparging in this setup is from bottom.

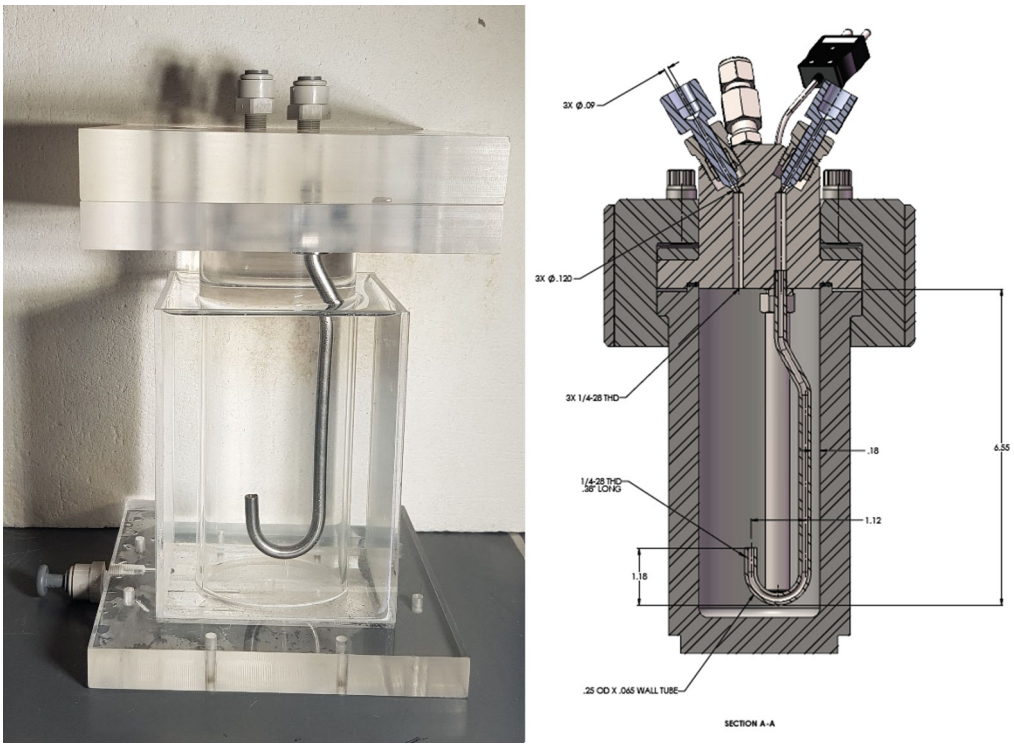

Figure 2.8: Comparison of the bubble column with flange and the molten salt xenon experiment. Sparging in this setup is from the tube.

2.3.1. Visualization Experiment

Flow visualization with high-speed camera is widely used in two-phase flow experiments. Visualization can provide qualitative information of bubble movement, morphology and phase distribution in a wide region, which are often useful for flow regime identification and could infer more accurate measurement. Research is also made for quantitative measurement of bubble velocity and size distribution with visualization, though the accuracy is often limited. In this section, the technique for time averaged void fraction measurement with flow visualization is introduced. The results are compared with existing experimental correlation as a validation.

2.3.1.1. Experiment Procedure

The flow visualization experiments were carried out by sparging in air bubbles at different flowrates from the bottom. Photron's high speed camera MINI AX100 is used in our experiments to generate the high-speed videos. A frame rate of 250 fps is found to be enough for bubble identification, tracking, and accurate measurement of void fraction. The maximum measurement time of the camera is limited by its internal memory. Therefore, a higher frame rate is not favorable since the total averaging time will be reduced.

Proper illumination of the flow region is important to obtain a high-quality video for post processing. First is the selection of light source. At high frame rate, the change in luminosity of the incandescent light bulb caused by the AC power cycle becomes noticeable, which is troublesome for post processing. One solution is to use DC power with the incandescent light bulb,

which would be quite expensive. The other choices are halogen photography lamps [97] and LED light [59]. The frequency of LED light is controlled by the electronic circuits driving them, which can create a vast array of refresh rates. Some high-power LED light is free of this flickering, and SANSI's 5500lm flood light is used in our experiment. Apart from the selection of light, the illumination method is also important. To obtain a uniform background, the light source is usually placed behind a light screen and the flow region. The uniform backlight would be reflected at the bubble interface, making the boundary of the recorded bubbles appears darker than the center and the background. This distinction in brightness would allow for bubble identification. The intensity of backlight should be such that this contrast is maximized.

Prior to the bubble injection, a short video of the background was taken for contrast. The background image is calculated by averaging the pre-recorded high-speed video, which could minimize the disturbance of the light source. This background image is important for bubble identification, which will be discussed in the next section.

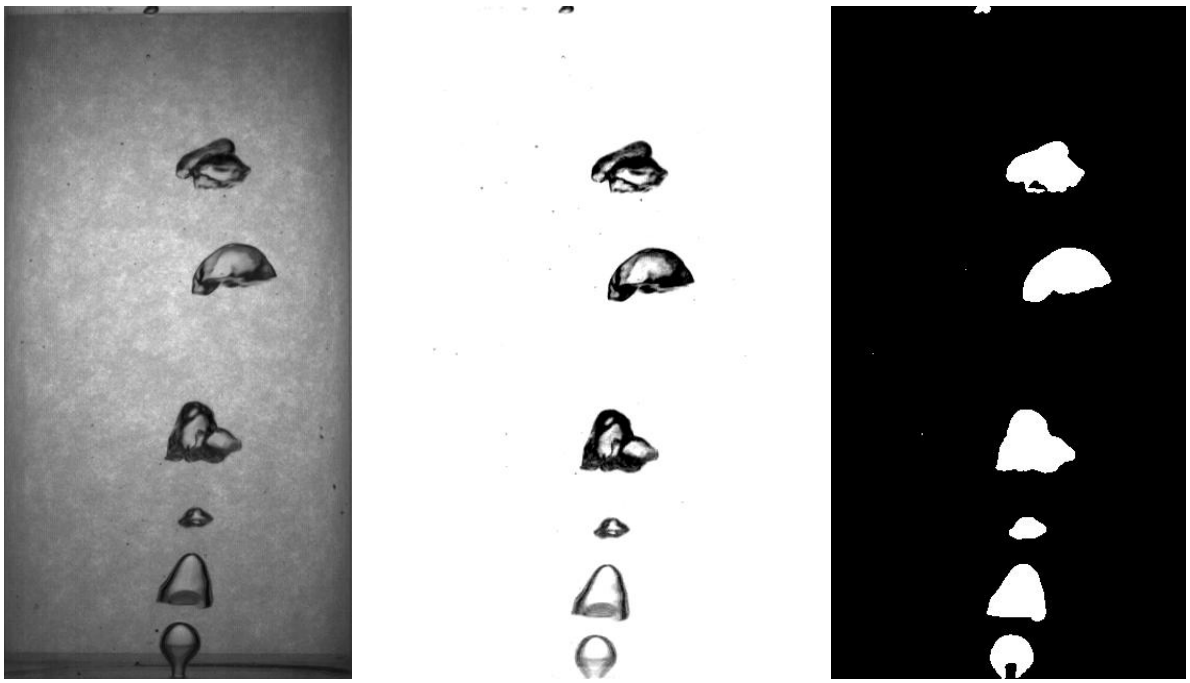

Original Image Enhanced Image Binarized and Filled Image

Figure 2.9: Processing procedures before bubble identification. Enhanced image is obtained by subtraction of background and applying adaptive histogram equalization. The final image is obtained by inversing the enhanced image, performing binarization based on continuity and filling the identified regions.

2.3.1.2. Image Processing

The videos are processed in MATLAB to extract the information of bubbles. These videos are taken in 8bit grayscale model. The brighter pixel possesses a higher intensity, and the overall level scale is from 0 to 255.

First, the videos are cropped to the region where the water box exists, since the refraction through the wall could be corrected. Then, the pre-recorded background is subtracted from each frame. Since the background is brighter than the bubbles, this subtraction is not directly the subtraction of the grayscale value. The easiest way to perform this subtraction is to first inverse the images and then make the subtraction. The resulting figure is inversed again to restore the normal grayscale. This process produces a figure with white background and black bubbles. The inverse of the grayscale is given as

$$I_{inv} = |255 - I|. \tag{2-14}$$

After the background extraction, the images are enhanced by adaptive histogram equalization. This algorithm performs histogram equalization on small areas of the figure to minimize the noise processed. The target histogram is a flat distribution, and the equalization would distribute the local intensity value more evenly. After the enhancing, the figure is subjected to a Gaussian filtering to remove the noise in the original figure and those created in the enhancing. At last, the obtained grayscale image is reversed and converted to binary image, where morphology algorithm becomes available. The objects in the binarized figures are identified by continuity of the white regions. Since the bubble is only darker at its boundaries, these objects often form a ring, rather than a solid region after binarization. It is then necessary to fill these hollow regions. The final filled bubbles and other images are shown in Figure 2.9.

In the binarized image, bubbles are identified by continuity of the white regions. With this identification, the equivalent diameter, estimated volume and the averaged void fraction in this region could be calculated with the following equations,

$$r_{eq} = \sqrt{\frac{A}{\pi}}, \qquad V_{est} = \frac{4}{3}\pi r_{eq}^3, \qquad \alpha_g = \frac{\overline{\sum V_{est}}}{V_{volume}}. \tag{2-15}$$

The existence of cap, spherical and elliptical bubbles at the same time poses difficulties for the identification, since distinguishing them in 2D images are not easy, as shown in Figure 2.9. Therefore, instead of specifically sorting the bubbles into groups, they are all treated equally in the processing. This is compensated by the usage of equivalent diameter. The total area is used to derive the equivalent diameter, which preserved the cross-sectional area. The error associated with this method is analyzed in the Appendix C. The measurement results for volume averaged void fraction or gas hold up at different gas flowrates and the convergence behavior of this method are shown in Figure 2.10 and Figure 2.11.

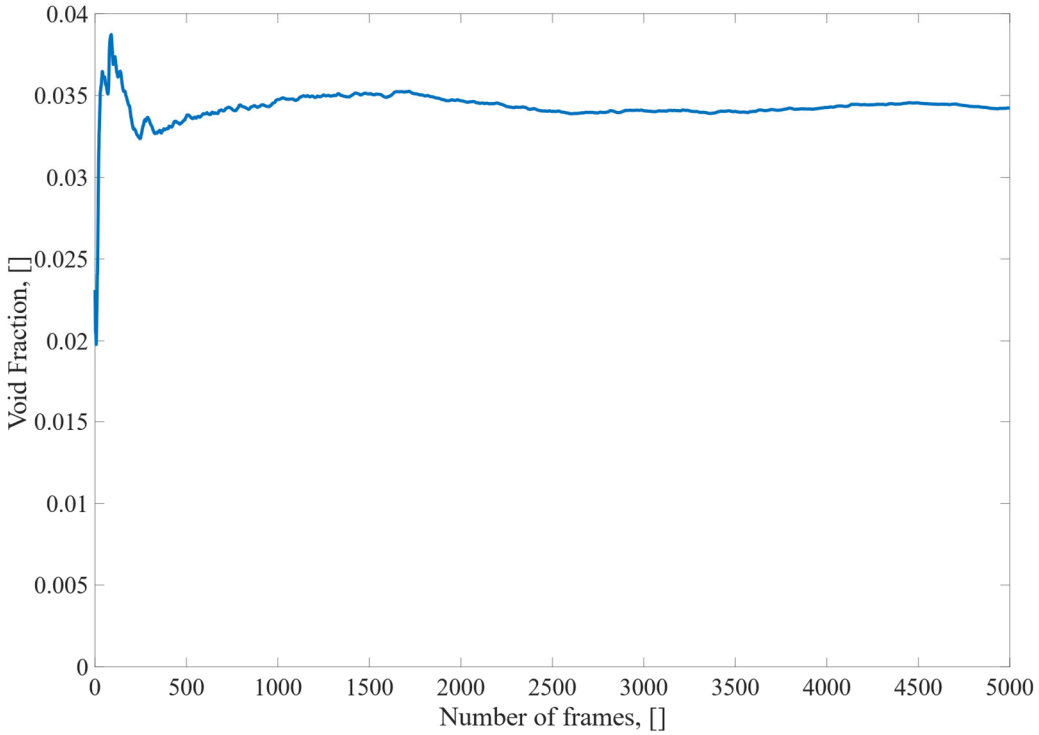

Figure 2.10: The convergence of measured void fraction at $J_g = 1\ cm/s$ with increasing frames using Equation 2-15. The small fluctuations are considered in the uncertainty analysis.

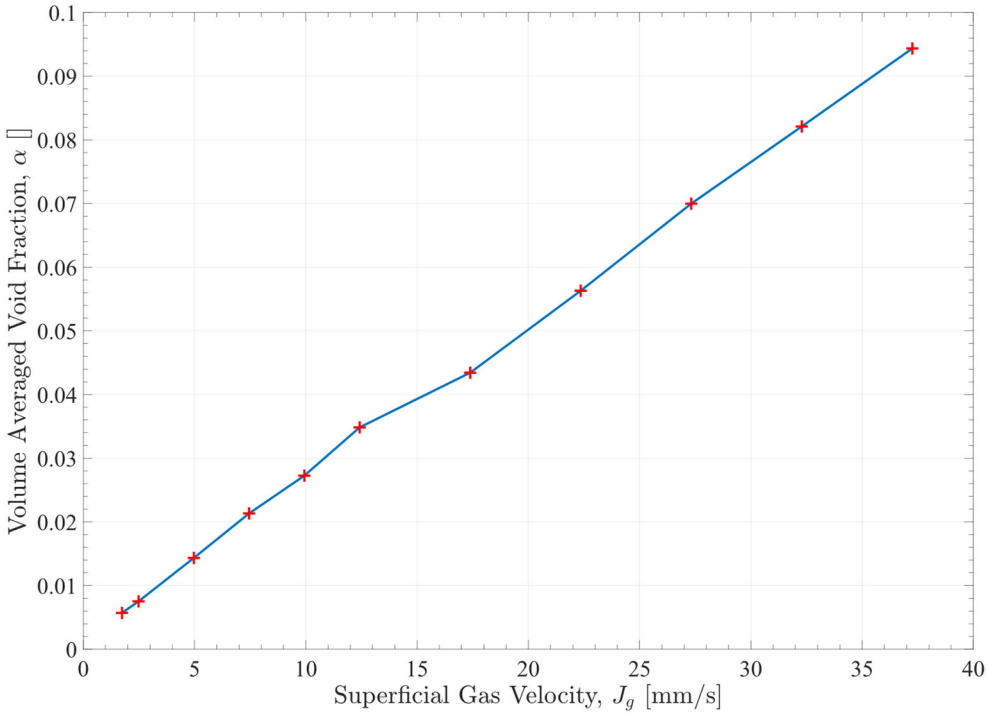

Figure 2.11: The average void fraction at different superficial gas velocity using Equation 2-15.

The most favorable feature of flow visualization experiment is that it provides information of the flow over the whole region. In addition to the measurement of volume averaged void fraction, a 2D void fraction profile could be obtained with the visualization method. From the identified bubble diameter and estimated volume, the instant bubble distribution can be projected onto a 2D plane and average it in time. This is exactly the definition of void fraction used in Eulerian-Eulerian two-fluid model [98]. The projection is done by weighing the bubble by its relative dimension in the projection direction. This will produce a 2D line averaged void fraction profile that could be compared with 3D simulation data after line averaging, as shown in Figure 2.12. The procedure is described by the equations below, the bubble diameter is multiplied with a distribution function η considering the shapes of individual bubbles,

$$\alpha_{2D}(x,y) = \overline{\left(\phi(x,y,t)\frac{\eta(x,y,t)d_{eq}}{L_{proJ}}\right)},$$

$$\eta(x,y,t) = \frac{\sqrt{r_m^2 - r^2}}{r_m}, \qquad r_m = max(r)\,. \tag{2-16}$$

ϕ is the phase indicator function, which is one at where bubbles are present and zero otherwise. r is the distance to the corresponding bubble centroid. η is geometric factor that takes account of the shape of the bubble. The choice of η is important to produce the correct projected void fraction profile. The form of η is taken from the simple geometric relationship of a sphere, as shown in Figure 2.13. However, the bubbles in the experiment are not perfect spheres, and how to choose the scaling factor becomes troublesome. If the equivalent diameter is chosen to be the scaling factor, the off-center value will drop very fast, causing a significant underestimation of the void fraction. Therefore, the farthest point from the center r_m is chosen as R. The 2D void fraction profile and the volume averaged void fraction obtained with the method described in this section will serve as an important way to validate the CFD model.

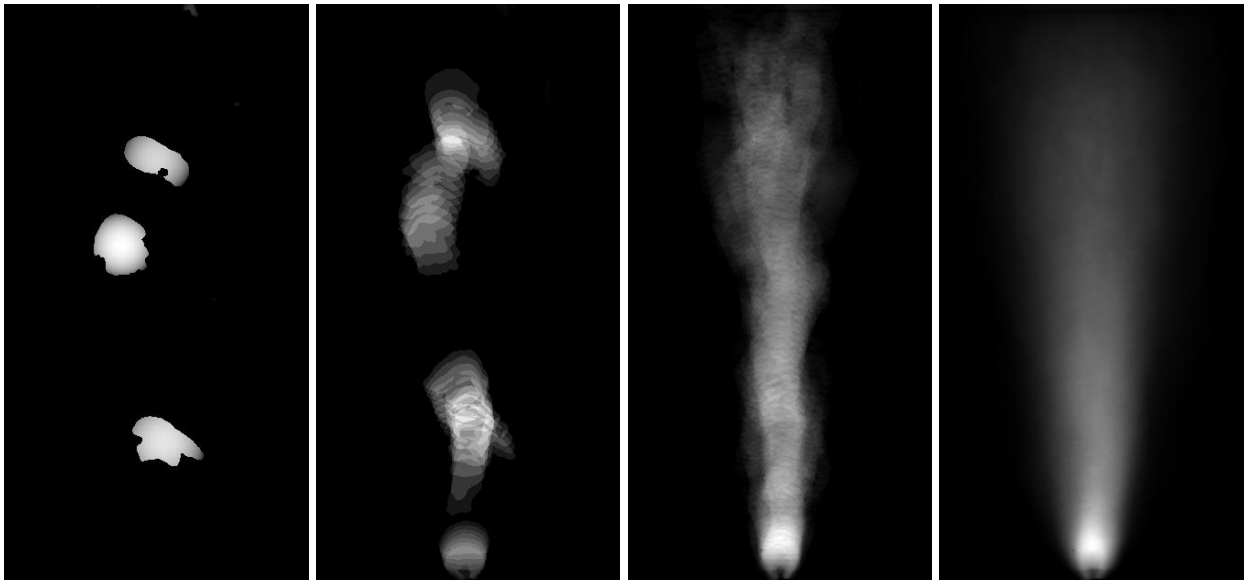

First Frame | 0.25s Averaging | 2.5s Averaging | 50s Averaging

Figure 2.12: Process of obtaining the 2D void fraction profile using geometric projection and time averaging.

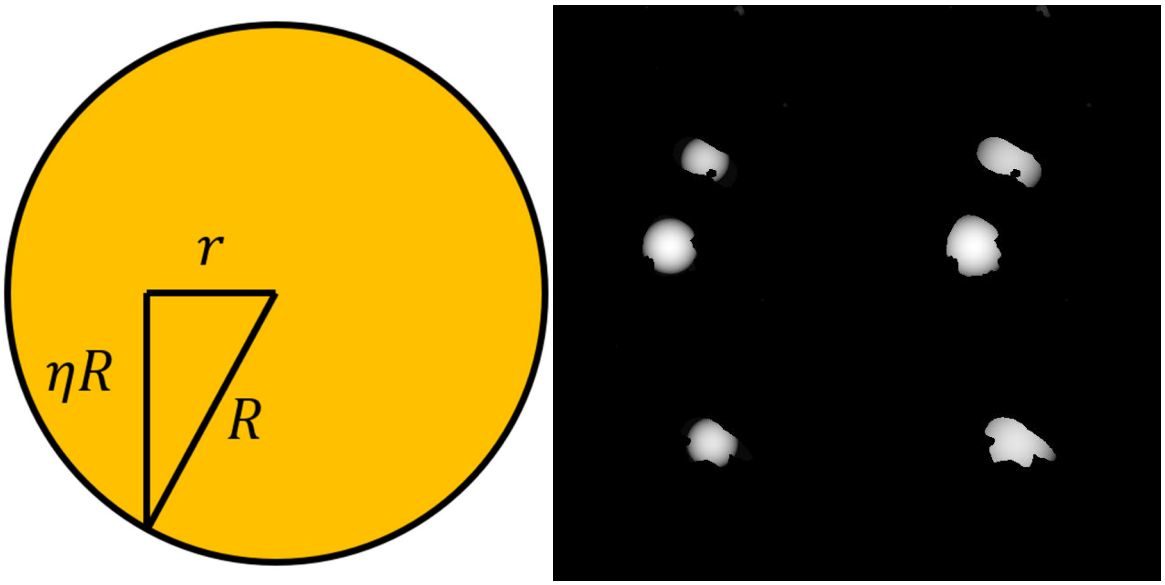

The Geometric Meaning of η Using $d_{eq}/2$ Using r_m

Figure 2.13: The geometric interpretation of the distribution function η used in the projection. The two figures to the right show the effect of using different length scale. The max diameter r_m is selected in the distribution function.

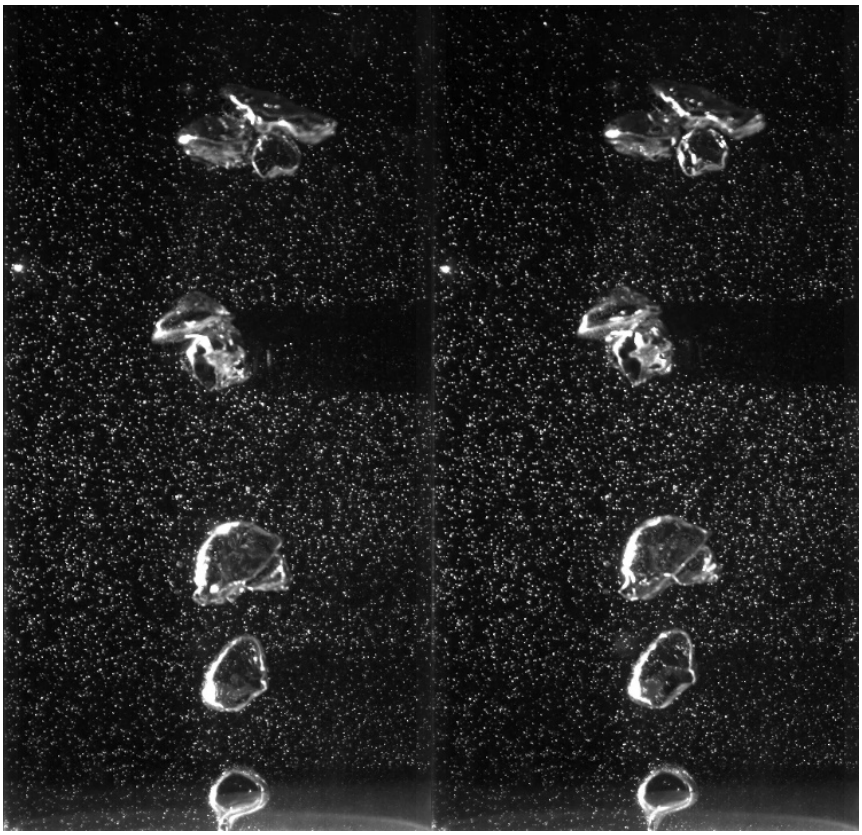

Figure 2.14: A pair of images taken 0.002s apart in the PIV experiment at $J_g = 2.5mm/s$.

2.3.2. Particle Image Velocimetry (PIV) Experiment

In this section, the widely used particle image velocimetry for velocity measurement is discussed. PIV is extensively used in single phase flow experiment to measure the flow velocity. The velocity in a wide region could be measured simultaneously with acceptable error for most engineering purposes. However, when being used in two-phase flow, additional challenges are introduced by the existence of interfaces. The problem and the corresponding solution are discussed in the following section.

2.3.2.1. Experiment Procedure

The Particle Image Velocimetry (PIV) experiment is carried out basically under the same condition as the flow visualization experiment. The geometry, flow condition and camera setup are the same. The LED illumination is now replaced with a thin laser sheet across the center of the flow region, which is perpendicular to the high-speed camera. The laser sheet is created by Dantec Dynamics' RayPower system. The laser is synchronized with the high-speed camera so that the shutter and the laser pulse is triggered at the same time, which also stabilizes the laser intensity for longtime operation and therefore increase the accuracy. Each pair of frames is taken in 0.002s, and the frame-pairs are taken at 20Hz. An example of the frame-pair (double frame) is shown in Figure 2.14. The short interval between frames ensures the accuracy of the PIV measurement. On the other hand, the low frequency of recorded frame-pairs or double-frames allows averaging over a long period, which is important to obtain a time averaged profile.

Before sparging starts, proper amount of polyamide seeding particles of $50\mu m$ are added to the water. These particles will reflect the laser and appear to be a white dot in the obtained video. In order to increase the contrast of the image and therefore the accuracy, the laser should be the only light source in the experiment, unless other light sources could be filtered out by special optical fiber. Too many or too few particles could increase the uncertainty of the measurement.

2.3.2.2. Data Processing

The video is first processed to prepare the PIV data, then analyzed with adaptive PIV algorithm using Dynamic Studio by DANTEC. The images are divided into different interrogation areas, which could be overlapping or not overlapping. Each particle in the image is best to occupy more than two pixels for best performance of the algorithm. This requires a proper choice of particle diameter. The number of particles is best above 10 in each interrogation area. Limited number of particles could increase the uncertainty of the measurement. On the other hand, if the particles are too dense in the flow, the uncertainty is also increased. In addition to the number of particles, the interrogation area should also be smaller or at least comparable to the characteristic length of the flow. The liquid velocity is obtained by analyzing the cross-correlation between the frames. The resulting vector could be viewed as the average velocity within each interrogation area.

The difficulty with two-phase PIV experiment is the existence of phase boundary, which also causes reflection in addition to the seeding particles. To minimize the error introduced by the reflected laser at bubble interface, the large bright area in the obtained video is masked, as shown in Figure 2.15. The details of this method are discussed in the work by Sathe et al. [58].

The velocity field and contour plot at $J_g = 2.5\ mm/s$ after time averaging are shown in Figure 2.16. The constant presence of bubbles near the inlet region introduces considerable error to the final vector map. For this reason, a transparent box is put in the figure. Another interesting measurement is the turbulence kinetic energy of the liquid phase, which could be estimated as

$$k^t = \frac{1}{2}\left(2u'^2 + v'^2\right). \tag{2-17}$$

Here, the fluctuation is assumed to be axially symmetric to arrive at the turbulence energy in Figure 2.17. However, since the interrogation area is larger than the smallest turbulence scale, a part of the energy spectrum might be lost, therefore underestimating k^t. On the other hand, the uncertainty of the measurement will inevitably be included in the real fluctuation, which would likely cause a small overestimation in the result.

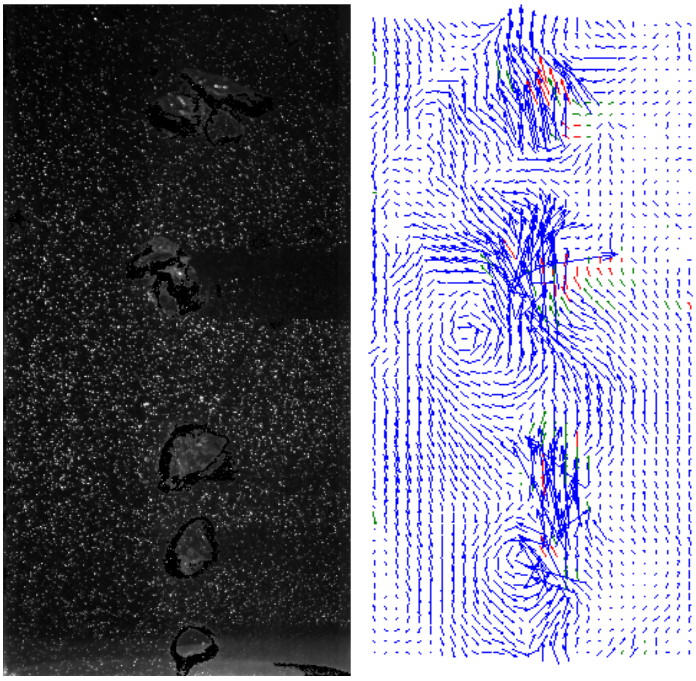

Figure 2.15: The masked image and corresponding velocity field. The large bright areas are masked to reduce the error caused by the reflection at the bubble interface.

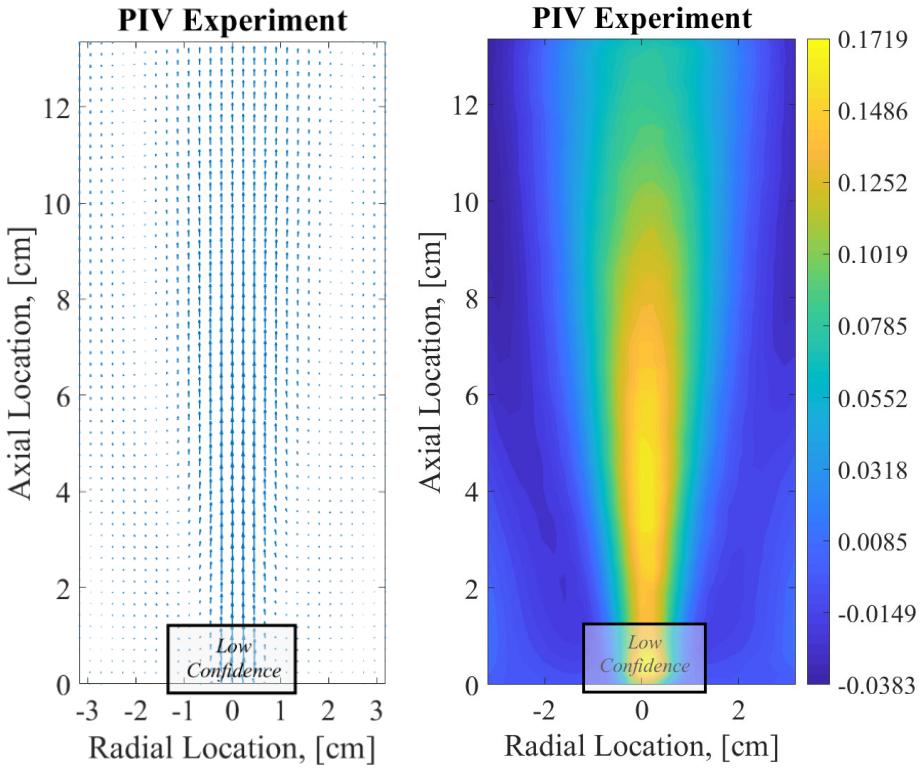

Figure 2.16: Averaged velocity field and axial velocity from PIV experiment at $J_g = 2.5mm/s$. A box is placed at the inlet region where large uncertainty exists.

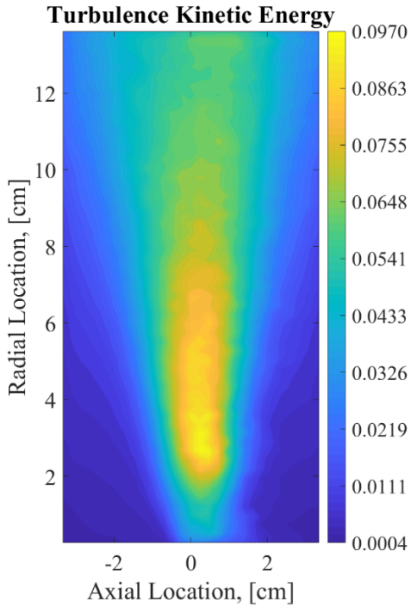

Figure 2.17: Nominal turbulent kinetic energy of the liquid phase k^t at $J_g = 2.5mm/s$. The error from the reflection could contributes to the measured k^t, and portion of k^t could be lost due to the unresolved length scale.

2.3.3. Dissolved Oxygen Charging Experiment

The mass transfer between the liquid and the gas is directly related to the removal of gaseous fission product. Various models are proposed for the mass transfer rate. In this section, the mass transfer coefficient is measured by monitoring the concentration of dissolved oxygen in the sparging process. The results from this experiment would be used to inspect and calibrate the mass transfer model in our simulation.

2.3.3.1. Experiment Procedure

The mass transfer coefficient of the bubble column is measured by monitoring the oxygen concentration in water. The dissolved oxygen (DO) concentration is measured by YSI's optical DO probe, FDO 4410 IDS. It has an accuracy of 1.5% and measurement range of 0 to 20 mg/L and a 95% response time of 45 second. The probe has a semi-permeable membrane where dissolved oxygen could enter. The presence of oxygen in this membrane with embedded luminescent dye would change the characteristic of the luminescence when blue excitation light is applied. This change is proportional to the partial pressure of oxygen in the membrane, which could be measured by the optical sensor in the probe. Since oxygen is not consumed in this process, there is no requirement for minimum flow rate as required by galvanic dissolved oxygen probe, meaning the disturbance of the flow could be minimized in the sparging process by placing it in the low velocity region.

In the experiment, the deionized water is first boiled for around a minute to remove the dissolved oxygen. Then the low dissolved oxygen water is sealed and stored in the bubble column for several hours for it to cool down to room temperature, therefore the solubility change due to temperature variation during experiment is negligible. When the temperature reaches the room temperature, the DO probe is inserted into the water from above, and air is sparged into the liquid from the bottom to charge it with oxygen, as shown in Figure 2.9. The injection is from a compressed air supply line, which is controlled by an in-line pressure regulator and followed by a panel-mounted air flowmeter. The concentration of the dissolved oxygen is measured every second. Considering the response time of the probe, certain lag would exist in the measurement result. The result taken at each second should be viewed as a time averaged value over the characteristic response time. However, this would not change the time constant of the process if the process is much longer than the response time.

Based on Henry's law and two-resistance model [32], the absorbed oxygen per unit volume per unit time could be expressed as:

$$\Gamma_{l,i} = ka_i\left(K_{l,i}^{\rho}\rho_g Y_{g,i} - \rho_l Y_{l,i}\right). \tag{2-18}$$

By assuming a homogenous oxygen concentration over the volume, the average oxygen concentration could be calculated by

$$C_{l,O^2} - C_{l,O^2}^{*} = \frac{\left(C_{l,O^2}^{0} - C_{l,O^2}^{*}\right)e^{-ka_i t}}{1-\alpha}. \tag{2-19}$$

The ka_i term in the exponential function is the volumetric mass transfer coefficient and could be calculated based on our dissolved oxygen concentration data. Variables C_{l,O^2}^{*} and C_{l,O^2}^{0} are the equilibrium and initial concentration of the dissolved oxygen. The inverse of the volumetric mass

transfer coefficient has the dimension of time, which is the characteristic time for this process. The equilibrium concentration is determined by running the sparging process at high flow rate long enough until the concentration of dissolved oxygen becomes stable. For the low volume fraction cases encountered in the validation experiments, the $(1-\alpha)$ factor could be ignored.

2.3.3.2. Results and Discussion

Figure 2.18 shows the result obtained at a flowrate of 1 SCFH, corresponding to J_g of $2.5\mathrm{mm/s}$. In this figure, the concentration of oxygen changes with time in an exponential manner, as expected for a homogeneous solution. Thus, this concentration could be properly described by the expression that is derived above. By fitting the curve with least square method, the important volumetric mass transfer coefficient ka_i could be extracted. It is possible to decouple the mass transfer coefficient k from the interfacial area a_i. Though, without accurate and direct measurement of a_i, this decoupling is not so beneficial with the huge uncertainty caused by it. Thus, it is more justified to keep the parameter directly measured from the experiment, without any empiricism introduced.

A series of experiments are carried out at different gas flowrates where three experiments at the same condition were done to check the repeatability of our method. To validate the experiment procedure, the experiment data are also compared with the existing experiment correlation given by Hikita et al. [99]. The results are shown in Figure 2.19. From the comparison, the agreement between the correlation and our result is acceptable relative to the uncertainty of the measurement. The minor difference between our measurement and previous experiment could result from the different geometries used in the experiments, which could not be reflected by superficial velocity alone. The repeatability of the experiment results is an indication of the small uncertainties of our method. A detailed uncertainty analysis is included in the Appendix C.

From Figure 2.19, it is shown that the mass transfer coefficient increases with the superficial velocity of the gas. The slope, on the other hand, is decreased with increasing gas flowrate.

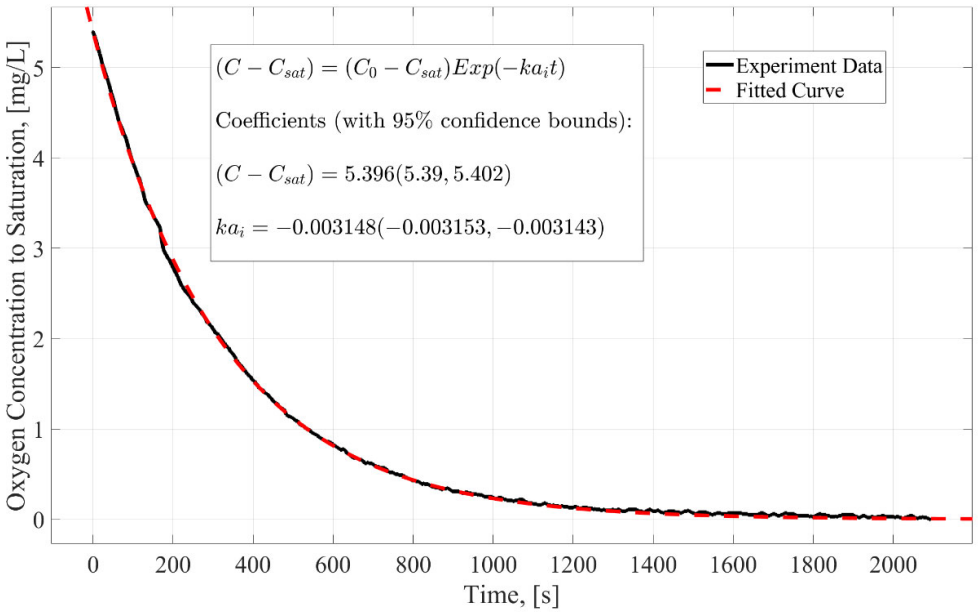

Figure 2.18: Dissolved oxygen concentration versus time during air bubble sparging. The trend agrees well to an exponential function, indicating that the volumetric mass transfer coefficient ka_i could be properly defined.

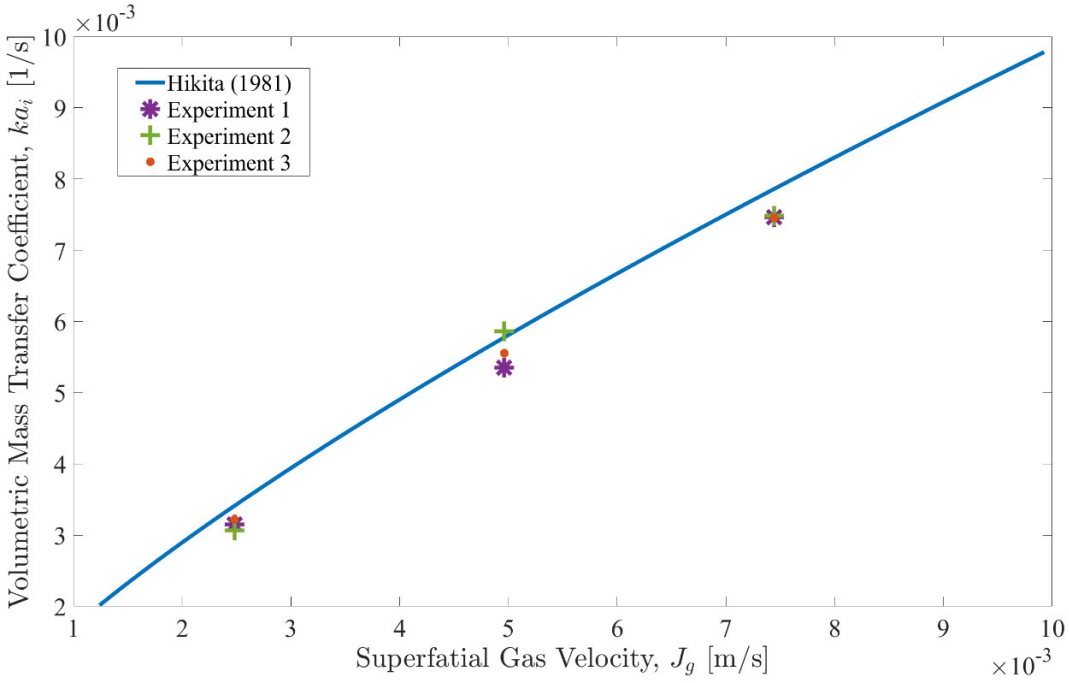

Figure 2.19: The measured ka_i at different superficial velocity J_g. The blue line is the predicted ka_i at different J_g using correlation by Hikita et al. [99]. Multiple experiments are carried out at the same condition to verify the repeatability of the experimental procedure.

2.4. Validation of Two-Phase Species Transfer Model

In order to choose the proper constitutive relation, grid size and boundary condition, a series of sensitivity studies are carried out. The simulation results for different cases are compared with each other. Some experiment results are also used here for reference. The base case based on these sensitivities study is summarized in Table 2.8. The simulation is carried out in ANSYS FLUENT 19.2 with user defined functions.

Table 2.8: Summary of models and setups for base case in the sparging simulation.

Setups	Turbulence	Bubble Induced Turbulence	Mesh	Outlet	Bubble Diameter
	RNG $k-\epsilon$ [100] Tchen's Theory [101]	Sato [66]	Medium	Free Surface	Gaddis [102]
Forces	Drag	Lift	Turbulent Dispersion	Virtual Mass	Wall Lubrication
	Ishii and Zuber [98]	Tomiyama [69]	Lopez de Bertodano [70]	Not Included	Not Included

2.4.1. Benchmark Study

In a previous report [64], a sensitivity study is performed and compared with experiment data. The model is then validated against experiment data and some simple cases with analytical solutions.

2.4.1.1. Validation of Void Fraction Prediction

From the visualization experiment, the volume averaged void fraction could be measured. These volume averaged quantities play an important role in determining the proper CFD model. The sensitivity studies are performed at $2.5\mathrm{mm/s}$ superficial gas velocity, and in this section, the model performance at different gas flowrates is examined. The result is considered acceptable, with a maximum error of 3%, as shown in Figure 2.20. The simulation agrees quite well with the measured void fraction. The result from the visualization experiment is compared with the plot given by Krepper et al. [103] to validate the experimental procedure. The difference between our measurements and reported data could be attributed to the difference in the geometries of the bubble columns, which couldn't be reflected by superficial velocity alone.

In addition to the volume averaged measurement, the simulation is also compared with the localized void fraction data. The radial void fraction profile at three axial locations from experiment and simulation are compared in Figure 2.21, Figure 2.22 and Figure 2.23.

From the comparison, general agreement between the simulation and experiment for the radial profile is found. It appears that the void fraction near the inlet is underpredicted, yet overpredicted at the top. The underprediction near the inlet may be related to the absence of virtual mass force, without which the velocity of the gas phase could be overestimated. From axial location of 6.8 cm to 9 cm, the predicted void fraction profiles are essentially unchanged, yet the

measured void fraction profile becomes more dispersed. This may indicate the deficiency of the models for lateral forces. Nevertheless, the agreement of averaged void fraction and radial profile between experiment and simulation is considered acceptable, with a maximum local error less than 25%. This validation study puts confidence in the simulation result at low gas flowrate, which is the flowrate to be used in the molten salt experiment.

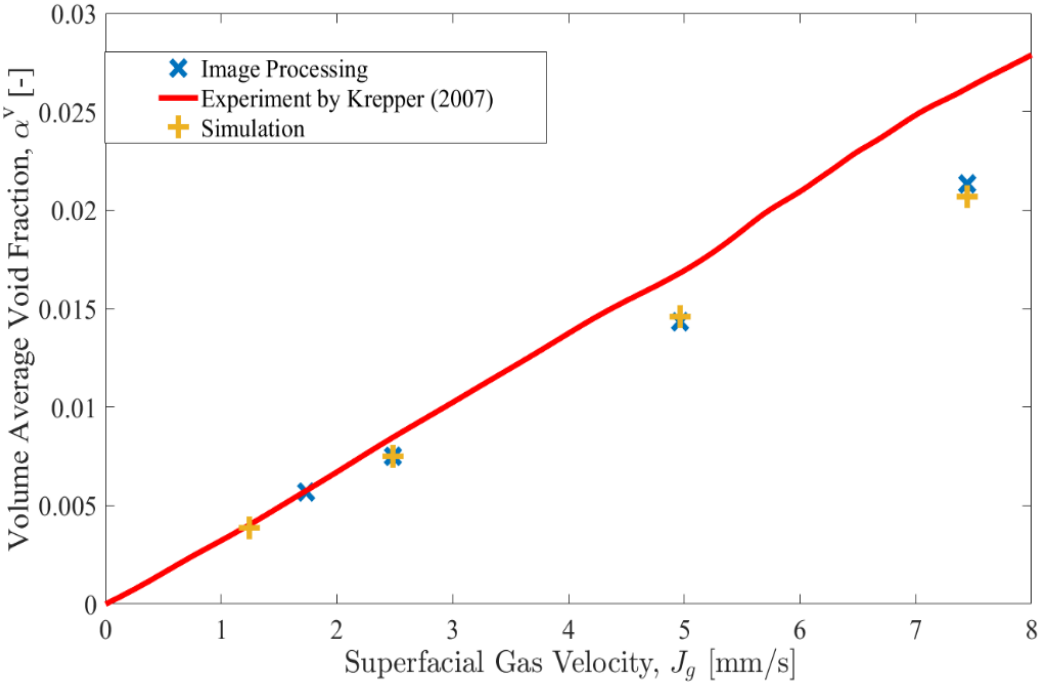

Figure 2.20: Comparison of predicted averaged void fraction from the simulation with experiment. The experimental result from Krepper et al. [103] is included as a validation of the visualization method used in this study.

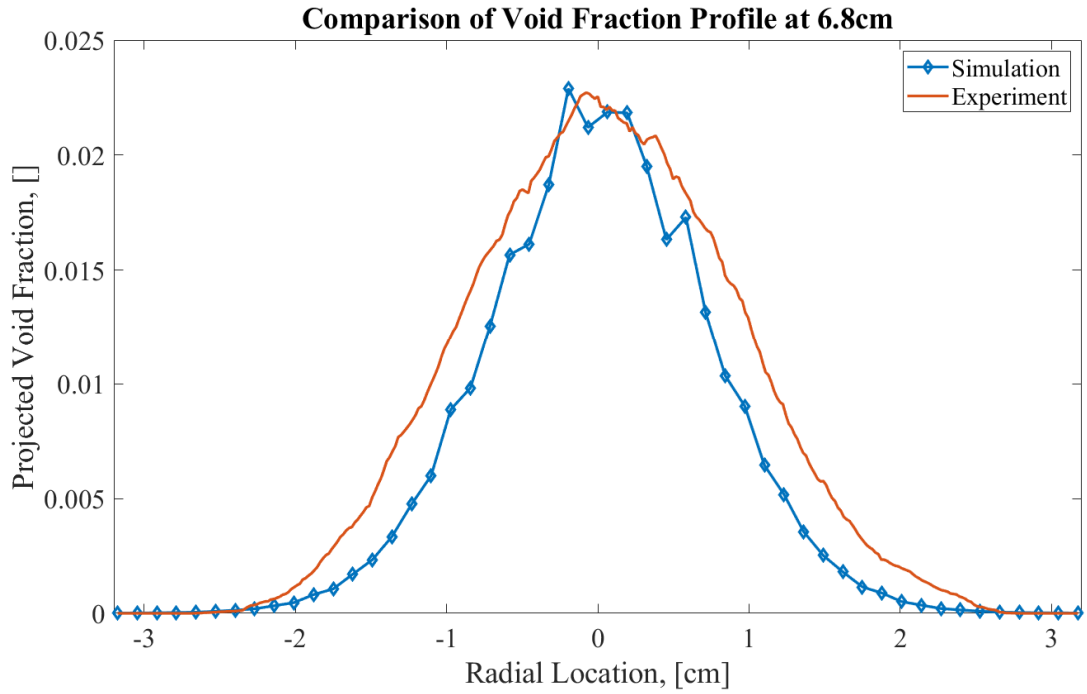

Figure 2.21: Comparison of void fraction profile from the experiment and simulation 6.8 cm above the inlet at 2.5 mm/s J_g.

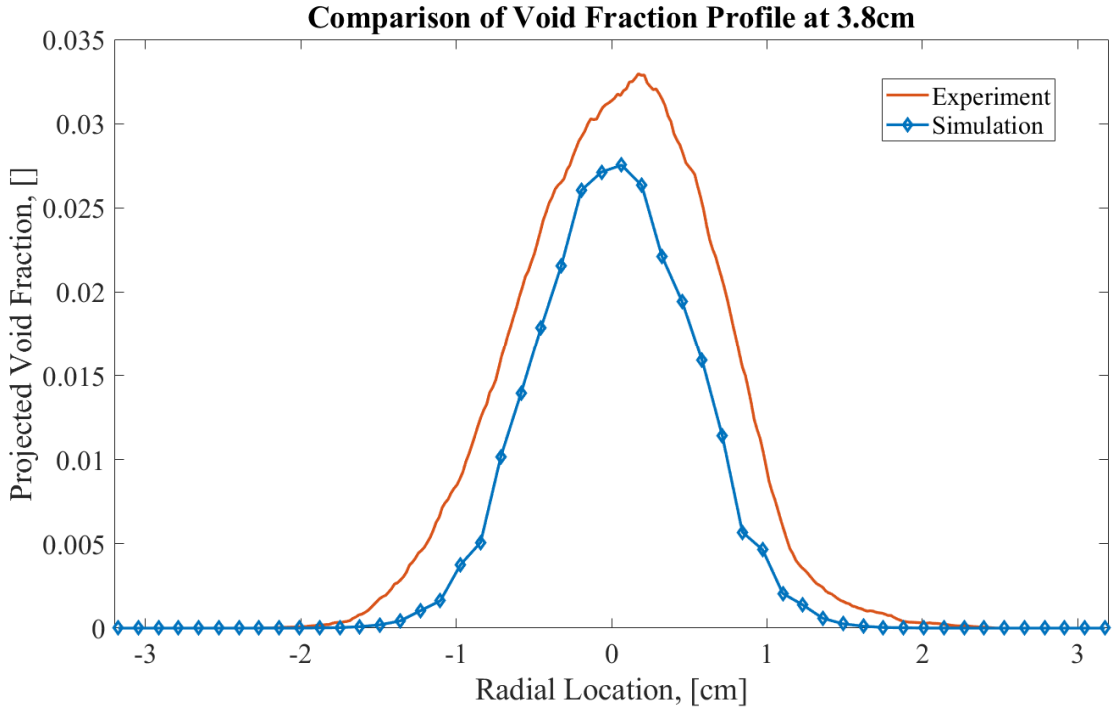

Figure 2.22: Comparison of void fraction profile from the experiment and simulation 3.8 cm above the inlet at 2.5 mm/s J_g.

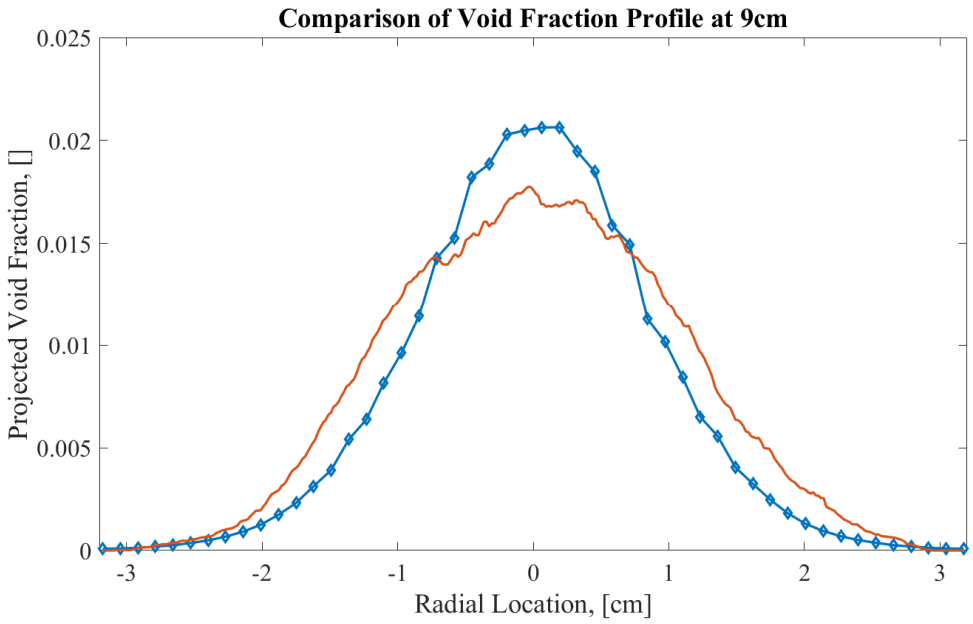

Figure 2.23: Comparison of void fraction profile from the experiment and simulation 9 cm above the inlet at 2.5 mm/s J_g.

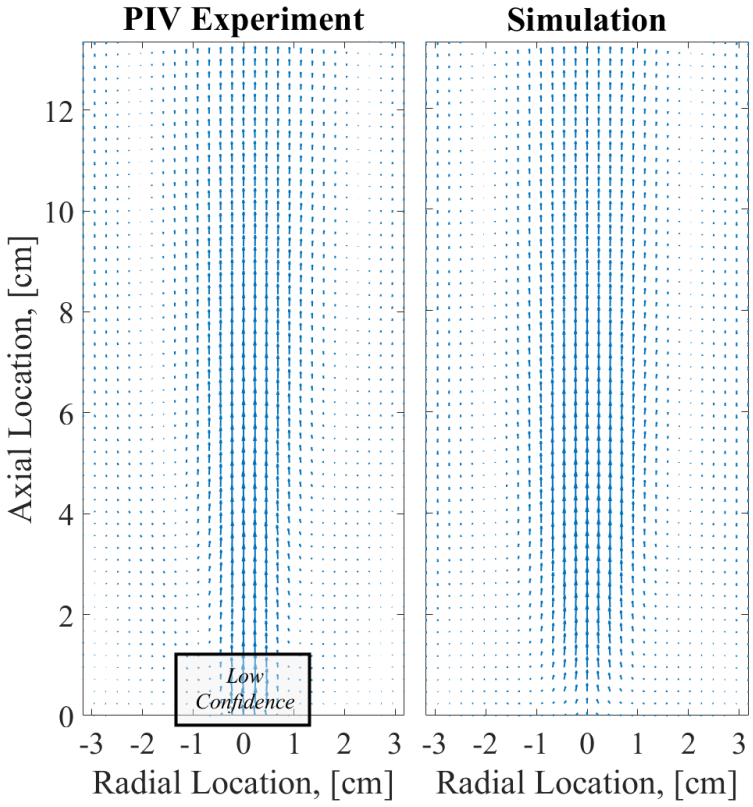

Figure 2.24: Comparison of liquid velocity vector plot at 2.5 mm/s J_g. The box is placed at the inlet region where uncertainty is large.

2.4.1.2. Validation of Liquid Velocity Prediction

The void fraction prediction of our CFD model has been validated against visualization experiments. The other important flow variable to be considered is the velocity field. The reliable measurement technique for gas velocity on a plane has not been well established, but the liquid velocity field could be measured quite accurately with the Particle Image Velocimetry (PIV) method introduced in Section 2.3.1.2. Since the diffusivity in gas phase is much larger than that in the liquid, the concentration of fission product in the gas phase should be more uniform, which makes the gas velocity profile less important in the fission product removal process. However, the concentration profile of the fission product is highly related to the circulation in the liquid phase. The velocity field at 2.5mm/s J_g is shown in Figure 2.24. General agreement between the experiment and simulation is found, though the movement near the inlet is not well captured. To quantitively compare the velocity, contour plots of the axial velocity are shown in Figure 2.25.

From the contour plot, the velocity field near the inlet is not well captured. However, further from the inlet, the agreement between the experiment and simulation improves. This could be seen from Figure 2.26, Figure 2.27 and Figure 2.28. Above axial location of 6.8 cm from the inlet, the simulation and measurement closely match. The measured velocity profile appears to be asymmetric, which is caused by the reflection at the bubble interface. The laser illuminates from the left, and a bright area will be produced in the PIV image when the laser is directly onto the bubbles. Post processing of the images removes most of these highlighted areas, but still causes a small overprediction of the velocity field on the right. In conclusion, our CFD model could satisfactorily predict the velocity field when strong localized effect is not present. Even when these effects are present, general agreement is still achieved.

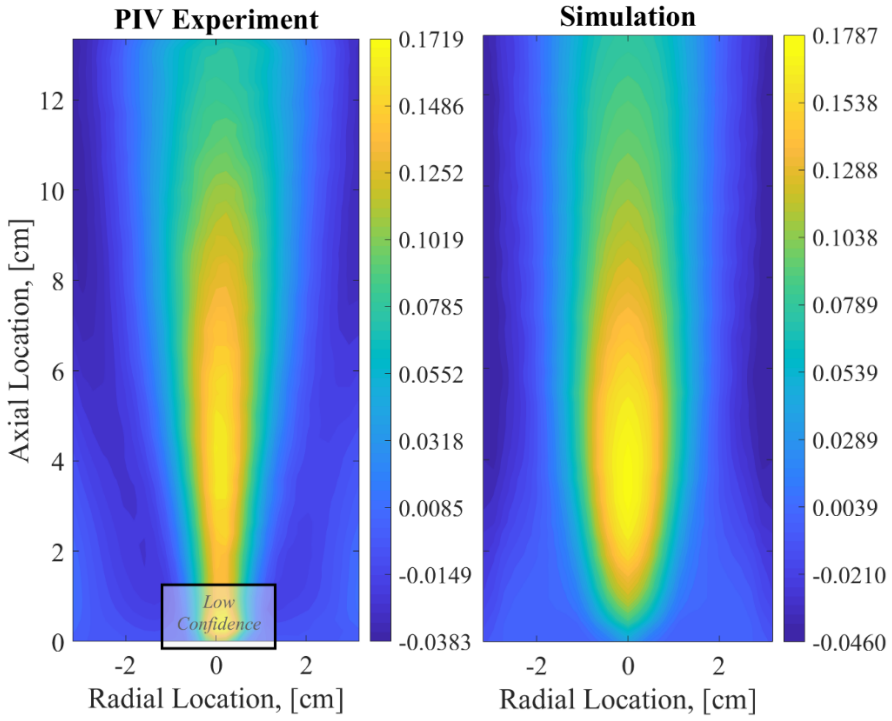

Figure 2.25: Comparison of liquid axial velocity contour, U_z [m/s], at 2.5 mm/s J_g.

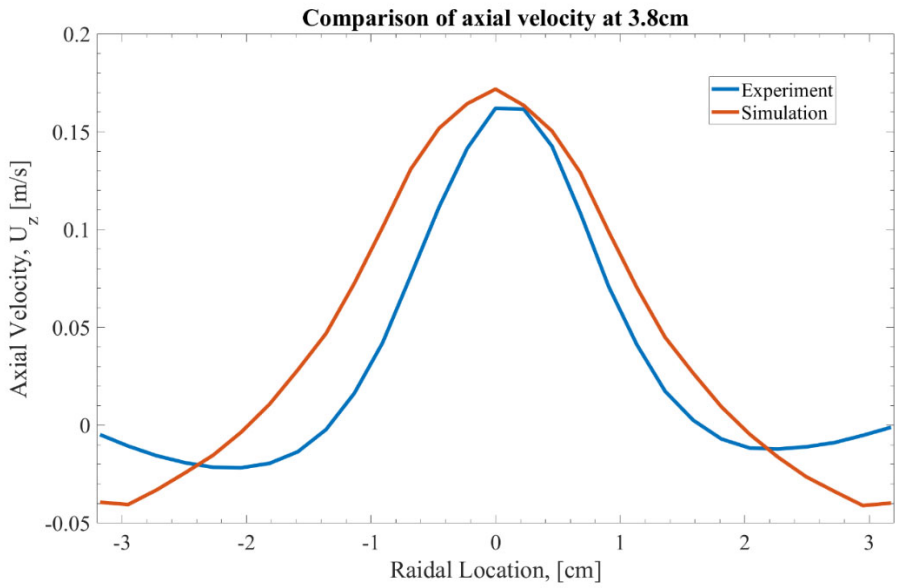

Figure 2.26: Comparison of liquid axial velocity profile at 3.8 cm above the inlet at 2.5 mm/s J_g.

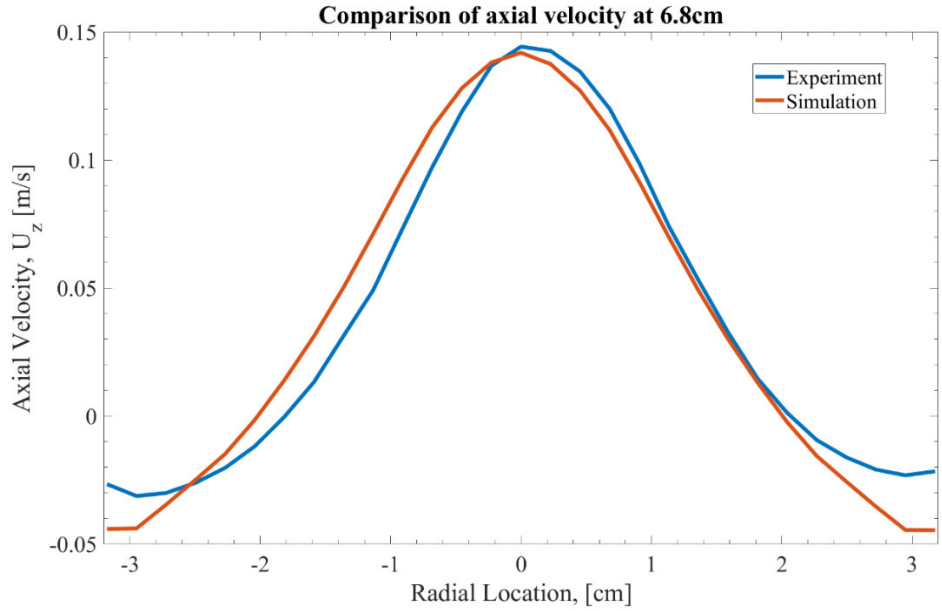

Figure 2.27: Comparison of liquid axial velocity profile at 6.8 cm above the inlet at 2.5 mm/s J_g.

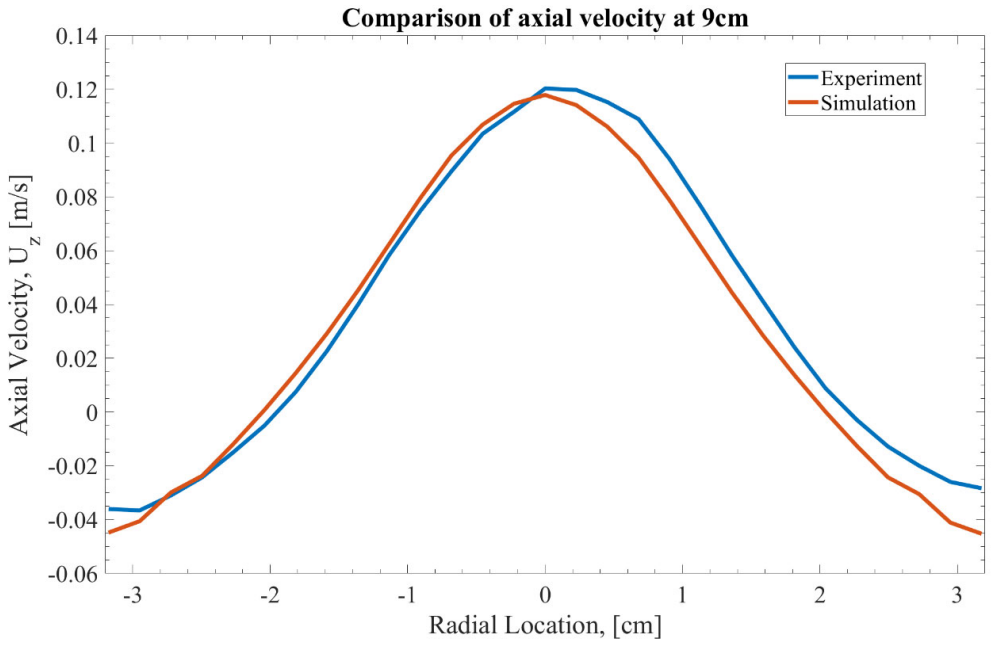

Figure 2.28: Comparison of liquid axial velocity profile at 9 cm above the inlet at 2.5mm/s J_g.

2.4.1.3. Verification of Pure Diffusion

Before the mass transfer validation, the diffusion within the phase is tested first, which is necessary to capture the species concentration distribution within the liquid. Considering the simple 1D diffusion problem, the concentration on one side of the cylinder is prescribed and the flux on the other is zero, and the initial concentration in the liquid is zero. The solution could be solved analytically if axial symmetry is assumed.

$$\frac{\partial C}{\partial t} = \frac{D\partial^2 C}{\partial z^2}, \qquad C(z=0) = C^0, \qquad \frac{\partial C}{\partial z}(z=H) = 0,$$

$$C - C_0 = -C^0 \sum_{n=0} \sin[(n+0.5)\left(\frac{\pi z}{H}\right)] \frac{e^{\left[\frac{-(n+0.5)^2\pi^2 Dt}{H^2}\right]}}{(2n+1)\pi}. \tag{2-20}$$

The simulation result is compared with the analytical solution in Figure 2.29. The concentration of xenon on one side is set to be the equilibrium concentration at 900 psi, which is the maximum xenon pressure in the molten salt experiment. From the result, the simulation and analytical result agree exactly with each other, confirming that species diffusion in the simulation is well captured.

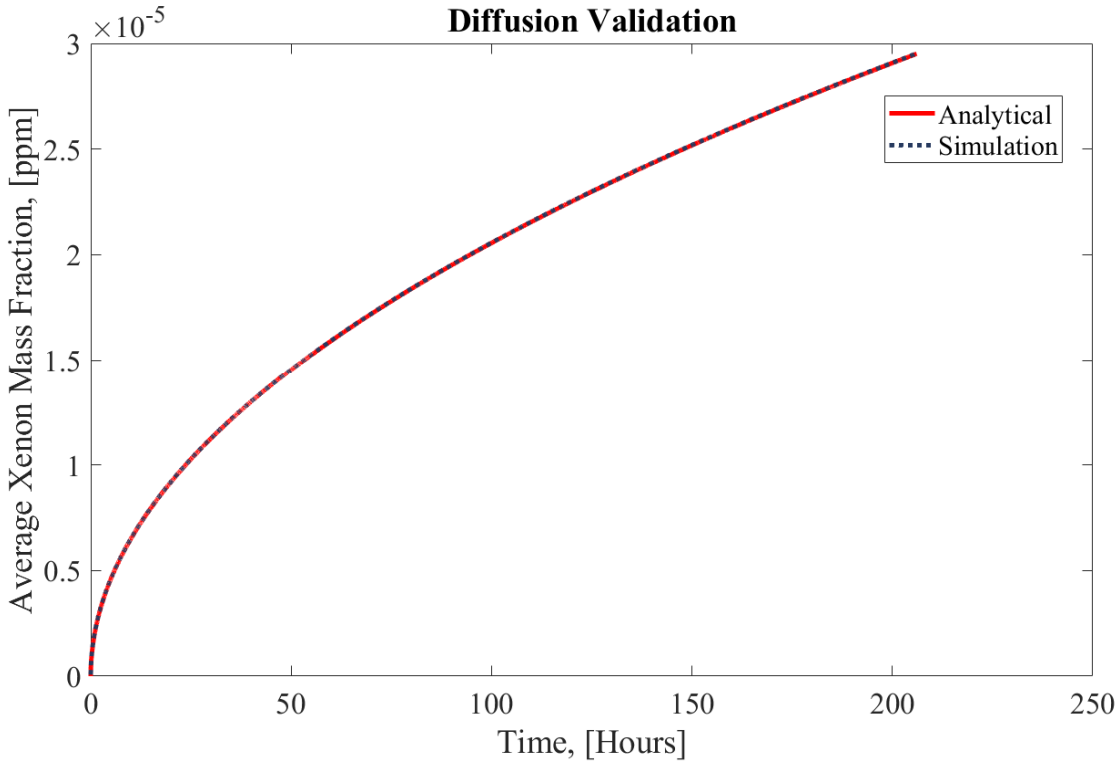

Figure 2.29: Comparison of simulation and analytical result for averaged xenon concentration due to pure molecular diffusion as a verification of the code.

2.4.1.4. Verification of Species Transport at Free Surface

Due to the chemical equilibrium at the interface, the species concentration has a sharp change at the thin liquid-gas interface. This sharp change becomes a discontinuity when the interface is modeled without thickness. This discontinuity, in the case of dispersed interface, could be smoothed, resulting in an extra source term in the species transport equation. The rate of species transfer could be related to other flow variables and the concentration gradient between phases. However, at the free surface, this is no longer available, as discussed in a previous report [104]. Therefore, special treatment of mass transfer at the interface is needed. The Continuous Species Transfer (CST) model is implemented and tested for a simple 1D problem [105, 106]. At the beginning, the liquid is charged with certain gas and in contact with another gas at the free surface. Assuming axial symmetry, the problem becomes 1D. The equilibrium ratio at the interface is arbitrarily set as 3, and the diffusivity is different for gas and liquid. The concentration distribution could be easily derived by solving the 1D equation as

$$3C_{l,i} = C_{g,i}, \qquad D_l \nabla C_{l,i} = D_g \nabla C_{g,i}, \qquad \frac{\partial C}{\partial t} = D\nabla^2 C. \tag{2-21}$$

The comparison of the CST model and the dispersed flow model is shown in Figure 2.30. From the result, it is shown that since the dispersed model does not capture the jump of species concentration and thus the diffusion-driven species transfer process, yet the CST model gives a much better prediction. When the liquid is in motion and convection also contributes to the mass transfer, the dispersed model may result in less accuracy. However, if one wants to study the

capability of cover gas to extract xenon from the salt, the CST model or other method needs to be applied in order to correctly capture the physics. However, for most of the conditions where a dispersed phase is presented, the mass transfer at the free surface is much slower. For simplicity of the simulation, this additional model is not used unless only free surface mass transfer is considered.

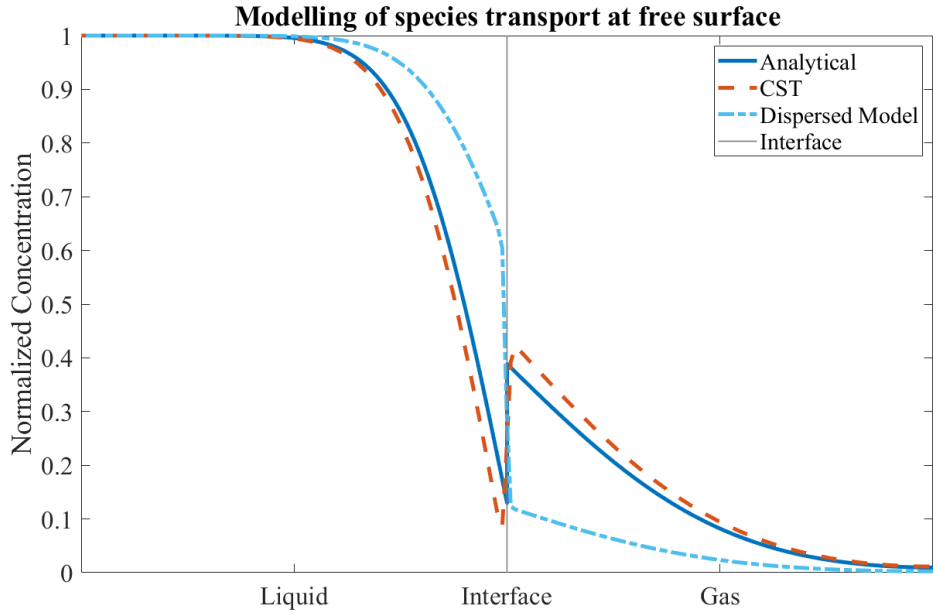

Figure 2.30: Comparison of different species transport models at free surface to theoretical solution. The Continuous Species Transfer (CST) model is required if resolving the concentration profile near the interface is required.

2.4.1.5. Validation of Mass Transfer Prediction

Now that the velocity and void fraction prediction have been validated, the last step is to validate the mass transfer model. In Section 2.2, the modeling of species transport is discussed in detail. The constitutive model for mass transfer coefficient needs to be related to local flow variables. The Ranz-Marshall model is based on analogy between mass transfer and heat transfer, utilizing the widely used Nusselt number correlation developed by Ranz et al. [107]. The surface renewal model and surface renewal stretch model are based on the similar assumption that the mass transfer near the bubble surface was controlled by renewal rate of the liquid at the bubble surface [71, 73], which is then related to liquid turbulent dissipation rate. The simulation result obtained with different models is listed in Table 2.9. The prediction of the Surface Renewal Stretch model is quite close to the experiment result. However, the prediction given by Ranz-Marshall model [107] is far from the experiment, which is likely to be related to the difference of phase interface, since liquid drops were used in the original work by Ranz et al. [107] and we shouldn't expect identical behavior to a gas bubble. The mobility of the interface and the turbulence near it might have a big impact on the result, considering the good performance of the Surface Renewal Stretch model. Moreover, the analogy between heat and mass transfer itself may not hold, with the large difference between thermal diffusivity and mass diffusivity. The two surface renewal models, on the other hand, yield much better result. The performance of the surface renewal model is

greatly improved, and the surface renewal model almost yield the exact experimental result without any calibration, indicating the physics related to this phenomenon is captured.

Table 2.9: Simulation and experimental volumetric mass transfer at $U_g = 2.5mm/s$.

	Ranz-Marshall [107]	Surface Renewal [71]	Surface Renewal Stretch [73]	Experiment
$ka_i(1/s)$	2.07×10^{-4}	1.2×10^{-3}	3.2×10^{-3}	3.1×10^{-3}

The concentration profiles of the dissolved oxygen during the simulation for the surface renewal stretch model at $U_g = 2.5mm/s$ are shown in Figure 2.31. At the beginning of the simulation, the mass fraction of oxygen in the liquid phase is set as 4.2×10^{-6}. The colormap is fixed to show the comparison between different profiles. The concentration of the dissolved oxygen becomes basically uniform after 30s. Though local variation of the concentration still exists, it is negligible compared with the concentration change. This observation justifies the volume averaged method described in Equation 2-19, where the liquid is treated as a homogeneous solution. The direct comparison of the dissolved oxygen concentration is shown in Figure 2.32. Good agreement is found between the simulation and experiment in the 114.7s simulation time. From the result, the CFD model could predict the mass transfer process within the uncertainty of the experiment measurement.

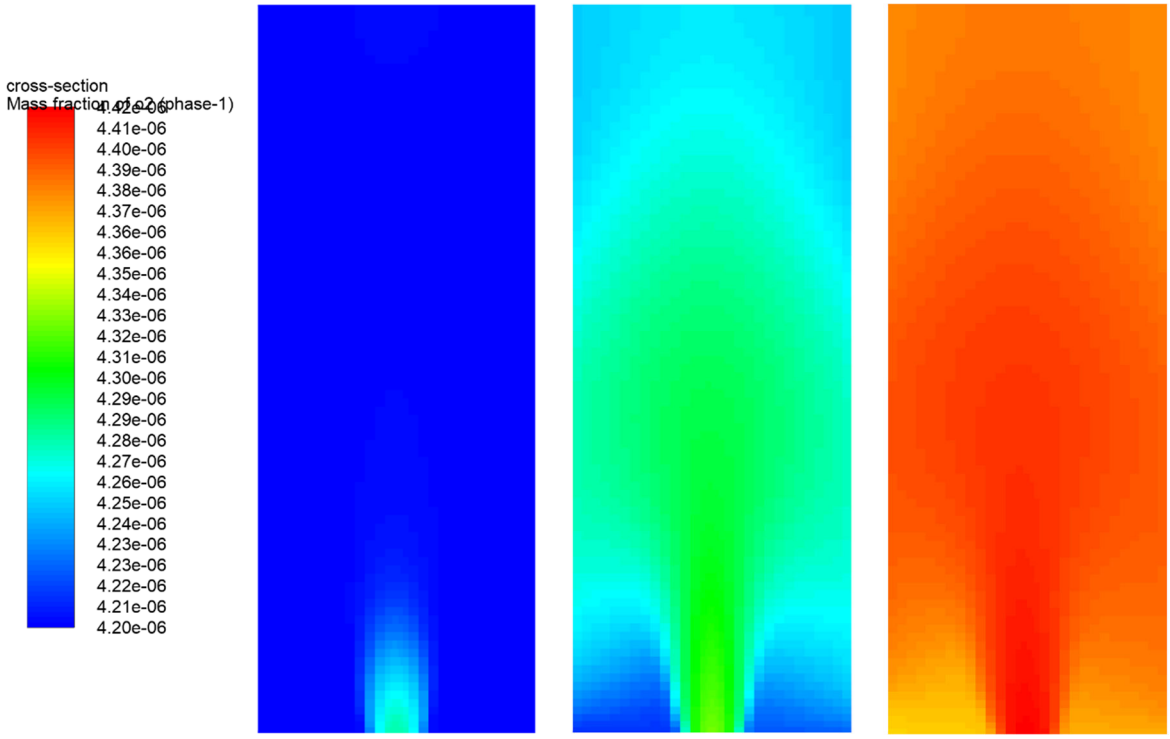

Figure 2.31: Concentration profile of oxygen during the simulation at the axial cross section for 2.5 mm/s J_g.

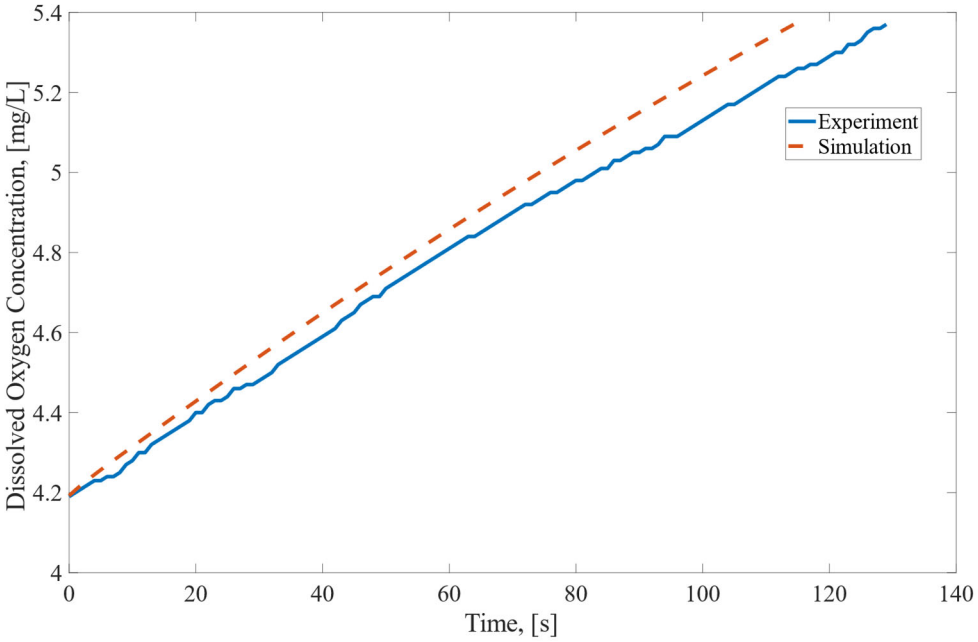

Figure 2.32: Comparison of the dissolved oxygen concentration versus time for the simulation and experiment at 2.5 mm/s J_g.

2.5. Simulation and Analysis of Xenon Removal in Pipe Flow

After the bubbles are created at the bubble generator, they are carried along by the salt flow. Dissolved xenon would escape from the salt and enter the gas phase, until equilibrium is reached. This process is referred to as xenon removal in this chapter. Xenon removal takes place at the gas-liquid interface, where the rate of the removal is mainly controlled by the amount of available interface in the two-phase flow. In the MSRE, the xenon removal is not confined in a certain part of the primary loop. Circulating helium bubbles are thought to exist in the experiments. Before the helium bubbles are eventually removed from the salt flow, a relatively long period of removal is achieved to ensure the equilibrium is reached. However, circulating bubbles are not neutronically favorable in the reactor and could also put the fuel pump at risk. Therefore, in the prototypical xenon removal system, the helium bubbles are designed to be injected and removed in the same bypass of the fuel salt. This design raises the concern of whether the removal time is long enough to remove most of the dissolved xenon in one pass. To address this concern, the mass transfer process must be studied. Here, horizontal pipe flow is considered in the prototypical design due to its simplicity.

Preliminary experiments and analysis were performed to study this two-phase mass transfer mechanism after the MSRE experiment [28, 48], the focus of which were to obtain an experimental correlation for mass transfer in bubbly pipe flow that could be further used in the design of a xenon removal system. Peebles [28] carried out a theoretical study for the mass transfer between liquid and spherical bubbles, along with a 1D analysis of the xenon removal process in pipe flow. Peebles stated that the knowledge at that time was not enough to arrive at a firm analysis of the mass transfer process, and further experiments were suggested. Subsequently, a more thorough study was carried out by Kress [48] in terms of the mass transfer in pipe flow. Dissolved oxygen

experiments are carried out to measure the mass transfer coefficient in such flow, arriving at an engineering model to relate the mass transfer coefficient to Reynolds and Schmidt number. Comparisons are made with existing experiments and correlations, yet no analysis is made on the xenon removal and the design of the removal system with the more accurate model.

In this section, the xenon removal process is revisited with CFD simulation, and a more sophisticated 1D analysis based on the simulations is carried out to evaluate the design requirements.

2.5.1. CFD Simulation of Xenon Removal and Regression Model

The CFD simulations are carried out with STAR CCM+ 2019. The models used in the simulations are discussed in Section 2.2. The considered geometry is a straight horizontal pipe. More complex shapes could be used to enhance mixing and xenon removal, however they would increase the pressure head required for the system and thus put more burden onto the fuel pump. In the considered geometry, as shown in Figure 2.33, uniform velocity and void fraction profiles are prescribed at the inlet (left), and the outlet (right) is set as a pressure outlet. The mass concentration of xenon at the inlet is set to 1ppm. The simulations are carried out under different flowrates, pipe diameters, salt properties, and bubble diameters. The length of the pipe is set to be long enough to have a developed flow.

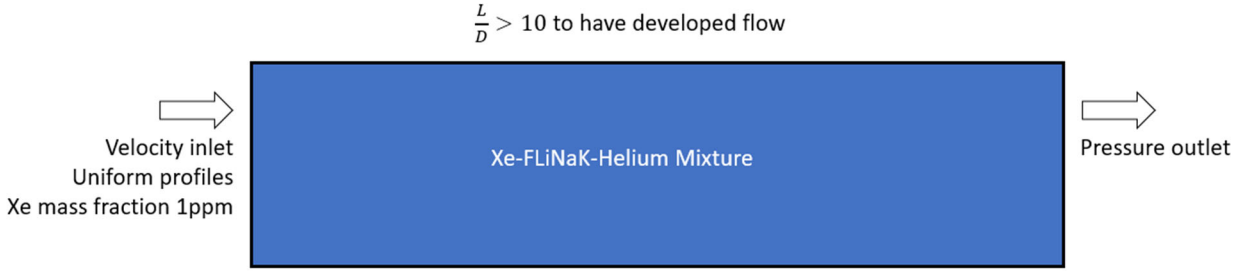

Figure 2.33. Setup of the xenon migration simulations in a horizontal pipe.

In Figure 2.34, the simulation profiles for the case with conditions listed in

Table 2.10 are shown. In the void fraction profile, it is observed that the void fraction is concentrated at the top of the pipe, and the radial profile becomes steady after a short distance from the inlet. Bubbles at the top quickly coalesce with each other, however, as long as the void fraction is kept low, the constant bubble diameter assumption is still valid. In the liquid velocity profile, it is shown that the profile becomes steady after a similar distance from the inlet as with the void fraction profile, therefore the uniform inlet condition would not affect the results in the developed region. Alternatively, the xenon mass fraction profile in the salt is constantly changing along the pipe as xenon escapes from the salt into the gas phase.

Table 2.10. Conditions for the simulation case shown in Figure 3.2.

Parameter	Value
Inlet velocity, $[m/s]$	3.00
Inlet gas void fraction, []	0.05
Pipe diameter, $[cm]$	10.16
Bubble diameter, $[mm]$	0.65
Salt viscosity, $[mN \cdot s]$	3.00

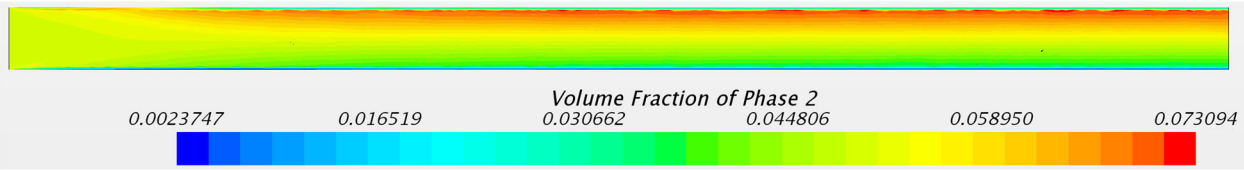

Void fraction of the gas phase.

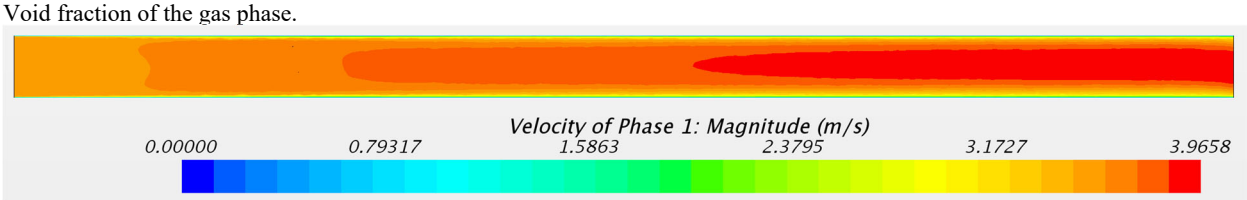

Velocity magnitude of the liquid phase.

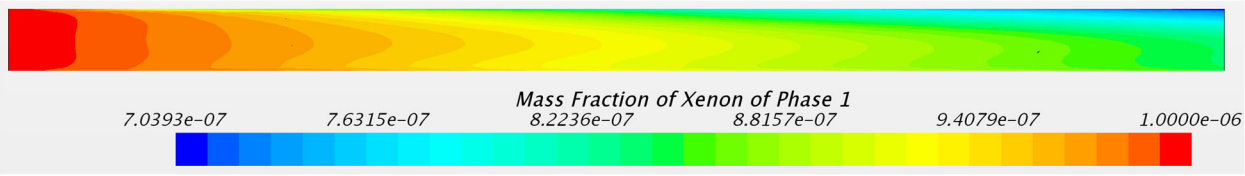

Xenon mass fraction of the liquid phase.

Figure 2.34. Simulation profiles of the flow variables in horizontal pipe flow. The length of the pipe is set long enough for the volume fraction and velocity profile become steady.

In order to compare the xenon removal rates under different conditions, statistics of the simulation data should be created. Because the velocity profile and void fraction profile get developed, the averaged xenon removal rate in the developed region would be a good representation of the process. Additionally, the xenon removal rate is proportional to the available interfacial area. Under the low void fraction condition, the interfacial area term should be separable from the removal rate, and only to be provided as an operational parameter. Therefore, the void weighted area averaged Sherwood number $\langle\langle Sh\rangle\rangle$ is calculated for each case in the developed region.

$$\langle\langle Sh\rangle\rangle = \frac{\int_A Sh\,\alpha\, dA}{\int_A \alpha\, dA}, \qquad Sh = \frac{kD}{D_m}. \tag{2-22}$$

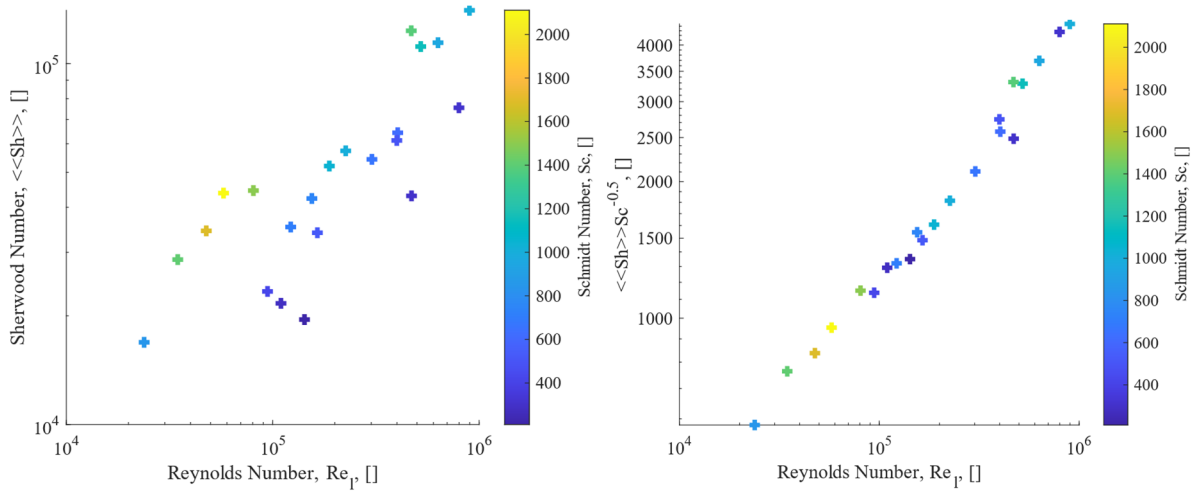

Figure 2.35. $\langle\langle Sh\rangle\rangle$ and its dependence on Sc and Re_l for simulations with Surface Renewal Stretch model. The basic scaling $\langle\langle Sh\rangle\rangle \sim Sc^{0.5} Re_l^n$ is determined from the result.

In the CFD simulation, the local xenon removal rate is calculated based on the local flow variables. However, from the system analysis perspective, the distributions of these local variables are not known. Engineering models must be created to relate $\langle\langle Sh\rangle\rangle$ to other macroscopic or averaged variables. From dimensional analysis, the $\langle\langle Sh\rangle\rangle$ could be dependent upon the pipe Reynolds number, Schmidt number, void fraction, and normalized bubble diameter. In Figure 2.35, the averaged results with respect to Reynolds number and Schmidt number are plotted. A logarithmic correlation between the Reynolds number and the $\langle\langle Sh\rangle\rangle / Sc^{0.5}$ is shown in the figure, indicating a regression model in the form of,

$$\langle\langle Sh\rangle\rangle \sim Sc^{0.5} Re_l^n. \tag{2-23}$$

The front factor in the regression model could be a constant or a weak function of other parameters from the good linearity shown in Figure 2.35. The effect of other parameters on the averaged Sherwood number is shown in Figure 2.36. Neither of these parameters causes a clear deviation from the linear dependence and thus to the front factor. Therefore, a constant front factor could be used.

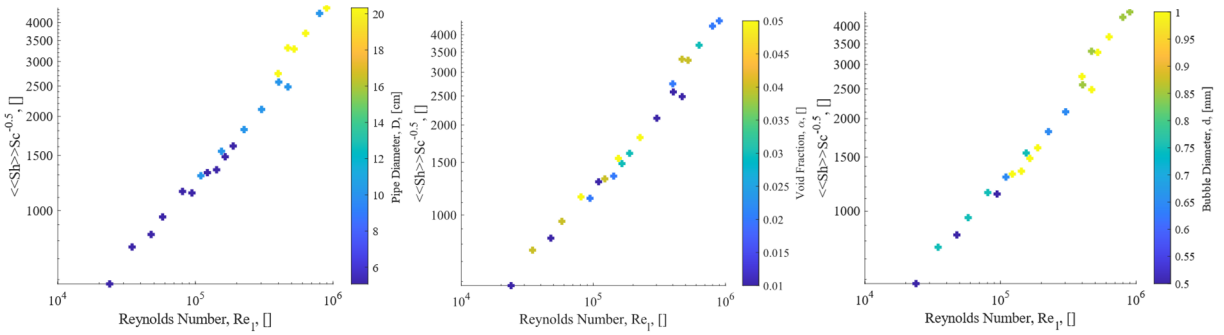

Figure 2.36. $\langle\langle Sh\rangle\rangle$ and its dependence on pipe diameter, void fraction and bubble diameter for simulations with Surface Renewal Stretch model. No clear dependence is found between Sherwood number and these parameters.

With the discussion above, a regression model is created to relate the averaged Sherwood number. The results are shown in Figure 2.37, where the dashed lines are the ±30% error line. The final expression of the regression model is,

$$\langle\langle Sh\rangle\rangle = 2.07 Re_l^{0.555} Sc^{0.5}, Sh = \frac{kD}{D_m}, Re_l = \frac{\rho j_l D}{\mu},$$

$$10^6 > Re_l > 2 \times 10^4, 2200 > Sc > 200. \tag{2-24}$$

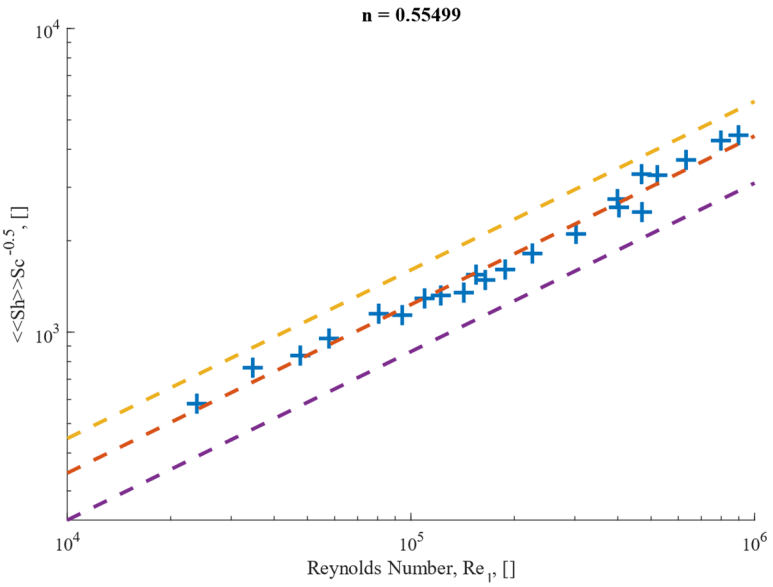

Figure 2.37. Regression model of $\langle\langle Sh\rangle\rangle$ using least square fitting for simulations with Surface Renewal Stretch model. The lines are ±30% error lines and fitted correlation.

The above regression model is constructed from the CFD simulation using localized correlation and flow variables. Since the profiles are available, the covariance between different variables describing the difference in flow distributions could be obtained when required in the 1D analysis. This is an advantage over the engineering models from the experiments where the measurement of the detailed profiles without disturbance to the flow is difficult. However, the accuracy of the constitutive relations dictates the accuracy of the simulation. The mass transfer model used in the simulations shown in Figure 2.35 through Figure 2.37 is the surface renewal stretch model developed and validated in an upward bubble column [73, 108]. It is selected in the previous milestone report because it corresponds to the on-going molten salt sparging experiment at University of Illinois. This model relates the local mass transfer coefficient to the turbulent dissipation rate and the molecular diffusivity of the dissolved gas. Since it does not distinguish between different scales and origins of the turbulence, which would be quite different under different flow conditions, considerable error could exist when applied to the horizontal pipe flow. To overcome this problem, different models or model coefficients should be used for different flow regimes.

According to the discussion above, another set of simulations is carried out with the small eddy cell model proposed by Lamont [71]. A similar regression analysis is performed, and the results are shown in Figure 2.38 and Figure 2.39.

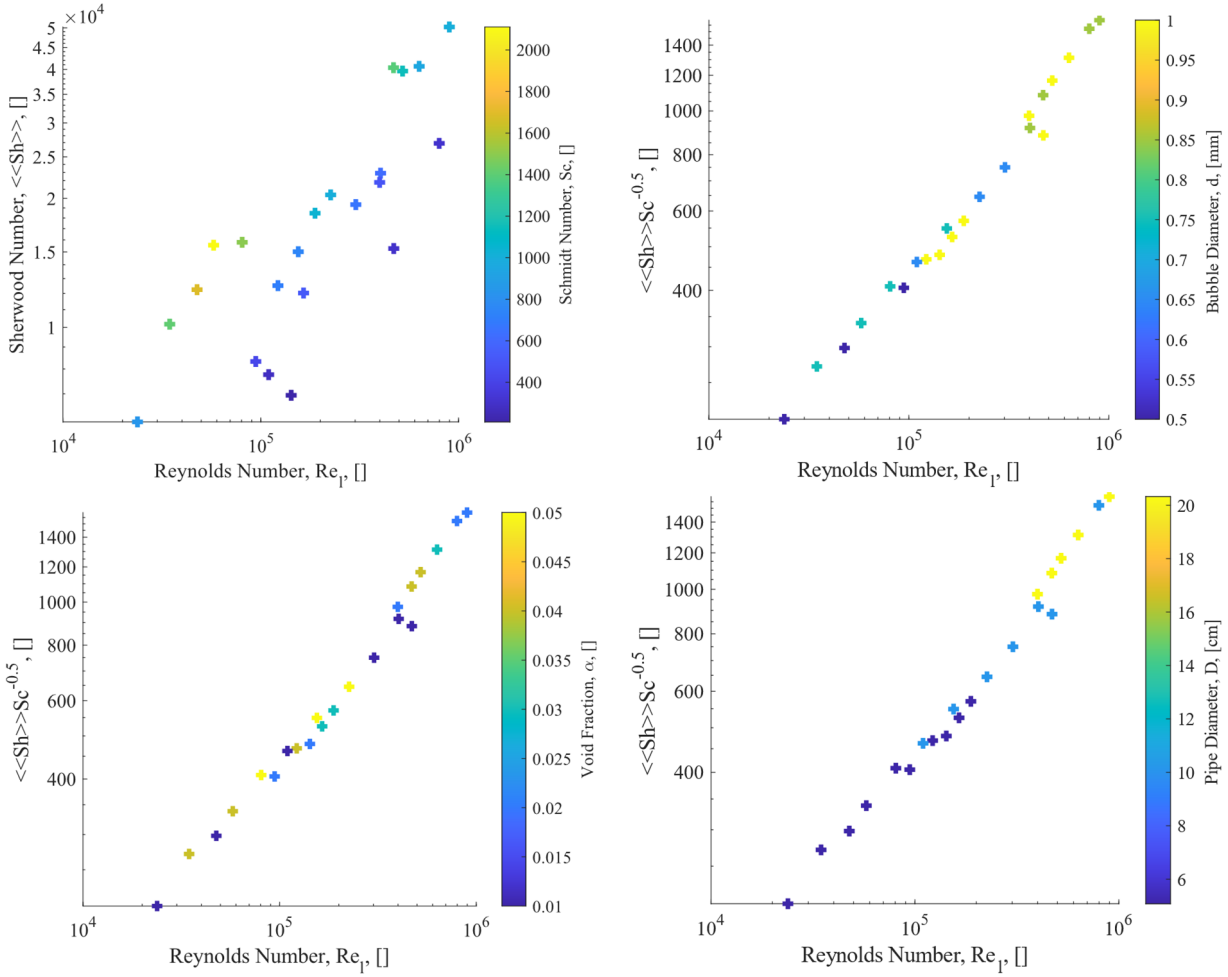

Figure 2.38. $\langle\langle Sh \rangle\rangle$ and its dependence on Sc, Re, bubble diameter, pipe diameter and void fraction for simulations with Small Eddy Cell model. The basic trends are similar to those with the Surface Renewal Stretch model.

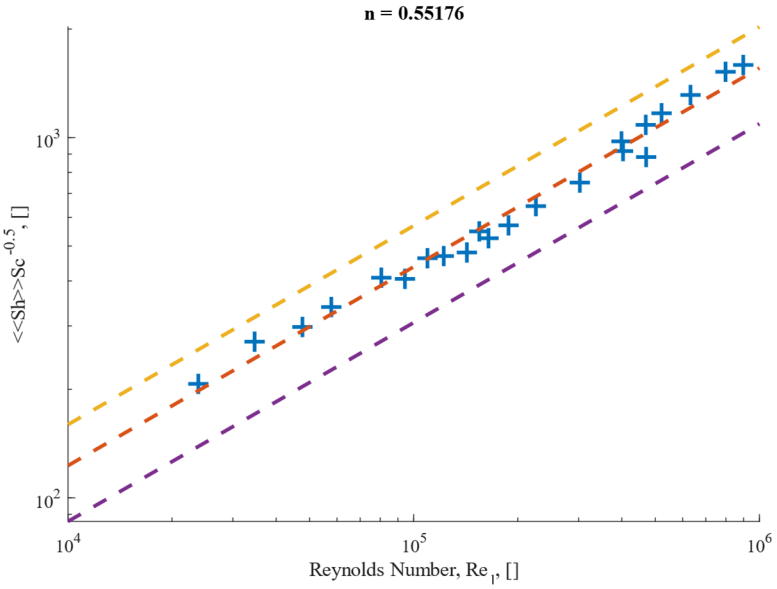

Figure 2.39. Regression model of $\langle\langle Sh\rangle\rangle$ using least square fitting for simulations with Small Eddy Cell model. The lines are ±30% error lines and fitted correlation.

The corresponding regression model from the small eddy cell model is,

$$\langle\langle Sh\rangle\rangle = 0.763 Re_l^{0.552} Sc^{0.5}, Sh = \frac{kD}{D_m}, Re_l = \frac{\rho j_l D}{\mu},$$
$$10^6 > Re_l > 2 \times 10^4, 2200 > Sc > 200. \quad (2\text{-}25)$$

When comparing the results between the different models, the difference mostly lies in the front factor, while the dependence on Re_l and Sc is essentially unchanged. This indicates that the hydrodynamics of the flow is decoupled from the species transfer between the phases. When compared with engineering model from experiments as shown in Figure 2.40 [64, 71, 109, 110], the small eddy cell model gives a reasonable prediction, while the surface renewal stretch model shows a large overprediction. All correlations are evaluated at $Sc = 800$. Some of the correlations involve more parameters, such as bubble diameter to pipe diameter ratio and void fraction, which are calculated according to the mean values of their applicable ranges. Moreover, the difference between the experimental correlations is noticeable, which is likely related to the uncertainties of the interfacial area measurement. For the following sections, the regression model from the small eddy cell model is used.

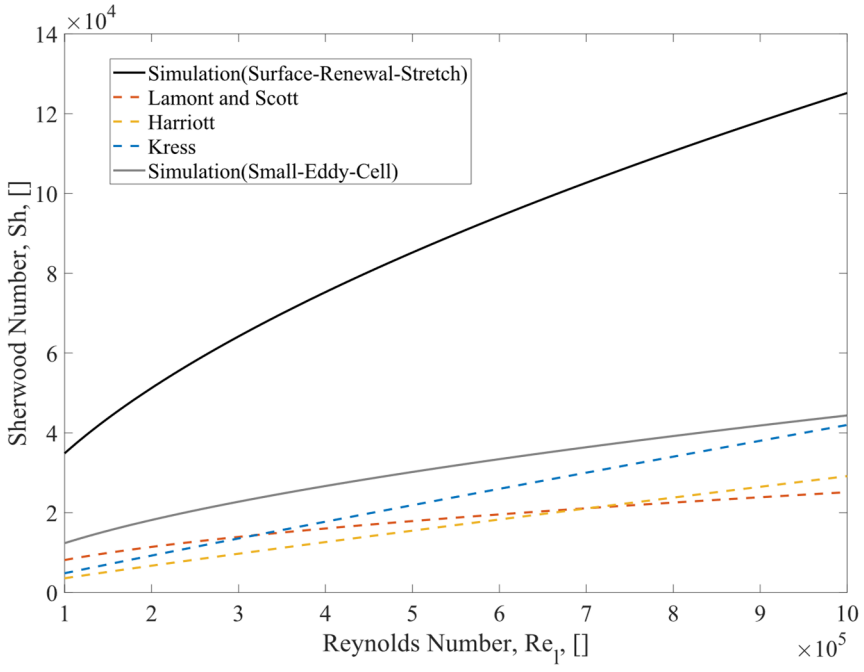

Figure 2.40. Comparison of the regression models from simulation and experiments. The Surface Renewal Stretch model seems to overpredict the species transfer in pipe flow condition. The correlation based on Small Eddy Cell model is used for subsequent analysis.

