Multiphysics Simulation and Verification of Molten Salt Reactors using GeN-Foam

A thesis submitted by
Mario A. Ponce Tovar

in partial fulfilment of the requirements
for the degree of
Master of Nuclear Engineering

conferred by both
Université Paris-Saclay
Universitat Politècnica de Catalunya

with supervision from
Andreas Vigand Schofield
Chief Physicist, Seaborg Technologies
Copenhagen, Denmark

... And it never failed that during the dry years, the people forgot about the rich years; and when the wet years returned, they lost all memory of the dry years. It was always that way.

EAST OF EDEN
JOHN STEINBECK

Abstract

In light of the growing need for detailed analysis of up-and-coming reactor technologies, several codes are under development with the aim of simulating multiphysics phenomena. These codes are largely unproven and unqualified, and as such they must be tested prior to being used for this purpose. This work takes on the task of assessing a particular OpenFOAM-based multiphysics code (GeN-Foam) for potential use in deterministic safety analysis of the Compact Molten Salt Reactor (CMSR) developed by Seaborg Technologies. It begins with an evaluation of the code against the MSFR benchmark developed at CNRS. The code is found to pass the benchmark without issue; thereby its fidelity with respect to several MSR-specific phenomena is established insofar as the benchmark results can be considered accurate. Next, a quite recent neutronic verification scheme is applied to the code, based on the extraction of the point-kinetic component of a transient solution. The code is found to be partially verified in this respect, with one issue encountered in the amplitude of the transient response to a sinusoidal reactivity perturbation. On the other hand, the phase response can be considered verified. Finally, the code is applied to the CMSR. A set of multiphysics simulations are performed of this full-core 3-D porous media model. A mesh independence study is performed on the fully-coupled steady-state solution and a sensitivity study is also performed with regard to the calculation of β_{eff} with DNP transport. Transient simulations are performed of bulk reactivity insertion, loss of moderator cooling, and unprotected loss of flow. Results fall well within reasonable expectations and the general safety benefits of molten salt reactors are exhibited. With this, GeN-Foam is positively assessed for continued use in the design and licensing campaign of the CMSR.

Keywords: *molten salt reactor, multiphysics, verification, GeN-Foam, CFD.*

Résumé

À cause du besoin croissant d'analyse détaillée des nouvelles technologies de réacteurs avancés, plusieurs programmes sont en cours de développement dans le but de simuler des phénomènes multiphysiques. Ces programmes ne sont pas encore prouvés ou qualifiés, ils doivent ainsi être testés avant d'être utilisés pour cette fin. Cette thèse a pour but d'évaluer un programme multiphysique basé sur OpenFOAM (GeN-Foam) en vue d'une utilisation potentielle dans l'analyse de sécurité déterministe du réacteur compact à sel fondu (CMSR) développé chez Seaborg Technologies. Cette thèse commence par une évaluation du programme par le *benchmark* MSFR développé au CNRS. Le programme passe l'évaluation référence sans problème. Ainsi, sa fidélité vis-à-vis de plusieurs phénomènes spécifiques aux réacteurs à sel fondu est établie dans la mesure où les résultats de référence peuvent être considérés comme exacts. Ensuite, un processus de vérification neutronique assez récent est appliqué au programme, à base de l'extraction de la composante cinétique ponctuelle d'une simulation transitoire. Le programme s'avère partiellement vérifié à cet égard, avec un problème rencontré dans l'amplitude de la réponse transitoire contre une perturbation de réactivité sinoisoïdale. Néanmoins, la réponse dans la phase peut être considérée comme vérifiée. Enfin, le programme s'applique à la simulation du CMSR. Un ensemble de simulations multiphysiques s'effectue à partir de ce modèle de média poreux 3-D à cœur complet. Une étude d'indépendance du maillage de calcul est réalisée sur la solution en régime stabilisé et entièrement couplé. Une analyse de sensibilité se réalise également par rapport au calcul de β_{eff} avec le transport des précurseurs de neutrons retardés. Des simulations transitoires sont réalisées pour l'insertion homogène de réactivité, la perte de refroidissement du modérateur, et la perte de débit non protégée. Les résultats sont bien dans les attentes raisonnables, et les avantages généraux de sécurité des réacteurs à sel fondu sont démontrés. Ainsi, GeN-Foam s'évalue positivement pour poursuivre son utilisation dans la campagne de conception et de license du CMSR.

Mots clés: *réacteur à sel fondu, multiphysique, verification, GeN-Foam, CFD.*

Contents

Chapter 1

Introduction

1.1 Motivation

In 2018, the UN Intergovernmental Panel on Climate Change released a special report stipulating that greenhouse-gas emissions must be reduced to effectively zero before 2050 in order to avoid the most disastrous effects of climate change (IPCC, 2018). Furthermore, it is foreseen that to have any chance of achieving this, an initial reduction of 45% from 2010 levels must be reached by 2030. The report lays out four different pathways to achieving this goal, and notably, a substantial increase in nuclear energy is included in all of them. This is because, while the study includes the potential for improvements in energy efficiency and more responsible consumption, it recognizes the reality of *Jevons' paradox* — the idea that improvements in efficiency tend to increase consumption to match initial resource use — as well as the growing need for energy in industrializing nations.

This landmark report is paired with a shift in the environmentalist movement from the misanthropic[1] ideal of "returning to nature" to a progressive view which recognizes that the challenge of preserving nature is necessarily wrapped up with countless social issues. As humanity hurtles towards the first climate goal with little progress in carbon reduction, it is becoming increasingly clear that more dramatic measures will need to be taken than installing efficient light-bulbs and solar rooftops. The new focus is on how to produce massive amounts of clean energy and harness it steadily in a way that out-competes fossil-fired power.

At the same time, the nuclear industry is positioning itself to meet the projected demand for clean generating capacity. This effort can be broadly divided into two basic categories; the first, a renewed construction of large scale multi-unit power plants using established technology, and the second, a planned wave of small modular reactors (SMRs) employing cutting-edge technologies which deliver "walk-away" safety among other benefits. It is expected that both categories will be needed and are envisioned to complement each other (as seen in the nuclear roadmaps of the UK, USA, and China, for example). In the latter category, however,

---

[1]In the author's opinion.

significant R&D advances are still needed to bring these reactors to reality. Thus, there are multiple ongoing efforts to improve and validate the design of such reactors while accounting for physical phenomena that are unique to them. These efforts are split between traditional research institutions as well as several spin-off and start-up companies that promise to bring on a new paradigm, with unencumbered thinking and a go-to-market mindset. One of these start-ups is Copenhagen-based Seaborg Technologies, the host of this project.

Perhaps the most important challenge to the near-term deployment of advanced nuclear reactors is licensing. On the one hand, the future of the nuclear regulatory framework is nebulous, with national regulators so far mainly dealing with solid-fueled and water-cooled reactors at the commercial level. On the other, there is a distinct lack of qualified or experimentally benchmarked simulation codes which can be used to make the safety case for advanced reactors. These codes are still under active development, and it remains to be seen how or when they will be validated. Given the pressing need for these technologies, however, it is not acceptable to wait until such codes are qualified before starting to employ them for reactor design. Advanced reactor developers must themselves take part in the verification and validation (V&V) process, ensuring that their simulations are as accurate as possible while working with code developers and research institutions. In the context of a multi-faceted development and licensing strategy, several codes must be assessed for potential use. Therefore, the present work focuses on the assessment and application of a particular multiphysics simulation code which is still under active development: GeN-Foam.

1.2 Objectives and Layout

Seaborg Technologies is developing a molten salt reactor (MSR) with circulating fuel: the CMSR. Therefore, it is essential for the chosen simulation codes to correctly handle the treatment of coolant and fuel as one and the same. Given the novelty of such codes, several options are being considered. In connection with the SEALION project, Seaborg has been developing its own multiphysics suite based on a transient fission matrix implementation (Henriroux et al., 2019), which was recently benchmarked (Groth-Jensen et al., 2021). Simultaneously, the company is looking outward at other developers, with a particular eye toward internally coupled codes as a complement to its own externally coupled approach. GeN-Foam is a code which seems promising for this pursuit, as it is being developed with liquid-fueled MSRs in mind (Fiorina et al., 2015).

In terms of its contribution to the host company, this project is a broad assessment of GeN-Foam as a viable tool for the design and licensing campaign of the CMSR. Since GeN-Foam has not yet been officially released, and the process of V&V by the developers has only just begun (De Oliveira and Mikityuk, 2018), this project takes on part of the verification before culminating in a set of simulations of the CMSR.

Concretely, the main goals of this work are:

1. To reproduce the results from a widely accepted benchmark for molten salt reactors (the CNRS-MSFR benchmark) using the latest version of GeN-Foam. This serves two functions: to show that this version of GeN-Foam is (very likely) free of serious flaws in the implementation of the relevant mathematical models including both neutronics and thermal-hydraulics as well as the coupling scheme, and to show that the author of this work gained sufficient competence in the use of GeN-Foam/OpenFOAM to execute the benchmark properly.

2. To perform a more rigorous neutronic verification of GeN-Foam in the form of a point-kinetic decomposition, to further demonstrate that the diffusion neutronics model as implemented in GeN-Foam gives results in line with theoretical expectations. The implication is then that if the mathematical model is an accurate representation of reality, then GeN-Foam is capable of predicting real-world results.

3. To apply the lessons learned from the benchmark and verification steps to a full-core model of the CMSR, adding on additional complexities, including: egress of delayed neutron precursors, Doppler feedback, and more realistic transient/safety analysis.

In terms of content, this document is constrained by an imposed page limit. As a result, it contains only a limited theoretical introduction to the relevant aspects of reactor physics and thermal-hydraulics. Furthermore, the literature review process which gave the author context for this work is omitted. The reader is expected to have an understanding of the principles of nuclear engineering and reactor analysis at least at the senior undergraduate level, although opportunities are occasionally taken to remind the reader of certain key aspects. More advanced topics that pertain to this work are explained in greater detail, e.g. the transport of delayed neutron precursors and the notion of adjoint flux, while others that are only tangentially related may be mentioned but not explained (e.g. certain numerical methods or CFD specifics).

What follows is the overall layout for navigating this work: This chapter has covered the general background and motivation for the project. In the following pages, chapter 2 presents a brief history and general principles of molten salt reactors and an introduction to MSR multiphysics, as well as a description of Seaborg's CMSR. Next, chapter 3 gives an overview of the software and hardware used to perform the simulations presented in this work. It also explains the fundamental aspects of the mathematical models employed in the software. Chapter 4 is the main body of this work; it describes the methodology adopted to achieve the above objectives and discusses the results obtained. For the sake of clarity, only the efforts that directly led to achieving the main objectives of this work are included. Finally, chapter 5 closes out the report with a brief summary and ideas about future work. Throughout the text, various symbols are used in equations which may not be explained if they are considered canonical. If in doubt, the reader can refer to the list of symbols, and likewise the acronyms. Certain figures which are not directly discussed are relegated to the appendix.

Chapter 2

Molten Salt Reactors

2.1 Origins and Revival

From early on in the history of nuclear power, when the first human-induced and sustained fission chain reaction took place under the stands of a football field in Chicago (Fermi, 1952), it was foreseen that a whole host of different fission reactor types could exist, arising from various combinations of fuel, coolant, and moderator (Weinberg, 1976). From aqueous-homogeneous reactors to plasma-bound "gas core" reactors, these span a wide range of potential uses in research, propulsion, as well as commercial power and heat generation. Several of these concepts were promptly studied and tested, and out of these, many were abandoned due to political, economical, or technological reasons (Rofer, 2015). And with a nuclear industry booming on the basis of light-water reactors, there was little financial incentive to pursue alternative technologies through the remainder of the 20th century.

Jumping to the present day, the situation looks much different. Previously abandoned ideas are being reconsidered in response to the global challenge of climate change as well as the internal challenges of the nuclear industry. Seeing this on the horizon, the Gen-IV International Forum (GIF) was conceived to identify and develop the next wave of commercial nuclear reactor technologies, presenting radical departures from currently deployed Gen-III designs (GIF, 2002). Based on criteria involving sustainability, economics, safety, reliability, as well as proliferation resistance, six general technologies were ultimately chosen for continued R&D among the member countries. These are, in descending order of technology readiness level (TRL): the sodium-cooled fast reactor, very-high-temperature reactor, supercritical-water reactor, molten salt reactor[2] (MSR), gas-cooled fast reactor, and the lead-cooled fast reactor (Gougar et al., 2015 — ORNL).

Molten salt reactors are one of the most attractive Gen-IV technologies due to their expected

---

[2]There are actually two main types of MSRs: solid-fueled (molten-salt cooled) and liquid-fueled (fissile material dissolved in the coolant salt). Here, the TRL advantage over gas-cooled and lead-cooled reactors applies only to the solid-fuelled MSR concept, which is estimated by the same ORNL report to have a decade-long headstart on its liquid-fueled counterpart.

high level of inherent safety, owing to their strong negative feedback and high accident tolerance. An especially popular safety feature is the inclusion of a drain tank which would evacuate the fuel from the reactor core and thereby ensure sub-criticality and decay heat removal in the event of a severe accident (EVOL, 2013). Additionally, it may be possible to retain the vast majority of volatile fission products in the fuel salt. This improved safety could be a boon for the business case as well, since the severity of potential accidents is vastly mitigated. In the long-term, MSRs may also be the most practical way to implement a thorium fuel cycle, which is an attractive proposition from the point of view of waste management and sustainability.

MSRs were first conceived as part of the Aircraft Reactor Experiment (ARE) at ORNL, which lasted just a few days and had the eventual goal of enabling an aircraft to fly almost indefinitely. A 2.5 MW_{th} prototype reactor was built for this experiment in 1954, containing circulating fuel salt and a beryllium oxide moderator (Bettis et al., 1957). Then in the late 1960's, under the direction of Alvin Weinberg, lessons from the ARE were applied to the Molten Salt Reactor Experiment (MSRE) — a 7.4 MW_{th} graphite-moderated reactor which operated for five years (Haubenreich and Engel, 1970). It served as a proof-of-concept for a future "chemist's reactor" envisioned by Weinberg, in which a ^{233}U-fueled core would be surrounded by a ^{232}Th salt blanket that would be extensively reprocessed[3], finally closing the fuel cycle (Siemer, 2015). In both reactors, a fluoride salt was used as the basis for the fuel salt eutectic, and many challenges surrounding corrosion were solved[4], while uncontrollable swelling of the graphite moderator hindered the prospect of commercialization (Siemer, 2015).

Nowadays, the reference model conceived by the European partners of GIF is known as the Molten Salt Fast Reactor (MSFR). It picks up more or less where the MSRE left off, aiming to avoid the problem of graphite swelling by foregoing a moderator altogether, and otherwise building on the priceless experience gained with fluoride salts at ORNL. It also adopts the breeding blanket strategy and thorium fuel cycle (Siemer, 2015). There is an ongoing effort at the EU (Euratom) level to develop the MSFR concept and ratchet up its TRL. Recent collaborations have resulted in notable improvement in MSR modelling capabilities, as well some work in chemistry, corrosion control and the design of support systems (SAMOSAFER Consortium, 2020). That being said, these projects have the stated goal of commercialization in the 2050's — a worrisome timeline in the frame of the climate crisis.

2.2 The CMSR

The Compact Molten Salt Reactor (CMSR) is a FUNaK (uranium fluoride-based salt) fueled, thermal-spectrum MSR under development at Seaborg Technologies (Cooper, 2020). It derives its "compact" moniker from its molten sodium hydroxide moderator which enables a greater

---

[3]The possibility of thorium breeding is widely recognized, but it is often overlooked that ^{233}Pa has a prohibitive absorption cross-section paired with a long half-life of 27 days (Hébert, 2016). This necessitates the separation of the ^{233}Pa while it decays.

[4]A nickel-based alloy, Hastelloy-N, was developed to resist attack by fluorides in a pure fuel salt. However, such alloys are not presently classed as structural materials in the ASME code (ASME, 2021). For this reason, current MSR developers are pursuing cladding strategies for components exposed to the fuel salt.

mean power density than graphite-moderated designs (thanks to a high concentration of hydrogen), while avoiding graphite's aforementioned swelling issues. The CMSR is a "single fluid" design in the sense that there is no breeding blanket or separate coolant salt in the primary circuit. Unlike the MSFR, efforts are made to commercialize the design in the near-term: discarding ambitions of a thorium fuel cycle, online refuelling, and more. It further diverges from both the MSRE and MSFR by employing a multitude of fuel tubes in a moderator-filled vessel rather than directly exposing the fuel salt to the moderator or having it flow unguided through the vessel. Thus the geometry of the reactor core more closely resembles a CANDU reactor than anything else (a conceptual rendering is shown in Fig. A.1).

The CMSR is currently in the *basic design* phase — the engineering phase between conceptual and detailed design. Seaborg aims to license and commercialize the CMSR within the decade, setting a blistering pace by the standards of the nuclear industry. This requires an intense modelling effort to provide adequate safety analysis to the regulators engaged. And while the CMSR offers potential advantages over competing designs, it also presents unique challenges for simulation. In particular, from a thermal-hydraulics point of view, the CMSR core is sufficiently different from other MSRs so as to significantly alter the approach for fluid dynamics. Resolving the flow through over two hundred fuel tubes is challenging, along with turbulent phenomena at opposing plena. As such, sophisticated CFD models must be employed in tandem with system and sub-channel codes to paint an accurate picture of the whole primary loop. Then, as with all liquid-fueled MSRs, a detailed multiphysics analysis is needed.

2.3   MSR Multiphysics

Nuclear reactors are complex systems due to their coupling of multiple fields of science and engineering. At a most fundamental level, one must consider the interplay between thermal-hydraulics (the coolant fluid motion and ensuing heat transfer from the nuclear fuel) and neutronics (fission power and DNP generation/distribution within the core). Feedback phenomena such as Doppler-broadening, thermal expansion, and void formation result in the heat transfer process influencing the fission chain reaction and vice-versa. To enable the safe and reliable operation of nuclear plants, these phenomena must be examined in computer simulations, experiments, and commissioning tests. Above all, one must ensure that feedback coefficients are negative during normal operation as well as all postulated accidents[5].

