

Nuclear Heat Transfer and Passive Cooling

Study B: Fuel Assembly CFD and UQ for a Molten Salt Reactor



Volume 1

Introduction to the Technical Volumes and Case Studies



Volume 2

Convection, Radiation and Conjugate Heat Transfer



Volume 3

Natural Convection and Passive Cooling



Volume 4

Confidence and Uncertainty



Volume 5

Liquid Metal Thermal Hydraulics



Volume 6

Molten Salt Thermal Hydraulics



Study A

Liquid Metal CFD Modelling of the TALL-3D Test Facility



Study B

Fuel Assembly CFD and UQ for a Molten Salt Reactor



Study C

Reactor Scale CFD for Decay Heat Removal in a Lead-cooled Fast Reactor



Study D

System Code and CFD Analysis for a Light Water Small Modular Reactor

Authors: Graham Macpherson Frazer-Nash Consultancy
Contributors: Lewis Farrimond Frazer-Nash Consultancy
Carolyn Howlett Frazer-Nash Consultancy
Bo Liu The University of Sheffield
Adam Owens Moltex Energy
Case Studies Lead: Graham Macpherson Frazer-Nash Consultancy
Approver: Richard Underhill Frazer-Nash Consultancy

Document Number: FNC 60148/50159R
Issue and Date: Issue 1, December 2021

Legal Statement

This document presents work undertaken by Frazer-Nash Consultancy Ltd and funded under contract by the UK Government Department for Business, Energy and Industrial Strategy (BEIS). Any statements contained in this document apply to Frazer-Nash Consultancy and do not represent the views or policies of BEIS or the UK Government. Any copies of this document (in part or in full) may only be reproduced in accordance with the below licence and must be accompanied by this disclaimer.

This document is provided for general information only. It is not intended to amount to advice or suggestions on which any party should, or can, rely. You must obtain professional or specialist advice before taking or refraining from taking any action on the basis of the content of this document.

We make no representations and give no warranties or guarantees, whether express or implied, that the content of this document is accurate, complete, up to date, free from any third party encumbrances or fit for any particular purpose. We disclaim to the maximum extent permissible and accept no responsibility for the consequences of this document being relied upon by you, any other party or parties, or being used for any purpose, or containing any error or omission.

Except for death or personal injury caused by our negligence or any other liability which may not be excluded by an applicable law, we will not be liable to any party placing any form of reliance on the document for any loss or damage, whether in contract, tort (including negligence) breach of statutory duty, or otherwise, even if foreseeable, arising under or in connection with use of or reliance on any content of this document in whole or in part.

Unless stated otherwise, this material is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License. You may copy and redistribute the material in any medium or format, provided you give appropriate credit, provide a link to the license and indicate if changes were made. If you remix, transform, or build upon the material, you may not distribute the modified material. You may not restrict others from doing anything the license permits.



Preface

Nuclear thermal hydraulics is the application of thermofluid mechanics within the nuclear industry. Thermal hydraulic analysis is an important tool in addressing the global challenge to reduce the cost of advanced nuclear technologies. An improved predictive capability and understanding supports the development, optimisation and safety substantiation of nuclear power plants.

This document is part of *Nuclear Heat Transfer and Passive Cooling: Technical Volumes and Case Studies*, a set of six technical volumes and four case studies providing information and guidance on aspects of nuclear thermal hydraulic analysis. This document set has been delivered by Frazer-Nash Consultancy, with support from a number of academic and industrial partners, as part of the UK Government *Nuclear Innovation Programme: Advanced Reactor Design*, funded by the Department for Business, Energy and Industrial Strategy (BEIS).

Each technical volume outlines the technical challenges, latest analysis methods and future direction for a specific area of nuclear thermal hydraulics. The case studies illustrate the use of a subset of these methods in representative nuclear industry examples. The document set is designed for technical users with some prior knowledge of thermofluid mechanics, who wish to know more about nuclear thermal hydraulics.

The work promotes a consistent methodology for thermal hydraulic analysis of single-phase heat transfer and passive cooling, to inform the link between academic research and end-user needs, and to provide a high-quality, peer-reviewed document set suitable for use across the nuclear industry.

The document set is not intended to be exhaustive or provide a set of standard engineering 'guidelines' and it is strongly recommended that nuclear thermal hydraulic analyses are undertaken by Suitably Qualified and Experienced Personnel.

The first edition of this document set has been authored by Frazer-Nash Consultancy, with the support of the individuals and organisations noted in each. Please acknowledge these documents in any work where they are used:

Frazer-Nash Consultancy (2021) Nuclear Heat Transfer and Passive Cooling,
Study B: Fuel Assembly CFD and UQ for a Molten Salt Reactor.