2.5.2. Improved Reduced Order Analysis

In the previous section, a regression model is constructed from CFD simulations to calculate the Sherwood number in the pipeline. The advantage of this approach is the additional information of local distributions of xenon concentration, liquid velocity, void fraction, and mass transfer coefficient. These flow variables have different profiles in the developed region, and when performing a reduced order analysis, these profiles are required to account for the local distribution of the flow. The motivation behind a reduced order analysis is to provide a less computationally expensive model for use in system analysis tools. A simple one-dimensional analysis is performed by Peebles [28], where no consideration is made for the profiles of the flow. This study uses a similar approach in deriving the one-dimensional two-fluid model to arrive at a one-dimensional mass transfer equation.

Starting for the phase averaged species transport equation,

$$\frac{\partial \alpha_k X_{k,i}}{\partial t} + \nabla \cdot \left(\alpha_k X_{k,i} \boldsymbol{u_k}\right) = \nabla \cdot \left(\alpha_k \left(D_k^t + D_{k,i}\right) \nabla X_{k,i}\right) + \Gamma_{k,i}. \tag{2-26}$$

To simplify the equation, only steady state conditions are considered, and the transient term on the left-hand side is ignored. When the variation of the system parameters is slow compared with the characteristic time in the xenon removal system, the system transient could be accounted for by changing the parameters in the steady state solution (referred to as quasi-steady state). Furthermore, the axial component of the diffusion term on the right-hand side is generally negligible compared with the interphase mass transfer term and could also be ignored. With these

assumptions, and performing area average over the cross-sectional area, the averaged equation in the liquid phase is derived as,

$$\frac{d}{dz}\left(\langle \alpha_l X_{l,i} u_l \rangle\right) = \langle \Gamma_{l,i} \rangle, \qquad \langle \Gamma_{l,i} \rangle = \langle k a_i \left(K_{l,i}^{\rho} X_{g,i} - X_{l,i}\right) \rangle. \tag{2-27}$$

The $\langle\ \rangle$ notation in the equation represent the area averaged quantities. It should be noted that,

$$\langle \alpha_l X_{l,i} u_l \rangle \neq \langle \alpha_l \rangle \langle X_{l,i} \rangle \langle u_l \rangle, \tag{2-28}$$

because of the cross-sectional distribution of the variables. The difference in the distribution of the variables results in a covariance as,

$$C_{X_l u_l} = \frac{\langle \alpha_l X_{l,i} u_l \rangle}{\langle \alpha_l \rangle \langle \langle X_{l,i} \rangle \rangle \langle \langle u_l \rangle \rangle}. \tag{2-29}$$

The $\langle\langle\ \rangle\rangle$ notation represents void weighted variables that are obtained in the CFD simulations. Because of the complexity in modeling these terms, it is desirable to reduce the number of covariances in the reduced equation. This could be done by performing a change of variables. The term $\langle \alpha_l X_{l,i} u_l \rangle$ physically represents the area averaged mass flux of xenon in the liquid phase, which is the physical quantity that should be tracked in the analysis. Therefore, this term is combined to define a new variable as,

$$\langle G_{l,i} \rangle \equiv \langle \alpha_l X_{l,i} u_l \rangle. \tag{2-30}$$

The second term in the area averaged equation is the interface mass transfer term $\langle \Gamma_{l,i} \rangle$, and should be related to the parameters of the system and the newly defined variable $\langle G_{l,i} \rangle$ to arrive at a solvable form,

$$\langle \Gamma_{l,i} \rangle = \langle k a_i \left(K_{l,i}^{\rho} X_{g,i} - X_{l,i}\right) \rangle. \tag{2-31}$$

First, the interfacial area term could be written as,

$$a_i = \frac{6\alpha_g}{d_{sm}}. \tag{2-32}$$

Plugging this into the source term, the following form is obtained,

$$\langle \Gamma_{l,i} \rangle = \frac{6}{d_{sm}} \left(K_{l,i}^{\rho} \langle \alpha_g X_{g,i} k \rangle - \langle \alpha_g X_{l,i} k \rangle\right). \tag{2-33}$$

Now, the area averaged quantities are separated using the intermediate covariances,

$$C_{X_g k} = \frac{\langle \alpha_g X_{g,i} k \rangle}{\langle \alpha_g \rangle \langle \langle k \rangle \rangle \langle \langle X_{g,i} \rangle \rangle}, \; C_{X_l k} = \frac{\langle \alpha_g X_{l,i} k \rangle}{\langle \alpha_g \rangle \langle \langle k \rangle \rangle \langle \langle X_{l,i} \rangle \rangle}. \tag{2-34}$$

Using these definitions, the source term is transformed into,

$$\langle \Gamma_{l,i} \rangle = \frac{6\langle \langle k \rangle \rangle}{d_{sm}} \left(K_{l,i}^{\rho} C_{X_g k} \langle \alpha_g \rangle \langle \langle X_{g,i} \rangle \rangle - C_{X_l k} \langle \alpha_g \rangle \langle \langle X_{l,i} \rangle \rangle\right). \tag{2-35}$$

In this way, the mass transfer coefficient is extracted outside of parentheses. The next step is to transform the variables from $\langle \alpha_g \rangle, \langle \langle X_{g,i} \rangle \rangle, \langle \langle X_{l,i} \rangle \rangle$ to $\langle G_{l,i} \rangle$. In order to do this, the liquid velocity must be included. Two more intermediate covariances are defined as,

$$C_{X_g u_g} = \frac{\langle \alpha_g X_{g,i} u_g \rangle}{\langle \alpha_g \rangle \langle\langle X_{g,i} \rangle\rangle \langle\langle u_g \rangle\rangle}, \qquad C_{X_l u_l} = \frac{\langle \alpha_l X_{l,i} u_l \rangle}{\langle \alpha_l \rangle \langle\langle X_{l,i} \rangle\rangle \langle\langle u_l \rangle\rangle}. \tag{2-36}$$

Plugging these into the equation, the results are,

$$\begin{aligned} \langle \Gamma_{l,i} \rangle &= \frac{6\langle\langle k \rangle\rangle}{d_{sm}} \left(K_{l,i}^{\rho} \frac{C_{X_g k}}{C_{X_g u_g}} \frac{\langle \alpha_g X_{g,i} u_g \rangle}{\langle\langle u_g \rangle\rangle} - \frac{C_{X_l k}}{C_{X_l u_l}} \frac{\langle \alpha_g \rangle \langle \alpha_l X_{l,i} u_l \rangle}{\langle \alpha_l \rangle \langle\langle u_l \rangle\rangle} \right) \\ &= \frac{6\langle\langle k \rangle\rangle}{d_{sm}} \left(K_{l,i}^{\rho} \frac{C_{X_g k}}{C_{X_g u_g}} \frac{\langle G_{g,i} \rangle}{\langle\langle u_g \rangle\rangle} - \frac{C_{X_l k}}{C_{X_l u_l}} \frac{\langle \alpha_g \rangle \langle G_{l,i} \rangle}{\langle \alpha_l \rangle \langle\langle u_l \rangle\rangle} \right). \end{aligned} \tag{2-37}$$

From the conservation of mass, the summation of the xenon mass fluxes in the liquid and the gas phase equal to a constant $\langle G_i \rangle_0$, the inlet xenon mass flux, which is a parameter from the system. Using this, and combining the covariances, the equations become,

$$\langle \Gamma_{l,i} \rangle = \frac{6\langle\langle k \rangle\rangle}{d_{sm}} \left(K_{l,i}^{\rho} C_g \frac{\left(\langle G_i \rangle_0 - \langle G_{l,i} \rangle \right)}{\langle\langle u_g \rangle\rangle} - \frac{\langle \alpha_g \rangle}{1 - \langle \alpha_g \rangle} \frac{C_l \langle G_{l,i} \rangle}{\langle\langle u_l \rangle\rangle} \right),$$

$$C_g = \frac{C_{X_g k}}{C_{X_g u_g}} = \frac{\langle \alpha_g X_{g,i} k \rangle \langle\langle u_g \rangle\rangle}{\langle \alpha_g X_{g,i} u_g \rangle \langle\langle k \rangle\rangle}, \qquad C_l = \frac{C_{X_l k}}{C_{X_l u_l}} = \frac{\langle \alpha_l \rangle \langle \alpha_g X_{l,i} k \rangle \langle\langle u_l \rangle\rangle}{\langle \alpha_g \rangle \langle \alpha_l X_{l,i} u_l \rangle \langle\langle k \rangle\rangle}. \tag{2-38}$$

At this stage, there are still variables of $\langle\langle u_g \rangle\rangle$, $\langle\langle u_l \rangle\rangle$, $\langle\langle k \rangle\rangle$, and $\langle \alpha_g \rangle$ that are left in the equation. The mass transfer coefficient $\langle\langle k \rangle\rangle$ would be calculated from the regression model derived in the previous section, or from the literature models. If the equation is solved with a system analysis code similar to RELAP5-3D, then $\langle\langle u_g \rangle\rangle$, $\langle\langle u_l \rangle\rangle$, and $\langle \alpha_g \rangle$ are readily available from the solution of the code, due to the fact that the hydrodynamics is essentially decoupled from the xenon removal process.

If the xenon removal equation needs to be solved along, more engineering models are required in addition to the regression model of $\langle\langle Sh \rangle\rangle$. In principle, during the interfacial mass transfer process, the void fraction and velocity of the different phases will be constantly changing, meaning the mass transfer equation cannot be decoupled from other equations. However, for the xenon removal cases, because the xenon concentration in the salt is relatively low, it can be assumed that other variables are not affected by xenon removal. Therefore, these variables can be related to the superficial velocities of the liquid phase and the gas phase using the drift flux model.

$$\langle\langle u_g \rangle\rangle = C_0 \langle j \rangle + \langle\langle u_{gj} \rangle\rangle, \qquad \langle j \rangle = \langle j_g \rangle + \langle j_l \rangle.$$

$$\langle \alpha_g \rangle = \frac{\langle j_g \rangle}{\langle\langle u_g \rangle\rangle}, \qquad \langle\langle u_l \rangle\rangle = \frac{\langle j_l \rangle}{1 - \langle \alpha_g \rangle}. \tag{2-39}$$

In the equations above, the superficial velocities $\langle j_g \rangle$, $\langle j_l \rangle$, and $\langle j_l \rangle$ are supplied as known parameters. The parameters C_0 and $\langle\langle u_{gj} \rangle\rangle$ in the drift flux model are functions of the superficial velocities. The void fraction could also appear in the C_0 and $\langle\langle u_{gj} \rangle\rangle$ models, in which case the solution is implicit. Nevertheless, with the input of the drift flux model and the assumptions of the decoupling between mass transfer and hydrodynamics, the xenon removal equation can be solved alone.

For some engineering models for horizontal pipe flow, the value of C_0 and $\langle\langle u_{gj}\rangle\rangle$ are constant [111, 112]. In this case, the equation could be explicitly integrated to arrive at a 0D model of the xenon removal. First, the equilibrium xenon mass flux is defined,

$$\langle G_{l,i}\rangle_\infty = \frac{\langle G_{l,i}\rangle_0}{\left(1+\frac{\langle\alpha_g\rangle}{\langle\alpha_l\rangle K_{l,i}^{\rho}}\frac{C_l\langle\langle u_g\rangle\rangle}{C_g\langle\langle u_l\rangle\rangle}\right)}. \tag{2-40}$$

This is the lowest possible xenon mass flux in the liquid phase at one pass. With this definition, the equation for the xenon mass flux in the liquid phase is solved as,

$$\langle G_{l,i}\rangle - \langle G_{l,i}\rangle_\infty = \left(\langle G_{l,i}\rangle_0 - \langle G_{l,i}\rangle_\infty\right)e^{-\frac{6\langle\langle k\rangle\rangle C_l L\langle\alpha_g\rangle}{d_{sm}\langle\langle u_l\rangle\rangle\langle\alpha_l\rangle}}. \tag{2-41}$$

Defining the removal efficiency based on the fraction of transferred xenon mass flux, the above equation is simplified as,

$$\beta = \left(1-e^{-\frac{L}{L_{xe}}}\right)\beta_\infty, \qquad L_{xe} = \frac{d_{sm}\langle\alpha_l\rangle\langle\langle u_l\rangle\rangle}{6\langle\langle k\rangle\rangle\langle\alpha_g\rangle C_l}.$$

$$\beta = 1-\frac{\langle G_{l,i}\rangle}{\langle G_{l,i}\rangle_0}, \qquad \beta_\infty = 1-\frac{\langle G_{l,i}\rangle_\infty}{\langle G_{l,i}\rangle_0}. \tag{2-42}$$

β is the removal efficiency, β_∞ is the maximum removal efficiency, and L_{xe} is the characteristic removal length of the process. L_{xe} could be viewed as the required length of the horizontal pipe for xenon removal. Based on the expression of L_{xe}, the required length increases with bubble diameter and liquid velocity, in the meantime decreases with mass transfer coefficient and void fraction. Because $\langle\langle k\rangle\rangle$ is also a function of liquid velocity through the Reynolds number, the characteristic length L_{xe} is a weak function of the liquid velocity. Apart from the expression of L_{xe}, another interesting finding is that, unlike C_l, the value of C_g only affects the maximum removal efficiency but does not change the characteristic length of the removal.

There is a seemingly problem in the expression for L_{xe}. As $\langle\alpha_g\rangle \to 0$, the required length for equilibrium becomes infinite. However, if there is only a single tiny bubble in the liquid, it is expected that the chemical equilibrium would be instantly reached. The problem lies in an assumption made in deriving the equation. It is assumed that the hydrodynamics is decoupled from the species transfer. This assumption requires that the volume of xenon in the gas phase is small compared with helium, even at equilibrium. Because of the low concentration of xenon in the salt, this assumption is valid even for very low void fraction cases, but not for the limiting case of a single tiny bubble.

2.5.3. Covariance of Horizontal Flow

As discussed in the last section, during the averaging process, covariances between different variables appear due to the difference in the distribution. The simulation data could be used to construct the required covariances. In this section, the covariances from different simulation cases are presented, along with the sensitivity to different parameters.

Before presenting the simulation data, it is beneficial to understand the covariances physically. The physics behind C_g and C_l is the effect of xenon distribution on its transport. If a uniform profile of xenon is prescribed, then C_l and C_g both reduce to one by definition. A better way to write C_g and C_l is,

$$C_g = \frac{\frac{\langle \alpha_g X_{g,i} k \rangle}{\langle \alpha_g \rangle \langle\langle k \rangle\rangle}}{\frac{\langle \alpha_g X_{g,i} u_g \rangle}{\langle \alpha_g \rangle \langle\langle u_g \rangle\rangle}}, \qquad C_l = \frac{\frac{\langle \alpha_g X_{l,i} k \rangle}{\langle \alpha_g \rangle \langle\langle k \rangle\rangle}}{\frac{\langle \alpha_l X_{l,i} u_l \rangle}{\langle \alpha_l \rangle \langle\langle u_l \rangle\rangle}} \approx \frac{\langle \alpha_g X_{l,i} k \rangle}{\langle \alpha_g \rangle \langle\langle k \rangle\rangle \langle X_{l,i} \rangle} = \frac{\langle \alpha_g k X_{l,i} \rangle}{\langle \alpha_g k \rangle \langle X_{l,i} \rangle}. \tag{2-43}$$

Then, from the expression, the meaning of C_g and C_l is the ratio of different weighted averages of xenon concentrations in the gas phase and the liquid phase, respectively. The approximation for C_l assumes low void fraction and strong turbulence, in which case $\frac{\Delta \alpha_l}{\alpha_l}$ is small, allowing α_l to be treated as a constant. The covariance C_g is essentially the concentration weighted by the removal rate, divided by the concentration weighted by gas flow. Since $X_{g,i}$ is larger where the removal rate is higher, and smaller where the gas flows is faster, it is expected that $X_{g,i}$ is larger than one. On the other hand, the inverse is expected for C_l, meaning it should be smaller than one. These predictions are confirmed by the simulation results.

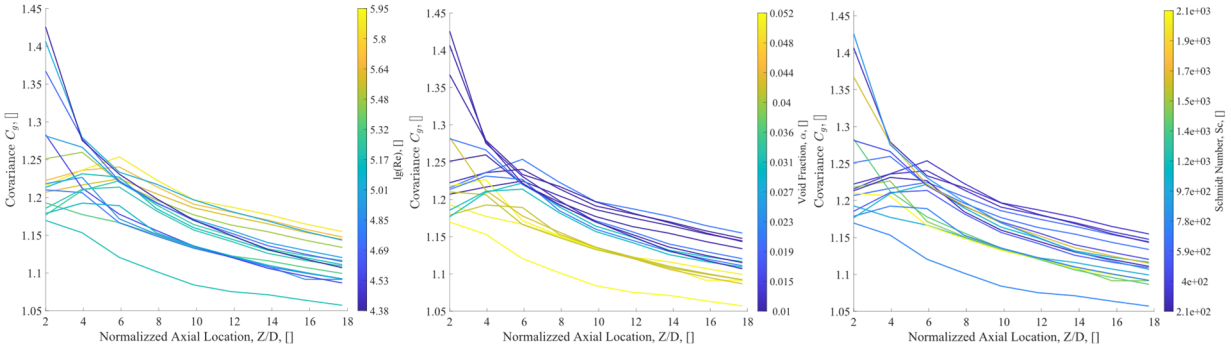

Figure 2.41. Sensitivity of C_g with Reynolds number, void fraction and Schmidt number versus z/D. The lines are scattered, indicating inappropriate scaling.

First, the simulation data used in the previous regression analysis are used again to extract the covariances. In Figure 2.41, the covariance C_g is plotted against the normalized axial location $\frac{Z}{D}$ where the color indicates the value of different parameters. From the results, it is noted that C_g is constantly changing along the axial location. This result is likely due to the constantly changing profiles of the xenon concentration since the other profiles do reach a steady state distribution. In addition, the value of C_g apparently increases with Reynolds number and decreases with void fraction. On the other hand, no clear dependence is shown for the Schmidt number. Because the variation of the covariances most likely comes from the xenon concentration profile, normalization with $\left(1 - \frac{\langle\langle X_{l,i} \rangle\rangle}{\langle\langle X_{l,0} \rangle\rangle}\right)$ could reveal more information on the distribution of C_g instead of plotting C_g against $\frac{Z}{D}$. Following this idea, the covariances are plotted against $\left(1 - \frac{\langle\langle X_{l,i} \rangle\rangle}{\langle\langle X_{l,0} \rangle\rangle}\right)$. The results are

shown in Figure 2.42. The drawback of the new representation is that it would be more difficult to relate the x-axis in the figures to the physical flow development.

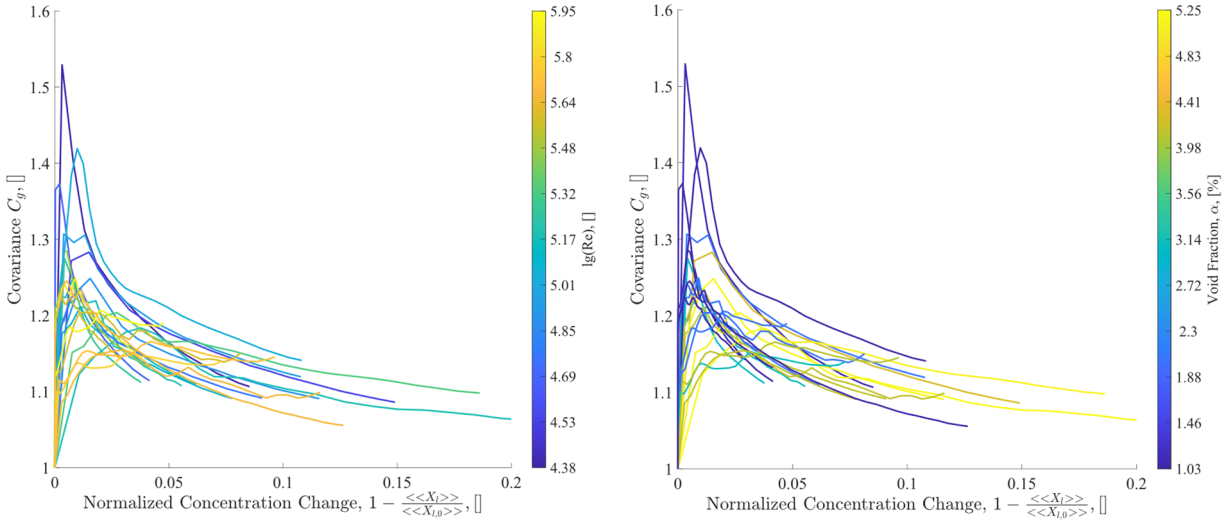

Figure 2.42. Sensitivity of C_g with Reynolds number and void fraction versus $\left(1-\frac{\langle\langle X_{l,i}\rangle\rangle}{\langle\langle X_{l,0}\rangle\rangle}\right)$. The scaling is improved compared with Figure 2.41.

As shown in the figure, the covariances for different cases are more concentrated and share a similar trend with the normalized axial location defined with xenon concentration. More importantly, the apparent dependence with void fraction and Reynolds number disappears, which is absorbed in the normalization with the ratio of $\frac{\langle\langle X_{l,i}\rangle\rangle}{\langle\langle X_{l,0}\rangle\rangle}$. Though the variation of the covariance between cases is still considerable, the change of C_g with the ratio of $\frac{\langle\langle X_{l,i}\rangle\rangle}{\langle\langle X_{l,0}\rangle\rangle}$ seems to follow an exponential trend after the entrance region. The trend for C_g with the parameter, $1-\frac{\langle\langle X_{l,i}\rangle\rangle}{\langle\langle X_{l,0}\rangle\rangle}$, is a quick rise followed by an exponential drop. The rising seems to be related to the development of the hydrodynamics of the flow. The maximum value of C_g is related to the ratio between species transfer and the mixing within the liquid phase.

Similarly, in Figure 2.43, the covariance C_l is plotted against the ratio of $\frac{\langle\langle X_{l,i}\rangle\rangle}{\langle\langle X_{l,0}\rangle\rangle}$ and different parameters. Similar to the figures of C_g, a clea trend with void fraction and Reynolds number is not observed. This result is an indication of the proper choice of the variable. However, C_l does not show a converging pattern in these figures. When the ratio of $\frac{\langle\langle X_{l,i}\rangle\rangle}{\langle\langle X_{l,0}\rangle\rangle}$ is small, the influence from C_l could become significant. To have a better understanding of these covariances, more simulation cases with extended pipes and increased mass transfer coefficient are carried out. The results for these extended cases are shown in Figure 2.44 and Figure 2.45. Here, the sensitivity of different parameters is examined for C_g and C_l respectively by keeping other parameters unchanged.

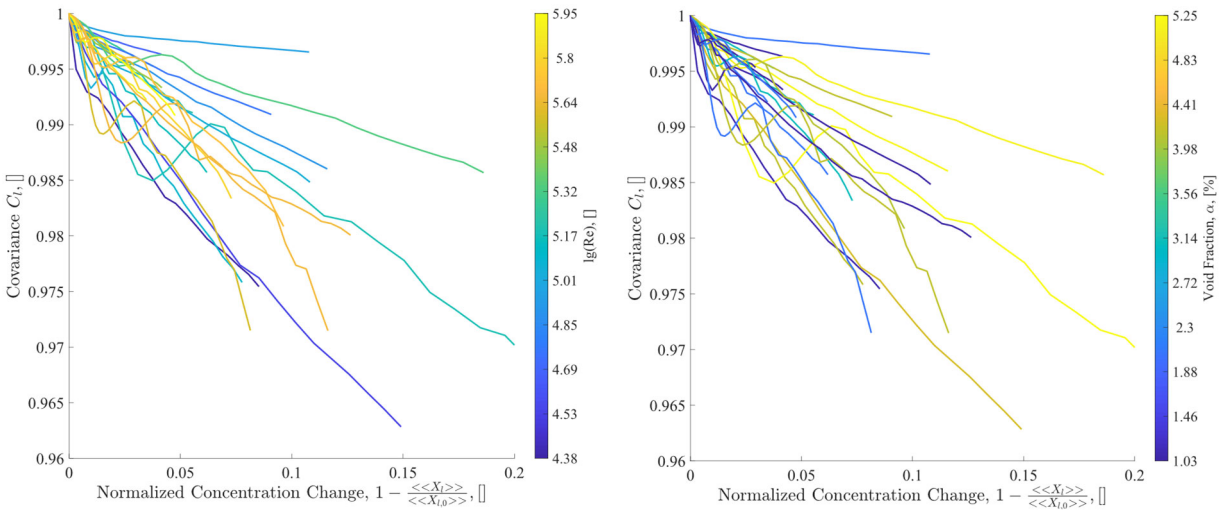

Figure 2.43. Sensitivity of C_l with Reynolds number and void fraction versus $\left(1-\frac{\langle\langle X_{L,i}\rangle\rangle}{\langle\langle X_{L,0}\rangle\rangle}\right)$. No convergence behavior is observed, indicating the pipe lengths in the existing simulations are not long enough for the studying C_l. Extended simulations are carried out as discussed later in this section.

As shown in Figure 2.44(a), the covariance C_g shares the similar trend between different cases, though the peak locations and magnitudes are not the same. The initial rise of C_g is related to the development of the profiles. After the profiles are established, the concentration profile becomes flatter and flatter as the removal goes on, indicating a decreasing C_g. It converges to unity at $\frac{\langle\langle X_{L,i}\rangle\rangle}{\langle\langle X_{L,0}\rangle\rangle}=0$ where xenon is completely removed. In Figure 2.44(b) to Figure 2.44(f), the sensitivity studies of different parameters are shown. The increase in viscosity causes the value of C_g to increase, but the peak location remains the same. When viscosity is increased, the dissipation near the wall becomes more significant, increasing the xenon removal, while the gas flow distribution is less affected, and thus the value of C_g increases. On the other hand, the effect of liquid velocity seems to be negligible for C_g, which is expected given that the liquid velocity in the turbulence regime is almost flat. The distinct trends in Figure 2.44(b) and Figure 2.44(c) and the above discussion show that the value of C_g also depends on dimensionless groups other than Re. The other three figures are related to the species transfer. The increase in diffusivity shifts the peak slightly to the right, but the magnitude appears to be unchanged. The increase in bubble diameter shifts the peak to the left and reduces the magnitude. The increase in void fraction moves the peak to the right and reduces the magnitude of C_g. Comparing the effects from hydrodynamic related and species transport related parameters, it appears that the hydrodynamic related parameters only change the magnitude of C_g, while the species transport related ones affect both the magnitude and the location of the peak.

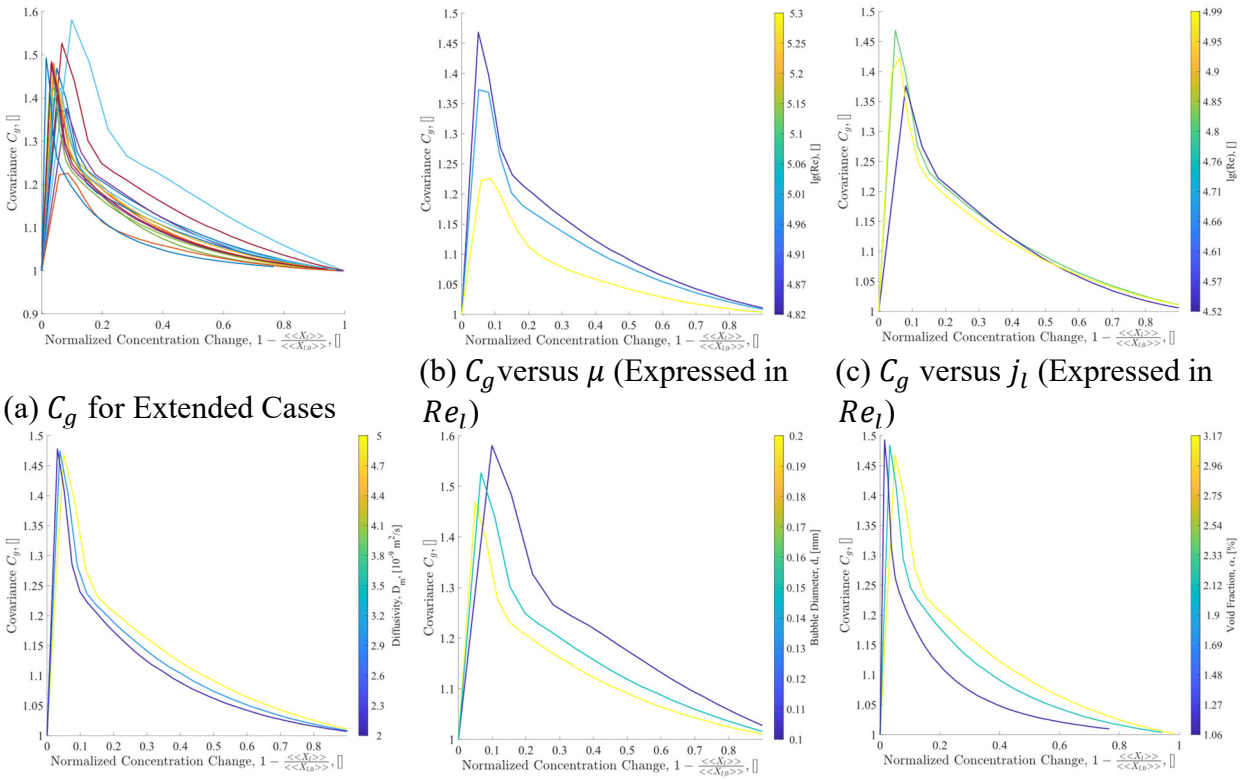

(a) C_g for Extended Cases

(b) C_g versus μ (Expressed in Re_l)

(c) C_g versus j_l (Expressed in Re_l)

(d) C_g versus Diffusivity

(e) C_g versus Bubble Diameter

(f) C_g versus Void Fraction

Figure 2.44. C_g from extended cases and sensitivity of parameters.

In Figure 2.45, the results for C_l under the same conditions are plotted. Compared with Figure 2.43, the covariances in the extended cases show a converging pattern after the initial "valley" of the curve. Unlike C_g, the covariance C_l does not converge to one as the dissolved xenon is removed. For cases with low flow velocity and high species transport rate, the value of C_l could become relatively low, and thus has a big impact on the physical process. From Figure 2.45(b), it is noted that the effect from viscosity is negligible. Unlike the case for C_g, the increase in viscosity increases the removal rate near the wall while simultaneously improving the mixing. The net result is therefore relatively small. From Figure 2.45(c), the increase in liquid velocity reduces the distributional effect as the overall mixing is improved. Compared with Figure 2.44, the effects of viscosity and velocity are inverted. In Figure 2.45(d), the increase in xenon diffusivity enhances the distributional effect. Initially, this result might seem counter-intuitive, but it is a great confirmation to the proposed physical explanation of the C_l. In turbulent flow, the mixing is dominated by turbulent diffusion, which is modeled through $D_{l,t} = \frac{\mu_t}{\rho_l Sc_t}$. Alternatively, the species transport between the phases is controlled both by turbulence And molecular diffusivity. Therefore, the increase in xenon diffusivity increases the xenon removal rate (or mass transfer coefficient k), when the mixing in the liquid is essentially unchanged, resulting in a smaller C_l.The bubble diameter and void fraction are related to the interfacial area a_i. When a_i increases, the value of C_l is smaller and vice versa.

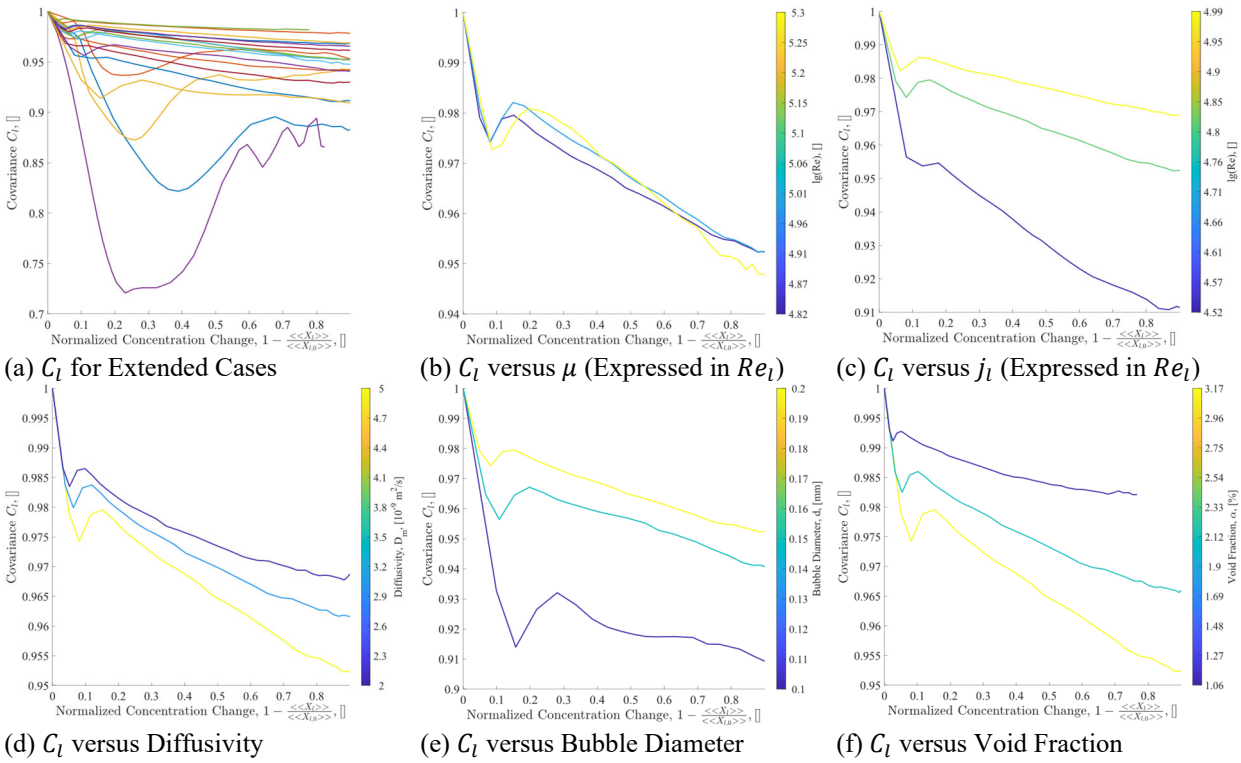

(a) C_l for Extended Cases (b) C_l versus μ (Expressed in Re_l) (c) C_l versus j_l (Expressed in Re_l)

(d) C_l versus Diffusivity (e) C_l versus Bubble Diameter (f) C_l versus Void Fraction

Figure 2.45. C_l from extended cases and sensitivity of parameters.

From the discussion C_l, it is noted that the effect from bubble diameter, void fraction, and diffusivity are almost equivalent through the change of the volumetric mass transfer coefficient ka_i. For hydrodynamic related variables like viscosity and velocity, the effect is more complicated because the mixing within the liquid is also changed. Therefore, it is beneficial to plot C_l with ka_i for the bubble diameter, void fraction, and diffusivity cases, as shown in Figure 2.46. The results for C_g are also included for comparison. As expected, the value of C_l decreases with increased ka_i and preserves the same functional form. For C_g, the effect of changing these parameters is different because of the different physics involved.

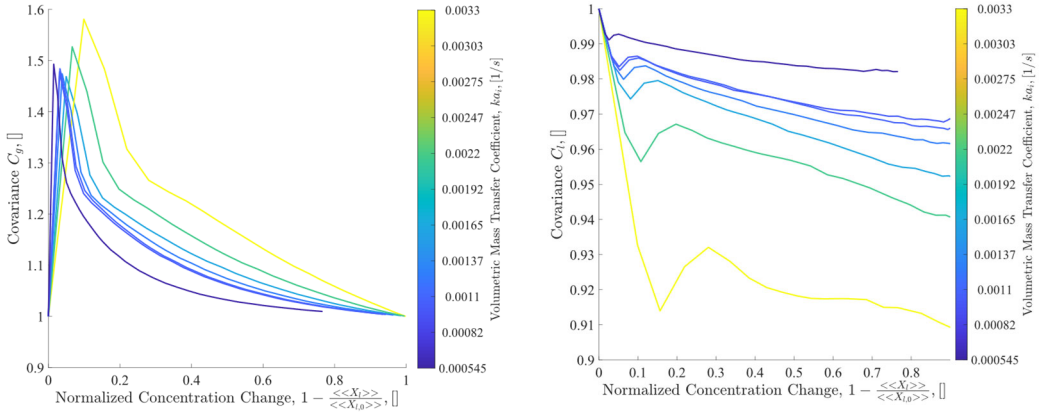

Figure 2.46. C_l and C_g at different ka_i. The effect of bubble diameter, void fraction and mass transfer coefficient can be roughly grouped together using ka_i.

It can be concluded that, the value of C_l and C_g are directly related to the xenon concentration profiles in the liquid and the gas phase. The qualitative change could be well predicted with physical reasoning. However, because of the complexity associated with two-phase flow mass transfer, proposing a model for these variables is challenging. For the scope of this chapter, such attempt is not made, but the discussion given in this section would be useful for future effort.

2.5.4. Analysis of Xenon Removal in Horizontal Pipe using Reduced Order Model

As discussed in Section 2.5.2, with the assumption of drift flux model and decoupling of hydrodynamics with species transfer, the removal efficiency could be solved as,

$$\beta = \left(1 - e^{-\frac{L}{L_{xe}}}\right)\beta_\infty, \qquad L_{xe} = \frac{d_{sm}\langle\alpha_l\rangle\langle\langle u_l\rangle\rangle}{6\langle\langle k\rangle\rangle\langle\alpha_g\rangle C_l},$$

$$\beta_\infty = 1 - \frac{1}{\left(1 + \frac{\langle\alpha_g\rangle}{\langle\alpha_l\rangle K_{l,i}^{\rho}}\frac{C_l\,\langle\langle u_g\rangle\rangle}{C_g\langle\langle u_l\rangle\rangle}\right)}. \tag{2-44}$$

In the following sections, the drift flux model proposed by Kocamustafaogullari and Huang [111] is adopted,

$$\langle\langle u_g\rangle\rangle = C_0\langle j\rangle + \langle\langle u_{gj}\rangle\rangle = 0.99\langle j\rangle - 0.09[m/s]. \tag{2-45}$$

The equilibrium ratio $K_{l,i}^{\rho}$ is related to the solubility of xenon in the fuel salt, a quantity which has not been measured experimentally for FLiNaK. A theoretical model is developed by Blander et al. [72]. The equilibrium ratio $K_{l,i}^{\rho}$ is predicted to be 1.1×10^{-5} at 600 ℃ in FLiNaK based on the model. However, based on its prediction on the solubility of other noble gases that have experimental data, large uncertainty is expected. The mass transfer coefficient $\langle\langle k\rangle\rangle$ is calculated based on the regression model from the CFD simulation results with the small eddy cell model in Section 2.5.1. In this section, sensitivity studies of different parameters are presented. The range of conditions of the sensitivity studies are listed in Table 2.11. The reference salt flowrate is taken as 10% of the loop flow rate of MSBR. The temperature is used to calculate the properties of the salt, using the correlations listed in

Table 2.12.

Table 2.11. Range of conditions in the sensitivity study.

Parameter	Range	Reference Value
d_{sm}, [mm]	$0.1 - 3$	0.5
Q_{salt}, [m^3/s]	$0.05 - 0.5$	0.1
D, [cm]	$10 - 40$	20.32
Q_{He}, [$\%Q_{salt}$]	$0.1 - 5$	2
$K_{l,i}^{\rho}$, [10^{-5}]	$1.1 - 21.2$	10
T_{salt}, [K]	$850 - 1100$	923
D_m, [$10^{-9}m^2/s$]	$0.5 - 5$	2.5

Table 2.12. Properties of FLiNaK used in the sensitivity study.

Property	Equation	Source
$\rho, [kg/m^3]$	$2579.3 - 0.624T$	Janz [77]
$\mu, [\text{Pa} \cdot \text{s}]$	$3.877 \times 10^{-5} e^{\frac{4327}{T}}$	Kubíková et al. [82]

2.5.4.1. Sensitivity of C_g and C_l

In Section 2.5.3, the covariances from the averaging process are calculated for the different simulation cases. In this section, their effect on the xenon removal process is explored in more detail. From the expression for the removal efficiency, it is noted that C_g only changes the maximum removal efficiency β_∞. At the reference condition given in Table 2.11, assuming $C_l = 0.9$, the maximum removal efficiency is shown in the red line of Figure 2.47(a). Because of the low solubility of xenon, the value of β_∞ does not change much with C_g. The change in C_g is only noticeable when the gas flowrate is relatively low compared with the salt flowrate. However, at this condition, the required pipe length for the removal would be extremely large as shown in Figure 2.47(b). Therefore, for the removal of xenon, the value of C_g is essentially negligible. For removal of fission product with higher solubility, it might need to be considered. On the other hand, the value of C_l always has a noticeable effect on the xenon removal process. The effect of C_l at the reference condition is shown in Figure 2.47(c) with $C_g = 1.2$. From the results, as the value of C_l deviates from unity, the removal of xenon becomes slower and slower. The physics behind this trend is discussed thoroughly in the previous section. Based on the calculation, it is expected that C_l could have up to a 20% influence on the removal rate. Though it is not the most influencing factor, its effect should be considered for an accurate engineering model. For the rest of this section, the value of C_l is prescribed as 0.9 and C_g is set as 1.2 for simplicity.

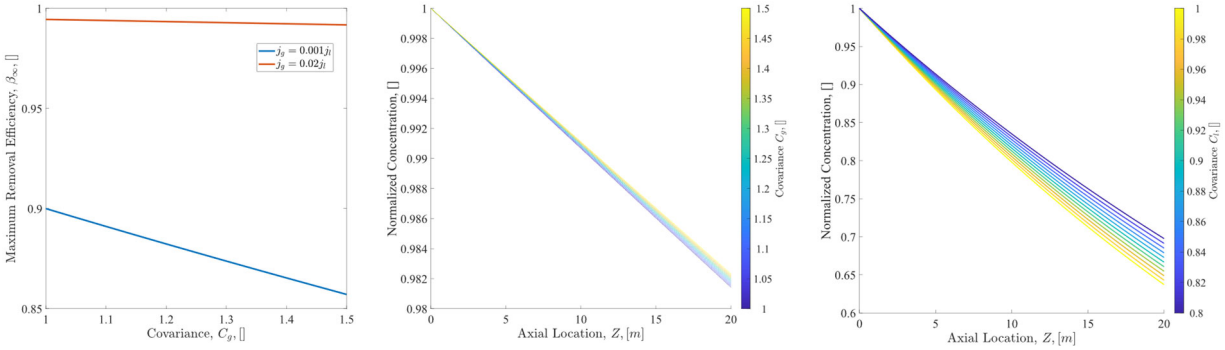

(a) Maximum Efficiency vs. C_g (b) Concentration at $j_g = 0.1\% j_l$ vs. C_g (c) Concentration vs. C_l.

Figure 2.47. Sensitivity of xenon removal to C_g and C_l.

2.5.4.2. Sensitivity of d_{sm} and Void Fraction

The rate of xenon removal is proportional to the available interfacial area between the gas and the liquid phase, a quantity which is directly related to the void fraction and the bubble mean

diameter. In Table 2.11, the range of helium gas flowrate is given which is used to calculate the area averaged void fraction reported in this section. This change of variables is made because void fraction is a more general quantity for presenting the results, yet it cannot be directly controlled, and thus is not considered as a parameter in Table 2.11.

Figure 2.48(a) shows that the normalized concentration changes significantly with bubble diameter. For large diameters, the concentration barely changes over the process. For smaller bubbles, the removal is relatively efficient, with about 90 percent of the xenon removed in 20 meters. The bubble size from different injection methods could be quite different. In the separator experiments performed by Gabbard [29], bubble diameter as small as 0.025 mm from shearing at the pump is reported. For a venturi type of bubble generator [113], bubbles with sizes below 0.25 mm are generated with scaling fluid at the relevant liquid flowrate. At low void fraction, the bubble coalescence is limited, and the small bubble diameter could be sustained for a relatively long time. For an efficient xenon removal system, the generation of small bubbles below 0.5 mm is crucial. Figure 2.48(b) presents the effect of void fraction on xenon removal. At the design reference condition, the void fraction should be as large as 10% to achieve an efficient removal. However, this calculation assumes a constant bubble diameter which might be invalid at a higher void fraction. Even if the bubble diameter does not increase significantly, a high gas flowrate would result in more off-gas to be processed, limiting the performance of the removal system.

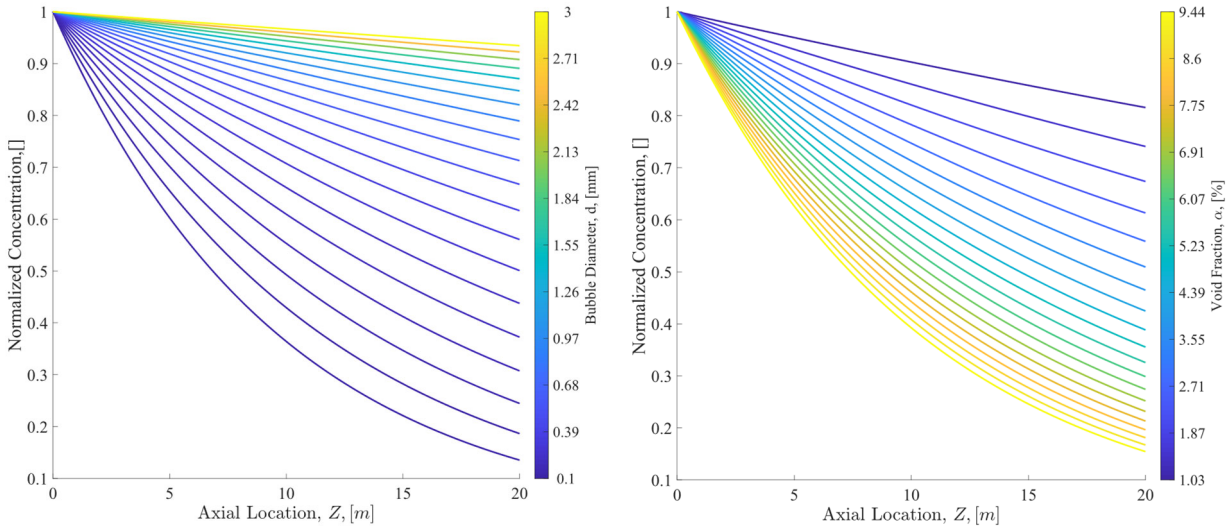

(a) Bubble Diameter, d.

(b) Void Fraction, α.

Figure 2.48. Sensitivity of xenon removal to bubble diameter and void fraction.

2.5.4.3. Effect of Pipe Diameter and Salt Flowrate

The current design of the xenon removal system considers diverting a fraction of the salt flow from the primary loop. Therefore, the overall removal is limited by the fraction of diverted salt flow. The requirement from the reactor physics analysis is expressed in terms of fraction of xenon removed between the outlet and inlet of the reactor, rather than the efficiency of the removal system. Therefore, in the sensitivity analysis, the salt flowrate in the removal system is the considered parameter, rather than the averaged velocity. The reference salt flowrate of $0.1 m^3/s$

corresponds to 10% of the loop flow rate in the MSBR design [10]. When varying the pipe diameter, the salt flowrate is fixed, and the average liquid velocity is changed accordingly.

In Figure 2.49(a), the xenon removal rate per unit length increases with pipe diameter at a fixed salt flowrate. As the pipe diameter increases, the viscous effect from the wall becomes weaker, and the mass transfer coefficient k decreases, however, the residence time of the salt in the system is increased. The net result from the process is an increased xenon removal rate within the system and the lower pressure head required to drive the salt flow. However, increasing pipe diameter results in increased costs due to the larger salt inventory and bigger system.

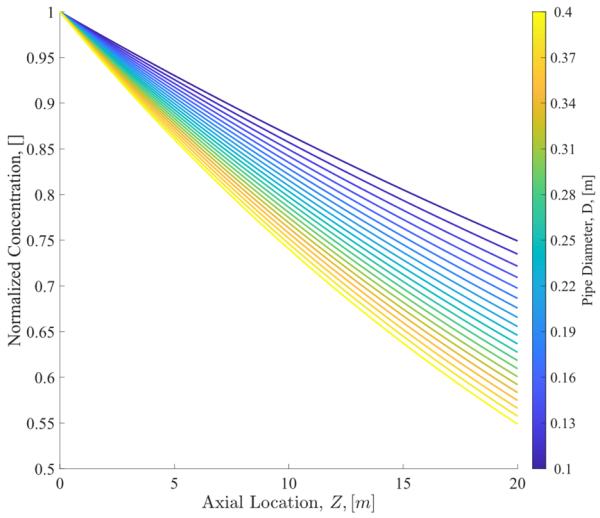

(a) Pipe Diameter.

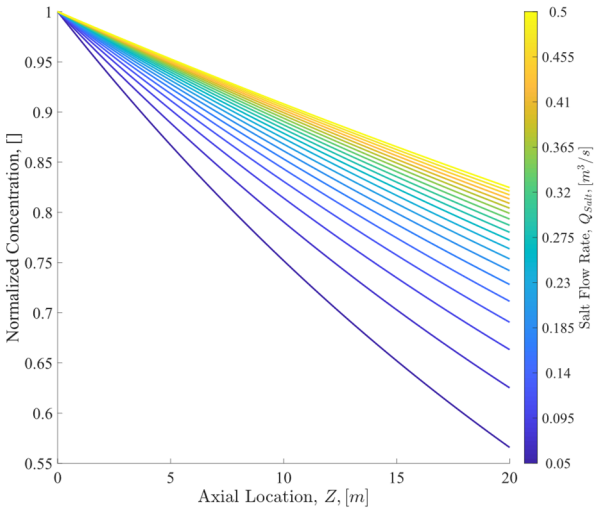

(b) Salt Flowrate.

Figure 2.49. Sensitivity of xenon removal to pipe diameter and salt flowrate.

A higher salt velocity result in a shorter residence time, but also a higher mass transfer coefficient. From the results in Figure 2.49(b), the net effect is a reduction in the removal efficiency. It is noted that as the salt flowrate increases, the removal efficiency becomes less dependent to the flowrate. This result indicates a removal system rather independent of the salt flowrate, although its size might make the idea impractical.

As a conclusion, the salt flowrate and system dimensions are closely related to the actual design of the reactor and the requirement onto the removal system. The removal efficiency will not be the only criteria, and economical factors should be considered.

2.5.4.4. Sensitivity of Xenon Diffusivity and Equilibrium Ratio

Xenon diffusivity and equilibrium ratio control the diffusion of xenon within the salt and the chemical equilibrium at the phase boundaries. These effects are decoupled from the hydrodynamics and are only controlled by the temperature and possibility the composition of the fuel salt. The experimental data for these quantities is not yet available in the literature. Therefore, large uncertainties are expected for the estimations given in Table 2.11, necessitating a sensitivity study.

The equilibrium ratio given in Table 2.11 and Figure 2.50(a) is the equilibrium ratio between the mass concentration of xenon in the liquid phase to the gas phase. The ratio is on the order of

10^{-4}, meaning the concentration of xenon in the gas phase would be much higher than the dissolved xenon concentration in the liquid phase. At equilibrium, the transport at the phase boundaries is balanced, and further removal could not be achieved. In principle, a smaller equilibrium ratio would increase the xenon removal efficiency, as shown in Figure 2.50(a). However, because of the low solubility of xenon, this effect is almost negligible compared to other factors. Similar to the discussion of C_g, the effect from the equilibrium ratio is only relevant for very low void fraction cases that are not likely to be adopted in the design.

Figure 2.50(b) shows that the removal efficiency increases with the diffusivity of xenon, which is expected. Even though the mixing within the salt is dominant by turbulence, the transfer at the interfaces still scales with the molecular diffusivity. The experimental data on the xenon diffusivity in fluoride salt used in MSRs is not available, though the diffusivity for inert gases in liquid is usually in the order of $10^{-9}m^2/s$, which is considered in this study. From the results, the removal efficiency changes significantly with the diffusivity at the reference conditions. Compared with other thermal physical properties [32], the uncertainty from xenon diffusivity is predominant. Therefore, the major uncertainty source in the current analysis lies in the xenon diffusivity in the salt. The large uncertainty related to the diffusivity prevents quantitative application of the current model to the design of the removal system.

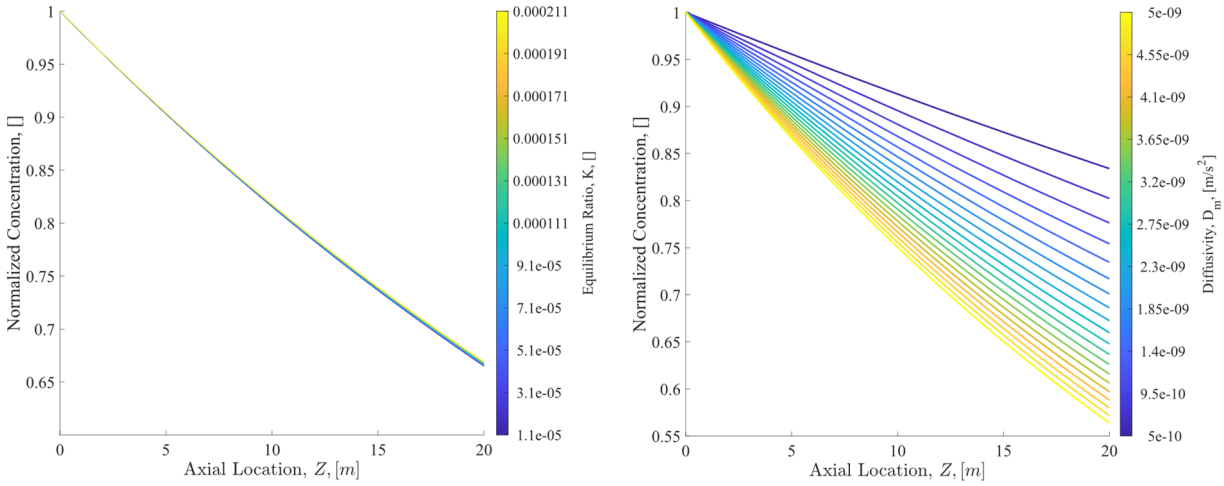

(a) Equilibrium Ratio.

(b) Xenon Diffusivity.

Figure 2.50. Sensitivity of xenon removal to equilibrium ratio and xenon diffusivity.

2.5.4.5. Effect of Salt Temperature

During the operation of the reactor, the temperature of the salt could be subjected to fluctuations. The effect of the temperature change is shown in Figure 2.51. Higher temperature results in higher removal efficiency, although the change is almost negligible. It should be noted that the calculation only considers the change of density and viscosity. The temperature effect from equilibrium ratio and diffusivity is not considered due to the lack of experimental data. However, it is expected that the equilibrium ratio does not have a significant influence on the results based on the observations in Figure 2.50 On the other hand, based on the Stokes–Einstein equation, diffusivity would increase with the temperature of the salt, and thus increase the removal efficiency, contributing to the trend shown in Figure 2.51. In fact, based on the Stokes–Einstein

equation and using the viscosity correlation from Kubikova et al. [82], the diffusivity as a function of temperature is,

$$D = \frac{D_0 T}{T_0} e^{4327\left(\frac{1}{T_0}-\frac{1}{T}\right)}. \tag{2-46}$$

If the value of T_0 is set as the reference temperature 923 K, then the ratio of $\frac{D}{D_0}$ varies between 0.6 to 2.5 in the temperature range. Therefore, xenon diffusivity is likely to be the dominant factor when temperature is changed judging from Figure 2.50 and Figure 2.51.

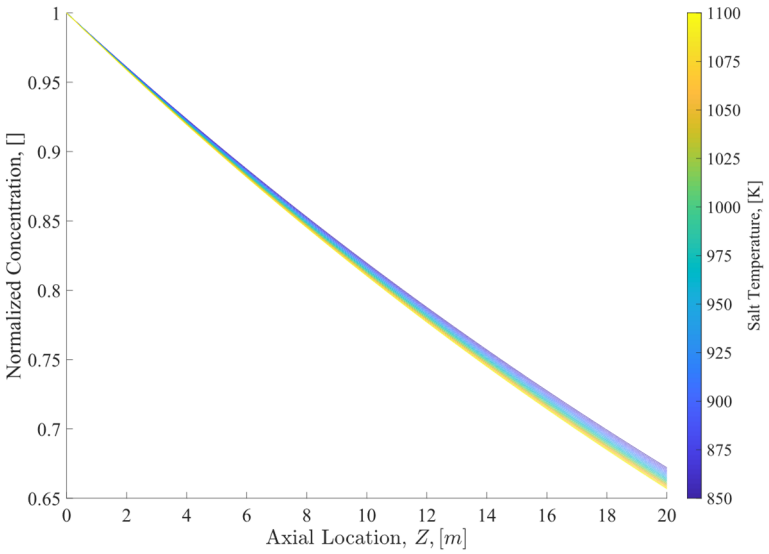

Figure 2.51. Sensitivity of xenon removal to temperature (only density and viscosity effect considered).

2.5.4.6. Effect of Circulating Helium Bubbles

In order to increase the removal of xenon, an alternative passage is to use a much longer removal length by allowing circulating helium bubbles to exist in the loop. In the MSRE, the circulating helium bubbles are assumed to form from the gas entrainment at the salt pump [26]. It is reasonable to consider a certain amount of circulating helium bubbles in the primary loop, contributing to the removal of xenon. For MSR design that does not include a separate xenon removal system, analysis of these circulating bubbles is crucial. For the separate xenon removal system considered in the current study, the influence of these circulating bubbles is closely related to the diverted flowrate. If only 10% of the salt flow is diverted from the main flow, on average a bubble would go through the primary loop 10 times before it enters the removal system and gets separated from the salt. Under this condition, it might be reasonable to assume the chemical equilibrium is reached between the phases when these circulating bubbles are removed. When equilibrium is assumed for the circulating bubbles, their effect can be described using the following equation,

$$\beta_{total} = \beta_{\infty,cb} + \left(1 - \beta_{\infty,cb}\right)\beta_{inject}. \tag{2-47}$$

Here, $\beta_{\infty,cb}$ represents the maximum removal efficiency for the circulating bubbles, β_{inject} is the removal efficiency in the pipeline removal section and β_{total} is the total removal efficiency compared with the xenon flowrate where no bubbles are present.

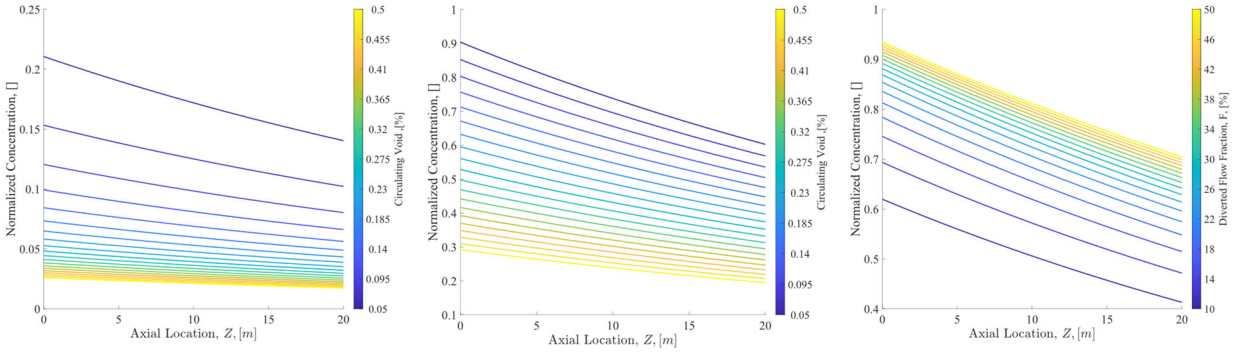

(a) Equilibrium Assumption. (b) Non-Equilibrium. (c) Diverted Flow Fraction.

Figure 2.52. Effect of circulating bubbles using estimated correlations. (a)Assuming equilibrium concentration is reached when the bubbles are removed. (b)Estimating the concentration of xenon in the removed bubbles. The difference with (a) is pronounced at lower circulating void fraction meaning that Equation 2-47 is not appropriate. (c)The diverted flowrate determines the averaged residence time of bubbles in the primary loop and thus changes the effect of circulating bubbles.

The results shown in Figure 2.52(a) suggest that a small fraction of circulating void could remove most of the dissolved xenon. However, for very low void fraction cases, the equilibrium assumption does not necessarily hold. To incorporate the nonequilibrium condition, the removal efficiency equation is modified as,

$$\beta_{total} = \overline{\beta_{cb}} + \left(1 - \overline{\beta_{cb}}\right)\beta. \tag{2-48}$$

$\overline{\beta_{cb}}$ is the averaged removal efficiency of the circulating bubbles at the bubble generator. β is the combined removal efficiency of the circulating bubbles and the newly injected bubbles in the pipeline removal section of the xenon removal system. The mean diameter of the circulating bubbles is likely to be much larger than the newly injected bubbles. Moreover, the void fraction of the circulating bubbles is also generally lower compared with the injected bubbles. Therefore, it is reasonable to assume $\beta \approx \beta_{inject}$, meaning the effects from the different bubbles are still decoupled. Furthermore, it is assumed that the "effective length" of the primary loop is around $L_{pr} \approx 100$ m, the bubble diameter of the circulating bubbles is 2 mm, the superficial velocity, and the mass transfer coefficient is the same to that in the removal section. Under these conditions, the value of $\overline{\beta_{cb}}$ could be calculated by applying a removal length of $\frac{L_{pr}}{F}$ (F being the fraction of diverted salt flow) to the removal efficiency equation.

For 10% of diverted flow, the modified results are shown in Figure 2.52 (b). From Figure 2.52(b), the equilibrium assumption is not valid, even for high void fraction cases. Nevertheless, Figure 2.52(a) and Figure 2.52(b) show that at low diverted flow condition, the effect from the circulating bubbles is profound. For 0.2% of circulating void, about 40% of the generated xenon would migrate to the gas phase before entering the removal system. When the diverted fraction is increased, the effect from these circulating bubbles would reduce due to the shorter contact time. Assuming the pipe diameter of the removal system changes accordingly at different diverted flow

fractions such that the superficial velocity is held constant, then the effect of diverted flow could be calculated by changing the factor F, as shown in Figure 2.52(c). When the diverted fraction increases, the removal efficiency drops because of the shorter contact time between the circulating bubbles and the salt.

As a conclusion, the existence of circulating bubbles could greatly improve the xenon removal from the salt to the gas phase. The drawback is that these circulating bubbles are not favorable in terms of neutronics and could potentially damage the primary loop pump. The optimized design involves considerations of many aspects of the reactor operation and requires further analysis.

2.6. Simulation and Analysis of the Bubble Separator

After the pipeline removal, most of the dissolved xenon has migrated to the helium bubbles. These bubbles should be separated from the salt before the diverted flow reenters the main flow, since bubbles in the core could cause problems to the reactor dynamics and potentially damage the fuel salt pump. Even in some designs where circulating bubbles are anticipated, the void fraction should still be kept low. The removal of helium bubbles is carried out at the bubble separator. Various conventional designs of bubble separators exist in the literature, including the horizontal separator, vertical separator, spherical separator, and cyclone separators. These separators usually require a large reservoir for the phase to separate under the effect of gravity and the pressure gradient. This reservoir could be expensive considering the increase in fuel salt inventory. The MSBR bubble separator design, on the other hand, does not require an additional reservoir, and the bubble separation is performed in-line with the pipeline removal. The bubble separation is achieved by introducing rotational movement to the flow through stational vanes [29]. As part of the development effort of the MSBR, experiments were performed by Gabbard [29] to investigate the bubble separation efficiency under different gas flowrates and bubble diameters. The axial and tangential liquid velocity profiles are also reported. Recently, Yin et al. carried out a series of experiments in a vertical bubble separator with a similar design to the MSBR separator [114-117]. In their experiments, the conditions are set in the range where complete separation of the bubbles is achieved, where a large bubble diameter is used. The ‘critical pressure’, defined as the pressure of the off-gas line where complete separation is achieved, is reported for different operation modes and boundary conditions. The liquid entrainment ratio is also studied since an excess amount of entrained liquid could be troublesome to handle in the actual salt system [114]. Subsequently, the velocity profiles of the liquid are measured with modern technique of particle image velocimetry (PIV) [116]. The single bubble trajectory is also evaluated in their work [115, 117].

To extend the knowledge obtained from these experimental studies to the actual molten salt system, CFD simulation is useful in the preliminary development phase of this technology. In this section, the modeling of the rotational bubble separator is presented. After comparison and validation with the existing experimental data, the sensitivities of different parameters are examined. Following the sensitivity and validation studies, the CFD simulation is extended to the molten salt system. From the simulation results, an engineering regression model is constructed to predict the bubble separation efficiency under different conditions, providing guidance to the development of a bubble separator suitable for the molten salt reactor. The simulations are carried out in STAR CCM+ using Idaho National Laboratory’s high performance computing system.

2.6.1. Setup of the 3D Simulation of Gas Separator

The governing equations and constitutive relations in the CFD simulations are briefly discussed in section 2.2. The Eulerian two-fluid model is used in the simulation and coupled with adaptive multi-size group model to account for the coalescence of the bubbles at the center of the separator. Rotational flow is created by the angled vanes at the entrance of the separator. The pressure gradient resulting from the liquid rotation pushes the bubbles towards the center where they coalesce. For some conditions, the coalescence of bubbles forms a "void core" at the center of the pipe [29, 114]. At the outlet of the separator, vanes of reversed angle eliminate the rotational movement and recover the pressure. Two gas outlets are located at the center of the vane hubs upstream and downstream. In the CFD simulation, uniform gas and liquid velocities are prescribed at the inlet. At the liquid and gas outlet, a pressure outlet condition is used with the pressure difference from experimental data. The simulation is transient in its form, but each simulation is run until the separation efficiency becomes steady in time. For some cases where periodical behavior is observed, the results reported are the time averaged quantities over the period. The geometry is modeled according to the dimensions given in the references [29, 118]. Two major simplifications are made for effective simulation. The tubing connected to the gas outlet is not modeled. Additionally, the vanes are not modeled explicitly. Instead, the rotational movement is created by introducing source terms in the momentum equation in the region where the vanes are located. The source term is evenly distributed along the axis, but different along the radial direction. If the flow is restricted so that the velocity along the radius is zero, then from conversation of momentum, the resulting tangential velocity from the source terms at the exit of the swirling vane region is,

$$u_t = u_z tan\alpha, \tag{2-49}$$

where α is the angle of the swirling vane. The drawback of this method is that these source terms do work on the flow. The static pressure changes across the swirling and recovery vanes are not created. According to the experimental data, this pressure change is around 1 meter of liquid pressure head in the range of conditions considered in the simulations. It is possible to enforce a rough energy conservation by applying a countercurrent force, although this is not performed in the simulation since the pressure change from the swirling vanes is relatively small compared with the pressure drop between the inlet and gas outlet line. With these simplifications, the simulation cases are almost axially symmetric. The only nonsymmetric effect is the gravity, which is perpendicular to the bubble separator in the MSBR design.

In Figure 2.53, the geometry and meshing pattern used in the simulation are shown. The pipe diameter is 4 inches at the inlet and reduces to 3.5 inches at the outlet. This design is used to retain the rotational velocity along the pipe.

Figure 2.53. Mesh and geometry for 3D simulation of the bubble separator.

In Figure 2.54, the grid independence study is shown. The indicator is the separation efficiency of the separator, defined as the gas flowrate through the outlet divided by the gas flowrate at the inlet. This choice is made because the separation efficiency is the most crucial parameter of the bubble separator. From the result, the separation efficiency fluctuates around 84% over 600 thousand cells. The characteristic length of the mesh at 650 thousand cells is 2.5 mm, about 1/40 of the pipe diameter.

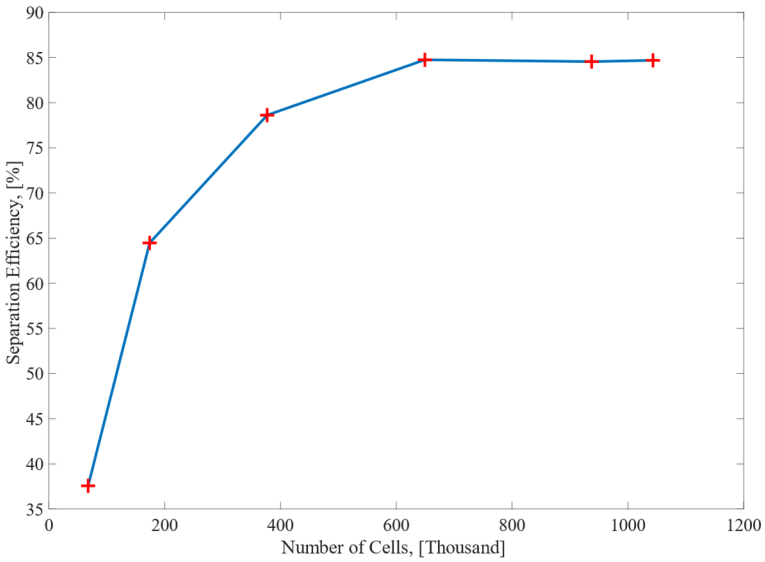

Figure 2.54. Grid independence study of the bubble separator. The fourth one is selected as the base size of the simulations.

In principle, 650 thousand cells or more should be chosen in the simulation. To accelerate the simulation process and maintain a similar level of mesh resolution, the characteristic length of 2.2 mm is retained, but the mesh is merged based on the gradient and magnitude of the void fraction. The mesh after the adaption is shown in Figure 2.55.

Figure 2.55. Mesh of adaptive size used in the CFD simulation. The mesh size is reduced using mesh adaption. The size of the mesh in the center is retained to resolve the void core region.

After the adaption, the removal efficiency changes to 82.8%, a result which is slightly lower than the higher resolution meshes. To assess the local effect of the adaption, the void fraction profiles for different meshes at the cross section downstream are plotted in Figure 2.56. Only a small change is observed compared with the original meshes for both the removal efficiency and the local void fraction profile. For the void fraction, the profile from the adaptive mesh agrees well with the highest resolution mesh away from the center. At the void center, the result is closer to the case with lower resolution. Therefore, the adaption is believed to be legitimate and is used in the 3D simulation of the bubble separator for validation against the data in the literature at various conditions [29, 116].

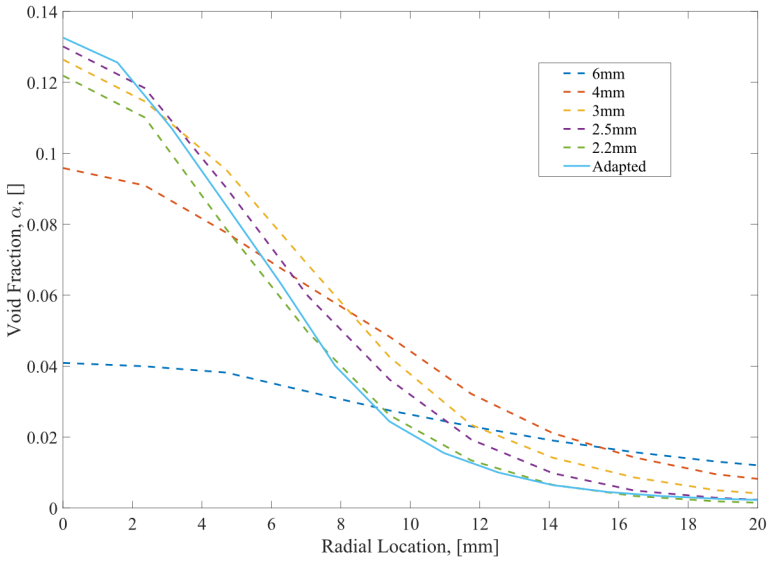

Figure 2.56. Comparison of void fraction profiles from different meshes. The adaptive mesh only slightly changes the void fraction profile and is used in the subsequent simulations.

Apart from the grid independence study, the effect of total bubble groups in the size group model should also be inspected. As shown in Figure 2.57, the separation efficiency is essentially unchanged with more groups. Therefore, three adaptive bubble size groups are used in the simulation.

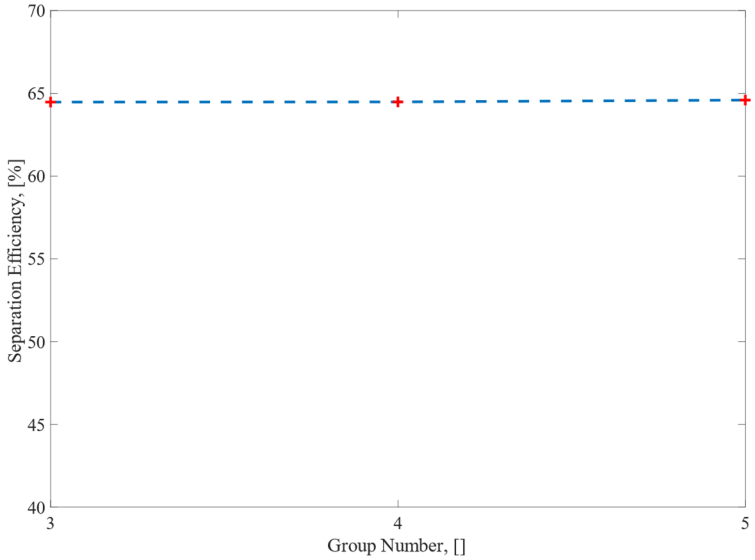

Figure 2.57. Comparison of separation efficiency at different total bubble groups. Three group is sufficient to resolve the problem.

The flow fields at the cross section along the axis from the adapted mesh are shown in Figure 2.58. The pressure decreases from the wall to the center, and thus the bubbles are pushed towards the center and coalescence with each other, forming a void core region. Most of the pressure drop is near the gas outlet because of the high liquid velocity. The pressure profile away from the

boundary is only slightly affected by the outlet pressure and mostly determined by the rotational movement of the flow. However, the liquid entrainment rate into the gas removal line is mostly controlled by the pressure at the gas outlet. The removal efficiency of the separator is also affected by the pressure condition at the boundary.

The void fraction profile generally agrees with the shape of the void core reported in literature, though it should not be considered as the instantaneous distribution of the gas bubbles. The void core has smaller diameter than the apparent diameter of the void fraction profile shown in the contour. Because of turbulence, the void core is not perfectly stable and would fluctuate slightly in space. The void fraction profile corresponds to the temporal density of the void core and thus has a larger diameter.

The liquid velocity profiles are divided into the tangential velocity and axial velocity. The radial velocity is negligible in the simulation and is therefore not included. From the velocity profile, it is observed that the tangential velocity first increases towards the center then decreases. This agrees with the basic vortex structure in viscous fluids, like the Rankine vortex. For the axial velocity, the velocity also first increases, then decreases towards the center, with the centerline velocity being non-zero. In fact, near the upstream gas outlet, the liquid velocity becomes countercurrent because of the liquid entrainment to the outlet.

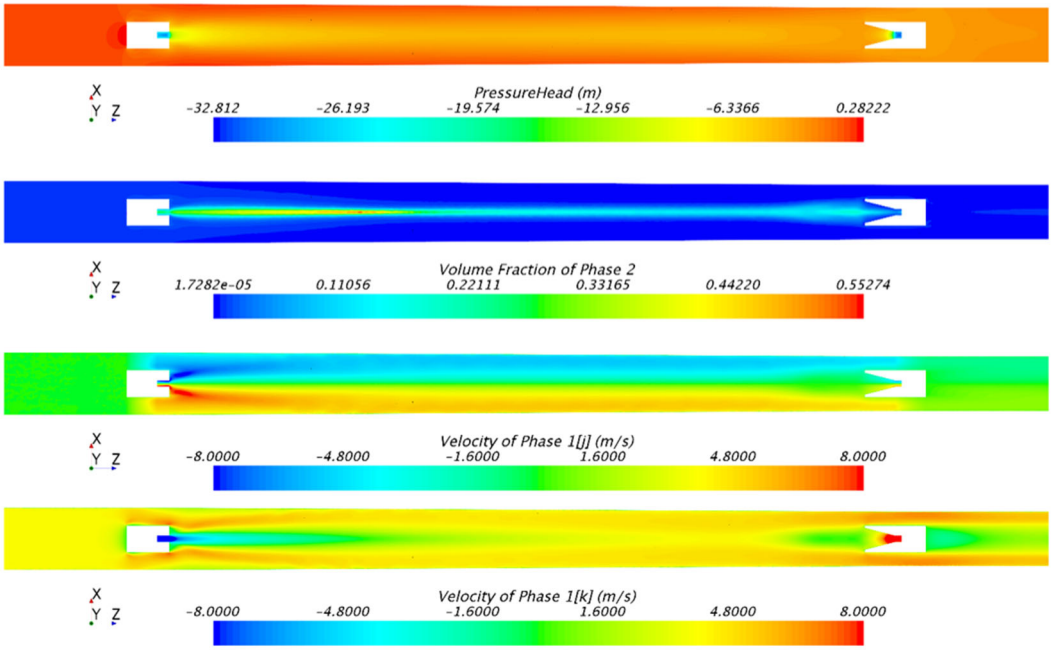

Figure 2.58. Void fraction, pressure, and liquid velocity contours at cross section.

In this section, the geometry, meshing, model independence, and grid independence study of the bubble separator simulation are introduced. The general trends of flow fields from the simulation are discussed. In the next section, the simulation results are compared with experimental data under various conditions.

2.6.2. Validation of 3D Simulation with Data in Literature

In the report from ORNL [29], experiments performed with a water and calcium chloride solution as a scaling fluid are reported for different separator configurations and boundary conditions. The results include separation efficiency, pressure profiles, and velocity profiles. For the separation efficiency, experiments with large and small bubbles are performed. Large bubbles are defined to have a diameter estimated between 0.508 to 0.762 mm. The reported removal efficiency for large bubbles is around 95% at different gas superficial velocities. The small bubbles have diameters ranging from 0.0254 to 0.127 mm. The removal efficiency for the small bubbles increases with gas superficial velocity and is considerably lower than that for the large bubbles. The parameter reported in the experiment includes gas flowrate and inlet void fraction. It is not clear what pressure is used to calculate the gas flowrate. Using the reported gas flowrate, the value of $\frac{j_g}{j_l}$ is about 5 times of the reported inlet void fraction. Due to these discrepancies, the inlet void fraction is chosen in the comparison with simulation data.

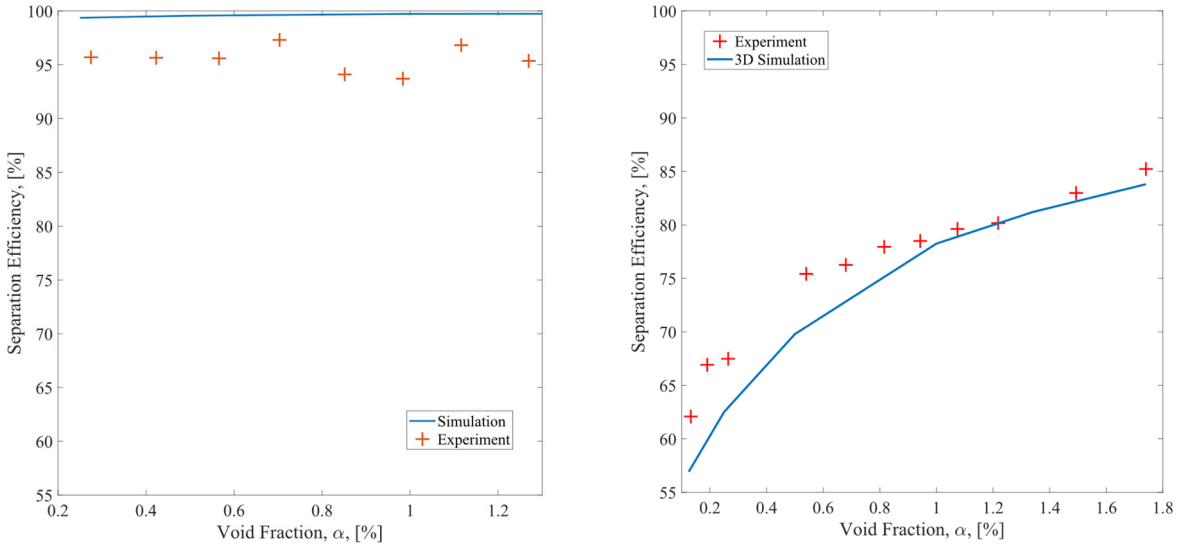

(a) Large Bubbles (b) Small Bubbles

Figure 2.59. Comparison of separation efficiency between simulation and the ORNL experiment [29].

The comparisons of the simulation results for the large bubbles and the small bubbles are shown in Figure 2.59. The large bubble diameter used in the simulation is 0.68 mm according to the report. Figure 2.59(a) shows that the simulation gives a slightly higher prediction of removal efficiency. For all the inlet void fractions, the removal efficiencies are about 99%, yet the efficiencies from the experiments are around 95%. This result might be related to the simplifications made with the separator geometry, as a completely symmetric configuration should have less escaped bubbles. For the simulation with small bubbles, the bubble diameter at the inlet is set at 0.17 mm, slightly larger than the reported value. The reason for doing this is because only the separator section of the experiment is modeled. The coalescence after the bubble generator, which is the pump in the experiment, but before the inlet of the simulation is not included. Thus, a larger bubble diameter is likely to be present at the inlet. From Figure 2.59(b), the separation efficiency increases with inlet void fraction, which is due to the increase in bubble coalescence. Good agreement between the simulation and the experiment is found. Even though not showing in

the range of condition considered here, it is expected that at higher void fraction, the separation efficiency will start to reduce because of the flow limit of in the outlet pipeline.

Apart from the separation efficiency, the velocity and pressure profiles at the cross-section perpendicular to the axis are also reported. Although these local quantities are not directly related to the function of the bubble separator, comparisons made with them reflect the extent of underlining physics that is captured by the CFD model. The velocity is measured through wedge type of velocity probe which uses the local pressure at two taps on the probe to calculate the velocity parallel to it. It is unclear how accurate this method is for two-phase flow given the complexity related to two-phase pressure drop. In the high void fraction region near the center, the measurement might have significant uncertainties. Moreover, the disturbance caused by the intrusive method is not considered in the simulation. The static pressure is also measured through the probe. The radial profiles near the upstream gas outlet from the experiment and the simulation are shown in Figure 2.60.

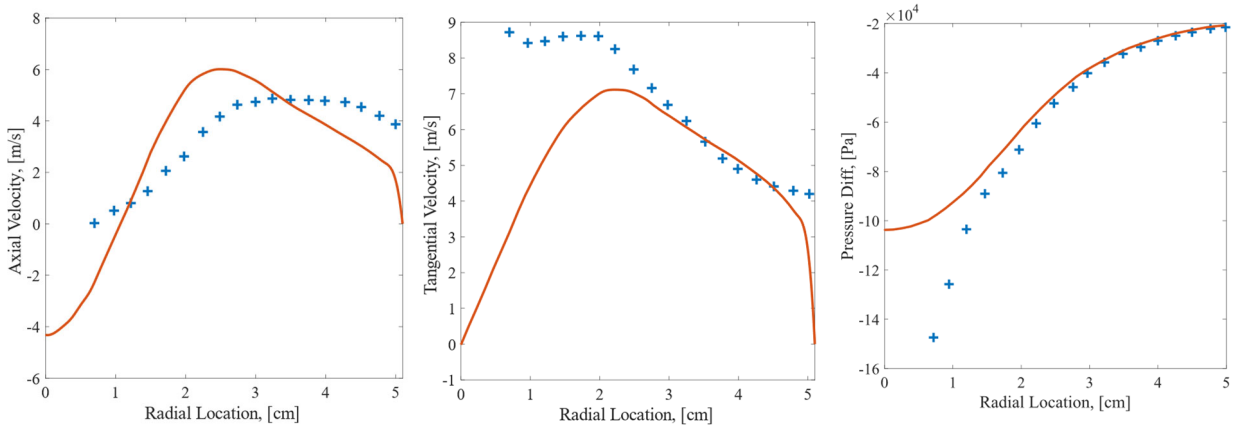

(a) Comparison of Axial Velocity (b) Comparison of Tangential Velocity (c) Comparison of Pressure

Figure 2.60. Comparison of the velocity and pressure profiles between the simulation and the ORNL experiment [29].

From the figures, the pressure profile from the simulation agrees well with the experimental result near the wall. In the void core region, the measurement from the experiments is significantly lower than the simulation. This might be because the void core is not described as a continuous phase, but rather approximated with dispersed flow model. The axial velocity profile of the simulation shares the same trend with the experimental data, though the shape of the profile is different. This is mostly likely from the way the swirling vanes are modeled. Because of the angle of the vanes and the rotational movement, the liquid is pushed towards the wall, and it is expected that the liquid velocity near the wall would be higher than a simple annulus channel as in the simulation. Since the measurement is made near the swirling vane, the effect from this effect is carried along and reflected in the comparison. For the tangential velocity, good agreement is found between 2 cm and 4 cm. Additionally, the experimental tangential velocity near the center does not decrease as the radial location decreases. This result is unexpected given that the tangential velocity at the vortex center should be zero for viscous flow. A possible reason for this result is due to the disturbance from the probe which could enhances the rotational movement in the unoccupied region.

Except for the experiments done at ORNL, following the recent development of the molten salt reactor, experiments on this special type of bubble separator are performed by Yin et al. [114-118] and Qian et al. [119, 120] for the Chinese Thorium Molten Salt Reactor (TMSR) program. The performed experiments include bubble behavior, entrained liquid flowrate, flow visualization and particle image velocimetry (PIV). A key difference with the ORNL data is that the separator efficiency is not studied since the experiments were performed with large bubbles where complete removal is usually achieved.

In the 2015 publication by Yin et al. [118], flow visualization and entrained liquid flowrate measurement are performed under a different operational mode. A corresponding simulation for the operational mode with both upstream and downstream removal is carried out with an outlet pressure difference of 23 kPa and Reynolds number of 77576. The mesh and void fraction profile are shown in Figure 2.61. The mesh is of similar resolution as for the ORNL experiment, and therefore a grid independence study is not performed again for the new mesh. From the void fraction profile, the characteristic is essentially the same. The overall removal efficiency is 97.6%, very close to complete removal in the experiment.

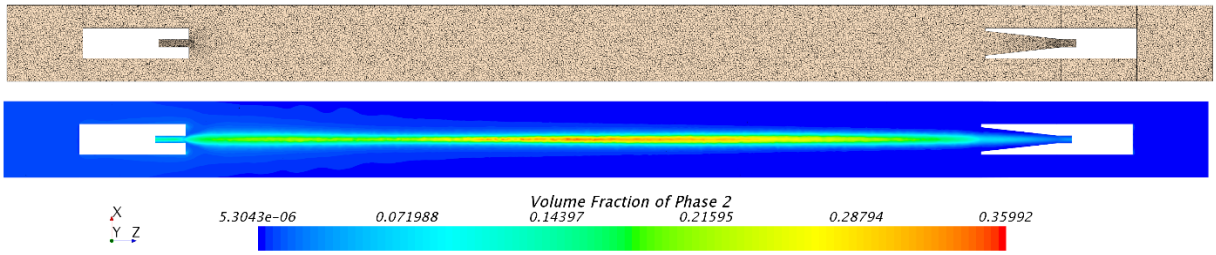

Figure 2.61. The mesh and void fraction profile from simulation based on conditions given in Yin et al. [118].

In the 2019 publication of Yin et al. [116], velocity profiles at the axial cross section are obtained through PIV. It is noted that the pressure boundary condition is not clearly given in the paper, only that the outlet pressure is maintained within 0.25 MPa. In the simulation, 0.25 MPa is taken as the pressure difference between the inlet and the gas outlet. The comparison is shown in Figure 2.62 for different measurement ports.

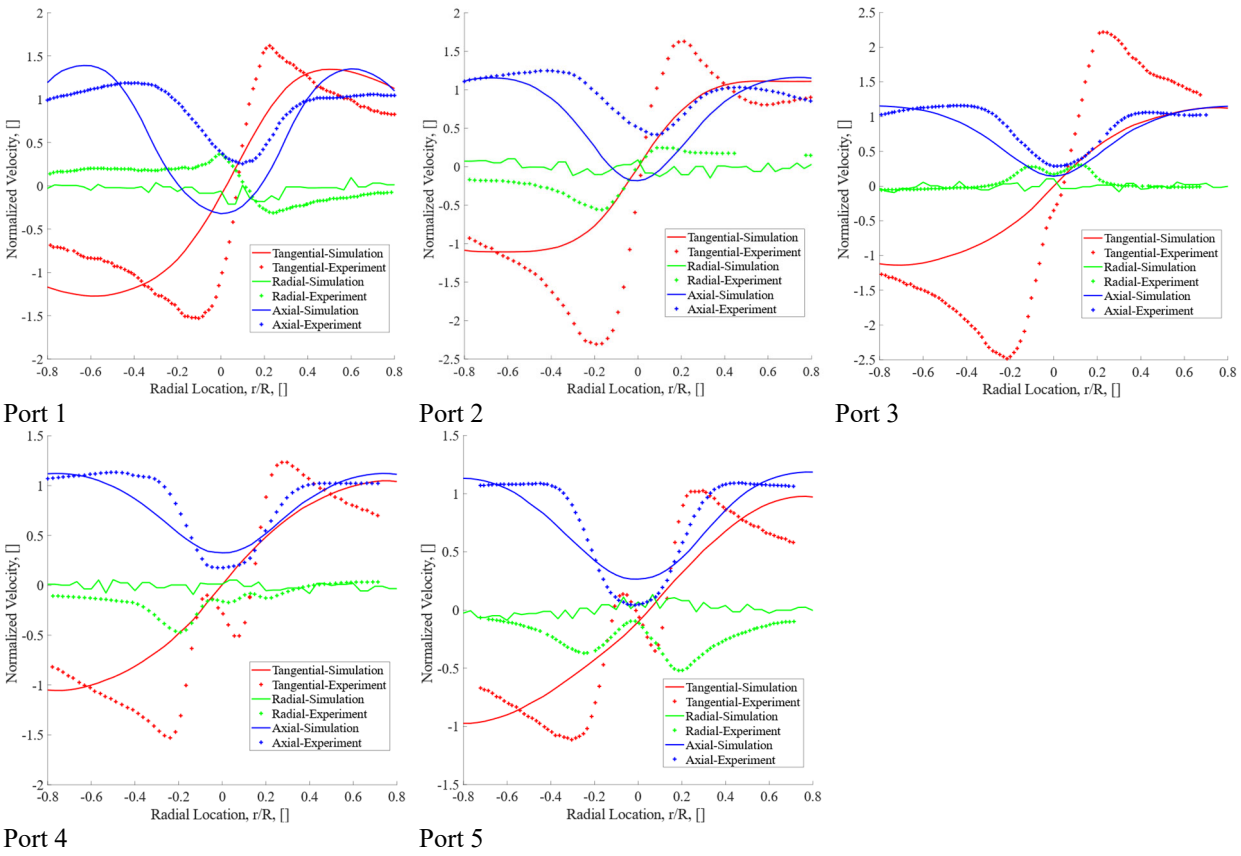

Figure 2.62. Comparison of velocity profiles in bubble separator between the simulation and data from Yin et al. [116].

The dashed lines are the data from the experiment, while the solid lines are profiles obtained in the simulation. The color represents the different components of velocity. At port 2, a portion of the experimental data for radial velocity is blocked in the original figure. The measurement ports are evenly distributed along the flow direction with Port 1 two pipe diameters away from the swirling vane. The most notable difference is the radial velocity. In the simulation, the radial velocity is almost zero everywhere in the separator because of the symmetrical setup. The positive radial velocity is towards the wall, and the negative velocity is towards the center. However, in the experiment, a strong asymmetrical effect is observed with radial velocity towards one side of the pipe at Port 1 and Port 2. From Port 3 to Port 5, the radial velocities in the experiments become roughly symmetric, but are changing constantly along the flow. This indicates that the symmetrical setup of the current simulation may not be sufficient to capture the delicate radial movement, even though the removal efficiency agrees. For the axial velocity, good agreement is found between the simulation and experiments at Port 3, 4 and 5. For Port 1 and Port 2, the axial velocities from the simulation are lower than the experimental results. Especially, at Port 1, the axial velocity from the simulation is countercurrent and towards the gas outlet, yet the experimental data is concurrent and away from the outlet. This result could be related to the difference in the boundary condition and the asymmetrical effect. For the tangential velocity, only rough agreement is found between the simulation and experiment, except for Port 3 where the magnitude of the tangential velocity is about the same, though the peak shape and location are different. The experimental tangential

velocity at Port 3 is significantly higher than the simulation result and the experimental tangential velocity at Port 2. This is unexpected as the frictional force should only reduce the rotational velocity, and the normal force from the wall does not contain any tangential components to enhance the rotation. At Port 4 and Port 5, the rotational velocity changes its direction from the center to the wall and is not reflected in the simulation.

From the validation study presented in this section, it is concluded that the CFD model and the simplified geometry could satisfactorily predict the separation efficiency in the separator at various conditions. However, for the local velocity profile, the agreement between the experiments and the simulation is limited. The general trend of the velocity is captured, although some fine features of the flow are missing from the CFD simulation, because of simplified physics at the void center and the simplified separator geometry. Nevertheless, in the current study, the interest is to properly predict the separation efficiency. After the validation, simulations should be performed to examine the sensitivity of different design parameters.

2.6.3. 2D Axisymmetric Simulation of Gas Separator

As discussed in the previous section, the simulation conditions are essentially axisymmetric except for gravity. When compared with the centripetal acceleration, the gravity perpendicular to the axis is negligible and thus could be ignored. Moreover, most of the flow fields in the CFD simulation are also axisymmetric. Therefore, the 3D simulation could be simplified to a 2D axisymmetric simulation in principle. When conditions are assumed for the simulation, the flow variables are assumed to be functions of r, z, and t in the cylindrical coordinate system. This does not mean the rotational movement is absent, but that u_θ is independent of θ, and that no net transport perpendicular to the 2D axisymmetric plane exists in the simulation. The motivation of reducing the simulation from 3D to 2D is for preforming a large number of simulation cases within a reasonable time, which is important for the sensitivity study and design optimization.

Due to the modified solution strategy, the momentum source term is not used in the 2D axisymmetric simulation. As a result, the swirling vane and recovery vanes are not included, and the swirling flow is prescribed at the inlet. The turbulent dispersion force model is changed to Simnion's model [121] which is readily available in STAR CCM+. Moreover, the Tomiyama drag and Tomiyama lift model require the input of gravity acceleration, which could not be included for a horizontal axisymmetric separator. Simplifications are made based on observations with the 3D simulation. The drag is mostly important in the low void fraction region where the bubble diameter is small. In the high void fraction region, the drag is hindered by the coalescence and symmetric interfacial area expression. Therefore, the drag model could be changed to Schiller-Naumann model which is suitable for small bubbles. [122]

On the other hand, the lift coefficient calculated from the Tomiyama model is shown in Figure 2.63. Because of the strong shear introduced by the rotational movement, the lift coefficient is essentially a constant equal to 0.288 in the separator. Therefore, this constant lift coefficient is used in the 2D axisymmetric model.

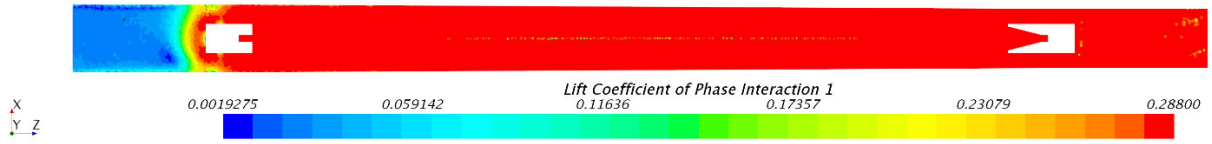

Figure 2.63. Tomiyama lift coefficient from the 3D simulation of bubble separator. The value is constant in the swirling section which allows the model to be simplified.

An example mesh and void fraction contour from the 2D simulation are shown in Figure 2.64. From the results, it appears that the void fraction contour shares the same pattern as seen in the 3D simulation. The major difference is that the void fraction could reach one at the center of the separator, while in 3D simulation, this is not observed.

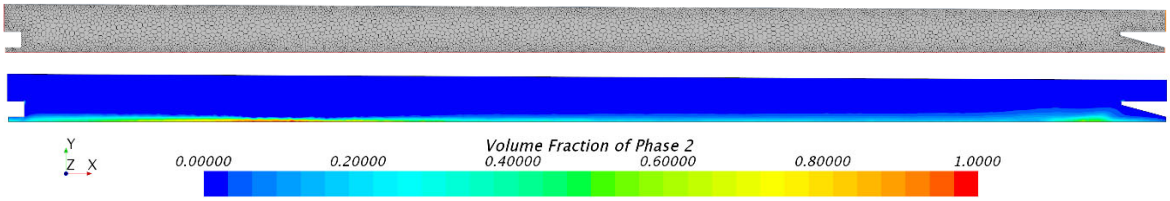

Figure 2.64. Typical mesh and void fraction profile from 2D simulation of bubble separator.

Whether the 2D simulation is a proper simplification of the 3D simulation depends on whether it could satisfactorily predict the separation efficiency as the 3D simulation. Therefore, simulations are carried out under the same condition as in the validation study in Figure 2.59(b). The result of the comparison is shown in Figure 2.65. The 3D and 2D simulation almost give the same prediction for low void fraction conditions. For higher void fraction, the 2D simulation gives slightly higher separation efficiency than the 3D simulation, which could be related to the higher void fraction at the void center.

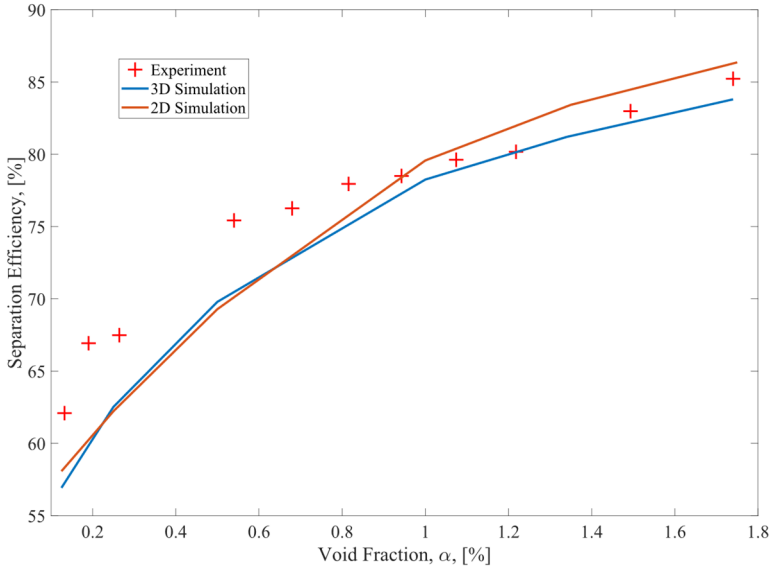

Figure 2.65. Comparison of predicted separation efficiency between 3D and 2D simulation.

Because of the good agreement between the 3D and 2D simulation, it is concluded that the 2D axisymmetric simulation could be used for the sensitivity study and design parameter evaluation presented in the next section.

2.6.4. Sensitivity of Design Parameters

In order to design a bubble separator that has good performance under various operational conditions, it is necessary to understand the effect of different parameters on the separation efficiency. In this section, sensitivity study from the 2D simulation is presented. The reference conditions used in the sensitivity study are summarized in Table 2.13.

Table 2.13. Reference conditions used in the sensitivity study.

Parameter	Range
Superficial Velocity	$5.534\ m/s$
Viscosity	$0.00423\ \text{Pa} \cdot \text{s}$
Density	$2004\ kg/m^3$
Void Fraction	1%
Bubble Diameter	0.127 mm
Pipe Diameter	15.24 cm
Gas Outlet Diameter	1.59 cm
Vane Slope	1

The parameters discussed in this section could be divided into two categories. The first category is design parameters which could not be changed intentionally during the operation. This includes the pipe diameter, outlet diameter, vane slope and bubble diameter. The bubble diameter, depending on how bubbles are generated, could vary depending on the gas flowrate or liquid velocity. Nevertheless, it is still viewed as a design parameter since it is unlikely to be directly controllable for most of the bubble generator designs. The other category is the operational parameters that are expected to be changing during reactor operation. This includes the salt flowrate, helium flowrate and salt temperature. The design parameters should be selected in a way that the efficiency of the separator is high enough within the range of operation parameters encountered in normal operation.

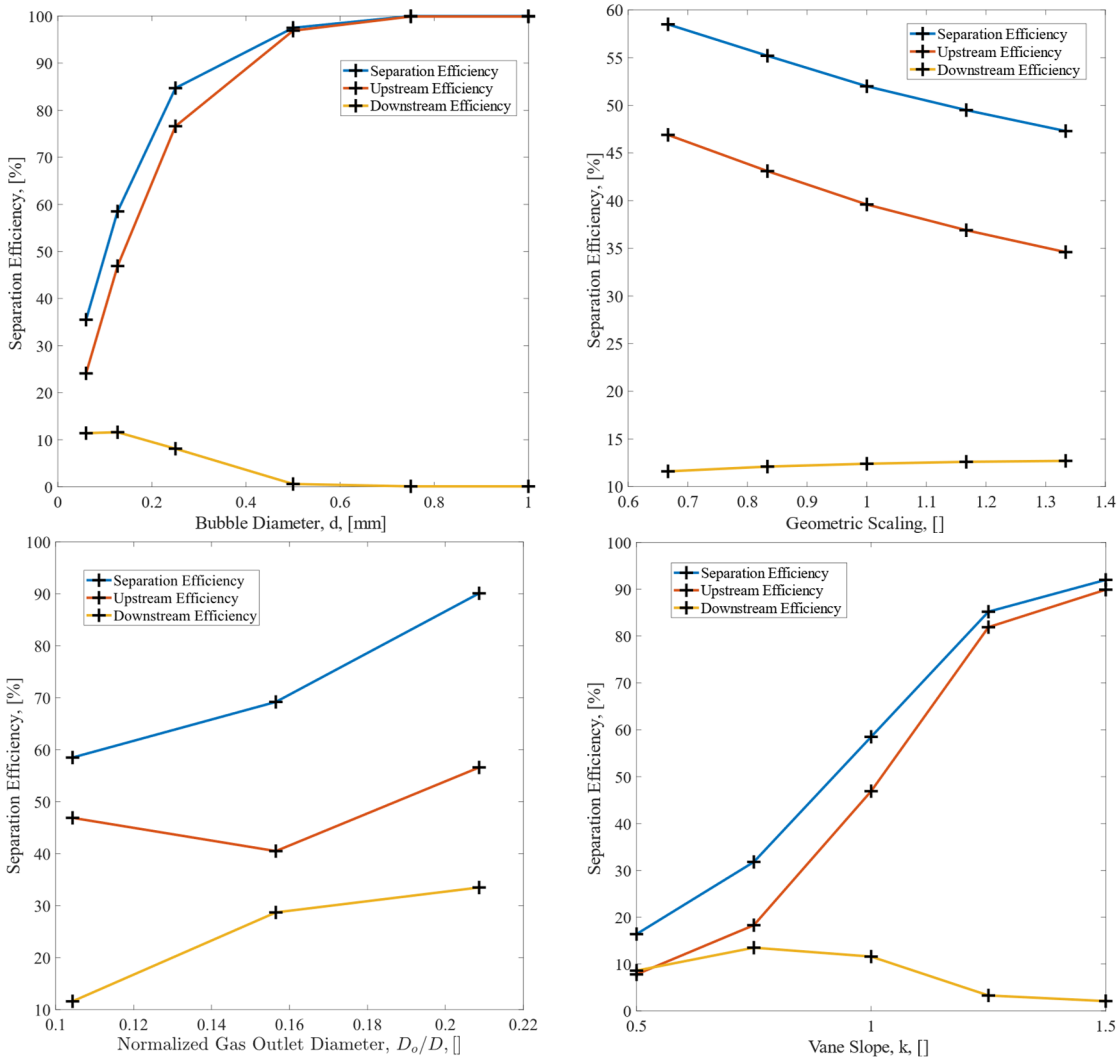

Figure 2.66. Effect of different design parameters on separator performance.

Figure 2.66 shows the effect of different design parameters. The separation efficiency is divided into the upstream efficiency and downstream efficiency, describing the flow through the different outlets individually. The reason for dividing the separation efficiency is the very different behavior of the upstream and downstream removal. The separation efficiency increases quickly with the bubble diameter. For bubble diameters larger than 0.5mm, the separation efficiency is almost 100% at the reference condition. As the upstream efficiency approaches 100%, the downstream efficiency decreases to zero. However, larger bubbles are not preferred because they reduce the available interfacial area for xenon removal. The choice of bubble diameter should be an optimization of removal efficiency and separation efficiency. By increasing the length scale of the separator, the removal efficiency upstream almost linearly decreases. On the other hand, the downstream efficiency slightly increases because of the reduced upstream efficiency. The separation efficiency increases linearly with the gas outlet diameter because of the void fraction profile. However, the entrained liquid flow increases with the outlet diameter and should be minimized. The slope of the vane controls the rotational movement in the separator. As the

tangential velocity increases, the pressure gradient that pushes the bubbles towards the center becomes stronger, and therefore the separation efficiency is increased.