For certain conditions and designs, it is possible to ensure safety criteria are met using conservative approximations, which greatly simplifies deterministic safety analysis (DSA) efforts. However, it has become apparent as the nuclear industry has progressed that some assumptions which seem conservative at first turn out not to be; and ironically, it takes a more detailed analysis to prove such an assumption, since unforeseen phenomena may or may not arise. A classic example is the modelling of Loss-of-Coolant Accidents in LWRs: counter-intuitively, small-break LOCAs may be more likely to lead to core meltdown than

---

[5]Case in point: the 1981 Chernobyl accident can be partially attributed to a design flaw in this respect (Kortov and Ustyantsev, 2013).

large-break LOCAs (Sun et al., 2020). Furthermore, a model which is conservative for one criterion may not be conservative for another — further complicating the conservative approach. This leads to the other DSA strategy: best-estimate plus uncertainty (BEPU), which entails simulating a given scenario with the most accurate models possible and establishing uncertainty bounds on the results. Not only is BEPU becoming increasingly important for established technologies, but the arrival of Gen-IV reactors demands BEPU analysis for most cases because conservative assumptions are harder to justify with less familiar phenomenology. Quite often, the BEPU approach implies the close coupling of neutronics, thermal-hydraulics and/or thermo-mechanics, as this is required to obtain a best-estimate result.

A fundamental consideration for the coupled analysis of different physical processes is the relative time-scale. If two interconnected processes occur at similar rates, there is no chance of modelling them in isolation with any accuracy as long as they are allowed to influence each other. On the other hand, if one process has a much shorter time-scale than the other, then the feedback it receives from the slow process can be neglected over a comparably short period. This is often the case when modelling the reactivity and heat transfer in a conventional reactor. If the fuel and coolant are separate entities, then only the fuel temperature is relevant over short time-scales, as the Doppler feedback within the fuel is sure to curb any sudden power excursion well before the coolant/moderator has time to absorb heat and expand. This is greatly beneficial for modelling solid-fueled reactors as this saves a great deal of effort in thermal-hydraulic modelling of the coolant. On the other hand, liquid-fueled molten salt reactors do not enjoy this partial decoupling of fuel and coolant physics because the fuel and coolant are one and the same, thus thermal expansion of the fuel salt plays a greater role than it otherwise would. Furthermore, the movement of fuel results in heightened non-linearity of the feedback (Cammi et al., 2011) due to transport of delayed neutron precursors (DNPs), which go on to emit delayed neutrons with different importance from the prompt neutrons.

It is evident that the interplay between neutronics and thermal-hydraulics is essential for modelling MSRs. At minimum, multiple coupling iterations are needed between the two in order to arrive at a steady-state solution. Away from steady-state, however, thousands upon thousands of iterations are required to maintain fidelity when the reactor power is allowed to vary freely. Depending on the transient scenario in question, time-steps may need to be reduced to milliseconds or microseconds, with several iterations needed to compute each one. And when these transients take tens or hundreds of seconds to stabilize, this entails a major computational effort. Great care must therefore be taken to facilitate the convergence of coupling solutions at each time-step so as to speed up calculations.

Two main strategies are generally employed for multiphysics coupling: monolithic (e.g. JFNK – He et al., 2020) and partioned coupling (operator splitting – MacNamara and Strang, 2016). These refer to the simultaneous and separate solutions of each physical model's governing equations, respectively. Either strategy requires iteration to arrive at a stable result, although simultaneous schemes are expected to be more stable as they can be fully implicit. Meanwhile, partitioned schemes are straightforward to implement and do not necessarily perform any worse. They can further be implemented either internally or externally (by interfacing different programs). In any case, multiphysics code developers and users must strive to use efficient, robust, and accurate numerical schemes.

Chapter 3

Computational Tools

This chapter describes the software tools used to obtain the results presented in this work, as well as the mathematical models they are based on. The aim is twofold: to provide a brief theoretical background on the governing equations, and to enable the reader to reproduce the results by specifying which features of the respective codes were used and which versions were installed.

Computations have been carried out both locally and on Seaborg's private computing cluster. The latter was used for the heaviest calculations which benefited from substantial parallelization, while comparably light-weight calculations which are included in this work were performed on the author's personal computer.

3.1   OpenFOAM

Open-source Field Operation And Manipulation (OpenFOAM, 2014; Weller et al., 1998) is a C++ toolkit for developing customized numerical solvers for partial differential equations (PDEs), which first emerged in the late 1990's and has been developed long since. While it is widely viewed as a CFD toolkit, it is also suitable for many other kinds of problems including electrodynamics, pricing of financial derivatives, and more. OpenFOAM ships with an extensive set of pre-made thermo-fluid solvers which can be employed directly, modified, or unplugged and packaged elsewhere. An application programming interface (API) is also available for high level application development of tailor-made solvers for new or modified physics.

OpenFOAM employs the finite volume method for the discretization of PDEs. In this approach, integrals containing a divergence term are converted to surface integrals over the surface enclosing each volume (by way of the divergence theorem) before being evaluated (Moukalled, Mangani, and Darwish, 2015). Notably, this allows for the use of unstructured meshes, which are prevalent in typical OpenFOAM applications. When defining mathematical models, OpenFOAM permits users to write the governing equations in roughly "hand-written"

form, which it then solves based on its object-oriented structure. OpenFOAM further allows straightforward parallelization via user-defined geometric partitioning through the included `decomposePar` utility.

Owing to its long history, there are many versions of OpenFOAM which can be downloaded and installed today. For the purpose of this thesis, the distribution used was OpenFOAM-v2012 (December 2020 release from ESI/OpenCFD), which was compiled locally from the source code without modification.

3.2   GeN-Foam

Generalized Nuclear Foam (GeN-Foam) is an OpenFOAM-based solver overseen by the Nuclear Technical Committee of the OpenFOAM collaboration. It is jointly developed by the Paul Scherrer Institute (PSI) and EPFL in Switzerland (Fiorina et al., 2015), and is designed to leverage the open-source nature of OpenFOAM along with the considerable competence in its use within the nuclear industry and beyond. It also implements the parallelization schemes made available in OpenFOAM seamlessly.

GeN-Foam is currently hosted on Gitlab, and there are two main branches for users to choose from: the `master` and `develop` branches. Since it is currently in beta, both branches are subject to frequent merges. While this work began with the more stable `master` branch, it was eventually necessary to switch to the `develop` branch as important bug fixes were implemented during the course of the present work. The git commit which corresponds to the final version used for this work has the following `SHA` key: `b08a83f3af18f6808132be11921eff764b976998`. As with Gen-Foam, the source code obtained from the branch was compiled without additional changes.

In comparison to the many existing OpenFOAM-based solvers, GeN-Foam is highly complex. For all the physics concerned, the governing equations are explicitly defined to suit the nuclear-focused application (including thermal-hydraulics). And on top of introducing neutronic and thermo-mechanical physics, there are also multiple models available for both neutronics and thermal-hydraulics, of which only those that were employed in the simulations for this work are described in this chapter.

3.2.1   Thermal-hydraulic Model

The single-phase thermal-hydraulics model in GeN-Foam has the ability to model sub-structures using coarse-mesh porous media, as well as structure-free regions with standard Reynolds-Averaged Navier Stokes (RANS) methods. The porous media model functions as a generalized sub-channel model since it relies on correlations to capture micro-scale effects, but in three dimensions instead of one. It also permits the interfacing of coarsely meshed porous regions with structure-free regions of more fine-scale CFD (for example, to capture flow dispersion

and eddy formation in the upper and lower plena).

Porous Media Formulation

In the porous media approach, complex sub-structures in a flow region are homogenized by defining an average porosity factor, γ, between 0 and 1:

$$\gamma = \frac{V_\gamma}{V} \ . \tag{3.1}$$

The RANS method employed in GeN-Foam averages over sub-regions of the pore volume, V_γ (in addition to the usual time-averaging), to obtain the governing equations at macro-scale (Clifford, 2013). This gives rise to unknown terms which describe the interaction between the fluid and porous structure. Then the RANS equations for compressible flow are modified by γ, and appropriate source/sink terms are added for momentum and energy. These extra terms depend on correlations to close the collection of equations in the model, thus ensuring that conservation laws are upheld. Setting a porosity of 1.0 gives back the original RANS equations (Fiorina et al., 2015).

The equations for porous media are now presented with quite standard notation, $\boldsymbol{u}$ being the velocity vector field, t being time, and so forth. Starting with mass continuity:

$$\frac{\partial \gamma \rho}{\partial t} + \boldsymbol{\nabla} \cdot (\gamma \rho \boldsymbol{u}) = 0 \ . \tag{3.2}$$

As usual any local non-conservation of mass (density) must be accounted for by convective transport (the divergence implying a local source or sink). Next, the momentum equation including a sub-scale momentum-sink denoted by $\boldsymbol{F}_{ss}$:

$$\frac{\partial (\gamma \rho \boldsymbol{u})}{\partial t} + \boldsymbol{\nabla} \cdot (\gamma \rho \boldsymbol{u} \otimes \boldsymbol{u}) = \boldsymbol{\nabla} \cdot (\mu_t \boldsymbol{\nabla} \boldsymbol{u}) - \boldsymbol{\nabla}(\gamma p) + \gamma \boldsymbol{F}_g + \gamma \boldsymbol{F}_{ss} \ , \tag{3.3}$$

which accounts for any local change in momentum as a result of convective acceleration (the Bernoulli effect), divergence of shear stress (diffusion), pressure drop, or gravity. The sink term $\boldsymbol{F}_{ss}$ originates from the pressure drop associated with the sub-scale structure that motivates the porous medium representation. In a sub-channel of a solid-fuel reactor, this would represent the fuel, cladding, and structural elements. Meanwhile in the CMSR, it can represent the region of the core occupied by the moderator (through which the fluid of interest, the fuel, does not flow). The pressure drop is then obtained from a correlation as in:

$$\boldsymbol{F}_{ss} = \boldsymbol{\kappa}(\boldsymbol{u_D}) \cdot \boldsymbol{u_D} \ , \tag{3.4}$$

where $\boldsymbol{\kappa}$ is calculated along a given principal direction with a Darcy friction factor f_D (Incropera et al., 2006):

$$\kappa(u_D) = \frac{f_D \rho u_D}{2 D_h \gamma^2} \ , \tag{3.5}$$

in which u_D is the magnitude of the average Darcy velocity:

$$\boldsymbol{u_D} = \gamma \boldsymbol{u} \ , \tag{3.6}$$

and the friction factor follows the correlation:

$$f_D = A_{f_D}(\text{Re})^{B_{f_D}} \ , \tag{3.7}$$

where A_{f_D} and B_{f_D} are constants obtained via experiment, and Re is the Reynolds number obtained in the usual manner.

Finally we have the modified energy equation:

$$\frac{\partial \gamma \rho e}{\partial t} + \boldsymbol{\nabla} \cdot \big(\boldsymbol{u}\gamma(\rho e + p)\big) = \boldsymbol{\nabla} \cdot \big(\gamma k_T \boldsymbol{\nabla} \text{T}\big) + \gamma \mathbf{F}_{ss} \cdot \boldsymbol{u} + \gamma \dot{\text{Q}}_{ss} \ , \tag{3.8}$$

where e is the specific internal energy and k_T is the thermal conductivity. This describes the local energy change in terms of the convective and conductive heat flux within the fluid, work done by the sub-scale drag force, and heat source/sink, $\dot{Q}_{ss}$, which corresponds to convective heat transfer with the sub-scale structure:

$$\dot{Q}_{ss} = A_V h \left(T_{ss} - T\right) \ , \tag{3.9}$$

where the heat transfer coefficient, h, is computed from the Nusselt number correlation:

$$\text{Nu} = A_{\text{Nu}}\text{Re}^{B_{\text{Nu}}}\text{Pr}^{C_{\text{Nu}}} + D_{\text{Nu}} = \frac{hD_h}{k_T} \ , \tag{3.10}$$

with D_h being the hydraulic diameter and Pr the Prandtl number.

Turbulence Modelling

Using RANS, turbulence effects can only be captured via the inclusion of additional models and closure relations, rather than arising naturally from direct numerical simulation (DNS). Within the frame of the porous media approach, particular attention must be paid to this modelling. Generally speaking, k-ϵ models are the standard for calculating mean flow characteristics. Here k is the turbulent kinetic energy and ϵ is the rate of dissipation of that energy. Then k and ϵ can be computed from a system of governing equations which relies on further correlations (the so-called k-ϵ transport equations). In porous media, however, it is evident that turbulence is affected by the presence of a sub-scale structure. In this case, the prediction of k and ϵ becomes a major challenge. Normally these would be calculated from the transport equations, but GeN-Foam avoids directly solving for these in the porous region. It instead forces the convergence of k and ϵ to user-input values which stem from pre-determined values at the outlet of the porous region. Since turbulence was not a major focus of this work, the k-ϵ equations (Launder and Sharma, 1974, Fiorina et al., 2015) are omitted here.

3.2.2 Neutronic Model

A true representation of the neutron balance in a given volume involves a modified Boltzmann equation, known as the neutron transport equation. The solution of this partial differential

equation lies in a seven-dimensional phase-space: three dimensions in space, three in velocity[6], and one in time. The solution must be finely resolved in all these dimensions, so that an accurate numerical treatment over an industrially relevant domain is still out of reach with modern computing capabilities (Duderstadt and Hamilton, 1976) — except perhaps at exascale (Gunow, 2018). The alternative is to approximate the neutron transport equation by one of many well-established methods. The most common of these is the diffusion approximation, as is described in introductory nuclear engineering textbooks (though only a heuristic derivation is usually given[7]). It is a well understood approximation that is known to be robust, meaning that it gives quite accurate results while being consistently faithful to the physical phenomena of interest. Furthermore, when combined with homogenized group constants obtained from a higher order model (lattice calculation or Monte Carlo) and appropriate boundary conditions, the diffusion equation can deliver outstanding results (Fiorina et al., 2016).

Multigroup Diffusion

The diffusion equation reduces the phase-space of the neutron transport solution to just five dimensions by treating the angle-integrated flux, which depends on kinetic energy rather than a 3-D velocity vector. However, because the multiple reaction cross-sections vary across neutron energies, it is essential to consider multiple energy groups even when applying the diffusion approximation. This also allows local effects from material heterogeneities to be resolved; for instance, the thermal peak occurring near a reflective boundary. Within a given energy group, i, the neutron balance is well-approximated by the familiar expression:

$$\frac{1}{v_i}\frac{\partial \varphi_i}{\partial t} = \nabla \cdot D_i \nabla \varphi_i + \sum_j \nu \Sigma_{f,j} \frac{(1-\beta_{\text{eff}})}{k_{\text{eff}}} \varphi_j \chi_{p,i} - \Sigma_{r,i}\varphi_i + \sum_k \lambda_k C_k \chi_{d,i} + \sum_{j \neq i} \Sigma_{j \to i}\varphi_j \ , \ (3.11)$$

where in a given energy group, the angle-integrated flux becomes the scalar flux. The above expression considers the spatial streaming (leakage), birth of fission neutrons in the given group arising from each of the other groups, j, and from within, removal from group i due to scattering or absorption, in-scattering from other groups (mainly down-scattering), and the delayed neutron source that belongs to the group. The delayed neutron source is far from constant, since the DNP concentrations change dynamically in response to changes in temperature, burn-up, and movement of the fuel. Thus the multigroup diffusion equation must be coupled with an additional equation which gives the DNP concentration as a function of time and space.

Delayed Neutron Precursor Concentration

In the case of a liquid-fueled reactor, the typical equation for the DNP concentration must be modified with an additional advection term, $\boldsymbol{\nabla} \cdot (\boldsymbol{u_D}C_k)$, which accounts for DNPs being

---

[6]Or equivalently, two in angle and one in speed.

[7]The biggest culprit being *Lamarsh, 1983*. For a rigorous derivation from the Boltzmann equation, the reader is referred to *Duderstadt and Hamilton, 1976*.

transported throughout the core, as well as a diffusive term, $\nabla \cdot (D_m \nabla C_k)$ (Fiorina, 2013), which is important for slowly-decaying DNPs and slow-moving flows:

$$\frac{\partial C_k}{\partial t} + \boldsymbol{\nabla} \cdot (\boldsymbol{u_D} C_k) = \nabla \cdot (D_m \nabla C_k) + \frac{\beta_{\text{eff},k}}{k_{\text{eff}}} \sum_j \nu \Sigma_{f,j} \varphi_j - \lambda_k C_k \ , \qquad (3.12)$$

where we also have the fission source of DNPs and the decay sink that yields the actual delayed neutrons that contribute to Eq. 3.11.

The transport of a given molecular species in a fluid is characterized by the dimensionless Schmidt number:

$$\text{Sc} = \frac{\mu}{\rho D_m} = \frac{\text{viscous diffusivity}}{\text{molecular diffusivity}} \ , \qquad (3.13)$$

where μ and ρ take on their usual meanings, and D_m is the mass diffusion coefficient as found in Fick's law of diffusion[8]. Like the Prandtl number, the Schmidt number is fundamentally a property of the given fluid (and suspended species) rather than a property of the flow. In GeN-Foam the Schmidt number is used to compute the diffusion coefficient using Eq. 3.13.

Group Constant Implementation

For the homogenized group constants, GeN-Foam performs an interpolation between user-defined "reference" (nominal) values and "perturbed" values from a handful of pre-computed cases. For example, one may set the `nuclearData` file in a GeN-Foam simulation assuming one fuel temperature and also set the `nuclearDataFuelTemp` file for an elevated temperature. The same idea applies for the coolant temperature or density via the `nuclearDataRhoCool` and `nuclearDataTCool` files. Between all these, the user may account for both Doppler feedback and density feedback with respect to the macroscopic cross-sections, emission spectra (χ_p and χ_d), delayed neutron fractions, and DNP decay constants. GeN-Foam gives the user full liberty to obtain the group constants in any way they see fit, as they are read in plain-text ASCII format.

3.2.3   Multiphysics Coupling

GeN-Foam employs a partitioned solution (operator splitting) for coupling of thermal-hydraulics, neutronics, and thermo-mechanics. This involves "inner iterations" in which the respective physics are set to converge to an acceptable tolerance before they are fed into one another. That is, as opposed to monolithic coupling in which all equations are handled simultaneously. The latter approach is favoured in other codes such as Moltres (Lindsay et al., 2018), due to its potential gains in stability and speed. However, this comes at the cost of increased software complexity and can be a development barrier if a fully implicit solution is not offered by the underlying solver package. This is the case with the standard release of OpenFOAM, which only offers such implicit solutions in a non-mainstream distribution. Thus, some gains

---

[8]The same Fick's law which is adopted for the heuristic derivation of the neutron diffusion equation.

in computational speed and elegance are sacrificed by the GeN-Foam developers in favour of maintaining compatibility while taking advantage of the large developer community tied to the standard version of OpenFOAM.

The actual implementation of the coupling is done in a customized `PIMPLE` loop – a combination of two ubiquitous CFD algorithms, `PISO` and `SIMPLE`, normally used for transient and steady-state flow, respectively. As far as OpenFOAM is concerned, the multiphysics simulations are always "transient", i.e. the time-dependent PDEs are solved. Here one distinguishes between reactor physics terminology and CFD terminology[9]. Fields from separate sub-solver meshes are mapped based on a cell volume weighted algorithm. In transient reactor physics simulations, one of several convergence acceleration schemes can be applied as well.

