

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Design, Analysis, and Optimization of Breed-and-Burn Reactor Cores

Permalink

<https://escholarship.org/uc/item/8xt8m141>

Author

Keckler, Christopher Thomas

Publication Date

2020

Peer reviewed|Thesis/dissertation

Design, Analysis, and Optimization of Breed-and-Burn Reactor Cores

by

Christopher T Keckler

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering – Nuclear Engineering

and the Designated Emphasis

in

Computational and Data Science and Engineering

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Associate Professor Massimiliano Fratoni, Co-chair
Professor of the Graduate School Ehud Greenspan, Co-chair
Professor Alper Atamtürk

Spring 2020

Design, Analysis, and Optimization of Breed-and-Burn Reactor Cores

Copyright 2020
by
Christopher T Keckler

Abstract

Design, Analysis, and Optimization of Breed-and-Burn Reactor Cores

by

Christopher T Keckler

Doctor of Philosophy in Engineering – Nuclear Engineering
and the Designated Emphasis in
Computational and Data Science and Engineering

University of California, Berkeley

Associate Professor Massimiliano Fratoni, Co-chair

Professor of the Graduate School Ehud Greenspan, Co-chair

Breed and Burn (B&B) reactors are a special class of fast reactor that promises enormous sustainability and proliferation-resistance benefits compared to current commercial nuclear power systems. Utilizing a very tight neutron economy in the fast spectrum, B&B reactors are able to sustain criticality in the equilibrium state without either fuel enrichment or reprocessing. As with many advanced reactor systems, however, B&B reactors currently suffer from a number of uncertainties and infeasibilities hampering their practical deployment. This dissertation aims to better understand and propose solutions to some of these primary uncertainties. Utilizing state-of-the-art computational tools and techniques, the properties of a prototypical B&B core are investigated and quantified in high fidelity, enabling deeper understanding of the physical phenomenon driving their unique characteristics. Subsequently, three of the most pressing current uncertainties in B&B cores are investigated in detail: (1) reactor orificing, (2) transient performance and inherent safety, and (3) the tradeoff between material integrity and resource utilization.

To understand the feasibility of adequately cooling fuel assemblies over their in-core lifetimes, new mathematical optimization techniques are developed and applied to devise optimal orificing strategies subject to a variety of important operational constraints, the results of which demonstrate how certain methods of orificing are better suited to B&B operation than others. To understand B&B transient performance, the first open-literature transient models are developed and analyzed for a wide variety of hypothetical accident scenarios, finding challenges associated with the unfavorable reactivity feedback mechanisms inherent to B&B core designs. In order to improve transient performance, a novel passive safety system is incorporated into the core and its behavior simulated using a newly-coupled code system. Finally, the topic of material damage is explored by first using newly-proposed methods for uncertainty quantification on the material damage levels. Once the uncertainties have been understood, the study moves to exploring the tradeoff between material damage

and resource utilization when low-enriched uranium is used as a feed fuel. In order to improve this tradeoff, new two-tiered reactor systems are designed and analyzed which enable significantly higher resource utilization than current light water reactor systems without requiring fuel reprocessing or physical/chemical separations. Each of these explorations allows for some of the uncertainties around B&B reactors to be clarified, and this thesis therefore enables improved viability of implementing the breed-and-burn mode of operation as part of future nuclear power systems.

Contents

Contents	i
1 Introduction	1
1.1 Motivation for B&B Reactors	1
1.2 Previous Work and the State of the Art	2
1.3 Outstanding Issues and Challenges	7
1.4 The Structure of this Dissertation	11
2 Overview of the Reference Core	12
2.1 Design Strategy	12
2.2 The Design Process Using ADOPT	13
2.3 Converting the Reference Core to Hexagonal Geometry	17
2.4 Overview of Core Performance Characteristics	27
3 Assembly Orificing	36
3.1 Introduction to the Orificing Problem	36
3.2 Assessment of Static Integral Orificing	41
3.3 Assessment of Static Grid Plate Orificing	43
3.4 Assessment of Variable Orificing	53
3.5 Summary of Orificing Strategies	55
4 Nominal Transient Performance	57
4.1 Transient Modeling Approach	57
4.2 Nominal Performance	62
4.3 Summary of Nominal Transient Performance	74
5 ARC System Incorporation	75
5.1 ARC System Overview	75
5.2 ARC System Performance	80
5.3 Introduction of Modified ARC System	95
5.4 Modified ARC System Performance	102
5.5 Other Impacts of ARC System Inclusion	123

5.6	Technology Gaps	138
5.7	Summary of ARC System Performance in the B&B Core	142
6	Material Damage and B&B Reactors	145
6.1	Review of B&B Core Design Strategies with a Focus on DPA	145
6.2	Nuclear Data Uncertainties in Quantified DPA	151
6.3	Enriched B&B Reactors	163
6.4	Double-Cladding of eB&B Discharged Fuel	174
6.5	The eB&B/SHR System	187
6.6	Potential for Three-Stage Fuel Cycles	213
6.7	Fuel Cycle Performance Intercomparison	216
6.8	Summary of the eB&B Two-Tiered Systems	222
6.9	Future Possibilities for eB&B Cores	223
7	Conclusions	225
	Bibliography	229
A	Quantifying Errors Introduced by r-z Models	241
A.1	Inherent Approximations of r-z Models	241
A.2	Previous Investigations	242
A.3	Comparison Procedure	243
A.4	Finding an Optimal hex-z Approximation	244
A.5	Results	248
A.6	Summary of r-z/hex-z Comparison	254
B	Reduction of Core Volume	255
B.1	Reduction in the Number of Fuel Batches	255
B.2	Elimination of Radial Reflector Assemblies	260
B.3	Summary of Reduced Volume Performance	261

Acknowledgments

I would like to acknowledge all who have helped me over the course of my graduate work, to whom I owe my gratitude. My time during graduate school has been full of helpful, thoughtful, and kind people who have offered advice, feedback, and direct help on both my immediate work and on my future career. In lieu of a listing of such people, which is both tacky and prone to accidental omissions, I will trust that those who belong on this list know who they are and understand the extent to which I am grateful. If in doubt as to whether or not you belong on this list, chances are that you do. I truly owe each of you a debt which I can only hope to pay back throughout my career, something that I sincerely plan to do.

I am also obliged to acknowledge the organizations who have funded me throughout my time in graduate school. The Department of Energy Office of Nuclear Energy (DOE) supported me for the first three years under their Nuclear Energy University Programs (NEUP) Fellowship, while the final two were funded by the Nuclear Regulatory Commission Graduate Fellowship. In addition, the grant from DOE NEUP to support the investigation of the ARC system provided me with funds to attend conferences and spend a summer at Argonne National Laboratory, both of which were instrumental in my development as a nuclear engineer.

Preface

Before embarking on this dissertation, a few words concerning the contributions of others are in order. The work in this dissertation is based upon the work of many others who have come before me in the Department of Nuclear Engineering at the University of California, Berkeley, mainly those who have worked under Professor Greenspan on B&B design. Out of the many, two names deserve particular recognition: Florent Heidet and Staffan Qvist. Both Florent and Staffan laid the vast majority of the groundwork which resulted in the reference core outlined in Chapter 2, without which I would have had little to base the remainder of my work upon. Staffan also was the original inventor of the ARC system and was responsible for conceptualizing the methods with which to simulate its response, which is the focus of much of Chapter 5. This foundation enabled me to pick up the ARC system project quickly, to perform the various design studies, and to make significant contributions of my own to further improve the ARC system design.

In addition, three visiting students completing their Masters program theses have also made significant contributions to the concepts explored in this dissertation. Yves Robert, from the Institut National des Sciences et Techniques Nucleaires, Cadarache, worked with me on the optimization of reactor orificing, the focus of Chapter 3. After I conceptualized the optimization methods and laid out the techniques and plan for the study, Yves performed many of the numeric calculations and made significant theoretical contributions which have enhanced the quality of those analyses. His work is closely intertwined with my own, and for this reason some of what is presented in that chapter may rightfully be claimed jointly by both of us. Similarly, Carlo Oggioni from the Polytechnico di Milano worked with me on analysis of the modified ARC system after its design and conceptualization was completed by myself. Taking the numerical methods and plan from me, Carlo performed many simulations and did much analysis to understand the fundamental nature of the modified ARC system, which in many instances surprised us, and helped to guide the work presented in this thesis. Although they resemble Carlo's in many ways due to his fundamental role in helping to guide the work, ultimately the results presented in in Chapter 5 are my own work. Finally, Robert Mardus-Hall from the University of New South Whales was the first to make the suggestion that a Tesla valve, or something with similar characteristics, could potentially benefit the ARC system. This suggestion led to a huge thrust of effort related to resolving a number of issues in the original ARC system, the topic of Chapter 5, and was a key insight that helped me to understand many of the unique properties of the ARC system in general.

Chapter 1

Introduction

1.1 Motivation for B&B Reactors

The need for vast amounts of energy to meet increasing worldwide demand is firmly established, borne out by energy and population trends over the past century (see, for instance, [121, 130]). The requirement that this energy be carbon-free to simultaneously address climate change has also been firmly established within the international community [139]. However progress on this front has been markedly slower to materialize, with new capacity from coal and natural gas power stations coming online faster than all carbon-free sources combined [41] and the transportation sector remaining essentially completely reliant on gasoline and diesel [32]. The potential role of nuclear energy in this massive energy transition has been expounded upon at length by many authors, with the general consensus that nuclear energy will be needed in future clean energy systems to complement other sources of carbon-free energy such as hydro, wind, and solar (see, for instance, [88, 107, 41]). As such, this dissertation takes it as a necessity that nuclear energy must be maintained *at least* at current capacity levels throughout the world, if not significantly expanded.

However, much of the United States' current nuclear generating capacity is set to go offline in the coming decades as older plants reach their operating lifetimes. By the year 2040, 48% of currently operating power plants in the nation will have reached the end of their lifetimes (assuming all of them receive and fully utilize 20 year extensions on their original licenses). While the emerging trend towards a *second* license extension would delay this massive loss of generation capacity [3], even still only two power reactors are under construction in the United States and only a handful more are even under consideration at the time of writing. In addition, recent experience with building new reactors, particularly in the United States and Europe, has demonstrated the economic challenges associated with continued reliance on large Light Water Reactors (LWRs), which have dominated both historic and recently added nuclear capacity. Such economic challenges call into question the viability of continuing to build LWRs and motivate the exploration of alternative systems. The central question concerning future nuclear power systems is thus – what form should

the world’s future nuclear energy infrastructure take?

Many alternative answers to this important question have been proposed. While the worldwide nuclear energy industry has remained focused on LWRs, the research community has been increasingly looking towards non-LWRs, or radically different LWRs in the case of light-water cooled Small Modular Reactors (SMRs) (see, for instance, [120]), to improve aspects such as economics, sustainability, and safety. A broad overview of the major non-LWR systems under investigation can be found in the literature of the Generation IV International Forum (GIF) [111]. Among the six broad categories proposed by GIF, three of them are focused specifically on fast-spectrum systems (the Sodium Cooled Fast Reactor (SFR), Gas Cooled Fast Reactor (GFR), and Lead Cooled Fast Reactor (LFR)), with at least two of the other categories having fast-spectrum variants (the Molten Salt Reactor (MSR) and Supercritical Water Cooled Reactor (SCWR)). This focus on the fast spectrum reflects the widely known operational benefits enabled by the use of fast neutrons, namely the ability to breed fissile fuel from low-grade material and to efficiently burn higher actinides resulting from thermal-spectrum systems (see, for instance, [149]). These potential benefits have been known for decades, having been most recently articulated in the exhaustive DOE Fuel Cycle Screening and Evaluation study which recommended fast-spectrum reactors with continuous recycle as the most promising option for improving the sustainability of the nuclear fuel cycle [152]. Such a result was hardly surprising, since the gains in fuel utilization and reductions in waste streams have been understood since the very early days of nuclear power dating all the way back to Enrico Fermi and the decision to build EBR-I and EBR-II [150].

That being said, the fact that continuous recycling has never been fully integrated into the civilian nuclear industry clearly implies that there are significant barriers to implementation. Two such barriers have been widely understood to be the cost of recycling fuel and the proliferation concerns associated with it. One fuel cycle, evaluated as part of the DOE study to be ‘potentially promising’ but of lower benefit than the continuous recycle systems, is the use of fast critical reactors to drive natural uranium to very high burnup [152]. In these systems, referred to in this dissertation as Breed-and-Burn (B&B), reasonably high fuel resource utilization can be achieved *without* the need for recycling or enrichment in the equilibrium state, enabling a highly proliferation resistant fuel cycle with large improvements as compared to current LWR systems. These B&B systems are the broad focus of this dissertation. In this chapter, the fundamental physics behind B&B cores is outlined. Future chapters are aimed at either designing, optimizing, or analyzing B&B cores to better understand their unique properties and improve their feasibility for eventual commercial implementation.

1.2 Previous Work and the State of the Art

B&B reactors are a special class of fast reactor that holds promise for greatly simplifying the fuel cycle and improving proliferation resistance while achieving high resource utilization as compared to current LWRs [43]. Using a very hard neutron spectrum and a tight neutron

economy, B&B cores breed enough fissile material from low-grade feed fuel to sustain criticality without the use of fuel recycling or enrichment in the equilibrium state. While fuel is in the breeding state, it absorbs excess neutrons leaking from fuel that is in the burning state. Eventually fuel in the breeding state builds up enough fissile material to become a net producer of neutrons, at which point the fuel transitions into the burning state and drives breeding in other feed fuel. As operation continues, the buildup of fission products and the depletion of fuel material limits the amount of time that fuel can remain in the burning state, after which it again becomes a neutron sink and should be discharged from the reactor. This leads to an upper bound in terms of burnup around 50-60%FIMA, beyond which some form of fuel conditioning must be employed for continued use. This lifecycle can be elegantly explained using the neutron balance concept shown in Equation 1.1 where N is the number of excess neutrons, N_{HM} is the number density of heavy metal, $\bar{\nu}$ is the average number of neutrons produced per fission, k is the multiplication factor for the system, and BU is burnup represented in %FIMA.

$$N = N_{HM} \int_0^{BU} \bar{\nu}(BU') \left(1 - \frac{1}{k(BU')} \right) dBU' \quad (1.1)$$

Using Natural Uranium (NU) or Depleted Uranium (DU) as feed fuel, a typical neutron balance might look something like Figure 1.1, where k has been evaluated using an infinite unit cell model of a representative B&B assembly. As burnup begins, the neutron balance immediately turns negative due to the low fissile content and the absorption of neutrons in ^{238}U . As these initial neutrons are absorbed, however, they convert ^{238}U into ^{239}Pu through the well known breeding process after undergoing two β^- decays.¹ This breeding eventually causes there to be enough fissile material such that the neutron balance turns around and begins climbing again. If adequately designed, the neutron balance will eventually cross zero and turn positive at the ‘minimum required burnup,’ denoted as BU^* , at which point the fuel has given back as many neutrons to the reactor as it initially absorbed. This value is paramount because it is the absolute minimum burnup level which is required for a critical reactor to be possible, and thus BU^* represents the criterion that must be satisfied in order to operate in the breed-and-burn mode of operation. In a well-designed B&B system, $BU^* \approx 20\%FIMA$.

After achieving the minimum required burnup, the neutron balance may continue to increase as more fissile material is bred. Eventually, however, due to the buildup of fission products and the burnout of heavy metal content, the neutron balance will again turn around until hitting zero for a third time at BU^{max} . Burnup beyond this level will lead to the inability to achieve criticality, and thus this point represents the upper bound on burnup which can be achieved in a system without an external neutron source. In a typical B&B core, $BU^{max} \approx$

¹While in principle the same concept could utilize ^{232}Th instead of ^{238}U , two aspects makes this alternative more difficult to successfully achieve: (1) the lower number of neutrons produced per fission, and (2) the fast fission cross-section of ^{232}Th is both smaller and has a higher threshold than the same for ^{238}U .




Figure 1.1: Example neutron balance for a generic B&B assembly.

60-70%FIMA, although it is never achieved in practice due to concerns related to irradiation damage of structural materials in the core.

Without having sufficiently low losses and a hard enough spectrum², the neutron balance may never become positive making it such that breed-and-burn cannot be achieved without an external source of neutrons. This concept is demonstrated in Figure 1.2 where the neutron leakage fraction is increased³. As the fraction of wasted neutrons increases, two important changes occur: (1) BU^* increases, and (2) BU^{max} decreases. This makes it such that it is both more difficult to achieve self-sustaining breed-and-burn operation and attainable burnups are inherently lower, driving B&B core designers to minimize wastage of neutrons in the system above all else.

From examining the neutron balance, it should be obvious that it will not be possible to achieve a critical reactor using only NU or DU without an initial supply of higher reactivity fuel. Instead, as with any fast reactor system, enriched fuel needs to be supplied for the initial startup core(s) to supply the initial deficit of excess neutrons. Thus while B&B cores can be sustained without the use of fuel enrichment in the *equilibrium* state, they will require enrichment infrastructure during the initial deployment phase. It is noted, however, that schemes have been devised where excess neutrons from one B&B core are able to breed the fissile material required for the startup of further B&B cores, a technique known as ‘spawning,’ so that enriched fuel is needed only for the very first B&B core [43, 144]. Thus the breed-and-burn fuel cycle is viewed as enhancing proliferation resistance because it

²The harder spectrum enables the highest possible value of η , as discussed further in Chapter 6.

³similar results apply for increasing proportions of neutrons lost to any source, including parasitic losses in structural materials or control elements.




Figure 1.2: Neutron balance for a generic B&B assembly as the proportion of leaking neutrons increases.

would enable a gradual phase-out of enrichment capabilities from the civilian nuclear power industry. Additionally, the ability to operate without enrichment should provide direct reductions in fuel cost.

While the lifetime of fuel can be limited by the neutron balance, typically damage to core structural materials leads the designer to discharge fuel at its minimum required burnup, closer to 20-30% FIMA. Assuming the use of NU feed fuel, a burnup of 20% FIMA corresponds to a heavy metal utilization of 20%. In comparison to a typical Pressurized Water Reactor (PWR), which reaches only around 0.6% heavy metal utilization [79], B&B reactors represent an enormous improvement from the perspective of resource utilization. While the B&B fuel cycle is not as resource efficient as a continuous recycle strategy, which could theoretically reach near 100% resource utilization, the fact that fuel recycling is not required is a major benefit that again promotes proliferation resistance in addition to reducing costs.

Within the class of solid-fueled B&B cores (alternative liquid fuel systems are discussed in Chapter 6), there are generally two variants of the breed-and-burn concept – the standing-wave and the traveling-wave. Early on in B&B development, the traveling-wave variant saw heavier interest and a more thorough theoretical development (see, for instance, [8, 28, 125, 102, 129, 122, 26]). The traveling-wave variant functions by keeping fuel stationary as a breed-burn fission wave propagates through the fuel with time, as depicted in Figure 1.3. As fertile material ahead of the wave is converted to fissile fuel through breeding, the fission

wave progresses forward in a self-sustaining and self-regulating process that leaves spent fuel in its wake. Much theoretical work on the stability and propagation speeds of such fission waves has been developed.

However the traveling-wave variant suffers from a number of inherent drawbacks, despite its operational elegance. For one, the fuel in a traveling-wave reactor needs to be extraordinarily tall in order to reduce leakage as much as possible while allowing space for the fission wave to propagate with time, if its full potential is to be achieved. This tall fuel induces a pressure drop much higher than a typical fast spectrum system that challenges operational feasibility. More importantly, though, is the inefficient use of neutrons that results from axial propagation of the fission wave. As the fission wave progresses, half of all axially-leaking neutrons leak in the direction of the spent fuel wake, where they may cause fissions but are highly underutilized from the perspective of breeding. The result is a minimum required burnup that is typically $>40\%$ FIMA, which is probably the biggest reason why the traveling-wave variant is viewed as less feasible than the standing-wave variant.




Figure 1.3: Diagram of the traveling-wave variant of breed-and-burn operation.

The standing-wave variant, on the other hand, requires fuel to be moved as the fission wave stays approximately localized in space, as depicted in Figure 1.4. While this complicates operations, as it requires a large amount of fuel assemblies to be periodically shuffled, it overcomes both of the primary difficulties of the traveling-wave variant while staying within the paradigm of traditional SFR core design. Therefore, it is almost universally accepted that the standing-wave variant is the more feasible between the two, and throughout this dissertation it should be assumed that B&B refers to the standing-wave variant unless otherwise specified.

Considerably less theoretical development has focused on the standing-wave variant, with instead the vast majority of work being focused on direct simulations and design optimization. Included in the long line of B&B core studies are those from the University of Cali-

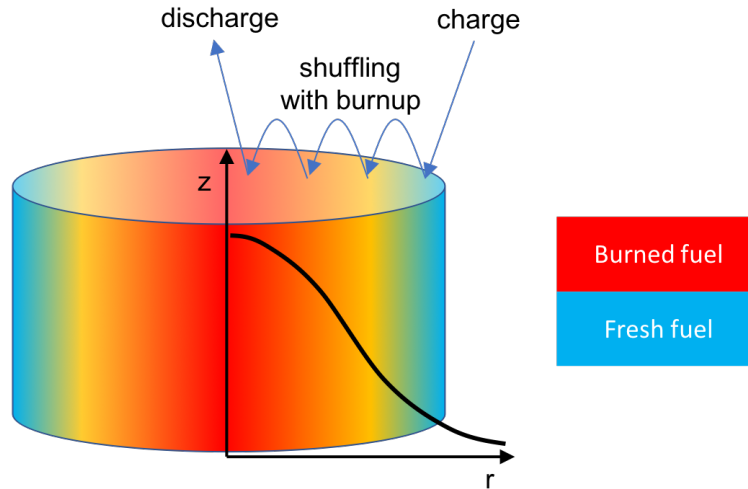


Figure 1.4: Diagram of the standing-wave variant of breed-and-burn operation.

fornia, Berkeley [54, 29, 114, 156], the private venture TerraPower, LLC [57, 39, 144], the Massachusetts Institute of Technology [110, 155], and the Tokyo Institute of Technology [59, 84], among others [158, 159, 50, 49, 52]. Excellent summaries of the lineage of B&B designs can be found in, for instance, [87] or [119], and the reader is directed there for further details on the historical development. Through this long lineage, the B&B concept has reached a point of relative maturity from the perspective of neutronics and design strategies, although a variety of outstanding issues remain, as discussed in Section 1.3. The state of design maturity is discussed in further detail in Chapter 2, where the reference core of this dissertation is outlined.

1.3 Outstanding Issues and Challenges

While the B&B concept has seen huge amounts of development over the past couple decades, both from academic and industrial researchers, a number of significant challenges remain. The most widely discussed and probably largest knowledge gap concerns irradiation damage from fast neutrons. Because B&B operation using NU or DU feed fuel requires relatively high levels of burnup, the associated damage to in-core structural materials (particularly the fuel cladding) is correspondingly high – in the vicinity of 500 displacements per atom (DPA) [65]. The highest damage levels experimentally quantified to date in a fast neutron field were in the Fast Flux Test Facility, where the ferritic-martensitic steel HT9 was taken just in excess of 200 DPA [143]. However, since that date, the United States (as well as many countries friendly to the United States) has lost its operating fast reactors, thus putting on hold any plans to further extend the experimental database for material damage under

fast fluence. While researchers at TerraPower have concentrated effort on extending the experimental database using ion irradiations [48], ultimate validation of any ion beam or theoretically-derived results will be impeded by the unavailability of sustained, high levels of fast fluence. While the United States currently is in the early phases of developing a new fast-spectrum test reactor [2], validation of material damage up to 500 DPA will need to wait not only for the facility to be licensed and built, but also for the required fluence to be accumulated, something that may take on the order of a dozen years and is not guaranteed to be successful. This discrepancy between the minimum cladding damage required by breed-and-burn operation and the experimental database is viewed as one of the most important obstacles currently limiting the feasibility of B&B reactors.

Other challenges related to materials are issues of swelling and creep in both fuel and structural components. From the perspective of fuel, the very long residence times and accumulated burnups can lead to excessive swelling and pressure buildup in the fuel pins. Typically this can be accommodated through techniques which essentially make the design less aggressive, for instance using a lower fuel smear fraction or utilizing thicker cladding with longer gas plena, and indeed burnups on the order of 20% have been achieved in other SFR designs using these techniques. The primary way in which B&B cores are different from other high-burnup SFRs is that each design choice which makes the overall design less aggressive also makes the possibility of breed-and-burn operation more difficult to achieve. Whereas in typical SFR designs the enrichment level can be slightly increased to compensate for less-aggressive design choices, for B&B this is not an option and instead less-aggressive designs are translated into higher minimum required burnups which make the material damage problems even worse. Thus, there is a delicate balance between design choices to accommodate the high levels of burnup in B&B cores.

Similarly, the understanding of metallic fuel and cladding swelling at such high levels of burnup is subject to uncertainties [142, 48, 62] which combine such that reasonably large uncertainties exist in predictions of important core characteristics. It has been demonstrated that the phenomenon of axial fuel swelling in B&B cores has a large impact on criticality [52] which will be directly translated into increases in the minimum required burnup. Although not demonstrated, it is expected that the same phenomenon will drive uncertainties in other important core parameters such as the reactivity feedback coefficients and the control system worth. These uncertainties force the designer to account for them with margins, which again drive up the minimum required burnup and are counter to the breed-and-burn mission.

Other important challenges associated with the novel aspects of B&B cores center on thermal performance during both steady-state and transient conditions. A fundamental part of the thermal design of any reactor core is ensuring that each assembly will have enough flow allocated to it so that it can maintain sufficient margin to its thermal design limits. Such a task of ‘reactor orificing,’ while not always straightforward, is generally not exceptionally difficult to accomplish because it is aligned with the widely-employed design philosophy that the radial power distribution should be made as flat as realistically possible. However it will be shown that such radially-flat power profiles are in direct contradiction with breed-and-burn operation such that B&B cores have highly-peaked power profiles. This fact, coupled

with the necessity that all assemblies have the same geometric design so that they can be shuffled freely throughout the core, makes it uncertain as to whether an orificing strategy can be proposed that will conform to constraints related to manufacturing, operational complexity, degradation of above-core structures, and plant thermodynamic efficiency, as well as others [57, 155]. While this has never been shown to be an infeasibility for B&B cores, it is an area of uncertainty which has yet to be resolved in the community.

Additionally, even if acceptable steady-state thermal-hydraulic performance can be achieved, the transient performance and stability of B&B cores has been subject to widespread question due to reactivity feedback mechanisms which are severely degraded by the breed-and-burn mode of operation in a few ways [117]. The most strongly degraded reactivity feedback mechanism is the coolant void reactivity, which is driven to be strongly positive due to the low leakage and hard spectrum required to sustain breed-and-burn operation. While on the one hand the coolant void reactivity feedback has become strongly positive, other negative feedback mechanisms with the potential to keep the system stable are reduced in either magnitude or certainty. For instance, the very hard spectrum required to obtain sufficient breeding reduces the effectiveness of the doppler broadening mechanism by allowing very little overlap between the flux spectrum and the resonance region in U-238 (the isotope primarily responsible for driving the Doppler broadening effect [149]), as shown in Figure 1.5.

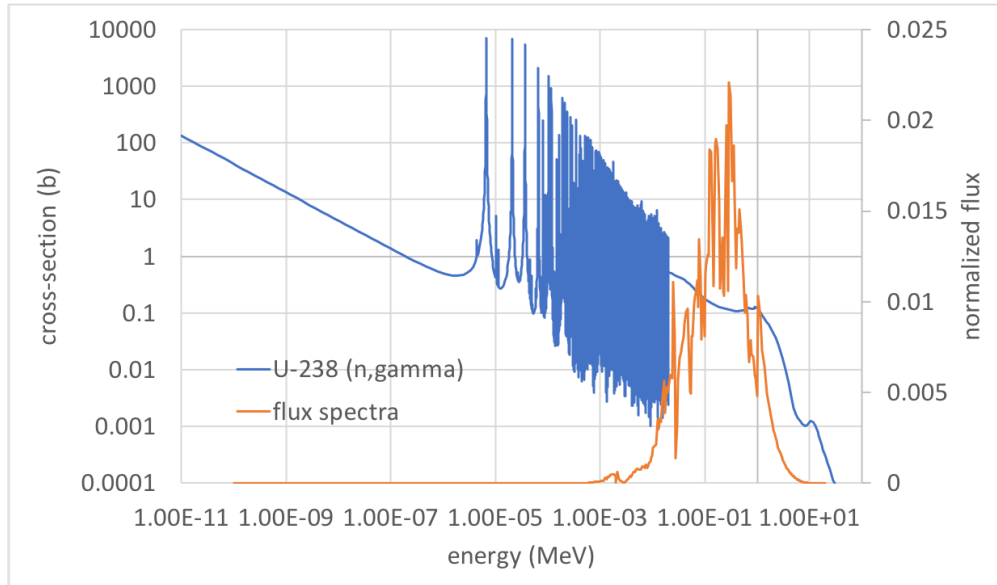


Figure 1.5: U-238 (n,γ) cross-section along with a typical neutron spectrum for a B&B core.

Other examples of degraded or uncertain negative feedback mechanisms include both the axial expansion and control rod driveline feedbacks. In the case of the axial expansion feedback, this may become uncertain at high burnups due to the high levels of swelling,

porosity, and cracking in highly irradiated fuel (though this is typically more of a concern in ceramic fueled cores) or the potential for the fuel to have mechanically bonded with the cladding, both of which may impede fuel thermal expansion making it so the axial expansion mechanism is not entirely reliable [149]. In the case of the control rod driveline, again excessive fuel swelling plays a role in making the distribution of fuel somewhat uncertain and more dilute than in fresh fuel so that an expansion of the control rod drivelines has less marginal impact. Further, the reactivity swing of B&B cores is intentionally low which reduces the use of control rods during steady-state operation and thus limits the worth of the control rod driveline feedback mechanism in general. Even when control rods are inserted, their marginal sensitivity is characteristically low in the axial periphery due to the unique axial distribution of fissile material that builds up as a result of the breed-and-burn process, reducing this feedback mechanism for much of the cycle. These various factors contribute to potentially making the transient response of B&B cores worse than might be seen in other SFRs. With the exception of the very brief transient results presented by TerraPower [57], to the author's knowledge no other B&B transient evaluations have been produced in the open literature, representing an uncertainty in the feasibility of commercially implementing these systems.

One of the last major outstanding issues related to the commercial implementation of B&B reactors is the transition to equilibrium and its impact on the economics and sustainability of breed-and-burn operation. The vast majority of studies that quantify the benefits of breed-and-burn operation focus on the equilibrium operational regime. In such a regime, the benefits as compared to other commercial systems are well understood. However, a newly-built fast-spectrum reactor cannot be critical on such low-grade fuel as NU, and thus a source of higher worth fuel is needed during the initial startup phase of every B&B core. Using Enriched Uranium (EU) for each startup core will have profound effects on the life-cycle resource utilization of breed-and-burn operation to the point that the benefits over a PWR will be severely degraded, and thus the spawning strategy mentioned previously is imperative to the overall attractiveness of B&B reactors. This has two implications: (1) in order to derive any benefits from breed-and-burn operation, many more than just one B&B reactor must be built and operated in coordination to supply the necessary high-grade discharge fuel that is needed to start up new cores, and (2) the integrity of fuel assemblies must be assured for fluences *even higher* than the minimum required burnup of a single core. Both of these present significant challenges which have so far been largely taken for granted by the (academic) breed-and-burn community, but will be vitally important for any realistic system. In addition, the need for EU enriched to levels higher than 5w% will be a challenge from both policy and infrastructure perspectives, although this is an issue common to nearly all SFR designs, and as such some progress is being made on this front from other commercial entities interested in High Assay Low Enriched Uranium (HALEU) [94].

1.4 The Structure of this Dissertation

It has been briefly highlighted that B&B reactors offer promise for improving the economics, sustainability, and proliferation resistance of the overall nuclear energy landscape. It has also been outlined that there are a number of outstanding issues that may ultimately impede their commercial implementation. The focus of this dissertation is to try to address a number of these issues within the context of various computational studies. While not all outstanding challenges are addressed, this dissertation attempts to provide novel solutions to some of the most pertinent problems with the hopes of improving the commercial viability of B&B reactors.

Beyond this chapter, the dissertation begins in Chapter 2 by providing an in-depth discussion of the reference B&B core utilized in the remaining chapters. The design strategy is discussed in relation to current trends in B&B core design and a number of important characteristics are presented, although additional details are further discussed in later chapters for cases where this would be more appropriate.

Chapter 3 then takes a fresh look at the problem of B&B reactor core orificing, proposing new methods and performing various analyses to understand the limits and possibilities in achieving adequate cooling during steady-state operation. Chapters 4 and 5 then focus on the performance of B&B cores during transients. Chapter 4 looks at both protected and unprotected transients in the reference core and quantifies some of the sensitivities on that transient performance. Chapter 5 extends that work to incorporate a novel passive safety system into the B&B design, examining how it impacts inherent safety and quantifying its impacts on steady-state operations.

Moving away from transient performance, Chapter 6 finally focuses on various aspects related to material damage in B&B cores. First, the uncertainty in the predicted DPA derived from underlying nuclear data uncertainties is quantified for the first time using a newly proposed methodology. Afterwards, the focus turns towards understanding the trade-offs between resource utilization and material damage in B&B cores with slightly enriched feed fuel. To improve the tradeoff between resource utilization and material damage while remaining within current material damage constraints, two novel methodologies for further utilizing the discharged fuel are proposed and investigated from the perspective of neutronic design. Finally, some conclusions and ideas for future work are proposed in Chapter 7.

Chapter 2

Overview of the Reference Core

The majority of studies in this dissertation utilize a reference core as a baseline for comparisons and to demonstrate various methods or design techniques. The same reference core is utilized across chapters, when possible, to avoid confusion. This chapter serves as an introduction to the reference core, and should be consulted as the basis for all statements made henceforth throughout this dissertation unless otherwise specified.

2.1 Design Strategy

In general, the reference core is a prototypical standing-wave B&B design. It has been designed in an aggressive manner to achieve the lowest possible material damage within the context of a ‘linear-assembly’ design, as defined by Petroski et al. [110] (i.e. the core is composed of fuel assemblies which are made up of classical fuel pins arranged in a lattice). Metallic U-10Zr and sodium are used as fuel and coolant, respectively, as these have been proven many times over to enable the highest breeding ratio and lowest minimum required burnup in large, high power-density cores (see, for instance, [114, 54]). A smear density of 75% is assumed for the fuel to accommodate swelling with burnup, and DU (0.25 w% U-235) is assumed as the feed material. The ferritic-martensitic steel HT9 is assumed for all in-core structural materials, due to it having the largest experimental database under fast neutron fluence. HT9 has been irradiated up to 200 DPA in FFTF [143], exhibiting acceptable swelling behavior and limited degradation of mechanical properties, and is thus a common choice for high burnup cores.

The active core height is taken to be 300 cm so as to reduce axial leakage as much as possible while maintaining an acceptable pressure drop. The batch-wise shuffling scheme, shown in Figure 2.1, is derived from that of Heidet [55]. It is designed to be near-optimal for a large-core B&B design in terms of reducing the minimum required burnup. Such a shuffling scheme simultaneously promotes efficient breeding while maintaining acceptable criticality behavior by minimizing radial leakage and directing excess neutrons inwards towards low-grade fertile fuel at the core center. This shuffling scheme results in confining the majority

of power production to an annular ring towards the radial mid-region of the core, as outlined in greater detail in Section 2.4.




Figure 2.1: Shuffling scheme employed in the reference B&B core.

With just a handful of exceptions, the design employs only proven technology. The first major exception is the assumption that current metal alloys will be able to withstand the irradiation damage required by the breed-and-burn mode of operation while maintaining structural integrity. This is not strictly a design *choice*, but rather a consequence of the current limited state of knowledge for material behavior under fast neutron irradiation combined with the requirements imposed by the minimum required burnup. Therefore this assumption is necessary to even consider the possibility of breed-and-burn operation in the first place.

The second possible exception is the potentially questionable assumption of fuel venting, which has accumulated considerably less operational experience than the typical large fission gas plenum designs employed in most SFRs. Nevertheless, fuel venting was employed in the Dounreay Fast Reactor (DFR) for many years [92]. It is also being pursued by TerraPower, LLC in their B&B design [57] and for the control rods in the China Experimental Fast Reactor (CEFR) [128], and is thus at least partly accepted by the SFR community. The assumption of fuel venting is not so much an unknown from the perspective of whether or not it *can* be done, but rather from the perspective of how it *should* be done and what the consequences will be on reactor operations. The benefits derived from fuel venting enable enhanced performance from a number of perspectives which greatly increase the feasibility of B&B cores, including: (1) the elimination of the fission gas plenum and its associated height, enabling lower pumping power requirements and a tighter lattice pitch; (2) the reduction of the necessary cladding thickness which reduces parasitic absorptions; and (3) the removal of volatile fission products which would otherwise serve as neutron poisons in high-burnup fuels. Further discussion of fuel venting in relation to B&B core performance is provided in [119].

2.2 The Design Process Using ADOPT

The majority of calculations used to design the reference B&B core are performed and driven by the SFR design code ADOPT [118], written in Python and specifically tailored to the

design of B&B cores. The key insight to the usefulness of ADOPT is the relatively small design space available for the breed-and-burn mode of operation. Because breed-and-burn operation requires an extremely tight neutron economy to even be feasible, nearly all design decisions *must* be made in the direction that benefits the neutron economy, with all other considerations (i.e. transient performance, material damage, operational complexity, etc.) viewed as secondary. This makes the core design process rather deterministic and thus amenable to an automated procedure, which is precisely the role that ADOPT was designed to fill.

ADOPT takes a limited number of user-inputs, such as power level, guesses for the radial and axial peaking factors, and assembly performance constraints, and performs design iterations over two levels, as depicted in Figure 2.2. The first inner level iterates over certain assembly design parameters, such as pin diameter and pitch, using coarse-level thermal-hydraulic calculations to determine a design which maximizes the fuel volume fraction, thus ensuring a core with high breeding ratio. After converging the assembly dimensions, ADOPT takes other user-inputs, such as the number of assemblies and the batch-wise shuffling scheme, and drives whole-core Monte Carlo coupled transport/depletion iterations in Serpent [86]. This second set of inner iterations is repeated while altering the cycle length to achieve a critical equilibrium cycle. The equilibrium power profile is then checked for convergence against the previous outer iteration, with outer iterations continuing until a converged design is achieved. All Serpent calculations associated with the ADOPT design process employ an r-z representation, where individual assemblies of a given fuel batch are smeared together into cylindrical annuli with properly averaged, smeared constituent materials. This r-z approximation offers savings in computational memory and meshing efforts. The process used to convert this to a more realistic hexagonal core lattice and the associated impacts on predicted core performance are discussed in Section 2.3.

Utilizing ADOPT makes parametric studies both simple and easy to perform, enabling rapid exploration of sensitivities and design options. Such capabilities are utilized extensively in Chapter 6, where the impacts of fuel enrichment and type are investigated in support of reducing material damage. For the reference core in this study, however, the design space explorations of Qvist [114] are primarily used in order to guide the design process. All design parameters and constraints provided to ADOPT are listed in Table 2.1a.

Due to the computational expense of designing and analyzing shuffling schemes on an assembly-level, the neutronic calculations in this dissertation are all on a batch-averaged basis, meaning that the all of the fuel within a given batch is burned with the same batch-averaged flux and their compositions are the same with burnup. This is justified by the large neutron mean free path of the fast spectrum, and is investigated further in Appendix A. Control assemblies are not included in ADOPT calculations due to the difficulty of properly modeling their effects within the context of an r-z model (also discussed in Appendix A). In all ADOPT calculations, the active core is divided into 9 axial segments of equal height, and the core is radially surrounded by a single row of reflector assemblies followed by another row of B₄C shield assemblies. Below the core is a 5 cm thick B₄C shield, while above the core is a 10 cm gas plenum/venting device followed by a 5 cm B₄C shield. Each shield is




Figure 2.2: Levels of iterations driven by ADOPT in the design process for a B&B core.

composed of 90w% enriched B₄C (smear density of 98%) clad in HT9, with volume fractions of 50.8%, 11.9%, 25.3%, 7.0%, and 4.1% for shield, clad, active coolant, duct, and inter-assembly coolant, respectively. The core is surrounded on all sides by coolant. A diagram of the r-z neutronic model is provided in Figure 2.3. The resulting general core performance characteristics are provided in Table 2.1b, while more exhaustive details are provided in Section 2.4.

Table 2.1: User-specified and ADOPT-calculated design parameters for the reference core.

(a) User-specified			(b) ADOPT-calculated		
Parameter	Value	Units	Parameter	Value	Units
Core Power	3500	MW	Equilibrium Cycle Length	2.51	EFY
Inlet Temperature	355	° C	Fuel Residence Time	30.13	EFY
Outlet Temperature	510	° C	Rod Pitch	1.3609	cm
Number of Assemblies	504	-	Rod Outer Diameter	1.2217	cm
Number of Control Assemblies	0	-	Clad Thickness	0.6109	mm
Assembly Duct Flat-to-Flat	18.5	cm	Pins per Assembly	169	-
Duct Thickness	0.35	cm	Fuel Volume Fraction	38.1	%
Inter-Assembly Gap	0.20	cm	Gap Volume Fraction	12.7	%
Number of Batches	12	-	Cladding+Wire Volume Fraction	12.7	%
Assemblies per Batch	42	-	Coolant Volume Fraction	29.4	%
			Duct Volume Fraction	7.0	%

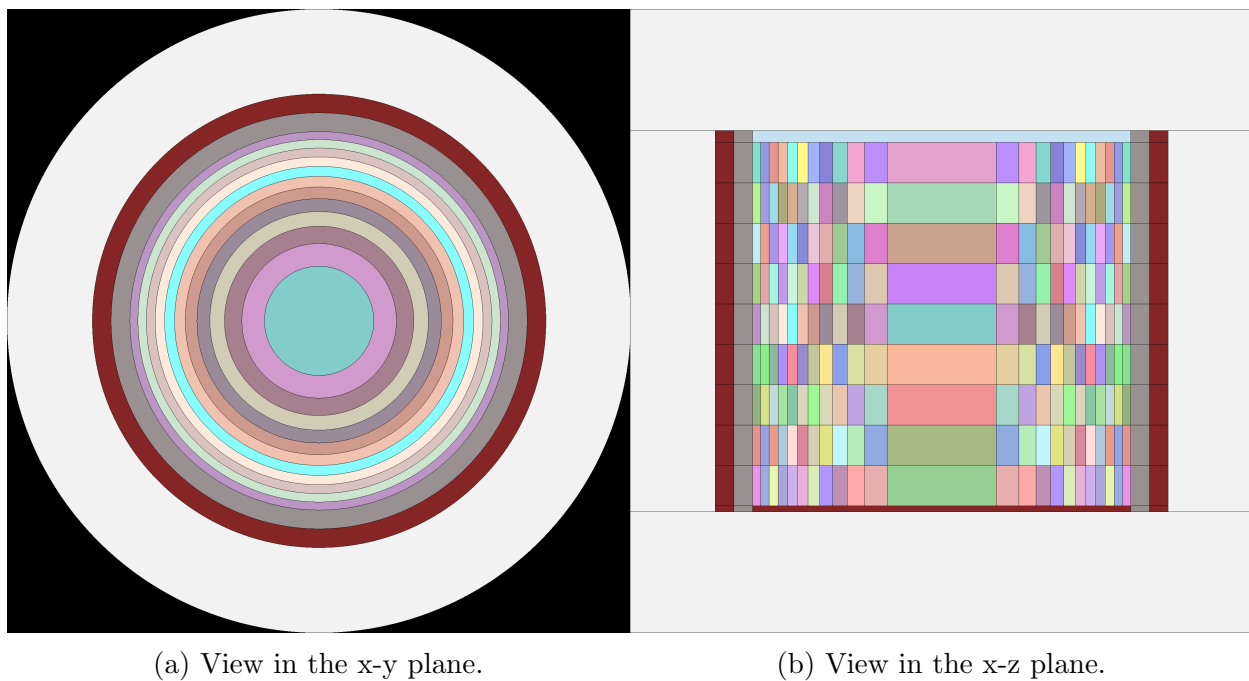


Figure 2.3: Neutronic model employed by ADOPT for the reference B&B core, where each color represents a unique material.

2.3 Converting the Reference Core to Hexagonal Geometry

For many preliminary design studies, the results obtained through ADOPT's r-z modeling are sufficient. However, the detailed analysis undertaken in this dissertation requires neutronic models where assemblies are represented on a hexagonal basis representative of the actual core lattice. The process of converting the coarse design provided by ADOPT into a more realistic neutronic model is not necessarily straightforward, as there are many variants which could be represented by the same r-z model.

Luckily, properties of the fast spectrum make it such that the detailed arrangement of assemblies within the core does not, to first order, hugely impact important core performance characteristics, such as criticality or discharge burnup. Nevertheless, the arrangement of assemblies within B&B cores is important in at least a couple regards. The first is related to the placement of assemblies in a broad, batch-level sense. It has been demonstrated that the shuffling patterns employed in B&B cores can have a profound impact on the minimum required burnup, and thus arranging and shuffling assemblies in general accordance with a well-designed fuel management scheme is paramount. The second is related to the specific placement of control assemblies, which displays markedly stronger sensitivity in terms of control system worth and impacts on radial power distribution than might be seen in other fast spectrum systems. This point is demonstrated in some detail in Appendix A. For the reference core of this dissertation, these two points were considered during the process of developing a hexagonal representation of the core which was originally designed with ADOPT in Section 2.2. The process used in this development is outlined below.

Discussion of General Considerations

The breed-and-burn process leads to a positive feedback where fissile material is preferentially bred in localized regions of high power, potentially leading to excessive power peaking, both radially and axially. Various studies have attempted to minimize radial peaking by utilizing different shuffling schemes. For instance, the systematic study by Zheng et al. [158] examined both “out-in” and “in-out” shuffling schemes as a function of refueling period and found that the “out-in” scheme performed better in terms of reducing radial peaking. Zheng shows that, utilizing a combination of “out-in” shuffling and refueling periods near 3 years, the radial peaking can be brought down to near 2.5. It has been shown by Heidet [54] that the peaking factor can be further brought down by employing shuffling schemes other than simple unidirectional schemes (this is the basis for the shuffling scheme depicted in Figure 2.1). Hou et al. explores this possibility by using a simulated annealing optimization to determine optimal shuffling patterns in both 2D and 3D shuffling paradigms [65]. They find that a radial peaking factor as low as 1.64¹ can be achieved in a 2D shuffled B&B core

¹It should be noted that this radial peaking factor was determined using, in the author's opinion, a coarse neutronic mesh, and likely represents and underestimation of the actual radial peaking factor that

by employing a “bounce” shuffling scheme, although this comes at the expense of slightly increased radial leakage probability as compared to “out-in” shuffling. A qualitative overview of the benefits and drawbacks of different shuffling patterns in standing-wave B&B cores is given by Qvist et al. in [119].

The resulting power profiles for all examined shuffling schemes are highly radially localized, with most of the power generated in only a handful of batches. This radially-localized power production implies that the specific location of control assemblies could be more important in B&B cores to obtain adequate shutdown margin as opposed to, for instance, homogeneous core layouts, where the power profile is comparatively flat and slowly varying. In general, control assemblies must be placed in regions of high importance in order to have sufficient worth. For B&B cores, these high importance regions may be small in size. By placing control assemblies in high worth regions, the neutron spectrum is softened and parasitic losses in sodium and steel are increased, which penalizes the core neutron balance and implies a higher required burnup.

In using ADOPT for preliminary design, the locations of particular assemblies are obscured and the explicit placement of control assemblies is neglected, with the implicit assumption that the control system design can be performed upon establishing core feasibility sans control system. Justification for this design approach in homogeneous fast cores is provided by, for example, results by Hoffman et al. for the Advanced Burner Reactor [60]. In their study, the core design is modified in a number of ways to achieve conversion ratios from 0.25 to 0.75, all while successfully employing the same control system layout. A similar study performed by Xiao et al. [153] shows that the control assembly layout for a particular homogeneous breeder core can remain unchanged as the breeding ratio is increased from 1.0 to 1.4. Results such as these justify using a simplified or assumed control system model, or even neglecting the control system entirely, in preliminary core design studies to narrow the design space and expedite the design process. Although the majority of studies on standing-wave B&B systems have utilized r-z neutronic models without control system inclusion, no equivalent results justifying this practice have been reported specifically for B&B cores.

Limited investigations on the subject have been reported by Heidet and Greenspan [56]. In their studies on a medium-sized B&B core, they employ hybrid models which include explicit control system layouts superimposed on an r-z fuel representation to quantify shutdown margin. Their study finds that this hybrid r-z model does not significantly alter the predictions of a pure r-z model. However, Girardin et al. [40] has shown that r-z models are inadequate to characterize control systems in large fast cores in the first place, particularly when significant shadowing effects between control assemblies are present. Girardin et al. demonstrates a difference of more than 3000 pcm in total control system worth between hex-z and r-z models of a 2400 MWth homogeneous gas-cooled fast reactor, representing a significant mischaracterization of the control system. Therefore, the limited studies of control system characterization in B&B cores using r-z models might be viewed with skepticism, and the following study explores the design space for control assembly layouts in the

would be seen if power were tallied on a finer mesh.

reference core using neutronic models that explicitly represent the hexagonal core lattice.

Procedure

Once an r-z equilibrium core is designed with ADOPT, the material compositions from each batch are transferred into models which explicitly represent the hexagonal geometry of the core lattice, with each hexagonal position representing an assembly with smeared composition. To understand the impacts of the control assembly layouts, four hexagonal models are generated while adhering to the following goals:

1. Maintain sixth-core symmetry
2. Attempt to keep all assemblies within a batch adjacent to each other
3. Keep the number of assemblies in each batch equal
4. Attempt to approximate cylindrical batches
5. Incorporate 19 control assemblies into the core

The layouts of each of the four candidate designs are shown in Figure 2.4, where the average position of the control assemblies is gradually moved away from the core center moving from Layouts (a) to (d). Each model uses 24 equal-length segments to capture axial power, flux, and burnup gradients. In the following Monte Carlo calculations, only neutron transport is simulated to save on computational expense, as the power deposited by photons does not impact the worth of the control system. Calculations are run with all rods fully withdrawn (with the exception of calculations in support of determining control system worth), with enough particle histories to allow for resolution of k_{eff} down to less than 10 pcm and assembly powers down to about 100 kW.

The assumed control assembly design is largely taken from those used in the metallic-fueled Advanced Burner Reactor designed by Argonne National Laboratory, outlined in [133]. Each control assembly contains 7 absorber pins arranged in a hexagonal lattice, surrounded by two concentric hexagonal ducts made of HT9. Each absorber pellet is 4.57801 cm in diameter and is clad with 0.71625 mm thick HT9. The absorber material is B_4C , and throughout the study the enrichment in ^{10}B is varied from natural (19.9%) up to an assumed maximum of 90%.

Impacts of Control System Layout on Core Performance

The following sections explore how the various control system layouts depicted in Figure 2.4 impact the B&B core's performance, including the power profile, numerous integral neutronic parameters, and the control system worth. The different core layouts are explored by taking the equilibrium material compositions and cycle length from the ADOPT calculations and running a depletion sequence over a single burnup cycle. Further calculations to determine

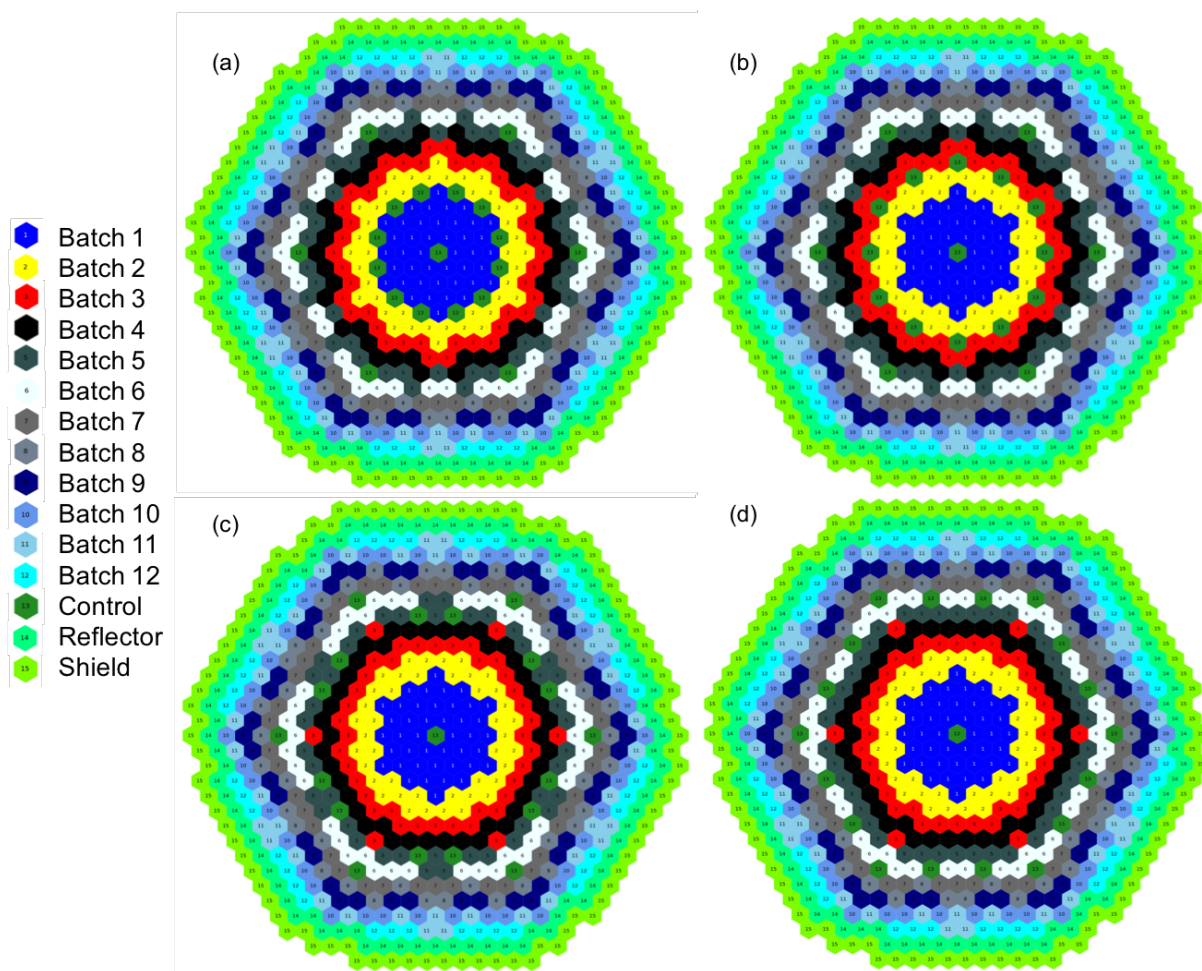


Figure 2.4: Assembly layouts for each of the four hexagonal core layouts consider. Moving from (a) to (d), control assemblies are gradually moved further away from the core center.

the impacts on the equilibrium cycle length are not performed. The resulting trends are ultimately used to decide upon a control system for the reference core that is used throughout the remainder of this dissertation.

Power Distribution

Figure 2.5 shows the uncontrolled radial power distribution on an assembly level for Layout (a), axially integrated over each assembly. The general shape of the power distribution is similar for all four layouts, and therefore plots of the other layouts are left out for brevity. Because no gamma transport is modeled and no energy release from capture reactions is assumed, the power production in the control assemblies is exactly zero. For all layouts, the majority of power production takes place in a ring about four assemblies thick towards the

radial midpoint of the core, while the power production in the core center and periphery drops off rapidly. This is due to the shuffling scheme, which puts fuel that has seen the most fluence (i.e. has built up the most plutonium) near the radial midpoint. This power profile is fundamentally different than what is seen in a typical SFR design and contributes directly to the sensitivity of control rod worth to control assembly positions seen later.




Figure 2.5: Uncontrolled fission power at Beginning of Equilibrium Cycle (BOEC) in each assembly calculated by Serpent 2 for Layout (a). Powers in MW.

Table 2.2 shows a detailed breakdown of the fraction of total power produced in each batch for all four layouts, along with the differences between them. Between layouts, the power within each batch may vary by up to 7% of total core power, or about 250 MW. The largest changes take place in batches that immediately border the high power ring. These large power variations highlight the nature of ADOPT as only a *preliminary* design tool, since the power distribution known to ADOPT does not include the impacts of the specific control system layout, implying that the true equilibrium core characteristics would require an iteration involving control assembly layouts designed using hexagonal core models.

Integral Neutronic Characteristics

Table 2.3 gives a few key neutronic parameters and the maximum differences between them for each core layout. The difference between the criticality of Layouts (b) and (d) is nearly

Table 2.2: Percentage of total core power produced in each batch for the four layouts depicted in Figure 2.4 at BOEC, along with the batch-wise differences between core layouts.

Batch	1	2	3	4	5	6	7	8	9	10	11	12
Average (%)	3.93	19.41	25.56	20.88	15.50	8.95	3.87	1.17	0.47	0.13	0.09	0.03
Range (%)	0.34	4.10	2.66	7.06	3.69	1.36	0.34	0.10	0.04	0.01	0.01	0.00
Range (MW)	12.06	143.45	93.16	247.21	128.98	47.53	11.86	3.60	1.44	0.40	0.29	0.11

§5, which is more than half of the cycle reactivity swing. The decreased reactivity of Layout (b) results from the softer spectrum, depicted in Figure 2.6 for Batch 3, and increased parasitic neutron absorption induced by control assemblies located directly in the region of highest power production. This slightly softer spectrum leads to less fast fission in fissionable isotopes and a very slightly lower $\bar{\nu}$.

Table 2.3: BOEC core neutronic parameters for each of the four layouts depicted in Figure 2.4.

	k_{eff}	$\bar{\nu}$	Conversion Ratio	β_{eff}	Neutron Lifetime (s)
Layout (a)	$1.01680 \pm 5.0E-05$	$2.90673 \pm 9.1E-06$	$1.57589 \pm 1.89E-04$	$3.71037E-03 \pm 1.65E-05$	$1.92187E-07 \pm 1.88E-10$
Layout (b)	$1.00724 \pm 7.4E-05$	$2.90616 \pm 6.4E-06$	$1.59338 \pm 2.39E-04$	$3.68316E-03 \pm 1.26E-05$	$2.04997E-07 \pm 1.82E-10$
Layout (c)	$1.01699 \pm 6.4E-05$	$2.90660 \pm 4.7E-06$	$1.57143 \pm 2.67E-04$	$3.69055E-03 \pm 1.54E-05$	$1.94055E-07 \pm 1.80E-10$
Layout (d)	$1.02440 \pm 1.0E-04$	$2.90718 \pm 5.7E-06$	$1.55490 \pm 2.18E-04$	$3.67045E-03 \pm 2.30E-05$	$1.87833E-07 \pm 1.07E-10$
Max Diff	0.01726	0.00102	0.03848	0.03992E-03	0.17164E-07
Max Diff	1.68%	0.04%	2.41%	1.08%	8.37%

Table 2.3 shows that the difference in conversion ratio, defined as the ratio of fertile nuclide capture to fissile nuclear absorption, is not large. It is interesting that the conversion ratio is largest in Layout (b), whereas this layout has the softest neutron spectrum and the lowest $\bar{\nu}$. This is possibly due to the placement of the control assemblies further from the blanket fuel in Layout (b), while the three other layouts allow parasitic absorption in the added sodium and structure of the control assembly locations of neutrons leaking into the blanket batches.

In terms of kinetic behavior, the position of control assemblies only makes about a 1% difference in β_{eff} between the different layouts. However, the neutron lifetime does change by as much as nearly 10% between Layouts (b) and (d), being longest in Layout (b) and steadily decreasing as the control assemblies are moved away from the high power region. This is consistent with the softer neutron spectrum in the high power region of Layout (b) and may also contribute to varying reactivity coefficients between the different designs, though this is not explicitly quantified here.




Figure 2.6: Normalized neutron spectra at the core midplane in batch 3 for Layouts (b) and (d) along with the difference between the two.

Control System Worth

Figure 2.7 shows k_{eff} at BOEC with control rods fully inserted as a function of ^{10}B enrichment for all four layouts. It is clear that the core can never be controlled for Layout (d) no matter how highly enriched the absorber material is. Additionally, the slope with increasing enrichment for Layout (d) is very small, implying that even if a better neutron poison than boron carbide were used, it is unlikely to make a difference since the sensitivity is very low when the control system it is located so far from the high-importance region.

Table 2.4: Total control system worths at BOEC for the layouts depicted in Figure 2.4 with absorber material enriched to 90%.

Layout	(a)	(b)	(c)	(d)
Control System Worth (%)	9.9	27.2	12.5	5.5

The total system worth for each layout with 90% enriched absorber material at BOEC is given in Table 2.4. Between Layouts (c) and (d), the control system worth is cut in half, even though the control assemblies are moved outwards by only one ring. Similarly,




Figure 2.7: k_{eff} with rods fully inserted as a function of increasing boron enrichment for all four core layouts.

between Layouts (a) and (b), the control system worth is increased threefold just by moving the control assemblies by only 1.5 rings, making the difference between an uncontrollable core and one with significant margin. This demonstration makes clear that the position of the control rods can have a very large impact on the performance characteristics of a B&B core. The only layout that is likely to have enough control system worth to cover the requirements of reactivity gain, temperature defect, reactivity fault, and uncertainties is Layout (b), which unfortunately is the layout with the lowest uncontrolled k_{eff} . This tradeoff should be expected since the region near the control assemblies is the area of highest importance.

Selecting a Control System

In order to understand the worth required by the control system, a study of the contributing components was undertaken, with Figure 2.8 showing a summary. The largest contribution to the requirement is from the cycle reactivity gain, which is roughly \$11. The next highest contribution is from uncertainties such as those due to nuclear data, material composition, and the flux distribution. Without performing a detailed sensitivity/uncertainty study, the contribution from uncertainties is assumed to be at most \$3.0, based on studies done in support of the Clinch River Fast Breeder Reactor [112]. The temperature defect, while

smaller than in many other SFRs due to the very weak doppler feedback highlighted in Section 2.4, is the next highest contributor with a requirement just over \$2.0. The final two contributions to the control system requirement are the reactivity fault and the stuck-rod criteria, each of which are equal to the worth of the highest worth control assembly (quantified subsequently to be \$1.66).

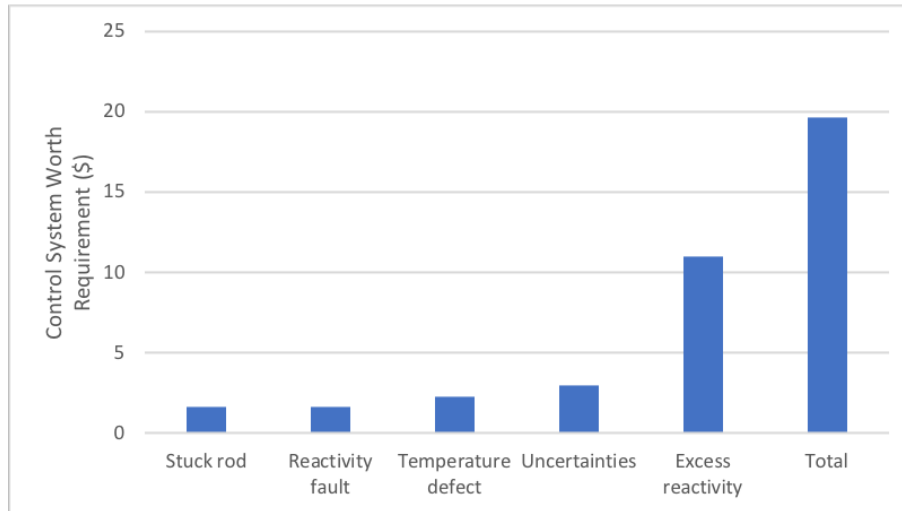


Figure 2.8: Contributions to the total control system worth requirement.

Combined, these components lead to a total control system worth requirement of roughly \$20, with requirements being somewhat lower at BOEC than EOEC due to the core's positive reactivity swing. Out of the four layouts depicted in Figure 2.4, only Layout (b) is able to muster enough worth to be feasible, even with the absorber material enriched up to 90%, and therefore Layout (b) is chosen for the reference core. This layout is repeated as Figure 2.9 in larger form to aid the reader. Still, however, the enrichment level of the absorber material needs to be specified. Figure 2.10 shows the worth of Layout (b) as a function of absorber enrichment throughout the cycle. Natural enrichment is found to be sufficient to provide the necessary worth, and therefore natural enrichment is selected due to the economic penalty associated with isotopic enrichment. It is noted as well that enrichment up to 90% would allow for the control system to have enough worth such that two redundant systems could be formed using the same layout, as may be required by a regulator. For the purposes of this study, however, natural enrichment levels in the control system are used throughout.




Figure 2.9: Core map of the reference core, repeated from Figure 2.4b in larger form.




Figure 2.10: Total control system worth for Layout (b) as a function of absorber enrichment at various points in the cycle.

2.4 Overview of Core Performance Characteristics

With the hexagonal core layout determined, this section explores the unique performance characteristics of B&B cores. These characteristics lay the foundation for many of the outstanding challenges highlighted in Chapter 1, and thus a thorough understanding is highly relevant to discussions throughout this dissertation.