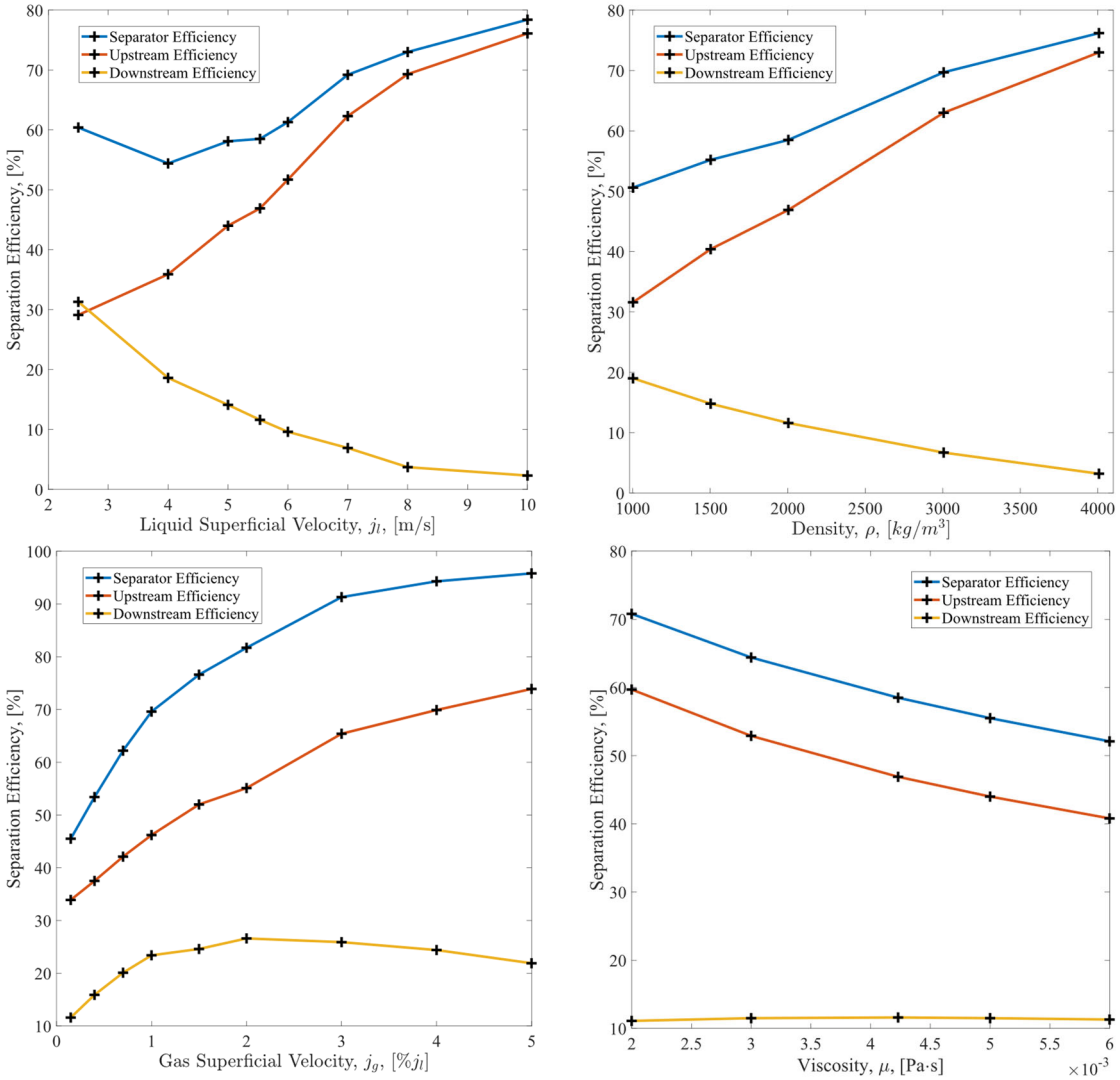

Figure 2.67. Effect of different operational parameters on separator performance.

In Figure 2.67, the effect of operational parameters is presented. As the liquid superficial velocity increases, the upstream efficiency becomes higher because of the increased centrifugal forces. The downstream efficiency, on the other hand, is reduced. As the temperature of the salt increases, the density of the salt decreases along with the viscosity. The net effect is a slight increase in the separation efficiency. The last operational parameter is the gas superficial velocity. In the range of inlet void fraction considered in the simulation, the separation efficiency increases with the gas flowrate, though the slope becomes lower and lower as the inlet void fraction increases. It is expected that at certain inlet void fraction, the amount of gas exceeds the capacity of the separator, and the efficiency starts to reduce.

In this section, the effect of different parameters on the separation efficiency is discussed, though the physical explanation of the different behaviors of the upstream and downstream

efficiency is not given. To understand the underlining physics behind the sensitivity study, a simplified model based on physical reasoning is required. In the next section, a regression model is developed based on observations of the sensitivity study and additional physical assumptions.

2.6.5. Regression Model

As discussed earlier, due to the computational cost of performing CFD analysis, either a reduced order model or regression model should be constructed for use in the system analysis code and the design optimization process. There are several important criteria of the bubble separator, among which the separation efficiency is the most vital. For large bubble diameter conditions, the separation efficiency could easily reach almost 100% [29, 114], and make the separation efficiency an irrelevant quantity. However, these conditions are not favorable for the xenon removal process. To limit the overall system size and thus the cost of the xenon removal system, a small bubble diameter is preferable, even though it would reduce the bubble separation efficiency. To understand the impact of different parameters on the bubble separation, sensitivity studies are carried out. To further facilitate the design optimization process and for the system analysis aspect, a regression model for bubble separation efficiency should be constructed.

From the sensitivity study and the experimental data, several observations are made on the behavior of the bubble separation efficiency. These observations pose certain requirement on the format of the regression model. These observations are summarized as,

- Separation efficiency could be divided into upstream and downstream efficiency, each has different trend with respect to system parameters.
- When the upstream efficiency reaches 100%, the downstream efficiency would be 0% since no bubbles are left to be collected.
- Due to the velocity field at the pipe center, where both concurrent and countercurrent flow exist, the downstream efficiency also slightly influences the upstream efficiency.
- The upstream efficiency increases with density and liquid velocity, decreasing with viscosity.
- The downstream efficiency decreases with liquid inertia, which could be the related to the total pressure difference in the system.
- The separation efficiency increases with void fraction.
- The separation efficiency increases with gas outlet dimeter, and when the outlet diameter is relatively large, the efficiency would approach 100%.
- Most of the bubbles are removed from the upstream outlet.
- With these observations, some further assumptions are made to prescribe the format of the regression model.
- There exists a nominal void center diameter D_{void} in the separator. Its relative magnitude with the gas outlet diameter is the deterministic factor of the separation efficiency, which takes the form of,

$$\frac{D_0}{D_o + D_{void}}. \tag{2-50}$$

- Most of the parameters only appears in the correlation of D_{void}. Considering the sensitivity study, the format is assumed to have the shape,

$$D_{void} = \frac{c_1 D \left(\frac{D}{d_b}\right)^{c_2}}{(1 + k^{c_3} Re_D)(100\alpha)^{c_4}}. \tag{2-51}$$

- The downstream efficiency decreases to 0% when upstream reaches 100%, therefore an additional modifier exists,

$$\left(1 - \frac{D_0}{D_o + D_{void}}\right). \tag{2-52}$$

- The downstream efficiency is controlled by liquid inertia and outlet pressure. Therefore, a front factor also exists,

$$\hat{\eta} = \left(\frac{\rho j_l^2 D_c^2}{D_o^2 \Delta p} b_1 + 1\right)^{-1}. \tag{2-53}$$

- The upstream efficiency is slightly influenced by the downstream efficiency, through the front factor of $\hat{\eta}$, with a format,

$$\frac{1}{1 + a_2\, \hat{\eta}}. \tag{2-54}$$

Using these assumptions, the final expected form of the regression model is expressed as,

$$\hat{\epsilon} = \frac{1}{1 + a_2\, \hat{\eta}} \left(\frac{D_0}{D_o + D_{void}}\right) + a_1\, \hat{\eta} \left(1 - \frac{D_0}{D_o + D_{void}}\right) \left(\frac{D_0}{D_o + D_{void}}\right),$$

$$\hat{\eta} = \left(\frac{\rho j_l^2 D_c^2}{D_o^2 \Delta p} b_1 + 1\right)^{-1}, \qquad D_{void} = \frac{c_1 D \left(\frac{D}{d_b}\right)^{c_2}}{(1 + k^{c_3} Re_D)(100\alpha)^{c_4}}. \tag{2-55}$$

With the format constructed, the next step is to determine the coefficients in the regression model. Seven coefficients exist in the model, and the expression does not possess a linear form. This makes it difficult to determine these coefficients through direct optimization using least square fitting. To solve this problem, the features of the expected format are used. The efficiency could be divided into upstream and downstream efficiency as,

$$\hat{\epsilon}_u = \frac{1}{1 + a_2\, \hat{\eta}} \left(\frac{D_0}{D_o + D_{void}}\right),$$

$$\hat{\epsilon}_d = a_1\, \hat{\eta} \left(1 - \frac{D_0}{D_o + D_{void}}\right) \left(\frac{D_0}{D_o + D_{void}}\right). \tag{2-56}$$

If a_1 and a_2 are prescribed, then the unknowns in the equations are D_{void} and $\hat{\eta}$, which could be solved numerically. Then $\frac{1}{\hat{\eta}}$ is plotted against $\frac{\rho j_l^2 D_c^2}{D_o^2 \Delta p}$, which is expected to have a linear trend from the expected format of $\hat{\eta}$,

$$\frac{1}{\hat{\eta}} = \frac{\rho j_l^2 D_c^2}{D_o^2 \Delta p} b_1 + 1. \tag{2-57}$$

Linear regression is performed for the above equation, where the coefficient b_1 and the goodness of the fitting are recorded. The optimized coefficients a_1, a_2 and b_1 are then determined based on the goodness of the fitting. When this process is performed with the simulation data, an additional filtering is carried out to remove any data points that have a solved $\hat{\eta}$ smaller than 0.1, which are not well behaved for the modeling of $\hat{\eta}$ given the negligible downstream removal efficiency for those cases. The result of this optimization is shown in Figure 2.68.

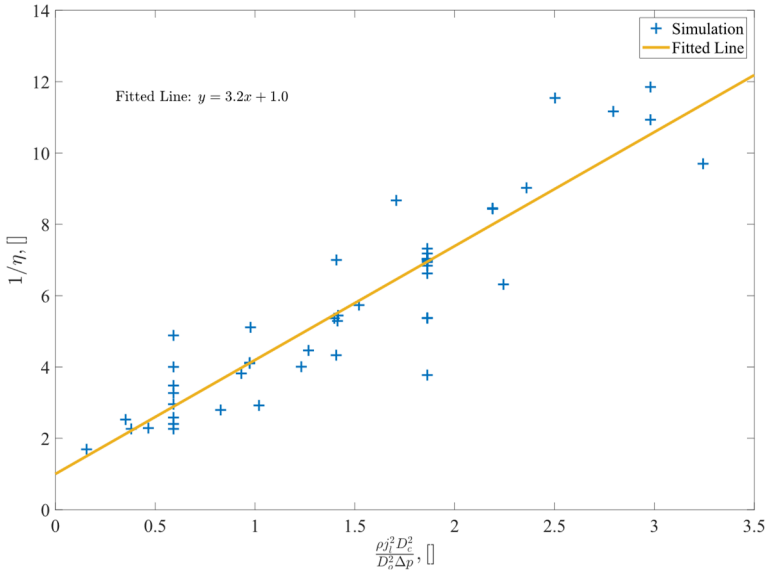

Figure 2.68. Regression of $1/\hat{\eta}$ and $\rho j_l^2 D_c^2 / D_o^2 \Delta p$.

The final values of these coefficients are,

$$a_1 = 3.26, \qquad a_2 = 0.23, \qquad b_1 = 3.20. \tag{2-58}$$

With a_1, a_2 and b_1 determined, the next step is to finalize the coefficients for D_o. Here, instead of using the solved D_o along with the solved $\hat{\eta}$, the correlation of $\hat{\eta}$ is plugged into the original equation for the overall separation efficiency,

$$\hat{\epsilon} = \frac{1}{1 + 0.23\,\hat{\eta}} \left(\frac{D_0}{D_o + D_{void}} \right) + 3.26\,\hat{\eta} \left(1 - \frac{D_0}{D_o + D_{void}} \right) \left(\frac{D_0}{D_o + D_{void}} \right),$$

$$\hat{\eta} = \left(\frac{\rho j_l^2 D_c^2}{D_o^2 \Delta p} 3.2 + 1 \right)^{-1}. \tag{2-59}$$

Using the simulation results for the separation efficiency, the only unknown in this equation is D_{void}, which can be easily solved. The reason of this roundabout is to reduce the error generated from prescribing the format of $\hat{\eta}$, while the minor influence from other parameters is not considered. The expected format of D_{void} is,

$$D_{void} = \frac{c_1 D \left(\frac{D}{d_b}\right)^{c_2}}{(1 + k^{c_3} Re_D)(100\alpha)^{c_4}}. \tag{2-60}$$

Four coefficients exist in the expression for D_{void}. Similarly, instead of performing a direct optimization, a stepped approach is used. First, c_3 and c_4 are assigned to 1, and optimization is performed for c_1 and c_2. The result is shown in Figure 2.69(a), where the dashed lines are the ±30% error lines and the 1:1 line between the simulation and the model prediction. Cases considered to be outliers are marked in Figure 2.69(a). Most of these outliers are from the simulations for the sensitivity study of k, the slope of the rotational vanes. This indicates the prescribed value for c_3 is not appropriate. By adjusting the value of c_3 to different integers, it is determined that $c_3 = 4$. With this modification, the modified prediction is shown in Figure 2.69(b). Similarly, the outliers are marked, and by inspecting these cases, they are associated with the sensitivity study of void fraction. This means the order of magnitude on the void fraction denominator c_4 should be reduced. After changing c_4 to 0.5, the results are further improved as shown in Figure 2.69(c). With these adjustments, the optimization is performed again, and the results are shown in Figure 2.69(d), where the error lines are changed to ±15%. From the results, the prediction of the regression equation is mostly within $\pm 15\%$ for the conditions considered in the simulation, and the trends of different parameters are mostly captured. The final form of the regression equation is,

$$\hat{\epsilon} = \min\left[\frac{1}{1 + 0.23\,\hat{\eta}}\left(\frac{D_0}{D_o + D_{void}}\right) + 3.26\,\hat{\eta}\left(1 - \frac{D_0}{D_o + D_{void}}\right)\left(\frac{D_0}{D_o + D_{void}}\right), 1\right].$$

$$\hat{\eta} = \left(\frac{\rho j_l^2 D_c^2}{D_o^2 \Delta p}\,3.2 + 1\right)^{-1}, \qquad D_{void} = \frac{4.89 D \left(\frac{D}{d_b}\right)^{1.27}}{(1 + k^4 Re_D)\sqrt{100\alpha}}. \tag{2-61}$$

The minimizing function is included for the cases with large bubble diameter where the prediction of the regression equation could be slightly large than 1. The range of conditions used to construct the regression modeling terms of MSBR system parameters are summarized in

Table 2.14.

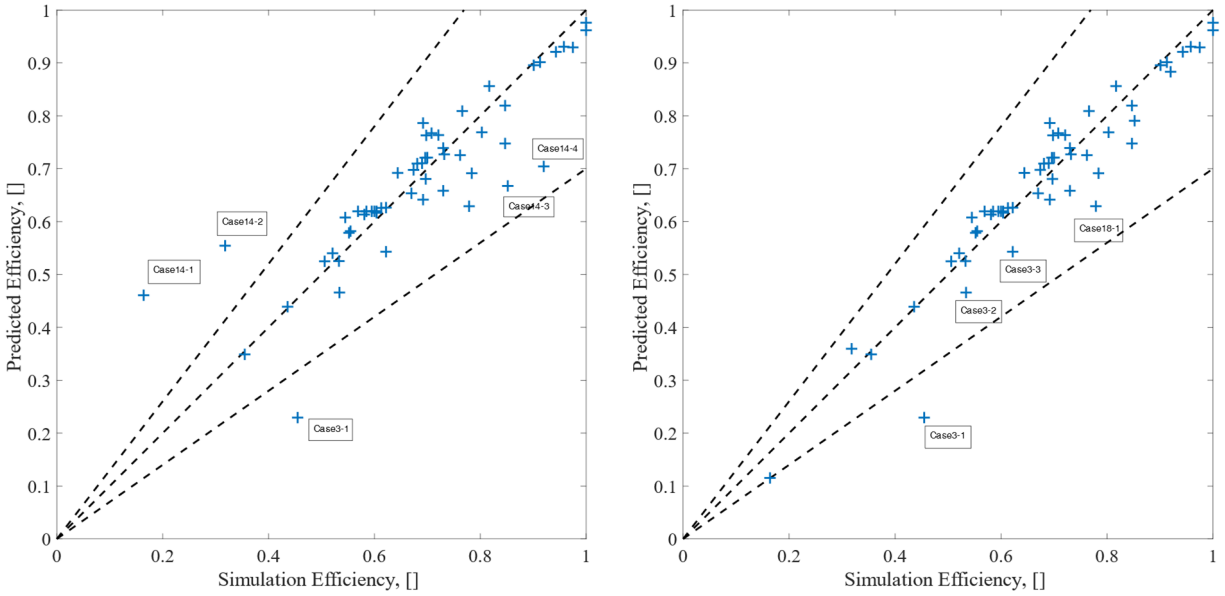

(a) Initial Test.

(b) Resolving Vane Slope.

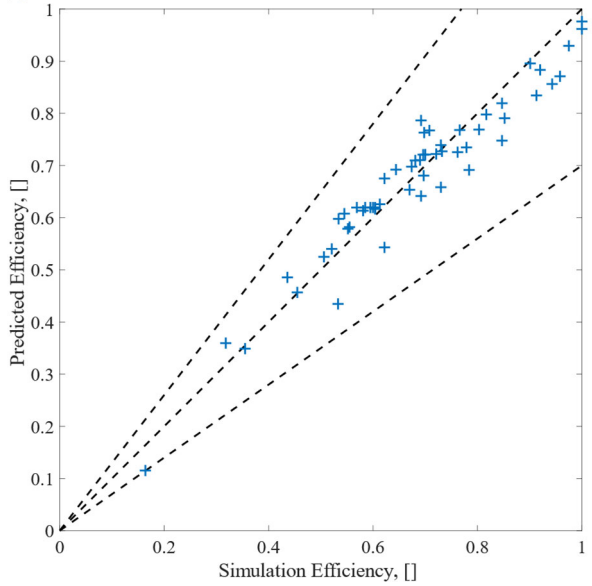

(c) Resolving Void Fraction.

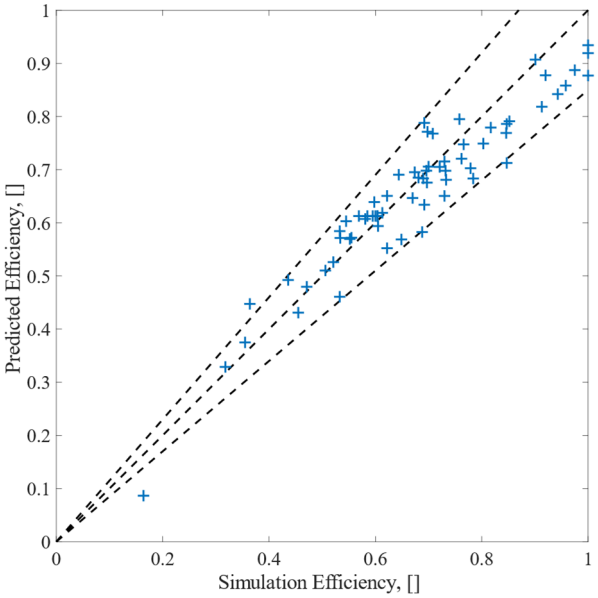

(d) Final Result.

Figure 2.69. Steps for determining the coefficients of D_{void}.

Table 2.14. Range of conditions covered in the analysis

Parameter	Range
Flowrate	2% to 23% of MSBR loop flow rate
Viscosity	850 K to 1100 K*
Density	850 K to 1100 K*
Void Fraction	0.15% to 5%
Bubble Diameter	0.06 mm to 1 mm

* The temperature is according to the properties of FLiNaK. The actual range of properties used in the simulations are beyond this range. The temperature only indicates that the operational temperature is covered by the simulations.

To evaluate the how the regression equation could be extended to the conditions that are not used to construct the model, additional simulations with salt flowrates from 23% to 64% of MSBR loop flow rate are carried out. The results are shown in Figure 2.70, where the new cases are colored by orange.

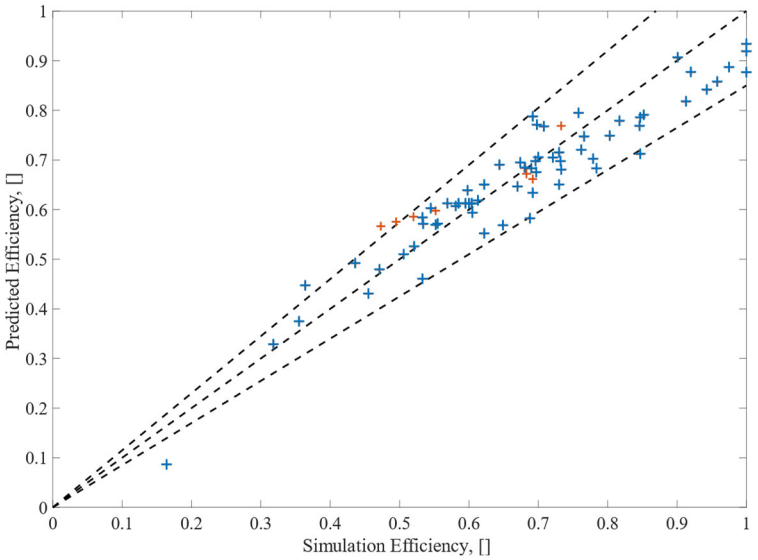

Figure 2.70. Evaluation of the regression model with independent simulation cases. The red markers are additional simulation cases which are not used in determining the coefficients.

From the figure, the predictions of the regression model are still close to the simulation, mostly with errors less than 15%. From these results, it is concluded that the regression equation has captured the general behavior of the bubble separator and could be used to estimate the performance of the separator under different conditions.

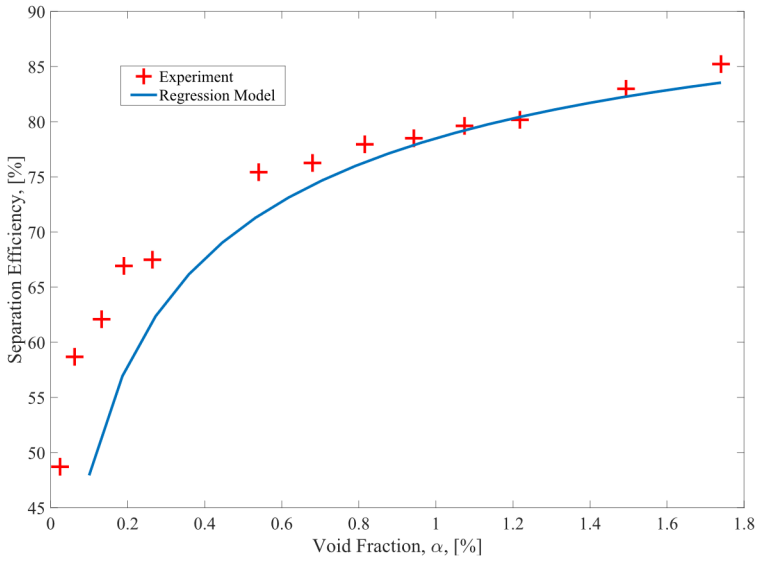

Figure 2.71. Evaluation of the regression model with experimental data.

In Figure 2.71, the regression model is also compared with the experimental data for small bubbles shown in Figure 2.59(b). Good agreement is found between the experimental data and the predicted results. Underprediction is found for small void fraction cases, but for void fraction over 1%, the error is within a few percent.

There are some restrictions of the current regression model that should be noted. The gas outlet pressure Δp in the simulations, defined as the pressure difference between the salt outlet and the gas outlet, is set above the point where reserved flow and possible flow instabilities could happen. The variation of the pressure difference in the simulation cases is within a small range around the reported pressure head given in literature and has not been extensively tested. Physically speaking, when the gas outlet pressure is low enough, a stable void center could not be retained, and the separation efficiency would decrease with the pressure. Since this regime is undesirable in the operation of the separator, prediction of the separation efficiency in this regime is less interesting, and not considered in the current regression model. It is expected that overpredictions would happen when the model is extended to low pressure cases that are beyond the simulation range of conditions. To account for this, a correction factor related to pressure and liquid inertia might be added to the front of the current formula.

2.7. Lessons Learned

- Grid independence study should always be performed before performing sensitivity study for models. Some of the initial sensitivity studies for CFD simulations are performed before the grid independence study to determine appropriate models. Though the conclusions from these sensitivities study hold true, but the simulations are repeated for the final mesh to ensure the results are all consistent.
- Parametric control and post-processing method should be mature before performing massive simulation cases. Some of the early simulation cases don't use parametric control and have to be somewhat manually handled to extract the required data in a case-by-case manner.

2.8. Summary, Conclusions, and Future Work

A CFD model is proposed for simulation of the xenon removal process in molten salt with inert gas sparging. In order to predict the mass transfer rate, the velocity and phase distribution of each phase are needed. Therefore, the Eulerian two-fluid model coupled with species transfer is used in the simulation to capture the local velocity, phase distribution and species concentration at the same time.

For accurate prediction of the process, a careful evaluation of the constitutive relations and material properties is required. Therefore, the interfacial forces, turbulence model, material properties, mesh size, bubble diameter and boundary conditions are first evaluated through sensitivity study in a cylindrical bubble column which is covered in a separate report [64]. The geometry is chosen according to the molten salt sparging experiment described in Appendix A.

In order to validate the CFD model, three sets of experiment are done in an air-water bubble column with the same geometry used in the simulation. Since the specific interaction model is scaled, the model validated with air-water experiment is applicable to molten salt in principle. The

void fraction profile is measured with a visualization experiment, the velocity profile is obtained with Particle Image Velocimetry and the mass transfer coefficient is measured by tracking the dissolved oxygen concentration. Through the validation, it is concluded that the CFD model could satisfactorily predict the volume averaged void fraction. The local void fraction and velocity profile a few centimeters away from the inlet could be predicted, though the agreement immediately near the inlet is limited. The mass transfer coefficient from the simulation is compared with the experimental data. It is concluded that with calibration from the experiment, the mass transfer could be predicted. However, adjustment of the coefficient based on flow conditions is still required because of the difference in turbulence characteristics.

With the CFD model established and validated against various cases, it is concluded that the model can reasonably capture the underlining physics of the two-phase mass transfer process, and therefore as an engineering tool for designing a commercial scale xenon removal system. In this remaining chapter, the prototypical xenon removal system is studied with the validated CFD model and reduced order analysis. The system is divided into three components – the bubble generator, the pipeline for xenon removal and the bubble separator. The bubble generator determines the bubble diameter that is present in the pipeline and bubble separator, which is not modeled explicitly. Rather, the bubble diameter is treated as a design parameter that is provided to the analysis. The pipeline is where the dissolved xenon migrates to the gas phase. CFD simulation of the migration process is performed, from which a regression equation of the mass transfer coefficient is constructed and compared with literature data. Based on the simulation results, a 1D reduced order model is derived from area averaging of the 3D governing equation. The distributional effect of different flow variables is assessed through covariances. This 1D reduced order model is then applied to the actual molten salt system to assess the sensitivity of different parameters. The effect of circulating bubbles is analyzed with a conclusion that they are beneficial when the portion of diverted flow into the xenon removal system is low. Apart from the pipeline removal, the bubble separator is also studied with CFD simulations. A validation study is made with two sets of experimental data in the literature. It is concluded that the CFD simulation, both 3D and 2D, could satisfactorily predict the separation efficiency at the bubble separator, though the local pressure and velocity profile are less accurate. Based on the sensitivity study of the separator, a regression model for the separation efficiency is constructed using physical reasoning. The regression model agrees well with the simulation and experimental separation efficiency in the range of conditions considered in this chapter.

There are several limitations of the current study. The CFD model is only rigorously validated against the air-water experiments, with known properties of oxygen solubility and diffusivity. Though theoretical estimation for xenon solubility and empirical value for diffusivity in molten salt are reported and used in the current work, the accuracy of these values is still questionable. Moreover, the models validated with air-water experiments may still need to be evaluated with experiment in molten salt. Though the models are written in dimensionless form, its dependence on the material properties may not be sufficiently validated, since most of the validations are carried out in air-water or steam-water flow. Another potential issue is that the flow in a commercial scale xenon removal system is much more complicated than the flow presented in the validation experiment. The model validation may not be completely applicable in various conditions of the removal system. In fact, when the model is applied to the actual removal system, modifications described in the corresponding sections are made to accommodate the difference in flow conditions. Though comparisons are also made with experimental data relevant to the

prototypical removal system, the targeted quantities are averaged system performance rather than the detailed flow fields.

The CFD model does not consider the compressibility of the gas phase, yet the pressure difference in the bubble separator is enough to change the gas volume by 50%. The regression model of the bubble separator has not been tested against a wide range of pressure, and overestimation is expected at lower pressure difference. For the simulation of xenon removal, the diffusivity and solubility of xenon are only estimated value with large uncertainties. Though the results present in this chapter could enhance the understanding of the xenon removal system, the available experimental knowledge is not enough for the model to be used in designing an actual xenon removal system.

With the discussion above, various experiments in molten salt are favorable. A carefully designed xenon experiment is needed to obtain the solubility and diffusivity of xenon in molten salt, which are directly related to the mass transfer process. Without concrete measurement of these properties, the uncertainty of simulation would be relatively large. Moreover, experiments of the bubble separator with quantified uncertainties are also needed for proper validation of the CFD model in the removal system.

The simulation could also be extended to incorporate compressibility effects that are important in the prototypical system. The sensitivity of boundary conditions and constitutive models might be investigated again in the prototypical system again considering the difference in flow conditions.

CHAPTER 3: FUEL CYCLE SIMULATION AND ANALYSIS

In this chapter, fuel cycle simulation is performed to analyze the effect of fission product removal in TAP MSR and MSBR. The present work introduces the open-source reprocessing simulation package, SaltProc [123], which couples with the continuous-energy Monte Carlo depletion calculation code, Serpent 2 [124], for fuel composition dynamics analysis in various MSRs taking into account a realistic, physics-driven model of an online fuel reprocessing system.

3.1. Introduction to Fuel Burnup and Online Reprocessing Simulations in MSRs

All liquid-fueled MSR designs involve various levels of online fuel processing. Minimally, noble gaseous fission products (e.g., Kr, Xe) escape from the fuel salt during routine reactor operation and must be captured. Other systems might be used to enhance the removal of those elements. Most designs also call for the removal of rare earth metals from the core since these metals act as neutron poisons. Some designs suggest a more elaborate list of elements to process (Figure 3.1), including the temporary removal of protactinium from the salt or other regulation of the actinide inventory in the fuel salt [125]. Fresh fuel salt with dissolved fissile and/or fertile material (e.g., ^{233}U, ^{232}Th, low-enriched uranium (LEU), a transuranic vector from LWR spent nuclear fuel (SNF)) make up the salt mass loss caused by poison removal and conserves the total mass in the primary loop.

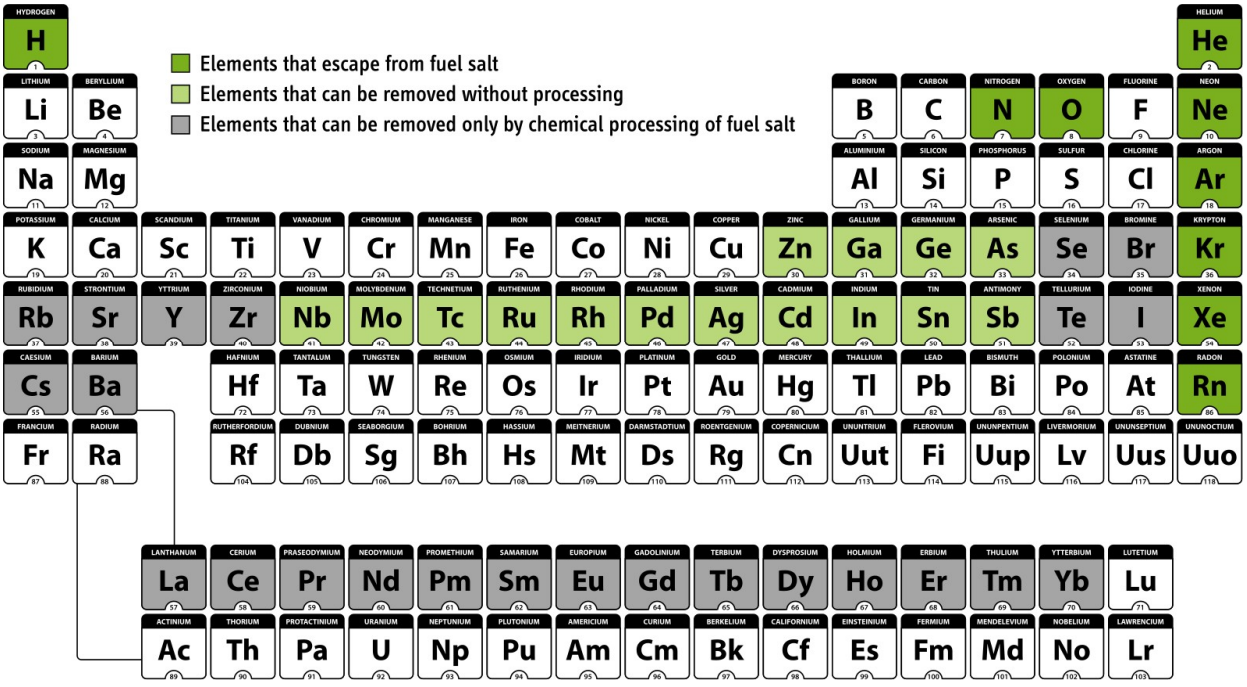

Figure 3.1: Processing options for MSR fuels (reproduced from Ahmed et al. [125]).

Most liquid-fueled nuclear reactor concepts adopt continuous separations and feeds: the core material is circulated to or from the core at all times (continuously) or specific intervals (batch-wise). In contrast, in a solid-fueled reactor, fission products and actinides remain within the initial fuel material throughout its time in the core.

The ability to perform online fuel salt reprocessing improves the potential neutronics performance of liquid-fueled reactors. First, liquid-fueled reactors can operate with relatively low excess reactivity because fissile material can be continuously added to the core. Second, continuously removing fission products, including strong absorbers (poisons), can significantly improve fuel utilization and decrease parasitic neutron absorption. Third, online reprocessing decreases the amount of decay heat, dissipating after shutdown. Finally, for a breeder3 excess of fissile material might be continuously extracted from the core and used to startup new reactors. Nevertheless, the removal of each element from the liquid fuel salt presents a unique challenge in terms of chemical separation, storage, and disposal of the separated materials.

Contemporary nuclear fuel depletion software lacks continuous fuel salt reprocessing modeling. To handle material flows in potential online removal and feed of liquid-fueled systems, early MSR simulation methods at Oak Ridge National Laboratory (ORNL) integrated neutronics and fuel cycle codes (i.e., Reactor Optimum Design (ROD) [126]) into operational plant tools (i.e., Multiregion Processing Plant (MRPP) [127]) for MSR fuel reprocessing system de- sign. Extensive research in fast and thermal MSR analysis has yielded specialized tools for burnup calculations in liquid-fueled nuclear systems [128-134]. Table 3.1 presents a list of recent efforts, along with the main features of the employed methods and software.

Two main online reprocessing simulation approaches have been demonstrated in the literature: batch-wise and continuous. In the batch-wise approach, the burnup simulation stops at a given time and restarts with a new liquid fuel composition (after removal of discarded materials and addition of fissile/fertile materials).

ORNL researchers have developed ChemTriton, a Python script for SCALE/TRITON, which employs the batch-wise approach to simulate a continuous reprocessing and refill for either single or multiple fluid designs. ChemTriton models salt treatment, separations, discharge, and refill using SCALE/TRITON depletion simulation over small time steps to simulate continuous reprocessing and deplete the fuel salt [135, 136].

Table 3.1: Tools and methods for liquid-fueled MSR fuel salt depletion analysis.

	Nuttin et al., 2005 [136]	Aufiero et al., 2013 [130]	Betzler et al., 2018 [135]	Present work
Neutronics software	MCNP REM	Serpent 2	SCALE6.2 ORIGEN-S	Serpent 2
	stochastic	stochastic	deterministic	stochastic
Geometry	unit cell	full-core 3D	unit cell	full-core 3D
Removal/feed	continuous	continuous	batch-wise	batch-wise
Separation efficiency	fixed, must be defined by user before simulation			function of many parameters
Fuel reprocessing plant	single component, “black” box model			realistic multi-component model
Reactivity control	continuous adjustment of fissile material injection		batch injection of fissile material	periodical adjustment of geometry and fissile material injection
Safety parameters evolution	thermal feedback	not considered	thermal feedback	thermal feedback, void reactivity coefficient, control rod worth

In the continuous approach, accounting for removal or addition of material presents a greater challenge since it requires adding a term to the Bateman equations. Both ORIGEN [137] and the Serpent burnup routine [138] solves a set of the Bateman equations using one- group averaged flux and transmutation cross sections obtained from a transport calculation. The Bateman equations describe the rate of change of each isotope, i, due to neutron induced reactions and decay processes [139]:

$$\frac{dN_i}{dt} = \sum_{m=1}^{M} l_{im}\lambda_m N_m + \phi \sum_{m=1}^{M} f_{im}\sigma_m N_m - (\lambda_i + \phi\sigma_i + r_i - f_i)N_i + F_i \mid i \in [1, M] \qquad (3\text{-}1)$$

(1) (2) (3) (4) (5) (6)

where

N_i = number density of nuclide *i* $[cm^{-3}]$

M = number of nuclides [−]

l_{im} = fraction of decays of nuclide *m* that result in formation of nuclide *i* [−]

λ_i = radioactive decay constant of nuclide *i* $[s^{-1}]$

ϕ = neutron flux, averaged over position and energy $[cm^{-2} \cdot s^{-1}]$

f_{im} = fraction of neutron absorption by nuclide *m* leading to the formation of nuclide *i* [−]

σ_m = average neutron absorption cross section of nuclide m $[cm^2]$ ri = continuous removal rate of nuclide i from the system $[s^{-1}]$

f_i = continuous feed rate of nuclide *i* $[s^{-1}]$

F_i = production rate of nuclide *i* directly from fission $[cm^{-3} \cdot s^{-1}]$.

The terms on the right-hand side of the equation represent:

(1) production of species i as a result of the decay of all the nuclides present;

(2) production of species i as a result of neutron capture by all nuclides present;

(3) loss of nuclide i through its own decay;

(4) loss of nuclide i as a result of neutron capture;

(5) loss of nuclide i through continuous removal from the system;

(6) gain of nuclide i as a result of continuous feed to the system.

Nuttin et al. developed an in-house depletion code called Rules for Evolution calculations with MCNP (REM), which directly couples with MCNP [140] to simulate fuel salt material evolution in a simplified MSBR-like liquid-fueled system. That work directly integrated the Bateman differential equations using neutron flux from MCNP, tracking all the isotopes available in the data library, and controlling reactivity to maintain reactor criticality [141].

In a similar vein, Aufiero et al. extended Serpent 2 for continuous reprocessing simulations by adding an explicit pseudo-decay term representing fission product removal (-$N_i r_i$ term in

Equation 3-1) for each target poisonous nuclide [130]. The developed extension directly accounts for the effects of online fuel reprocessing on depletion calculations and features a reactivity control algorithm. The extended version of Serpent 2 was assessed against a dedicated version of the deterministic ERANOS-based EQL3D procedure in [128] and applied to analyze the MSFR fuel salt isotopic evolution.

More recently, Betzler et al. added to SCALE/TRITON continuous removals capability for depletion simulation [134]. Similar to Aufiero et al. this extended SCALE/TRITON directly adds feed and removal terms in the burnup matrix and solves it using existing ORIGEN capabilities. TRITON's continuous reprocessing capability was validated against the batch-wise script ChemTriton for single-channel Molten Salt Reactor Experiment (MSRE)-like model. Unlike ChemTriton, this new capability will be available for all SCALE users in the 6.3 release. However, at the moment, it is undergoing extensive testing and validation procedures and unavailable for external users.

Some of the tools listed in Table 3.1 used significant approximations that may lead to inaccurate fuel evolution predictions and others unavailable for external users. This work introduces an open-source simulation package, SaltProc, which expands the capability of the continuous-energy Monte Carlo Burnup calculation code, Serpent 2, for depletion calculations of liquid-fueled MSRs.

Most of the existing tools in the literature represented the fuel salt reprocessing plant as an invariable "black box" model, which removes target elements all at once with a fixed efficiency, determined by the user before starting the depletion simulation. Typically, such a "black box" model is characterized by a vector of removing elements and their extraction efficiencies:

$$\begin{bmatrix} N_0^b \\ \vdots \\ N_e^b \\ \vdots \\ N_E^b \end{bmatrix} \times \begin{bmatrix} \epsilon_0 \\ \vdots \\ \epsilon_e \\ \vdots \\ \epsilon_E \end{bmatrix} = \begin{bmatrix} N_0^a \\ \vdots \\ N_e^a \\ \vdots \\ N_E^a \end{bmatrix} \tag{3-2}$$

where

N^b = number density vector before reprocessing [cm^{-3}]

N^a = number density vector after reprocessing [cm^{-3}]

ϵ = extraction efficiency [−] vector for all elements e in (0, E).

The main issues related to static "black box" model assumptions in the literature neglect:

Time varying extraction. Realistically, long-term reactor operation will require a time dependent extraction efficiency vector. The current tools treat separation efficiency as constant.

The impact of operational parameters on separation efficiency. In reality, the extraction efficiency depends on temperature, power level, current fuel salt isotopic composition, and material mass flow rate. Gas solubility in the salt is inversely proportional to the salt temperature; hence, the extraction efficiency expected to be lower for the higher temperature of the salt.

Discrete component performance and dynamics in the multi-component system. All reprocessing plant components are treated as a single "black box" component in existing simulation tools. However, the fuel salt in a reprocessing plant undergoes many separate components (e.g., helium bubbling, nickel mesh filter, etc.) that target specific elements. Some of these components can be connected in series, parallel, or series-parallel. The "black box" model (only single process) requires extensive pre-simulation analytic work from the user to calculate the lumped separation efficiency vector before a simulation is run and cannot be adjusted during the simulation. Additionally, treating the processing system as a single "black box" neglects dynamics related to relative component flow rates. Finally, the discrete waste streams from each component are not tracked separately in "black box" tools. However, this information is necessary for fuel reprocessing system optimization.

Based on the review of existing tools, a few possible directions are identified for the improvement of MSR tools:

Reproducibility/availability. Serpent is the only contemporary nuclear reactor physics software that can perform depletion calculations that can take into account online fuel salt reprocessing regimes. However, this built-in online reprocessing routine is undocumented: the discussion forum for Serpent users is the only useful source of information at the moment. Other mentioned tools are available for internal users only. These issues can be a barrier to reuse research software and to reproduce scientific results. Thus, a new, open-source, reproducible tool for fuel processing simulation would assist in the production of reproducible research in the area of liquid-fueled reactor modeling.

Realistic fuel reprocessing system model. Significant approximations in fuel reprocessing parameters deteriorate fuel salt composition predictions since the evolution of safety parameter accuracy is strongly dependent on fuel salt composition. A realistic fuel reprocessing system model will allow reprocessing component parameter optimization, increase the fidelity of fuel and waste stream composition calculations, and advance reprocessing system design.

Variable extraction efficiency. Most research efforts in the literature (except Nuttin et al.7) assume ideal 100% extraction efficiency of all removed elements, which stayed constant during the whole reactor lifetime. Realistically the efficiency is time-dependent and changes with respect to operational parameters: temperature, power level, salt composition, etc. Thus, the ability to set up dynamic separation efficiency must be added in MSR simulation tools to advance depletion calculations.

Reactivity control. Reconfigurable moderator configuration in the TAP core presents a challenge because of the core geometry changes with time. The reactivity control module, which adjusts the core geometry to maintain criticality, is an exceptional capability for simulating new, more advanced MSR concepts and short-term transients.

Safety characteristics evolution during reactor operation. The MSR fuel salt accumulates FPs and transuranic elements, which significantly shift the neutron energy spectrum. This spectrum shift might worsen the core safety during operation. The impact of the fuel salt evolution on the MSR safety parameters must be carefully investigated and reported.

The main goal of the current work is to develop a generic, open-source tool, SaltProc, capable of simulating a wide range of liquid-fueled systems, including multi-fluid and multi-region

designs, and validate it against existing modeling efforts. Additionally, SaltProc enables poison extraction simulation based on a realistic physics-based fuel processing model.

The structure of this chapter is as follows. Section 3.2 details online reprocessing modeling and the proposed computation tool architecture. In an attempt to avoid the pitfalls of a "black box" understanding and to identify method limitations at an early stage, governing equations and working principles are stated and discussed. Section 3.3 covers SaltProc demonstration and validation efforts with a focus on the TAP MSR, taking into account adjustable moderator configuration. In addition, Section 3.3 gives the safety parameter overview and its evolution during the TAP lifetime-long reactor operation. SaltProc demonstration for short-term depletion calculations and evaluation of load-following potential of the TAP MSR and the MSBR are presented in Section 3.4 and 3.5. Moreover, the safety parameters dynamics during short-term transients have been evaluated at the end of Section 3.4 and 3.5. The final section summarizes this work's contribution to the nuclear community, and a conclusion is offered together with an outlook for future work on the topic.

3.2. Online reprocessing Modeling Approach

In this section, the modeling strategy of the fuel processing system is presented. First an overview on the design of the system is given, identifying the employed processing approach and related engineering models. Then the software design of the SaltProc code is presented.

3.2.1. Fuel Salt Reprocessing Overview

Removing specific chemical elements from a molten salt is a complicated task that requires intentional design (e.g., chemical separations equipment design, fuel salt flows to equipment). This section contains a brief overview of a generic MSR fuel salt reprocessing system; modeling such systems is the focus of the current dissertation.

3.2.1.1. Gas separation system

Gaseous fission products (e.g., Xe) must be removed from the fuel salt to avoid reactor poisoning, especially during startup and power maneuvering. This is particularly true for ^{135}Xe, with its strong neutron capture cross section ($\approx 10^6 \sim 10^7\ b$ in a thermal energy range). ^{135}Xe is produced directly from fission in about 0.2% of ^{235}U fissions ($\gamma_{135_{Xe}}$), but an even larger fraction of ^{135}Xe is produced by the decay of ^{135}I and ^{135}Te (Table 3.2). ^{135}I and ^{135}Te yields from fission are $\gamma_{135_I} = 3.1\%$ and $\gamma_{135_{Te}} = 3.3\%$, respectively. Thus, the total ^{135}Xe production from fission is about 6.6% of fissions (of ^{235}U), most of this is from ^{135}I and ^{135}Te decay. Noble gases (e.g., tritium, xenon, and krypton) can be removed from the fuel salt as follows:

(a) a bubble generator injects helium bubbles in the salt stream;

(b) noble gases migrate to the helium bubbles due to their insolubility in the salt [1];

(c) and a gas separator discharges the fission-product-rich bubbles from the salt to the off-gas system.

Table 3.2: ^{135}Xe production sources and principal rate constants involved (reproduced from Kedl et al. [26]).

^{135}Xe gain mechanism	Principal rate parameters involved
Direct from fission yield $\gamma_{^{135}Xe} = 0.0022$	$\Sigma_f \gamma_{^{135}Xe} \phi$ (for ^{235}U fission)
^{135}I decay yield $\gamma_{^{135}Xe} = 0.031$, it decays to ^{135}Xe with $\tau_{1/2} = 6.68\ h$	$\Sigma_f \gamma_{^{135}I} \phi$ (for ^{235}U fission)
^{135}Te decay yield $\gamma_{^{135}Xe} = 0.033$, it decays to ^{135}I with $\tau_{1/2} = 19\ s$	$\Sigma_f \gamma_{^{135}Te} \phi$ (for ^{235}U fission)

Figure 3.2 shows the key pathways for xenon production, accumulation, and removal in a typical MSR. Figure 3.3 shows the conceptual design of the MSBR gas separation system. In that system, helium bubbles of a specific size are introduced in a salt stream via the primary pump bowl. These bubbles absorb noble gases before being separated from the salt by a gas separator. ORNL suggested that the MSBR off-gas system would inject d 0.508mm helium bubbles in the pump bowl, redirect 10% of the fuel salt flow through a bubble separator to remove the bubbles, and then return the flow into the pump suction. Robertson et al. reported that the helium bubble size was approximately 25% of the throat width (blue circle on Figure 3.4) and was independent of the gas flow rate [1]. Consequently, it is possible to regulate the helium bubble size by changing the throat width in the bubble generator.

To realistically model the gas separation system, a mathematical model is needed to describe noble gas extraction efficiency during reactor operation. Particularly, a model of xenon extraction efficiency as a function of sparger design parameters is needed to accurately model the ^{135}Xe removal in a fuel salt depletion simulation. The gain and loss terms for ^{135}Xe dissolved in the fuel salt are listed in Table 3.2 and

Table 3.3. The removal efficiency for the xenon in the pump bowl was measured during Molten Salt Reactor Experiment (MSRE) operation.

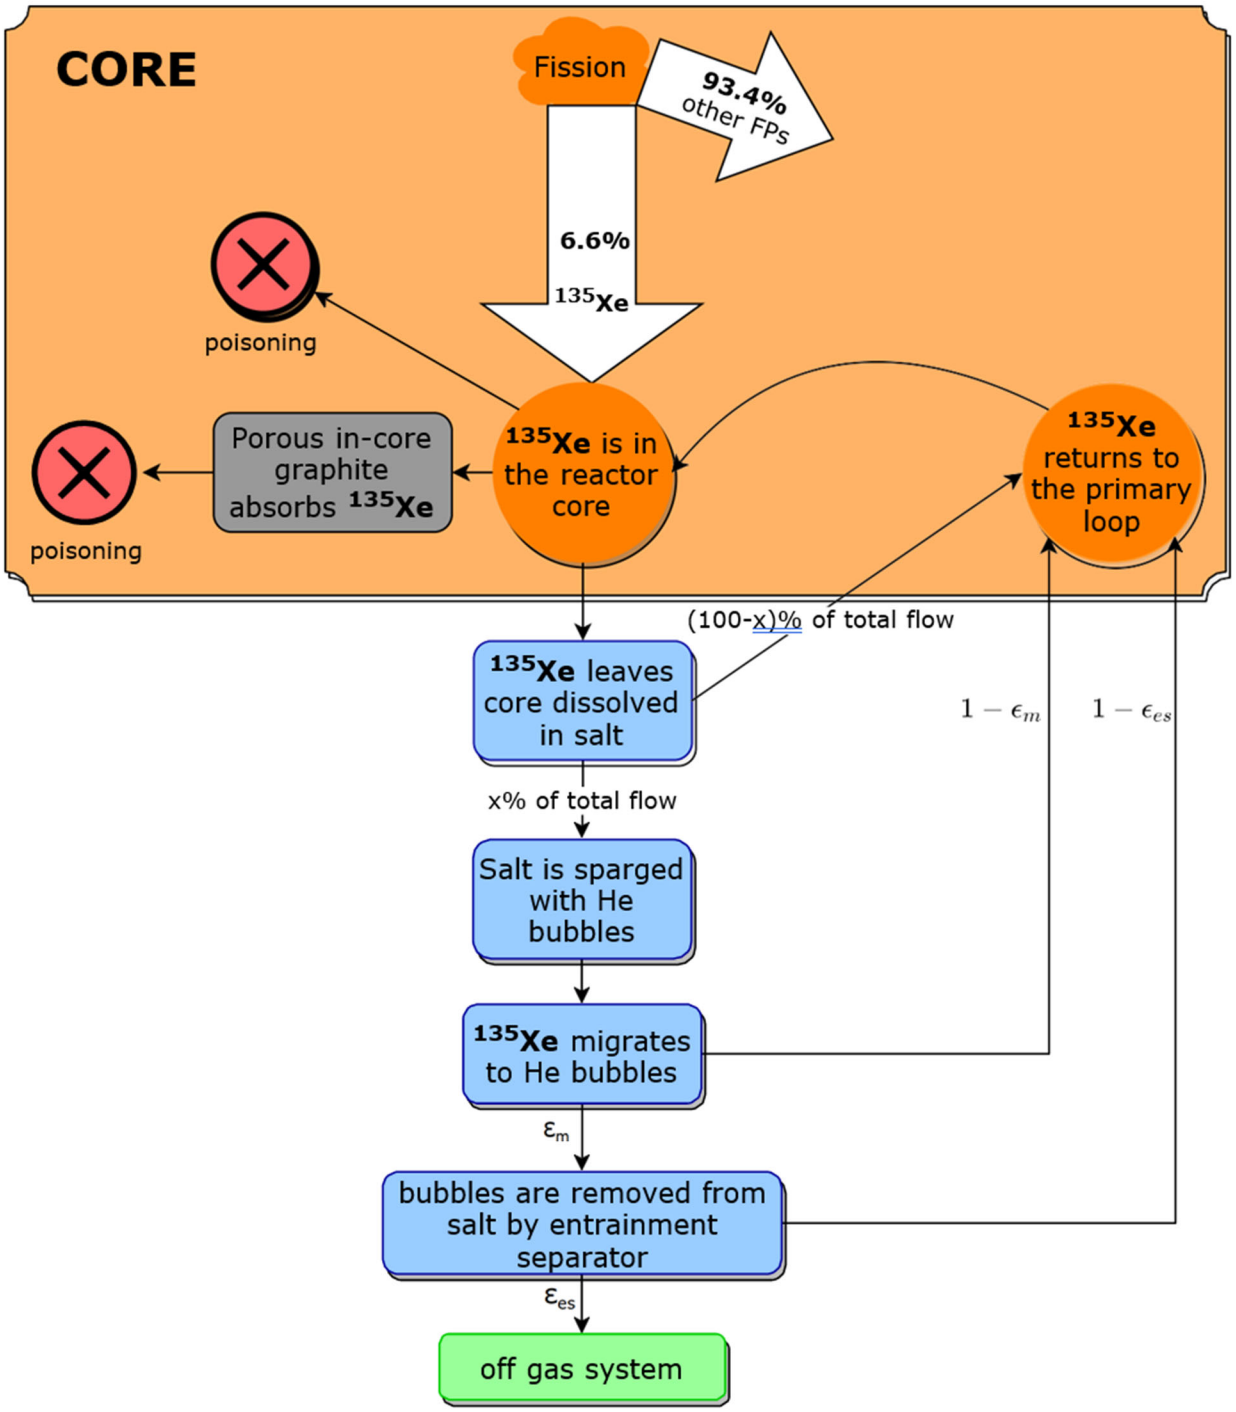

Figure 3.2: Schematic of ^{135}Xe circulation in a generic MSR. x is the fraction of fuel salt flow from the pump discharge redirected to the gas separation system, while ϵ_m and ϵ_{es} are the efficiencies of migration (of ^{135}Xe to the helium bubbles in the sparger) and separation (of gas in the entrainment separator), respectively. The orange color represents the fuel salt in the primary loop, the blue color represents the gas separation system, and the gray color is the moderator in the core. Fission yields assume ^{235}U fission only.

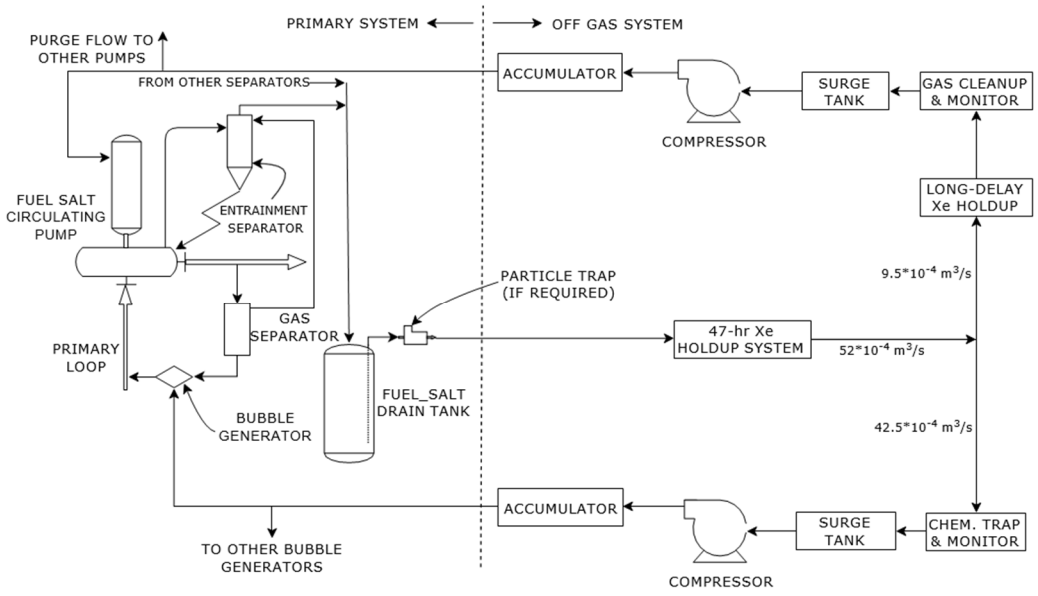

Figure 3.3: Schematic flow diagram of the MSBR gas separation system (reproduced from Robertson et al. [1]).

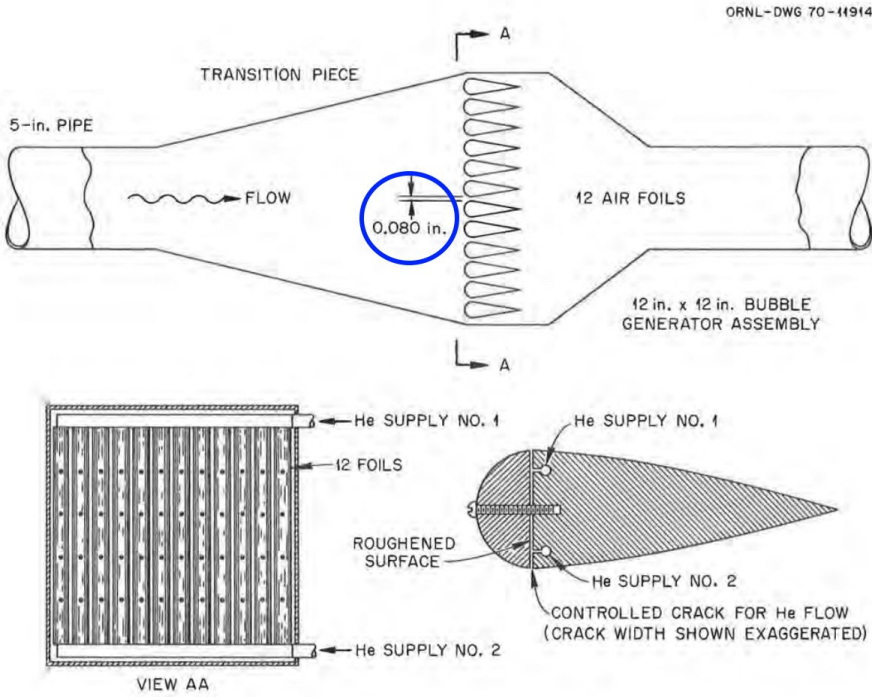

Figure 3.4: Preliminary concept of an MSBR bubble generator (reproduced from Robertson et al. [1]). The blue circle shows throat width, which determines bubble size.

Table 3.3: ^{135}Xe loss terms and principal rate constants involved (reproduced from Kedl et al. [26]).

^{135}Xe loss mechanism	Principal rate parameters involved
Decay of dissolved ^{135}Xe ($\tau_{1/2} = 9.1\ h$)	Decay constant(λ)
^{135}Xe burnup dissolved ^{135}Xe burnup as it passes through the core	Neutron flux (ϕ)
^{135}Xe migrated to helium bubbles	Removal efficiency (ϵ_m)
^{135}Xe transferred into circulating He bubbles; this xenon will eventually be burnup, decay, or stripped via bubble separator	Mass transfer coefficient (h), decay constant (λ), neutron flux (ϕ), bubble removal efficiency (ϵ_{es})

However, the technical report ORNL-4069 by Kedl-Houtzeel only stated its range (from 50 to 100%) and concluded, "It is probably a complex parameter like the circulating-void fraction and depends on many reactors operational variables" [26]. Peebles et al. in ORNL-TM-2245 has reported xenon removal efficiency (ϵ_{Xe}) in a gas separation system as a function of many parameters [28]:

$$\epsilon_m = \frac{1 - \epsilon^{-\beta}}{1 + \alpha} \tag{3-3}$$

where

$$\alpha = \frac{RT}{H} \cdot \frac{Q_{salt}}{Q_{He}} \tag{3-4}$$

$$\beta = \frac{K_L a A_C L(1 + \alpha)}{Q_{salt}} \tag{3-5}$$

$$\begin{aligned} R &= universal\ gas\ constant\ [L \cdot Pa \cdot mol^{-1} \cdot K^{-1}] \\ T &= salt\ temperature\ [K] \\ H &= Henry's\ law\ constant\ for\ solute\ gas\ [Pa \cdot mol^{-1} \cdot L] \\ Q_{salt} &= volumetric\ salt\ flow\ rate\ [m^3/s] \\ Q_{He} &= volumetric\ helium\ flow\ rate\ [m^3/s] \end{aligned}$$

Most of the input parameters for that correlation are obvious and easy to obtain from the system component design. The mass transfer coefficient for transferring xenon into helium bubbles (KL) can be estimated experimentally, but published information is currently insufficient to inform an accurate mathematical model appropriate for Computational Fluid Dynamics (CFD). Thus, Peebles et al. reported the mass transfer coefficient correlation for the MSBR salt (LiF-BeF_2-ThF_4-UF_4) but for a limited case. While it is out of the scope of this work to accurately estimate mass transfer coefficient, this work seeks to provide a tool which would allow the user to specify any mathematical model for a separation efficiency. Provided a mass transfer coefficient, the user can incorporate it into the model. Equation 3-5 would apply to other noble gases (e.g., Kr), but Henry's law constant (H) varies by element. As a result, the obtained mathematical model for gas removal efficiency might be employed to inform a realistic physics-based fuel reprocessing model in SaltProc.

3.2.1.2. Insoluble fission product filtering

Approximately 40% of FPs have gaseous elements in their decay chains. Some of the non-gaseous FPs produced in the MSR core (e.g., noble and semi-noble metals) have negligible solubility in the molten salt. Some fraction of noble and semi-noble solid FPs plate out onto the internal surfaces of the primary loop equipment, complicating their removal [143]. The remaining noble and semi-noble metals can be removed along with corrosion products using a mechanical filtration system, which "consists largely of a high surface area mechanical filter, likely a nickel mesh, to promote deposition of suspended, undissolved fission and corrosion products," according to Holcomb et al. [144]. The filter is manufactured from porous metal, has limited capacity, and needs periodic replacement. The filter replacement must be done using remote-controlled equipment due to high radiotoxicity of undissolved FPs and residual fuel salt remaining on the filter [145].

The historic MSRE program provided basic information design and performance of the large mechanical filter. Figure 3.5 shows the piping layout of the filter, storage, and processing tanks. The filter pressure vessel is made of high-nickel alloy (Inconel) and accommodates 40-μm pore size sintered Inconel fibers. This large molten salt filter had a total filtering area of 0.8 m^2 and was designed to filter approximately 1 kg of the molten salt per minute, but the removal efficiency has never been reported. Also, the design of the filter, the filter holder, and the remotely operated equipment for the filter replacement for commercial-scale MSR designs presents a significant engineering challenge [145].

In this work, ideal and constant separation efficiency in the filtering system is assumed. However, a physics-driven mathematical formula is implemented in the current version of SaltProc with adjustable parameters.

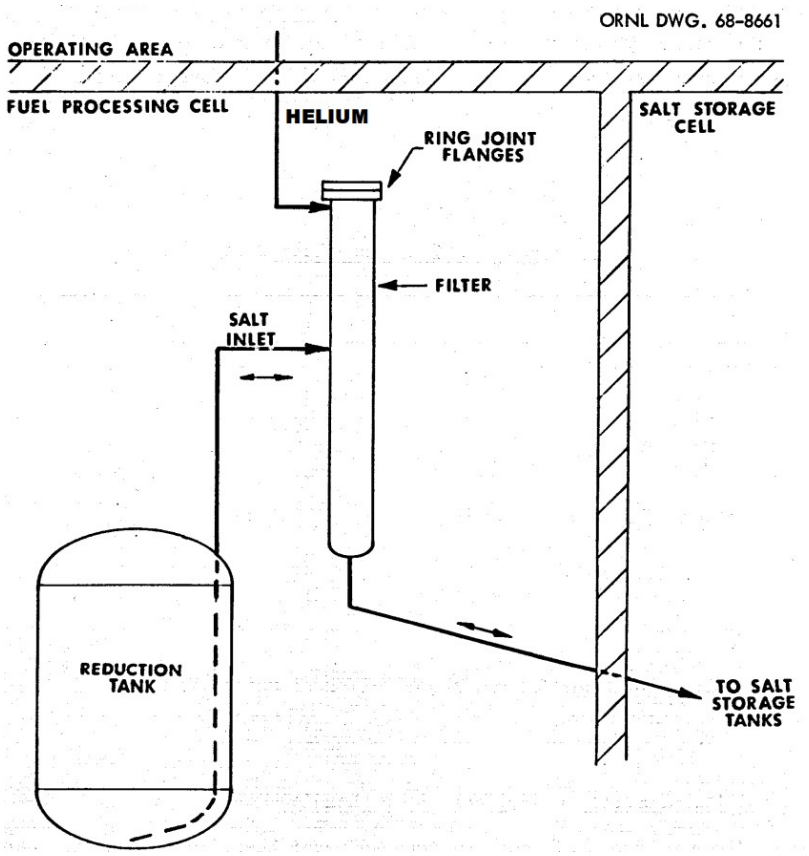

Figure 3.5: Schematic flow diagram of the large molten salt mechanical filter designed and operated during the MSRE (reproduced from Lindauer et al. [146]).

3.2.1.3. Fuel chemical processing facility

In addition to noble gases, noble metals, and semi-noble metals, the fuel salt reprocessing system must extract other FPs such as the lanthanides. These absorb fewer neutrons than ^{135}Xe, but their removal is crucial to guarantee normal operation. Unfortunately, lanthanides have relatively high solubility in the carrier salt and must be removed by chemical extraction. In thorium-fueled MSR designs, ^{232}Th in the fuel salt absorbs thermal neutrons and produces ^{233}Pa, which then decays into the fissile ^{233}U (Figure 3.6). Protactinium presents a challenge since it has a large absorption cross section in the thermal energy spectrum. Accordingly, ^{233}Pa is continuously removed from the fuel salt into a tank in which ^{233}Pa decays to ^{233}U without poisoning the reactor. This feature allows the thorium-fueled MSR to avoid neutron losses to protactinium, keeps FPs on a trace level, and increases the efficiency of ^{233}U breeding.

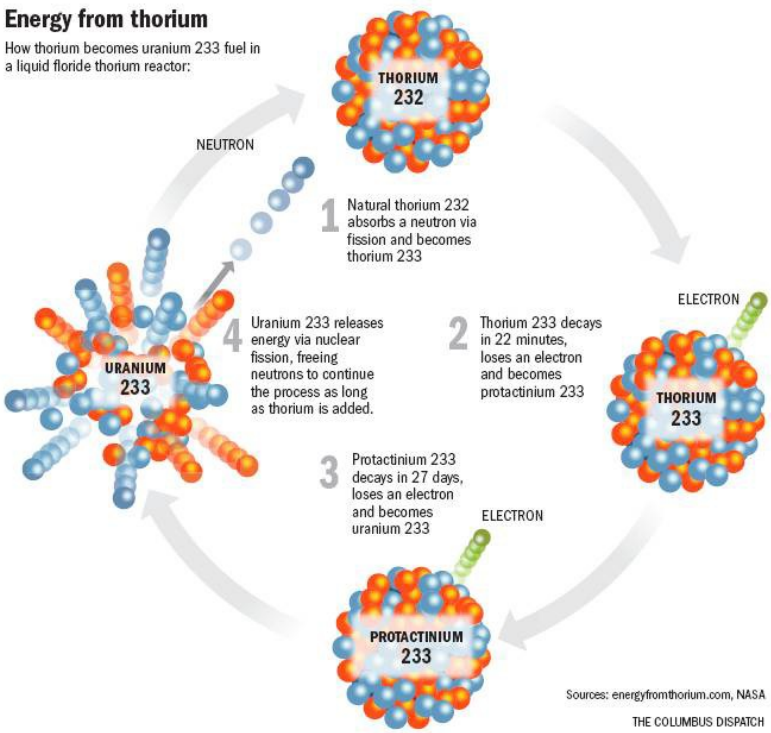

Figure 3.6: Production of ^{233}U from ^{232}Th (reproduced from Sorensen [147]).

Many authors report that a liquid-liquid reductive extraction process is the best option for removing protactinium and soluble FPs from molten fluoride salts [148-150]. In that process, the protactinium or lanthanides can be selectively stripped from the salt into liquid bismuth due to different chemical potentials. Moreover, the MSRE experience indicated that the extraction could be carried out rapidly and continuously [151].

The principal scheme of the MSBR reprocessing facility concept is shown in Figure 3.7. The fuel salt is first temporarily stored for cooling and decay of the shortest-lived fission products, then it is directed to the primary fluorinator. There, most of the uranium is removed by fluorination to UF6. After that, the salt is routed to an extraction column where it is combined with a mixture containing metallic bismuth, lithium, and thorium reductants.

The remaining uranium and protactinium are reductively extracted to a bismuth solution, leaving a salt that only contains fission products dissolved in carrier salt (base composition LiF-BeF_2-ThF_4). The salt then goes through a reduction column where UF_6 is reduced to UF_4, preparing it for return to the reactor. BeF_2 and ThF_4 are also added, and all residual bismuth is removed from the salt. After a final cleanup step and valence adjustment, the purified salt returns to the reactor [147, 152].

The bismuth, accommodating some uranium and protactinium, is routed to a hydrofluorination column where metallic solutes in the bismuth are oxidized into their fluoride forms in the presence of a decay salt[5]. The decay salt, containing UF_4, PaF_4, and ThF_4, passes into

[5] The decay salt contains UF_4, PaF_4, ThF_4 and FPs. Uranium produced after ^{233}Pa decay is extracted and directed back into the reactor. Decay salt is the precursor for the waste salt as it was periodically discarded every 220 days [1].

a decay tank where ^{233}Pa decays to ^{233}U. The uranium generated by protactinium decay is removed through fluorination to UF_6 and directed to the reduction column to refuel the purified fuel salt. A hydrofluorinator and a fluorinator can remove approximately 95% of the uranium from the stream [1].

To maintain or adjust the fissile material concentration in the reactor (and, consequently, control the reactivity), ^{233}U is added from the protactinium decay tank to fully processed salt on its way back to the reactor. Adding fissile material is performed by sparging the salt with UF_6 and hydrogen to produce UF_4 in the salt and HF gas [1].

After these separation steps, the fuel salt stream from the protactinium isolation system contains only traces of protactinium and uranium but contains practically all of the rare earths. A fraction of this salt stream is redirected to a reductive extraction process for removing rare earths. The principal scheme of a rare earth removal system is shown in Figure 3.8. A molten salt flow that contains rare earth fluorides is fed to the center of an extraction column. The salt flows countercurrent to a liquid bismuth stream, which contains thorium and lithium. In the upper part of the column, the rare earths are reduced and transferred to the downflowing liquid metal stream. Below the feed point, the rare earth concentration is increased in the salt and metal streams in order to produce a concentration high enough for disposal [148].

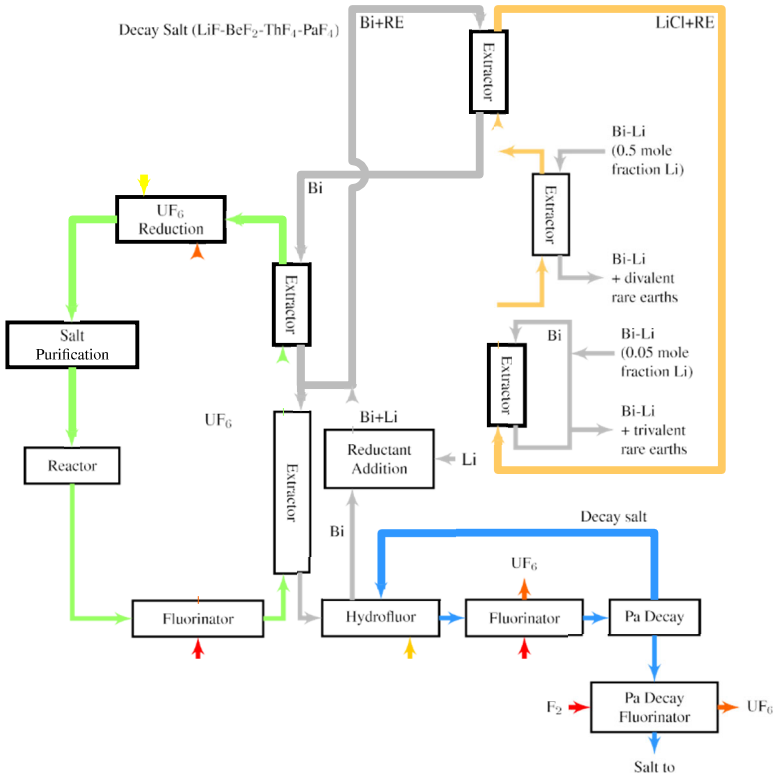

Figure 3.7: Simplified block diagram of chemical processing scheme for a single-fluid MSBR (reproduced from Sorensen [147]). RE represents the rare earth elements extracted from the salt.

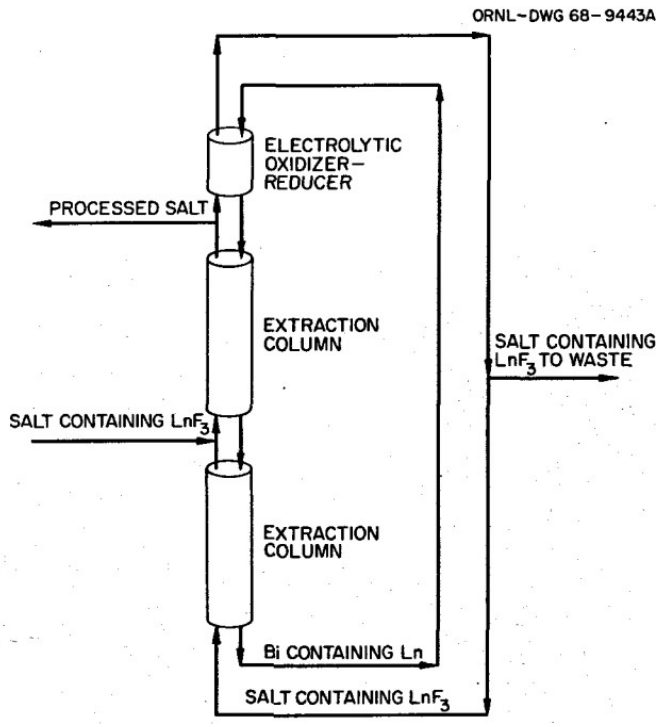

Figure 3.8: Rare earth removal from a fuel salt by reductive extraction (reproduced from Briggs et al. [148]).

While it is out of the scope of this work to derive the accurate chemistry-based mathematical formula for rare earths and protactinium separation efficiency, this work seeks to provide a flexible tool that is able to simulate chemical processes in significant detail concerning vital system design parameters.

3.2.2. Simulation Tool Design and Capabilities

The first version of the SaltProc tool for calculating MSR fuel composition evolution, taking into account an online reprocessing system, was developed in 2018 as a part of a M.S. thesis [123, 153]. The tool was designed to expand Serpent 2 depletion capabilities for modeling liquid-fueled MSRs with online fuel reprocessing systems. SaltProc v0.1 uses HDF5 [154] to store data and the PyNE Nuclear Engineering Toolkit [155] for Serpent 2 output file parsing and nuclide naming. SaltProc v0.1 is an open-source Python package that uses a batch-wise approach to simulate continuous feeds and removals in MSRs. In the batch-wise simulation, the fuel salt is processed in the processing system once after each depletion timestep. As long as the depletion time step is much longer than the circulation time V_{core}/Q_{salt}, which is 18.2 s for the TAP MSR at design condition, this approach is essentially equivalent to periodically activating the processing system in continuous operation of the reactor because of the relatively large time constant of fission product accumulation.

SaltProc v0.1 only allows 100% separation efficiency for either specific elements or groups of elements at the end of the specific “cycle time”. Capabilities of the developed tool, working with the Monte Carlo software Serpent 2, were demonstrated using the full-core MSBR design for

a simplified case with ideal removal efficiency (100% of mass for target elements removed) [156]. The SaltProc v0.1 architecture and the principal structure were not designed for flexible implementation of sophisticated online reprocessing systems, including realistic variable extraction efficiencies. In the current work, SaltProc v0.1 was completely refactored using Object-Oriented Programming (OOP) to create a comprehensive generic tool to realistically model complex MSR fuel reprocessing systems while taking into account variable extraction efficiencies, time-dependent core geometry, and the mass balance between the core and the reprocessing plant.

3.2.2.1. Software architecture

The SaltProc v1.0 Python toolkit couples directly with Serpent 2 input and output files, to couple the reprocessing system to depletion calculation. Python 3 OOP standard features are used to create a flexible, user-friendly tool with great potential for further improvement and collaboration. Figure 3.9 shows the SaltProc v1.0 class structure which includes 4 main classes:

Depcode. *Depcode* class contains attributes and methods for reading the user's input file for the depletion software, initial material (e.g., fuel and/or fertile salt) composition, principal parameters for burnup simulation (e.g., neutron population and number of cycles for Monte Carlo neutron transport) and running the depletion code.

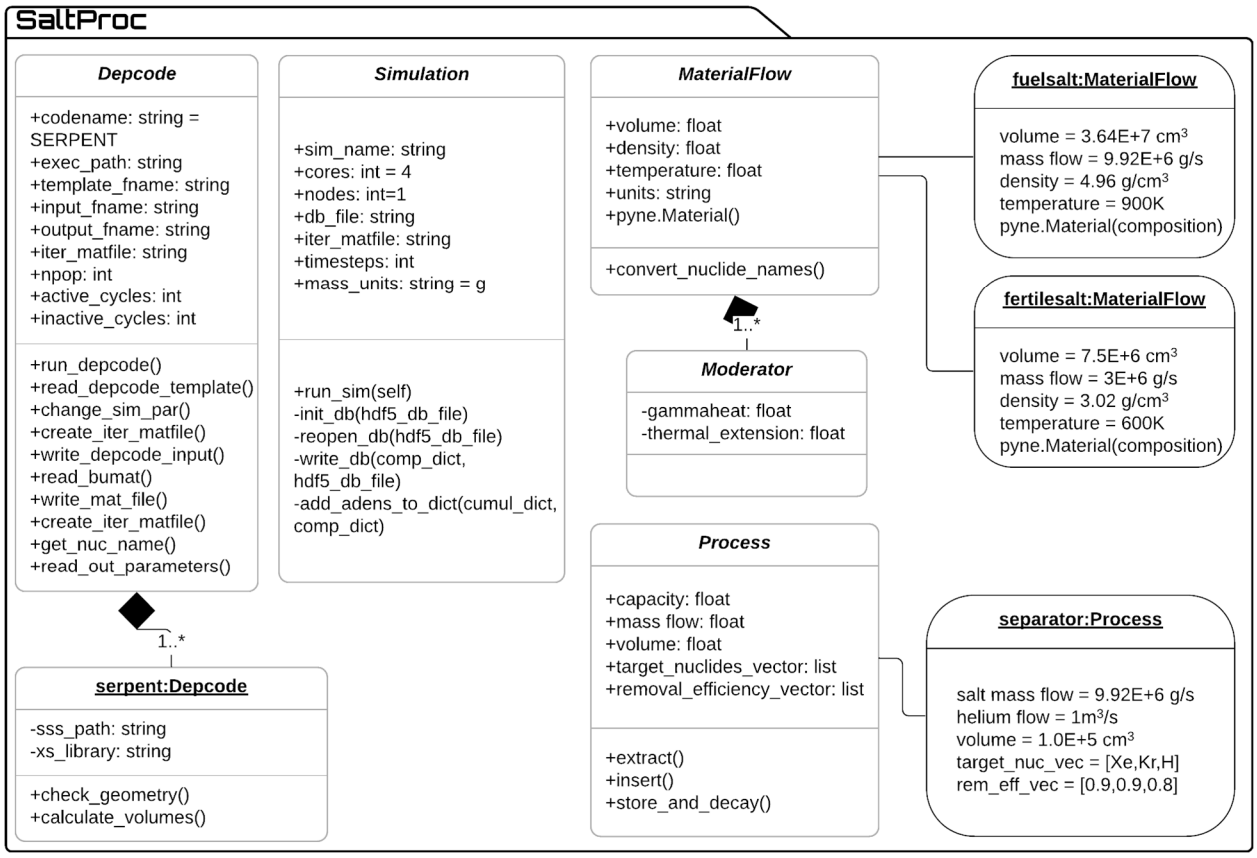

Figure 3.9: SaltProc v1.0 Python package class diagram in UML notation with examples of object instances.

Simulation. Simulation class runs a depletion step, creates, and writes an HDF5 database, tracks time, and converts isotopic composition vector nuclide names from a depletion code format to human-readable format.

MaterialFlow. Each MaterialFlow object represents the material flowing between Process objects (Figure 3.10 and Figure 3.11). All instances of this class contain an isotopic composition vector stored in PyNE Material object [155], mass flow rate, temperature, density, volume, and void fraction. Existing PyNE Material capabilities convert the units of the isotopic composition vector (e.g., from the atomic density provided by Serpent to a mass fraction or absolute mass in desired units) and decay the material (i.e., model the MSBR protactinium decay tank). The main idea of the MaterialFlow object is to pass detailed information about the salt starting at the MSR vessel outlet throughout reprocessing components (Processes), which modify the MaterialFlow object before depleting the material in the next depletion step.

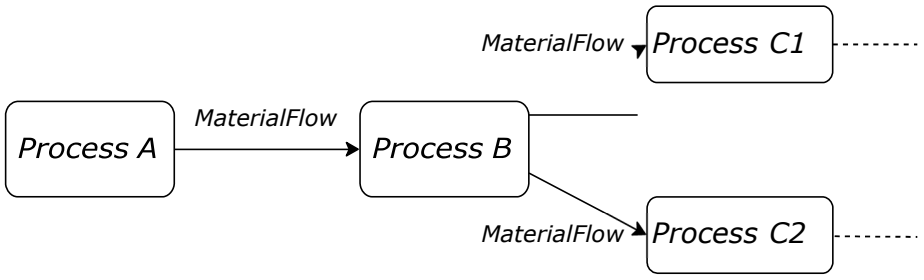

Figure 3.10: Schematic for passing material data between fuel processing system components shown for a general case.

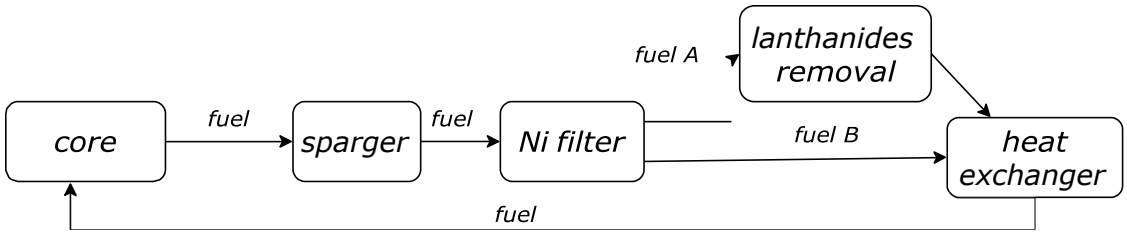

Figure 3.11: Schematic for passing material data between fuel processing system components shown for the TAP concept.

Process. Each Process object represents a realistic fuel processing step characterized by its throughput rate, volumetric capacity, extraction efficiency for each target element (can be a function of many parameters), waste streams, and other process-specific parameters. The feed Process injects fresh fuel salt MaterialFlow directly into the reactor core (e.g., adding fissile material with a specific mass flow rate to MaterialFlow after performing all removals).

Such a class structure provides outstanding flexibility in simulating various MSR fuel processing system designs. A library of various MaterialFlow is created (e.g., fuel salt flow, fertile salt flow, refueling salt flow) and Process (e.g., helium sparging facility, gas separator, nickel filter) object examples to help a user to create a model of a desired reprocessing scheme quickly. At runtime, the user should connect Process objects in series, parallel, or both with MaterialFlow

objects to form a comprehensive reprocessing system. To make the reprocessing system definition self-explanatory and straightforward, a standardized graph description language, dot, is employed, which is widely used in computer science for describing undirected and directed graphs [157]. A directed graph is a set of objects (vertices) connected by edges, where the edges have a direction associated with them. In the context of this work, a vertex is Process object (a component of the fuel processing systems), and a directed graph edge is MaterialFlow object (salt flow).

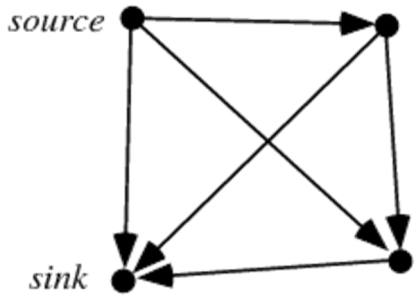

Figure 3.12: General example of directed graph with a source (e.g., a reactor outlet) and sink (e.g., a reactor inlet).

The reprocessing plant structure described with dot can be simply plotted using Graphviz [158] and those plots can be used for analysis, optimization, and publication purposes. The user also has the flexibility to create custom objects with desired attributes and methods and contribute back to the code package using GitHub (https://github.com/arfc/saltproc).

3.2.2.2. Tool flowchart

Figure 3.13 illustrates the online reprocessing simulation algorithm coupling SaltProc v1.0 and Serpent. A json-compatible user input file for SaltProc contains depletion software parameters such as paths to the depletion software executable, neutron population and number of criticality cycles, and total heating power. Additionally, the input file contains reprocessing system parameters such as structure of reprocessing system, capacity and efficiency of the system components, and molten salts thermophysical properties. To perform a depletion step, SaltProc v1.0 reads a user-defined Serpent template file. This file contains input parameters such as the path to a nuclear data library, material isotopic composition at startup, burnup calculation parameters, and boundary conditions. SaltProc v1.0 fills in the template file and runs Serpent single-step depletion.

After the depletion calculation, SaltProc v1.0 reads the depleted fuel composition file into the MaterialFlow object (core outlet in Figure 3.13). This object contains an isotopic composition vector, total volume of material, total mass, mass flow rate, density, temperature, void fraction, etc. For the simplest reprocessing case, if all fuel processing components are connected in series (100% of total material flow goes through a chain of separation components), the core outlet object flows sequentially between Processes, and each Process removes a mass fraction of target elements with specified extraction efficiency. After that the removed material mass is compensated by fresh fuel salt to maintain the salt inventory in a primary loop. Finally, the resulting isotopic composition after reprocessing is stored in the HDF5 database and dumped in a new composition file for the

next Serpent depletion run. SaltProc v1.0 also stores the isotopic composition before reprocessing and waste stream from each fuel processing component in the HDF5 database.

For a more general case with multiple concurrent extraction processes, separate MaterialFlow objects are created for each branch with a user-defined mass flow branching percentage (e.g., 90% of total mass flow rate flows via left branch and 10% throughout a right branch). The total mass and isotopic composition vector for each MaterialFlow object are calculated as a fraction of incoming core outlet flow. Then each MaterialFlow object is passed via a cascade of Processes to separate selected chemical elements with specific efficiency. Finally, the left-hand-side MaterialFlow object is merged with the right-hand-side, and similarly to the previous case, fresh fuel salt feed compensates for mass losses in the Processes and keeps the fuel salt mass in the primary loop constant.

The UML diagram (Figure 3.9) allows the user to model a complex, multi-zone, multi-fluid MSR operation and is sufficiently general to represent myriad reactor systems. SaltProc v1.0 only stores and changes the isotopic composition of the fuel stream, which makes it a flexible tool to model any geometry: an infinite medium, a unit cell, a multi-zone simplified assembly, or a full-core. This flexibility allows the user to perform simulations of varying fidelity and computational intensity. SaltProc v1.0 is an open-source tool available on GitHub. Although the user needs Serpent 2.1.31 installed on his computer to use SaltProc. The tool leverages unit tests and continuous integration crucial for software sustainability [159, 160]. The documentation automatically generated using Sphinx [161] is available here: https://arfc.github.io/saltproc/.

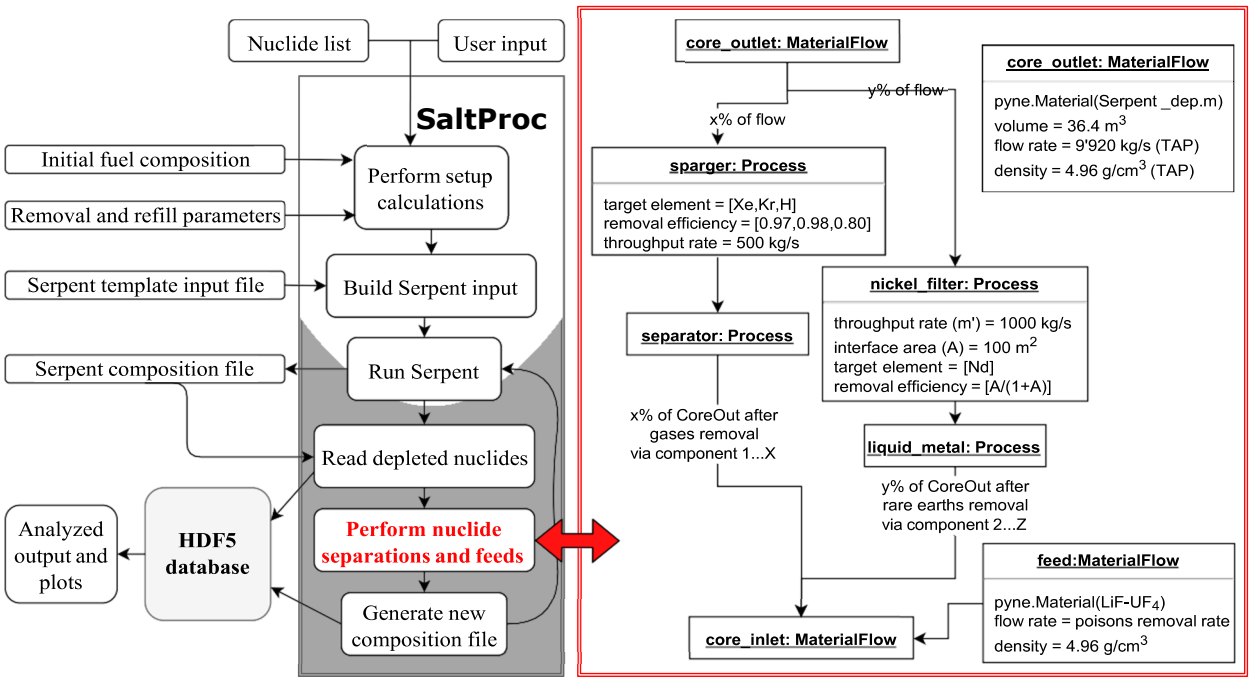

Figure 3.13: Flow chart for the SaltProc v1.0 Python package.

In summary, the development approach of SaltProc v1.0 is focused on producing a generic, flexible and expandable tool that extends Serpent 2 Monte Carlo code for advanced in-reactor fuel cycle analysis as well as simulate many online refueling and fuel reprocessing systems.

3.2.2.3. Reactivity control module

Simulation of specific MSR concepts requires changing the reactor core geometry during lifetime-long operation modeling. For instance, the TAP concept aims to increase the core lifetime by using continuous fresh fuel feeds, removing FPs, and reconfiguring moderator rod assemblies to compensate for negative reactivity insertion due to fissile material burnup. The concept proposes maintaining reactivity in the long term by replacing stationary moderator assemblies with denser lattices to increase the moderator-to-fuel ratio [18]. SaltProc v1.0 can switch from one core geometry to another core geometry (e.g., with larger moderator-to- fuel ratio) to mimic moderator or absorber movement if the effective multiplication factor, k_{eff}, falls below 1. This unique capability allows SaltProc v1.0 to analyze the fuel cycle performance of any liquid-fueled MSR system, including advanced designs with a moving moderator (e.g., TAP MSR).

3.2.3. Concluding Remarks

This section presented an overview of a fuel salt reprocessing in MSRs. It described various components of the plant and the physical or chemical mechanism responsible for neutron poison extraction from the salt. General core physics aspects and Serpent 2 depletion software capabilities have then been discussed. It also introduced SaltProc, a Python package developed and used to simulate continuous feeds and removals in various MSR designs.

In the following sections, SaltProc v1.0 will be demonstrated and validated for the TAP MSR, load following analysis are performed for both TAP MSR and MSBR. For a similar validation of the code with MSBR, the reader can refer to the thesis of Rykhlevskii [162].

3.3. Lifetime-long Depletion of Transatomic Power MSR

This section presents the TAP MSR core lifetime-long (25 years) depletion simulation with moderate time resolution (3-day depletion step) and a constant, 100% power level. The results obtained with SaltProc v1.0 are compared with full-core TAP depletion analysis by Betzler et al. [18] with assumed ideal removal efficiency (100% of the target isotope is removed). This validation effort showed that the SaltProc v1.0 solution matches the case with ideal extraction efficiency. Finally, this section presents a lifetime-long fuel salt depletion simulation for the case with a realistic, physics-based mathematical model for noble gas removal efficiency, which provides fuel isotopic composition evolution during 25 years of the TAP MSR operation. Additionally, this section presents safety and operational parameters evolution during operation. Detailed insights about fuel salt composition and neutron spectrum dynamics obtained herein will be used in the following sections to investigate TAP reactor poisoning during load-following.

3.3.1. TAP System Model

In Section 1.3, a general description of the TAP reactor design is given. In this section, the TAP core and fuel salt reprocessing system models for demonstrating SaltProc v1.0 are described in detail.

3.3.1.1. Serpent 2 Full-core Model

Nested and lattice geometry types, as well as transformation capabilities of Serpent [124], are employed to represent the TAP core. Figure 3.14 shows the XY section of the whole-core model at the expected reactor operational level when all control rods are fully withdrawn. Figure 3.15 and Figure 3.16 show a longitudinal section of the reactor. This model contains the moderator rods with their silicon carbide cladding, the pressure vessel, and the inlet and outlet plena (Table 3.4). Fuel salt flows around square moderator assemblies consisting of lattices of small-diameter zirconium hydride rods in a corrosion-resistant material. The salt volume fraction for Figure 3.14 is 0.917204, which means the modeled core is under-moderated and has an intermediate neutron spectrum. Quarter-core configurations of the TAP core with various salt volume fractions, used in the current work to maintain criticality for a reasonable operational period (20 years), are listed in Table D.1, Figures D.1, and D.2 in Appendix D.

To represent the reactivity control system, the model has:

a) control rod guide tubes made of nickel-based alloy;

b) control rods represented as boron carbide (B_4C) cylinders with a thin Hastelloy-N coating;

c) air inside guide tubes and control rods.

The control rods must be able to suppress excess reactivity at the BOL when the core configuration is the most reactive, and the neutron spectrum is the hardest. The control rod design shown on Figure 3.14, Figure 3.15, and Figure 3.16 is comprised of a cluster of 25 rods that provide a total reactivity worth of 3922 ± 10 pcm at the BOL.

The control rod cluster is modeled using the TRANS Serpent 2 feature, which allows the user to change the control rod position during the simulation easily. The current work assumed that all control rods are fully withdrawn from the core (Figure 3.16), but the user can use reactivity control capabilities in SaltProc v1.0 to change control rod position during operation. In this dissertation, all figures of the core were generated using the built-in Serpent plotter.

The neutron population per cycle and the number of active/inactive cycles were chosen to obtain a balance between minimizing uncertainty for a transport problem (28 pcm for k_{eff}) and simultaneously minimizing computational time.

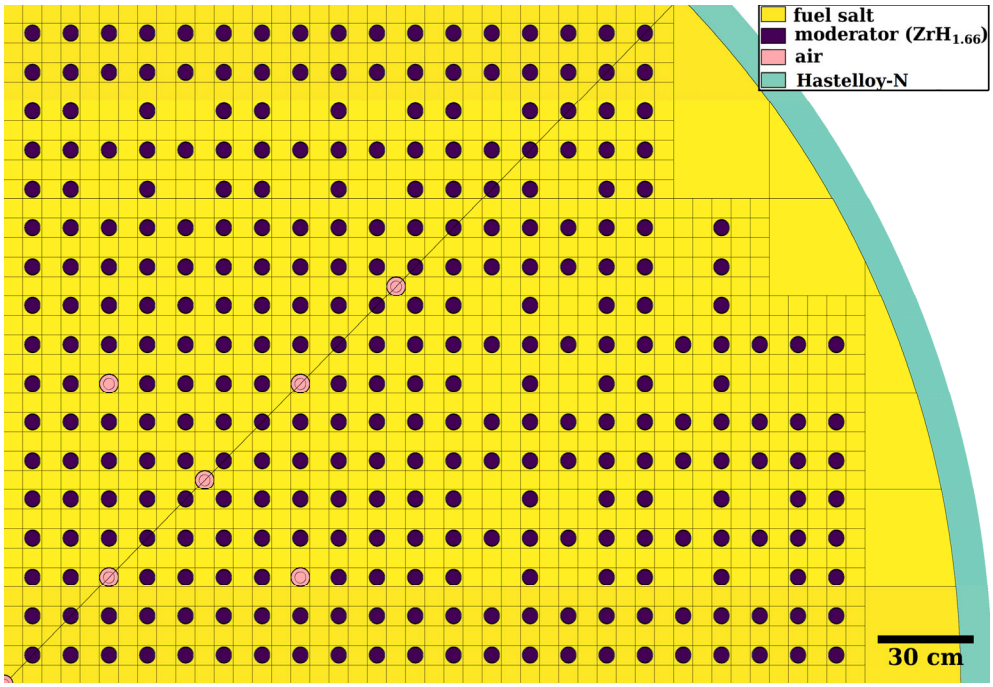

Figure 3.14: An XY section of the TAP model at horizontal midplane with fully withdrawn control rods at BOL (347 moderator rods, salt volume fraction 0.917204) [163, 164].

Table 3.4: Geometric parameters for the full-core 3D model of the TAP (reproduced from Betzler et al. [18]).

Component	Parameter	Value	Unit
Moderator rod	Cladding thickness	0.1	cm
	Radius	1.15	cm
	Length	3	m
Moderator assembly	Pitch	3	cm
	Array	5×5	rods×rods
	Pitch	15	cm
Core	Assemblies	268	assemblies/core
	Inner radius	1.5	m
	Plenum height	25	cm
	Vessel wall thickness	5	cm

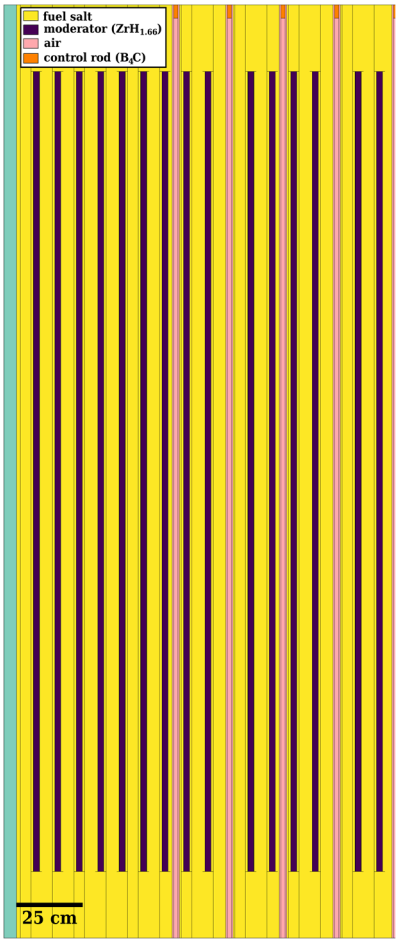

Figure 3.15: 45° XZ section of the TAP core model [163, 164].

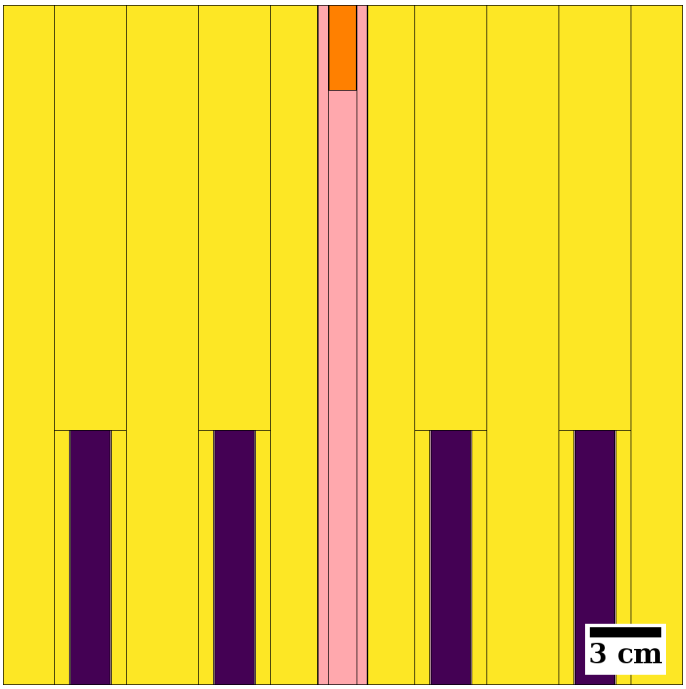

Figure 3.16: Zoomed XZ section of the top of the moderator and control rods in the TAP model.

3.3.1.2. Model of the Fuel Reprocessing System

The original TAP reprocessing system design (Figure 1.5) and neutron poison removal rates (

Table 3.3) are thoroughly analyzed to determine a suitable reprocessing scheme for the SaltProc v1.0 demonstration (Figure 3.17). This section presents two demonstration cases: with ideal and realistic, non-ideal gas removal efficiency. Realistic noble gas removal efficiency is based on the physical model for noble gas extraction efficiency discussed in Section 3.2.1.1.

Arrows on Figure 3.17 represent material flows, percent represents a fraction of total mass flow rates; ellipses represent fuel reprocessing system components; boxes represent waste streams; the diamond shows refuel material flow (UF_4, 5 wt% of ^{235}U). The efficiency of gas migration to helium bubbles (ϵ_m) and efficiency of gas bubbles separation from the salt (ϵ_{es}) are different for various demonstration cases and discussed in more detail in Sections 3.3.2.1 and 3.3.2.2 Efficiency of noble metal extraction in the nickel filter (Figure 1.5, orange block) and semi-noble metals/rare earths (RE) in the liquid metal extraction system (Figure 1.5, green block) is assumed fixed and equal 100% and 57%, respectively.