For the case of liquid fuel, the particular simulation requirements are signalled by a user-defined flag: `liquidFuel`. This flag sets the fuel and coolant temperatures to be equal and assigns the heat generated by the fuel to be internal to the coolant. If this flag is set `true`, particular attention must be paid to the numerics to ensure the stability of the calculation. Thus an `upwind` interpolation scheme with `Gauss` discretization is applied to the divergence term in Eq. 3.12. This numerical scheme is only accurate to first-order, but is the most stable in OpenFOAM and is the most likely to lead to converged results (at the expense of speed and accuracy). Nonetheless, this is the approach suggested by the developers of GeN-Foam, which is sensible given the non-linearity of the DNP feedback.

3.3   Serpent2

Serpent2 is a continuous-energy neutron transport code based on the Monte Carlo method, and is developed at VTT Technical Research Centre of Finland (Leppänen et al., 2015). It is primarily used for k-eigenvalue, burn-up/depletion and source calculations, yielding results at the modern limit of accuracy. Unfortunately, achieving good *precision* comes at great cost to computational speed, due to the stochastic nature of these calculations. Thus, deterministic codes are still in great demand as they can generate adequate results in a fraction of the time. Monte Carlo codes are impractical for multiphysics transient simulation since countless time-steps are required; however, they may be combined with deterministic codes in a number of ways to improve the accuracy of such deterministic multiphysics calculations. In this work, a recent version of Serpent2 (2.1.32) was used to generate the group constants which are taken as inputs by GeN-Foam.

3.3.1   Group Constant Generation

For full-core calculations, deterministic neutronics codes rely on homogenized cross-sections (fission, capture, scattering) and other group constants as inputs. The accuracy of these

---

[9]In reactor physics, a transient simulation is when the power is not fixed as it would be in an eigenvalue calculation. That is to say, the time-dependent neutron transport problem is solved from initial conditions.

calculations is highly dependent on these inputs, therefore they must be generated for each core geometry, composition, and operating state being considered. These cross-sections and group constants can be generated by higher-order lattice physics codes (e.g. WIMS, APOLLO) or by Monte Carlo codes — this was part of the original motive for the development of Serpent in the first place (Leppänen, 2007). For cross-sections, the idea is to collapse them into the specified energy groups and spatial regions while preserving the reaction rate, $\Sigma\Phi$ (Leppänen, Pusa, and Fridman, 2016). Then for a given energy group, i, a discrete homogenized cross-section is obtained as an average of the continuous cross-section weighted over the flux:

$$\Sigma_i = \frac{\int_{E_i}^{E_{i-1}} \int_V \Sigma(\mathbf{r}, E)\Phi(\mathbf{r}, E)\ \mathrm{d}E\mathrm{d}\mathbf{r}}{\int_{E_i}^{E_{i-1}} \int_V \Phi(\mathbf{r}, E)\ \mathrm{d}E\mathrm{d}\mathbf{r}}\ . \tag{3.14}$$

Meanwhile, the diffusion coefficient in each group is obtained by an altogether different procedure, since it has no transport-equation analogue. It is thus derived directly from the B_1 or P_1 equations, depending on whether leakage correction is desired or not (Leppänen, Pusa, and Fridman, 2016). And finally, the delayed neutron parameters, β and λ, are calculated via the so-called Meulekamp method (Meulekamp and Marck, 2017).

Each of the perturbed states used for feedback calculations in GeN-Foam must be prepared independently. Once this is done in Serpent2, the conversion of the generated outputs to the GeN-Foam format is facilitated by a provided Octave script, `autoOctaveSerpentToFoamXS.m`.

3.4   Auxiliary Software

Meshing Tools

A number of meshing tools were used over the course of this work, since they each have their own subtleties. For the geometry of the 2D-MSFR in the CNRS-MSFR benchmark, OpenFOAM's `blockMesh` utility was the ideal tool for generating perfectly orthogonal meshes with exponential grading near the walls. Then, for initial test cases of flow through a porous channel, Gmsh was used to generate the mesh from a CAD model. However, problems were later encountered with refining this sort of mesh while retaining low non-orthogonality (and numerical stability, as a result). Therefore, the final meshes used in this work for the CMSR cases were generated using ANSYS SpaceClaim, which is a CAD modeller and mesh generator in one.

Post-processing Tools

As is the general custom when working with OpenFOAM, results of simulations were viewed and selectively extracted using the `paraFoam` uitility, which is based on the ParaView suite. Results were other-wise post-processed using Octave and Python (with `matplotlib` and `numpy` libraries). Octave was particularly needed as it is the prescribed tool for dealing with Serpent2 output files.

Chapter 4

Methodology & Results

4.1   MSFR Benchmark

In order to gain confidence in GeN-Foam's ability to model the relevant physical phenomena of an MSR, as well as the author's own ability to employ GeN-Foam for that purpose, a series of tests were performed against the MSFR benchmark originally developed at CNRS in Grenoble, France. A recent publication demonstrated the performance of various multiphysics codes developed by leading institutions against this benchmark (Tiberga et al., 2020). The objective in this work is then to reproduce the benchmark results with GeN-Foam while gaining an understanding of the physics involved. Going through the benchmark also serves as a smooth transition to the full-core CMSR simulations as it has most of the same fundamental features, while leaving out complications of porous media, turbulence, and DNP egress/re-entry.

The simulation codes employed in the original benchmark are based on a variety of methods, both in terms of numerical schemes and mathematical models. The TU Delft code, for example, discretizes the domain with finite elements, while the rest use finite volumes. Meanwhile, their neutronic solvers use different approximations of the transport equation including the diffusion equation (PSI), the simplified-P_3 equations (CNRS), and discrete ordinates (TUD).

It should also be noted that the PSI results correspond to a defunct version of GeN-Foam from almost two years ago. As mentioned earlier in this work, GeN-Foam is under continual development, and significant changes have been made to the source code in that time frame. In particular, the thermal-hydraulic model has been re-worked to accommodate two-phase flow simulations. Thus it is valuable to compare the present results against the old PSI result, since — until now — it remained to be seen if the changes had caused significant deviation from previous benchmark results (that is, if any serious bugs or modelling errors had been introduced into the code). And indeed, over the course of this work, a handful of bugs were discovered in collaboration with the GeN-Foam developers, including one important error in the initialization of the thermal-hydraulic model[10].

---

[10]GeN-Foam was patched in response to this, and the new version was used in the results of this work.

In the sections that follow, the benchmark is presented and the pertaining results from this work are described (and labelled *GeN-Foam* in the plots). This is done while refraining from extended discussion since the physical phenomena are largely explained in *Tiberga et al., 2020*.

4.1.1   Benchmark Description

Contrary to what its name might suggest, the CNRS-MSFR benchmark does not consist of a model of the reference MSFR (EVOL, 2013). It does, however, consist of a fast-spectrum molten salt reactor: a quasi-2D square cavity filled with fuel salt and topped with a moving lid. The salt composition is chosen such that the reactor is critical with the prescribed dimensions and vacuum boundary conditions. The fluid flow is laminar, and the Boussinesq approximation with a linear expansion coefficient is used to model buoyancy. Six energy groups and eight DNP groups are applied, with group constants computed in Serpent2 from the JEFF-3.1 cross-section library. Different physical models are tested individually and eventually coupled in steady-state and transient cases. These cases consist of pre-defined steps, which aim to break up the benchmark into digestible pieces that facilitate debugging. They are summarized here from *Tiberga et al., 2020*:

0 Single physics

    0.1 *Fluid dynamics only*: velocity field with moving lid.

    0.2 *Neutronics only*: fission source with static fuel salt.

    0.3 *Energy balance only*: temperature field with fixed flow and heat source fields from 0.1 and 0.2. In addition, a constant-temperature heatsink is introduced of the form:

$$q'''(\mathbf{r}) = \zeta \left(T_{\text{ext}} - T(\mathbf{r}) \right) \ , \tag{4.1}$$

    where ζ is a volumetric heat transfer coefficient and T_{ext} is a constant.

1 Steady-state coupling

    1.1 *Coupled neutronics and fluid dynamics*: fission source with circulating fuel salt.

    1.2 *Coupled neutronics and heat transfer*: power coupling with a fixed flow field.

    1.3 *Coupled neutronics and thermal-hydraulics with natural circulation*: buoyancy testing with a static lid and free flow field.

    1.4 *Coupled neutronics and thermal-hydraulics with forced circulation*: full coupling with a moving lid and free flow field.

2 Transient coupling

    2.1 *Small-amplitude reactivity perturbation*: starting from the steady-state solution of 1.4, a perturbation to the heat transfer coefficient is introduced in the form of a sinusoid varying with time:

$$\zeta(t) = \zeta_0 + 0.1\zeta_0 \sin(2\pi f t) \ , \tag{4.2}$$

thereby the heat removal $q'''(\mathbf{r}, t)$ also varies sinusoidally. The system response is expected to be linear for a small amplitude in $\rho(t)$ (Tiberga et al., 2020).

To actually perform the benchmark, the MSFR cavity was first modelled in Serpent2 and group constants were extracted[11]. To this end, the ENDF-VII cross-section library was used, as it was readily available on the Seaborg cluster. This in turn forced the generation of six precursor groups, as opposed to the eight groups in the original benchmark.

4.1.2 Single Physics

As per the specifications of the CNRS benchmark, the flow field, fission rate density, and temperature field were independently computed. Nearly identical results were obtained as those from the original benchmark. Fig. 4.1 shows the fission rate density in a static flow field, computed as:

$$\int_E \Sigma_f \Phi \, \mathrm{d}E = \sum_j \Sigma_{f,j} \varphi_j \ . \tag{4.3}$$

The effects of the vacuum boundary condition can be mainly observed in the non-zero fission rate at the edges. This boundary condition was implemented by applying the more general *albedo* boundary condition within GeN-Foam, which can impose a reflective condition based on the reflecting material surrounding the core. The albedo is defined as the ratio of incoming to outgoing neutron current:

$$\alpha = \left. \frac{|\mathbf{J}^-|}{|\mathbf{J}^+|} \right|_+ , \tag{4.4}$$

and the condition is applied in diffusion theory (Fiorina et al., 2016) according to:

$$\frac{D_i}{\varphi_i} \nabla \left(\varphi_i \right) = -\frac{1}{2} \left(\frac{1 - \alpha}{1 + \alpha} \right) = \Gamma \ , \tag{4.5}$$

where the constant Γ is the actual value input in the `defaultFlux` boundary condition file. In the case of a vacuum boundary, the albedo is simply zero, making $\Gamma = -\frac{1}{2}$.

Next a basic thermal-hydraulic test was performed using the static fission source of step 0.2. As a result of the moving lid above the cavity, the flow field has a steep gradient at the top of the domain. This results in an asymmetric velocity profile along the y-axis, which in turn influences the temperature profile. Although the fission source is symmetric, its heat is carried unevenly as shown in Fig. 4.2. The presence of the external constant-temperature heatsink (Eq. 4.1) is what allows the temperature field to come to a steady state.

---

[11]The group constants are actually provided in *Tiberga et al., 2020*, but they were reproduced for the purpose of verification and competence-building.
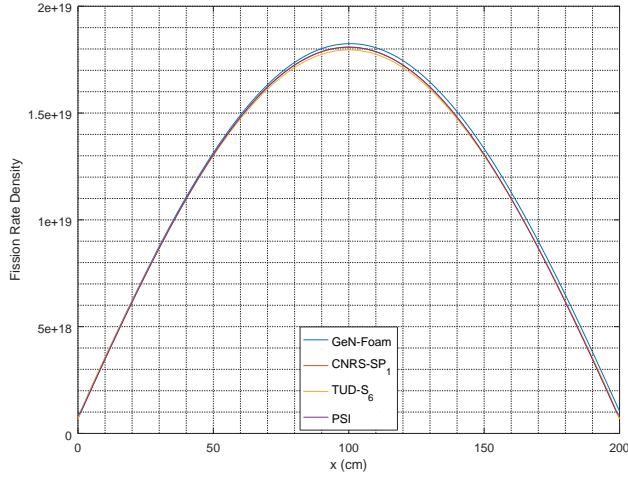
Figure 4.1: Step 0.2, fission rate density along the x-axis.
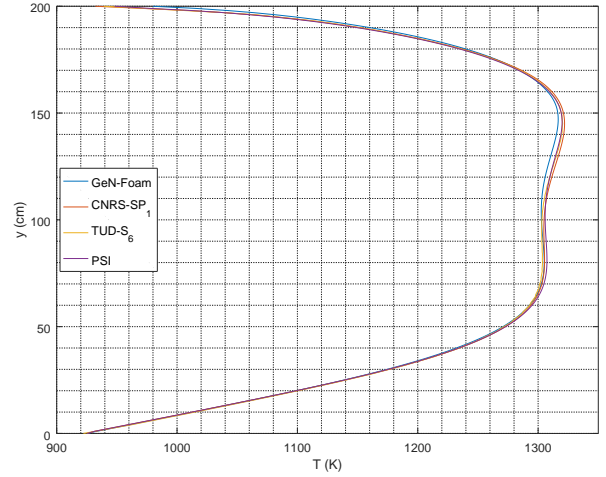
Figure 4.2: Step 0.3, temperature profile along the y-axis.

4.1.3 Steady-state Coupling

Following the single-physics test cases, a serious of increasingly complex coupling tests were performed in steady-state. First, the flow field computed in step 0.1 was applied to the neutronic solver to establish a one-way coupling which produced a tilted delayed neutron source as shown in Fig. 4.3. This source, which can be recognized from the DNP coupling term in Eq. 3.11, is integrated over all precursor and energy groups.

From here, the one-way coupling was expanded to include heat transfer from the fission source with moving fuel. The fission rate density was again computed according to Eq. 4.3, although the fission cross-section had to be adjusted according to the temperature field to account for changes in fuel density. This correction was done according to the linear thermal strain model used in the simulations:

$$\frac{\Delta V}{V} = \beta_V \Delta T \ , \tag{4.6}$$

where V is the volume, β_V is the volumetric expansion coefficient, and T is the temperature in Kelvin. From this, the adjusted cross-section is:

$$\Sigma'_f = \Sigma_f \frac{1}{1 + \beta_V \Delta T} \ . \tag{4.7}$$

Next, the difference in fission rate density between step 1.2 and step 0.2 is shown in Fig. 4.4, rising to at most 6% near the centre of the cavity. The temperature profile is also shown for step 1.2 in Fig. 4.5, but now along the x-axis.

Following this, the fully coupled simulation was tested under pure buoyancy conditions, i.e. without a moving lid. Natural circulation was observed and the delayed neutron source was obtained (Fig. 4.6). As seen, the source is symmetric along the horizontal axis, which corresponds to the symmetry of the flow field — two opposing oval-shaped currents meeting at the vertical axis (see Fig. A.2).
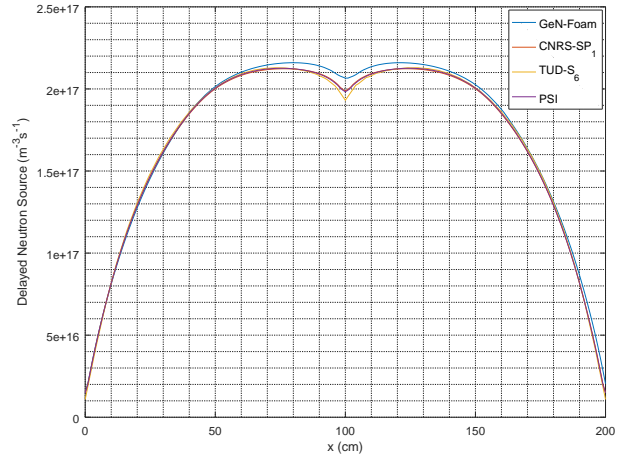
23
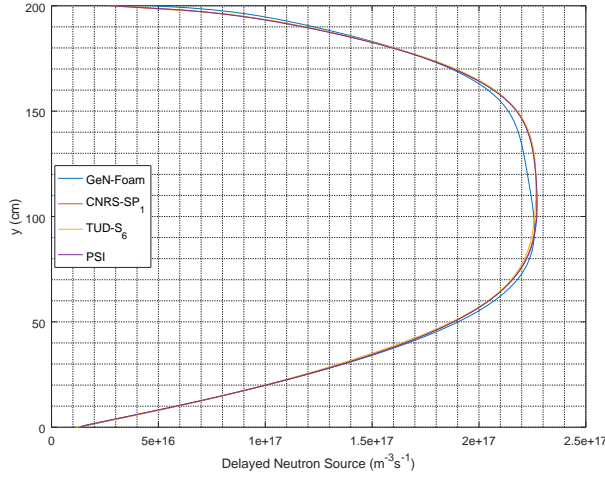
Figure 4.3: Step 1.1, delayed neutron source along the y-axis.
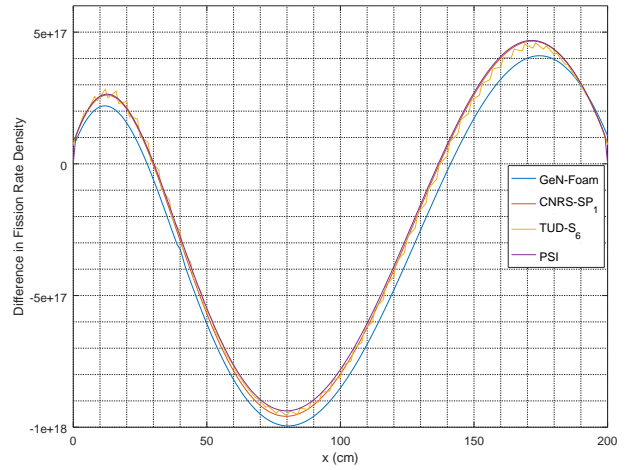
Figure 4.4: Step 1.2, change in fission rate density along the x-axis.
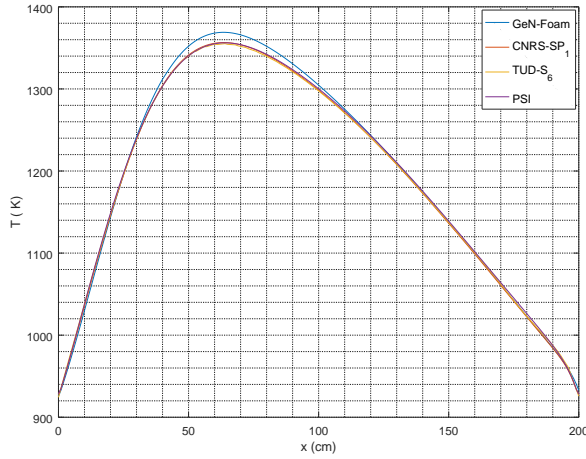
Figure 4.5: Step 1.2, temperature profile along the x-axis.

Figure 4.6: Step 1.3, delayed neutron source along the x-axis.