The uncontrolled criticality evolution over the equilibrium cycle is provided in Figure 2.11. As previously indicated, the core has a strongly positive reactivity swing, which is somewhat stronger at BOEC when the fertile/fissile ratio is higher and the breeding ratio can be made exceptionally high. Such efficient breeding results from the combination of three main aspects: (1) a very high fertile content loading, (2) a very hard neutron spectrum, and (3) specifically tailored shuffling schemes. The reactivity swing is relatively strong, especially in comparison to some SFR cores designed specifically for near-zero reactivity swings in order to minimize control requirements over their lifetimes (see, for instance, [105]). Although it leads to higher worth requirements for the control system, the strong breeding is beneficial from the perspective of achieving a lower minimum required burnup and is thus pursued in the reference core.




Figure 2.11: Criticality (including Monte Carlo uncertainties) and breeding ratio evolution over the equilibrium cycle in the reference core.

The very high breeding ratio leads to quite high Pu-239 concentrations, particularly in fuel batches which have experienced significant fluence. A map of Pu-239 concentration throughout the core is given in Figure 2.12. As is a common feature of breed-and-burn operation, the concentration of Pu-239 is severely axially-peaked, with the peaking becoming

less severe in batches near discharge as the Pu-239 content is burned out. At most, the Pu-239 enrichment level is around 10 atom%. This is somewhat lower than what might be loaded into a typical continuous-recycle SFR system, and is enabled due to the low core leakage.




Figure 2.12: Enrichment of fuel in Pu-239 locally throughout the reference core at BOEC. Enrichment values are in atom%.

The unique spatial distribution of Pu-239 leads to remarkable power distributions, especially compared with those of typical SFR designs. Figure 2.13 shows the evolution of the power profile in the reference core over the equilibrium cycle. The high-powered annular region shifts with burnup towards the radial center, with the central batches increasing significantly their share of total power production. A detailed breakdown of the power production in each batch at BOEC and EOEC is provided in Table 2.5. Over the course of the cycle, Batches 3 and 4 consistently have the largest share of total power production. Batches 7-12, while representing 50% of the total fuel mass, contribute only 7% of total power.

Due to their very low power production, the outer batches could be considered as a series of radial blanket assemblies in that they exist solely to catch neutrons leaking radially from the core. Contrary to typical SFR blanket assemblies, however, their geometry cannot be optimized for low power operation due to the fact that they will eventually be shuffled into locations of high power density, and thus must be designed subject to peak power constraints. This introduces issues related to thermal efficiency of the reactor coolant system. In addition, over just a single burn cycle the power-per-assembly can shift dramatically, as demonstrated




Figure 2.13: Axially-integrated radial power distribution in the reference core throughout the equilibrium cycle.

Table 2.5: Batch-integrated power fractions at BOEC and EOEC of the reference core, along with the average change in a given assembly of each batch over the cycle.

Batch	1	2	3	4	5	6	7	8	9	10	11	12
BOEC Power (%)	3.23	16.34	24.03	23.31	18.09	9.06	4.14	1.18	0.44	0.11	0.06	0.02
EOEC Power (%)	7.59	19.74	22.48	19.42	14.88	8.98	4.39	1.59	0.65	0.16	0.11	0.03
Average Assembly Power Swing (%)	136.00	20.93	-6.33	-16.55	-17.72	-1.22	5.66	35.11	48.33	47.58	58.65	46.90
Average Assembly Power Swing (MW)	3.64	2.84	-1.26	-3.20	-2.67	-0.09	0.19	0.34	0.17	0.04	0.03	0.00

in Table 2.5. Particularly in lower-powered batches, the power-per-assembly for an average assembly in a batch can more than double as the cycle progresses. These power swings add additional complexity to the issue of adequately cooling each assembly, and such issues will be discussed extensively in Chapter 3.

In a similar manner, the breed-and-burn strategy leads to a buildup of fissile content towards the axial core center which drives a highly peaked axial power profile, as shown in Figure 2.14 for each batch. Significant deviation from the chopped-cosine profile of a

core with axially-constant enrichment is obvious. Provided in Table 2.6 is the same data summarized in the form of the axial peaking factors. It is seen that as burnup increases, the axial peaking factor initially increases. As burnup continues, however, the peaking factor begins to decrease as fissile content burns out and fission products build up. These competing effects result in power profile being most highly peaked in Batch 1, where enough fluence has been achieved in order to have built up significant fissile content, but not enough has been achieved to help flatten the power profile. This subsequent decrease in the axial peaking factor is beneficial because it eases the cooling requirements in the highest powered batches.

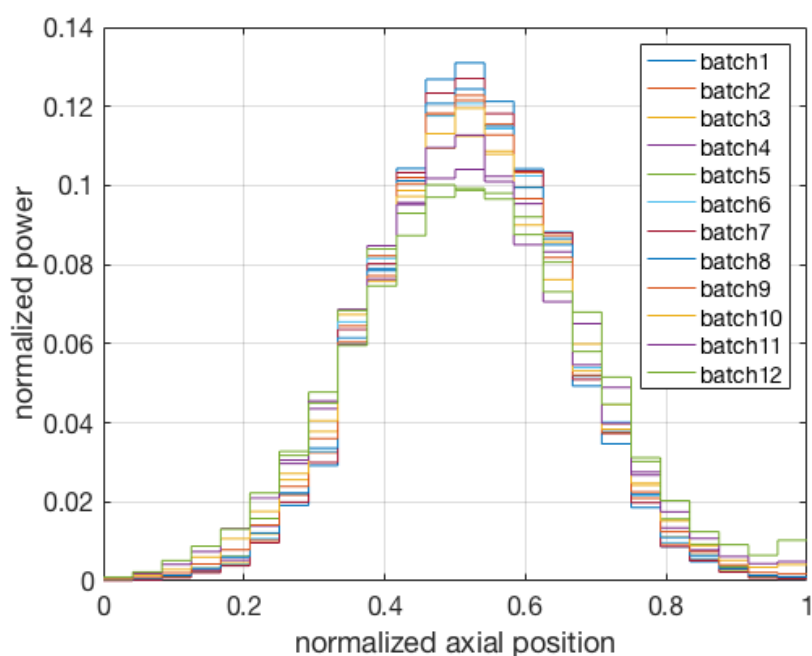


Figure 2.14: Radially-integrated power in each batch as a function of axial position in the reference core at BOEC.

Table 2.6: Axial peaking factors for each batch of the reference core at BOEC.

Batch	1	2	3	4	5	6	7	8	9	10	11	12
Axial Peaking Factor	3.14	2.91	2.69	2.49	2.38	2.89	3.05	2.98	2.94	2.86	2.70	2.40

In comparing the radial power and flux profiles, Figure 2.15 shows that the two generally follow the same distribution, although the flux profile is less peaked than the power profile due to the high fertile content depressing power production in regions of otherwise high




Figure 2.15: Flux and power profiles for the reference core on a batch-wise basis at beginning of cycle.

flux (i.e the core center). Regardless, the flux depression at the core center acts as an internal blanket which serves to make the core relatively loosely coupled, from a neutronics perspective. This impedes the communication of flux perturbations laterally throughout the core, and introduces the possibility of radial instabilities which would otherwise be unexpected in a more classical SFR design.

These coupled burnup/power distributions lead to a burnup profile at EOEC that superficially resembles the Pu-239 content variation, as shown in Figure 2.16. It is noted that the burnup profile is more axially-peaked than the Pu-239 profile, as would be expected based on the fact that Pu-239 is the primary fissile isotope driving criticality and the flux distribution. The discharge burnup peak is shifted slightly towards the top of the core, where the coolant temperature is higher and thus its density is lower, leading to fewer parasitic absorptions. In agreement with the power profiles shown previously, it is noted that the majority of burnup takes place in Batches 2-5. Upon discharge, the peak burnup is 31.5%FIMA, while the average burnup is 12.4%FIMA. This large difference between peak and average burnups is an inherent limitation of the classical breed-and-burn mode of operation when pursuing extremely low-leakage design, and has been the primary motivation for studies considering 3D shuffling.

As might be anticipated from the unique power profiles, leakage in the B&B core is extremely low. Table 2.7 provides the axial and radial leakage probabilities in the reference core, showing that essentially no neutrons are lost beyond the thick radial blankets. As




Figure 2.16: Burnup profile at EOEC of the equilibrium cycle in the reference core. Burnups are represented in %FIMA.

discussed in Chapter 1, such low leakage enables breed-and-burn operation with the lowest possible material damage, and is thus a driving aspect of the fundamental design strategy.

Table 2.7: Axial and radial leakage probabilities from the active core region at BOEC.

Axial	Radial
0.2%	0.0%

Such low leakage combined with strong axial gradients in fissile content contribute to interesting control rod worth curves. Using the previously determined control system, Figure 2.17 shows the control rod worth as a function of insertion depth. In comparison to a typical SFR, the curve is notably sharp. In addition, the curve displays two regions of extremely flat incremental worth at the top and bottom of the active core region as a result of the very low fissile content there. This has important consequences both on the operation of the core and on the response to rod withdrawal accidents, as will be discussed in Chapter 4.

Also as a result of the combined low leakage and unique fissile content distribution, interesting material worth distributions are seen. Figures 2.18a-2.18c display the worths of the primary core constituents in each batch as a function of axial position, as calculated using the Generalized Perturbation Theory (GPT) capabilities of Serpent [14]. Each of the




Figure 2.17: Impact on core criticality as the control system is inserted from the top of the core at EOE. All rods are assumed to be uniformly inserted.

material worth profiles follows the power profile in general shape, with worths generally being highest in Batch 3. It is significant to note that the coolant *void* worth (i.e. the negative of the coolant worth) is positive throughout the entirety of the core. This is sharply in contrast with a typical SFR, where the void worth would typically turn negative towards the edges of the core as the importance of leakage increases in comparison to parasitic absorption and spectral hardening. Because leakage is so low in the first place, the slight increase in leakage when coolant voids has minimal impact. This leads to a strongly positive total core void worth of \$12.03 and \$11.95 at BOEC and EOE, respectively, which has a profound impact on the transient performance of the system, as will be explored in detail in Chapter 4.

The cladding worth shows similar behavior with slightly lower magnitude due to the small cross-sections of the HT9 constituents. The fuel worth, however, is notably less axially-peaked than either the coolant or cladding worths, especially in higher burnup batches. This is because as the fuel accumulates burnup, its fission product inventory provides a negative worth to counter the positive worth of the fissile content, with this effect being most pronounced at the axial center where power is concentrated.

Many of these performance characteristics, namely the high burnup and strong power peaks, contribute to excessive accumulation of material damage in the in-core structural materials (i.e. cladding). As discussed in Chapter 1, this is one of the primary infeasibilities of the breed-and-burn mode of operation, subject to our current limited understanding of

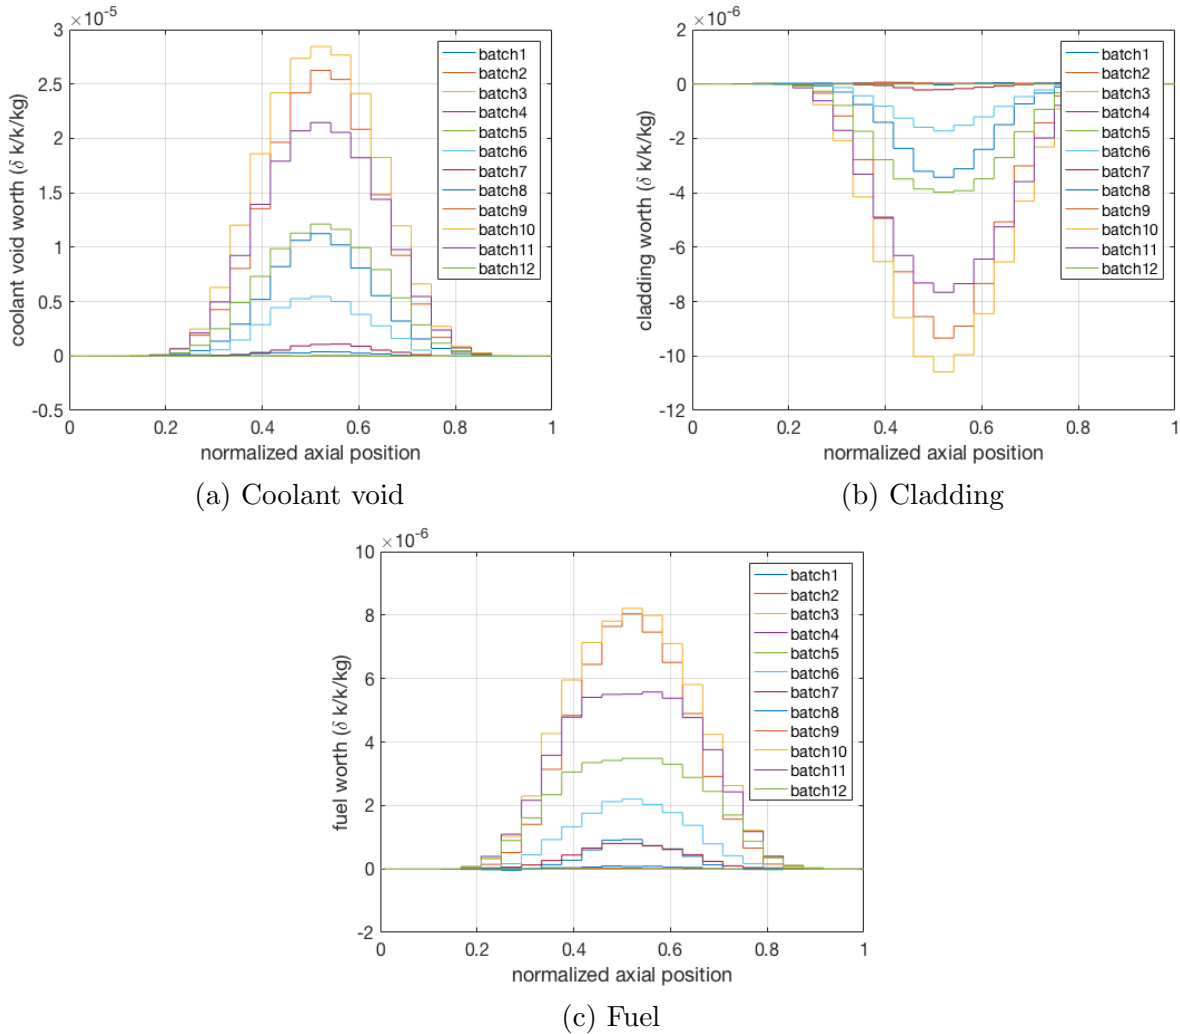


Figure 2.18: Distribution of material worths in the reference core at EOEC.

irradiation damage in reactor components. Figure 2.19 shows the distribution of cladding damage accumulation in terms of DPA at the axial midplane (the location where material damage is the most exacerbated). The peak discharge DPA is determined to be 556 DPA in this particular design, which is in line with contemporary designs. As expected, there is a direct correlation between the power and material damage within each batch, resulting in the majority of cladding damage taking place in Batches 2-5.

Summary of Reference Core Performance

This section has provided a broad overview of important performance characteristics in the reference core so as to frame discussions in the remainder of this thesis. It has been




Figure 2.19: DPA accumulation in each batch of the reference core at the axial midplane.

demonstrated that the B&B core displays unique characteristics fundamentally driven by the shuffling scheme employed, resulting in exceedingly low leakage and unique fissile content distributions. This is reflected in distinctive power profiles and material worth distributions which play a fundamental role in the challenges faced in the commercialization of B&B technology.

Chapter 3

Assembly Orificing

It has generally been understood since the inception of the standing-wave breed-and-burn concept that its very nature will make allocating flow among assemblies more difficult than in a typical SFR. Although acknowledged, very little effort has been devoted to understanding this problem in the public domain. This chapter focuses on this problem, henceforth termed the ‘orificing’ problem, proposing methods to determine optimal orificing schemes and investigating sensitivities to important constraints specifically for the case of B&B cores, using the reference core as a case study.

3.1 Introduction to the Orificing Problem

In any SFR, individual assemblies are typically orificed in order to distribute the total core flow appropriately between different assemblies [149]. Orificing involves placing a passage of prescribed dimension in the flow path such that a pressure loss is induced, reducing the flow to some desired level. Prescribing the flow rates allocated to each assembly, or the ‘orificing scheme,’ is an important task of the reactor designer that can have profound impacts on plant efficiency and achievable power levels. In designing such an orificing scheme, many competing constraints need to be met simultaneously, making the task a potentially complicated optimization process that must account for evolving core characteristics with burnup. For these reasons, orificing should be considered as a fundamental aspect of any new core design study.

In the case of standing-wave B&B cores, the orificing task may be exceptionally difficult due to inherent characteristics of the breed-and-burn process. Namely, the presence of large numbers of very low-powered assemblies, which are often overcooled at early points in the cycle due to their large positive power swings, can severely depress the average coolant outlet temperature, making it difficult to achieve the desired average outlet temperature without having high-powered assemblies exceed their thermal limits. This issue may potentially be handled in a variety of ways, but ultimately the ability to design sufficient orificing schemes for B&B cores has not been addressed in the open literature. Therefore, this chapter

examines different methods of orificing and analyzes their suitability to standing-wave B&B reactors.

Constraints on Assembly Orificing

When determining a core's orificing scheme, a variety of constraints must be accounted for simultaneously. In general, each of these constraints may be imposed at all or only select times during the fuel management cycle, and thus the constraining values may evolve with time. This section discusses the origins and implications of the major constraints for B&B core designs. The specific constraining values and methods used to enforce them within a numerical design environment are outlined in Section 3.3.

1. **Adequate cooling of the assemblies:** The primary concern of the orifice designer is to ensure that all assemblies are cooled adequately during their residence time in the core. The most basic requirements are to keep the coolant, cladding, and fuel temperatures away from their respective limiting temperatures with some margin. For metallic-fueled cores, the most limiting constraint is usually the peak inner-cladding temperature, whereas for oxide-fueled cores the peak centerline temperatures may be most limiting [149]. It is also important to account for temporal changes in assembly characteristics as material is bred and assemblies are shuffled, such as changes in material properties and power levels. Temporal variations are especially pertinent in B&B cores, where assembly powers may double or even triple over the equilibrium cycle (discussed in more detail in Chapter 2). Previous methods for orificing design, such as that proposed by Heidet [53], are not able to effectively handle such large variations in power generation over a cycle.
2. **Plant thermodynamic efficiency:** From an economic perspective, it is vital to keep the mixed outlet temperature near the value required by the power conversion system throughout the cycle. Without this concern, all assemblies could simply be cooled according to the demands of the peak-powered assembly and orificing would not be required. In large B&B cores, where low-powered breeding assemblies constitute a major fraction of the core volume, it may be difficult to compensate for the depressed outlet temperatures from these assemblies while maintaining acceptably low temperatures in the higher-powered assemblies. Depending on the balance-of-plant systems, some tolerance may be allowed around the desired mixed outlet temperature, but too large of variations will lead to inefficiencies and suboptimal performance. In addition, the ability to cope with large changes in assembly powers over the cycle is imperative from the perspective of this constraint as well.
3. **Maximum coolant flow rate:** To meet the cooling demands of an assembly, it is necessary to match the flow rate to the assembly power. However this cannot be done to arbitrarily high power levels due to constraints on the coolant velocity and pressure drop. Excessively high flow velocity can lead to flow-induced vibrations of core

components [108] and, especially in the case of Lead Fast Reactors (LFRs), accelerated corrosion [149]. Additionally, a constraint on the coolant flow rate is induced by the core pressure drop, which is limited by the size of the primary coolant pumps and the economic tradeoffs of the system design. Because pressure drop is proportional to coolant density, designs with very heavy coolants such as lead or lead-bismuth may be constrained more by pressure drop than concerns of flow-induced vibration. In the case of sodium coolant, however, flow-induced vibration constraints are often more limiting. Regardless, maintaining each assembly's flowrate below the relevant constraint is important to the integrity of reactor components.

4. **Adjacent assembly outlet temperatures:** Due to the high thermal conductivity of liquid metals, thermal striping of above-core structures is a concern when outlet temperatures of adjacent assemblies are significantly different [5]. As coolant of different temperatures mixes in the outlet plenum, vortices form at the interface of the two flows. As these vortices impact the upper internals, the high thermal conductivity of sodium allows for high-frequency temperature oscillations to be imparted to the structure, which can lead to rapid thermal expansion/contraction cycles and eventual failure of the upper internals. This phenomenon is a key challenge in B&B cores, where adjacent assemblies may have widely differing power levels [57]. Through orificing strategies which limit the differences in outlet temperatures between adjacent assemblies throughout the cycle, thermal striping may be mitigated.

Methods of Orificing

Multiple methods of orificing have been proposed and used in the past. The three methods explored in this paper, each of which having varying degrees of real-world operational experience, are referred to as 'static integral orificing,' 'static grid plate orificing,' and 'variable orificing.' This section outlines each of these methods.

1. **Static integral orificing:** The majority of SFRs have used the static integral orificing method, depicted in Figure 3.1. In this method, orifices are included as part of each assembly's inlet nozzle, and as such, are an integral part of each assembly. Thus, as assemblies are charged/discharged from the core or possibly shuffled, the orifice remains with the assembly, resulting in each assembly receiving roughly the same amount of flow over its residence time no matter how its location or power level may change. This method allows for moderate flexibility in off-equilibrium operation or for changes in the core design strategy, as when new assemblies are introduced to the core the orifices may be changed as well.
2. **Static grid plate orificing:** A method less commonly used in SFRs is static grid plate orificing. This method was employed, for instance, in EBR-II [81] and is currently the method employed in General Electric's Boiling Water Reactors (BWRs) [24]. Static grid plate orificing is accomplished by limiting the apparent orifice of each assembly




Figure 3.1: Diagram of one method of achieving static integral orificing.

using the grid plate, an example of which is depicted diagrammatically in Figure 3.2. Because orifices are tied to a location and not to specific assemblies, this method allows for flow rates of shuffled assemblies to vary between cycles. As such, this method generally allows for larger assembly power swings and may be particularly suited to B&B operation. However, the grid plate is a fundamental component of the reactor vessel internals, and as such it is likely both difficult and expensive to replace. Therefore, static grid plate orificing is viewed as less flexible than static integral orificing in the sense that once a fuel-management scheme has been chosen, it may be difficult to deviate while staying within the bounds of the associated orificing strategy, making it potentially difficult to cope with off-equilibrium operation. However, nothing prevents the possibility of a hybrid static integral/grid plate orificing method in which aspects of both schemes are employed, either where different regions of the core employ different methods or where both methods are stacked on top of each other for a given assembly/position combination. Such hybrid schemes may enable some amount of flexibility to be incorporated into a grid plate orificing scheme.

3. **Variable orificing:** The final method examined in this study is variable orificing. Although not a fast reactor, this method has been employed in the Hallam Sodium Graphite Reactor (SGR) built by Atomics International [19]. Additionally, it has been proposed multiple times for both SFRs and Gas-cooled Fast Reactors (GFRs) [149,




Figure 3.2: Diagram of one method of achieving static grid plate orificing.

96], with a patent having recently been filed by TerraPower, LLC [7]. In this method, a means to physically change each orifice is located either above or below the core, in the cases of Hallam and TerraPower, respectively, and controlled remotely. In the method proposed by TerraPower, the apparent orifice to each assembly is varied by rotating a sleeve over top of the assembly inlet nozzle such that the holes in the sleeve either align or anti-align with the holes in the nozzle, as depicted in Figure 3.3. Either during operation or between cycles, the orifice may be changed to accommodate varying power levels. This method is extremely flexible to allow for both large power swings or off-equilibrium operations, but introduces a new, complicated system into the reactor vessel that may have important safety and economic implications. Regardless, it may be extremely well-suited to B&B operation, and therefore is explored in this study.




Figure 3.3: Diagram of one method of achieving variable orificing.

3.2 Assessment of Static Integral Orificing

While static integral orificing is historically the most widely used, this section uses bounding analytical techniques to rigorously demonstrate its unsuitability for realistic B&B designs. This is done by showing that the constraint on assembly outlet temperatures combined with the large number of overcooled assemblies does not allow for constraints on the core-average outlet temperature to be met.

Take Q to be a matrix of assembly powers where element Q_{ijk} represents the integrated power of assembly $j \in \{1, \dots, n_a\}$ within batch $i \in \{1, \dots, n_b\}$ at timestep $k \in \{1, \dots, n_t\}$ of the equilibrium cycle, where n_b is the number of batches, n_a is the number of assemblies per batch, and n_t is the number of timesteps representing discretized points in the equilibrium cycle to capture the power evolution. B&B shuffling schemes are typically such that all assemblies within a given batch will be shuffled into every other batch before being discharged. While the details of an assembly's power swing over its residence time will depend on its specific shuffle path, the minimum possible power swing is obtained if the highest-powered assembly of the lowest-powered batch is eventually shuffled to the position of lowest power in the highest-powered batch, i.e.

$$\min_i \left[\max_j \left[\min_k Q_{ijk} \right] \right] \rightarrow \max_i \left[\min_j \left[\max_k Q_{ijk} \right] \right]$$

Because flow rates must be allocated to satisfy temperature constraints at peak-power conditions, an assembly undergoing this specific shuffle will achieve the smallest possible

amount of overcooling while in its period of low-power operation. To bound the analysis, it is assumed that every assembly within each batch can achieve this same minimum amount of overcooling, and thus each batch is represented by just this single minimally-overcooled assembly. Knowledge of the assembly-level shuffling scheme is not needed in such a bounding analysis.

If assembly-to-assembly heat transfer is neglected, the average outlet temperature of an assembly can be calculated as

$$T_{out,ijk} = T_{in} + \frac{Q_{ijk}}{\dot{m}_{ijk}c_p} \quad (3.1)$$

where T_{in} is the core inlet temperature, \dot{m}_{ijk} is the flow rate of assembly j in batch i at time k , and c_p is the heat capacity of the coolant, assumed to be averaged over the coolant temperature range. Taking a flow-weighted average, the core-average outlet temperature at time k is

$$\bar{T}_{out,k} = \frac{\sum_i \sum_j T_{out,ijk} \dot{m}_{ijk}}{\sum_i \sum_j \dot{m}_{ijk}} \quad (3.2)$$

For the case of static integral orificing in a B&B core with our bounding assumption, $\dot{m}_{ijk} \equiv \dot{m} \forall \{i, j, k\}$, i.e. all assemblies have the same flow rate at all times, corresponding to the minimally-overcooled assembly. The bounding flow rate of this analysis is,

$$\dot{m} = \frac{\alpha}{\Delta T_{max} c_p} \quad (3.3)$$

where $\alpha = \max_i [\min_j [\max_k Q_{ijk}]]$ and ΔT_{max} is the constraint imposed on the maximum average temperature increase over an assembly. The core-average outlet temperature can then be calculated subject to specific values of ΔT_{max} as

$$\bar{T}_{out,k} = \frac{\sum_i \left[T_{in} + \frac{q_{ik}}{\alpha} \Delta T_{max} \right]}{n_b} \quad (3.4)$$

where

$$q_{ik} = \begin{cases} \min_j Q_{ijk} & \text{if } i = \arg \max_i [\min_j [\max_k Q_{ijk}]] \\ \max_j Q_{ijk} & \text{else} \end{cases} \quad (3.5)$$

Figure 3.4 shows the relationship between \bar{T}_{out} and ΔT_{max} at BOEC and EOEC for the reference core. It is found that in order to achieve the desired core-average outlet temperature of 510°C, the assembly temperature increase constraint must be, at a minimum, nearly 350°C, even in this best-case bounding analysis. A temperature increase of this magnitude is likely unacceptable because it will lead to an assembly outlet temperature of roughly 700°C in the highest powered assemblies, implying peak inner cladding and coolant temperatures in excess of their respective limits accounting for hot spot peaking factors and margins required

for anticipated transient conditions. Therefore, static integral orificing is judged as unable to meet the relevant constraints of B&B cores. Further discussion on this topic is provided briefly in Appendix B in relation to the possibility of reducing the number of batches in the B&B core.




Figure 3.4: Core-average outlet temperature as a function of the assembly-average temperature increase constraint at BOEC and EOEC.

3.3 Assessment of Static Grid Plate Orificing

While there has been significantly less real-world experience with static grid plate orificing in SFRs, the nonapplicability of static integral orificing makes static grid plate orificing potentially more attractive to B&B cores than it might otherwise be. Additionally, the ability for assemblies to receive new orifices between shuffles may make the static grid plate method particularly suited to B&B operation. Therefore, this section examines the static grid plate method in detail for the reference core during equilibrium operation. This section begins by introducing a new optimization technique used to design the orificing scheme, then proceeds to apply the optimization technique using a variety of constraints to determine the sensitivity of the results. All optimizations are performed using the CPLEX solver, developed by IBM [70].

Optimization Technique

It has been shown in previous work [77] that traditional orificing design methods may be insufficient to handle the challenging characteristics of B&B cores. The same study briefly introduced a new method using Mixed-Integer Linear Programming (MILP) to find optimal orificing schemes that is sufficiently flexible to handle a wide range of constraints, objective functions, and both spatial and temporal variations. The proposed formulation is repeated and expanded in the current study, with the model outlined below. Because there is no benefit to differentiating between assemblies of different batches in this formulation, the subscripts i and j , previously used to denote a particular assembly of a particular batch, respectively, are combined into the single subscript $i \in \{1, \dots, n_a n_b\}$.

$$\begin{aligned}
 & \min_{\delta_{i\ell}} \sum_{\ell=1}^{n_x} \beta_{\ell} : \\
 & \textcircled{1} \quad \Delta T_{max} \geq \frac{Q_{i,k}}{c_p \dot{m}_i} \quad \forall i, \forall k \\
 & \textcircled{2a} \quad \sum_{i=1}^{n_a n_b} \dot{m}_i \geq \frac{\sum_i \frac{Q_{i,k}}{c_p}}{\bar{T}_{out} + \gamma - T_{in}} \quad \forall k \\
 & \textcircled{2b} \quad \sum_{i=1}^{n_a n_b} \dot{m}_i \leq \frac{\sum_i \frac{Q_{i,k}}{c_p}}{\bar{T}_{out} - \gamma - T_{in}} \quad \forall k \\
 & \textcircled{3} \quad v_{max} \geq \frac{\dot{m}_i}{\rho A_f} \quad \forall i \\
 & \textcircled{4a} \quad \sum_{\ell=1}^{n_x} (\delta_{i,\ell} \Delta T_{i,k,\ell}) - \sum_{\ell=1}^{n_x} (\delta_{i',\ell} \Delta T_{i',k,\ell}) \leq \xi \quad \forall i, \forall i' : i \perp i', \forall k \\
 & \textcircled{4b} \quad \sum_{\ell=1}^{n_x} (\delta_{i',\ell} \Delta T_{i',k,\ell}) - \sum_{\ell=1}^{n_x} (\delta_{i,\ell} \Delta T_{i,k,\ell}) \leq \xi \quad \forall i, \forall i' : i \perp i', \forall k \\
 & \textcircled{5} \quad \dot{m}_i = \sum_{\ell=1}^{n_x} \delta_{i,\ell} x_{\ell} \quad \forall i \\
 & \textcircled{6} \quad \sum_{\ell=1}^{n_x} \delta_{i,\ell} = 1 \quad \forall i \\
 & \textcircled{7} \quad \delta_{i,\ell} \in \{0, 1\} \quad \forall i, \forall \ell \\
 & \textcircled{8} \quad \beta_{\ell} \geq \delta_{i,\ell} \quad \forall i, \forall \ell
 \end{aligned}$$

The variables of the model are $\delta_{i,\ell} \in \mathbb{Z}_2^{n_a n_b \times n_x}$ and $\beta \in \mathbb{R}^{n_x}$. Constraint $\textcircled{1}$ limits the average temperature rise within an assembly to less than the user-specified ΔT_{max} . Constraint $\textcircled{2}$ forces the mixed average outlet temperature to be within a user-specified tolerance γ of

the desired outlet temperature \bar{T}_{out} . Constraint (3) limits the average flow velocity within an assembly to be less than the user-specified v_{max} , where ρ and A_f are the coolant density and assembly flow area, respectively. Constraint (4) limits the difference between outlet temperatures of adjacent assemblies, where the symbol ‘ \perp ’ is meant to indicate ‘adjacent-to’ and ξ is the user-specified constraining value. Constraints (5), (6), and (7) force the model to choose a single flow rate in each assembly from a pool of pre-specified flow rates, as further discussed below. And finally, constraint (8), in conjunction with the objective function, leads the model to optimize for the smallest number of unique flow rates among core assemblies (referred to as the smallest number of ‘orifice groups’). This is chosen as the objective function in the static grid plate case to avoid practical issues with having an overly complicated grid plate.

In conjunction with constraint (5), the user must specify $x \in \mathbb{R}_{++}^{n_x}$, termed the ‘discretized flow vector’, a vector of discretized flow rates which may be adopted in any of the core assemblies. This discretized vector allows for the physics to be linearized by facilitating the building of tables for $\Delta T_{i,j,k}$, which otherwise would be nonlinear functions of \dot{m}_i in Constraints (4). While not strictly required, the discretization of x also allows for tables of material properties, such as c_p and ρ , to be built using temperature-dependent data corresponding to conditions in each individual assembly. This use of temperature-dependent properties allows for more rigorous physics to be incorporated into the model, and is thus employed for all calculations in this study.

With constraint parameters specified, the task of the optimization solver is to use the binary $\delta_{i,j}$ variables to select which values of x should be chosen for each assembly to provide the minimum number of orifice groups while adhering to the specified constraints. Because the user may specify data from multiple points in the cycle for a single calculation (or even points from various burn cycles if considering non-equilibrium conditions), the MILP method is very flexible to account for changes in power profiles and shuffling schemes.

Requiring the user to specify the values in x may become a source of suboptimality if n_x is too small. Therefore a ‘mesh refinement study,’ such as the example shown in Figure 3.5, should be performed to ensure that the results are independent of n_x . Although it would be safest to simply use a huge number for n_x , Figure 3.5 also shows that the time taken to solve the optimization problem increases nearly exponentially as n_x gets larger. Such behavior is typical of MILP formulations, and motivates the user to make n_x as small as possible while still achieving sufficiently converged results. In this study for the static grid plate method, all calculations take $n_x = 200$, using logarithmically-spaced elements between 0.01 and 100 kg/s for x .

Quantification of Assembly-to-Assembly Heat Transfer

Because the optimization model is restricted to simplified, linearized physics, a number of approximations are embedded, the most significant of which is the neglect of assembly-to-assembly heat transfer. Predicting outlet temperatures considering assembly-to-assembly




Figure 3.5: Results of an example orificing optimization mesh refinement study using logarithmically-spaced elements between 0.01 and 100 kg/s for x .

heat transfer is highly nonlinear and quickly becomes too complicated for optimization approaches such as MILP. Because its presence will serve to smear peak temperatures closer to the average, neglecting assembly-to-assembly heat transfer is expected to be conservative from the perspective of minimizing the number of orifice groups. However it is important to understand the error introduced through this approximation to ensure that the results are still meaningful.

To quantify the amount of heat transfer taking place between adjacent assemblies, the Finite Element (FE) code SAM from Argonne National Laboratory is used [68]. Within SAM, a model of two thermally-connected assemblies is created as depicted in Figure 3.6. Each assembly is discretized into 24 axial segments and concentric hexagons to capture the within-assembly temperature gradients. Each hexagonal channel is modeled as a 1D pipe with energy exchange between adjacent channels according to a mixing model proportional to the channel interface area and the difference between channel enthalpies. Each assembly is surrounded by HT9 duct, adiabatic on all sides except at the interface between the two assemblies, where heat transfer is permitted via conduction only across the inter-assembly gap. The method of discretization and the decision to allow only conduction in the inter-assembly gap were inspired by the studies of Yang and Joo [154] and employed to save on computational costs.

Figure 3.7 shows how the difference in outlet temperatures between the two assemblies varies as the power-to-flow ratio is changed in one of the two assemblies, as calculated by SAM. The plot shows the difference when assembly-to-assembly heat transfer is considered compared to when it is neglected in two cases: (1) power is held constant while flow is reduced




Figure 3.6: Diagram of the discretization method used in the study of assembly-to-assembly heat transfer with SAM. 3.6a shows the discretization used for a single assembly, while 3.6b shows how two assemblies are connected together. Diagram is not to scale, and actual FE models model the geometry as described in Chapter 2.

in one of the two assemblies, and (2) flow is held constant while power is increased in one of the two assemblies. The results indicate that the difference between cases (1) and (2) is negligible, and only the power-to-flow ratio is important. Additionally, the results show the impact of assembly-to-assembly heat transfer to be nearly linear over the domain of interest. From Figure 3.7, it is concluded that neglecting assembly-to-assembly heat transfer results in an error of roughly 5% for the geometry of the reference core. Further studies have shown this result to be insensitive to perturbations in duct and gap thicknesses, although a full sensitivity study is outside the scope of this analysis.

While this study examines the case of two assemblies in isolation, any real system will consist of many more adjacent assemblies and a much more complicated pattern of assembly-to-assembly heat transfer, which will serve to further smooth the differences between adjacent assemblies' outlet temperatures. Moreover, as ξ is reduced, the absolute error introduced by the model is reduced. In light of this, the error derived from the neglect of assembly-to-assembly heat transfer in the MILP formulation is judged as acceptable.

Results and Sensitivities

The nominal constraints used for the static grid plate orificing method are given in Table 3.1. The value of v_{max} is taken from Heidet [54] and is primarily to account for flow-induced vibration, which for the geometry in this design is more constraining than a pressure drop constraint of 1 MPa. The values of ΔT_{max} , γ , and ξ are derived from a variety of sources, including statements made by TerraPower [57], but because these constraints are flexible, they are subsequently examined as part of a sensitivity study.

Figure 3.8 shows the results of the optimization. With the nominal constraints, it is found

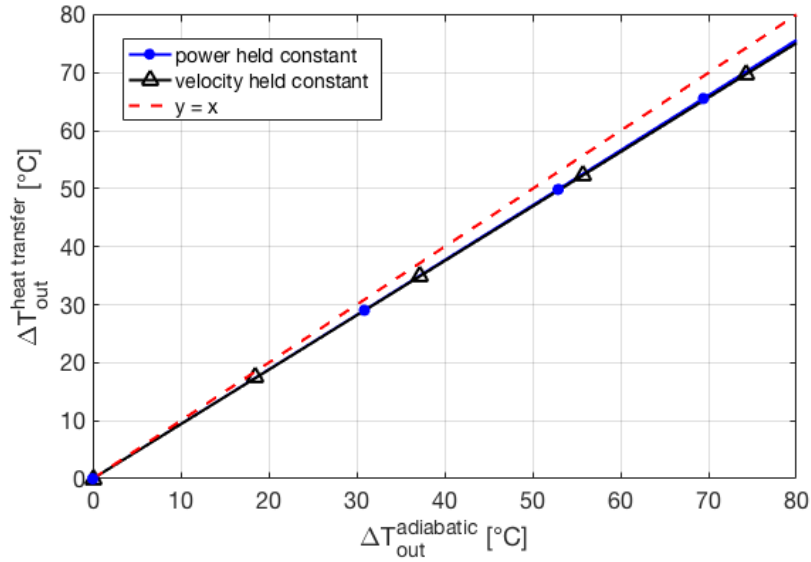


Figure 3.7: Difference between outlet temperatures as the power-to-flow ratio is varied in the scenario depicted in Figure 3.6 when assembly-to-assembly heat transfer is considered, as a function of the same temperature difference when assembly-to-assembly heat transfer is neglected.

Table 3.1: Nominal constraints used for the static grid plate orificing optimization.

Symbol	Parameter	Value	Units
ΔT_{max}	Max coolant temperature rise	210	$^{\circ}C$
T_{in}	Inlet coolant temperature	355	$^{\circ}C$
\bar{T}_{out}	Target average outlet temperature	510	$^{\circ}C$
γ	Tolerance on target outlet temperature	5	$^{\circ}C$
v_{max}	Max coolant velocity	12	m/s
ξ	Max difference in outlet temperature of adjacent assemblies	50	$^{\circ}C$

that 10 orifice groups will be required for equilibrium conditions, laid out as shown in Figure 3.9. Comparing against the batch layout shown in Figure 2.9, stark similarities are seen. For instance, Batch 1 has a one-to-one correspondence with the central orifice group. Similarly, the control assemblies are all within a single orifice group. The region containing Batches 2-5 are largely grouped into a single orifice group, which makes sense considering that Batches 2-5 contain the vast majority of power production and all experience relatively large power swings from BOEC to EOEC. Conversely, because the region towards the core periphery has many assemblies at very low powers, the constraints on adjacent outlet temperature forces

multiple orifice groups between the annular high-flow region and the outer low-flow region to ‘step down’ the outlet temperature gradient.

While these trends could be used to guide the designer in an ad-hoc method for determining the orificing, devising a strategy with so few groups is a task complicated in both space and time. Additionally, due to the mathematical structure of MILP, the optimal solution returned by the solver is guaranteed to be a global minima, meaning no amount of designer ingenuity can further reduce the number of orifice groups, within the bounds of the problem specified. For this reason, the proposed optimization method is thought to be particularly useful.




Figure 3.8: Results of the optimization in the static grid plate case as the constraints are varied. As one constraint is varied, the others are kept at the nominal values of Table 3.1.

Figure 3.8 also shows the sensitivity study results. The number of orifice groups are particularly insensitive to variations in γ , meaning it is likely not necessary to allow the mixed outlet temperature to vary much from the desired value. Conversely, the results are particularly sensitive to ξ . As ξ is decreased from its nominal value, the number of orifice groups increases sharply. In fact, when $\xi \leq 30^\circ\text{C}$ the problem becomes infeasible, no matter how many orifice groups are used. This fact is highly relevant to thermal striping considerations and should be considered in the design of above-core structures. Similarly, as ΔT_{max} is tightened, the number of orifice groups sharply increases. This is because, with lower maximum outlet temperatures, it becomes more difficult to compensate for the depressed outlet temperatures of assemblies when they are at low power. However, whereas




Figure 3.9: The layout of orifice groups using static grid plate orificing along with the nominal constraints of Table 3.1 for the equilibrium cycle.

the number of orifice groups could be decreased by decreasing ξ from its nominal value, ΔT_{max} sensitivity shows a lower limit of 10 groups for $\Delta T_{max} \geq 190^\circ C$, indicating marginal gains to be made by loosening this constraint. This analysis shows that it is indeed feasible to utilize static grid plate orificing in B&B reactors, although the number of orifice groups may need to be somewhat larger than that used by other SFR designs.

Demonstration of Assembly-Specific Constraints

In the studies so far, all constraints were applied uniformly to all assemblies, when in reality different assemblies have experienced different levels of burnup, endowing them with differing material properties which may induce differing constraints. For instance, assemblies with very low burnup may have an open fuel-cladding gap, changing the heat transfer characteristics as compared to more heavily burned assemblies. Because static grid plate orificing allows assemblies to receive new orifices between shuffles, the constraints defining the strategy can be considered according to time- and burnup-dependent assembly characteristics. This aspect of static grid plate orificing is further explored in this section.

Because the optimization model applies constraints only to average outlet temperatures, when determining the constraint limits it is necessary to draw relationships between peak and average temperatures. A more detailed analysis might employ a coupled fuel-

performance/thermal-hydraulics solver to draw these relations, but in the demonstrative analysis of this study, a simplified approach is taken. SAM is used stand-alone to model a single assembly with burnup-dependent fuel conductivity and gap sizes, and peak temperatures as a function of the assembly average outlet temperature are determined for each batch. All assemblies of a given batch are assumed to have the same burnup, and the highest-powered assembly of each batch is used to determine the batch-constraining values.

The burnup-dependent fuel thermal conductivity is determined as a function of the weight fraction of plutonium according to Equation 3.6, taken from the IFR handbook [61]. In this relation, W_Z and W_P are the weight fractions of zirconium and plutonium, respectively. To account for the effect of increasing porosity, Equation 3.7 can be used, where k_0 is the result of Equation 3.6, P is the fractional porosity, and β is an empirical factor, taken to be 2 as recommended in the source.

$$k_0[W/m/K] = A + BT + CT^2 \quad : \quad \begin{aligned} A &= 17.5 \left(\frac{1 - 2.23W_Z}{1 + 1.61W_Z} - 2.62W_P \right) \\ B &= 1.54 * 10^{-2} \left(\frac{1 + 0.061W_Z}{1 + 1.61W_Z} + 0.90W_P \right) \\ C &= 9.38 * 10^{-6} (1 - 2.70W_P) \end{aligned} \quad (3.6)$$

$$\frac{k}{k_0} = \frac{1 - P}{1 + \beta P} \quad (3.7)$$

For batches with burnup less than 2%FIMA, the gap is assumed to remain open as recommended by Hartanto et al. [52], and the porosity is assumed to be linearly related to burnup. In batches where the gap has closed, the porosity is assumed to be equal to the fuel smear density. With these simplifying assumptions, the batch-dependent fuel thermal conductivity is determined as shown for select temperature values in Table 3.2.

Table 3.2: Assumed fuel thermal conductivity (W/m/K) at select temperatures for each batch, along with fuel burnup.

Batch	1	2	3	4	5	6	7	8	9	10	11	12
Burnup (%FIMA)	2.3	12.5	19.8	26.5	31.0	6.84	3.77	0.7	0.2	0.1	0.0	0.0
623 K	15.3	14.2	14.0	14.1	14.1	14.5	15.0	16.1	16.5	20.8	22.3	23.2
823 K	19.1	17.8	17.6	17.7	17.7	18.2	18.6	19.9	20.3	25.6	27.4	28.5
1023 K	23.2	21.8	21.6	21.7	21.8	22.2	22.8	24.1	24.6	31.1	33.3	34.5

With this framework, SAM is used to determine peak temperatures in each batch as a function of the average outlet temperature, as shown in Figure 3.10 for the fuel, cladding,

and coolant. For the case of peak coolant temperatures, the relationship with average outlet temperature is only a function of the power-to-flow ratio, so all batches follow the same curve. It is seen that the fuel-cladding eutectic limit is the most constraining for every batch. Due to the very high thermal conductivity of metallic fuel, even after significant burnup, peak fuel temperatures are much less restrictive than peak cladding temperatures.




Figure 3.10: Peak temperatures versus average outlet temperature for the peak-power assembly in each batch, along with constraint imposed by fuel melting, fuel-clad eutectic reactions, and coolant boiling.

Using constraints derived by taking the lowest average outlet temperature at which a given batch violates one of the three constraints shown in Figure 3.10, the static grid plate orificing optimization is repeated for two more cases. In Case A, all assemblies are con-

strained according to the requirements of the most constraining batch, while in Case B each batch is constrained according to its own burnup-dependent requirements. In order to save on computational effort, the mixed outlet temperature and adjacent outlet temperature difference constraints are tightened to $\gamma = 1^\circ C$ and $\xi = 40^\circ C$ for both cases. Table 3.3 shows that, by considering batch-wise material properties, the number of orifice groups may be reduced substantially, highlighting one of the additional benefits of static grid plate as compared to static integral orificing.

Table 3.3: Minimum number of orifice groups using a static grid plate orificing scheme for: (A) all assemblies being constrained uniformly according to the most limiting batch, and (B) each batch being constrained according to its own burnup-dependent characteristics.

Case	Number of Groups
A	25
B	20

3.4 Assessment of Variable Orificing

The final method investigated in this study is variable orificing. This section performs a similar analysis as was done for the case of static grid plate orificing to determine the tradeoffs between the two methods.

Optimization Technique

The optimization procedure used for this analysis is largely the same as was outlined in Section 3.3, with one key difference – because the orifices can now be changed, the concept of an ‘orifice group’ loses its meaning. Therefore it does not make sense to minimize the number of orifice groups within the paradigm of variable orificing, and the objective function should be changed. Although any variable or constraint within the model could be chosen as the objective due to its linear structure, this analysis chooses to minimize ξ , the maximum difference in outlet temperatures between adjacent assemblies. Such an objective function is chosen to help minimize thermal striping concerns, one of the key challenges for B&B designers. Additionally, the sensitivity study in Section 3.3 revealed ξ to be the most sensitive parameter, indicating that the strongest benefit could be derived by minimizing this parameter. With this change of objective function, the new optimization model formulation becomes:

$\min_{\delta_{i\ell}} \xi :$

$$\textcircled{1} \quad \Delta T_{max} \geq \frac{Q_{i,k}}{c_p \dot{m}_i} \quad \forall i, \forall k$$

$$\textcircled{2a} \quad \sum_{i=1}^{n_a n_b} \dot{m}_i \geq \frac{\sum_i \frac{Q_{i,k}}{c_p}}{\bar{T}_{out} + \gamma - T_{in}} \quad \forall k$$

$$\textcircled{2b} \quad \sum_{i=1}^{n_a n_b} \dot{m}_i \leq \frac{\sum_i \frac{Q_{i,k}}{c_p}}{\bar{T}_{out} - \gamma - T_{in}} \quad \forall k$$

$$\textcircled{3} \quad v_{max} \geq \frac{\dot{m}_i}{\rho A_f} \quad \forall i$$

$$\textcircled{4a} \quad \sum_{\ell=1}^{n_x} (\delta_{i,\ell} \Delta T_{i,k,\ell}) - \sum_{\ell=1}^{n_x} (\delta_{i',\ell} \Delta T_{i',k,\ell}) \leq \xi \quad \forall i, \forall i' : i \perp i', \forall k$$

$$\textcircled{4b} \quad \sum_{\ell=1}^{n_x} (\delta_{i',\ell} \Delta T_{i',k,\ell}) - \sum_{\ell=1}^{n_x} (\delta_{i,\ell} \Delta T_{i,k,\ell}) \leq \xi \quad \forall i, \forall i' : i \perp i', \forall k$$

$$\textcircled{5} \quad \dot{m}_i = \sum_{\ell=1}^{n_x} \delta_{i,\ell} x_\ell \quad \forall i$$

$$\textcircled{6} \quad \sum_{\ell=1}^{n_x} \delta_{i,\ell} = 1 \quad \forall i$$

$$\textcircled{7} \quad \delta_{i,\ell} \in \{0, 1\} \quad \forall i, \forall \ell$$

In theory, variable orificing could be used to change individual orifices on-the-fly, even during full-power operation, although it may be challenging from logistical and regulatory points of view to operate in this way. Additionally, such a scheme would obviate the role of the orifice designer. Therefore, this study makes the assumption that orifice changes are only executed between cycles and again optimizes the orificing scheme for the equilibrium cycle. In the case that on-the-fly changes are allowed, the above model could be used as part of an optimal control system to control the orifices, as it still takes into account all constraints which must be considered when making changes to the orificing.

Results and Sensitivities

Figure 3.11 shows the results for the variable orificing case. As the constraints are varied around their nominal values, the results are very insensitive. In fact, as the constraints are loosened, ξ does not improve at all, implying that the value $\xi = 31^\circ\text{C}$ is the absolute lower limit achievable in the reference B&B core. This result agrees with the results of Figure 3.8, where no solution could be found for $\xi \leq 30^\circ\text{C}$. The implication is that thermal

striping may only be mitigated through a reduction of ξ up to a certain point, and any further mitigation will need to come from other sources. Regardless, if the lower bound is acceptable, Figure 3.11 shows a very large design space available through the variable orificing method. Because the orifices may be changed in the variable orificing paradigm, the constraints applied to individual assemblies may be varied as assemblies accumulate burnup, and thus similar benefits as demonstrated in Section 3.3 are expected to apply to the variable orificing case as well.




Figure 3.11: Results of the optimization in the variable orificing case as the constraints are varied. As one constraint is varied, the others are kept at the nominal values of Table 3.1.

3.5 Summary of Orificing Strategies

This chapter has investigated three different methods of orificing and their applications to a prototypical B&B core. Through bounding analytical calculations, it was demonstrated that static integral orificing is likely not suitable. By using a MILP optimization technique, the static grid plate and variable orificing methods were investigated in detail. While the MILP method utilizes simplified physics models, supplemental FE thermal-hydraulic calculations have demonstrated the introduced error to be minimal.

It was found that the two latter orificing methods are well-suited to B&B operation, largely because they allow for the orifice of a given assembly to be changed between shuf-

fling operations. In the static grid plate method, it is thought that as few as 10 orifice groups may be able to accommodate the equilibrium state power evolution using realistic constraints, although strong sensitivities to constraining parameters, particularly the difference between adjacent assembly outlet temperatures and the temperature rise across the core, were seen which may cause the number of required orifice groups to be nearly 20 if sufficiently tight constraints are imposed. Similar sensitivities were not seen in the variable orificing method, although it was discovered that a lower bound of approximately 30°C exists on the difference between adjacent assembly outlet temperatures. In comparison to the static grid plate method, variable orificing offers significantly higher flexibility to accommodate off-equilibrium operations, at the expense of higher operational complexity. In order to gain some amount of flexibility in the static grid plate method without the operational complexity of variable orifices, a hybrid method where static integral and grid plate orifices could be proposed. For both methods, one major benefit is that constraints may be applied specifically to each assembly, which may open up the orificing design space. In either case, the feasibility for orificing B&B reactors using conventional means has been affirmed and a new computational tool for designing orificing strategies has been demonstrated.

For the remainder of this dissertation, static grid plate orificing is assumed due to its operational simplicity. The orifice grouping depicted in Figure 3.8 are utilized in all further discussions and calculations having to do with B&B transient performance in Chapters 4-5.

Chapter 4

Nominal Transient Performance

A crucial aspect in the commercialization of any reactor system is an understanding of how the system will respond to a variety of transient initiators. So far in the open literature, detailed assessments of the inherent safety of B&B cores have been virtually absent. On the one hand, metal fueled SFR cores have a long history of excellent transient performance which bodes well for B&B designs. However, a number of B&B design aspects are significantly different than for a typical metal SFR (namely their large size, strongly positive coolant void worth, and unique power profiles), which makes it unclear if previous SFR transient results will translate over to B&B designs. This chapter provides an exploration of the transient performance in the reference B&B core, examining both the nominal performance and a number of sensitivities to important assumptions.

4.1 Transient Modeling Approach

Numerical Model

For the modeling of transient behavior, the SAS4A/SASSYS-1 code system (referred to as SAS for short) is used [25]. SAS has a long history of use in SFR transient analyses for both metallic and oxide cores, and has been benchmarked against a number of experiments performed in EBR-II and FFTF, including unprotected loss of flow, unprotected loss of heat sink, and unprotected rod withdrawal (sometimes referred to as ‘transient overpower’) scenarios (see, for instance, [134, 135]). This history of successful application in similar reactor systems makes SAS the ideal candidate for assessing the performance of the B&B core.

SAS uses a limited number of representative core channels and pipe-network models in order to predict the whole-plant transient behavior. Each core channel represents the average behavior of the assemblies represented by that channel using a single flow channel model (as compared to, for instance, a sub-channel model of an entire assembly), as depicted in Figure 4.1. In this model, all fuel pins are smeared into a single average fuel pin, with

representative amounts of cladding, coolant, and structure (wire-wraps are smeared into the cladding). As heat is generated in the fuel, 1D conduction is allowed to the radially-adjacent nodes from the fuel, through the cladding, and into the coolant, with further conduction allowed into the structure material (represented by a 1D slab). As the coolant flows upward, conduction is allowed back to the solid components, allowing for a psuedo-2D treatment within each channel. Although the model allows for heat generation in the cladding and coolant, all power is assumed to be generated in the fuel. Below and above the fueled region are reflectors and plena, which are treated in a similar manner as the fueled region, except without any heat generation.




Figure 4.1: Channel model used to represent the core heat transfer in the transient model.

Assemblies are grouped into 7 representative channels for the simulations, as depicted in Figure 4.2. Grouping was done on the basis of burnup characteristics, reactivity coefficients,

flowrates, and power production, attempting to put assemblies of similar characteristics together. Therefore, the channel layout used in the SAS model closely resembles both the core power profile and orificed flow map depicted in Figures 2.13 and 3.9, respectively. Fuel assemblies are grouped into 4 channels to capture the power gradient across the core, with two additional channels devoted to the peak-power and peak power-to-flow assemblies to capture peak temperatures during the transient. A final channel was allocated for reflector and control assemblies to ensure that their low power and importance levels are not smeared into channels representing fueled assemblies. The number of radial and axial nodes in each component are provided in Table 4.1.




Figure 4.2: Grouping of assemblies into channels for the transient model of the B&B core.

The primary loop and balance of plant systems are taken to be the same as that for the Advanced Burner Reactor (ABR) designed at Argonne National Laboratory [80], scaled up to meet the higher power level of the B&B core. An overview of the plant systems is presented in Figure 4.3. The primary loop consists of four pumps taking suction from the cold pool and directing the flow through the core, where it is heated. Core flow exits into the outlet plenum, flows over the control rod drives, and then flows through the shell-side of the Intermediate Heat Exchangers (IHXs) before again entering the cold pool. While Figure 4.3 only shows a single intermediate loop, in actuality there are four identical loops. These four loops provide the primary heat removal capability to the core. As the intermediate loop flow

Table 4.1: Nodalization of the core channel model for the B&B core.

Region	Axial Nodes	Radial Nodes
Upper Reflector	5	2
Gas Plenum	4	1
Fuel	26	10
Cladding	26	3
Coolant	40	1
Structure	40	2
Lower Reflector	5	2

traverses through the IHXs, it accepts heat from the primary loop flow. The intermediate flow is then pumped up and out of the pool to the steam generators, where it gives up its heat to produce steam before being pumped back down to the IHXs. The water-side of the steam generators is not modeled in SAS, and instead boundary conditions are applied to the secondary loop at the steam generators to impose either a constant temperature drop or a constant outlet temperature, depending on the particular transient.




Figure 4.3: General plant layout for the reference core.

In the absence of forced circulation in the intermediate loops, reactor heat is removed through the Direct Reactor Auxiliary Cooling System (DRACS). Again, although only one DRACS loop is depicted in Figure 4.3, there are actually two independent DRACS loops.

Each DRACS loop is designed to remove 0.5% of core power at nominal conditions when operating at full capacity. The DRACS loops operate similarly to the intermediate loop, with primary flow traversing through the shell-side of the DRACS Heat Exchangers (DHX) as it moves from the outlet plenum back to the cold pool. However, unlike the intermediate loop, the DRACS loop utilizes NaK as the working fluid. While the DRACS loops are only meant to remove heat during an accident scenario, they operate continuously via natural circulation of the NaK working fluid. To minimize losses during normal operation, dampers are placed on the air-side of the NaK-to-air heat exchangers that act as the ultimate heat sinks for the DRACS loops. In the case of an accident scenario, these dampers may be opened to allow the DRACS loops to operate at full capacity.

The core power evolution is calculated using a point kinetics approach. The transient model accounts for reactivity feedback from axial expansion of the fuel, cladding, and structure; radial expansion of the grid plate; doppler broadening with fuel temperature; coolant density with coolant temperature; and thermal expansion of the control rod drives. The axial expansion feedbacks are calculated by using the temperature of each axial node to thermally expand the fuel, cladding, and structure. Then, the worth curves for each constituent (Figures 2.18a-2.18c) are used to calculate a reactivity response. Because the constituent worths in each channel may vary, the spatial-dependence is accounted for on a channel-by-channel basis. Similar methods are used to calculate the radial expansion feedback, where a 0D temperature for the grid plate is combined with the temperature of the upper load pads to determine how far the core has expanded. The control rod driveline feedback is determined according to the differential expansion of the control rod drivelines and the core barrel. As core flow washes over the control rod drivelines, a 0D temperature is calculated and the drivelines are thermally-expanded. Similarly, a 0D model is used to determine the temperature of the core barrel, from which the grid plate is suspended. The differential expansion determines how far the control rods are inserted with time, which is used in conjunction with a control rod worth curve (Figure 2.17) to determine a reactivity impact. The assembly bowing feedback mechanism is not accounted for. Details of the numerical implementation of the thermal-hydraulics and reactivity feedback models are found in the SAS theory manual [25].

Transient Initiators and Assumptions

Three different transients are examined in this dissertation – the (1) Unprotected Loss of Heat Sink (ULOHS), (2) Unprotected Loss of Flow (ULOF), and (3) Unprotected Transient Overpower (UTOP). For all transients, the core begins at nominal steady-state conditions, as outlined in Chapter 2. Due to the rapid swelling of metallic fuel at low levels of burnup [104], all channels are assumed to have closed fuel-clad gaps. The core shutdown system is assumed to fail and a single DRACS loop designed for capacity of 0.5% of core power is passively engaged by opening the dampers on the NaK-to-air heat exchanger immediately upon transient initiation. The ambient air temperature for the heat exchanger is assumed to be 27°C. The second DRACS loop is assumed to fail. All transients are simulated at EOE

where the worth vested in the control rods is greatest. The assumptions that differ between the three transient scenarios are outlined in Table 4.2.

Table 4.2: Assumptions specific to each of the three transients examined in the B&B core.

	ULOF	ULOHS	UTOP
Primary Loop Heat Rejection	Nominal	Zero	Nominal
Pump Speed	20 s halving time	Nominal	Nominal
Control System	No scram	No scram	\$1.66 inserted at 0.005 \$/s
Thermal Stratification	Cold and hot pools	None	None

In the ULOHS, heat rejection through the steam generators is assumed to immediately cease, with core flow remaining at nominal and thermal stratification assumed to be negligible because the pumps remain operating at full capacity. In the ULOF, the pump motors are assumed to trip and core flow decreases with the impeller inertia until reaching a lower threshold value, at which time the rotors are assumed to lock in place and the transition to natural circulation occurs. The same flow halving time of 20 s is assumed for all pumps, both primary and intermediate¹. Because of the severe decrease in flow as the transient progresses, thermal stratification models for the cold and hot pools are used. In the UTOP, a reactivity insertion equivalent to the worth of the highest-worth rod is withdrawn, with full flow and no thermal stratification assumed. The rate of control rod withdrawal is assumed to be 0.005 \$/s, as was used in the safety analyses for the IFR [146], with subsequent sensitivity studies performed on this parameter. In both the ULOF and UTOP, the steam generator outlet temperature is assumed to remain constant.

4.2 Nominal Performance

Unprotected Transients

Due to a very strong positive coolant reactivity feedback, the nominal transient performance in the B&B core is poor, as summarized in Table 4.3. Both the ULOF and UTOP transients are seen to boil, with boiling initiation taking place within 20 minutes and 3 minutes for the ULOF and UTOP, respectively. To understand the mechanisms driving this result, the transient evolution for each of these three scenarios will now be explored in more detail.

Results of the ULOHS are given in Figure 4.4. Immediately upon transient initiation, the core power level begins to decrease due to a strongly negative radial expansion feedback, which acts promptly due to the core gridplate being the first component that the

¹This coastdown time is in line with those assumed by other researchers for metallic SFRs (see, for instance, [157, 33])


Table 4.3: Time of boiling initiation in the B&B core for various unprotected transients.

	ULOHS	ULOF	UTOP
Time of boiling (s)	-	1284	150

elevated-temperature inlet flow encounters. With little delay, however, the positive coolant reactivity feedback begins to counter, leading to a net reactivity that is only just slightly negative, especially as the core barrel heats up in response to the elevated coolant temperatures, eventually leading to a positive CRDL feedback. After remaining in a state of slight subcriticality, the core eventually approaches a pseudo-equilibrium near 4000 s in line with the reactor decay heat level, with elevated coolant and cladding temperatures due entirely to the loss of heat rejection capability, but with reduced peak fuel temperatures as compared to steady-state. In general, this transient is both smooth and benign, resulting in temperatures well within margins of any detrimental consequences. However, it is noted that the coolant feedback is the strongest overall feedback mechanism, keeping the net reactivity near zero for much of the transient even as the majority of remaining feedbacks are negative. Such a strong role played by the coolant density feedback will be a common thread throughout all B&B transient analyses.


The transient evolution in the ULOF is presented in Figure 4.5 until the time of boiling. Contrary to the ULOHS scenario, the first feedback mechanism to act is the coolant density due to the increased power-to-flow ratio as the pumps begin to coast down, which leads to a small, temporary power increase and an associated fuel temperature peak. Rapidly after that, the increased core temperatures act to turn the remaining feedbacks negative, bringing core power down at a somewhat slower pace than the core flow reduction. This mismatch between power and flow leads to an increase in coolant and cladding temperatures, but the peak fuel temperature simultaneously decreases due to the reduction in core power. Eventually just after 200 s, core temperatures reach a temporary steady-state as the power and flow continue to decrease in unison. Throughout this, most of the negative reactivity is driven by the CRDL associated with the increased coolant outlet temperature. Just as the core is beginning to level out at a power near 10% of nominal, the low mass flow through the tripped pumps leads to a sudden stop of the rotors at almost 1000 s in the primary loop, as the pump inertia decreases below the strength of the frictional losses on the impeller. Such ‘pump jamming’ causes a sudden drop in core flow, which leads to a spike in coolant temperature and a resulting strong reactivity insertion from the reduced coolant density. Ultimately, this sets off a series of strong diverging temperature-reactivity oscillations which lead to coolant boiling around 1300 s. It is emphasized that if the coolant void reactivity were not so strong, these oscillations would be much smaller in magnitude and would likely die out, as is seen in other systems like the ABR [46]. The role of the pump design in the progression of the ULOF is investigated further in Section 4.2.