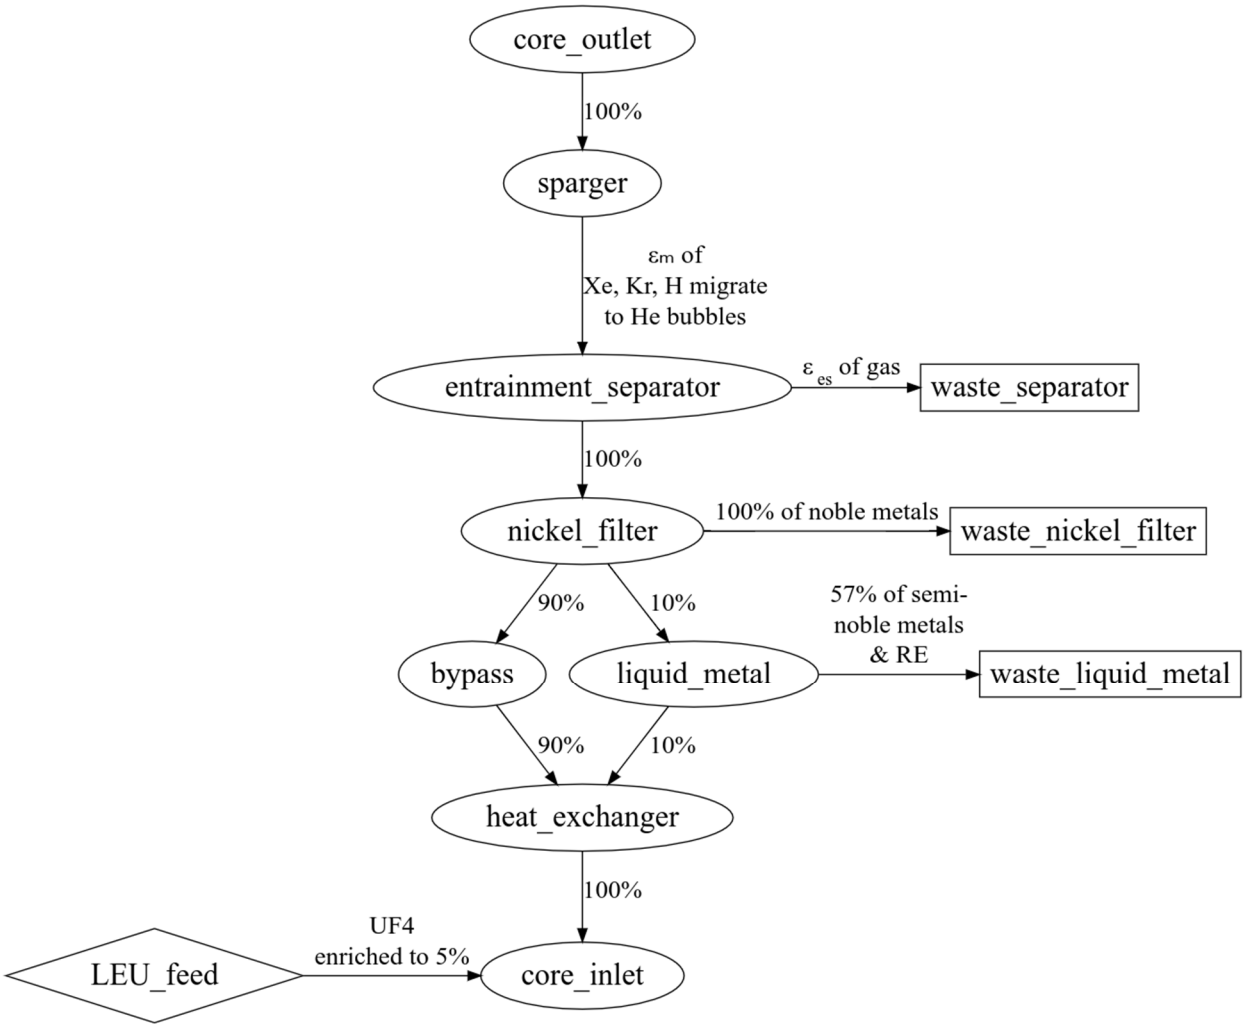

Figure 3.17: TAP reprocessing scheme flowchart used for the demonstration of SaltProc v1.0.

The gas removal components (sparger/contactor and entrainment separator) are located in-line because the estimated full loop time[6] for the fuel salt is about 20 seconds and approximately equal to the cycle time (

Table 3.3). To extract volatile gases every 20 seconds, the gas removal system must operate with 100% of the core throughout flow rate (in-line gas removal system). In this section, the efficiency of noble gas migration to helium bubbles and the efficiency of bubble removal from the salt by the entrainment separator (ϵ_m, ϵ_{es} on Figure 3.17, respectively) are selected separately for each demonstration case.

The nickel filter in the TAP concept is designed to extract noble/semi-noble metals and volatile fluorides (

[6] Full loop time is the time taken by a particle of the coolant to make one full circle in the primary loop.

Table 3.3). Similar to volatile gases, noble metals must be removed every 20 seconds and, hence, the filter should operate at 100% of the flow rate through the core. The nickel filter removes a wide range of elements with various effective cycle time (

Table 3.3).

Lanthanides and other non-noble metals have a lower capture cross section than gases and noble metals. These elements can be removed via a liquid-metal/molten salt extraction process with relatively low removal rates (cycle time 50 days). This is accomplished by directing a small fraction of the salt mass flow leaving the nickel mesh filter (10% of the flow rate throughout the core) to the liquid-metal/molten salt component of the reprocessing system, in which lanthanides are removed with a specific extraction efficiency to match the required cycle time (

Table 3.3). The remaining 90% of the salt mass flow is directed from the nickel filter to the heat exchangers without performing any fuel salt treatment.

The removal rates vary among nuclides in this reactor concept, which dictate the necessary resolution of depletion calculations. To compromise, a 3-day depletion time step was selected for the long-term demonstration case based on a time step refinement study by Betzler et al. [18] A complimentary time step refinement study is presented in Section 3.3.2.1 to determine the impact of temporal resolution on the depleted composition calculation.

3.3.2. Long-term Depletion Demonstration and Validation

3.3.2.1. Constant, Ideal Extraction Efficiency Case

To validate SaltProc v1.0, a lifetime-long depletion calculation is performed with ideal extraction efficiency. This case was selected to repeat fuel salt depletion as close as possible to the ChemTriton simulation for the full-core TAP reactor by Betzler et al. The following assumptions and approximations are made in their work [18]:

Effective cycle times as prescribed by the Transatomic Power Technical White Paper [15] (

a) Table 3.3) with 100% noble gas removal efficiency; hence, ϵ_{es} and ϵ_{m} in the reprocessing model (Figure 3.17) are both set to 1.0.
b) 5% LEU feed rate is equal to the rate of fission product removal.
c) 3-day depletion step.
d) Quarter-core, 3-D model with vacuum boundary conditions.
e) Delayed neutron precursor drift was neglected.

These assumptions are adopted for code-to-code verification of SaltProc v1.0 against ChemTriton. The ENDF/B-VII.1 [166] nuclear data library is used for this case to be consistent with Betzler's work. Unfortunately, some crucial details have not been reported in [18]: (1) exact core geometries for various moderator rod configurations except startup configuration; (2) the excess reactivity at startup; (3) the library from which S(α, β) tables for thermal scattering in zirconium hydride are obtained. This section presented my best effort to repeat Betzler's simulation using the same input data to validate SaltProc for the TAP concept.

Effective multiplication factor dynamics

Figure 3.18 and Figure 3.19 demonstrate the effective multiplication factor obtained using SaltProc v1.0 with Serpent. The k_{eff} was obtained after removing fission products and adding feed material at the end of each depletion step (3 days for this case). SaltProc v1.0 updated the moderator rod configuration to the next configuration (e.g., from 1388 rods per core to 1624 rods per core) once the predicted value of k_{eff} at the end of the next depletion step dropped below 1. This algorithm mimics regular maintenance shutdown when the TAP core excess reactivity is exhausted, and moderator rod assemblies should be reconfigured to operate the next cycle.

An optimal number of moderator configurations (cycles) is found to be 15 (see Appendix D). Fewer cycles would improve capacity factor but need larger excess reactivity at the Beginning of Cycle (BOC), which is strictly limited by reactivity control system worth. More cycles would require more frequent moderator rod reconfigurations, which worsens the capacity factor. The interval between the first and second moderator configuration was only 12 months, the shortest interval between moderator configuration updates. For the operation interval between 2 and 16 years after startup, the intervals between shutdowns for moderator rod updates were 18-26 months. However, towards the EOL, the intervals between moderator rod reconfigurations dropped to 13 months. Overall, the average interval between regular shutdowns for the core reconfiguration was 18 months, which exactly matches the refueling interval for conventional LWRs and is consistent with Betzler et al. (16 months) [18].

The k_{eff} fluctuates significantly as a result of the batch-wise nature of the online reprocessing approach used. Loading the initial fuel salt composition with 5% LEU into the TAP core leads to a supercritical configuration with an excess reactivity of about 3200 pcm (Figure 3.18). Without performing any fuel salt reprocessing and spectrum shifting, the core became subcritical after 30 days of operation [164]. SaltProc calculates an operational lifetime of 22.5 years, after which the fuel salt reached a total burnup of 81.46 MWd/kgU. The end of an operational lifetime is achieved when the minimum SVF is obtained, as restricted by the moderator geometry parameters (e.g., moderator rod diameter, rod pitch, the internal diameter of the reactor vessel). Table 3.5 compares obtained results with Betzler et al. [18].

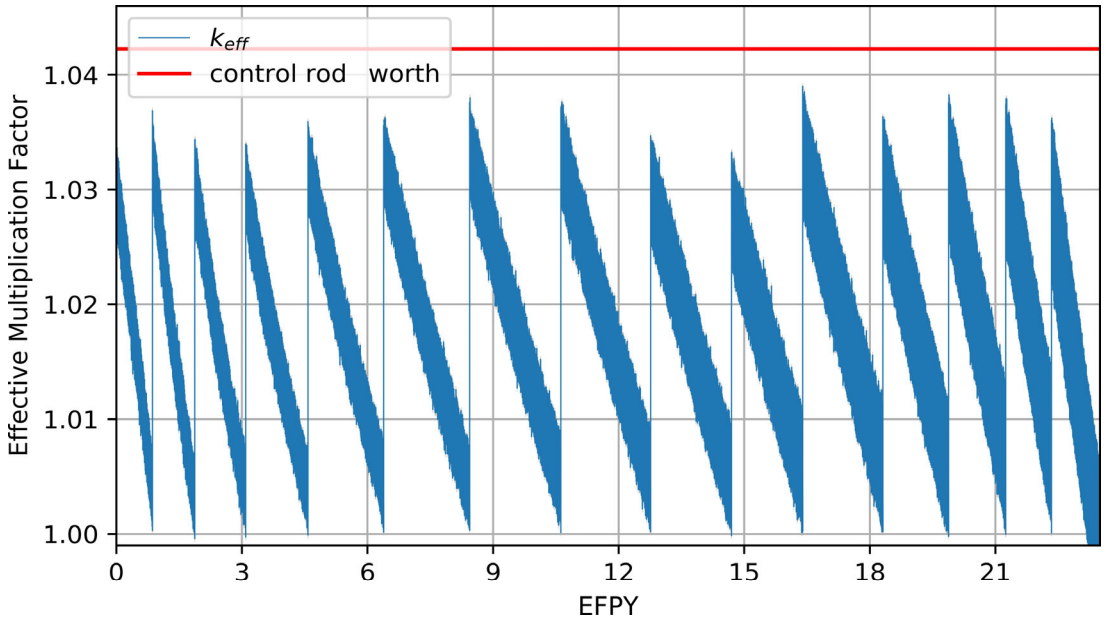

Figure 3.18: Effective multiplication factor dynamics during 23.5 years of operation for the full-core TAP core model for the case with an ideal removal efficiency of fission product. Confidence interval σ=28 pcm is shaded.

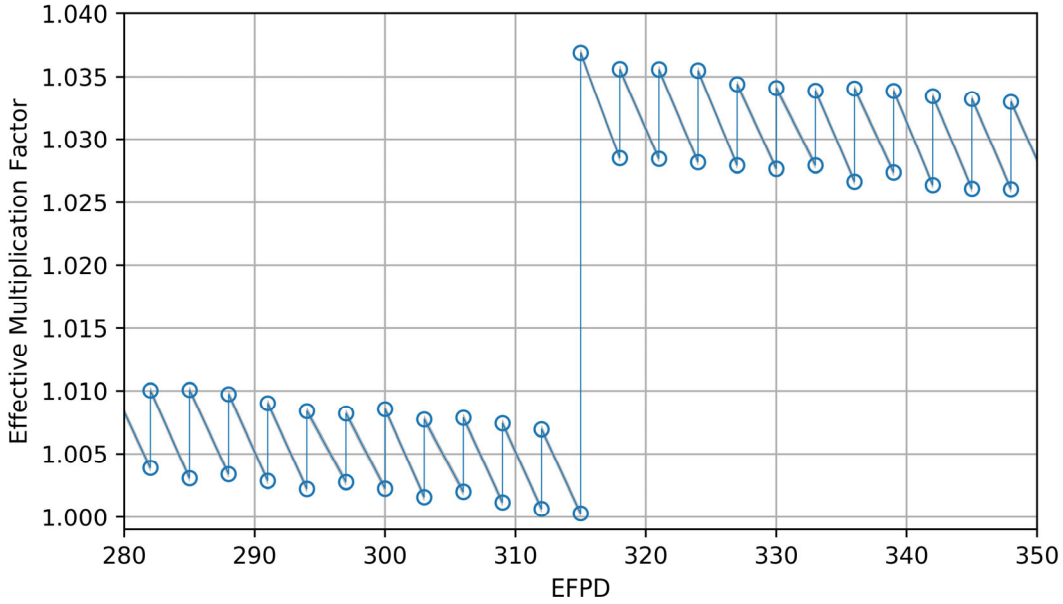

Figure 3.19: Zoomed effective multiplication factor for the interval from 280 to 350 EFPD while transitioning from Cycle #1 (startup geometry configuration, 347 moderator rods, SVF=0.91720353) to Cycle #2 (SVF=0.88694). Confidence interval σ=28 pcm is shaded.

Overall, SaltProc-calculated operational lifetime and burnup are lower than the reference by approximately 22% and 17%, respectively. A better match in the operational lifetime between

SaltProc v1.0 and ChemTriton can be obtained if a detailed moderator configuration description of Betzler's model will be available in the future.

Table 3.5: Comparison of main operational parameters in the TAP reactor between the current work and Betzler et al. [18].

Parameter	Current work	Betzler, 2017 [18]
Operational lifetime [y]	22.5	29.0
Discharge burnup [MWd/kgU]	76.30	91.9
Average moderator reconfiguration interval [months]	18	16

Fuel salt isotopic composition dynamics

Figure 3.20 show that continuous LEU feed into the TAP reactor is not sufficient to maintain the fissile ^{235}U content of the core, as the uranium enrichment steadily decreases from 5% at the BOL to 1% at the EOL. However, during the first 13 years of operation, the TAP MSR breeds fissile ^{239}Pu and ^{241}Pu, reaching a peak of total fissile plutonium inventory of 2.15 t (Figure 3.22). Figure 3.21 shows that a significant amount of non-fissile plutonium (^{238}Pu, ^{240}Pu, and ^{242}Pu) and uranium (^{236}U) builds up in the reactor during operation and negatively impacts criticality of the reactor. ^{239}Pu and ^{241}Pu are major contributors to the fissile material content of the core, keeping it critical during the second half of the operational lifecycle. The total ^{239}Pu inventory in the core rises during the first 11 years of operation due to the harder neutron spectrum. After 11 years, the softer spectrum breeds less ^{239}Pu from ^{238}U, and more of ^{239}Pu is progressively burned. Obtained results are in good agreement with results in ORNL Report by Betzler et al. (

Table 3.6) [18].

A lifetime-long SaltProc depletion calculation requires a 5% LEU feed rate of 460.8 kg per year to maintain the fuel salt inventory in the primary loop, which is consistent with the reference. Table 3.7 shows the main fuel cycle performance parameters calculated using SaltProc and compared with the reference. Normalized per GWth-year, the TAP concept requires about 5.23 t of fuel compared with 4.14 t reported by Betzler et al. SaltProc calculated waste production normalized per GW_{th}-year is 5% less than reported by ORNL. Potentially, the TAP can operate with LWR SNF as the fissile material feed. The heavy metal component of LWR SNF has a lower fissile material weight fraction than 5% enriched uranium and adds less fertile ^{238}U to the fuel salt, potentially reducing the operational lifetime. Nevertheless, in the case of using waste material (e.g., transuranium elements from LWR SNF) in this fueling scenario, the TAP concept has superior waste reduction metrics.

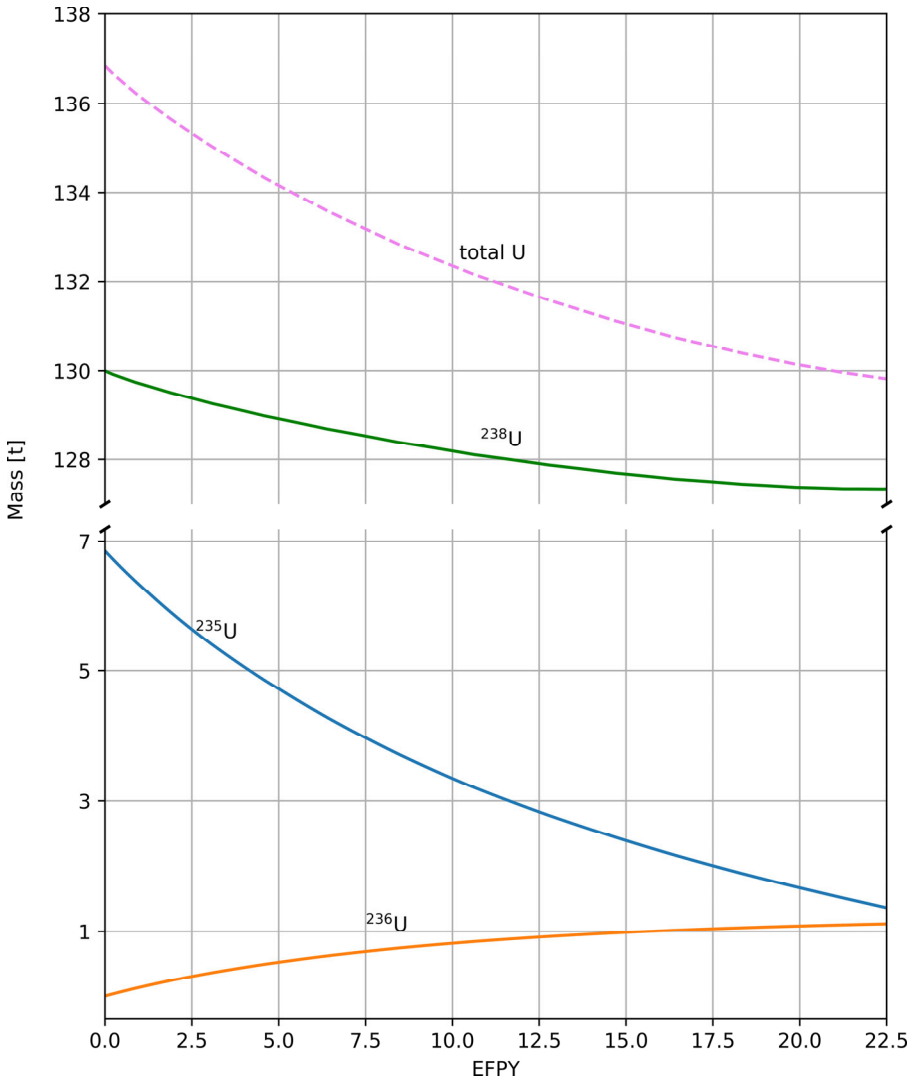

Figure 3.20: SaltProc-calculated uranium isotopic fuel salt content during 22.5 years of operation.

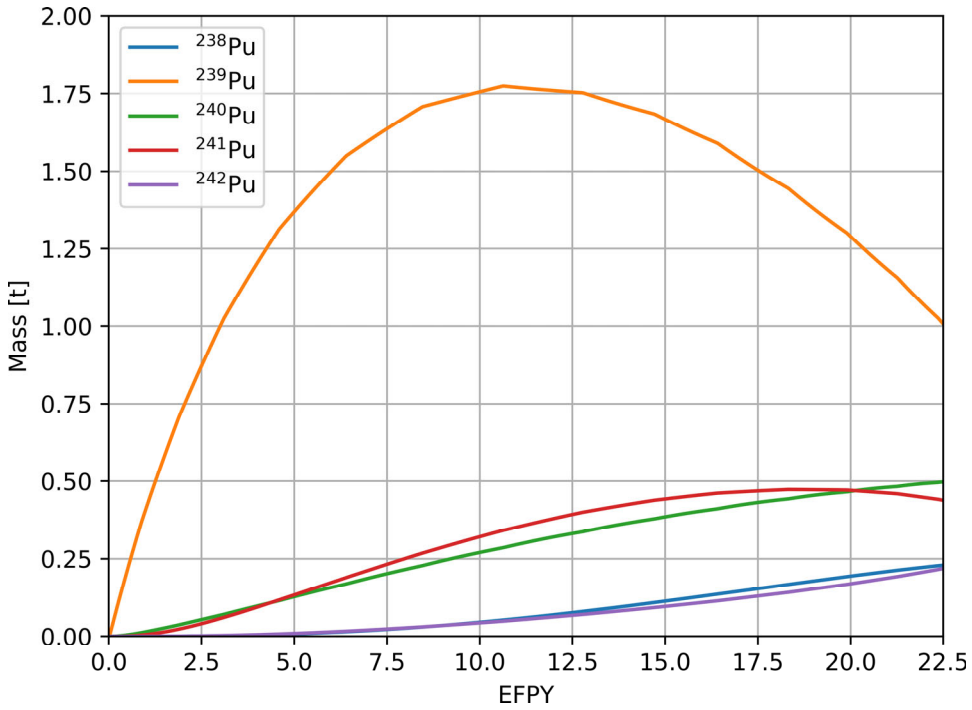

Figure 3.21: SaltProc-calculated plutonium isotopic fuel salt content during 22.5 years of operation.

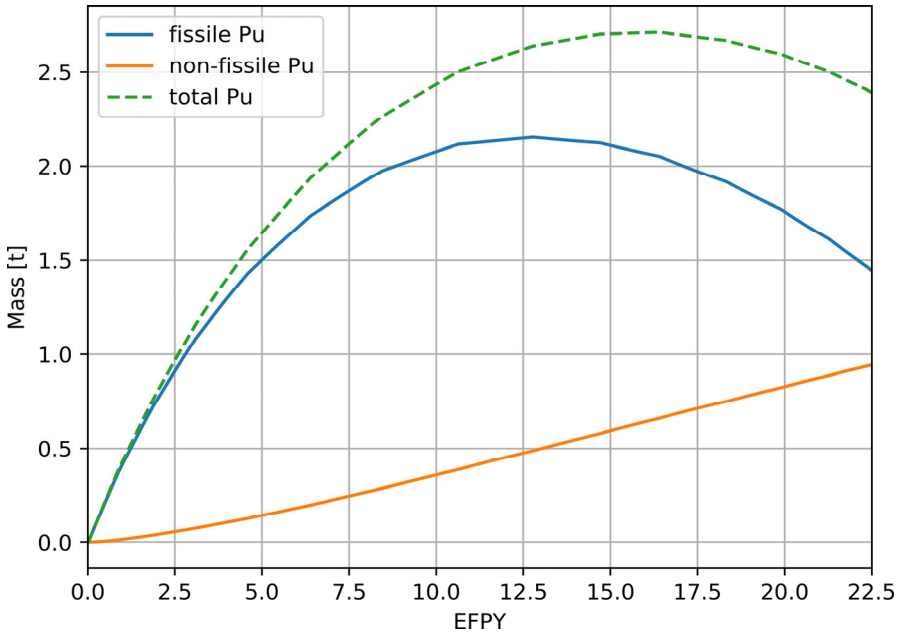

Figure 3.22: SaltProc-calculated fissile and non-fissile plutonium fuel salt content during 22.5 years of operation.

Table 3.6: Comparison of major heavy isotopes inventories at the EOL in the TAP reactor between the current work and Betzler et al. [18].

	Isotope	Current work mass [kg]	Betzler, 2017 mass [kg]	Δm [%]
Fissile	^{235}U	1299	1160	11%
	^{239}Pu	942	995	-5%
	^{241}Pu	427	435	-2%
	Total	2668	2590	3%
Non-fissile	^{236}U	1123	1200	-6%
	^{238}U	127,353	132,400	-4%
	^{238}Pu	235	280	-16%
	^{240}Pu	503	1000	-50%
	^{242}Pu	230	310	-26%
	Total	129,444	135,190	4%

SaltProc and compared with the reference. Normalized per GWth-year, the TAP concept requires about 5.23 t of fuel compared with 4.14 t reported by Betzler et al. SaltProc calculated waste production normalized per GWth-year is 5% less than reported by ORNL. Potentially, the TAP can operate with LWR SNF as the fissile material feed. The heavy metal component of LWR SNF has a lower fissile material weight fraction than 5% enriched uranium and adds less fertile ^{238}U to the fuel salt, potentially reducing the operational lifetime. Nevertheless, in the case of using waste material (e.g., transuranium elements from LWR SNF) in this fueling scenario, the TAP concept has superior waste reduction metrics.

Table 3.7: Comparison of normalized by GWth-year total fuel load and actinide waste from the TAP reactor obtained in the current work and Betzler et al. [18].

Parameter	Current work	Betzler, 2017 [18]
5% LEU feed rate [kg/y]	460.8	480.0
Loaded fuel [t per GW$_{th}$-y]	5.23	4.14
Waste [t per GW$_{th}$-y]	3.57	3.74

Neutron energy spectrum

Significant thermalization of the neutron spectrum is observed as moderator rods are added into the core configuration (Figure 3.23). At startup, the neutron spectra from the current work and Betzler et al. are matched well because the core geometry, its SVF, and initial fuel composition in these two simulations are similar. The Pearson correlation coefficient[7] r_{BOL} = 0.91115, which

[7] Pearson correlation coefficient is calculated by the following formula:

$$r = \frac{\sum_{i=1}^{N}(\Phi_i^{ref} - \overline{\Phi^{ref}})(\Phi_i - \overline{\Phi})}{\sqrt{\sum_{i=1}^{N}(\Phi_i^{ref} - \overline{\Phi^{ref}})^2 \sum_{i=1}^{N}(\Phi_i - \overline{\Phi})^2}}$$

where

$$\Phi_i^{ref}, \Phi_i = neutron\ flux\ for\ i^{th}\ energy\ bin\ reported\ in\ the\ reference\ and\ the\ current\ work\ [n/cm^2 \cdot s]$$

indicates a strong, positive association between the spectra at the BOL (see Figure 3.23, upper plot). At the EOL, SaltProc/Serpent-calculated spectrum is more thermal than reported by Betzler et al. [18], but the correlation coefficient r_{EOL}= 0.90987 shows that the spectra are still extremely strongly related (see Figure 3.23, lower plot).

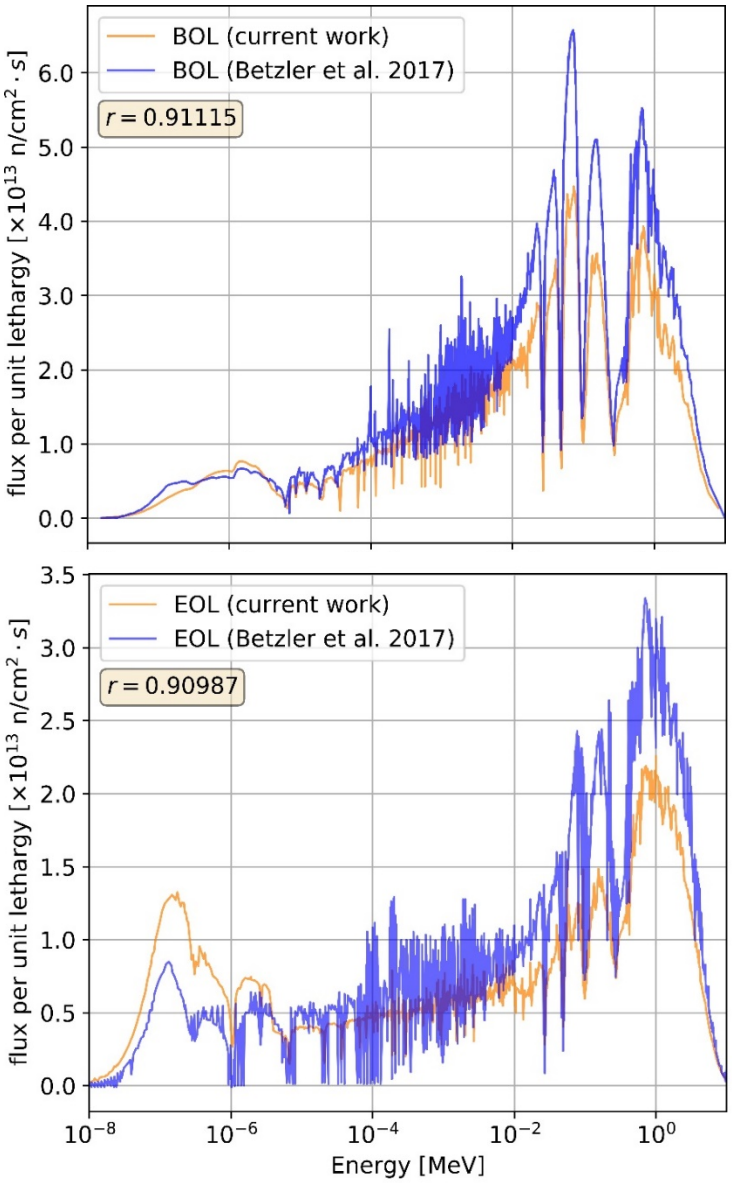

Figure 3.23: Neutron flux energy spectrum at the BOL (upper) and the EOL (lower) obtained using SaltProc/Serpent (orange) compared with ChemTriton/Shift (blue) [18].

The harder spectrum at the BOL tends to significantly increase resonance absorption in ^{238}U and decrease the absorptions in fissile and construction materials. Thus, the softer spectrum in the

$\overline{\Phi^{ref}}, \overline{\Phi}$ = *neutron flux avergaed over N energy bins reported in the reference and current work* $[n/cm^2 \cdot s]$

N = *number of neutron energy bins* $[-]$.

current work compared with Betzler et al. led to fewer resonance captures[8] of neutrons by ^{238}U, hence, less ^{239}Pu bred from ^{238}U. Therefore, the SaltProc/Serpent calculation in the current work underpredicts the destruction (i.e., fission and capture) of ^{235}U and overpredicts the destruction of ^{238}U (see

Table 3.6). Finally, the softer neutron spectrum leads to more fissions in fissile plutonium isotopes (^{239}Pu and ^{241}Pu) which also decreases non-fissile plutonium (

Table 3.6) and total actinide waste production (Table 3.7).

Time step refinement

The results shown in this section are obtained from SaltProc calculations with a uniform depletion time step of 3 days. The duration of the time step was chosen after performing a parametric sweep to determine the longest depletion time step that provides suitable calculation accuracy. A longer time step potentially reduces the SaltProc calculation costs, providing results faster for lifetime-long (25-year) simulations.

Figure 3.24 shows k_{eff} evolution obtained with 3-, 6-, 12-, and 24-day depletion time intervals for a 25-year simulation. The interval between moderator configuration updates was assumed similar for all four cases for consistency. The multiplication factor at the BOC for each moderator configuration reduced with increasing time step duration. At the End of Cycle (EOC) for each geometry, k_{eff} 1.0 for a 3-day time step but drops below 1.0 to 0.9980, 0.9972, and 0.9948 for 6-, 12-, and 24-day step, respectively. The decrease is because more poisonous FPs (e.g., ^{135}Xe) are produced in the core during longer depletion intervals. With longer time steps, a large concentration of poisons is obtained at the end of the depletion step when those poisons are being removed, resulting in substantial criticality growth.

Figure 3.25 and Figure 3.26 show that the longer time steps appropriately capture uranium depletion (1% difference even for a 24-day time step), but the observed difference in fissile ^{239}Pu mass is significant when the depletion interval is 6 days or longer (0.5% difference for 6-day step). Using a 6-day depletion interval leads to overprediction of ^{239}Pu production by 5 kg at the EOL (Figure 3.26). The use of a 6-day time step caused an overprediction of total plutonium production by 9.6 kg. Notably, significant quantity for plutonium currently in use by the IAEA is 8 kg (80% ^{238}Pu) [167]. Thus, a 6-day depletion interval or longer leads to significant error in the predicted plutonium inventory at the EOL (larger than 1 significant quantity).

Increasing the depletion time interval significantly reduces computational cost but also deteriorates the accuracy of depletion calculations (i.e., 24-day step gave 4 speedups but causes about 1.5% error in ^{239}Pu mass prediction). Calculations using a depletion time step of 6 days or more demonstrated a significant difference in calculated k_{eff} (i.e., 300 pcm for 6-day) and depleted mass (e.g., 0.34% error in ^{235}U predicted mass for 6-day) from those using a 3-day depletion step. In the current work, a 3-day depletion step was selected to adequately predict the mass of major heavy isotopes in the fuel salt during 25 years of the TAP reactor operation.

[8] The energy range for ^{238}U resonance neutron capture is between 10^{-5} and 10^{-2} MeV.

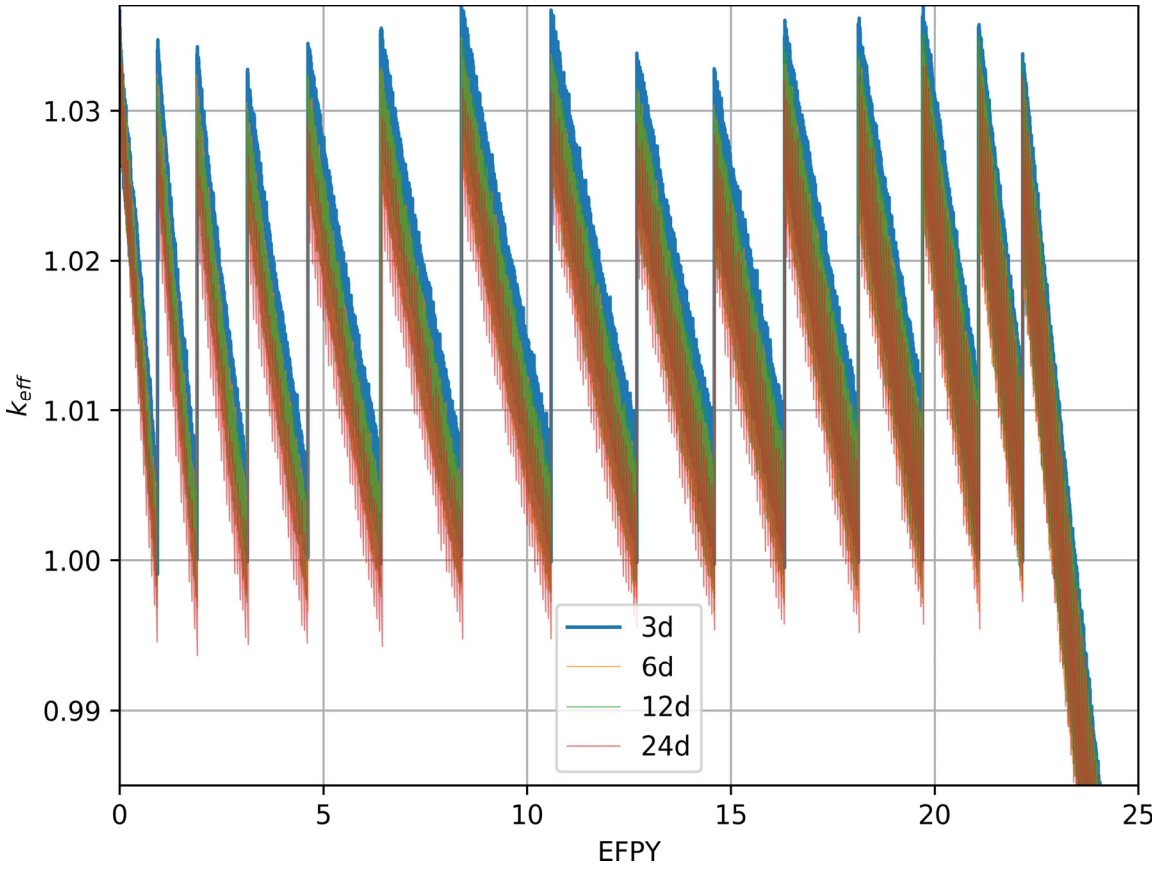

Figure 3.24: SaltProc-calculated effective multiplication factor (k_{eff}) during operation for different depletion time step sizes.

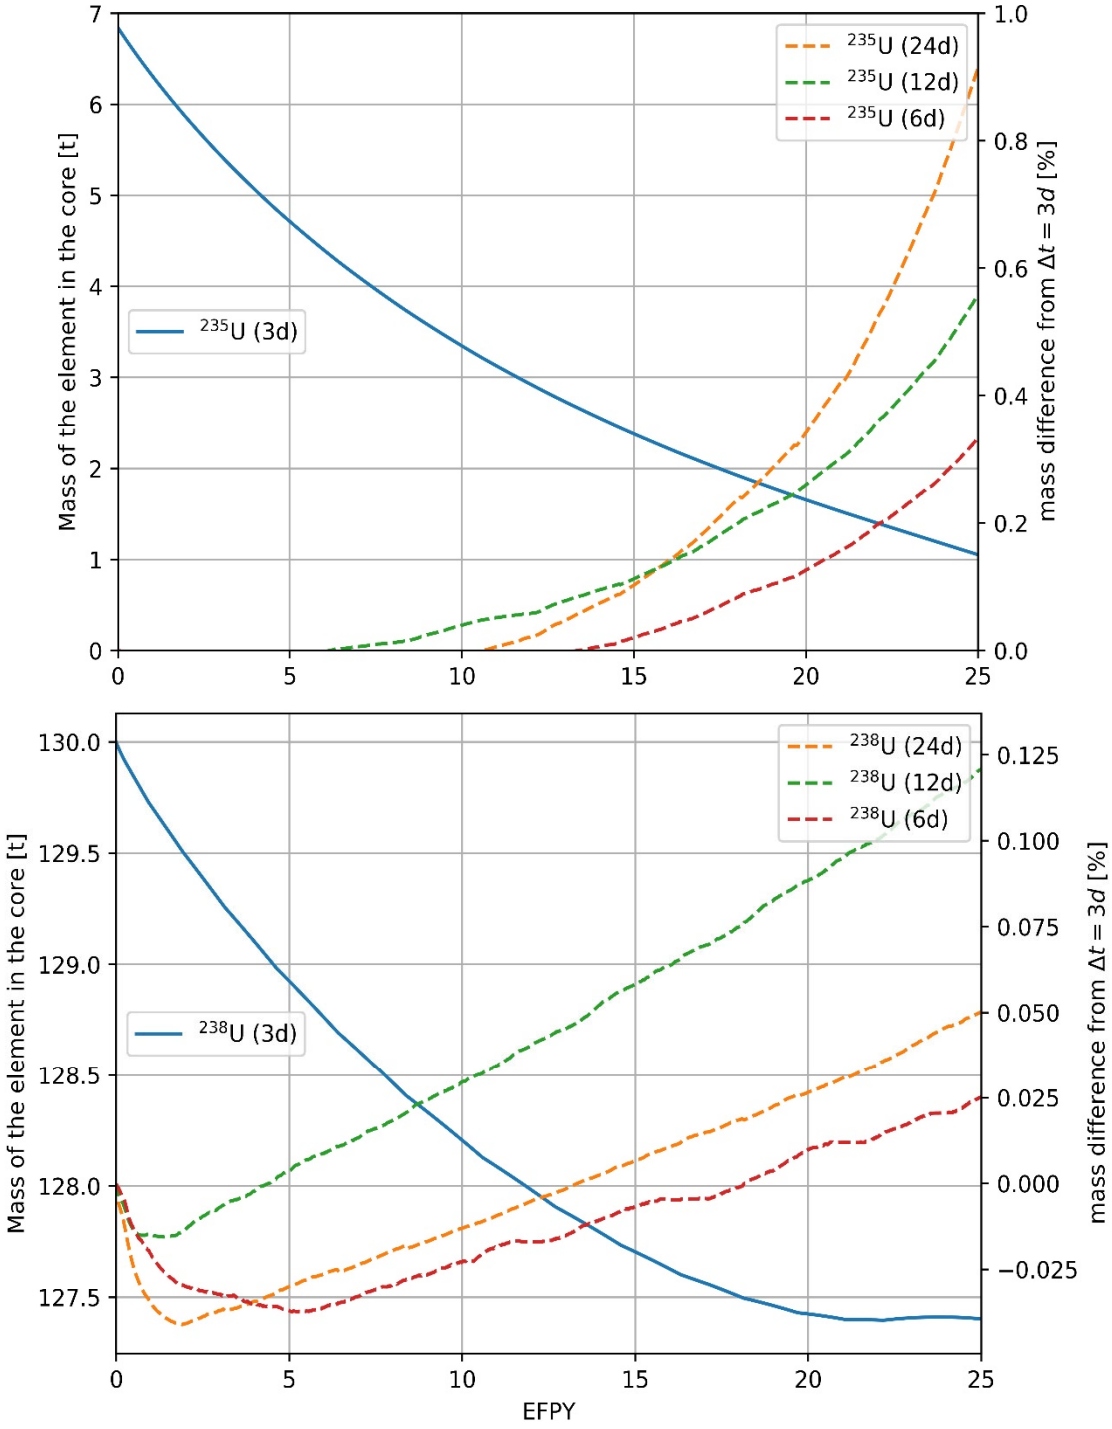

Figure 3.25: SaltProc-calculated ^{235}U (upper) and ^{238}U (lower) content during operation for different depletion time step sizes.

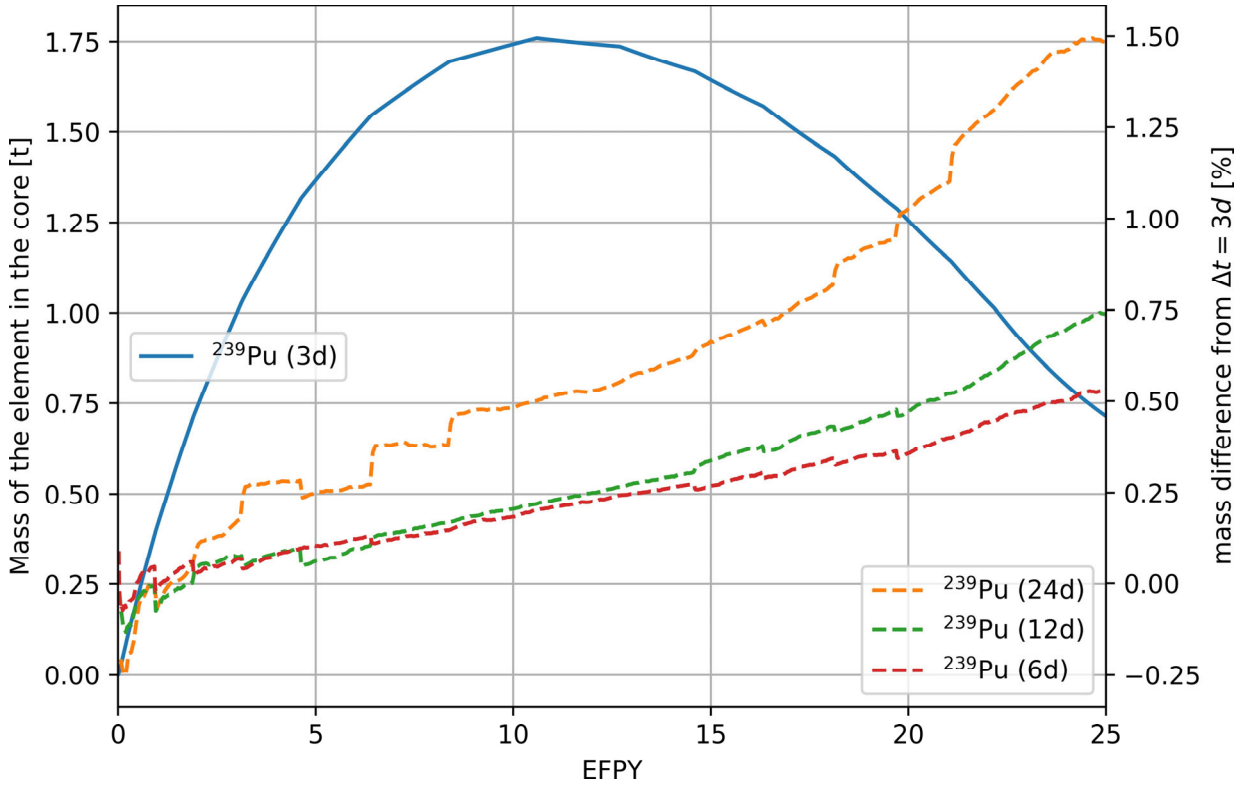

Figure 3.26: SaltProc-calculated ^{239}Pu content during operation for different depletion time step sizes.

3.3.2.2. Realistic Extraction Efficiency Case

This section demonstrates SaltProc v1.0 for lifetime-long depletion simulation similar to Section 3.3.2, but with realistic, physics-based correlations for noble gas removal efficiency. For the demonstration case herein, efficiency of xenon, krypton, and hydrogen extraction are determined using the model by Peebles et al. (Equation 3-3) discussed earlier in Section 3.2.1.1. The gas-liquid interfacial area per unit volume (a) to inform Equation 3-3 is a function of salt/gas flow rates and gas bubble diameter [168]:

$$a = \frac{6}{d_b} \frac{Q_{He}}{Q_{He} + Q_{salt}} \tag{3-6}$$

where

$$Q_{salt} = volumetric\ salt\ flow\ rate\ [m^3/s]$$
$$Q_{He} = volumetric\ helium\ flow\ rate\ [m^3/s]$$
$$d_b = helium\ bubble\ diameter\ [m].$$

Additionally, the following parameters inform Equation 3-3 for the prototypic sparger: (1) salt volumetric flow rate throughout the sparger $Q_{salt} = 2$ m^3/s; (2) sparging gas (helium) volumetric flow rate $Q_{He} = 0.1$ m^3/s; (3) helium bubble diameter $d_b = 0.508$ mm as advised by ORNL [1]; (4) sparger length L = 11 m; (5) sparger diameter D = 0.4 m (sparger cross section $A_c = 0.126\ m^2$).

The liquid phase mass transfer coefficient (K_L) selection presents a challenge since published information to inform Equation 3-3 is applicable for only laboratory-scale conditions [32]. Peebles et al. stated that Equation 3-3 is valid for K_L in a range from 1 to 100 ft/hr (from 0.0847 to 8.477 mm/s) [28]. For the demonstration case herein, 25-year depletion calculations are performed for KL of 0.0847, 2.1167, and 8.4667 mm/s to investigate the effect of noble gas removal efficiency on lifetime-long fuel depletion calculations.

The extraction efficiency is gas specific because solubility in the salt (Henry's law constant) is different for various gases. Table 3.8 reports the dimensionless Henry's law constant and corresponding calculated efficiency of noble gas (Xe, Kr, H) migration to the helium bubbles (ϵ_m) in the prototypic sparger for various mass transfer coefficients. Total separation efficiency (Table 3.8, last three columns) refers to the efficiency of extraction target gaseous elements after performing helium sparging in the sparger followed by separation of noble- gas-reach bubbles from the salt in the axial-flow centrifugal bubble separator [29].

Table 3.8: The noble gas extraction efficiency at working temperature T=627 ℃ calculated using Equation 3-3 [28] assuming salt volumetric flow rate $Q_{salt} = 2\ m^3/s$, helium volumetric flow rate $Q_{He} = 0.1\ m^3/s$, helium bubble diameter $d_b = 0.508\ mm$, and sparger volume $V = 1.4\ m^3$. The liquid phase mass transfer coefficient is varied in validity range [0.0847, 8.4667] mm/s.

Element	Henry's law constant (K_H) [-]	Efficiency of migration to He bubbles (ϵ_m) for K_L [mm/s]			Efficiency of total separation (ϵ)* for K_L [mm/s]		
		8.4667	2.1167	0.0847	8.4667	2.1167	0.0847
Xe	5.70E-05	0.963	0.5639	0.0327	0.9149	0.5357	0.031
Kr	2.80E-04	0.9595	0.563	0.0327	0.9115	0.5349	0.031
H	3.90E-03	0.9066	0.5499	0.0326	0.8613	0.5224	0.0309

*With axial-flow centrifugal bubble separator by Gabbard et al., which allows the bubble separation efficiency $\epsilon_{es} = 0.95$ [29]. Thus, total gas removal efficiency (ϵ) can be calculated as follows $\epsilon = \epsilon_m \times \epsilon_{es}$.

Effective multiplication factor dynamics

Figure 3.27 and Figure 3.28 demonstrate the effective multiplication factor dynamics (k_{eff}) during 25 years of operation with 15 various moderator rod configurations (cycles) described in Appendix D. SaltProc v1.0 coupled to Serpent calculated k_{eff} after removing fission products and feeding 5% LEU at the end of each depletion step (3 days as was determined in Section 3.3.2.1). Notably, the core went subcritical during the first cycle (startup moderator rod configuration) after 330 and 318 days for K_L = 8.4667 and 0.0847 mm/s, respectively.

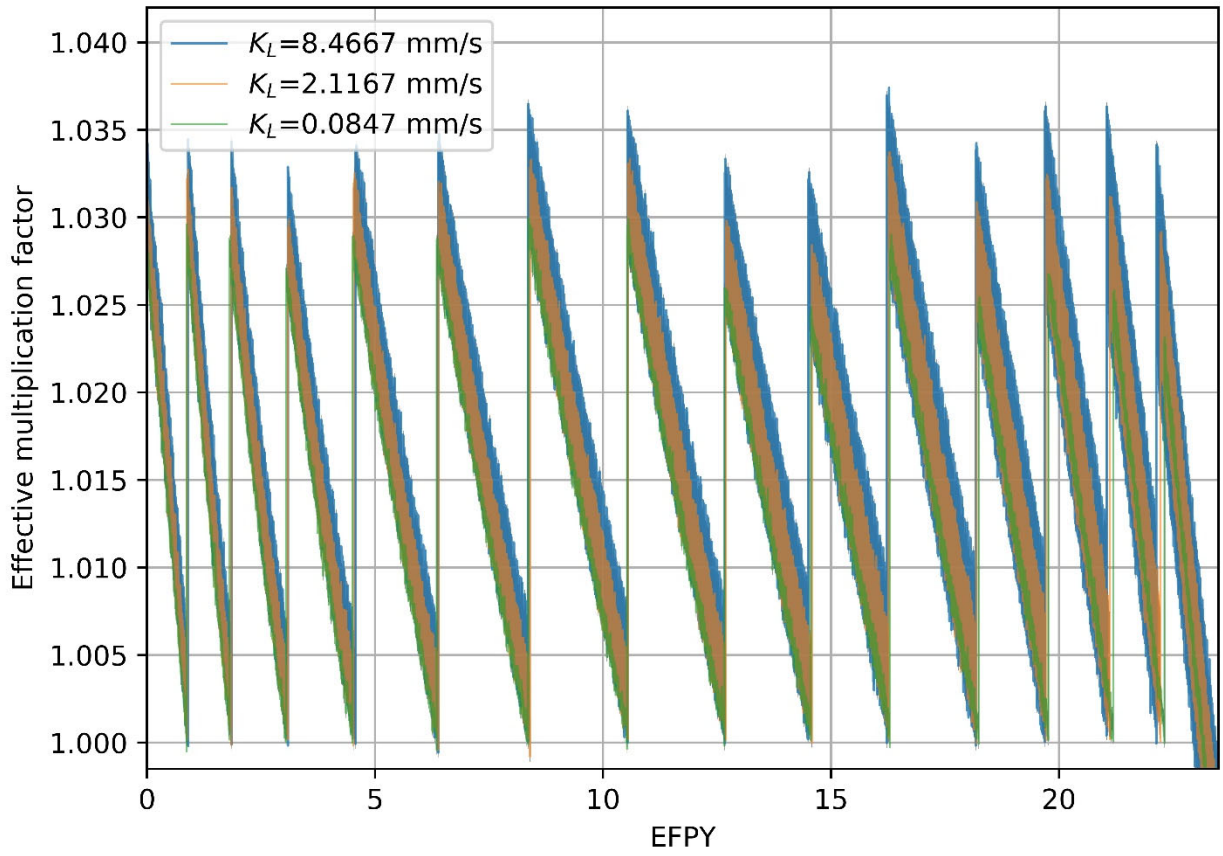

Figure 3.27: Effective multiplication factor dynamics for the full-core TAP core model during 25 years of operation for the case with a realistic removal efficiency of fission product and various mass transfer coefficients. Confidence interval σ=28 pcm is shaded.

A reduced mass transfer coefficient worsens the neutron poison efficiency, which shortens the interval between shutdowns for moderator rod updates. Additionally, the presence of unremoved poisons in the core suppresses the effective multiplication factor after moderator reconfiguration (≈500 pcm lower for K_L=0.0847 mm/s than for K_L=8.4667 mm/s at the BOL and 1100 pcm at the EOL). Overall, noble gas removal provides significant neutronics benefits (fewer neutrons are lost in strong absorbers such as ^{135}Xe), better fuel utilization, and enables longer moderator rod reconfiguration intervals.

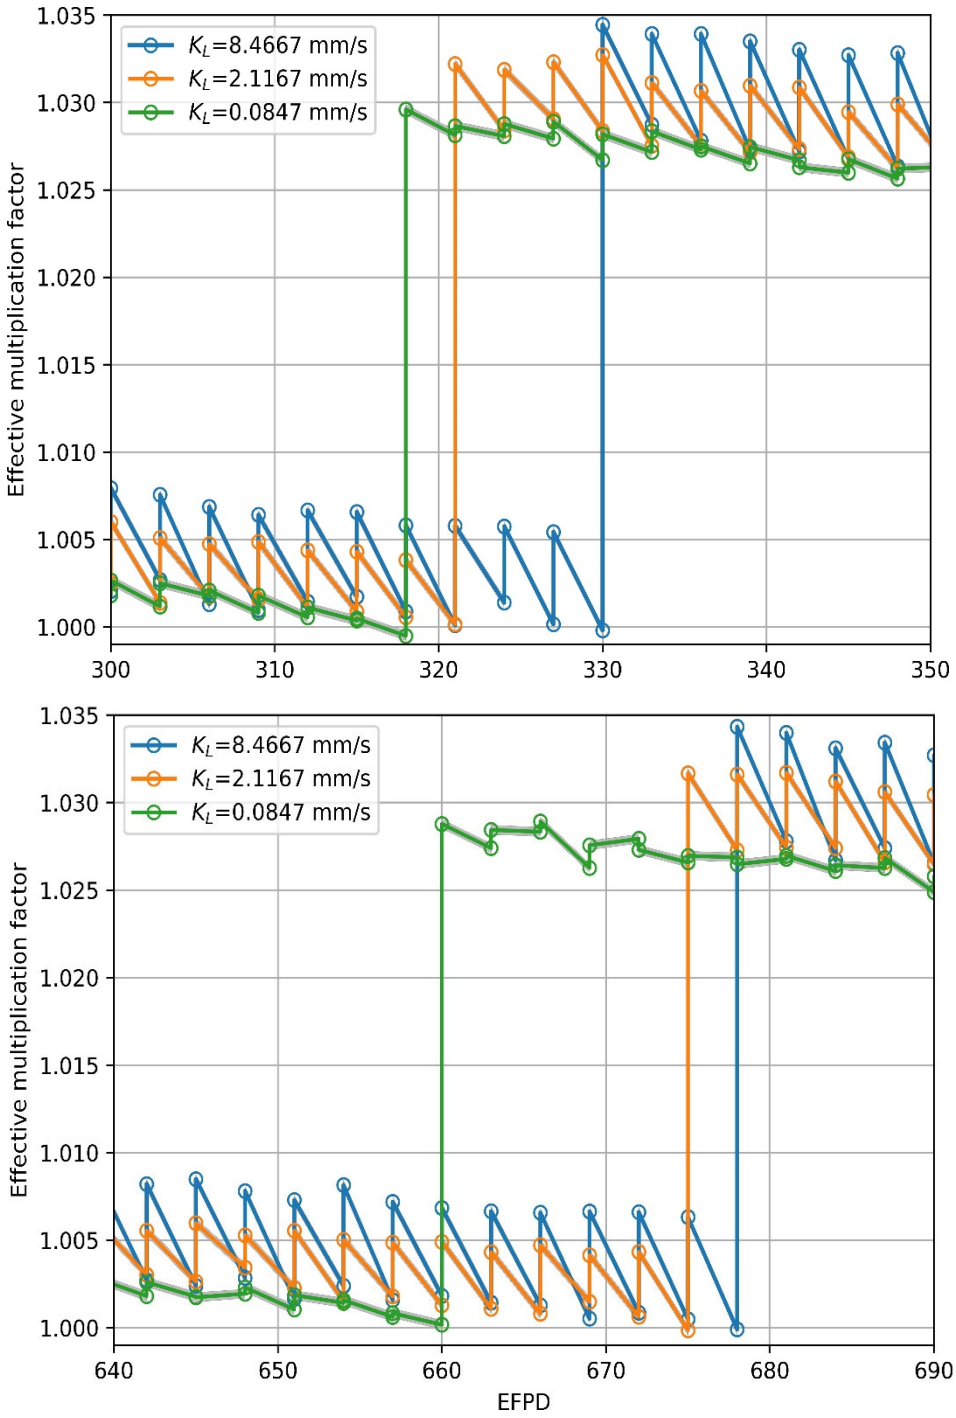

Figure 3.28: Zoomed effective multiplication factor dynamics while switching from Cycle #1 (startup geometry configuration, 347 moderator rods, SVF=0.917) to Cycle #2 (SVF=0.887) (upper panel) and from Cycle #2 to Cycle #3 (SVF=0.881) (lower panel) for various mass transfer coefficients (K_L). Confidence interval σ = 28 pcm is shaded.

Neutron spectrum

Figure 3.29 shows the normalized neutron flux spectrum for the full-core TAP core model in the energy range from 10^{-9} to 15 MeV. The neutron energy spectrum at the EOL is harder than at the BOL due to moderator-to-fuel ratio growth during reactor operation caused by periodic moderator rod reconfigurations. The TAP reactor spectrum is harder than in a typical LWR and correlates well (Pearson correlation coefficient 0.8) with the TAP neutronics white paper [15] and ORNL reports [16, 18]. The liquid phase mass transfer coefficient (K_L) and, consequently, noble gas removal efficiency (E), has a negligible effect on the spectrum in the fast range (between 10−2 and 10 MeV) at the EOL.

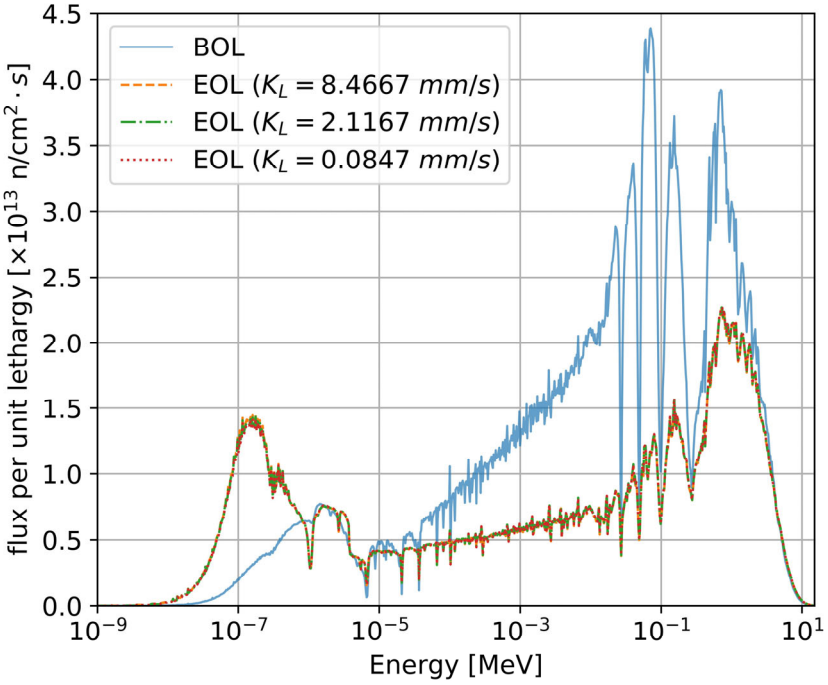

Figure 3.29: The neutron flux energy spectrum normalized by unit lethargy at the BOL and EOL for the case with a realistic removal efficiency of fission product and various mass transfer coefficients.

However, Figure 3.30 demonstrates a notable difference in the thermal range of the spectrum due to the enormous ^{135}Xe absorption cross section ($\sigma_{a,^{135}Xe} = 2.6 \times 10^6\ b$). Figure 3.35 shows that ^{135}Xe mass in the core at the EOL for the case with low noble gas removal efficiency (K_L = 0.0847 mm/s) is significantly larger than for the case with high removal efficiency (K_L = 8.4667 mm/s) which leads to higher neutron loss due to absorption in xenon. Overall, noble gas removal from the fuel salt alters the neutron spectrum.

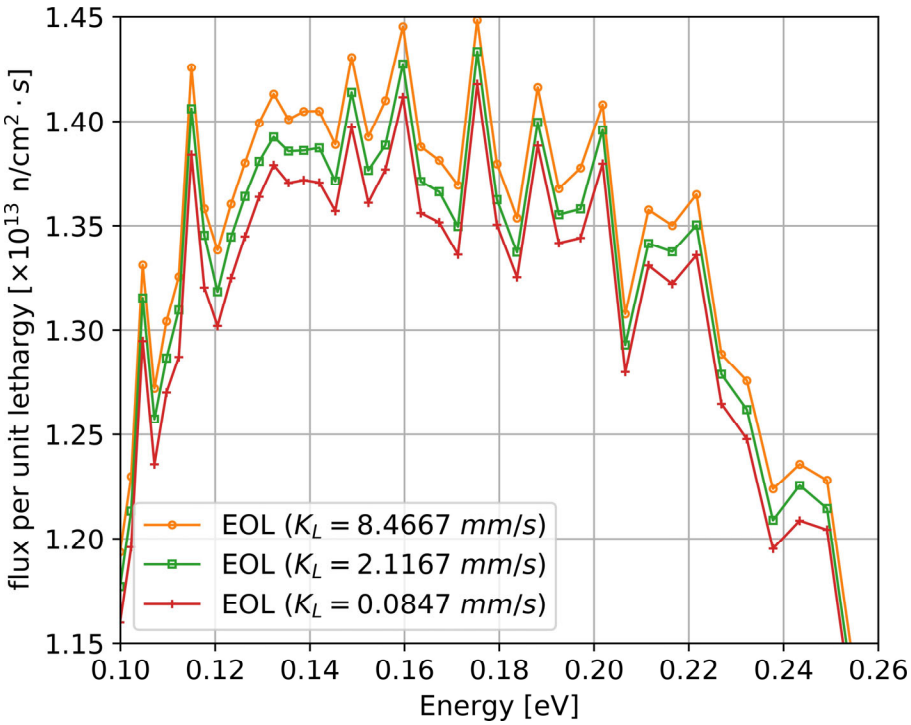

Figure 3.30: The neutron flux energy spectrum normalized by unit lethargy EOL zoomed in the thermal energy range.

Fuel salt isotopic composition evolution

The time-dependent isotopic compositions obtained with different noble gas extraction efficiencies behave very similarly. For ^{235}U predicted mass, the difference between K_L = 8.4667 mm/s (e.g., 91.5% of ^{135}Xe is removed) and K_L = 0.0847 mm/s (e.g., 3.1% of ^{135}Xe is removed) is within 0.2% for the first 14 years and rises rapidly to 1.15% over the remaining 10 years (Figure 4.24). The simulations with a mass transfer coefficient smaller than 8.4467 mm/s retain more ^{235}U during operation because more neutrons are parasitically absorbed by the noble gas, which leads to a lower fission rate. The relative mass difference in ^{238}U is small (Figure 3.32), but the absolute difference is approximately 50 kg at the EOL, with low removal efficiency corresponding to a reduced EOL inventory of ^{235}U.

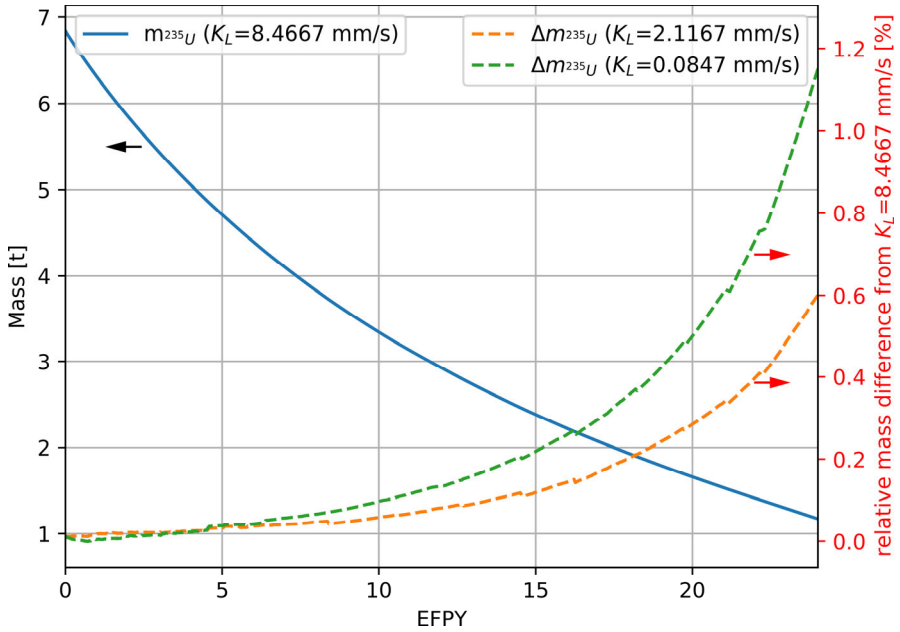

Figure 3.31: SaltProc-calculated mass of ^{235}U in the fuel salt during 25 years of operation for K_L = 8.4667 mm/s compared with less effective noble gas removal.

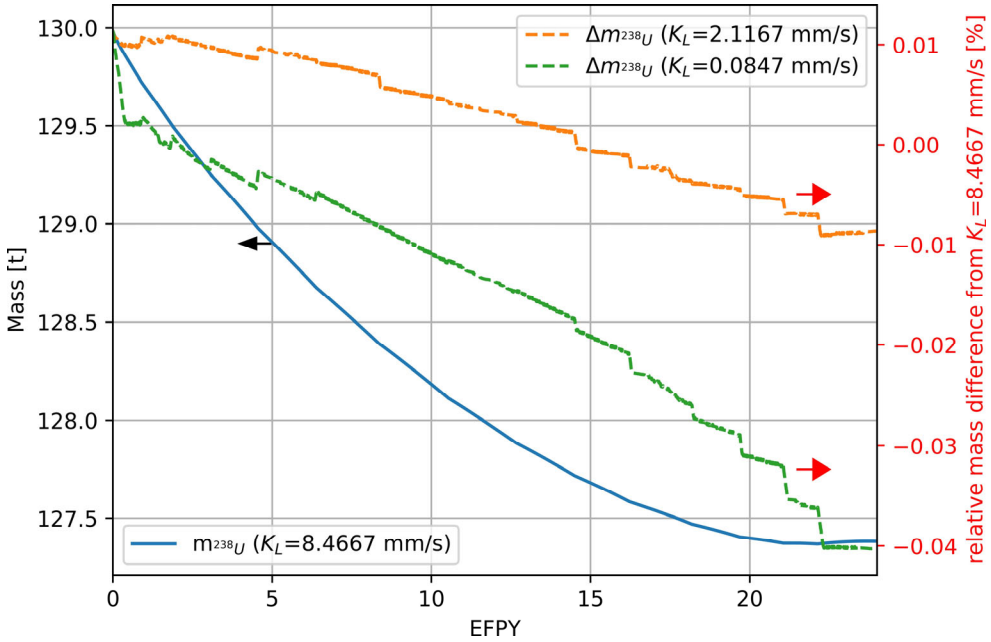

Figure 3.32: SaltProc-calculated mass of ^{238}U in the fuel salt during 25 years of operation for K_L = 8.4667 mm/s compared with less effective noble gas removal.

Differences in the plutonium production between cases with different gas removal efficiencies are much greater. Over 3% more ^{239}Pu mass is generated in the case with K_L = 0.0847 mm/s than with K_L = 8.4667 mm/s (Figure 3.33). The greater mass of neutron poison (^{135}Xe) in the core leads to a harder spectrum (Figure 3.30), which results in a faster rate of destruction of ^{238}U and increased breeding of fissile ^{239}Pu.

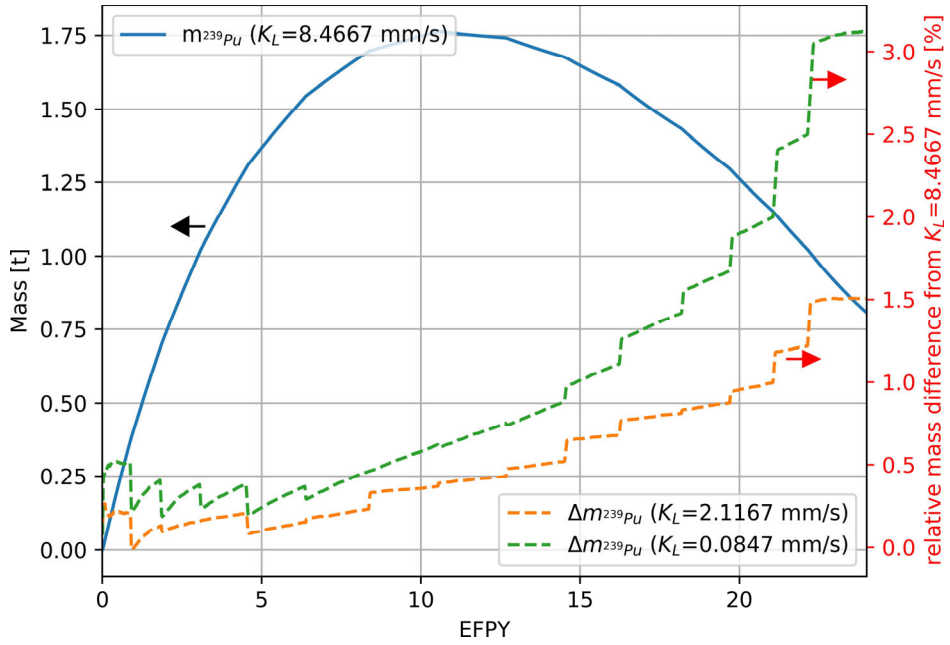

Figure 3.33: SaltProc-calculated mass of ^{239}Pu in the fuel salt during 25 years of operation for K_L = 8.4667 mm/s (91.5% of ^{135}Xe is removed) compared with less effective noble gas removal.

Figure 3.34 demonstrates ^{135}Xe mass dynamics in the TAP core during 25 years of operation for various mass transfer coefficients. Jumps in ^{135}Xe mass every few years reflect the spectral shifts due to moderator rod reconfiguration. In contrast, the mass of ^{135}I, which is the primary direct precursor of ^{135}Xe, is approximately 18 g and stays almost constant over 25 years.

Figure 3.35 shows ^{135}Xe mass at the end of each depletion time step before and after performing the fuel salt reprocessing procedure in SaltProc v1.0. ^{135}Xe concentration in the core after performing FP removals behaves as expected and is consistent with calculated extraction efficiencies in Table 3.8. Notably, the ^{135}Xe production rate increases during the first seven years of operation and then decreases rapidly to 17 g during the remaining 17 years as the spectrum thermalizes during operation.

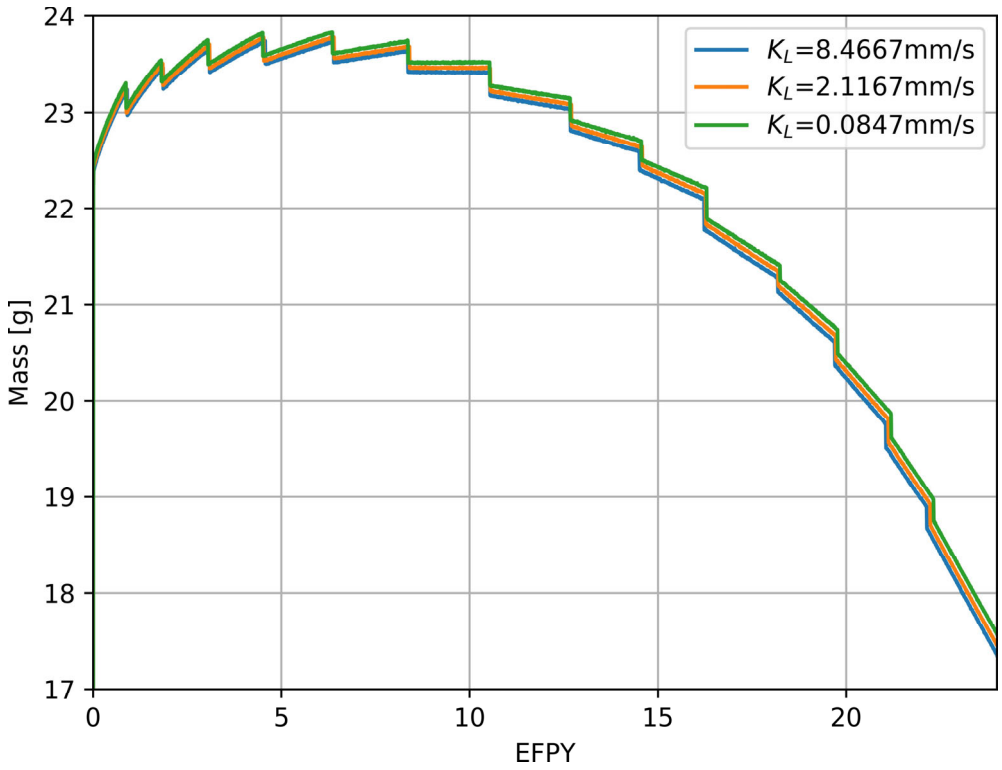

Figure 3.34: SaltProc-calculated mass of ^{135}Xe in the fuel salt during 25 years of operation for the case with a realistic removal efficiency of fission product and various mass transfer coefficients (K_L).

An analytic verification is also performed to confirm SaltProc v1.0 correctness by comparing the mass of ^{135}Xe to the expected mass after performing removals after each depletion step with realistic efficiency (Table 3.8). The expected mass of a reprocessed isotope is calculated as follows:

$$m_a = m_b \times (1 - \epsilon_m) \times (1 - \epsilon_{es}) \tag{3-7}$$

where

$$m_a = mass\ of\ the\ isotope\ after\ applying\ removals\ and\ feeds\ [g]$$
$$m_b = mass\ of\ the\ isotope\ right\ before\ reprocessing\ [g]$$
$$\epsilon_m = efficiency\ of\ the\ isotope\ migration\ to\ helium\ bubbles\ [-]$$
$$\epsilon_{es} = entrainment\ separator\ extraction\ efficiency\ [-].$$

This simple check showed that the SaltProc-calculated mass of ^{135}Xe (Figure 3.35) matches the expected mass exactly. Thus, SaltProc v1.0 extraction module correctly removes target isotopes with a specified extraction efficiency. Finally, this correctness check is added as SaltProc v1.0 unit test.

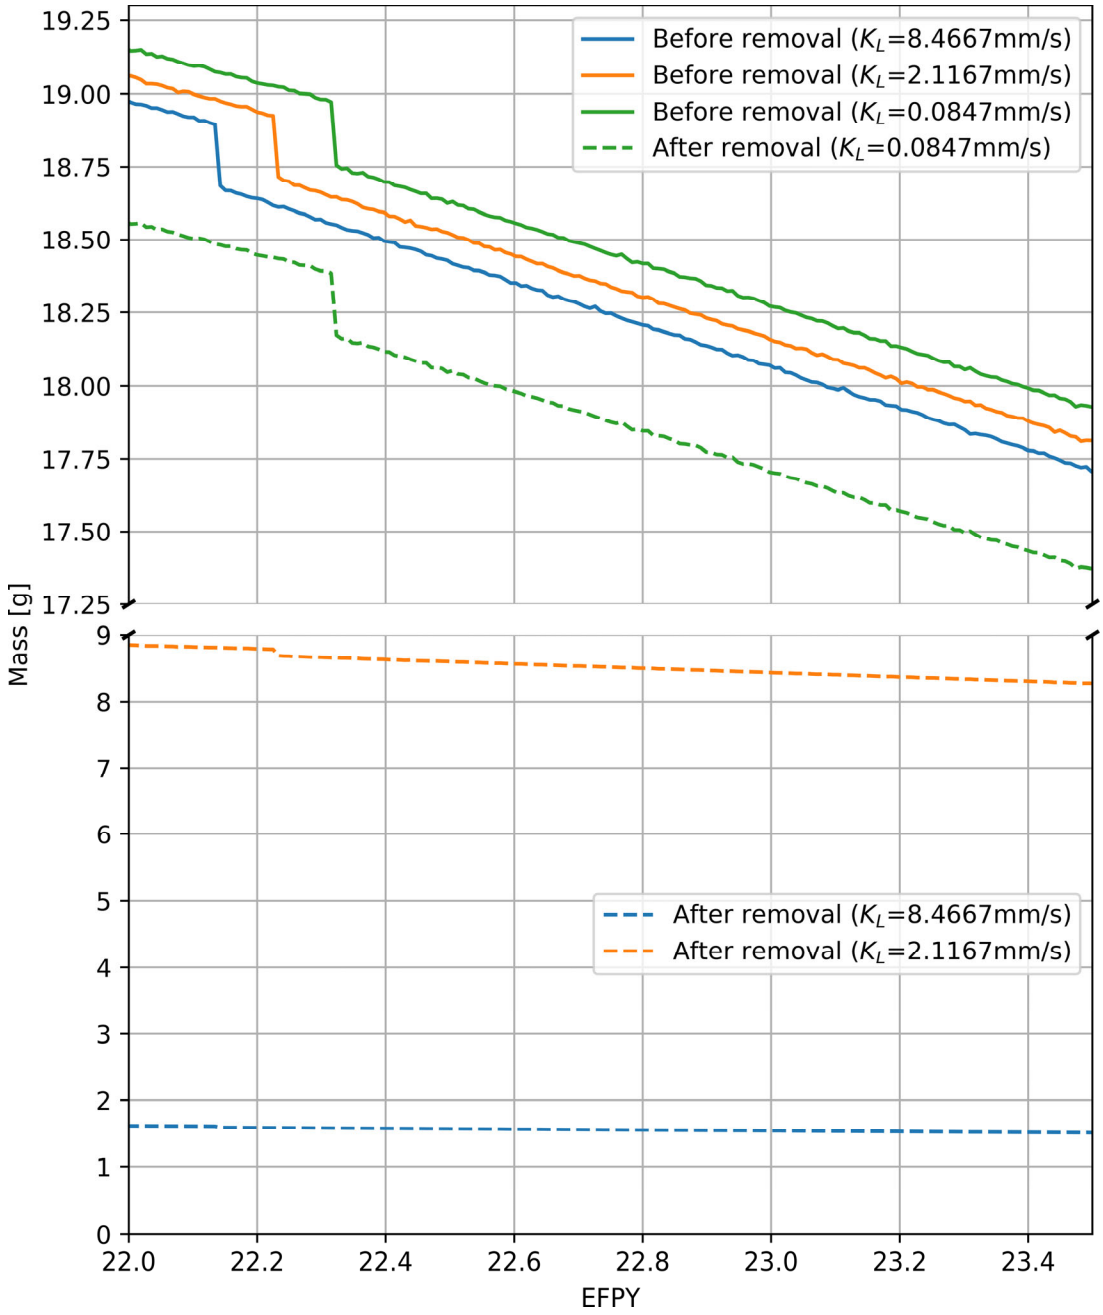

Figure 3.35: SaltProc-calculated mass of ^{135}Xe in the fuel salt during the last 18 months of operation for various mass transfer coefficients (K_L) at the end of each depletion step before and after performing the salt treatment.

3.3.3. Safety and Operational Parameters

The previous section (Section 3.3.2) reported fuel salt composition evolution during 25 years of TAP MSR operation. The inventory of fissile ^{235}U decreased with time, while the inventories of fissile, ^{239}Pu and ^{241}Pu, increased. At the same time, many poisonous actinides (e.g., ^{236}U, ^{240}Pu, ^{242}Pu) built up in the core, shifting the neutron energy spectrum. Moreover, the TAP design

assumes an intentional spectrum shift by adding more moderator rods during operation. In this section, analysis is performed on how such neutron spectrum shift affects major safety and operational parameters such as temperature and void coefficients of reactivity, total control rod worth, and other reactor kinetic parameters.

3.3.3.1. Temperature Coefficient of Reactivity

The main physical principle underlying reactor temperature feedback is an expansion of heated material. When the fuel salt temperature increases, the density of the salt decreases, but at the same time, the total volume of fuel salt in the core remains constant because it is bounded by the vessel. When the moderator rod temperature increases, the density of zirconium hydride decreases, reducing space between moderator rods and displacing fuel salt from the core. Another physical principle underlying temperature feedback is the Doppler broadening of the resonance capture cross section of the ^{238}U due to thermal motion of target nuclei in the fuel. The Doppler effect arises from the dependence of the capture cross sections on the relative velocity between neutron and nucleus. The Doppler coefficient of reactivity of thermal reactors is always negative and instantaneous.

The temperature coefficient of reactivity, α, quantifies reactivity changes due to temperature change in fuel and moderator component of a reactor core. The $\alpha_{T,j}$ represents the temperature coefficient of reactivity of a component j (fuel, moderator, or isothermal) and can be calculated as:

$$\alpha_{T,j} = \frac{\partial \rho}{\partial T_j} \; [pcm/K] \tag{3-8}$$

where

$$\rho = \frac{k_{eff} - 1}{k_{eff}} \times 10^5 \; [pcm] \tag{3-9}$$

$k_{eff} =$ *effective multiplication factor corresponding to T of component j* $[-]$
$\partial T_j =$ *change in average temperature of component j* $[K]$.

If the temperature change is assumed to be uniform throughout the core, the temperature coefficient of reactivity is usually called Total or Isothermal Temperature Coefficient (ITC), $\alpha_{T,ISO}$, and can be defined as the change in reactivity per unit of temperature change:

$$\alpha_{T,ISO} = \frac{\Delta \rho}{\Delta T} \; [\mathrm{pcm/K}] \tag{3-10}$$

where

$\Delta \rho =$ *change in reactivity* $[pcm]$
$\Delta T_j =$ *change in average temperature of the core* $[K]$.

However, fuel and moderator temperature are rarely equal because fuel heats up much faster than the moderator; thus, the fuel temperature coefficient ($\alpha_{T,F}$ or FTC) and the moderator temperature coefficient ($\alpha_{T,M}$ or MTC) must be calculated separately. In the base case simulation in this work, the fuel salt and the moderator temperatures are fixed at 900K. To determine $\alpha_{T,F}$, The fuel salt temperature is perturbed from 800K to 1000K in increments of 50K while fixing the

moderator temperature at 900K (base case). Likewise, $\alpha_{T,M}$ is calculated by perturbing the moderator temperature from 800K to 1000K with 50K increments, while fixing the fuel temperature at 900K.

The range of temperature perturbation for the temperature coefficient calculation has been selected based on operational parameters. The TAP MSR operates in the range of 773-973K (500-700°C), which is far below the salt boiling point of approximately 1473K [14]. The salt freezes below 773K [169]. At the other end of the temperature spectrum, the temperature higher than 973K passively melts a freeze plug, which drains the fuel salt from the reactor vessel to the drain tanks. The drain tanks have a subcritical configuration with a large free surface area to readily dissipate heat by passive cooling [14]. Thus, calculating temperature coefficients in the temperature range from 800 to 1000K captured the outcomes of most accident transients.

To determine the temperature coefficients, the cross-section temperatures for the fuel and moderator were changed in the range of 800-1000K. For $\alpha_{T,F}$ calculation, changes in the fuel temperature impact cross section resonances (Doppler effect) as well as the fuel salt density. The density of fuel salt changes with respect to temperature as follows [170]:

$$\rho_{salt}(T) = 6.105 - 12.72 \times 10^4 T\ [K] \quad [g/cm^3] \tag{3-11}$$

The uncertainty in the salt density calculated using Equation 3-11 is approximately 0.036 g/cm^3 at 900K. In contrast, when the moderator temperature changes, the density, cross section temperature, and the geometry also change due to thermal expansion of the solid zirconium hydride ($ZrH_{1.66}$) rods. Accordingly, the new moderator density and sizes are calculated using a linear temperature expansion coefficient [171]:

$$\alpha_L = 2.734 \times 10^{-5}\ [K^{-1}] \tag{3-12}$$

Using this thermal expansion data, it is taken into account that the displacement of the moderator surfaces by generating corresponding geometry definitions for each Serpent calculation. That is, $\alpha_{T,M}$ calculation takes into account the following factors:

- thermal Doppler broadening of the resonance capture cross sections in $ZrH_{1.66}$;
- hydrogen S(α, β) thermal scattering data shift due to moderator temperature change;
- density change due to moderator thermal expansion/contraction;
- corresponding geometric changes in the moderator rod diameter and length.

By propagating the k_{eff} statistical error provided by Serpent 2, the corresponding uncertainty in each temperature coefficient is obtained using the formula:

$$\delta\alpha_T = \left|\frac{1}{T_{i+1} - T_i}\right| \sqrt{\frac{\delta k_{eff}^2(T_{i+1})}{k_{eff}^4(T_{i+1})} + \frac{\delta k_{eff}^2(T_i)}{k_{eff}^4(T_i)}} \tag{3-13}$$

where

$$k_{eff} = effective\ multiplication\ factor\ corresponding\ to\ T_i\ [-]$$
$$\delta k_{eff} = statistical\ error\ for\ k_{eff}\ from\ Serpent\ output\ [pcm]$$
$$T_i = perturbed\ temperature\ in\ the\ range\ of\ 800 - 1000K.$$

Notably, other sources of uncertainty are neglected, such as design parameter uncertainty, cross section measurement error[9], and approximations inherent in the equations of state providing both the salt and moderator density dependence on temperature.

Figure 3.36 shows reactivity as a function of fuel, moderator, and total temperature for the TAP MSR at the BOL and EOL. At startup, reactivity change with temperature clearly fits linear regression (R-squared [10] is 0.9, 0.99, and 0.98 for fuel, moderator, and isothermal case, respectively). Also, while the linear relationship between reactivity and moderator temperature worsens toward the EOL, an $R^2>0.7$ still indicates a strong linear association between ρ and T (R^2 is 0.99, 0.87, and 0.74 for fuel, moderator, and total case, respectively). It is determined the temperature coefficient of reactivity separately for each component (fuel, moderator, and isothermal) using the slope of the linear regression for each.

Table 3.9 summarizes temperature coefficients of reactivity in the TAP core calculated at the BOL and EOL. The fuel temperature coefficient is negative throughout operation and becomes stronger toward the EOL as the spectrum thermalizes due to additional, retained fission products and actinides building up in the fuel salt. The MTC and ITC are both strongly negative at startup. However, the MTC became weakly positive toward the EOL due to the same spectral shift. To better understand the dynamics of temperature coefficient evolution, the temperature coefficients are calculated for 15 distinct moments during operation to cover all moderator rod configurations described in Appendix D.

Table 3.9: Temperature coefficients for the TAP reactor at the BOL and EOL.

Coefficient	BOL [pcm/K]	EOL [pcm/K]
FTC	-0.350 ± 0.050	-0.868 ± 0.045
MTC	-1.134 ± 0.050	+0.746 ± 0.045
ITC	-1.570 ± 0.050	-0.256 ± 0.045

Figure 3.37 shows temperature coefficient evolution for the TAP reactor during 25 years of operation and takes into account the spectral shift due to moderator rod reconfigurations. The fuel temperature coefficient is almost constant for 19 years but decreases for the last 6 years (configurations with 1498 and 1668 moderator rods in the core). In contrast, the moderator temperature coefficient decreases from 1.134 pcm/K to 2.280 pcm/K during the first 11 years and then increases up to 0.746 pcm/K at the EOL. The moderator temperature increase at startup pushes thermal neutrons to higher energies, nearly up to the lowest ^{238}U resonances in the capture cross section. After 11 years, similar moderator temperature increases shifts neutrons into the same energy region, but this time that energy range is populated not only with ^{238}U but also with low-lying resonances from the actinides and fission products.

Additionally, the moderator temperature coefficient increases after 11 years of operation because there is twice as much moderator in the core at 11 years compared to the BOL. The

[9] Chapter 7 of [162] presents uncertainty quantification method for propagating cross section measurement uncertainty throughout depletion calculations. While it is out of scope of this work to estimate nuclear-data related uncertainty of the temperature feedback coefficient, method from [162] can be adopted for the future work to perform such calculations.

[10] Coefficient of determination (R^2) is a statistical measure of how well measured data fit linear regression line.

moderator temperature increase causes fuel salt displacement due to the thermal expansion of the moderator rods, which has a particularly strong effect when the salt volume fraction is less than 75%. That is, when moderator heats up, the moderator-to-fuel ratio increases due to thermal expansion of zirconium hydride, which in turn leads to positive change in reactivity.

Finally, the isothermal temperature coefficient dynamics are similar to the MTC: the ITC decreases from 1.57 pcm/K to 2.66 pcm/K first 13 years of operation. After that, the ITC grows rapidly up to 0.256 pcm/K at the EOL. Overall, the ITC remains negative throughout operation but became relatively weak after 25 years of operation (comparing with conventional Pressurized Water Reactor (PWR), which has an isothermal temperature coefficient of $\alpha_{T,ISO} \approx -3.08$ pcm/K [172]).

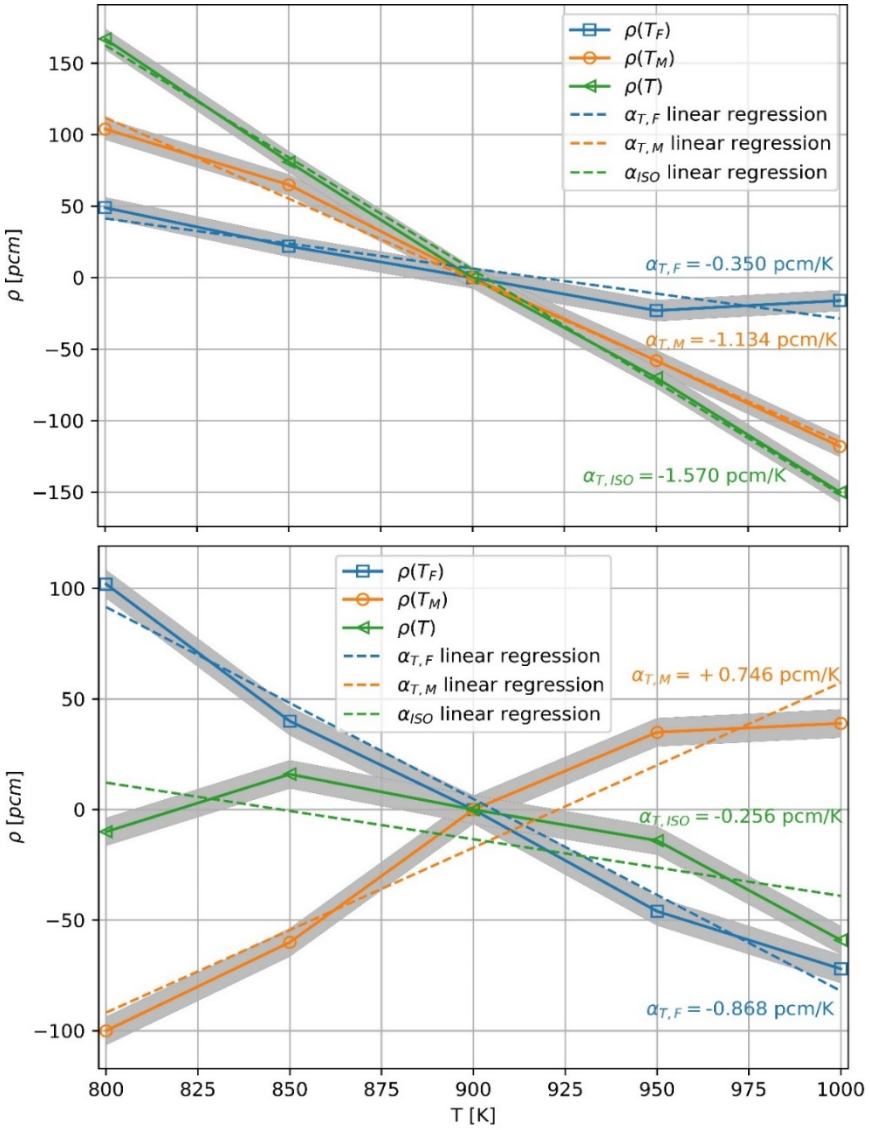

Figure 3.36: Serpent-calculated reactivity as a function of fuel salt (blue), moderator (orange), and both fuel/moderator (green) temperature at BOL (upper) and EOL (lower). The uncertainty ±σ region is shaded.

3.3.3.2. Void Coefficient of Reactivity

The effect of fuel voids (i.e., bubbles) on reactivity is evaluated by reducing the fuel salt density from the base value (0% void) assuming helium volume fraction in the salt varies between 0 and 2%. The temperatures of both the fuel salt and the moderator are held constant at 900K. Because a decrease in the salt density causes an increase of moderator- to-fuel ratio, an increase in the helium volume fraction (voids) increases reactivity as shown in Figure 3.38. However, the slope of the line (void coefficient of reactivity (αv)) decreases toward EOL due to the gradually increasing volume of moderator in the core (the volume fraction of the fuel salt at the EOL is less than 54%).

Figure 3.38 shows the void coefficient evolution during 25 years of operation, taking into account 15 moderator rod reconfigurations. The positive void coefficient of reactivity, though not ideal, does not compromise the reactor safety, if fuel density change resulted be coupled to a change in temperature. And, while some void fraction fluctuations may happen due to gaseous fission product production, their generation rates are usually almost constant. However, a large volume of sparging gas (helium) can be accidentally introduced into the TAP core in case of the bubble separator malfunction. Thus, the bubble separator must have backup safety mechanism to avoid sudden positive negativity insertion in case of the separator failure, particularly at the BOL. These observations from calculating reactivity coefficients should be taken into account in the TAP MSR accident analysis and safety justification.

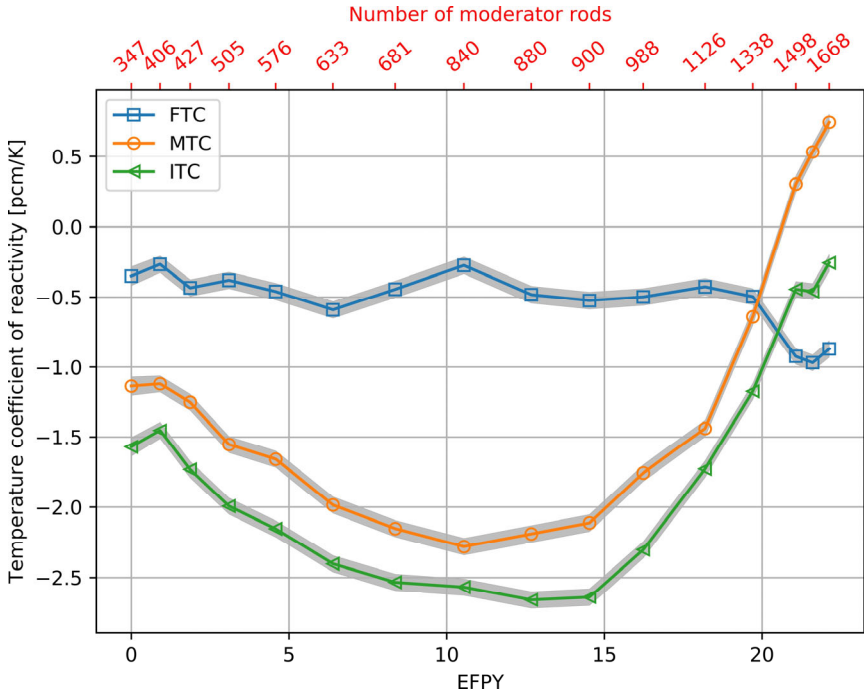

Figure 3.37: Serpent-calculated fuel, moderator, and isothermal temperature coefficients of reactivity as a function of time and number of moderator rods in the TAP core. The uncertainty ±σ region is shaded.

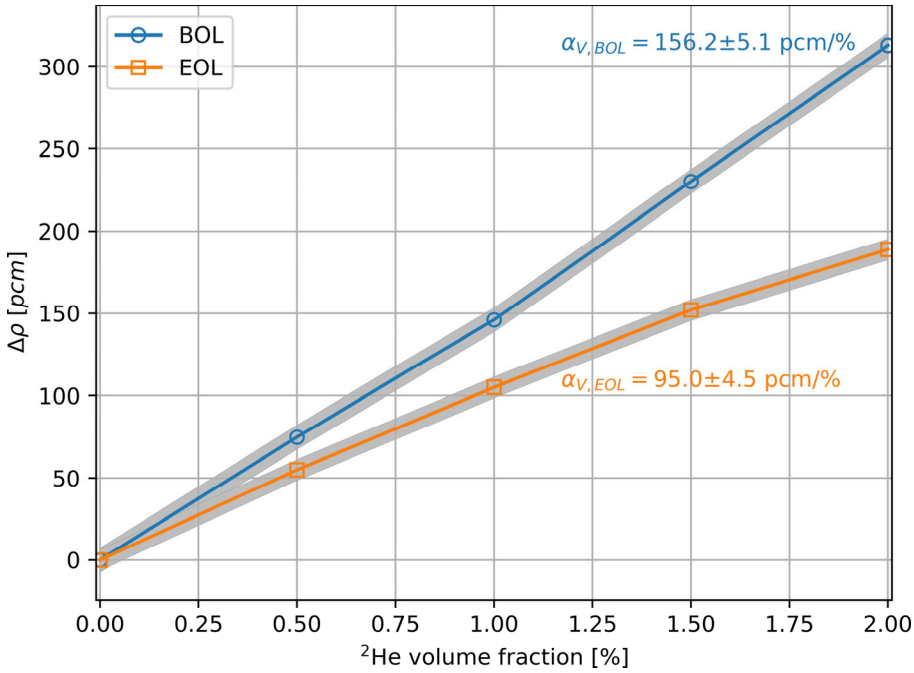

Figure 3.38: Serpent-calculated reactivity as a function of void volume fraction [%] in the fuel salt. The uncertainty ±σ region is shaded.

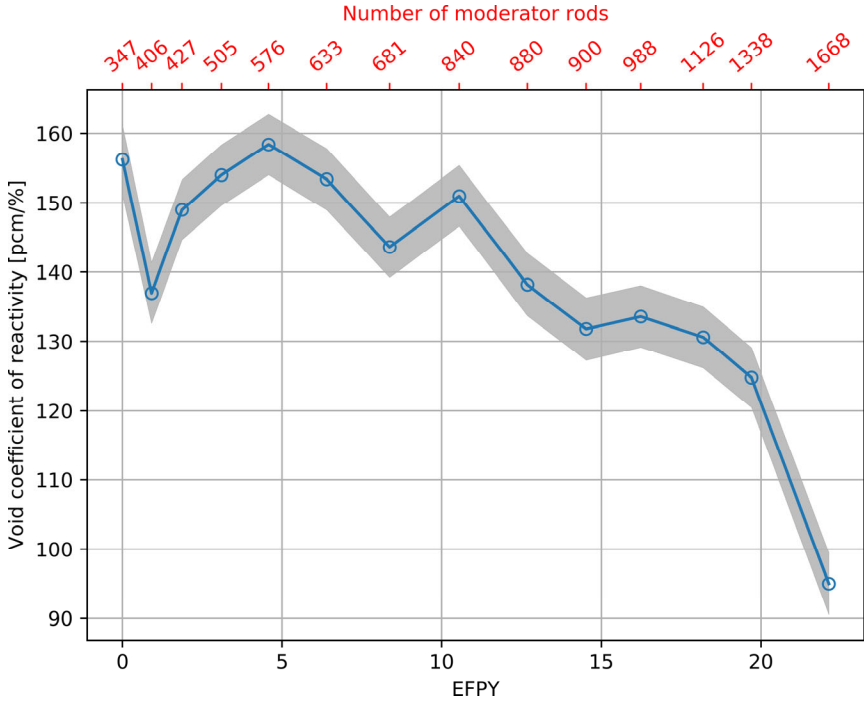

Figure 3.39: Serpent-calculated void coefficient of reactivity as a function of time and number of moderator rods in the TAP core. The uncertainty ±σ region is shaded.

3.3.3.3. Reactivity Control Rod Worth

In the TAP concept, control rods perform two main functions: to shut down the reactor at any point during operation by negative reactivity insertion and to control the excess of reactivity after moderator rod reconfiguration during regular maintenance. In an accident, the control rods would be dropped down into the core. The total control rod worth (ρ_{CRW} or CRW) is calculated for various moments during 25 years of operation to evaluate neutron spectrum shift influence on the CRW.

The reactivity worth of all control rods is defined as:

$$\rho_{CRW}(pcm) = \frac{k_{eff}^{W} - k_{eff}^{I}}{k_{eff}^{W} k_{eff}^{I}} \times 10^{5} \tag{3-14}$$

$$\rho_{CRW}(\$) = \frac{1}{\beta_{eff}} \frac{k_{eff}^{W} - k_{eff}^{I}}{k_{eff}^{W} k_{eff}^{I}} \tag{3-15}$$

where

$k_{eff}^{W} = effective\ multiplication\ factor\ when\ all\ rods\ are\ fully\ withdrawn$
$k_{eff}^{I} = \ effective\ multiplication\ factor\ when\ all\ rods\ are\ fully\ inserted$
$\beta_{eff} = effective\ delayed\ neutron\ fraction.$

The statistical error of the reactivity worth is obtained using formula:

$$\delta\rho_{CRW}(pcm) = \sqrt{\frac{\left(\delta k_{eff}^{W}\right)^{2}}{\left(k_{eff}^{W}\right)^{4}} + \frac{\left(\delta k_{eff}^{I}\right)^{2}}{\left(k_{eff}^{I}\right)^{4}}} \tag{3-16}$$

$$\delta\rho_{CRW}(\$) = \frac{1}{\beta_{eff}} \sqrt{\frac{\left(\delta k_{eff}^{W}\right)^{2}}{\left(k_{eff}^{W}\right)^{4}} + \frac{\left(\delta k_{eff}^{I}\right)^{2}}{\left(k_{eff}^{I}\right)^{4}} + \frac{\left(\delta\beta_{eff}\right)^{2}\left(k_{eff}^{W} - k_{eff}^{I}\right)^{2}}{\beta_{eff}^{2}\left(k_{eff}^{W} k_{eff}^{I}\right)^{2}}} \tag{3-17}$$

where

$$\delta k_{eff}^{W}, \delta k_{eff}^{I}, \delta\beta_{eff} = statistical\ errors\ from\ Serpent\ output.$$

Figure 3.41 demonstrates control rod worth evolution during 25 years of the TAP reactor operation. The cluster of 25 control rods made of boron carbide (B_4C) provided a reactivity worth of 5.059 ± 0.014 \$ at the BOL. However, spectral shift due to additional moderator rods toward the EOL leads to significant change in ρ_{CRW}. Adding more moderation near control rods increases ρ_{CRW} due to the local neutron spectrum thermalization (see transition from 347 to 427 moderator rods, Figure A.1). In contrast, adding moderator rods far away from the control rod positions leads to ρ_{CRW} degradation (see transition from 427 to 505 moderator rods, Figure A.1). On the one hand, the spectrum thermalizes and many fission product poisons exhibit larger absorption cross sections in the thermal energy range. On the other hand, higher actinides (particularly, isotopes of plutonium) are accumulated in the fuel salt which deteriorates control rod worth. Overall, ρ_{CRW} decreases to 4.472 ± 0.015 \$ at the EOL.

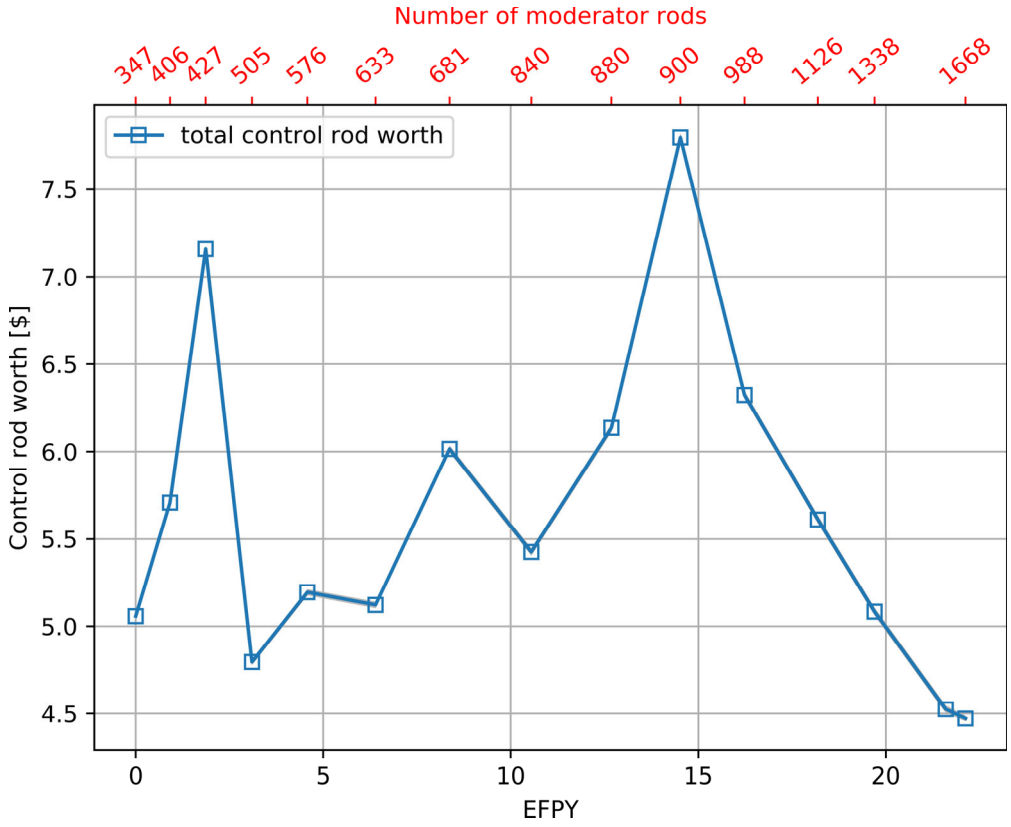

Figure 3.40: Serpent-calculated total control rod worth as a function of time and number of moderator rods in the TAP core. The uncertainty ±σ region is shaded.