In these coupled cases, the results from this work are again a very close match with those of the original benchmark participants, though some deviations are now noticeable. These likely stem from the calculation of the delayed neutron source as a result of DNP transport (Fig. 4.3), and the discrepancy is attributed to the different cross-section data used for generating the homogenized group constants and the statistical uncertainty with which they were generated.

4.1.4 Transient Coupling

In preparation for transient coupling (step 2.1), a final steady-state case was prepared (step 1.4) consisting of full coupling but with a moving lid, so that both forced flow and buoyancy-driven flow were present. The transient case was then prepared according to the prescribed perturbation of the volumetric heat transfer coefficient. This perturbation was implemented

using the `fvOptions` feature of OpenFOAM which allows changes to be made in the balance equations even after a solver such as GeN-Foam is compiled, all through the designated input file. This capability was enabled in GeN-Foam (by the developers) via the inclusion of an `fvOptions` term in the source code energy equation.

The transient was then simulated starting from the steady state of step 1.4, and repeated with various perturbation frequencies. The power wave was seen to follow the same frequency as the given ζ-perturbation in all cases, with the average power, phase, and amplitude also tending to asymptotic values. The relative gain and the phase shift between the γ-wave and the power wave was calculated after this stable oscillation was reached. Following the convention of the benchmark, the gain is defined as:

$$\text{Gain} \ = \frac{\left(P_{\max} - P_{\text{avg}}\right)/P_{\text{avg}}}{\left(\zeta_{\max} - \zeta_{\text{avg}}\right)/\zeta_{\text{avg}}} \ , \tag{4.8}$$

which is just an expression for *how much* the power oscillates in proportion to a ζ-perturbation of a given strength. It is noted that the denominator of Equation 4.8 is constant, as only the frequency is varied between transient cases.

While for a linear-order perturbation, the system response (power) oscillates at the same frequency as the input, the amplitude of its variation does depend on the frequency. Therefore, this portion of the benchmark aims to capture this dependence which is fundamentally tied to the reactor kinetics (further discussed in section 4.2.1). The same is true for the phase shift, which is calculated from a sinusoidal fit matching the asymptotic behaviour ($t \to \infty$).

The gain and phase of the power-response to the ζ-oscillation depends on the DNPs, which attenuate the dynamic response of the system. For a given DNP group, its reactivity contribution relative to prompt neutrons is only slightly affected by the shift in delayed emission spectrum, but its decay constant strongly dictates how much a temperature oscillation affects the interactions of the emitted delayed neutrons as a result of varying cross-sections. This influence is a matter of relative time-scales, and having multiple precursor groups with different decay constants further complicates this effect.

In Fig. 4.7, the gain is observed to drop off from a maximum (nearly unity) as the ζ-frequency is increased. The amplitude of the power wave goes from being proportional to that of the ζ-wave to being comparably small. This is an intuitive result, as there is some delay to be expected between a given change in ζ and a change in temperature; therefore it is possible for ζ to oscillate too quickly for its effect to be felt in the temperature field. In this manner, the temperature change tends to the integral of the ζ-perturbation (which is nil) and in turn, the reactivity and power are no longer altered.

Meanwhile, Fig. 4.8 shows the phase shift starting near zero and approaching a maximum absolute value of 90° as the perturbation frequency is increased. As mentioned in *Tiberga et al., 2020*, this limit is related to the derivative of the temperature wave, with the peak temperatures and reactivities occurring at the inflections of the ζ-perturbation. Thus the power response approaches an inverted cosine wave, as one would expect from a prompt-only point kinetics model. Indeed, as the frequency is increased, eventually none of the DNPs notice the reactivity oscillation.
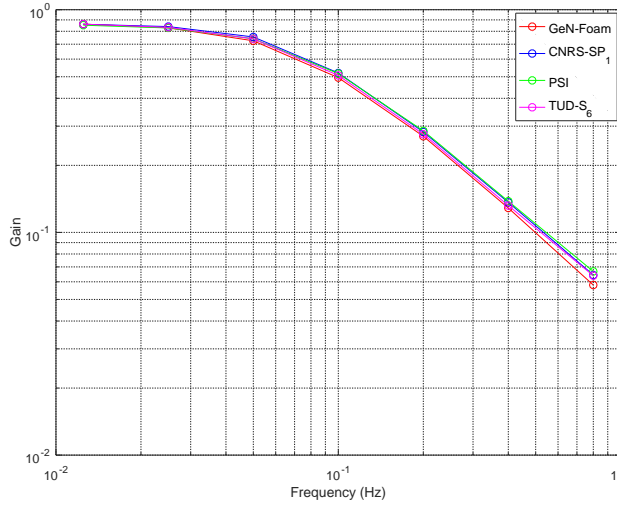
Figure 4.7: Step 2.1, ζ-power gain.
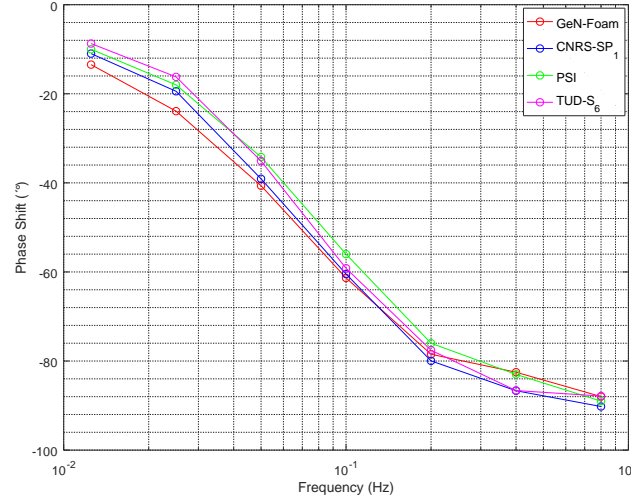
Figure 4.8: Step 2.1, ζ-power phase shift.

4.2 Neutronic Verification

In connection with the previously mentioned SEALION project, a neutronic verification scheme targeted toward MSRs has recently been developed at Chalmers University of Technology (Demazière, Mylonakis, and Vinai, 2021), based on an extension of a similar scheme for solid-fueled reactors (Demazière, Dykin, and Jareteg, 2017). Both schemes involve the manipulation of the zero-power transfer function (that is, the purely neutronic transfer function) of the simulated core in response to a reactivity perturbation. This approach is perfectly general in the sense that it applies to any reactor geometry and neutronic simulation method. This raises an opportunity to verify the capability of GeN-Foam to produce correct[12] neutronic solutions. Thus, an effort was made to implement the verification method in GeN-Foam, lending further credibility to the results obtained in this work. Given the overall time constraint, only the case of static fuel was considered and it is the one described hereafter. That being said, it lays the foundation for continued verification with the flowing-fuel scheme.

4.2.1 Theoretical Foundation

In nuclear reactors, the system response to a stationary fluctuation (one oscillating in time but not in space) is composed of two components: a point-kinetic one, describing the amplitude variation of the shape function (the fundamental mode which is assumed to maintain its shape), and a component describing the deviation from point-kinetics, that is, the evolution of the shape function attributed to the harmonic solutions of the transport equation. This is true even for strongly heterogeneous systems; only the relative contribution of the point-kinetic

---

[12]In the context of V&V, the aim of verification is not demonstrate that the solution matches reality, but that it reflects the imposed mathematical description of reality. This is particularly relevant in the absence of experimental data, which is often the case for MSRs.

component varies between systems.

As was shown by *Demazière, Dykin, and Jareteg, 2017*, it is possible to isolate the point-kinetic response as well as the associated transfer function and use it to verify the neutronic solver. For static-fueled reactors, this is done by comparing the point-kinetic transfer function to a well-defined analytical form. In the case of MSRs, there is no such form as a result of the movement of DNPs, but the method is nonetheless extensible based on the same principles.

In general, one can frame the point-kinetic response as:

$$\frac{\delta P(\omega)}{P_0} = G_0(\omega)\delta\rho(\omega) \ , \tag{4.9}$$

where $G_0(\omega)$ is the point-kinetic transfer function, $\delta\rho(\omega)$ is the input perturbation to the reactivity, and the left-hand side of the equation is the point-kinetic system response that results in terms of power (that is, the integrated response with no spatial dependence). The general idea of the verification scheme is then to extract the point-kinetic response directly from the transient simulation, along with the reactivity perturbation, and finally obtain the transfer function, $G_0(\omega)$.

In the case of a reactor with static-fuel, the point-kinetic component of the zero-power transfer function is expressed (for many DNP groups) exactly as (Bell and Glasstone, 1970):

$$G_0(\omega) = \frac{1}{i\omega\left(\Lambda_0 + \sum_k \frac{\beta_k}{i\omega+\lambda_k}\right)} \ . \tag{4.10}$$

This analytical form has a well-documented behaviour: it has a characteristic plateau in the amplitude at a value $\frac{1}{\beta}$, with a width that scales as the inverse of the mean generation time, Λ_0. Similarly, the bode plot of the phase shifts to the right as Λ_0 decreases. And as with the transient response in section 4.1.4, the DNPs are responsible for the overall shape.

Next, if Λ_0 can be estimated from the simulation, then the analytical transfer function is constrained with only the angular frequency, ω, as a free variable. Indeed, Λ_0 can be extracted (in two groups) by:

$$\Lambda_0 = \frac{\int_V \left[\frac{1}{v_1(\mathbf{r})}\phi_{1,0}^\dagger(\mathbf{r})\varphi_{1,0}(\mathbf{r}) + \frac{1}{v_2(\mathbf{r})}\phi_{2,0}^\dagger(\mathbf{r})\varphi_{2,0}(\mathbf{r})\right]\mathrm{d}\mathbf{r}}{\int_V \left[\nu\Sigma_{f,1}(\mathbf{r})\phi_{1,0}^\dagger(\mathbf{r})\varphi_{1,0}(\mathbf{r}) + \nu\Sigma_{f,2}(\mathbf{r})\phi_{2,0}^\dagger(\mathbf{r})\varphi_{2,0}(\mathbf{r})\right]\mathrm{d}\mathbf{r}} \ , \tag{4.11}$$

where v_j is the neutron speed corresponding to the j^{th} energy group and $\phi_{j,0}^\dagger$ is the adjoint flux at transient onset. This expression gives the ratio of two quantities: the neutron population weighted by the adjoint flux, $\phi_j^\dagger$, in each energy group (recall that $\varphi_j = n_j v_j$), and the fission neutron production rate in each group, also weighted by the adjoint flux. Since both are computed over the entire reactor core, this is in fact the time it would take to replace the steady-state neutron population with another of equal size, i.e. to go from one "generation" to the next, all while neglecting delayed neutrons. A key idea here is that the adjoint flux is directly proportional to the *neutron importance function* (Bell and Glasstone, 1970), hence

its usefulness as a weighting function: it allows for the space-dependent neutron reactivity worth to be accounted for. And, since it appears in both the numerator and denominator, the proportionality factor has no impact on the result.

Then, the point-kinetic system response in the simulation can be obtained from:

$$\frac{\delta P(\omega)}{P_0} = \frac{\int \left[\frac{1}{v_1(\mathbf{r})} \phi_{1,0}^\dagger(\mathbf{r}) \delta\varphi_1(\mathbf{r},\omega) + \frac{1}{v_2(\mathbf{r})} \phi_{2,0}^\dagger(\mathbf{r}) \delta\varphi_2(\mathbf{r},\omega) \right] \mathrm{d}\mathbf{r}}{\int \left[\frac{1}{v_1(\mathbf{r})} \phi_{1,0}^\dagger(\mathbf{r}) \varphi_{1,0}(\mathbf{r}) + \frac{1}{v_2(\mathbf{r})} \phi_{2,0}^\dagger(\mathbf{r}) \varphi_{2,0}(\mathbf{r}) \right] \mathrm{d}\mathbf{r}} \ , \tag{4.12}$$

that is the relative neutron noise, $\frac{1}{v_j(\mathbf{r})} \delta\varphi_j(\mathbf{r},\omega)$, weighted by the adjoint flux in each energy group and normalized to the total neutron population also weighted by importance. As with Eq. 4.11, it stands to reason that flux perturbations in regions inhabited by the most reproductive neutrons are weighted higher in determining the system response.

Finally, the transfer function corresponding to the simulation can be obtained by dividing the above equation by $\rho(\omega)$. The phase and amplitude of the transfer function at various frequencies can be compared to the analytical form, and the error calculated to verify the code. In practice, since most transient reactor physics codes operate in time-domain, one can either convert these operations between domains via numerical Fourier transforms, or perform the analysis directly in time-domain by concentrating on the asymptotic behaviour of the system. The latter option is valid because the system described by point-kinetics is known to be linear[13], hence there is a single dominating frequency that governs the system response, with the remainder rapidly decaying over time. In time domain, then:

$$\frac{\delta P(t)}{P_0} = \frac{\int_V \left[\frac{1}{v_1(\mathbf{r})} \phi_{1,0}^\dagger(\mathbf{r}) \delta\varphi_1(\mathbf{r},t) + \frac{1}{v_2(\mathbf{r})} \phi_{2,0}^\dagger(\mathbf{r}) \delta\varphi_2(\mathbf{r},t) \right] \mathrm{d}\mathbf{r}}{\int_V \left[\frac{1}{v_1(\mathbf{r})} \phi_{1,0}^\dagger(\mathbf{r}) \varphi_{1,0}(\mathbf{r}) + \frac{1}{v_2(\mathbf{r})} \phi_{2,0}^\dagger(\mathbf{r}) \varphi_{2,0}(\mathbf{r}) \right] \mathrm{d}\mathbf{r}} \ . \tag{4.13}$$

Meanwhile, the reactivity perturbation over time is obtained from first-order perturbation theory (Bell and Glasstone, 1970) as:

$$\delta\rho(t) = \frac{1}{F(t)} \int_V \left\{ \left[\delta\nu\Sigma_{f,1}(\mathbf{r},t) - \delta\Sigma_{a,1}(\mathbf{r},t) - \delta\Sigma_{\mathrm{rem}}(\mathbf{r},t) \right] \phi_{1,0}^\dagger(\mathbf{r}) \psi_1(\mathbf{r},t) \right.$$
$$\left. + \delta\nu\Sigma_{f,2}(\mathbf{r},t) \phi_{1,0}^\dagger(\mathbf{r}) \psi_2(\mathbf{r},t) + \delta\Sigma_{\mathrm{rem}}(\mathbf{r},t) \phi_{2,0}^\dagger(\mathbf{r}) \psi_1(\mathbf{r},t) - \delta\Sigma_{a,2}(\mathbf{r},t) \phi_{2,0}^\dagger(\mathbf{r}) \psi_2(\mathbf{r},t) \right\} \mathrm{d}\mathbf{r} \ , \tag{4.14}$$

where Σ_{rem} is the scattering removal (down-scattering) cross-section and where:

$$F(t) = \int_V \left[\nu\Sigma_{f,1}(\mathbf{r},t) \phi_{1,0}^\dagger(\mathbf{r}) \psi_1(\mathbf{r},t) + \nu\Sigma_{f,2}(\mathbf{r},t) \phi_{1,0}^\dagger(\mathbf{r}) \psi_2(\mathbf{r},t) \right] \mathrm{d}\mathbf{r} \ . \tag{4.15}$$

The shape function, ψ_j, corresponds to the principal assumption of point kinetics:

$$\varphi_j(\mathbf{r},t) = P(t) \psi_j(\mathbf{r},t) \cong P(t) \psi_j(\mathbf{r}) \ . \tag{4.16}$$

And the asymptotic behaviour of $\frac{\delta P(t)}{P_0}$ can be compared with $\delta\rho(t)$ based on the phase and amplitude. It is this phase and amplitude which must match that of Eq. 4.10 for different input perturbation frequencies, in order for the simulation code to be verified.

---

[13]There would be no guarantee of this if one was dealing with the entire system response, rather than just the point-kinetic component.

4.2.2 Implementation

The verification scheme as just described was implemented in GeN-Foam using the two-dimensional reactor model from the CNRS-MSFR Benchmark. A number of simplifications were made to the model to facilitate the verification process, without sacrificing the integrity of the verification. To begin with, since GeN-Foam is capable of solving the multigroup diffusion equation for an arbitrary number of energy groups, the number of groups was reduced to two, so as to match the derivation of the reactivity perturbation as presented in Eq. 4.14 (thereby avoiding the need to derive and implement it in six groups). As GeN-Foam generates the energy groups on-the-fly using the same code when solving the diffusion problem, a verification with any number of groups (greater than one) is expected to be equally significant.

One can further simplify the task by choosing the solution of the steady-state and transient simulations to be spatially symmetric, thereby facilitating the process of performing the necessary numerical integrations and saving considerable disk space. Thus, the same type of transient simulations were performed as in step 2.1 of the benchmark (a sinusoidal variation of the volumetric heat transfer coefficient, ζ, was imposed), except that fluid flow was switched off. In this manner, any effects of a moving lid or buoyancy on the temperature distribution were negated. At the same time, this satisfied the stipulation that the verification scheme as laid out in section 4.2.1 is only valid for static fuel.

The forward flux, adjoint flux, and fuel temperatures were all extracted for each case. The forward flux and temperatures had to be extracted at every time-step, while the adjoint flux only had to be extracted at steady-state. Only the values along the x-axis of the core geometry were extracted, having reduced the problem to one dimension via the spatial symmetry. Then, the point-kinetic system response was computed as in Eq. 4.13 via numerical integration. The reactivity perturbation was also computed as in Eq. 4.14, with the density-adjusted cross-sections calculated from the temperature field and volumetric expansion coefficient (Eq. 4.7). Finally, the phase and amplitude between the two were obtained via sinusoidal fitting in time-domain, as explained.

In Fig. 4.9 and Fig. 4.10, the perturbation and point-kinetic response are plotted for the case of a low-frequency and mid-frequency perturbation, respectively. Both responses are seen to be monochromatic, as expected for a linear system. Meanwhile, the frequency-dependence of the relative phase and amplitude is showcased, with the low-frequency case having a noticeably higher phase at-a-glance. The low-frequency case exhibits hardly any non-asymptotic behaviour, while the mid-frequency case requires a sinusoidal fit at $t > 10$ s.