Contents

1	Introduction	1
2	Planning the Analysis	3
2.1	Reactor and Conditions to be Analysed	3
2.2	Planning the Analysis	5
2.2.1	Moltex Fuel Assembly	6
2.2.2	Extent and Boundary Conditions for Fuel Assembly CFD Model	10
2.2.3	Sources of Input Data	12
2.2.4	Thermal Hydraulic Phenomena of Importance	17
2.2.5	Modelling Strategy	20
3	Performing the Analysis	24
3.1	Buoyancy Affected Vertical Channel Flow	24
3.2	Closed Vertical Tubes With a Heat Generating Fluid	28
3.3	Resolved Rod Assembly	33
3.3.1	Geometry and Mesh	33
3.3.2	Case Setup	36
3.3.3	Obtaining a Converged Solution	38
3.3.4	Interpretation of Results	41
3.3.5	Turbulence Modelling in the Fuel Pin	43
3.4	Single Fuel Pin	46
3.5	Single Coolant Subchannel	47
3.5.1	Comparison to Full Assembly Model	47
3.5.2	Variation in Flow and Derivation of Correlations	53
3.5.3	Solid Fuel Approximation	58
3.6	Transverse Effective Conductivity	59
4	Application of the Results	63
4.1	Porous Assembly	64
4.1.1	Porous Model Implementation	67
4.1.2	Comparing Porous and Resolved Assembly Models	71
4.2	Surrogate Model of Fuel Conditions	78
4.3	Uncertainty Quantification of Single Coolant Channel	80
4.3.1	Selection of Uncertain Input Parameters	81
4.3.2	Sensitivity Analysis and Uncertainty Quantification Method	83
4.3.3	Uncertainty Quantification Results	84
4.3.4	Sensitivity Analysis Results	89
5	References	92
6	Nomenclature	95
7	Abbreviations	97
	Appendix A Single Fuel Pin Turbulence Assessment	98

1 Introduction

This case study is concerned with modelling a fuel assembly for a Molten Salt Reactor (MSR), including heat generating fluids and radiative heat transfer. To provide realism and improve the utility of this case study, it has been developed in collaboration with Moltex Energy¹.

As for all the case studies in this series, this analysis provides a ‘worked example’ of a specific modelling task to illustrate the modelling approaches described in the technical volumes. Therefore, it is recommended that this case study is read in conjunction with the following technical volumes:

Volume 2: Convection, Radiation and Conjugate Heat Transfer

Volume 3: Natural Convection and Passive Cooling

Volume 4: Confidence and Uncertainty

Volume 6: Molten Salt Thermal Hydraulics

Many publicly available examples of thermal hydraulic analyses are performed as part of benchmark or validation studies, as demonstrated in Study A (Liquid Metal CFD Modelling of the TALL-3D Test Facility) and Study D (System Code and CFD Analysis for a Light Water Small Modular Reactor). In these cases the geometry and boundary conditions are usually simple and well defined, and comparison or validation data (in the form of both high-fidelity simulation results and experimental measurements) is available. However, for many industrial thermal hydraulic modelling engineers, this does not generally represent the situations in which they often need to work.

Industrial simulations are built using the best information and methods available at the time, while recognising that there are unknowns and uncertainties that remain, and that additional confidence building measures would be necessary to make a high-significance decision using it. Simulations of this type can be used to inform design decisions, and will form part of the process of deciding what additional experiments or measures are necessary to build this confidence. This case study and Study C (Reactor Scale CFD for Decay Heat Removal in a Lead-cooled Fast Reactor) are examples of this type of analysis.

The objectives of this case study focus on two areas:

- Using Computational Fluid Dynamics (CFD) to create heat transfer and pressure drop correlations applicable to a porous medium representation of a fuel assembly. Correlations suitable for steady-state normal operation, and for quasi-steady-state buoyancy driven natural circulation will be developed, including an assessment of their uncertainty.
- Quantification of the uncertainty in the CFD model predictions due to uncertainty in material properties. There is generally a poor state of knowledge of the thermophysical properties

¹ www.moltenergy.com

Introduction

of molten salts – they are not known with high accuracy, and change with temperature and salt composition. An assessment of the extent to which variation in them affects the figures of merit will be demonstrated, which can be used to target where efforts to develop better knowledge are best placed, and to determine which properties do not merit substantial additional investigation.

While the methods demonstrated by this case study use a specific MSR design as an example, fuel assemblies for many reactor designs could be analysed in a similar manner. The layout of the described reactor is broadly in line with the Moltex Stable Salt Reactor – Wasteburner (SSR–W) design. However, to enable open publication, various aspects of the geometry, component performance and analysis scenario are fictitious. Therefore, no quantitative data presented in this study (either input data or analysis results) should be taken as illustrating the actual design or performance of the SSR–W.