Overall, the design of the reactivity control system is sufficient to shut down the TAP reactor at the BOL. However, the spectral shift, moderator rod reconfigurations, and the change in the salt composition during operation drive the total control rod worth below excess reactivity, violating reactor safety (insufficient shutdown margin). Thus, the number of control rods, their position, or material selection must be revised to make sure that the TAP reactor could be safely shut down at any moment during operation. For example, europium oxide (Eu_2O_3) might be a better absorbing material for the control rods [173].

3.3.3.4. Reactor Kinetic Parameters

Most of the neutrons produced in fission are prompt (99%). But less than 1% of neutrons are later emitted by fission products that are called the delayed neutron precursors (DNP). The term "delayed" means, that the neutron is emitted due to β-decay with half-lives in the range from few milliseconds up to 1 minute. Even though, the number of delayed neutrons per fission neutron is quite small (1% for most fissile isotopes), they play an essential role in the nuclear reactor control. Delayed neutrons presence changes the dynamic time response of a reactor to reactivity change from 10^{-7} s to 10 s, making it controllable by reactivity control system such as control rods. In nuclear library JEFF-3.1.2, delayed nuclear precursors are divided into 8 groups, each with different characteristic half-life, λi. The delayed neutron fraction, β_i, is defined as the fraction of all fission neutrons that appears as delayed neutrons in the i^{th} group.

It is crucial to study kinetic parameter dynamics because the fuel salt composition changes with time and new actinides appear in the fuel, which alters the emission of delayed neutrons. Figure 3.41 and Figure 3.42 show precursor-group-wise delayed neutron fraction (DNF, β_i) and decay constant (λ_i) evolution during 25 years of TAP MSR operation. The effective delayed neutron fraction (β_{eff}) in the TAP core decreased dramatically from 7.245×10^{-3} (0.5%) at the BOL to 4.564×10^{-3} (0.6%) at the EOL (37%).

Similarly, the effective precursor decay constant (λ_{eff}) slipped slightly from 0.481 s^{-1}(±0.8%) to 0.468 s^{-1}(±1.1%) during 25 years of operation. During operation, the concentration of ^{235}U decreases, and the concentration of fissile plutonium isotopes (e.g., ^{239}Pu) increases. Notably, ^{239}Pu emits about 2.5 times fewer delayed neutrons than ^{235}U; delayed neutron yields are 0.00664 and 0.01650 for the ^{239}Pu and ^{235}U, respectively. Thus, as fuel salt burnup increases, delayed neutron emission is controlled by plutonium isotopes (e.g., ^{239}Pu and ^{241}Pu) and decreases with time. All decay constants show a slight decrease toward the EOL due to the reactor spectrum hardening. This 37% decline in the effective delayed neutron fraction and 3% decline in the effective precursor decay constant must be taken into account in the TAP design accident analysis and safety justification.

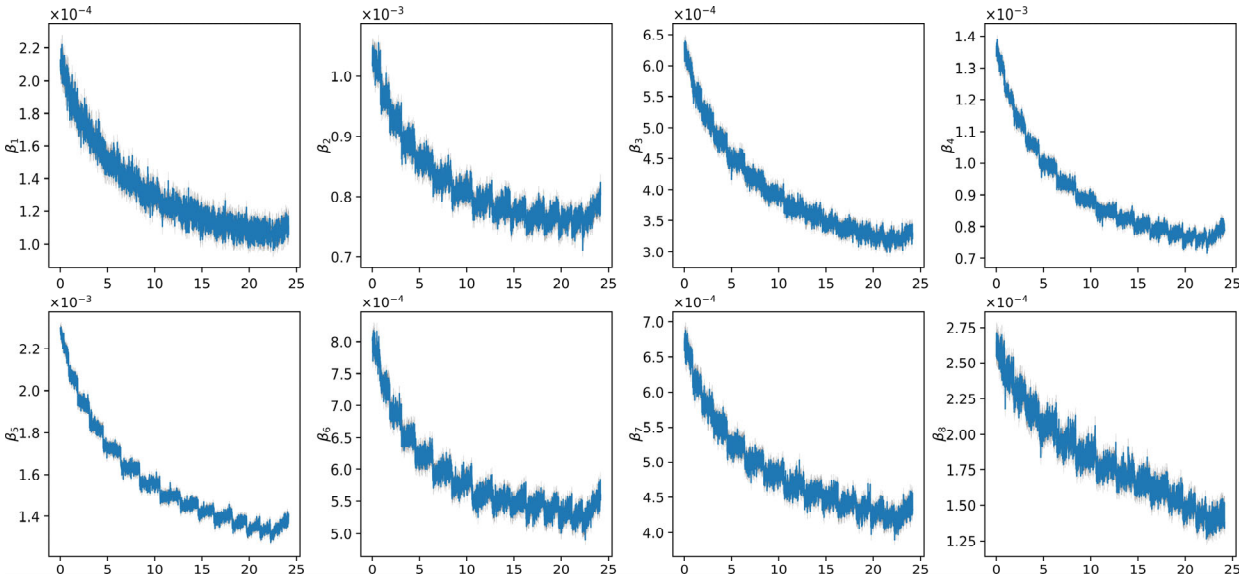

Figure 3.41: Evolution of the precursor-group-wise delayed neutron fraction (β_i) as a function of time for the TAP MSR. The uncertainty ±σ region is shaded.

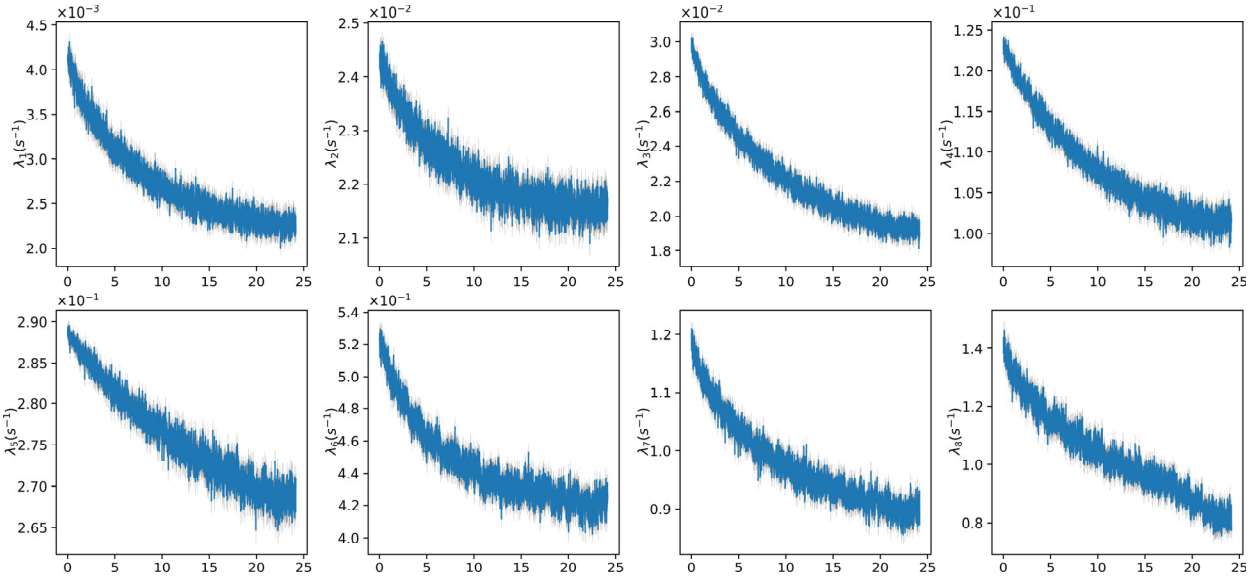

Figure 3.42: Evolution of the precursor-group-wise decay constant (λ_i) as a function of time for the TAP MSR. The uncertainty $\pm\sigma$ region is shaded.

3.3.4. Concluding Remarks

This section demonstrated SaltProc v1.0 capabilities for lifetime-long fuel salt depletion simulations applied to the TAP MSR. Section 4.1 summarized the TAP MSR core and fuel salt reprocessing system details that inform the SaltProc model (Section 4.2).

Section 4.3.1 presented lifetime-long depletion simulations with SaltProc v1.0. The 25- year simulation assumed ideal removal efficiency (e.g., 100% of target neutron poison is being removed at the end of each depletion step). This validation effort demonstrated good agreement with a reference ORNL report [18]. Full-core 3D SaltProc/Serpent analysis showed that spectrum hardening over the first 13 years of operation produces a sufficient amount of fissile plutonium to achieve the fuel salt burnup of 76.3 MWd/kgU after 22.5 years of operation. SaltProc-calculated inventories of major heavy isotopes at the EOL are consistent with results in the literature. The difference in mass between SaltProc and the reference was only 3% and 4% for fissile (^{235}U, ^{239}Pu, ^{241}Pu) and non-fissile (^{236}U, ^{238}U, ^{238}Pu, ^{240}Pu, ^{242}Pu) isotopes, respectively. Finally, the SaltProc-calculated feed rate is 460.8 kg of UF_4 per year, which consistent with 480 kg/y reported by Betzler et al. [18]

The time step refinement study in Section 3.3.2.1 showed that accurate uranium isotopic content predictions could be obtained with a relatively long depletion time step (6- or 12- day). However, the significant absolute difference in plutonium mass at the EOL (10 kg for a 6-day step) could be a safeguards issue, as this represents more than one significant quantity (8 kg) over the reactor lifetime. Overall, to get accurate plutonium isotopic content without raising proliferation issues, a 3-day depletion time step must be used.

Section 3.3.2.2 demonstrated SaltProc v1.0 for a 25-year depletion simulation with a realistic, physics-based noble gas removal efficiency. When identifying a reasonable mathematical model for realistic gas removal efficiency (ϵ), the liquid phase mass transfer coefficient (K_L)

demonstrated a strong correlation with ϵ. Thus, SaltProc simulations using different K_L in validity range from 0.0847 to 8.4667 mm/s (corresponding ^{135}Xe removal efficiency $\epsilon \in [0.031,0.915]$) showed that the larger liquid phase mass transfer coefficient and corresponding higher noble gas extraction efficiency provided significant neutronics benefit, better fuel utilization, and longer time between shutdowns for moderator rod reconfiguration. Notably, the larger mass transfer coefficient also provides a slightly more thermal neutron spectrum because poisonous FPs (^{135}Xe) absorb fewer thermal neutrons. In the following sections, the results of these realistic depletion simulations will be used for short-term transient simulations and safety parameter analysis.

Finally, this section demonstrated safety and operational parameter evolution during 25 years of the TAP MSR operation. In general, the safety of the reactor worsens with time due to actinides and FPs accumulating in the fuel salt. Shifting neutron spectrum from epithermal to thermal by periodically adding more moderator rods also has a negative influence on crucial safety an operational characteristic. These observations must be taken into account in the TAP MSR designing, accident analysis, and safety justification.

3.4. Load-following and Safety Analysis of TAP MSR

In order to be competitive in the current domestic energy market, MSRs may need the flexibility to follow net load on the grid. Such load-following operation has the potential to increase the commercial competitiveness of nuclear power dramatically. Due to the increasing penetration of renewables into the electric grid, base-load operation carries the risk of correspondingly frequent negative electric energy pricing. Thus, responsiveness to net electricity demand is essential to market relevance for new designs [174]. This section presents a validation demonstration applying SaltProc v1.0 to simulate fuel salt depletion with online reprocessing during short-term transient to evaluate load-following capabilities of the TAP MSR.

3.4.1. Technical Aspects of Load Following with Nuclear Reactors

The physical constraints limiting power variations in conventional LWRs include [175]:

- thermal strain and stress to fuel materials[11];
- fuel burnup (low excess reactivity at the EOC);
- ^{135}Xe poisoning (iodine pit);
- reactivity thermal feedback (change in the temperature of the primary coolant and fuel causes negative reactivity insertion which limits power regulation capabilities).

Each of these physical effects is currently under active international research. This section focuses only on the fission product poisoning, especially the “iodine pit”. The “iodine pit”, also called the “iodine hole” or “xenon pit”, is the reactor’s inability to start a few hours after the reactor power decreases due to peak of ^{135}Xe concentration in the core. The ^{135}Xe is the strongest known neutron absorber ($\sigma_{a,^{135}Xe} = 2.6 \times 10^6\ b$) with a half-life $\tau_{1/2}$ = 9.17h and yield for ^{235}U fission about 6.6%. Figure 3.43 shows the entire decay chain, which characterizes ^{135}Xe gain and loss

[11] This constrain does not apply to circulating-fuel MSRs because the fuel is into a liquid form.

channels. The vast majority of ^{135}Xe (6.4%) is produced from ^{135}I decay ($\tau_{1/2}$ = 6.6h). About half of ^{135}I is produced directly from fission and half from ^{135}Te decay ($\tau_{1/2}$ = 19s) [176].

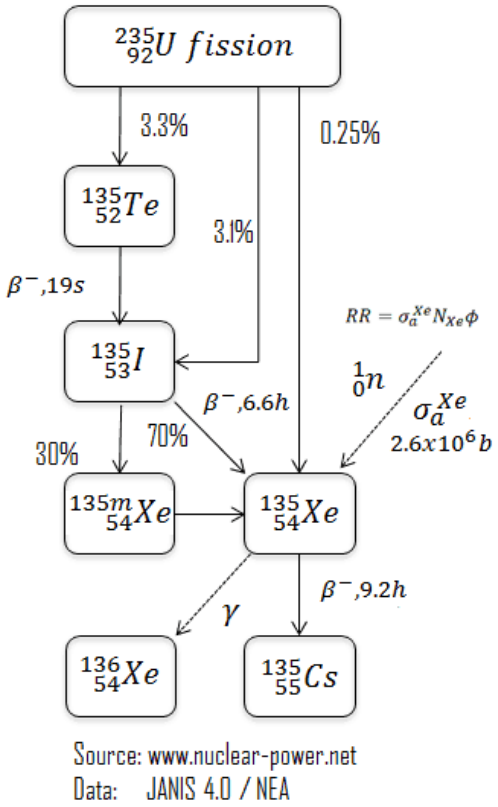

Figure 3.43: Mechanisms of ^{135}Xe gain and loss in the reactor core (reproduced from [176]).

Under normal operating conditions, ^{135}Xe is transmuted to ^{136}Xe ('burned out') in the reactor core as it is produced. So, while it harms the neutron economy, balancing the reactor controls can compensate for its effect. The burnout of ^{135}Xe for an operating reactor can be described as follows:

$$^{135}Xe + {}^{1}n \rightarrow {}^{136}Xe\ (stable) \tag{3-18}$$

Because ^{135}Xe is produced partially from the ^{135}I decay, the ^{135}Xe concentration directly depends on the ^{135}I concentration. Therefore, the iodine and xenon rate of change can be described as follows

$$\frac{dI(t)}{dt} = \gamma_I \Sigma_f \phi - \lambda_I I \tag{3-19}$$

$$\frac{dX(t)}{dt} = \lambda_I I + \gamma_X \Sigma_f \phi - \lambda_X X - \sigma_{a,X} \phi X \tag{3-20}$$

where

$$I, X = number\ density\ of\ {}^{135}I,\ {}^{135}Xe\ [cm^{-3}]$$
$$\gamma_I, \gamma_X = effective\ yield\ of\ {}^{135}I,\ {}^{135}Xe\ [fission^{-1}]$$
$$\lambda_I, \lambda_X = decay\ constant\ of\ {}^{135}I,\ {}^{135}Xe\ [s^{-1}]$$

$$\Sigma_f = macroscopic\ fission\ cross\ section\ of\ ^{235}U\ [s^{-1}]$$
$$\phi = neutron\ flux\ [cm^{-2}s^{-1}].$$

The difficulty comes when the reactor power is reduced, and there are fewer neutrons to burn the ^{135}Xe out, so its concentration increases and further suppresses reactor power. In this case, the core takes some time to recover from the power reduction impact of ^{135}Xe. This response to changing power levels, particularly from higher to lower power, dramatically slows the reactor's response to power demand [177].

In a liquid-fueled MSR, gaseous fission products (e.g., xenon) can be dynamically removed from the fuel salt by the gas separation system (see Section 3.2.1.1). Thus, xenon gas, including problematic ^{135}Xe, can be removed from the fuel salt outside the reactor core to eliminate its negative impact on the core neutronics. If the gas separation system can remove the vast majority of xenon, it is possible to alter the reactor power output in a wide range with very brief required recovery time. Overall, ^{135}Xe removal during reactor operation would potentially allow precise and flexible dynamic control of the reactor power level to follow power demands, typically referred to as 'load following.'

This section presents modeling and simulation of load following transient operation of the TAP MSR. This study focuses on the ^{135}Xe/^{135}I balance in the TAP core and its effect on reactor performance. In this section, short-term (24 hours) depletion is simulated with the core power changing in the [0,100%] range for xenon removal efficiency (ϵ_{Xe}) varied between 0 and 0.915 (see Table 3.8).

This section also demonstrates an analysis of reactor load-following capability for various moderator configurations and fuel salt compositions to bound the necessary efficiency of the gas removal system to ensure load-following operation.

3.4.2. TAP MSR Load Following Analysis

All of the analysis herein used SaltProc v1.0 with the full-core 3-D model of the TAP MSR developed using Serpent 2 (see Section 3.3.1. The multi-component, online reprocessing system model with realistic noble gas removal efficiency described in Section 3.3.2.2, is used to simulate fission product removal and fresh fuel injection during the anticipated transient. To simulate transients with time-dependent power generation, a new capability is added to SaltProc v1.0 to perform fuel salt depletion with variable time step size and power level[12] during each depletion step. The depletion calculation in the load following regime captures the effects of ^{135}Xe poisoning and illuminates the benefit of using an online gas removal system in the TAP concept.

3.4.2.1. Power Load Curve Selection Approach

The load and generation must be continuously and almost instantly balanced in an electric power system. This is a physical requirement independent of the market structure. Regulation and load following (in the real-time energy market they are provided by the intra-hour workings) are

[12] For simplicity, the reactor power level is adjusted by changing the normalization factor in Serpent (set power P[W]). This simplification assumes that spatial and energy distribution of the neutron flux remains constant and only the magnitude of the flux changes with time. That is, control rod movement and the corresponding change in the flux spatial and energy distribution are not treated here.

the two services required to maintain a balance between power generation and power load. Figure 3.44 demonstrates the morning ramp-up decomposed into the total load (green), smooth load-following ramp (blue), and regulation (red). The smooth load- following slowly rises from 3566 MW to 4035 MW over 3 hours. Regulation compensates for high-frequency fluctuations in the load around the underlying trend within the 55MW range. In the PJM region (Delaware, Illinois, Indiana, Kentucky, Maryland, Michigan, New Jersey, North Carolina, Ohio, Pennsylvania, Tennessee, Virginia, West Virginia, and the District of Columbia), New York, and New England, the 5-min ramping capability of a generator is required for the regulation, while in Texas and California it is 15-min and 10-min, respectively [178].

In this context regulation refers to the use of online generation or storage that is equipped with automatic generation control and can change output quickly (MW/min ramp rate) to compensate for the minute-to-minute fluctuations in customer loads and correct for unintentional fluctuations in power generation [178]. Typical natural gas peaking plants can ramp at or above 10% of their capacity per minute [142]. Elite combustion engine peakers (Wärtsilä) can ramp up from 10% to 100% load (or down) in less than one minute [179]. Hydropower plants also typically have accurate, high-speed ramping capability suitable for regulation [178].

Conventional nuclear power plants (Generation III/III+) can be used for load-following (blue curve on Figure 3.44) but have limited maneuverability. For example, the German Konvoi reactors are designed for 15,000 cycles with daily power variations from 100% to 60% power level with ramp rate up to 2%/min [180], which is by order of magnitude slower than fossil-fueled plants. The MSRs must enable daily power variation with a much more flexible range (from 100% to 0% and from 0% to 100%) and ramp rate up to 10%/min to compete with these generators. The physical constrains limiting power variation range and ramp rate in nuclear reactors was listed in Section 3.4.1. The current section of the dissertation focuses only on the ^{135}Xe poisoning effect. Other physical effects, such as thermal strain and stress in structural materials are not treated here.

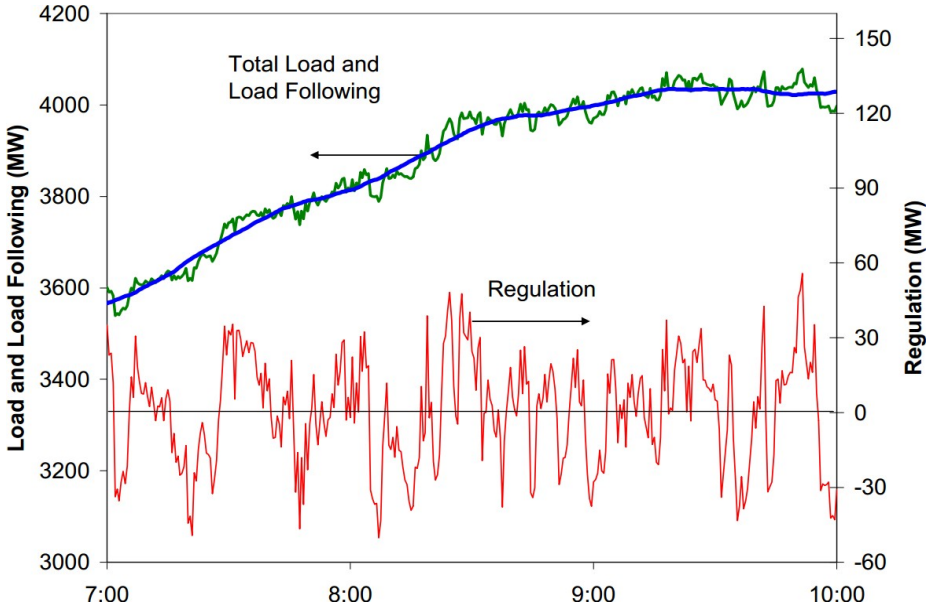

Figure 3.44: Regulation (red) compensates for minute-to-minute fluctuations in system total load (green), load following (blue) compensates for the inter- and intra-hour ramps (reproduced from [178]).

Performing a depletion calculation with SaltProc v1.0 to mimic the load-following maneuvering shown in Figure 3.44 would require a very fine time step (e.g., 15-minute step). To simulate power change with the desired rump rate (0.1 Hot Full Power (HFP)/min), depletion time step resolution of less than a 1-minute is needed. Such fine resolution requires thousands of depletion time steps to simulate 12-hour transient involving unreasonable computational costs. Instead, this section presents simulations with a 1-hour time step to investigate the impact of gaseous fission product removal on the reactor response to power demands.

The most challenging power transient for conventional LWRs from the viewpoint of xenon poisoning is well-defined in the literature. If after the ^{135}Xe concentration reaches equilibrium (40-50 hours after startup with fresh fuel), the reactor power was decreased from 100% to 0% (e.g., the reactor is tripped), the ^{135}Xe concentration and corresponding negative reactivity insertion would reach maximum in about 10-11 hours after shutdown [181, 182]. Notably, the time after shutdown when ^{135}Xe concentration reaches a maximum strongly depends on the reactor neutron energy spectrum.

Thus, to demonstrate SaltProc v1.0 capabilities for a short-term transient with the reactor power change and to investigate load-following capabilities of the TAP reactor with a focus on the xenon poisoning, the following worst-case power load profile is selected:

a) operate on 100% of HFP long enough to reach ^{135}I/^{135}Xe equilibrium;

b) instantaneous power drop from 100% to 0%;

c) shutdown for t_X^{max} [hours] to reach the ^{135}Xe concentration extremum;

d) restart the reactor instantly from 0% to 100% power level and operate on 100% for a few hours.

Or in math formulation:

$$P(t) = \begin{cases} 100\%, & t < t_{eq} \\ 0\%, & t_{eq} \le t \le t_{eq} + t_{\mathrm{X}}^{\max} \\ 100\%, & t > t_{eq} + t_X^{max} \end{cases} \tag{3-21}$$

where

$$P(t) = reactor\ power\ level\ [\%]$$
$$t_{eq} = time\ after\ startup\ to\ reach\ ^{135}Xe\ equilibrium\ concentration\ [h]$$
$$t_X^{max} = time\ after\ shutdown\ when\ ^{135}Xe\ concentration\ peaks\ [h].$$

This postulated worst-case transient could be considered as backing up solar power with nuclear on a high-solar-penetration grid (e.g., in California). Any other power load profile (i.e., blue load-following line shown in Figure 3.44) would demonstrate a significantly milder xenon poisoning effect because of the power demand change in the [0,100%] range realistically is not instantaneous. That is, if the TAP MSR would be able to maintain criticality in the described stress test (e.g., k_{eff} 1.0 during all stages of the transient), then it is capable of following a realistic load curve.

The local extremum of xenon concentration can be described as follows

$$\frac{dX(t)}{dt} = 0 \tag{3-22}$$

The system of Ordinary Differential Equations which consist of Equations 3-19, 3-20 and 3-22 must be solved to calculate when the ^{135}Xe concentration reaches maximum.

If the ^{135}I and ^{135}Xe concentrations at shutdown is I_0 and X_0, respectively, the time after shutdown when the ^{135}Xe concentration peaks is given by:

$$t_X^{max} = \frac{1}{\lambda_X - \lambda_I} \log\left(\frac{\lambda_X(\lambda_I[X_0 + I_0] - \lambda_X X_0)}{\lambda_I^2 I_0}\right) \tag{3-23}$$

Since the ^{135}I and ^{135}Xe concentrations at shutdown in the TAP core are expected to be different at the BOL and EOL due to significant spectral shift, t_X^{max} is recalculated for each case to obtain the worst possible xenon poisoning effect. The ultimate goal of this effort is to evaluate the timing and impact of problematic fission product removal (i.e., xenon removal) on maximum negative reactivity insertion.

3.4.2.2. Results and Analysis

The TAP full core depletion analysis was performed using SaltProc v1.0. A 1-hour depletion time step captures rapid changes in reactivity and isotopic composition during the transient. Figure 3.45 demonstrated the effective multiplication factor evolution during postulated worst-case transient when the reactor is tripped for 11 hours (typical time to reach maximum ^{135}Xe concentration in conventional LWRs) and then restarted. The gas removal system for that demonstration case was inactive to enhance the xenon poisoning effect. At the beginning of the transient (initial conditions), the reactor operated for 8448 days (23 years), and all moderator rods are inserted in the core (see 1668 rods configuration in Figure A.2). The negative effect of xenon poisoning is expected to be the greatest at the EOL when the core has the most thermal neutron spectrum. The multiplication factor decreases by 64 pcm during the first two hours after shutdown (^{135}Xe concentration reached its maximum) and then increases by 242 pcm because ^{135}Xe loss due to decay overcame its gain from ^{135}I decay. The k_{eff} increase accelerated after reactor power turned back to 100% due to ^{135}Xe burnout. Figure 3.45 clearly indicates that the time after shutdown when the ^{135}Xe reaches its extremum (t_X^{max}) is significantly shorter for the TAP reactor than for LWRs (11 hours).

Using ^{135}I and ^{135}Xe number densities at the 8448th day of operation (the 10th day before the EOL) from long-term realistic analysis (see Section 3.3.2.2) and Equation 3-23, the xenon peak time is calculated for the TAP MSR with all moderator rods inserted: t_X^{max} = 2.76h.

To estimate maximum negative reactivity insertion due to xenon poisoning, the transient simulation is repeated with the finer time resolution (15 minutes instead of 1 hour) and shutdown time of 2.75 hours (e.g., the time between the shutdown and power ramp-up to 100% is equal t_X^{max}).

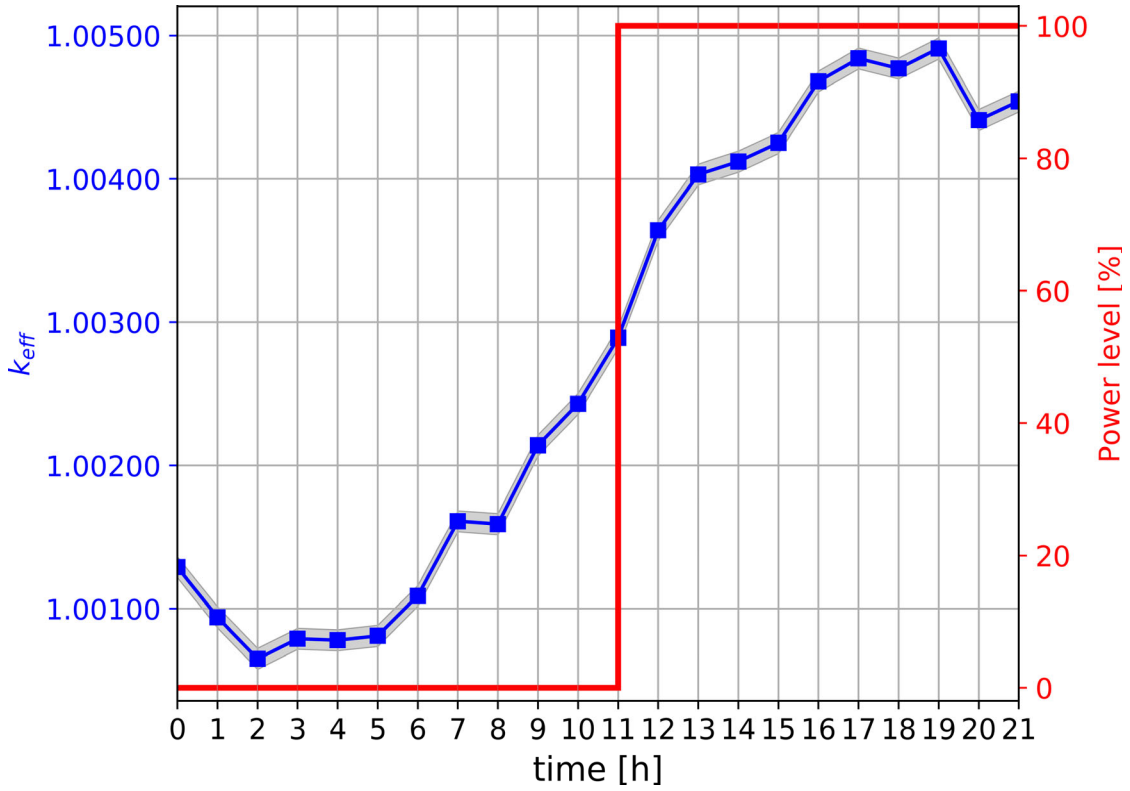

Figure 3.45: The effective multiplication factor dynamics for an 11-hour shutdown (well-known xenon peak time for LWRs) for the TAP reactor, 10 days before the EOL (all moderator rods inserted), the gas removal system is turned off. Uncertainty (σ = 7 pcm) is shaded.

Figure 3.46 shows that the effective multiplication factor dropped by 70 pcm during the first 2.75 hours after shutdown as predicted by Equation 3-23. After power ramps up from 0% to 100%, k_{eff} returned to its initial value (1.00151) in 75 minutes. The imbalance between ^{135}I production and ^{135}Xe burnout is the main reason for this positive reactivity boost. Notably, maximum negative reactivity insertion due to ^{135}Xe buildup after shutdown in the PWR (1500 pcm) is two orders of magnitude greater than in the TAP MSR (70 pcm). Thus, the TAP reactor with inactive gas removal system remains critical throughout worst-case power change even during the 8448th of operation (the 10th day before k_{eff} drops below 1) when operative excess reactivity is low (151 pcm 70 pcm). If the shutdown happens during the last 9 days of the TAP reactor operation, then the operator would not be able to restart it until t_X^{max} = 2.76h after shutting down.

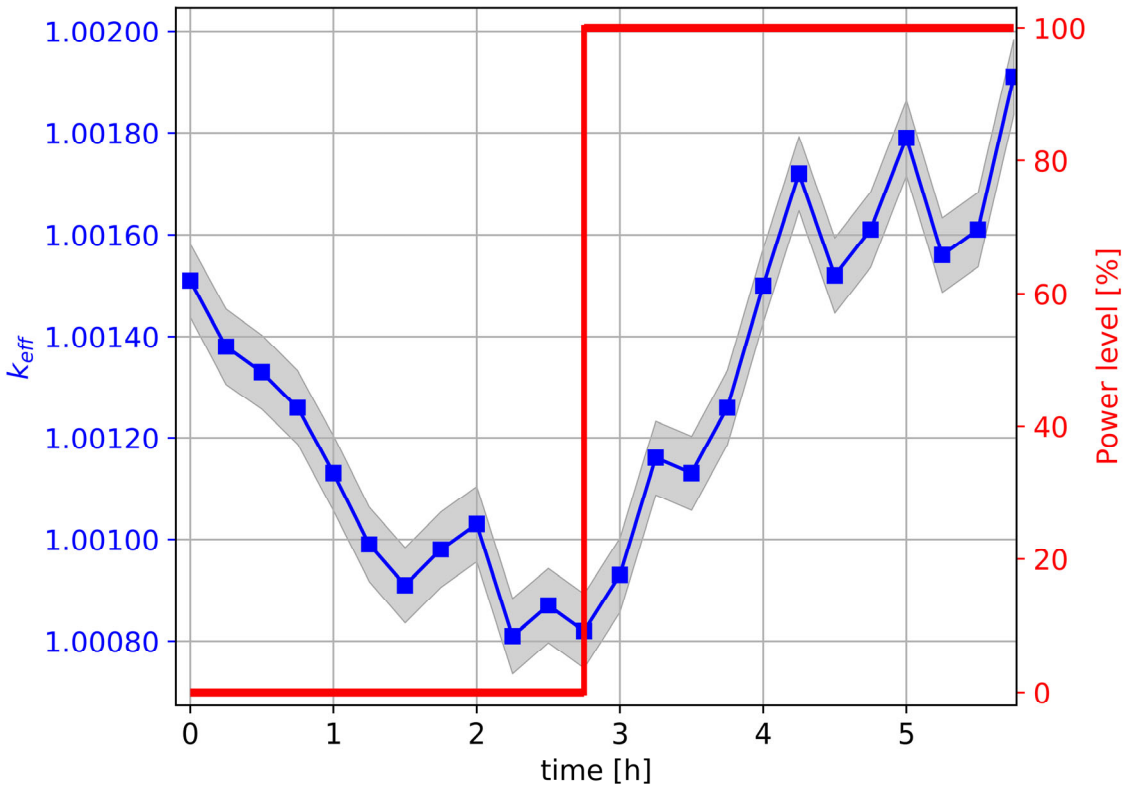

Figure 3.46: The effective multiplication factor dynamics for the worst-case load curve (2.75-hour shutdown) for the TAP reactor, 10 days before the EOL (all moderator rods inserted), the gas removal system is turned off. Uncertainty (σ = 7 pcm) is shaded.

The analysis of the fuel composition evolution provides clearer information about the ^{135}Xe/^{135}I equilibrium and the core state. Figure 3.47 shows changes in the number density of isotopes influential to the TAP core neutronics throughout the transient. The ^{135}I/^{135}Xe number density ratio after reaching xenon equilibrium is equal to 1.0. After shutdown, ^{135}I decays to ^{135}Xe that is not burned up. The ^{135}I decay caused xenon concentration to increase by 4% from equilibrium after 2.75 hours due to a shorter ^{135}I half-life ($\tau_{1/2,^{135}I} = 6.6\ h$ = 6.6h vs. $\tau_{1/2,^{135}Xe} = 9.17\ h$). Thus, during the first 2.75 hours, ^{135}Xe gain from ^{135}I decay slightly overcame ^{135}Xe decay loss. In sum, the ^{135}Xe peak is almost negligible (+4%) even in the worst-case load profile scenario due to a lower ^{135}I/^{135}Xe concentration ratio at the equilibrium: 1.0 and 2.3 for the TAP reactor and PWR, respectively [183].

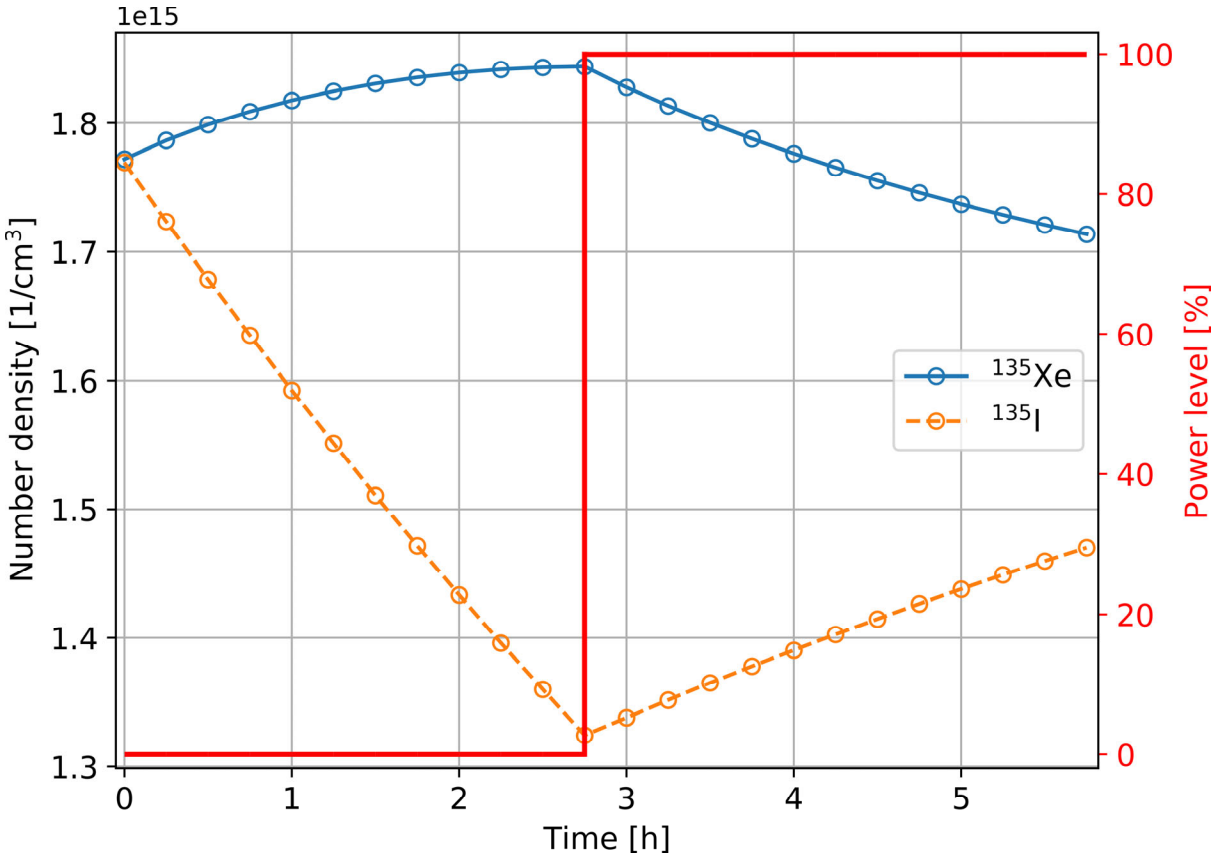

Figure 3.47: Number density of ^{135}Xe and its direct precursor ^{135}I for the worst-case load curve (2.75-hour shutdown) for the TAP reactor, 10 days before the EOL (all moderator rods inserted), the gas removal system is turned off.

Table 3.10 shows that even without gas removal the TAP reactor experienced insignificant effect of xenon poisoning during the transient, however, the effect worsened toward the EOL. The fuel composition and multiplication factor evolution analysis described earlier are repeated to evaluate impact of the reactor spectrum (geometry #1 has a significantly harder spectrum than geometry #15) and the fuel salt composition on the effect of xenon poisoning and on the reactor's potential ability to follow load. The effect of xenon poisoning worsens toward the EOL because the ^{135}Xe concentration peak is larger for the most thermal core configuration (all moderator rods inserted, the largest moderator-to-fuel ratio). Right after the final moderator configuration update (switch from geometry #14 to #15), the xenon concentration peak is slightly larger than at the 297th day of the cycle. The fissile ^{235}U, ^{239}Pu, and ^{241}Pu concentration decreasing during last cycle due to burnup, while poisonous actinides (e.g., ^{238}Pu, 240Pu, 242Pu, 236U) concentration increases which impacts ^{135}I/^{135}Xe number density ratio and, consequently, ^{135}Xe concentration peak value. Notably, such phenomena are not observed for the BOL (geometry #1, SVF=0.903) or Middle of Life (MOL) (geometry #8, SVF=0.766).

Table 3.10: Effect of ^{135}Xe poisoning after shutdown for the TAP reactor operation with inactive gas removal system ($\epsilon_{Xe} = 0$). Stochastic uncertainty $\sigma_\rho = 7$ pcm.

Geometry	SVF [-]	Time after moderator configuration update [d]	Operative excess reactivity (ρ_0) [pcm]	Analytically predicted ^{135}Xe peak time (t_X^{max}) [h]	Maximum relative ^{135}Xe concentration change [%]	Maximum reactivity change after shutdown ($\Delta\rho$) [pcm]
1	0.903	3	3542	0.749	0.33	-10
1	0.903	288	405	0.5	0.14	-15
1	0.903	315	165	0.484	0.13	-4
8	0.766	3	3014	0.688	0.36	-10
8	0.766	390	1529	0.722	0.39	0
8	0.766	777	204	0.751	0.42	0
15	0.536	3	2263	2.528	3.32	-57
15	0.536	153	1160	2.647	3.69	-60
15	0.536	297	129	2.758	4.07	-70

Overall, the TAP MSR could be restarted after shutdown even without gas removal in worst-case initial conditions: the most thermal moderator configuration, low operative excess reactivity at the end of the burnup cycle, instantaneous power drop, and ^{135}Xe concentration at its extremum. To investigate the benefits of online fission gas removal on the xenon poisoning effect, the postulated transient simulation is repeated for different moments in time (e.g., BOL, MOL, EOL) with a fully operational gas removal system ($\epsilon_{Xe} = 0.915$).

Figure 3.48 demonstrates a more notable xenon poisoning effect for the case with high gas removal efficiency than for the no-removal case. The reactivity drops by 100 pcm during the first hour after shutdown. The gas removal system keeps ^{135}Xe concentration very low by continuously extracting 91.5% of xenon isotopes. Simultaneously, the online reprocessing system extracts ^{135}I very slowly (cycle time is 60 days); hence, ^{135}I/^{135}Xe concentration ratio is significantly greater than for the no-removal case (11.0 vs. 1.0). According to Equation 3-23, ^{135}Xe concentration should reach local extremum in about 11 hours after shutdown, but this equation disregards online reprocessing. The depletion simulation performed using SaltProc v1.0 demonstrated that ^{135}Xe concentration peaked in one hour after the shutdown and caused the reactivity drop by 100 pcm (see Figure 3.48). Afterward, the reactivity restored quickly (2 hours) to its initial value because the gas removal system extracts 91.5% of xenon every hour. Overall, ^{135}Xe loss due to its decay and online gas extraction is more significant than ^{135}Xe gain due to ^{135}I decay throughout the transient.

Table 3.11 shows that the TAP reactor with high gas removal efficiency experienced small effect of xenon poisoning during the transient, and it also worsened toward the EOL. Similar to the analysis with inactive gas removal system, maximum negative reactivity insertion due to xenon poisoning worsens toward the EOL because ^{135}Xe absorption cross section drops dramatically as energy grows. Notably, the maximum ^{135}Xe concentration peak is significantly greater for an excellent gas removal efficiency ($\epsilon_{Xe} = 0.915$) than for the no-removal case ($\epsilon_{Xe} = 0$): +197% and +4%, respectively. Despite greater ^{135}Xe concentration peak, negative change of reactivity after shutdown for the $\epsilon_{Xe} = 0.915$ case is slightly deeper than for the $\epsilon_{Xe} = 0$ case: -100 pcm and -70 pcm, respectively. The reason for this is the neutron energy spectrum in the TAP MSR, which is harder than in conventional light-water thermal reactors. As we know, fast reactors are unaffected by xenon poisoning because the absorption cross section of ^{135}Xe in the fast spectrum is insignificantly larger than absorption cross section of other fission products [182, 184]. The TAP concept has intermediate spectrum which softens towards the EOL. Finally, the effect of xenon poisoning in TAP MSR is almost negligible and can be easily compensated by control rod movement, while in well-studied PWR it presents a challenge (1500 pcm) [183].

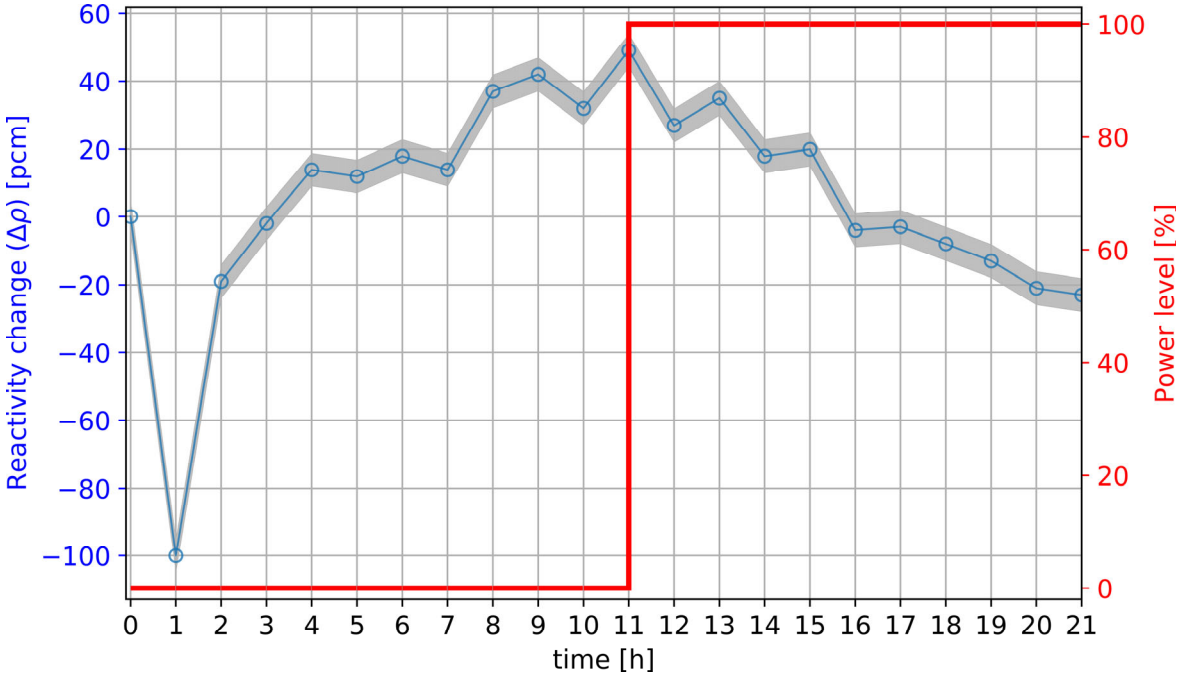

Figure 3.48: Reactivity dynamics during an 11-hour shutdown for the TAP reactor, 10 days before the EOL (all moderator rods inserted), the gas removal system operates with efficiency ϵXe = 0.915. Uncertainty (σ = 5 pcm) is shaded.

Additionally, the neutron spectrum of both reactors are analyzed to understand the difference between ^{135}I/^{135}Xe gain and loss for the TAP MSR and PWR. Figure 3.49 demonstrates the neutron flux energy distribution normalized by unit lethargy for both reactors. The TAP reactor spectrum at the BOL (SVF=0.903) is much harder than for the PWR due to a lack of moderation in the TAP core and its type ($ZrH_{1.66}$ instead of light water). The harder neutron spectrum leads to weaker ^{135}Xe transmutation because the capture cross section declines rapidly with energy (see Figure 3.49, lower plot, solid red line, energy range from 10^{-7} to 10^{-4} MeV). As a result, the ^{135}I/^{135}Xe number density ratio is 0.78 for the TAP MSR at the BOL, which is significantly lower than that

for the PWR with fresh fuel (2.3). Thus, ^{135}Xe gain from ^{135}I decay cannot overcome ^{135}Xe loss due to decay, and no xenon concentration peak is observed at the BOL (

Table 3.10, first three rows).

Table 3.11: Effect of ^{135}Xe poisoning after shutdown for the TAP reactor operation with the high ^{135}Xe removal efficiency (ϵ_{Xe} = 0.915). Stochastic uncertainty σ_ρ = 5 pcm.

Geometry	SVF [-]	Time after moderator configuration update [d]	Operative excess reactivity (ρ_0) [pcm]	^{135}I/^{135}Xe concentration ratio before shutdown [-]	Maximum relative ^{135}Xe concentration change [%]	Maximum reactivity change after shutdown ($\Delta\rho$) [pcm]
1	0.903	9	3344	8.96	174	-50
1	0.903	171	1930	8.76	173	-40
1	0.903	324	570	8.66	172	-38
8	0.766	3	3570	8.87	175	-61
8	0.766	366	2150	8.9	174	-40
8	0.766	762	762	8.93	175	-33
15	0.536	9	3370	11.07	194	-105
15	0.536	90	2771	11.17	195	-108
15	0.536	303	1265	11.42	197	-100

The TAP MSR neutron spectrum thermalizes toward the EOL due to additional moderator rod insertion. Figure 5.7 shows that the TAP core spectrum at the EOL (after 23 years of operation, all moderator rods inserted into the core) is thermal and similar to the PWR spectrum. However, the ^{135}I/^{135}Xe inventory ratio for the PWR with fresh fuel is significantly greater than for the TAP core at the EOL despite similar spectra (2.3 and 1.0, respectively). The reason for that difference is the different fissile content. Results in Section 3.3.2 shown that toward the EOL fissile ^{235}U is being substituted with fissile ^{239}Pu and ^{241}Pu. More specifically, instead of 6.8 t of ^{235}U at startup, at the EOL, the fuel salt contains 1.3 t of ^{235}U, 1 t of ^{239}Pu, and 0.5 t of ^{241}Pu. That is, the fuel salt fissile inventory in the TAP MSR at the EOL contains 46 wt% of ^{235}U, 36 wt% of ^{239}Pu, and 18 wt% of ^{241}Pu.

Table 3.12 shows ^{135}I and ^{135}Xe yields from thermal fission for all fissile isotopes contained in the fuel salt. At the BOL, ^{135}I and ^{135}Xe in the TAP reactor and PWR are produced from ^{235}U fission. The ^{135}I isotope production rate per thermal fission stays approximately the same during 23 years of operation because ^{135}I yield is very close for all considered fissile isotopes. However,

the rate of ^{135}Xe production directly from fission for fissile plutonium isotopes is significantly greater than for the ^{235}U (e.g., 5 times greater for ^{239}Pu and 8 times greater for ^{241}Pu). Thus, a greater ^{135}Xe production rate toward EOL with approximately the same ^{135}I production rate leads to a smaller ^{135}I/^{135}Xe concentration ratio. Overall, ^{135}I/^{135}Xe number density ratio increasing from 0.78 to 1.0 during 25 years of the TAP MSR operation, which leads to a more massive ^{135}Xe concentration peak after shutdown and worsens xenon poisoning effect.

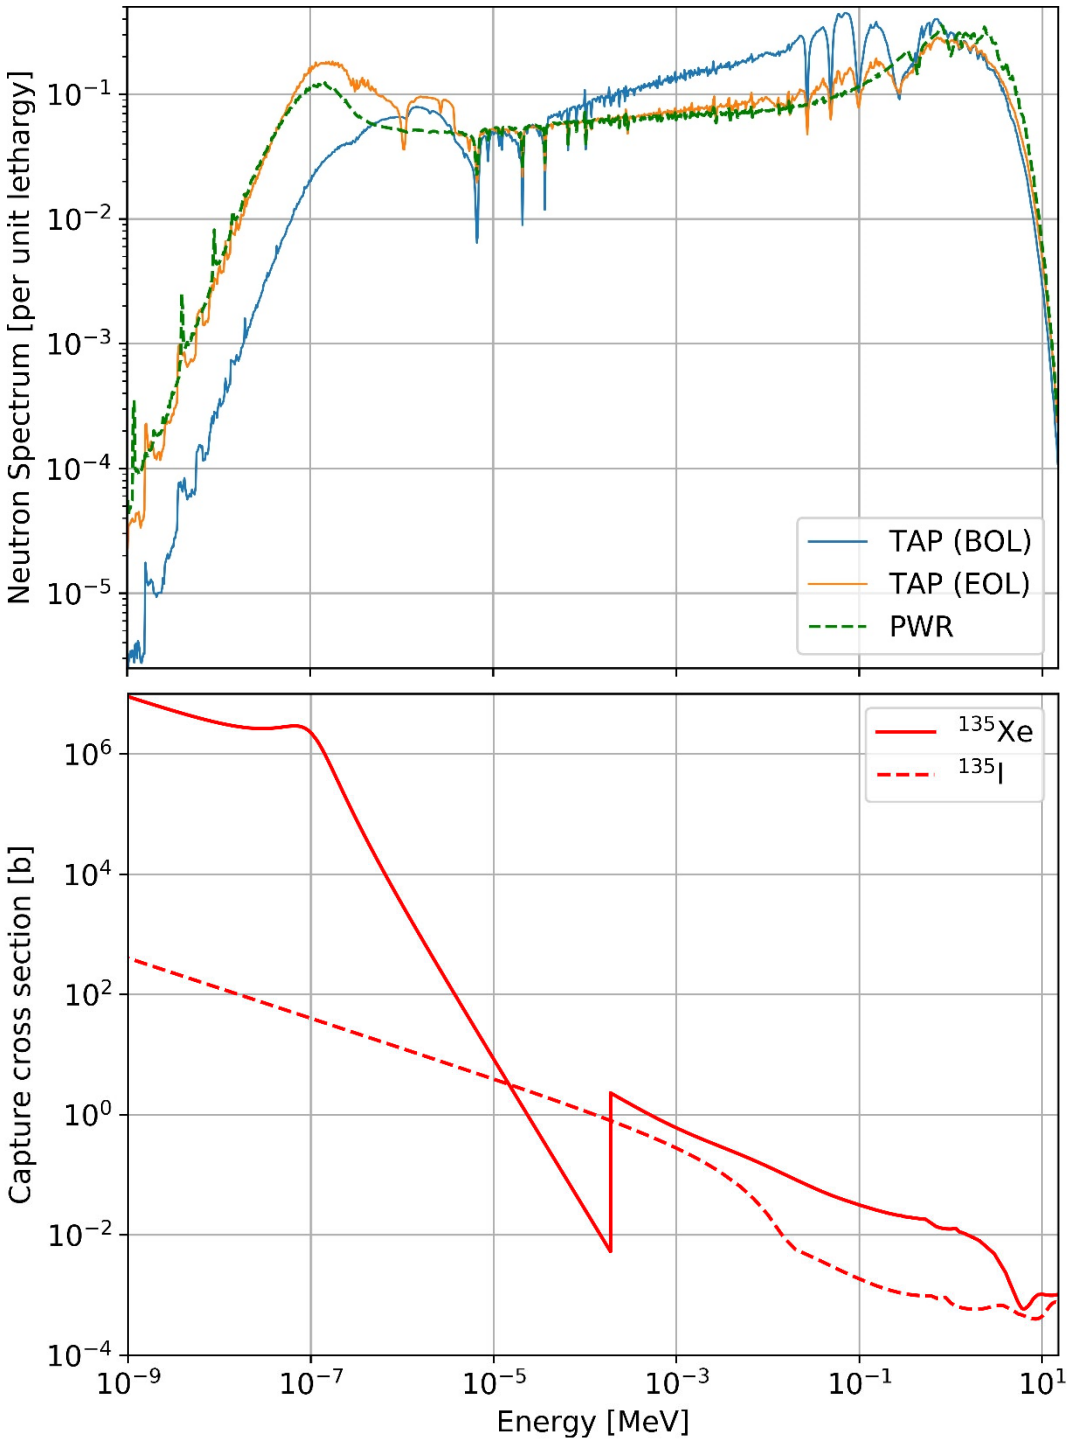

Figure 3.49: Neutron spectra normalized by lethargy for the PWR and TAP (upper) and ^{135}I, ^{135}Xe caption cross section (lower) [183].

Table 3.12: Fission product yields (isotopes per fission) from thermal fission [185].

Isotope	^{235}U	^{239}Pu	^{241}Pu
^{135}I	0.0639	0.0633	0.0684
^{135}Xe	0.0022	0.0103	0.0017

In conclusion, a negligible xenon poisoning is observed in the TAP reactor during the anticipated transient because it has a relatively hard neutron energy spectrum even at the most thermal core configuration (all moderator rods are inserted into the core). The harder spectrum gives a small $^{135}I/^{135}Xe$ concentration ratio which leads to a low ^{135}Xe concentration peak after the shutdown. Notably, the fission gas removal with high efficiency did not significantly change the xenon poisoning effect because the ^{135}Xe absorption cross section fell dramatically as neutron energy grows. Overall, the TAP reactor can effectively load- follow even without fission gas removal.

3.4.3. Safety and Operational Parameters Evolution during Load Following

To analyze the impact of the load-following transient on the TAP concept safety, the safety and operational parameters are calculated at various moments during postulated earlier worst-case power change transient (0% power level for 11 hours, instantaneous power boost to 100%, and then 10 hours on 100% power level) using methodology from Section 3.3.3. The combination of fuel and moderator temperature feedback coefficients must remain negative, and the reactivity worth of control rods must be sufficient to shut down the reactor throughout the transient. Ideally, the reactor is more controllable if major safety and operational parameters remain stable and unaffected by the substantial power level change.

3.4.3.1. Temperature Coefficient of Reactivity

Figure 3.50 shows the temperature feedback coefficient evolution for the TAP reactor during the power change transient. The Fuel Temperature Coefficient ($\alpha_{T,F}$) became less negative during the first hour of the transient due to a slight spectrum hardening because the ^{135}Xe concentration peak changes the Doppler effect in the fuel salt. After turning the power back on, all three temperature coefficients of reactivity remain stable because the fuel salt composition remain almost unchanged. Overall, the isothermal temperature coefficient, $\alpha_{T,ISO}$ remains negative and strong throughout the postulated transient and fluctuates slightly within stochastic error range $\sigma_{\alpha_{T,ISO}} \pm 0.043\ pcm/K$.

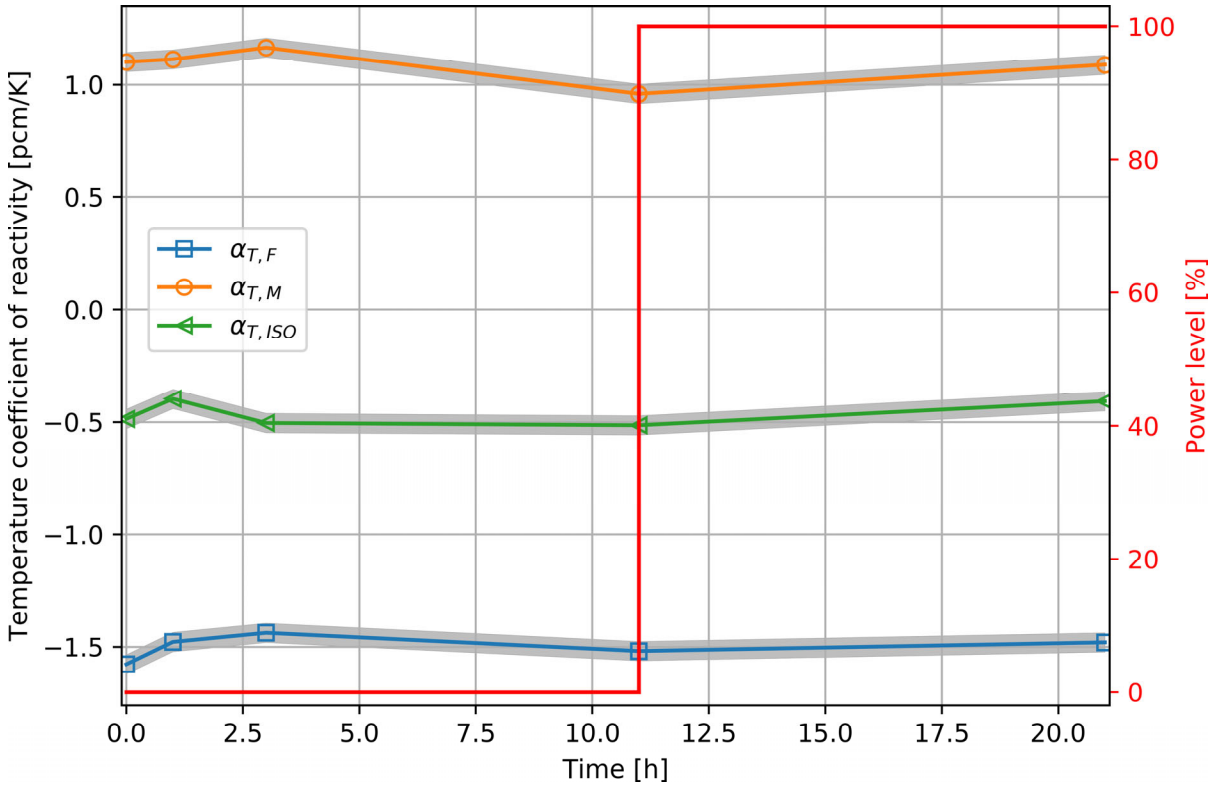

Figure 3.50: Temperature feedback coefficients during the postulated transient for the TAP reactor, 10 days before the EOL (all moderator rods inserted), the gas removal system operates with efficiency $\epsilon_{Xe} = 0.915$. The uncertainty, σ, is shaded.

3.4.3.2. Void Coefficient of Reactivity

Figure 3.51 demonstrates the void coefficient of reactivity evolution during the postulated transient. The αV remains almost constant throughout the postulated transient. All observed changes in the void coefficient of reactivity are due to the stochastic nature of the Monte Carlo method ($\sigma_{\alpha_V} \pm 4\ pcm/\%$).

3.4.3.3. Reactivity Control Rod Worth

Figure 3.52 demonstrates the control rod worth evolution during the power change transient. The control rod worth remains almost constant and sufficient to shut down the reactor throughout the postulated transient. During the first three hours of the transient, the control rod worth decreases from 1998.9±8.9 pcm to 1988.3±8.9 pcm due to a slight spectrum hardening caused by ^{135}Xe concentration raise. Overall, the control rod worth changes are insignificant and lie within the stochastic error range ($\sigma_{CRW} \pm 8.9\ pcm$).

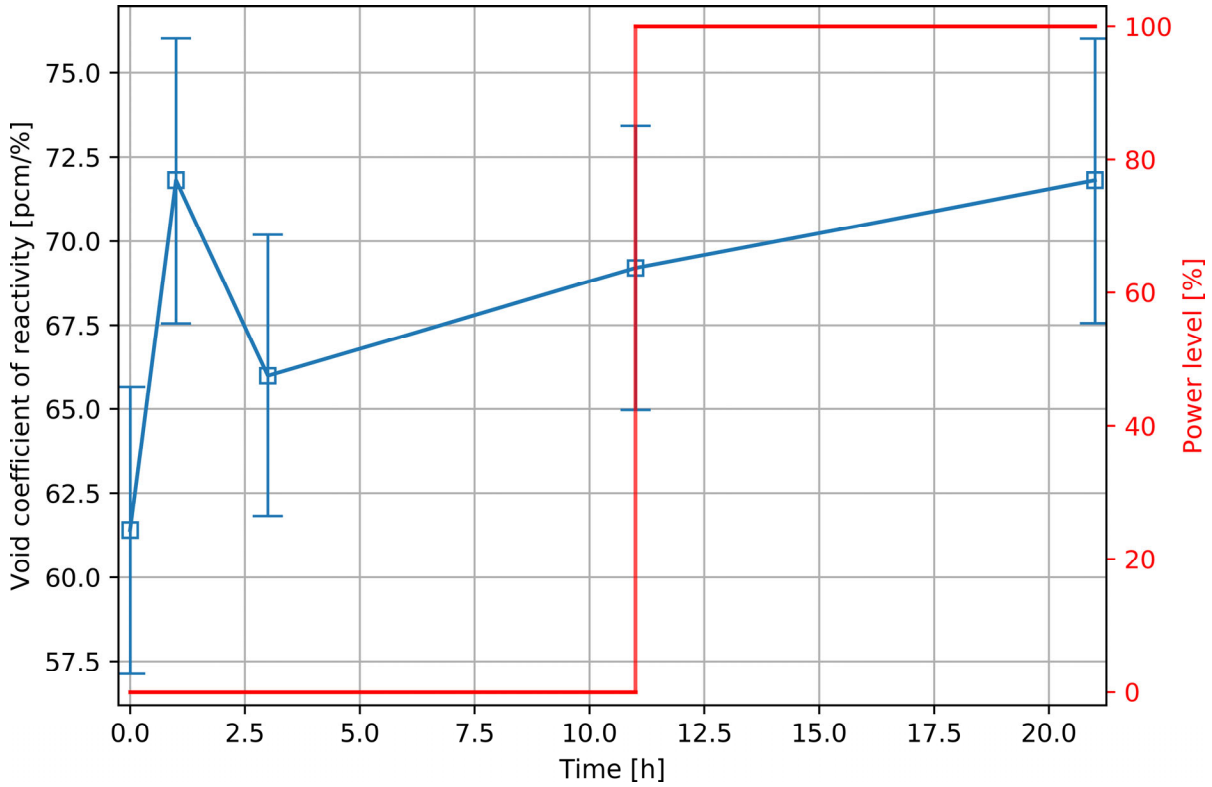

Figure 3.51: Void coefficient of reactivity as a function of time during postulated transient for the TAP reactor, 10 days before the EOL (all moderator rods inserted), the gas removal system operates with efficiency $\epsilon_{Xe} = 0.915$.

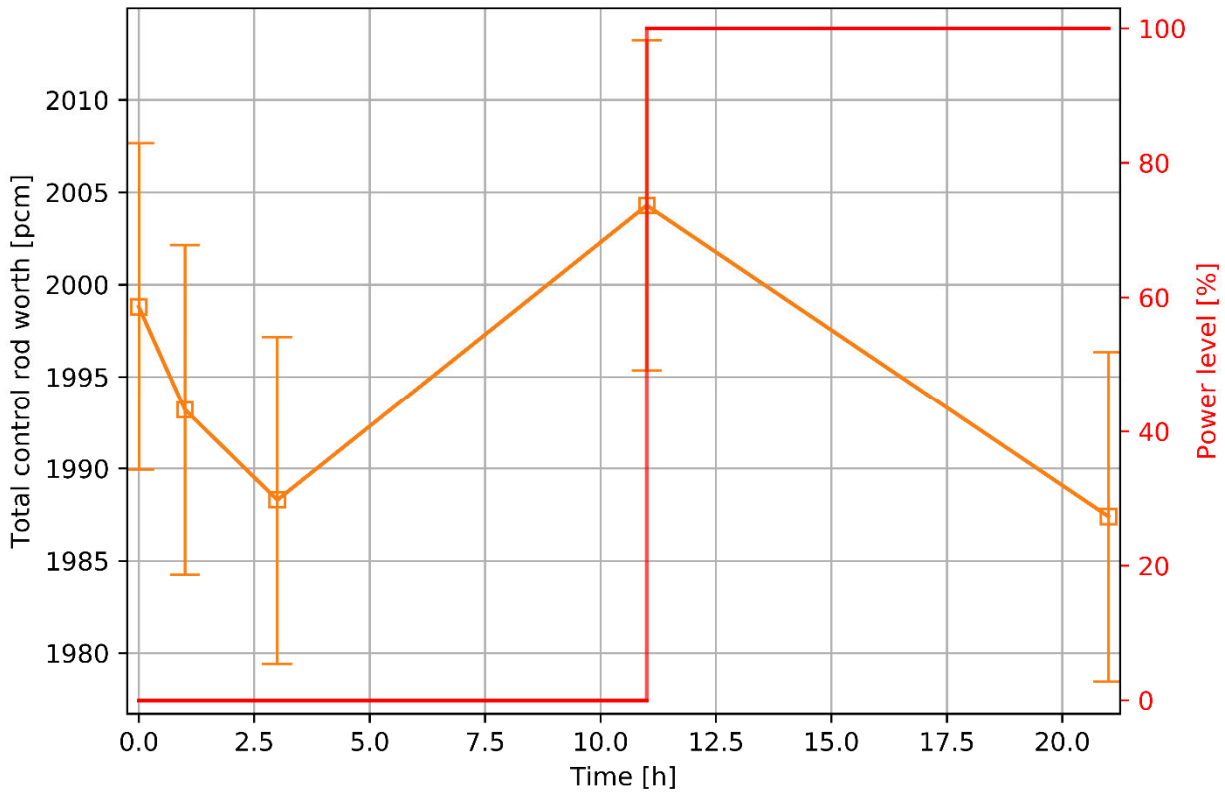

Figure 3.52: Total control rod worth as a function of time during postulated transient for the TAP reactor, 10 days before the EOL (all moderator rods inserted), the gas removal system operates with efficiency $\epsilon_{Xe} = 0.915$.

3.4.4. Concluding Remarks

This section demonstrated the short-term depletion simulations for the TAP reactor with the core power level variation in the range of [0, 100%] using SaltProc v1.0 and Serpent. two different noble gas removal scenarios are considered: (1) no gas removal (e.g., $\epsilon_{Xe} = 0$), and (2) fully operational gas removal system (e.g., $\epsilon_{Xe} = 0.915$). The results in the literature reported that negative xenon poisoning effect for conventional LWRs reaches its extremum $\Delta\rho$ 1500 pcm in approximately 11 hours after shutdown. Such a vast reactivity drop complicates the LWRs load-following.

For the case with no gas removal ($\epsilon_{Xe} = 0$), ^{135}Xe concentration peaks about 45 and 165 min after the shutdown at the BOL and EOL, respectively. The xenon concentration peaks sooner for the harder core configuration (e.g., at the BOL, SVF=0.9) because ^{135}Xe absorption cross section drops dramatically as neutron energy grows above 0.1 eV, thus, ^{135}Xe burn out faster in a harder spectrum. Thus, the harder spectrum leads to a smaller ^{135}I/^{135}Xe concentration ratio and, consequently, lower xenon concentration peak after shutdown. Without gas removal (e.g., ^{135}Xe loss after shutdown due to decay only) xenon concentration at the BOL remains almost constant ($\Delta N_{^{135}Xe} = 0.33\%$), and no effect of xenon poisoning was observed ($\Delta\rho$ = -10±7 pcm). However, at the EOL, when all moderator rods are in and the neutron spectrum is more thermal, a more

significant effect of poisoning is observed: ^{135}Xe concentration increased by 4% with corresponding negative reactivity insertion of -70 pcm. For the case with very effective noble gas removal (ϵ_{Xe} = 0.915), the time when ^{135}Xe concentration peaks cannot be predicted analytically, because, after shutdown, the gas removal system removes a major fraction of xenon gas at the end of each depletion step. Moreover, the ^{135}I/^{135}Xe ratio is significantly greater (e.g., between 8.66 and 11.42) than for the non-removal (ϵ_{Xe} = 0) case. Thus, ^{135}I decay leads to a substantial increase in ^{135}Xe concentration right after shutdown (up to +200% at the EOL), and corresponding reactivity drop (-108±5 pcm at the EOL). However, after the first hour, reactivity increases quickly because the gas removal system extracts most of the xenon every 1 hour (the selected SaltProc v1.0 depletion time step). The true effect of xenon poisoning for the TAP reactor with active gas removal is expected to be even less severe because the real system would remove noble gases continuously, not discretely as simulated by SaltProc (e.g., xenon would be removed with ϵ_{Xe} = 0.915 every moment, not once per hour). Overall, more realistic results for load-following transients can be obtained with better time resolution. Though the ideal depletion time step should be closer to full loop time (20 seconds), such fidelity would require an enormous computation burden.

Finally, this section demonstrated that the TAP reactor maintains required safety margins during postulated transients. The temperature feedback coefficients, void coefficient of reactivity, and control rod worth all remain within stochastic uncertainty throughout the transient. Small elevation in total temperature coefficient and void coefficient of reactivity during the first hour after shutdown is due to the ^{135}Xe concentration raise and corresponding short-term neutron spectrum hardening. In conclusion, the TAP MSR, even without gas removal, is capable of the safe restart after reducing power from 100% to 0% even when ^{135}Xe concentration peaks. While this work has confirmed neutronics feasibility of resilience against the iodine pit, separate thermomechanical structural analysis is needed to confirm that structural materials could withstand such dramatic core power fluctuations.

3.5. Load-following and Safety Analysis of Molten Salt Breeder Reactor

The previous section has shown that the TAP MSR is unaffected by xenon poisoning during power variation because it has a relatively fast neutron energy spectrum. While long-term performance metrics such as fuel utilization would definitely benefit from online removal of poisonous fission products, the gas removal system is not necessary to ensure safe TAP system operation during a short-term power drop and restart transient. However, Section 3.4 clearly demonstrated a strong impact of the noble gas removal on the reactor neutronics during power adjustments. Thus, another liquid-fueled MSR design with thermal spectrum (not epithermal like in the TAP core) was considered to investigate the benefits of the online gas removal for load-following operation.

This section presents fuel salt depletion analysis with SaltProc during a short-term power transient to evaluate load-following capabilities of the graphite-moderated molten salt reactor design with a thermal neutron energy spectrum - Molten Salt Breeder Reactor (MSBR). The details of the MSBR design, the full-core Serpent model, and the results of long-term depletion simulation with SaltProc were described in a previous report [162]. The load- following transient postulated in Section 3.4.2.1 is simulated. To investigate the effect of noble gas removal efficiency on the

load-following operation, three various regimes of the gas removal system operation are considered:

a) no gas removal (ϵ_{Xe} = 0.0);

b) moderate gas removal efficiency (ϵ_{Xe} = 0.536);

c) high gas removal efficiency (ϵ_{Xe} = 0.915).

Then, major safety and operational parameters are calculated for all three regimes at various moments of the transient to ensure that the critical safety margins are maintained. Finally, the TAP MSR and MSBR behavior during the postulated load-following transient are compared.

3.5.1. Depletion Analysis Results

The methodology described previously in Section 3.4 is used for the MSBR full-core depletion analysis with SaltProc v1.0 with a 30-minute depletion time step to capture rapid changes in reactivity. Equation 3-23 predicted the time after shutdown when ^{135}Xe concentration peaks (t_X^{max}) in the range from 6.8h (ϵ_{Xe} = 0.0, 30 years after startup) to 7.5h (ϵ_{Xe} = 0.915, BOL). The t_X^{max} for the MSBR is longer than for the TAP reactor (2.75h) due to a much more thermal neutron energy spectrum. To be consistent throughout different gas removal regimes while investigating load-following capabilities of the MSBR, the following transient (power load profile) very similar to the transient chosen in Section 3.4 is selected:

a) operate on 100% of HFP to reach ^{135}I/^{135}Xe equilibrium (at least 3 days from the startup).

b) instantaneous power drop from 100% to 0%.

c) shutdown for t_X^{max} 7.5h to reach the ^{135}Xe concentration extremum.

d) instant restart from 0% to 100% power level and operate on 100% for 5 hours.

3.5.1.1. Reactivity Dynamics

Figure 3.53 and Figure 3.54 show the effective multiplication factor and reactivity dynamics for the various gas removal efficiencies in the MSBR during the transient, described earlier. For the no-removal case (Figure 3.53, upper panel), the effective multiplication factor dropped after t_X^{max} = 7.5h by 1457 pcm and 1035 pcm at BOL and after 15 years of full-power operation, respectively. Thus, the Equation 3-23 correctly predicted the moment when the xenon poisoning effect maximized in the no-removal case (ϵ_{Xe} = 0). After the power ramp-up from 0% to 100%, the effective multiplication factor returned to its initial value in about 3 hours. Notably, maximum negative reactivity insertion due to ^{135}Xe buildup after the MSBR shutdown is very similar to the PWR (both at startup): 1457 pcm and 1500 pcm [183], respectively. Additionally, the xenon poisoning effect diminished toward the EOL because the ^{135}Xe concentration peak is more significant for the softer thermal spectrum (the MSBR spectrum becomes harder during operation due to plutonium and other strong neutron absorbers accumulation in the fuel salt). Finally, the effect of ^{135}Xe poisoning is almost the same after 15 and 30 years of operation because the fuel salt composition reaches its equilibrium after about 16 years of full-power operation (see Section 3.3.2 of Rykhlevskii's thesis [162]).

The middle and lower plots in Figure 3.54 show reactivity change during the MSBR shutdown for 7.5 hours and following power ramp up to 100% for moderate (ϵ_{Xe} = 0.536) and high (ϵ_{Xe} = 0.915) removal efficiency, respectively. In contrast with no gas removal, reactivity dropped during the 30-minutes interval after shutting down by 161 pcm and 189 pcm for moderate and high removal efficiency, respectively. Afterward, the reactivity boosts by 1494 pcm and 2608 pcm for ϵ_{Xe} = 0.536 and 0.915, respectively. Such reactivity change happens because the gas removal system extracted 53.6% and 91.5% of xenon mass at the end of the 30-minute depletion step. The more effective xenon removal leads to greater positive reactivity jump, as expected. Notably, the reactivity stabilizes at approximately 2500 pcm level about 5 hours after the shutdown because the ^{135}Xe loss due to its decay and online gas removal equalizes ^{135}Xe gain from ^{135}I decay. Overall, the online gas removal from the fuel salt even with moderate efficiency is beneficial to the core neutronics and significantly reduces the xenon poisoning effect (-161±10 pcm instead of -1494±10 pcm). Moreover, the very high removal efficiency (ϵ_{Xe} = 0.915) is unnecessary to significantly reduce the effect of xenon poisoning and enable the load-following capability of the MSBR.

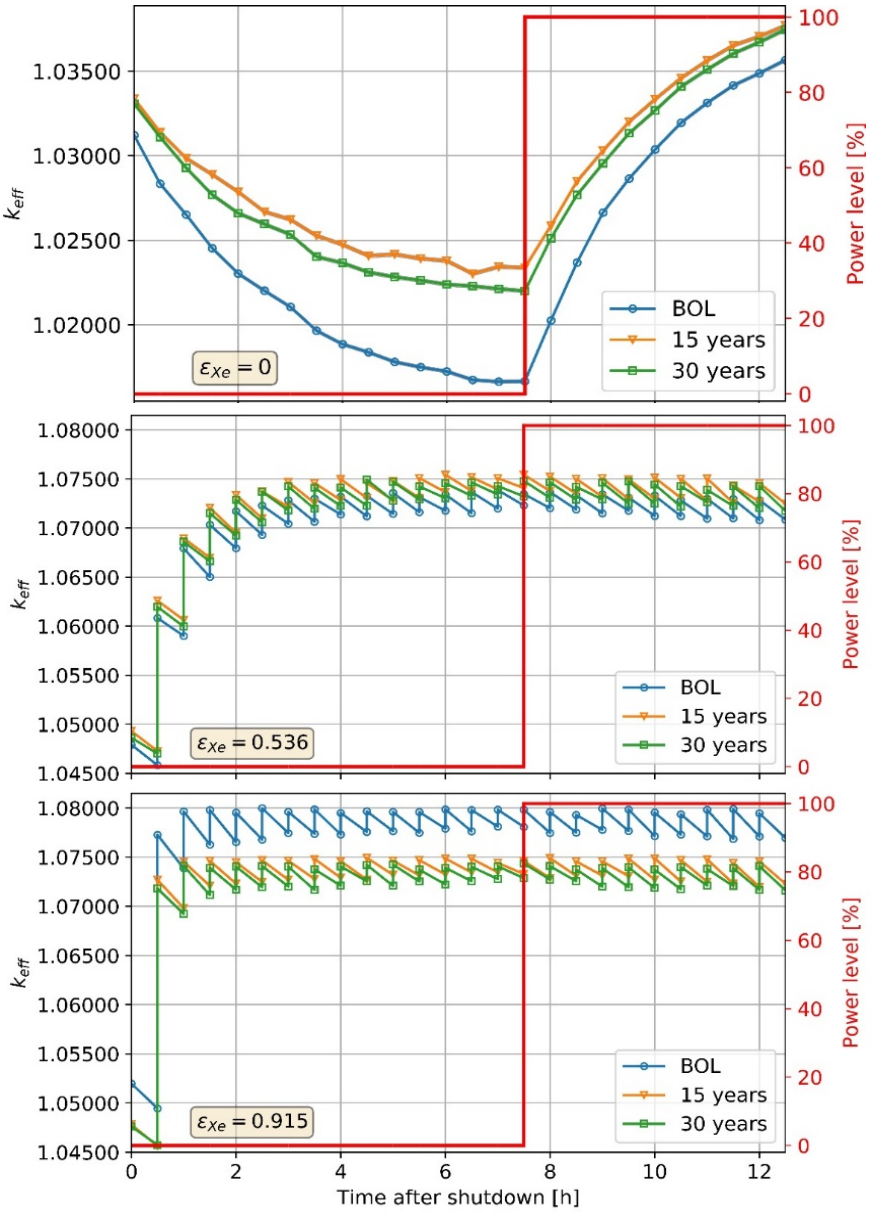

Figure 3.53: SaltProc-calculated evolution of the effective multiplication factor during the postulated load-following transient for various regimes of the gas removal system operation. The uncertainty (σ = 10 pcm) is shaded.

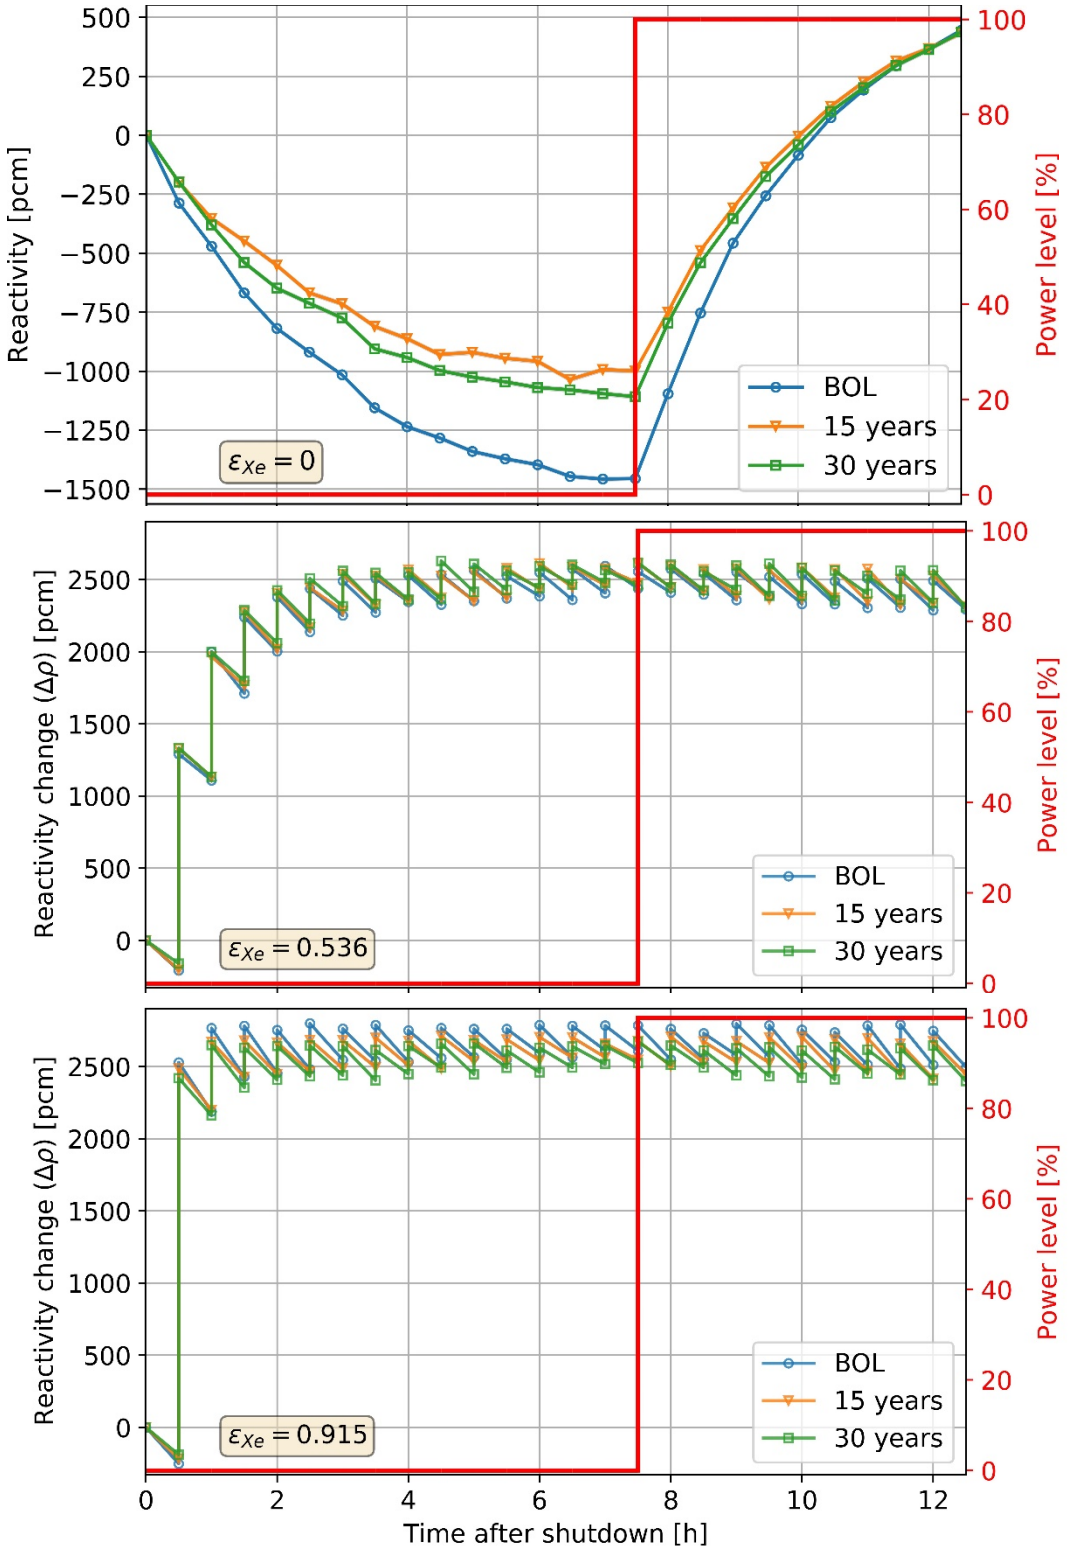

Figure 3.54: SaltProc-calculated evolution of the reactivity during the postulated load-following transient for various regimes of the gas removal system operation. The uncertainty (σ = 10 pcm) is shaded.

3.5.1.2. Fuel Salt Composition Evolution

Figure 3.55 shows ^{135}Xe and ^{135}I mass dynamics evolution during the postulated transient for various gas removal efficiencies. The ^{135}I/^{135}Xe concentration ratio at the beginning of the transient for the no-removal case is 2.45 and 2.03 at the BOL and after 30 years of full-power operation, respectively. Because the ^{135}I/^{135}Xe concentration ratio is greater at startup, the ^{135}Xe concentration peak is 11% higher than at the EOL, which is consistent with the TAP MSR results. However, a larger ^{135}Xe concentration does not necessarily worsen the xenon poisoning effect (Figure 3.54) because the spectrum hardens toward EOL and the ^{135}Xe absorption cross section slumps with higher neutron energy (see Figure 3.49). For the high gas removal efficiency regime, the ^{135}I/^{135}Xe concentration ratio is 2.47 and 2.08 at the BOL and after 30 years of full-power operation, respectively. For the BOL and a 30-year case, the ^{135}Xe concentration peaked only by 8% at the end of a first 30- minute depletion step, which caused a 189-pcm negative reactivity insertion. Afterward, the concentration of ^{135}Xe dropped quickly because the gas removal system extracted most of the fission gas. The ^{135}Xe concentration in the fuel salt before the shutdown is approximately 7 times greater than after the power turned back on, which caused significant reactivity growth by 2550 pcm. Surprisingly, the removal of 12 g of ^{135}Xe from t = 30min to t = 60min caused an impressive 2600-pcm positive reactivity insertion (217 pcm/g^{135}Xe reactivity worth).

Such large fluctuations in the ^{135}Xe concentration are observed due to the batch-wise nature of SaltProc simulations (e.g., the fraction of target poison is being removed discretely at the end of each depletion step). Realistically, the gas removal system would extract gas from the fuel salt continuously, which would result in a much smoother change in the concentration and, accordingly, in the reactivity. Notably, for both BOL and EOL, the ^{135}Xe mass stabilized at 1 g in about 3-4 hours after the shutdown and then inclined slowly (60 mg EFPH) after power ramp-up from 0 to 100%. That is, when the reactor returns to a full-power level, the ^{135}Xe concentration during a few days will be significantly lower than before the load-following transient. Thus, fewer thermal neutrons will be parasitically absorbed in the fission gas. As a result, long-term fuel cycle performance metrics such as fuel utilization and core lifetime would benefit enormously from a very low ^{135}Xe concentration in the core after the postulated transient. In other words, the transient cleans up the fuel salt, but more analyses are required to evaluate all benefits of this finding.

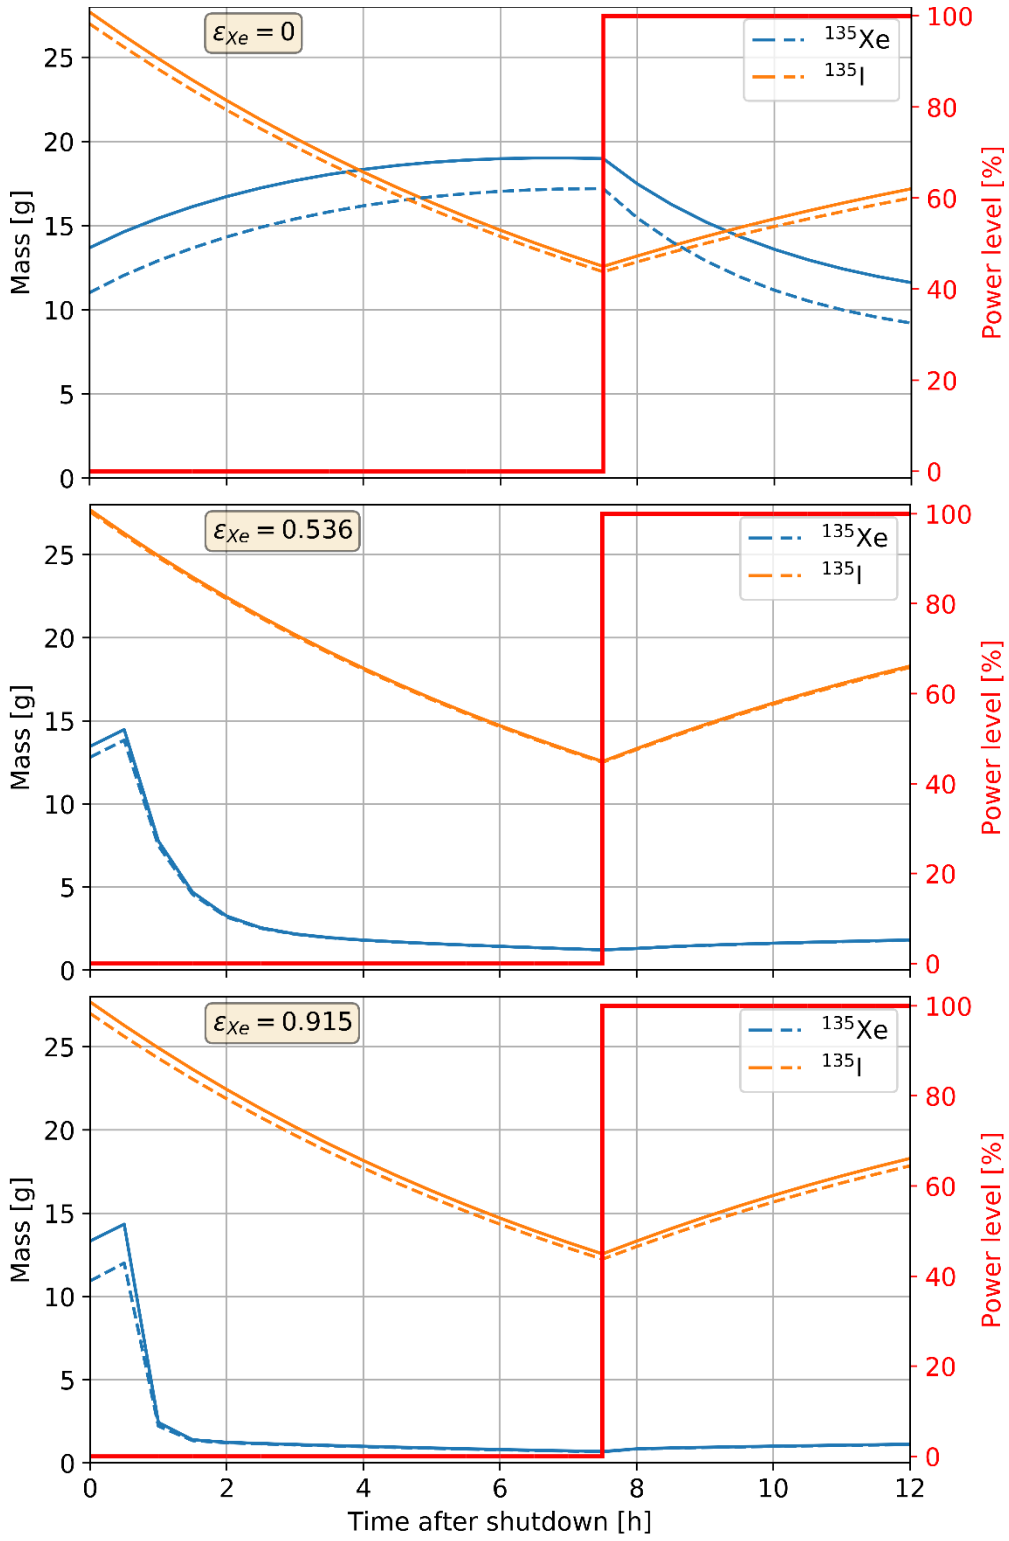

Figure 3.55: Comparison of ^{135}Xe and ^{135}I isotopic content at the BOL (dashed line) and after 30 years of operation (solid line) for various gas removal regimes.

In the case of moderate gas removal efficiency, the fission product concentration changes very similarly to a high removal efficiency case. The $^{135}I/^{135}Xe$ concentration ratio is 2.15 and 2.06 at the BOL and after 30 years of full-power operation, respectively, and caused a 7.5% hike in ^{135}Xe concentration. Surprisingly, a significantly lower gas removal efficiency ($\epsilon_{Xe} = 0.536$ instead of 0.915) provided comparable benefits to the core neutronics during the postulated load-following transient. Similarly, to the $\epsilon_{Xe} = 0.915$ case, the ^{135}Xe mass stabilized at 1.5 g about 5 hours after the shutdown and then increased slowly (165 mg EFPH) after power ramp-up from 0 to 100%. In conclusion, a simpler and cheaper gas removal system with extraction efficiency $\epsilon_{Xe} = 0.536$ is sufficient to suppress the xenon poisoning effect to an acceptable level (-161 pcm) and improve the load-following capability of the MSBR.

3.5.1.3. Neutron Spectrum

Figure 3.56 shows that the MSBR spectrum after 30 years of operation (solid line) is harder than at the startup (dashed line). Compared with the MSBR, the TAP MSR spectrum is significantly harder even when all moderator rods are inserted to the core. Notably, the MSBR spectrum has a clear peak in the thermal energy region, but flat neutron energy dependence in intermediate and fast energy region, which is quite common for thermal reactors. In contrast, the TAP core spectrum at the EOL has a high peak in the fast and lower peak in the thermal energy region, which is typical for epithermal/intermediate reactors. This is the main reason why for the postulated load-following transient, a significant xenon poisoning effect is observed in the MSBR , and negligible xenon impact is observed in the TAP MSR (see Section 3.4).

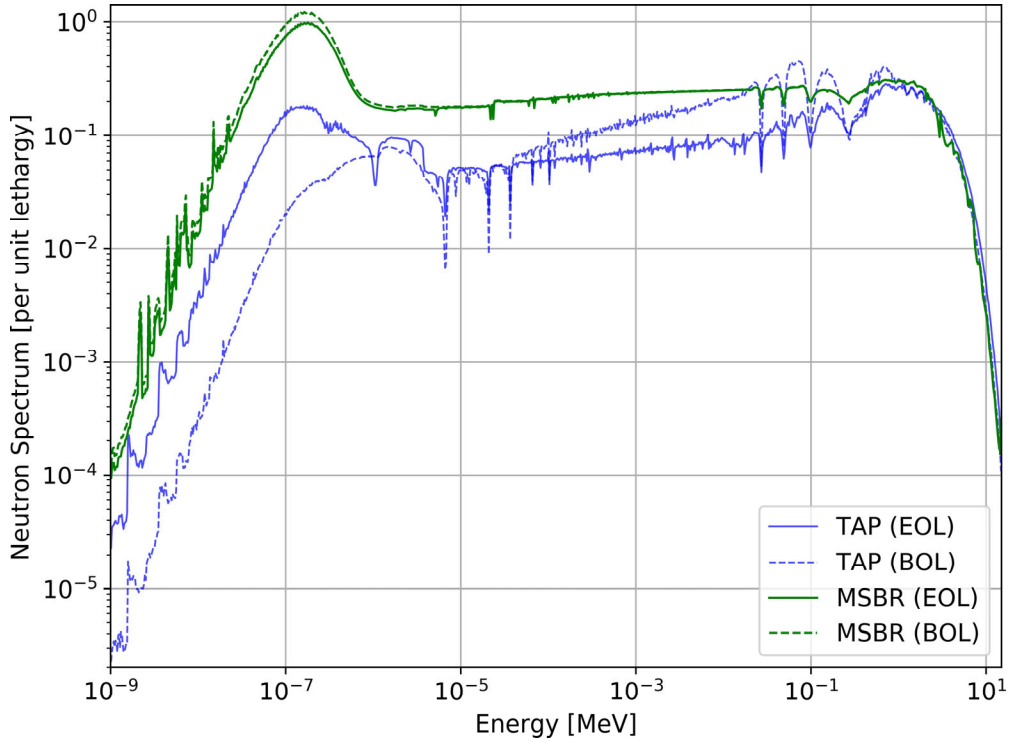

Figure 3.56: Neutron spectra normalized by lethargy for the MSBR and TAP at various moments during operation. The neutron flux uncertainties σ_Φ are 0.6% and 0.18% for the TAP reactor and MSBR, respectively.

Any graphite-moderated liquid-fueled MSR conceptual design [13] would potentially demonstrate similar benefits from an online noble gas removal from the fuel salt.

3.5.2. Safety and Operational Parameters

The significant change of strong absorber concentrations in the fuel slightly shifts the core spectrum, potentially impacting the reactor's safety. Since rapid changes in the fuel salt composition cannot be allowed to compromise critical safety margins, major safety and operational parameters are calculated at various moments throughout the postulated transient using approaches from Sections 3.3.3 and 3.4.3. The total temperature coefficient of reactivity (α_{ISO}) must remain negative, and the total control rod worth (CRW) must be sufficient to trip the reactor throughout the postulated transient. Ideally, major safety and operational parameters should stay almost constant because the changes in those parameters would require fast response from the reactor control systems (i.e., control rod jerk in response to a CRW change).

3.5.2.1. Temperature Coefficient of Reactivity

Figure 3.57 shows the temperature feedback coefficient dynamics for the MSBR during the transient for various gas removal efficiencies (ϵ_{Xe} = 0.536 and 0.915). The Fuel Temperature Coefficient ($\alpha_{T,F}$) becomes less strong at the beginning of the transient for all cases. The reason for this is a slight spectrum hardening due to the ^{135}Xe concentration peak that changed the Doppler broadening of resonances. After that, the magnitude of $\alpha_{T,F}$ slowly increased due to a steady incline in the ^{135}Xe concentration.

The isothermal temperature coefficient, α_{ISO}, is -0.36±0.09 pcm/K at the beginning and remains stable during the first 30 minutes of the transient for the moderate removal efficiency case. Then, as the gas removal system reduces ^{135}Xe concentration in the core, α_{ISO} becomes even more negative: -1.52±0.09 pcm/K when the ^{135}Xe mass stabilized at 1.5 g in about 5 hours after the shutdown. After power ramp-up from 0% to 100%, α_{ISO} also remains stable since the ^{135}Xe mass increasing very slowly. On the whole, another exciting benefit from the online gas removal is improved passive safety (more powerful temperature feedback coefficient) throughout and, possibly, a few days after the postulated transient due to low concentration of the ^{135}Xe in the fuel salt.

For the high gas removal efficiency regime (ϵ_{Xe} = 0.915), the isothermal temperature coefficient worsens from -0.54 ± 0.09 pcm/K to approximately -0.22 ± 0.09 pcm/K during first the 30 minutes after shutdown. Afterward, however, once the gas removal system extracted a significant fraction of the ^{135}Xe from the fuel salt, α_{ISO} recovered, becoming significantly more negative (-1.39 and -1.56 pcm/K at the BOL and after 30 years of operation, respectively) due to the spectrum softening. In brief, the temperature feedback in the MSBR becomes stronger when neutron poisons concentration in the fuel decreases. As a result, flattening the ^{135}Xe concentration curve improves the MSBR passive safety.

[13] Integral Molten Salt Reactor (IMSR) from Terrestial Energy [186], Molten Salt Demonstration Reactor (MSDR) from Oak Ridge National Laboratory [187], Liquid fluoride thorium reactor (LFTR) from Flibe energy [188], etc.

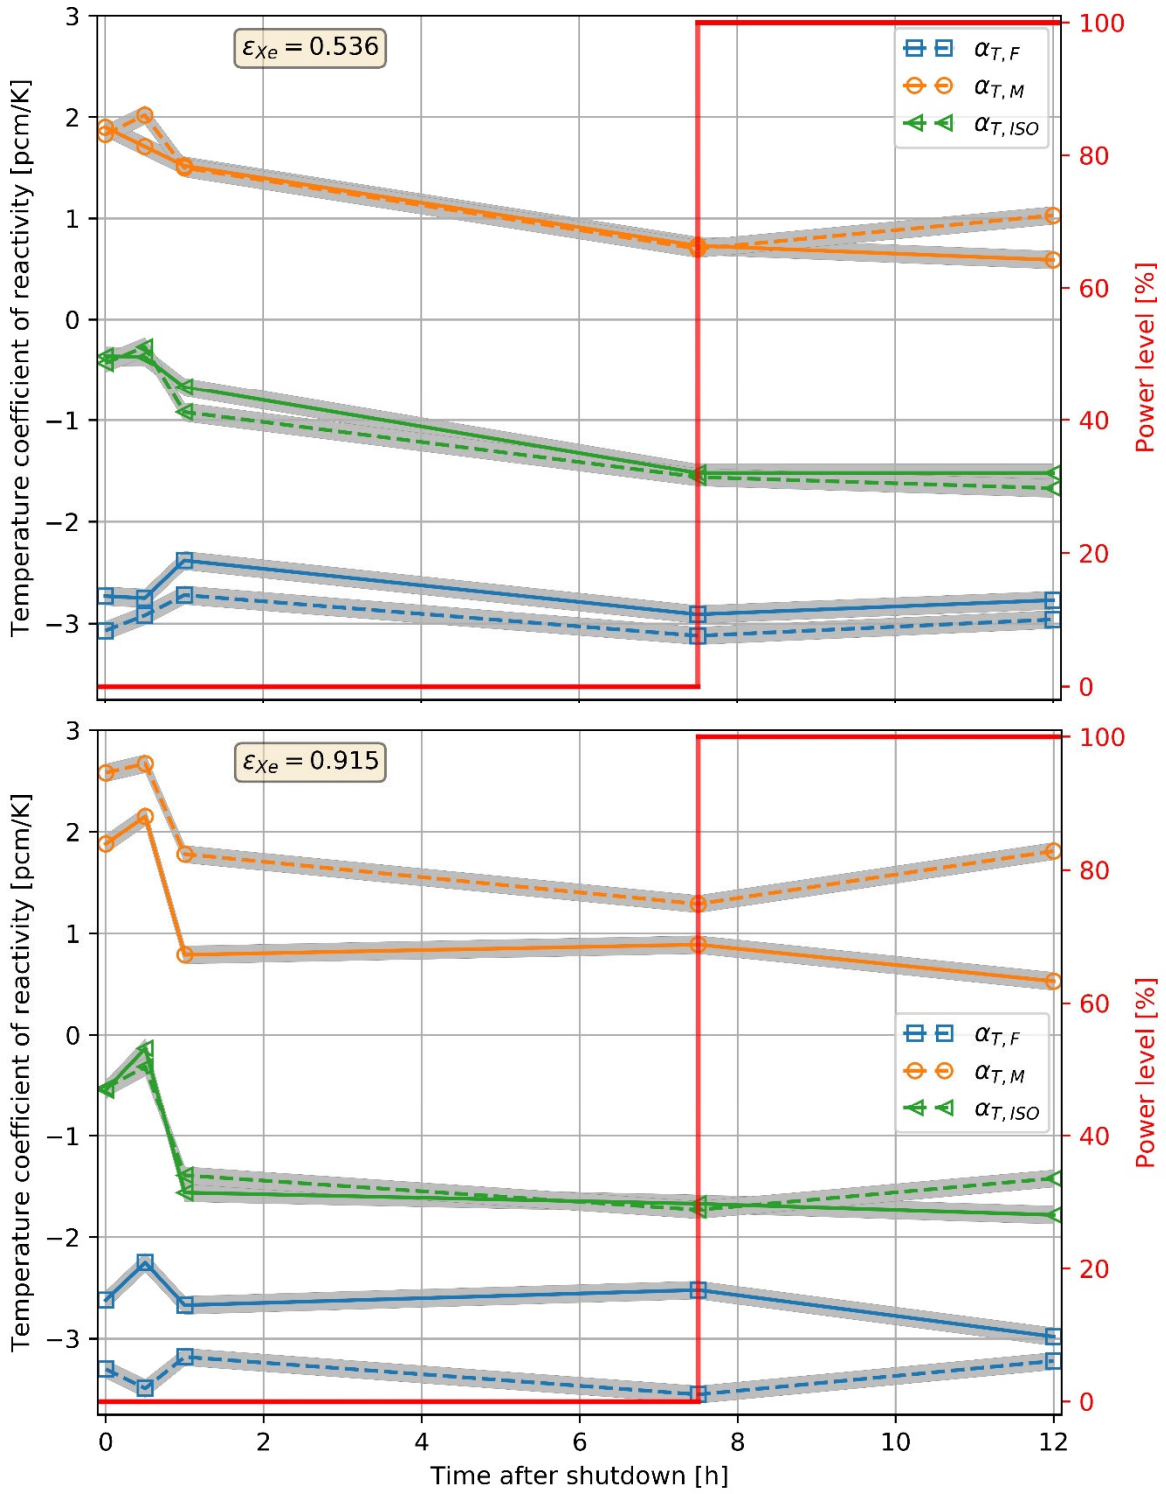

Figure 3.57: Temperature feedback coefficients during the postulated transient for the MSBR operating with moderate ($\epsilon_{Xe} = 0.536$, upper) and high ($\epsilon_{Xe} = 0.915$, lower) gas removal efficiency at the BOL (dashed line) and after 30 years of operation (solid line). The uncertainty, $\pm\sigma$, is shaded.