Based on this line of analysis, the phase and amplitude of the simulated point-kinetic transfer function are compared with its analytical form (Eq. 4.10) in Fig. 4.11 and Fig. 4.12, respectively. The analytical phase and amplitude are computed according to:

$$\arg\left(G_0(\omega)\right) = \arctan\left(\frac{\Im\left(G_0(\omega)\right)}{\Re\left(G_0(\omega)\right)}\right) \ , \ |G_0(\omega)| = \sqrt{\Im\left(G_0(\omega)\right)^2 + \Re\left(G_0(\omega)\right)^2} \ . \qquad (4.17)$$

In addition, the alternate analytical forms which correspond to singular DNP group scenarios are shown, describing what the transfer function would be if all the DNPs shared the decay
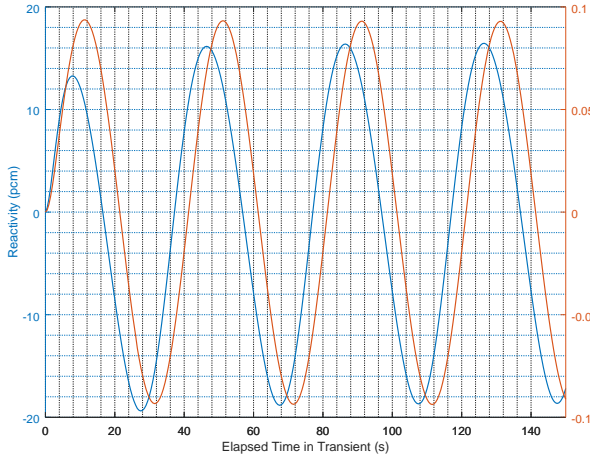
Figure 4.9: Point-kinetic system response vs. reactivity perturbation, $f = 0.025$ Hz.
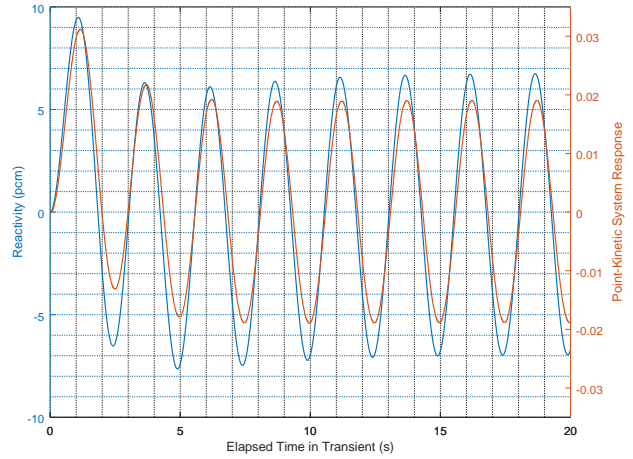
Figure 4.10: Point-kinetic system response vs. reactivity perturbation, $f = 0.4$ Hz.

constant of the slowest-decaying or fastest-decaying group. Therefore, these correspond to "slow" and "fast" bounds for the transfer function. Eq. 4.11 was used to obtain the mean generation time used in these functions. As expected, all of them exhibit a plateau with a height of $\frac{1}{\beta}$, or about 154 pcm. The hollow points on these figures pertain to the cases which were simulated in this work. To obtain these results, the strength of the ζ perturbation had to be increased so as to produce a discernible reactivity perturbation even at high frequencies. This turned out to be a limitation in the extent to which the verification scheme could be carried out, since at frequencies past 100 Hz it was not possible to generate a significant reactivity perturbation without having a non-physical (briefly negative) ζ.

It is observed that the phase of the simulated transfer function is well in line with the true analytical form. This is a key result as it confirms that the delicate dynamics imposed by the presence of DNPs are properly simulated in GeN-Foam, at least with respect to the delay between a reactivity perturbation and the system response that follows. One should note that Fig. 4.11 shows the phase difference with the ρ-wave, as opposed to the ζ-wave as in Fig. 4.8; therefore, the correct estimation of $\delta\rho$ is critical. The phase-behaviour can be roughly divided into three regions: At low frequencies, the full effect of the DNPs is manifested, as any change in power is dependent on their collective delay. Then the magnitude of the phase decreases as one approaches middle frequencies ($f \approx 10$ Hz — corresponding to the period of the shortest-lived DNP group) wherein the DNPs are outpaced (as explained in 4.1.4) and the prompt neutrons are the ones to respond, which they can do more quickly. Finally, as the frequency is further increased up to 100 Hz, the prompt neutrons themselves begin to lag behind the perturbation, depending on their energy. A harder spectrum results in a shorter mean generation time, meaning that the frequency required to induce prompt neutron lag is greater (hence the width of the phase curve being determined by Λ_0).

Meanwhile, the amplitude result is not as ideal. While the shape of the desired amplitude curve is maintained, there is a consistent offset corresponding to a relative error of around 60%. Since the analytical form has been demonstrated to be correct (at least in the plateau region)
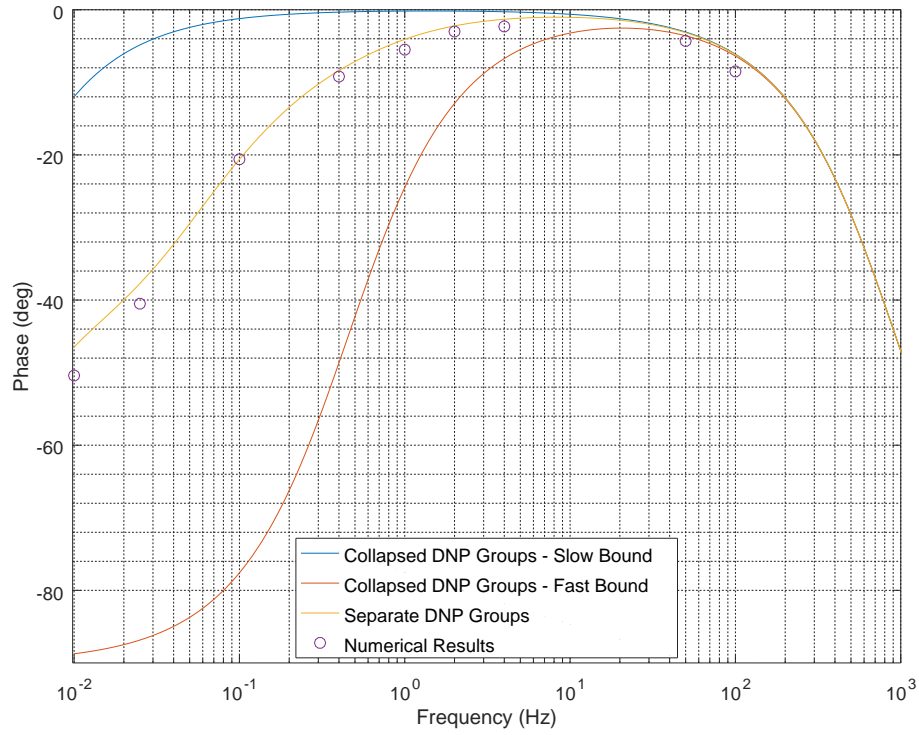
Figure 4.11: Phase of the point-kinetic transfer function.
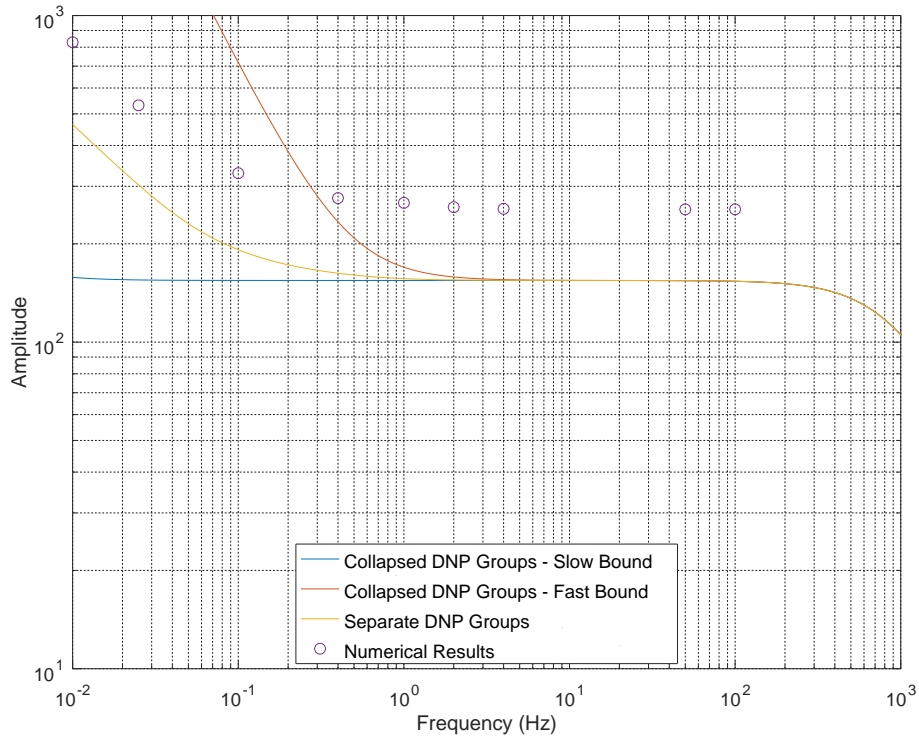
Figure 4.12: Amplitude of the point-kinetic transfer function.

and the point-kinetic response matches the observed relative power evolution (as expected for a bulk reactivity perturbation), there are two possible explanations: there may be some error with the calculation of the reactivity perturbation (most likely), or with the actual reactivity response to the induced temperature change within GeN-Foam (doubtful, given the correct MSFR benchmark results). In the transients modelled, the value of β was relatively constant, largely uninfluenced by the varying temperature — therefore this is discounted as the potential source of error. One possible way to correct the error would be to directly input the reactivity perturbation rather than producing it from a ζ perturbation; however, there is presently no way of doing this in GeN-Foam.

Although the results of the verification attempt are imperfect, they are still promising as everything points to a mistake in the data analysis as opposed to GeN-Foam itself. It would be surprising if the neutronics model were somehow able to model the phase response correctly but not the amplitude response. Bearing this in mind, and together with the success in passing the benchmark, a great deal of confidence is gathered in the capabilities of GeN-Foam.

4.3 Application to the CMSR

Having justified the utility and accuracy of GeN-Foam to a great extent, the attention now turns to a model of the CMSR active core. This model is the most complex considered in this work, and is the culmination of the GeN-Foam know-how acquired throughout. The aim is to obtain some of the basic quantitative results such as the desired outlet temperature at steady-state, to investigate the accuracy and sensitivity of the DNP transport, and to demonstrate some of the intrinsic safety features of the CMSR in transient scenarios.

4.3.1 Model Setup

For the purpose of this work, a simplified model of the CMSR was used from an early iteration of conceptual design. This serves to accelerate the Monte Carlo calculations necessary for group constant generation (compared to the true geometry) while protecting company IP. The basic parameters of this model are shown in Table 4.1, while approximate thermophysical parameters of the FUNaK salt[14] are shown in Table 4.2.

A Serpent2 input deck was thus prepared having the geometry shown in Fig. 4.13. In this illustration, the fuel salt is shown in red while the moderator is blue. There are also six empty channels in white, where control or shutdown rods might otherwise be placed. Neutron poison was added until criticality was reached. Then, the entire region inside the reactor vessel was grouped under one Serpent2 `universe`, from which the homogenized group constants were generated with 10^6 neutron histories, with a statistical error of at most 5%. While potentially significant, this must be weighed against the much larger uncertainty posed by the thermophysical data of FUNaK (Pedersen, 2020). Separate group constants were also

---

[14]NB: the boiling point of FUNaK is approximated by the boiling point of FLiNaK.
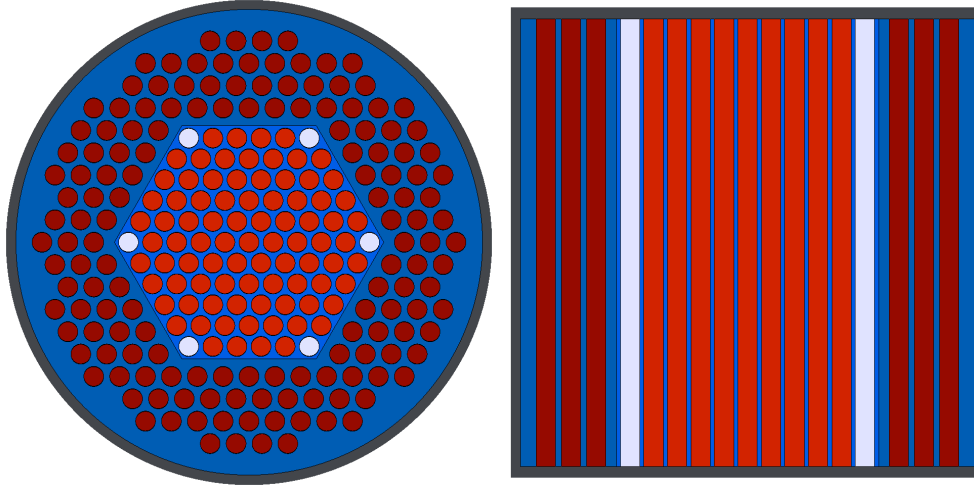
Figure 4.13: Top and side cut views of the simplified CMSR core as modelled in Serpent2.

extracted for the perturbed cases of increased fuel temperature and reduced coolant density[15], changing only the temperature in the former and only the density in the latter.

In GeN-Foam, the simplified CMSR was modelled as a single region (mesh zone) of porous media, in which the thermal-hydraulic conditions and Serpent2-generated group constants were applied. From a thermal-hydraulic point of view, the inspiration for this comes from the use of porous media for modelling complex structures such as heat exchangers, which has been known to yield acceptable results (Milani Shirvan et al., 2016). For this porous media, the porosity was set according to the total internal cross-section of the fuel tubes (the wetted area) in proportion to the total area inside the vessel. In this manner, the geometry of the core was completely homogenized not only for the flow but also the neutronics — a significant approximation yielding mixed results. This approach allowed the author to rapidly prepare the corresponding computational mesh and execute a variety of relevant simulations for the CMSR while exploring the limitations of the porous media model.

Standard boundary conditions were applied much like the MSFR benchmark cases: vacuum

---

[15]While there is no physical distinction between the fuel and coolant, this allows individual contributions of Doppler and density feedback to be isolated.

Table 4.1: Basic parameters of the simplified CMSR core.

Parameter	Value	Unit
Power	250	MW$_{\text{th}}$
Height	2.0	m
Radius	1.2	m
Flow Rate	2900	kg/s
Porosity	0.41	
Inlet Temperature	823	K

Table 4.2: Thermophysical properties of FUNaK eutectic salt.

Property	Value	Unit
Melting Point	763	K
Boiling Point	$\sim$1700	K
Expansion Coefficient	2.725×10^{-4}	K^{-1}
Dynamic Viscosity	1.324×10^{-2}	Pa·s
Heat Capacity	920	J/kg·K
Thermal Conductivity	0.86	W/m·K

conditions for the flux and the usual fluid-dynamic condition of no-slip at the walls. In addition, however, since this model has an inlet and outlet, this introduced some new conditions. The flow rate, for example, was set as an inlet velocity condition (0.466 m/s) corresponding to an assumed density of 4250 kg/m^3, the design porosity, and mass flow. This velocity corresponds to the actual velocity of the fuel through a given tube (whereas older versions of GeN-Foam would require the Darcy velocity to be set instead). The outlet velocity, meanwhile, was set to `zeroGradient`. The outlet pressure was set to a reference value of 10^5 Pa and the inlet pressure was allowed to be back-propagated according to frictional and gravitational pressure drops using the `fixedFluxExtrapolatedPressure` condition. Turbulent parameters of k and ϵ were left at their default values, including a turbulent intensity of 0.05 and a mixing length of 0.001 m.

4.3.2 Steady-state Solution

Mesh Independence

The approach to steady-state after the initial flow field is already established can be seen in Fig. 4.14 with respect to k_{eff}, the average fuel temperature, and the maximum fuel temperature. In these nominal steady-state simulations, the power is held constant at 250 MW, thus preventing any effect of reactivity on the thermal-hydraulic outcome. On the other hand, there is a one-way coupling between the fuel temperature and the multiplication factor, as seen by the shape of the k_{eff} curve mirroring the temperature. These values are compared between computational meshes of various sizes: the same mesh was refined in the axial and radial directions, but not in the azimuth (to take advantage of symmetry). It is observed that all three variables seem to converge to a mesh-independent value as the mesh grows more refined (though convergence is not entirely reached). This is also demonstrated in Fig. 4.16, where the exponential convergence is clear. Such convergence is a good indicator of the numerical accuracy of the simulations, along with the decaying residuals shown in Fig. A.3.

Fuel Temperature

By the time the solution is approximately mesh-independent, the average temperature throughout the core is seen to level off at approximately 867 K. Meanwhile, maximum temperature rises just past 1000 K at the core outlet. Fig. 4.18 shows the profile of the temperature along the central axis, which takes on the shape of the axial integral of the power distribution. Just focusing on the outlet temperature, Fig. 4.19 shows the temperature distribution along the whole diameter, along with a dashed line representing the average temperature computed over the full outlet area (seen in Fig. A.6). The average temperature at the outlet reaches 941 K, which exceeds the design specification of 923 K (Klinkby, 2019) stemming from the steady-state expression:
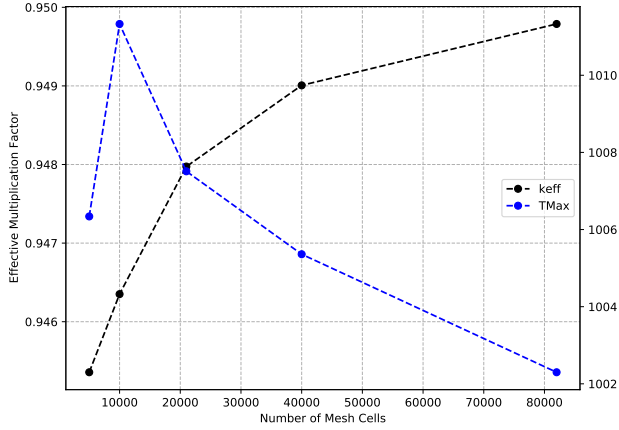
$$T_{\text{outlet,avg}} = T_{\text{inlet,avg}} + \frac{\dot{Q}}{\dot{m}c_p} = 916.7 \text{ K} . \tag{4.18}$$
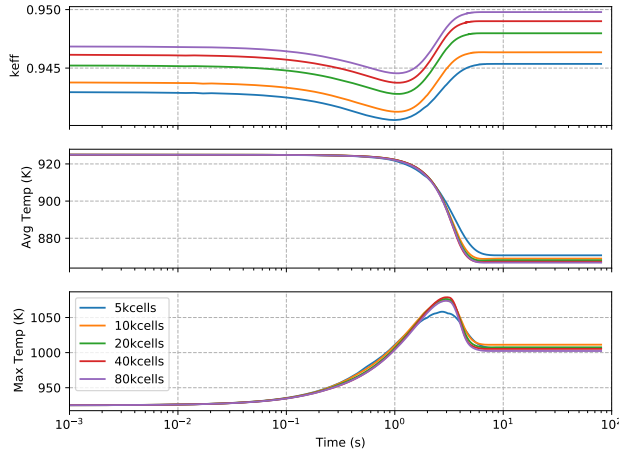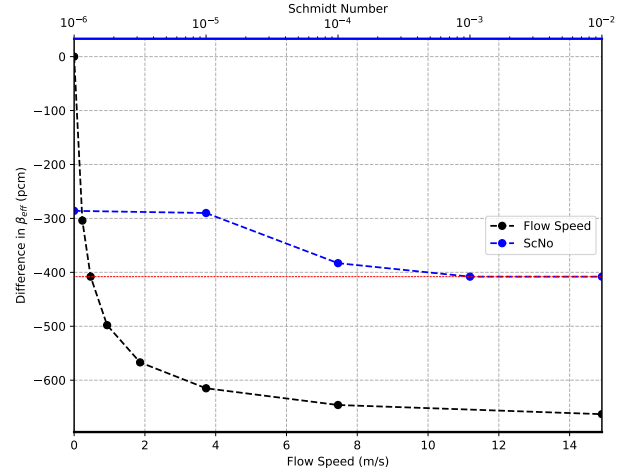
Figure 4.14: Approach to steady-state with varying mesh size.
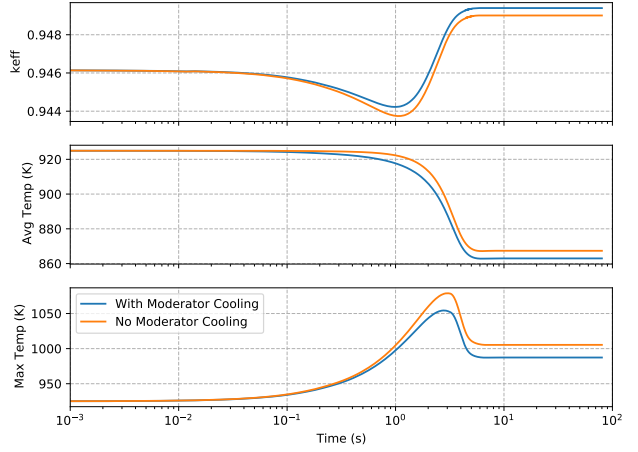
Figure 4.15: Approach to steady-state with and without moderator cooling.