The ULOF evolution is qualitatively in agreement with the limited results reported by



(a) Power and flow

(b) Reactivity



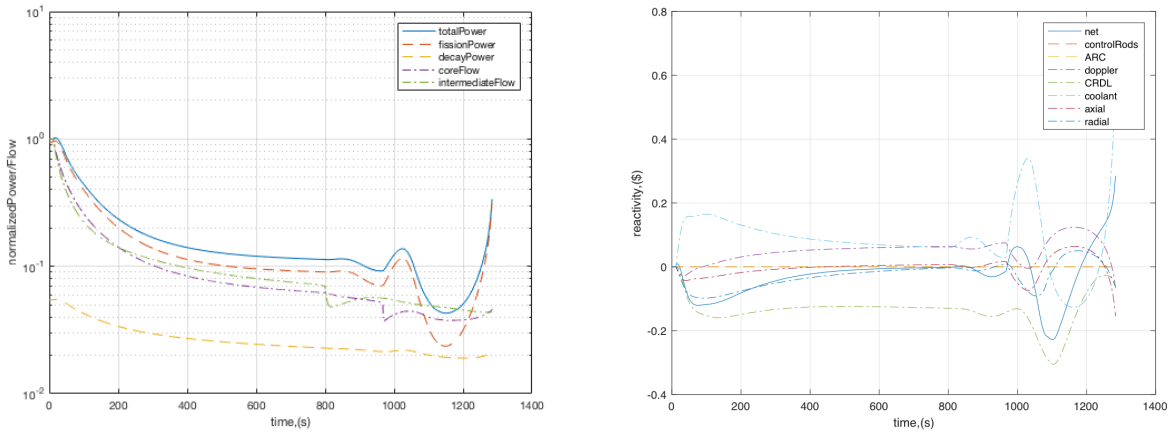
(c) Temperatures

Figure 4.4: Transient response to the ULOHS scenario in the nominal B&B core.

TerraPower, LLC for their prototype B&B core design, dubbed the TWR-P [57]. They report results only for the first 1000 s of a ULOF accident, although details such as the pump coast-down halving time and the presence of rotor-locking are not provided. For the first 1000 s, however, the two sets of results are generally in agreement, although they report a peak inner cladding temperature in excess of the fuel-cladding eutectic temperature during the first few minutes of the transient, whereas such behavior is not predicted here².


The transient evolution in the UTOP is presented in Figure 4.6 until the initiation of

²The author can only guess that the flow halving time employed by TerraPower, LLC is slightly less than what is used in this dissertation, since they see higher temperatures and slightly stronger reactivity feedbacks than what are shown in Figure 4.5b



(a) Power and flow

(b) Reactivity




(c) Temperatures

Figure 4.5: Transient response to the ULOF scenario in the nominal B&B core.


boiling. The reactivity evolution is dominated by the withdrawal, with other feedbacks being delayed in their response due to thermal inertia of components. This leads to a positive net reactivity for the duration of the transient and an associated monotonic power increase up to roughly 120% of nominal by the time coolant boiling occurs at 150 s. In addition to the reactivity from the rod withdrawal, the coolant reactivity feedback provides the only other positive contribution to the net reactivity. The strongest negative reactivity component comes again from the CRDL, which is particularly strong at the point in the cycle that these transients are simulated (EOEC) due to the core's positive reactivity swing, which means that the rods are inserted to roughly the core midplane at the time of transient initiation. The temperature evolution is also monotonically increasing in response to the smooth power

increase. It is noted that the peak fuel temperatures are seen to be the highest in the UTOP out of the three unprotected transients. Although peak fuel temperature does not *quite* reach melting by the time boiling takes place, the fuel temperature is significantly higher than the fuel-clad eutectic temperature (variously reported to be somewhere between 650-750°C [136, 156, 78, 126, 153, 60]) for much of the transient. In addition, it is likely than an uncertainty assessment might find that local fuel temperatures do in fact reach the melting point before coolant boiling is achieved. Regardless, the point is moot because coolant boiling is seen to occur within only a handful of minutes of transient initiation.



(a) Power and flow

(b) Reactivity



(c) Temperatures

Figure 4.6: Transient response to the UTOP scenario in the nominal B&B core.

This section has found the nominal transient performance to be unacceptable from the perspective of passive safety, with the main challenge being the strongly positive coolant

reactivity feedback, although the impacts of assumptions on the results have not been investigated. The next section will investigate these assumptions to understand if the B&B core may be made to have acceptable unprotected transient performance if the assumed plant design were modified.

Unprotected Transient Sensitivities

Pump Coast Down Time

One assumption that has a large impact on the ULOF scenario is the inertia of the primary pumps, which impacts the flow halving time. The nominal scenario assumes the pump coast down halving time to be 20 s, although this value could be changed in practice either by selecting a different pump or by modifying the pump with flywheels to increase their inertia. To evaluate the sensitivity to pump inertia, a study was performed in which the pump coast down time was parametrically varied away from 20 s. Figure 4.7 shows the time of boiling initiation in the ULOF scenario as a function of primary pump coast down time. Boiling is encountered with all of the examined coast down times, however increasing the coast down time can push off boiling by a few minutes, if halving times as large as 30 s can be physically realized. The two jumps in Figure 4.7 result from the transient progressing slowly enough for the control rod driveline and the radial expansion feedbacks to play effective roles in delaying the runaway boiling behavior. Regardless, the nominal ULOF performance is poor even with increased coast down times as a result of the oscillatory behavior set off by the transition from forced- to natural-circulation.

Control Rod Driveline Feedback

The control rod driveline feedback is an uncertain parameter that has a large impact on the transient progression and boiling time, particularly at EOEC when the feedback is strongest. Figure 4.8 shows the time of boiling in the ULOF and UTOP as a function of the control rod drive thermal expansion coefficient. As the expansion coefficient increases from its nominal value of $2^{\circ}C^{-1}$, the rods are inserted further into the core for a given change in outlet temperature, implying a stronger feedback mechanism. While this method is purely a numerical technique for uniformly increasing the strength of the feedback, it could in fact be implemented in a real design by changing the material which the rod drives are made out of. The trend for the UTOP scenario is similar to the sensitivity to the flow-halving time – as the control rod driveline feedback increases in strength, the time of boiling can be pushed off, although never by more than a few minutes. However the trend in the ULOF case is less clear, showing limited improvement at first, followed by slight performance degradation as the control rod driveline feedback increases in strength. Regardless, no realistic increases in control rod driveline are able to prevent boiling in either the ULOF or UTOP.




Figure 4.7: Time of boiling initiation in the ULOF transient of the B&B core with varying primary pump coast down halving times.

Rate of Reactivity Insertion

Because boiling occurs so early in the UTOP, a sensitivity study was performed on the rate of reactivity insertion. Because sodium reactors are operated near atmospheric pressure, control rods are not rapidly ejected as in light water reactors. Instead, the speed of control rod withdrawal during a UTOP would be largely the result of the control rod drive motor design. The nominal value of $0.005 \text{ } \$/\text{s}$ was taken from studies of the Integral Fast Reactor (IFR) [146], although in reality this can be viewed as a design parameter, representing requirements on the control system design which must be met in order to achieve adequate passive safety. Figure 4.9 shows the time of boiling in the UTOP as the reactivity insertion rate is varied by 20%. The results indicate that reducing the insertion rate makes a minimal impact, with boiling pushed off by less than a minute in the slowest withdrawal case. In addition, Figure 4.9 shows the impact of assuming a constant rate of reactivity insertion versus a constant rate of withdrawal. While many studies make the simplifying assumption that reactivity is added at a constant rate during a UTOP, in reality the velocity of rod withdrawal is constant and the rate of reactivity insertion varies according to the differential rod-worth curve. By considering the true variation of reactivity insertion with time, the time of boiling worsens, but only marginally. No matter, boiling occurs within only a couple of minutes of transient initiation with all examined reactivity insertion rates.




Figure 4.8: Time of boiling initiation in the B&B core with varying control rod drive thermal expansion coefficients.

Rod Withdrawal Reactivity Magnitude

Since boiling was seen in the UTOP scenario no matter the rate of reactivity insertion, it is anticipated that the only method of avoiding boiling would be to reduce the magnitude of the reactivity insertion. This could be enacted by redesigning the control system to spread its worth over a larger amount of rods such that no single rod is holding down too large of worth. This possibility is investigated by systematically reducing the magnitude of the withdrawn rod's worth until boiling is no longer induced. As might be anticipated from Figure 4.6b, the withdrawn rod worth needs to be reduced down to roughly \$0.60 in order to avoid boiling. Even still, at this reduced rod worth the temperature margins are slim, as shown in Figure 4.10, and boiling is only avoided because the boiling point is increased substantially at the elevated pressures seen during the transient. Further, the cladding peak temperatures greatly exceed the fuel-cladding eutectic temperatures for all times beyond roughly one minute, meaning significant cladding wastage is expected.

As was demonstrated in Chapter 2, redesign of the control system to limit the worth of each control assembly (i.e. by adding more control assemblies) would have a large negative




Figure 4.9: Time of boiling initiation in the UTOP of the B&B core with varying reactivity insertion rates.

impact on the neutron balance and lead to higher material damage, and is thus undesirable. In addition, the results of the current sensitivity study indicate that even if the the rod worth were reduced down as far as $\$0.60$, although coolant boiling and fuel melting could *perhaps* be avoided, significant wastage of the fuel cladding would still result. Therefore it is concluded that inherent safety in the UTOP scenario of the B&B core may not be achievable.

Locked Rotor Loss Coefficient

It was seen in Section 4.2 that boiling in the ULOF is caused by the transition from forced-to natural-circulation, when a sudden drop in flow rate occurs. This flow transition occurs when the pump impellers have so little inertia that they suddenly lock in place due to friction. Following rotor-locking, flow through the pump must traverse around the locked rotor, which partially blocks the flow passage and induces a pressure loss. This added pressure loss is the cause of the sudden drop in flow, and the magnitude of the drop is dependent on the pressure loss induced by the pump. Because this induced pressure loss may vary in practice, a final sensitivity study was performed on the loss coefficient of the pump under locked-rotor conditions. Figure 4.11 shows the time of boiling as the loss coefficient of the primary pump is reduced from its nominal value of 2.5 (the value used in the ABR design). As the loss coefficient is reduced, the time of boiling increases significantly until eventually boiling is avoided altogether for coefficients of less than 0.9. Such low loss coefficients allow for the magnitude of the sudden flow decrease to be reduced, which then reduces the strength




Figure 4.10: Core temperatures with time for the UTOP scenario in the B&B core when the rod withdrawal is limited to $\beta_{0.60}$.

of the sudden coolant reactivity insertion and subsequent oscillations, as shown in Figure 4.12. Although this study has shown that boiling can be avoided altogether by reducing the loss coefficient sufficiently, it is not clear that such small loss coefficients can be achieved in reality. If inherent safety in the ULOF scenario is highly desired, such loss coefficients should be viewed as a design criteria for the primary pumps.

Protected Transients

Since the previous section has shown that inherent passive safety may be difficult to achieve in the reference B&B core, this section investigates the requirements that must be imposed on the control system in order to avoid negative consequences during *protected* transients (i.e. where the control system enacts a scram). These transients are abbreviated with a ‘P’ for ‘Protected’, as in PLOHS, PLOF, and PTOF. In all cases, the full temporal evolution of the transient response is omitted for brevity. Instead, the results of each transient are presented in terms of peak temperatures in Figure 4.13 as the delay between transient initiation and control system scram is increased. Time to coolant boiling is not depicted because it is never encountered in the protected transients.

With regards to both peak fuel and cladding temperatures, the PLOHS accident is insensitive to the delay. This is because peak temperatures in the PLOHS accident (and ULOHS accident) occur at the initial steady-state due to strongly negative radial and axial expan-




Figure 4.11: Time of boiling initiation in the ULOF of the B&B core with varying locked rotor loss coefficients.

sion reactivity feedbacks. As the lower gridplate is the first component to be exposed to the increased inlet temperature, it actuates more rapidly than other feedback mechanisms to give a net negative reactivity insertion.

Both the PLOF and PTOF results have strong dependencies on the delay before scram. While peak temperatures in the PTOF increase linearly as the scram delay increases, peak temperatures in the PLOF flatten off due to the exponential decay of core flow. In terms of avoiding fuel melting, the PTOF is the most challenging, as this transient is driven directly by power, which reflects most readily in fuel temperature increases. Regardless, even a 60 second delay between TOP initiation and scram is adequate to avoid fuel melting. In both cases, avoiding fuel-cladding eutectic reactions requires a more rapid response from the control system. In order to avoid such reactions, the PLOF dictates a shorter delay time of 20 seconds, as shown in Figure 4.13b. In all three transients, the reactor can easily be brought under control in response to the postulated transients provided that the reactor control system is functioning.




Figure 4.12: Net and coolant reactivity components in the UTOP of the B&B core with varying locked rotor loss coefficients.




Figure 4.13: Peak temperatures in each protected transient as the delay between initiation and scram is increased.

4.3 Summary of Nominal Transient Performance

In this section, a variety of transients and their assumptions have been investigated in the reference core. It has been found that the fuel-cladding eutectic reactions during a PLOF is the most restrictive out of the protected accidents, requiring the control system to scram within 20 s of transient initiation. On the other hand, unprotected transients pose a much larger challenge in the B&B core, with coolant boiling and elevated fuel/cladding temperatures seen in both the ULOF and UTOP scenarios. While a variety of sensitivities have been investigated, only the pump rotor loss coefficient is seen to qualitatively alter the behavior of the transient progression, potentially allowing for boiling to be avoided in the ULOF if the pumps can be designed properly. All other investigated parameters lead to only marginally improved performance. No design changes were found to enable inherent passive safety in the UTOP. It is likely that the only method of enabling passive safety in the UTOP is simply to reduce the worth of the reactivity insertion upon rod withdrawal, which would only be possible through a complete redesign of the reactor control system to limit the worth of the individual rods. These results highlight the difficulty in achieving complete passive safety in such large-core SFR designs.

It should be noted, however, that inherent safety is not necessarily a requirement for commercialization from the regulator. In fact, the most common reactor systems employed in the world, the PWR and BWR, do *not* display favorable inherent safety characteristics and instead rely upon many highly complex auxiliary systems in order to mitigate the consequences of the accidents proposed in this study. Therefore it is emphasized that the failure of the B&B system to achieve inherent passive safety does not exclude it as a feasible system. This result simply means that a B&B core's safety performance may not be as favorable as the excellent performance exhibited by smaller SFRs.

Chapter 5

ARC System Incorporation

The previous chapter examined the transient performance in the reference core and found it challenging to achieve inherent safety in the case of either ULOF or UTOP accidents. The purpose of this chapter is to explore a new method of improving transient performance by employing Autonomous Reactivity Control (ARC) systems. After introducing the basic premise of the ARC system, the transient models from Chapter 4 are modified to include the ARC system response and transient improvements are assessed. Finally, a number of other aspects concerning ARC system inclusion are discussed and quantified.

Much of the work in this chapter is derived from a broader project recently completed at UC Berkeley. In the broader project, many additional aspects of ARC system design and implementation are investigated by both myself and others. For more information on the background and physical implementation of ARC systems, the reader is referred to the summary report of that project [46].

5.1 ARC System Overview

The long history of SFR research throughout the world has led to a mature understanding of factors which drive their transient performance. Out of all of the feedback mechanisms in a well-designed SFR core, only the coolant voiding effect has the potential to be a significant positive feedback with respect to increasing temperature, potentially leading to an unstable system¹. Although it has been demonstrated many times over that a positive void coefficient can be tolerated as long as the other feedback mechanisms sufficiently negative to dominate the overall reactor response, this potential for instabilities has generated a huge amount of discussion in the SFR community, ultimately prompting designers to devise a number of strategies to improve transient performance. In general, there are three types of methods that have been used: (1) leakage-based, (2) moderation-based, and (3) absorption-based. An overview and thorough discussion of each of these methods can be found in [117], and the

¹Other mechanisms such as cladding expansion can have a positive effect as well, but to a much more limited extent.

reader is referred there for further background. Also in the same paper is a convincing argument that none of these classical methods are applicable to B&B cores due to their unique constraints, such as the need for an extremely tight neutron economy and the requirement that assemblies are shuffled uniformly through the core.

Further, because the very feasibility of breed-and-burn operation depends upon having the highest possible breeding, its design is driven almost entirely to advantage neutron economy with little attention given to other aspects, resulting in extremely low leakage and a higher void coefficient than almost any other SFR design. Thus the core to which none of the mitigation strategies outlined above are applicable is in fact the core that requires them the most. For these reasons, the ARC system was invented by Qvist during the course of his PhD specifically to accommodate the needs of low-leakage SFR designs such as the B&B [114].

The ARC system is based loosely on the design of the Lithium Expansion Module (LEM) [74] with a number of key differences to increase its viability. In a LEM, a liquid column of lithium absorber is passively inserted and retracted from the core in response to changes in core temperature. Such a system allows for an easily reversible response to any initiating event without significantly impacting the neutron economy or spectrum during nominal operations. However, the system as originally proposed suffers from practical complications that would make it difficult to implement, such as the need to have a column of liquid lithium stratified on top of a gas bubble – a hydrostatically unstable configuration. The ARC system uses the same concept but rearranged into a workable geometry in order to successfully implement the mission of LEMs.

The ARC system in its standard configuration is installed as a modification to a conventional SFR fuel assembly design, as shown in Figure 5.1. The system consists of two reservoirs located at the top and bottom of the assembly (shown in more detail in Figure 5.2), and two concentric tubes that link the reservoirs. The inner tube is open at both ends and connects the insides of both reservoirs, while the outer tube is open at the bottom (connected to the lower reservoir) and at the top connects to a closed gas-filled reservoir. During operation, the upper reservoir is completely filled with a liquid (henceforth the “expansion” liquid), while the lower reservoir contains the same expansion liquid and, floating on top of it, a separate immiscible liquid (henceforth the “absorber” liquid). The remaining free volume between the two concentric tubes in the closed system is filled with an inert gas. The outer ARC-tube has the same outer dimension as a fuel rod. Installing an ARC-tube therefore implies replacing one of the fuel rods in the assembly.

During an accident/transient scenario in the reactor, the ARC system is intended to respond in the following way:

1. Some event raises the temperature in the core, which heats up the coolant
2. The heated coolant flows to the top of the assembly and transfers heat to the expansion liquid in the upper reservoir




Figure 5.1: An ARC-equipped fuel assembly (not to scale). The components relating to the ARC system (upper and lower reservoir and tubes) are in approximate scale to one another, while the midsections of the assembly (that contains the fuel bundle) have been reduced in axial length by about a factor of 10 for greater clarity. Courtesy of Qvist [46].

3. The expansion liquid in the upper reservoir thermally expands. Since the reservoir is completely filled and sealed at the top, this expansion is directed down the inner ARC-tube that connects the two reservoirs
4. As expansion liquid enters the lower reservoir from the upper reservoir (through the inner ARC-tube), the level of absorber liquid rises into the active core, while compressing the inert gas above
5. The absorber liquid, which has a high neutron capture cross-section, introduces negative reactivity by absorbing neutrons in the core, which in turn causes a reduction in power and temperature
6. As the core cools down, the temperature of the expansion liquid starts to fall. Thermal




Figure 5.2: CAD model of an ARC-equipped fuel assembly, courtesy of Qvist [46].

contraction combined with the pressure of the inert gas again lowers the axial level of the absorber liquid until the system reaches a stable critical configuration

A schematic view of how the liquid levels change in the lower ARC reservoir as the temperature increases is shown in Figure 5.3. Properly designed, the ARC-system is meant to act as a thermostat in the core, autonomously controlling temperature without the need for any operator action, electrical systems, or moving mechanical parts. The actuation responds to temperature and relies solely on the laws of physics, and is therefore an inherent feedback mechanism.

The selection of ARC system working fluids is paramount to achieving a workable design and a number of requirements need to be considered:

1. The expander fluid needs to have a sufficiently large thermal expansion coefficient to actuate the system in response to a realistic change in temperature
2. The absorber fluid needs to have a sufficiently large neutron absorption cross-section in the fast spectrum to provide a strong enough reactivity insertion
3. The expander fluid and inert gas need to be relatively neutronically inert to avoid parasitic absorptions
4. To maintain proper stratification, the expander fluid should have the highest density at operating temperatures, followed by the absorber fluid and then the inert gas
5. The fluids should not be overly corrosive at the expected temperatures

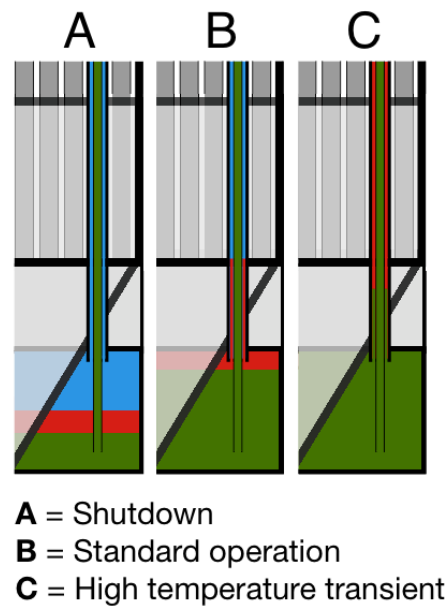


Figure 5.3: Schematic view of the lower ARC reservoir at different states courtesy of Qvist [46]. Green, red, and blue are the expander, absorber, and inert fluids, respectively.

6. The three fluids should not be soluble within each other to a large extent
7. Reaction products from any of the fluids should behave acceptably in the environment from chemical and radiological considerations
8. The three fluids should be chemically compatible with other core components (namely sodium)
9. The three fluids should remain in their proper state of matter over the expected temperature range

After an exhaustive screening process outlined in [46], the fluids were selected as shown in Table 5.1. Although experimental data on the K-Li system is sparse at operational temperatures, these fluids were found to meet all of the requirements sufficiently to move forward with further assessment of ARC system performance. The practical issues which may arise from K-Li mixing, including loss of lithium, tritium production, and a neutronic penalty, are addressed in Section 5.5.

Other Engineered Passive Safety Systems

It should be mentioned that other systems meant to accomplish similar goals have been proposed, collectively referred to as ‘engineered passive safety systems.’ One of these has

Table 5.1: Selected materials for the three ARC system working fluids.

Fluid	Material
Expander	Potassium
Absorber	Lithium
Gas	Helium or Argon

already been mentioned – the LEM. Others include the Gas Expansion Module (GEM) [147], Curie point actuated scram rods [101], rupture-seal actuated movable absorbers [145], floating absorbers which respond to changes in fluid density [51, 10], and control rods with enhanced driveline expansion mechanisms [51, 34]. An extensive overview of the majority of these systems, including their benefits, drawbacks, and applicability to B&B cores, is provided in [113]. It is finally mentioned that at least a handful of these systems have been successfully experimentally tested, including GEMs in FFTF [127], Curie point systems in JOYO [138], and floating absorbers in BR-10 [15]. It is envisioned that the ARC system may provide improvements over all of the aforementioned systems by actuating in response to *any* transient initiator (as compared to GEMs which respond only to a LOF), being compatible with B&B operations, and allowing for simple un-actuation after the transient has been arrested.

5.2 ARC System Performance

Beyond the basic conceptual design outlined previously, the ARC system is a complicated component involving many design variables. Luckily all of these unique variables are reflected in terms of ARC system performance as only four important characteristics. Table 5.2 outlines the relationship between these design variables and ARC performance characteristics.

In order to assess the potential of the ARC system, it is necessary to perform a sweep over each of the four characteristics in Table 5.2 to find the combination of characteristics which allows for the greatest benefits to be gained. This is the main task of this section. First though, the methods used to computationally model the ARC system actuation are outlined. Then the parametric study is performed, the results of which are used to assess the potential for the ARC system to improve B&B transient performance. It is noted that, while this chapter concerns only the reference B&B core, another core, the oxide variant of the ABR, has been examined as well as part of the overall ARC system project to understand how ARC systems behave in cores with different properties. This second core differs from the B&B in many ways including the use of oxide fuel as compared to metallic fuel, being an actinide burner as compared to a breeder, and being much smaller in both physical dimension and power level [80]. Details of the ARC system performance in this other core will be published

Table 5.2: Relationship between ARC system design variables and performance characteristics.

Performance Characteristic	Variable	Design Variables
Maximum System Worth	w	<ul style="list-style-type: none"> - Flow area in the inner ARC tubes - Number of fuel pins replaced with ARC tubes - Enrichment of absorber fluid in ${}^6\text{Li}$
Coolant temperature span required to fully actuate the system	S	<ul style="list-style-type: none"> - Volume of upper reservoir - Flow area in inner ARC tubes within the active core - Flow area in outer ARC tubes - Number of fuel pins replaced with ARC tubes - Volume of lower reservoir - Height of active core
Coolant temperature increase required to bring the absorber fluid to the bottom of the active core	ΔT_{act}	<ul style="list-style-type: none"> - Volume of upper reservoir - Flow area in inner ARC tubes below the active core - Flow area in outer ARC tubes - Number of fuel pins replaced with ARC tubes - Volume of lower reservoir - Distance between lower reservoir and active core
Rate of temperature increase of the upper reservoir in response to a change in the coolant outlet temperature	τ	<ul style="list-style-type: none"> - Geometric design of the upper reservoir - Material constituents of the upper reservoir

in a companion paper.

Computational Modeling Techniques

Whole-core transient behavior is modeled using SAS in the same way as in Chapter 4, although a number of changes to the model were required to reflect the addition of the ARC system and its unique transient response. Modifications were made to reflect three aspects: (1) the taller core and associated pressure drop, (2) the additional thermal inertia associated with the addition of the ARC reservoirs, and (3) the reactivity response of the ARC system as the coolant outlet temperature changes. Aspects (1) and (2) are implemented rather simply by just adding additional axial nodes to the model with material properties, flow areas, and geometries corresponding to the ARC system components. The more major modifications are related to aspect (3) and discussed in detail in the following paragraphs. Although the implementation of the ARC system entails the removal of some amount of fuel and replacement with ARC tubes, this modification is neglected in terms of its impacts on transient performance and the material worth distributions. This is justified based on the small amount of fuel that is removed (1 out of 169 pins in each assembly) and the negligible level of neutron interactions within the expander fluid.

Due to lack of native capabilities within SAS to predict the fluid dynamics, heat transfer, and neutronics associated with the ARC system response, the decision was made to use simplified models within SAS along with a number of high-fidelity external calculations in

order to model the transient behavior. This approach, outlined diagrammatically in Figure 5.4, has a couple of advantages, including avoiding the need to make extensive source code modifications to SAS and the ability to assess a wide range of ARC systems at reasonable computational expense. As indicated, the control system module within SAS is used to combine the disparate pieces into the overall reactivity insertion with time.

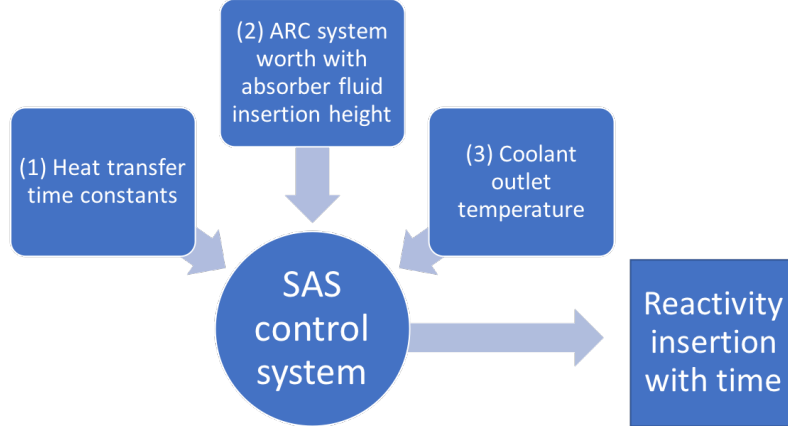


Figure 5.4: Overview of how the ARC system response is modeled in SAS.

In the context of the current ARC system modeling, the control system takes three inputs and manipulates them in order to determine the ARC system reactivity insertion as a function of time. The first input to the control system is an integral parameter representing the heat transfer time constant from the coolant flow to the ARC upper reservoir. This time constant, denoted as τ and referenced previously in Table 5.2, has a huge amount of information embedded within in, including the upper reservoir geometry, coolant flow regime, and material properties, and is used in conjunction with Equation 5.1 to predict the ARC upper reservoir temperature in response to a change in the driving function (the coolant temperature). This equation corresponds to a so-called ‘lag-compensator’, where T_{UR} is the average temperature of the expander fluid in the upper reservoir and T_c is the coolant outlet temperature.

$$T_{UR}(t) + \tau \frac{dT_{UR}}{dt} = T_c(t) \quad (5.1)$$

This differential equation has a simple solution as provided in Equation 5.2.

$$T_{UR}(t) = T_c(t) + (T_{UR}(0) - T_c(t))e^{-\frac{t}{\tau}} \quad (5.2)$$

To determine τ it is necessary to either perform experiments or detailed Computational Fluid Dynamics (CFD) calculations. This study utilizes CFD calculations performed as part of the broader ARC system project by Gradecka [46]. The cited reference quantifies τ by simulating a step-change in the coolant outlet temperature and the upper reservoir

temperature response. Using Equation 5.2 along with the simulated response data, a value of τ can be fitted which represents the time-lag between a change in the coolant outlet temperature and the responding upper reservoir temperature. This process is repeated as the geometry of the upper reservoir is varied in order to optimize the heat transfer efficiency. Through this extensive study, a value of τ as low as 1.3 s was identified. In the context of the parametric studies performed over the ARC system design variables, 1.3 s is taken as the lower bound on τ (representing the absolute fastest heat transfer possible in the system), while it was assumed that τ is unbounded from the top².

In the case of a LOF, the decreasing flow rate with time acutely impacts the heat transfer efficiency between coolant and upper reservoir. An example of this τ variation with flow is presented in Figure 5.5, showing that resistance to heat transfer may increase nearly 4x for the natural circulation flowrates which might be expected during a LOF. This phenomenon is important because it reduces the speed with which the ARC system is actuated. In order to account for this, a table representing τ at different flowrates as determined via CFD calculations is constructed and passed to SAS.




Figure 5.5: The impact of flow rate on the value of τ as predicted through CFD calculations. Courtesy of Gradecka [46].

To account for this, the lag-compensator model in Equation 5.1 needs to be modified into that of a ‘variable lag-compensator’, as described by Equation 5.3.

$$T_{UR}(t) + \tau(t) \frac{dT_{UR}}{dt} = T_c(t) \quad (5.3)$$

The variable lag-compensator has a more complicated analytical solution as in Equation 5.4, which is newly implemented into SAS with the assumption that τ remains constant over each time step.

²A design with a very large τ could be achieved employing insulating materials and sub-optimal geometries, if it were desirable.

$$T_{UR}(t) = T_{UR}(0)e^{-\int_0^t \frac{dt'}{\tau(t')}} + e^{-\int_{t_0}^t \frac{dt'}{\tau(t')}} \int_0^t \frac{T_c(t')}{\tau(t')} e^{\int_{t_0}^{t'} \frac{dt''}{\tau(t'')}} dt' \quad (5.4)$$

Once a value of τ is determined, it is still necessary to know $T_c(t)$ in order for SAS to predict the expander fluid temperature. This value is pulled into the control system from the core thermal-hydraulic models at the core outlet. In reality, each channel has a unique outlet temperature which will cause the ARC systems in each separate assembly to be actuated at slightly different times from each other. This aspect is neglected, and instead the coolant outlet temperature is taken from only Channel 2 (see Figure 4.2) since this channel contains the vast majority of core power. Sensitivity studies to the channel which drives the response have shown virtually no impact chiefly for three reasons: (1) the ARC system operates based on a *change* in temperature, (2) the outlet temperature profile over the core is relatively flat due to orificing, and (3) the kinetics modeling uses a point kinetics approach and thus local flux perturbations are not accounted for.

Using these two pieces of information, SAS is able to predict the upper reservoir temperature with time, but still needs to translate this to a reactivity insertion. This is done using a mass balance to determine the height of the absorber fluid with time and a pre-calculated table of the ARC reactivity worth with height. Because the upper reservoir is filled solid with nearly-incompressible expander fluid, any temperature increase is directly translated into an increase in volume and a corresponding increase in the insertion height nearly instantaneously (in fact at the speed of sound).³ The functionality translating ARC system height to reactivity insertion magnitude is determined using static criticality calculations where the absorber height is increased in steps to determine the impact on core criticality. The shape of this curve is exactly the same as the control rod worth curve in Figure 2.17, except the ARC system is inserted from the bottom whereas the control rods are inserted from the top. The overall calculational routine as described above is provided in detailed as a flow chart in Figure 5.6.

Parametric Studies

The same three transients as were examined in Chapter 4 and summarized in Table 4.2 are re-examined, this time with the ARC system included into the core. A parametric study over the ARC system design space was performed for each transient, as detailed in the following subsections.

ULOHS

The feasible design space for improving the ULOHS was found to be very large. Figure 5.7 summarizes the results of the parametric study in the ULOHS in terms of margin to

³It is noted that this increase in liquid height is virtually independent of the inert gas pressure because the compressibility of the gas is so much larger than that of either the expander or absorber fluids.




Figure 5.6: Flow chart of the computational scheme used to determine the ARC system reactivity insertion with time in SAS. Items highlighted in yellow are user inputs.

coolant boiling gained over the case without ARC system when $\tau = 1.3$ s. It is noted that no parametric studies on τ are here performed because τ is already taken to be its most aggressive value, and increasing τ would only worsen performance. The response surface is found to be very smooth over the design space, being in general monotonic and lacking any sort of bifurcation behavior. This enables flexibility in designing the ARC system to accommodate the needs of other transients and makes the design process straightforward.

ARC designs which most improve the ULOHS performance generally have small S , large w , and small ΔT_{act} – in other words the most aggressive designs. This is expected because a faster actuation causes core power to decrease more rapidly, thus enabling lower coolant temperatures when the system reaches equilibrium because the integrated power input is reduced. Because the ULOHS generally has peak fuel temperatures at the beginning of transient nominal conditions, reducing the peak fuel temperature through the use of the ARC system is not possible. While it appears that greater benefits might be achievable by further increasing w , decreasing S , or decreasing ΔT_{act} , the parametric study has been restricted to $w \leq \$2$, $S \geq 50^\circ\text{C}$, and $\Delta T_{act} \geq 10^\circ\text{C}$ based on engineering judgement as to the limitations of the physical design. Within these bounds, roughly 65°C of margin to coolant boiling can be gained as compared to the nominal ULOHS.




Figure 5.7: Margin to coolant boiling gained in the ULOHS scenario of the B&B core with ARC system inclusion as a function of ARC system design parameters.

The detailed temporal evolution for the ULOHS scenario is depicted in Figure 5.8. As is the case without ARC system inclusion, at no point does reactor power exceed the steady-state level. Immediately as the coolant inlet temperature increases, core reactivity turns negative for the same reasons as before. However whereas in the nominal ULOHS scenario core reactivity hovered very close to zero starting at around 1000 s, the actuation of the ARC system keeps reactivity thoroughly negative throughout the transient and allows for core power to be reduced more strongly. The ARC reactivity insertion is very smooth, although somewhat slow due to the time that it takes for the outlet temperature to increase enough to affect sufficient thermal expansion of the expander fluid. In fact, peak coolant temperatures do not begin to increase until roughly 1000 s into the transient. For the B&B ULOHS it would be very difficult to force a faster ARC system actuation due to the very high thermal inertia of such a large core and the slow nature of the ULOHS sequence.

UTOP

Contrasting with the ULOHS results, the parametric study for the UTOP at first revealed no ARC designs which could avoid boiling. This is due to a series of oscillations induced by the ARC system response in the very early stage of the transient, as depicted in Figure 5.9 for the extreme case with $w = \$2.0$, $S = 100^\circ C$, $\Delta T_{act} = 10^\circ C$. The oscillations appear to be caused by the strong, rapid insertion of the ARC system as the control rod is withdrawn. As the ARC

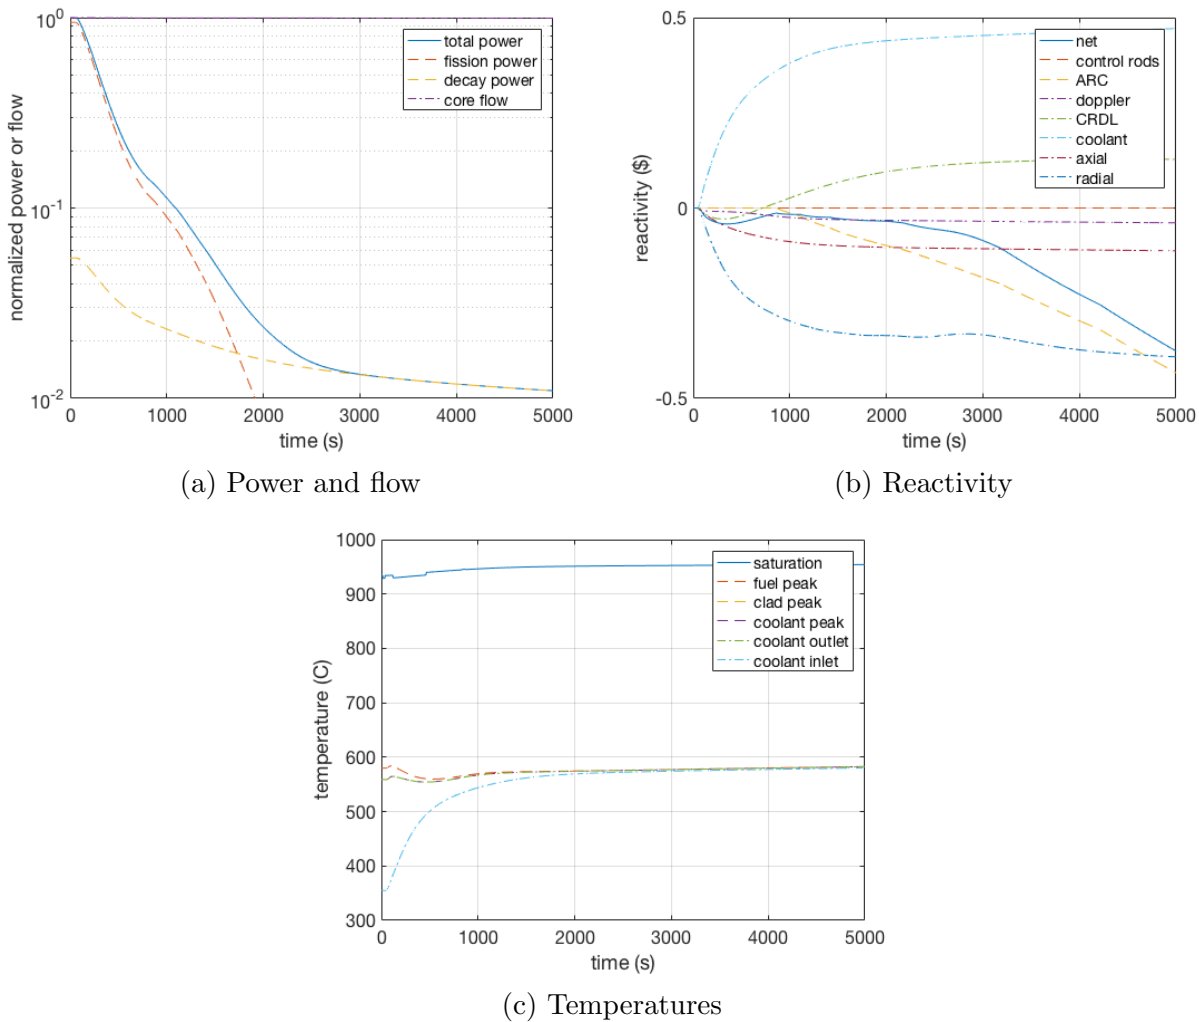
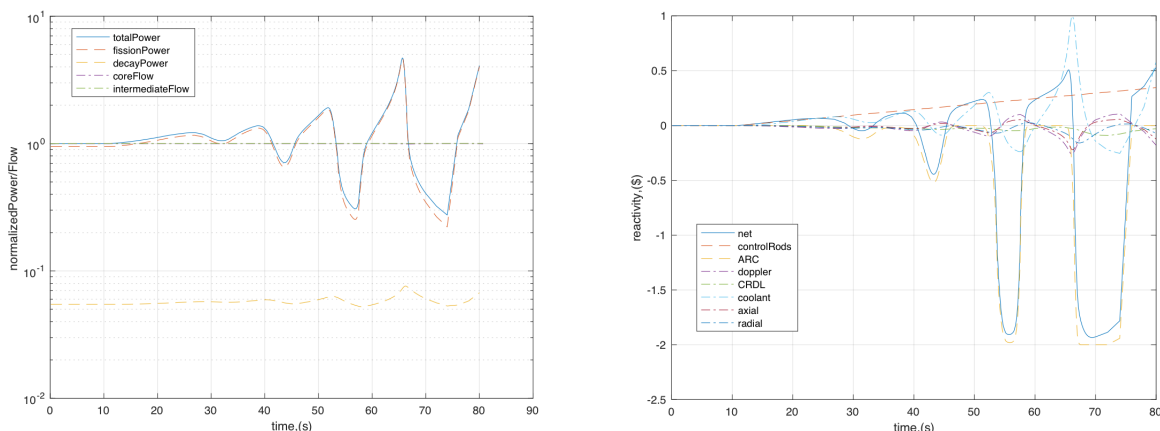


Figure 5.8: Evolution of the ULOHS with ARC system inclusion.


system provides a strong negative reactivity insertion, net reactivity is temporarily brought negative and core temperatures begin to decrease, which causes the feedback mechanisms to relax. Meanwhile the control rod is still in the process of withdrawing, and so as the feedback mechanisms relax the reactivity is again driven positive by the rod withdrawal, which leads to core temperature increases and another strong ARC system actuation. This cycle repeats, each time getting stronger as the various feedback mechanisms act slightly out of phase with each other until eventually core power increases so strongly that boiling is initiated with the core at roughly 400% power. In addition, both the fuel melting and fuel-cladding eutectic temperatures are exceeded. It is further noted that boiling is seen to occur roughly one minute *earlier* in the UTOP scenario when the ARC system is included, as compared to the

nominal core without ARC system. These results highlight that the ARC system has the potential to actually degrade transient performance if not properly designed.



(a) Power and flow

(b) Reactivity



(c) Temperatures

Figure 5.9: Evolution of the UTOP with ARC system inclusion.

This unfavorable behavior is made possible by the high thermal conductivity of the metallic fuel in the B&B core, which leads to tight coupling between fuel and ARC system temperatures. Because the ARC system is delayed in its response as compared to the coolant feedback, this tight coupling can lead to diverging oscillatory behavior, which ultimately induces coolant boiling. Similar oscillatory behavior was not seen in the oxide ABR [46] likely due to oxide fuel's lower thermal conductivity, which impedes the communication of high fuel temperatures to the coolant, thereby reducing the coupling between fuel temperature

and ARC system response.⁴ In order to break the coupling in the B&B core, a different method is investigated – increasing τ such that the heat transfer between coolant and upper reservoir is degraded. This method was found to smooth oscillatory behavior in the oxide ABR, although those oscillations were of different origin and much weaker in magnitude.

Figure 5.10 shows how the ARC system response in the UTOP scenario is impacted by τ . As τ is increased, the oscillation magnitude decreases until they are removed almost entirely with $\tau = 300$ s. At $\tau = 250$ s the transient progresses without boiling, indicating that this value of τ decouples the system enough for acceptable transient performance. Although no studies have been performed to determine how to make τ this large, it is assumed achievable by using a material with low thermal conductivity and/or altering the shape of the upper reservoir.




Figure 5.10: ARC system reactivity insertion in the UTOP as τ is increased.

Using this larger τ , passive safety in the UTOP scenario is realized, as depicted for the example case of $w = \$2.0$, $S = 100^\circ\text{C}$, and $\Delta T_{act} = 10^\circ\text{C}$ in Figure 5.11. Some oscillatory behavior still occurs in the first 500 s of the transient, however the oscillations are damped as the rod withdrawal completes and eventually a new steady-state is reached with core temperatures less than 100°C above steady-state values. Attempting to completely remove the oscillations by further increasing τ is not effective from the standpoint of reducing peak

⁴In addition, the oxide ABR did not require as aggressive of ARC system designs to achieve significant safety benefits due to the more favorable baseline transient performance.

temperatures. Because peak temperatures occur early in the rod withdrawal sequence (for the cases where diverging oscillations have been removed), making τ larger than it needs to be causes slower ARC system actuation and higher peak temperatures during the initial transient phase. Even with the remaining oscillations, peak coolant temperatures are kept significantly far from boiling. Therefore the ARC system should be designed with heat transfer characteristics only weak enough so that diverging oscillations may be avoided, and no weaker.


Figure 5.12 shows the parametric study for $\tau = 300$ s, the smallest value (with some margin) which enabled the avoidance of boiling. The margin to boiling is maximized by ARC designs with high w and low S , where roughly 160°C of margin to coolant boiling can be achieved. This should be compared to the nominal core without ARC system inclusion, for which boiling occurred within a few minutes of transient initiation unless the highest worth control assembly could be limited to below ~ 0.60 . Therefore the inclusion of the ARC system provides a substantial improvement in the UTOP performance. Boiling is not able to be avoided for $w \leq \sim 1.0$, implying that an aggressive ARC design is necessary to obtain passive safety in the UTOP. This conclusion is in line with those for the ULOHS in that the most aggressive ARC system provides the greatest benefits, with the caveat for the UTOP scenario that an overly aggressive design can lead to strong diverging oscillations.

Since τ needs to be increased to achieve acceptable UTOP behavior, a study was performed to determine the impact of higher τ on the ULOHS performance. Figure 5.13 shows that as τ is increased up to 300 s, the impact on peak temperatures in the ULOHS is negligible as long as ΔT_{act} is sufficiently small. Therefore the larger value required by the UTOP transient should also be acceptable from the perspective of ULOHS passive safety.

ULOF


Similar to the UTOP, no ARC system designs with $\tau = 1.3$ s were able to prevent boiling in the ULOF. The cause of boiling in this case is similar in concept to that which initially led to boiling in the UTOP. In both scenarios, a driving force leads to continually strengthening ARC system actuations, followed by rapid decreases in core temperatures which lead to relaxation of the feedback mechanisms and corresponding positive reactivity insertions. Whereas in the UTOP the driving force was the rod withdrawal, for the ULOF the driving force is the reduction in core flow. In both scenarios, the driving force is independent of the ARC system, meaning that its progression is not impacted by the actuation of the ARC system. These two scenarios are contrasted with the ULOHS, where the driving force (the coolant temperature increase) is much more strongly impacted by the ARC system actuation by the reduction in coolant temperatures with more aggressive designs.

With the ARC system included, the ULOF oscillatory behavior is in fact worsened as compared to the case without the ARC system, as shown in Figure 5.14 for the case with $\tau = 1.3$ s, $w = 2.0$, $\Delta T_{act} = 10^\circ\text{C}$, and $S = 50^\circ\text{C}$. Unfortunately, no relaxation of either τ nor S was able to eliminate the oscillations enough to avoid boiling. If the ARC system feedback is in fact retarded enough to avoid oscillatory behavior in the early phases of



(a) Power and flow

(b) Reactivity



(c) Temperatures

Figure 5.11: Evolution of the UTOP with ARC system inclusion and $\tau = 300$ s.

actuation, boiling is eventually triggered at the transition from forced- to natural-circulation in the same way as was seen in the ULOF without the ARC system. Although increasing the pump coast down time showed ULOF gains over the nominal case, this method does not help to avoid boiling because it only delays the time of the flow transition, which ultimately leads to boiling.

In the case with the ARC system included, boiling at the flow transition can still be avoided when the pump rotor loss coefficient is sufficiently small, as shown in Figure 5.15, but because τ needs to be increased in order to avoid inducing oscillations during the early coast down phase, it is found that the ARC system essentially does not actuate during the flow transition at all. This implies that the inclusion of the ARC system does not *benefit*




Figure 5.12: Margin to coolant boiling gained in the UTOP scenario of the B&B core with ARC system inclusion as a function of ARC system design parameters.

the ULOF, it simply does not further impede it, provided that the pump is designed with a sufficiently small loss coefficient in the first place. For certain, it is found that the presence a non-smooth flow transition challenges ARC system performance when the coolant density feedback is so strongly positive.

Summary of ARC System Performance

Although protected transients showed acceptable behavior, unprotected transient performance without ARC system inclusion was poor, with both the UTOP and ULOF leading to boiling. By including ARC systems, significant gains were achieved in both the ULOHS and UTOP scenarios. With aggressive ARC system designs, the margin to coolant boiling was increased by roughly 65°C in the ULOHS. Additionally, by increasing τ to 300 s, oscillations that were present with lower values of τ could be damped and boiling was able to be avoided in the UTOP as well, with more than 150°C margin to boiling now achievable. However, no ARC system design could be found to mitigate the oscillations in the ULOF at the transition from forced- to natural-circulation, and therefore complete passive safety was not enabled if the pump design is such that a non-smooth flow transition takes place. In light of this, the optimal ARC system design is identified as the one which most improves the ULOHS and UTOP performance: $w = \$1.75$, $S = 50^\circ\text{C}$, $\Delta T_{act} = 10^\circ\text{C}$, and $\tau = 300$ s. Subsequent




Figure 5.13: Peak coolant temperatures in the ULOHS as a function of ARC design parameters with increasing τ .

sections will investigate further methods for improving the performance of the ARC system so that all three unprotected scenarios can achieve inherent safety.




Figure 5.14: Reactivity feedback evolution for the ULOF with ARC system inclusion.




Figure 5.15: Coolant and ARC reactivity feedbacks in the ULOF with $\tau = 300$ s for various locked pump rotor loss coefficients.

5.3 Introduction of Modified ARC System

While the previous section identified some major benefits that could be achieved by incorporating ARC systems into the B&B core, it was not able to enable fully-passive inherent safety performance. Particularly in the ULOF scenario, the tendency for the ARC system to form oscillations could not be avoided due to the rapid transitions during both the early coast-down phase and the later flow-transition phase. While these results indicate that the proposed ARC system may not be suitable for the B&B core *as is*, this section will introduce a modification to the ARC system design aimed at reducing oscillations. Such a new design, if successful, would enable passive safety in all three transients as well as unlock improved performance in the UTOP by enabling more aggressive ARC system designs. The focus of this section is to overview the modified ARC system design and assess its performance.

It has been noted that the ARC-induced oscillations involve a complicated cycle where strongly negative reactivity insertions result in rapidly reduced core temperatures and subsequent un-actuation. It is thus deduced that one of the weak points of the ARC system is its ability to un-actuate both quickly and out of phase with other feedback mechanisms at undesirable points in time. In principle, the only difference between a scram and the ARC system actuation is the permanence of a scram. Because the actuation state of control rods after a scram is not influenced by the transient progression, as long as the scram is strong enough it is able to overpower any oscillatory behavior in the core and force a shutdown condition. If the ARC system could be made more similar in nature to a scram, the potential for diverging oscillations could be avoided.

This is not to say, however, that we wish to turn the ARC system actuation into a scram – in fact one of the largest benefits of the ARC system as compared to other passive safety features is its ability to un-actuate after a transient has been arrested. This ability enables easy recovery which is unmatched by more scram-adjacent mechanisms like Curie point activated control rods or rupture seal devices. Additionally, the ability to un-actuate is important for other missions of the ARC system, such as the passive load-following abilities investigated by Svanstrom [137]. Therefore the nature of the ARC system that allows subsequent un-actuation after a transient is something that should be preserved.

To compromise between the standard ARC system and a system with more of a scram-like response, a pseudo-one-way valve is proposed to be incorporated into the ARC system. Such a valve would be incorporated into one of the ARC tubes and would impose a different pressure loss in each direction of flow. Conceptually, the purpose would be to allow a very low resistance to forward actuation while imposing a significant resistance to un-actuation so that the ARC system may act quickly to arrest a transient without being able to rapidly oscillate back and forth even as core temperatures may be rapidly decreasing. The next few subsections outline the conceptual design of an ARC system modified to incorporate such a pseudo-one-way valve.

Valve Design Options

In order to accomplish the goals of the modified ARC system, two variations on the pseudo-one-way valve have been identified. The first solution is perhaps the more elegant of the two, a so-called Tesla valve, as diagrammed in Figure 5.16. The Tesla valve was invented by Nicola Tesla in the early 1900s [141] and is not so much a ‘valve’ in the traditional sense of the word – it can neither fully block flow nor be adjusted by an operator to meter flow – but is more appropriately described as a leaky check valve. For these reasons, the Tesla valve has seen limited industrial use. The basic idea behind the Tesla valve is to arrange a series of flow passages such that as fluid flows in one direction the channel geometry has minimal impact, while in the other direction the flow is forced back upon itself to induce vortices and pressure losses strong enough to impede throughput. In its original conception, the Tesla valve was intended for applications in which reverse flow would be undesirable but complete elimination of reverse flow would lead to unacceptable pressure buildup.

In relation to the current study, the Tesla valve possesses a number of characteristics that make it uniquely suited to the goals of the modified ARC system. Most importantly, the Tesla valve would allow for resistance to un-actuation of the ARC system, helping to reduce the potential for oscillations. Second, the Tesla valve would have minimal impacts on the speed of the forward actuation if properly designed, and therefore would not be detrimental to the goals of the ARC system. Finally, the Tesla valve incorporates no moving parts whatsoever and is a completely passive system with high resistance to failure.




Figure 5.16: Conceptual diagram of the Tesla valve taken from the original patent [141].

The second identified option is an orificed check valve as diagrammed in Figure 5.17. Such a valve would operated in the following manner:

- During forward actuation, the poppet is unseated and pressed against a mesh grate while flow traverses the relatively large channel with only a very small pressure loss
- During reverse actuation, the poppet seats and flow is able to traverse only through a small calibrated orifice which induces a large pressure loss and restricts the flowrate




Figure 5.17: Conceptual diagram of the orificed check valve.

In comparison to the Tesla valve, the orificed check valve has both advantages and disadvantages. The primary disadvantage is that the orificed check valve incorporates a moving part – the floating poppet. While this component presents extremely limited opportunity for failure, it is somewhat more complicated than the static Tesla valve and would require thorough testing in a representative environment. In addition, some sort of sealant will likely be required where the poppet seats, and this material must be chosen to be compatible with liquid metal flow. A few design considerations could promote resilience of the orificed check valve to failure, such as making multiple orifice holes to ensure operability in the unlikely chance of an orifice blockage and sizing the poppet to be small enough so that it cannot block the inner ARC tube in the case of failure of the perforated screen.

The primary advantages of the orificed check valve are (1) the ease of manufacturing and (2) the ease of designing the valve to achieve the desired loss during un-actuation. In relation to manufacturability, the orificed check valve would be straightforward, with similar designs already in production from multiple vendors. While it is envisioned that the manufacture of Tesla valves is not beyond capability, it is definitely more complicated and there are essentially no current vendors. The second benefit concerning the ability to design the valve is more crucial to the feasibility of the modified ARC system concept. While we speak of a ‘rapid’ or ‘strong’ ARC system actuation, in reality the flow velocities seen in the system will be on the order of 1 m/s or less. In order to achieve sufficient pressure losses to impede the reverse actuation at these low velocities, the loss coefficient will have to be extraordinarily large and it is not at all clear if Tesla valves will be able to meet the requirement. In addition, performing the analysis and design of a Tesla valve to meet the requirements would involve computationally expensive CFD analyses which are outside of the scope of this dissertation. The orificed check valve, on the other hand, is straightforward to design to the requirements of the modified ARC system using established correlations for flow through constricted channels, at least to first order.

For any concentrated pressure loss, Equation 5.5 can be used, where K is the loss coefficient (determined from correlations), ρ is the fluid density, and v is the fluid velocity. To understand the maximum flow loss that might be possible using the orificed check valve option, the Handbook of Hydraulic Resistance is used to find K [71].

$$\Delta p = K \frac{\rho v^2}{2} \quad (5.5)$$

A lower bound on the diameter of the orifice (and thus an upper bound on the loss coefficient) can be estimated as the most restrictive of the following three scenarios:

1. The diameter of the hole approaches the size of the fluid molecules
2. The diameter of the orifice approaches the limits of manufacturing and associated tolerances
3. Capillary forces at the interface between two fluids induce a pressure much larger than the flow driving pressure

Approaching the extremely small diameters of scenario (1), much smaller than the nanometer scale, is considered unrealistic because a hole of such a size would not be manufacturable, and thus scenario (2) would supersede. While the manufacturing limitations on the orifice diameter are not extremely well-defined, it is expected that orifices of *at least* 0.1 mm could be easily manufactured based upon the widespread availability of drill bits of this size. The limitations associated with scenario (3) can be analyzed using Equation 5.6, where γ is the surface tension between the expander and absorber fluids, θ is their contact angle, r_c is the interface radius, and p_c is the resulting pressure differential generated by capillary action.

$$p_c = \frac{2\gamma \cos(\theta)}{r_c} \quad (5.6)$$

Experimental studies by Alchagirov et al. indicate that at temperatures relevant to ARC system operation, the wetting angle between liquid lithium and liquid potassium is very near 0° , indicating extremely good wetting between the two materials as a result of their weak interaction [9]. This small wetting angle implies $\theta \approx 90^\circ$ in the context of the the two fluids inside of a capillary tube (i.e. the valve orifice) and capillary pressures near zero, which would mean that capillary forces and their associated resistance to flow should not have an impact on the lower limit for the orifice diameter. Regardless, data on the lithium-potassium-steel system are sparse, and other means for ensuring capillary forces do not impede flow through the orifice should be devised. One such method would be to ensure that the orificed check valve never sees multi-phase flow, i.e. the flow passage through the valve only ever sees liquid potassium such that there is never a fluid interface formed inside of the orifice. This would preclude the possibility for capillary forces to act in the small-diameter orifice, and thus it would be important to place the orifice check valve in a location where single-phase flow can

be guaranteed. Making this design choice would then make the lower bound on the orifice diameter restricted solely by manufacturing constraints.

The loss coefficient for a small-diameter orifice of the type envisioned for the orificed check valve can be approximated using the correlations of Idel'chik for a thick-edged orifice in a straight channel [71]. For bulk liquid potassium flow at ~ 1 m/s through an orifice of diameter 0.1 mm, $Re \approx 400$, and thus the relation for $Re \leq 10^5$ is used as reproduced in Equation 5.7, where $Re = \frac{\rho v D_h}{\mu}$ is the Reynold's number, ρ is the fluid density, D_h is the hydraulic diameter of the orifice, μ is the fluid dynamic viscosity, F_1 is the flow area outside the orifice, F_0 is the flow area inside the orifice, ℓ is the thickness of the orifice plate, and K_ϕ , K_0 , ϵ_0 , and λ are tabulated parameters found in Idel'chik.

$$K = \left(K_\phi + \epsilon_0^{-Re} K_0 + \lambda \frac{\ell}{D_h} \right) \left(\frac{F_1}{F_0} \right)^2 \quad (5.7)$$

Using representative parameters in Equation 5.7, it is found that the maximum loss coefficient through the calibrated orifice in the orificed check valve can be as high as roughly $K_r = 10^8$ when using an orifice diameter of 0.1 mm. This value will be used as the upper bound in this study of the modified ARC system.

In addition, either valve will result in some finite pressure loss during forward actuation. In general it is desirable to keep this forward loss as small as possible in order to enable a quick actuation. Ball check valves can have widely varying loss coefficients depending on the specific geometry, and so any value assumed for the proposed geometry would be, at best, a guess for this ill-defined geometry. This study assumes $K_f = 50$ as a lower bound, which is on the higher end for similar ball check valves (see, for instance, [95]), and it is assumed that if desirable, the forward loss coefficient can always be made larger.

Addition of Cover Gas

No matter which option is chosen for the pseudo-one-way valve, it is necessary to make a further modification to the ARC system to enable the valve to have the intended effect – the addition of a cover gas in the upper reservoir on top of the expander fluid. To understand why this is necessary, consider the one-dimensional incompressible Navier-Stokes equations as outlined in Equation 5.8, where z is the spatial coordinate, t is time, C_p is the isobaric heat capacity, T is the fluid temperature, and q''' is a heat source. For this thought experiment, consider that the liquid portion of the ARC system is filled with just a single fluid (potassium) in order to simplify the analysis without loss of generality. This is justified by the two fluids both being liquid metals with fluid properties of similar magnitude, as shown in Table 5.3 for representative ARC system temperatures.

Table 5.3: Select fluid properties for the ARC system working fluids at representative temperatures.

	Potassium	Lithium
Density (kg/m ³)	723	512
Viscosity (Pa-s)	1.64E-4	6.45E-4

$$\begin{aligned}
\frac{\partial \rho}{\partial t} + \frac{\partial \rho v}{\partial z} &= 0 \\
\rho \frac{\partial v}{\partial t} + \rho v \frac{\partial v}{\partial z} &= -\frac{\partial p}{\partial z} - \rho g - \frac{f}{D_h} \frac{\rho v |v|}{2} - K \frac{\rho v^2}{2} \\
\rho C_p \frac{\partial T}{\partial t} + \rho C_p v \frac{\partial T}{\partial z} &= q''' \\
\rho &= \rho(T)
\end{aligned} \tag{5.8}$$

In the original ARC system, for a given change in upper reservoir temperature there is a specified fluid velocity entering the inner ARC tube based upon mass conservation. Let us consider this velocity as a boundary condition, neglecting the dynamics in the upper reservoir itself and omitting this region from the considered domain. Let us also consider that the ARC tubes and lower reservoir are isothermal and no heat is added to this portion of the domain, since thermal expansion in those components is not the major motive for actuation. In such a case, the energy equation is superfluous. In addition, because the equation of state is only a function of fluid temperature and we have assumed $\frac{\partial T}{\partial t} = 0$, we can say that $\frac{\partial \rho}{\partial t} = 0$. This simplifies the mass equation to simply $\frac{\partial v}{\partial z} = 0$. Substituting into the momentum equation yields the following scalar equation:

$$\rho \frac{\partial v}{\partial t} = -\frac{\partial p}{\partial z} - \rho g - \frac{f}{D_h} \frac{\rho v |v|}{2} - K \frac{\rho v^2}{2} \tag{5.9}$$

In Equation 5.9, in fact all terms are fully specified from the boundary condition except the pressure gradient. This implies that no matter the value for K associated with the pseudo-one-way valve, the flow velocity *cannot* be impacted – it can only lead to excessively high pressures as the value of K is increased for a given upper reservoir temperature increase. In essence, the fluid velocity is driven entirely by mass conservation, with no influence from the system pressure gradients. This can be understood more physically by thinking of the ARC system fluid dynamics in terms of a resistance network. Because the upper reservoir is initially filled completely with expander fluid, there is an infinite resistance to upward flow and zero ‘capacitance’ due to the incompressible nature of the fluid. No matter how large of a resistance is placed below the upper reservoir, it will be smaller than infinity and thus fluid will always be forced through the valve with the same speed for a given driving

force. During the un-actuation sequence, the same principles apply but in reverse such that a thermal contraction of the expander fluid in the upper reservoir will induce a boundary condition on the fluid velocity which is not influenced by the addition of a one-way valve. Thus to allow the valve to have the desired impact, some ‘capacitance’ needs to be added to the system in the form of a cover gas. Such a cover gas will allow for changes in internal energy of the upper reservoir to be stored as changes in cover gas pressure such that the fluid velocity is then determined in part by momentum conservation as compared to purely mass conservation.

While this analysis made the assumptions that there was only a single liquid in the system and the system could be considered as isothermal, the same conclusions would apply if these assumptions were relaxed and more complex analyses were applied. The treatment of both expander and absorber fluids would require a jump discontinuity to be added to account for the differing densities of the fluids, but because both fluids are essentially incompressible this would make no difference. The fact that the ARC system is not isothermal would also not impact the general conclusion if it is accepted that changes in system pressure cannot *induce* changes in fluid temperature.

Summary of Conceptual Design

The conceptual design of the modified ARC system is presented in Figure 5.18. The pseudo one-way valve is placed at the bottom of the inner ARC tube, immersed deeply enough in the expander fluid that only expander fluid is expected to flow through the valve. The additional cover gas in the upper reservoir is stratified on top of the expander fluid such that there are now two separate gas regions in the upper reservoir. To keep these clear, the original gas chamber that is stratified on top of the absorber fluid is referred to as the ‘gas reservoir’, while the new gaseous region stratified on top of the expander fluid is referred to as the ‘cover gas.’ It is emphasized that these regions are separated by steel liner and thus do not directly interact with each other.




Figure 5.18: The modified ARC system separated from the fuel assembly (not to scale). Pseudo one-way valve is shown at the bottom of the inner ARC tube.

5.4 Modified ARC System Performance

This section assesses the modified ARC system's transient performance with intent to optimize its design to achieve inherent safety in all three proposed transients. However, first the computational modeling techniques need to be heavily adapted to the physics of the modified ARC system. Once a new computational scheme is in place, parametric studies over the new design variables are performed.

This section compliments a similar study that was performed to determine the benefits of the modified ARC system in an oxide ABR core [100]. In that study, it was found that the modified ARC system could provide significant benefits over the original design by essentially enabling more aggressive actuations. Given the fundamental differences between the oxide ABR and the reference B&B core of this dissertation, it is necessary to reevaluate these benefits specifically for the case of the B&B core, as is done in this section.

Computational Modeling Techniques

In the modified ARC system there no longer exists a direction functionality between the upper reservoir temperature and the ARC system insertion height. This is because the addition of the cover gas allows for both mass *and* momentum conservation to determine the system response, obviating the mass-conservation-only methods used in previous analyses.

Instead, the full set of Navier-Stokes equations must be solved to predict the internal fluid dynamics of the modified ARC system, something that is not possible within SAS itself. Therefore the analysis of the modified ARC system utilizes SAS coupled to SAM as discussed in more details later.