Overall, the combination of fuel and moderator thermal feedback coefficients, α_{ISO}, remains negative throughout the postulated transient. Moreover, a simpler and cheaper gas removal system with extraction efficiency $\epsilon_{Xe} = 0.536$ provided more predictable thermal feedback coefficient dynamics throughout the transient due to a more gradual change in the ^{135}Xe concentration.

3.5.2.2. Void Coefficient of Reactivity

Figure 3.58 demonstrates the void coefficient of reactivity evolution during the postulated transient. In contrast with the TAP MSR, the void coefficient of reactivity after 30 years of full-power operation is substantially higher than at the startup for both gas removal regimes. The reason for this is the hardening of the MSBR spectrum toward EOL, which is the opposite of the TAP MSR spectrum evolution. Thus, an unexpected void insertion due, for example, to a gas separation system failure in the MSBR would have more severe consequences for the EOL.

For the high gas removal efficiency, αV fluctuates during the postulated transient between 42 and 61 pcm void% at the BOL and between 87 and 102 pcm/void% after 30 years of operation. The ^{135}Xe concentration spike caused corresponding α_V drop due to the short-term spectrum hardening. Then, αV quickly recovers to its initial value. Similarly, to the temperature feedback coefficient, the moderate gas removal efficiency provided more predictable αV dynamics throughout the transient. Additionally, a small αV fluctuation during the transient at the EOL for the case with $\epsilon_{Xe} = 0.536$ ($\Delta\alpha_V \approx 25$ pcm/void%) would simplify the gas separator backup safety mechanism. Overall, all observed changes in the void coefficient of reactivity throughout the load-following transient for all cases are within the 3-σ range ($\sigma_{\alpha_V} \pm 5\ pcm/\%$). These observations should be taken into account in the MSBR accident analysis and safety justification.

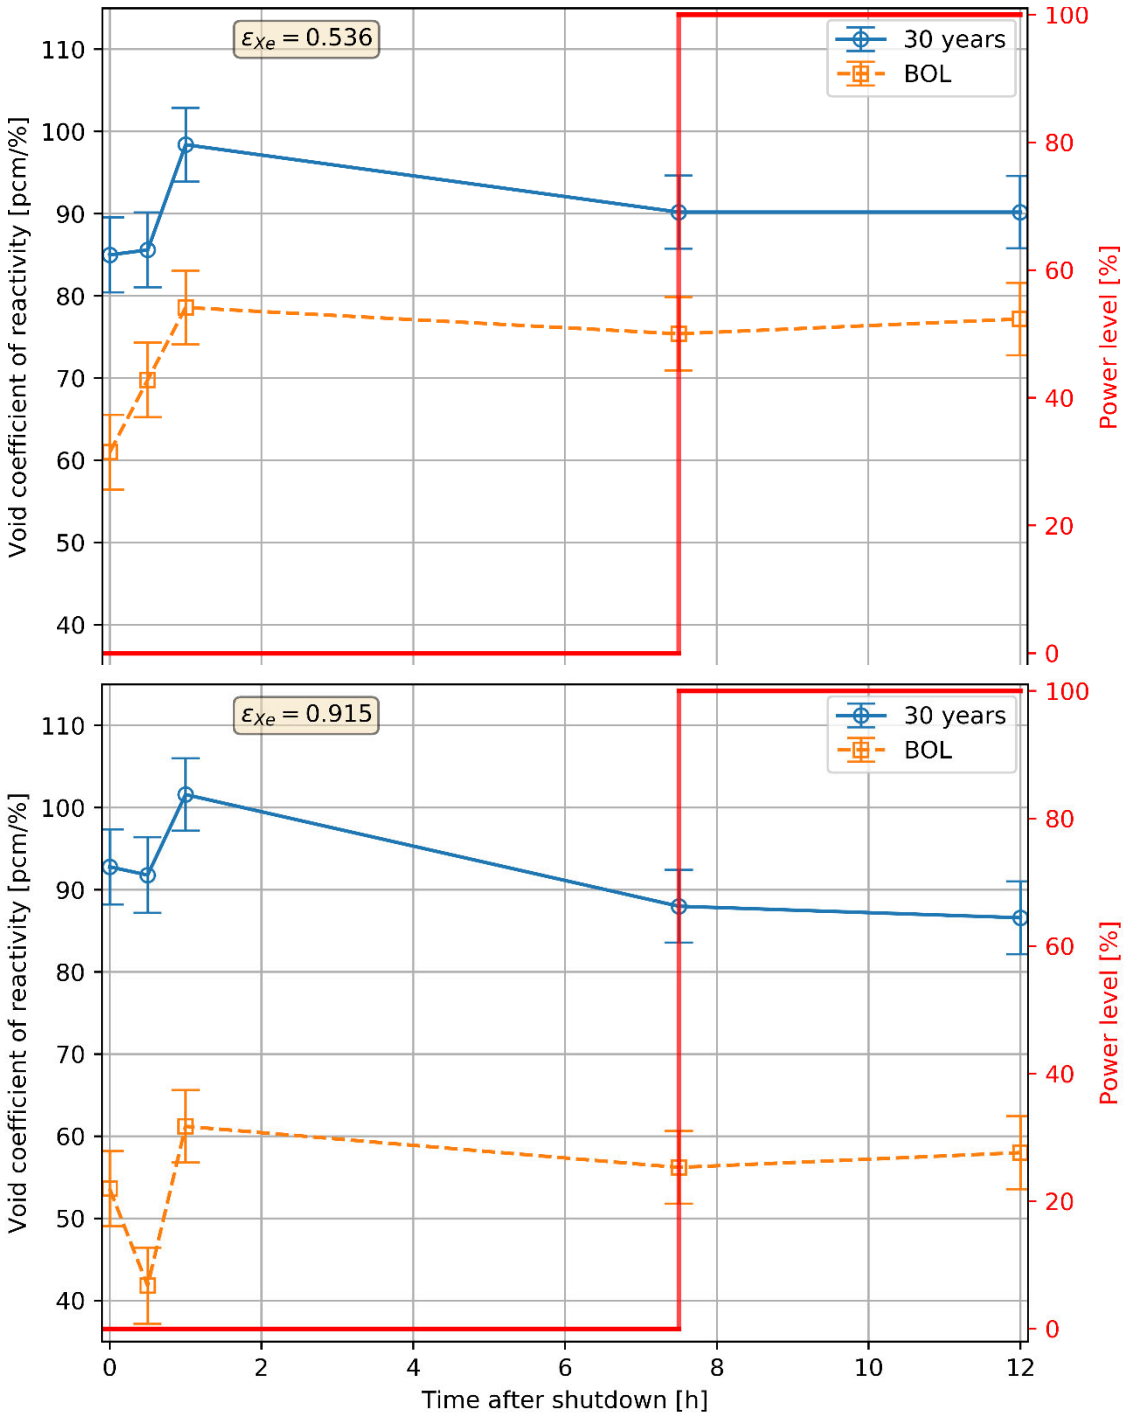

Figure 3.58: Void coefficient of reactivity as a function of time during postulated transient for the MSBR operating with moderate (ϵ_{Xe} = 0.536, upper) and high (ϵ_{Xe} = 0.915, lower) gas removal efficiency at the BOL (dashed line) and after 30 years of operation (solid line).

3.5.2.3. Reactivity Control Rod Worth

Figure 3.59 shows the control rod worth evolution during the postulated transient. For the high gas removal efficiency regime after 30 years of full-power operation, the control rod worth dropped by 46±9 pcm during the first 30 minutes after the shutdown. This happens due to a short-term spectrum hardening related to the ^{135}Xe concentration peak. In the next 30 minutes, the CRW recovers to its initial value and keeps increasing throughout the transient because the gas removal system steadily reduces the ^{135}Xe concentration in the core. Notably, the control rod worth is greater at the BOL because the absorption cross section of ^{10}B (used as an absorber in the control rods) declines rapidly with energy. Overall, the control rod worth benefits from the MSBR spectrum softening toward EOL.

For the moderate gas separation efficiency regime, the control rod worth remains almost constant during the first hour after shutdown. Afterward, the CRW increased by 4% due to the spectrum softening caused by the increased ^{135}Xe concentration. As for other safety parameters, the control rod worth also benefits from a less effective gas removal system due to smother xenon concentration dynamics and a more predictable neutron spectrum shift.

Unfortunately, the total control rod worth is insufficient to shut down the reactor throughout the postulated transient for both medium and high removal efficiency ($\epsilon_{Xe} = 0.536$ and $\epsilon_{Xe} = 0.915$). The reactivity change during the transient is up to 2600 pcm, while the total control rod worth is only about 1250-1425 pcm. The MSBR was designed with only two graphite and two boron-carbide rods located in the center of the core (see Figure 1.6) for operative reactivity control and relied heavily on fissile feed adjustment as a primary reactivity control mechanism. However, the fissile feed cannot be adjusted quickly, and nuclear regulations require control rods to have sufficient worth to shut down the reactor safely at any time. Therefore, the control rod design in the MSBR must be reexamined to ensure the total control rod worth of at least 3000 pcm to ensure safety during the transient with rapid power change.

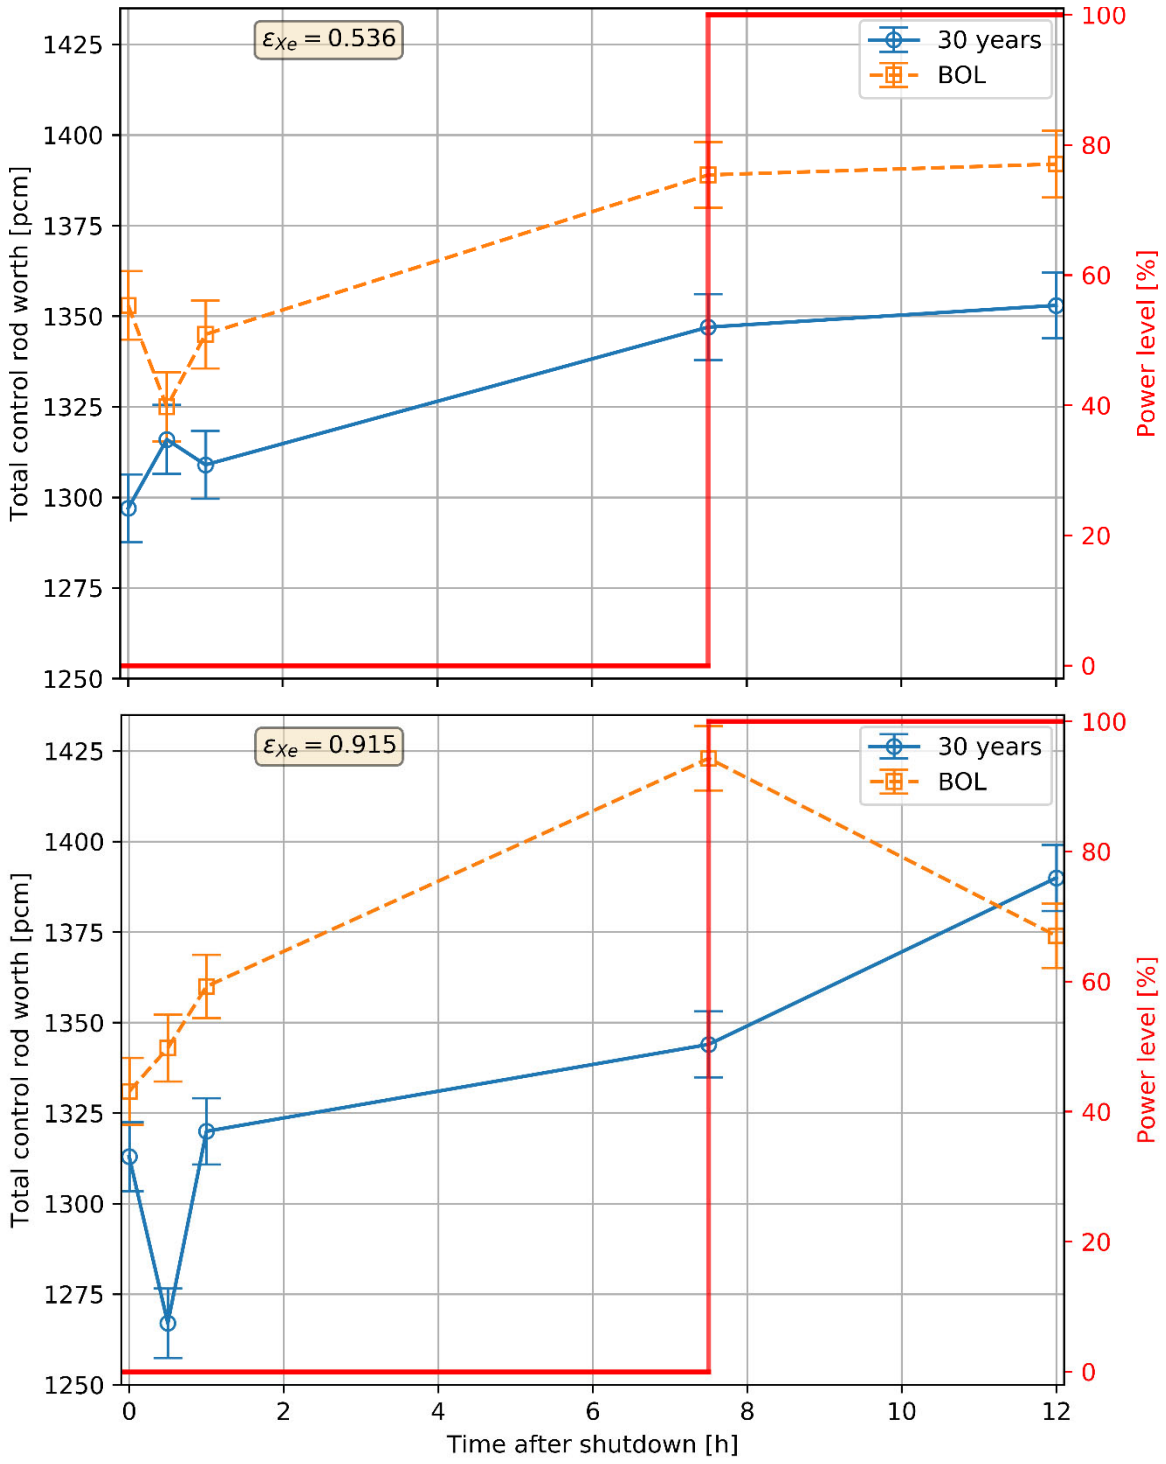

Figure 3.59: Total control rod worth as a function of time during postulated transient for the MSBR operating with moderate ($\epsilon_{Xe} = 0.536$, upper) and high ($\epsilon_{Xe} = 0.915$, lower) gas removal efficiency at the BOL (dashed line) and after 30 years of operation (solid line).

3.5.3. Additional Load Following Transients Simulations

In addition to the load follow analysis performed in Section 3.5.1.1, more simulations are performed to design a feasible sparging system. Reactor core behavior at different load-following transients are evaluated and decisions are drawn on what removal efficiency is needed to maintain the criticality in the MSBR when the sparging system is enabled.

Two critical load following scenarios were investigated:

- The first worst-case scenario simulates 8 hours at full power after an 8-hour shut-down, providing maximum Xe poisoning effect in the reactor.
- The second worst-case scenario considers a short period load-following for maxi-mum Xe accumulation over time. In this scenario, the reactor runs at full power for one hour after an hour of shutdown, and this repeats several times.

First, sensitivity analysis for different load-following transients is carried out. The results from Section 3.5.1.1 and Figure 3.60 point out that MSBR cannot do load-following without gas removal at BOL (30 days), MOL (15 years), and EOL (30 years) as the effective multiplication factor decreases with the start of the shutdown. Online gas removal from the fuel salt even with moderate efficiency significantly reduces the xenon poisoning effect, yet very high removal efficiency seems unnecessary to negate the negative effect of xenon poisoning. Load-follow at EOL is the worst for k_{eff} and consequently considered for sensitivity analysis.

- First Scenario

Figure 3.61 shows the results of the first scenario (single load-follow) for k_{eff}. After a lifetime of operation at $\epsilon_{Xe} = 0.536$, single load-following was attempted. In this transient, for the base case geometry, the reactor can recover from the Xe poisoning effect after $\epsilon_{Xe} = 26.8\%$. Generally, increasing gas removal efficiency increases excess reactivity. If higher efficiency is used, then the reactor recovers excess reactivity quicker, within a few hours.

As to the breeding ratio depicted in Figure 3.62, single load-following transient results in a gradual decrease. Increasing the gas removal efficiency slightly lowers the breeding ratio during the load-follow.

For the delayed neutron fraction (β_{eff}) given in Figure 3.63, no significant change is observed. Instead, β_{eff} fluctuates in a narrow range due to the statistical deviation.

Multiple consecutive load-following transients cause sharp changes in salt composition. As can be clearly seen in Figure 3.64, k_{eff} begins fluctuating with Xe buildup and burndown period. It is concluded from the result that to keep the reactor stable, gas removal efficiency at least $\epsilon_{Xe} = 53.6\%$ is required.

With the same load-following period, the number of transients is increased. It is observed that a higher gas removal efficiency (at least $\epsilon_{Xe} = 76.9\%$) is needed to keep the reactor stable. Therefore, these results indicated that as the number of power ramps increases, a higher gas removal efficiency is required for stable reactor behavior.

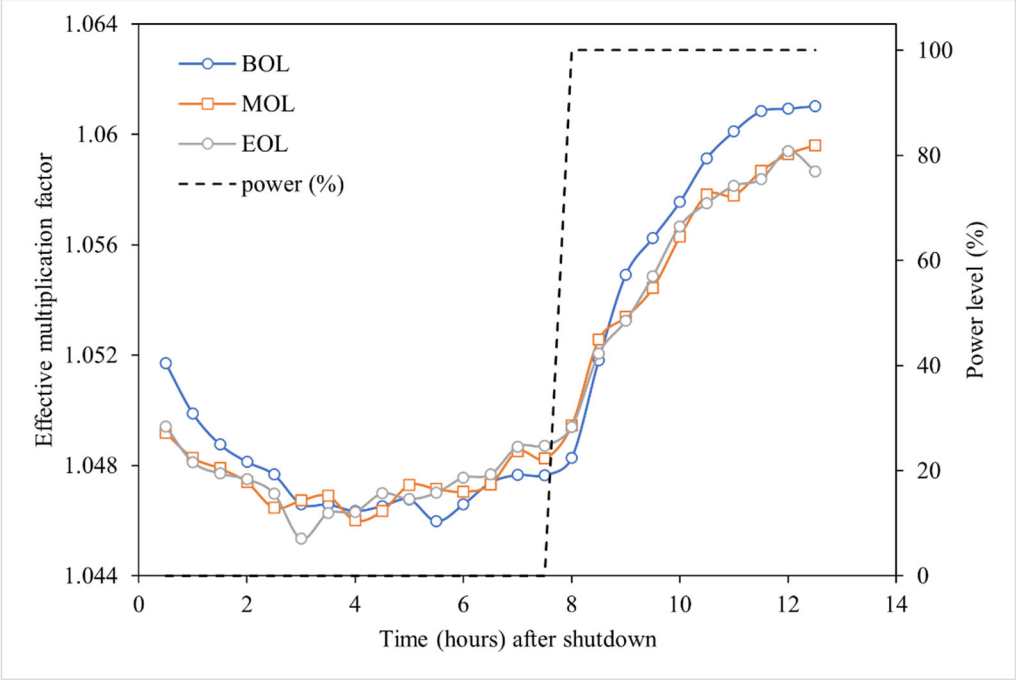

Figure 3.60: Load follow is attempted at BOL (30 days), MOL (15 years) and EOL (30 years) without gas removal system. Uncertainty in k_{eff} is 25 pcm. 30 mins time resolution.

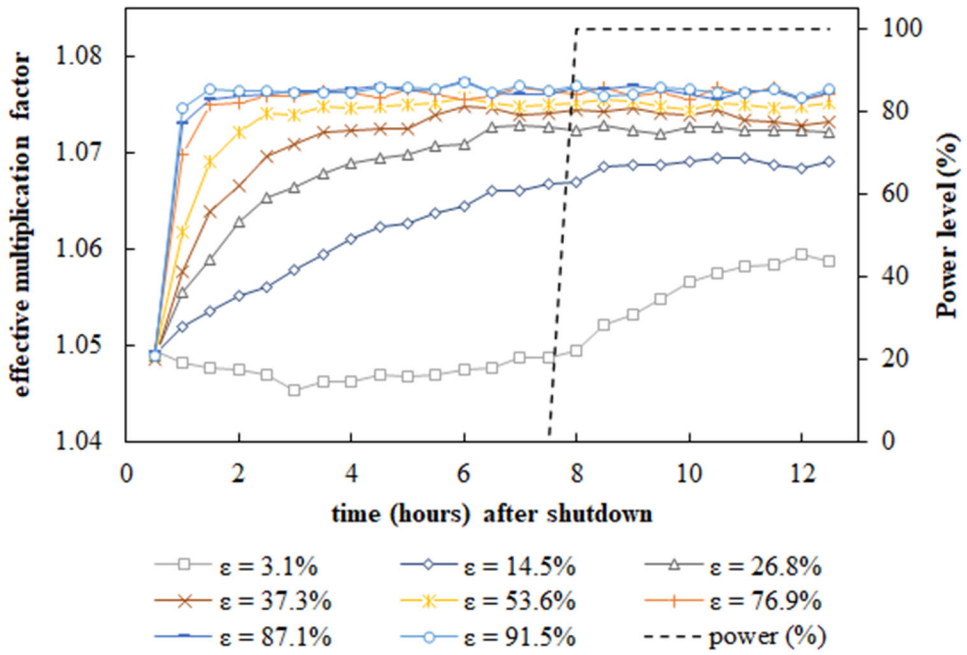

Figure 3.61: After a lifetime of operation at ϵ_{Xe}= 0.536, load follow is attempted at EOL. Above shows k_{eff} during load follow transient for various total Xe removal efficiencies (ϵ_{Xe}) over time after shutdown.

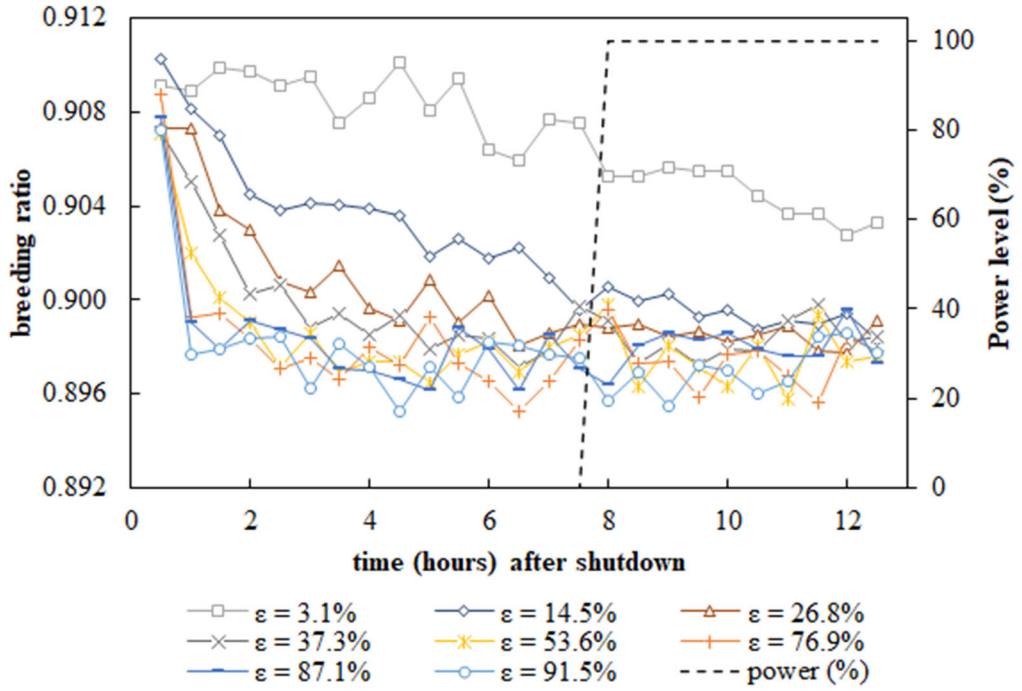

Figure 3.62: After a lifetime of operation at ϵ_{Xe}= 0.536, load follow is attempted at EOL. Above shows breeding ratio during load follow transient for various total Xe removal efficiencies (ϵ_{Xe}) over time after shutdown.

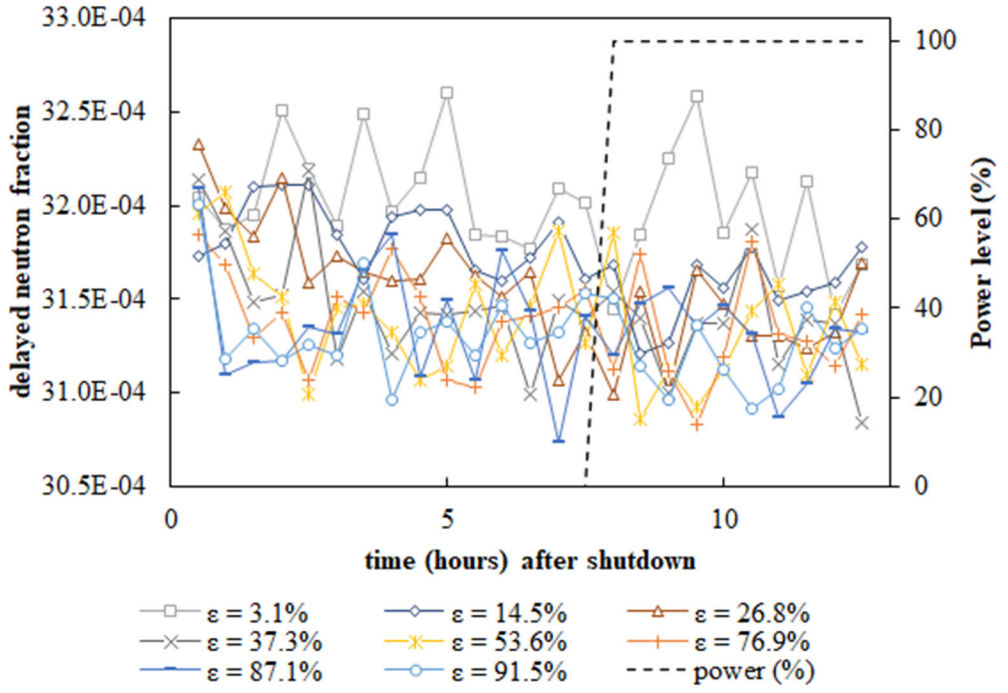

Figure 3.63: After a lifetime of operation at ϵ_{Xe}= 0.536, load follow is attempted at EOL. Above shows β_{eff} during load follow transient for various total Xe removal efficiencies (ϵ_{Xe}) over time after shutdown.

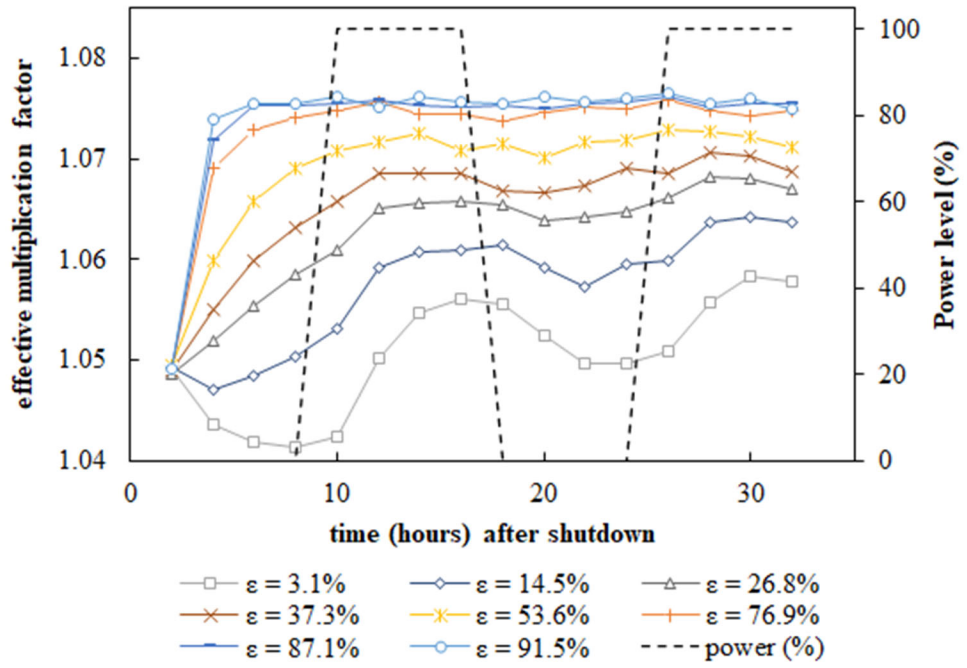

Figure 3.64: After a lifetime of operation at ϵ_{Xe}= 0.536, load follow is attempted at EOL. Above shows k_{eff} during multiple load-follow transients for various total Xe removal efficiencies (ϵ_{Xe}) over time after shutdown.

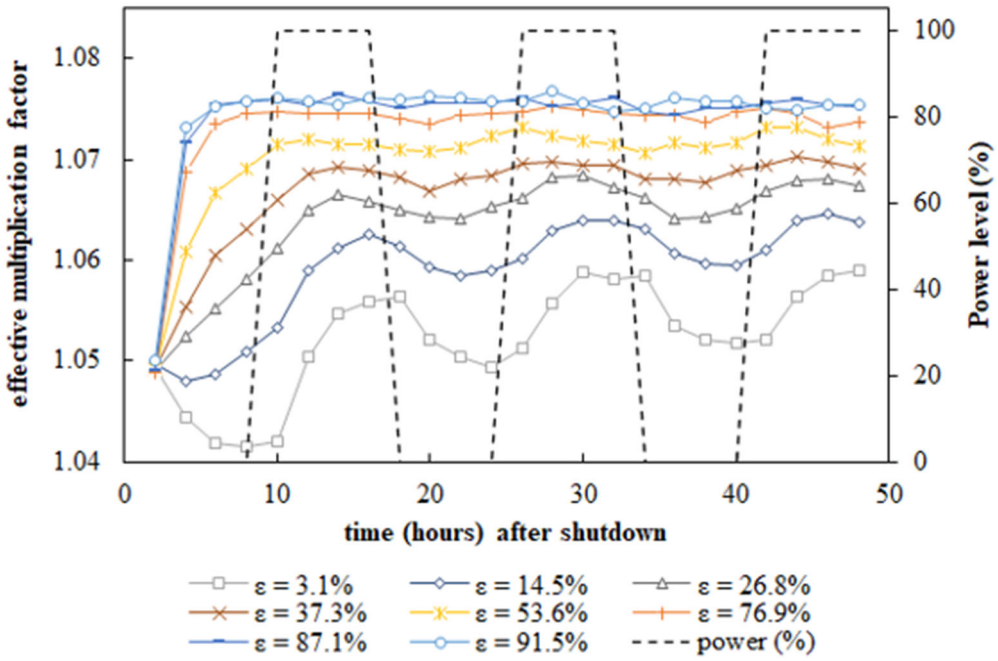

Figure 3.65: After a lifetime of operation at ϵ_{Xe}= 0.536, load follow is attempted at EOL. Above shows k_{eff} during multiple load-follow transients for various total Xe removal efficiencies (ϵ_{Xe}) over time after shutdown.

- Second Scenario

Unlike the previous load-following transients, in this part, short period load-following transients (second scenario) are studied here, and four consecutive power ramps are implemented. As can be seen Figure 3.66, a quick recovery from shutdown is observed even with low gas removal efficiency.

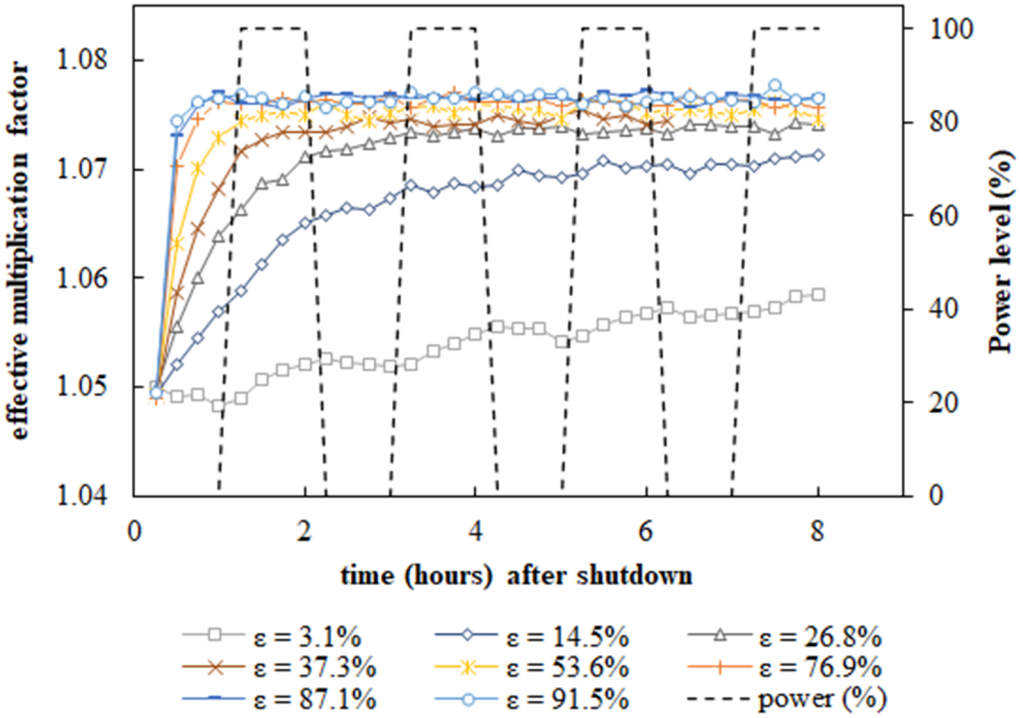

Figure 3.66: After a lifetime of operation at ϵ_{Xe}= 0.536, load follow is attempted at EOL. Above shows k_{eff} during multiple load-following transient for various total Xe removal efficiencies (ϵ_{Xe}) over time after shutdown.

3.5.4. Concluding Remarks

This section demonstrated SaltProc v1.0 capabilities to simulate the short-term depletion with the power variation from 0% to 100% for the MSBR. The methodology from Section 3.4 is applied to investigate the xenon poisoning effect in the MSBR for three various gas removal system regimes: (1) no gas removal, (2) moderate gas removal efficiency, and (3) high gas removal efficiency.

When the gas removal system is inactive, the ^{135}Xe concentration peaked in about 7.5 hours after shutdown, which caused the reactivity drop by 1457 and 1035 pcm for the startup and equilibrium fuel salt composition. Such a negative effect of the xenon poisoning is consistent with other thermal reactor designs (i.e., -1500 pcm for PWR [183]). In contrast with results for the TAP MSR in Section 3.4, the MSBR demonstrated a significant negative impact of the ^{135}Xe concentration spike on the core neutronics after shutdown. The reason for that is significantly greater initial ^{135}I/^{135}Xe concentration ratio: 2.45 and 1.0 for the MSBR and TAP reactor at the

BOL, respectively. Thus, the ^{135}Xe peak is significantly higher for the MSBR than for the TAP reactor: +56% and +0.33%, respectively. Finally, the ^{135}Xe parasitically absorbs substantially more neutrons in the thermal (MSBR) than in the epithermal (TAP MSR) neutron spectrum, which amplifies the xenon poisoning effect when the spectrum softens. In contrast with the spectrum thermalization toward EOL in the TAP reactor, in the MSBR, the neutron spectrum hardens toward EOL due to plutonium and other strong absorbers accumulation in the fuel salt. Thus, for the MSBR, the xenon poisoning effect becomes less severe toward EOL.

The online gas removal in the MSBR demonstrated an impressive positive impact on the core neutronics. The gas removal system operation almost eliminated the effect of xenon poisoning by removing the vast majority of ^{135}Xe during the first hour after the shutdown. During the first 30-minute interval, the reactivity dropped by 161 and 189 pcm for moderate and high removal efficiency, respectively. Afterward, the reactivity raised by 2700 pcm for both efficiencies in a few hours because the ^{135}Xe inventory fell from 14 to 1-2 g. Indeed, the ^{135}Xe loss due to decay and active gas removal significantly overcame its gain from the ^{135}I decay only (no fission happens; thus, no new ^{135}I is produced). Notably, the amplitude of the reactivity swing after shutdown is more significant for the BOL when the xenon reactivity worth is greater due to the softer neutron spectrum. Finally, significantly lower gas removal efficiency (ϵ_{Xe} = 0.536 instead of 0.915) provided comparable benefits to the MSBR core neutronics during the postulated load-following transient. Then, this section demonstrated that the MSBR maintains necessary safety margins throughout the postulated load-following transient. Thus, the temperature coefficient of reactivity and the total control rod worth worsen slightly during the first 30 minutes of the transient when the ^{135}Xe concentration peaked, causing corresponding neutron spectrum hardening. After that, the fast ^{135}Xe concentration decline improved all safety and operational parameters among the cases. Unfortunately, the reactivity worth of two control rods made of boron carbide (B_4C) is insufficient to compensate for huge reactivity change after the shutdown. Even though the total control worth rises throughout the transient, the reactivity system design is unfeasible for load-following and must be redesigned.

Based on the analysis shown in Section 3.5.3, MSBR can, without difficulty, operate under load follow transient with low gas removal efficiency, unless the shutdown or low power operation period is too long, typically greater than 4 hours. Recovery time depends directly on gas removal efficiency and load-following period. In conclusion, the xenon poisoning impact on the core neutronics is much stronger in the MSBR than in the TAP MSR. Thus, the MSBR without gas removal is incapable of flexible restart after reducing power from 100% to 0%. However, online gas removal, even with moderate separation efficiency, helps eliminate the iodine pit problem and enable the load-following capability of the MSBR without compromising its safety. Another benefit from the online gas removal is a stronger thermal feedback.

The work determined that the gas removal system should have a smart control coupled with reactivity control and power regulation systems. Such a system must boost the separation efficiency right before and during the first few minutes after power drop to flatten the ^{135}Xe peak. Then, the control system should reduce the removal efficiency to avoid a sizeable positive reactivity insertion due to a fast ^{135}Xe concentration drop. Finally, a more detailed study of power-changing transients must be performed using SaltProc v1.0 with better time resolution (i.e., a 1-min interval) to understand better how to adjust the gas removal efficiency during power adjustments.

3.6. Conclusions and Future Works

Liquid-fueled nuclear reactors offer several advantages over their traditional solid-fueled counterparts, which makes them a promising option for nuclear fuel cycle closure while offering improved inherent safety. Simulating such systems presents a challenge because existing reactor physics software for fuel burnup historically has been developed for traditional, solid-fueled reactors.

This section demonstrated a flexible, open-source tool, SaltProc, for simulating fuel depletion in a wide range of circulating-fuel (e.g., liquid fuel circulating throughout the primary loop) nuclear reactors that takes into account unique features of such systems: online fuel reprocessing and refueling. SaltProc extends the continuous-energy Monte Carlo burnup calculation code, Serpent 2, for the simulation of material isotopic evolution in any nuclear reactors with circulating, liquid fuel with the main focus on the liquid-fueled MSRs. This work demonstrates a clear contribution to the nuclear engineering community by providing a tool for fuel depletion calculations in any generic nuclear system with circulating fuel.

The need for this work has been shown by a summary of the current state of the art of MSR depletion simulator capabilities. The literature review in Section 3.1 concluded that most MSR depletion simulators typically assume ideal (rather than realistically constrained) poison removal rates for the nuclear system performance modeling. Moreover, most of the simulators assumed constant extraction efficiency vectors, which must be determined by the user in the input file and cannot be a function of other parameters. SaltProc is capable of modeling the peculiarities of MSRs, namely: complex, multi-component reprocessing system structure and realistic extraction efficiency of fission product described as a function of many parameters. Furthermore, SaltProc can maintain reactor criticality by adjusting the reactor core geometry. In addition to fundamental simulation capabilities, SaltProc has a scalable design and allows the development of additional advanced capabilities in the future.

Validation simulations for the TAP MSR have demonstrated the SaltProc capability to model reactors with adjustable moderator configuration. Results for a realistic multi- component model of the fuel salt reprocessing system with assumed ideal removal efficiency are validated with full-core TAP depletion analysis by Betzler et al. [18]. In the realistic reprocessing system with non-ideal removal, the fuel salt composition is strongly influenced by the neutron spectrum hardening due to presence of neutron poisons (e.g., ^{135}Xe) in the core. Thus, more effective noble gas extraction efficiency significantly reduced neutron loss due to parasitic absorption, which led to better fuel utilization and extended core lifetime.

SaltProc is used to perform short-term depletion analysis with power maneuvering in the P ϵ [0, 100%] range to investigate load-following capability in the TAP MSR and MSBR designs. Online gaseous fission product removal significantly improved the load-following capability of the MSBR by reducing the reactivity worth of xenon poisoning from 1457 pcm to 189 pcm. A negligible effect of xenon poisoning is observed in the TAP MSR because its neutron energy spectrum is relatively hard even for the most thermal core configuration (all moderator rods are inserted). Thus, the TAP MSR can effectively load-follow even without continuous gas removal.

Once fuel salt composition evolution was obtained for various MSR designs and power levels, major safety and operational parameters are analyzed at different moments during operation. Specifically, changes in temperature and void coefficients of reactivity and total control rod worth

were evaluated for the TAP concept and MSBR for two timeframes: lifetime-long full-power operation and short-term load-following transient. On a long timescale, the safety parameters worsened during full-time operation for both considered reactor designs due to a significant spectral shift. For the load-following transient, the combination of fuel and moderator temperature coefficient remained strongly negative throughout the transient for both reactors. Notably, the MSBR safety benefited from continuous fission gas removal, while the TAP MSR safety and operational parameters remained stable due to its harder spectrum. Unfortunately, the total control rod worth was insufficient to shut down the MSBR due to a considerable reactivity swing during the load-following transient. Thus, the reactivity control system of the MSBR must be redesigned to ensure safe power maneuvering. Finally, for scientific reproducibility, HDF5 databases generated with SaltProc in this work are published in Illinois Data Bank [189].

For future work, continued research into SaltProc-Serpent and related topics could progress in many different directions. First of all, other liquid-fueled MSR designs with on-site fuel salt reprocessing system should be modeled using SaltProc to improve the cross-code validation portfolio. For example, SaltProc can be validated with a recently published effort for the Chinese Single-fluid Double-zone Thorium Molten Salt Reactor (SD-TMSR) [173].

Next, optimization of reprocessing parameters (e.g., time step, feeding rate, removal rate for various fission product groups) could establish the best fuel utilization, breeding ratio, or safety characteristics for various designs. This might be performed with a parameter sweeping outer loop, which would change an input parameter by a small increment, run the simulation, and analyze output to determine optimal configuration. Alternatively, the existing RAVEN optimization framework [190] might be employed for such optimization studies.

Only the batch-wise online reprocessing approach has been treated in this work. The batch-wise operation used in this chapter (one processing per batch timestep) is valid as long as the timestep is longer than the loop circulation time, V_{core}/Q_{salt}. When this scheme of simulation is used for continuous operation and processing, the batch time should be set to V_{core}/Q_{salt} in principle. However, because the characteristic time of the isotope changes in the core is much longer than the loop circulation time, it is possible to approximate the process with a longer batch timestep, using a modified removal efficiency instead of the system efficiency. In every batch timestep, the salt could go through the processing system $Q_{salt}\Delta t_{batch}/V_{Core}$ times. Assuming an exponential process with constant parameters, the modified efficiency is,

$$\epsilon_{continue} = 1 - (1 - \epsilon_m \epsilon_{es})^{\frac{Q_{salt}\Delta T_{batch}}{V_{core}}}. \quad (3\text{-}24)$$

Moreover, Serpent 2 was recently extended for continuous online fuel reprocessing simulation [130]. This extension could be employed for immediate removal of fission product gases (e.g., Xe, Kr), which have a strong negative impact on core lifetime and breeding efficiency. Thus, using the built-in Serpent 2 Monte Carlo code online reprocessing & refueling material burnup routine would significantly speed up computer-intensive full-core depletion simulations.

Additional physical models for fission product extraction efficiency will enrich the capabilities of SaltProc.

CHAPTER 4: COUPLED REACTOR CORE SIMULATION

4.1. Introduction

In Chapter 3, the load following operations of the TAP MSR and MSBR are evaluated through the overall reactivity with long time steps. Valuable information on the effects of fission product removal is obtained. Having excess reactivity in the reactor during the load following operation is the necessary condition, yet a load-following operation is an inherently transient event, which means that there may be changes in the reactor core that can compromise safety or adversely affect reactor operations but are unobserved in steady operations. Therefore, it is important to study the evolution of the reactor parameters in order to evaluate the safety of the reactor under these load-following conditions. This can be achieved by modelling the transient neutronics and thermal-hydraulics behavior of the reactor core.

In this section, the load-following behavior of the Transatomic Power Molten Salt Reactor (TAP MSR) is studied under different conditions, such as ^{135}Xe accumulation, uniform simple advection flow versus non-uniform flow, and gas entrainment in the reactor core. Since determining the limits of the TAP MSR core is one of the goals of this task, it is worthwhile to first establish a conservative estimate of the reactor safety constraints so that it can be determine that how far the reactor parameters are from these safety constraints during the load-following operations. In the next few sections, the design and dimensions of the modelled TAP MSR core are introduced, and its potential safety constraints are highlighted.

4.1.1. TAP MSR Primary Loop and Reactor Core

Figure 4.1 [14] shows the primary loop of the TAP MSR, which consists of a reactor vessel, a passive safety freeze valve, a primary heat exchanger, a set of fuel salt pumps, and a fuel salt processing system. The reactor vessel is made of Hastelloy and contains the fuel salt, moderator rods, and the control rods. The control rods penetrate the reactor vessel from the top and their degree of insertion can be altered in order to adjust the reactivity of the reactor.

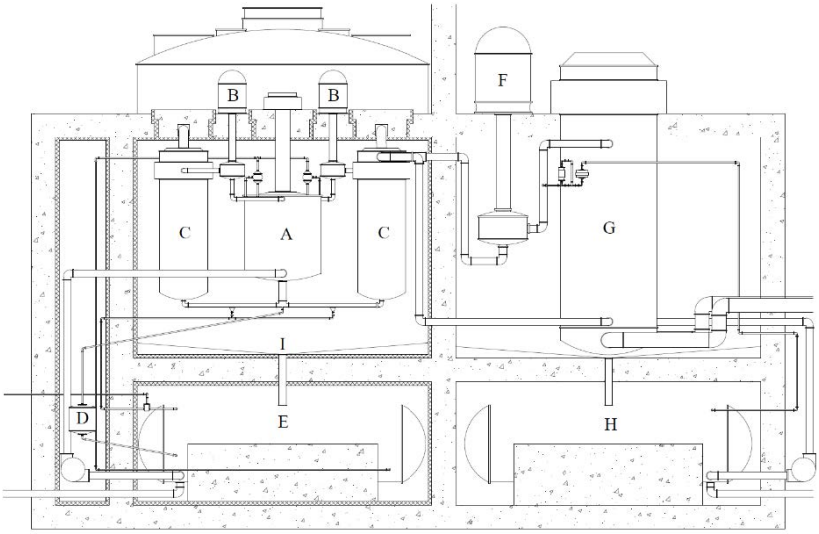

Figure 4.1: Simplified schematic of the TAP MSR primary and intermediate loops, reproduced from the Transatomic Power Technical White Paper V2.1[14]. The figure legend is reproduced as follows: A) Reactor Vessel, B) Fuel Salt Pumps, C) Primary Heat Exchangers, D) Freeze Valve, E) Primary Loop Drain Tank, F) Intermediate Loop Salt Pump, G) Steam Generator, H) Intermediate Loop Drain Tank, I) Fuel Catch Basin.

Molten fuel salt first enters from the bottom of the reactor vessel and undergoes fission as it moves upwards into the active region of the reactor, where geometrical and material conditions allow for a fission reaction to be sustained. As the fuel salt flows through the reactor core, it heats up before exiting from the top. The hot fuel salt then enters the primary heat exchanger where thermal energy is transferred to the cooler intermediate loop, resulting in a decrease in fuel salt temperature. Thereafter, the fuel salt is treated by a processing system where fission products are removed and where the composition of the fuel salt is controlled, before returning to the reactor vessel. The flow of the fuel salt in the primary loop is maintained by the fuel salt pumps, which are used to adjust the flow rate of the fuel salt. There is also a freeze valve, which can drain the fuel salt during an overheating event.

Since the transient analysis in this study is specifically focused on the core of the TAP MSR, the reactor core is considered in isolation from the rest of the primary loop.

4.1.2. Core Geometry and Material Properties

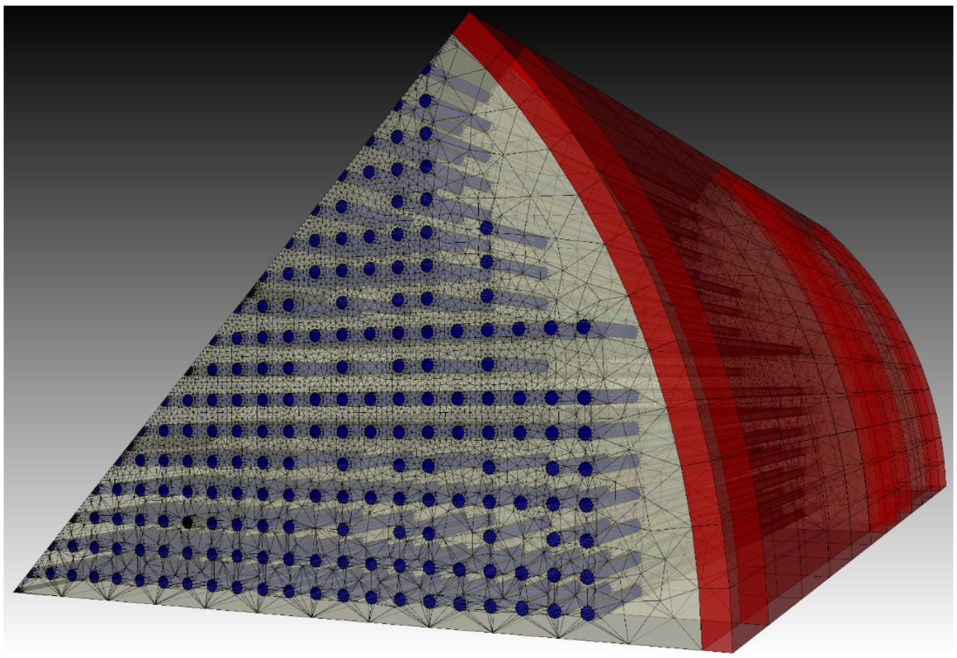

Figure 4.2: Mesh geometry used in this study for the Moltres simulations. Taking advantage of the symmetry present in the design, the constructed mesh geometry was $1/8^{th}$ of the actual reactor core, with reflective boundary conditions at the interfaces. The color representations are as follows: Red – Reactor Vessel Wall, Light Yellow – Fuel Salt, Dark Gray – Control Rods, Blue – Fuel Salt radially co-located with the Moderator Rods.

The mesh geometry of the TAP MSR core used in the core simulations is shown in Figure 4.2. The simulated core consists of a reactor vessel wall modelled as a cylindrical shell that contains the fuel salt, moderator rods, and control rods. In the actual TAP MSR core design, the top and bottom of the reactor vessel are tapered. However, these tapered regions of the core were excluded from the simulation geometry in order to reduce model complexity. Within the reactor vessel, the moderator and control rods are arranged in regular lattice patterns [Figure 4.3]. The moderator rods consist of Zirconium Hydride ($ZrH_{1.66}$) encased in Silicon Carbide while the Control Rods are composed of Boron Carbide (B_4C) and enter the reactor vessel from the top through vertical guide tubes. These guide tubes, as well as other support structures, were not modelled in order to reduce model complexity.

The moderator and control rods arrangements were based on the Oak Ridge National Lab report (ORNL/TM-2017/475) [18]. The control rod material was changed from 70%-Gd_2O_3–30%-Al_2O_3 to B_4C in order to increase control worth, which is consistent with the choice of material made in the fuel cycle simulation discussed in Chapter 3. In addition, the number of moderator rods was increased from 333 to 347 per quarter core, in order to increase the excess reactivity at start-up (BOL) and to reduce moderator rods shuffling frequency.

The relevant material properties are presented in

Table 4.1. Some of the material properties at high temperatures (~900K) were approximated due to a lack of available data.

Table 4.1: Tabulation of the material property values used in the simulations in this study.

	Material	Density ($10^3 kg \cdot m^{-3}$)	Heat Capacity, c_p ($JK^{-1}kg^{-1}$)	Thermal Conductivity, k_{th} ($W \cdot m^{-1}K^{-1}$)
Vessel Wall	Hastelloy	8.86	395.8[14]	20.9 [191]
Fuel Salt[15]	$LiF\text{-}UF_4$ [16]	(6.105-0.001272T) [170]	746.7 [192]	5 [17]
Moderator	$ZrH_{1.66}$	5.66	726.7 [193]	19.8 [18] [193]
Cladding	SiC	3.21	1200 [194]	77.8 [194]
Control Rod	B_4C	2.52	1800 [195]	10 [19] [195]

The reactor dimensions used in this study are presented in

Table 4.2 [18]. The central 300 cm portion of the reactor where the control rods can travel between the parked position (withdrawn) and fully inserted position is the active region of the reactor. Axially above and below this active region are two 25 cm plenum regions, making the total reactor height 350 cm [Figure 4.4].

[14] Estimated using Dulong-Petit law.

[15] The viscosity of the fuel salt used in this work was estimated from Janz [170] to be 0.0209 Pa·s.

[16] 27.5 mol% UF_4, 5% $U\text{-}^{235}$ enrichment.

[17] Thermal conductivity was estimated. T in the density formula refers to temperature in Kelvins.

[18] Thermal conductivity and heat capacity estimated at 900 Kelvins.

[19] Conservative estimates for thermal conductivity and heat capacity.

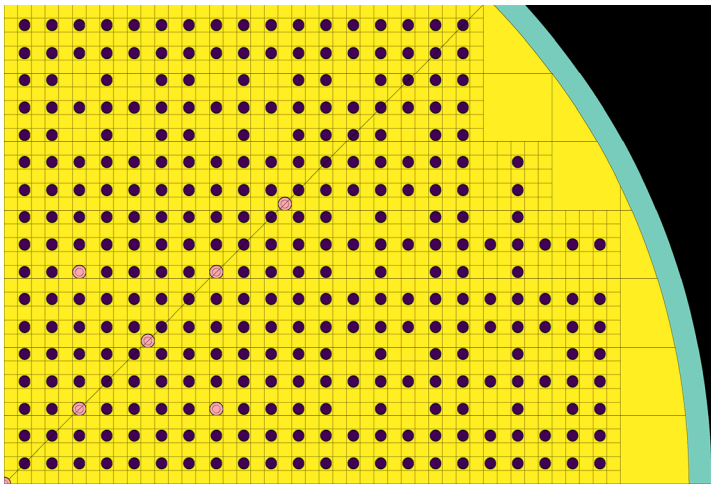

Figure 4.3: Axial cross section of the TAP MSR core illustrating the lattice like pattern of the moderator (violet), control rods (pink) and fuel salt (yellow). This configuration was based on the Oak Ridge National Lab report (ORNL/TM-2017/475) and contains 347 moderator rods per quarter as opposed to the original 333 moderator rods per quarter. The number of moderator rods was increased in order to increase the excess reactivity at the beginning of life (BOL) of the reactor so that reconfiguration of the moderator rods can happen less often and in a more practical operation timeframe.

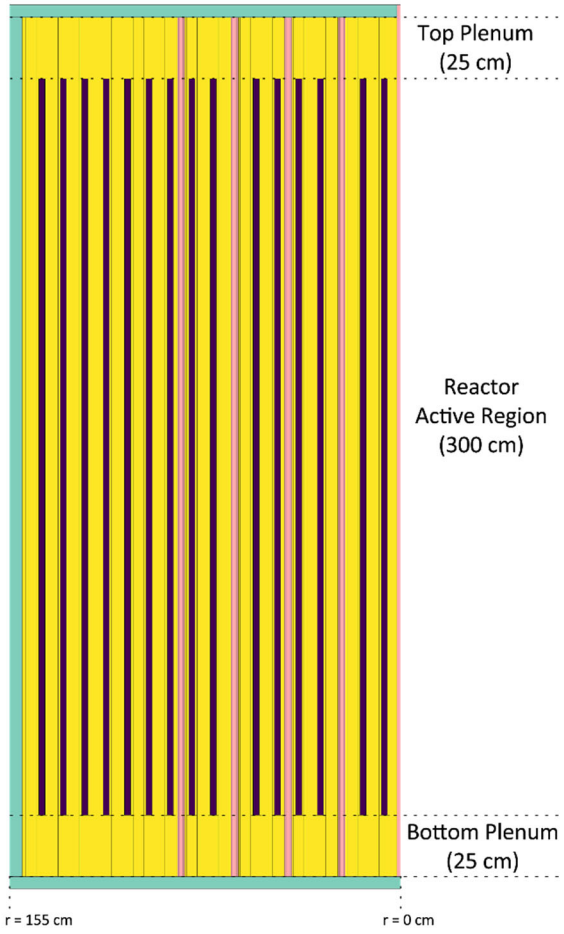

Figure 4.4: 45° RZ cross section of the TAP MSR core denoting the top and bottom plenums as well as the active reactor region. The figure is a view of the TAP MSR core from a 45 degree angle similar to a view from the left side of Figure 4.2 or the diagonal line of Figure 4.3.

Table 4.2: Dimensions of the reactor, referenced and extended from the Oak Ridge National Lab report (ORNL/TM-2017/475) [18].

Component	Parameter	Value

Moderator Rod	Cladding thickness	0.10	cm
	Radius	1.15	cm
	Length	300	cm
	Pitch	3.0 cm	
Control Rod	Radius	1.25	cm
	Length	325	cm
	0% Insertion Position (Parked)	25 cm from top edge 325 cm from top edge	
	100% Insertion Position		
Core	Inner Radius	150	cm
	Vessel Wall Thickness	5	cm
	Plenum Height	25	cm
	Active Reactor Region	300	cm
	Total Reactor Height	350 cm	

4.1.3. Reactor Core Constraints

The overheating of reactor core components is considered to be the most significant constraint in the TAP MSR core. While the fuel salt (LiF-UF4, 27.5 mol% UF_4), moderator ($ZrH_{1.66}$), moderator cladding (SiC) and control rods (B_4C) can tolerate relatively high temperatures, the loss of tensile strength of Hastelloy, which makes up the reactor vessel wall and control rod guide tubes, limits the overall maximum temperature of the reactor core. It was estimated by Yoshioka and Kinoshita [196] that the maximum temperature for Hastelloy should be approximately 900 °C to 1000 °C in order to avoid a significant loss of tensile strength, which may lead to structural deformation or breakage. Therefore, in this work, 900 °C is used as a temperature constraint for the Hastelloy components.

Another temperature limit is the thawing temperature of the salt freeze valves. In the TAP MSR design, the freeze valves located at the bottom of the reactor vessel are designed to thaw at a temperature of around 700 °C. Therefore, a temperature constraint of 700 °C is also imposed at the bottom of the reactor vessel. Nonetheless, such overheating is unlikely to occur unless there is a loss of flow caused by an external accident scenario. Under normal operating conditions, and possibly under natural circulation, the flow would continuously sweep the hot fuel salt upwards from the bottom of the vessel.

Finally, it is noted that the outlet salt temperature of the TAP MSR was designed to be around 650 °C. Considering that the fuel salt melting temperature is 490 °C, the reactor inlet temperature is chosen to be 550 °C, giving a temperature difference of approximately 100 °C when the reactor operates at its rated power level ($1250MW_{th}$). This was achieved by setting the salt flow velocity at 53 cm/s, which was estimated from the energy balance of the fuel salt using the average heat capacity of the salt[20].

[20] In Section 4.3 and beyond, the salt flow velocity was increased to 65 cm/s to adjust for the slight exceeding of target outlet temperature observed in the results of Section 4.2.

4.2. TAP MSR Core Neutronics (Base Case)

The purpose of the base case TAP MSR core neutronics simulation is to determine the bounding conditions of the reactor at the (relatively) high- and low-power states and whether the reactor can rapidly ramp to and from these power states using the fresh fuel salt and fuel salt with equilibrium ^{135}Xe (base cases) through a three-day depletion simulation, similar to what is done in Chapter 3. The work in this section involves the development of tools and methodology required to conduct the simulation and the base cases help to demonstrate that the methodology produces physically reasonable results.

4.2.1. Key Simulation Procedures and Methodology

The neutronics simulation of the TAP MSR core can be summarized in the following key steps:

1) First, the neutronics cross sections are calculated for each material present in the reactor core (fuel salt, moderator, moderator cladding, control rods, and Hastelloy) using Serpent 2 [124].
2) Then, the required reactor geometries and meshes are generated for the MOOSE-based code, Moltres [197-199], using the mesh generation software, Trelis. These geometries were identical to the ones used in the Serpent 2 simulations, with the exception of a change in control rod positions.
3) At the same time, the Moltres input files are generated, which described the reactor conditions to be simulated in Moltres, such as the inlet temperature, flow rate, and material properties. Both the Moltres input file(s) and the Trelis input file(s) were created with the help of a Python script which ensured that the Moltres and Trelis input files were consistent.
4) Finally, Moltres is used to solve for the temperature and neutron flux profiles of the reactor core, based on the parameters provided in the neutronics cross section files, mesh files, and the Moltres input files.

In the case of transient simulations, such as the insertion or withdrawal of control rods, step 2 to 4 are repeated with small step changes in the control rod positions within the mesh.

4.2.1.1. Neutronics Cross Section Generation Using Serpent 2

To generate the required neutronics cross sections for the Moltres code, the geometry of the TAP MSR core, along with the material comprising the volumes in the geometry, was first described within a Serpent 2 input file. In addition to the geometry of the reactor core, the Serpent 2 input file also contained information regarding the isotopic composition of each reactor core materials, which were categorized into separate "universes". Then the Monte Carlo simulation is performed with a neutron population of 10,000 for 1000 criticality cycles and 200 inactive cycles at the rated power of 1250 MW_{th}, using the Joint Evaluated Fission and Fusion Nuclear Data Library (JEFF-3.1.2) [200]. The output of the Monte Carlo simulation was then formatted into the cross-section files required by Moltres. These Moltres cross section files described the average macroscopic cross section for each material present in the TAP MSR core for a particular reactor

configuration at the expected temperature range. It should be noted that the Serpent 2 simulation only need to be run once for each unique reactor configuration, such as a particular moderator rod configuration or fuel salt composition.

4.2.1.2. Moving-Mesh Technique for Control Rod Movement

The mesh files used by Moltres were generated using Trelis and were stored in the *Exodus-II* format. They described the reactor core geometry and the details of the mesh such as the volumes, surfaces and positions contained in each *Block* and *Surface* entity, which would be used by Moltres to identify the elements and nodes to assign material properties and boundary conditions.

On the other hand, the Moltres input files defined the type of problem to be solved. They contained information such as the fuel salt flow rate, core inlet temperature, materials present (along with their properties), the neutron and temperature diffusion kernels (along with the neutron multiplication kernels), the simulation time, boundary conditions, as well as the reference to the respective neutronics cross section files discussed earlier.

While Moltres can easily solve for the time-dependent evolution of a reactor using a static (constant) mesh, there was no straightforward way to simulate the movement of reactor components without changing the mesh geometry. Therefore, to simulate the movement of reactor components, such as the insertion or withdrawal of control rods, an approach developed in this thesis was to generate multiple mesh files (and their corresponding Moltres input files) with minute changes in the reactor geometry at small timesteps, simulate the reactor behavior for each pair of Moltres and mesh input files, and finally restarting the simulation for the next pair of input files in a chronological order.

As it would be difficult to keep track of the Moltres and mesh input file pairs, the Moltres and Trelis (mesh generation) input files were generated using a single Python script that was capable of interpolating between user defined "key frames". Each key frame described a particular reactor geometry (moderator and control rod positions and degree of insertion), fuel salt flow rate, core inlet temperature, and simulation time. Defining two or more key frames and the number of timesteps between each pair of key frames allows the script to interpolate between the parameters and generate the corresponding Trelis and Moltres input files for each timestep.

For example, defining two key frames that are 100 timesteps apart, with control rods fully withdrawn and inserted, would generate 101 Trelis input files where the control rods are inserted at 1% increment. The Trelis input files would then be used to generate the Exodus-II meshes to be used in Moltres. Simultaneously, a set of Moltres input files would be generated where their parameters (salt flow rate, core inlet temperature and simulation time) were interpolated accordingly.

The main impetus for generating new mesh file for each timestep, instead of an instantaneous reactor geometry change, was that the timescales for the reactor transients were at the same order of magnitude as the rate of change in reactor configuration (i.e., the movement of control rods). Therefore, it is not appropriate to simulate an instantaneous change between the initial and final reactor configurations without introducing unacceptable amounts of inaccuracy in the load-following transient simulation.

4.2.1.3. Key Physics in Moltres Simulation

Moltres is a MOOSE-based Finite Element (FE) Partial Differential Equation (PDE) solver that is built specifically to simulate a neutron diffusion (and multiplication) system with advective fuel, e.g., a Molten Salt Reactor (MSR). While the use case is typically focused on MSRs, Moltres can be used for the simulation of many other advanced reactor systems owing to its modular design as a MOOSE-based application.

In the Moltres simulation of the TAP MSR performed in this work, several key physics govern the behavior of the reactor. They are:

1) Neutron diffusion in all materials and neutron multiplication in the fuel salt.
2) Delayed Neutron Precursor (DNP) generation in the fuel salt and subsequent decay, which produces delayed neutrons.
3) Temperature diffusion in all materials.
4) Convection of thermal energy in the flowing fuel salt, which is modelled as an advection of temperature.
5) Convection of DNP in the flowing fuel salt, which is modelled as an advection of DNP concentration.

Expressing the above physics in mathematical form, the following equations are obtained, which are reproduced from [198]:

1) Neutron diffusion and multiplication

$$\frac{1}{v_g}\frac{\partial \phi_g}{\partial t} - \nabla \cdot \mathrm{D}_g \nabla \phi_g + \Sigma_g^r \phi_g = \sum_{g \neq g'}^{G} \Sigma_{g' \to g}^S \phi_{g'} + \chi_g^p \sum_{g'=1}^{G} (1-\beta) \nu \Sigma_{g'}^f \phi_{g'} + \chi_g^d \sum_{i}^{I} \lambda_i C_i \,. \quad (4\text{-}1)$$

From left to right, the terms represent the time rate change of neutron flux, the spatial diffusion of neutron flux, the removal of neutron flux via absorption and out-scattering, the neutron flux gained from in-scattering, neutron multiplication (fission) and finally the neutron flux gained from delayed neutrons, respectively.

In this work, only 2 groups of neutron fluxes considered: fast and thermal, with the *g-indices* being 1 and 2. At the same time, 8 DNP groups are included with *i-indices* being 1 through 8. One simplifying assumption in the Moltres simulation is that the delayed neutrons are only generated as fast neutrons (i.e., group 1). It should also be noted that the fission term only applies to the fuel salt region.

2) DNP generation, decay, and advection

$$\frac{\partial C_i}{\partial t} = \sum_{g'=1}^{G} \beta_i \nu \Sigma_{g'}^f \phi_{g'} - \lambda_i C_i - \frac{\partial}{\partial z} u C_i. \quad (4\text{-}2)$$

In this equation, the terms represent the time rate change of DNP concentration, the DNP generation from fission, the loss of DNP concentration due to decay, and the advection of DNP concentration due to the flow of the fuel salt. The advection term represents the uniform upwards movement of the DNP concentration due to the upwards flow of the fuel salt. However, the last

term can be changed to one where the DNP concentration follows a velocity field instead of a uniform upwards flow ($\nabla \cdot \vec{u}C_i$).

3) Temperature diffusion and advection

$$\rho_f c_{p,f} \frac{\partial T_f}{\partial t} + \nabla \cdot \left(\rho_f c_{p,f} \vec{u} \cdot T_f - k_f \nabla T_f\right) = Q_f. \quad (4\text{-}3)$$

In this equation, the terms represent the per unit volume time rate change of thermal energy (or temperature if dividing the whole equation by $\rho_f c_{p,f}$), the convection of thermal energy, the spatial diffusion of thermal energy, and the volumetric heat generation.

Similar to the DNP equation, the convection term is reduced to advection if the salt flow is a uniform upward flow. This convection/advection term would only apply to the fuel salt region as only the fuel salt is flowing in the reactor. Similarly, it is common for the volumetric heat generation term to apply only to the core region. Nonetheless, Moltres allows for the case where the energy generated is further distributed into gamma and neutron irradiation heating in the other materials. In this work, the simpler former approach is selected in order to reduce computational cost.

4.2.1.4. Calculation of Neutron Multiplication using Moltres

Besides solving for the neutron flux and temperature profiles of the reactor, Moltres is configured to solve for the approximate neutron multiplication factor (k_{eff}) of the reactor at each timestep via an eigenvalue problem approach. Moltres does so by first taking a simulation output and extracting the temperature profile of the particular timestep(s) of interest in order to determine the macroscopic cross sections of each element in the mesh [Figure 4.5]. With the cross sections of each element known, Moltres then solves the neutron diffusion equation for the largest k-eigenvalue which is the neutron multiplication factor (k_{eff}) of the reactor.

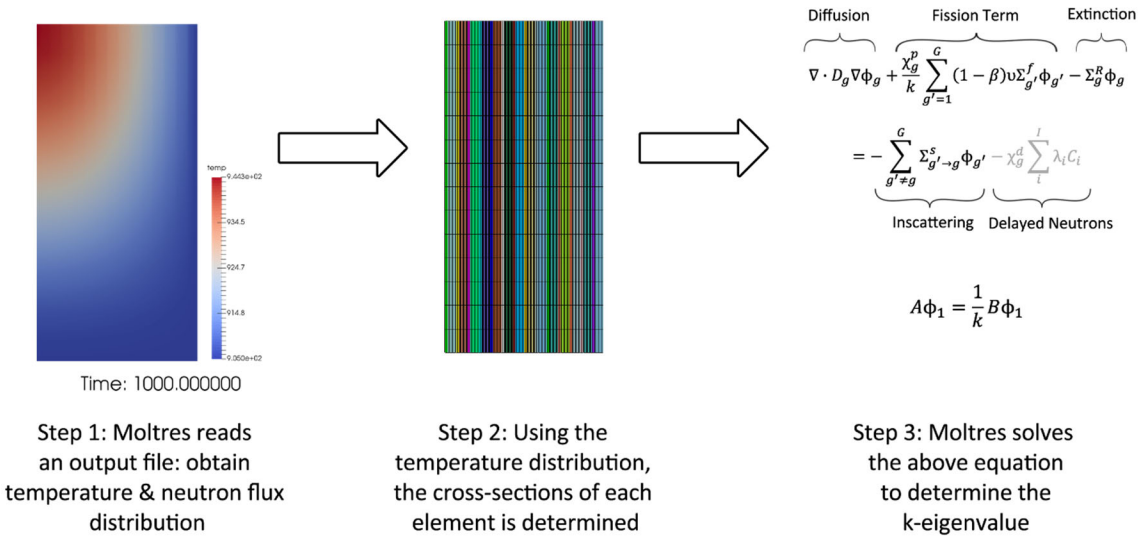

Figure 4.5: Schematic of the method used to determine the k-eigenvalue of the reactor at a particular state. The example shown was from an MSRE simulation which was used to test the Moltres k-eigenvalue method.

During steady state operations, this k-eigenvalue is expected to be a constant that is close to 1. However, during reactor transients, such as the movement of control rods, the k-eigenvalue would change over each timestep. The change of k-eigenvalue over time reflects how the reactivity of the reactor core behaves in response to a transient operation.

Additionally, by solving for the k-eigenvalue a second time using a mesh geometry with control rods fully inserted, the difference between the two k-eigenvalues can be used to provide an indication of the shutdown margin of the reactor, thereby giving us another way to evaluate the neutronics safety of the load-following operations.[21]

4.2.1.5. Simulation Approach for Load-Following and Effects of ^{135}Xe

The Moltres simulations were ran on both the Idaho National Lab's (INL) *Falcon 1* as well as the National Center for Supercomputing Applications' (NCSA) *Blue Waters* computer clusters due to the appreciable computing resources required to run the simulations.

The general strategy for determining the operating conditions at the low-power and high-power (~1250 MW_{th}) states involved performing an initial simulation with fresh fuel ("Fresh" case) to obtain a steady state solution, before restarting from this solution with perturbed control rod positions to obtain the desired reactor power. From here, the operational envelope for load-following was determined by performing transient power ramp simulations between the two reactor states. These procedures were repeated for a fuel salt containing ^{135}Xe at equilibrium levels ("Xenon" case).

[21] This shutdown margin study was ultimately not performed in this work after we had changed the control rod material to B_4C. The new control rod material was found to be performing 'too well' as a neutron absorber such that shutdown margin was not a concern for the fuel compositions used in this work. Nonetheless, the tools and methodology are in place for such a study and can be used evaluate other fuel compositions or moderator rod configurations.

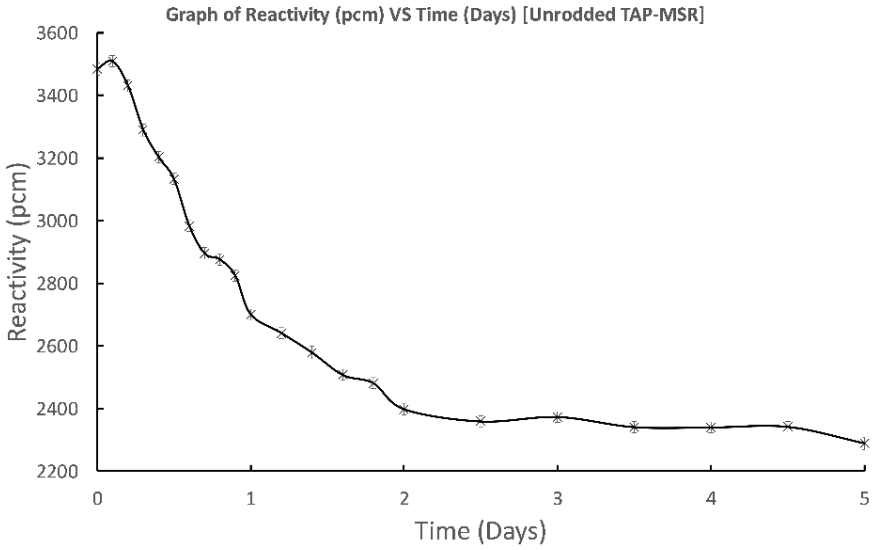

Figure 4.6: Reactivity dynamics of the TAP MSR at a uniform temperature of 900 K with all control rods withdrawn and with the fuel salt composition at the beginning of life (BOL). The reactivity was calculated in Serpent 2 to determine the time when fission product concentrations, including ^{135}Xe, would reach quasi-equilibrium. The reactivity reached a steady state after around 2 days but the burnup time for the extraction of fuel salt composition was chosen to be 3 days to introduce a suitable margin. From the graph, ^{135}Xe and other fission products introduced a reactivity decrease of around 1100 pcm compared with fresh fuel salt at BOL.

The composition of the fuel salt is obtained with equilibrium ^{135}Xe concentration by performing a burnup calculation in Serpent 2 until the reactivity of the TAP MSR stabilizes. The stabilization of the reactivity signifies that the buildup of neutronically important fission products, including ^{135}Xe, has reached quasi-equilibrium concentrations. From Figure 4.6, this burnup period corresponds to about 3 days. Then the results from the Fresh and Xenon cases are compared in order to determine the effects of ^{135}Xe and other fission products on the load-following behavior.

4.2.2. Results and Discussion

4.2.2.1. Reactor Boundary Conditions

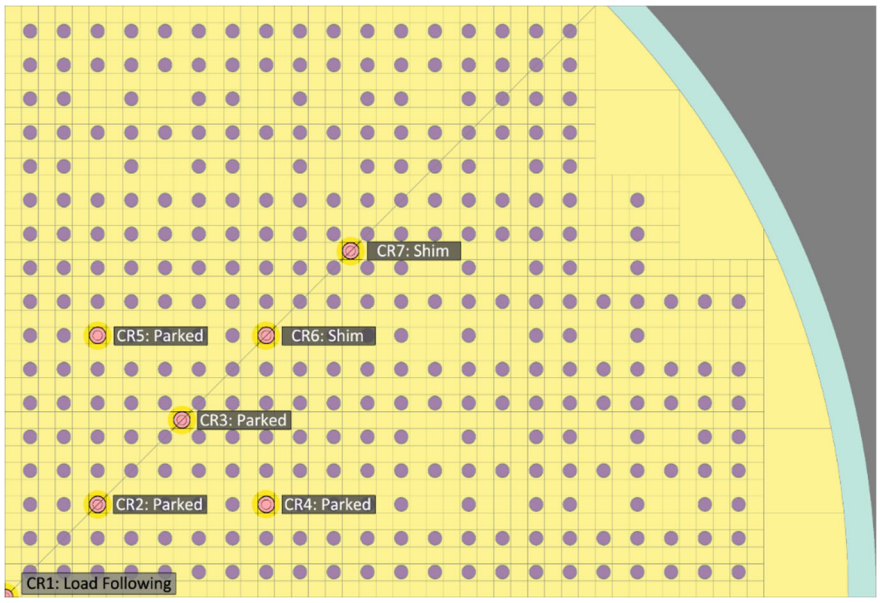

Figure 4.7: Control rod configuration for the TAP MSR simulation with fresh fuel salt. The central rod, CR1, was used for power adjustments (load-following), CR2 to CR5 were parked within the top plenum while CR6 and CR7 were 100% inserted and acted as shim rods. For the Xenon case, CR6 was instead parked in the top plenum to counteract the lower reactivity of the fuel salt.

The control rod configuration for the Fresh case is shown in Figure 4.7. In this configuration, the control rods CR2 to CR5 were parked at 0% insertion (i.e., fully withdrawn, residing only in the upper plenum region), the two outermost control rods, CR6 and CR7, were fully inserted and acted as shims (100% insertion), while the central control rod, CR1, was used for the adjustment of reactor power (i.e., load-following). For the Xenon case, CR6 was parked in the upper plenum instead of acting as a shim rod in order to counteract the lower reactivity (-1100 pcm) due to the presence of fission products including ^{135}Xe.

The reason for parking some control rods at 0% insertion was due to a preliminary simulation finding that the B_4C control rods were too strongly absorbing to enable practicable reactor control if all control rods were moved simultaneously. Therefore, the control rod deployment strategy shown in Figure 4.7 is chosen, where only the central rod was moved during power maneuvers while the outer control rods acted as shims (100% insertion) to remove the excess reactivity and to allow the full range of the reactor power to be controlled using only the central control rod.

From these configurations, the bounding reactor parameters are obtained for the Fresh and Xenon cases and the resulting parameters at the high- and low-power states are summarized in Table 4.3. At the same time, the temperature and neutron flux profiles are presented in Figure 4.8.

Table 4.3: Summary of reactor parameters at the high- and low-power states of the Fresh and Xenon cases. In the table, the power level is given with respect to the rated power of the TAP MSR (1250 MW_{th}).

Case	Power Level (%)	Central Rod Insertion (%)	Average Outlet Temperature (°C)	Maximum Local Temperature (°C)
Fresh (Low)	0.68	84	550.8	-
Fresh (High)	78.77	60	639.9	794.54
Xenon (Low)	0.67	75	550.8	-
Xenon (High)	94.53	55	663.4	856.9

It should be noted that the initial high-power states in both cases were terminated at a quasi-steady state due to the considerable amount of computing resources required[22]. Therefore, the high-power states shown were only an approximation of the true steady state. The high-power state for the Fresh case was only producing about 79% of the rated power but it was decided not to further fine tune the CR1 position due to the computing resources required.

Comparing the Fresh and Xenon cases, the temperature and neutron flux profiles were largely similar for both cases. However, the temperature and neutron flux profiles in the Xenon case had a larger spread in the radial direction due to the removal of CR6, which acted as a shim rod in the Fresh case.

[22] The values for the high-power states in Table 4.3 were updated with the results from the power up ramps after it was observed that the end of the power up ramp simulations produced a better steady state.

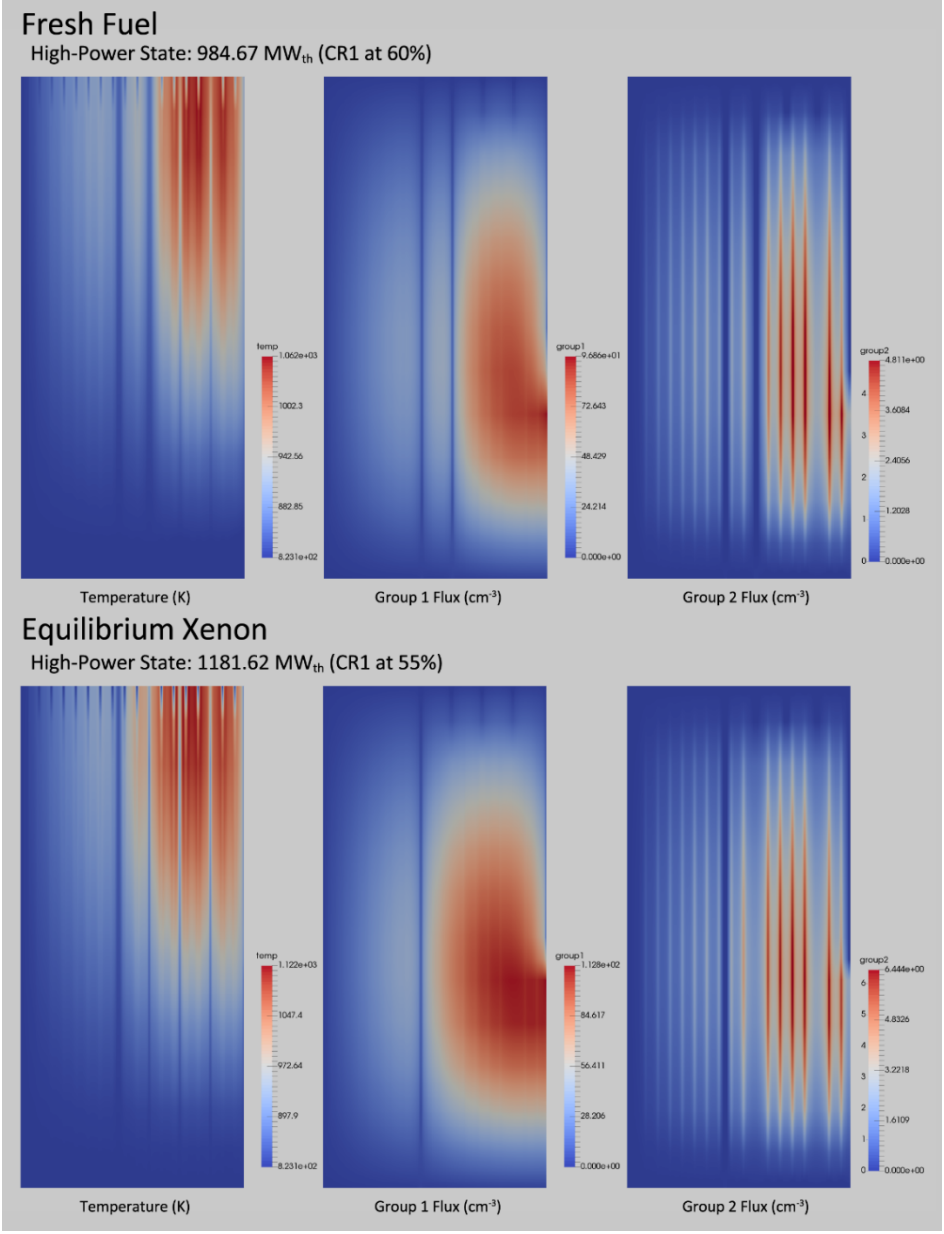

Figure 4.8: Temperature and neutron flux distrubutions of the TAP MSR at the high-power state using fresh fuel salt (top) and equilibrium Xenon (bottom). This view was in the 45° RZ cross section, similar to that in Figure 4.4. The slight inhomogeneity in the temperature of the high-power states was due to the system not reaching a true steady state, owing to the considerable amount of computing resources required to proceed. The units of the scales used were in 10^{13} #/cm^3 for the neutron fluxes and Kelvins for the temperature.

For the Xenon case, the average outlet temperature was in excess of 650 °C at 94.53% of the rated power. This suggests that the heat capacity of the fuel salt had been overestimated and that the flow rate should be increased in order to lower the outlet temperature while maintaining the rated power. This was corrected in the later sections of this work.

4.2.2.2. Reactor Power Ramps

With the approximate high-power and low-power states known, the power ramp simulations were then performed between these two states using a control rod (CR1) movement rate of 50 cm/s and 143 and 119 intermediate mesh files for the Fresh and Xenon cases, respectively (i.e., a CR1 movement of 0.5 cm in 0.01 seconds between each intermediate mesh file). The reactor power, reactivity, and maximum temperature during the transient operations are shown in Figure 4.9.

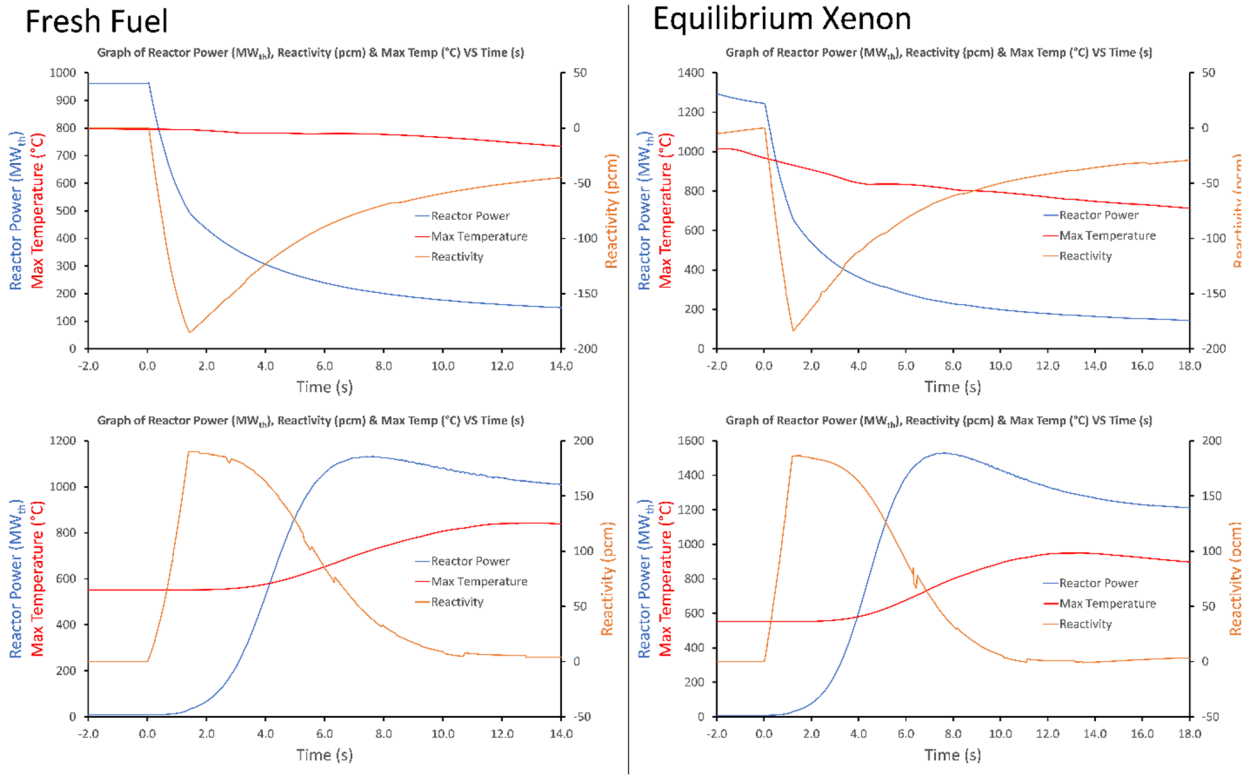

Figure 4.9: Reactor power, reactivity and maximum temperature during power down ramps (top) and power up ramps (bottom) for the Fresh (left) and Xenon (right) cases. The movement of control rods started at t = 0.0 s and completed at approximately t = 1.5 s. For power down ramps, the reactor power decreased asymptotically after the completion of control rod movement while for power up ramps, the reactor power briefly overshot the target level before stabilizing.

During power down ramps, the reactivity of the TAP MSR reached -185 pcm and -183 pcm, for the Fresh and Xenon cases respectively, before recovering towards criticality as the hot fuel salt was replaced by cooler and higher density fuel salt. The maximum power change rates were about -45% and -51% of the rated reactor power per second at the beginning of the down ramp and the time to reach 90% of the target power levels were around 35 s and 23 s for the Fresh and Xenon cases, respectively. This implies the power down ramp rates of approximately -120% (Fresh) and -220% (Xenon) of the rated reactor power per minute.

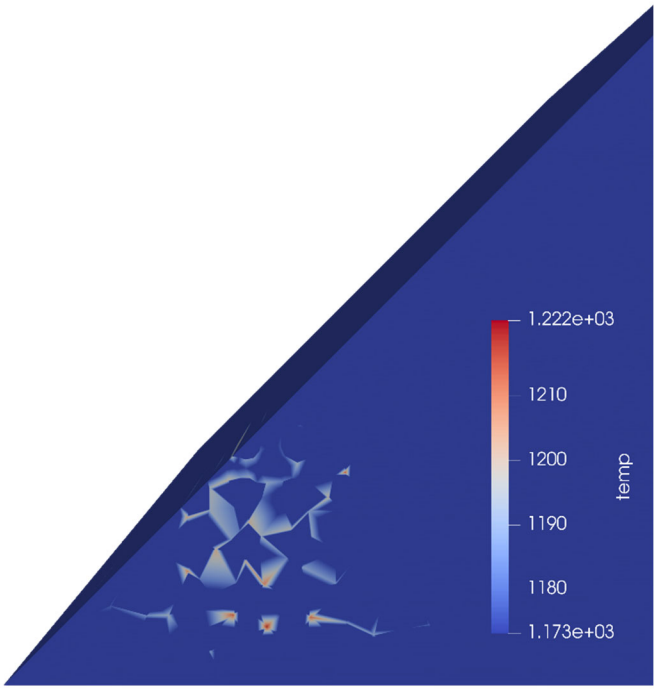

Figure 4.10: A top down view of the temperature profile of the TAP MSR during the timestep when the maximum temperature occurred. The regions in dark blue were either at or below 900 °C while the regions in light blue to red were nominally higher than 900 °C, up to a maximum of 949.6 °C. The shape and irregularity of the non-conforming regions suggest that the temperature values in these regions were artefacts of the coarse mesh used and the components in the TAP MSR might not necessarily have exceeded the temperature constraints.

Meanwhile, during power up ramps, the reactivity of the TAP MSR reached +190 pcm (Fresh) and +186 pcm (Xenon) by the completion of the control rod movement. The maximum power change rate was about +30% (Fresh) and +40% (Xenon) of the rated reactor power per second, and the time to reach the steady target reactor power levels was about 25 s for both cases. This implies a power up ramp rate of approximately +185% (Fresh) and +225% (Xenon) of the rated reactor power per minute.

During power up ramps, the reactor power briefly overshot the target level by around +18% (Fresh) and +30% (Xenon) before decreasing to the desired steady state power level. Such overshoot was expected and even required in order to increase the average fuel salt temperature to that in the high-power state. However, because of this power overshoot, the maximum local temperature should be monitored to determine if the temperature limit has been exceeded. It was observed that the maximum local temperature was 842.6 °C for the Fresh case but 949.6 °C for the Xenon case, which was higher than the temperature constraint for Hastelloy components (900 °C) as discussed in Section 4.1.3.

While this could be due to the fact that the power ramp in the Xenon case was closer to 100% of the rated reactor power, an examination of the temperature profile [Figure 4.10] suggests that the high temperature was a numerical artefact due to the coarseness of the mesh. Therefore, the

components in the simulated TAP MSR might not necessarily have exceeded the temperature constraints.

In the following sections, the control rod configuration and salt flow rate were refined which eliminated this overheating numerical issue in the simple advection scheme.

4.3. Thermal Hydraulics Core Simulation

The purpose of the thermal-hydraulics core simulation was to improve the fidelity of the fuel salt flow from a simple advection model in Section 4.2 to an Incompressible Navier-Stokes (INS) model. The work in this section and Section 4.4 helped to determine the differences in the salt flow profiles between the two methods and their effects on the reactor temperature as well as whether the temperature constraints could be met.

4.3.1. Key Simulation Procedures and Methodology

The thermal-hydraulics simulation of the TAP MSR core was performed using the following key steps:

1) First, a steady state simulation of the fuel salt flow is performed using the Incompressible Navier-Stokes (INS) physics submodule in MOOSE. Because the control rods in the TAP MSR were housed in guide-tubes, the fuel salt flow was largely isolated from the varying control rods positions. This meant that one single salt flow simulation could be reused for different control rods configurations, assuming that the differences in neutronics behavior do not significantly affect the salt flow.
2) Then, the neutronics of the TAP MSR is simulated with the uniform advection scheme using Moltres, similar to the steps taken in Section 4.2. The neutronics simulations determined the high- and low-power states and the transient ramp behavior between these states, using fresh fuel salt and fuel salt with equilibrium ^{135}Xe.

The neutronics solution, which determined the heat generation profile, was then combined with the steady state fuel salt flow (velocity field) solution to obtain a new *uncoupled* thermal-hydraulics solution.

4.3.1.1. Incompressible Navier-Stokes Equation

To simulate the Incompressible Navier-Stokes flow of the fuel salt, the *Navier-Stokes* physics module [201] from MOOSE is used for the simulation kernels, with PSPG and SUPG stabilizations[23] enabled and with a stabilization factor of 1/3 for the PSPG and 1/2 for the SUPG stabilization schemes. The body force physics (gravity) was not included as the INS simulation was performed using uniform fuel salt temperature (i.e., uniform fuel salt density); as the reactor was vertically oriented, the inclusion of gravity would only increase axial pressure drop without changing the local flow solution.

The incompressible N-S equations before stabilization are:

[23] PSPG: Pressure-Stabilized Petrov-Galerkin; SUPG: Streamline-Upwind Petrov-Galerkin

1) Mass balance (incompressible flow):

$$\nabla \cdot \vec{u} = 0. \tag{4-4}$$

This term represents the incompressible flow, which is the result of setting the material derivative of density to be zero in the Navier-Stokes mass balance equation.

2) Momentum balance:

$$\rho \frac{d\vec{u}}{\mathrm{dt}} + \rho(\vec{u} \cdot \nabla \vec{u}) = -\nabla P + \mu \nabla^2 \vec{u} + \rho \vec{g}. \tag{4-5}$$

From left to right, the terms represent the per unit volume time rate change of momentum, the convection of momentum, the pressure force, viscous stress force, as well as the gravitational body force. In this section, the gravitational acceleration was set to zero which resulted in zero contribution from the last term. However, in Section 4.4, the gravitational acceleration was set to -9.81 ms^{-2} $\hat{z}$ with the inclusion of a $\rho \alpha \vec{g}\left(T - T_{ref}\right)$ term in order to capture the Boussinesq body force.

The reactor geometry used in the simulation of the INS salt flow was identical to that used in the neutronics simulation (simple advection). However, a slight change is made in the INS simulation mesh so that the fuel salt could flow freely within the top and bottom plenums, which is shown in Figure 4.11. For the boundary conditions, all internal surfaces were given the no-slip boundary conditions (zero velocities at the surfaces) while the reactor inlet and outlet were given an axial velocity (59.4 cm/s) which was adjusted to match the simple advection flow velocity of 65 cm/s in the central reactor region[24].

[24] The INS inlet and outlet boundary condition of 59.4 cm/s takes into account the absence of moderator rods in the plenums and the effect of no-slip boundary condition due to the vessel wall so that the salt flow rate can match that in the simple advection flow.

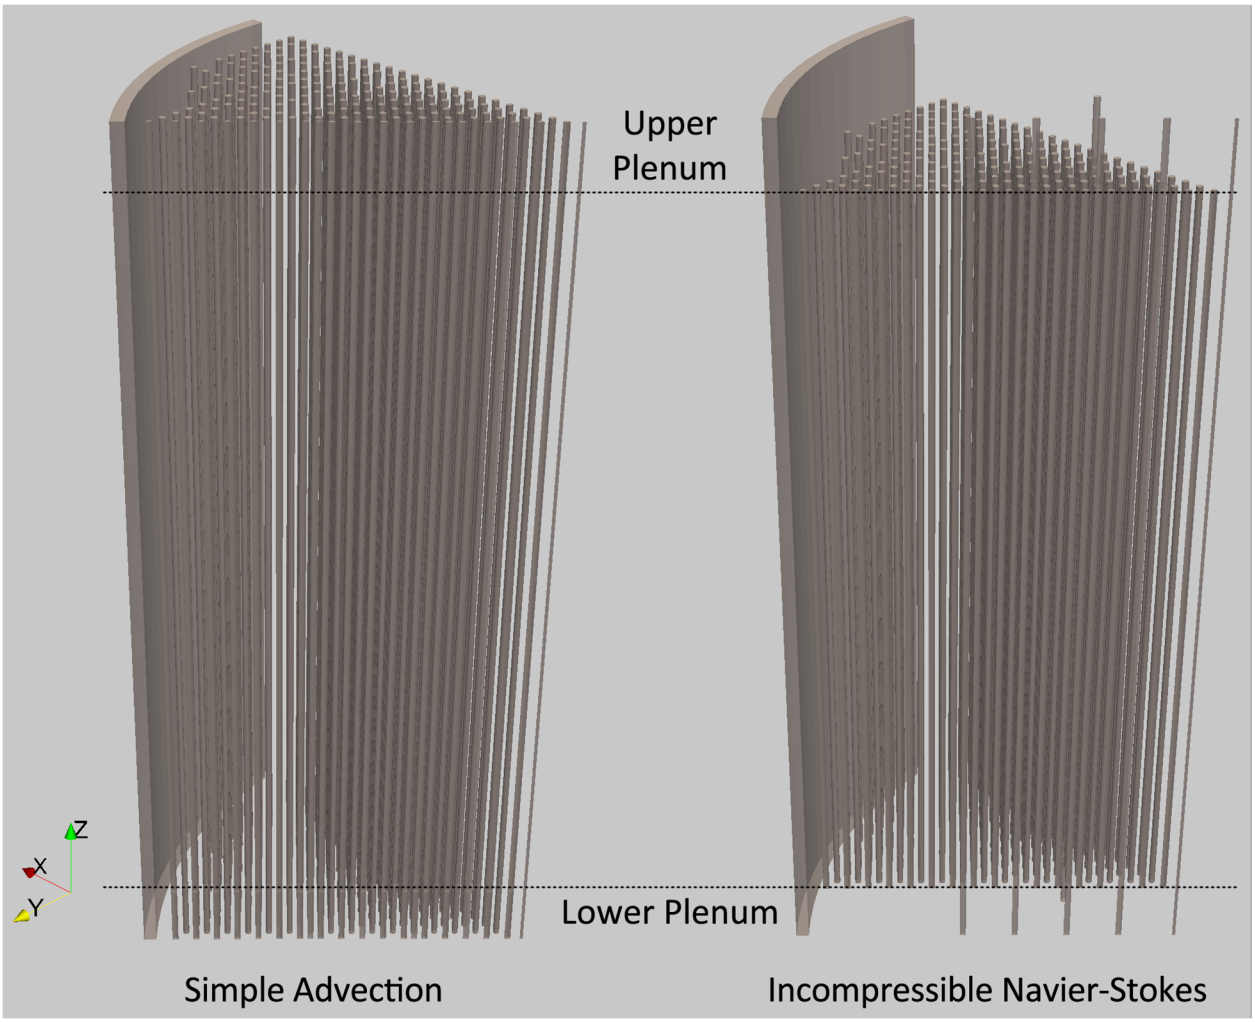

Figure 4.11: Mesh geometry used in the Simple Advection simulation and the Incompressible Navier-Stokes simulation. The figure shows the 'walls' of the reactor features with the flowing fuel salt depicted as the transparent space. In the Simple Advection geometry, the regions axially above and below the moderator rods were treated as no-flow regions while that in the Incompressible Navier-Stokes allow fuel salt to flow freely.

For the symmetry sides of the mesh geometry, the reflective boundary condition of $\vec{u} \cdot \hat{n} = 0$ is applied. In addition to the velocity boundary conditions, the pressure at the reactor outlet was set to zero to 'pin' the pressure solution numerically.

The key output of the INS simulation is the velocity field of the fuel salt in the core, which is used in the simulation approach discussed in Section 4.3.1.3.

4.3.1.2. Simple Advection Neutronics Simulation and Optimization of Control Rod Maneuver Strategy

The neutronics simulation scheme used in this work was similar to that described in Section 4.2. However, an improvement was made in the way that the shim control was deployed.

From our simulations in Section 4.2, it is determined that a 'black mode' control rod configuration, where all control rods were inserted at the same axial position, resulted in a reactivity change that was too sensitive to rod insertion to allow practical reactor control. Therefore, shim rods are used to reduce the excess reactivity.

In Section 4.2, the outermost control rods were inserted fully to reduce the core reactivity down to a controllable range, which was managed by the central (load-following) control rod. Meanwhile, the rest of the 'unused' control rods were parked (withdrawn) in the upper plenum. In this section however, the control rod deployment is optimized by distributing the shim control evenly amongst the non-central control rods. This involved placing the non-central control rods at the same axial position to absorb the excess reactivity before using the central control rod for load-following.

The difference in neutron flux profiles resulting from the different shim controls is illustrated in Figure 4.12. In Figure 4.12, the axial views were taken from the diagonal side of the reactor in the RZ plane while the radial views were taken from the cross section of the reactor near the region of highest neutron flux.

In the previous shim control method, because there were two fully inserted shim rods, the neutron flux was being suppressed to form a cross-shaped region. If it is required that the reactor in both configurations produce the same amount of power, the heat generation in the previous configuration would be confined to the cross-shaped region and run the risk of overheating. Additionally, the neutron flux shape caused by the fully inserted shim rods would cause the moderator rods behind them to become ineffective.

In the current optimized control rod deployment because all control rods other than the central rod were placed at the same axial position, the neutron flux at the fissioning region of the core was unhindered by the control rods in the radial direction. This helped to spread out the neutron flux more evenly, leading to more uniform heat generation and resulting in a lower maximum temperature.