Figure 4.16: Mesh independence study.

**Figure 4.17:** DNP transport-induced difference in β_{eff} from the case of static fuel.

However, this expression does not account for the varying density of the fluid, which affects the mass flow when the inlet boundary condition is expressed in terms of velocity. In this case, the mass flow is decreased at steady-state, thus raising the outlet temperature.

As a separate note, neither the simulations in Fig. 4.14 nor this formula take into account any heat transfer to the moderator. Neither do they account for radiative heat transfer, and all fission power is assumed to be preserved in the fuel salt — when in reality some of it would be deposited in the moderator, some in the structural materials, and some outside the core. Between all of these simplifications, the calculated average outlet temperature is expected to over-estimate the true value.

The high peak temperature is a potential source of concern, as corrosion can be significantly accelerated by temperature (as is generally true for chemical reactions). That being said, Fig. 4.16 shows that the peak temperature is quite sensitive to the size of the mesh, and so further refinement is needed to better estimate it. In addition, the peak temperature depends strongly
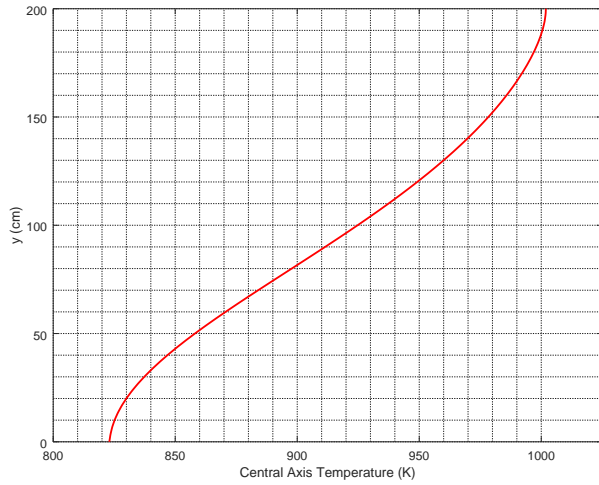
Figure 4.18: Temperature profile along the vertical axis of the CMSR.
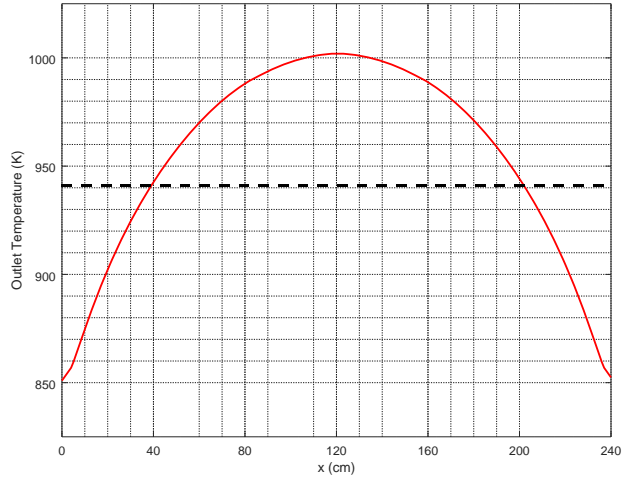
Figure 4.19: Temperature profile at the outlet of the CMSR.

on the turbulent parameters of the simulation, which have not been tuned to any significant extent. Although, a test case was run with purely laminar flow and a steeper radial profile was obtained leading to a higher peak temperature with almost the same average temperature. In any case, a porous media model such as this one is certainly not the best tool for estimating the peak temperature, particularly when the neutronic meshing does not resolve the individual fuel tubes.

Moderator Cooling Impact

Heat exchange to the moderator was added after the initial steady-state cases were tested. To achieve this, the sub-scale heat structure feature of GeN-Foam was used — another benefit of the porous media approach[16]. The correlation in Eq. 3.10 was applied to determine the heat transfer coefficient between the fuel salt and the structure, which has a surface area equal to the combined inner surface of all the fuel tubes. The coefficients of this correlation were not tuned, thus the results of this addition are to be judged qualitatively. The temperature of the structure was set constant at 800 K, so as to represent a circulating moderator with significant cooling. Such a condition might be imposed in the CMSR design if the corrosion from NaOH becomes a concern at the temperature of the fuel.

The impact of this moderator cooling, shown in Fig. 4.15, is fairly significant. The difference in average fuel temperature between the case without moderator cooling (corresponding also to Fig. 4.14) and the case with cooling is about 4 K, while the difference in peak temperature reaches 18 K.

---

[16]A more detailed analysis would require conjugate heat transfer at the interface and additional thermal-hydraulic modelling of the moderator.

Flux and Multiplication Factor

From the nominal Serpent2 calculation, the k-eigenvalue of the CMSR was calculated to be 1.00019 ± 0.00012 with the collision estimator. However, since all geometric detail was then smoothed out in the process of group constant generation, the k-eigenvalue computed in GeN-Foam was approximately 0.95 — a difference of over 5000 pcm. This discrepancy is not totally surprising, as it is known that reactivity can be significantly enhanced by the heterogeneity of a reactor core (Lamarsh, 1983). This is confirmed by a simple modification of the GeN-Foam case: if the `liquidFuel` option is turned off, k_{eff} rises to almost unity, since GeN-Foam assumes that the fuel and coolant are two separate entities unless stated otherwise[17]. There is also the added effect of temperature feedback from thermal expansion and Doppler broadening, which is not accounted for in the baseline Serpent2 simulation;however, the evolution of k_{eff} in Fig. 4.14 shows that this has a comparably minor effect (the initial temperature of 925 K in these simulations corresponds to the temperature in the baseline Serpent2 deck).

The flux was computed using a vacuum boundary condition, as in the case of the MSFR benchmark. The cases shown so far made use of two energy groups, which are plotted along the core radius and vertical axis in Fig. 4.20 and Fig. 4.21. These exhibit the shape of the analytical solution of the diffusion equation for a finite homogeneous cylinder: A Bessel function of the first kind, J_0, along the radius and a cosine along the axis. However, there is a notable shift in the axial peak due to the uneven temperature distribution induced by the flow. As with the MSFR benchmark, the flux at the boundaries is increased by the inclusion of an extrapolated distance, and the lack of radial mesh grading near the boundary results in a non-physical discontinuity in the derivative between the outermost cell and the next one in.

While there are no material discontinuities to be resolved in this CMSR model with different energy groups, it is nonetheless expected that having more groups will result in a more accurate estimation of the k-eigenvalue, due to the presence of resonances at various energies. Thus, a case was tested with eight energy groups and the results are compared in Table 4.3. The impact of the switch to eight groups is clearly significant ($\sim$150 pcm), but not as much as the impact of implementing power and temperature feedback ($\sim$250 pcm). Note that the fully coupled case has a lower k_{eff} because the temperature in the core is otherwise assumed to be 925 K everywhere, so the average temperature is lower in this case (Fig. 4.14). In addition, the neutron flux distribution in the eight-group case is shown in Fig. A.4 and Fig. A.5.

---

[17]The discrepancy arises regardless of whether temperature feedback is considered, and is thus independent of internal vs. external heating.
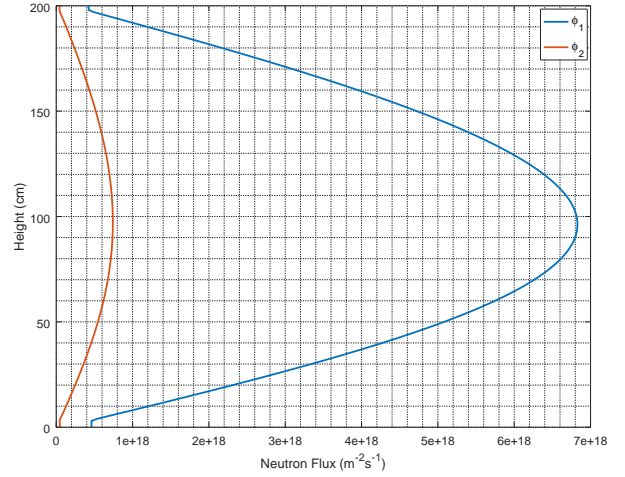
Table 4.3: Effective multiplication factor with/without power coupling and different energy groups.

Case	k_{eff}
Full coupling, 2-group	0.949788
Full coupling, 8-group	0.948072
Partial coupling, 2-group	0.946874
Partial coupling, 8-group	0.945468
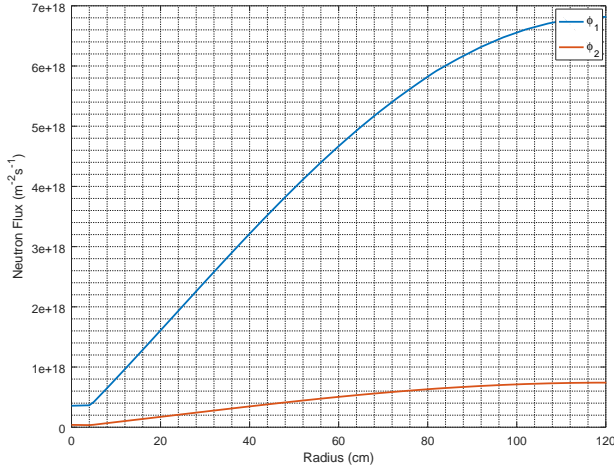
Figure 4.20: Radial flux profile
in the CMSR core.

Figure 4.21: Axial flux profile
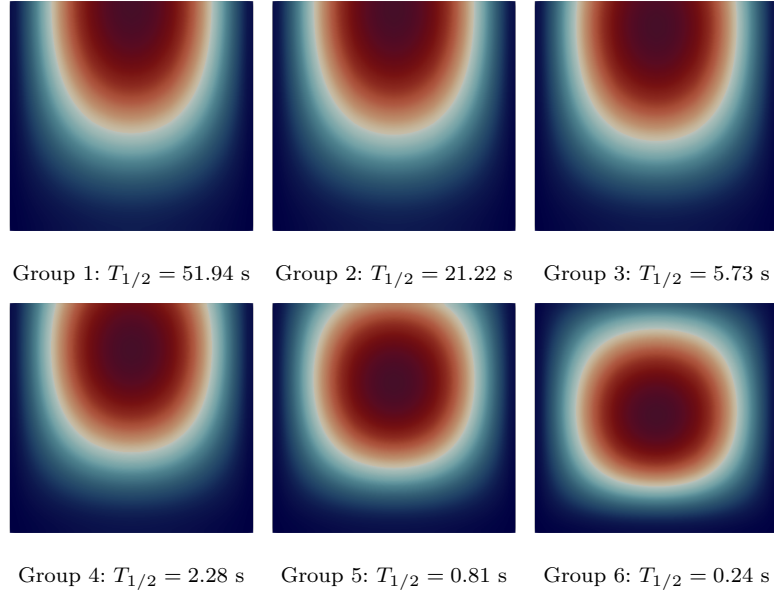in the CMSR core.

DNP Transport

Delayed neutron precursors have a disproportionate effect on the reactivity and dynamics of nuclear reactors, owing to the fact that the fission chain reaction at criticality is limited by the production of said delayed neutrons. In the case of liquid-fueled MSRs, DNPs are transported through the core and primary loop after they are born from fission, both due to the flow of the fuel salt as well as Brownian motion (recall the diffusion term in Eq. 3.12). This has important consequences for both normal operation and accident scenarios. During normal operation, the DNPs continually exit the core and so their contribution to the effective delayed neutron fraction, β_{eff}, is greatly reduced[18]. In this context, it is important to distinguish between the usual definition of β_{eff}, which only considers the difference in emission spectra between prompt and delayed neutrons, and the more general definition (used in this work) which also includes changes in importance due to incongruity with the flux distribution.

In simulations of the CMSR, the general behaviour of the DNPs is shown in Fig. 4.28. These images capture the essence of DNP transport in the regime of forced convection: slowly decaying DNPs are readily carried downstream where their contribution to β_{eff} is minimized, while the most rapidly decaying DNPs are hardly affected by the flow stream, since they decay shortly after being born. A more exact representation of this effect is shown in Fig. 4.29, where the source of delayed neutrons from each DNP group ($\lambda_k C_k$) is compared between the actual CMSR case and a case with static fuel (temperature held constant between the two). The integral delayed neutron source inside the core is greatly reduced and in turn so is the reactivity effect, β_{eff}.

For the simulations of the CMSR presented so far, the Schmidt number was set to an arbitrarily high value to impose a complete dominance of convective transport over diffusive transport.

---

[18]One considers delayed neutrons born outside the active core to be inconsequential, as only a small fraction of those will arrive back in the active core.

Group 1: $T_{1/2} = 51.94$ s   Group 2: $T_{1/2} = 21.22$ s   Group 3: $T_{1/2} = 5.73$ s

Group 4: $T_{1/2} = 2.28$ s   Group 5: $T_{1/2} = 0.81$ s   Group 6: $T_{1/2} = 0.24$ s

Figure 4.28: Relative concentration of each DNP group at steady-state.

There is a present lack of information about the diffusion coefficient of DNPs in molten salts, particularly in the FUNaK salt used for the CMSR. Therefore, the Schmidt number can only be approximated within a few orders of magnitude (if one assumes the diffusion coefficient is somewhat in line with similar species in similar fluids, or at best derives a theoretical value from kinetic theory). Nevertheless, it is expected that diffusion will play a negligible role in the case of forced convection with high mass flow, as is the case in normal operation. It is valuable, then, to confirm this hypothesis — seeing to what degree the Schmidt number has an impact on β_{eff}.

A further study was thus undertaken to vary the flow speed of the fuel through the core, and separately vary the Schmidt number while holding the inlet speed constant at the nominal condition of 0.466 m/s. Meanwhile, energy-coupling was turned off so as to maintain the same temperature field and thereby isolate the effect of DNP transport on β_{eff}. The results of this study are shown in Fig. 4.17, with the difference in reactivity taken as the difference in β_{eff}. The result of varying the flow speed is as expected; β_{eff} decreases as an exponential decay. This decrease can be seen to converge to the total value of $\beta_{\text{eff,static}}$ (approximately 680 pcm as computed in Serpent2). On the other hand, the effect of varying the Schmidt number is more nuanced, as there are two asymptotic values of $\Delta\beta_{\text{eff}}$ rather than just one. When the Schmidt number is increased, this naturally leads to the dominance of convective transport, as was originally sought in previous CMSR simulations. On the other hand, it was initially expected that continually decreasing the Schmidt number would lead to a reduction of $\Delta\beta_{\text{eff}}$ to nil. However, the opposing dominance of diffusive transport leads to a homogenous distribution of DNPs inside the core, which alters their contribution to β_{eff} based solely on the neutron importance within the core. The dashed red line in Fig. 4.17 identifies the cases at 0.466 m/s inlet speed with dominant convective transport.

Another consideration is that, while most DNPs that exit the core decay quickly, some have the opportunity to re-enter at the inlet, and thus still contribute to β_{eff}. Therefore, the calculation
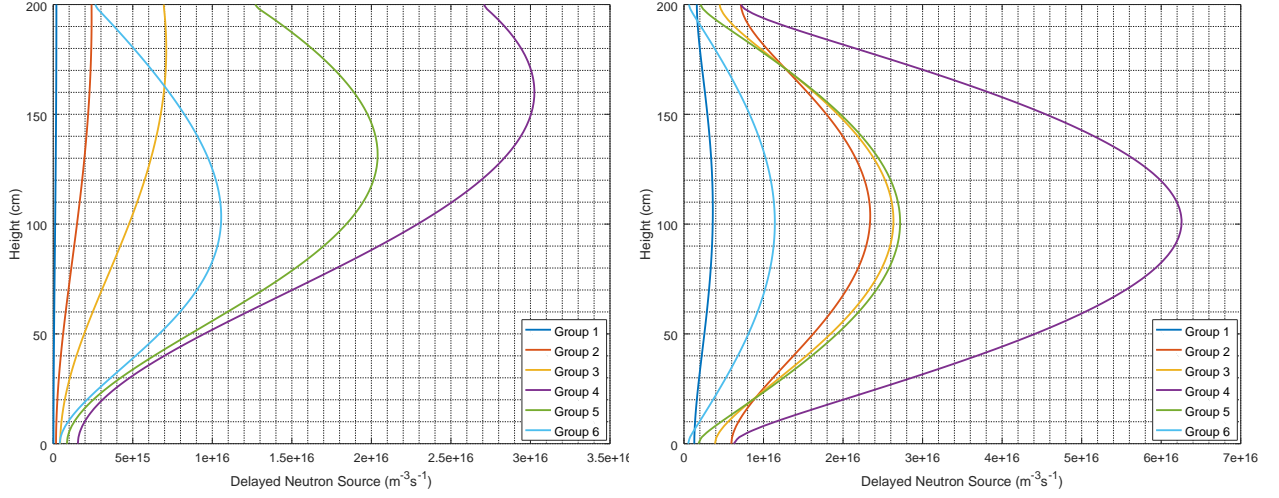
Figure 4.29: Delayed neutron source along the central axis of the CMSR with flow (left) and without flow (right).

of β_{eff} in such reactors depends on the in-core and ex-core circulation time, not just just the time-to-egress. A simple analytical expression has previously been derived for the loss of reactivity (difference in β_{eff} in this study) as a result of DNP transport (Cammi et al., 2011):

$$\Delta\beta_{\text{eff}} = \Delta\rho = \sum_k \frac{\beta_k \lambda_k}{\lambda_k + \frac{1-\exp\left(-\lambda_k \tau_{\text{ext}}\right)}{\tau_{\text{int}}}} \ . \tag{4.19}$$

Though DNP re-entry is not considered in this work, the above formula is still applicable if one considers the limit of an infinite ex-core residence time. Meanwhile, the in-core residence time is readily approximated since the axial velocity is nearly constant. With this, the reactivity loss is found to be approximately 350 pcm, compared to the simulated result of 408 pcm (Fig. 4.17). If one were to account for re-entry, for instance with an ex-core time three times longer than the in-core time, the resulting β_{eff} loss would be 317 pcm — a difference of about 10%.