SAM Modeling

SAM solves the single-phase, incompressible, one-dimensional Navier-Stokes equations outlined in Equation 5.8. However, the fluid domain internal to the ARC system is in fact not single-phase (due to the presence of three fluids), incompressible (due to the inert gas), nor one-dimensional (due to the bulk volume in the lower reservoir). Nevertheless, SAM has been chosen for this analysis for a number of reasons.

First, while the inert gas volumes in the upper reservoir are compressible, the fluid *dynamics* in these regions is of very little interest. The gas volumes are important mainly as pressure boundary conditions and thus can be treated in a simplified manner within SAM (to be elaborated later). Second, even neglecting the gas regions, the domain is still two-phase due to the presence of both liquid potassium and lithium. From the perspective of the ARC system, however, the liquid potassium is much more important to the overall response firstly because there is much more of it than lithium and secondly because it is the fluid which thermally-expands and drives the actuation. Because these two liquid metals have roughly similar fluid properties, to first order the modified ARC system behavior will be well characterized if the domain is assumed to be entirely single-phase potassium. Considering the other uncertainties arising from poorly characterized material properties at elevated temperatures and the ill-defined geometry, the induced error is judged to be of minor significance. Finally, the assumption of one-dimensional flow is justified again because the actual fluid dynamics within the lower reservoir are of minimal interest in predicting the height of the absorber column. As long as the bulk mass balance within the lower reservoir can be computed in a zero-dimensional sense, the mass entering the active core can be predicted. In light of the alternative – multi-phase, compressible CFD calculations which would be prohibitively expensive for parametric analyses – the approximations made to enable the use of SAM are justified for the time being.

The computational model utilized in SAM is represented in Figure 5.19 with the previously outlined assumptions. Because (1) the volume of expander fluid is largest in the upper reservoir and (2) the upper reservoir is highly optimized for efficient heat transfer, it is assumed to be the only region to exchange heat with the bulk coolant flowing through the core. The heat exchange is effected using Newton’s law of cooling, where the heat transfer coefficient, h , is adjusted in order to achieve the desired value of τ . The upper reservoir itself is represented by a zero-dimensional liquid volume that is connected to a cover gas, where the liquid level is tracked and the volumes of each component are able to change with time based on the thermal expansion of the expander fluid and the pressure of the cover gas. The cover gas is modeled using the ideal gas law. The inner ARC tube is modeled as a one-dimensional pipe with connections to the upper reservoir and lower reservoir at top

and bottom. Between the inner ARC tube and the lower reservoir is placed the pseudo-one-way valve, which is represented using Equation 5.5 with different loss coefficients in each direction of flow. The lower reservoir is represented by a zero-dimensional voluminous liquid junction which uses a balance over each connection to determine the mass, momentum, and energy fluxes. Finally, the outer ARC tube and gas reservoir are modeled together in a single zero-dimensional component which uses the ideal gas law to calculate the pressure of the cover gas in conjunction with a mass balance to determine the liquid height. Such a zero-dimensional component neglects the pressure head from the cover gas, which is justified based upon its low density in comparison to the liquid.




Figure 5.19: The modified ARC system as modeled in SAM. Note the lack of absorber fluid.

To demonstrate the capabilities of the SAM model, results of a standalone toy problem where the driving temperature (coolant temperature) undergoes a step change of 50°C are presented in Figure 5.20. The impact of the valve forward loss coefficient (K_f) on the actuation is shown in terms of the maximum internal ARC system pressure and the actuation height (H_{ARC}) is shown. In these initial tests, the diodicity of the valve, $D = \frac{K_r}{K_f}$, is equal to 1. The value of K_f is seen to have a very large impact on the response for such a sharp increase in the coolant driving temperature. Unless $K_f \geq \sim 10^4$, large pressure oscillations which are slow to dampen are seen to occur, as the closed system acts largely like a manometer and the internal fluids move back and forth with a period of roughly 30 s. This behavior highlights the importance of having a forward loss coefficient higher than the lower bounds discussed in Section 5.3, without which the ARC system will overshoot and oscillate around

its equilibrium actuation level until damped by frictional losses. However it is noted that such oscillations are not necessarily detrimental if the top of the absorber column remains above the active core, in which case no oscillations will be reflected in the reactivity response.


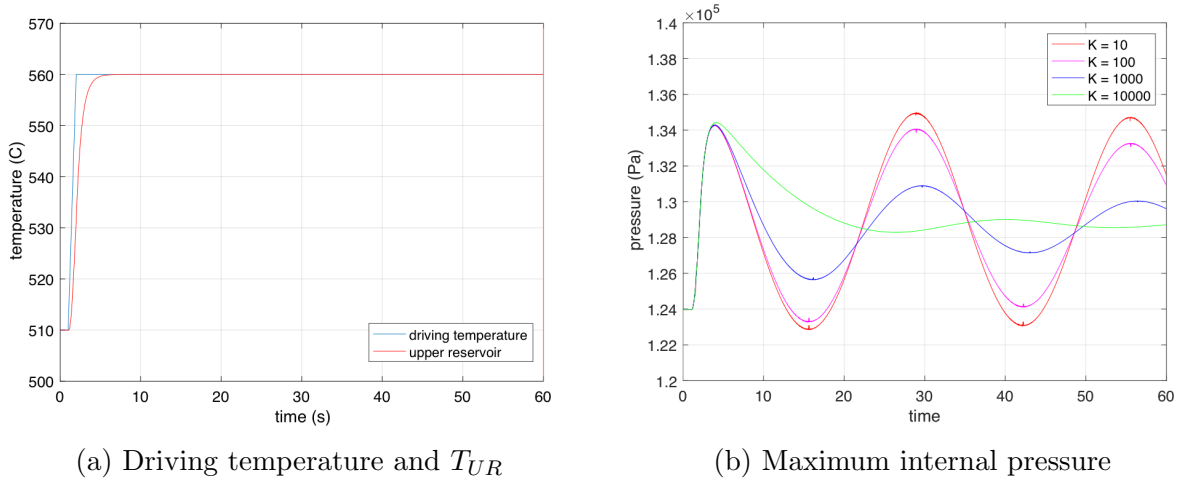


Figure 5.20: Response of the modified ARC system to a step change in coolant temperature for different values of K_f and $\tau = 1.3$ s.

To demonstrate the impact of having $K_r \geq K_f$, a new problem is presented in Figure 5.21, where the step increase in coolant temperature is followed by a step decrease. The impacts of increased diodicity can be seen in two aspects: (1) smaller oscillations following forward actuation and (2) a longer time to un-actuate following a decrease in coolant temperature.

Finally, the impacts of the cover gas volume (V_{CG}) can be seen in Figure 5.22. Because the cover gas acts as a buffer on system pressure, a larger cover gas leads to a slower actuation and a lower equilibrium insertion height.

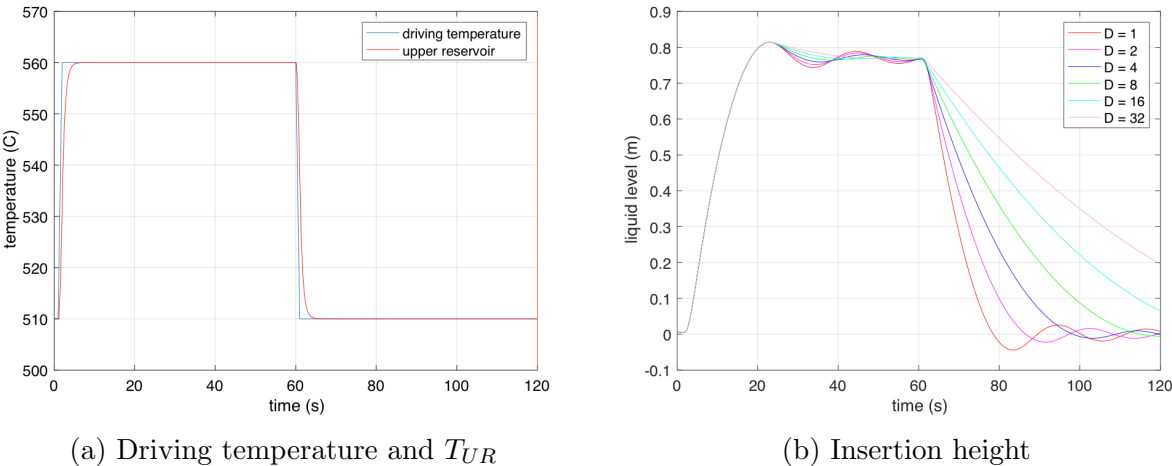


Figure 5.21: Response of the modified ARC system to a step increase and subsequent decrease in coolant temperature for different values of K_f and $\tau = 1.3$ s.




Figure 5.22: Height of the ARC system actuation in response to a step change in coolant temperature for $K_f = 10^4$, $D = 1$, and different values of V_{CG} .

These results demonstrate the important physical effects taking place within the modified ARC system and how they are impacted by design choices. In addition, they highlight the complicated nature of the response, underscoring how difficult it would be to use simplified

models within SAS to represent the system actuation. To summarize, the impacts of the new design variables of the modified ARC system are:

- K_f influences the speed of forward actuation and the overshoot beyond the equilibrium insertion height
- K_r influences the damping rate of any oscillations and the speed of reverse actuation
- V_{CG} influences the equilibrium insertion depth and the speed of forward actuation

SAS-SAM Coupling

Now that the details of the SAM model have been laid out, a flow chart of how SAS and SAM are coupled is provided in Figure 5.23. The coupling uses a Picard iteration paradigm in which a driver script written in Python coordinates the execution of each code and facilitates information exchange between them. Using the restart capabilities of each code, each time step is iterated until convergence. No modifications to either the SAS or SAM source code are required. The information exchanged between the two codes is as follows:

1. The driver extracts the coolant flowrate (F) and outlet temperature (T_{out}) calculated by SAS
2. The flowrate is converted offline to a heat transfer coefficient (h) corresponding to the value of τ at the current flowrate
3. h and T_{out} are passed to SAM, which uses them to predict the actuation height (H_{ARC}) at the current timestep
4. H_{ARC} is converted offline to a corresponding reactivity insertion using a table of w versus H_{ARC}
5. w is passed to SAS, which calculates an updated T_{out}
6. Iterations are repeated until T_{out} is sufficiently converged

Because the coolant outlet temperature at a given time step is only weakly influenced by the current reactivity insertion due to the time that it takes for the coolant to flow to the top of the active core, it is rare that multiple Picard iterations are performed on a given time step. Instead, the time step refinement plays a much more major role in converging the coupled-code system. Figure 5.24 shows the convergence study on two important variables during the early phase of a ULOF: the ARC system insertion height and the coolant outlet temperature. It can be seen, particularly for the outlet temperature predictions, that convergence is reached with a $\Delta t \approx 10^{-1}$ s, below which little change in the results is seen. At larger time steps (i.e. $\Delta t = 2$ s), the outlet temperature can be underpredicted by more than 60°C following the initial coast-down phase. Based on this convergence study and to strike a compromise

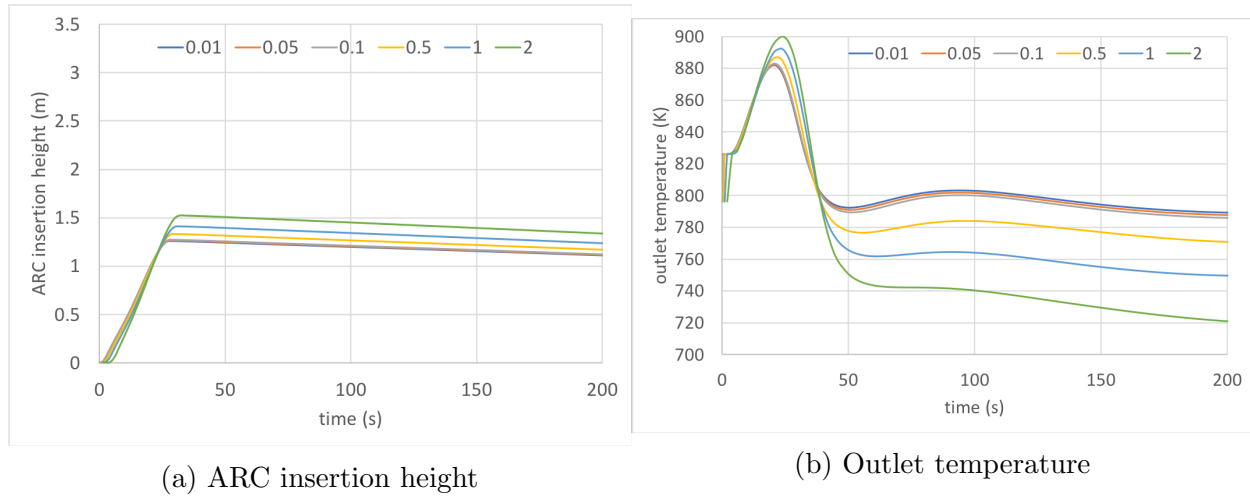


Figure 5.24: Predictions of the ARC insertion height and coolant outlet temperature as the time step size is refined in a ULOF transient. Time steps are in seconds.

the total ARC system worth fixed at the value determined to be optimal for the standard ARC system design, $w = \$1.75$. In addition, the purpose of the modified ARC system is to enable hyper-aggressive actuations, and thus smaller actuation worths are counter to the system's intended purpose.

In reality, the two most important design variables are likely the volume of the upper reservoir cover gas (V_{CG}) and the reverse loss coefficient (K_r), and thus this section takes them as the focus of the parametric study. However, other design parameters also play a role in the system transient response. For instance, it was shown in Figure 5.20c that the potential for overshoot and subsequent oscillations in the ARC system is high unless K_f is sufficiently large. In reality, the potential for overshoot is determined by a complicated interplay of many design variables, and if K_r is sufficiently large oscillations should be minimized no matter. But in order to minimize this potential, K_f is taken to be 10^4 , somewhat on the larger side. Other design variables not varied in this parametric study include the ARC tube inner radii (1.9 and 5.5 mm for the inner and outer tubes, respectively), ARC tube cladding thicknesses (0.6 mm each), heat transfer from coolant to upper reservoir ($\tau = 1.3$ s), distance between lower reservoir and active core (5 cm), amongst others.

ULOHS

Figure 5.25 shows the impact of V_{CG} on the progression of the ULOHS transient. Similar to the original ARC system ULOHS progression, shown in Figure 5.8, the system does not actuate until ~ 1000 s. Since the ARC system is not needed in the early transient phase to keep net reactivity negative, the early transient progresses in a favorable direction regardless of the ARC system design. Therefore the ARC system actuation has little impact on the

overall ULOHS progression *trend*, but is reflected mostly in the peak temperatures that are reached. As V_{CG} is made larger, the ARC system actuation becomes slower such that with volumes $\sim 900 \text{ cm}^3$ the net reactivity is hardly impacted. Smaller values of $V_{CG} \sim 50\text{-}100 \text{ cm}^3$ enable stronger actuations and reduced peak temperatures by roughly $40\text{-}50^\circ\text{C}$ as compared to systems with larger V_{CG} , and roughly $30\text{-}40^\circ\text{C}$ as compared to the ULOHS transient without ARC system inclusion.

However, in comparison to the ULOHS progression with the original ARC system (see Figure 5.8), the modified ARC system is found to perform slightly worse, with peak temperatures $\sim 20^\circ\text{C}$ higher in the case with the smallest cover gas volume examined. This is due to the cover gas, even in the smaller volumes investigated in Figure 5.25, subduing the strength of the ARC system reactivity insertion even though w is the same in both cases. Since the cover gas itself provides some capacitance for pressure buildup, any volume will serve to slow down the ARC system actuation, which is in fact a drawback of the modified ARC system. As the volume is reduced, the behavior approaches that of the original ARC system.

For the case of the ULOHS, no oscillatory behavior is seen because the transient is sufficiently slow and monotonic. This lack of oscillatory behavior, while of course desirable, does not allow for the true purpose of the *modified* ARC system and pseudo-one-way valve to be highlighted. Thus the ULOHS with the modified ARC system will actually see a slight degradation of performance as compared to the original ARC system. Because inherent safety is still trivially achieved, this is a small price to pay for the potential benefits enabled in the UTOP and ULOF transients discussed next. Since the system never sees oscillatory behavior, K_r will have no impact on the transient progression and thus a parametric study over this variable is not performed.

UTOP

Figure 5.26 shows the impact of V_{CG} on the progression of the UTOP transient. Similar to the ULOHS, the impact of V_{CG} on the modified ARC system performance is mainly to change the equilibrium reactivity insertion strength and the resulting peak temperatures. The impact on peak temperatures is much stronger for the UTOP however, with peak temperatures brought down by 250°C moving from cover gas volumes of 900 cm^3 down to 50 cm^3 . Such sensitivity is in line with the overall strength of the ARC system insertion, with the UTOP reaching full insertion relatively early in the transient, while the ULOHS only ever saw insertions up to $\sim \$0.1$.

Similar to what was seen in the more-aggressive UTOP transients with the original ARC system, the early phase of the transient during the rod withdrawal shows strong temperature oscillations. These oscillations are shown in more detail in Figure 5.27. However, contrary to the original ARC system with $\tau = 1.3 \text{ s}$, the modified ARC system is able to damp these oscillations in reactivity-space with the pseudo-one-way valve such that diverging oscillations do not form. The effect of the pseudo-one-way valve can be seen in the ARC reactivity response of Figure 5.27b, where very strong negative reactivity insertions are followed by much

slower, temporary reverse actuations. The incorporation of the pseudo-one-way valve leads to this “chugging-type” actuation sequence, which although is a non-smooth, rather complicated phenomenon, is far better than diverging oscillators and enables rapid responses from the ARC system as long as V_{CG} is made sufficiently small. Since the modified ARC system is made to be very sensitive, small temperature-reactivity oscillations are seen throughout the transient, even as the core settles at a new equilibrium. These small oscillations are due to the rapid actuation of the ARC system in response to small increases in temperature which in fact overshoot, thus leading to small decreases in temperature and subsequent re-actuation of the system. Although these oscillations are small in magnitude and non-divergent, they represent one of the drawbacks of having such a sensitive ARC system as compared to one with a more delayed, slower onset.

Figure 5.26 also demonstrates that unacceptable performance will be seen if V_{CG} is made too large. First, the onset of boiling is seen towards the end of the rod withdrawal sequence when $V_{CG} = 900 \text{ cm}^3$. More restrictive, however, is likely the peak cladding temperatures which are reached at the equilibrium state for $V_{CG} \geq 300 \text{ cm}^3$. In these cases, the fuel-cladding eutectic temperature can be greatly exceeded, at which point the integrity of the cladding in peak assemblies may be severely degraded if the reactor is kept at this state for any significant amount of time. This puts a hard limit on how large the cover gas volume can be made for acceptable performance.

Similar to the ULOHS behavior, the optimal modified ARC system does not perform better than the optimal original ARC system, but rather the modified version approaches similar behavior to the original system as V_{CG} is made smaller. This is because in the UTOP with the original ARC system, acceptable non-oscillatory behavior was able to be achieved anyways by making τ large enough to effectively decouple the temperature oscillations from the reactivity insertion. The pseudo-one-way valve accomplishes effectively the same goal via a different mechanism, which then allows τ to be made small again. While this does not provide immediate benefits to the UTOP progression as compared to the original ARC system, the potential benefit of allowing τ to again be very small is that the system actuation may be rapid enough to handle the initial coast-down period in the ULOF, as investigated in the next section.

With $V_{CG} = 100 \text{ cm}^3$, Figure 5.28 shows the impact of K_r . As K_r is decreased through multiple orders of magnitude, the magnitude of the reverse actuations increases notably, as expected. Whereas the ARC behavior with $K_r = 10^8$ can be described as ‘chugging’, as the loss coefficient is reduced the behavior looks more and more like oscillations which continue essentially indefinitely once the system is actuated. While these oscillations are not seen to diverge and peak temperatures remain within their bounds, the large, fast temperature fluctuations in the case where $K_r = 10^6$ will have negative implications for the integrity of structural components in the core and should be avoided. Therefore this study puts a lower bound on the acceptable value of K_r and a constraint on the design of the pseudo-one-way valve.

ULOF

Figure 5.29 shows the impact of V_{CG} on the progression of the ULOF transient. Compared to either the ULOHS or UTOP, the progression of the ULOF is much more strongly impacted by the ARC system design. To start, the early phase of the ULOF transient is highly impacted by V_{CG} . As flow ramps down and temperatures initially increase sharply, ARC system designs with sufficiently small V_{CG} see a negative reactivity insertion strong enough to rapidly reduce temperatures to levels in fact *below* their initial steady-state values. While this drives the ARC system to un-actuate, the impact of the pseudo-one-way valve slows the reverse actuation such that oscillations do not develop. This is in contrast with the original ARC system, which developed diverging oscillations during the coast-down phase when τ was small enough (see Figure 5.14). This result is the first direct confirmation that the modified ARC system can provide benefits by enabling aggressive actuations while avoiding oscillations.

No matter the impacts during the early transient phase, behavior tends to converge to the same trend in the vicinity of ~ 1000 s when the transition to natural circulation takes place. This flow transition leads to rapid temperature spikes and a second major ARC system actuation, this one being much stronger than the first. Still, however, thanks to the pseudo-one-way valve, diverging temperature-reactivity oscillations can be avoided if V_{CG} is small enough such that the associated reactivity insertion is sufficiently large to overcome the strongly-positive coolant density feedback. For cases with sufficiently small V_{CG} , core temperatures stabilize, at least temporarily, far below their steady-state conditions.

Ultimately, the fact that temperatures following the transition to natural circulation stabilize *below* their steady-state values is the cause of the rapid temperature and power ramps which lead to boiling and/or fuel melting in all cases. This assertion can be understood from Figure 5.30. During the flow transition, the temporarily increased power-to-flow ratio leads to temperature increases and subsequent strong actuation of the ARC system. The strong, sustained actuation essentially causes the fission power to drop to zero, with all power then being decay heat which is low enough such that the power-to-flow ratio stabilizes near 0.5. This low power-to-flow ratio leads to reduced core temperatures such that the majority of reactivity feedback mechanisms turn positive and without the ARC system, net reactivity would be strongly positive.

Initially as the ARC system slowly disengages, temperatures do not increase because net reactivity is still strongly negative. However, as the un-actuation proceeds, eventually net reactivity turns positive and fission power begins to rise, but because the fission power level is initially so drastically low, this power increase is not reflected in core temperatures. As the fission power level continues to increase over the span of multiple minutes, it passes through orders of magnitude until eventually reaching levels on par with the decay heat, at which point power is increasing so quickly that the negative feedback mechanisms associated with core temperature increases cannot act in time to arrest the transient. While this result has only been presented for the specific ARC system design with $V_{CG} = 100 \text{ cm}^3$, the same type of behavior is present irrespective of the cover gas volume, with the only difference being the

time that boiling is initiated, as shown in Figure 5.29.

Such results imply that core temperatures below their steady-state values while the ARC system is strongly actuated is an unstable situation which should be avoided, since it allows for a state that would have positive reactivity if the ARC system were taken away. To avoid this situation while still allowing the ARC system design to be aggressive enough such that boiling during the early phase of the ULOF can be avoided, either the rough transition from forced- to natural-circulation should be smoothed such that the ARC system does not actuate so strongly and the oscillation potential is avoided in the first place, or the value of K_r should be further increased such that the reverse actuation and the resulting core temperature increases are slow enough for other reactivity feedback mechanisms to stabilize the power at some intermediate ARC insertion level. In light of the fact that K_r is already taken as the assumed upper limit in the results reported in Figure 5.29, this second method is not viewed as viable. In addition, the first method would be counter to the purpose of the modified ARC system in the first place, since it was shown in Figure 5.15 that the *original* ARC system could avoid diverging oscillations in the ULOF if the flow transition were made sufficiently smooth anyways.

Sensitivity to Steam Generator Boundary Conditions

Since it is undesirable to have core temperatures below their steady-state values while the ARC system is strongly actuated, this section investigates the boundary conditions on the steam generators to understand their impact on putting the core in this unstable configuration. Within SAS, there are two options for specifying the behavior of the steam generators as a boundary condition on the secondary loop: (1) specify the secondary coolant temperature drop across the steam generator with time, or (2) specify the secondary coolant temperature at the outlet of the steam generator with time. All previous calculations in this dissertation have utilized option (1), where the steam generator temperature drop is assumed to remain constant at its steady-state value throughout the transient. This might be the case, at least approximately to first order, if the tertiary loop were unchanged during the transient. However this scheme leads to the reduced core inlet temperatures which ultimately bring temperatures below their steady-state conditions, through the reduction in power and the reduced power-to-flow ratio shown in Figure 5.30.

To avoid this unstable configuration, the steam generator boundary condition (2) is attempted, where the outlet temperature of the steam generator is held constant at its steady-state value. This could be accomplished, for instance, through a steam bypass that adjusts the feedwater rate to the steam generators as the reactor power level decreases. While this would require more complicated steam generator controls, it is likely a more realistic option for very low power regimes since the other boundary condition could artificially inflate the reactor power level by demanding more heat be rejected than what is actually being produced. In keeping the steam generator outlet temperature constant, the core inlet temperature is actually slightly increased as the transient progresses due to the reduced heat

rejection ability in the intermediate heat exchanger under low-flow conditions. Thus, the unstable configuration can potentially be avoided.

Figure 5.31 shows the results of the ULOF transient with various upper reservoir cover gas volumes and the modified steam generator boundary condition. It is found that the ARC system enables inherent passive safety in the ULOF under the new steam generator boundary conditions, although the transient evolution is found to be quite complicated. Again, the pump coast down period leads to rapidly increasing coolant temperatures. In cases where the upper reservoir cover gas is sufficiently small, the ARC system is actuated rapidly causing temperatures to quickly decrease. This leads to a subsequent un-actuation of the ARC system, although the pseudo-one-way valve prevents oscillations from forming in response. As the transient progresses, eventually the flow transition occurs at ~ 1000 s and the ARC system is strongly actuated in response to the resulting temperature spike. In this case, however, the modified steam generator boundary condition makes it so that the resulting reactivity insertion is smaller than it was with the previous boundary condition, as displayed in Figure 5.30. This is because the modified steam generator boundary condition leads to an increased coolant inlet temperature which has a large impact on making the radial expansion feedback mechanism more strongly negative. Because of this increased contribution, the ARC system no longer has to hold down as much negative reactivity to keep the system subcritical, and thus ARC system actuations are seen to be, in general, smaller.

As a direct result of the increased core inlet temperature, the core is not put into the unstable configuration which led to the rapid power spikes displayed in Figure 5.29. This enables the core to survive the ULOF transient without reaching even the fuel-cladding eutectic temperature, showing a very robust response to such a severe transient. As in previously examined transients, smaller cover gas volumes allow for reduced peak temperatures, particularly during the very early coast down phase where the temperature spikes due to a power-flow mismatch. In this early phase, peak coolant temperatures can be brought down by nearly 40°C between different cases examined in this study. While this is important margin, in reality the strongest benefit that the modified ARC system brings to the ULOF transient is the ability to achieve inherent passive safety in the first place, with the peak temperatures being a secondary driver of the cover gas volume.

Summary of Modified ARC System Performance

This section has assessed the performance of the modified ARC system in the reference B&B core using a new coupling between the transient dynamics code SAS and the thermal-hydraulics code SAM, wherein SAS is used to simulate the whole-plant response and SAM is used to simulate the ARC system internal fluid dynamics. It has been found that the transient response to the three unprotected accident scenarios is highly sensitive to the volume of the upper reservoir cover gas, particularly in the UTOP event. In general as the cover gas volume increases, the sensitivity and strength of the ARC system goes down and peak temperatures increase due to the pressure capacitance of the cover gas. This

phenomenon makes it so the ULOHS and UTOP scenarios actually perform slightly worse than those with the original ARC system, although this discrepancy can be minimized by using sufficiently small cover gas volumes. However, the inclusion of the pseudo-one-way valve allows for very strong and rapid ARC system actuations without inducing oscillations, which benefits the ULOF scenario most strongly. Using this modified ARC system and assuming a proper design of the steam generator control systems, the ULOF accident can be handled with inherent passive safety, a strong improvement over the original ARC system and the B&B core sans ARC system. It should be mentioned, however, that these results appear strongly dependent on the steam generator behavior during the transient, which has a large impact on the core inlet temperature. If improperly designed, the hyper-sensitive ARC system actuation has the potential to bring the reactor to an unstable configuration which ultimately leads to very strong, although delayed, power spikes and onset of boiling.

This improved transient performance, while attractive from the perspective of inherent safety, comes at the cost of significantly more complicated transient progressions. In comparison to the smooth, monotonic behavior displayed by most other feedback mechanisms, the modified ARC system, if made strong enough to counter the rapid dynamics in the ULOF scenario, leads to a ‘chugging’ actuation sequence in both the UTOP and ULOF scenarios. While this is preferable to diverging oscillations, the pressures, forces, and temperatures encountered in such actuation sequences may be both undesirable and difficult to predict with high accuracy.

In addition, the physical design of the ARC system becomes more complicated with the modified ARC system, as the operation of the pseudo-one-way valve requires the addition of a second cover gas of relatively specific volume. In the case that somehow the volume of the cover gas changes – perhaps through leakage, accidental mixing with the other cover gas during the manufacturing process, or the production of gases within the ARC system itself – the actuation properties of the system may change dramatically in ways which are difficult for the reactor operator to diagnose. Further, this study has shown that the pseudo-one-way valve must have exceptionally large loss coefficients in order to effectively damp oscillations due to the relatively low fluid velocities inside the ARC system. Although it is believed that valves with sufficiently large loss coefficients could, in principle, be manufactured, they will likely have orifices of very small diameters which are potentially prone to plugging or closure with time, which though would not have negative *safety* implications, would make it more difficult to un-actuate the system. Therefore, while the modified ARC system has been shown to enable aggressive actuations without inducing diverging oscillations, its additional complexities may challenge the ultimate desirability.







Figure 5.25: Peak temperatures and reactivity evolutions for the ULOHS scenario with modified ARC system inclusion for various upper reservoir cover gas volumes.




(a) Net reactivity




(b) ARC reactivity



(c) Peak fuel temperature



(d) Peak cladding temperature



(e) Peak coolant temperature

Figure 5.26: Peak temperatures and reactivity evolutions for the UTOP scenario with modified ARC system inclusion for various upper reservoir cover gas volumes.

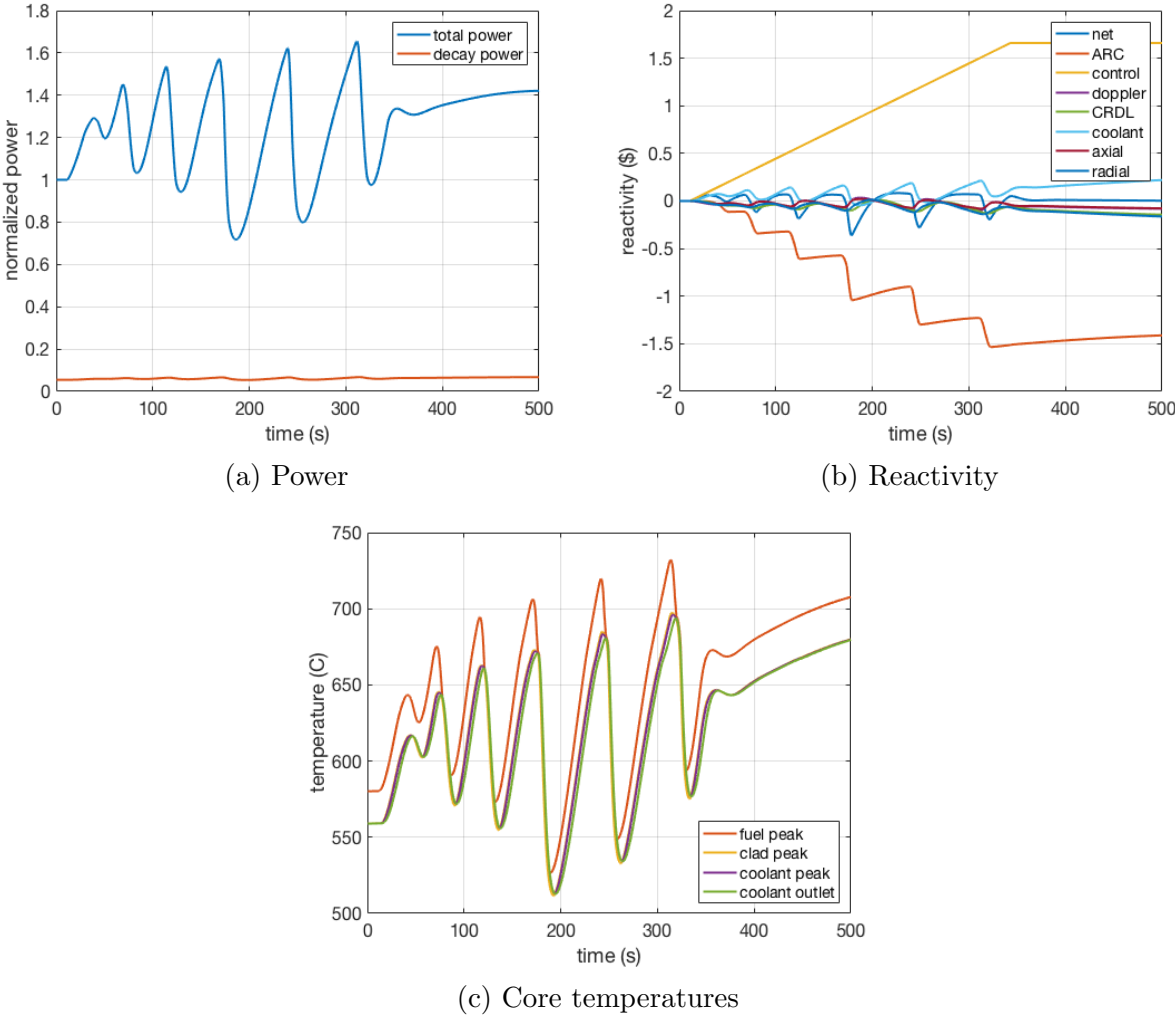


Figure 5.27: Early phase of the UTOP transient with modified ARC system inclusion and $V_{CG} = 100 \text{ cm}^3$, highlighting the damped reactivity oscillations.

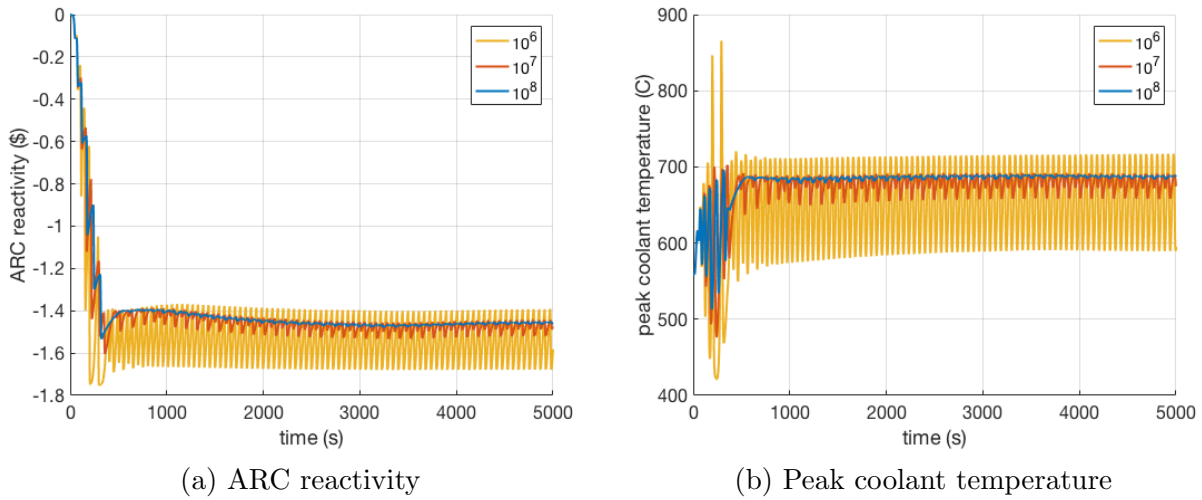





Figure 5.28: Impact of the reverse loss coefficient in the pseudo-one-way valve in the UTOP scenario.




(a) Net reactivity




(b) ARC reactivity



(c) Peak fuel temperature




(d) Peak cladding temperature




(e) Peak coolant temperature

Figure 5.29: Peak temperatures and reactivity evolutions for the ULOF scenario with modified ARC system inclusion for various upper reservoir cover gas volumes.



(a) Reactivity feedbacks



(b) Power and flow

Figure 5.30: Reactivity feedback and power response to the ULOF with modified ARC system inclusion and $V_{CG} = 100 \text{ cm}^3$.

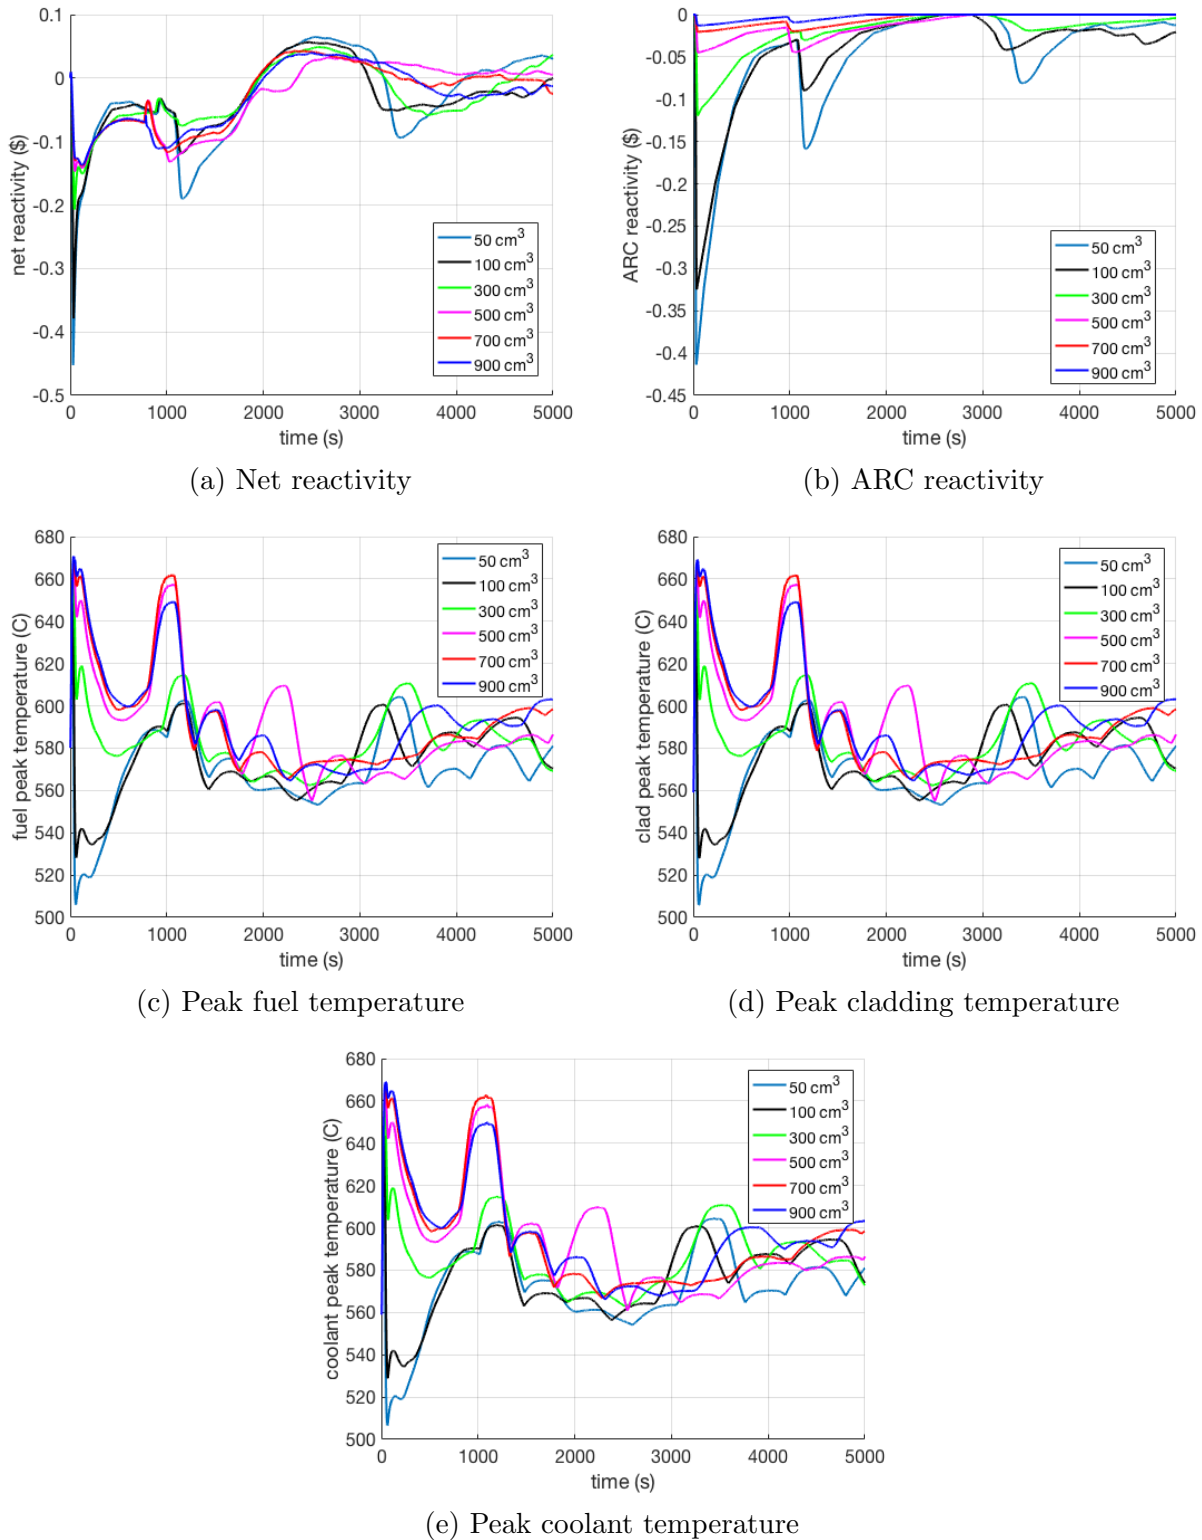


Figure 5.31: Peak temperatures and reactivity evolutions for the ULOF scenario with modified ARC system inclusion for various upper reservoir cover gas volumes with steam generator outlet temperature held constant.

5.5 Other Impacts of ARC System Inclusion

It has been shown that ARC system inclusion has the potential to enhance SFR performance. However it is expected that inclusion will also induce a neutronic penalty stemming from the removal of fuel and the addition of non-fuel in-core materials. The purpose of this section is to investigate this penalty, while also examining the amount of tritium produced in ARC system components and the amount of absorber depleted over the ARC system's residence time. To make this study general, a wide range of ARC system designs are investigated, allowing for sensitivities to ARC design parameters to be identified. While this study assumes the original ARC system as described in Section 5.1, there are not anticipated to be major differences in the neutronic impact resulting from the modified ARC system of Section 5.3.

Computational Models

The reference core outlined in Chapter 2 is used in this study. It is assumed that each fuel assembly is equipped with an ARC system of equivalent design. Therefore it is not expected that a significant change in the radial flux profile will be induced, and impacts on the radial power profile are not examined. Instead, this study is interested in the quantification of k_{eff} . Appendix A demonstrates the adequacy of r-z models at predicting k_{eff} for B&B cores, and therefore this study employs r-z models for all criticality calculations to save on computational expense. All calculations except those to quantify tritium production are performed with the ENDF/B-VII.1 cross-section library. Calculations to quantify tritium production are performed using the JEFF-3.1 library due to the lack of fission product yields for tritium in the ENDF library.

To quantify the various impacts, modifications are made to an equilibrium cycle model of the reference core without ARC system inclusion. After making modifications to represent different aspects of ARC system inclusion, the model is run again using the equilibrium cycle characteristics of the reference core. New equilibrium cycle calculations are not performed. It is natural to measure the neutronic penalty in terms of the increase in equilibrium cycle length, but because there is a direct relation between core excess reactivity and cycle length, quantifying the k_{eff} penalty is equivalent to finding the resulting cycle length increase. In light of this relationship and the longer run times of equilibrium calculations, the neutronic penalty is quantified in terms of k_{eff} evolution rather than equilibrium cycle characteristics.

Determination of Soluble Lithium Distribution

Although the three ARC system fluids are largely immiscible, each fluid is soluble within the others to a small degree. This is especially important in the direction of lithium being dissolved in potassium or helium, as this allows for lithium to be present in the active core, which may lead to a neutronic penalty or the production of tritium. To quantify these impacts, first the steady-state distribution of soluble absorber fluid within the core must be determined.

Within the ARC system, the absorber fluid is in contact with both the expander fluid and the inert gas. The geometry near these interfaces is represented in Figure 5.32a. The lithium can take two paths to enter the active core region via diffusion: (1) upwards into the helium gas or (2) first downwards into the lower reservoir of potassium and then upwards through the inner concentric tube. Due to geometric uncertainties associated with path (2), a conservative assumption is employed: the distance of initial downward diffusion of lithium into potassium is negligible and the lithium is assumed to be located immediately below the inner concentric tube, only having to diffuse upwards to reach the active core. In reality, the downward diffusion step will act as a buffer delaying the poison from entering the active core. This assumption allows for both pathways to be treated as independent 1-dimensional diffusion problems of lithium into the active core, as depicted in Figures 5.32b and 5.32c.

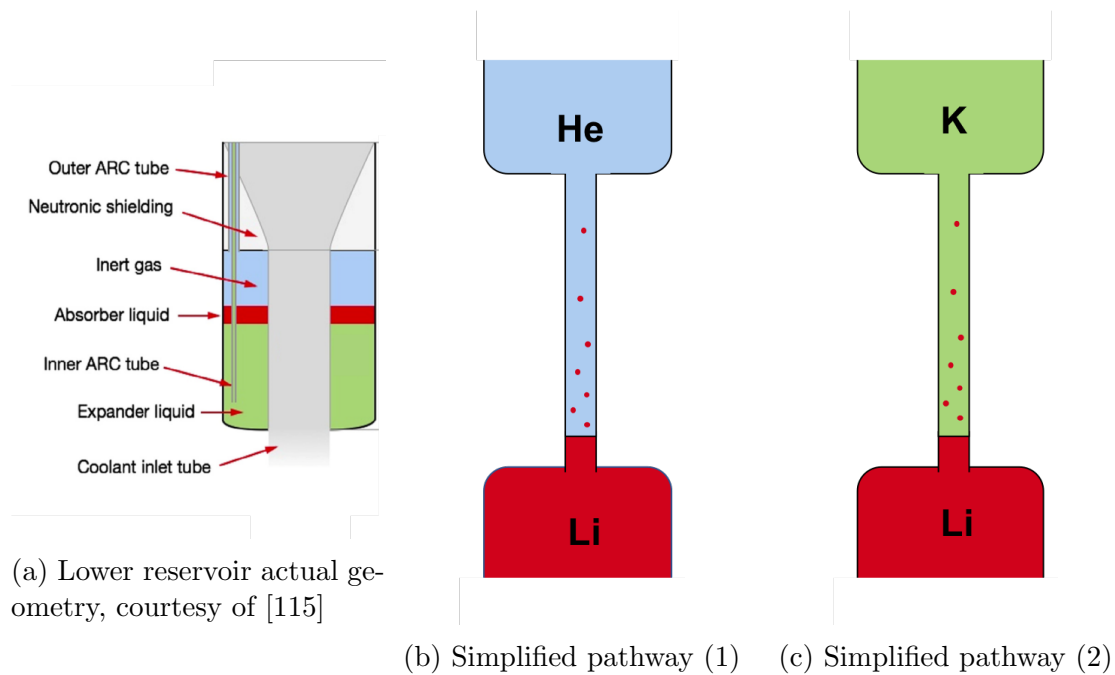


Figure 5.32: Actual and simplified ARC system geometries used to study the impact of dissolved lithium in the core, fictitiously depicting the migration of lithium upwards into the active core.

During operation, the in-core lithium profile will reach a steady-state according to a balance of lithium sources and sinks. Including 1-dimensional diffusional fluxes and a volumetric neutron flux sink, this balance is represented by Equation 5.10, where N is the molar concentration of the mixture, N_{Li} is the molar concentration of lithium, $x_{Li} = N_{Li}/N$, D_i is the diffusion coefficient of lithium through either liquid potassium or gaseous helium, ϕ is the neutron flux, and σ is the microscopic cross-section for lithium depletion. Due to the

small diameter of the ARC system tubes (a few millimeters), convective mass transport is assumed negligible.

$$\frac{dN_{Li}(z)}{dt} = -\frac{d}{dz}N(z)D(z)\frac{dx_{Li}(z)}{dz} - \phi(z)\sigma(z)N_{Li}(z) \quad (5.10)$$

For the case that the dissolved lithium fraction is very small, we can approximate $N \approx N_i$, where N_i is the molar concentration of either potassium or helium. Due to the low solubility of lithium in the other ARC fluids, the diffusion coefficient can be made independent of position if evaluated at an average system temperature. Neglecting time dependence, we then have a steady-state reaction-diffusion equation.

$$0 = -D\frac{d^2N_{Li}(z)}{dz^2} - \phi(z)\sigma(z)N_{Li}(z) \quad (5.11)$$

To determine the relative magnitudes of the competing diffusion and depletion mechanisms, Equation 5.11 can be nondimensionalized by introducing dimensionless variables $\bar{N}_{Li} = N_{Li}AL$ and $\bar{z} = z/L$, where L and A are a characteristic length and area, respectively. Substituting these variables into Equation 5.11, the nondimensional equation is shown below:

$$0 = -\frac{d^2\bar{N}_{Li}}{d\bar{z}^2} - \frac{\phi\sigma L^2}{D}\bar{N}_{Li} \quad (5.12)$$

In this equation, a dimensionless term similar to the Thiele Modulus is formed as $\frac{\phi\sigma L^2}{D}$. This term represents the ratio of lithium depletion via neutron transmutation to lithium replenishment via diffusion. Evaluating this dimensionless ratio with representative values gives an estimate of the relative strengths of the two competing processes.

The flux can be determined for the B&B core with Serpent 2. Doing so, one finds that the flux may vary from about 10^{13} 1/cm²/s in the core periphery regions up to about 10^{15} 1/cm²/s in the high fission-density region of the core.

The lithium depletion cross-section can be determined similarly. Natural lithium is composed of two isotopes with depletion modes shown in Figure 5.33. All transmutation pathways, with the exceptions of ${}^6\text{Li}(n,t){}^4\text{He}$, ${}^7\text{Li}(n,\gamma){}^8\text{Li}$, and ${}^6\text{Li}(n,\gamma){}^7\text{Li}$, are threshold reactions above roughly 5 MeV. As seen in Figure 5.34, the neutron flux at these high energies is comparatively little, implying threshold reactions can be ignored. Additionally, the ${}^7\text{Li}(n,\gamma){}^8\text{Li}$ and ${}^6\text{Li}(n,\gamma){}^7\text{Li}$ cross-sections are neglected because they are roughly 4 orders of magnitude smaller than that for ${}^6\text{Li}(n,t){}^4\text{He}$. Finally, assuming 100% enrichment in ${}^6\text{Li}$ leads to a conservative assessment, as this isotope has both higher transmutation cross-section and will produce more tritium. Evaluating this cross-section over the spectrum of the B&B core, the transmutation cross-section is found to be roughly 1 b.

The characteristic length scale L is taken as the system diffusion length, $L = 2\sqrt{Dt}$, where t is the assembly in-core residence time. Making this substitution, it is found that the dimensionless ratio is independent of the diffusion coefficient, making the relative strength




Figure 5.33: Transmutation cross-sections in both isotopes of lithium.

between flux depletion and mass diffusion dependent only on time, neutron flux level, and the transmutation cross-section.

Using the values determined above, a range for the dimensionless ratio is shown in Figure 5.35 for two different neutron flux levels representing positions of high flux (10^{15} $1/\text{cm}^2/\text{s}$) and positions of low flux (10^{13} $1/\text{cm}^2/\text{s}$). In regions of low flux, the steady-state profile is mainly governed by mass diffusion. For regions of high flux, the balance is dominated by mass diffusion in earlier times, but dominated by flux depletion at later times. However, for a B&B core, much of an assembly's lifetime is spent in regions of lower flux level as fissile material is being bred. Therefore, for much of an assembly's lifetime, the lithium concentration profile will be dominated by mass diffusion.

This analysis implies that the steady-state lithium profile is not dominated by flux depletion, meaning that there will be some amount of lithium present within the expander and inert gas ARC fluids. Although Equation 5.11 can be iteratively solved according to the specific flux profile of this particular B&B core, the nondimensional analysis implies that the neutron flux depletion can be neglected in many cases. Any lithium concentration profile derived under this assumption is conservative in the sense that it will over-predict the amount of lithium in the core, giving upper bounds on both the neutronic penalty and the amount of in-core tritium produced. Making this assumption, Equation 5.11 becomes:

$$0 = \frac{d^2 N_{Li}(z)}{dz^2} \quad (5.13)$$

Integrating twice gives a linear relation, and the constants of integration are found




Figure 5.34: Normalized neutron spectra at two radial points within the B&B core. Batch 5 corresponds to a position close to the high flux zone and has experienced 12 burnup cycles, while Batch 10 corresponds to a position close to the core periphery and has experienced 2 burnup cycles.

through the boundary conditions. The concentration of lithium immediately at the fluid interface can be determined using thermodynamic data if the two fluids are in thermal equilibrium, providing a boundary condition at the lower end. For both scenarios, each chamber of the upper reservoir (filled with either potassium or helium) is assumed to be an infinite sink for soluble lithium, providing a boundary condition at the upper end. This infinite sink assumption is validated subsequently in Section 5.5. Finally, the lithium profile is determined as shown in Equation 5.14, where S_i is the equilibrium solubility of lithium in fluid i and L_{ARC} is the axial length of the ARC system tubes.

$$N_{Li}(z) = -\frac{S_i N_i}{L_{ARC}} z + S_i N_i \quad (5.14)$$

The solubility of liquid lithium in liquid potassium at equilibrium can be obtained from the equilibrium phase diagram, whereas the solubility of liquid lithium in gaseous helium can be obtained from the vapor pressure of lithium. Although the outlet temperature for the B&B system under consideration is outside of the range of validity by 10°C , Equation




Figure 5.35: Dimensionless ratio as a function of residence time in the core for two different values of neutron flux.

5.15 from Bale [16] is used to determine the solubility of lithium in potassium due to lack of data at higher temperatures.

$$S_{Li \rightarrow K} [\text{atom}\%] = 691.86 \times \exp -\frac{4229.9}{T + 273} \quad : \quad 63 \leq T \leq 500^\circ C \quad (5.15)$$

The equilibrium solubility of lithium in potassium is 0.82 atom% at the inlet temperature, while the partial pressure (and thus solubility) of lithium into gaseous helium is near zero [73]. Because the lithium profile is a maximum at the fluid interface, the low solubility indicates that lithium dissolution in helium will be negligible, allowing for the in-core inventory of lithium to be evaluated as exclusively that which is dissolved in potassium.

Neutronic Impact

To identify the root causes of any neutronic penalty, the study is done in steps as follows:

1. Fuel is removed from a number of pins per assembly and replaced with void to evaluate the impact of lower fuel volume fraction

2. A range of combinations of expander fluid, inert gas, and steel are added to the inside of the voided fuel pins to evaluate the impact of increased parasitic absorption by ARC in-core materials
3. Absorber fluid is added to the expander fluid to evaluate the impact of soluble absorber diffusion into the active core

The results of these steps are outlined in the following subsections.

Removal of Fuel

In the standard ARC system design, multiple fuel pins may be replaced with multiple sets of concentric tubes in order to achieve a strong enough negative reactivity insertion. This impact is quantified by reducing the fuel volume fraction in each assembly to represent the loss of one, two, and three pins worth of fuel with the remaining volume fractions left unchanged. The reduced fuel volume fraction is made up by introducing a corresponding void volume fraction. The impact on k_{eff} of these modifications is shown in Figure 5.36.




Figure 5.36: k_{eff} evolution in the equilibrium cycle with the fuel removed from a number of pins.

The decrease in core reactivity is approximately linear with the number of fuel pins removed, and the maximum penalty is roughly 700 pcm for the case with three pins removed. The impact of this reactivity penalty on the minimum required burnup (and hence material damage accumulation) can be approximated using the relationship between BU^* and the fuel

volume fraction derived by Qvist and Greenspan in [116]. For removal of 1 pin's worth of fuel, the fuel/coolant volume ratio is reduced by 0.6%. For a core leakage probability $\sim 1\%$, the results of Qvist imply an increase in the BU^* of $\sim 0.1\%$ on top of the $\sim 30\%$ burnup peak already accumulated at discharge, which would imply an additional $\sim 2-3$ DPA per fuel pin replaced by ARC tube. For most ARC system designs it is expected that only a handful of fuel pins will need to be replaced by ARC tubes; therefore, this quantification is expected to cover the majority of cases that a designer may wish to explore.

Addition of ARC Tubes

This section introduces the in-core ARC materials (i.e. expander fluid, inert gas, and steel) into the model to determine their neutronic impact. In reality, some amount of absorber fluid will be soluble in the other fluids and will therefore also be present in the active core – this impact is quantified in Section 5.5.

In practice, a designer may wish to tailor the actuation characteristics by tuning the volume fractions of each ARC system constituent. To keep this study general, a variety of ARC system designs are investigated in terms of these variable volume fractions. To be conservative, it is assumed that three fuel pins have been replaced with ARC tubes. Then, a series of criticality calculations are performed while varying the volume fractions within the ARC tubes of steel, expander fluid, and inert gas. The minimum volume fraction for any of the three constituents is assumed to be 0.1. The assumed materials for this study are HT9 steel, liquid potassium, and gaseous helium for structural materials, expander fluid, and inert gas, respectively. The ideal gas law at a pressure of 1 atm is used to determine the number density of helium atoms. Density correlations for HT9 and potassium are taken from [115] and repeated here, along with that for lithium (the absorber fluid):

$$\rho_{HT9} \left[\frac{kg}{m^3} \right] = 7824 - 9.288 \times 10^{-5}T - 0.1859T^2 \quad : \quad 298 \leq T \leq 923 \text{ K} \quad (5.16)$$

$$\rho_K \left[\frac{g}{cm^3} \right] = 0.8415 - 2.172 \times 10^{-4}T - 2.70 \times 10^{-8}T^2 + 4.77 \times 10^{-12}T^3 \quad : \quad T \leq 1250^\circ C \quad (5.17)$$

$$\rho_{Li} \left[\frac{kg}{m^3} \right] = 562 - 0.1T \quad : \quad 400 \leq T \leq 1300 \text{ K} \quad (5.18)$$

Compared to the removal of fuel, the addition of the ARC system materials results in a smaller k_{eff} penalty. The worst neutronic performance results from designs with higher steel and expander volume fractions due to their higher densities and cross-sections compared to the inert gas, leading to a neutronic penalty of at most 100 pcm. This study shows that, even with three fuel pins replaced by ARC tubes, a large design space exists to tune ARC systems without incurring significant neutronic penalties.

Soluble Absorber

To assess the neutronic impact of the soluble lithium, a bounding criticality calculation is performed with the ARC design that introduces the most in-core lithium: three fuel pins replaced by ARC tubes of 80, 10, and 10 vol% expander, steel, and gas, respectively. The lithium is 100% enriched in ${}^6\text{Li}$ and dissolved in the expander fluid according to Equation 5.14. The Serpent model allows for depletion of fuel materials but keeps the concentration of lithium constant to reflect the assumption that lithium replenishment is faster than depletion. The criticality calculations show a neutronic penalty from the soluble lithium so small that it cannot be distinguished within the 20 pcm uncertainty of the calculations, showing that the neutronic penalty from in-core dissolved lithium is negligible.

Tritium Production

In-Core Tritium Production

Although the neutronic penalty from the in-core lithium has been shown to be small, in-core lithium may still be undesirable from the perspective of tritium production. This isotope is of particular concern due to its high mobility and its role as a biological hazard. In-core tritium can be produced through a number of pathways, including (1) as a fission product, (2) in control material through the ${}^{10}\text{B}(n,2\alpha)t$ reaction, (3) in helium through the ${}^3\text{He}(n,p)t$ reaction, and (4) in lithium through both the ${}^6\text{Li}(n,\alpha)t$ and ${}^7\text{Li}(n,n\alpha)t$ reactions. Pathways (1) and (2) will be present independent of ARC system inclusion, whereas pathways (3) and (4) will be enhanced by ARC system inclusion because neither He nor Li are normally present in significant quantities. Due to the low density of gases and the low relative abundance of ${}^3\text{He}$ (0.0002%), the main pathway for increased tritium production is likely the introduction of ARC system lithium. Since it is assumed that the lithium fluid will be significantly enriched in ${}^6\text{Li}$ and the ${}^7\text{Li}(n,n\alpha)t$ reaction is a threshold reaction above 2.5 MeV [17], the ${}^6\text{Li}(n,\alpha)t$ reaction will likely be the main contributor.

To quantify the increased tritium production, tritium inventories between BOEC and EOEC as calculated by Serpent are compared. This is done both with and without ARC system inclusion to isolate the tritium produced specifically by the ARC system. In both cases, the control system is fully withdrawn, so tritium produced in the control elements is not considered. Calculations including lithium employ the profile specified by Equation 5.14.

Tritium production in the reference core is found to be 2.075×10^4 Ci/yr, most of which results from the fission process, with fission yields of 0.0103% and 0.0142% for fissions on ${}^{238}\text{U}$ and ${}^{239}\text{Pu}$, respectively. The addition of the ARC system brings the tritium production level up to 2.117×10^4 Ci/yr, or a $\sim 2\%$ increase, although it is emphasized that this is an over-estimation, as it assumes a steady-state lithium profile independent of neutron flux.

To put this in context, NRC data shows that the average tritium *discharge* from a PWR in the US fleet (i.e. that which is released from the plant as liquid effluent) is 725 Ci/yr [37], most of which results from the addition of small amounts of lithium into the coolant for chemistry control. While the production in the B&B core is larger than the PWR

discharge, these numbers are not directly comparable because more tritium is produced in the PWR without being discharged, primarily within the fuel, and is thus not counted in the liquid effluent. Data from the Canadian Nuclear Safety Commission indicates that a typical CANDU will produce 6.5×10^5 Ci/yr mostly through the (n,γ) reaction on deuterium in the heavy water moderator, of which 8.9×10^3 Ci/yr will be released as liquid or gaseous effluents [72]. This puts the B&B tritium production level roughly 30 times lower than the CANDU, which is widely-deployed throughout the world. Studies of other potentially attractive advanced reactor systems have predicted the tritium production to be 1.7×10^6 Ci/yr for a 240 MW Fluoride-salt-cooled High-temperature Reactor (FHR) [123] and 9.6×10^3 Ci/yr for a 600 MW Very High Temperature Reactor (VHTR) [36], where the majority of tritium is produced from small concentrations of ^6Li in the coolant and graphite, respectively. On a per-megawatt basis, this puts the B&B core approximately 2.5 and 1000 times lower than the VHTR and FHR, respectively. Therefore it is concluded that the production of in-core tritium from the ARC system is minor.

Ex-Core Tritium Production

In addition to that which has diffused into the active core, a large source of lithium lies below the core in the lower reservoir. The net tritium production in the lower reservoir with time is determined through Equation 5.19, where ϕ is the neutron flux in the lower reservoir region, σ_{Li} is the tritium production cross-section from lithium, σ_T is the cross-section for tritium burnout, λ_T is the decay constant of tritium, N_T is the number density of tritium, and N_{Li} is the number density of lithium at 355°C according to Equation 5.18.

$$\frac{dN_T}{dt} = \phi\sigma_{Li}N_{Li}(t) - \lambda_T N_T(t) - \phi\sigma_T N_T(t) \quad (5.19)$$

Tritium burnout occurs through $(n,2n)$ and $(n,3n)$ reactions, both of which are threshold reactions above 8 MeV. Therefore they are conservatively neglected in this analysis. Because there is a large amount of lithium in the system and no lithium production, the number density of lithium is conservatively taken as a constant. Making these assumptions, Equation 5.19 can be solved for N_T as:

$$N_T(t) = \frac{\phi\sigma_{Li}N_{Li}}{\lambda_T} [1 - e^{-\lambda_T t}] \quad (5.20)$$

Due to uncertainties in the physical geometry of the lower reservoir, both ϕ and σ_{Li} are tallied by Serpent 2 immediately below the lower shield without including the lower reservoir in the model. Therefore these values do not account for self-shielding effects in the lithium layer and are conservative in the sense that the neutron flux will be over-predicted.

Equation 5.20 can be used to determine the concentration of tritium in each assembly at BOEC and EOEC. Taking the difference of these values, integrating over all assemblies, and converting to units of Ci/yr, Figure 5.37 shows the tritium production in the lower reservoir as a function of lower shield thickness. In each case, the lower reservoir ^6Li volume

is assumed to be equal to two times the volume of the ARC system tubes to account for the fact that during actuation some amount of absorber fluid will accumulate in the upper reservoir, requiring the lithium volume to be larger than the volume of the tubes.




Figure 5.37: Amount of tritium produced in the lower reservoir and lower shield as the lower shield thickness is increased, as well as the nominal amount produced in-core.