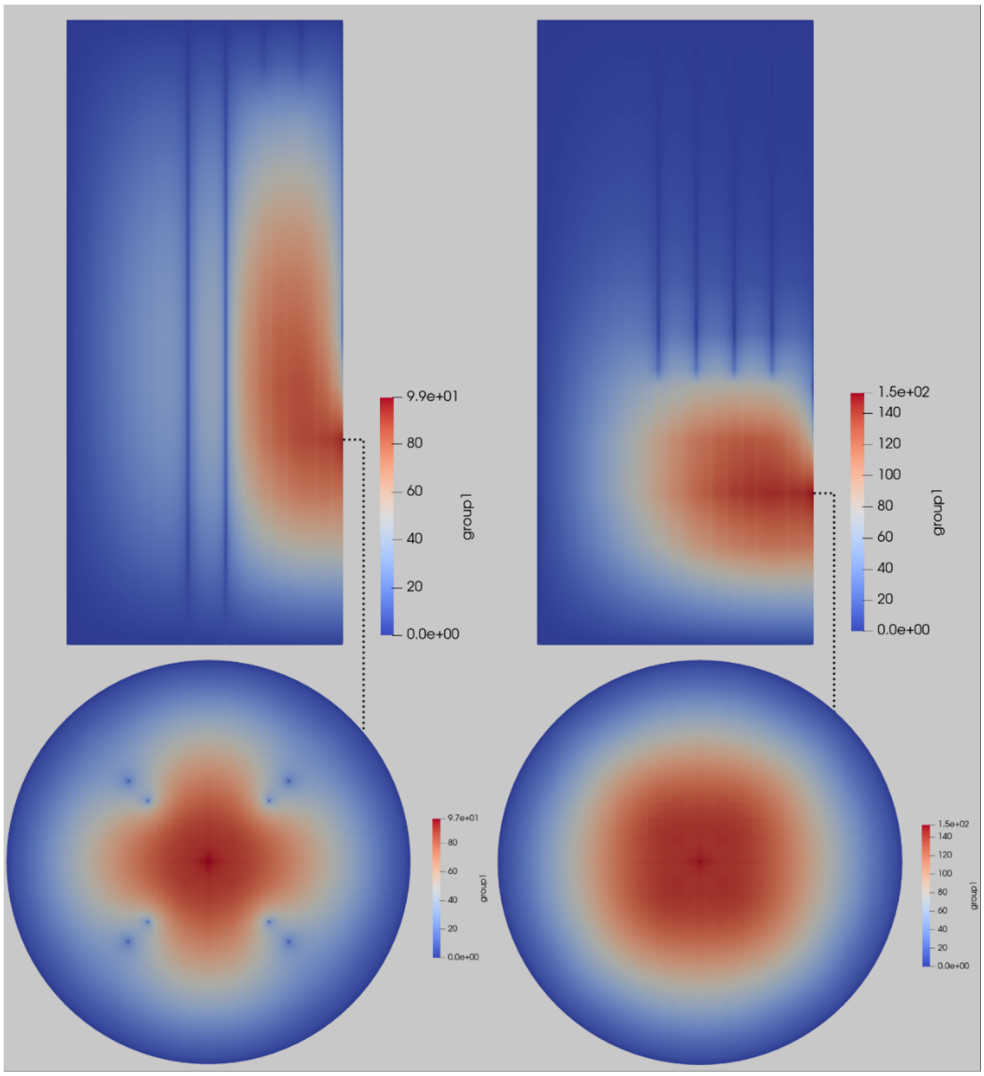

Figure 4.12: Neutron flux profiles (Group 1) resulting from the shim control methods used in the Section 4.2 (left) and the current section (right). The rectangular views at the top were taken from the diagonal side of the reactor in the RZ plane while the circular views were taken from top-down of the cross section of the reactor near the region of highest neutron flux. The difference in maximum neutron flux values in this figure was partly due to the difference in reactor power (left: 78.8% VS right: 88.0%). The units of the scales used were in 10^{13} #/cm^3.

4.3.1.3. Uncoupled Thermal-Hydraulics Simulation for Heat Advection

The heat generation profile from the neutronics simulation was combined with the INS velocity field solution to yield the *uncoupled* thermal-hydraulics solution. This was achieved by generating heat (temperature) using the neutron flux profiles obtained in the neutronics simulation and advecting the generated heat with the velocity field obtained in the INS simulation. The heat (temperature) advection was performed with SUPG stabilization as solving for scalar advection via the Finite Element approach would produce spurious and unphysical overshoots and undershoots if numerical stabilization was not used. Nonetheless, while greater amounts of

stabilization would lead to lower artefacts and faster convergence, care was taken to use the minimum stabilization required for the elimination of overshoots and undershoots in order to avoid introducing unnecessary numerical errors.

It should be noted that the above method only yields the *uncoupled* thermal-hydraulics solution, where there is no feedback between the neutronics and INS physics. This approach is chosen instead of a fully coupled method due to the prohibitively high computational cost required in the fully coupled method. This topic will be elaborated further in the discussion in section 4.3.2.4.

4.3.2. Results and Discussion

4.3.2.1. INS Solution

The velocity and pressure solutions of the INS simulation are shown in Figure 4.13. In the figure, the axial velocity of the fuel salt in the central bulk region between the moderator and control rods reached approximately 90 cm/s (upwards) before decreasing towards zero near the reactor surfaces. The velocity in the bottom plenum was largely uniform except at the regions upstream of the moderator rods, where the fuel salt decelerated upon approaching the moderator rods. In the top plenum, the fuel salt formed separate high and low velocity channels downstream of the moderator rods before mixing near the top boundary due to the imposed velocity boundary condition.

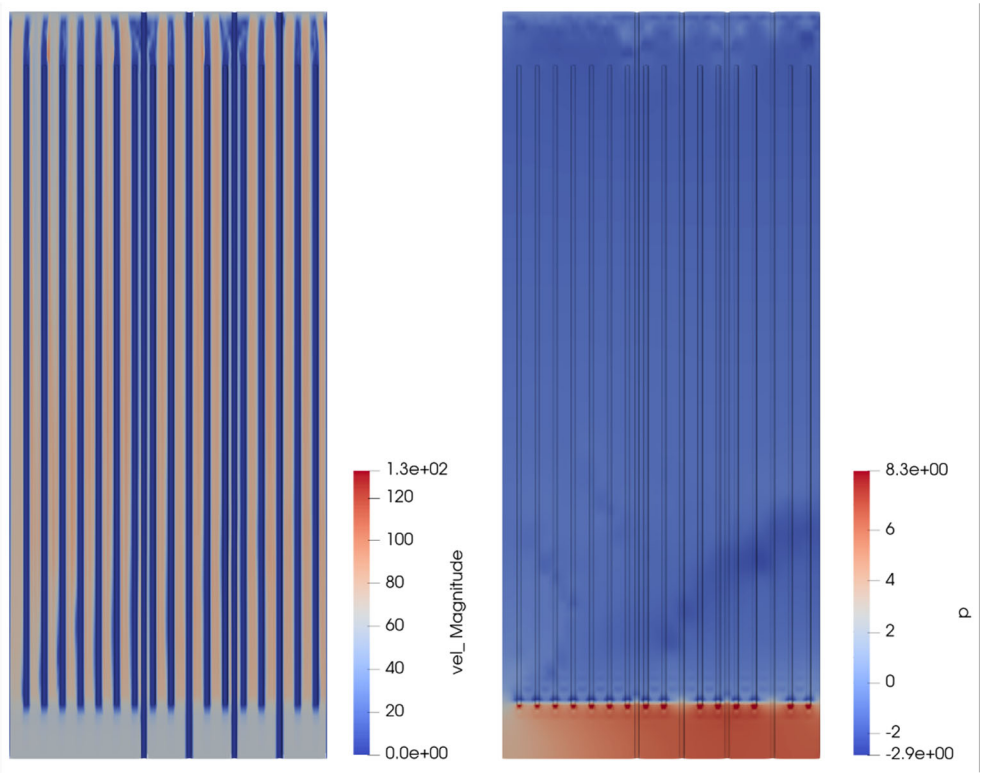

Figure 4.13: Velocity (left) and pressure (right) profiles from the diagonal side of the reactor in the RZ plane. In the figure, only the fuel salt was shown, with the reactor features (vessel wall, control

rods, moderator rods) removed. The units of the scales used were cm/s and 10^2 Pa for the velocity and pressure profiles respectively. The fuel salt flow in the figure was directed upwards.

Due to the coarse mesh size used in the 1/8th core model, only a laminar-like velocity field solution is obtained with approximately quadratic velocity profiles between reactor surfaces with no-slip boundary conditions. Nonetheless, the Reynolds number ($Re \approx 3300$) for the equivalent simple advection velocity suggests that the flow would have turbulent characteristics.

The discrepancy between the expected turbulent flow characteristics and laminar-like solution in the current INS velocity solution would result in two important consequences:

1) The actual velocity profile of the fuel salt between reactor surfaces would be more uniform compared with the current INS solution due to the presence of turbulent mixing. This flattening of the velocity profile would bring the axial velocity closer to that of the simple advection case. The turbulent mixing would also reduce the size of the low velocity zone around the reactor surfaces.
2) The increased mixing from turbulence would increase the heat transfer rate between the reactor surfaces and the bulk of the fuel salt. This would reduce the maximum local temperature of the reactor core by allowing hot fuel salt to advect more readily away from the reactor surfaces.

For the pressure profile, there was a distinct high-pressure zone in the bottom plenum before transitioning to a low-pressure zone once the flow passed upwards into the reactor active region. The small amounts of negative pressure values were the result of some undershooting of the pressure solution caused by an insufficient numerical stabilization. While the pressure profile suggests a sharp pressure drop of approximately 650 Pa as the fuel salt crosses upwards of the bottom plenum, it should be noted that the simulation was performed without gravitational acceleration. In practice, the hydrostatic pressure due the weight of the fuel salt ($\rho \approx 5000\ kg/m^3$) would be much greater than this pressure drop.

4.3.2.2. Simple Advection Neutronics Solution

The bounding neutronics solutions for the fresh fuel and equilibrium ^{135}Xe cases ("Fresh" and "Xenon" cases for brevity) with simple advection are presented in

Table 4.4.

Table 4.4 summarizes the percentage power level (of the rated 1250 MW_{th}), position of the shim and central control rods, average core outlet temperature and the maximum local temperature, at the high-power and low-power states. In these steady state solutions, the axial and radial neutron

flux profiles were similar to that shown in the Figure 4.12 (current shim control methods in the right diagrams).

Table 4.4: Summary of the reactor conditions for simulations performed with the simple advection scheme using Fresh Fuel and Equilibrium ^{135}Xe as the fuel salt. The table lists the reactor power level, shim and central control rod insertion as well as the average outlet and maximum local temperatures at the low and high power states. The up ramp and down ramp load-following operations were performed to and from these states for the respective Fresh and Xenon cases.

Case	Power Level (%)	Shim Insertion (%)	Central Rod Insertion (%)	Average Outlet Temperature (°C)	Maximum Local Temperature (°C)
Fresh (Low)	7.70	57	81	557	563
Fresh (High)	88.0	57	70	641	723
Xenon (Low)	5.97	39	69	555	561
Xenon (High)	69.5	39	62	622	704

From these high- and low-power steady states, the load-following simulations were performed for the Fresh and Xenon cases similar to that done in the previous section. The results of the power up ramp operations are shown in Figure 4.14.

For both the Fresh and Xenon cases, the movement of the central control rod in the power up ramp operation completed in around half a second. Since the graphs are approximately scaled with respect to their final power level, they provide a rough indication of the transient response time for the two cases. In Figure 4.14, the reactor response in terms of the rise in reactor power and maximum temperature was slower in the Xenon case compared with that in the Fresh case. Because both transients occurred on the order of a few seconds, the difference in response time would produce negligible impact on power load-following operations. However, the faster power response in the Fresh case could lead to a faster and greater rise in local temperatures due to a more rapid heating of the fuel salt. This is observed in Figure 4.6 where there is a greater and faster rise in maximum local temperature in the Fresh case as compared with the Xenon case. Additionally, since the region with active fissions was confined within the 'un-shimmed' region, which was smaller in the Fresh case (129 cm axially) as compared with the Xenon case (183 cm axially), the resulting neutron fluxes and the associated volumetric heat generation was higher in the Fresh case compared with the Xenon case.

Since the goal is to determine whether the thermal safety constraints of the reactor could be met during load-following operations, the more conservative Fresh case is used to provide the heat generation profile for the uncoupled thermal-hydraulics solution.

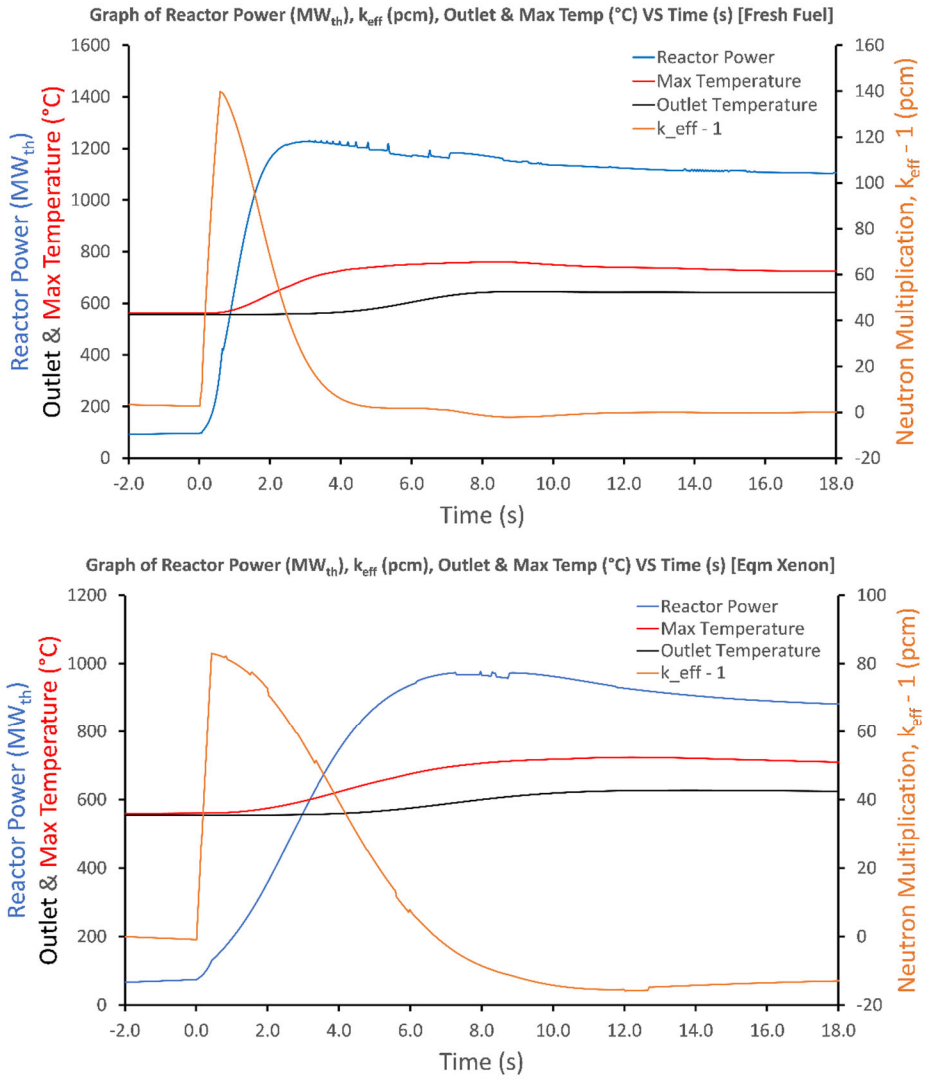

Figure 4.14: Graphs of Reactor Power, Neutron Multiplication Factor (k_{eff}, presented as $k_{eff} - 1$), Outlet and Maximum Temperatures during power up ramp operations for Fresh Fuel (top) and Equilibrium ^{135}Xe (bottom). In the graphs, the movement of the central control rod started at t = 0.00 s and completed at 0.66 s and 0.42 s for the Fresh and Xenon cases respectively.

4.3.2.3. Uncoupled Thermal-hydraulics Solution

The temperature profile in the uncoupled thermal-hydraulics solution is shown in Figure 4.15. In the figure, it is observed that the highest temperatures occurred along the walls of the reactor features and the maximum local temperature reached around 1700 K (~1400 °C). As discussed in Section 4.3.2.1, this was due to the low fuel salt velocity along the walls of these reactor features, causing heat generated at these regions to be trapped by the stagnant flow (i.e., heat in these regions mostly dissipated via thermal diffusion instead of convection). As mentioned in Section 4.3.2.1, additional mixing due to turbulence in the fuel salt flow (not included in this simulation) should help in heat removal.

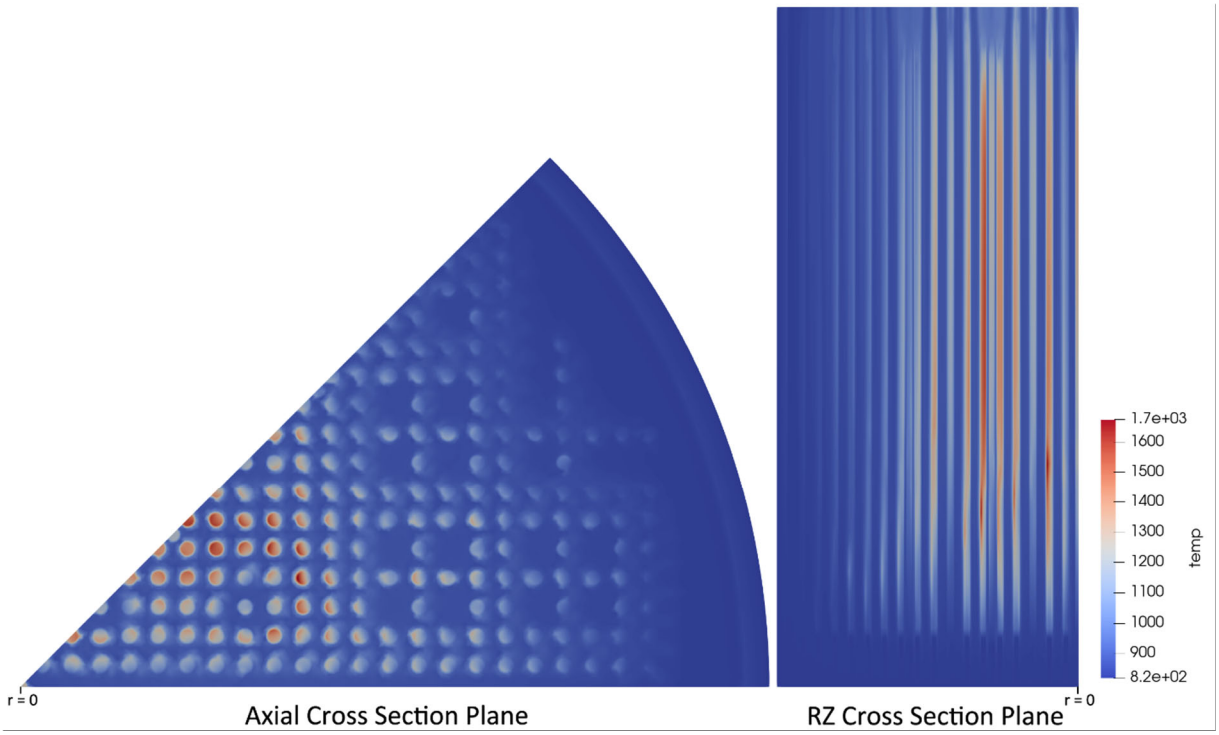

Figure 4.15: Axial (left) and RZ (right) cross section planes of the temperature profile solved using the solutions from the neutronics and INS simulations at t = 600 s (10 mins). The greatest local temperatures occurred along the walls of the reactor features which correspond to the regions of low flow velocities in the INS solution. In the figure, the temperature scale shown was in Kelvins.

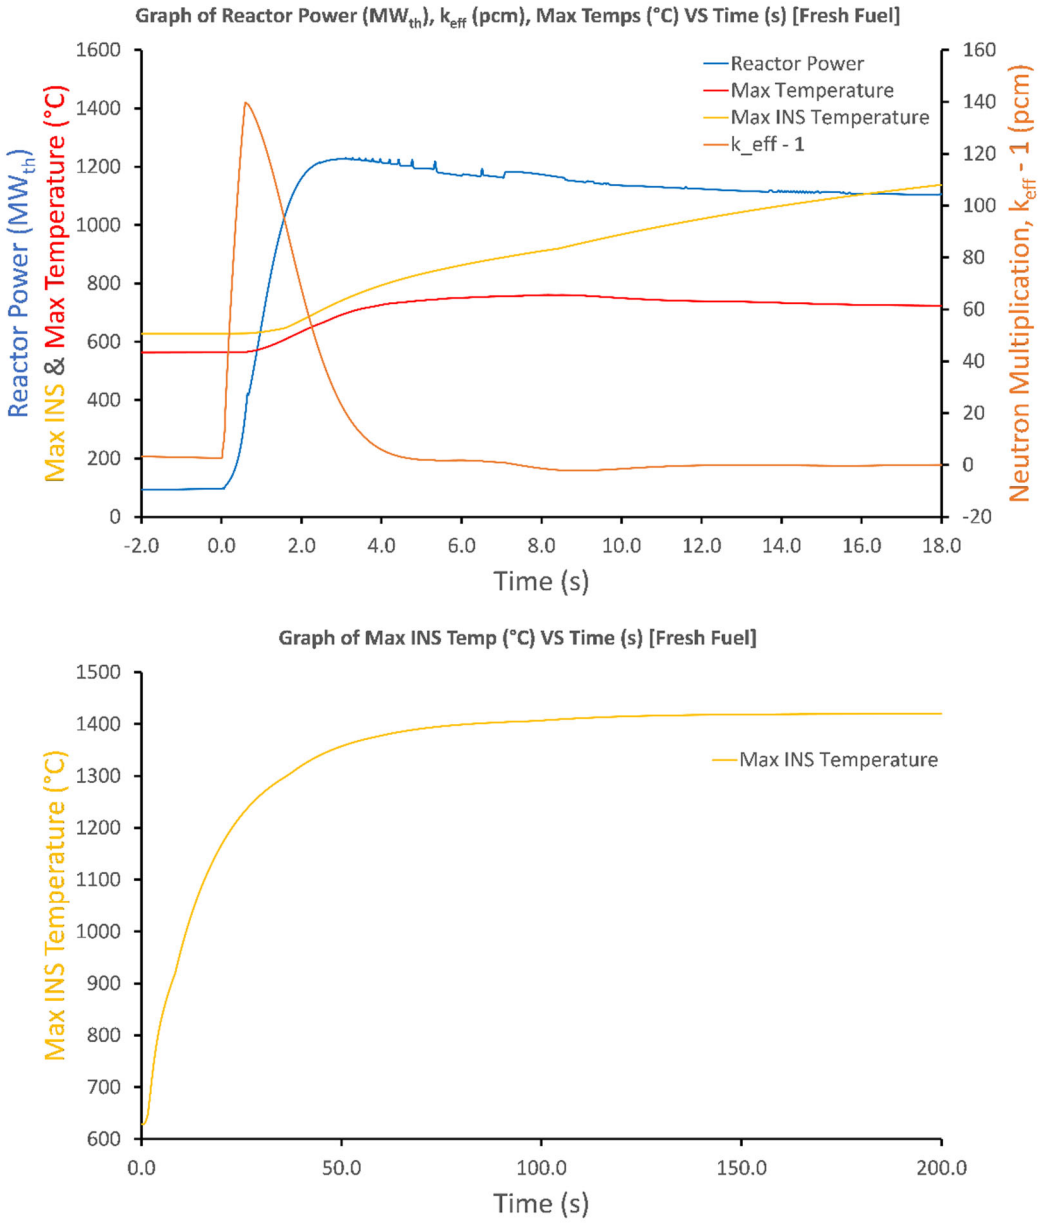

Figure 4.16: Graph of the maximum local temperature in the combined thermal-hydraulics solution plotted with the simple advection solution (top) for the first 18 seconds, as well as a long run plot for 200 seconds (bottom). The maximum local temperature stabilized around 1420 °C at the end of the bottom graph.

The maximum local temperature in the uncoupled thermal-hydraulics solution is shown in Figure 4.16. In the top graph, the response time of the maximum local temperature in the combined solution was as fast as that in the simple advection case. However, while the maximum temperature in the simple advection case stabilized after around 20 s, the maximum local temperature in the combined solution continued increasing. From the bottom graph in Figure 4.16, the temperature stabilized around 200 s, reaching 1420 °C.

The maximum local temperature of 1420 °C obtained here clearly exceeded the temperature constraint of 900 °C for the Hastelloy components. This suggests that the methods used in this section might be inadequate and warrants an alternative approach in obtaining a temperature solution for the evaluation of the safety constraint.

4.3.2.4. Degree of Physics Coupling and Associated Difficulty

In this section, the heat generation profile from the neutronics simulation is combined with the INS velocity field solution to yield the *uncoupled* thermal-hydraulics solution, where heat generated in the fissioning region of the core moved through the core by following the velocity field. The simulation has *no coupling* between the neutronics and INS physics because the neutronics and velocity fields were solved separately without feedback from each other.

Extending from this approach would be a *one-way coupling* where either the INS velocity field solution is used in solving for the neutronics solution or vice versa. The determination of whether to use the neutronics solution or the INS solution as the independent variable depends on the relative sensitivities of the two different physics involved.

Finally, if both the neutronics and INS physics kernels are included in the MOOSE/Moltres input file, both the neutronics and INS physics simultaneously can be solved, and a fully coupled (two-way) solution is obtained. This was the original intended approach of this study but was found to be unachievable due to the great computational cost required.

The difficulty in moving from the *uncoupled* approach towards higher degrees of physics coupling can be summarized by the following two important factors:

1) **INS mesh size**: The INS simulation requires a very fine mesh in order to solve for the velocity field of the TAP MSR accurately. This is in contrast to the mesh required by Moltres, which can tolerate relatively coarse meshes. In this work, the original mesh is refined by a significant degree in order for the INS simulation to produce a reasonable solution. Such mesh refinement can be seen in Figure 4.17, where the mesh interval near the periphery had increased by a factor of 6 and the axial interval had increased from once every 30 cm to once every 2.5 cm, a factor of 12. As a result of the mesh refinement, the degrees of freedom (DOF) in the INS simulation increased from 0.33 M to 4.9 M (and to 2.5 M in an intermediate mesh with 5 cm axial interval).

   Using 1152 cores on the INL Falcon HPC, the INS simulation took about 110 hours to reach steady state (using the 5 cm axial mesh; 127 k core-hours) and an additional 6 hours for the stabilization and refinement step (using the 2.5 cm axial mesh; 7 k core-hours).

2) **Number of Variables**: The INS problem contained 4 non-linear variables (pressure and the x-, y-, and z- velocity components) while the neutronics problem contained 11 non-linear variables (2 neutron flux groups, temperature, and 8 delayed neutron precursor groups, DNP). To fully couple the neutronics and INS physics into one simulation, the neutronics kernels would have to share the fine mesh required by the INS simulation. Alternatively, a slightly weaker coupling approach could be taken with the use of the MOOSE MultiApp system where the INS physics submodule solves for 13 non-linear variables (pressure, three velocity components, temperature

advection, and 8 DNP advection) before down-sampling and passing the temperature and DNP solutions to Moltres to solve for the neutron fluxes on a coarser mesh.

The former approach would increase the degrees of freedom tremendously to an estimated 9.3M when using the 5 cm axial mesh (15 non-linear variables) while the latter approach would produce less accurate solutions without necessarily being less costly (13 non-linear variables on fine mesh). Since the computational cost increases non-linearly with the increase in DOF, the simulation with the estimated 9.3M DOF would be prohibitively expensive to solve.

In this section, moving from 2.5 M DOF to 4.9 M DOF resulted in an increase in computational time from 11 minutes to 40 minutes per timestep. If it is assumed that the computational cost increases in a quadratic manner with respect to DOF, a simulation with 9.3 M DOF would require 160 minutes per timestep on 1152 cores, which would be an exceedingly costly endeavor.

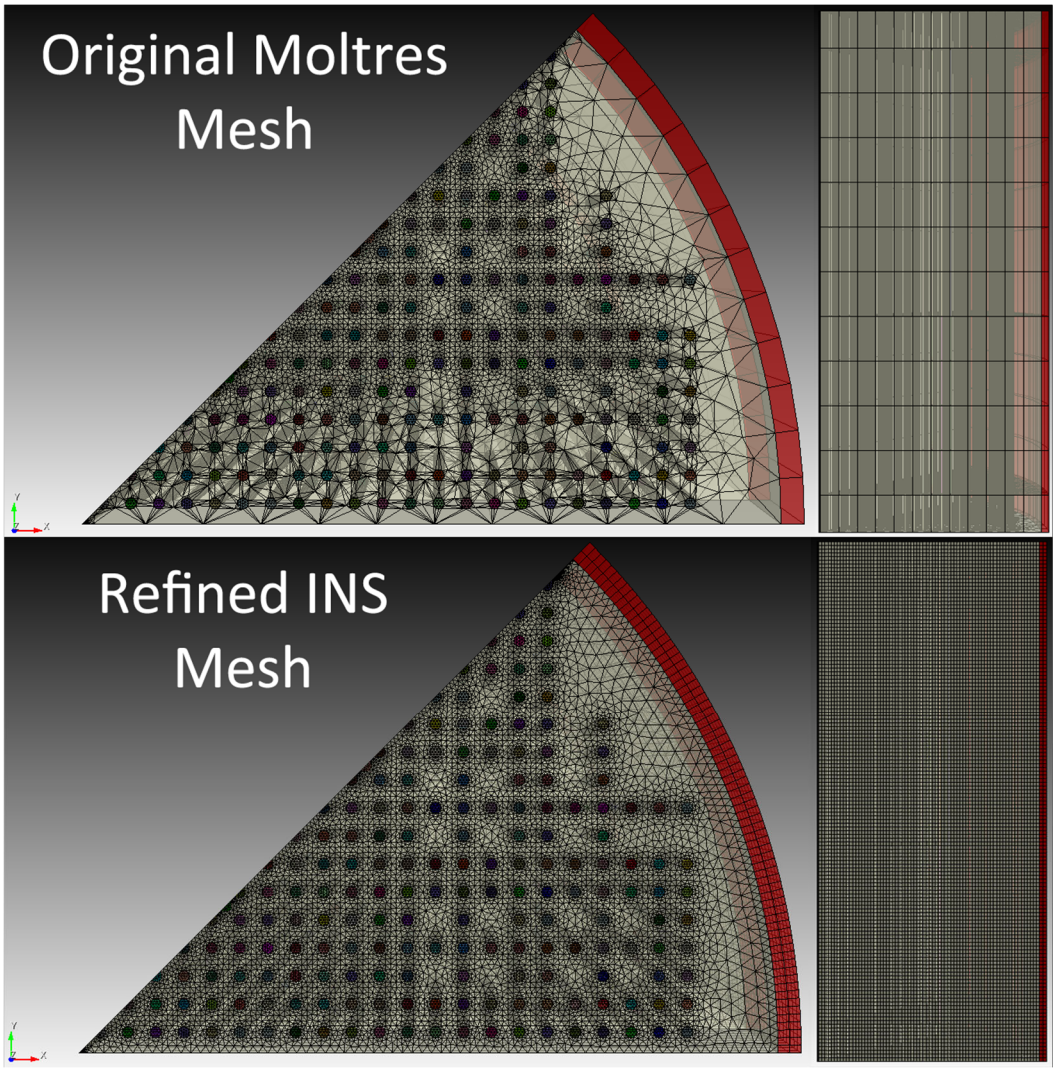

Figure 4.17: Axial (left) and RZ (right) cross section planes of the TAP MSR mesh, showing the difference between the original Moltres mesh and the refined INS mesh. The refined mesh shown here has a perimeter interval that is 6 times that of the original mesh, and an axial interval that is 12 times compared with the original mesh.

In light of the above, further physics coupling involving the $1/8^{th}$ core geometry is unlikely to be practicable. In order to proceed with a higher degree of physics coupling, the size of the problem had to be reduced. This was done in Section 4.4 by simulating only a unit cell channel.

4.4. Unit Cell Approach for Thermal Hydraulic Simulation

The unit cell approach was developed to study the localized temperature distribution without being severely limited by the computational resource requirements of the $1/8^{th}$ core simulation.

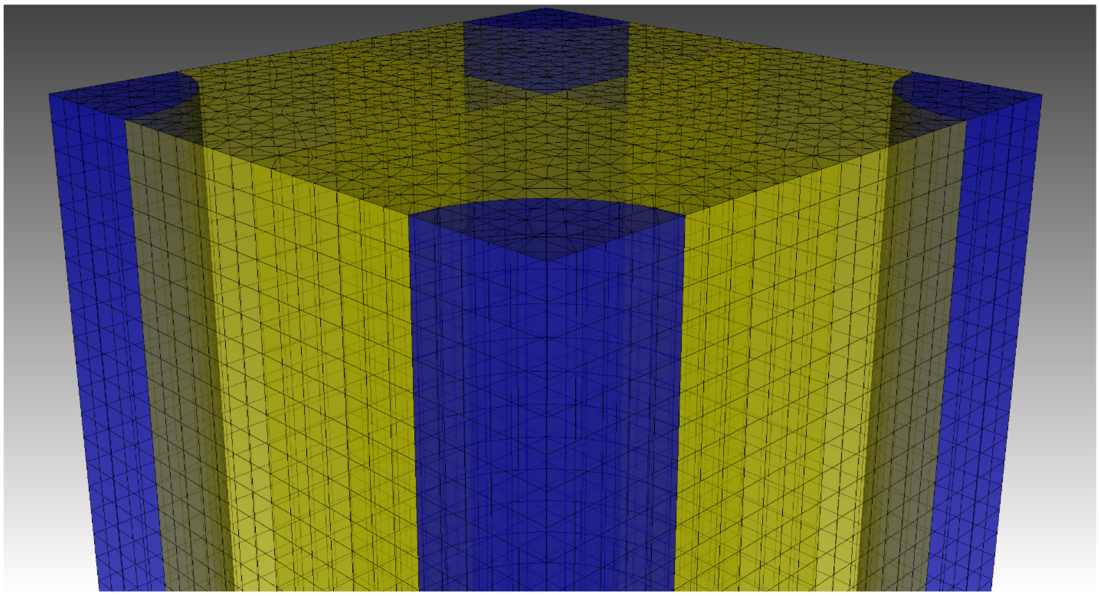

Figure 4.18: Side view of the full unit cell mesh. The fuel regions continuing axially from the moderator rods are shown in blue while the fuel channel is colored yellow. Due to the symmetry present in the unit cell, only one corner of the unit cell mesh was used in the simulations.

4.4.1. Methodology

The unit cell approach involved the following key steps:

1) First, a simplified mesh that represents a typical fuel channel in the TAP MSR is generated as shown in Figure 4.18. This fuel channel consists of a central fuel region surrounded by four moderator rod quarters in each corner. Due to the symmetry present in the fuel channel, the actual mesh used was only one quarter of the unit cell and reflective boundaries were specified at the symmetry surfaces.
2) Axially, the unit cell consists of the top and bottom plenum regions, in addition to the central region that is partially occupied by the moderator rod quarters. To simplify the mesh further, the cladding material (SiC) was also replaced by the material of the moderator rod ($ZrH_{1.66}$).
3) Then, the appropriate simplified boundary conditions are specified. For heat generation, the conservative neutron flux values of $100 \times 10^{13}/cm^3$ (Group 1) and $4 \times 10^{13}/cm^3$ (Group 2) were used to determine the volumetric heat generation of $196\ W/cm^3$, which was specified for the region between the top edge of the bottom plenum (z = -150 cm) and the bottom edge of the shim control rods (z = -21 cm). [See Figure 4.19, z = 0 cm represents the middle of the core.]
4) These values correspond to the typical values for the high-power state in the Fresh case. The velocity at the inlet and outlet was set to 59.4 cm/s while moderator surfaces were given the no-slip boundary condition.
5) As with the earlier approach, the steady state velocity profile is solved first (30 s simulation time) before introducing heat generation physics, in order to avoid overheating stagnant fuel.

The unit cell mesh size was also much finer than the 1/8th core mesh, with an element size of 2.5 mm as compared with 25 mm for the 1/8th core mesh (a 10 × resolution increase).

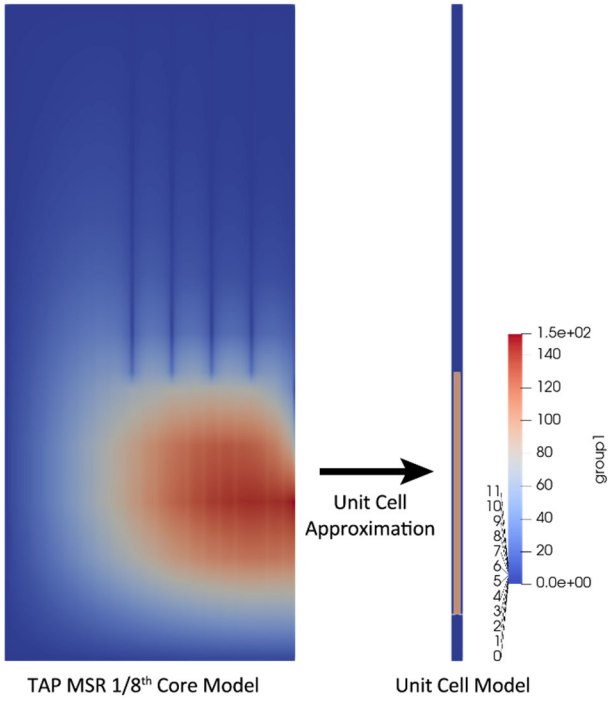

Figure 4.19: Diagram showing the approximation of Group 1 neutron flux in the unit cell model from the Fresh case high-power state simulated using the 1/8th core model.

After obtaining the steady state velocity field solution at 30 s simulation time, the heat generation physics is turned on (INS physics kernels remain enabled), and the simulation solved for temperature advection, as well as the velocity and pressure fields in the presence of gravity.

Since the density of the fuel salt is temperature dependent, the simulation is similar to the Boussinesq approach where the buoyancy force caused by the lower density of the hotter fuel salt is considered. In addition to this, the temperature dependent viscosity of the fuel salt is also included where an increase in fuel salt temperature leads to a decrease in viscosity[25].

4.4.2. Unit Cell Model Results

Figure 4.20 shows the evolution of the fuel average temperature, core exit temperature and maximum local temperature from the unit cell simulation. In the figure, the average fuel temperature increased approximately linearly before reaching a steady temperature of 613 °C at about 70 s. However, the maximum local temperature exhibited a notable fluctuation between 750 °C and 785 °C and this fluctuation appears to have influenced the core exit temperature, which also showed a fluctuation between 630 °C and 640 °C.

[25] Viscosity data from [170], fitted to this relation: $\mu(T) = (1.076111581E-2) \times \left(\frac{T(K)}{1000}\right)^{-4.833548134} Pa \cdot s$

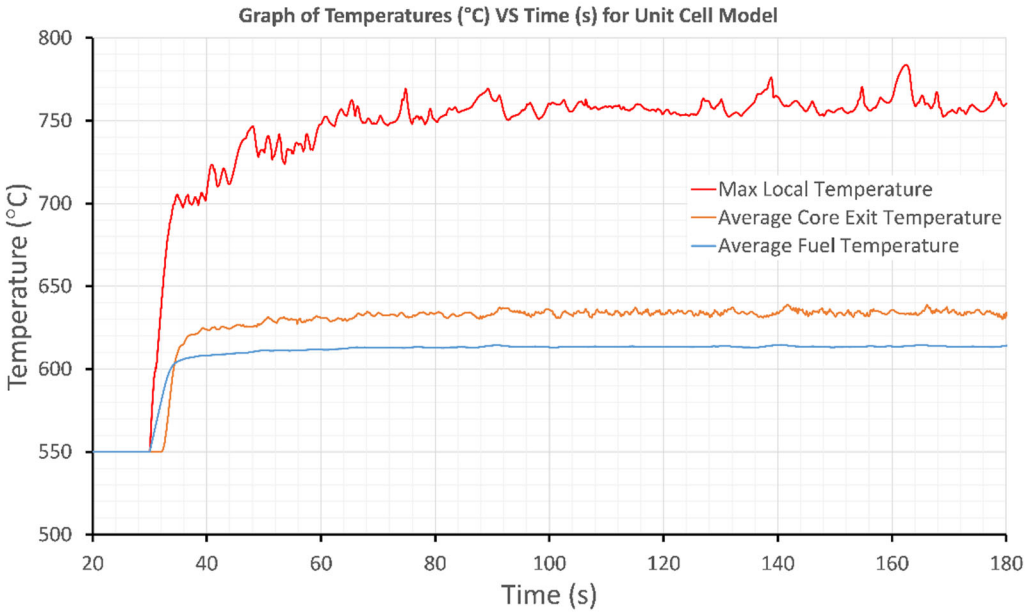

Figure 4.20: Graph of the maximum local temperature, average core exit temperature and average fuel salt temperature from the unit cell simulation. While the average fuel salt temperature increased fairly smoothly to a steady temperature, the maximum local temperature and average core exit temperature exhibited notable fluctuations.

An examination of the simulation results revealed that the maximum temperature occurred at the top end of the heated region ($z \approx -21$ cm) near the moderator rod surface where the axial fuel velocity was the lowest. This result was expected, as the slowest moving fuel salt near the surface would continuously gain heat until the end of the heated region where its thermal energy (and temperature) would be the greatest. However, the fluctuation in the maximum temperature in this region was unexpected.

Further examination of the results uncovered a region of flow instability near the bottom end of the moderator rod. From Figure 4.21, there was a sharp rightwards deflection of the fuel salt flow at the bottom of the moderator rod where the upwards flowing fuel salt meets an abrupt flow obstruction (moderator rod). This deflected flow continued for some distance before joining the bulk salt flow and in the region just above this flow deflection, two separating vortices can be clearly seen. The vortices are an example of the Kelvin-Helmholtz (KH) instability, which is caused by the velocity difference between the deflected fuel salt and the salt just above the deflection. In addition to this velocity difference, non-linear effects from the temperature dependent density and viscosity helps to perturb the salt flow, which further drives the instability.

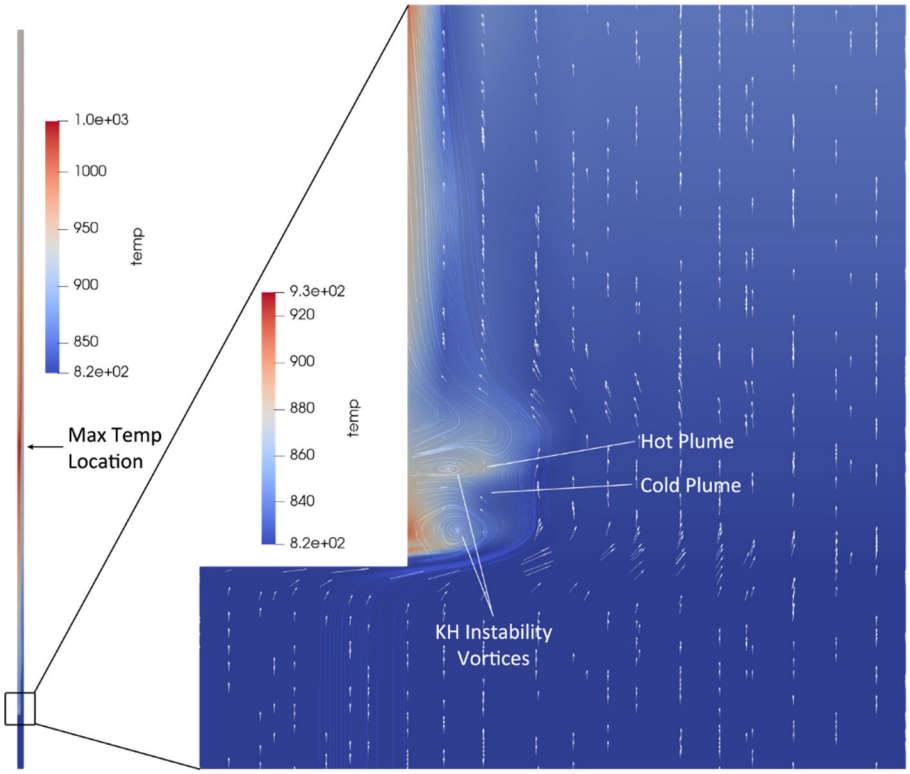

Figure 4.21: Unit cell channel (left) showing the maximum temperature location as well as the zoomed-in 45° cross section view near the Kelvin-Helmholtz (KH) instability region (right), with the x- and y-axes scaled upwards by a factor of 10. The moderator rod has been excluded from the view so that the boundary can be seen clearly. The white arrows depict the direction and magnitude of the salt flow velocity while stream tracer lines help to visualize the salt flow. In the figure, the hot and cold plumes as well as the KH instability vortices can be clearly observed.

As the flow was unstable, counter-rotating vortices would continually emerge and separate from the instability region before being swept upwards. When this happens, the fuel salt in this region would have fluctuating velocities, leading to an increased or decreased residence time in the heated region, which creates alternating hot and cold "plumes". When a hot plume reaches the end of the heated region, it causes the maximum temperature to rise, and conversely for a cold plume.

In addition to the plume effect, it should also be noted that unlike the laminar-like result discussed in Section 4.4.3, the velocity field in the unit cell approach is no longer strictly in the axial direction. This means that the fuel salt near the moderator rod surface would be able to advect into the bulk fuel salt flow when it acquires some radial velocity.

These effects lead to a fluctuation of the maximum local temperature as shown in Figure 4.20 and more importantly, result in a disruption of the continuous buildup of thermal energy by mixing otherwise stagnant fuel salt, thereby limiting the maximum local temperature. From the unit cell simulation with conservative volumetric heat generation, the peak maximum local temperature was no greater than 783 °C, which is below the temperature constraint of 900 °C specified in Section 4.1.3.

4.4.3. Mesh Resolution in a Unit Cell

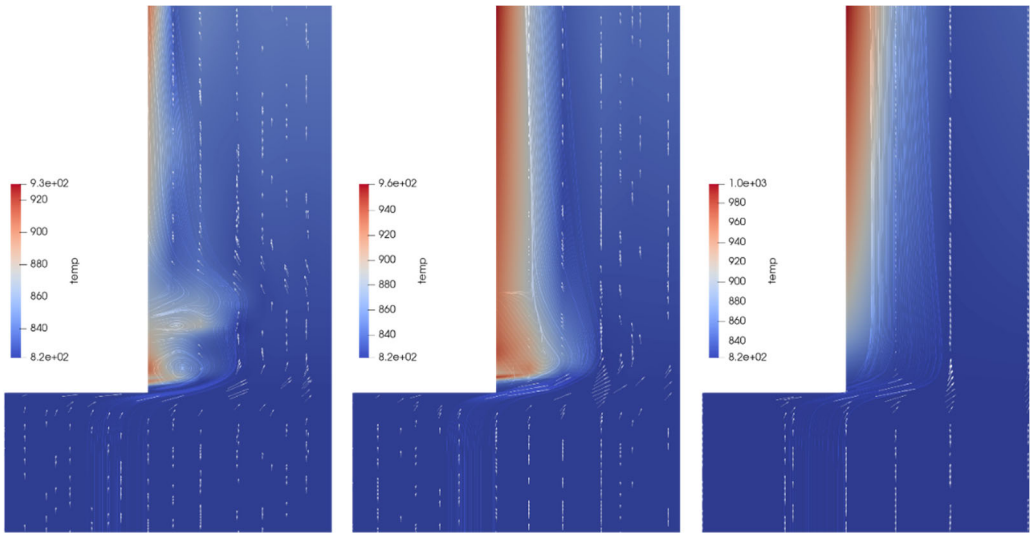

Figure 4.22: Unit cell approach simulated with three different mesh sizes of 2.5 mm, 3.33 mm, and 5.0 mm. The increase in mesh size resulted in a decrease in spatial resolution which led to the disappearance of fine flow details and features such as the vortices from the KH instability.

The mesh resolution has a very significant impact on the ability to simulate fine flow details such as those encountered in flow instabilities. This is illustrated in Figure 4.22, which shows the unit cell approach simulated with mesh sizes of 2.5 mm, 3.33 mm, and 5.0 mm. From the figure, as the mesh size increased, the flow became more regular, and the vortices diminished before disappearing entirely in the 5.0 mm mesh. The disappearance of fine flow details was due to the mesh size being larger than the length scale of the flow features, such that there was not enough spatial resolution to resolve these flow features. Conversely, it is expected that a decrease in mesh size would increase the spatial resolution and enhance the level of detail, resulting in a more complex and turbulent flow profile.

An important consequence of the increase in mesh resolution was the appearance of a mixing effect that helped to disrupt the continuous buildup of heat in the fuel salt. This helped to limit the maximum local temperature of the fuel salt and was the reason for the notable difference in results between the unit cell approach and the $1/8^{th}$ core approach.

4.5. Investigation on Gas Entrainment

Besides the thermal-hydraulics simulations to investigate the temperature safety of the TAP MSR, the effects of gas entrainment on the neutronics behavior of the TAP MSR core are investigated. The source of the gas entrainment into the fuel salt would be the sparging process, which involves the introduction of Helium gas to absorb ^{135}Xe and other fission gasses before a gas separation step to remove the gas from the fuel salt. In the case of an incomplete gas removal, tiny bubbles would remain trapped in the fuel salt that enters the reactor core, potentially affecting its load-following behavior.

In order to study the effects of the gas entrainment, the up-ramp load-following operation is simulated under the four conditions described in Table 4.5.

Table 4.5: List of cases simulated in the gas entrainment study.

Case	Salt Composition	Helium Volume Fraction
1 – B0	BOL +6 days	0%
2 – B1	before sparging	1%
3 – A0	BOL +6 days	0%
4 – A1	after sparging	1%

The salt composition is chosen to be that at 6 days after BOL in order to allow ^{135}Xe (and other fission products) to build up to a steady state. This is represented by Case 1 (B0), which is just before sparging, with 0% Helium volume fraction. Immediately after the start of gas sparging, assuming an incomplete gas removal giving rise to 1% gas entrainment, the fuel salt composition is effectively the same as that before sparging but with a slightly lower density. This is represented by Case 2 (B1).

After running the sparging process for a sufficient amount of time, almost all ^{135}Xe is removed but still with 1% gas entrainment. This is represented by Case 4 (A1). Thereafter, helium gas flow is stopped, and the gas separator remains operating for a sufficient period of time so that all Helium gas is removed, which is represented by Case 3 (A0). These four cases represent the extreme scenarios in the amount of gas entrainment and thereby allow us to bound the design and operation of the gas sparging system.

4.5.1. Methodology

The gas entrainment investigation involved the following key steps:

1) First, the fuel salt composition at 6 days after BOL is generated, before and after sparging, using SaltProc [156] developed in Task 2 of the MEITNER project.
2) Then, Serpent 2 is used to calculate the neutronics cross sections required by Moltres for the four cases. For the 1% Helium gas entrainment cases, we assumed homogeneous mixing of the gas into the fuel salt and we modified the density of the fuel salt accordingly to take into account the presence of the 1% Helium gas at 900 K[26].
3) Finally, the neutronics simulation methodology described in Section 4.3.2 is used to obtain the high- and low-power states before performing the up-ramp load-following simulations. In both 1% gas entrainment cases, the same control rod insertion positions

[26] The formula used to calculate the gas-entrained fuel salt density is $\rho_{f,new} = 0.99\rho_f + 0.01\rho_{He}$, where we have used $\rho_{He} = 0.0526\ kg/m^3$ for the density of Helium at 900 K.

are used as their corresponding 0% entrainment cases in order to determine the differences caused only by the gas entrainment.

4.5.2. Results and Discussion

The results of the high- and low-power states corresponding to the four cases are summarized in Table 4.6.

Table 4.6: Summary of the reactor conditions for simulations performed with the simple advection scheme using fuel salt compositions in the four cases. The table lists the reactor power level, shim and central control rod insertion as well as the average outlet and maximum local temperatures at the high- and low-power states.

Case	Power Level (%)	Shim Insertion (%)	Central Rod Insertion (%)	Average Outlet Temperature (°C)	Maximum Local Temperature (°C)
1 – B0 (Low)	8.78	49	76	558.7	566.8
1 – B0 (High)	83.4	49	64	635.5	723.9
2 – B1 (Low)	6.74	49	76	556.9	565.0
2 – B1 (High)	111.2	49	64	666.8	788.5
3 – A0 (Low)	3.39	57	75	553.3	556.2
3 – A0 (High)	90.7	57	61	644.1	735.2
4 – A1 (Low)	22.9	57	75	570.1	591.4
4 – A1 (High)	118.0	57	61	675.4	797.6

From Table 4.6, it can be seen that while Case 1 (B0) and Case 3 (A0) were operating well within the reactor's operation envelope, the entrainment of sparging gas into the fuel salt generally resulted in higher reactor power levels and operating temperatures.

Comparing Case 1 (B0) and Case 2 (B1), the 1% gas entrainment led to a slightly lower power level (-2.04%, relative to rated power) at the low-power configuration but a significant increase in power level (+27.8%) at the high-power configuration. Similarly, between Case 3 (A0) and Case 4 (A1), gas entrainment leads to the higher power level changes of +19.5% and +27.3%, respectively. This implies that gas entrainment, resulting in lower fuel salt densities, would introduce an overall positive reactivity to the reactor core (i.e., a positive void coefficient). The

positive void coefficient had been previously reported in [18] and is due to an increase in moderator-to-fuel ratio in the undermoderated TAP MSR[27] [15].

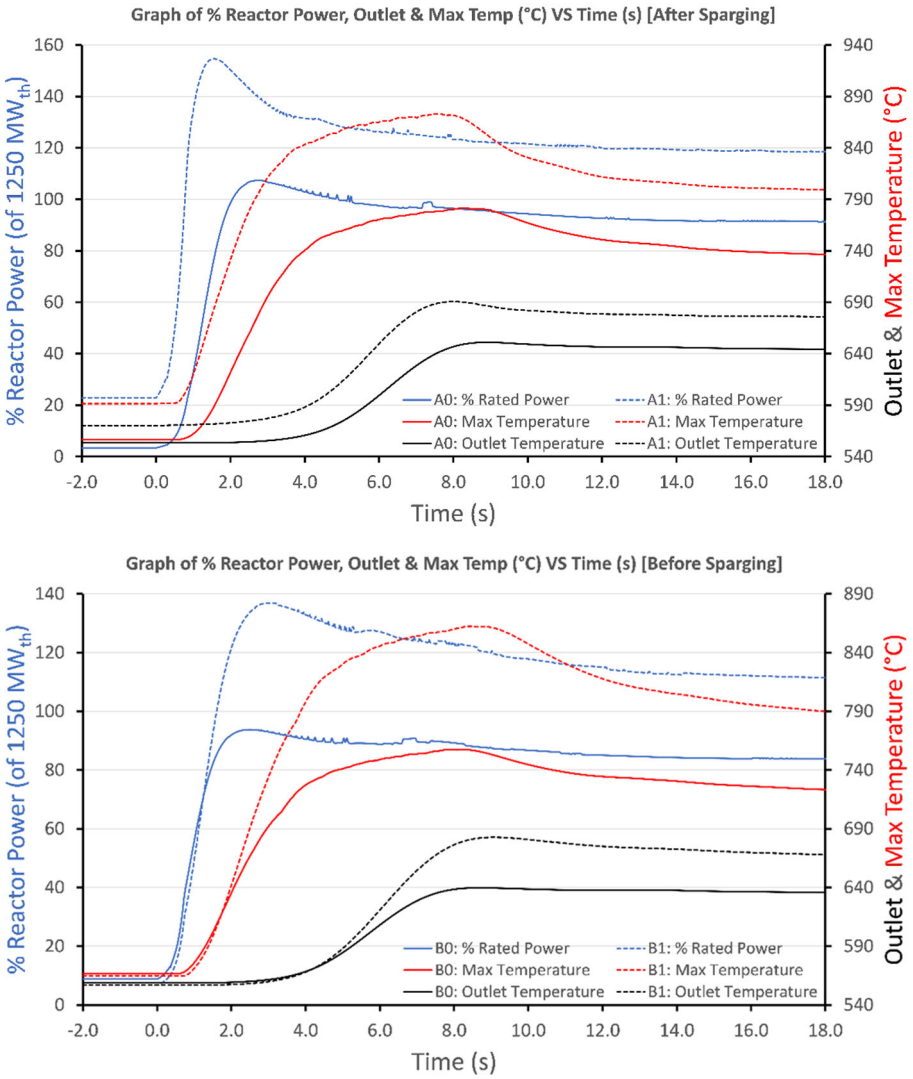

Figure 4.23: Graphs of % Reactor Power, Outlet Temperature and Maximum Temperatures during power up ramp for the salt compositions 'before' sparging and after sparging (removal of ^{135}Xe and other fission gasses). In the graph, the cases with 1% gas entrainment are plotted with dotted lines. The small spikes in the % Reactor Power were caused by a difference in the number of non-linear iteration solves in some time steps, due to the use of adaptive time stepping.

The results from the power up ramp simulations for these four cases are plotted in Figure 4.23. From the figure, all three important parameters (reactor power level, core outlet temperature and maximum local temperature) were higher in the 1% gas entrainment cases, as expected.

[27] The TAP MSR was designed to begin operation with a low moderator-to-fuel ratio so that additional moderator rods can be gradually inserted to counteract the effects from the build-up of neutron absorbing fission products.

Nonetheless, all four cases exhibited the same general up ramp behavior, which was a sharp rise in power level and temperatures before decreasing towards a new steady state due to negative feedback. The key takeaway from the power up ramp simulations is that the reactor core's response time to a load-following operation is remarkably fast, on the order of seconds, even with the presence of gas entrainment.

4.5.3. Gas Entrainment and Reactor Safety

Gas entrainment in the fuel salt introduces an overall positive reactivity to the reactor core. This means that with an intermittent gas sparging system, the sparging process should not be started when the reactor is operating close to its full rated power level. As discussed previously, the presence of gas entrainment did not significantly affect the speed of the reactor response. This suggests that the increase in reactor power level due to gas entrainment could be as fast as the rate of gas introduction into the system and that power reduction maneuvers should be undertaken in anticipation of a gas sparging operation[28].

The above issue can be reduced or even avoided by (i) limiting the fraction of fuel salt that is diverted to the gas sparging process and/or (ii) performing continuous gas sparging. In the former approach, the impact of starting the gas sparging process is decreased by limiting the rate of gas introduction into the system, thereby allowing the operator enough time to adjust for the changing power level. In the latter approach, the system is allowed to operate continuously with 1% gas entrainment (or any other efficiency limit of the gas separator) such that any additional removal of entrained gas would result in a decrease in reactor power level. Since it is expected that the gas sparging system could be taken offline for maintenance, a combination of both approaches would be prudent.

4.6. Potential Extensions and Conclusions

4.6.1. Potential Extensions

4.6.1.1. Looping of Delayed Neutron Precursors in Fuel Salt

In this work, the Moltres simulations focused specifically on the core of the reactor in order to study the neutronics behavior of the reactor during load-following. A potential extension would be to consider the rest of the primary loop and the effect it may have on the load-following behavior of the reactor core.

[28] It should be noted that while Table 4.4 and Figure 4.23 showed an overloading of the TAP MSR under gas entrainment, power reduction maneuvers would have been undertaken in anticipation of gas sparging and such reactor overload should not occur in practice.

One important consideration is the looping of delayed neutron precursors where the fuel salt would spend some time outside of the reactor core before being reintroduced into the bottom of the core. The looping of the precursors is expected to increase the reactivity of the reactor as previously generated neutron precursors are being reintroduced into the TAP MSR core. There will also be temporal effects on load-following due to the finite amount of time required for the fuel salt to complete a recirculation cycle. Therefore, it may be worthwhile to explore the effects of the looping of delayed neutron precursors on the TAP MSR load-following behavior in future work. Such work may be performed using a system-wide simulation code such as the System Analysis Module (SAM) [202] developed at the Argonne National Laboratory.

4.6.1.2. Neutron Spectrum and Load Following

According to the work presented in Chapter 3, as the TAP MSR operates in the epithermal neutron spectrum at the beginning of life (BOL), the effects of ^{135}Xe and ^{135}I on changing power level is very much diminished and the iodine pit that occurs in conventional thermal-spectrum reactors does not exist in the TAP MSR at BOL. Therefore, the presence of ^{135}Xe and ^{135}I in the TAP MSR would not significantly affect the load-following capability (ramping rates) of the TAP MSR during BOL, as was the case in this work.

However, the TAP MSR is expected to operate in an increasingly thermal neutron spectrum over the lifetime of the reactor as fission products build up and more moderator rods are inserted to maintain criticality. Previous research on a molten salt breeder reactor [156] had showed a notable change in control rod worth over the lifetime of the reactor as the neutron spectrum changes. As a result, it is expected that the effects of ^{135}Xe and ^{135}I on the load-following capability would also become more prominent for the TAP MSR over its lifetime. Hence, it may be worthwhile to explore the load-following behavior of the TAP MSR using fuel salt compositions at various stages of the reactor lifetime. This can allow us to fully understand the effects of the changing neutron spectrum and ascertain that sufficient shutdown margin can be maintained.

These fuel salt compositions could be obtained using the SaltProc [123] code that is discussed extensively in Chapter 3, which can simulate the online removal of targeted fission products while the fuel is being depleted. Such simulation of online fuel reprocessing is important in order to accurately predict the fuel salt composition over the lifetime of the reactor as more moderator rods are inserted to maintain criticality.

4.6.2. Lessons Learned

- For the use of thermodynamic simulations (INS) to determine bounding scenarios (e.g., determine whether maximum temperature would exceed material temperature), more insights can be gained from a fine mesh simulation of a unit cell than a coarse mesh simulation of the full core.
- Using no-slip BC (which is a valid condition) with coarse mesh will cause the simulation to arrive at a laminar solution even with a large Reynold's number, which is incorrect. With a fine mesh, turbulent flow characteristics can be obtained.

- At the TAP MSR salt flow speed (substantial Reynold's number of Re $\approx$ 3300), the turbulent-like characteristics cause notable mixing. This allows the use of simple advection to be a good approximation when performing neutronics calculations (e.g., with Moltres).
- For large simulation problems, it is often useful to benchmark the computing resource required under various conditions before performing the actual simulation. E.g., two 0.04s timesteps computes faster than a single 0.08s timestep in the ¼-core Moltres simulations.

4.6.3. Conclusions

In this work, a method is developed to simulate the neutronics behavior of the TAP MSR under load-following conditions. Based on the simulations, it is observed that power overshoots are expected during rapid up ramp load-following operations. While this did not result in the exceeding of material temperature constraints under the improved control rod deployment, it would be prudent to minimize the stress to the reactor components during up ramps, such as by ramping up to approximately 80% of the rated reactor power before completing the power ramp at a lower rate.

The temperature safety of the TAP MSR is also studied by simulating its thermal-hydraulics behavior using the $1/8^{th}$ core model and the unit cell model. While the $1/8^{th}$ core model predicted a laminar-like flow, which led to a maximum local temperature in excess of 1400 °C, the unit cell model was able to simulate the fuel salt velocity to a much higher fidelity and predict the presence of an unstable recirculation region. This unstable flow disrupts the continuous buildup of heat, leading to a lower maximum local temperature of 783 °C under a conservative volumetric heat generation rate. With the inclusion of turbulence effects expected in reality, this maximum local temperature is expected to decrease even further.

In the process of performing the thermal-hydraulics study in this work, an improvement in the control rods deployment is also made, which led to a greater uniformity of the neutron flux and heat generation profile. This provided an important improvement to the thermal safety of the reactor.

Finally, an investigation on the effects of gas entrainment is performed, and it is found that gas entrainment in the fuel salt introduces an overall positive reactivity to the reactor core. In response to this, two approaches are considered to mitigate the safety issue arising from the positive void coefficient. An important point to be noted is that the reactor core's load-following response time is on the order of seconds regardless of the presence of gas entrainment or ^{135}Xe buildup, and that this response time is much faster than that required for power grid load-following.

In light of the above, based on the simulations performed in this study, it is concluded that the TAP MSR core is able to perform load-following operations while doing so without exceeding its thermal safety constraints.

CHAPTER 5: CONCLUSIONS AND OUTLOOKS

As one of the Gen-IV nuclear reactor concepts, the molten salt reactor receives increasing development interests in the recent years. One distinguishing feature of the liquid-fueled molten salt reactor is its improved ability to operate in a load-following mode with the unique online fission product removal system. Conceptual designs by the Transatomic Power, Terrestrial Energy, and the Flibe Energy include the ability of load-following as one of their trademarks. The project team is awarded to design a fuel processing system that enables liquid-fueled MSRs to load follow, using the TAP MSR as the prototype. Since the development of liquid-fueled molten salt reactors globally is mostly at the level of general conceptual designs, with some initial efforts to build a 2MWt testing reactor TMSR-LF1, the efforts towards designing the fuel processing system are concentrated on the experiments, modeling, tool development that would facilitate the fuel processing system development. Still, general conclusions and recommendations of the processing system are given through careful analysis using the available data in the literature which are summarized below.

During the project period, four interconnected aspects of the development of the fission product removal system in liquid-fueled molten salt reactors are investigated. These aspects are basic properties of dissolved fission product and fission product removal experiment, the simulation and analysis of the fission product removal system, the fuel cycle simulation, and the coupled reactor neutronics and thermal hydraulics transient simulation.

In Chapter 2, simulations and analyses of a prototypical xenon removal system initially considered in the MSBR are performed. Engineering models that predict the system performance are constructed based on the CFD simulations. These models provide the overall removal efficiency of the processing system that are valuable for the work described in Chapter 3. The constitutive models used in the CFD simulations are carefully inspected with validation experiments performed by the project team. The validated CFD models are used in the simulation of a prototypical xenon removal system. Comparisons are made between the simulations and existing experimental data of the processing system, where good agreements are found for the system level parameters.

In Chapter 3, fuel cycle simulations and analyses are carried out using SaltProc developed by the project team and Serpent 2. Lifelong depletion simulations of TAP MSR with ideal and realistic fission product removal are performed. The processing system is carefully reviewed, and the corresponding models are built into SaltProc for general use of the code. This depletion simulation is compared with existing fuel cycle simulation of TAP MSR for validation of the tools and methods developed in this project. After the validation, load-following simulations of TAP MSR and MSBR are studied, focusing on the xenon poisoning effect. From the results, it is concluded that the xenon poisoning effect in the TAP MSR design will not affect the load-following operation due to its hard neutron spectrum. For the MSBR, which is a thermal reactor, the removal of xenon is essential if the power should be quickly changed.

In Chapter 4, coupled thermal hydraulic and core neutronics simulations are performed for the TAP MSR. The thermal limit of the reactor core is studied under various condition, especially for quickly increasing to full power from low power conditions. The simulation considers the actual movement of the control rod through dynamic meshing. Based on the current simulations and core modeling, it is concluded that under designed flowrate, the power ramping will not exceed

the temperature limit of the core. Also, the power of the TAP MSR could be changed within a minute from 1% to almost 100%, which is well above the target power change rate for load-following operation. Moreover, it is observed that power overshoots are expected during rapid up ramp load-following operations, which may require a staged power ramping strategy for smooth and safe operation.

From the findings discussed in each chapter, it is shown that removing the fission product continuously from the liquid-fueled thermal MSR is important to enable load-following operation. Based on the fuel cycle analysis and the coupled reactor core simulation, the safety limits are generally not exceeded during the operation, except that the control rod design of the MSBR is not sufficient at the end of the fuel life. Moreover, in Chapter 3, the batch wise fuel cycle simulation is performed with prescribed removal efficiency. Based on the findings described in Chapter 2, the required removal efficiency could be achieved with a reasonable sized removal system.

The solubility and diffusivity of xenon in molten salt used in the current work is of large uncertainty, whereas an economical sizing design of the removal system is highly dependent on the diffusivity. On the other hand, the solubility only has limited effect for an active processing system using inert gas sparging. The available experimental data of the removal system, especially for the bubble separator, are still limited. The engineering model developed in Chapter 2 still has large uncertainties, partly due to the limited validation sources.

The fuel cycle analysis uses batch-wise approach with a much larger timestep compared with the “cycle time”. Though some discussions are given in the chapter conclusion on how to relate the batch-wise simulation and the corresponding removal efficiency to continuous operation, some temporal characteristics of the core are still not captured. Moreover, the models built based on the work of Chapter 2 have not been incorporated into the code for realistic operation.

The coupled core simulation evaluated the temperature constraint of the core, which is not exceeded during the load-following simulation. However, other thermal constraints, for example the resulting thermal stress from the fast power ramping, are not considered. Moreover, not all the effects that prevent the load-following operation, discussed in Section 3.4.1, are considered, both in the fuel cycle analysis and the coupled core simulation.

With these limitations in mind, the general conclusion that removal of xenon is beneficial to the load-following operation of a MSR could still be drawn. Though these issues need to be addressed before the design of the processing system proceed to the commercialized stage.

APPENDIX A: EXPERIMENTAL INVESTIGATION OF FISSION GAS REMOVAL

A.1. Introduction

The experimental component of this project attempted the development of effective means of removing ^{135}Xe from the molten salt fuel. The approach was to use bubbles of inert gas to purge the dissolved ^{135}Xe and other inert and volatile gases from the molten salt. This process has been studied in the molten salt reactor experiment (MSRE) program where it was shown to be an important mechanism for controlling reactivity due to ^{135}Xe removal. In that system, fine distributions of bubbles were injected as a natural consequence of the molten salt pumping process. The bubble fraction was controlled by changing the pump speed and by controlling the cover gas pressure. The MSRE used either He or Ar during various phases of their experiments as a cover gas, and this cover gas was mixed into the molten salt to form a fine, distributed bubble structure.

The effectiveness of this approach for purging ^{135}Xe is due to the fact that Xe is highly insoluble in molten salt. The availability of the bubble distribution aided the release of the Xe to the cover gas plenum volume and reduced the amounts of Xe entrained in the molten salt and returned to the reactor core. Some differences were found in the behavior of Ar versus He through the experiments with each did not precisely cover the same sets of experimental conditions. One important consideration is the He is moderately soluble in molten salt and Ar is not. Thus, experiments with Ar showed a monotonic behavior with gas bubble size, whereas the experiments with He did not.

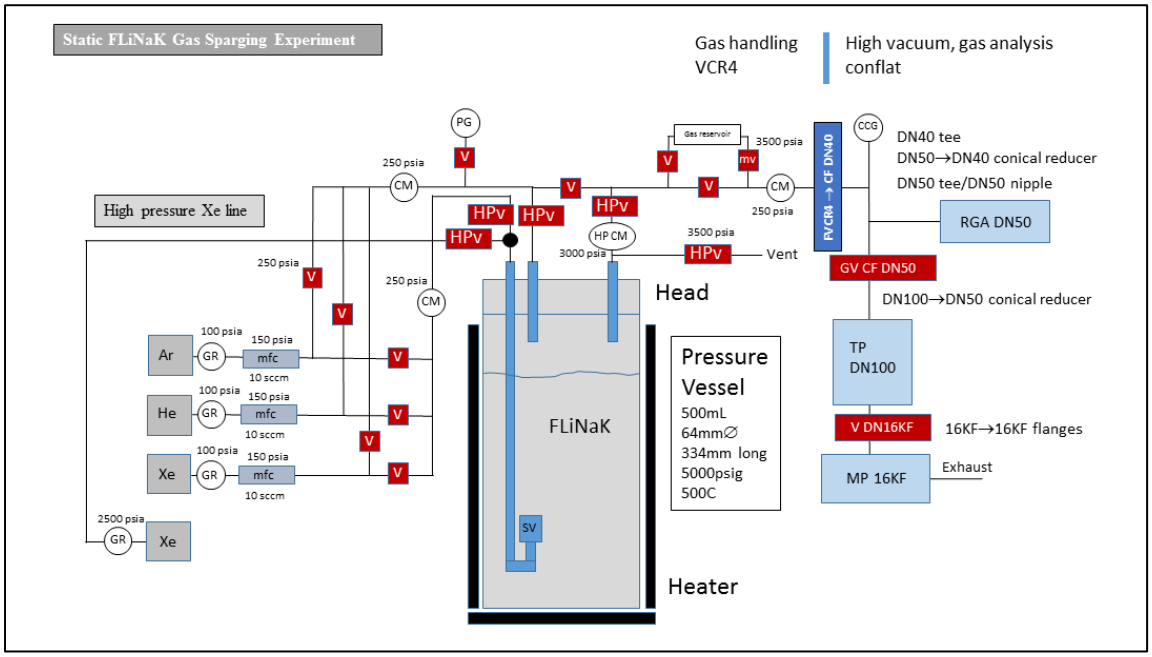

Figure A1: Schematic diagram of the UIUC static molten salt gas sparging experiment. The main components are the pressure vessel, residual gas analyzer, gas handling manifold, and the pumping station.

A.2. Gas Sparging Experiment Design and Construction, FLiNaK Supply.

Gas sparging experiment was designed and constructed. A schematic diagram is shown in Figure A1. Photographs of the sparging experiment are shown in Figure A2.

The design of the gas sparging experiment used an Inconel pressure vessel to heat the molten salt, provide controlled gas follow, and perform static pressure measurements. The quantification of gas sparging was based on a residual gas analyzer (RGA) that would measure the amount gas species released in absolute (moles) quantities.

Figure A2: Photographs of the UIUC static molten salt gas sparging experiment. From left to right: The gas handling manifold; the pressure vessel and heater; the pump station and RGA.

Two batches of eutectic FLiNaK (LiF-NaF-KF::46.5-11.5-42.0 mol%) were provided to UIUC by ANL (10 kg) and ORNL (2 kg). ANL performed a comparison analysis of these two salt batches [ANL/CFCT-20/11]. Overall, the ORNL batch was deemed more pure based on X-ray diffraction (XRD) analysis and ICP-MS, with fewer trace impurities. The most significant impurity in the ORNL FLiNaK was calcium, while Fe and Ni were the primary trace impurities in the ANL FLiNaK. X-ray diffraction analysis of the ORNL material was clean, with all reflections indexed with the three fluoride compounds that comprised the salt. On the other hand, the XRD analysis of the ANL salt exhibited reflections that could not be attributed to fluoride compounds. The melting temperature was determined using differential scanning calorimetry and the two salt batches behaved similarly: ANL T_{melt}= 455.4 ºC, ORNL T_{melt}= 457.2 ºC. The compositions of the two salt batches are shown in Table A1.

Table A1: Comparison of FLiNaK salt compositions from ANL/CFCT-20/11.

	Salt Composition (mol %)		
	LiF	NaF	KF
Eutectic - FLiNaK	46.5	11.5	42
ORNL - FLiNaK	44.7	11.2	44.1
ANL - FLiNaK	47.1	11.1	41.8

A.3. Gas Sparging Experiment Calibration.

The gas sparging experiment required calibration of the RGA. The measurement of moles of gas species (Kr in this case) is obtained from the integrated peak area correlated to the known Kr moles. The calibration procedure used a static gas mixture of Kr plus He, with the pumping station evacuating this mixture across the RGA.

The calibration curve for Kr plus He is shown in Figure A3. One important aspect of the calibration is that it is independent of He pressure; this is demonstrated in Figure A3. An example of the raw Kr RGA signal data versus time and RGA GUI showing the resolution of the Kr isotopes are also shown in this figure.

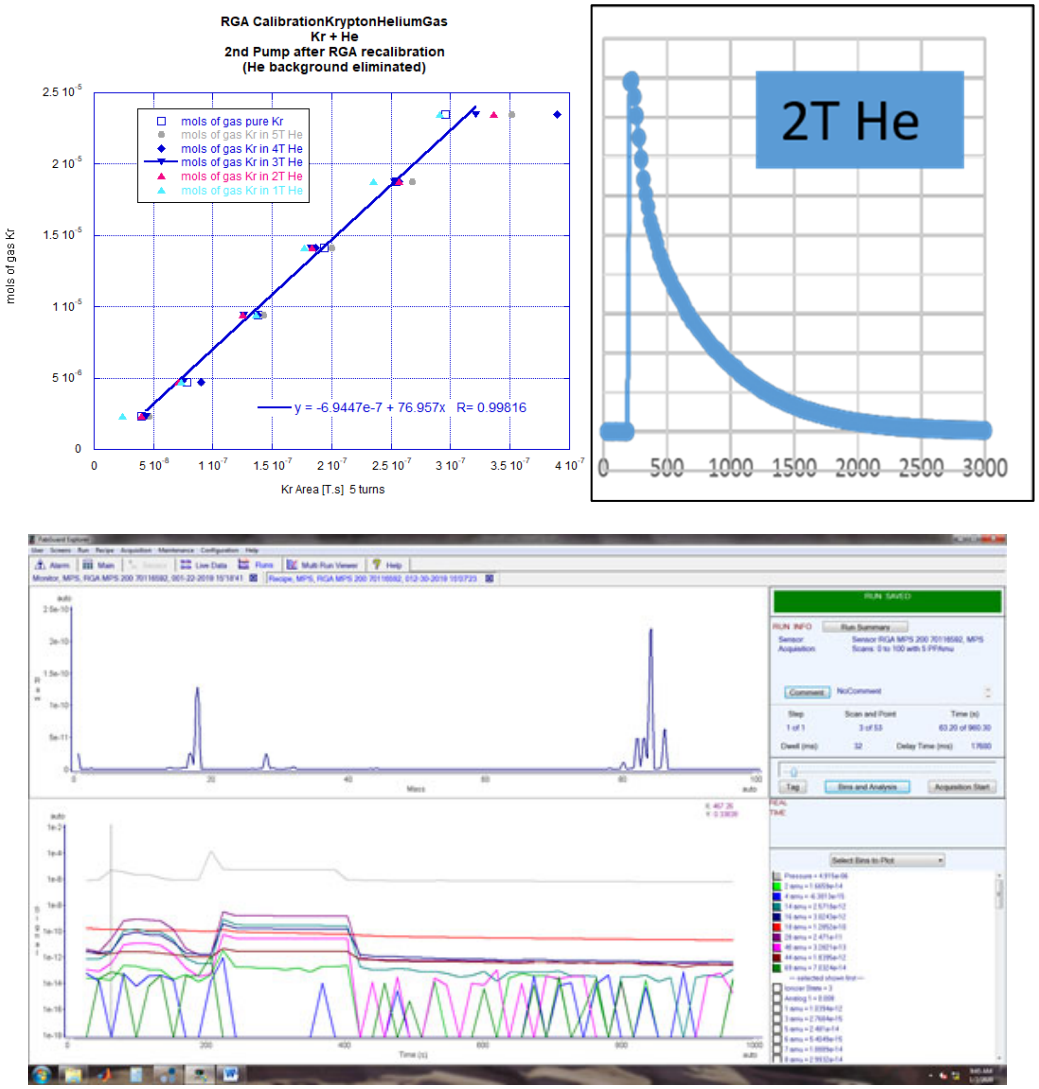

Figure A3: Upper left: Calibration curve for Kr gas sparged with He. The x axis is the integrated RGA peak area and the y axis is the moles of Kr gas. The best-fit linear line represents the RGA calibration curve. Upper right: Raw data for the Kr signal from the RGA versus time (x axis in units of seconds) for 2 Torr of He. Bottom: RGA GUI showing resolution of the Kr isotopes.

A.4. Gas Sparging Experimental Data.

The gas sparging experiment was not successful. Several problems arose that are detailed in the next section, *Lessons Learned*. Figure A4 shows an example of the RGA calibration. These are Kr release data versus time for 500 °C FLiNaK and the total integrated peak area has been converted to moles Kr using the linear relationship shown in Figure A3. The Kr release data shown in Figure A4 are not typical. Data associated with Kr release from molten FLiNaK are presented below.

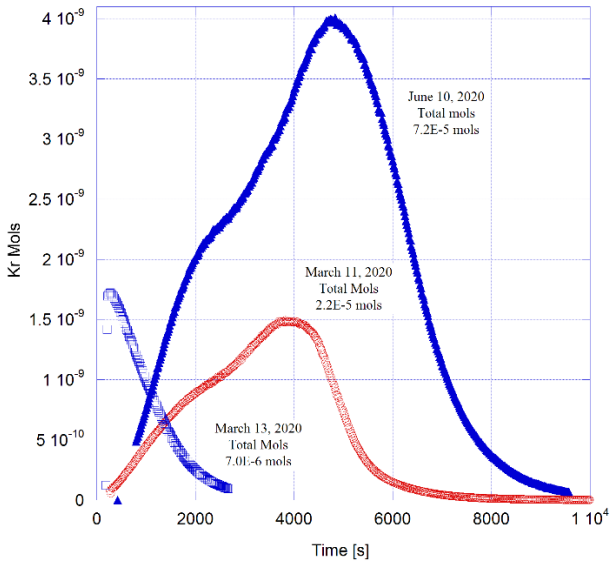

Figure A4: Krypton release data using helium gas flow. The calibration curve in Figure A3 was used to determine total moles.

Figure A5 shows a demonstration of sparging and the effect of overpressure on Kr release for 500 °C FLiNaK (eutectic ANL batch). The left-hand figure shows the Kr-84 mass signal versus time with and without He sparging. An increase in Kr removal is seen with He sparging. The sparging gas flow was low, 0.2 SCCM. These data are fit with the sum of two exponential decays and the decay constants are consistent with greater Kr removal during He sparging (larger m2 and m4 absolute values). The right-hand figure shows the Kr-84 mass signal versus time for two Kr cover gas overpressures. A greater release signal is observed for the higher overpressure, an effect due to greater gas solubility with greater overpressure. Note that the difference in ordinate scales between Figures A5 is due to normalization. Normalization does not affect the fitted time constants.

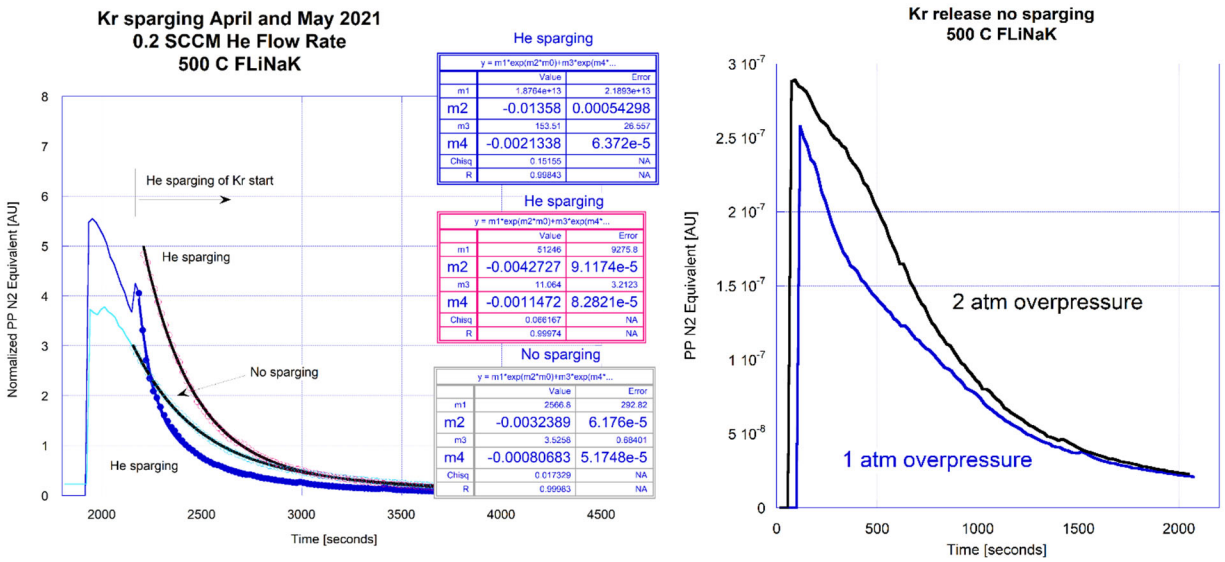

Figure A5: (left) Three sets of Kr-84 mass release data are shown, two with He sparging as labeled and one without sparging. Approximately 200 s into the initial release He sparging at 0.2 SCCM was started, as indicated. All data sets are fit with the sum of two exponential decays. Fitting constants m2 and m4 are the exponential time constants. (right) Effect of Kr cover gas overpressure on measured Kr release without gas sparging. Greater peak area is observed at greater overpressure.

The exponential time constants obtained from the Figure A5 fitting are tabulated in Table A2. Time constants with greater absolute magnitude (more negative) represent faster kinetics. The time constants for two He sparging measurements are greater than that for the no sparging case. As explained in the next section, these data represent Kr release via diffusion from the molten FLiNaK; the sparging is actually gas flow across the cover gas space, not through the sparging tube.

Table A2: Fitted exponential time constants for Kr release from 500 °C molten FLiNaK.

	m2 [1/s]	m4 [1/s]
sparging 1	-0.0136	-0.00213
sparging 2	-0.0043	-0.00115
no sparging	-0.0032	-0.00081

A.5. Lessons Learned from the Experimental Work Scope

Several technical issues related to the behavior of molten FLiNaK and the static gas sparging experiment were encountered during the course of the experimental work scope for this project. Overall, molten FLiNaK is challenging to work with when proper handling facilities are not available. Specifically, a large dedicated glove box is required for remote handling and *in situ* experiments (our experiment is *ex situ* in that it is on a bench top). *This is the most important*

lesson learned: The use of a dedicated glove box is critical for successful experimental work using molten salts.

The technical issue encountered are described below and represent lessons learned:

1. The important issue encountered early in the work scope was the inability to seal the pressure vessel containing FLiNaK properly in the glove box. This required the FLiNaK to be melted and then the pressure vessel cover to be mounted and tightened remotely, within our glove box. It was not possible to tighten sufficiently and this led to improper seating of the cover. The pressure vessel head cover was machined to increase the clearance and this allowed for proper seating. Consideration was given to melting the FLiNaK in a secondary vessel and using pressurized gas to transport the molten salt into the pressure vessel. This would have required heating the transfer line in excess of the melting temperature. Further, the molten salt could have solidified at cold spots along the transfer line or within the fixture into the pressure vessel. Given the eventual ability to seal the pressure vessel we did not pursue remote transfer of the molten salt. *Lesson learned: Remote transfer of the molten FLiNaK would have led to greater experimental success*.
2. In addition to improper seating of the pressure vessel head cover, the thermocouple well was designed to use a Ni metal gasket for sealing. This proved to leak and we eventually welded the thermocouple well to the head cover.
3. Another significant issue was the attack of the FLiNaK on the sparging tube. This caused several tubes to unseal and/or catastrophically fail. The consequence of this was that the sparging experimental data presented above represent the effect of sweeping out the cover gas with He and the concomitant effect of Kr release via diffusional transport from the molten FLiNaK. Eventually, we used a sparging tube with greater wall thickness.
4. The initial sparging tube was designed to screw into the pressure vessel head cover and tightened with a nut. This designed was not air tight and led to a gas pathway into the cover gas space, bypassing the sparging tube. The final iteration used a sparging tube welded to the pressure vessel head cover. *Lesson learned: Welding all attachments to the pressure vessel head cover would have led to greater experimental success.*
5. Even with better pressure vessel head seating, a welded thermal well, and a welded sparging tube, the pressure vessel was not air tight, especially above the FLiNaK melting temperature. This was clear from the RGA signal, which indicated significant N_2 and O_2 above 460 °C. These two gas signals were not observed at ambient temperature. The consequence of this leakage was contamination of the molten FLiNaK. *Lesson learned: The use of a large glove box with a thermal well would prevent air contamination*.
6. Loading and removing FLiNaK from the pressure vessel was very cumbersome since a glove box was required to prevent exposure of the solid FLiNaK to air. The contamination of the molten FLiNaK appeared to facilitate a chemical reaction of the FLiNaK with the Inconel pressure vessel. This made removal of the solid FLiNaK extremely difficult. *Lesson learned: The use of a large glove box with a thermal well would have enabled us to avoid contamination and likely reduced the number of times the FLiNaK needed to be replaced.*

7. The vapor pressure of FLiNaK, while low, resulted in the coating of salt on many internal surfaces, some remote from the pressure vessel. *Lesson learned: This seems to be an unavoidable consequence of working with molten salt*.

At the time of submission of this report we are still performing FLiNaK gas sparging experiments. The goal is to obtain valid gas sparging data suitable for publication.

APPENDIX B: PHYSICAL AND CHEMICAL ANALYSIS OF THE FLINAK RECEIVED FOR THIS PROJECT

This appendix includes the property measurements of the FLiNaK salt used in our experiment, which is made by Argonne National Laboratory. For a complete discussion on the procedures, one can refer to the original reports [21, 22]. In addition, ORNL also provides two batch of FLiNaK, which is compared to the salt from ANL by X-ray diffraction by the ANL resource team, which is included in APPENDIX E.

B.1. Density of FLiNaK

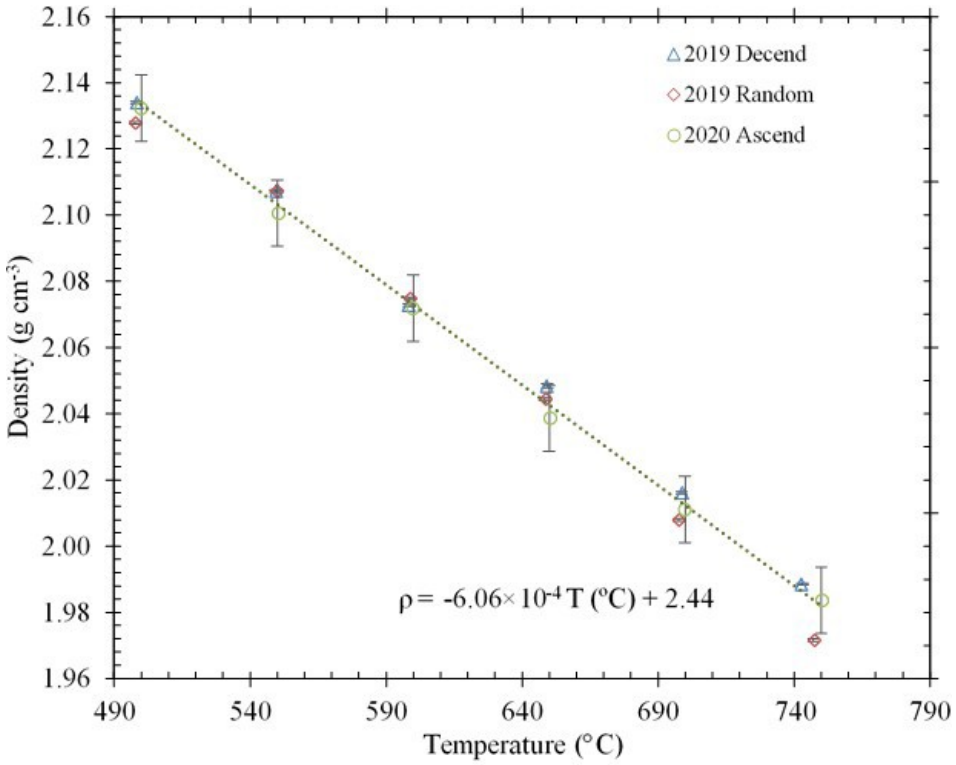

Figure B1: Measured Density of FliNaK at Argonne at Different Temperature Orders, taken from Rose et al. [21].

The density measurement is made by using the displacement method based on Archimedes principle, shown in Figure B1. The density is calculated by immersing a small bob with known mass into the molten salt and measuring the weight through the connected wire and a balance. Surface tension effect is also considered. In Figure B2, the results are compared with data in literature.

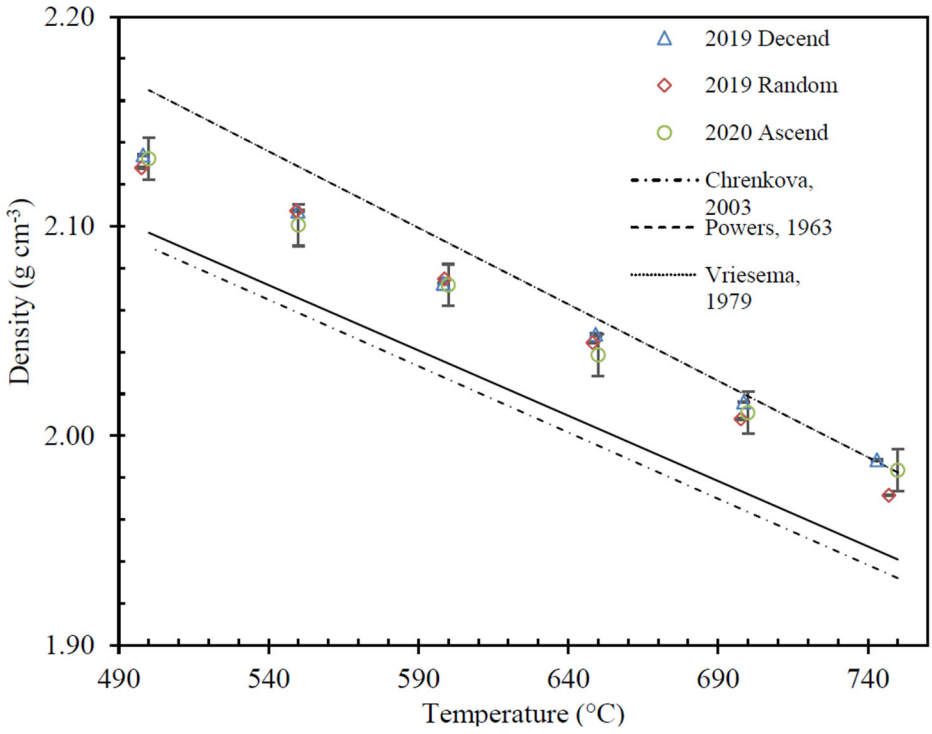

Figure B2: Measured Density of FliNaK Compared with Literature Values, taken from Rose et al. [21].

B.2. Viscosity of FliNaK

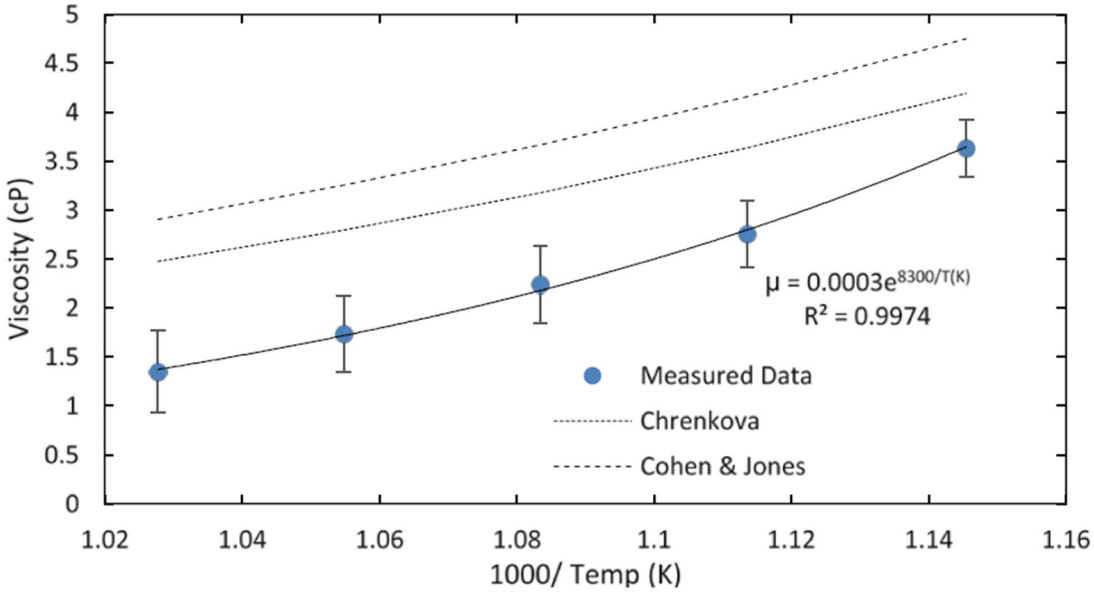

Figure B3: Measured Viscosity of FliNaK at Different Spindle Speeds at 650 °C, taken from Rose et al. [21].

The viscosity of FliNaK is measured by a rotating cylinder method, with customized spindle on a Brookfield viscometer which provides measurement of torque and rotational speed. The results are shown in Figure B3 and compared with literature data.

B.3. Thermal Diffusivity of FliNaK

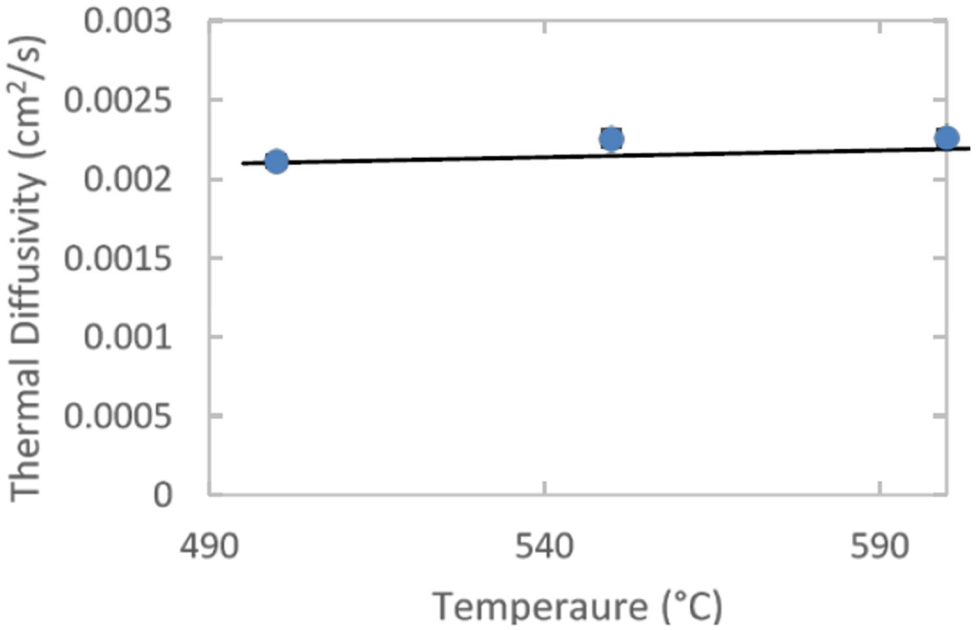

Figure B4: Measured Thermal Diffusivity of FliNaK, taken from Rose et al. [21].

The thermal diffusivity is measured with a laser flash analyzer, DLF1200 from TA Instruments.

B.4. Heat Capacity of FliNaK

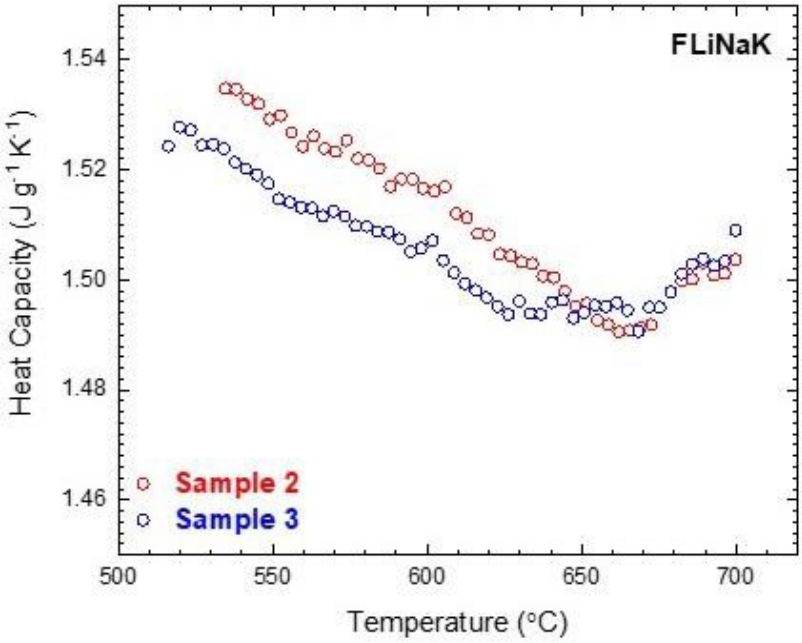

Figure B5: Calculated Heat Capacity of FliNaK For Samples S2 and S3, taken from Lichtenstein et al. [22].

The heat capacity of FliNaK is measured through differential scanning calorimetry. The initial decrease of the heat capacity is not expected and is argued to be related to the eutectic phase transition occurring during the measurement.

APPENDIX C: UNCERTAINTY OF THE VALIDATION EXPERIMENTS FOR GAS SPARGING

C.1. Uncertainty for Visualization Experiment

The void fraction in the experiment is derived by estimating the volume of each bubble. The possible uncertainty sources are: 1) uncertainty from the gas flow meter, 2) the error in the bubble identification process, and 3) the error from the volume estimation method.

1) Error from the gas flow meter

The quantity being measured is the void fraction at a point, and the averaged void fraction. Since the void fraction is relatively small, a linear dependence of gas flow rate could be assumed. This assumption is reasonable, judging from the experiment result. Therefore, the void fraction could be written as the flowrate, q, times some other function. The uncertainty is therefore additive. This is also true for the other three uncertainty sources.

$$\alpha = qf, \qquad \frac{d\alpha}{\alpha} = \sqrt{\left(\frac{dq}{q}\right)^2 + \left(\frac{df}{f}\right)^2}.$$

The uncertainty of the gas flow meter is 5%.

2) Error in the bubble identification process

In obtaining the individual bubbles, the images are processed. Some bubbles are not fully recognized and lose a part of their shape. Sometimes, two bubbles are recognized as one, since they are connected. By looking at each frame, roughly speaking, around 10% bubbles lost a small fraction, and around 10% bubbles become connected. Since this error is highly related to the error in the volume estimation, a rigorous evaluation is not easily performed. Instead, this error is considered an error to the correct bubble volume. In this way, the error for bubble identification could be expressed as

$$\frac{V_{ob}}{V_{ext}} = \left(\frac{V_{ob}}{V_{ap}}\right)\left(\frac{V_{ap}}{V_{ext}}\right), \qquad \frac{dV_{ob}}{V_{ext}} = \sqrt{\left(\frac{dV_{ob}}{V_{ap}}\right)^2 + \left(\frac{dV_{ap}}{V_{ext}}\right)^2}.$$

The relative error shown here, after multiplied with its probability of occurrence, equals to $\frac{df}{f}$. For most cases, the bubbles lost 20% of its cross-sectional area. Assuming the bubble is spherical, this means

$$r_{ap} = \sqrt{\frac{0.8A}{\pi}}, \qquad V_{est} = \frac{4}{3}\pi r_{ext}^3, \qquad V_{ap} = \frac{4}{3}\pi r_{ext}^3 \times (0.8)^{\frac{3}{2}}, \qquad \frac{V_{ap}}{V_{est}} = 0.72.$$

Since only 10% of bubble has this problem, the final uncertainty should be 2.8%. Considering the contribution from worst cases, this uncertainty is increased to 5%.

One the other hand, for the bubbles in contact with each other, the overlapped cross section is around 40% of their cross-sectional area. For these bubbles, assuming they are identical, the following correlations are found,

$$r_{ap} = \sqrt{1.6A/\pi}\,, \qquad V_{est} = \frac{8}{3}\pi r_{ext}^3, \qquad V_{ap} = \frac{4}{3}\pi r_{ext}^3 \times (1.6)^{\frac{3}{2}}, \qquad \frac{V_{ap}}{V_{est}} = \frac{(1.6)^{\frac{3}{2}}}{2} \approx 1.$$

This means, for most cases, the overlapping does result in significant uncertainty. For the extreme cases, where there is no overlapping or completed overlapped, the error becomes

$$\frac{V_{ap}}{V_{est}} = \frac{(2)^{\frac{3}{2}}}{2} = 1.41, \qquad \frac{V_{ap}}{V_{est}} = 0.5.$$

In a way, the two errors are balancing each other. Since only the time averaged measurement is used, this error will be much smaller than the error from losing a part of the bubble shape.

3) Error in estimating the bubble volume.

This error comes from the fact that bubbles are not perfect spheres, but they are still assumed to be sphere in the volume estimation method. The error related with this error source could be estimated by considering the typical distorted bubble shape. The cap and elliptical bubbles with prescribed dimensions are analyzed with this method.

To begin with, for a random oriented surface element in space, the projected area onto a plane is proportional to its surface area as

$$< A_{proj} > = \frac{1}{4} < A_{surf} >.$$

Considering an ellipsoid with $a = 1, b = 2, c = 4$, which is already a highly distorted bubble, the derivation goes as

$$A_{surf} \approx \left[\frac{(ab)^{1.6} + (bc)^{1.6} + (ac)^{1.6}}{3}\right]^{\frac{1}{1.6}} 4\pi, \qquad < V_{est} > = \frac{4}{3}\pi \left[\frac{< A_{surf} >}{4\pi}\right]^{\frac{3}{2}},$$

$$V_{el} = \frac{4}{3}\pi abc, \qquad \frac{V}{< V_{est} >} = 0.704.$$

which means even for highly distorted bubbles, the overestimation of our measurement is still within an acceptable range. The error associated with cap bubbles could be analyzed with a similar method, which could lead to an underestimation of the bubble volume around the same level of uncertainty. The different errors for cap bubble and ellipse bubble compensate with each other, which would reduce the overall error of our measurement. Since the measurement is averaged in time, and the occurrence of cap shape and ellipsoid bubbles are similar, the error should be greatly reduced. To be conservative, assuming a gaussian distribution for the error, and using the 3σ principle to reduce the maximum error by 3. This is not rigorously defined but meant to give an upper limit for this error source.