While Eq. 4.19 is highly useful, it does not account for the variation of neutron importance within the core. One can directly calculate the β_{eff} (rather than the reactivity loss from DNPs leaving the core) with the following formula (Aufiero et al., 2014):

$$\beta_{\text{eff},i} = \frac{\int_V \sum_{j=1}^6 \phi_j^\dagger \chi_{d,j} \lambda_i c_i \ d\mathbf{r}}{\int_V \sum_{j=1}^6 \phi_j^\dagger \chi_{d,j} \sum_{k=1}^8 \lambda_k c_k \ d\mathbf{r} + \int_V \sum_{j=1}^6 \phi_j^\dagger \chi_{p,j} \sum_{j'=1}^6 \varphi_{j'} \left(\nu\Sigma_f\right)_{j'} \ d\mathbf{r}} \ . \tag{4.20}$$

Similar to Eq. 4.11 and Eq. 4.12, this expression uses the adjoint flux as a proxy for neutron importance. Thus it can be interpreted as the ratio of the importance-weighted delayed neutron source from a given precursor group to the total importance-weighted neutron production rate. Moreover, it is perfectly general regardless of the in-core or ex-core residence time, since it depends on the real spatial concentration of DNPs in the core. The total β_{eff} was then computed using this method, giving 275 pcm for the case without any temperature feedback (as in Fig. 4.17). This corresponds to a difference in β_{eff} of approximately: $\Delta\beta_{\text{eff}} = 275 \text{ pcm} - 680 \text{ pcm} = -405 \text{ pcm}$, which is almost exactly the value obtained by directly

comparing reactivities, thus indicating that the latter method is a reliable indicator of the effect of flow speed and Schmidt number on β_{eff}. Going further, Eq. 4.20 can be applied to a case *with* thermal coupling. Adjusting the terms in said equation for the temperature field, this results in a β_{eff} of 272 pcm rather than 275 pcm — a marginal difference which shows that the partially coupled case is a good approximation.
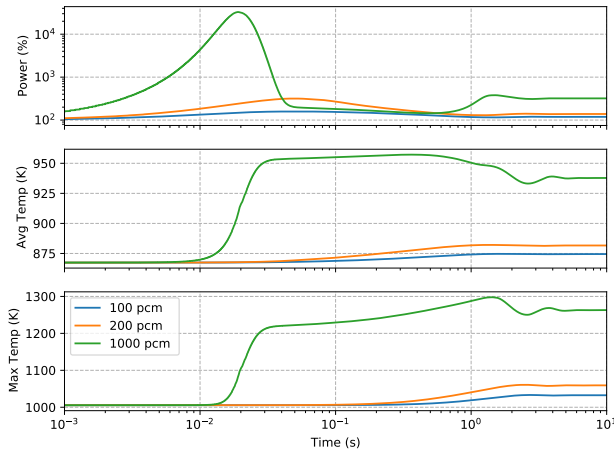
4.3.3 Transient Response

Unprotected Transient Over-Power

Like many reactor physics codes, GeN-Foam allows a bulk reactivity insertion to be directly imposed in a transient simulation. By default, a transient begins by adopting the last-known estimation of k_{eff} as the value to use in the production term of the diffusion equation. Thus, if a transient begins from steady-state, the production and loss rate are perfectly balanced (by definition) and no changes occur in reactor power. Alternatively, the user can choose the value of k_{eff} to be used for the simulation, so that a desired starting reactivity can be set in spite of the material properties. For example, if a user sets the starting value to unity, then the transient is allowed to proceed as it actually would based on the last-known k_{eff} at steady-state. Or, as was done in this work, one can tune the imposed k_{eff} to set a particular initial reactivity.

In Fig. 4.30, reactivities of 100 pcm, 200 pcm, and 1000 pcm are inserted in three separate cases, without insertion of control rods (sometimes referred to as a UTOP transient). The first two are both insertions less than β_{eff}, and as such the response is expected to be strongly attenuated by delayed neutrons. On the other hand, the 1000 pcm case is comfortably above β_{eff}, such that the prompt response becomes apparent. Not only does the power and temperature climb much higher in this case, but the power peak occurs much earlier. It is only because of the combined effects of Doppler-broadening and density feedback that the power does not climb orders of magnitude further. In all cases, the fuel temperature exhibits some instabilities in the first few seconds, which are also reflected in the reactor power.

Unprotected Loss of Flow

A ULOF is a postulated accident which pertains to liquid-fueled MSRs (Fiorina et al., 2014), and will most likely be present in the design basis of the CMSR. The DSA campaign therefore demands an understanding of the progression and consequences of a ULOF — the loss of circulation in the primary loop without control rod insertion. This has two simultaneous effects: an increase in the importance of DNPs (reactivity goes up), and an increase of temperature in the fuel (reactivity goes down). However, being that the DNPs are transported by the same fuel which is slowed down, the reactivity insertion from an increase in the delayed neutron fraction in the core is expected to occur on the same time-scale at which the fuel temperature is increased. The question is then whether the negative temperature feedback is sufficient to overcome the increase in β_{eff} throughout the transient.

Figure 4.30: Transient response to bulk reactivity insertion.
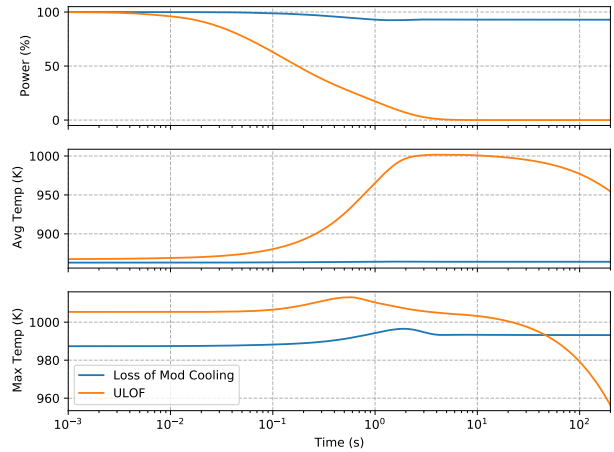
Figure 4.31: Transient response to changes in heat exchange.

A limiting case of ULOF is now considered, in which the flow is immediately brought to a complete halt. Though, one should note that this is not necessarily the worst-case scenario as the immediate loss of flow results in the steepest and greatest increase in fuel temperature, which also leads to the most negative reactivity. There may yet be some slower loss-of-flow scenario that yields a comparable insertion of reactivity from DNPs while having less temperature feedback (due to the different behaviour and influence of the DNP groups) which may result in a temporarily supercritical state. On the other hand, the scenario considered here gives the fastest positive reactivity insertion from accumulating DNPs, though it still takes time for this accumulation to occur. It is observed in Fig. 4.31 that the negative reactivity from the increase in fuel temperature overwhelms any DNP effects, as the power does not surge upward at any point in time. The power instead drops rapidly to zero (note that decay heat is not modelled) — the reactor shuts itself down. It should also be noted that it is impossible for this sort of accident to result in a prompt-critical state since the delayed neutron source is the only contributor to positive reactivity, so that any insertion is in a sense compensated by an increased margin to prompt-criticality.

Moderator Cooling Failure

As a point of comparison to the ULOF, a much less severe postulated accident also based on loss of heat transfer has been simulated. From the onset of the transient, the moderator cooling which was shown in Fig. 4.15 was instantly turned off, causing the fuel temperature to rise. In this scenario, the steady-state solution reverts to the case presented in Fig. 4.14. The transition between these two steady-states is thus shown in Fig. 4.31. As compared to the ULOF, the rise in fuel temperature is rather gradual and the power decreases slowly to match. The peak temperature also occurs slightly later than the ULOF since it corresponds to the new steady-state. Overall, this type of accident seems to be of low consequence, with no adverse effects besides a modest temperature increase.

Chapter 5

Conclusion

This work has followed a natural progression from testing and verification of GeN-Foam to its successful application toward multiphysics modelling of the CMSR. In addition to the computed results, this served as a general assessment of GeN-Foam with respect to its potential to inform the design of the CMSR going forward. Overall, the results are quite satisfactory, and encourage the future use of GeN-Foam for MSR modelling.

At the start, it was shown that the current version of GeN-Foam is fully capable of reproducing the results of the CNRS-MSFR benchmark. There were some minor deviations arising from inconsistent choices of cross-section data, but these do not detract from the overall agreement in the results. On the contrary, the level of agreement is remarkable considering that the benchmark was performed with almost no communication with the original participants. With this, the capability of GeN-Foam to accurately model fully coupled temperature-density feedback and DNP feedback in both steady-state and transient conditions was demonstrated, insofar as the accuracy of the original benchmark results can be accepted as true.

Using the benchmark model as a base, a verification of the neutronics solver in GeN-Foam was then performed based on the analytical form of the point-kinetic transfer function. This required considerable understanding of reactor dynamics and analysis thereof in the frame of control theory. The verification was a success for the most part, with only a deviation in the absolute value of the amplitude of the transfer function. With this undertaking, it was shown that the diffusion neutronics model in GeN-Foam is implemented correctly, and in a more rigorous sense than with the benchmark.

Finally, having built up confidence in GeN-Foam's accuracy and effectiveness as a simulation tool, a new model was prepared of the CMSR. This model added on a great deal of complexity compared to prior cases. With no comparable reference results, its steady-state solution was analysed in depth, extracting relatively converged flow, flux, temperature, and DNP concentration fields. Particular focus was placed on DNP transport and its sensitivity to the simulation parameters, with the main finding that the Schmidt number need not be known to high accuracy in such regimes of strongly forced convection. Appropriate feedback effects were observed at steady-state with fixed power. Selected transient scenarios were also simulated,

exhibiting the expected behaviour which reflects the inherent safety of MSRs when the power is allowed to vary. No unexpected positive reactivity feedback was observed at any point.

While the results produced in this work are encouraging, there is plenty of room for further improvement. The MSFR benchmark could be re-done with matching cross-section libraries and higher-precision group constant generation, to be sure of the cause of the minor discrepancies observed. The neutronic verification could also be corrected by implementing a direct reactivity oscillation in GeN-Foam through some modification of the source code (though this would in a sense defeat the purpose of V&V), or by further study of the reactivity perturbation estimation as a result of the temperature fluctuation (including possibly a series of Monte Carlo calculations). From this point, the full verification scheme for liquid fuel could also be applied as per *Demazière, Mylonakis, and Vinai, 2021* to subject GeN-Foam's treatment of DNP transport feedback to the same rigour. In addition, as GeN-Foam continues to be developed, a strict V&V procedure requires such benchmarks and verification schemes to be re-applied any time a new version is used to obtain safety-related results.

In terms of the DSA campaign of the CMSR, the model implemented so far would need to be significantly improved to be viable. A more representative geometry would need to be described by the neutronic mesh, so as to preserve the heterogeneous properties of the core. The neutronic model could also be enhanced by using one or more of the higher-order approximations of the transport equation available in GeN-Foam (SP_3, S_N) which would become useful for modelling heterogeneities and local reactivity insertions (e.g. spurious control rod withdrawal). The resulting flux distribution and k-eigenvalue would then need to be benchmarked against an equivalent Monte Carlo result.

On the thermal-hydraulic side, the porous media model could continue to be used as long as greater attention is paid to the turbulence modelling implementation, ensuring that the length scale of the discrete fuel tubes is considered. There is also potential for the two-phase thermal-hydraulics model of GeN-Foam to be used for simulating the effect of noble gases (emitted as fission products) as well as helium bubbling on the reactivity feedback of the fuel. Components outside of the active core could also be added, including the upper and lower plena, the heat exchanger to the secondary side, and the rest of the primary loop. This would allow GeN-Foam to be used for high-level system analysis with a holistic understanding of pressure drops throughout, temperature gradients, and DNP transport including re-entry. Armed with a more sophisticated model of the CMSR and primary loop, it would also be possible to simulate postulated accidents in much greater detail. For example, a ULOF could be fully modelled with a gradual ramp-down of the primary pumps (using GeN-Foam's `momentumSource` feature) and the resulting effects throughout the loop.

Between its continued development in the coming years and an increased competence in its use, GeN-Foam is a great candidate for multiphysics analysis of the CMSR. Although its open-source nature may detract from its licensing potential, it could nonetheless play a supporting role in the DSA campaign as well, provided that appropriate V&V steps are taken.

Acknowledgements

First and foremost I would like to thank my supervisor, Andreas Vigand Schofield for his guidance and valuable feedback, which he gave freely over the course of this work. Moreover, *mange tak* to the whole of Seaborg Technologies for hosting me during the five months of my thesis. The support I have received throughout has been invaluable, both in relation to this project and in my next steps. Special thanks to Tim Ruscoe for crucial OpenFOAM discussion and generous meshing assistance, Jacob Groth-Jensen for facilitating use of the Seaborg cluster, Esben Klinkby for additional comments on this thesis, and to Mateusz Pater for his past multiphysics work which was a personal inspiration for me to go this route.

I must also recognize the incredible work of countless people that this thesis relies on, particularly the modern movers and shakers in the field — some of whom I had the fortune of having as instructors and mentors. The developers of OpenFOAM and GeN-Foam deserve particular mention for making their software free and open. Thank you especially to Carlo Fiorina and the EPFL doctoral students for offering tech support with GeN-Foam. Thanks also to Christophe Demazière for providing manuscripts about the neutronic verification scheme along with helpful advice.

Finally, I cannot thank my friends and family enough for keeping me going through these historic times. To Narcís, Pau, Krzysztof, Ana, and Oscar, thanks for the tapas, the Zoom calls, and the laughs at UPC. To my Canadian friends, thanks for checking in during lockdowns and for keeping in touch even though plans to visit were thwarted. Thanks also to Arman and the Seaborg interns past and present, for the regular opportunities to decompress. Thank you Elise, for making the last year a lot brighter. And thank you to my parents and brother, for showing me the meaning of tenacity, and for supporting me as I ventured into what became a slightly unconventional study-abroad experience.

September 2021
Copenhagen, Denmark

Symbols

Neutronics

$\mathbf{J}$	neutron current
α	neutron albedo
β	delayed neutron fraction
χ_d	delayed emission spectrum
χ_p	prompt emission spectrum
λ	decay constant
Λ_0	mean generation time
ν	mean fission neutron yield
Φ	angle-integrated flux
φ	scalar flux
$\phi^\dagger$	adjoint flux
ψ	point-kinetic shape function
ρ	reactivity
Σ	macroscopic cross-section
C	DNP concentration/number density
D	neutron diffusion coefficient
D_m	mass diffusion coefficient
E	neutron kinetic energy
k_eff	effective multiplication factor
P	point-kinetic amplitude or neutronic power
v	neutron speed

Thermal-hydraulics

$\boldsymbol{F}_g$	gravitational force
$\boldsymbol{F}_{ss}$	sub-scale drag force
$\boldsymbol{u}$	fuel velocity
$\boldsymbol{u_D}$	fuel Darcy velocity
β_V	thermal expansion coefficient
γ	volumetric porosity
μ	dynamic viscosity
ρ	density
ζ	volumetric heat transfer coefficient
A_V	sub-scale surface area
c_p	specific heat capacity at constant pressure
D_h	hydraulic diameter
e	specific internal energy
f_D	Darcy friction factor
h	surface area heat transfer coefficient
k_T	thermal conductivity
$\dot{m}$	mass flow rate
p	pressure
$\dot{Q}$	thermal power
$\dot{Q}_{ss}$	sub-scale heat source/sink
T	fuel salt temperature
T_ext	external temperature
T_{ss}	sub-scale temperature
Nu	Nusselt number
Pr	Prandtl number
Re	Reynolds number
Sc	Schmidt number

Acronyms

ARE	Aircraft Reactor Experiment
ASME	American Society of Mechanical Engineers
BEPU	Best Estimate Plus Uncertainty
CAD	Computer-Aided Design
CANDU	Canada Deuterium Uranium
CFD	Computational Fluid Dynamics
CMSR	Compact Molten Salt Reactor
CNRS	Centre National de la Recherche Scientifique
DNP	Delayed Neutron Precursor
DNS	Direct Numerical Simulation
DSA	Deterministic Safety Analysis
ENDF	Evaluated Nuclear Data File
EPFL	École Polytechnique Fédérale de Lausanne
FOAM	Field Operation and Manipulation
GIF	Generation-IV International Forum
JEFF	Joint Evaluated Fission and Fusion File
JFNK	Jacobian-Free Newton-Krylov
LOCA	Loss of Coolant Accident
LWR	Light Water Reactor
MSFR	Molten Salt Fast Reactor
MSRE	Molten Salt Reactor Experiment
ORNL	Oak Ridge National Laboratory
PDE	Partial Differential Equation
PSI	Paul Scherrer Institute
RANS	Reynolds-Averaged Navier-Stokes
SEALION	Seaborg External multiphysics Architecture for Licensing and IP development of Nuclear reactors
SMR	Small Modular Reactor
TRL	Technology Readiness Level
TUD	Technische Universiteit Delft
ULOF	Unprotected Loss of Flow
V&V	Verification & Validation

Bibliography

ASME (2021). *Boiler and Pressure Vessel Code Section III – Rules for Construction of Nuclear Facility Components*.

Aufiero, Manuele et al. (2014). "Calculating the effective delayed neutron fraction in the Molten Salt Fast Reactor: Analytical, deterministic and Monte Carlo approaches". In: *Annals of Nuclear Energy* 65, pp. 78–90. ISSN: 03064549. DOI: 10.1016/j.anucene.2013.10.015. URL: http://dx.doi.org/10.1016/j.anucene.2013.10.015.