Figure 5.37 indicates that tritium production attenuates approximately exponentially with shield thickness. In order for tritium production in the lower reservoir to be below that of the active core, the thickness of the lower shield needs to be increased from its nominal value of 5 cm up to roughly 11 cm, although this would still lead to a doubling of the overall tritium production as compared to the case without ARC system inclusion. Regardless, even without the additional shielding, the yearly tritium generation is estimated to be at least 6 times lower than the previously-mentioned data for a typical CANDU reactor [72], and is thus expected to be acceptable from a regulatory perspective.

Also shown in Figure 5.37 is tritium production within the lower shield, determined in the same manner as was used for the lower reservoir but replacing variables in Equation 5.20 where appropriate. Because B_4C shield material can lead to tritium through $^{11}B(n,t)^9Be$ and $^{10}B(n,2\alpha)t$ reactions, it is important to ensure that the increased shield thickness does not negate the reduced tritium production in the lower reservoir. Figure 5.37 indicates that tritium production in the lower shield is always orders of magnitude below the in-core production, and is therefore a small price to pay for reducing tritium in the lower reservoir.

This section demonstrates that tritium production in ex-core ARC system components may be larger than that of in-core ARC components. This ex-core production can be effec-

tively managed by placing shielding between the active core and the ARC lower reservoir to meet the needs of the designer, however a tradeoff exists between tritium production and increased core length.

Depletion of Lithium Reservoir Inventory

With both diffusion and transmutation of lithium taking place, it is important to ensure that the lower reservoir maintains enough inventory to assure adequate actuation strength. This section investigates this concern.

Lower Reservoir Depletion Through Mass Diffusion

The amount of lithium leaving the lower reservoir can be determined according to Fick's first law:

$$J_{Li}(z) = -D \frac{dN_{Li}}{dz} \quad (5.21)$$

Under the assumption that the lithium profile is dominated by diffusion, the concentration gradient can be evaluated with Equation 5.14. Although this assumption was conservative for estimating the neutronic penalty and tritium production, it is not conservative for estimating the amount of lithium diffusing from the lower reservoir. Figure 5.35 shows that the flux depletion mechanism may play a larger role than mass diffusion in high fission-power regions. Because flux depletion reduces the lithium concentration, it will increase the diffusion of lithium into the active core. This may be accounted for by modifying the setup to have the infinite sink for lithium atoms at the core midplane rather than the top of the core, representing a neutron flux that is strong enough to remove all lithium atoms at this location. In this way, the flux of lithium atoms from the lower reservoir will be doubled.

To determine the absolute rate of lithium leaving the lower reservoir, the interfacial area needs to be specified. Figure 5.38 shows the rate of lithium loss from the lower reservoir due to mass diffusion as a function of the inner ARC tube diameter for a range of diffusion coefficients, demonstrating the need for an accurate diffusion coefficient. Unfortunately no experimental data could be found for lithium diffusion in either liquid potassium or gaseous helium. For diffusion between liquids, the diffusion coefficient can be estimated using the semi-empirical Wilke-Chang correlation, which has been found to predict diffusion coefficients to within roughly 20% for a variety of organic fluids, although experience with liquid metals is not well characterized. However, diffusion coefficients are generally on the order of 10^{-10} - 10^{-8} cm²/s for a wide variety of liquid-liquid diffusion systems [42].

To assess the significance of these results, it is necessary to determine how much lithium will be lost during an assembly's lifetime. For the core under consideration, each assembly is in-core for 12 cycles of 2.5 years. Using the highest diffusion rate from Figure 5.38 over the entire assembly residence time, 3.3×10^{19} atoms of lithium are predicted to diffuse out of the lower reservoir. Converted to volume at the temperature of the lower reservoir, this




Figure 5.38: Rate of diffusion of lithium atoms from the lower reservoir into the active core as a function of the radius of the inner ARC tube for different assumed diffusion coefficients assuming an infinite sink for Li atoms at mid-core. All curves assume three ARC tubes per assembly.

corresponds to $5.6 \times 10^{-4} \text{ cm}^3$. The minimum lithium volume stored in the lower reservoir can be estimated by the volume of the inner ARC tubes, calculated to be 1590 cm^3 for the case of a single assembly with three ARC tubes. From this comparison, the comparatively small volume of lithium that will be lost to diffusion is considered negligible.

Lower Reservoir Depletion Through Neutron Transmutation

Most of the lithium nuclei lost to transmutation will be split into nuclei of tritium and helium, both of which may take on gaseous forms. This study assumes that the gaseous nuclei generated are either kept in solution in the lower reservoir or accumulated in the inert-gas region of the ARC system such that they do not interfere with ARC system operation.

The amount of lithium lost to neutron transmutation can be determined in a manner similar to that which was used to quantify the ex-core tritium production in Section 5.5, except using the total transmutation cross-section as compared to the tritium production cross-section. The number of lithium atoms lost to transmutation over the course of an assembly's residence time is predicted in Figure 5.39. As the lower shield thickness is increased, the amount of lithium lost to transmutation decreases exponentially. With a 5 cm thick lower shield, the volume of lithium lost is nearly 7 cm^3 . This is nearly five orders

of magnitude more than the volume lost to mass diffusion, although this amount can be brought down by more than an order of magnitude by increasing the lower shield thickness by 10 cm. Regardless, even with only a 5 cm thick lower shield, the volume of lithium lost to transmutation is more than two orders of magnitude smaller than the estimate for the lower reservoir lithium inventory of 1590 cm^3 . This result implies that the shielding requirement between the core and the lower reservoir is likely less limited by the loss of lithium than it is by the production of tritium determined in Section 5.5.




Figure 5.39: Rate of lithium transmutation due to neutron flux and its associated volume in the lower reservoir as a function of the lower shield thickness.

Validation of Infinite Sink Assumption

In previous analyses, two different assumptions were employed, each conservative for the analysis to which it was applied:

1. Neutron flux depletion mechanism is less important than mass diffusion mechanism (Section 5.5)
2. Neutron flux depletion mechanism dominates mass diffusion mechanism at the core center (Section 5.5)

In Case (1), it was assumed that the upper reservoir could act as an infinite sink for lithium atoms. This assumed boundary condition needs to be validated. In Case (2), the boundary condition is set to give an upper bound estimate, and no validation is necessary.

Using Equations 5.14 and 5.21, the amount of lithium in the upper reservoir can be determined as a function of assembly residence time, as shown in Figure 5.40 for a range of upper reservoir volumes and the same setup used in Section 5.5 (i.e. three ARC tubes per assembly, $D = 10^{-8}$ cm²/s, and inner ARC tube radius $r = 0.75$ cm). Also depicted are the solubility limits of lithium in potassium at 510°C and 200°C (upper reservoir temperatures at standard operation and refueling, respectively) as determined by Equation 5.15.




Figure 5.40: Atom percent of lithium dissolved in upper reservoir potassium as a function of time for various upper reservoir volumes, along with the solubility limits at operating and refueling temperatures.

Figure 5.40 shows how the volume of the upper reservoir can make an impact on the validity of the assumption that the upper reservoir acts as an infinite sink. From previous studies [46], typical upper reservoir sizes are likely to be on the order of 500 cm³. This puts the amount of lithium dissolved in the upper reservoir more than 2 orders of magnitude below even the refueling temperature solubility limit, providing validation to the infinite sink assumption used in Section 5.5.

Summary of Other Impacts

This section has examined a number of potentially negative aspects of ARC system inclusion in a the reference B&B core. The neutronic penalty resulting from loss of fuel and addition of ARC system materials was evaluated to be negligible, with the loss of fuel material being most significant. The production of tritium within in-core ARC system materials was evaluated to be only a few percent of the amount produced in the nominal core, whereas the production within ex-core ARC system lithium was determined to be of possible concern. With regards to ex-core tritium production, it has been shown that the addition of shielding material between the core and the ARC system lower reservoir can mitigate this issue if a slight increase in core height, beyond that which is already required to accommodate the upper and lower reservoirs, is permitted. Finally, the amount of lithium lost from the lower reservoir to diffusion and transmutation was evaluated to be negligible in comparison to the amount stored in the lower reservoir. In light of these evaluations, the practicality of the ARC system with regards to neutronic performance has been validated.

5.6 Technology Gaps

Throughout this work, either due to unknown information or particularly complicated physics, a number of assumptions have been made. Some of these assumptions are vitally important to the conclusions drawn, while others have more of a secondary impact or are only important for particular reactor systems. This section outlines these assumptions and frames them in the context of future work which should be performed before ARC systems can be utilized. In the end, the technology gaps are rated with respect to their importance and the difficulty of resolving them.

- **Solubility of Reaction Products in ARC System Fluids:** As the ARC system operates, lithium will be transmuted to tritium and helium. Most of the current study has assumed that these reaction products will remain in solution of either the lithium or potassium fluids, but the actual behavior is poorly characterized. If tritium and helium do stay in solution, it is possible that they will impact the mutual solubilities of lithium and potassium, necessitating further study to ensure that mass diffusion of lithium does not increase. Otherwise if helium and tritium do not stay in solution, it is important to understand where they end up. The best case would be that they end up entirely in the gas reservoir, but they may also make their way into the expander fluid chamber of the upper reservoir, where a small gas bubble could form. This may have severe implications on the actuation characteristics of the ARC system that could invalidate its utility. Therefore this phenomenon needs to be studied in detail. Such a study should account for temperatures at shut down as well, when the reaction products would be most likely to precipitate out.

- **Diffusion Coefficients of ARC System Fluids:** The diffusion coefficients of lithium into potassium and vice-versa are poorly characterized. As outlined in Section 5.5, these parameters may have a large impact on the amount of lithium lost from the lower reservoir. Although the current analysis shows it is unlikely to be a cause for concern, better characterization of the system diffusion coefficients will allow for more definitive conclusions to be drawn about absorber fluid depletion.
- **Solubility of Lithium in Potassium at Elevated Temperatures:** The solubility values used in this study do not fully cover the range of temperatures expected under transient conditions. In the case that an ARC system is actuated for an extended period of time, an elevated solubility may allow for non-negligible amounts of lithium to diffuse into the expander fluid. While it is likely that excessive amounts of lithium dissolved in the expander fluid will fall out of solution when temperatures are again reduced, this mechanism is not fully understood. The assumption is that any precipitate will fall back into the bulk lithium inventory, but it is possible that some amount of precipitate will remain in other regions of the ARC system, where it may lead to a reactivity penalty, additional tritium production, and decreased ARC system worth. The first step to understanding this behavior is to better understand the solubility of lithium in potassium at higher temperatures.
- **Impact of ARC Tube Swelling on System Actuation:** As the ARC systems remain in the core, swelling is expected to take place in the ARC-tubes due to neutron irradiation. This swelling may reduce the diameters of both the inner and outer ARC-tubes and lead to different actuation characteristics compared to “fresh” ARC systems. It is important to account for this for multiple reasons. First, the small diameters of the ARC-tubes raises the possibility for capillary forces to interfere with ARC system dynamics. Second, the reduced ARC-tube diameters resulting from swelling will likely lead to faster actuation. Because it was shown in the transient studies of this work that extremely rapid actuations may lead to oscillatory behavior, it is important to understand the swelling behavior so that oscillations can be avoided.
- **Loading Fluids into ARC System:** While the procedure for assembling an ARC system has been extensively explored, the procedure for loading the ARC fluids into the system during assembly is still not entirely clear. Because both lithium and potassium are solid at room temperature, they may be loaded in solid form and allowed to stratify once melted. However the inert gas must be loaded at pressure, which may be complicated by the need to perform accurate welds once the gas has been loaded. Additionally, while the fluids will remain properly stratified once they are in the proper configuration, it may be difficult to ensure that none of the inert gas or absorber fluid sit on top of the expander fluid within the upper reservoir when they are initially loaded. This concern may require the fuel assembly to be first flipped upside down while the ARC fluids are in a liquid state to allow the upper reservoir to be emptied of inert gas, and then turned right-side-up in a controlled fashion so that the inert

gas and absorber fluid stratify on top of the expander fluid in the outer ARC tube and gas reservoir while the upper reservoir remains entirely filled with expander fluid. Ultimately this is not viewed as a barrier to realization of the ARC system concept, but the current study has not devised a solution to this particular issue.

- **Impact of ARC System Inclusion on Fuel Recycling:** Because ARC systems may be desired in a wide variety of SFR cores, the prospect of recycling fuel in assemblies that contain ARC systems should be investigated. Most recycling processes begin with some form of chopping and dissolution of each assembly, typically performed in open air. Typically this would be preceded by a thorough scrubbing of each assembly to remove sodium from the exterior due to its reactivity with air and water. For assemblies with ARC systems included, large reservoirs of potassium, which is also reactive with air and water, now must be dealt with during the chopping and dissolution phases. This may complicate the process and may necessitate the draining of each ARC system prior to reprocessing.
- **Use of Point Kinetics:** In most SFR cores, the use of point kinetics is justified due to the tight neutronic coupling between different regions of the core. However, for very large cores with many blanket assemblies, point kinetics may not be as well suited. This study assumed that point kinetics was a valid approximation for the B&B core, but it is possible that some transients, such as a localized flow-blockage, may induce significant asymmetries into the response. In cases such as these, ARC systems may not actuate uniformly across the core, and spatial dynamics of the transient may be important to accurate predictions. Another especially important case would be the breach of an ARC system leading to localized reactivity transients, which have the potential to induce spatial reactivity oscillations. This has not been investigated in this work.
- **Simultaneous Actuation Assumption:** In the current analysis, it was assumed that all ARC systems actuate simultaneously in response to a core heat-up. However, because each assembly will have a somewhat different outlet temperature, the actuation of each ARC system will not be simultaneous for a number of reasons. First, the absorber level within each ARC system will be different at the beginning of any actuation due to the steady-state variation in outlet temperatures. Second, because some transients may be more spatially-localized than others, it is possible that a heat-up in one section of the core may lead to the actuation of some systems much more rapidly than others that are further away from the initiating event. Finally, because some assemblies may be very over-cooled during steady-state operation, it is possible that these assemblies may not actuate nearly as strongly as others with a lower power-to-flow ratio. These uncertainties should be quantified in order to fully understand the dynamics of any transient with the ARC system included.

- **Impact of Lower Reservoir Temperature on Actuation:** During both the ULOF and UTOP transients, the inlet temperature does not change appreciably. However the defining characteristic of the ULOHS is a change in the inlet coolant temperature, some of which will be communicated through conductive and convective heating to the lower reservoir. As such, the expander fluid located in the lower reservoir will thermally expand, in addition to the expansion taking place in the upper reservoir. This thermal expansion has not been quantified or accounted for in the transient simulations performed thusfar, meaning the transient simulations likely underestimate the speed of ARC system actuation during the ULOHS scenario. While this may be the case, it is not expected that this will have important consequences due to the fact that the ULOHS was consistently the least challenging transient. Additionally, the kinetic behavior during the ULOHS exhibited no signs of oscillatory or unstable behavior even in very aggressive ARC system designs, and therefore it is anticipated that a slightly faster actuation is unlikely to induce oscillatory effects.
- **Impact of ARC Tube Axial Temperature Distribution:** Similar to the previous point, the axial temperature distribution of the ARC system is not accounted for in any of the transient simulations. This will have an impact on multiple aspects, but most importantly the density of the absorber fluid and the speed of actuation. Because the absorber fluid starts below the core, it will be at the temperature of the inlet coolant. However as a transient progresses, the absorber fluid will enter the active core where it will heat up. In rapid transients, this heating will likely be slow in comparison to the speed at which the fluid enters the core, and so impacts are likely to be small in the initial phase of the transient. In slower transients, however, the rate of heat transfer to the absorber fluid may be large enough that thermal expansion of the absorber fluid should be accounted for in determining the level of the fluid in the core. Additionally, during transients that terminate with elevated core temperatures, the heat up of the absorber fluid will cause it to have a lower density, and thus its effectiveness as a neutron poison will diminish. Depending on how high the core temperatures get, this effect may be more or less pronounced.

To summarize, a number of gaps exist concerning transport and thermodynamic properties of the ARC system fluids, which are particularly important if the ARC system is to be included in long-lived cores. These gaps may only be resolved through experimental measures and will likely require significant resources. A separate category of gaps may be resolved through more detailed modeling, including the impacts of neglecting axial temperature gradients within the ARC system fluids and the use of point kinetics. These gaps may be resolved without any extra experimentation, although some will require extensive modeling efforts. Another category deals with aspects which require engineering and process design. These gaps, including the loading of system fluids and the impact on fuel recycling, will require additional analysis similar to that performed in this project to determine the ARC system manufacturing process.

For each of the gaps, Table 5.4 provides a preliminary rating to characterize two aspects. First, the importance of the gap is rated on a scale of 1-3, with 1 signifying very high priority and 3 signifying relative unimportance. Second, the gaps are rated with respect to the amount of effort it will take to resolve them, with 1 signifying a gap that will be relatively straightforward to resolve and 3 signifying a gap that will require extensive effort. Therefore, a gap with ratings of 1 in both categories will be of both high importance and low effort, and a gap with ratings of 3 in each category will be of great difficulty and minor significance. These ratings may be used to determine directions for future work.

Table 5.4: Technology gaps rated according to their importance and ease of resolution.

Gap	Importance	Difficulty
Solubility of reaction products	1	3
Diffusion coefficients	3	3
Solubility at elevated temperatures	2	3
ARC tube swelling	2	1
Loading of system fluids	1	1
Impact on fuel recycling	3	2
Use of point kinetics	2	2
Simultaneous actuation	2	2
Lower reservoir temperature	3	3
ARC tube axial temperature distribution	3	2

5.7 Summary of ARC System Performance in the B&B Core

This chapter has investigated the use of the novel Autonomous Reactivity Control (ARC) system to improve the transient performance of the reference B&B core. After demonstrating the lacking passive safety performance of the B&B core in Chapter 4, this chapter begins by briefly introducing the principles and design of the classical ARC system. Then, after discussing the techniques used to incorporate the unique system response into the SAS transient models, a series of parametric studies are performed over the ARC system design variables for each of the ULOHS, UTOP, and ULOF accident scenarios. The ULOHS is found to be easily improved through ARC systems with a wide design space, with those having the strongest absolute worth providing the most additional margin to coolant boiling and fuel melting. While the UTOP scenario initially displayed diverging oscillations induced by the ARC system, these were shown to be diminished by making the ARC system less sensitive to rapid fluctuations in coolant temperature. Such a method enabled inherent passive safety to

be achieved in the UTOP scenario, while not significantly degrading the performance of the ULOHS. However, the passive safety in the ULOF scenario remained unattainable for two reasons: (1) if the ARC system were made less sensitive in the same manner as the UTOP, the rapid progression caused by the transition from forced- to natural-circulation could not be mitigated, or (2) if the ARC system were allowed to remain highly sensitive, it would lead to oscillations during the initial pump coast down phase, in addition to being incompatible with achieving inherent safety in the UTOP.

This difficulty in simultaneously achieving aggressive and sensitive actuations while avoiding diverging oscillations led to the development of the *modified* ARC system, wherein a pseudo-one-way valve is incorporated into the ARC system to allow rapid forward actuations but retarded reverse actuations. In order for this valve to properly function, it was demonstrated that a second cover gas needs to be added to the upper reservoir. Designs incorporating these two changes were then developed, and possible options for the valve were suggested along with a discussion of their limitations. To assess the potential for this modified ARC system, new computational methods were developed to explicitly simulate the internal ARC system fluid dynamics with SAM and couple them to the whole-core transient simulations in SAS. Using these new capabilities, parametric studies over the upper reservoir design and valve loss coefficients were performed, in addition to sensitivity studies on the boundary conditions. It was shown that the modified ARC system actuation is strongly sensitive to the cover gas volume in the upper reservoir, with larger gas volumes leading to impeded actuations both in terms of speed and strength, driving designs to have smaller gas volumes in general. Using this modified ARC system, inherent passive safety in the ULOF scenario was finally obtained, although this result is sensitive to the boundary conditions assumed on the reactor's steam generators. While the modified ARC system was shown to enable passive safety in all three transient scenarios, the actuation sequence and ARC system design are both complicated significantly as compared to the original ARC system, and the new system needs to be more fully assessed for new failure modes.

After demonstrating the potential benefits to transient performance, the ARC system is then evaluated from the perspective of its impacts during steady-state reactor operations. In general, the reactivity impact of ARC system inclusion was found to be small-to-negligible. Similarly, the amount of lithium lost from the lower reservoir through transmutation reactions and diffusion into the expander fluid were evaluated to be negligible. While the ARC system was found to lead to additional tritium production from transmutation of the lithium absorber fluid, this amount has been quantified as relatively small in comparison to the amount already being produced in the reference core, and especially in comparison to other reactor systems such as the CANDU. Further, the tritium production in the lower reservoir can be potentially mitigated through the incorporation of small amounts of shielding material between it and the active core, with little impact on reactor characteristics. Finally, the chapter concludes by discussing various important assumptions or potential technology gaps that may have an impact on the viability of ARC system implementation. Many questions related to chemistry and thermodynamics of the system fluids are of vital importance and need to be studied further. Other questions related to the modeling assumptions and

techniques used in this dissertation are pointed out in that they may enable a higher-fidelity understanding of system behavior and may highlight specific phenomenon not able to be simulated using the current techniques, but these are not expected to change the general conclusions of this study.

Chapter 6

Material Damage and B&B Reactors

Virtually no study of B&B cores would be complete without discussion of material damage. This chapter is devoted to assessing the current state of B&B core design in relation to material damage constraints, gaining a deeper understanding of how accurately we are able to quantify this important concern, and proposing methods which may enable bypassing the material damage constraints while still enjoying some of the benefits of the breed-and-burn fuel cycle. In the end, two new fuel cycles are proposed and investigated in detail to assess their benefits as compared to current LWR and classical B&B fuel cycles.

6.1 Review of B&B Core Design Strategies with a Focus on DPA

As briefly discussed in Chapter 1, it has long been recognized that one of the primary infeasibilities of B&B reactor systems is the necessarily high neutron fluences and associated damage to structural materials [43], and thus much effort has been expended on overcoming the issue. To do so, there are two possible options: (1) extend the experimental database for structural materials beyond the current limit of 200 DPA, or (2) alter B&B design strategies to reduce material doses to within the current database.

Extending the experimental database has three primary issues that make it fundamentally difficult to resolve: (1) the United States (and in fact the western world) does not have the experimental facilities to irradiate materials in a fast flux¹, (2) even if the capability existed, it would take many years to accumulate the high doses that are necessary, and (3) there is no guarantee that the outcome of such experimental campaigns would be successful from the perspective of enabling B&B operation as it is currently understood. Therefore nearly all effort from the research community has been along the route of altering our design strategies to enable breed-and-burn operation with lower required material doses, with some exceptions

¹It was discussed in Chapter 1 that the United States is currently in the early phases of designing a new experimental SFR facility that would satisfy this need.

including the use of ion-beam experiments to simulate material damage at accelerated rates [48] or performing material irradiations in Russian SFRs. The ability to achieve material damage below current constraints is hampered, however, by working within historical core-design paradigms so as not to introduce other infeasibilities in the process of solving those related to material damage. This section outlines the myriad design strategies which have been investigated so far to varying degrees of success.

Pseudo-Classical Strategies

Traditional methods involve using highly-optimized, but conventional fuel shuffling strategies and aggressive assembly designs to reduce fuel leakage and keep the spectrum as hard as possible (this is the method employed for the reference core outlined in Chapter 2). The underlying notion behind these methods is that breeding and doubling times are improved as the spectrum hardens, and better breeding enables shorter cycle lengths. Studies employing these classical methods have looked into a wide variety of fuel/cladding/coolant combinations, including sodium, lead, and gas coolants in addition to metallic, oxide, nitride, and carbide fuel forms (see, for instance, [54, 114]). It has consistently been found that the combination of sodium coolant, HT9 cladding, and zirconium-alloy metallic fuel provides the smallest minimum required burnup while remaining within the accepted design bounds². By eliminating the gas plenum to enable a very tight lattice pitch and using an optimized radial shuffling scheme, this type of method has resulted in damage levels as low as ~ 530 DPA [65]. Further gains through such conventional methods are expected to be limited.

Hou et al. have extended this notion to include 3D shuffling schemes, where fuel may be moved axially as well as radially throughout the core [65]. Whereas low-leakage B&B designs with conventional shuffling schemes result in very high axial power peaking and low utilization of fuel at the axial periphery, 3D shuffling aims to reduce axial peaking levels while improving the average fuel utilization, thus decreasing peak cladding damage. Although such shuffling schemes inherently increase neutron leakage, this method has proven to be highly successful, leading to minimum required peak cladding damages as low as ~ 350 DPA. The technical feasibility of such 3D shuffling appears promising; however, the practicality may be questioned as it introduces a number of complications into a standard SFR system, such as requiring more than double the amount of shuffling operations at every refueling. Even still, further gains from 3D shuffling are expected to be limited, and thus there is still a long way to go before reaching the 200 DPA constraint.

Incorporation of Moderating Material

Osborne and Deinert have suggested the addition of moderating material to soften the spectrum [103]. The motivation for this is derived from the observation that the DPA cross-

²For instance, it is possible to achieve lower minimum required burnups by utilizing unalloyed metallic fuel, but concerns of fuel-cladding chemical interactions and fuel swelling, coupled with the limited available data for lower alloys, dictate that U-10Zr fuel should be used.

section increases roughly exponentially at high energies (see Figure 6.1 later in this chapter). Thus, the suggestion is that softening the spectrum only slightly may reduce DPA while still allowing for sufficiently high breeding. They demonstrate this idea on a long-lived SFR which differs in a number of ways from the classical B&B design, including the use fuel enriched up to 15w% with plutonium and the absence of shuffling. By (1) doping the fuel with beryllium, (2) including beryllium in the reflectors, and (3) replacing the region of peak cladding damage with a diluent, they were able to reduce peak cladding damage in a metallic-fueled core from 273 to 230 DPA, although no investigation was performed to determine which of the three design changes drove this reduction. The study showed that core reactivity was adversely impacted, though the use of plutonium-enriched fuel was able to compensate for this. Follow-up studies by Kuwagaki et al. have shown that doping the fuel in a conventional B&B design with beryllium does not lead to a reduction in peak cladding damage [83]. While the effective DPA cross-section is reduced through spectral softening, the negative impact on reactivity dominates and actually leads to higher cladding damage requirements. Therefore, the use of fuel uniformly doped with a moderator is not viewed as an effective method of reducing cladding damage in breed-and-burn systems.

Similar studies by Kuwagaki et al. have attempted to divide the core into two separate regions which are separated by a row of moderator assemblies [82]. This strategy is motivated by the realization that excess neutron production can be maximized by fissioning in the fast spectrum, while absorption cross-sections are conversely maximized in the thermal spectrum. Other researchers have come to this conclusion and proposed moderating neutrons leaking from the fast-spectrum region into the thermal-spectrum region by incorporating a moderating zone, but its actual implementation into a realistic core design is difficult from multiple perspectives. First, it is difficult to moderate neutrons sufficiently on their way from the fast- to the thermal-region without devoting excessive core volume to moderator assemblies, which produce no power and parasitically absorb neutrons. The addition of moderator assemblies also locally taints the spectrum in the fast-region, degrading the ability to fast fission. Finally, the fact that fission cross-sections at thermal energies are many times higher than those at fast energies means that the thermal region must be limited to low flux levels as compared to the fast region to avoid the production of power peaks in thermalized assemblies. This ultimately limits the usefulness of the thermalized zone because the low flux levels make it difficult to breed sufficient fissile material. This combination of inherent limitations makes the realization of such a strategy challenging, and the author knows of no instances where it has been achieved to any real degree of success.

Studies out of the Karlsruhe Institute of Technology by Chen et al. have investigated the possibility of using very low density water as coolant in a travelling-wave B&B core, with the implicit assumption that such a coolant would enable reduced material damage because it would soften the spectrum [26]. They show that breed-and-burn operation can in fact be achieved in their design, which utilizes oxide fuel of 100% theoretical density (TD). However no actual quantification of material damage is provided. Because the neutron spectrum is still quite hard, the fuel residence time is very long, and the design uses the traveling-wave strategy, the claims of reduced material damage are difficult to believe without

being substantiated. In addition, the feasibility of such a design is also questionable due to the assumption of oxide fuel with 100% TD. In reality such high density fuel is not manufacturable and will suffer from intense Fuel-Cladding Mechanical Interaction (FCMI), which will force a realistic design to implement fuel of roughly 80-85% TD. Oxide fuels with such low densities have been shown by other researchers to be incapable of breed-and-burn operation in the fast spectrum [54], and therefore it is questionable whether the proposed design holds merit.

Incorporation of Limited Separations/Recycling

Because the primary benefits of breed-and-burn operation result from the elimination of enrichment and recycling from the fuel cycle, researchers have generally avoided using them to improve B&B core performance. To reduce the costs of fuel recycling while still realizing some of the benefits, some researchers have proposed using a limited form of fuel ‘conditioning’ (see, for instance, [56, 75]). In one version of reconditioning, discharged fuel is processed to replace the cladding without chemically processing the fuel itself. Before the replacement with new cladding, the volatile fission products are allowed to de-gas, allowing for a reduction of parasitic absorptions and a relief of fuel pin pressurization, and enabling more aggressive fuel pin designs and extended neutronic lifetime of the fuel. Most importantly, though, this method allows for the cladding to be periodically replaced when it reaches its damage limit, allowing for the theoretical maximum burnup to be approached while effectively side-stepping the issue of cladding damage. While the technical feasibility of the suggested form of fuel reconditioning can probably be established based on extrapolating the experience with the fuel cycle facility of EBR-II [58], its acceptability from economic, environmental, and proliferation perspectives is yet to be proven. In particular, this method may be viewed as similarly challenged from a political perspective due to the fact that the cladding is breached at all.

Studies by Zhang have further generalized the B&B concept by coupling a *portion* of the core to a recycling scheme in the Seed-and-Blanket (S&B) design [156]. These studies have shown the possibility of achieving fuel utilization comparable to B&B designs while satisfying fuel damage constraints. At the same time, the S&B core offers improvements over conventional SFR designs with *full* recycling by significantly reducing the volume of fuel that must be processed. This concept holds great promise for realizing the benefits of a B&B fuel cycle, at least partially, but relies on fuel recycling capabilities which are currently unavailable in the United States. Therefore, while the S&B concept avoids the material damage infeasibility which plagues pure B&B designs, it introduces another issue that, while not technical in nature, still poses a significant challenge to realization.

Di Sanzo has investigated the possibility of increasing burnup in B&B cores through various fuel processing techniques, although in this case the investigations were not done with the aim of satisfying the cladding DPA constraint [29]. Instead, di Sanzo investigates techniques meant to push fuel burnup beyond the limits imposed by either reactivity or material damage concerns under the assumption that current materials can withstand the

~ 500 DPA required by a typical B&B core design (i.e. the classical breed-and-burn operation is assumed possible, and now we want to push the burnup even further). Such methods include: (1) double-cladding, where fuel is removed after 20%FIMA and a second layer of cladding is added to each fuel pin so that it may maintain its structural integrity for another round of burnup in a second-tier B&B core; (2) limited-separations reprocessing schemes where fuel is de-clad and vented of gaseous fission products before being re-clad and put back into the reactor for additional burnup; and (3) recycling discharged B&B fuel into oxide form for further use in an LWR to burn down the remaining plutonium content. None of these methods are able to (or even attempt to) reduce the material damage, but they are mentioned because they are closely related to methods investigated subsequently in Section 6.5 of this thesis. Of particular relevance is the recognition that discharged B&B fuel may contain significant residual fissile material which is of value if the right methods can be devised to utilize it.

Alternatives to Traditional Fuel Assemblies

With the notion that essentially all reactor types of the classical design paradigm have been exhausted, some researchers have begun investigating less traditional core designs, including those with pebble or liquid fuels. Greenspan et al. did preliminary investigations into the design of a pebble-bed B&B core, which would utilize metallic fuel pebbles clad in HT9 that continuously circulate throughout the core [44]. The idea behind this is to distribute the burnup such that the burnup peaks could be eliminated, similar to the motivations behind 3D shuffling of fuel assemblies. This concept suffered from the fact that the buildup of fission gases will lead to unacceptably high pressures inside of the pebbles when they are pushed to the burnups required for breed-and-burn operation, which can only be overcome by reducing the smear density to untenable levels. In addition, the pressure drop induced by flowing coolant through a pebble bed of reasonable power density that is tall enough to sustain breed-and-burn operation would be unrealistically high. For these reasons, the pebble-bed B&B core was abandoned in favor of the 3D shuffled core mentioned previously.

Two different general variations of the liquid-fuel version of the B&B concept have been proposed. The first is closely related to the classical Molten Salt Reactor (MSR), however these versions eliminate the moderator to enable a fast spectrum. Such designs typically employ chloride salts in order to avoid moderation in other salts such as FLiBe or FLiNaK, and are thus called Molten Chloride Fast Reactors (MCFRs). These systems have been investigated by a number of researchers (see, for instance, [63, 90, 13]), and are closely related to the second reactor concept being pursued by TerraPower, LLC [1]. MCFRs offer the potential advantages of strongly-negative reactivity feedbacks, low excess reactivity, and the complete elimination of problems associated with solid fuels (i.e. FCMI, fission gas buildup, degradation of mechanical properties, etc.). The general conclusion is that MCFR are feasible using the U-Pu cycle if the chloride salt is enriched in ^{37}Cl and the core volume is somewhat large in comparison to other MSR systems. However, significant uncertainties must be overcome before commercial implementation may be realistic, particularly having

to do with corrosion chemistry and solubility of fission products at the burnups required for breed-and-burn operation. These uncertainties will require research efforts on par with those required by the material damage challenges of traditional solid-fuel B&B. Additionally, the material damage on the core barrel is expected to be very high in MCFR systems due to the fact that it is difficult to shield core components from the flux of liquid fuel as it flows throughout the reactor (at least in single-fluid designs). This is challenging not only from the perspective of DPA, but also from the buildup of helium expected to take place in the nickel alloys usually chosen in molten-salt systems as a result of their corrosion resistance. Therefore while MCFR-type strategies are an alternative to the classical B&B design which is feasible from neutronic and thermal-hydraulics perspectives, they as well suffers from a lack of experience with the necessary materials.

An alternative liquid-fuel system has been proposed by Moltex Energy and investigated by Kasam which utilizes two separate molten salt fluids in the core [76]. The first fluid holds dissolved fuel and is contained in ‘assemblies’ located in the active core, while the other is fuel-free and acts as the heat transport system. Such a separation of fuel and coolant offers many advantages over traditional MSR systems, and in the case of the Moltex design, does not require two fluid circulation systems because the fuel salt does not circulate outside of the active core. Advantages over single-fluid designs include the ability to shuffle fuel assemblies throughout the core to best promote breeding, the ability to shield important structural components from radiation damage or replace them upon excessive fluence, and the elimination of radioactivity from the ex-core primary loop. It has been found that such a design is likely feasible from the perspective of neutronic performance, although this design in fact suffers from the same uncertainties in molten salt material properties that hamper the MCFR design. In general this concept appears promising but has seen less much less development from the design community thusfar, and therefore its viability is not as clearly defined.

Summary of Alternative Design Strategies

The appeal of the breed-and-burn fuel cycle has led to many efforts to devise a system which is feasible based on currently-understood material constraints. Essentially all reactor design paradigms have been investigated at this point, in addition to investigations into both fast and thermal spectra. From all of these efforts, no design has yet been found which enables the benefits of the breed-and-burn fuel cycle while remaining within the bounds of current technology, and therefore new B&B design strategies are still a topic of active research. Later in this chapter, two new strategies will be identified to enable some of the benefits from the breed-and-burn fuel cycle in the short-term.

6.2 Nuclear Data Uncertainties in Quantified DPA

To support a deeper understanding of how well previous studies have succeeded in reducing material damage in B&B cores, it is important to understand the uncertainties associated with DPA predictions. To date, such an uncertainty quantification has not been performed. This section develops methods for quantifying such uncertainties and applies them to the reference core to understand for the first time the impacts that nuclear data uncertainties have on predictions of material damage.

Introduction

Two common measures to quantify material damage and predict fuel performance under irradiation in fast reactors are Displacements Per Atom (DPA) and fast fluence (i.e., fluence of neutrons with energy greater than 0.1 MeV). Fast fluence is widely employed as a measure of material damage (see, for instance, [80] or [4]), nevertheless such metric is not capable to discern between different fast neutron spectra due to its underlying assumption that neutron damage is proportional only to the number of fast neutrons, independent of the neutron spectrum within the fast region. On the other hand, DPA can be used as a spectrum-sensitive correlation parameter to predicted performance of systems based upon data measured under similar, but different, spectra or with different particles, such as ion beams [20]. For this reason, DPA is likely a more reliable measure of radiation damage to be used in feasibility studies of fast reactors. However, the use of DPA as a correlation parameter is predicated upon an accurate prediction of the flux spectrum. For small design changes meant to improve material damage, it is important to ensure that the predicted improvements are larger than the uncertainties associated with those predictions; therefore, in order to truly understand the benefits of novel design modifications as compared to a reference design, the uncertainties that contribute to inaccuracies in predicted DPA must be quantified. Once the uncertainties are understood, the impacts of design changes on material damage can be more completely assessed.

This section examines uncertainties in predicted DPA that result from uncertainties in cross-section data. This is accomplished by using Generalized Perturbation Theory (GPT) and covariance data to quantify uncertainties in the predicted neutron spectrum, the radial flux profile, and the core flux level while assuming that the cycle length remains unchanged. Using fine-group DPA cross-sections which quantify the number of displacements generated in a given material as a function of incoming neutron energy, uncertainties are then propagated through the calculation of effective one-group DPA cross-sections and accumulated DPA on a batch-averaged level. The suggested methodology is then applied to the reference core outlined in Chapter 2.

DPA Theory and Methodology

DPA is not a quantity that can be readily measured, and as such its quantification relies upon a number of theoretical arguments and computations. Different methods for calculating DPA have been proposed. This study uses the method endorsed by ASTM for low-alloy steels [12]. Whereas uncertainties in most other cross-sections non-linearly impact the results of neutronics calculations, the ASTM method uses a DPA cross-section that is essentially a correlation parameter and does not feedback into the calculation of the neutron spectrum or other neutronic parameters. Therefore, for sufficiently similar systems (i.e. systems with broadly similar neutron spectra as might be indicated by, for instance, the phrases “fast spectrum” or “thermal spectrum”), uncertainties in the fine-group DPA cross-section are not necessarily important as long as all analysts use the same values and the results obtained using such values are able to be meaningfully related to quantities of interest (i.e., material lifetimes, swelling rates, degradation of mechanical properties, etc.). Such methodology does not necessarily ensure accurate quantification of the underlying phenomena leading to DPA accumulation, but rather ensures that all analysts are speaking a common language when using DPA as a correlation parameter to predict material properties under irradiation. In the ASTM method, the fine-group DPA cross-section is provided as data, depicted in Figure 6.1. This fine-group data is generated using the NRT (Norgett-Robinson-Torrens) model to convert damage-energy to number of displacements with a displacement threshold of 40 eV and an efficiency of 80% [98]. Although it is well documented that the NRT model has limitations related to the replacement of displaced atoms into lattice vacancies during the recombination phase of a displacement event, its widespread use by researchers throughout the nuclear community motivates its use in the current study [97].

Regardless, the DPA cross-section increases nearly exponentially at high energies, implying that small changes in the fast portion of the predicted neutron spectrum may have a large impact on the effective one-group DPA cross-section value. For such studies to be reliable, accurate predictions of the spectrum-averaged DPA cross-section are necessary. This observation, along with the relatively high uncertainties in nuclear data at high energy in general, is the basis for performing the present sensitivity and uncertainty study.

In a reactor core where fuel is subdivided into multiple batches (i.e. groups of fuel assemblies loaded at the same time and managed similarly throughout their in-core lifetimes), the average accumulated DPA in a fuel assembly can be calculated as in Equation 6.1, where w is the effective one-group DPA cross-section, \hat{w} is the group-wise DPA cross-section, ϕ is the total reactor flux, $\hat{\phi}_\ell = \phi_\ell/\phi$ is the ratio of total flux in batch ℓ to reactor total flux, $\hat{\phi}_\ell^i = \phi_\ell^i/\phi_\ell$ is the fraction of batch ℓ 's flux in energy group i , t_ℓ is the time spent in batch ℓ , n_b is the number of fuel batches in the core (i.e., number of cycles), and n_r is the number of response energy groups.

$$DPA = \sum_{\ell=1}^{n_b} \phi_\ell t_\ell w_\ell = \phi \sum_{\ell=1}^{n_b} \left(\hat{\phi}_\ell t_\ell \sum_{i=1}^{n_r} (\hat{w}^i \hat{\phi}_\ell^i) \right) \quad (6.1)$$




Figure 6.1: DPA cross-section for low-alloy steels as a function of energy.

If \hat{w}^i and t_ℓ in Equation 6.1 are taken as completely certain, uncertainties in the accumulated DPA result from uncertainties in the neutron spectrum ($\hat{\phi}_\ell^i$), the batch-wise flux level ($\hat{\phi}_\ell$), and the total flux level (ϕ). The following subsections outline how these uncertainties are treated.

Uncertainties in Neutron Spectrum

Determination of spectrum uncertainties is accomplished using the GPT capabilities in Serpent [14], which employs a collision-history approach to estimate sensitivities, combined with a set of covariance matrices. GPT, as compared to Standard Perturbation Theory (SPT), refers to the ability to calculate sensitivities to responses other than only k_{eff} , such as ratios of linear and bilinear functions of the forward and adjoint fluxes. As will be shown, the DPA can be viewed as a ratio of linear functions of the forward flux. Additionally, studies by Hou et al. [67] have highlighted how large spectral variations across the core, owing to both strong location and burnup dependencies, require special care in generating group constants for deterministic transport studies of B&B cores, motivating the use of Monte Carlo transport. Therefore, the GPT capabilities within Serpent are especially suited to the problem at hand. Using GPT, Serpent calculates sensitivity responses for a ratio of fluxes as in Equation 6.2, where E_i is some slice in energy space and b_ℓ is a given batch of fuel.

$$r_{i,\ell} = \frac{\int_{b_\ell} \int_{E_i} \phi(r, E) dE dr}{\int_{b_\ell} \int_0^\infty \phi(r, E) dE dr} \quad (6.2)$$

By using a series of response functions with different integration bounds, sensitivity coefficients over the entire neutron spectrum can be determined for each batch of fuel. However, because $\hat{\phi}_\ell^i$ is a relative quantity, it is subject to the normalization condition $\sum_i \hat{\phi}_\ell^i = 1$, which implies $\sum_i d\hat{\phi}_\ell^i/dx = 0$, where x is any arbitrary perturbation parameter. This constraint is unknown to Serpent and is accounted for using Algorithm 1 as a post-processing step, where n_p is the number of perturbation energy groups and x represents any perturbed cross-section.

```

obtain  $\hat{\phi}_\ell^i$  and  $\frac{d\hat{\phi}_\ell^i}{dx_j}$  from Serpent results;
for  $i = 1:n_r$  do
     $\hat{\phi}_\ell^i = \hat{\phi}_\ell^i + \sum_{j=1}^{n_p} \frac{d\hat{\phi}_\ell^i}{dx_j} \Delta x_j : \Delta x_j \rightarrow 0$ 
end
for  $i = 1:n_r$  do
     $\hat{\phi}_\ell^i = \hat{\phi}_\ell^i / \sum_{i=1}^{n_r} \hat{\phi}_\ell^i ;$ 
     $\frac{d\hat{\phi}_\ell^i}{dx} = (\hat{\phi}_\ell^i - \hat{\phi}_\ell^i) / \Delta x : \Delta x \rightarrow 0;$ 
     $\frac{d\hat{\phi}_\ell^i}{dx_j} = \frac{d\hat{\phi}_\ell^i}{dx_j} * \frac{d\hat{\phi}_\ell^i}{dx} / \frac{d\hat{\phi}_\ell^i}{dx} ;$ 
end
Result:  $\frac{d\hat{\phi}_\ell^i}{dx_j} \forall (i, j)$ 

```

Algorithm 1: Constraining sensitivities such that $\sum_i \hat{\phi}_\ell^i = 1$.

If the sensitivity of $\hat{\phi}_\ell^i$ to a perturbation in perturbation-group k of cross-section j is written as $s_{j,k}^{\hat{\phi}_\ell^i}$, the sensitivity vector for a given response is denoted as $\mathbf{S}_j^{\hat{\phi}_\ell^i} = [s_{j,1}^{\hat{\phi}_\ell^i}, s_{j,2}^{\hat{\phi}_\ell^i}, \dots, s_{j,n_p}^{\hat{\phi}_\ell^i}]^\top$. Provided with a covariance matrix $\mathbf{C}_{(j,j')}$ of the same perturbation-group discretization, the variance of $\hat{\phi}_\ell^i$ for the reaction pair (j, j') can be calculated to first-order with Equation 6.3 (the ‘sandwich rule’), resulting in n_r uncertainties for each reaction pair in each fuel batch.

$$(\sigma_{(j,j')}^{\hat{\phi}_\ell^i})^2 = (\mathbf{S}_j^{\hat{\phi}_\ell^i})^\top \mathbf{C}_{(j,j')} \mathbf{S}_{j'}^{\hat{\phi}_\ell^i} \quad (6.3)$$

Uncertainty in Batch-wise Flux Levels

Nuclear data uncertainties may impact the flux profile such that the relative flux within each batch is uncertain. This is quantified again using Serpent GPT capabilities to tally the sensitivity on the ratio of fluxes shown in Equation 6.4, where Ω indicates all fuel batches.

$$r_\ell = \frac{\int_{b_\ell} \int_0^\infty \phi(r, E) dE dr}{\int_\Omega \int_0^\infty \phi(r, E) dE dr} \quad (6.4)$$

Similar to before, the resulting sensitivity coefficients should satisfy the normalization condition $\sum_{\ell} \hat{\phi}_{\ell} = 1$, and therefore another normalization algorithm similar to Algorithm 1 is employed. Once the normalized coefficients are obtained, the sandwich rule is again applied to determine the uncertainty, $\sigma_{(j,j')}^{\hat{\phi}_{\ell}}$, associated with a given reaction pair, (j, j') .

Uncertainty in Absolute Flux Levels

The normalization relating power level and total flux is given in Equation 6.5, where ϵ is the amount of energy released in fission, Σ_f is the one-group macroscopic fission cross-section, and V is the reactor volume.

$$P = \epsilon\phi\Sigma_f V \quad (6.5)$$

For a reactor with a given power level, uncertainty in the neutron spectrum could lead to uncertainties in the absolute flux level by changing the one-group fission cross-section. Any changes in the predicted reactor flux level will have a direct impact on accumulated DPA predictions, and thus this mechanism should be accounted for if changes in the neutron spectrum lead to changes in the spectrally-averaged fission cross-section, as might be expected, for example, in the energy range near the U-238 fission threshold. Sensitivity coefficients for Σ_f are evaluated again using the combination of GPT and covariance data, where in this case the response ratio tallied in Serpent is given in Equation 6.6.

$$r = \frac{\int_{\Omega} \int_0^{\infty} \phi(r, E) \Sigma_f(r, E) dE dr}{\int_{\Omega} \int_0^{\infty} \phi(r, E) dE dr} \quad (6.6)$$

Using the sandwich rule, uncertainties in the one-group fission cross-section, $\sigma_{(j,j')}^{\Sigma_f}$, can be evaluated for each reaction pair, (j, j') . This uncertainty can then be propagated to an uncertainty in the flux level using a first order expansion of Equation 6.5.

Propagation to Uncertainty in Accumulated DPA

To propagate these three sources of uncertainty to the calculation of batch-wise one-group DPA cross-sections and accumulated DPA, it is necessary to expand Equation 6.1. Neglecting cross-correlation terms between the three sources, propagating uncertainty in Σ_f to ϕ , and pulling t_{ℓ} out of the summation (because it is the same for every batch in the reference core), the first-order uncertainty propagation equation for the accumulated DPA is found to be:

$$\sigma_{(j,j')}^{DPA} = t_\ell \phi \Sigma_{\ell=1}^{n_b} \hat{\phi}_\ell \Sigma_{i=1}^{n_r} \hat{w}^i \hat{\phi}_\ell^i \star$$

$$\sqrt{\left(\frac{\Sigma_f}{\sigma_{(j,j')}}\right)^2 + \frac{\Sigma_{\ell=1}^{n_b} \left[\left(\hat{\phi}_\ell \Sigma_{i=1}^{n_r} \hat{w}^i \hat{\phi}_\ell^i \right)^2 \left(\left(\frac{\sigma_{\hat{\phi}_\ell^i}}{\hat{\phi}_\ell} \right)^2 + \left(\frac{\sqrt{\Sigma_{i=1}^{n_r} (\hat{w}^i)^2 (\sigma_{\hat{\phi}_\ell^i}^i)^2}}{\Sigma_{i=1}^{n_r} \hat{w}^i \hat{\phi}_\ell^i} \right)^2 \right) \right]}{\left(\Sigma_{\ell=1}^{n_b} (\hat{\phi}_\ell \Sigma_{i=1}^{n_r} (\hat{w}^i \hat{\phi}_\ell^i)) \right)^2}} \quad (6.7)$$

The number of perturbation energy groups, n_p , is taken to be the same as the SCALE 56-group covariance matrices used in this study [89], which are mainly based on the ENDF/B-VII.1 nuclear data library. The number of response groups, n_r , is taken according to the 640-group discretization used in the ASTM DPA standard. n_b is the number of fuel batches in the reference core, outlined in Chapter 2. Sensitivities are calculated to perturbations of only the top ten most abundant isotopes in the core to save on computational cost, listed in Table 6.1. For each isotope, sensitivities are calculated for each reaction channel with a corresponding covariance matrix in the SCALE covariance set, which includes variances for various reaction channels as well as select covariances between different reactions of a given isotope. All GPT calculations are performed with ENDF/B-VII.1 cross-sections using 4.8×10^9 neutron histories. Any sensitivities with Monte Carlo standard deviation estimates higher than 10% are discarded, although this was seldom the case for sensitivities of any relevance.

Table 6.1: Top ten most abundant isotopes in the core used in perturbation calculations listed in order of decreasing abundance.

Rank	Isotope
1	U-238
2	Fe-56
3	Na-23
4	Cr-52
5	Fe-54
6	Pu-239
7	Fe-57
8	Cr-53
9	C-12
10	Cr-50

Results

Uncertainty in Accumulated DPA

Figure 6.2 shows the top reaction channels leading to uncertainty in the fission cross-section, which will be reflected as uncertainty in the core flux level and resulting DPA accumulation. The largest uncertainty comes from the ^{238}U inelastic cross-section, which leads to nearly 1.5% relative uncertainty on the fission cross-section. Other reaction channels lead to significantly smaller relative uncertainties of less than 0.5%. With the exception of the ^{238}U capture cross-section, the top contributors to uncertainty in the fission cross-section are either elastic or inelastic reactions.




Figure 6.2: Uncertainty in the whole-core fission cross-section from the top reaction channels.

Uncertainties in the normalized spatial flux profile, $\hat{\phi}_e$, are not presented graphically because their values were found to be negligible. The highest relative uncertainty values in each batch, all less than 0.01%, result again from the ^{238}U inelastic scattering cross-section. The next highest uncertainties occur for the ^{239}Pu fission cross-section in batches with the highest flux levels, although these are less than half of those found for ^{238}U . Such uncertainty values are much smaller than those found for the total fission cross-section, and therefore it is concluded that the propagated uncertainty in the accumulated DPA will be minimally influenced by uncertainties in the spatial flux profile.

Figure 6.3 shows the top reaction channels leading to uncertainty in the DPA cross-section for Batch 1, as an example. Results for other batches show similar trends, changing

somewhat between batches due to differing isotopics and localized spectral variations. Again, the highest uncertainties result from ^{238}U inelastic scattering uncertainties, with the next most important sources of uncertainty all being either elastic or inelastic scattering cross-sections. For all batches, the absolute uncertainties in DPA cross-sections are quantified to be less than 2 b, with the overall DPA cross-section being roughly 250-300 b. These uncertainties are therefore similar in importance to those found for the whole-core fission cross-section.




Figure 6.3: Uncertainty in the Batch 1 DPA cross-section from the top channels.

To understand where uncertainties in the one-group DPA cross-section come from, Figure 6.4 depicts the neutron spectrum along with the group-wise spectrum uncertainty associated with the ^{238}U inelastic scattering cross-section. With the exception of a few isolated peaks, spectral uncertainties resulting from uncertainties in the ^{238}U inelastic scattering cross-section are mainly concentrated in the regions from 0.1-20 keV and 1-10 MeV, where the spectral uncertainties are roughly 5-10%. Because both the flux level and the DPA cross-section are very low in the region from 0.1-20 keV, this portion of the spectrum does not contribute significantly to the DPA uncertainty. The DPA cross-section is highest in the region from 1-10 MeV, which leads to appreciable uncertainties in the one-group DPA cross-section, although this is tempered by the declining flux at these high energies. For cores with harder spectra, it is expected that this high energy region with ^{238}U inelastic scattering cross-section uncertainties could lead to notably higher uncertainties in the one-group DPA cross-section. This notion is backed up by examining the uncertainty in the

DPA cross-section as a function of the local fast flux in each batch across the core, as shown in Figure 6.5. Because the core is large and many of the low-powered batches see mainly leakage flux from the high-powered batches, a significant spectral variation is seen across the core. Figure 6.5 shows that as the fast flux fraction in a batch increases, the uncertainty in the one-group DPA cross-section associated with that batch also generally increases.




Figure 6.4: Normalized neutron flux and uncertainty from $(\sigma_{inel}^{U238}, \sigma_{inel}^{U238})$ in Batch 1.

Figure 6.6 shows the DPA uncertainty accumulated over the equilibrium cycle associated with different reaction pairs. Because $(\sigma_{inel}^{U238}, \sigma_{inel}^{U238})$ contributes the most uncertainty to both the fission and DPA cross-sections, it is responsible for the largest portion of the accumulated absolute uncertainty, roughly 10 DPA. This is followed by $(\sigma_{inel}^{Na23}, \sigma_{inel}^{Na23})$, which contributes only slightly more than 2 DPA. With the exception of $(\sigma_{cap}^{U238}, \sigma_{cap}^{U238})$, the top ten reaction pairs are either elastic or inelastic scattering reactions. The majority of DPA uncertainty accumulation takes place in Batches 2-5, where a combination of radial flux peaking and harder spectrum makes uncertainties from the ^{238}U inelastic cross-section particularly effective. These factors also contribute to making $(\sigma_{inel}^{Pu239}, \sigma_{inel}^{Pu239})$ the eighth highest contributor, even though this isotope is only significantly present in a few batches. In total, these reaction channels contribute only slightly more than 10 DPA uncertainty to the total accumulated DPA. Considering that peak damage in this core is on the order of 500 DPA, the relative DPA uncertainty resulting from nuclear data and the associated spectral uncertainties is estimated to be about 2%.




Figure 6.5: Uncertainty in the one-group DPA cross-section from $(\sigma_{inel}^{U238}, \sigma_{inel}^{U238})$ as a function of local fast flux ($E \geq 0.1$ MeV) in each batch across the core, along with a linear fit.

Comparison with Fast Fluence

For comparison, a similar study is undertaken to determine the uncertainty on the peak fast fluence in the core. This is accomplished by substituting the expression in Equation 6.8 as the weighting function used in Equations 6.1 and 6.7 instead of the DPA cross-section. In this study, the fast fluence threshold, E_f , is taken as the typical value of 100 keV.

$$\hat{w}(E) = \begin{cases} 0 & \text{for } E < E_f \\ 1 & \text{for } E \geq E_f \end{cases} \quad (6.8)$$

Before presenting the fast fluence uncertainties, a comparison between accumulated DPA and fast fluences are given for each batch in Figure 6.7, plotted as a function of the average neutron energy in the batch determined as in Equation 6.9.

$$\bar{E}_\ell = \frac{\int_{b_\ell} \int_0^\infty E \phi(r, E) dE dr}{\int_{b_\ell} \int_0^\infty \phi(r, E) dE dr} \quad (6.9)$$

Because the spectrum in B&B cores varies significantly among batches due to strongly varying fissile compositions and proportions of leakage flux between batches, the average neutron energy in a batch varies from 0.1 to 0.5 MeV. Over this range, the ratio of DPA




Figure 6.6: Uncertainty in the accumulated DPA for the top ten reaction pairs.

to fast fluence accumulation follows a nearly linear trend, showing the spectrum-sensitivity of the DPA metric as compared to fast fluence. This comparison highlights the potential drawbacks of using only fast fluence to quantify material damage in fast-spectrum systems and underscores the discrepancy in constraints when designing a system to meet either DPA or fast fluence constraints.

A summary of the fast fluence uncertainty results are presented in Figure 6.8. Comparing with the accumulated DPA uncertainties of Figure 6.6, again the uncertainty is dominated by the inelastic scattering cross-section of ^{238}U . In general, the relative contributions to the accumulated uncertainty are similar to that of the DPA uncertainties, with the exception of inelastic scattering of ^{23}Na , which contributes notably less uncertainty to the accumulated fast fluence than it does to the accumulated DPA. An additional difference between the uncertainties of the two measures of radiation damage is that the relative uncertainty in the accumulated fast fluence is slightly less than the relative uncertainty in the accumulated DPA, being roughly 1.6% as compared to 2.0%. This is in agreement with the previous observation that DPA uncertainties are highest in the high-energy range, but because the fast fluence metric does not distinguish between neutrons of different energies within the high-energy range, this increased uncertainty is not reflected in the fast fluence uncertainty. These results indicate that fast fluence may be predicted more reliably than DPA for B&B systems, although the margin is slight and even DPA predictions have relatively small uncertainties.




Figure 6.7: Ratio of accumulated DPA to accumulated fast fluence in each batch as a function of the average neutron energy in the batch, along with a linear fit. The numbers next to each datapoint indicate the corresponding fuel batch.

Summary of DPA Uncertainties

This study has shown that the accumulated DPA uncertainty in the reference core due to uncertainties in the nuclear data are small. The dominant mechanism is the inelastic scattering cross-section of ^{238}U , leading to 2% uncertainty in the accumulated DPA. The ^{238}U inelastic cross-section has such a large impact compared to other reaction channels due to the high number density of ^{238}U throughout the core combined with the high spectrum uncertainties at energies above 1 MeV where the DPA cross-section is highest. Additionally, ^{238}U inelastic scattering is an important slowing-down mechanism and its cross-section at high energy has relatively high uncertainty. For the core examined, the sharply decreasing flux at energies above 1 MeV keeps this uncertainty in check, but in cores with harder spectra this mechanism may lead to higher DPA uncertainties. In general, nearly all of the accumulated DPA uncertainty attributable to uncertain cross-section data results from either elastic or inelastic scattering reactions. These reaction channels lead to uncertainties mainly in the core flux level and the spectrally-averaged DPA cross-section, with minimal impacts on the radial flux profile.




Figure 6.8: Uncertainty in the accumulated fast fluence for the same reaction pairs as Figure 6.6.

6.3 Enriched B&B Reactors

In the same way that Zhang [156] generalized the B&B fuel cycle to include limited amounts of fuel recycling, the B&B fuel cycle may be generalized in an orthogonal direction to include fuel enrichment. Fuel enrichment has not been widely explored to date largely because of the negative impact on resource utilization, although there is nothing inherent about the breed-and-burn mode of operation that eliminates the possibility for enriched fuel. In fact, TerraPower is planning to use enriched fuel in their prototype reactor which will operate in breed-and-burn mode to a limited extent [57]. The enrichment gap between ‘pure’ B&B designs without enrichment and typical PWRs is such that there is significant margin between the resource utilization of the two systems which may still be exploited using limited amounts of enrichment. Furthermore, enrichment capabilities are widely available throughout the world at relatively low cost and no political baggage. Thus, using slightly enriched uranium as feed fuel may be able to simultaneously meet cladding damage constraints, improve resource utilization as compared to current LWR technologies, and avoid introducing unworkable technology into the fuel cycle. This section will explore this possibility, examining the tradeoff between material damage and resource utilization in B&B cores that employ enriched fuel, termed ‘eB&B’ cores.

It should first be noted that enriched feed fuel for B&B cores has been explored to some extent by Yarsky, who assumed the use of 5w% feed material in the conceptual design of

a gas-cooled B&B core [155]. Yarsky's designs employ either UC fuel or UN fuel enriched in the isotope ^{15}N , and as a result the peak material damage is still near 375 DPA even with the use of enriched feed material. It will be shown in this study that using fuels with higher heavy-metal densities can enable lower levels of material damage for similar feed enrichment levels. In addition, this section will explore the tradeoff between feed enrichment and resource utilization in much greater detail than what is presented by Yarsky.

Methodology

The different cases examined in this study are compared on the basis of their critical equilibrium cycle performance as calculated by ADOPT. All cores are designed according to the same assumptions as provided in Table 2.1a and outlined in Chapter 2, with the exception of the fuel type and enrichment levels (which are parametrically varied) and the target k_{eff} (taken as 1.00 in this study, but 1.03 in Chapter 2 to account for the lack of control system in ADOPT calculations)³. No studies of control system design are undertaken in this section.

This study primarily quantifies accumulated cladding damage through the DPA metric, although additional quantifications of material damage are presented in terms of accumulated fast fluence ($E \geq 100$ keV) since many researchers use a fast fluence limit of 4×10^{23} n/cm². Equilibrium resource utilization, U , is quantified in terms of the heavy metal mass fissioned in the reactor as compared to the amount needed to be mined during the equilibrium cycle. This can be expressed as in Equation 6.10 for the case where only a single stream of fuel at a given enrichment is used in the system, where e_f is the feed enrichment, e_n is the enrichment of the natural uranium feedstock, and e_d is the enrichment of the depleted uranium waste resulting from enrichment. For reactor systems where multiple feed streams or multiple enrichment levels are utilized, the more general form of the expression must be used. e_n and e_d are assumed to be 0.7w% and 0.25w%, respectively, consistent with the study performed by Kim et al. from which the reference PWR fuel utilization is taken [79]. This equation implies a $1/e_f$ drop in fuel utilization for a given burnup, highlighting the importance of keeping the feed enrichment as small as possible to achieve resource utilization better than that of a PWR.

$$U = \frac{HM_{fissioned}}{HM_{mined}} = \frac{e_n - e_d}{e_f - e_d} \frac{HM_{charge} - HM_{discharge}}{HM_{charge}} \quad (6.10)$$

First, the reference core design is modified to use feed fuel ranging from depleted to 7w% enriched uranium with a constant axial enrichment profile. Because the use of enriched fuel de-emphasizes the importance of minimizing leakage, the study then repeats a number of

³It is noted that a conference paper presented by the author at GLOBAL 2019 in Seattle covered the same subject matter, although the results in that paper differ from those which will be presented in this thesis. The reason for the difference is that the conference paper mistakenly utilized unalloyed metallic fuel even though the paper says otherwise, when it should have utilized U-10Zr fuel as envisioned for the B&B core due to concerns related to fuel-cladding eutectic reactions and swelling under irradiation. The results in this dissertation utilize the correct fuel composition and thus supersede those from the conference paper.

cases with decreasing core heights from 3.0 to 1.5 m, allowing for higher average burnups. The study then explores the possibility of grading the enrichment level as a function of axial position to avoid placing enriched fuel in regions of low importance and to flatten the axial power profile. The axial-grading patterns investigated, determined using a trial-and-error approach, are shown in Figure 6.9.

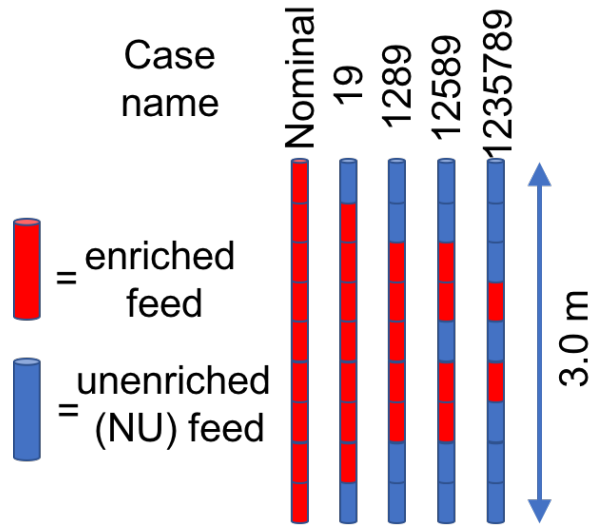


Figure 6.9: Unoptimized axial enrichment gradings investigated in this study.

Finally, the study explores the use of different fuel types. While metallic fuel is currently preferred within the B&B design community due to its high heavy-metal density and its limited moderating power, the use of fuel enrichment may alleviate some of these concerns. Therefore, the usage of oxide and carbide fuels, typically ruled out due to their reactivity penalty compared with metallic fuel, are explored to allow for slightly softer spectra and thus reduced DPA accumulation. Temperature- and burnup-dependent properties for both metallic and oxide fuels are already incorporated into ADOPT, but properties for carbide fuel are not. Considerably less experience has been accumulated with carbide fuel, and as such the properties used in the current study are for pure UC fuel at a fixed temperature of 1500°C as provided in [149].

Nominal Case with Enriched Fuel

Figure 6.10a shows peak cladding damage as the feed enrichment of the reference core is increased uniformly. As feed enrichment increases, peak DPA drops nearly linearly at a rate of roughly 60 DPA/%. Fast fluence constraints are found to be more restrictive than DPA constraints in this case, requiring feed enrichments of 6.7 and 6.0w%, respectively. This result is qualitatively in line with the results of Zhang [156] and represents a discrepancy

that must be resolved by the nuclear materials community in order for B&B design space to be fully characterized. Because DPA quantification allows for a more nuanced reflection of the neutron spectrum, it is used for the remainder of this study, although it should be kept in mind that fast fluence constraints may dictate higher feed enrichments.