Combining the error from the three processes, a conservative estimation of uncertainty in measured void fraction is given as,

$$\frac{d\alpha}{\alpha} = \sqrt{5\%^2 + 5\%^2 + 10\%^2} = 12.2\%.$$

C.2. Uncertainty of the Mass Transfer Coefficient

The uncertainty of mass transfer coefficient measurement could come from the 1) gas flowrate, 2) the oxygen probe and 3) the curve fitting uncertainty.

1) Error from the gas flow meter is described previously and results in 5% uncertainty.

2) Error from the oxygen probe: The oxygen probe has an accuracy of 1.5% based on manufacturer specifications.

3) The uncertainty from curve fitting: The uncertainty from curve fitting is negligible compared with the other two error sources, on the order of 0.1%.

The combined uncertainty for mass transfer coefficient measurement would be

$$\frac{dk}{k} = \sqrt{5\%^2 + 1.5\%^2} = 5.2\%.$$

C.3. Uncertainty of the PIV Experiment

The uncertainty of PIV measurements, particularly in two-phase flow, remains an open research area. Reliable method for determining the uncertainty for instantaneous measurement is still under development. The result of a PIV experiment is obtained by calculating the cross correlation of the two frames at each interrogation area, which is hard to be directly related to other quantities [203]. Instead, an easier way of estimating the error is to relate the uncertainty of measurement to the cross-correlation plane. Using the Peak Height Ratio method introduced by Charonko et al. [204], instantaneous errors at each interrogation area is obtained. This uncertainty should be interpreted as the uncertainty of measurement, rather than systematic or instrumental error, since it is commonly accepted that the PIV technique itself is bias free [203]. However, the result used in this document is the time averaged velocity, therefore these errors are averaged with the volume weighted root mean square as

$$\bar{\epsilon} = \frac{\sqrt{\sum(u_i \epsilon_i)^2}}{\sqrt{\sum u_i^2}}.$$

The distribution of uncertainty at different points is shown in Figure A.1. From the result, an uncertainty of 0.1 is reasonable for the instantaneous measurement. The uncertainty above 0.1 is mostly located in the region close to the inlet, which was already reported as low confidence. With repeating measurement of the same quantity, the uncertainty will be reduced by a factor of $\sqrt{n}$, where n is the number of measurements made. The problem here is the time averaged profile is not obtained with one frame pair, but thousands of frames, and the meaning of the number of measurements n is obscured. In Figure A.2, the time averaged axial velocity at 6.9 cm above the inlet is shown. It appears 2000 frame pairs are enough to obtain the time averaged velocity, which is 400s in time. Since over 4000 frame pairs is recorded in the experiment, it could be argued that two measurements were carried out in the process. Assuming the uncertainty of 0.1 is carried along from the instantaneous measurement to the time averaged measurement, an uncertainty of $\frac{10}{\sqrt{2}}\% \approx$

7% is estimated. From the local maximum and minimum of the measurement, with a variation of 5%, the uncertainty of 7% seems to be a good estimation.

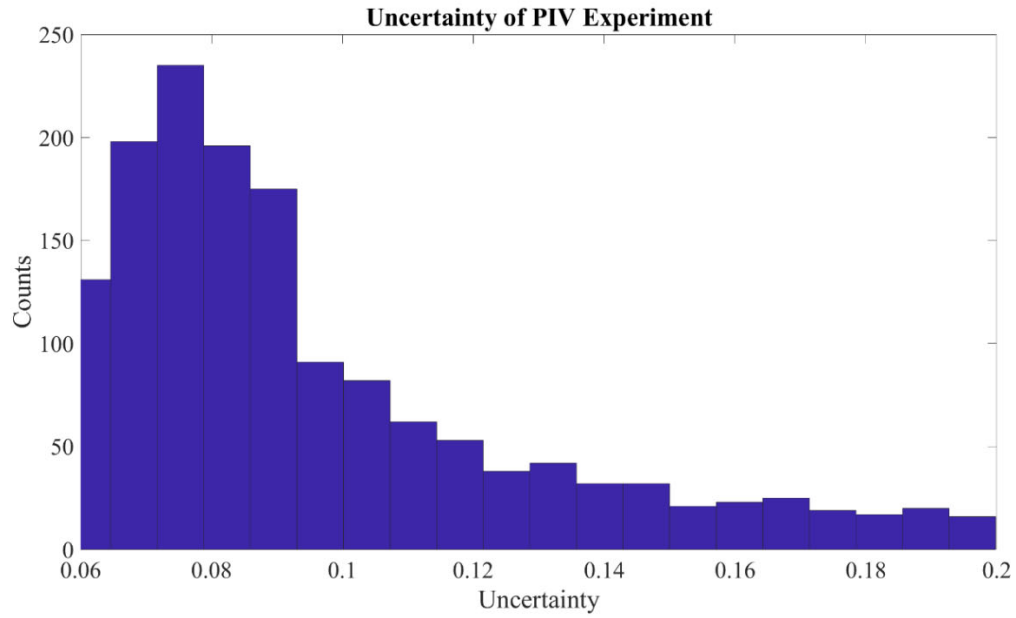

Figure C1: Averaged Uncertainty.

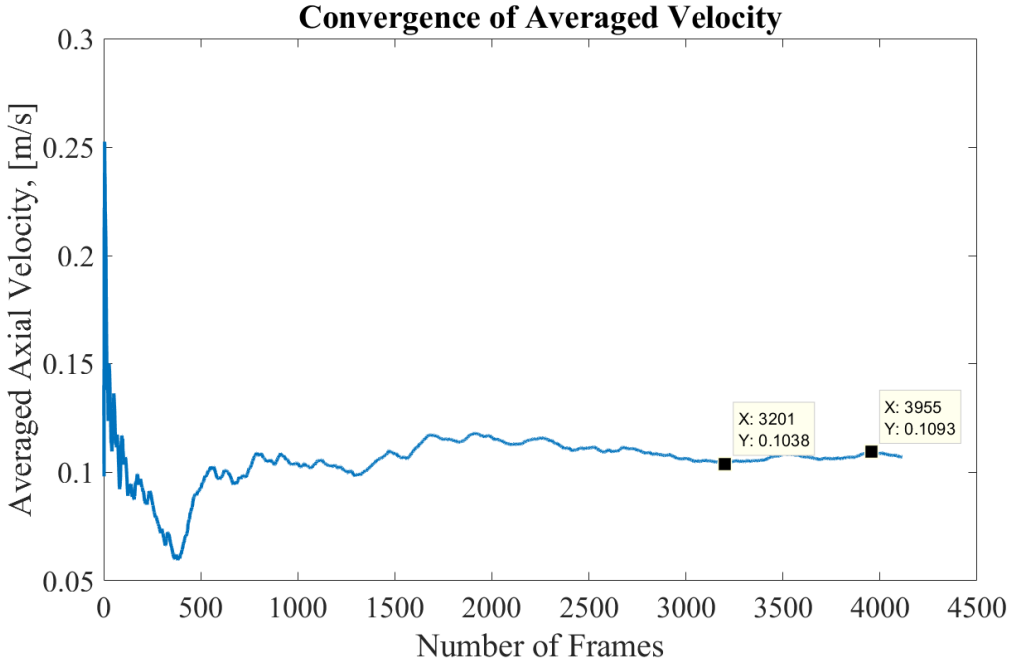

Figure C2: Convergence of Averaged Axial Velocity at axial location of $6.9cm$.

APPENDIX D: RECONFIGURABLE MODERATOR IN TAP CORE

Table D.1: Geometric details for the full-core 3D model of the TAP with various moderator rod assembly configurations.

Case	Number of $ZrH_{1.66}$ rods in the quarter core	SVF	Moderator-to fuel ratio
1 (BOL)	347	0.917204	0.09027
2	406	0.903126	0.10727
3	427	0.898115	0.11344
4	505	0.879503	0.137
5	576	0.862563	0.15933
6	633	0.848962	0.17791
7	681	0.837509	0.19402
8	840	0.799571	0.25067
9	880	0.790026	0.26578
10	900	0.785254	0.27347
11	988	0.764257	0.30846
12	1126	0.731329	0.36737
13	1338	0.680744	0.46898
14	1498	0.642567	0.55626
15 (EOL)	1668	0.602004	0.66112

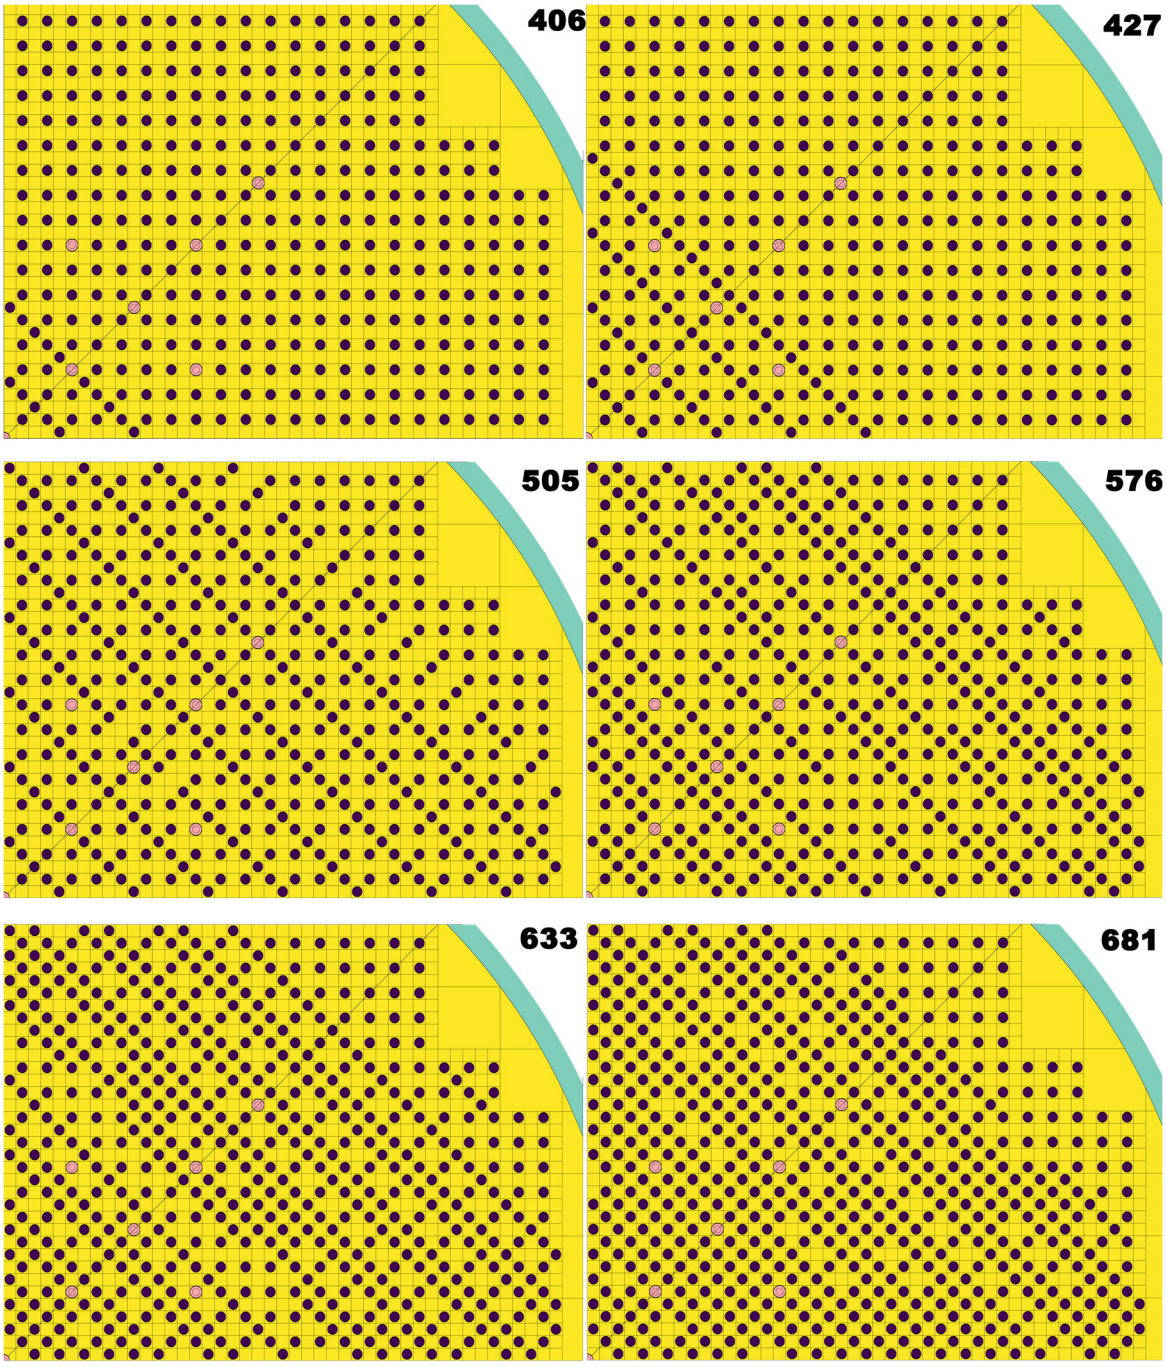

Figure D.1: An XY section of the TAP model at horizontal midplane for the first six years of operation (excluding startup moderator rods configuration) with the SVF between 0.91 and 0.84. The number in the top-right corner of each figure indicates the number of moderator rods in the case.

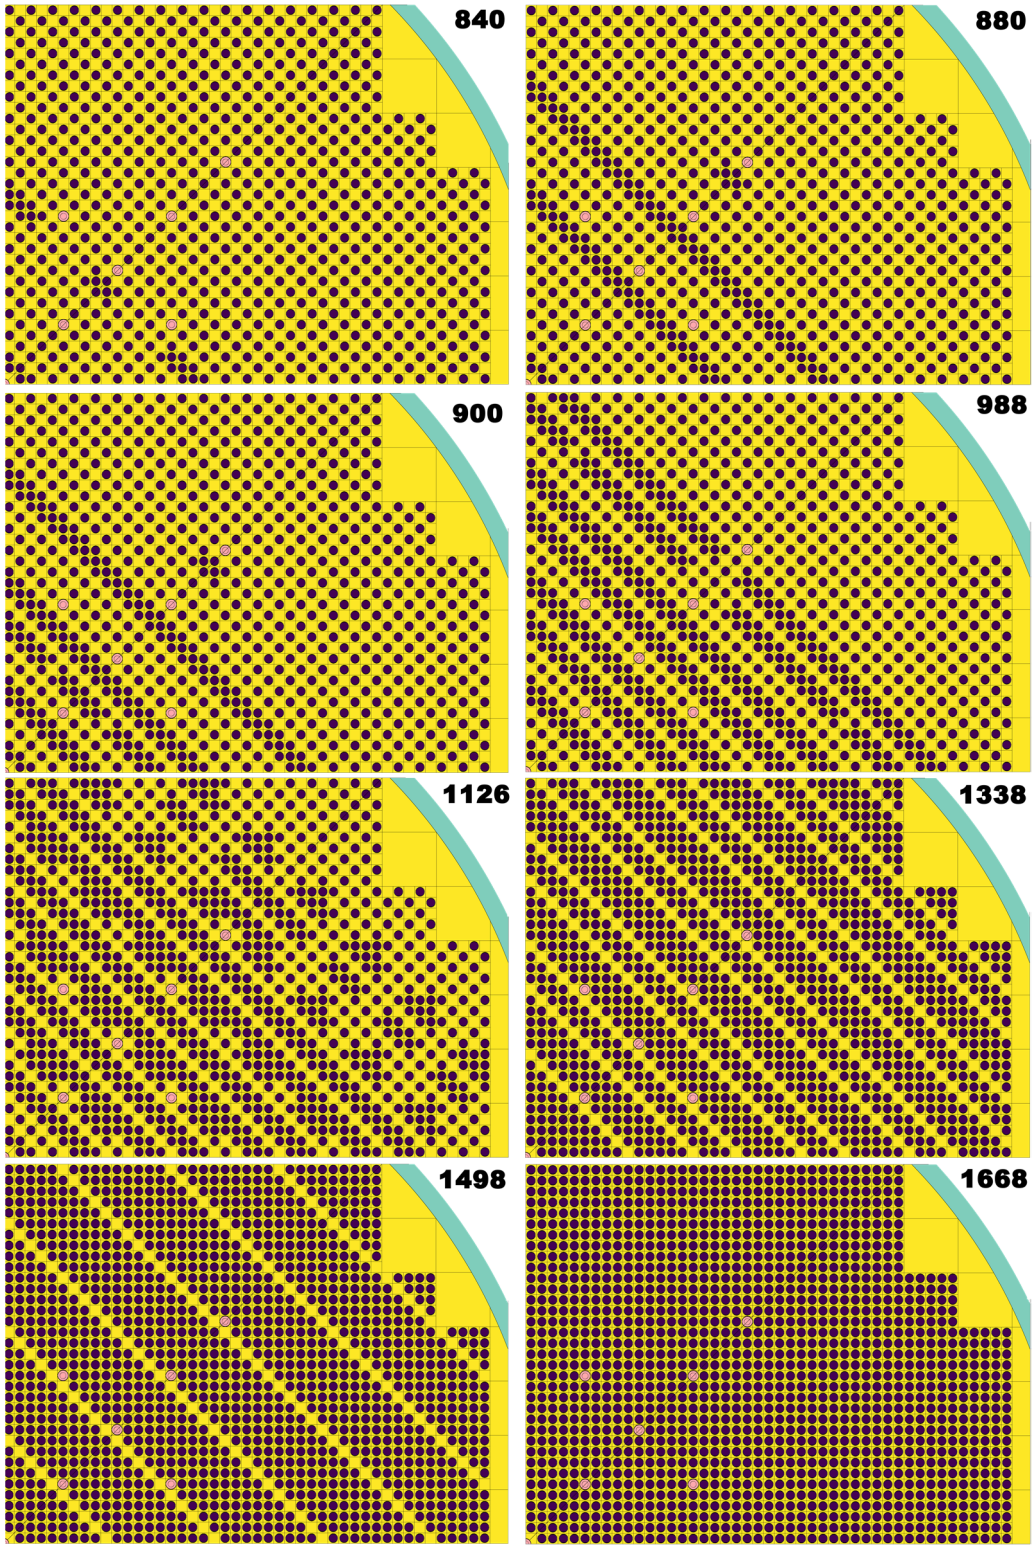

Figure D.1: An XY section of the TAP model at horizontal midplane the SVF between 0.8 and 0.6. The number in the top-right corner of each figure indicates the number of moderator rods in the case.

APPENDIX E: COMPARISON REPORT ON ANALYTICAL RESULTS FOR FLINAK

The content of this appendix is provided by the resource team at Argonne National Laboratory. The investigators are M.A. Rose, E. Wu, T. Lichtenstein, and J. Krueger.

A common fluoride coolant salt mixture LiF-NaF-KF (46.5-11.5-42.0 mol%), which is commonly referred to as FLiNaK, is being prepared at Argonne National Laboratory for use in experiments to be conducted at UIUC. University researchers will conduct sparging experiments to determine if removal of select fission gasses by sparging is feasible. Oak Ridge National Laboratory also possesses a supply of FLiNaK salt that was prepared by an outside vendor previously. Analyses were conducted at Argonne to determine the melting point and purity of both salts.

E.1. Salt Preparation and Analysis Methods

Reagent salts were purchased from commercial suppliers at 99.9% or higher purity and characterized prior to use. Material certificates obtained from manufacturers to provide information on potential impurities are provided in Appendix E1. Individual reagent salts were unpackaged inside of a glovebox maintained at below 10 ppm O_2 and 1 ppm H_2O. Reagent salts were then held at elevated temperatures below their melting points in cleaned nickel crucibles for more than 3 hours to release any adsorbed water. Samples of reagent salts were then taken for analysis. Salt reagents are analyzed to determine their elemental composition and purity by x-ray diffraction (XRD) and wet chemical methods. XRD is used to investigate oxygen or water contamination by looking for oxides, oxyfluorides/chlorides or hydrates that may have been generated during storage or handling. The results of XRD analyses of reagent LiF, NaF, KF, and UF_4 are provided in Figure E1.

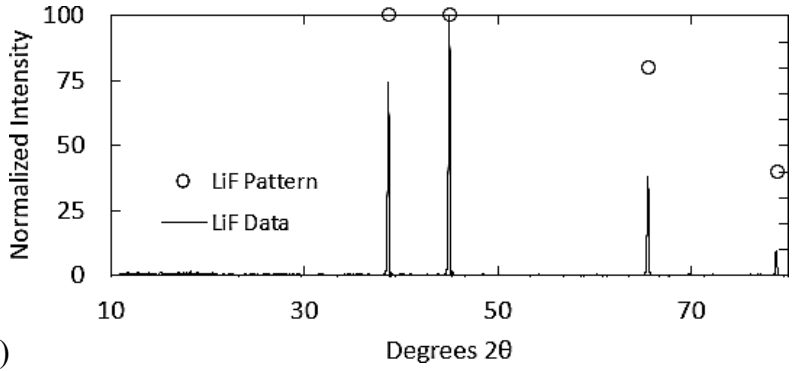

(a)

Figure E1: Measured X-ray Diffraction Patterns and Peak Intensity compared with Literature Data for Reagent Salts (a) LiF, (b) NaF, (c) KF, and (d) UF4.

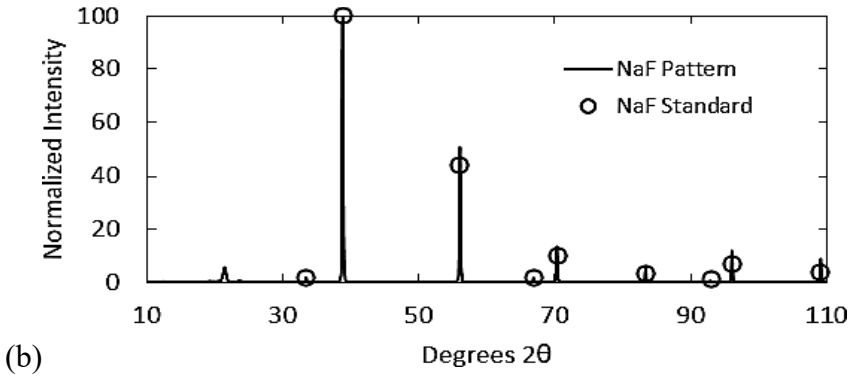

(b)

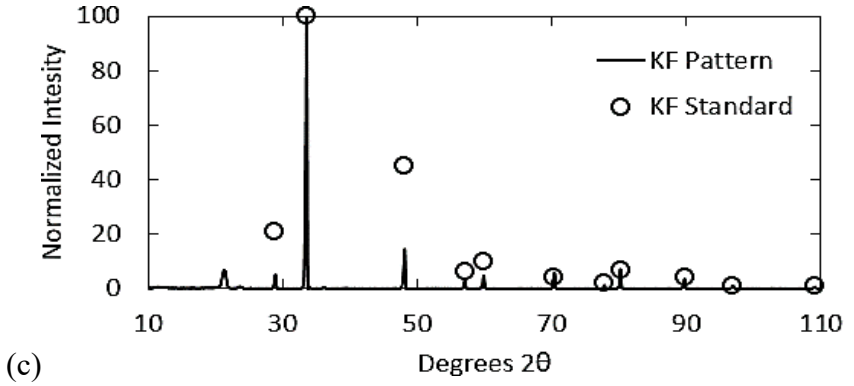

(c)

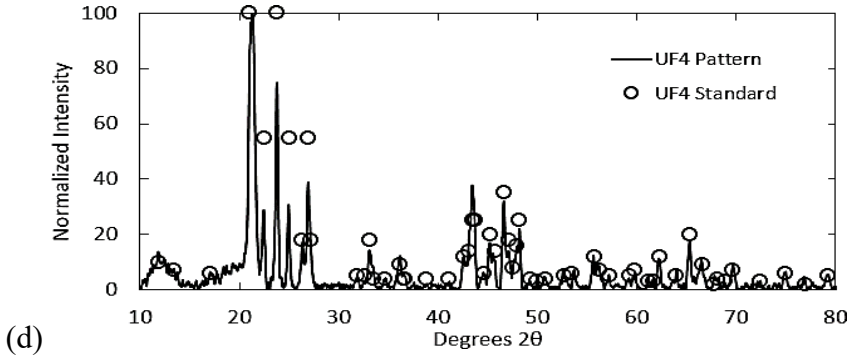

(d)

Figure E1: (cont.).

Eutectic mixtures were prepared by weighing and mixing appropriate quantities of salt reagents to create the target eutectic compositions. Each mixture was melted, cooled and ground to a fine powder. This procedure was done at least twice to promote homogeneity. Salt mixtures were analyzed to determine their elemental composition and purity by differential scanning calorimetry (DSC), XRD and wet chemical methods. DSC is very sensitive to small changes in composition and gives the best indication of how close a salt mixture is to its eutectic composition. X-ray diffraction was used to check for oxygen and water contamination that may have occurred during the preparation of the salt mixture. Wet chemical methods include ICP-MS and ICP-OES analysis.

The temperature measured by DSC is calibrated based on the reported melting points of five pure metal standards (Sn, Zn, Al, Ag, and Au) that span the range from 231.9 to 1064.2 °C; the

melting point of each metal is measured twice. Figure E2 shows the results for calibration prior to the analyses of ORNL-FLiNaK expressed as ΔT = T(measured)-T(nominal). The repeatability was excellent for duplicate measurements made with each metal, but most measured temperatures differed slightly from the reported values. The red curve shows the quadratic fit used by the PROTEUS software to adjust the measured temperatures based on the measured melting temperatures of the standards. The difference between the curve and each of the standard values represents the accuracy of measurements made at and near that temperature. In this case, the accuracy is well within 1 °C for temperatures up to about 700 °C and within about 2 °C at higher temperatures.

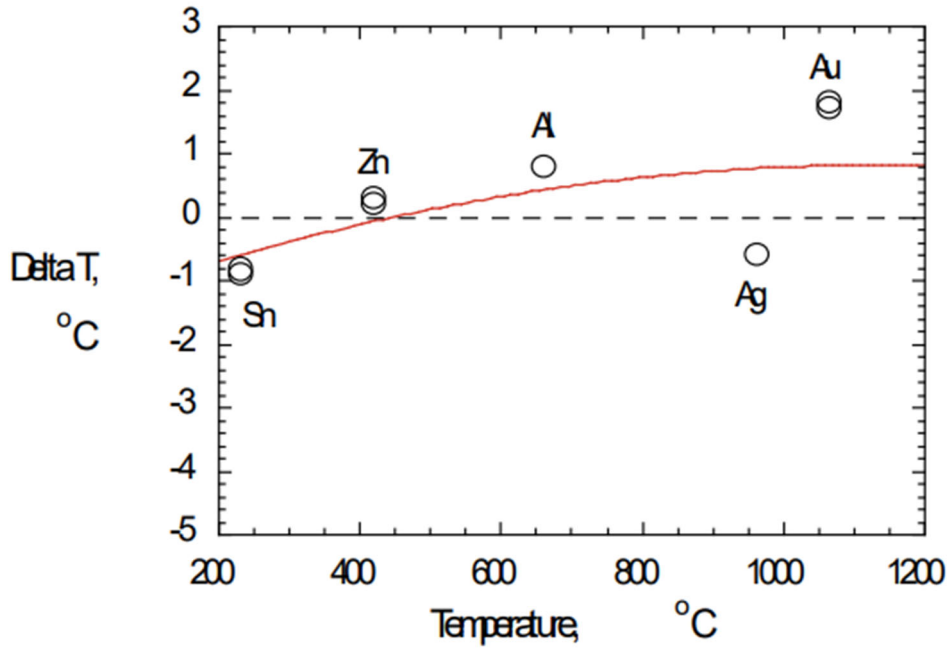

Figure E2: Calibration curve for DSC prior to analyses of the salts.

All calibration results are being tracked to detect significant changes in performance and long- term bias of the DSC and can be seen in Appendix E2, Table B.1. A summary of the mean calibration metal melting temperatures are shown in Table 1 along with values for two standard deviations (2s) to quantify the precision of the measurements. Calibrations made at different times resulted in slightly different calibrations curves, but the accuracies were always within about 1 °C and 2 °C for temperatures below and above 700 °C, respectively. No bias is observed through about 2.5 years of operation and the deviations between reported and measured melting temperature remain within 2 °C for temperatures up to 660 °C and within about 4 °C at higher temperatures. These differences are taken into account in the calibration curve used for measurements of the salt specimens.

Table E1: Mean Melting Temperatures for Calibration Metals Measured by DSC over 2.5 Years.

Calibration Metal	Sn	Zn	Al	Ag	Au
Reported Melting T, °C	231.9	419.6	660.3	961.8	1064.2
Mean Measured Melting T, °C	231.0	419.8	660.8	959.3	1062.4
2s, °C	0.8	2.0	0.9	2.2	3.8

E.2. Analyses of ORNL-FLiNaK

About 3 g of FLiNaK salt was provided by Oak Ridge National Laboratory for analysis from a batch of salt prepared by an outside vendor for them. This is referred to as ORNL-FLiNaK for convenience. The results of elemental composition measurements using ICP-MS, phase analysis using X-ray diffraction and melting point determination by using DSC are reported herein.

a. Wet Chemical Analysis

Elemental compositions were measured by dissolving three samples (68, 73 and 93 mg) in high- purity acid solution and analyzing the solutions for the major constituents Li, Na, and K and trace element impurities with ICP-MS. The individual sample measurements were converted to mole percent based on the cation analyses and are given in Table E2 with the mean values and one standard deviation (s). Calculations of mole percent compositions from the reported cation analyses are shown in Appendix E3.

Table E2: Result of Major Element Analysis of ORNL-FLiNaK by ICP-MS Composition (mol%).

Sample ID	LiF	NaF	KF
ORNL A	44.8	11.2	44.0
ORNL B	45.3	11.1	43.6
ORNL C	44.2	11.2	44.6
Average	44.8	11.2	44.1
s	0.5	0.1	0.5

The individual sample measurements, averages and standard deviations for impurities in the three samples are shown in Table E3. The ICP-MS results are typically accurate to within ±10 % of the measured concentration. The differences in the measured concentrations of the major constituents are well within this range, but variances in several trace elements exceed 10%. This probably indicates those elements are not uniformly distributed in the salt within the volumes represented by these sample sizes.

Table E3: Result of Trace Element Analysis of ORNL-FLiNaK by ICP-MS Trace Elements (ppm).

Sample ID	Mg	Ca	Cr	Mn	Fe	Ni	Rb	Cs
ORNL A	30	160	2.6	1.9	23	1.1	14	0.17
ORNL B	31	210	2.7	2.2	21	1.2	13	0.14
ORNL C	34	180	2.6	2.2	33	1.6	13	0.10
Average	32	180	2.7	2.1	25	1.3	13	0.14
s	2	24	0.1	0.1	6	0.3	0	0.04

b. DSC Analysis

Three portions of the salt (24.76, 20.06, and 24.05 mg) were encapsulated in hermetically sealed gold cells under inert atmosphere for thermal analysis. A Netzsch standard thermal analyzer (STA) with thermogravimetric analyzer (TGA) was operated as a DSC to quantify the melting behavior of the salt. Figure 3 shows the signals upon heating the salt in replicate DSC scans at 5 K/min that represent the behavior observed for all specimens. One large eutectic peak was observed, the onset of which indicates the melting temperature. Small peaks are observed in the tail of each eutectic peak that probably indicate melting of anisotropic crystals. The Netzsch software Proteus was used to determine the onset temperature of the eutectic peak from each of the scans and the results are listed in Table 4. The range of measured melting temperatures was 457.2 ± 0.2 °C.

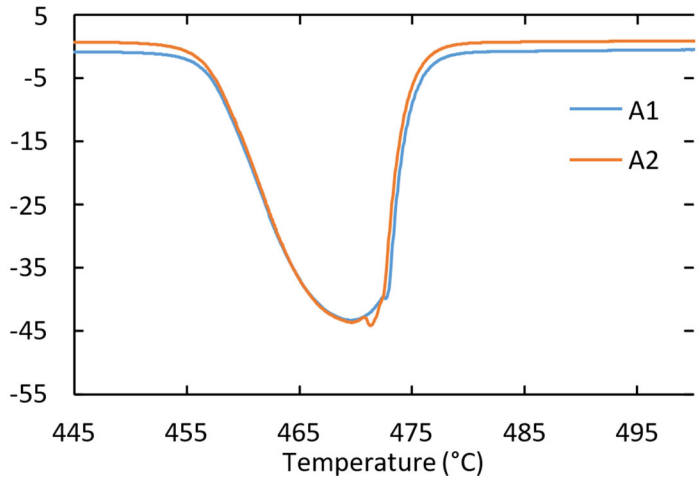

Figure E3: Representative DSC Scans of ORNL-FLiNaK.

Table E4: Eutectic Melting Temperatures for ORNL-FLiNaK.

Sample	Run	Eutectic Melting Temperature (°C)
A	1	457.1
	2	457.0
	3	457.2
	4	457.3
B	1	457.2
	2	457.5
C	1	457.0
	2	457.5
Average		457.2
s		0.2

c. XRD Analysis

Three small samples of the salt (89.3, 73.5 and 81.4 mg) were fixed on glass slides using double- sided tape and covered with Kapton film for X-ray diffraction analysis. The resulting patterns, shown in Figure E4, were compared with possible oxides, oxyfluorides and hydrates

available in the International Center for Diffraction Data (ICDD) database to verify that the salt had not been contaminated during preparation. The XRD patterns from the three analyses are the same and consistent with a ternary LiF-NaF-KF system. There is some noise in the 10-25° 2-theta region due to the Kapton film and tape used to secure the sample, but no unindexed peaks were observed. This indicates any contaminants are present in low enough amounts to be undetectable by XRD (typically about 3%).

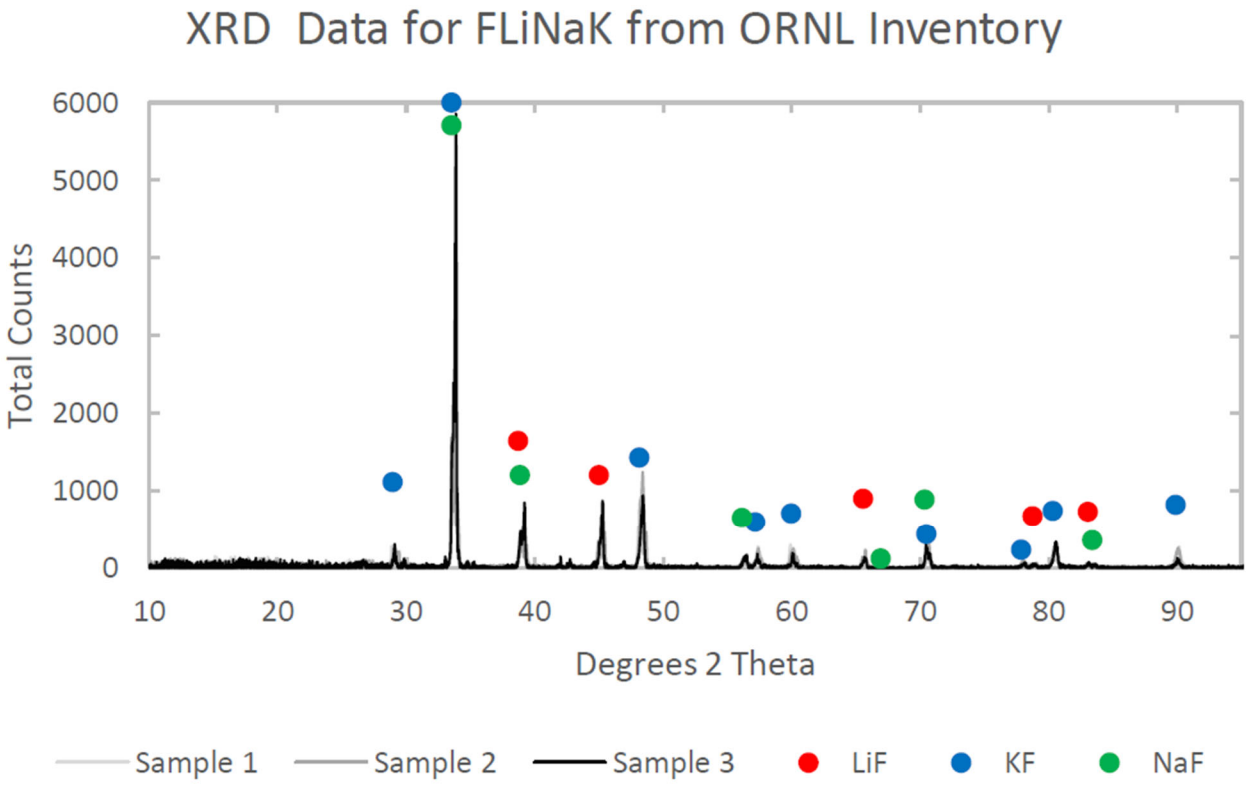

Figure E4: X-Ray Diffraction Patterns for three replicate samples of ORNL-FLiNaK.

E.3. Analyses of ANL-FLiNaK

A batch of FLiNaK salt was prepared at Argonne (hereafter referred to as ANL-FLiNaK) from reagent LiF, NaF, and KF under an inert atmosphere in a nickel crucible. The NaF and KF were purchased from VWR™ with 99.99% and 99% purity, respectively. The LiF was available at Argonne and originally purchased from Applied Engineering Materials. 600 g of LiF was added to 234.7 g NaF and 1206.1 g KF and melted at approximately 600 °C. The resulting salt was cooled, and crushed mechanically.

d. Wet Chemical Analysis

Elemental composition was determined by dissolving one sample (53.8 mg) in high-purity acid solution and analyzing the resulting solution with ICP-MS. Separate analyses were used to quantify the major constituents and impurities. The results for the major elements are provided in Table E5 and the results for the minor elements are shown in Table E6. The ICP-MS results are typically accurate to within ±10 % of the measured concentration.

Table E5: Result of Major Element Analysis of ANL-FLiNaK by ICP-MS Composition (mol%).

LiF	NaF	KF
47.1%	11.1%	41.8%

Table E6: Result of Trace Element Analysis of ANL-FLiNaK by ICP-MS Trace elements (ppm).

Mg	Ca	Cr	Mn	Fe	Ni	Rb	Cs
24	76	6	3	117	179	44	2

e. DSC Analysis

A portion of the salt was hermetically sealed in a gold cell within the inert atmosphere glovebox for thermal analysis. The same Netzsch STA with TGA used to analyze the ORNL-FLiNaK was operated as a differential scanning calorimeter to characterize the melting behavior of the ANL-FLiNaK salt. Figure E5 shows the signals from replicate DSC scans run at 5 K/min, which closely overlap. The small feature seen in both scans at approximately 450 °C indicates the composition is slightly off eutectic. This is consistent with the measured impurity levels. The occurrence of the small peak does not interfere with determination of the eutectic melting temperature.

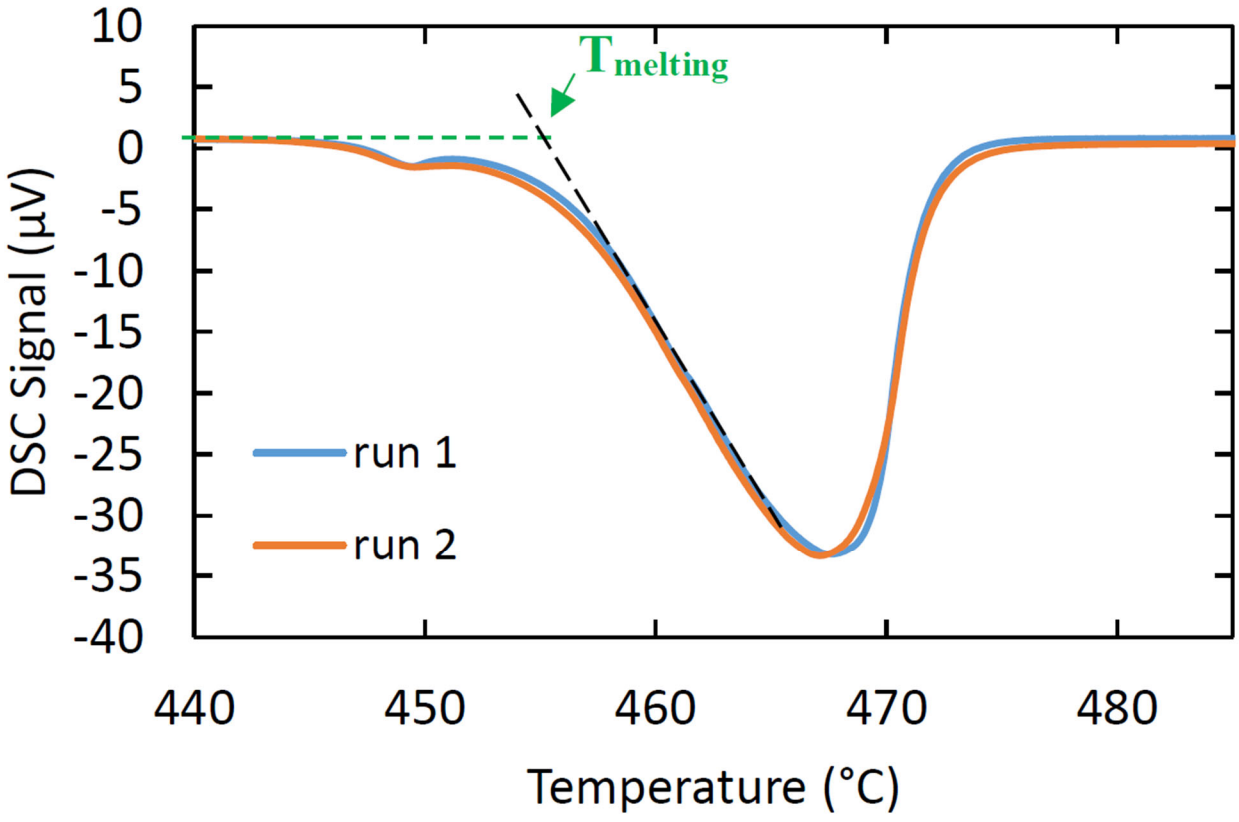

Figure E5: DSC Scans of ANL-FLiNaK Salt.

The interpretation of DSC curves requires analysis of the shape of the curve with respect to several variables to determine if a deviation from background is in fact due a transition reaction.

These include effects of salt composition, instrument artifacts and environmental factors. Once a deviation from the flat background of the DSC curve is identified as a transition reaction, the associated peak can be quantified to determine transition temperature. The approach outlined in the textbook *Differential Scanning Calorimetry* [205] is used by the PROTEUS software to determine the eutectic melting temperature from the observed peaks. This same approach is used to determine the melting temperature of the calibration metals and the salts. In this approach, a flat background is drawn underneath the observed peaks and a line tangent to the linear portion of the melting curve for each event is extrapolated to meet the background line. This approach is illustrated in Figure E5. The onset temperatures determined by the Netzsch software Proteus are listed in Table E7.

Table E7: Eutectic Melting Temperatures for Dissolution of ANL-FLiNaK.

Sample	Run	Eutectic Melting Temp.(°C)
A	1	455.6
	2	455.2
	Average	455.4
	s	0.3

f. XRD Analysis

A portion of the salt was held on a glass slide with double-sided tape and covered with Kapton film for XRD analysis. The resulting pattern, shown in Figure E6, was compared with patterns for relevant oxides, oxyfluorides, and hydrates, that could have formed by reactions with oxygen and moisture in the glovebox atmosphere using the International Center for Diffraction Data (ICDD) database to verify the salt had not been contaminated during preparation. Patterns for metals found in significant quantities in the ICP-MS data and their oxides were also checked against the XRD data. Peaks at 41, 44 and 53.5 ° two-theta are likely due to Ni, Fe and NiO, but there is no evidence of oxides, oxyfluorides or hydrates that indicate contamination by the glovebox atmosphere. The peaks below 27° are attributed to the tape and Kapton film used to secure the sample to the slide. When sample sizes are small, signal intensity is reduced and the effect of the film and tape are more pronounced. For example, maximum signal intensity for this sample of ANL-FLiNaK was ~110 counts, whereas the samples for ORNL-FLiNaK had a maximum signal intensity of >5000 counts. Despite the extensive peak matching done here, unindexed peaks still persist. A more extensive database of XRD patterns is needed to identify the remaining unindexed peaks.

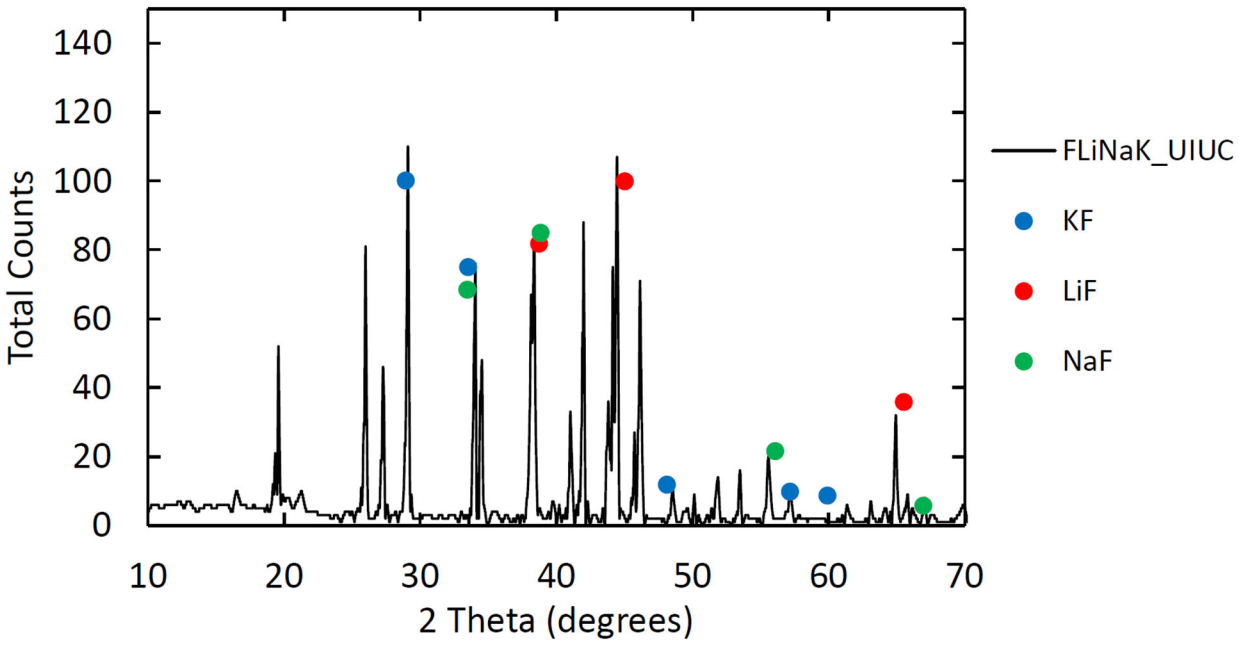

Figure E6: X-Ray Diffraction Pattern for ANL-FLiNaK Salt.

E.4. Comparing Property Measurements for ORNL-FLiNaK and ANL-FLiNaK

A brief literature review of property measurements made for the FLiNaK system was used to assess the values measured for the ORNL and Argonne salts. Thermal analysis of the eutectic LiF-NaF-KF (FLiNaK) was performed by Bergmann and Dergunov in 1941 [206]. They observed eutectic melting of FLiNaK at 454 °C, most likely using cooling curves measured during differential thermal analysis (DTA). Hoffman reported a eutectic melting temperature of 456.7 °C, but didn't describe the apparatus that was used [92]. Rogers et al. performed DSC on eutectic FLiNaK and concluded eutectic melting occurred at 462 ± 2 °C based on multiple heating curves [94]. Cheng et al. performed DSC and determined eutectic melting occurred at 457 °C from heating curves and at 453 °C from cooling curves [6]. Cheng et al. provided DSC thermograms but did not report the heating and cooling rates. It is established that molten salts often experience super cooling that leads to artificially low melting temperatures being determined from cooling curves. It is, therefore, preferable to use heating curve data to determine the melting temperature.

In general, the eutectic melting of the FLiNaK system has been measured to be between 453 and 462 °C. Although most authors use 454 °C as the accepted value, the value of 457°C determined by Cheng et al. from heating curves is probably more accurate [207]. The eutectic melting temperatures measured for the salt from ORNL inventory was 457.2 ± 0.2 °C, which agrees with the Cheng et al. value of 457 °C within this uncertainty. The melting temperature for the salt prepared at Argonne was determined to be 455.4 ±0.3 °C. The melting temperatures of the Argonne prepared salt were slightly depressed from this value, as expected due to the higher impurity content (e.g., Ni, Fe and NiO) and slightly off eutectic composition.

The eutectic composition and compositions measured for ORNL-FLiNaK and ANL-FLiNaK are shown in Table E8. The measured compositions of both ORNL-FLiNaK and ANL-FLiNaK differed slightly from the eutectic composition: the ORNL-FLiNaK had slightly lower

LiF and NaF contents and a slightly higher KF content, whereas the ANL-FLiNaK had a slightly higher LiF content and slightly lower NaF and KF contents. The analytical uncertainties in the cation analyses are magnified slightly in the calculated mol % values. The DSC curves indicate the ORNL-FLiNaK is very near the eutectic composition, whereas the observation of a low-temperature peak in the DSC analysis of the ANL-FLiNaK indicates it is slightly off the eutectic. Nevertheless, the difference between the mean values for the melting points of the ORNL- FLiNaK and ANL-FLiNaK is less than 2 °C, which is near the accuracy of the DSC.

Table E8: FLiNaK Salt Compositions.

	Salt Composition (mol %)		
	LiF	NaF	KF
Eutectic - FLiNaK	46.5	11.5	42
ORNL - FLiNaK	44.7	11.2	44.1
ANL - FLiNaK	47.1	11.1	41.8

Analyses also showed that the ORNL-FLiNaK salt has fewer trace impurities than the ANL-FLiNaK salt. The major contributor of trace impurities in the salt from the ORNL inventory was calcium. The major contributors to the trace impurities in the salt prepared at Argonne are iron and nickel, which is expected because the salt was prepared in a nickel crucible and removed from the crucible by using steel tools. Because most MSR systems will be constructed of nickel and stainless-steel alloys, the salt made at Argonne represents that expected to be present in an MSR system.

Appendix E1: Material Certifications

Specification

1.05686.0050 Lithium fluoride 99.99 Suprapur®

	Specification	
Purity (metallic)	≥ 99.99	%
Ba (Barium)	≤ 5.0	ppm
Ca (Calcium)	≤ 2.0	ppm
Cd (Cadmium)	≤ 0.5	ppm
Co (Cobalt)	≤ 0.5	ppm
Cs (Cesium)	≤ 20	ppm
Cu (Copper)	≤ 0.5	ppm
Fe (Iron)	≤ 0.5	ppm
K (Potassium)	≤ 10	ppm
Mg (Magnesium)	≤ 0.5	ppm
Mn (Manganese)	≤ 0.5	ppm
Na (Sodium)	≤ 10	ppm
Ni (Nickel)	≤ 0.5	ppm
Pb (Lead)	≤ 0.5	ppm
Rb (Rubidium)	≤ 10	ppm
Sr (Strontium)	≤ 5	ppm
Zn (Zinc)	≤ 0.5	ppm

Evelyn Allmann
Responsible laboratory manager quality control

This document has been produced electronically and is valid without a signature.

Figure E7: Certificate of Analysis for LiF.

SIGMA-ALDRICH®

sigma-aldrich.com

3050 Spruce Street, Saint Louis, MO 63103, USA
Website: www.sigmaaldrich.com
Email USA: techserv@sial.com
Outside USA: eurtechserv@sial.com

Certificate of Analysis

Product Name:
Sodium fluoride – anhydrous, powder, 99.99% trace metals basis

Product Number:	**450022**
Batch Number:	**MKBV9645V**
Brand:	ALDRICH
CAS Number:	7681-49-4
MDL Number:	MFCD00003524
Formula:	FNa
Formula Weight:	41.99 g/mol
Quality Release Date:	17 JUL 2015

NaF

Test	Specification	Result
Appearance (Color)	White	White
Appearance (Form)	Powder	Powder
Gravimetric Analysis % F With Lead Acetate		45.3 %
ICP Major Analysis Confirms Sodium Component	Confirmed	Conforms
Purity 99.99% Based On Trace Metals Analysis	Meets Requirements	Meets Requirements
Trace Metal Analysis	≤ 150.0 ppm	31.3 ppm
Aluminum (Al)		0.6 ppm
Barium (Ba)		0.5 ppm
Calcium (Ca)		1.5 ppm
Cesium (Cs)		5.0 ppm
Iron (Fe)		14.2 ppm
Potassium (K)		6.6 ppm
Lithium (Li)		0.9 ppm
Magnesium (Mg)		1.3 ppm
Manganese (Mn)		0.3 ppm
Rubidium (Rb)		0.4 ppm
Strontium (Sr)		< 0.1 ppm

Ali Ataei

Ali Ataei, Manager

Figure E8: Certificate of analysis for NaF.

SIGMA-ALDRICH®

sigma-aldrich.com

3050 Spruce Street, Saint Louis, MO 63103, USA
Website: www.sigmaaldrich.com
Email USA: techserv@sial.com
Outside USA: eurtechserv@sial.com

Certificate of Analysis

Product Name:
Potassium fluoride – anhydrous, powder, ≥99.9% trace metals basis

Product Number:	**449148**
Batch Number:	**MKBX2566V**
Brand:	ALDRICH
CAS Number:	7789-23-3
MDL Number:	MFCD00011398
Formula:	FK
Formula Weight:	58.10 g/mol
Quality Release Date:	31 DEC 2015

KF

Test	Specification	Result
Appearance (Color)	White	White
Appearance (Form)	Powder	Powder
ICP Major Analysis	Confirmed	Confirmed
Confirms K Component		
Purity	Conforms	Conforms
>=99.9% Based On Trace Metals Analysis		
Trace Metal Analysis	≤ 1000.0 ppm	614.7 ppm
Boron (B)		1.4 ppm
Calcium (Ca)		3.0 ppm
Chromium (Cr)		2.1 ppm
Cesium (Cs)		12.7 ppm
Iron (Fe)		8.4 ppm
Lithium (Li)		1.9 ppm
Manganese (Mn)		0.9 ppm
Sodium (Na)		527.4 ppm
Rubidium (Rb)		57.0 ppm

Michael Grady, Manager
Quality Control
Milwaukee, WI US

Figure E9: Certificate of analysis for KF.

U. S. Department of Energy
New Brunswick Laboratory

New Brunswick Laboratory
Certified Reference Materials
Certificate of Analysis

CRM 17-B

Uranium (Normal) Tetrafluoride - UF_4

Total Uranium	75.87%
U^{+4}	74.5%
UO_2 (Ammonium Oxalate Insoluble)	1.0%
Fe	0.003%
Ni	0.002%
Mo	0.0001%
V	<0.00005%

July 1961
New Brunswick, New Jersey

Clement J. Rodden
Area Manager

Figure E10: Certificate of analysis for UF_4.

Appendix E2: DSC Calibration Data

Table E9: DSC Calibration Data over Past 2.5 Years.

	Measured Melting Temperatures (°C)				
Date of Measurement	Sn	Zn	Al	Ag	Au
08/18/17	231.2	418.2	660.7	959.4	1063.1
08/18/17	230.5	418.8	660.7	959.5	1062.9
01/02/18	231.3	419.9	661.3	959.6	1062.0
01/02/18	231.2	420.0	661.5	959.7	1062.1
01/24/18	230.4	419.4	660.3	957.8	1059.9
01/24/18	230.4	419.4	661.1	958.0	1059.9
09/18/18	230.8	421.7	660.7	958.8	1062.1
09/18/18	230.8	421.6	660.7	958.8	1062.3
02/01/19	231.6	419.3	660.1	958.6	1061.5
02/01/19	231.6	419.5	660.0	958.6	1061.5
12/01/19	231.0	419.8	661.1	961.2	1065.9
12/01/19	231.1	419.9	661.1	961.2	1066.0
Mean Measured T_m	231.0	419.8	660.8	959.3	1062.4
2s	0.8	2.0	0.9	2.2	3.8
ΔT_m	1.2	3.5	1.5	3.4	6.1
Nominal T_m	231.9	419.6	660.3	961.8	1064.2

Appendix E3: Calculation of Salt mol % from Reported Cation wt.%

Wet chemical analysis by ICP-MS reports the major element analyses as a weight percent of the total sample. Table E10 provides the results for analysis of three ORNL-FLiNaK samples as an example. These are converted first to grams of each element in the sample by multiplication of the sample weight by the reported weight percent. As compiled in Table E11, the mass of each element is converted to moles of each element in the sample by dividing the mass of each element by the atomic weight of that element. Because the Li, Na, and K all have 1:1 stoichiometric ratios with their fluoride salt forms, the number of moles of each salt constituent is equal to the number of moles of the cation. The number of moles of each constituent can be divided by the total number of moles to represent the analytical results as the mole% of the major constituents.

Table E10: Reported ICP-MS Results.

Sample ID	Sample weight [g]	Li [wt%]	Na [wt%]	K [wt%]
ORNL A	0.0676	7.57	6.27	41.9
ORNL B	0.0733	7.59	6.17	41.2
ORNL C	0.0929	7.97	6.67	45.2

Table E11: Conversion from Mass of Each Element to mol% of Salt Constituents.

	Mass of Elements (grams)			Moles of Salt				mole % of salts		
Sample ID	Li	Na	K	LiF	NaF	KF	Total	LiF	NaF	KF
ORNL A	0.00512	0.00424	0.0283	7.38E-04	1.84E-04	7.24E-04	1.65E-03	44.8%	11.2%	44.0%
ORNL B	0.00556	0.00452	0.0302	8.01E-04	1.97E-04	7.72E-04	1.77E-03	45.3%	11.1%	43.6%
ORNL C	0.00740	0.00620	0.0420	1.07E-03	2.70E-04	1.07E-03	2.41E-03	44.2%	11.2%	44.6%

REFERENCE

[1] R.C. Robertson, Conceptual design study of a single-fluid molten-salt breeder reactor, Oak Ridge National Lab., Tenn., 1971.
[2] J.R. Engel, Experience with the molten-salt reactor experiment, Nuclear Applications and Technology 8(2) (1970) 118-136. https://doi.org/10.13182/NT8-2-118.
[3] R.R. Romatoski, L.W. Hu, Fluoride salt coolant properties for nuclear reactor applications: A review, Annals of Nuclear Energy 109 (2017) 635-647. https://doi.org/10.1016/j.anucene.2017.05.036.
[4] J. Serp, M. Allibert, O. Beneš, S. Delpech, O. Feynberg, V. Ghetta, D. Heuer, D. Holcomb, V. Ignatiev, J.L. Kloosterman, L. Luzzi, E. Merle-Lucotte, J. Uhlíř, R. Yoshioka, D. Zhimin, The molten salt reactor (MSR) in generation IV: overview and perspectives, Progress in Nuclear Energy 77 (2014) 308-319. https://doi.org/10.1016/j.pnucene.2014.02.014.
[5] P.R. Kasten, E.S. Bettis, R.C. Robertson, Design studies of 1000-Mw(e) molten salt breeder reactors, Oak Ridge National Lab., Tenn., 1966.
[6] J.R. Engel, H.F. Bauman, J.F. Dearing, W.R. Grimes, H.E. McCoy, W.A. Rhoades, Conceptual design characteristics of a denatured molten-salt reactor with once-through fueling, United States, 1980, p. Medium: ED; Size: Pages: 160.
[7] J.C. Gehin, J.J. Powers, Liquid fuel molten salt reactors for thorium utilization, Nuclear Technology 194(2) (2016) 152-161. https://doi.org/10.13182/NT15-124.
[8] W.B. Cottrell, H.E. Hungerford, J.K. Leslie, J.L. Meem, Operation of the Aircraft Reactor Experiment, Oak Ridge National Lab., Tenn., 1955.
[9] H.G. MacPherson, The molten salt reactor adventure, Nuclear Science and Engineering 90(4) (1985) 374-380. https://doi.org/10.13182/NSE90-374.
[10] M.W. Rosenthal, P.N. Haubenreich, F.B. Briggs, Development status of molten-salt breeder reactors, Oak Ridge National Lab., Tenn., United States, 1972.
[11] G. Locatelli, M. Mancini, N. Todeschini, Generation IV nuclear reactors: current status and future prospects, Energy Policy 61 (2013) 1503-1520. https://doi.org/https://doi.org/10.1016/j.enpol.2013.06.101.
[12] M.W. Rosenthal, P.R. Kasten, R.B. Briggs, Molten-Salt Reactors - History, Status, and Potential, Nuclear Applications and Technology 8(2) (1970) 107–117.
[13] W.M. Stacey, Nuclear reactor physics, John Wiley & Sons, 2018.
[14] Transatomic Power Corporation, Technical White Paper, 2016, http://www.transatomicpower.com/.
[15] Transatomic Power Corporation, Neutronics Overview, 2016, http://www.transatomicpower.com/.
[16] B.R. Betzler, J.J. Powers, A. Worrall, L. Dewan, S. Robertson, M. Massie, Two-Dimensional Neutronic and Fuel Cycle Analysis of the Transatomic Power Molten Salt Reactor, 2017.
[17] E. Strickland, Transatomic power building a "walk-away safe" reactor, IEEE Spectrum 51(7) (2014) 25–26.
[18] B.R. Betzler, S. Robertson, E.E. Davidson, J.J. Powers, A. Worrall, L. Dewan, M. Massie, Assessment of the Neutronic and Fuel Cycle Performance of the Transatomic Power Molten Salt Reactor Design, 2017.
[19] Transatomic Power Corporation , Transatomic Reactor Documentation, 2019, http://www.transatomicpower.com/.

[20] B.R. Betzler, J.J. Powers, N.R. Brown, B.T. Rearden, Implementation of Molten Salt Reactor Tools in SCALE, Proc. of GLOBAL International Fuel Cycle Conference, 2017.
[21] M.A. Rose, E. Wu, M.A. Williamson, Thermophysical Property Measurements: Improved Density, Viscosity and Thermal Diffusivity Methods, United States, 2020, p. Medium: ED.
[22] T. Lichtenstein, M.A. Rose, J. Krueger, E. Wu, M.A. Williamson, Thermochemical Property Measurements of FLiNaK and FLiBe in FY 2020, United States, 2020, p. Medium: ED.
[23] A.P. Fraas, A.W. Savolainen, Design report on the Aircraft Reactor Test, United States, 1956.
[24] R.C. Robertson, MSRE design and operations report. Part I. Description of reactor design, Oak Ridge National Lab., Tenn., 1965.
[25] D. Scott, W.P. Eatherly, Graphite and Xenon Behavior and their Influence on Molten-Salt Reactor Design, Nuclear Applications and Technology 8(2) (1969) 179-189. https://doi.org/10.13182/nt70-a28624.
[26] R.J. Kedl, A. Houtzeel, Development of A model for computing Xe 135 migration in the MSRE, Oak Ridge National Lab., Tenn., 1967.
[27] J.R. Engel, R.C. Steffy, Xenon behavior in the molten salt reactor experiment, Oak Ridge National Lab., Tenn., 1971.
[28] F.N. Peebles, Removal of Xe-135 from circulating fuel salt of the MSBR by mass transfer to helium bubbles, Oak Ridge National Lab., Tenn., 1968.
[29] C.H. Gabbard, Development of an axial-flow centrifugal gas bubble separator for use in MSR xenon removal system, Oak Ridge National Lab., Tenn., United States, 1972.
[30] F. Rubio, L. Bond, E. Blandford, Scaled experiment investigating sonomechanically enhanced inert gas sparging mass transfer, Nuclear Engineering and Design 324 (2017) 171-180. https://doi.org/10.1016/j.nucengdes.2017.08.016.
[31] Y. Kanai, K.-i. Fukunaga, K. Terasaka, S. Fujioka, Mass transfer in molten salt and suspended molten salt in bubble column, Chemical Engineering Science 100 (2013) 153-159. https://doi.org/10.1016/j.ces.2012.11.029.
[32] J. Chen, C.S. Brooks, CFD Simulation of xenon removal by helium bubble sparging in molten salt, 18th International Topical Meeting on Nuclear Reactor Thermal Hydraulics (NURETH-18), Portland, Oregon, 2019, pp. 134-156.
[33] H. Funahashi, K. Hayashi, S. Hosokawa, A. Tomiyama, Study on two-phase swirling flows in a gas–liquid separator with three pick-off rings, Nuclear Engineering and Design 308 (2016) 205-213. https://doi.org/https://doi.org/10.1016/j.nucengdes.2016.08.030.
[34] C. Zheng, W. Yang, G. Wang, G. Fan, C. Yan, X. Zeng, A. Liu, Experimental study on a new type of separator for gas liquid separation, Frontiers in Energy Research 7 (2019). https://doi.org/10.3389/fenrg.2019.00102.
[35] P.M. Bardet, P.F. Peterson, Options for scaled experiments for high temperature liquid salt and helium fluid mechanics and convective heat transfer, Nuclear Technology 163(3) (2008) 344-357. https://doi.org/10.13182/NT163-344.
[36] S. Piyush, E. Matt, S. Manohar, S. Phil, G. Thermal Hydraulics, Molten salts for high temperature reactors: University of Wisconsin molten salt corrosion and flow loop experiments – issues identified and path forward, 2010.
[37] K. Britsch, M. Anderson, P. Brooks, K. Sridharan, Natural circulation FLiBe loop overview, International Journal of Heat and Mass Transfer 134 (2019) 970-983. https://doi.org/https://doi.org/10.1016/j.ijheatmasstransfer.2018.12.180.

[38] D. Zhang, L. Liu, M. Liu, R. Xu, C. Gong, J. Zhang, C. Wang, S. Qiu, G. Su, Review of conceptual design and fundamental research of molten salt reactors in China, International Journal of Energy Research 42(5) (2018) 1834-1848. https://doi.org/10.1002/er.3979.
[39] M. Liu, D. Zhang, C. Wang, S. Qiu, G.H. Su, W. Tian, Experimental study on the heat transfer characteristics of fluoride salt in the new conceptual passive heat removal system of molten salt reactor, International Journal of Energy Research 42(4) (2018) 1635-1648. https://doi.org/10.1002/er.3959.
[40] M. Pourtousi, J.N. Sahu, P. Ganesan, Effect of interfacial forces and turbulence models on predicting flow pattern inside the bubble column, Chemical Engineering and Processing: Process Intensification 75 (2014) 38-47. https://doi.org/10.1016/j.cep.2013.11.001.
[41] H.A. Jakobsen, H. Lindborg, C.A. Dorao, Modeling of bubble column reactors: progress and limitations, Industrial & Engineering Chemistry Research 44(14) (2005) 5107-5151. https://doi.org/10.1021/ie049447x.
[42] M.V. Tabib, S.A. Roy, J.B. Joshi, CFD simulation of bubble column—An analysis of interphase forces and turbulence models, Chemical Engineering Journal 139(3) (2008) 589-614. https://doi.org/https://doi.org/10.1016/j.cej.2007.09.015.
[43] R. Krishna, J.M. van Baten, Mass transfer in bubble columns, Catalysis Today 79-80 (2003) 67-75. https://doi.org/10.1016/s0920-5861(03)00046-4.
[44] T. Wang, J. Wang, Numerical simulations of gas–liquid mass transfer in bubble columns with a CFD–PBM coupled model, Chemical Engineering Science 62(24) (2007) 7107-7118. https://doi.org/10.1016/j.ces.2007.08.033.
[45] D. Wiemann, D. Mewes, Calculation of flow fields in two and three-phase bubble columns considering mass transfer, Chemical Engineering Science 60(22) (2005) 6085-6093. https://doi.org/https://doi.org/10.1016/j.ces.2005.04.054.
[46] A.A. Kulkarni, Mass transfer in bubble column reactors: effect of bubble size distribution, Industrial & Engineering Chemistry Research 46 (2007) 2205-2211.
[47] R. Rzehak, E. Krepper, Euler-Euler simulation of mass-transfer in bubbly flows, Chemical Engineering Science 155 (2016) 459-468. https://doi.org/https://doi.org/10.1016/j.ces.2016.08.036.
[48] T.S. Kress, Mass transfer between small bubbles and liquids in cocurrent turbulent pipeline flow, Oak Ridge National Lab., Tenn., 1972.
[49] N. Kantarci, F. Borak, K.O. Ulgen, Bubble column reactors, Process Biochemistry 40(7) (2005) 2263-2283. https://doi.org/https://doi.org/10.1016/j.procbio.2004.10.004.
[50] F. Yoshida, K. Akita, Performance of gas bubble columns: Volumetric liquid-phase mass transfer coefficient and gas holdup, AIChE Journal 11(1) (1965) 9-13. https://doi.org/10.1002/aic.690110106.
[51] R. Krishna, P.M. Wilkinson, L.L. Van Dierendonck, A model for gas holdup in bubble columns incorporating the influence of gas density on flow regime transitions, Chemical Engineering Science 46(10) (1991) 2491-2496. https://doi.org/https://doi.org/10.1016/0009-2509(91)80042-W.
[52] N. Devanathan, D. Moslemian, M.P. Dudukovic, Flow mapping in bubble columns using CARPT, Chemical Engineering Science 45(8) (1990) 2285-2291. https://doi.org/https://doi.org/10.1016/0009-2509(90)80107-P.
[53] S. Degaleesan, Fluid dynamic measurements and modeling of liquid mixing in bubble columns, 1998.

[54] S.B. Kumar, D. Moslemian, M.P. Duduković, A γ-ray tomographic scanner for imaging voidage distribution in two-phase flow systems, Flow Measurement and Instrumentation 6(1) (1995) 61-73. https://doi.org/https://doi.org/10.1016/0955-5986(95)93459-8.
[55] V.V. Buwa, V.V. Ranade, Dynamics of gas–liquid flow in a rectangular bubble column: experiments and single/multi-group CFD simulations, Chemical Engineering Science 57(22) (2002) 4715-4736. https://doi.org/https://doi.org/10.1016/S0009-2509(02)00274-9.
[56] Y.A. Hassan, T.K. Blanchat, C.H. Seeley, R.E. Canaan, Simultaneous velocity measurements of both components of a two-phase flow using particle image velocimetry, International Journal of Multiphase Flow 18(3) (1992) 371-395. https://doi.org/https://doi.org/10.1016/0301-9322(92)90023-A.
[57] J. Basley, L.R. Pastur, F. Lusseyran, T.M. Faure, N. Delprat, Experimental investigation of global structures in an incompressible cavity flow using time-resolved PIV, Experiments in Fluids 50(4) (2011) 905-918. https://doi.org/10.1007/s00348-010-0942-9.
[58] M.J. Sathe, I.H. Thaker, T.E. Strand, J.B. Joshi, Advanced PIV/LIF and shadowgraphy system to visualize flow structure in two-phase bubbly flows, Chemical Engineering Science 65(8) (2010) 2431-2442. https://doi.org/https://doi.org/10.1016/j.ces.2009.11.014.
[59] A. Zaruba, E. Krepper, H.M. Prasser, B.N. Reddy Vanga, Experimental study on bubble motion in a rectangular bubble column using high-speed video observations, Flow Measurement and Instrumentation 16(5) (2005) 277-287. https://doi.org/https://doi.org/10.1016/j.flowmeasinst.2005.03.009.
[60] Y.M. Lau, K.T. Sujatha, M. Gaeini, N.G. Deen, J.A.M. Kuipers, Experimental study of the bubble size distribution in a pseudo-2D bubble column, Chemical Engineering Science 98 (2013) 203-211. https://doi.org/https://doi.org/10.1016/j.ces.2013.05.024.
[61] H.M. Letzel, J.C. Schouten, R. Krishna, C.M. van den Bleek, Gas holdup and mass transfer in bubble column reactors operated at elevated pressure, Chemical Engineering Science 54(13) (1999) 2237-2246. https://doi.org/https://doi.org/10.1016/S0009-2509(98)00418-7.
[62] S.S. öztürk, A. Schumpe, W.D. Deckwer, Organic liquids in a bubble column: Holdups and mass transfer coefficients, AIChE Journal 33(9) (1987) 1473-1480. https://doi.org/10.1002/aic.690330907.
[63] R. Lau, P.H.V. Lee, T. Chen, Mass transfer studies in shallow bubble column reactors, Chemical Engineering and Processing: Process Intensification 62 (2012) 18-25. https://doi.org/https://doi.org/10.1016/j.cep.2012.10.003.
[64] G. Huang, Y. Shentu, W. Chen, K. Cai, B. Dong, D. Wang, J. Yin, Experimental investigation on hydrodynamics and mass transfer of a novel fission gas removal equipment for molten salt reactor, Nuclear Engineering and Design 372 (2021) 110977. https://doi.org/https://doi.org/10.1016/j.nucengdes.2020.110977.
[65] R.I. Issa, P.J. Oliveira, Numerical prediction of phase separation in two-phase flow through T-junctions, Computers & Fluids 23(2) (1994) 347-372. https://doi.org/https://doi.org/10.1016/0045-7930(94)90045-0.
[66] Y. Sato, M. Sadatomi, K. Sekoguchi, Momentum and heat transfer in two-phase bubble flow—I. Theory, International Journal of Multiphase Flow 7(2) (1981) 167-177. https://doi.org/https://doi.org/10.1016/0301-9322(81)90003-3.
[67] A. Tomiyama, I. Kataoka, I. Zun, T. Sakaguchi, Drag coefficients of single bubbles under normal and micro gravity conditions, JSME International Journal Series B 41(2) (1998) 472-479. https://doi.org/10.1299/jsmeb.41.472.

[68] J.F. Richardson, W.N. Zaki, Sedimentation and fluidisation: Part I, Chemical Engineering Research and Design 75 (1997) S82-S100. https://doi.org/https://doi.org/10.1016/S0263-8762(97)80006-8.
[69] A. Tomiyama, H. Tamai, I. Zun, S. Hosokawa, Transverse migration of single bubbles in simple shear flows, Chemical Engineering Science 57(11) (2002) 1849-1858. https://doi.org/https://doi.org/10.1016/S0009-2509(02)00085-4.
[70] M.L. de Bertodano, Turbulent bubbly two-phase flow in a triangular duct, 1992.
[71] J.C. Lamont, D.S. Scott, An eddy cell model of mass transfer into the surface of a turbulent liquid, AIChE Journal 16(4) (1970) 513-519. https://doi.org/10.1002/aic.690160403.
[72] M. Blander, W.R. Grimes, N.V. Smith, G.M. Watson, Solubility of noble gases in molten fluorides. II. In the Lif–NaF–LF eutectic mixtures, The Journal of Physical Chemistry 63(7) (1959) 1164-1167. https://doi.org/10.1021/j150577a033.
[73] B. Jajuee, A. Margaritis, D. Karamanev, M.A. Bergougnou, Application of surface-renewal-stretch model for interface mass transfer, Chemical Engineering Science 61(12) (2006) 3917-3929. https://doi.org/https://doi.org/10.1016/j.ces.2006.01.026.
[74] I. Kataoka, M. Ishii, A. Serizawa, Local formulation and measurements of interfacial area concentration in two-phase flow, International Journal of Multiphase Flow 12(4) (1986) 505-529. https://doi.org/https://doi.org/10.1016/0301-9322(86)90057-1.
[75] H. Luo, Coalescence, breakup and liquid circulation in bubble column reactors, (1995).
[76] C. Martínez Bazán, J. RodrÍguez-RodrÍguez, G.B. Deane, J.L. MontaÑEs, J.C. Lasheras, Considerations on bubble fragmentation models, Journal of Fluid Mechanics 661 (2010) 159-177. https://doi.org/10.1017/S0022112010003186.
[77] G.J. Janz, R. Tomkins, Physical properties data compilations relevant to energy storage. IV. Molten salts: data on additional single and multi-component salt systems, National Standard Reference Data System, 1981.
[78] H. Poppendick, Physical property charts for some reactor fuels, coolants and miscellaneous materials, Oak Ridge National Lab., Tenn., 1952.
[79] W. Powers, S. Cohen, N. Greene, Physical properties of molten reactor fuels and coolants, Nuclear Science and Engineering 17(2) (1963) 200-211.
[80] S.I. Cohen, T.N. Jones, A summary of density measurements on molten fluoride mixtures and a correlation for predicting densities of fluoride mixtures, Oak Ridge National Lab., Tenn., 1954.
[81] M. Salanne, C. Simon, P. Turq, P.A. Madden, Heat-transport properties of molten fluorides: Determination from first-principles, Journal of Fluorine Chemistry 130(1) (2009) 38-44. https://doi.org/10.1016/j.jfluchem.2008.07.013.
[82] B. Kubíková, I. Macková, M. Boča, Phase analysis and volume properties of the (LiF–NaF–KF)eut–K2ZrF6 system, Monatshefte für Chemie - Chemical Monthly 144(3) (2013) 295-300. https://doi.org/10.1007/s00706-012-0886-2.
[83] S.I. Cohen, W.D. Powers, N.D. Greene, A physical property summary for fluoride mixtures, Oak Ridge National Lab., Tenn., 1956.
[84] B. Vriesema, Aspects of molten fluorides as heat transfer agents for power generation, TU Delft, Delft University of Technology, 1979.
[85] A. Merzlyakov, V. Ignatiev, S. Abalin, Viscosity of LiF–NaF–KF eutectic and effect of cerium trifluoride and uranium tetrafluoride additions, Nuclear Engineering and Design 278 (2014) 268-273. https://doi.org/10.1016/j.nucengdes.2014.07.037.

[86] J. Cibulková, M. Chrenková, R. Vasiljev, V. Kremenetsky, M. Boča, Density and viscosity of the (LiF + NaF + KF)eut (1) + K2TaF7 (2) + Ta2O5 (3) melts, Journal of Chemical & Engineering Data 51(3) (2006) 984-987. https://doi.org/10.1021/je050490g.
[87] K. Tørklep, H.A. Øye, Viscosity of the eutectic LiF-NaF-KF melt (FLINAK), Journal of Chemical and Engineering Data 25(1) (1980) 16-17.
[88] C. Ewing, J. Spann, R. Miller, Radiant transfer of heat in molten inorganic compounds at high temperatures, Journal of Chemical and Engineering Data 7(2) (1962) 246-250.
[89] D.F. Williams, L. Toth, K. Clarno, Assessment of Candidate Molten Salt Coolants for the Advanced High Temperature Reactor (AHTR), United States. Department of Energy2006.
[90] V. Khokhlov, I. Korzun, V. Dokutovich, E. Filatov, Heat capacity and thermal conductivity of molten ternary lithium, sodium, potassium, and zirconium fluorides mixtures, Journal of Nuclear Materials 410(1-3) (2011) 32-38. https://doi.org/10.1016/j.jnucmat.2010.12.306.
[91] M. Smirnov, V. Khokhlov, E. Filatov, Thermal conductivity of molten alkali halides and their mixtures, Electrochimica Acta 32(7) (1987) 1019-1026.
[92] H. Hoffman, J. Lones, Fused salt heat transfer. Part II. Forced convection heat transfer in circular tubes containing NaF-KF-LiF eutectic, Oak Ridge National Lab., Tenn., 1955.
[93] D. Williams, Assessment of candidate molten salt coolants for the NGNP/NHI heat-transfer loop, Oak Ridge National Lab., Tenn., 2006.
[94] D.J. Rogers, T. Yoko, G.J. Janz, Fusion properties and heat capacities of the eutectic lithium fluoride-sodium fluoride-potassium fluoride melt, Journal of Chemical & Engineering Data 27(3) (1982) 366-367. https://doi.org/10.1021/je00029a041.
[95] H.H. Uhlig, The solubilities of gases and surface tension, The Journal of Physical Chemistry 41(9) (1937) 1215-1226. https://doi.org/10.1021/j150387a007.
[96] J.P. Simonin, Effect of polarization on the solubility of gases in molten salts, J Chem Phys 134(5) (2011) 054508. https://doi.org/10.1063/1.3544374.
[97] J. Hiscock, C.Y. Ching, Gas bubble size and velocity measurements using high-speed images of gas-liquid flows, Journal of Flow Visualization and Image Processing 8(4) (2001).
[98] M. Ishii, T. Hibiki, Thermo-fluid dynamics of two-phase flow, Springer Science & Business Media2010.
[99] H. Hikita, S. Asai, K. Tanigawa, K. Segawa, M. Kitao, The volumetric liquid-phase mass transfer coefficient in bubble columns, The Chemical Engineering Journal 22(1) (1981) 61-69. https://doi.org/https://doi.org/10.1016/0300-9467(81)85006-X.
[100] V. Yakhot, S.A. Orszag, Renormalization group analysis of turbulence. I. Basic theory, Journal of Scientific Computing 1(1) (1986) 3-51. https://doi.org/10.1007/bf01061452.
[101] G. Gouesbet, A. Berlemont, A. Picart, Dispersion of discrete particles by continuous turbulent motions. Extensive discussion of the Tchen’s theory, using a two‐parameter family of Lagrangian correlation functions, The Physics of Fluids 27(4) (1984) 827-837. https://doi.org/10.1063/1.864711.
[102] E.S. Gaddis, A. Vogelpohl, Bubble formation in quiescent liquids under constant flow conditions, Chemical Engineering Science 41(1) (1986) 97-105. https://doi.org/10.1016/0009-2509(86)85202-2.
[103] E. Krepper, B.N. Reddy Vanga, A. Zaruba, H.-M. Prasser, M.A. Lopez de Bertodano, Experimental and numerical studies of void fraction distribution in rectangular bubble columns, Nuclear Engineering and Design 237(4) (2007) 399-408. https://doi.org/https://doi.org/10.1016/j.nucengdes.2006.07.009.

[104] J. Chen, Validation of a multiphase CFD model with mass transfer for xenon removal in molten salt reactor, University of Illinois, 2020.
[105] Y. Haroun, D. Legendre, L. Raynal, Volume of fluid method for interfacial reactive mass transfer: Application to stable liquid film, Chemical Engineering Science 65(10) (2010) 2896-2909. https://doi.org/10.1016/j.ces.2010.01.012.
[106] H. Marschall, K. Hinterberger, C. Schüler, F. Habla, O. Hinrichsen, Numerical simulation of species transfer across fluid interfaces in free-surface flows using OpenFOAM, Chemical Engineering Science 78 (2012) 111-127. https://doi.org/10.1016/j.ces.2012.02.034.
[107] W. Ranz, W.R. Marshall, Evaporation from drops, Chemical Engineering Progress 48(3) (1952) 141-146.
[108] J. Chen, C.S. Brooks, Experiments and CFD simulation of mass transfer and hydrodynamics in a cylindrical bubble column, Chemical Engineering Science 234 (2021) 116435. https://doi.org/https://doi.org/10.1016/j.ces.2020.116435.
[109] T.S. Kress, J.J. Keyes, Liquid phase controlled mass transfer to bubbles in cocurrent turbulent pipeline flow, Chemical Engineering Science 28(10) (1973) 1809-1823. https://doi.org/https://doi.org/10.1016/0009-2509(73)85063-8.
[110] P. Harriott, R.M. Hamilton, Solid-liquid mass transfer in turbulent pipe flow, Chemical Engineering Science 20(12) (1965) 1073-1078. https://doi.org/https://doi.org/10.1016/0009-2509(65)80110-5.
[111] G. Kocamustafaogullari, W.D. Huang, Internal structure and interfacial velocity development for bubbly two-phase flow, Nuclear Engineering and Design 151(1) (1994) 79-101. https://doi.org/https://doi.org/10.1016/0029-5493(94)90035-3.
[112] R. Kong, S. Kim, Characterization of horizontal air–water two-phase flow, Nuclear Engineering and Design 312 (2017) 266-276. https://doi.org/10.1016/j.nucengdes.2016.06.016.
[113] C.H. Gabbard, Development of a Venturi Type Bubble Generator for Use in the Molten-Salt Reactor Xenon Removal System, United States, 1972.
[114] J. Yin, Y. Ma, Y. Qian, D. Wang, Experimental investigation of the bubble separation route for an axial gas–liquid separator for TMSR, Annals of Nuclear Energy 97 (2016) 1-6. https://doi.org/10.1016/j.anucene.2016.06.018.
[115] J. Yin, Y. Qian, Y. Ma, D. Wang, Experimental study on the bubble trajectory in an axial gas-liquid separator applied for tritium removal for molten salt reactors, Nuclear Engineering and Design 320 (2017) 133-140. https://doi.org/10.1016/j.nucengdes.2017.05.015.
[116] J. Yin, Y. Qian, T. Zhang, D. Wang, Measurement on the flow structure of a gas-liquid separator applied in TMSR, Annals of Nuclear Energy 126 (2019) 20-32. https://doi.org/10.1016/j.anucene.2018.11.009.
[117] J. Yin, Y. Qian, Y. Ma, D. Wang, Y. Song, T. Zhang, H. Li, W. Liu, Numerical and experimental study on an isolated bubble in the swirling separator, Nuclear Engineering and Design 350 (2019) 107-115. https://doi.org/10.1016/j.nucengdes.2019.05.010.
[118] J. Yin, J. Li, Y. Ma, H. Li, W. Liu, D. Wang, Study on the Air Core Formation of a Gas–Liquid Separator, Journal of Fluids Engineering 137(9) (2015). https://doi.org/10.1115/1.4030198.
[119] Y. Qian, T. Zhang, J. Li, Y. Song, J. Yin, D. Wang, H. Li, W. Liu, Experimental study on the air core evolution and stability in the gas-liquid separator, 17th International Topical Meeting on Nuclear Reactor Thermal Hydraulics, NURETH 2017, 2017.
[120] Y. Qian, T. Zhang, J. Li, Y. Song, J. Yin, D. Wang, H. Li, W. Liu, Simultaneous PIV/PLIF and Pulsed Shadowgraphy Measurement of Gas-Liquid Flows in a Swirling Separator, Nuclear Technology 205(1-2) (2019) 272-280. https://doi.org/10.1080/00295450.2018.1486161.

[121] O. Simonin, P. Viollet, Modelling of turbulent two-phase jets loaded with discrete particles, Phenomena in multiphase flows (1990) 259-269.
[122] L. Schiller, A drag coefficient correlation, Zeit. Ver. Deutsch. Ing. 77 (1933) 318-320.
[123] A. Rykhlevskii, J.W. Bae, K. Huff, arfc/saltproc: Code for online reprocessing simulation of molten salt reactor with external depletion solver SERPENT, Zenodo (2018).
[124] J. Leppanen, M. Pusa, T. Viitanen, V. Valtavirta, T. Kaltiaisenaho, The Serpent Monte Carlo code: Status, development and applications in 2013, Annals of Nuclear Energy 82 (2014) 142–150.
[125] A. Ahmad, E.B. McClamrock, A. Glaser, Neutronics calculations for denatured molten salt reactors: Assessing resource requirements and proliferation-risk attributes, Annals of Nuclear Energy 75 (2015) 261–267.
[126] H.F. Bauman, G.W. Cunningham Iii, J.L. Lucius, H.T. Kerr, C.W. Craven, Jr., Rod: A Nuclear and Fuel-Cycle Analysis Code for Circulating-Fuel Reactors, 1971.
[127] C.W. Kee, L.E. McNeese, MRPP: multiregion processing plant code, 1976.
[128] C. Fiorina, M. Aufiero, A. Cammi, F. Franceschini, J. Krepel, L. Luzzi, K. Mikityuk, M.E. Ricotti, Investigation of the MSFR core physics and fuel cycle characteristics, Progress in Nuclear Energy 68 (2013) 153–168.
[129] R.J. Sheu, C.H. Chang, C.C. Chao, Y.W.H. Liu, Depletion analysis on long-term operation of the conceptual Molten Salt Actinide Recycler & Transmuter (MOSART) by using a special sequence based on SCALE6/TRITON, Annals of Nuclear Energy 53 (2013) 1–8.
[130] M. Aufiero, A. Cammi, C. Fiorina, J. Leppänen, L. Luzzi, M.E. 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 441(1–3) (2013) 473–486.
[131] D. Heuer, E. Merle-Lucotte, M. Allibert, M. Brovchenko, V. Ghetta, P. Rubiolo, Towards the thorium fuel cycle with molten salt fast reactors, Annals of Nuclear Energy 64 (2014) 421–429.
[132] J. Park, Y. Jeong, H.C. Lee, D. Lee, Whole core analysis of molten salt breeder reactor with online fuel reprocessing, International Journal of Energy Research 39(12) (2015) 1673–1680.
[133] B.R. Betzler, J.J. Powers, A. Worrall, Molten salt reactor neutronics and fuel cycle modeling and simulation with SCALE, Annals of Nuclear Energy 101(Supplement C) (2017) 489–503.
[134] B.R. Betzler, K.B. Bekar, W. Wieselquist, S.W. Hart, S.G. Stimpson, Molten Salt Reactor Fuel Depletion Tools in SCALE, Proc. of GLOBAL International Fuel Cycle Conference, American Nuclear Society, 2019.
[135] B.R. Betzler, S. Robertson, E.E. Davidson, J.J. Powers, A. Worrall, L. Dewan, M. Massie, Fuel cycle and neutronic performance of a spectral shift molten salt reactor design, Annals of Nuclear Energy 119 (2018) 396–410.
[136] J.J. Powers, T.J. Harrison, J.C. Gehin, A new approach for modeling and analysis of molten salt reactors using SCALE, Transactions of the American Nuclear Society, American Nuclear Society, 555 North Kensington Avenue, La Grange Park, IL 60526 (United States), 2013.
[137] I.C. Gauld, G. Radulescu, G. Ilas, B.D. Murphy, M.L. Williams, D. Wiarda, Isotopic Depletion and Decay Methods and Analysis Capabilities in SCALE, Nuclear Technology 174(2) (2011) 169–195.
[138] J. Leppänen, M. Pusa, Burnup calculation capability in the PSG2/Serpent Monte Carlo reactor physics code, Proc. M&C (2009) 3–7.
[139] N. Tsoulfanidis, The Nuclear Fuel Cycle, American Nuclear Society, 2013.
[140] C.J. Werner, MCNP - A General Monte Carlo N-Particle Transport Code, Los Alamos National Laboratory, 2017.

[141] A. Nuttin, D. Heuer, A. Billebaud, R. Brissot, C. Le Brun, E. Liatard, J.M. Loiseaux, L. Mathieu, O. Meplan, E. Merle-Lucotte, H. Nifenecker, F. Perdu, S. David, Potential of thorium molten salt reactorsdetailed calculations and concept evolution with a view to large scale energy production, Progress in Nuclear Energy 46(1) (2005) 77–99.
[142] K. Huff, Enabling Load-Following Capability in the TAP MSR, 2018.
[143] R.B. Briggs, Molten-Salt Reactor Program semiannual progress report for period ending July 31, 1964, 1964, pp. 397-397.
[144] D.E. Holcomb, R.A. Kisner, S.M. Cetiner, Instrumentation Framework for Molten Salt Reactors, 2018.
[145] J. McFarlane, P.A. Taylor, D.E. Holcomb, W. Poore Iii, Review of Hazards Associated with Molten Salt Reactor Fuel Processing Operations, 2019.
[146] R.B. Lindauer, C.K. McGlothlan, Design, Construction, and Testing of a Large Molten Salt Filter, Oak Ridge National Laboratory, Tech. Rep. ORNL-TM-2478 (1969).
[147] K. Sorensen, One-Fluid MSBR Chemical Processing, 2006.
[148] R.B. Briggs, Molten-salt reactor program. Semiannual progress report, 1969.
[149] S. Delpech, C. Cabet, C. Slim, G.S. Picard, Molten fluorides for nuclear applications, Materials Today 13(12) (2010) 34–41.
[150] X. Doligez, D. Heuer, E. Merle-Lucotte, M. Allibert, V. Ghetta, Coupled study of the Molten Salt Fast Reactor core physics and its associated reprocessing unit, Annals of Nuclear Energy 64(Supplement C) (2014) 430–440.
[151] M.E. Whatley, L.E. McNeese, W.L. Carter, L.M. Ferris, E.L. Nicholson, Engineering development of the MSBR fuel recycle, Nuclear Applications and Technology 8(2) (1970) 170–178.
[152] W.L. Carter, E.L. Nicholson, Design and Cost Study of a Fluorination–Reductive Extraction–Metal Transfer Processing Plant for the MSBR, 1972.
[153] A. Rykhlevskii, Advanced online fuel reprocessing simulation for Thorium-fueled Molten Salt Breeder Reactor, University of Illinois at Urbana-Champaign, 2018.
[154] The HDF Group, Hierarchical data format, version 5, (1997), https://portal.hdfgroup.org/display/HDF5/HDF5.
[155] A. Scopatz, P.K. Romano, P.P.H. Wilson, K.D. Huff, PyNE: Python for Nuclear Engineering, Transactions of the American Nuclear Society, American Nuclear Society, 2012, pp. 985–987.
[156] A. Rykhlevskii, J.W. Bae, K.D. Huff, Modeling and simulation of online reprocessing in the thorium-fueled molten salt breeder reactor, Annals of Nuclear Energy 128 (2019) 366-379.
[157] E.R. Gansner, S.C. North, E. Koutsofios, Drawing graphs with dot, AT&T Bell Laboratories (1996).
[158] J. Ellson, E.R. Gansner, E. Koutsofios, S.C. North, G. Woodhull, Graphviz and dynagraph – static and dynamic graph drawing tools, Graph Drawing Software, Springer-Verlag, 2003, pp. 127–148.
[159] H. Krekel, B. Oliveira, R. Pfannschmidt, F. Bruynooghe, B. Laugher, F. Bruhin, pytest: Python testing tool, 2004.
[160] The Travis CI API, https://github.com/travis-ci/travis-api.
[161] G. Brandl, Sphinx: Python Documentation Generator, 2009.
[162] A. Rykhlevskii, Fuel Processing Simulation Tool for Liquid-fueled Nuclear Reactors, University of Illinois at Urbana-Champaign, 2020.
[163] A. Chaube, D. O'Grady, A. Rykhlevskii, K.D. Huff, TAP MSR model for Serpent 2, Zenodo (2019).

[164] A. Rykhlevskii, K.D. Huff, Milestone 2.1 Report: Demonstration of SaltProc (UIUC-ARFC-2019-04), 2019.
[165] Transatomic Power Corporation, The Lost Moderator: The use of Zirconium Hydride in Reactor Application, 2018.
[166] M.B. Chadwick, ENDF/B-VII.1 Nuclear Data for Science and Technology: Cross Sections, Covariances, Fission Product Yields and Decay Data, Nuclear Data Sheets 112(12) (2011) 2887–2996-2887–2996.
[167] D.A. Close, IAEA Safeguards for the Fissile Materials Disposition Project, 1995.
[168] E. Sada, H. Kumazawa, C.H. Lee, H. Narukawa, Gas-liquid interfacial area and liquid-side mass-transfer coefficient in a slurry bubble column, Industrial & engineering chemistry research 26(1) (1987) 112–116.
[169] C.J. Barton, H.A. Friedman, W.R. Grimes, H. Insley, R.E. Moore, R.E. Thoma, Phase Equilibria in the Alkali Fluoride-Uranium Tetrafluoride Fused Salt Systems: I, The Systems LiF-UF4 and NaF-UF4, Journal of the American Ceramic Society 41(2) (1958) 63–69.
[170] G.J. Janz, G.L. Gardner, U. Krebs, R.P.T. Tomkins, Molten Salts: Volume 4, Part 1, Fluorides and Mixtures Electrical Conductance, Density, Viscosity, and Surface Tension Data, Journal of Physical and Chemical Reference Data 3(1) (1974) 1-115. https://doi.org/10.1063/1.3253134.
[171] S. Yamanaka, K. Yoshioka, M. Uno, M. Katsura, H. Anada, T. Matsuda, S. Kobayashi, Thermal and mechanical properties of zirconium hydride, Journal of Alloys and Compounds 293-295 (1999) 23–29.
[172] B. Forget, K. Smith, S. Kumar, M. Rathbun, J. Liang, Integral Full Core Multi-Physics PWR Benchmark with Measured Data, 2018.
[173] O. Ashraf, A. Rykhlevskii, G.V. Tikhomirov, K.D. Huff, Preliminary design of control rods in the single-fluid double-zone thorium molten salt reactor (SD-TMSR). Part I: Design development, material selection, and worth analysis, Annals of Nuclear Energy 152 (2021) 108035.
[174] U.S. Energy Information Administration, U.S. Electric System Operating Data. 2016, https://www.eia.gov/.
[175] A. Lokhov, Technical and economic aspects of load following with nuclear power plants, NEA, OECD, Paris, France (2011).
[176] Production and Removal of Xenon 135, 2020, https://www.nuclear-power.com/nuclear-power/reactor-physics/reactor-operation/xenon-135/.
[177] A. Lokhov, Load-following with nuclear power plants, NEA News 29(2) (2011) 18–20.
[178] B. Kirby, M. Milligan, Method and case study for estimating the ramping capability of a control area or balancing authority and implications for moderate or high wind penetration, 2005.
[179] Combustion Engine vs Gas Turbine- Ramp rate, 2020, https://www.wartsila.com/energy/learn-more/technical-comparisons/combustion-engine-vs-gas-turbine-startup-time.
[180] H. Ludwig, T. Salnikova, A. Stockman, U. Waas, Load cycling capabilities of german nuclear power plants (NPP), VGB powertech 91(5) (2011) 38–44.
[181] J.R. Lamarsh, A.J. Baratta, Introduction to nuclear engineering, Addison-Wesley Massachusetts, 1975.
[182] G.I. Bell, S. Glasstone, Nuclear Reactor Theory, Van Nostrand Reinhold Company, 1970.
[183] A. Rykhlevskii, D. O'Grady, T. Kozlowski, K.D. Huff, The Impact of Xenon-135 on Load Following Transatomic Power Molten Salt Reactor, Transactions of the American Nuclear Society, American Nuclear Society, 2019, pp. 1441–1444.

[184] S. Svanström, Load following with a passive reactor core using the SPARC design, 2016.
[185] A.L. Nichols, M. Verpelli, D.L. Aldama, Handbook of nuclear data for safeguards: database extensions, August 2008, IAEA, 2008.
[186] D. LeBlanc, Integral molten salt reactor, in: T.E. Inc (Ed.) 2015.
[187] E.S. Bettis, L.G. Alexander, H.L. Watts, Design Studies of a Molten-Salt Reactor Demonstration Plant, 1972.
[188] K. Sorensen, Liquid-Fluoride Thorium Reactor Development Strategy, Thorium Energy for the World, Springer, 2016, pp. 117–121.
[189] A. Rykhlevskii, K.D. Huff, SaltProc output for TAP MSR and MSBR online reprocessing depletion simulations, (2020).
[190] A. Alfonsi, C. Rabiti, D. Mandelli, J.J. Cogliati, C. Wang, D.P. Maljovec, P.W. Talbot, C.L. Smith, RAVEN Theory Manual, 2016.
[191] Hastelloy C-276 Technical Data, 2015, https://www.hightempmetals.com/techdata/hitempHastC276data.php.
[192] E. Capelli, O. Benes, R.J.M. Konings, Thermodynamic assessment of the LiF-ThF4-PuF3-UF4 system, Journal of Nuclear Materials 462 (2015) 43-53.
[193] S. Yamanaka, K. Yamada, K. Kurosaki, U. Masayoshi, K. Takeda, H. Anada, T. Matsuda, S. Kobayashi, Thermal properties of zirconium hydride, Journal of Nuclear Materials 294 (2001) 94-98.
[194] O. Nilsson, H. Mehling, R. Horn, J. Fricke, R. Hofmann, S.G. Muller, R. Eckstein, D. Hofmann, Determination of the thermal diffusivity and conductivity of monocrystalline silicon carbide (300-2300 K), High Temperatures High Pressures 29 (1997) 73-80.
[195] C. Wood, D. Emin, P.E. Gray, Thermal conductivity of boron carbides, Physical Review B 31(10) (1985) 6811 - 6814.
[196] R. Yoshioka, M. Kinoshita, 11 - Liquid fuel, thermal neutron spectrum reactors, in: T.J. Dolan (Ed.), Molten Salt Reactors and Thorium Energy, Woodhead Publishing, 2017, pp. 281–373.
[197] C.J. Permann, D.R. Gaston, D. Andr, R.W. Carlsen, F. Kong, A.D. Lindsay, J.M. Miller, J.W. Peterson, A.E. Slaughter, R.H. Stogner, R.C. Martineau, MOOSE: Enabling massively parallel multiphysics simulation, SoftwareX 11(2352-7110) (2020) 100430.
[198] A. Lindsay, G. Ridley, A. Rykhlevskii, K. Huff, Introduction to Moltres: An application for simulation of Molten Salt Reactors, Annals of Nuclear Energy 114 (2018) 530-540.
[199] A. Lindsay, K. Huff, Moltres: finite element based simulation of molten salt reactors, Journal of Open Source Software (2018).
[200] OECD/NEA, The JEFF-3.1.2 Nuclear Data Library, Tech. Rep. JEFF Report 24, OECD/NEA Data Bank, OECD/NEA, 2014.
[201] J.W. Peterson, A.D. Lindsay, F. Kong, Overview of the incompressible Navier-Stokes simulation capabilities in the MOOSE, Advances in Engineering Software 119 (2018) 68-92.
[202] R. Hu, SAM Theory Manual. Technical Report ANL/NE-17/14, Argonne National Laboratory, 2017.
[203] S. Bhattacharya, J.J. Charonko, P.P. Vlachos, Particle image velocimetry (PIV) uncertainty quantification using moment of correlation (MC) plane, Measurement Science and Technology 29(11) (2018) 115301. https://doi.org/10.1088/1361-6501/aadfb4.
[204] J.J. Charonko, P.P. Vlachos, Estimation of uncertainty bounds for individual particle image velocimetry measurements from cross-correlation peak ratio, Measurement Science and Technology 24(6) (2013) 065301. https://doi.org/10.1088/0957-0233/24/6/065301.

[205] G. Höhne, J. McNaughton, W. Hemminger, H.-J. Flammersheim, H.-J. Flammersheim, Differential scanning calorimetry, Springer Science & Business Media2003.
[206] A.G. Bergmann;, E.P. Dergunov, No Title, Dokl. Ac. Sc. URSS. 31 (1941) 753.
[207] J.-H. Cheng, P. Zhang, X.-H. An, K. Wang, Y. Zuo, H.-W. Yan, Z. Li, A Device for Measuring the Density and Liquidus Temperature of Molten Fluorides for Heat Transfer and Storage, Chinese Physics Letters 30(12) (2013) 126501. https://doi.org/10.1088/0256-307x/30/12/126501.

SELECTED NOMENCLATURE[29]

d	Bubble diameter
$\boldsymbol{u}_k$	Velocity of phase k
$\overline{\overline{\rho_k}} = \rho_k$	Density of phase k
n_i	Number density of bubble group i
μ	Dynamic viscosity
$Y_{k,i}$	Mass fraction of species i in phase k
d_{sm}	Sauter mean diameter
$\mathfrak{G}_k^T$	Turbulent stress tensor
g	Gravity acceleration
a_i, a	Interfacial area concentration
ν	Kinematic viscosity
$K_{k,i}^{\rho}$	Mass concentration equilibrium ratio for species i in phase k
$X_{k,i}$	Mass concentration of species i in phase k
K_L, k_l	Mass transfer coefficient
k	Mass transfer coefficient, vane slope or multiplication factor
$\Gamma_{l,i}$	Mass transfer of species i source term to phase k
Γ_k	Mass transfer source term to phase k
D_m	Molecular diffusivity
p	Pressure

[29] Only recurring symbols are included. Some symbols have different meaning but should be separable based on context.

$Re_l = \frac{j_l D}{\nu}$	Reynold number for pipe flow.
$Sc = \frac{\nu}{D_m}$	Schmidt number
$Sh = \frac{kD}{D_m}$	Sherwood number for pipe flow
$\overline{\overline{\boldsymbol{M_m}}}$	Source term for interfacial forces for mixture
$\overline{\overline{\boldsymbol{M_k}}}$	Source term for interfacial forces on phase k
j_g	Superficial gas velocity
j_l	Superficial liquid velocity
σ	Surface tension
T	Temperature
D_k^{t}	Turbulent diffusivity for species transport
ϵ	Turbulent dissipation rate
k^t	Turbulent kinetic energy
$Sc_t = \frac{\nu^t}{D^t}$	Turbulent Schmidt number
$\mathfrak{G}_k$	Viscous stress tensor
α_k	Void fraction of k phase
Q	Flow rate
$D_{k,i}$	Molecular diffusivity of species i in phase k
$G_{l,i}$	Mass flowrate of species i in phase k
u_{gj}	Drift velocity
β, ϵ_m	Removal efficiency, or efficiency of xenon migration
L_{xe}	Characteristic length for xenon removal

C_g/C_l	Covariances describing the distributional effect for mass transfer
D_{void}	Nominal void center diameter
D_o	Gas outlet diameter in bubble separator
D_c	Dimension related to the downstream gas outlet in bubble separator
$\hat{\eta}$	Factor describing the relative efficiency of downstream gas outlet
D	Pipe diameter
N	Number density of nuclides
$\hat{\epsilon}, \epsilon_{es}$	Bubble separation efficiency at the bubble separator.
τ	Half-life of nuclide
γ	Yield ratio of nuclear reaction
Σ_f	Macroscopic fission cross section
λ	Decay constant
ϕ	Neutron flux
H	Henry's law constant
ρ	Reactivity
α_T	Temperature coefficient of the reactivity