Bell, George and Samuel Glasstone (1970). *Nuclear Reactor Theory*. New York.

Bettis, E. S. et al. (1957). "The Aircraft Reactor Experiment—Design and Construction". In: *Nuclear Science and Engineering* 2.6, pp. 804–825. ISSN: 0029-5639. DOI: 10.13182/nse57-a35495.

Cammi, A. et al. (2011). "Transfer function modeling of zero-power dynamics of circulating fuel reactors". In: *Journal of Engineering for Gas Turbines and Power* 133.5. ISSN: 07424795. DOI: 10.1115/1.4002880.

Clifford, Ivor D (2013). "A Hybrid Coarse and Fine Mesh Solution Method for Prismatic High Temperature Gas-Cooled Reactor Thermal-Fluid Analysis". PhD thesis. Pennsylvania State University.

Cooper, Daniel (2020). "Next-generation nuclear reactor". In: *Materials World. Institute of Materials, Mining, and Engineering*, pp. 36–38. URL: https://www.iom3.org/resources/publications/materials-world.html.

De Oliveira, Rodrigo G.G. and Konstantin Mikityuk (2018). "Analytical solutions to a coupled fluid dynamics and neutron transport problem with application to GeN-Foam verification". In: *Annals of Nuclear Energy* 121, pp. 446–451. ISSN: 18732100. DOI: 10.1016/j.anucene.2018.07.036. URL: https://doi.org/10.1016/j.anucene.2018.07.036.

Demazière, C., V. Dykin, and K. Jareteg (2017). "Development of a point-kinetic verification scheme for nuclear reactor applications". In: *Journal of Computational Physics* 339, pp. 396–411. ISSN: 10902716. DOI: 10.1016/j.jcp.2017.03.020.

Demazière, C., A. Mylonakis, and P. Vinai (2021). "Development and test of a novel neutronic verification scheme for Molten Salt Reactors". In: *American Nuclear Society*. Vol. 124. 1, pp. 504–507. URL: https://www.ans.org/pubs/transactions/article-49641/.

Duderstadt, J.J. and L.J. Hamilton (1976). *Nuclear Reactor Analysis*, p. 123. DOI: 10.13182/nt76-a16299.

EVOL (2013). *EVOL Final Report*. Tech. rep., pp. 1–34. URL: https://cordis.europa.eu/docs/results/249/249696/final1-final-report-f.pdf.

Fermi, Enrico (1952). *Fermi's Own Story*. URL: https://www.energy.gov/sites/prod/files/TheFirstReactor.pdf.

Bibliography

Fiorina, Carlo (2013). "The Molten Salt Fast Reactor as a Fast-Spectrum Candidate for Thorium Implementation". In: March, p. 183. URL: https://www.politesi.polimi.it/bitstream/10589/74324/1/2013_03_PhD_Fiorina.pdf.

Fiorina, Carlo et al. (2014). "Modelling and analysis of the MSFR transient behaviour". In: *Annals of Nuclear Energy* 64, pp. 485–498. ISSN: 03064549. DOI: 10.1016/j.anucene.2013.08.003. URL: http://dx.doi.org/10.1016/j.anucene.2013.08.003.

Fiorina, Carlo et al. (2015). "GeN-Foam: A novel OpenFOAM® based multi-physics solver for 2D/3D transient analysis of nuclear reactors". In: *Nuclear Engineering and Design* 294, pp. 24–37. ISSN: 00295493. DOI: 10.1016/j.nucengdes.2015.05.035. URL: http://dx.doi.org/10.1016/j.nucengdes.2015.05.035.

Fiorina, Carlo et al. (2016). "Development and verification of the neutron diffusion solver for the GeN-Foam multi-physics platform". In: *Annals of Nuclear Energy* 96, pp. 212–222. ISSN: 18732100. DOI: 10.1016/j.anucene.2016.05.023. URL: http://dx.doi.org/10.1016/j.anucene.2016.05.023.

GIF (2002). "A technology roadmap for generation IV nuclear energy systems". In: *Nuclear Energy Research Advisory Committee and the . . .*, pp. 1–97. URL: http://www.gen-4.org/PDFs/GenIVRoadmap.pdf.

Gougar, H D et al. (2015). *Assessment of the Technical Maturity of Generation IV Concepts for Test or Demonstration Reactor Applications*. Tech. rep. INL/EXT-15-36427, pp. 1–33.

Groth-Jensen, J. et al. (2021). "Verification of multiphysics coupling techniques for modeling of molten salt reactors". In: *Annals of Nuclear Energy* 164, p. 108578. ISSN: 03064549. DOI: 10.1016/j.anucene.2021.108578. URL: https://doi.org/10.1016/j.anucene.2021.108578.

Gunow, Geoffrey Alexander (2018). "Full Core 3D Neutron Transport Simulation Using the Method of Characteristics with Linear Sources". PhD thesis. URL: https://dspace.mit.edu/handle/1721.1/119030%0Ahttp://hdl.handle.net/1721.1/119030.

Haubenreich, P. and J. Engel (1970). "Experience With the Molten-Salt Reactor Experiment". In: *Nuclear Applications and Technology* 8.2, pp. 118–136. ISSN: 0550-3043. DOI: 10.13182/nt8-2-118.

He, Qingming et al. (2020). "The JFNK method for the PWR's transient simulation considering neutronics, thermal hydraulics and mechanics". In: *Nuclear Engineering and Technology* 52.2, pp. 258–270. ISSN: 2234358X. DOI: 10.1016/j.net.2019.07.029. URL: https://doi.org/10.1016/j.net.2019.07.029.

Hébert, A. (2016). *Applied Reactor Physics*. 2nd ed. Montreal: Presses internationales Polytechnique. ISBN: 9782553016981.

Henriroux, Claire et al. (2019). "Development of Regulatory Compliant Coupled Neutronics-Thermohydraulics Models for Molten Salt Reactors". In: *International Conference on Mathematics and Computational Methods applied to Nuclear Science and Engineering*. Portland, pp. 1848–1856.

Incropera, Frank P et al. (2006). *Fundamentals of Heat and Mass Transfer*. New Jersey. DOI: 10.1016/j.applthermaleng.2011.03.022. arXiv: 1105-.

IPCC (2018). "IPCC report Global warming of 1.5°C". In: *IPCC* 2.October, pp. 17–20. URL: https://www.ipcc.ch/sr15/.

Klinkby, E.B. (2019). *Reactor Core System - System Design Description*. Tech. rep. Seaborg Technologies.

Kortov, V. and Yu Ustyantsev (2013). "Chernobyl accident: Causes, consequences and problems of radiation measurements". In: *Radiation Measurements* 55, pp. 12–16. ISSN: 13504487. DOI: 10.1016/j.radmeas.2012.05.015. URL: http://dx.doi.org/10.1016/j.radmeas.2012.05.015.

Lamarsh, John R. (1983). *Introduction to Nuclear Engineering*. 3rd Editio. Addison-Wesley Publishing Company. ISBN: 9780134570051.

Launder, B.E. and B.I. Sharma (1974). "Application of the energy-dissipation model of turbulence to the calculation of flow near a spinning disk". In: *Letters in Heat and Mass Transfer* 10.2, pp. 131–138. DOI: https://doi.org/10.1016/0094-4548(74)90150-7..

Leppänen, Jaakko (2007). "Development of a new Monte Carlo reactor physics code". PhD thesis. ISBN: 9513870189.

Leppänen, Jaakko, Maria Pusa, and Emil Fridman (2016). "Overview of methodology for spatial homogenization in the Serpent 2 Monte Carlo code". In: *Annals of Nuclear Energy* 96. ISSN: 18732100. DOI: 10.1016/j.anucene.2016.06.007.

Leppänen, Jaakko et al. (2015). "The Serpent Monte Carlo code: Status, development and applications in 2013". In: *Annals of Nuclear Energy* 82, pp. 142–150. ISSN: 18732100. DOI: 10.1016/j.anucene.2014.08.024.

Lindsay, Alexander et al. (2018). "Introduction to Moltres: An application for simulation of Molten Salt Reactors". In: *Annals of Nuclear Energy* 114, pp. 530–540. ISSN: 18732100. DOI: 10.1016/j.anucene.2017.12.025. URL: https://doi.org/10.1016/j.anucene.2017.12.025.

MacNamara, Shev and Gilbert Strang (2016). *Splitting Methods in Communication, Imaging, Science, and Engineering*. Ed. by Rowan Glowinski, Stanley J. Osher, and Wotao Yin. Springer, p. 95. ISBN: 978-3-319-41587-1. DOI: 10.1007/978-3-319-41589-5.

Meulekamp, Robin Klein and Steven C Van Der Marck (2017). "Calculating the Effective Delayed Neutron Fraction with Monte Carlo Calculating the Effective Delayed Neutron Fraction with Monte Carlo". In: *Nuclear Science and Engineering* 152.2, pp. 142–148.

Milani Shirvan, Kamel et al. (2016). "Enhancement of heat transfer and heat exchanger effectiveness in a double pipe heat exchanger filled with porous media: Numerical simulation and sensitivity analysis of turbulent fluid flow". In: *Applied Thermal Engineering* 109, pp. 761–774. ISSN: 13594311. DOI: 10.1016/j.applthermaleng.2016.08.116. URL: http://dx.doi.org/10.1016/j.applthermaleng.2016.08.116.

Moukalled, F., L. Mangani, and M. Darwish (2015). *The Finite Volume Method in Computational Fluid Dynamics*. Springer Publishing Company, p. 4. ISBN: 9783319168739.

OpenFOAM (2014). *The OpenFOAM 7 User Guide*. URL: https://www.openfoam.com/.

Pedersen, A.V. (2020). *Preliminary thermophysical properties relevant to CMSR thermal hydraulics modeling*. Tech. rep. Seaborg Technologies.

Rofer, Cheryl (2015). "Why did the US abandon a lead in reactor design?" In: *Physics Today*, pp. 1–12. DOI: 10.1063/pt.5.2029.

SAMOSAFER Consortium (2020). *SAMOSAFER*. URL: https://samosafer.eu/project/.

Siemer, Darryl D. (2015). "Why the molten salt fast reactor (MSFR) is the "best" gen IV reactor". In: *Energy Science and Engineering* 3.2, pp. 83–97. ISSN: 20500505. DOI: 10.1002/ese3.59.

Sun, Dabin et al. (2020). "An Improved Best Estimate Plus Uncertainty Method for Small-Break Loss-of-Coolant Accident in Pressurized Water Reactors". In: *Frontiers in Energy Research* 8.August, pp. 1–16. ISSN: 2296598X. DOI: 10.3389/fenrg.2020.00188.

Tiberga, Marco et al. (2020). "Results from a multi-physics numerical benchmark for codes dedicated to molten salt fast reactors". In: *Annals of Nuclear Energy* 142, p. 107428. ISSN: 18732100. DOI: 10.1016/j.anucene.2020.107428.

Weinberg, Alvin M (1976). "The Maturity and Future of Nuclear Energy". In: *American Scientist* 64.1, pp. 16–21.

Weller, H. G. et al. (1998). "A tensorial approach to computational continuum mechanics using object-oriented techniques". In: *Computers in Physics* 12.6, p. 620. ISSN: 08941866. DOI: 10.1063/1.168744.
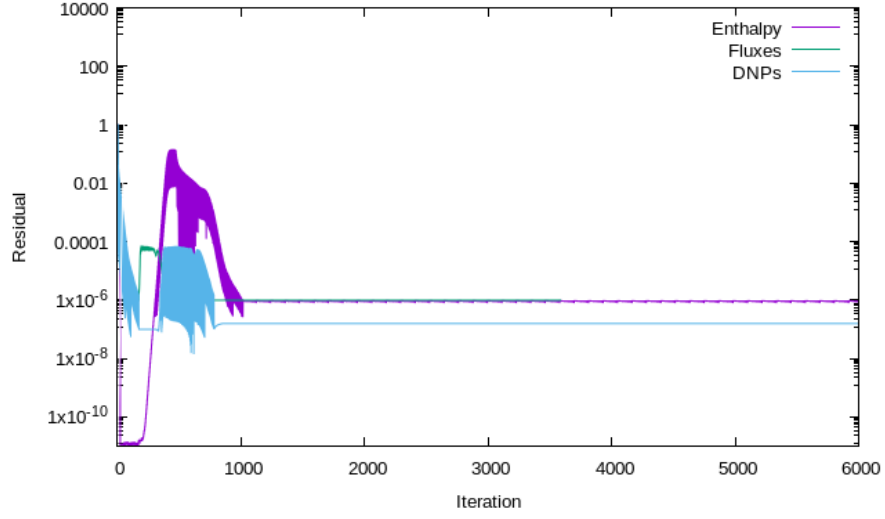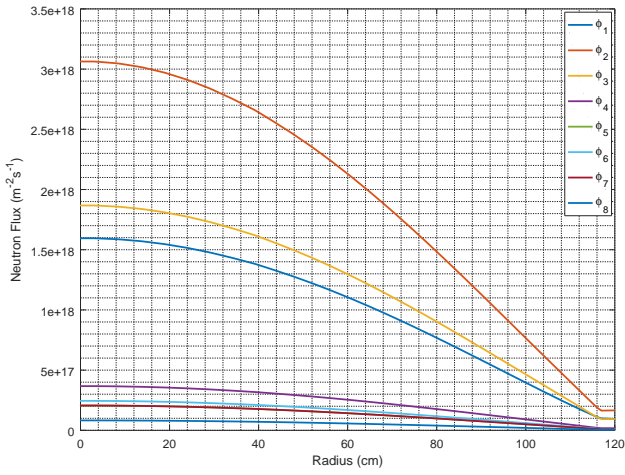
Appendix A

Supplementary Plots
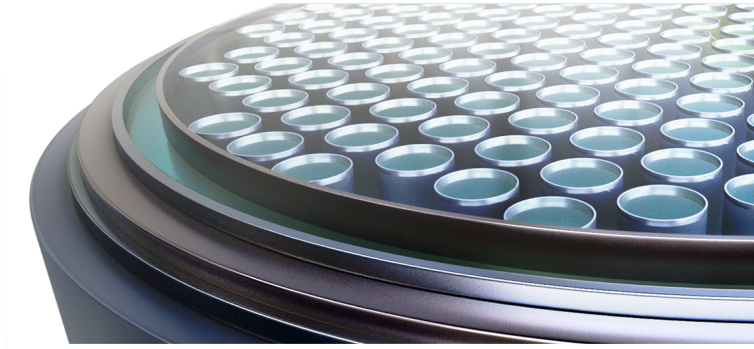
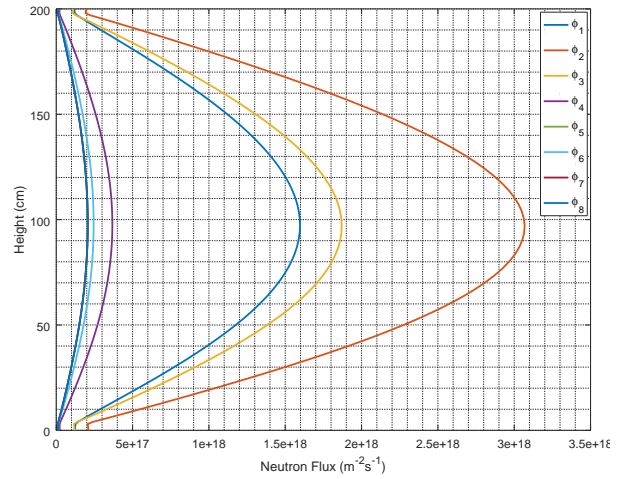**Figure A.1:** Conceptual view of the CMSR showing fuel tubes in the reactor vessel.
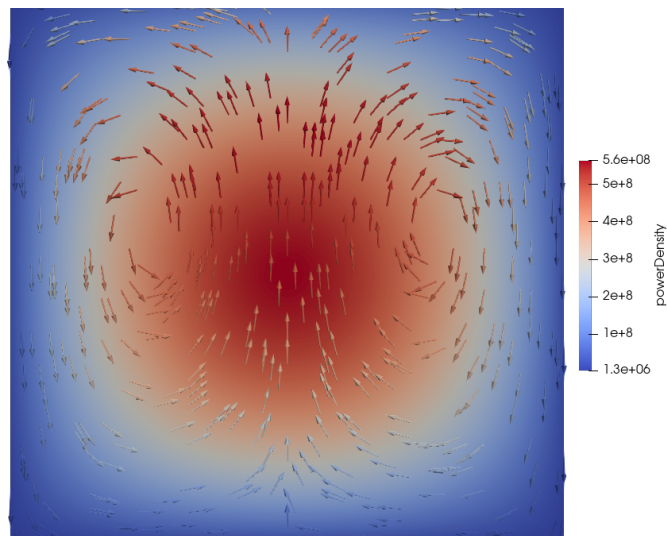
Figure A.2: Flow streamlines in the 2D-MSFR cavity for the case of buoyancy-driven flow (step 1.3), shown against the power distribution (W/m^3).
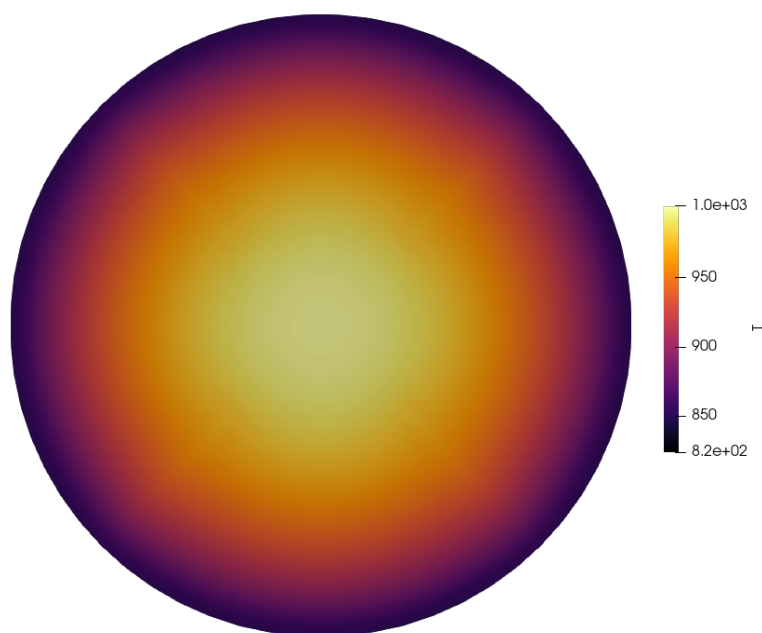
Figure A.3: Example of residuals converging to target in a CMSR simulation.

Figure A.4: Radial flux profile in the CMSR core, in 8 groups.

Figure A.5: Axial flux profile in the CMSR core, in 8 groups.

Figure A.6: Temperatures in a 2-D slice of the CMSR core outlet.