Figure 6.10: Peak cladding damage and resource utilization for the reference core with increasing feed enrichment.

Figure 6.11 summarizes the main parameters which lead to DPA reductions in terms of ratios between cases with NU and 6w% enriched feed. The dominant effect is the reduction in the minimum required cycle length that occurs as enrichment increases, which directly reduces accumulated cladding damage by reducing fluence. In addition, both the radial and axial flux peaking decrease with increasing feed enrichment, bringing down peak values. Conversely, the one-group DPA cross-section actually increases as enrichment is increased because the spectrum hardens. Figure 6.11 shows that this effect is much weaker than the other effects, leading to a net decrease in DPA accumulation as enrichment increases.

Figure 6.12 shows how the feed enrichment level impacts one of the key reactivity feedback mechanisms – the coolant void worth. At both BOEC and EOEC, the coolant void worth drops by a factor of ~ 2 moving from low to high enrichment levels, driven in large part by the increased delayed neutron fraction at higher enrichments. Because U-235 has a delayed neutron fraction of 0.0064 as compared to 0.0021 for Pu-239 [31], increasing the fraction of fissions taking place in U-235 can serve to slow down transients and make the importance of material perturbations less severe. Therefore the eB&B core concept may also admit benefits from the perspective of inherent safety in addition to bringing down accumulated material damage levels.

Figure 6.10b shows the resource utilization achieved as feed enrichment is increased, where the predicted $1/e_f$ drop is prominent. At enrichment levels high enough to meet peak




Figure 6.11: Ratio of select characteristics which contribute to accumulation of cladding damage for the case with NU as compared to 6w% feed fuel.

DPA constraints, resource utilization is slightly *lower* than that for a PWR. Part of the reason for this poor fuel utilization is that axial power peaking is very large in these cores, being as high as 2.4 in the case with NU feed. Therefore much of the fuel towards the axial periphery achieves very low burnup and consequently brings down the average discharge burnup. In the case of a standard B&B design with NU feed, this is not important because low burnup in low-grade feed material has a very small impact on overall resource utilization, but as enrichment increases this inefficiency becomes more significant. This is the motivation for examining shorter B&B cores in the next section.

Shortened Core Height

For modest reductions in core height, the increase in the enrichment level required to remain below the DPA constraint is minimal, as shown in Figure 6.13a. This is due to the highly axially-peaked power profile which leads to very low axial leakage levels, even in cores as short as 2.0 m, as demonstrated in Figure 6.14. For similar reasons, the peak flux level initially remains largely unchanged as the core height is decreased from its nominal value. This enables small increases in resource utilization, as depicted in Figure 6.13b, such that the 2.0 m eB&B core achieves resource utilization roughly equal to that of a PWR. However when the core height nears 1.5 m, the axial power profile begins to interact more strongly with the




Figure 6.12: Core coolant void worth and delayed neutron fraction as a function of feed enrichment level in the eB&B core.

leakage probability, degrading the neutron economy and leading to higher material damage levels which require higher feed enrichments to remain below material damage constraints.

It is found that breed-and-burn operation cannot be achieved with core heights below 2.0 m and feed enrichments below 3w% (which is why those points are omitted from Figures 6.13 and 6.14). The addition of axial reflectors to the 1.5 m core helps to a minor degree, enabling a lower feed enrichment of roughly 6.8w% and a corresponding resource utilization just slightly above that of a PWR. Regardless, the benefits obtained from shortening the core, while somewhat significant in a relative sense, do not enable resource utilizations appreciably better than current LWRs.

This section has studied the sensitivity to reducing the core height while keeping the power level the same, therefore leading to an increase in power density. An alternative possibility would be to reduce the core height while keeping the power *density* constant by derating the reactor power. Such a technique would enable higher fuel volume fractions because the coolant flow area could be made smaller under the reduced heat load, which could allow for a lower required feed enrichment due to the reduction in parasitic losses to coolant. Preliminary studies of this methodology, however, have shown only a very minor sensitivity reflected in the feed enrichment required to meet the DPA constraint. Since reducing the power level would be detrimental to the economics of the system, this methodology has not been pursued further.




Figure 6.13: Peak cladding damage and resource utilization for the shortened variants of the reference core with increasing feed enrichment.

Axially-Graded Feed Enrichment

Figure 6.15a shows that utilizing NU towards the axial periphery, as in Case 19, has very minimal impact on the peak DPA, as might be expected due to the very low burnup in those regions of the core. The resource utilization in this case is very similar to that of the 2.0 m tall core. Using NU in more axial segments, as in Case 1289, results in a noticeable increase in the enrichment level needed to meet the DPA constraint but also some benefits from the perspective of resource utilization such that the eB&B is able to have resource utilization just higher than PWR levels, similar in effect to the 1.5 m core with axial reflectors. Replacing enriched fuel from the axial center with NU, as in Case 12589, brings a sizable penalty in reactivity such that DPA constraints cannot be met with lower than 7w% enrichment levels, at which point resource utilization is degraded.

In addition to reducing the volume of enriched feed, the intent of cases without enriched fuel at the axial center (i.e. Cases 12589 and 1235789) was to reduce axial peaking, and thus bring down peak DPA values. In reality, this strategy has little impact on the axial peaking factor because the Pu-239 content builds up so strongly in this region due to there being the highest flux levels. This same effect was encountered by Yarsky in his investigations of axial enrichment grading in his gas-cooled B&B design [155], and thus it is confirmed that axial enrichment grading will not provide benefits from the perspective of axial peaking.

Alternative Fuel Types

The use of both oxide and carbide fuels is examined in this section due to the softened spectrum these fuels provide compared to metal fuel. Provided that fuel enrichment can




Figure 6.14: Total leakage probability for the shortened variants of the reference core with increasing feed enrichment.




Figure 6.15: Peak cladding damage and resource utilization for the reference core with increasing feed enrichment for axially-graded enrichment patterns. Indicated enrichment corresponds to axial regions as described in Figure 6.9.

make up for the loss in reactivity resulting from the softer spectrum, this may be a viable option. Figure 6.16 shows the spectrum in Batch 3 (the highest power batch) for the metallic, oxide, and carbide cores employing the enrichment levels required to achieve B&B operation. Table 6.2 gives the spectrally-averaged DPA cross-sections for the same scenarios along with the required enrichment levels, showing that the softer spectrum in the oxide core allows for a 12% reduction. Even though the DPA cross-section decreases, negative aspects of these alternative fuel types, such as the lower heavy-metal inventory and reduced breeding ratio, outweigh the positives. In the case of carbide fuel, 6w% feed fuel is required just to achieve breed-and-burn operation, at which point the accumulated material damage is still more than 330 DPA. In order to meet the DPA constraint, the fuel enrichment must be increased to $\sim 9\text{w}\%$, at which point fuel utilization is much lower than that of a PWR. In addition, the lack of operational experience with carbide fuel as compared to U-10Zr alloys makes it so there is no real impetus to move towards carbide fuels in the eB&B cores. In the case of oxide fuel, it appears that the DPA constraint cannot be met while operating in breed-and-burn mode, even if the feed enrichment is increased beyond 10w%. While breed-and-burn operation is achieved at 10w% enrichment, moving to higher enrichments makes the fissile content so high that the core is in fact purely a burner, and the design space to operate in breed-and-burn mode is extremely small. Therefore, it does not appear that either oxide or carbide fuels will allow for improvements in resource utilization subject to current DPA constraints.

Table 6.2: DPA cross-sections at the axial midplane of Batch 3 in the metal, oxide, and carbide cores at the enrichment levels required to achieve breed-and-burn operation.

	Metal	Carbide	Oxide
σ_{DPA} (b)	268.5	251.9	236.6
Enrichment (w%)	0.25	5	10

Discussion of eB&B Cores

This eB&B study has shown that the tradeoff between resource utilization and material damage is severe. While it is trivial that enriched fuel can enable breed-and-burn operation without exceeding material damage constraints, primarily through reductions in the necessary cycle length and flux peaking, the utilization of fuel resources is extremely degraded due to the enrichment process. Using either shorter cores or optimized enrichment patterns, resource utilization can be increased to levels only slightly higher than those in a PWR. While the DPA cross-section can be brought down by using other fuel forms that enable a softer spectrum, these options do not allow for improved performance from the perspective of resource utilization either. In addition, this study did not account for either (1) fuel




Figure 6.16: Neutron spectra in Batch 3 of each of the metal, oxide, and carbide cores at the enrichment levels required to achieve breed-and-burn operation.

requirements during the transition to equilibrium or (2) neutron losses to control materials during operation, both of which will serve to worsen fuel utilization such that improvements over a PWR are not expected.

While the use of enriched fuel alone is not able to simultaneously achieve acceptable cladding damage and high fuel utilization, the results of this study indicate that feed fuel enriched only slightly, perhaps to 1.5w%, may provide reductions in cladding damage of roughly 150 DPA while still allowing for fuel utilization up to 8x that of a PWR. Combining this with another methodology to reduce the cladding damage accumulation, such as the use of 3D shuffling, may enable cladding damage constraints and fuel utilization improvements to be met simultaneously. The results of this study also imply the incremental benefits which may be achieved if cladding damage constraints were to be relaxed. While relaxing cladding damage constraints from the current value of 200 DPA to the desired value of near 500 DPA is not guaranteed to succeed, a relaxation to ~ 300 DPA is much more realistic in the near-term and has a higher likelihood to succeed. Such a relaxation combined with the use of enriched feed fuel would allow for improvements in fuel utilization of 3-4x compared to a PWR. However, the true benefits of the breed-and-burn operation will only be achieved as the DPA constraint is relaxed beyond ~ 400 DPA, assuming a conventional core design paradigm.

Because the fuel achieves relatively low burnup for the level of enrichment, it is inferred

that significant fissile material is contained in the discharge fuel. Figure 6.17 demonstrates this point for fuel discharged from the nominal 7 w% eB&B core. As expected, the U-235 content is most depressed in the axial center where the majority of power has been produced, but even at this axial position the U-235 enrichment level is still approximately 2.5a%. More significantly, a large inventory of Pu-239 has built up at the axial center through the breeding process such that its concentration is approximately 6a% at discharge. Because the structural materials have reached their DPA limit, though, this fissile content cannot be burned without some sort of mitigation effort. This same realization is what motivated the double-clad strategy of di Sanzo [29], although in that case the fuel could not achieve breed-and-burn operation within material damage constraints in the first place because it was not enriched. In the case of the eB&B core, the motivation to use the discharged fuel is even greater because (1) the assemblies will have higher reactivity than those discharged from a typical B&B core and (2) the cost of fuel is much higher because of the need to enrich.

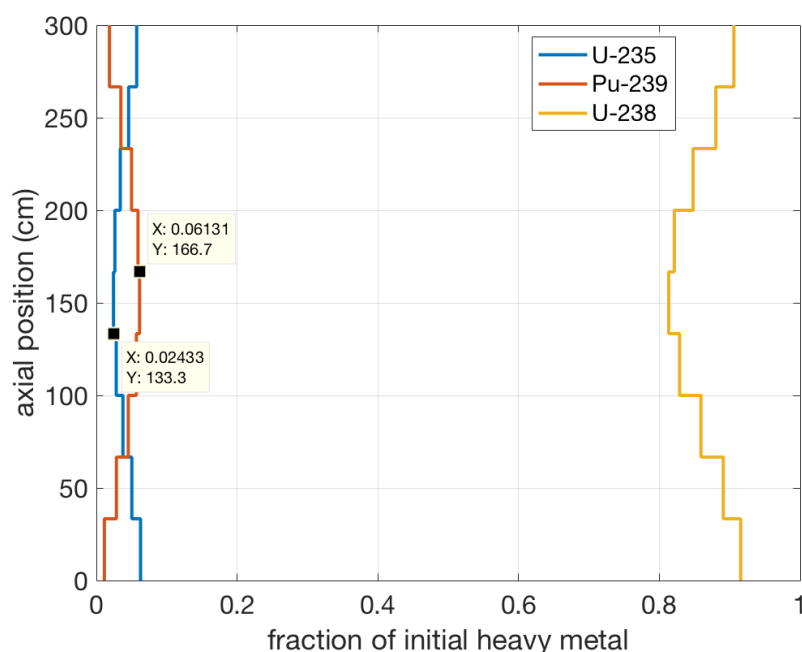


Figure 6.17: Fissile content of discharged 7w% eB&B fuel as a function of axial position.

In order to utilize the discharged eB&B fuel, there are a few options to bypass the material damage constraints. The first and most widely discussed option is to de-clad the fuel and utilize some form of minimally-invasive reconditioning technique to refurbish the fuel into a new assembly. Such a technique trivially bypasses the material damage constraint but is paradoxical as applied to the eB&B core because the ability to de-clad and refurbish fuel would obviate the need to enrich the fuel in the first place. The same goes for any methods which employ traditional reprocessing techniques.

The second option would be to double-clad the discharged eB&B fuel pins and reconstitute them into another assembly which could be further burned in the fast spectrum. This is similar to what has been investigated by di Sanzo, except in that case the first stage reactor was a typical B&B core fed with NU or DU that far exceeded the DPA constraint. The benefit to instead using an eB&B core as the first stage would be that the material damage constraint is never exceeded, while further burnup in the double-clad fuel would enable resource utilization higher than that achieved in the single-stage eB&B core. The results of di Sanzo imply that this strategy can often double the achievable burnup. Because such methodology doesn't breach the cladding, it avoids any proliferation concerns and is thus investigated further in Section 6.4.

The final strategy would be to put the fuel into a thermal spectrum reactor *as is* without any form of reprocessing. This helps to avoid the DPA constraint for two reasons: (1) thermal neutrons induce orders of magnitude lower material damage in the cladding material due to the lower energy transfer upon collision, and (2) flux levels in thermal reactors are typically orders of magnitude lower than those in fast reactors due to the higher fission cross-sections at low energies. To the author's knowledge, this strategy has never before been investigated, and is the focus of Section 6.5.

6.4 Double-Cladding of eB&B Discharged Fuel

In brief, the concept of double-cladding is as follows. In order to extend the life of individual fuel pins beyond their material damage constraints without removing the fuel from the cladding, it is possible to simply add a second layer of protection by encasing the full fuel pin in another layer of cladding [29]. This would entail (1) the deconstruction of each fuel assembly into its constituent fuel pins, (2) the encasing of each pin with a second layer of cladding, and (3) the reconstitution of a new fuel assembly with the double-cladded fuel pins.

While a detailed process for performing these tasks has yet to be devised, it is assumed it would need to be done under a sodium environment due to the fuel heat load. The difficulty of performing remote tasks under sodium has been seen in many operating SFRs in the past (see, for instance, the extensive design measures taken to enable recovery of dropped assemblies during refueling [81]), and therefore the possibility of performing such a complicated task is not at all clear. In addition, manufacturing processes for fuel pins adhere to strict tolerances so that the geometry of a finished assembly is known to high precision – such high precision of remotely-fabricated double-clad pins is perhaps unlikely. Fitting a new cladding tightly around a discharged fuel pin may be quite difficult without incorporating an appropriately sized gap between the old and new cladding layers which may strongly impact cooling requirements. Regardless, this study makes the assumption that such a complicated *engineering* operation can be performed and assesses the benefits from the perspective of *physics*. If the theoretical benefits enabled by this strategy justify them, it is assumed that proper remote handling methods can be developed.

Methodology and Assumptions

From a neutronics point of view, the addition of a second cladding layer is undesirable because it will lead to a lower fuel volume fraction (i.e. higher parasitic absorption and softer spectrum) and worsened thermal performance. The extent to which each of these aspects will be negatively impacted depends highly on the required thickness of the second cladding layer, something which is not clearly defined. In reality, the structural integrity of the old cladding at 200 DPA is not completely gone, and in combination with a second cladding layer the old cladding may still serve some structural function. Therefore the thickness of the second cladding layer should be determined using a detailed structural analysis of fuel pins with their discharge properties, particularly if the sensitivity to cladding volume is strong. In addition, the thickness of the second cladding layer should be designed for a target burnup (i.e. material damage level). In most cases this would be assumed to be another 200 DPA, after which the second cladding layer would reach its lifetime as well. However it is possible that the lifetime of the reconstituted fuel assemblies may be limited not by material damage but rather by reactivity, in which case the cladding could potentially be made thinner according to the expected assembly lifetime. In lieu of a detailed structural analysis of the fuel pins, this study makes the assumption that the second cladding layer is equal in thickness to the first and tightly hugs the first cladding layer. To quantify the impacts of this assumption, a sensitivity study is performed.

With this assumption, the remaining lattice parameters, such as pin pitch and number of pins per assembly, still need to be determined. Based on the results of di Sanzo, which indicate that the cooling requirements of a double-clad pin are not substantially different than the original single-clad pin [29], the coolant flow area in each subchannel is kept the same as in the nominal core while the pin pitch is adjusted accordingly. In addition, the coolant mass flow rate is kept unchanged from the nominal core. In this way, the most limiting constraint encountered in the eB&B designs, the coolant velocity, is unimpacted (see Chapter 2 for a discussion of the design constraints). For this preliminary study, the duct thickness is assumed to be the same as the nominal core, and the duct radius is adjusted so that it conforms to the larger pins. The number of pins per assembly is kept unchanged. The material volume fractions under these assumptions are provided in Table 6.3, and a sensitivity study around these values is subsequently performed. Schematics of both the original assembly and the double-clad assembly are provided in Figure 6.18.

The fuel composition is taken as the discharge fuel from the first-stage eB&B core. Because enrichment costs were minimized in the axially-graded Case 1289 (see Figure 6.9), the discharged fuel from this core with 6.7w% feed is used as input to the second-stage double-clad core. Other first-stage options are subsequently evaluated. The burnup potential of the second-stage core is preliminarily evaluated using the neutron balance method, outlined in Equation 1.1, where an axially-finite fuel assembly unit cell with 9 axial segments is used such that the isotopic variation is represented. The neutron balance method is used as compared to a full-core burnup study because the fuel management scheme suitable for such a second-stage core is not well defined at this stage in the investigation. Once the trend of

Table 6.3: Volume fractions of constituents in the unit cell models for single- and double-clad assemblies. Fuel volume fraction includes the fuel-clad gap and coolant fraction includes the inter-assembly gap.

	Single clad	Double clad
Fuel	0.535	0.464
Clad	0.125	0.229
Coolant	0.269	0.241
Duct	0.071	0.066

the neutron balance is understood, the fuel management scheme can be inferred. The power density is assumed to be the same as the core-averaged specific power (i.e. power per fuel mass) in the first-stage core, and the fuel is burned up until reaching a peak damage level of 200 DPA on the second cladding layer.

Results

Figure 6.19 shows the multiplication factor, neutron balance, and cladding DPA for the double-clad assembly along with that for a single-clad assembly so that the impact of the second cladding layer can be quantified. The fuel discharged from the first-stage eB&B core is found to have high reactivity, resulting in a multiplication factor ~ 1.11 for the case of the single-clad assembly at its Beginning of Life (BOL). Adding a second layer of cladding leads to a reactivity penalty of roughly 6000 pcm, although the fuel still remains supercritical. As a result, the neutron balance is consistently positive, meaning the concept of the ‘minimum required burnup’ is not relevant in the second-stage double-clad core.

Even for the double-clad fuel, reactivity is not the factor limiting the fuel’s lifetime, as the fuel remains supercritical for more than 30 years. Instead, the fuel lifetime is limited by the accumulation of cladding damage up to the 200 DPA limit, which occurs after roughly 13.5 and 14.5 years for the single- and double-clad cases, respectively. This is approximately the same length of time that the fuel is in the first-stage eB&B core, implying that double the amount of energy can be extracted from the fuel if a second-stage core is employed. It is interesting to note that the single-clad case would actually generate *less* energy than the double clad case because the DPA accumulation rates are higher in the single-clad case owing to a harder neutron spectrum. Because the fuel lifetime is limited by DPA rather than reactivity, the spectral softening from the second cladding layer actually ends up being a benefit, in this case. This suggests that the fuel lifetime might be made even longer if the spectrum were softened further, which is investigated subsequently.

The sensitivity of these results in terms of fuel lifetime with respect to perturbations in the cladding and coolant volume fractions is provided in Table 6.4 as approximated using a central difference scheme, where perturbed values are obtained through direct perturbations




Figure 6.18: Neutronic unit-cell models employed in Serpent for the investigation of double-clad fuel pins in a second-stage B&B core.

to the geometry. Because the fuel lifetime in the second-stage core is limited by DPA accumulation and not reactivity, the fuel lifetime actually *increases* as the cladding and coolant volume fractions go up. This is because, similar to what was seen when comparing the single- and double-clad lifetimes, the DPA accumulation rate is slightly lower with higher non-fuel volume fractions because the spectrum is softened. This also explains why the sodium volume fraction would have a larger sensitivity than the cladding, since sodium is a lighter atom and thus better moderator than those in steel. It is emphasized that this result holds only while the DPA accumulation is more limiting than the reactivity. Considering, as seen in Figure 6.19a, that the fuel is still supercritical even after 20 years, it is anticipated to be unlikely that enough cladding will be required for the double-clad fuel so as to make the lifetime limited instead by reactivity considerations. This result bodes well for the feasibility

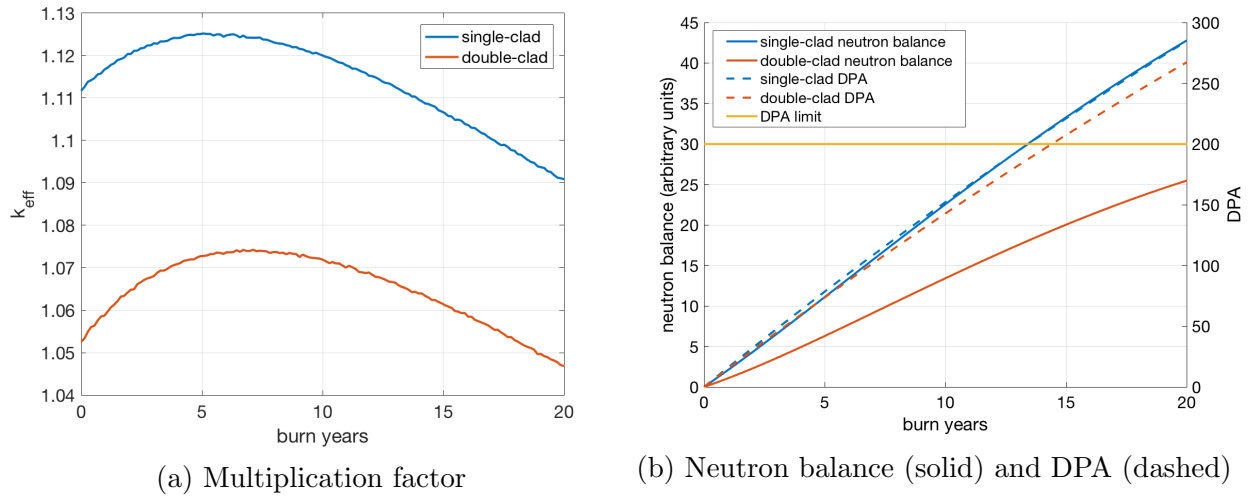


Figure 6.19: Multiplication factor, neutron balance, and material damage accumulation in the single- and double-clad cases using fuel discharged from the first-stage eB&B core.

of the double-clad concept.

Table 6.4: Sensitivity of the fuel lifetime to the cladding and coolant volume fractions.

	Clad	Coolant
Sensitivity ($\frac{\%lifetime}{\%volume}$)	0.10	0.40

In addition, the second cladding layer has an impact on the material worths, in particular the coolant void worth. Figure 6.20 shows how the coolant void worth is reduced in the double-clad case as compared with the single-clad case over the lifetime of the fuel in the second-stage core. While Figure 6.20 should not be taken as the absolute void worth in the second-stage core due to the fact that it was generated using a unit-cell model, these results indicate that the void worth in the double-clad case will be smaller by $\sim 10\%$. This is due to the fact that increased parasitic absorptions in the thicker cladding material will be more pronounced, lessening the importance of parasitic absorptions in the coolant. Although not a motivation for pursuing the double-clad second-stage core, these results do indicate that the second-stage core could have better inherent safety performance than the first-stage eB&B core, which was itself previously shown to have more favorable void worth compared to a classical (unenriched) B&B core (see Figure 6.12).




Figure 6.20: Coolant void worth determined using the unit-cell models of a second-stage core with single- and double-cladding. Results have been smoothed in post-processing to eliminate fluctuations due to statistical noise.

Sensitivity to DPA Constraint

As highlighted in Section 6.3, the resource utilization of eB&B cores is highly impacted by the enrichment levels needed to meet the DPA constraint. Higher DPA constraints result in less DU waste because the minimum required burnup can be achieved with lower-grade feed fuel. In addition, higher DPA means longer cycle lengths (i.e. more energy extracted from the same amount of fuel). In the case of the first-stage eB&B core this tradeoff is straight forward and has already been quantified. However, higher DPA constraints will also impact the second-stage core by altering the composition of the discharged eB&B fuel and enabling longer cycle lengths. This section investigates the sensitivity of the neutron balance in the second-stage core to the DPA constraint by relaxing the constraint from 200 to 300 DPA. Such damage levels, although perhaps not feasible right now, have a reasonable likelihood of being feasible in the future through further materials irradiation campaigns⁴.

From Figure 6.9, it is known that for the axially-graded Case 1289, a constraint of 300 DPA can be met with approximately 5.3w% feed fuel (as compared to 6.7w% for the 200 DPA constraint). Figure 6.21 demonstrates the different fissile isotopic compositions of the

⁴Such irradiation campaigns may not take that long, either, if further irradiation can be accumulated on the HT9 samples already taken to 200 DPA in FFTF.

discharged eB&B fuel with different DPA constraints, where the impacts of the initial enrichment grading pattern can be clearly seen. In general, the U-235 content of fuel discharged at 300 DPA is lower for two reasons: (1) the feed enrichment to the first-stage core is lower and (2) the fuel has a significantly longer residence time and thus greater burnout. However the longer lifetime also enables greater buildup of Pu-239. Higher burnout of Pu-239 at the axial center can be seen in the 300 DPA case, as might be expected as the breed-and-burn process progresses.

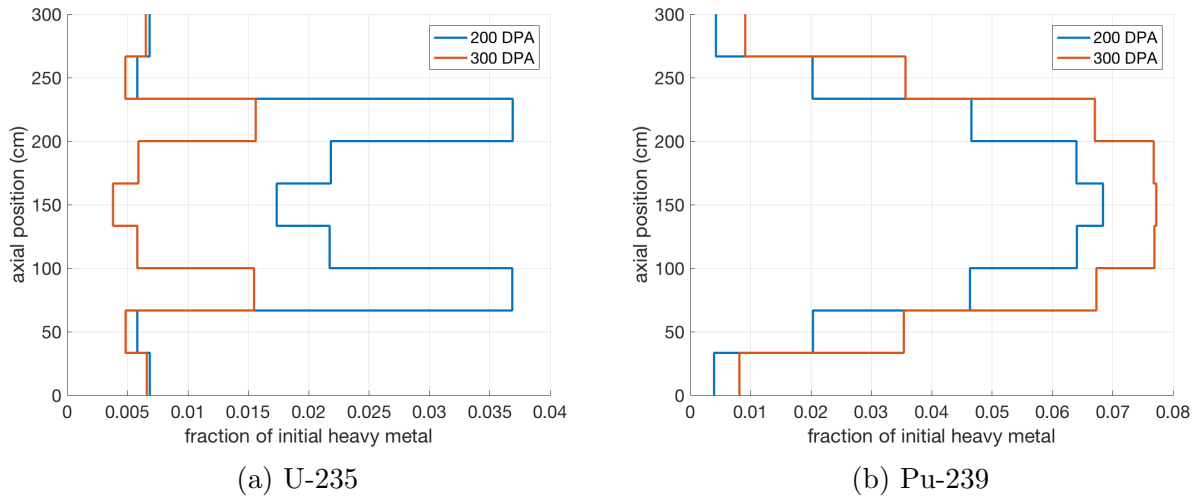


Figure 6.21: Axial concentration profiles for U-235 and Pu-239 at BOL in the second-stage core for various DPA constraints in the first-stage core.

Figure 6.22 shows the multiplication factor, neutron balance, and material damage associated with the two material damage constraints in the second-stage double-clad core. Ultimately the isotopic differences shake out such that the 300 DPA case has a higher BOL k_{eff} , driven largely by the Pu-239 content. The k_{eff} evolutions between the two cases are similar in the sense that both cases start significantly supercritical and remain supercritical throughout their lifetimes. However, for the 200 DPA case the fuel sees a reactivity increase during the first 7 years of burnup, while the 300 DPA case sees essentially a monotonic decrease over its life. While seemingly different, the two k_{eff} profiles are simply shifted versions of each other, where the 300 DPA case is a snapshot of the 200 DPA case at a later phase of burnup (albeit also shifted down from the reactivity penalty associated with the lower initial U-235 enrichment levels).

Table 6.5 shows the lifetime of an assembly in the first- and second-stage cores for the various DPA constraints. Once again it is found that the assembly lifetime is limited not by reactivity, but by material damage constraints. It is clear that extending the DPA constraint enables significantly more energy extraction, with a 50% increase in the DPA constraint leading to nearly a 200% increase in the assembly lifetime. Such extended lifetimes would




Figure 6.22: Multiplication factor, neutron balance, and material damage accumulation in the double-clad cases with increasing DPA constraints.

improve the resource utilization and economics of the eB&B concept on top of the benefits derived from the lower enrichment requirements in the first place.

Table 6.5: Assembly lifetimes in the first- and second-stage cores for various DPA constraints.

DPA Constraint	Feed enrichment (w%)	First-stage lifetime (yr)	Second-stage lifetime (yr)	Total assembly lifetime (yr)
200	0.7/6.7	10.0	14.4	24.4
300	0.7/5.3	17.5	27.9	45.4

It is interesting to note that the damage accumulation rate is slightly lower in the case with a 300 DPA constraint. This is due to the flux level being $\sim 4\%$ higher in the 200 DPA case due to there being less Pu-239. Pu-239 has not only a slightly larger fission cross-section than U-235 at high energies, but also its fission releases almost 3% more energy, both of which contribute to enabling a lower flux in the core for the same power generation.

Adding Natural Uranium Breeding Assemblies

The excess reactivity shown in Figure 6.19 for the double-clad case implies that further breeding may be possible if a second stream of fertile material were also fed into the core. This would be beneficial from the perspective of increasing the utilization of discharged

first-stage eB&B fuel by supplementing it with low-grade fuel, as well as possibly reducing the core reactivity swing. In this way, the amount of power generated from the discharged first-stage fuel can be extended while also improving safety characteristics.

To preliminarily understand the potential benefits from such a strategy, the unit cell model is modified so that some fraction of fuel pins are replaced with NU. This is meant to be a surrogate for a full-core situation in which some fraction of assemblies are entirely NU blankets, as opposed to a situation where a given assembly is mixed with both NU and discharged eB&B pins (which may be both more expensive to manufacture and difficult to adequately cool). Using the fact that neutrons in the fast spectrum have a long mean free path, the two types of pins are mixed into a single assembly for the purposes of this scoping study, as in Figure 6.23. The outer diameter of the NU pins is taken to be the same as the double-clad pins in the assembly, although the cladding thickness of the NU pins is taken as only the thickness of a *single* layer because these pins will only have to withstand a single pass of burnup. To fill the extra space inside the pins where the second cladding layer would otherwise be, the NU fuel pellets have a larger diameter.




Figure 6.23: Example of neutronic unit cell model used to quantify the impacts of substituting some fraction of discharged eB&B fuel with NU.

Figure 6.24 shows the impact of an increasing proportion of DU pins on the criticality evolution, neutron balance, and DPA accumulation. In general, the DPA accumulation on the cladding is impacted very little by the fraction of NU pins, and for all cases the 200 DPA limit is reached after roughly 14 years in the second-stage core.

In contrast, the neutron balance is highly impacted by the presence of NU pins, although it is interesting to note that the maximum reactivity is largely unimpacted. As the NU fraction is increased, the BOL k_{eff} decreases strongly, leading the neutron balance to be

initially negative for NU fractions higher than $\sim 10\%$. Regardless, for NU fractions below $\sim 34\%$, the strong increase in k_{eff} in the first couple years of burnup enables the neutron balance to quickly turn around such that a net-zero neutron balance can be achieved before the cladding reaches its DPA constraint. Such results potentially imply a significant extension of energy production from the discharged eB&B fuel. Because NU fuel has no DU waste stream or enrichment costs, this would represent an increase in resource utilization at very little cost.




Figure 6.24: Multiplication factor, neutron balance, and material damage accumulation in the cases with increasing proportions of NU in the second-stage double-clad core.

The fact that the fuel remains significantly supercritical even after burnup to 200 DPA in the second-stage core, particularly in the cases with higher fractions of NU pins, implies the possibility of attaining further burnup by adding a *third* layer of cladding, at least from the perspective of reactivity. From the perspective of manufacturing, the process would be very much similar to that which was used to double-clad the fuel the first time, and thus if it is possible to add a second layer of cladding in the first place, it is assumed possible to add a third. Other issues related to coolability of the fuel would need to be evaluated, including the myriad uncertainties associated with heat conduction through the various cladding and gap layers. Regardless, as a possibly attractive strategy, it is here evaluated from the perspective of a neutron balance, as shown in Figure 6.25.

k_{eff} is seen to increase only for the first few years in the third-stage core since much of the fuel has already been burned for more than 20 years, after which point it shows a strong monotonic drop as the heavy metal content burns out and fission products build up. On top of the limited breeding, the reactivity margins are slimmer than in the double-clad case due to the extra cladding volume fraction. These thin margins are highlighted by the neutron balance cases where the radial leakage probability is artificially increased by multiplying k_{eff}

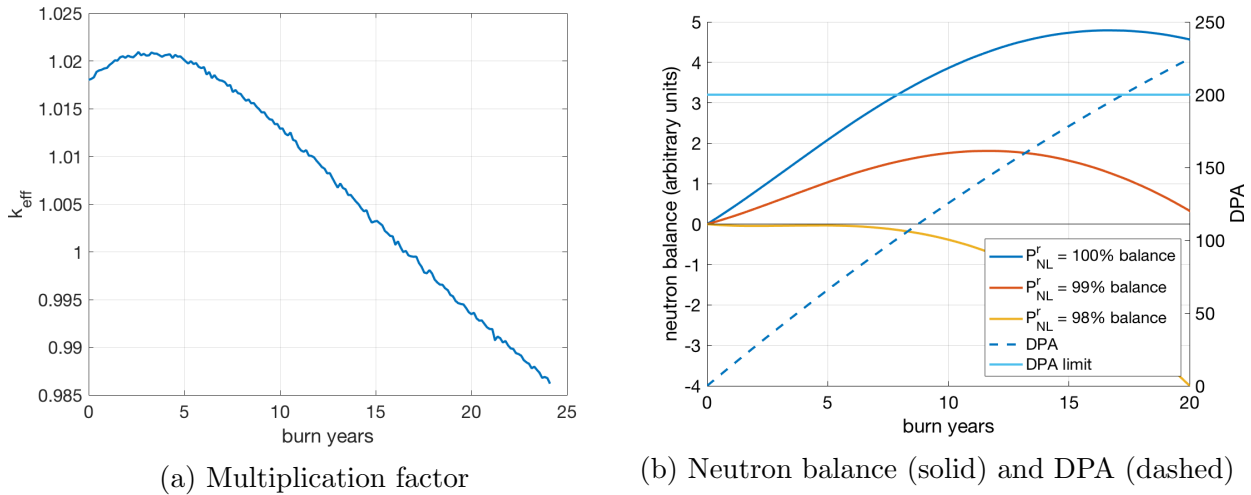


Figure 6.25: Multiplication factor, neutron balance, and material damage accumulation in the triple-clad case.

in Equation 1.1 by the radial non-leakage probability, P_{NL}^r . In order for the triple-clad third-stage core to be viable from the perspective of reactivity, radial leakage must be kept below 2%, but for designs with $P_{NL}^r \leq \sim 1\%$, the fuel lifetime will again be limited by irradiation damage to the cladding after roughly 17 years. It is interesting to note that once again the additional cladding layer softens the spectrum such that a notable decrease in the DPA accumulation rate is seen.

Attaining such low leakage is perhaps unlikely in a fast spectrum system without the use of many radial blankets (as in the original B&B core design, for instance), in which case the fuel lifetime becomes severely limited instead by reactivity. Considering the uncertainties associated with the required cladding thickness and overall assembly geometry in the triple-clad case, the the potential of such a strategy is questionable.

Softening the Spectrum

It was identified previously that the addition of the second cladding layer induced a slight softening of the spectrum which had an identifiable impact on the DPA accumulation rate such that the DPA-limited assembly lifetime was slightly increased. This section investigates the possibility that the lifetime of assemblies in the second-stage core may be elongated by *intentionally* introducing limited moderating material to the core such that DPA accumulation is reduced to the point that the core is limited instead by reactivity considerations (i.e. the neutron balance). This study is performed in a similar manner as depicted in Figure 6.23, except instead of replacing fuel pins in the assembly with NU, the pins are replaced by pins of solid moderator. As the number of fuel pins is reduced, the power density is

adjusted so that the specific power density remains unchanged. In order to evaluate the merits of different moderating materials, two solid moderators are investigated: (1) graphite and (2) ZrH.⁵ All results in this section pertain to the axially-graded enrichment pattern in the first-stage eB&B core.

Figure 6.26 shows the radially-averaged neutron spectrum at the axial midplane at BOL in the second-stage core for various numbers of pins replaced by graphite moderator. While the spectrum is softened, the amount of moderating material is small enough that the overall spectrum is still very much fast, and no issues associated with localized power peaking are anticipated.

Figure 6.27 shows how the neutron balance and DPA accumulation are impacted by the addition of moderating material.⁶ The addition of moderating material is found to be effective at reducing the DPA accumulation rate such that the assembly lifetime can be increased by nearly 30% for the case of graphite moderator. While ZrH moderator is also effective at reducing the DPA accumulation rate, ZrH has a larger neutronic penalty such that the assembly becomes limited by reactivity with a lifetime increase of only $\sim 18\%$. This differing neutronic impact between the two moderators is more clearly depicted in Figure 6.28, where the tradeoff between reactivity and DPA lifetime is shown to be nearly twice as bad as for graphite. However, the number of moderator pins needed to affect reduced DPA accumulation rates in the ZrH variant is roughly one third of that in the graphite variant, implying higher core power densities and lessened manufacturing requirements. In light of the increased fuel lifetimes and the little additional effort required, spectral-softening in the second-stage double-clad core is viewed as a viable option for further increasing sustainability without introducing reprocessing into the fuel cycle. However, improvements may be somewhat smaller than those from the strategy of adding additional NU pins to the double-clad core, as discussed further in Section 6.7.

Summary of Double-Cladding Strategy

While the engineering process for double-cladding discharged fuel pins from a first-stage eB&B core is not yet devised, the potential of such a two-tier system from the physics perspective is substantial. By double-cladding the fuel and putting it into a second-stage core, the lifetime of the fuel can be approximately doubled as compared to what it would be in the eB&B core alone, with its lifetime limited again by the accumulation of material damage in excess of the 200 DPA constraint. While the fuel discharged from the second-stage core after reaching the 200 DPA limit still has significant fissile content, the addition of a third cladding layer to enable another round of burnup is likely not possible from the perspective of reactivity. However, by supplementing the discharged eB&B fuel with a stream of DU, the lifetime power production from the fuel can be increased by $\sim 17\%$ with little

⁵For now these two moderating materials are only briefly introduced. A much more thorough discussion of each moderator is presented subsequently in Section 6.5.

⁶The neutron balance in this case is performed without correcting for radial leakage, and thus the reactivity lifetimes derived here should be viewed as upper bounds on attainable performance.




Figure 6.26: Normalized spectra at the core midplane in the second-stage double-clad core with various numbers of fuel pins replaced by graphite moderator.




Figure 6.27: Neutron balance (solid) and material damage (dashed) in the second-stage double-clad core with various amounts of graphite and ZrH moderator added. Different lines correspond to differing numbers of fuel pins replaced by moderator.




Figure 6.28: Neutron balance at discharge versus the DPA-limited assembly lifetime in the graphite and ZrH moderated second-stage double-clad cores.

marginal cost. An alternative method of achieving similar gains was also shown through the introduction of limited amounts of moderating material to reduce the material damage accumulation rate, thereby shifting the lifetime-limiting constraint from material damage to reactivity. Finally, the strong sensitivity to the DPA constraint implies the double-cladding concept may become much more attractive as the DPA constraint is relaxed. Therefore, the double-cladding strategy is viewed as promising and it is judged that further investigation is merited.

6.5 The eB&B/SHR System

Regardless of the benefits of a second-stage double-clad core, the process required to add a second layer of cladding to highly irradiated fuel is unknown and may be extremely difficult. In light of this, an alternative method is explored in this section – placing the discharged eB&B fuel into a second-stage *thermal* spectrum reactor without any form of fuel reprocessing, reconditioning, or double-cladding. The motivation for doing this is generally the same as that for the double-clad core, but in fact there are other potential benefits. The thermal spectrum could be additionally beneficial from the perspective of proliferation resistance by having a low conversion ratio (i.e. not breeding additional fissile material) and reducing the

quality of the discharge Pu vector. Further, the kinetic behavior can potentially be made to be more favorable in a thermal spectrum system by increasing the neutron lifetime and reducing the coolant void worth. These aspects as well as others are investigated in this section as the design of a thermal-spectrum reactor capable of accepting discharged eB&B fuel is refined.

Selection of Coolant, Moderator, and Duct Materials

The choice of major constituents is the first fundamental task in the design of any core. In this particular case, we are restricted by a number of factors mostly related to the fact that we are dealing with fuel pins discharged from a B&B core, which has specific design strategies and goals. This section discusses the general considerations used to preliminarily design the second-stage thermal-spectrum core.

Coolant Selection

A general overview of the properties of different coolants which have either real-world experience or are widely accepted can be found in the dissertation of Novak [99]. Table 6.6 summarizes the process used to select the coolant by eliminating a variety of common options. The majority of coolants, except perhaps molten salts, would be eliminated based solely on the high power density that B&B cores are designed to have along with the tight lattice pitch. Molten salts, however, are generally not compatible with metallic fuel from the perspective that they operate at temperatures which challenge the integrity of the cladding due to fuel-cladding eutectic reactions in the vicinity of 650°C. Because the fuel type and cladding material cannot be changed due to considerations of neutron economy in the first-stage eB&B core, molten salt coolant cannot be accommodated. This brief analysis leaves essentially only one option for the coolant in the second-stage core – sodium. Because this is the same coolant as in the first-stage core, issues of chemical and temperature compatibility are trivially avoided. Additionally, sodium can handle the specific powers that these pins were designed for, and therefore coolability is assured in the second-stage core if the specific power density is kept similar. Finally, complications having to do with operation at high pressures can be avoided and all of the same plant systems as in the first-stage core can be utilized. Therefore, the second-stage thermal-spectrum core is selected to have sodium coolant.

Moderator Selection

With the coolant selected to be a non-moderating material, an effective solid moderator needs to be added to the core. The most commonly utilized solid moderator is graphite, which has been used in dozens of reactors of various types including most notably the Magnox, RBMK, and AGR designs. One lesser-known design is the sodium-cooled graphite-moderated reactor (SGR), two of which actually operated in the United States during the 1960s: (1) the

Table 6.6: Various coolants and their infeasibilities for use in a second-stage core utilizing discharged eB&B fuel.

Coolant	Infeasibility
Gas (CO ₂ , He)	- High specific power density - Tight hexagonal pitch
Water (H ₂ O, D ₂ O)	- Residual sodium contamination from first-stage core - High specific power density - Tight hexagonal pitch
Molten Salts (FLiBe, FLiNaK, etc.)	- Relatively high melting temperature in comparison to limiting fuel/cladding eutectic temperatures
Lead	- Tall core and high pressure drop - High specific power density

Sodium Reactor Experiment (SRE) built and operated by Atomics International in Southern California [109, 106] and (2) the Hallam reactor built by Atomics International and operated by Consumers Public Power District in Nebraska [27, 30]. Between these two reactors, there are 4-5 years of operational experience in SGRs, similar to that for other reactor types of current interest such as pebble-beds or Molten Salt Reactors (MSRs). The operational experience from these reactors was in general favorable. While the SRE suffered a core melt accident, this was related to the particular material employed as a sealant in the reactor coolant pumps which has since been phased out and is not in any way an inherent drawback of SGR technology [11].

The SGR was initially pursued to improve the economic performance of commercial power reactors, with a few often cited potential benefits [35]: (1) low pressure operation would enable safety benefits and the elimination of heavy forged components, (2) higher temperature operation would enable increased thermal efficiency, (3) no chemically reactive combinations of materials in the core [19], and (4) negative power feedback coefficients [18]. This is in contrast with LWRs, which operate at extremely high pressures, relatively low temperatures, and have the potential for hydrogen generation at elevated temperatures due to reactions between water and zircaloy. While the operational experience with the two SGRs proved these principles, ultimately the SGR concept was not commercially pursued for reasons which are not entirely clear, but in general do not appear to be based on any fundamental technical issues.

While graphite is a leading candidate for the moderator in the second-stage core, graphite systems tend to be quite large due to their poor moderating ability in comparison to lighter elements such as hydrogen. The resulting low power density of graphite-moderated cores is an economic drawback of SGRs that could potentially be overcome by employing a different

moderating material. One alternative that has seen extensive utilization is ZrH^7 , which has been used for years in the TRIGA reactors built by General Atomics [22] and has been studied for use in LWRs to increase their power density [45, 47]. Other reactors which have operated with ZrH as a moderator include the SNAP [131] and HTRE-3 [93] reactors, although experience in both cases was very limited. The experience in combining sodium coolant with ZrH is limited to the TOPAZ class of nuclear rockets, which have very limited operational experience in Russia and are of such a drastically different design that their experience is likely not useful to the current design study [69]. Regardless, ZrH is investigated in this study as an alternative to graphite primarily as a means to increase the power density. This alternative design will be called the sodium-cooled hydride-moderated reactor (SHR).

To understand the tradeoffs between the two moderator options, a series of unit cell calculations are performed to assess the core power density and achievable burnup. However, first the limitations associated with having a solid moderator are discussed. To take an assembly out of the eB&B core and incorporate a moderator without altering the constitution of the assembly, the moderator is included as blocks surrounding the assembly duct as shown in Figure 6.29. Each moderator block is an annular hexagon surrounded on all sides by a moderator liner (i.e. cladding) to protect the moderator from coolant ingress (which can have a significant reactivity penalty) and loss of hydrogen (in the case of ZrH moderator).




Figure 6.29: Example of block geometry for incorporating moderator around the discharged eB&B fuel assembly.

Because this geometry is very heterogeneous in its distribution of moderator, there is potential for power peaks near the interface between moderator and fuel pins which would

⁷The stoichiometry between Zr and H can vary, but the most commonly used ratio is $\text{ZrH}_{1.6}$.

make it difficult to achieve adequate cooling. Figure 6.30 shows the intra-assembly radial peaking factor (i.e. the peak pin power divided by the average pin power) in the graphite-moderated unit cell as the number of pins (represented as the number of concentric rings of pins) is reduced from its nominal value of 8. With 8 rings of pins (i.e. 169 pins in total) the intra-assembly radial peaking is nearly 3.5 – far too large to be acceptable. In fact, in order to obtain reasonable radial power peaking the number of rings needs to be decreased down to just 2 (i.e. a 7-pin cluster), which is in fact the geometry that was employed in the Atomic International SGR designs. Such results imply that it will be necessary to disassemble the discharged eB&B assemblies into pins and reconstitute a smaller assembly.




Figure 6.30: Intra-assembly radial peaking factor for increasing number of pin rings in the assembly in the case of a graphite-moderated unit cell.

This requirement is less than ideal because the discharged assemblies will be highly irradiated and difficult to handle, likely requiring the process to be performed under sodium to maintain acceptable temperatures. Regardless, this process has two large benefits in comparison to the alternatives. First, it will not require any amount of chemical processing, separation of fissile material, or cladding breach as would be required by a traditional re-processing scheme or the less-invasive ‘reconditioning’ schemes. Second, it will still be much less complicated than double-cladding, which essentially includes the disassembly and re-constitution steps as sub-processes in the overall double-cladding process. Therefore, while the need to reduce the assembly size down to 7 pins is undesirable, it might still be seen as more viable than other methods which have been proposed. In addition, it is possible that

the assembly design for the first-stage eB&B core could be altered with this requirement in mind so that deconstruction and reconstitution is easy to achieve. While this study simply assumes an assembly design typical of traditional SFRs, modifying the design to better suit the transition between the first- and second-stage cores could greatly increase the viability of such a concept.

With the new 7-pin assembly clusters, a study is performed to determine the optimal moderator dimensions such that burnup can be maximized and the core remains under-moderated. In this study, assembly dimensions are held constant, with the assembly duct kept at its nominal thickness of 3.5 mm and assumed to be made of HT9, as the thickness of the moderator block is increased parametrically. The moderator liner is assumed to be 1 mm thick, since it does not have to maintain any load, and made of zircaloy to avoid parasitic absorptions that would be stronger in HT9. The fuel is taken as that discharged from the uniformly-enriched 3.0 m eB&B core with 6w% feed fuel. The neutronic model consists of a radially-infinite repeated lattice of axially-finite unit cells with 9 axial segments corresponding to the same material discretization used in the eB&B neutronic model. The fuel is burned at the same average specific power density as in the eB&B core (which is also the same as was used for the double-clad study), implying coolability as long as the radial and axial peaking factors in the second-stage core are kept within the bounds of those from the first-stage eB&B core, something which should be easy considering how high the radial peaking factor is in B&B cores in the first place.

Figure 6.31 shows the attainable cycle length such that the condition $k_{eff}^{ave} = \frac{\int k_{eff}(t)dt}{\int dt} = k_{eff}^{target}$ is satisfied as a function of the moderator block outer dimension (flat-to-flat radius). The target multiplication factor is parametrically varied to account for different levels of radial leakage that might be present in a finite-dimensional core. Stark differences are found between the two variants. First, it is noted that the SHR displays a prominent peak at the optimal moderator FTF = 3.6 cm, while the SGR shows a very broad peak (and thus a much wider design space) roughly centered at 13.0 cm. These two differences are easily understood based upon the slowing down properties of each moderator, namely that hydrogen has a mean lethargy gain per collision (ξ) of 1.0 while graphite's is 0.158, implying the marginal impacts of more graphite volume are much smaller than that for hydrogen. This is a general characteristic of all graphite-moderated reactors that implies very large core volumes. On the other hand, achievable cycle lengths are multiple years longer in the SGR owing to the lower absorption cross-sections of graphite than zirconium shown in Figure 6.32.

These lower cross-sections contribute directly to differences in the core multiplication factors, shown in Figure 6.33 for both variants at their optimal moderator geometries. For thermal-spectrum systems, the two variants have interesting reactivity evolutions. Namely, k_{eff} actually *increases* with time during the early phases of burnup in a manner similar to a breeder core, although each core has a conversion ratio near $\sim 0.2-0.3$. The lower absorption in the SGR allows for it to have the higher maximum k_{eff} , but the SHR has an advantage in that its BOL k_{eff} is roughly 5,000 pcm higher than in the SGR, making it such that the initial neutron deficit is lower in the SHR.




Figure 6.31: Attainable cycle length in the SGR and SHR cores as a function of the moderator dimensions for various k_{eff}^{target} values.

The cause of the reactivity increase with time during the early burnup phase can be better understood by examining the axial power evolution, as shown in Figure 6.34. When the fuel is first put into the second-stage thermal-spectrum core, its axial power profile is extremely unique. The majority of power is pushed to the axial edges during the first couple years of burnup, shifting to the more familiar cosine-esque profile after ~ 5 years. The cause of this power shift is the presence of large amounts of fission products near the fuel axial-midplane built up from the first-stage eB&B core. Because cross-sections are lower in the fast-spectrum, these fission products did not significantly burn out while the fuel resided in the eB&B core. However, when exposed to thermal neutrons, the fission product cross-sections are orders of magnitude higher such that they act as strong neutron poisons until they experience enough fluence to be burned out, in a similar manner to the burnable absorbers intentionally added to LWR fuel assemblies. It should be emphasized that in an actual core design with assemblies of different burnups located next to each other, the k_{eff} evolution and axial power profile may look different (as discussed in shortly), but this discussion provides insight into the expected full-core behavior.

The major factors driving the SGR to have lower reactivity at BOL can be better understood by examining the factors of the six-factor formula, as provided in Table 6.7. While the SGR performs significantly better from the perspective of resonance escape, p , (as expected from the moderator cross-sections in Figure 6.32) and slightly better from the perspective of thermal utilization, f , (again due to the lower moderator cross-sections), the SHR performs much better in the number of neutrons produced per absorption, η , and the fast fission factor, ϵ . Additionally, the SHR has lower axial leakage (radial leakage is not estimated in radially-infinite calculations) because the diffusion length is smaller.




Figure 6.32: Capture cross-sections for various isotopes in each of the moderators.

Table 6.7: Components of the six-factor formula for the SGR and SHR at BOL.

	η	f	p	ϵ	P_{NL}^f	P_{NL}^t
SGR	1.3097	0.7257	0.7646	1.2624	0.99	0.98
SHR	1.4666	0.7004	0.6705	1.4036	1.00	1.00
%diff ((SGR-SHR)/SGR)	-12.0	3.5	12.3	-11.2	-1.0	-2.0

An explanation for the higher values of η and ϵ in the SHR can be found by comparing the spectra between the two cores, shown in Figure 6.35. The SHR, while still displaying the expected thermal peak, has a notably higher fraction of fast neutrons than does the SGR. This poor moderation from ZrH as compared to graphite is largely the result of the unique moderating properties of ZrH [91]. On an atomistic scale, hydrogen bound in a zirconium lattice acts similar to a harmonic oscillator with discrete energy states, the smallest of which is at roughly 0.14 eV. This implies that the hydrogen in ZrH will upscatter any neutrons below this energy such that moderation below 0.14 eV is very poor in the absence of another moderating material. In ZrH systems built or proposed to date (i.e. TRIGAs and LWRs), the presence of water as coolant mitigates this effect such that it does not play an important role, but in the case of the SHR there is no other material to effectively moderate these low-energy neutrons. This unique characteristic of the SHR directly leads to the harder spectrum as compared to the SGR and enables a larger value of ϵ .




Figure 6.33: Evolution of the multiplication factor in the SGR and SHR unit-cell models for the optimal moderating geometries.

While the SGR performs better overall from the perspective of reactivity, the fact that it needs a larger moderator volume is a drawback as it reduces the core power density. Table 6.8 compares the power densities estimated for the SGR and SHR with those of other commercially-deployed reactor systems. Values for the PWR, BWR, CANDU, and HTGR were taken from Duderstadt & Hamilton [31] and values for the SFRs were taken from Kim et al. [80]. The B&B power density is taken from the whole-core calculations for the 3.0 m tall core (in the case of the shorter cores examined in Section 6.3 the power density will simply scale inversely with core height). The SGR and SHR power densities are estimated by assuming that the fuel will have the same *specific* power density as in the B&B core and that the radial peaking factor will be reduced from 2.2 in the B&B core to 1.5 (typical for a thermal-spectrum system). While SFRs generally have high power densities due to the efficient heat transfer of liquid sodium, the B&B core in fact has a relatively low power density due to the large number of low-powered blanket batches⁸. Comparatively, the SGR and SHR cores perform even worse due to the necessary moderator volumes. In addition, the fact that the moderator is separate from the coolant makes the power density lower than either a PWR or BWR. The SGR performs particularly poorly, with a power density 6x lower than even the worst of any widely-deployed power reactor system, the CANDU.

⁸See Chapter 3 and Appendix B for further discussion on the power density of B&B cores.




Figure 6.34: Axial power profile evolution in the SHR with burnup.

Based primarily on this discussion of power density, it is judged that the SGR will have a tough time competing economically. Therefore the remainder of this section focuses on evaluating the performance of the SHR, which has a power density sufficiently high to be potentially viable. It should be noted, however, that the SGR has certain advantages that are foregone by selecting the SHR. First, the SGR performed somewhat better from the perspective of reactivity, and thus selecting the SHR implies a slight penalty in the achievable burnup. This topic will be discussed further in a comparison of fuel cycle performance in Section 6.7. Secondly, and probably more significantly, the experience with graphite as a moderator is much more extensive than that with ZrH, especially with sodium as a coolant. Chemical interactions between sodium and graphite are known to be benign⁹, and the behavior of graphite under high fluence is at least *better* understood (though still quite uncertain).

While ZrH presents a much more attractive design from the perspective of power density, this choice does come with a number of drawbacks mainly related to unknowns of material integrity. The primary unknown is how ZrH will behave under high fluence. Other important unknowns have to do with the ability for hydrogen to migrate out of the zirconium lattice. Such concerns are what drives the need to contain ZrH within a liner, or else the hydrogen (i.e. the moderating ability) will be lost with time to the coolant [140]. Even without diffusing into

⁹Though this does not enable removal of the moderator liner in the SGR, as the logging of sodium in the graphite moderator has a significant reactivity penalty and should be avoided.




Figure 6.35: Neutron spectra averaged over the entire assembly in the SGR and SHR unit-cell models.

the coolant, the migration of hydrogen under temperature gradients within the moderator blocks could lead to pressurization of the moderator cans or hydrogen embrittlement of the moderator liner walls. In addition, while the ZrH delta-phase (i.e. $\text{ZrH}_{1.6}$) has been found not to swell under irradiation [23], temperature gradients could cause hydrogen migration such that local phase transitions occur which lead to dimensional changes. Therefore, while this thesis chooses to pursue the SHR due to the economic infeasibility of the SGR, many questions concerning the viability of ZrH remain. While it would be desirable to keep the ZrH moderator blocks in the core for the lifetime of the reactor, the possibility of replacing the moderator blocks at refueling could mitigate these concerns.

Duct Selection

Table 6.9 shows the fraction of neutrons absorbed in each constituent of the SHR at BOL and EOL (in this case taken as ~ 8 years according to Figure 6.31). Of importance is how many neutrons are absorbed in the duct as compared to the moderator liner, which have roughly similar volumes. The much higher absorption in the duct is because the duct is made of HT9, whereas the moderator liner is assumed to be made of Zircaloy based on the favorable experience with the SRE and Hallam reactors [30, 18, 11]. The duct was assumed to be made of HT9 simply because that was the material in the first-stage eB&B

Table 6.8: Power densities for a variety of power reactor systems.

System	Power Density (W/cm ³)
PWR	101
BWR	56
CANDU	12
SFR (metal)	303
SFR (oxide)	215
HTGR	8
B&B	75
SGR	2
SHR	32

core. While HT9's swelling resistance makes it a good option for fast-spectrum systems, the absorption cross-sections for steels are much higher in the thermal-spectrum than those of other structural materials, particularly Zircaloy.

Table 6.9: Fraction of neutrons absorbed each major core constituent in the SHR.

	BOL (%)	EOL (%)
Cladding	5.2	3.5
Coolant	1.5	1.4
Duct	17.2	15.9
Fuel	59.3	63.6
Moderator	15.2	14.1
Moderator Liner	1.7	1.6

Figure 6.36 shows the profound impact of replacing the HT9 duct with Zircaloy in the SHR unit-cell. The entire trend is shifted up by roughly 13,000 pcm such that the BOL k_{eff} is now significantly supercritical and the attainable cycle length is extended from 2.8 to 16.5 years (assuming $k_{eff}^{ave} = 1.03$). This result strongly drives the selection of Zircaloy over HT9 or other iron-based materials. Because assemblies are smaller and lighter in the SHR than the eB&B core, no issues associated with structural strength are anticipated in a Zircaloy duct if the thickness of the duct is kept the same, making the use of Zircaloy likely feasible. In some way, this provides a silver lining to the fact that the assemblies need to be deconstructed and rearranged into 7-pin clusters between their time in the eB&B and SHR cores, because

such a reassembly process enables the duct material to be swapped with no extra steps.¹⁰ While the cladding is still made of highly-absorbing HT9, the smaller volume would make the impact of switching the cladding material to Zircaloy relatively minor. Furthermore, the irradiation database for Zircaloy does not extend up to the 200 DPA that would be necessary in the first-stage eB&B core, precluding its use as cladding in the first place if de-cladding processes are to be avoided.




Figure 6.36: Multiplication factor evolution in the SHR with two different duct materials.

Summary of Materials Selection

The previous few sections have narrowed the design space for a second-stage thermal-spectrum core capable of accepting discharged eB&B fuel as its feed. Based mostly on excluding other possibilities for thermal-hydraulic reasons, the coolant was selected to be sodium. In selecting a moderator, practical issues such as the desire to avoid a two-fluid core and real-world operational experience led to examination of solid moderating materials. While graphite was found to be preferential from the perspective of reactivity, its exceptionally low power density ultimately led to the selection of ZrH as a solid moderator. Power

¹⁰It should also be noted that, depending on the lifetime of the ZrH moderator blocks, it might make more sense to eliminate the duct altogether and instead replace the moderator blocks integrally along with the fuel. Such a possibility may be minorly beneficial from the perspective of power density as well, but is not investigated further in this dissertation to declutter the analysis.

peaking concerns forced the design to employ only 7-pin clusters, which necessitates dismantling discharged eB&B fuel assemblies before they can be reconfigured and put into the second-stage core. Finally, the reactivity penalty associated with HT9 drives the duct in the SHR assemblies to be made of Zircaloy, which is enabled during the reconfiguration between cores.

Sensitivity to DPA Constraint

Similar to the double-clad case, it is anticipated that if the DPA constraint in the first-stage eB&B core were relaxed to 300 DPA, the fuel lifetime would change. The required feed enrichment would decrease and the residence time would increase, the impacts of which are not easily predicted for the SHR. The lower feed enrichment results in lower fissile content towards the axial periphery, but greater burnup in the axial center will increase plutonium buildup. Conversely, the higher burnup increases the fission product load towards the axial center, which will more strongly impact the SHR than the double-clad core. Figure 6.37a shows how this balance plays out in terms of the multiplication factor in the SHR.

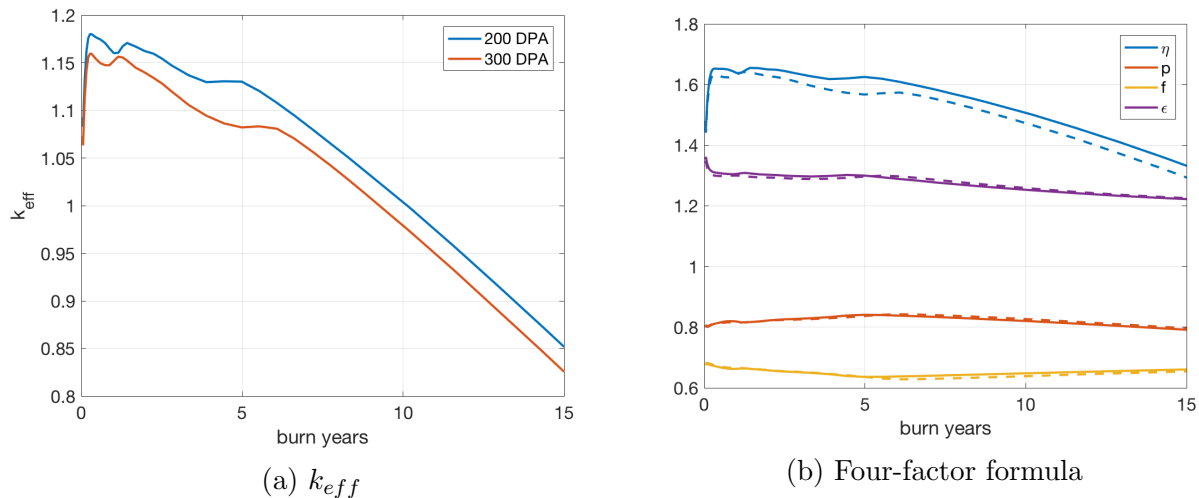


Figure 6.37: Evolution of k_{eff} and the four-factor formula parameters in the SHR with different DPA constraints in the first-stage eB&B core. Solid lines correspond to 200 DPA while dashed lines correspond to 300 DPA.

Relaxing the eB&B constraint to 300 DPA ends up having a ~ 2000 pcm reactivity *penalty*, the source of which can be determined from examining the components of the six-factor formula provided in Figure 6.37b. At BOL, the reactivity penalty is driven partially by a reduction in ϵ , which results from the lower fissile content towards the axial periphery in the 300 DPA case. As burnup progresses, the fission region shifts from the axial periphery, where primarily U-235 is being burned, to the core center, where primarily Pu-239 is being burned,

and the reactivity penalty in the 300 DPA case becomes mostly dominated by a reduction in η . Because Pu-239 has a significantly lower fission-to-capture ratio in the thermal spectrum, as shown in Figure 6.38, η decreases correspondingly¹¹. The shift of power towards the axial center also leads to a notable decrease in the thermal utilization factor as burnup progresses, caused by absorption in the larger mass of fission products built up from the first-stage core when burnup to 300 DPA is allowed.




Figure 6.38: Fission-to-capture ratio for U-235 and Pu-239 as a function of energy.

This decreased reactivity shortens the life in the second-stage SHR core, but the total fuel lifetime sees a slight increase due to the longer lifetime in the first-stage core, as shown in Table 6.10. In comparison to the strongly-positive sensitivity with the double-clad second-stage core, the eB&B/SHR system is relatively insensitive. Instead, relaxing the DPA constraint simply allows for the share of energy extracted in the first-stage eB&B core to be larger such that the appeal of the SHR becomes marginally less.

SHR Core Performance

To more realistically evaluate SHR core performance, a new model is constructed to partially account for the interplay between adjacent assemblies that would be seen in a full-core design. This new model contains three batches represented by a repeated unit cell of three assemblies, as presented in Figure 6.39. To achieve the equilibrium composition, criticality cycles are

¹¹While $\bar{\nu}$ is larger in Pu-239 than U-235 even in the thermal spectrum, its difference is not large enough to make up for the difference in the fission-to-capture ratio.

Table 6.10: Assembly lifetimes in the first- and second-stage cores for various DPA constraints.

DPA Constraint	Feed enrichment (w%)	First-stage lifetime (yr)	Second-stage lifetime (yr)	Total assembly lifetime (yr)
200	6.0	13.9	16.5	30.4
300	4.5	18.7	14.4	33.1

simulated while repeatedly discharging the oldest batch and charging fresh fuel in its place until the compositions from cycle n match those from cycle $n - 1$, all the while iterating over the cycle length to achieve a minimum cycle $k_{eff} = 1.03$. To keep the analyses tractable, only the case with axially-uniform enrichment in the first-stage eB&B core is examined presently.

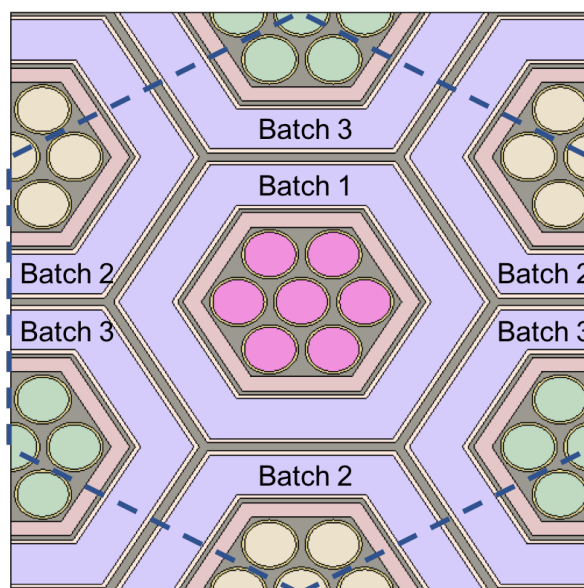


Figure 6.39: Schematic of the three-assembly unit-cell model used to simulate the interplay between batches in a full-core SHR. Dashed line indicates the bounds of the repeated unit cell.

Criticality

The converged k_{eff} evolution is shown in Figure 6.40, where the critical cycle for the three-batch core is found to be approximately 4.5 years (for a total assembly residence time of 13.5

years). During the first few days of burnup, a sharp drop is seen due to the buildup of Xe-135¹². Again, a strong increase in reactivity is seen during the early burnup phase as fission products from the eB&B core burn out until reaching a peak of roughly 1.09, somewhat lower than the peak from the single-assembly model due to the presence of older assemblies alongside newer ones. Using the k_{eff} evolution from the single-assembly model along with that from the three-assembly model, it is possible to estimate the critical cycle length for any number of batches through extrapolation of the trend in Figure 6.41. From this estimate, it is found that having ~ 5 batches should allow for the cycle length to approach that of an infinite batch system.




Figure 6.40: Multiplication factor evolution in the SHR as determined with the three-assembly model.

Power Profile

Figure 6.42 shows the axial power profile of the SHR determined using the three-assembly model. Comparing this with the single-assembly model axial power profile of Figure 6.34, it is seen to be somewhat tempered by the interplay between batches. For instance, the power peak at the bottom of the assembly at BOEC is lowered in comparison to that in

¹²While the fuel is assumed to be vented, the fraction of fission gases released with burnup from such highly burned fuel is not well understood, and therefore this study makes the assumption that Xe-135 and other gaseous fission products are not removed.




Figure 6.41: Linear extrapolation of the critical fuel residence time ($\min k_{eff} = 1.03$) versus number of batches along with the estimated cycle length for an infinite number of batches (i.e. $k_{eff}^{ave} = 1.03$).

the single-assembly case because the power profile in older batches has shifted more towards the axial center. A similar interaction also serves to flatten the cosine peak towards EOEC. Examining the axial power profile in each batch separately, as shown in Figure 6.43, reveals that the coupling between different batches is quite tight, as the presence of just one fresh batch forces the older batches to have an inverted axial power profile even though they would be strongly center-peaked on their own. This demonstrates the strength of fission products built up in the first-stage core and their “burnable absorber” effect in the second-stage SHR.

The unique power profile, although unintentional, is actually a strong benefit of the eB&B/SHR system. By having the fission products act as burnable absorbers, the power profile is pushed towards the axial periphery where fissile content that might otherwise be underutilized can be effectively burned. In addition, the core reactivity swing is reduced so that requirements on the control system are lessened. These benefits, particularly the burning of fuel towards the axial periphery, would be largely destroyed by standard control systems where elements are inserted either from the top or bottom of the core. Thus the SHR concept calls for some other form of control, for instance leakage-based control drums or the use of control rods that are always either fully removed or inserted through the full height of the core. In addition, it might be possible to consider adding further burnable absorbers to the axial center of the core within the moderator blocks to further limit the excess reactivity




Figure 6.42: Axial power profile integrated over all assemblies for various points in the equilibrium cycle as determined with the three-assembly model.

towards the beginning of cycle and drive higher burnup in the axial periphery. Such studies are outside the scope of this preliminary design analysis.

Reactivity Coefficients

While the second-stage double-clad core is assumed to have reactivity feedback mechanisms similar to those of the first-stage eB&B core, the SHR has completely different characteristics and thus a preliminary assessment of its unique kinetic parameters is in order. To approximately account for the interplay between assemblies of different batches, the three-assembly model is utilized in this assessment.

Before examining the specific reactivity coefficients, some generic kinetic parameters are presented in Table 6.11. Because the delayed neutron fraction is only minorly impacted by neutron energy, β_{eff} is mostly unchanged in the SHR as compared to first-stage eB&B core (see Figure 6.12) because the mix of fissioning isotopes is similar (i.e. mostly Pu-239 with a smaller fraction of U-235). Conversely, the neutron lifetime, ℓ , is dependent almost entirely on the energy of fissioning neutrons, and as such the thermal-spectrum SHR core has a neutron lifetime roughly two orders of magnitude larger than the eB&B cores examined in this dissertation (see, for instance, Table 2.3). This longer neutron lifetime generally provides a benefit from the perspective of slowing down transient events, although such benefits are




Figure 6.43: Axial power profile in each batch at BOEC as determined with the three-assembly model.

difficult to quantify explicitly.

Table 6.11: Kinetic parameters of the SHR core.

	BOEC	EOEC
β_{eff}	4.34×10^{-3}	3.96×10^{-3}
ℓ (s)	2.78×10^{-5}	2.83×10^{-5}

Table 6.13 shows the linearized reactivity coefficients for the SHR at BOEC and EOEC. All reactivity coefficients are calculated using direct perturbations of the material temperatures and/or densities with a Forward Euler finite difference approximation. To highlight some unique aspects of the SHR, reactivity coefficients are broken down into impacts resulting from changes in density and temperature separately. For density perturbations, the volumetric thermal expansion coefficients in Table 6.12 are used along with a temperature increase of 500 K to determine the perturbed densities. For temperature perturbations, an increase of 300 K is used for all materials without altering the material density. To calculate the radial expansion (grid plate) coefficient, the spacing of assemblies is adjusted according to a volumetric expansion of the lower gridplate for a temperature increase of 150 K.

Table 6.12: Volumetric thermal expansion coefficients used in the calculation of reactivity feedback coefficients.

Component	Material	$\frac{\Delta V/V}{K}$ (1/K)	Source
Cladding	HT9	5.518×10^{-5}	[85]
Coolant	Liquid Na	2.514×10^{-4}	[25]
Duct	Zircaloy	2.160×10^{-5}	[151]
Fuel ^a	U-10Zr	-	-
Moderator	ZrH	4.556×10^{-5}	[140]
Moderator Can	Zircaloy	2.160×10^{-5}	[151]

^a The volumetric expansion for the fuel is taken to be the same as for the cladding due to the high levels of burnup forcing the fuel into tight mechanical contact with the cladding.

As expected based on the absorption fractions in Table 6.9, the reactivity coefficients associated with cladding, coolant, duct, and the moderator liner are all relatively small, particularly their Doppler coefficients. For the same materials, the density coefficients with temperature are all slightly positive, as expected based on the fact that they are primarily parasitic absorbers from a neutronic perspective.

The only negative feedback mechanisms of moderate strength are those for the gridplate expansion, fuel doppler, and fuel density (which, even still, changes sign over the course of the burnup cycle). The gridplate expansion feedback is expected to be negative since it introduces more coolant into the assembly gaps and (oftentimes) allows for increased axial leakage. In the SHR, the increase in axial leakage is only barely seen because the axial leakage is so low in the first place. Instead, the dominant role is played by increased parasitic absorptions in the interstitial coolant, which is aided in its strength by the relatively large thermal expansion coefficient of HT9. The fuel doppler feedback, as usual, largely results from increased resonance absorption in U-238, leading the resonance escape probability at BOEC to decrease from 0.6627 to 0.6607 as the fuel temperature is increased by 300°C.

While the fuel density coefficient is negative at BOEC, it ends up being slightly positive by EOEC, being the only feedback mechanism to change sign over the burnup cycle. To understand why, the components of the four-factor formula are presented in Table 6.14. At BOEC, it is seen that as the fuel density increases (i.e. the fuel temperature decreases), the strongest impact is an increase in the number of fast fissions, and combined with the increases in η and f , these factors outweigh the decrease in the resonance escape probability. However, at EOEC the change in η with increasing fuel density becomes slightly negative such that the net change is more or less zero. This decrease in η at EOEC with fuel density increase is attributed to an energy self-shielding effect on the low Pu-239 resonance, as

Table 6.13: Reactivity coefficients associated with temperature and density changes in the SHR at BOEC and EOEC.

		$\alpha \times 10^6 \left(\frac{dk/k}{K} \right)$	
		BOEC	EOEC
Doppler	Cladding	-0.2	-0.3
	Coolant	0.0	0.0
	Duct	-0.3	-0.5
	Fuel	-8.2	-8.4
	Moderator	76.5	47.8
	Moderator Can	-0.3	0.3
Density	Cladding	2.2	1.8
	Coolant	3.3	3.2
	Duct	0.7	0.4
	Fuel	-5.8	0.9
	Moderator	-1.6	-2.6
	Moderator Can	0.6	0.7
Gridplate Expansion		-6.1	-5.5
Net		60.8	37.8

outlined in Figure 6.44. When the fuel density increases, the large, low-energy resonance has a tendency to reduce the number of neutrons in the thermal Maxwellian and instead increases absorptions in the upper resonance and the lower unresolved resonance regions, which have lower fission-to-capture ratios (see Figure 6.38). This effect is not prominent at BOEC because the proportion of fissions taking place in Pu-239 is much lower while the power profile is peaked towards the axial periphery.

By far the strongest feedback mechanism is the moderator Doppler, which is unexpected considering the Doppler broadening mechanism can only serve to increase absorptions. Furthermore, ZrH has relatively few resonances in the first place. But in fact the mechanism is so strong that it would only take a $\sim 60^\circ\text{C}$ increase in the moderator temperature to result in a $\$1.0$ reactivity insertion. The root of this mechanism can be better understood by examining the spectral change as the moderator temperature is increased, as shown in Figure 6.45. The increase in moderator temperature leads to an upward shift in the thermal Maxwellian due to the increase in upscattering. While this phenomenon is not unique to the SHR, what is unique is the role of Pu-239 in this system. As the Maxwellian shifts upwards, it moves closer to the low-lying resonance of Pu-239 near 0.3 eV which increases overall fission probability (i.e. f increases). Because there is so much Pu-239 in this thermal-spectrum system, the effect is quite strong. While this feedback mechanism is referred to as the ‘Doppler’ effect

Table 6.14: Components of the four-factor formula for the SHR at BOEC and EOEC as the fuel density is increased according to a 500°C temperature decrease.

		η	f	p	ϵ
BOEC	Nominal	1.5373	0.8104	0.6627	1.2953
	Perturbed	1.5402	0.8127	0.6500	1.3181
Difference (%)		0.2	0.3	-1.9	1.8
EOEC	Nominal	1.5292	0.8241	0.6461	1.2633
	Perturbed	1.5274	0.8263	0.6338	1.2837
Difference (%)		-0.1	0.3	-1.9	1.8




Figure 6.44: Spectrum change as the fuel density increases at EOEC in the SHR, along with the Pu-239 fission cross-section.

due to custom, it might more accurately be described as the ‘Maxwellian shift’ feedback.

It is noted that positive moderator Doppler coefficients were also seen in the SRE and Hallam reactors, although stability was achieved in those systems due to the differing time constants for temperature increases between the fuel and moderator [18, 19]. While a moderator temperature increase leads to a positive reactivity insertion, heat must be conducted through fuel, clad, coolant, duct, and moderator liner before the reactivity insertion is af-




Figure 6.45: Spectrum in the nominal SHR along with the difference when the moderator temperature is increased without density change.

ected. Conversely, fuel temperature increases are seen immediately upon power increase such that the negative fuel Doppler coefficient acts quickly to arrest transients before they progress to a stage at which the moderator temperature can be notably impacted. Thus while the isothermal temperature coefficient is positive, the *power* coefficient is negative. While it is expected that similar phenomena will prevail in the SHR, confirmation of such behavior can only come from mechanistic transient analyses, which are outside the scope of this dissertation.

Instead, Figure 6.46 provides a sensitivity study showing how the moderator Doppler feedback changes with the size of the moderator block. These results show a strong sensitivity and indicate that the moderator doppler coefficient can be made negative if the moderator block outer FTF is reduced to ~ 3.4 cm. Shown in the same plot is an estimate for the attainable equilibrium cycle length for the reduced moderator dimensions, demonstrating that reducing the moderator down to 3.4 cm would only result in a cycle length penalty of $\sim 3\%$. This small reduction in the cycle length would have minimal impact on the eB&B/SHR mission while having large benefits for core stability, and thus a reduction of the moderator dimensions away from its optimal geometry is likely warranted.

The fact that the moderator density coefficient is consistently negative confirms that the core is properly under-moderated, as was inferred from the studies of the optimal moderator block dimensions in Figure 6.31b. While this feedback mechanism also shows sensitivity




Figure 6.46: Moderator doppler coefficient estimated for cases with varying moderator block dimensions at BOEC, along with the net coefficient assuming other feedback mechanisms remain approximately unchanged. Also shown is the estimated cycle length using the single-assembly model and $k_{eff}^{ave} = 1.03$.

to the moderator block dimension, in general such sensitivity is far smaller than that for the moderator Doppler feedback, and the moderator density coefficient never approaches strengths large enough to mitigate the Doppler coefficient.

Use of Axially-Graded Enrichment Patterns

Because it was found previously that the eB&B core could be sustained with lower total enrichment requirements by axially-grading the enrichment pattern of the feed fuel (see Section 6.3), this section briefly investigates the eB&B/SHR system when utilizing axially-graded enrichment patterns. Furthermore, because the double-clad core was investigated using axially-graded enrichment patterns, performing a similar investigation for the SHR core allows for a more direct comparison between the two second-stage cores to be made. This section repeats some of the previous SHR analyses utilizing a single-assembly model of the roughly-optimized 1289 enrichment pattern (see Figure 6.9) with enrichment high enough to meet the 200 DPA constraint in the first-stage core (6.7w%, as compared to 6.0w% in the uniformly-enriched variant).

Figure 6.47a compares the k_{eff} evolution in the SHR for the uniform and axially-graded

cases. The reactivity penalty in the axially-graded enrichment pattern results in a decrease in the cycle length from 16.5 to 11.0 years. The cause of the reactivity penalty changes with time as the power profile shifts, resulting in the primary fissioning isotope changing with time and associated variations in η and ϵ (see the previous discussion of the DPA constraint sensitivity). Even with the reactivity penalty, however, the 11.0 year lifetime makes the axially-graded SHR potentially viable as an effective second-stage core. However, the fuel lifetime in the SHR is lower than it would otherwise be in a double-clad core by ~ 3.4 years, representing a tradeoff between the mechanical difficulty of the double-cladding process and the worsened resource utilization of the SHR.

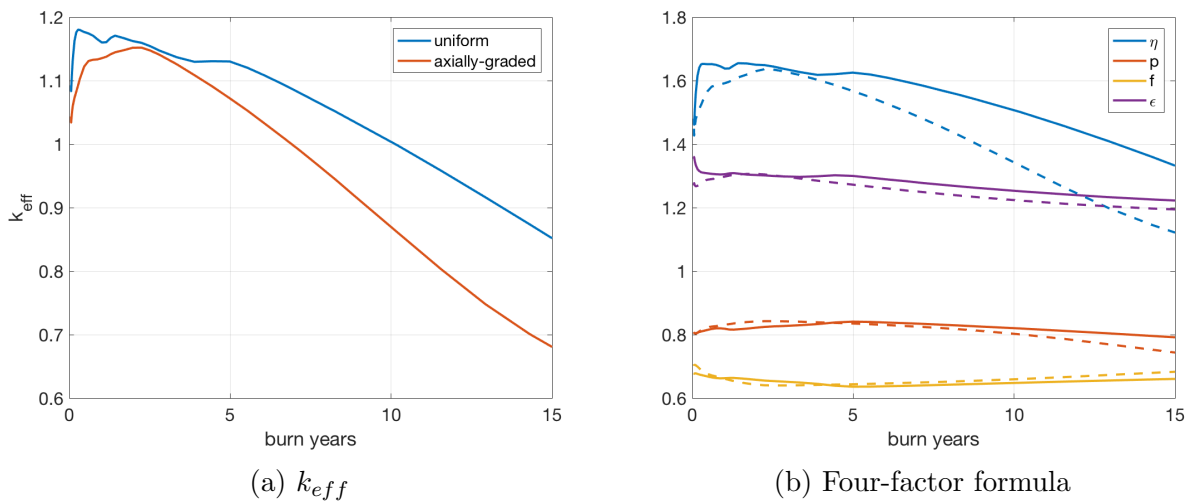


Figure 6.47: Evolution of k_{eff} and the four-factor formula parameters in the SHR with different feed enrichment patterns in the first-stage eB&B core. Solid lines correspond to uniform 6.0w% enrichment while dashed lines correspond to enrichment pattern 1289 with 6.7w% enrichment.

Summary of the eB&B/SHR System

As an alternative to the double-clad second-stage core, a new thermal-spectrum design based on sodium coolant, solid moderator, and zircaloy structural material has been proposed and preliminarily evaluated. While either graphite or ZrH moderator could potentially meet the needs of this system, ZrH is preferred on the basis of power density, and the resulting design is termed the SHR. However, graphite may offer somewhat improved burnup performance if its low power density can be tolerated in the SGR variant. The evolution of the power profile and multiplication factor in this second-stage core is seen to be heavily influenced by the presence of fission product buildup in the first-stage eB&B core such that uniquely favorable behavior results. As burnup progresses, the power profile shifts from being peaked towards

the axial periphery towards the more typical cosine-esque shape, enabling improved burnup of fissile content away from the core center. Additionally, the fission product inventory at BOEC acts similar to a burnable absorber in limiting k_{eff} such that control requirements are reduced as compared to other thermal-spectrum systems.

An analysis of the reactivity coefficients of the SHR show unique behavior stemming from the interplay between the thermal Maxwellian and the low resonance in Pu-239 which has the potential to make the moderator Doppler coefficient strongly positive for moderator geometries optimized for burnup. It has been shown that this reactivity coefficient can be mitigated by reducing the moderator dimensions slightly, which induces a relatively minor penalty in the achievable burnup. In the absence of this strongly-positive moderator Doppler coefficient, the reactivity response will otherwise be dominated by strongly-negative fuel Doppler and grid plate expansion coefficients such that favorable transient performance should be achievable.

By employing the SHR as a second-stage core, up to ~ 16.5 years of additional burnup can be extracted from the discharged eB&B fuel. This is similar to what can be achieved from a double-clad second-stage core, but in this case the fuel lifetime is limited by reactivity considerations whereas the double-clad core was limited by the accumulation of material damage. This difference leads to very different sensitivities to the material damage constraints such that if the DPA constraint were relaxed, the SHR becomes marginally less attractive than the double-clad core since the double-clad core has a much longer reactivity lifetime. However for the current 200 DPA constraint, the SHR offers similar advantages to the double-clad core with lower operational complexity at the stage of transferring fuel between the first- and second-stage cores.

6.6 Potential for Three-Stage Fuel Cycles

While the SHR was initially proposed as an alternative to the double-clad second-stage core, it is recognized that each system has unique properties that could be complementary. For instance, due to the higher fission product cross-sections in the thermal spectrum, burning the discharged eB&B fuel in the SHR core allows for higher burnout of fission products than would have resulted in the fast-spectrum. This higher burnout raises the possibility that fuel burned in the second-stage SHR core could potentially be useful again if put *back into* a fast-spectrum system, where the fission product cross-sections are lower and breeding on the remaining U-238 can be used to increase fissile content. Conversely, because the fuel lifetime in the second-stage double-clad core is generally limited by material damage constraints and not reactivity, there is the possibility that further energy may be extracted from discharged double-clad fuel if it were put into a thermal-spectrum system where irradiation damage concerns are lower. These possible three-stage systems are the subject of brief investigation in this section.

Possibility to Use SHR Discharge in Double-Clad Core

Because the cladding on discharged SHR fuel has already reached its DPA limit in the first-stage eB&B core, putting the fuel back into a fast-spectrum core would require the addition of a second cladding layer. This section investigates this possibility by taking the discharged fuel from the SHR core at various levels of burnup, adding a second layer of cladding, and examining the neutron balance in a fast-spectrum core using the same geometry as outlined in Section 6.4.

Figure 6.48 shows the multiplication factor evolution in this third-stage double-clad core along with the associated neutron balance for different residence times in the second-stage SHR. While k_{eff} consistently increases with burnup, the rate of increase is slow as is consistent with previous studies of fuel at higher burnups. For fresh fuel, the ability to achieve a net-zero neutron balance relies upon a rapid increase in k_{eff} at BOL that is not attainable with the highly-burned SHR discharge fuel. While the neutron balance in the third-stage core can be improved by discharging fuel from the SHR earlier, doing so reduces the attractiveness of the SHR in the first place. In fact, in order to achieve a net-zero balance in the third-stage double-clad core, fuel must be discharged from the SHR after only ~ 6 years. Once in the third-stage core, the lifetime is limited again by material damage constraints, which are reached after ~ 18 years in the best cases. This would therefore mean a total fuel lifetime of ~ 38 years for the three-stage system, as compared to 32.5 years in the eB&B/double-clad system with constant axial enrichment, representing a modest improvement at the expense of significant additional complexity.

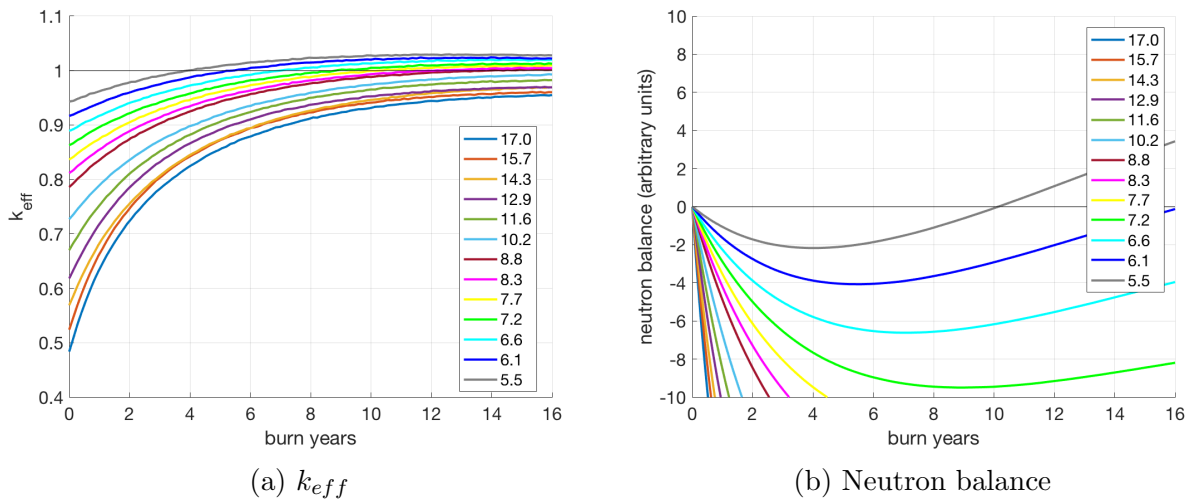


Figure 6.48: Evolution of k_{eff} and the neutron balance in a third-stage double-clad core utilizing discharged SHR fuel as feed, for various lifetimes in the SHR core. Lifetimes are given in years.

Possibility to Use Double-Clad Discharge in SHR Core

While the previous section explored the possibility of a three-stage system in the order of eB&B/SHR/double-clad, this section flips the last two stages so that a new three-stage system in the order of eB&B/double-clad/SHR is investigated. Such a three-stage system takes discharge assemblies from the double-clad core, decomposes them into 7-pin clusters, and adds moderator blocks to construct SHR assemblies as outlined previously in this chapter. In this case, however, the double-clad pins dictate a wider duct and thus the outer dimensions of the assembly are increased so as to keep the thickness of the moderator block unchanged.

Table 6.15 provides the assembly lifetimes in each stage for the case of uniform and axially-graded enrichment patterns as determined using the unit-cell model with $k_{eff}^{ave} = 1.03$. The total assembly lifetime in the uniform-enrichment case is not significantly different than what was found in the previous section for the eB&B/SHR/double-clad system, implying the three-stage system will have similar achievable burnup regardless of the order of the last two stages. The uniform case shows approximately 20% increase in assembly lifetime, as compared to only 13% for the axially-graded case, although it is likely that neither case shows large enough benefits to warrant the extra complexity.

Table 6.15: Assembly lifetimes in the first- (eB&B), second- (double-clad), and third-stage (SHR) cores for uniform and axially-graded enrichment patterns.

	Feed enrichment (w%)	First-stage lifetime (yr)	Second-stage lifetime (yr)	Third-stage lifetime (yr)	Total assembly lifetime (yr)
Uniform	6.0	13.9	18.6	6.5	39.0
Axially-graded	6.7/0.7	10.0	14.4	3.1	27.5

Summary of Three-Stage Fuel Cycles

This section has investigated the potential benefits of extending the previously-explored fuel cycles to have three-stages. By changing the order of the last two stages in the system, it was found that neither provides significantly higher potential than the other. While the fuel lifetime can be extended by $\sim 20\%$ by adding a third stage, the complexity of doing so would be very high. First, such a three-stage strategy would involve both the double-clad and SHR options, therefore incorporating the non-trivial difficulties of each. Second, because each stage is a unique reactor type, three separate reactors would be required in order to implement this strategy, involving high capital costs and requiring close coordination between operators. Therefore, the merits of the three-stage system do not appear to justify their costs, and this concept is not explored further.

6.7 Fuel Cycle Performance Intercomparison

Table 6.16 provides a comprehensive comparison of various fuel cycle metrics for the reactor systems investigated in this chapter. This section provides further analysis and discussion on the values in Table 6.16 so that the best system for various missions can be identified. All values are for equilibrium operations, although it is recognized that the transition to equilibrium could have very different impacts between systems. For B&B and eB&B systems, all values have been determined from full-core simulation data. For PWR systems, values have been taken from a combination of literature sources [152, 79, 38]. For double-clad, SHR, and SGR systems, values have been estimated using the maximum burnups found to be feasible in the preceding analyses using unit-cell calculations unless otherwise noted. Although these unit-cell calculations are expected to provide roughly accurate estimations of the tabulated data, some discretion is advised. Thermal efficiencies for all sodium-cooled systems has been assumed to be 42% [35], while the PWR has an assumed thermal efficiency of 33% [79].

Since all strategies employed in this chapter avoid fuel recycling or reconditioning, additional waste streams associated with those processes are avoided (see, for instance, [124]). In addition, the large infrastructure and cost associated with the recycling process will not be required. This provides an immediate benefit over traditional methods of obtaining improved resource utilization through recycling such as the PWR single-recycle. In addition, Table 6.16 shows that the PWR single-recycle strategy provides very limited gains over the PWR once-through strategy, with the exception of the discharged Pu mass, suggesting that this strategy is not particularly efficient at improving the sustainability of nuclear power generation in the first place.

In comparing the PWR with any of the single-stage eB&B core variants, only modest benefits are seen with two exceptions: (1) the power density can be greatly increased in the axially-shortened core and (2) the enrichment SWU requirement can be reduced by $\sim 45\%$ as compared to the once-through PWR. While resource utilization is consistently worse than in a PWR, fuel cost parameters such as the mining and SWU requirements range from marginally to significantly (in the axially-graded case) improved. Consistently, however, the discharge Pu content is multiple times larger across all single-stage eB&B core designs, which is to be expected based on the fact that these designs have been optimized to maximize the breeding ratio so as to reduce the minimum required burnup. In addition, the per-energy Pu production of an eB&B core is nearly twice that of a classical B&B core due to the fact that the cycle length is significantly shorter. Such large discharges of Pu could be a proliferation concern, especially considering the quality of the discharged Pu as shown in Table 6.17. While the PWR has a relatively large inventory of the troublesome isotopes Pu-238 and Pu-242, the Pu discharge vector from the eB&B core is quite pure, even more so than for the classical B&B.

Table 6.16: Fuel cycle summary of the various systems explored in the context of eB&B design studies. Bolded values correspond to the best performance in that category.

	Feed Enrichment (w%)	Mining requirement (tNU/GWe/yr)	Enrichment requirement (kSWU/GWe/yr)	Resource Utilization (%)	Power Density (W/cm ³)	Discharge Pu Production (kg/GWe/yr)
PWR						
Once-through	4.5	188.6	170.8	0.6	101	240
Single recycle in PWR	4.5	169.6	153.7	0.8	101/101	163
B&B	0.7	9.4	0	15.3	75	473
eB&B						
Uniform	6.0	165.1	161.4	0.5	75	737
Axially-shortened ^a	6.8	144.3	145.0	0.5	150	671
Axially-graded	0.7/6.7 ^b	158.1	96.6	0.5	75	801
eB&B/Double-clad						
Uniform	6.0	67.1	65.6	1.0	75/95 ^c	491
Axially-graded	0.7/6.7 ^b	71.4	43.5	1.3	75/95 ^c	529
Plus NU feed	0.7/6.7/0.7 ^d	53.7	32.7	1.7	75/95 ^c	463
Spectrally-softened ^e	0.7/6.7 ^b	60.6	37.1	1.5	75/85 ^c	447
eB&B/SHR						
Uniform	6.0	75.6	73.8	1.0	75/25 ^c	171
Axially-graded	0.7/6.7 ^b	75.3	46.0	1.2	75/25 ^c	202
eB&B/SGR						
Uniform	6.0	71.7	70.0	1.1	75/2 ^c	125
Axially-graded	0.7/6.7 ^b	67.3	41.1	1.4	75/2 ^c	152

^a For the 1.5 m tall, 3500 MW variant since it displayed the greatest benefits.^b For the unenriched and enriched zones, respectively.^c Second-stage power density assumes a full-core radial power peaking factor of 1.5 typical of traditional fuel management schemes.^d For the unenriched and enriched zones in the first-stage core, followed by enrichment for the supplementary feed in the second-stage core.^e Using graphite moderator.

Table 6.17: Pu production in various reactor systems in kg/GWe/yr.

	Pu-238	Pu-239	Pu-240	Pu-241	Pu-242	Total Pu
PWR ^a	6.0	123.8	56.8	35.8	17.3	240
B&B	2.0	380.2	81.8	8.1	1.7	473
eB&B ^b	2.4	676.6	56.5	3.1	0.2	739
Double-clad ^c	2.0	415.6	42.3	2.7	0.3	463
SHR ^d	6.5	51.0	81.6	21.5	9.9	171

^a Values from [79].

^b Values are for the uniform enrichment, 3.0 m tall case.

^c Values are for the axially-graded plus NU case.

^d Values are for the uniform enrichment case.

In contrast, Pu generation in the eB&B/SHR system is significantly lower, on par with that of a single-recycle PWR, due to effective burning of the fissile content in the SHR as shown in Figure 6.49. In addition, the Pu vector contains an even greater Pu-238 fraction than the PWR, with far less Pu-239. Though a fraction of the BOEC U-238 is converted to Pu-239 in the SHR, the Pu-239 content is effectively burned during the fuel's residence time. Overall, the core is found to have a conversion ratio of 0.29, which directly enables the more-favorable discharge Pu vector seen in Table 6.17. Burnup is peaked at more than 25%FIMA towards the axial center and is the lowest towards the periphery at $\sim 13\%$ FIMA. Furthermore, it is found that the incremental burnup in the SHR is larger towards the axial periphery than it is at the axial center, highlighting the effective burnup enabled in the second-stage SHR core.

The large discharge Pu mass in the eB&B cores could also be a drawback from the perspective of decay heat load and associated fuel storage implications, as shown in Figure 6.50. In this figure, decay heat load after discharge is compared against that for a PWR. PWR values are approximated using a radially-infinite unit-cell model of the 17×17 3.1w% enriched assemblies described in the BEAVRS benchmark [64]. The PWR assembly is burned for 2 cycles of 550 days until reaching a burnup of 41.7 MWd/kgHM at a power density of 4.2×10^{-2} W/kg (~ 90 W/cm³). In comparing the decay heat curve against literature data (mainly [79] and [6]), the values obtained from the unit-cell appear in line with established values. Further, the discharge Pu vector lines up favorably with the full-core literature values provided in Table 6.17, providing confidence that the model is delivering suitable results for the purposes of this comparison.

Between the B&B and eB&B cores, the qualitative trend is very similar, with the decay heat load for the eB&B core being shifted down as compared to the B&B core. Because the two cores have the same specific power density, this difference comes directly from the shorter cycle length of the eB&B core and the resulting lower levels of fission product and minor actinide buildup. The discharge SHR fuel, however, shows markedly different behavior.




Figure 6.49: Heavy metal content in the axially-uniform SHR at BOEC (solid) and EOEC (dashed) as a fraction of the initial heavy metal content in the first-stage eB&B core as predicted by the three-assembly model.

While the SHR is designed with the same average specific power density as the first-stage eB&B core, the decay heat immediately at discharge is largely influenced by the specific power density immediately at EOL, which can be significantly different than the average value. Therefore while Figure 6.50 indicates a significantly lower decay heat at discharge in the SHR, this value should be taken as approximate since the unit-cell SHR model cannot account for changes in power density over the cycle while fuel at EOL in the eB&B core will be close to its peak specific power density (see Figure 2.13 or Table 2.5). Nevertheless, once the impact of the short-lived fission products has diminished, the decay heat in the SHR shows a different profile than either the B&B or eB&B cores. Specifically, ~ 30 years after discharge the decay heat load of discharged SHR fuel actually overtakes that of the B&B core until again dropping below at ~ 400 years. The cause of this is the larger discharge masses of Pu-238 and Am-241, which each have α -decay energies of ~ 5.4 MeV. While Pu-238 is created almost entirely during full-power operation, the concentration of Am-241 is seen to increase with time after shutdown due to β^- decay of Pu-241, which, as shown in Table 6.17, is produced most strongly in the SHR. After ~ 400 years, decay heat in the B&B and eB&B fuel again overtake that of the SHR due primarily to the higher Pu-239 content, which has a half-life of 24,110 years and an associated α -decay of 5.2 MeV.

In comparing against the PWR values, again caution should be used during the early




Figure 6.50: Decay heat of discharged fuel from various systems in comparison to that from a PWR.

decay period due to the use of a unit-cell model. For the PWR, however, the average specific power density is nearly two times that of the B&B core, and thus the larger decay heat at discharge is reasonable. Regardless, in comparing decay heat loads after the initial discharge period, the PWR generally has lower heat load as would be expected due to the much lower burnups achieved. While Table 6.17 indicates that the PWR has very high per-energy Pu-238 production, the shorter residence time of the PWR makes it such that the total production of long-lived fission products and minor actinides is actually less than that of any of the other systems. Therefore, all systems investigated in this study have higher decay heat loads in the period after a handful of years.

None of the investigated systems are able to achieve breed-and-burn operation with enrichments lower than 5w%. That being said, the enrichment levels of all investigated options are well within the definition of HALEU and thus should be acceptable from the perspective of proliferation concerns. The systems with the lowest required peak enrichment (6.0w%) are those that employ the axially-uniform 3.0 m tall first-stage eB&B core, although other systems require only modestly higher enrichment levels near 7.0w%. This is not out of line in comparison to other advanced systems. For instance, modern LWRs such as the AP1000 employ enrichments up to 4.8w% [132], while various Generation IV concepts such as the FHR design by Kairos Power or the LFR design by LeadCold employ fuel enriched up to nearly 20w% [148].

From the perspective of resource utilization and enrichment costs, the clear winner is the eB&B/double-clad system supplemented with a secondary stream of NU pins, which achieves resource utilization nearly three times larger and enrichment costs roughly six times smaller than a PWR. This strategy performs particularly well because the NU feed has very little cost associated with it, so supplementing the discharged eB&B fuel with up $\sim 30\%$ NU pins directly enables higher burnup with essentially zero marginal costs. However, while this strategy performs particularly well, it is one of the more complicated from at least three perspectives: (1) it requires higher initial enrichment than some of the alternatives, (2) it employs the double-cladding procedure which is relatively complex and unknown, and (3) it utilizes fuel from two separate streams. As an alternative, the eB&B/SHR system with axially-graded enrichment performs at least twice as well as a PWR from the perspectives of mining costs and resource utilization, roughly four times better from the perspective of enrichment costs, and keeps the discharge Pu vector similar to that of a PWR. Since both strategies would enable significant benefits as compared to current PWR systems without introducing reprocessing, the decision of which to pursue should ultimately come down to an assessment of which technology is more viable. Moreover, because the benefits have been shown to be higher than those achievable in a single-recycling PWR, yet that strategy is already widely implemented in various European and Asian countries, it is expected that either proposed strategy could be attractive.

Although the SGR concept was previously disregarded due to its exceptionally low power density, it is listed in Table 6.16 to provide a lower bound on the mining and enrichment costs achievable from a second-stage thermal-spectrum core. Due to the lower absorption of graphite, the critical residence time for the SGR is able to be roughly 1 year longer than the SHR, which enables very modest improvements in per-energy mining costs, per-energy enrichment costs, resource utilization, and Pu production. Considering the low power density, these gains are not judged as strong enough to merit consideration of the SGR as an attractive alternative.

Finally, it is noted that the majority of benefits mentioned in this discussion would be enhanced if the DPA restriction were relaxed. As shown previously, there is a strong sensitivity to the DPA constraint in the double-clad core and a modest sensitivity in the SHR core, both of which would enable longer burn cycles, lower enrichment requirements, and higher utilization of fissile content from the first-stage eB&B core. Ultimately if the DPA constraint were relaxed enough to enable pure breed-and-burn operation (i.e. up to ~ 600 DPA), then it would make the most sense to utilize the classical B&B strategy to eliminate the need for a second core entirely. However, in the absence of such highly relaxed DPA constraints in the near future, the proposed two-stage systems may enable significantly enhanced performance in the interim with increasing benefits as the DPA experimental database is slowly extended.

6.8 Summary of the eB&B Two-Tiered Systems

This section has investigated a number of alternative systems that take the discharged fuel from an eB&B core for use as feed in a second-stage core, the two primary types being a double-cladded B&B core and a thermal-spectrum SHR. Within these two broad categories, various design strategies were pursued in order to deliver the largest benefits from a number of perspectives. In general, no single system out-performs the others in every quantified metric, and the optimal system will depend on the goals and constraints of the customer. However, a few generic conclusions can be made.

First, it has been confirmed that the sustainability of the nuclear fuel cycle can be improved significantly over currently deployed LWRs without using any form of fuel reprocessing while respecting constraints on material damage using the breed-and-burn mode of operation. In order to enable breed-and-burn operation within material damage constraints, feed fuel must be enriched upwards of 6w% in the so-called eB&B while employing the aggressive design strategies typical of classical B&B designs. Because such levels of enrichment lead to large losses in the enrichment process, the fuel must be taken to higher burnup using some form of second-stage core in order to show improvements as compared to current LWRs. In the double-cladded second-stage core, fuel lifetime is ultimately limited by again reaching the DPA constraint, while in the SHR second-stage core, reactivity loss is limiting. Utilizing an appropriately designed second-stage core, enrichment costs, mining costs, and resource utilization may be potentially improved by as much as roughly 3, 6, and 3 times, respectively. Furthermore, if the production of Pu is considered a proliferation risk, utilizing an SHR as a second-stage core is able to reduce the Pu production to lower than that of a PWR, on a per-energy basis.

These results demonstrate the potential promise of the newly-introduced two-stage systems. In fact, the obtained results show that performance even better than currently implemented single-recycle LWR systems can be obtained. However, a number of major simplifications in this study eventually need to be refined. First, the models used in this study were mainly simplified unit-cells. Although the unit-cells are expected to provide accurate trends and estimates, ultimately the true performance needs to be determined using optimized full-core designs. Second, in all cases the non-equilibrium operational phase has been neglected. It has been discussed that there may be a large discrepancy between the equilibrium and startup performance of traditional B&B cores, especially if the core needs to be started from EU. For the eB&B designs, this discrepancy is expected to be smaller since the equilibrium feed fuel is already EU and the cycle lengths are, in general, much shorter. Regardless, even for a PWR the transition cycles have an impact on the lifecycle performance that needs to be adequately understood. Other simplifications also exist, such as the neglect of cooling time between burnup in the first- and second-stage cores, although these types of simplifications are expected to have minimal impacts on the general conclusions. And of course, it should be finally mentioned that irradiation damage is not the only mechanism serving to degrade the fuel cladding. Time at temperature will lead to thermal creep and can enhance the migration of elements in the fuel/cladding system, both of which can degrade the structural

properties of the cladding. As the fuel in the proposed systems is envisioned to remain at temperature for more than 20 years, these mechanisms may be more challenging than in other high-burnup reactors with shorter lifetimes and thus need to be fully evaluated going forward.

6.9 Future Possibilities for eB&B Cores

The past few sections have investigated two different options for reusing discharged fuel from eB&B cores, both of which require the use of a two-reactor system. While such systems have potential for improving the sustainability of nuclear power generation, the requirement of building two separate, large-size power reactors to reap the benefits can be a barrier. For instance, not all regions can handle the addition of such large quantities of electricity, but locating the two reactors far from each other requires additional complexity from the perspective of transportation. Even in markets with large new capacity needs, two simultaneous capital-intensive projects may be beyond the capabilities of the local utility. Further, the tight coupling between the first- and second-stage reactors implies the need for close coordination between operation of the reactors, and it is easy to see how issues with operation of the first-stage reactor could impede fuel supply to the second. These issues motivate exploring the possibility of achieving benefits similar to the two-stage system using just a *single* core. While such studies are outside the scope of this dissertation, some possible options are discussed here to motivate further research.

The first proposed option is that a single reactor could be built with intention for two phases of operation. In the first phase, the eB&B core would operate as usual with discharge fuel kept in storage and gradually converted to the form required by the second-stage core. After some time operating in the first phase, all remaining fuel would be discharged and the core internals modified as necessary for the second phase of operation, which could correspond to either the double-clad or SHR strategy. By splitting the reactor lifetime into two phases like this, all of the plant systems only need to be constructed once and thus the capital cost burden would be much lower. In addition, coordination between the operations of first- and second-stage reactors is eliminated. This strategy would be challenging, however, due to the relatively long residence times of fuel in each of the cores. For each core, there will be a transition period during which resource utilization and fuel costs will be degraded before reaching equilibrium. Considering the fuel residence time in the eB&B core is on the order of 10-15 years, the transition to equilibrium will likely require a significant portion of the typical 80 year lifespan of the reactor. Adding in the transition for the second phase of operation, the amount of time left to operate in equilibrium may be limited under this strategy.

The second proposed option is to design an eB&B core which utilizes both single- and double-clad fuel in the same loading. While the addition of a second cladding layer has been shown to have a reactivity penalty, it has also been shown that the fuel still has enough reactivity to essentially double the burnup before again reaching the DPA limit. By double-

cladding the discharged single-clad fuel when it reaches the 200 DPA limit and putting it back into the core, the fuel is able to be a source of neutrons for longer, easing the required feed enrichment. Lower feed enrichments imply large improvements in fuel costs and resource utilization, and while it is expected that some amount of enrichment will still be required, this strategy would not require either a second reactor to be built or a multi-phase operational strategy over the reactor lifetime. The primary challenge associated with this strategy (on top of the challenges associated with the double-cladding process itself), would be designing a core lattice that could seamlessly accommodate the two different assembly designs. While this may require adapting traditional notions of assembly design, it is not expected to be insurmountable. Therefore this second option is preferred out of the two, and could offer significant promise in reducing the complexity and cost of the two-stage systems discussed previously in this chapter while simultaneously enhancing the benefits.

Chapter 7

Conclusions

This dissertation has examined a number of questions related to the design, analysis, and optimization of traditional B&B cores. After briefly overviewing the motivation for and fundamental principles behind B&B reactors, Chapter 2 spent some time examining the properties of a prototypical B&B reference core in detail. This analysis highlighted a number of prominent issues with B&B cores, such as (1) the ability to achieve high enough control system worth without strongly interfering with the breed-and-burn process, (2) the unique power profiles which lead to challenges in developing orificing schemes and adequate cooling, (3) the strong coolant void worth and its distribution throughout the core, and (4) the rapid buildup of material irradiation damage associated with the hard spectrum and long residence times. Throughout this chapter, these challenges have been analyzed and revealed in higher fidelity than had previously been available in the open literature. The development of a suitable control system configuration was resolved using tradeoff studies between various options, which ultimately resulted in relatively high worths-per-rod and was later shown to have important transient safety implications. Generally, the discussion in this chapter provided the basis on which the remaining chapters are developed.

Chapter 3 focused on the issue of reactor orificing. After discussing the orificing problem and its various constraints, the problem was framed as that of a mixed-integer linear program. Within this framework, three different methods of orificing were explored: (1) static grid plate, (2) static integral, and (3) variable. Using arguments based on bounding analyses of assembly power levels in each batch, it was shown that static grid plate orificing is likely not suitable for breed-and-burn operation mainly because the power swing of each assembly is too large, making it so that many assemblies will be over-cooled early in their life and forcing the outlet temperatures of high-powered assemblies to be unreasonably high. Using computational optimization methods, an extensive parametric study over the constraints within the static integral orificing paradigm showed that orificing schemes with as few as 10 unique groups are likely possible for steady-state operations. Further, the optimization method was shown to be highly general such that non-equilibrium operations can be handled simply. Finally, the method of variable orificing was also examined using a parametric study over the constraints, showing that a lower bound on the difference between outlet

temperatures of adjacent assemblies exists for the reference core. In general, this chapter has shown that acceptable orificing, although more complicated than for typical SFR designs, can be achieved in B&B cores – something that had not yet been conclusively demonstrated.

Chapter 4 then utilized the developed orificing scheme to build transient models of the reactor with the code SAS4A/SASSYS-1. This chapter at first investigated the response of the B&B core to unprotected transients, examining sensitivities to various operational parameters such as the primary loop pump properties, the rate of rod withdrawal during UTOP events, and the strength of different reactivity feedback mechanisms. While it was shown that ULOHS events may be handled with little difficulty, UTOP and ULOF events were shown to result in boiling within only a handful of minutes after transient initiation. In the case of the UTOP, the very strong reactivity insertion presented by the highest-worth rod resulted in boiling essentially independent of the investigated parameters, representing a tradeoff between the control system design in Chapter 2 and inherent safety performance. In the case of the ULOF, boiling is able to be avoided completely if the primary loop pumps are designed for a sufficiently smooth transition from forced- to natural-circulation, although the feasibility of such pump designs is unknown. These unfavorable results led to the investigation of protected transients and their sensitivities to the delay between transient initiation and scram, where it was found that the control system must scram within 20 s of LOF initiation to avoid potential cladding-eutectic reactions taking place. Regardless, although the unprotected transient performance was found to be lacking, the protected transient performance in the B&B core was found to be acceptable.

Chapter 5 takes as a starting point the inability to achieve inherent passive safety that was demonstrated in Chapter 4 and introduces the ARC system to potentially enable inherent passive safety. After describing the principles of the ARC system and the methods used to simulate its actuation, the three transients examined in Chapter 4 are again examined with the aid of the ARC system. Various parametric studies are done over the system design variables, which helped to identify the most sensitive variable to be the total system reactivity worth. While immediate benefits were found in the ULOHS in terms of reducing peak temperatures, response in the UTOP and ULOF scenarios was complicated by the tendency to form diverging oscillations. These oscillations were shown to be mitigated during the UTOP by making the ARC system less sensitive to temperature fluctuations, however attempting the same mitigation in the ULOF essentially made the ARC system unresponsive to the point of allowing boiling to occur. It was also shown that if the primary system pumps were properly designed the ULOF scenario could be handled, although in this case the ARC system provides essentially no benefits because its sensitivity must be reduced in order to avoid oscillations in the UTOP scenario.

In order to provide benefits to all examined transients without inducing diverging oscillations, a modification to the standard ARC system was proposed wherein a pseudo-one-way valve could be incorporated to allow aggressive actuations in the forward direction but retarded actuation in the reverse direction. To simulate this more complex design, a new code coupling methodology between SAS4A/SASSYS-1 and SAM was introduced and implemented. Using this new capability, the modified ARC system was incorporated into

transient simulations of the three unprotected transients. From these simulations, it was found that the ULOHS was essentially unimpacted by the new system as compared to the original, while more aggressive system designs were enabled without introducing oscillations in the UTOP or the ULOF. However, in the case of the ULOF, the tendency to put the reactor into an unstable regime which leads to delayed, but very rapid power spikes was consistently observed depending on the heat transport system boundary conditions, calling into question the potential usefulness of the modified ARC system in those scenarios. Regardless, significant improvements as compared to the original ARC system were enabled.

After examining the potential benefits from incorporating the ARC system into the B&B core, a variety of other operational aspects related to its inclusion were examined. It was found that a slight reactivity penalty will be induced by incorporating the ARC system, although the impacts of this were estimated to be insignificant. Through transmutation and diffusion of lithium out of the lower reservoir, some loss of absorber fluid is expected to take place, although in relation to the original volume this is expected to be negligible. Related to that, the production of tritium through transmutation of lithium was found to be potentially of concern, although in comparison to other existing reactor systems the levels are benign.

Finally, Chapter 6 focused on the accumulation of material damage, primarily quantified using the DPA metric, in B&B cores. After first presenting a literature review of design strategies to reduce material damage in these systems, it is concluded that nobody has yet succeeded in providing a viable path forward. Before moving towards proposing another alternative, the uncertainty in quantifying peak DPA deriving from uncertainties in nuclear data is investigated using a newly-proposed methodology. This study finds that the DPA uncertainties are generally small, mostly arising from uncertainties on the high-energy scattering cross-sections of U-238 and its impacts on correctly predicting the one-group fission cross-section.

After having quantified the uncertainties, the study introduces the concept of the eB&B core operating in the breed-and-burn mode of operation but utilizing feed fuel with limited enrichment. Through parametric analysis, it is found that breed-and-burn operation within the 200 DPA constraint can be achieved by employing feed fuel with 6-7w% enrichment. However it is found that at these enrichment levels, resource utilization is severely degraded through the enrichment process, with this result being relatively insensitive to various design optimization methods. In order to improve resource utilization and fuel cycle costs without introducing recycling or 'reconditioning', two alternatives for a second-stage core are introduced to utilize the discharged eB&B fuel as feed without violating constraints on material damage: (1) the double-clad B&B core and (2) the SHR core. Various strategies for most efficiently operating each second-stage core variant were explored and their sensitivities quantified. Between the two alternatives, both achieve similar additional burnup and reductions in fuel cycle costs, and ultimately the more desirable of the two should be decided by future detailed feasibility studies of both concepts. In general, though, it has been shown that resource utilization can be improved by roughly 3x as compared to traditional PWR fuel cycles, while other metrics such as mining and enrichment costs can be improved by roughly 3x and 6x, respectively, all without separation of fuel from cladding and thus while

maintaining the high proliferation resistance of the breed-and-burn fuel cycle and avoiding the costs associated with reprocessing. These studies, however, are all preliminary in the sense that they have used simplified neutronic models, and should be validated by subsequent studies.

Taken as a whole, the studies performed in this dissertation have furthered the state of knowledge around B&B reactors in a few ways. First, the feasibility of devising adequate orificing strategies has been confirmed. Second, a deeper understanding of the transient behavior and challenges to achieving inherent safety have been gained. Third, a thorough examination of the ARC system as applied to a typical B&B design has been performed, finding both impressive benefits and additional unforeseen challenges that still need to be resolved. Finally, new insight into the tradeoff between material damage and resource utilization has been developed which should help future designers understand where to focus their efforts in order to achieve efficient and feasible future breed-and-burn systems.

Bibliography

- [1] Aug. 2018. URL: <https://www.sustainability-times.com/low-carbon-energy/terrapower-presses-forward-on-next-gen-nuclear-energy/>.
- [2] Aug. 2019. URL: <https://www.federalregister.gov/documents/2019/08/05/2019-16578/notice-of-intent-to-prepare-an-environmental-impact-statement-for-a-versatile-test-reactor#p-36>.
- [3] Mar. 2020. URL: <https://www.nrc.gov/reactors/operating/licensing/renewal/subsequent-license-renewal.html>.
- [4] A. Abou-Jaoude, N.E. Stauff, and A. Erickson. “Performance and safety evaluation of a mixed-spectrum reactor design”. In: *Annals of Nuclear Energy* 126 (Apr. 2019), pp. 33–42. ISSN: 0306-4549. DOI: 10.1016/J.ANUCENE.2018.10.057.
- [5] M.J. Acton, G. Lenci, and E. Baglietto. “Structure-Based resolution of turbulence for sodium fast reactor thermal striping applications”. In: *International Topical Meeting on Nuclear Reactor Thermal Hydraulics 2015, NURETH 2015* 7 (2015), pp. 5416–5429.
- [6] Brian J Ade and Ian C Gauld. *Decay Heat Calculations for PWR and BWR Assemblies Fueled with Uranium and Plutonium Mixed Oxide Fuel Using Scale Prepared by*. Tech. rep. September. Oak Ridge, TN: Oak Ridge National Laboratory, 2011.
- [7] Charles Ahlfeld et al. *Nuclear Fission Reactor, Flow Control Assembly, Methods Therefore and a Flow Control Assembly System*. 2010. arXiv: arXiv:1208.5721.
- [8] A I Akhiezer et al. “Propagation of a nuclear chain reaction in the diffusion approximation”. In: *Physics of Atomic Nuclei* 62 (1999).
- [9] A B Alchagirov et al. “Wettability of Lithium by Liquid Sodium , Potassium , and Rubidium”. In: *High Temperature* 42.2 (2004), pp. 332–335.
- [10] Francisco J. Arias. “A negative reactivity feedback driven by induced buoyancy after a temperature transient in lead-cooled fast reactors”. In: *Nuclear Engineering and Technology* 50.1 (Feb. 2018), pp. 80–87. ISSN: 1738-5733. DOI: 10.1016/J.NET.2017.10.001.
- [11] R. Ashley et al. *SRE Fuel Element Damage: Final Report*. Tech. rep. Atomics International.

- [12] *ASTM E693-17, Standard Practice for Characterizing Neutron Exposures in Iron and Low Alloy Steels in Terms of Displacements Per Atom (DPA)*. 2017.
- [13] M. Aufiero, P. Rubiolo, and M. Fratoni. “Monte Carlo/CFD coupling for accurate modeling of the delayed neutron precursors and compressibility effects in molten salt reactors”. In: *Transactions of the American Nuclear Society* 116. June (2017). ISSN: 0003018X.
- [14] Manuele Aufiero et al. “A collision history-based approach to sensitivity/perturbation calculations in the continuous energy Monte Carlo code SERPENT”. In: *Annals of Nuclear Energy* 85 (Nov. 2015), pp. 245–258. ISSN: 0306-4549. DOI: 10.1016/J.ANUCENE.2015.05.008.
- [15] Y Bagdasarov et al. “Development of Passive Safety Devices for Sodium-Cooled Fast Reactors”. In: *Absorber materials, control rods and designs of shutdown systems for advanced liquid metal fast reactors*. Obninsk, Russian Federation: International Atomic Energy Agency, 1995.
- [16] C. W. Bale. “The K-Li (Potassium-Lithium) system”. In: *Bulletin of Alloy Phase Diagrams* 10.3 (June 1989), pp. 262–264. ISSN: 0197-0216. DOI: 10.1007/BF02877508.
- [17] R. Batchelor and J. H. Towle. “The interactions of neutrons with Li6 and Li7 between 1.5 and 7.5 MeV”. In: *Nuclear Physics* 47.C (1963), pp. 385–407. ISSN: 00295582. DOI: 10.1016/0029-5582(63)90881-2.
- [18] R.J. Beeley and J.E. Mahlmeister. *Operating experience with the sodium reactor experiment and its application to the hallam nuclear power facility*. Tech. rep. Canoga Park, CA: Atomics International.
- [19] R.J. Beeley et al. *The Sodium Graphite Reactor: Tomorrow’s Power Plant*. Tech. rep. Canoga Park, CA: Atomics International, 1960.
- [20] John D Bess et al. “Comparison of Displacement Damage Calculations Supporting MIMIC Analysis and Design for TREAT”. In: *Transactions of the American Nuclear Society*. Vol. 119. Orlando, FL, 2018.
- [21] David Blanchet et al. *AEN - WPRS Sodium Fast Reactor Core Definitions (version 1.2)*. Tech. rep. 2011. URL: <https://www.oecd-nea.org/science/wprs/sfr-taskforce/WPRS-AEN-SFR-Cores-V1.2.pdf>.
- [22] H Boeck et al. *History, Development and Future of TRIGA Research Reactors*. Vienna, Austria, 2016.
- [23] Jacopo Buongiorno, James W. Sterbentz, and Philip E. MacDonald. “Study of solid moderators for the thermal-spectrum supercritical water-cooled reactor”. In: *Nuclear Technology* 153.3 (2006), pp. 282–303. ISSN: 00295450. DOI: 10.13182/NT06-A3708.
- [24] *BWR/6: General Description of a Boiling Water Reactor*. Tech. rep. GE Nuclear Energy.

- [25] J.E. Cahalan and T. H. Fanning. *The SAS4A / SASSYS-1 Safety Analysis Code System*. Tech. rep. Lemont, IL: Argonne National Laboratory, 2012.
- [26] X Chen et al. “Boiling water cooled travelling wave reactor”. In: *Annals of Nuclear Energy* 134 (2019), pp. 342–349. ISSN: 0306-4549. DOI: 10.1016/j.anucene.2019.06.037.
- [27] J.D. Cochran and J.E. Owens. “Initial Operation of the Sodium Graphite Reactor at the Hallam Nuclear Power Facility”. In: *American Power Conference 26th Annual Meeting*. Chicago, IL, 1964.
- [28] Hugo van Dam. “Self-stabilizing criticality waves”. In: *Annals of Nuclear Energy* 27.16 (Nov. 2000), pp. 1505–1521. ISSN: 0306-4549. DOI: 10.1016/S0306-4549(00)00035-9.
- [29] Christian Di Sanzo. “Increasing Fuel Utilization of Breed and Burn Reactors”. PhD thesis. University of California, Berkeley, 2014.
- [30] R.W. Dickinson, R.C. Gerber, and C.L. Larson. “Hallam Nuclear Power Facility Prototype for Advanced Sodium-Cooled Power”. In: February (1962), pp. 1008–1011.
- [31] J Duderstadt and L Hamilton. *Nuclear Reactor Analysis*. Wiley-Interscience, 1976.
- [32] Bob Dudley. *BP Energy Outlook 2019 edition*. Tech. rep. British Petroleum, 2019.
- [33] F. E. Dunn and D. J. Malloy. “LMR Centrifugal Pump Coastdowns”. In: *Anticipated and Abnormal Transients in Nuclear Power Plants Topical Meeting*. Atlanta, GA, 1987.
- [34] M. Edelmann, G. Kussmaul, and W. Vath. “Improved Fast Reactor Safety by Passive Shut-Down”. In: *Progress in Nuclear Energy* 29 (1995), pp. 379–386.
- [35] J.A. Falcon and P.R. Keep. “A 200 MWE Prototype Sodium Cooled Reactor High Efficiency Power Plant”. In: *American Public Power Association 20th Annual Convention*. Cleveland, OH, 1963.
- [36] T. Flaspöehler and B. Petrovic. “Assessing Pathways to Tritium Production and its Detailed Spatial Distribution Throughout the VHTR”. In: *EPJ Web of Conferences* 106 (2016). ISSN: 2100014X. DOI: 10.1051/epjconf/201610603005.
- [37] *Frequently Asked Questions About Liquid Radioactive Releases*. Tech. rep. US Nuclear Regulatory Commission, 2018. URL: <https://www.nrc.gov/reactors/operating/ops-experience/tritium/faqs.html%5C#normal>.
- [38] E. Fridman and S. Kliem. “Pu recycling in a full Th-MOX PWR core. Part I: Steady state analysis”. In: *Nuclear Engineering and Design* 241.1 (2011), pp. 193–202. ISSN: 00295493. DOI: 10.1016/j.nucengdes.2010.10.036.
- [39] John Gilleland, Robert Petroski, and Kevan Weaver. “The Traveling Wave Reactor: Design and Development”. In: *Engineering* 2.1 (2016), pp. 88–96. ISSN: 20958099. DOI: 10.1016/J.ENG.2016.01.024.

- [40] G. Girardin et al. “Development and characterization of the control assembly system for the large 2400 MWth Generation IV gas-cooled fast reactor”. In: *Annals of Nuclear Energy* 35.12 (2008), pp. 2206–2218. ISSN: 03064549. DOI: 10.1016/j.anucene.2008.09.008.
- [41] Joshua Goldstein and Staffan Qvist. *A Bright Future: How Some Countries Have Solved Climate Change and the Rest Can Follow*. PublicAffairs, 2019, p. 288.
- [42] D Green and R Perry. *Perry’s Chemical Engineers’ Handbook*. McGraw-Hill Education, 2007.
- [43] Ehud Greenspan. “A Phased Development of Breed-and-Burn Reactors for Enhanced Nuclear Energy Sustainability”. In: *Sustainability* 4.12 (Oct. 2012), pp. 2745–2764. ISSN: 2071-1050. DOI: 10.3390/su4102745.
- [44] Ehud Greenspan et al. *A Pebble-Bed Breed and Burn Reactor*. Tech. rep. University of California, Berkeley, 2016.
- [45] E. Greenspan et al. “Hydride fuel for LWRs-Project overview”. In: *Nuclear Engineering and Design* 239.8 (2009), pp. 1374–1405. ISSN: 00295493. DOI: 10.1016/j.nucengdes.2008.11.023.
- [46] E Greenspan et al. *NEUP Project #15-8251 Final Report: Enhanced Performance Fast Reactors with Engineered Passive Safety System*. Tech. rep. Nuclear Energy University Programs, 2018.
- [47] JD Gylfe et al. *Evaluation of Zirconium Hydride as Moderator in Integral, Boiling Water-Superheat Reactors*. Canoga Park, CA, 1962.
- [48] M.J Hackett and G Povirk. *HT9 Development for the Traveling Wave Reactor*. 2012.
- [49] Donny Hartanto, Chihyung Kim, and Yonghee Kim. “An optimization study on the excess reactivity in a linear breed-and-burn fast reactor”. In: *Annals of Nuclear Energy* 94 (Aug. 2016), pp. 62–71. ISSN: 0306-4549. DOI: 10.1016/J.ANUCENE.2016.02.017.
- [50] Donny Hartanto and Yonghee Kim. “Alternative reflectors for a compact sodium-cooled breed-and-burn fast reactor”. In: *Annals of Nuclear Energy* 76 (Feb. 2015), pp. 113–124. ISSN: 0306-4549. DOI: 10.1016/J.ANUCENE.2014.09.048.
- [51] Donny Hartanto et al. “FAST and SAFE passive safety devices for sodium-cooled fast reactor”. In: *Transactions of the Korean Nuclear Society Spring Meeting*. May. Jeju, 2015.
- [52] Donny Hartanto et al. “Impacts of Burnup-Dependent Swelling of Metallic Fuel on the Performance of a Compact Breed-and-Burn Fast Reactor”. In: *Nuclear Engineering and Technology* 48.2 (Apr. 2016), pp. 330–338. ISSN: 1738-5733. DOI: 10.1016/J.NET.2015.12.009.
- [53] F Heidet, T.K. Kim, and T.A. Taiwo. “Orificing Strategy for Sodium-Cooled Fast Reactors”. In: *Transactions of the American Nuclear Society*.

- [54] Florent Heidet. “Maximum Fuel Utilization in Advanced Fast Reactors without Actinides Separation”. PhD thesis. University of California, Berkeley, 2010.
- [55] Florent Heidet and Ehud Greenspan. “Performance of Large Breed-and-Burn Core”. In: *Nuclear Technology* 181.3 (Mar. 2013), pp. 381–407. ISSN: 0029-5450. DOI: 10.13182/NT13-A15800.
- [56] Florent Heidet and Ehud Greenspan. “Superprism-Sized Breed-and-Burn Sodium-Cooled Core Performance”. In: *Nuclear Technology* 5450.October (2012). DOI: 10.13182/NT13-A15782.
- [57] Pavel Hejzlar et al. “Terrapower, LLC traveling wave reactor development program overview”. In: *Nuclear Engineering and Technology* 45.6 (2013), pp. 731–744. ISSN: 2234358X. DOI: 10.5516/NET.02.2013.520.
- [58] J C Hesson, M J Feldman, and L Burris. *Description and Proposed Operation of the Fuel Cycle Facility for the Second Experimental Breeder Reactor (EBR-II)*. Argonne, IL, 1963. DOI: <http://digital.library.unt.edu/ark:/67531/metadc11564/>.
- [59] Van Khanh Hoang and Toru Obara. “Impact of Neutron Spectrum Shift on Breed and Burn Reactor Concept”. In: *Energy Procedia* 131 (Dec. 2017), pp. 37–44. ISSN: 1876-6102. DOI: 10.1016/J.EGYPRO.2017.09.443.
- [60] Edward A. Hoffman, Robert N. Hill, and W.S Yang. *Preliminary Core Design Studies for the Advanced Burner Reactor over a Wide Range of Conversion Ratios*. Tech. rep. Lemont, IL: Argonne National Laboratory, 2006. DOI: 10.1007/s13398-014-0173-7.2. arXiv: 9809069v1 [arXiv:gr-qc].
- [61] G L Hofman et al. *Metallic Fuels Handbook*. Tech. rep. Argonne, IL: Argonne National Laboratory, 1985.
- [62] G.L. Hofman, L.C. Walters, and T.H. Bauer. “Metallic fast reactor fuels”. In: *Progress in Nuclear Energy* 31.1-2 (Jan. 1997), pp. 83–110. ISSN: 0149-1970. DOI: 10.1016/0149-1970(96)00005-4.
- [63] Boris Hombourger, Jiri Krepel, and Andreas Pautz. “Breed-and-burn fuel cycle in molten salt reactors”. In: *EPJ Nuclear Science Technology* 15 (2019), pp. 1–10.
- [64] N. Horelik et al. “Benchmark for Evaluation and Validation of Reactor Simulations (BEAVRS)”. In: *Proceedings of the International Conference of Mathematics and Computational Methods Applied to Nuclear Science and Engineering*. Sun Valley, ID: American Nuclear Society, 2013.
- [65] Jason Hou et al. “3D in-core fuel management optimization for breed-and-burn reactors”. In: *Progress in Nuclear Energy* 88 (2016). ISSN: 01491970. DOI: 10.1016/j.pnucene.2015.12.002.

- [66] Jason Hou et al. “In-Core Fuel Management Optimization for Breed-and-Burn Reactors With 3D Fuel Shuffling”. In: *PHYSOR 2016*. Vol. 88. May. Sun Valley, ID, 2016, pp. 3481–3492. ISBN: 9781510825734. DOI: <https://doi.org/10.1016/j.pnucene.2015.12.002>.
- [67] Jason Hou et al. “On Multi-Group Cross Sections for Breed-and-Burn Reactors”. In: *Transactions of the American Nuclear Society*. Anaheim, CA, 2014, pp. 4–7.
- [68] Rui Hu. *SAM Theory Manual*. Tech. rep. Argonne, IL: Argonne National Laboratory, 2017.
- [69] A Hyder. *Spacecraft Power Technologies*. Vol. 1. Imperial College Press, 2000, p. 476.
- [70] *IBM ILOG CPLEX Optimization Studio: CPLEX Parameters Reference*. Tech. rep. IBM Corporation, 2015, p. 594.
- [71] I E Idel’chik. *Handbook of Hydraulic Resistance: Coefficients of Local Resistance and of Friction*. Tech. rep. U.S. Atomic Energy Commission, 1966.
- [72] *INFO-0793: Tritium Releases and Dose Consequences in Canada in 2006*. Tech. rep. Canadian Nuclear Safety Commission, 2009.
- [73] D W Jeppson et al. *Lithium Literature Review: Lithium’s Properties and Interactions*. Tech. rep. Hanford, WA: Hanford Engineering Development Laboratory, 1978, pp. 1–111. DOI: 10.2172/6885395.
- [74] Mitsuru Kambe and Masaki Uotani. “Design and Development of Fast Breeder Reactor Passive Reactivity Control Systems: LEM and LIM”. In: *Nuclear Technology* 122.2 (May 1998), pp. 179–195. ISSN: 0029-5450. DOI: 10.13182/NT98-A2861.
- [75] Julia Abdul Karim, Jun Nishiyama, and Toru Obara. “Application of melt and refining procedures in the CANDLE reactor concept”. In: *Annals of Nuclear Energy* 90 (Apr. 2016), pp. 275–283. ISSN: 0306-4549. DOI: 10.1016/J.ANUCENE.2015.12.001.
- [76] Alisha Ana Kasam. “Conceptual design of a breed & burn molten salt reactor”. PhD thesis. University of Cambridge, 2018.
- [77] Chris Keckler et al. “A Mixed-Integer Linear Programming Method for Optimal Orificing in Breed-and-Burn Cores Chris Keckler”. In: *Transactions of the American Nuclear Society*. Philadelphia, PA, 2018.
- [78] Jintae Kim and Moosung Jae. “A study on reliability assessment of a decay heat removal system for a sodium-cooled fast reactor”. In: *Annals of Nuclear Energy* 120 (Oct. 2018), pp. 534–539. ISSN: 0306-4549. DOI: 10.1016/J.ANUCENE.2018.06.022.
- [79] Taek K Kim and T A Taiwo. “Systematic Evaluation of Uranium Utilization in Nuclear Systems”. In: (2010).
- [80] T.K. Kim et al. “Core design studies for a 1000 MWth Advanced Burner Reactor”. In: *Annals of Nuclear Energy* 36.3 (Apr. 2009), pp. 331–336. ISSN: 0306-4549. DOI: 10.1016/J.ANUCENE.2008.12.021.

- [81] Leonard J. Koch. *Experimental Breeder Reactor-II*. DOI: 10.1016/S0169-7552(98)00103-2.
- [82] K Kuwagaki. Private communication. 2020.
- [83] K. Kuwagaki, C. Keckler, and M. Fratoni. “Effect of Doping Fuel with Beryllium on Accumulated Radiation Damage in a Breed-and-Burn Reactor”. In: *Transactions of the American Nuclear Society*. Washington D.C., 2019. DOI: 10.13182/t30822.
- [84] Kazuki Kuwagaki, Jun Nishiyama, and Toru Obara. “Concept of Stationary Wave Reactor with Rotational Fuel Shuffling”. In: *Nuclear Science and Engineering* (June 2018), pp. 1–9. ISSN: 0029-5639. DOI: 10.1080/00295639.2018.1463744.
- [85] L. Leibowitz and R. A. Blomquist. “Thermal conductivity and thermal expansion of stainless steels D9 and HT9”. In: *International Journal of Thermophysics* 9.5 (Sept. 1988), pp. 873–883. ISSN: 0195-928X. DOI: 10.1007/BF00503252.
- [86] Jaakko Leppänen et al. “The Serpent Monte Carlo code: Status, development and applications in 2013”. In: *Annals of Nuclear Energy* 82 (Aug. 2015), pp. 142–150. ISSN: 0306-4549. DOI: 10.1016/J.ANUCENE.2014.08.024.
- [87] Roberto Lopez-Solis and Juan Luis François. “The breed and burn nuclear reactor: A chronological, conceptual, and technological review”. In: *International Journal of Energy Research* 42.3 (2018), pp. 953–965. ISSN: 1099114X. DOI: 10.1002/er.3854.
- [88] David MacKay. *Sustainable Energy – Without the Hot Air*. UIT Cambridge, 2009, p. 384.
- [89] W. J. Marshall et al. “Development and testing of Neutron cross-section covariance data for SCALE 6.2”. In: *ICNC 2015 - International Conference on Nuclear Criticality Safety* (2015), pp. 1213–1225.
- [90] Michael Vincent Martin et al. “Feasibility of a Breed-and-Burn Molten Salt Reactor”. In: *Transactions of the American Nuclear Society*. June. San Francisco, CA, 2017.
- [91] M Mattes and J Keinert. *Thermal Neutron Scattering Data for the Moderator Materials H₂O, D₂O, and ZrHx in ENDF-6 Format and as ACE Library for MCNP(X) Codes*. Vienna, Austria, 2005.
- [92] R.R. Matthews et al. “Performance and Operation of the Dounreay Fast Reactor”. In: *Third United Nations International Conference on the Peaceful Uses of Atomic Energy*. United Nations, 1964.
- [93] TK McCusker. *Final Report: Decontamination and Decommissioning of Heat Transfer Reactor Experiment Test Assemblies HTRE-2 and HTRE-3*. Idaho Falls, ID, 1989.
- [94] Edward McGinnis. *Testimony of Edward G. McGinnis, Principle Deputy Assistant Secretary for Nuclear Energy, U.S. Department of Energy before the U.S. House Committee on Energy and Commerce Subcommittee on Energy*. May 2018. URL: <https://docs.house.gov/meetings/IF/IF03/20180522/108339/HHRG-115-IF03-Wstate-McGinnisE-20180522.pdf>.

- [95] B Munson et al. *Fluid Mechanics*. Wiley, 2013.
- [96] A. M. Nolan. *An Overview of Environmental Control Aspects for the Gas-Cooled Fast Reactor*. Tech. rep. Richland, Washington: Pacific Northwest Laboratory, 1981.
- [97] Kai Nordlund et al. “Improving atomic displacement and replacement calculations with physically realistic damage models”. In: *Nature Communications* 9.1 (2018), pp. 1–8. ISSN: 20411723. DOI: 10.1038/s41467-018-03415-5.
- [98] M.J. Norgett, M.T. Robinson, and I.M. Torrens. “A proposed method of calculating displacement dose rates”. In: *Nuclear Engineering and Design* 33.1 (Aug. 1975), pp. 50–54. ISSN: 0029-5493. DOI: 10.1016/0029-5493(75)90035-7.
- [99] April Novak. “Multiscale Thermal-Hydraulic Methods for Pebble Bed Reactors”. PhD thesis. University of California, Berkeley, 2020. Chap. 6.
- [100] Carlo Oggioni et al. “SAS/SAM Coupled Analysis of Modified ARC System in SFR Transients”. In: *PHYSOR2020*. Cambridge, UK, 2020.
- [101] K Okada et al. “The Design of a Backup Reactor Shutdown System of DFBR”. In: *Absorber materials, control rods and designs of shutdown systems for advanced liquid metal fast reactors*. Obninsk, Russian Federation: International Atomic Energy Agency, 1995.
- [102] A.G. Osborne, G.D. Recktenwald, and M. Deinert. “Propogation of a solitary fission wave”. In: *Chaos* 22 (2012).
- [103] Andrew Osborne and Mark Deinert. “Reducing Irradiation Damage in a Long-Life Fast Reactor with Spectral Softening”. In: *Energies* 11.6 (2018), p. 1507. ISSN: 1996-1073. DOI: 10.3390/en11061507.
- [104] R. G. Pahl et al. “Irradiation behavior of metallic fast reactor fuels”. In: *Journal of Nuclear Materials* 188.C (1992), pp. 3–9. ISSN: 00223115. DOI: 10.1016/0022-3115(92)90447-S.
- [105] J. Park et al. “Design study of long-life small modular sodium-cooled fast reactor”. In: *International Journal of Energy Research* 41.1 (2017), pp. 139–148. ISSN: 1099114X. DOI: 10.1002/er.3609.
- [106] W. Parkins. *Design and Development of Components for the SRE*. Tech. rep. Downey, CA: North American Aviation, Inc., 1968.
- [107] Rauli Partanen and Janne Korhonen. *Climate Gamble: Is Anti-Nuclear Activism Endangering Our Future?* Cre8 Oy, 2017, p. 132.
- [108] Ben Pauw et al. “Characterizing Flow-Induced Vibrations of Fuel Assemblies for Future Liquid Metal Cooled Nuclear Reactors Using Quasi-Distributed Fibre-Optic Sensors”. In: *Applied Sciences* 7.8 (Aug. 2017), p. 864. ISSN: 2076-3417. DOI: 10.3390/app7080864.
- [109] E N Pearson. *SRE Systems and Components Experience: Core II*. Tech. rep. Canoga Park, CA: Atomics International, 1963.

- [110] Robert C. Petroski. “General Analysis of Breed-and-Burn Reactors and Limited-Separations Fuel Cycles”. PhD thesis. Massachusetts Institute of Technology, 2011.
- [111] Igor Pioro, ed. *Handbook of Generation IV Nuclear Reactors*. 1st ed. Woodhead Publishing, 2016, p. 940.
- [112] *Preliminary Safety Analysis Report*. Oak Ridge, TN, 1974.
- [113] S Qvist and E Bubelis. *Passive Shutdown Systems for Fast Neutron Reactors*. Tech. rep. Vienna, Austria: International Atomic Energy Agency, 2020.
- [114] Staffan A. Qvist. “Safety and core design of large liquid-metal cooled fast breeder reactors”. PhD thesis. University of California, Berkeley, 2013.
- [115] Staffan A. Qvist et al. “Autonomous Reactivity Control (ARC) – Principles, geometry and design process”. In: *Nuclear Engineering and Design* 307 (Oct. 2016), pp. 249–274. ISSN: 0029-5493. DOI: 10.1016/J.NUCENGDES.2016.07.018.
- [116] Staffan Qvist and Ehud Greenspan. “Design Space Analysis for Breed-and-Burn Reactor Cores”. In: *Nuclear Science and Engineering* 182.2 (Feb. 2016), pp. 197–212. ISSN: 0029-5639. DOI: 10.13182/NSE14-135.
- [117] Staffan Qvist and Ehud Greenspan. “Inherent safety of minimum - burnup breed & burn reactors”. In: *ICAPP 2012*. Chicago, 2012.
- [118] Staffan Qvist and Ehud Greenspan. “The ADOPT code for automated fast reactor core design”. In: *Annals of Nuclear Energy* 71 (2014), pp. 23–36. ISSN: 18732100. DOI: 10.1016/j.anucene.2014.03.013.
- [119] Staffan Qvist, Jason Hou, and Ehud Greenspan. “Design and performance of 2D and 3D-shuffled breed-and-burn cores”. In: *Annals of Nuclear Energy* 85 (2015), pp. 93–114. ISSN: 18732100. DOI: 10.1016/j.anucene.2015.04.007.
- [120] Jose Reyes Jr. “NuScale Plant Safety in Response to Extreme Events”. In: *Nuclear Technology* 178.2 (2012), pp. 153–163. DOI: 10.13182/NT12-A13556.
- [121] Richard Rhodes. *Energy: A Human History*. Simon and Schuster, 2018, p. 480.
- [122] V. D. Rusov et al. “On some fundamental peculiarities of the traveling wave reactor”. In: *Science and Technology of Nuclear Installations* 2015.1 (2015). ISSN: 16876083. DOI: 10.1155/2015/703069. arXiv: 1207.3695.
- [123] Raluca O Scarlat. *Hydrogen Isotope Transport in the Graphitic Fuel Elements of Fluoride-Salt Cooled High-Temperature Reactors (FHR)*. 2015.
- [124] Aleksandra Schwenk-Ferrero and Andrei Andrianov. “Comparison and screening of nuclear fuel cycle options in view of sustainable performance and waste management”. In: *Sustainability (Switzerland)* 9.9 (2017). ISSN: 20711050. DOI: 10.3390/su9091623.
- [125] Hiroshi Sekimoto, Kouichi Ryu, and Yoshikane Yoshimura. “CANDLE: The New Burnup Strategy”. In: *Nuclear Science and Engineering* 139.3 (2001), pp. 306–317. DOI: 10.13182/NSE01-01. eprint: <https://doi.org/10.13182/NSE01-01>.

- [126] E. R. Shemon, Y. Yu, and T. K. Kim. *Application of the SHARP Toolkit to Sodium-Cooled Fast Reactor Challenge Problems*. Tech. rep. Lemont, IL: Argonne National Laboratory, 2017.
- [127] PK Shen et al. “The U.S. Advanced Liquid Metal Reactor and the Fast Flux Test Facility Phase IIA Passive Safety Tests”. In: *ANP92 International Conference on Design and Safety of Advanced Nuclear Power Plants*. Tokyo, Japan, 1992.
- [128] L. Shikun, Z. Changshan, and X. Yingxian. “Diving-bell and double-vented B4C Control Rod Pin”. In: *Absorber materials, control rods and designs of shutdown systems for advanced liquid metal fast reactors*. Obninsk, Russian Federation: International Atomic Energy Agency, 1995.
- [129] Rijan Shrestha and Rizwan-uddin. “Modeling space-time evolution of flux in a traveling wave reactor”. In: *Annals of Nuclear Energy* 70 (Aug. 2014), pp. 90–95. ISSN: 0306-4549. DOI: 10.1016/J.ANUCENE.2014.03.011.
- [130] Vaclav Smil. *Energy and Civilization: A History*. The MIT Press, 2017, p. 568.
- [131] *SNAP Overview*. URL: https://www.etec.energy.gov/Operations/Major_Operations/SNAP_Overview.php.
- [132] *Status report 81 - Advanced Passive PWR (AP 1000)*. Tech. rep. Vienna, Austria: International Atomic Energy Agency Advanced Reactors Information Systems, 2011. URL: <https://www.iaea.org/NuclearPower/Downloadable/aris/2013/7.AP1000.pdf>.
- [133] N. E. Stauff et al. “Evaluation of the OECD / NEA / SFR-UAM Neutronics Reactivity Feedback and Uncertainty Benchmarks”. In: *FR17*. Yekaterinburg, Russia: International Atomic Energy Agency, 2017.
- [134] T. Sumner and A. Moiseyev. “Simulations of the EBR-II tests SHRT-17 and SHRT-45R”. In: *International Topical Meeting on Nuclear Reactor Thermal Hydraulics 2015, NURETH 2015 4* (2015), pp. 3044–3057.
- [135] Tyler Sumner. *BOP - 301 and BOP - 302R : Test Definitions and Analyses*. Tech. rep. Argonne National Laboratory, 2016.
- [136] Tyler Sumner et al. “Safety analysis of the FAST TEst reactor (FASTER) preconceptual design”. In: *Progress in Nuclear Energy* 108 (Sept. 2018), pp. 454–464. ISSN: 0149-1970. DOI: 10.1016/J.PNUCENE.2018.06.013.
- [137] Leo Svanström. “Load following with a passive reactor core using the SPARC design SPARC design”. PhD thesis. Uppsala Universitet, 2016.
- [138] Misao Takamatsu et al. “Demonstration of control rod holding stability of the self actuated shutdown system in joyo for enhancement of fast reactor inherent safety”. In: *Journal of Nuclear Science and Technology* 44.3 (2007), pp. 511–517. ISSN: 00223131. DOI: 10.1080/18811248.2007.9711316.

- [139] Core Writing Team, R.K. Pachauri, and L.A. Meyer. *Climate Change 2014 Synthesis Report*. Tech. rep. IPCC, 2014.
- [140] Kurt Terrani. “Incorporation of Hydride Nuclear Fuels in Commercial Light Water Reactors”. PhD thesis. University of California, Berkeley, 2010.
- [141] Nikola Tesla. *Valvular Conduit*. 1920.
- [142] M.B. Toloczko and F.A. Garner. “Variability of Irradiation Creep and Swelling of HT9 Irradiated to High Neutron Fluence at 400-600 C”. In: *Effects of Radiation on Materials: 18th International Symposium*. Conshohocken, PA: ASTM International, 1999. ISBN: 8466927344. DOI: 10.6100/IR336248.
- [143] M.B. Toloczko, F.A. Garner, and C.R. Eiholzer. “Irradiation creep and swelling of the US fusion heats of HT9 and 9Cr-1Mo to 208 dpa at $\sim 400\text{C}$ ”. In: *Journal of Nuclear Materials* 212-215 (Sept. 1994), pp. 604–607. ISSN: 0022-3115. DOI: 10.1016/0022-3115(94)90131-7.
- [144] N. Touran et al. “Technical Considerations and Capabilities of a Near-Term Deployable Traveling Wave Reactor Core”. In: *Transactions of the American Nuclear Society*. December. 2010.
- [145] Simon Vanmaercke et al. “Development of a secondary SCRAM system for fast reactors and ADS systems”. In: *Science and Technology of Nuclear Installations 2012* (2012). ISSN: 16876075. DOI: 10.1155/2012/351985.
- [146] D.C. Wade, R.A. Wigeland, and D.J. Hill. “The safety of the IFR”. In: *Progress in Nuclear Energy* 31.1-2 (Jan. 1997), pp. 63–82. ISSN: 0149-1970. DOI: 10.1016/0149-1970(96)00004-2.
- [147] J.B. Waldo et al. “Application of the GEM shutdown device to the FFTF reactor”. In: *Trans. Am. Nucl. Soc.; (United States)* 53 (Jan. 1986).
- [148] J Wallenius et al. “Unprotected Transients in Sealer: A Small Lead-Cooled Reactor for Commercial Power Production in Arctic Regions”. In: *Proceedings of the Physics of Reactors (PHYSOR) Conference*. Cancun, MX, 2018, pp. 1437–1448.
- [149] Alan Waltar, Donald Todd, and Pavel Tsvetkov, eds. *Fast Spectrum Reactors*. New York: Springer, 2012. ISBN: 9781441995711.
- [150] Catherine Westfall. *Civilian Nuclear Power on the Drawing Board: The Development of Experimental Breeder Reactor-II*. Tech. rep. Argonne National Laboratory, 2003.
- [151] C L Whitmarsh. *Review of Zircaloy-2 and Zircaloy-4 Properties Relevant to NS Savannah Reactor Design*. Tech. rep. Oak Ridge National Laboratory, 1962.
- [152] R Wigeland and et al. *Nuclear Fuel Cycle Evaluation and Screening – Final Report*. Tech. rep. Department of Energy, 2014.
- [153] Yunlong Xiao et al. “Neutronics studies on the feasibility of developing fast breeder reactor with flexible breeding ratio”. In: *Journal of Nuclear Science and Technology* 53.1 (2016), pp. 129–138. ISSN: 00223131. DOI: 10.1080/00223131.2015.1030464.

- [154] Won Sik Yang and Han Gyu Joo. “LMR core temperature calculation based on implicit formulation of the ENERGY model and a Krylov subspace method”. In: *Annals of Nuclear Energy* 26.7 (1999), pp. 629–640. ISSN: 03064549. DOI: 10.1016/S0306-4549(98)00074-7.
- [155] Peter Yarsky. “Core Design and Reactor Physics of a Breed and Burn Gas-cooled Fast Reactor”. PhD thesis. Massachusetts Institute of Technology, 2005, p. 248.
- [156] Guanheng Zhang. “Advanced burner reactor with breed-and-burn thorium blankets for improved economics and resource utilization”. PhD thesis. University of California, Berkeley, 2015. DOI: 10.1080/00295450.2017.1337408.
- [157] G. Zhang et al. “Uncertainty Quantification of ABR Transient Safety Analysis”. In: *Proceedings of ANS Best Estimate Plus Uncertainty International Conference (BEPU 2018)*. 2018.
- [158] M. Y. Zheng et al. “Performance of radial fuel shuffling sodium cooled Breed and Burn reactor core”. In: *Annals of Nuclear Energy* 96 (2016), pp. 363–376. ISSN: 18732100. DOI: 10.1016/j.anucene.2016.06.019.
- [159] Meiyin Zheng et al. “Preliminary design study of a board type radial fuel shuffling sodium cooled breed and burn reactor core”. In: *Nuclear Engineering and Design* 278 (Oct. 2014), pp. 679–685. ISSN: 0029-5493. DOI: 10.1016/J.NUCENGDES.2014.08.017.

Appendix A

Quantifying Errors Introduced by r-z Models

Since the calculations with ADOPT utilize r-z neutronic models, it is important to confirm that the predictions made on the basis of r-z modeling are sufficiently similar to those provided by more detailed, fully hexagonal models. This section will use two example r-z B&B cores along with geometrically-optimal hexagonal approximations to them in order to quantify differences in key performance parameters.

A.1 Inherent Approximations of r-z Models

The use of r-z models introduces approximations of two types: (1) geometry and (2) materials. Geometrical approximations derive from the fact that assemblies are grouped into volume-conserving cylindrical annuli. Therefore the hexagonal nature of both the internal core lattice and the exterior core boundary is lost. From the perspective of minimizing leakage, a cylindrical core is more optimal than the irregular shapes resulting from hexagonal lattices. However the impact of inner-core geometrical approximations is less clear. For large cores, the number of assemblies per batch may not be enough to make a continuous ring in the hexagonal setting, especially in batches towards the radial periphery. These batches will then consist of discontinuous assemblies that may be separated by multiple mean-free-paths. In r-z models, these within-batch gaps are closed, eliminating paths through the batch from the core center to the radial periphery. Additionally, assemblies in outer batches may be thinned, meaning that the thickness of a batch's annular representation is smaller than the flat-to-flat of its assemblies. These effects have complicated interactions, making it difficult to predict their overall impact through analytic means.

Another geometrical approximation introduced may be core symmetry. In hexagonal lattices, perfect core symmetry is difficult to obtain if the number of assemblies per batch is to be kept constant, as is desirable in B&B cores. Core symmetry is only possible if the number of assemblies per batch is both even and divisible by three. However for the

central batch, the number of assemblies minus one must be divisible by three for symmetry, if the central position is to be included. This paradox can be circumvented by cleverly placing control assemblies or leaving positions empty, but in general it is difficult to achieve perfect core symmetry for arbitrary combinations of core size and batches. The impacts of asymmetry would be readily apparent in predictions made with hex-z models, but r-z models hide asymmetry and thus lose some amount of fidelity.

A final geometric approximation results from the placement of control assemblies in specific lattice locations. Without a hexagonal lattice, different choices can be made to approximate the control system in an r-z model, but the impacts of these choices are unclear due to the fact that control assemblies are typically spaced apart from each other and do not form continuous regions. Additionally, the control assemblies are, by design, strong absorbers, and thus prone to self-shielding effects which require accurate modeling. The r-z control system representation has not been previously investigated in B&B cores, and is perhaps the most uncertain in its impacts.

Material approximations result from the fact that a given region is assigned a single material in most neutronic calculations. By grouping assemblies into a single region, they are smeared together and depleted at a uniform power. Depending on the core layout and power profile, this smearing within r-z models can potentially move materials of high importance into regions of lower importance or vice-versa, resulting in a loss of fidelity, an obfuscation of the power profile, and impacts on core criticality that are difficult to predict.

A.2 Previous Investigations

Previously, a couple of brief studies to validate the use of r-z models in B&B core analysis have been performed. Heidet [54] examined two cases: (1) an 8-batch, 571 assembly, sodium-cooled, U-Zr(10) fueled, 3000 MW_{th} core with melt-refining during the transition to equilibrium; and (2) a 16-batch, 445 assembly, sodium-cooled, U-Zr(6) fueled, 1000 MW_{th} core with an AIROX process at equilibrium. For the larger core, close agreement was found between k_{eff} predicted by r-z and hex-z models, but discrepancies of up to 15% were found in the radial power peaking factor, with the r-z prediction being less than that of the hex-z model. In the smaller core, similar trends prevail, with the predicted k_{eff} always higher in the r-z model. For the smaller core, detailed power fractions are provided that show agreement to within 0.5% of core power for all batches. While good agreement in general trends between the two models is seen, a few key aspects of the comparison have been left out: (1) it is not clear how the hexagonal layouts were determined, (2) it is stated if the materials in each assembly are batch-averaged in the hex-z models, and (3) no control assemblies are included for the larger core. Further, limited discussion is provided about the source of the errors.

A second comparison between r-z and hex-z models was performed by Hou et al. [66] in their studies of 3D shuffling in B&B reactors. Upon determining the optimal shuffling scheme through the use of r-z models, the equilibrium cycle is modeled again using a hex-z layout.

With this second model, the equilibrium cycle error is quantified, although details of the hex-z model core layout, assembly-level shuffling scheme, and assembly material compositions (i.e. batch-averaged versus assembly-specific) are not provided. Additionally, it is not clear if a control system is included in either model. Regardless, the comparison shows relatively good accuracy between the two models, with the r-z model underpredicting cycle length and peak cladding damage by roughly 5% and 2%, respectively. Converse to the results of Heidet, however, Hou et al. show a slight overprediction of the radial power peaking factor. No discussion of k_{eff} evolution or power profile is provided. Hou et al. make the claim that the errors stem primarily from the approximation that all assemblies within a batch are depleted at the same power level in r-z models, although no results are presented to substantiate this claim.

A.3 Comparison Procedure

Although r-z models approximate the reality of hex-z geometry, the present comparison flips the paradigm, taking the r-z model as the reference and the hex-z model as the approximation, using a hex-z model that is geometrically as close as possible to the r-z model. In this way, the results from the two models are compared and it is claimed that, in the *best case*, r-z models will result in some minimum amount of error. For each core, the comparison is carried out in a number of steps. First a small r-z B&B core is developed with ADOPT. The second step is to determine an optimal approximation of the r-z core using a hexagonal lattice of fuel assemblies. This is achieved through the linear programming formulation outlined in Section A.4. Using the determined hexagonal layout and characteristics from the r-z ADOPT models, a hex-z model of the same core is built in Serpent 2 utilizing batch-average materials to account for the impact of geometric, but not material, approximations. To quantify the impacts of material approximations, an additional hex-z model is built where each assembly is depleted separately. Because initial material compositions are taken from the r-z equilibrium calculation, the two hexagonal models are the same at Beginning of Equilibrium Cycle (BOEC). The two versions of the hex-z model are referred to as ‘hex-z-BA’ and ‘hex-z-AS’ for Batch-Averaged and Assembly-Specific fluxes, respectively.

Each eigenvalue calculation uses the same number of histories selected to ensure converged fission sources and assembly power uncertainties of less than 1%. Burnup calculations split the equilibrium cycle into three equal-length steps and use ENDF/B-VII.1 cross-section libraries. After all calculations complete, their results are compared to determine the impact of the modeling paradigm employed. Due to the way this comparison is performed and limitations of ADOPT, it is stressed that this comparison does not quantify the impacts that r-z models have on the determination of equilibrium cycle characteristics, but rather simply takes a given r-z core model and determines how accurate it is compared to a hex-z model of the same initial composition. To quantify the impact on equilibrium cycle characteristics would require knowledge of the shuffling scheme on an assembly-basis, which is outside the scope of this work.

A.4 Finding an Optimal hex-z Approximation

To find the hex-z layout that best approximates an r-z model, a mixed-integer linear programming (MILP) formulation has been developed. In this formulation, δ_{ik} is a binary variable indicating if assembly i is in batch k , n is the user-specified number of batches, $k \in \{1, \dots, n\}$ is the batch number, m is the user-specified number of assemblies per batch, $p = mn$ is the number of assembly positions, $i \in \{1, \dots, p\}$ is the assembly position number, d_{ik} is data specifying how far assembly position i is from batch k , $\ell \in \{1, \dots, 6\}$ is a 1/6 core wedge, and $\zeta_{i\ell}$ is data indicating if assembly position i is in wedge ℓ .

This formulation takes as input the geometrical parameters of an r-z model and returns the hexagonal layout that as closely approximates the r-z geometry as possible. The objective function consists of two terms: the first groups assemblies of a given batch close to the location of the batch in the r-z model, and the second promotes 1/6 core symmetry. Constraint (1) says that each batch must have m assemblies, while constraint (2) says that each position can only have a single assembly. Constraints (3) and (4) are used to promote core symmetry. Finally, constraint (5) says that each assembly may be placed in exactly one location. Using this formulation, the hexagonal layout that most closely matches an r-z geometry can be obtained without user-bias.

$$\begin{aligned}
 & \min_{\delta_{ik}} \sum_{k=1}^n \sum_{i=1}^p (\delta_{ik} d_{ik} (n - k + 1)) + \sum_{j=1}^6 s_{j,j+1} : \\
 \textcircled{1} \quad & \sum_{i=1}^p \delta_{ik} = m && \forall k \\
 \textcircled{2} \quad & \sum_{k=1}^n \delta_{ik} = 1 && \forall i \\
 \textcircled{3} \quad & s_{j,j+1} \geq |x_{j,k} + x_{j+1,k} - \frac{p}{6n}| && \forall k, \forall j \\
 \textcircled{4} \quad & x_{\ell k} = \sum_{i=1}^p \delta_{ik} \zeta_{i\ell} && \forall k, \forall \ell \\
 \textcircled{5} \quad & \delta_{ik} \in \{0, 1\}
 \end{aligned}$$

To study the effects of assembly thinning, two variants of the same core, denoted ‘C1’ and ‘C2’, are examined. The user-specified and ADOPT-calculated parameters of both cores are given in Tables A.1a-A.1b. The user-specified shuffling schemes and MILP-determined layouts are depicted in Figure A.1. Similar to the reference core outlined in Section 2.1, the core employs U-Zr(10) feed fuel of natural enrichment with a smear density of 75% and HT-9 cladding/duct with vented fuel pins.

Table A.1: User-specified and ADOPT-calculated design parameters for reference cores C1 and C2.

(a) User-specified			(b) ADOPT-calculated		
Parameter	Value	Units	Parameter	Value	Units
Core Power	1400	MW	Equilibrium Cycle Length (C1, C2)	1671, 675	EFPD
Inlet Temperature	355	° C	Fuel Residence Time (C1, C2)	32.05, 25.87	EFPY
Outlet Temperature	510	° C	Rod Pitch	1.3608	cm
Number of Assemblies	196	-	Rod Outer Diameter	1.2208	cm
Number of Control Assemblies	0	-	Clad Thickness	0.6104	mm
Assembly duct Flat-to-Flat	18.0	cm	Pins per Assembly	169	-
Duct Thickness	3.50	mm	Fuel Volume Fraction	38.04	%
Inter-Assembly Gap	2.00	mm	Gap Volume Fraction	12.68	%
Number of Batches (C1, C2)	7, 14	-	Cladding+Wire Volume Fraction	12.73	%
Assemblies per Batch (C1, C2)	28, 14	-	Coolant Volume Fraction	29.50	%
			Duct Volume Fraction	7.04	%

To quantify the impacts of r-z models on control system worth predictions, C1 is modified to include control assemblies using four different representations, shown in Figure A.2. The hex-z model (Model (d)) is determined first by placing seven control assemblies to promote azimuthal symmetry while obtaining acceptable control system worth for the core. Once control assemblies are placed, the remaining hexagonal lattice is populated using the MILP methodology. With the hex-z Model (d) as a baseline, three r-z models with various approximations for the control system are constructed. Model (a) augments the original r-z model to have an annular region of control material between Batches 2 and 3, with the control annulus conserving the volume of seven control assemblies. Model (b) is an improvement upon Model (a) by shifting some of the control material to a central cylindrical control region to represent the central control assembly. Finally, a hybrid approximation, inspired by Heidet [54], was configured as Model (c). Model (c) keeps the r-z representation for the fuel while placing control assemblies in the specific hexagonal layout of the hex-z model, adjusting batch radii as necessary to conserve volumes.

The control assemblies use natural B_4C as the absorbing material, with materials and volume fractions taken from the metallic ABR design [21]. Control system worth is predicted by fully inserting all control assemblies simultaneously at nominal operating conditions and taking the difference between uncontrolled and controlled k_{eff} . This procedure is repeated for each of Models (a)-(d) at different levels of burnup, and predictions of system worth are compared between models.

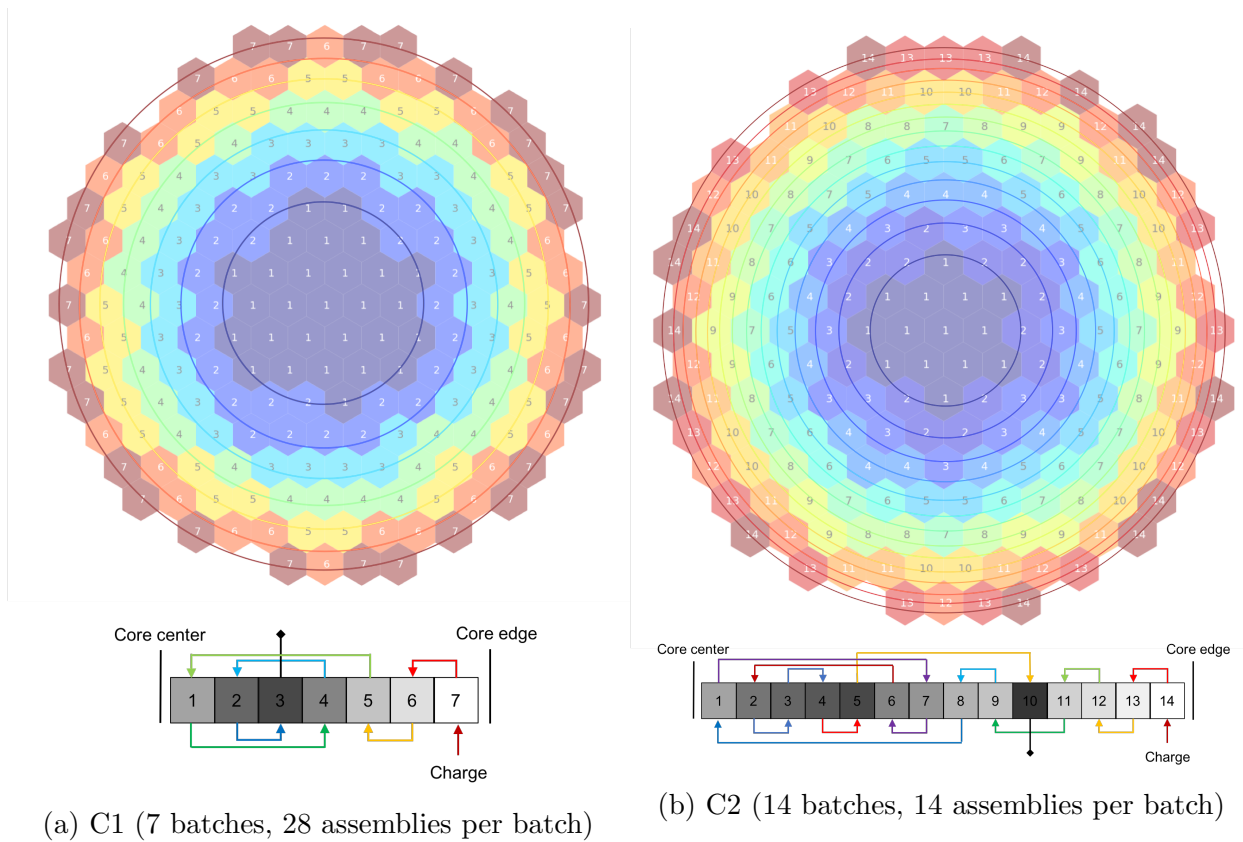


Figure A.1: Hexagonal layouts for C1 and C2 determined through the MILP optimization overlaid with the corresponding r-z models, and their shuffling schemes.

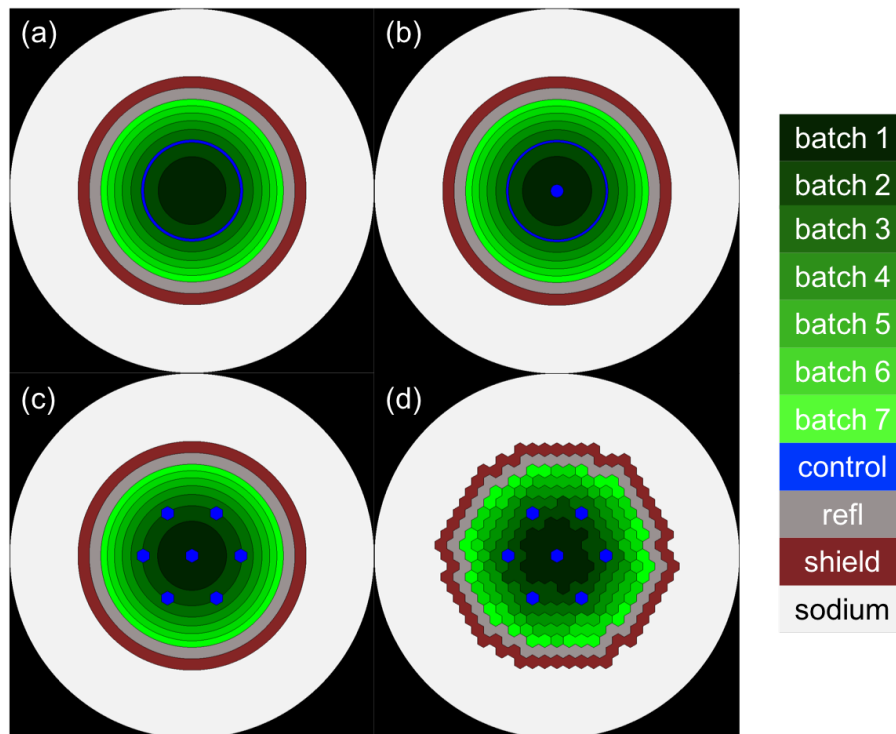


Figure A.2: The four different models used for predicting control system worth showing the approximations used in representing the control assemblies. (a) cylindrical absorber material placed between Batches 2 and 3, (b) cylindrical absorber material placed between Batches 2 and 3 as well as in the core center, (c) hexagonal control assemblies placed within an r-z model, (d) reference hex-z model

A.5 Results

Integral Neutronics Parameters

Figure A.3 shows the evolution of a variety of integral parameters in C1 predicted by both r-z and hex-z models, along with their percent differences. The trends are similar between C1 and C2, and therefore plots for C2 have been omitted. The differences are always less than 2.5% for all examined parameters, with β_{eff} and k_{eff} showing the largest and smallest differences, respectively. Figure A.3a shows that k_{eff} is always overpredicted by the r-z model, indicating a slightly better neutron economy in the r-z model. This result is consistent with results reported by Heidet [54] for the smaller core. In general, all models show acceptable agreement for the parameters investigated, with consistent prediction of trends and small absolute differences.

Power Distribution

For reference, Figure A.4 shows batch-integrated powers at each burn step as calculated with the r-z models. Figure A.5 then shows batch-integrated power percent differences between the r-z and hex-z-BA models. Differences are largest in the outermost batch, where the flux is lowest, power production is the smallest, and assemblies within the batch are most separated from each other. In high-power batches where batch assemblies are closely packed, percent differences do not exceed 5%, which is consistent with results of the smaller core investigated by Heidet [54]. Differences in the outer batches are much larger in C2 than C1, which might be expected due to the greater distances between assemblies of the same batch for C2. This trend is depicted in Figure A.6, which shows maximum percent differences in batch-integrated powers as a function of burnup. For both C1 and C2, the maximum difference decreases as the cycle progresses. This trend aligns with that for k_{eff} , suggesting that the r-z and hex-z models tend to converge with burnup.

Figure A.6 also shows a comparison between the hex-z-BA and hex-z-AS models. The maximum difference between the r-z and hex-z models is greater when assemblies are burned individually, as opposed to using batch-average compositions. However the impact of assembly-specific burnup is seen to be smaller than that of using hexagonal geometry in the first place, suggesting that geometric approximations are more important than material approximations in obtaining accurate predictions. This suggestion is further supported by examining the radial power peaking factors calculated using all three models, as shown in Figure A.7. The r-z peaking factors are calculated via two different methods: the coarse method uses powers integrated over each batch, while the fine method tallies power on a finer r-z mesh, with each mesh an annulus equal in volume to one fuel assembly. The hex-z peaking factors are calculated on an assembly basis. The radial peaking factor is underpredicted by r-z models by as much as 14% in both cores. Using a finer mesh in the r-z models provides improvements, but estimates are still off by as much as 6% and 11% in C1 and C2, respectively. Between hex-z-BA and hex-z-AS models, however, the difference is always less than 2%.

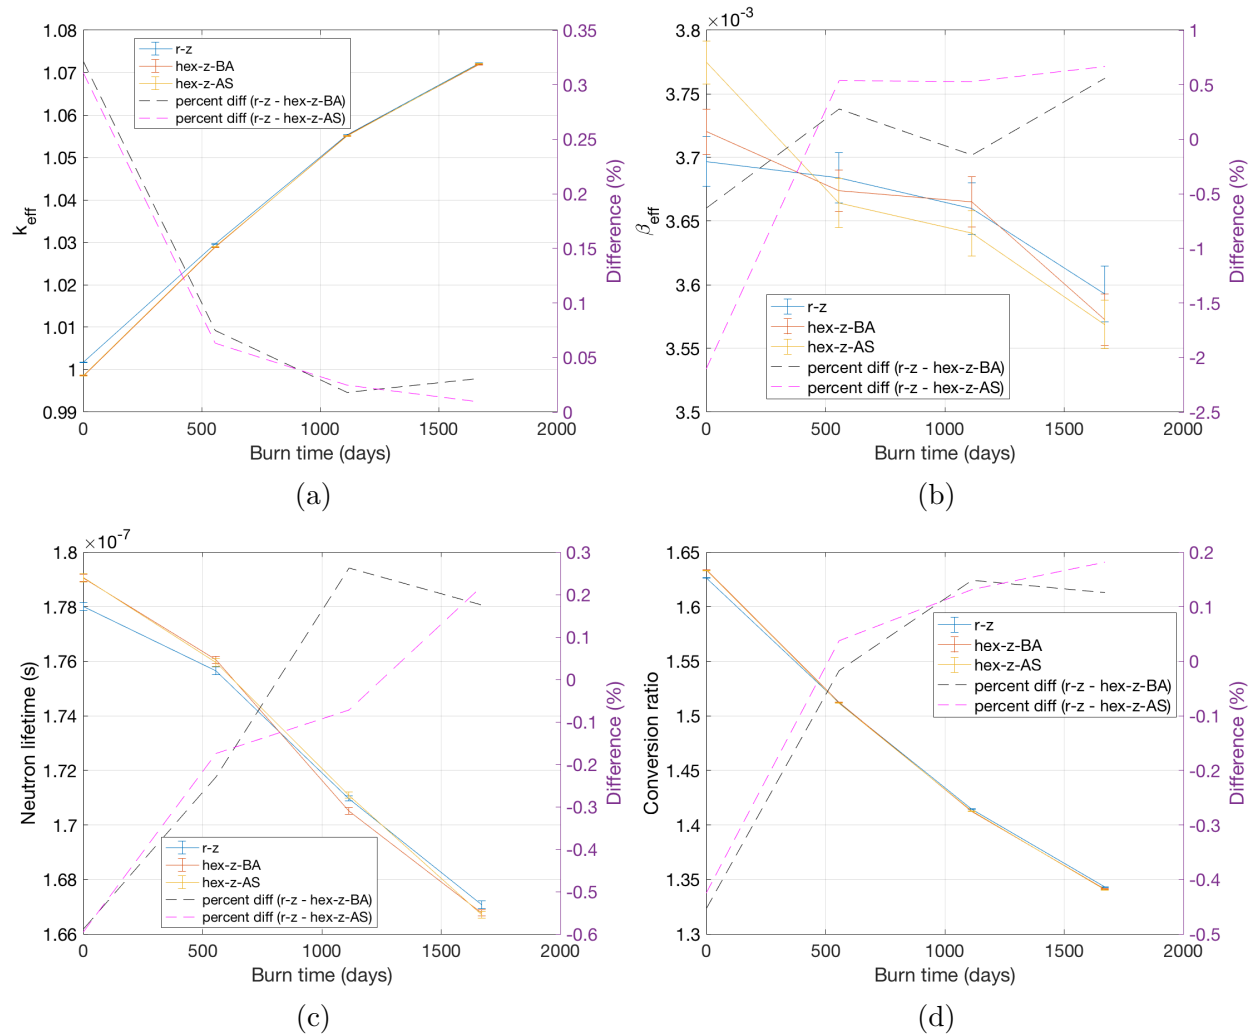



Figure A.3: Evolution of integral neutronics parameters for all variants of C1 and their percentage error.


Control System Worth

Figure A.8 shows the control system worth as predicted by Models (a)-(d) over the equilibrium cycle. All three r-z models overpredict the system worth at BOEC by at least \$6 (2236 pcm), but this difference decreases to roughly \$3 (1118 pcm) at End of Equilibrium Cycle (EOEC). These discrepancies are consistent, although somewhat less than, those seen by Girardin et al. [40].

Not only are the absolute values of control system worth mispredicted, but the trend with burnup is mischaracterized as well. Model (d) predicts system worth to slightly increase, whereas the other models predict significant decreases with burnup. Specifically, Model (a)




(a) C1




(b) C2

Figure A.4: Batch-integrated power in each batch over the four burn steps in C1 and C2 as calculated by the r-z models.



(a) C1



(b) C2

Figure A.5: Difference between batch-integrated power predicted by the r-z and hex-z-BA models for each batch over the four burn steps in C1 and C2. Difference is r-z – hex-z-BA.

shows a decrease by nearly half, resulting mainly from the lack of a central control assembly combined with the radially-inward power shift as the cycle progresses. It is expected that Model (c) would show the best agreement with the hexagonal layout of Model (d), and indeed both models exhibit slowly varying worths with increasing burnup. However both the trend and the absolute worth predicted by Model (c) are different than from Model (d).






Figure A.6: Absolute value of the maximum percent difference between batch-integrated power of the r-z, hex-z-BA, and hex-z-AS models over the four burn steps in C1 and C2.

To understand why the trend is mischaracterized, it is useful to examine system worth with all rods except the central rod inserted, as depicted in Figure A.9 for Models (c) and (d). In this case, the trends are very similar between the two models. Considering the results of Figure A.7a, the trend mischaracterization can be understood as the result of differences in power peaking. The hex-z model predicts sharper power peaking, which at BOEC falls in the region between the central rod and the main control assembly ring, whereas the r-z model spreads power further into Batches 4 and 5. This allows for the system worth to be artificially inflated at BOEC in the r-z models. But as the power shifts towards the core center with burnup, system worth becomes dominated by the central control assembly in both the r-z and hex-z cases.

Although the different trends can be accounted for, Figure A.9 still shows a significant difference in the absolute control system worth. This discrepancy is consistent with the results reported later in Section 2.3 showing a very strong sensitivity to control system layouts. Thus, what appears to be a mischaracterization in the trend of system worth is actually the result of a large overprediction at BOEC by the r-z models combined with a high sensitivity towards specific control system locations.



(a) C1



(b) C2

Figure A.7: Radial power peaking factors predicted by the various models over the four burn steps in C1 and C2.




Figure A.8: Total control system worth at hot full power as predicted by Models (a)-(d) for each burn step during the equilibrium cycle.




Figure A.9: Control system worth without the central rod inserted at hot full power as predicted by Models (c) and (d) for each burn step during the equilibrium cycle.

A.6 Summary of r-z/hex-z Comparison

This brief study has demonstrated the discrepancies that might be expected when using r-z neutronic models to predict the core performance of B&B cores. It has been shown that many integral parameters, such as k_{eff} and β_{eff} , can be predicted correctly to within a couple percent. Similarly, the batch-integrated power production can be adequately predicted with r-z models, except for perhaps in peripheral batches with very low power generation. However, two important aspects of core performance are not sufficiently characterized through r-z modeling: (1) the radial peaking factor, and (2) the control system worth. This thesis therefore undertakes a more rigorous study utilizing hexagonal neutronic models in the design of the control system for the reference core, as is explored in Section 2.3.

Appendix B

Reduction of Core Volume

In order to reduce radial leakage as much as possible, B&B cores are typically designed with an abundance of radial batches. Similarly the cores are made as tall as possible while satisfying pressure drop constraints. Such a design strategy serves to optimize the neutron economy, but may be at odds with other design concerns like economics and safety. The economic penalty arises mainly in relation to having a very large core with the associated material inventories. These costs could be directly reduced by making the core smaller in volume. The degraded safety performance associated with such low leakage arises from two sources: (1) the negative impact of low-powered assemblies on the average coolant outlet temperature and the consequentially higher-than-necessary outlet temperatures of peak-powered assemblies, and (2) the strongly positive coolant void feedback mechanism which results from the low leakage and drives transient performance in an unfavorable direction. These points are explored in more detail throughout this appendix in order to optimize the core design with consideration to aspects other than neutron economy.

B.1 Reduction in the Number of Fuel Batches

While it has been extensively discussed that the outer fuel batches produce very little power, their major purpose is in fact not to produce power but to produce fissile material, primarily Pu-239. Figure 2.12 shows that the maximum local enrichment levels of Pu-239 in Batches 12 and 11 are roughly 0.14% and 0.42% at EOEC, respectively. These low levels of enrichment even after two burn cycles may call into question the real benefit of having those batches in the first place, since fissile content is built up much more quickly in subsequent burn cycles as the fuel is moved closer to the region of high flux. Further evidence in support of eliminating these outer batches comes from Figure 2.18c where it is shown that the worth of the fuel in these batches is essentially zero, even at the axial center where fissile content is the highest.

On the one hand it may be true that these outer batches provide very little benefit, but as long as they do not hurt core performance from any other perspective, their small beneficial

contribution to the tight neutron economy may be justified. Figure B.1 demonstrates that this notion is not true from the perspective of reactor orificing. Repeating the analysis of Section 3.2, except in this case removing Batches 11 and 12, shows how reactor orificing could benefit from removing the low-powered radial fuel batches. In order to achieve a core-average outlet temperature of 510°C , the maximum outlet temperature for an assembly can be reduced by almost 60°C . Not only does this potentially increase the viability of using a static grid plate orificing strategy, it also demonstrates that safety margins could be improved and highlights a tradeoff between neutron economy and reactor safety.




Figure B.1: Core-average outlet temperature as a function of the assembly-average temperature increase constraint at BOEC in the for the reference core, as well with Batches 11 and 12 removed.

Other potential benefits associated with removing one or more fuel batches include:

1. Increase in core power density
2. Reduction in the number of shuffle operations during refueling
3. Reduction in up-front fuel costs

In order to understand the viability of removing these two outermost batches, the reference core is redesigned with less batches, using the same number of assemblies per batch but simply removing a batch's worth of assemblies. ADOPT is used in order to find the critical

cycle length while all other design inputs as outlined in Table 2.1a are kept constant. Figure B.2 shows how both the equilibrium cycle length and the core leakage are impacted by this change. By removing only Batch 12, the core leakage increases somewhat, but the impact on the equilibrium cycle length is very slight. Removing both Batches 11 and 12, however, results in a radial leakage probability of nearly 1% and an increase in the equilibrium cycle length of almost 7.5%, which will be directly reflected as increases in the minimum required burnup and associated material damage levels.




Figure B.2: Critical equilibrium cycle length and radial leakage probability as the number of fuel batches is reduced in the reference core.

Because radial leakage increases, it is expected that a slight improvement in the void worth may result. Figure B.3 actually shows this trend to be the opposite of what was expected, with total core void worth strongly *increasing* as the number of batches is decreased. To understand this behavior, the various terms of the six-factor formula are investigated, as in Equation B.1 where $\eta = \frac{\nu\Sigma_f^F}{\Sigma_a^F}$ and $f = \frac{\Sigma_a^F}{\Sigma_a^F + \Sigma_a^{NF}}$. This equation can be simplified into the three-factor formula because we are dealing with a fast spectrum, and therefore the resonance escape probability, fast fission factor, and non-leakage probability for thermal neutrons are meaningless.

$$k = \eta f \epsilon p P_{NL}^{fast} P_{NL}^{thermal} = \eta f P_{NL}^{fast} = \frac{\nu\Sigma_f^F}{\Sigma_a^{TOT}} P_{NL}^{fast} = \frac{\nu\Sigma_f}{\Sigma_c + \Sigma_f} P_{NL} \quad (\text{B.1})$$




Figure B.3: Total core void worth at BOEC in the reference core as the number of batches is reduced.

Using Equation B.1 we can arrive at a closed-form expression for the coolant density reactivity coefficient in terms of quantities which can be tallied using Monte Carlo perturbation theory. Taking the derivative with respect to sodium density, Na , and taking $\frac{\partial P_{NL}}{\partial Na} \approx 0$ (to be justified later),

$$\frac{\partial k}{\partial Na} = \frac{\partial(\frac{\nu\Sigma_f}{\Sigma_c + \Sigma_f})}{\partial Na} P_{NL} + \frac{\partial P_{NL}}{\partial Na} \frac{\nu\Sigma_f}{\Sigma_c + \Sigma_f} \approx \frac{\partial(\frac{\nu\Sigma_f}{\Sigma_c + \Sigma_f})}{\partial Na} \quad (\text{B.2})$$

Now taking ν to be a constant over the energy range of interest (also justified later), this derivative can be evaluated using the quotient rule. Written in terms of sensitivity coefficients where S_j^i represents the sensitivity of response i to perturbation j ,

$$S_{Na}^k = \frac{\frac{\partial k}{\partial Na}}{\frac{k}{Na}} = \nu \frac{\Sigma_f \Sigma_c (S_{Na}^{\Sigma_f} - S_{Na}^{\Sigma_c})}{(\Sigma_c + \Sigma_f)^2} \quad (\text{B.3})$$

Each of the terms in Equation B.3 can be estimated using the GPT capabilities in Serpent. Their values are provided in Table B.1 for the nominal core along with the reduced-batch cores.

To justify the assumption that $\frac{\partial P_{NL}}{\partial Na} \approx 0$, Table B.2 shows the values of S_{Na}^k tallied directly by Serpent as compared to those calculated using Equation B.3 and the values in

Table B.1: Talled contributions to the coolant void reactivity at BOEC in the nominal and reduced-batch cores.

	$S_{Na}^{\Sigma_f} \times 10^2$	$S_{Na}^{\Sigma_c} \times 10^1$	$\Sigma_f \times 10^2$	$\Sigma_c \times 10^2$	ν
Nominal	-5.405	-1.138	2.29	4.17	2.907
11 Batches	-5.682	-1.210	2.21	4.05	2.907
10 Batches	-6.117	-1.303	2.12	3.89	2.908

Table B.1, along with the differences between them. In general the difference between the two methods of calculation is less than 5%, indicating good agreement and justifying the usage of the approximations in Equation B.3.

Table B.2: Difference between S_{Na}^k tallied directly by Serpent as compared to those calculated using Equation B.3.

	Serpent-calculated	Equation B.3	Error (%)
Nominal	-3.922×10^{-2}	-3.980×10^{-2}	-1.47
11 Batches	-4.462×10^{-2}	-4.261×10^{-2}	4.50
10 Batches	-4.615×10^{-2}	-4.589×10^{-2}	0.56

Using the values in Table B.1 it is possible to assess what drives the void coefficient to become stronger in cores with less batches if we can find the derivatives of the sensitivity relation in Equation B.3 (i.e. the sensitivities on the sensitivity). Then by plugging in representative numbers for the nominal core and estimating the change resulting from a representative perturbation in each parameter as a finite difference (i.e. $\frac{\partial S_{Na}^k}{\partial \Sigma_f} \approx \frac{\Delta S_{Na}^k}{\Delta \Sigma_f} \rightarrow \Delta S_{Na}^k = \Delta \Sigma_f \times \frac{\partial S_{Na}^k}{\partial \Sigma_f}$), the parameter with the largest impact can be found. This process is carried out in Table B.3.

As a check on the accuracy of this method, the values in Table B.3 are summed, converted to total void worth in units of dollars, and compared with the values in Figure B.3. The two methods give a difference in total void worth from 12 to 10 batches that are within 0.4% of each other, indicating very good agreement and providing confidence in this methodology for breaking down the void reactivity coefficient effect as the number of batches is reduced.

Table B.3 indicates that the change in $\frac{\partial \Sigma_c}{\partial N_a}$ between the nominal core and the reduced-batch cores is what most strongly leads to the increasing coolant void worth, with the change in $\frac{\partial \Sigma_f}{\partial N_a}$ working in the opposite direction with slightly smaller magnitude. As shown in Table B.1, these derivatives become more strongly negative as the number of batches decreases, likely because the spectrum softens very slightly (average neutron energy as calculated by Equation 6.9 is 0.397, 0.394, and 0.392 MeV for the nominal, 11-batch, and 10-batch cores,

Table B.3: Sensitivities on the coolant density feedback to the parameters which influence this effect.

Variable, x	Expression for $\frac{\partial S_{Na}^k}{\partial x}$	$\Delta S_{Na}^k(x_{10\text{-batch}} - x_{nominal})$
$S_{Na}^{\Sigma_f}$	$\nu \frac{\Sigma_f \Sigma_c}{(\Sigma_c + \Sigma_f)^2}$	-4.73×10^{-3}
$S_{Na}^{\Sigma_c}$	$-\nu \frac{\Sigma_f \Sigma_c}{(\Sigma_c + \Sigma_f)^2}$	1.09×10^{-2}
Σ_f	$\nu \Sigma_c (S_{Na}^{\Sigma_f} - S_{Na}^{\Sigma_c}) \frac{\Sigma_c - \Sigma_f}{(\Sigma_c + \Sigma_f)^3}$	-4.04×10^{-4}
Σ_c	$\nu \Sigma_f (S_{Na}^{\Sigma_f} - S_{Na}^{\Sigma_c}) \frac{\Sigma_f - \Sigma_c}{(\Sigma_c + \Sigma_f)^3}$	7.76×10^{-4}
ν	$\frac{\Sigma_f \Sigma_c (S_{Na}^{\Sigma_f} - S_{Na}^{\Sigma_c})}{(\Sigma_c + \Sigma_f)^2}$	~ 0

respectively) and thus the neutrons are closer towards the thermal region where cross-sections vary much more rapidly than in the fast region. This table also indicates that the changing magnitudes of Σ_f and Σ_c between cores, while they do have an impact, are much smaller than their derivatives. Finally, the impact of ν is virtually zero because its value changes very little between cores.

This discussion has highlighted that, while there may be benefits related to the ease of orificing using the static integral method and core power density, reducing the number of fuel batches can have unintended consequences on other aspects of operation. The primary ways in which degraded performance may result is in the increase in required equilibrium cycle length and an increase in the core void worth. The first aspect was anticipated, and it can be limited in a straightforward manner to an acceptable value if only a single fuel batch is removed, bringing the core from 12 to 11 batches. However the increase in void worth with fewer batches (and higher radial leakage) was unexpected and represents a potential reason to keep the number of fuel batches at 12. Such a tradeoff does not have a clear best answer and must ultimately be chosen by the designer according to broader design goals.

B.2 Elimination of Radial Reflector Assemblies

Because the radial reflector assemblies are not shuffled through the core, their power levels do not change significantly over the cycle and thus they do not negatively impact orificing or safety in the same way that the low-powered batches do. However, in such a large core, the number of radial reflector assemblies is large, being nearly 15% of all assemblies in the core (roughly double the amount of assemblies in a single batch). Further, the radial reflectors serve a very limited purpose on account of the many blanket assemblies between them and the power producing region. Therefore this section investigates the impact of potentially eliminating the radial reflectors and having only a single row of shield assemblies outside the core.

Table B.4 shows the impact of removing the radial reflector assemblies on radial leakage is very slight, with the radial leakage still being far below 1%. The resulting impact on the critical cycle length is negligible (within statistical inaccuracy of such calculations) since breeding in the outermost batch is already so low. Additionally, the impact on the total core void worth of removing the reflector is found to be beneficial, leading to a reduction in total void worth of approximately 6%. Such results indicate that removing the radial reflector from the reference core provides essentially only benefits, resulting in savings by reducing the necessary size of the core barrel, eliminating the costs associated with the reflector assemblies (purchase, monitoring, replacement, grid plate accommodation, etc.), and improving transient performance.

In addition, calculations were performed to see if the results of this section could be combined with those of the previous by eliminating both Batch 12 and the radial reflector simultaneously. These results are also presented in Table B.4, showing a somewhat larger impact on radial leakage while still being below 1%. The larger leakage in this case does end up leading to a slightly longer equilibrium cycle length by approximately 5%, on par with that when both Batches 11 and 12 were removed in Section B.1. In addition, the coolant void worth ends up increasing again when removing Batch 12 for similar reasons as were outlined in detail in the previous section. Therefore the benefits of removing both the reflector and Batch 12 are not as immediately clear as they are for removing only the reflector.

Table B.4: Radial leakage probability for the reference core and with all radial reflector assemblies removed.

	Radial Leakage (%)		Total Void Worth at BOEC (\$)
	BOEC	EOEC	
Nominal	0.20	0.23	11.82
No Reflectors	0.32	0.37	11.07
No Reflectors + 11 batches	0.66	0.73	12.63

B.3 Summary of Reduced Volume Performance

This short study has examined the potential benefits to be had by reducing the core volume in two ways: (1) reducing the number of fuel batches and (2) removing the radial reflector assemblies. Reducing the number of fuel batches allows for easier orificing and associated reductions in peak temperatures, as well as an increase in core power density, however it can lead to increased leakage (i.e. worse neutron economy and higher cladding damage) and worsened core void worth. This latter point was unexpected because it goes against the general thinking that higher leakage implies more favorable kinetic properties, and thus was the subject of additional analysis which determined the cause to be a change in the

effective one-group capture cross-section sensitivity with less fuel batches. While removing two batches has notable negative impacts on both neutron economy *and* void worth, removing just a single batch has minimal impact on neutron economy and a smaller (though notable) impact on void worth. Due to the competing effects of this design change, its general desirability is unclear and subject to the goals of the designer.

Because of the many radial blanket assemblies, removing the radial reflector results in a very slight increase in neutron leakage which provides a slight benefit to void reactivity through an increase in leakage without being a detriment to the neutron economy. In light of the benefits of removing the reflector assemblies from the perspectives of power density and cost, this option is likely attractive.