2 Planning the Analysis

2.1 Reactor and Conditions to be Analysed

The SSR–W design uses a molten salt with dissolved fissile material contained within closed vertical fuel pins, a hexagonal array of which are arranged into a fuel assembly that has a layout similar to that used in liquid metal cooled reactors. The assemblies are immersed in a coolant salt in a pool-type reactor configuration. During normal operation, the coolant salt is driven through the fuel assemblies and primary heat exchangers by pumps.

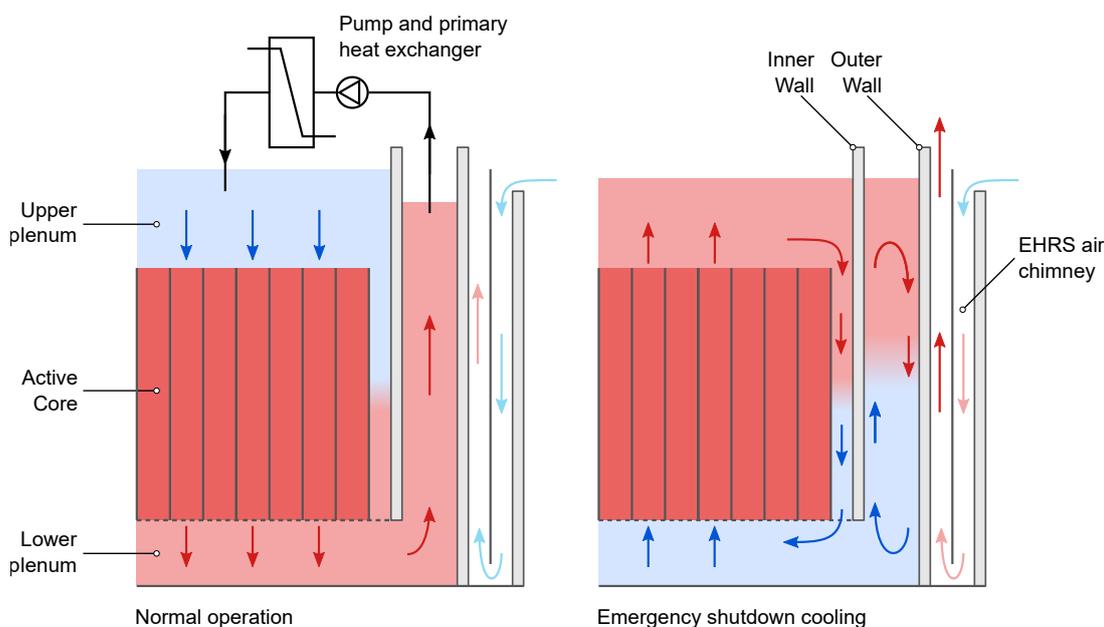


Figure 2.1: Schematic layout of SSR–W showing expected flow directions.

Under the conditions of a postulated Station Blackout (SBO) event (loss of off-site and on-site power), the reactor will ‘trip’ (fission will stop), but the loss of power will cause the forced coolant flow and heat removal by the primary heat exchangers to stop. The reactor will then undergo a transient to establish decay heat removal by natural circulation, rejecting decay heat to a passive Emergency Heat Removal System (EHRS) air chimney.

A feature of the SSR–W concept design assessed here is that during normal operation, the flow of coolant is from the upper plenum to the lower plenum. To establish natural circulation, this direction must reverse. A schematic of the reactor and flows in these two conditions is shown in Figure 2.1. In the passive cooling conditions, the coolant circulates through the fuel assemblies in the active core to the upper plenum and returns to the lower plenum via bypass paths. The decay heat is

Planning the Analysis

transferred through the inner wall to a circulating cavity, and then through the outer wall to the air chimney.

Under all conditions, the available margin between the maximum fuel salt and pin material temperatures and their limits should be known, and it is necessary to demonstrate that the EHRS is able to meet its required performance, while maintaining sufficient margins. The flow rates and temperatures in the reactor under the required range of normal operation, fault and decay heat removal conditions for design and licensing purposes can only be practically predicted by a low-fidelity representation: a system code model or porous medium CFD representation of the core¹. These techniques require correlations to determine the flow losses and heat transfer characteristics of coolant flow in the fuel assemblies. These could be determined from representative and/or prototypical experiments, literature correlations or CFD analysis, potentially using a combination of them all to provide comparison.

At the concept design stage of a novel reactor, it is likely that sufficiently specific experimental data will not be available for the geometry, fluids or flow and heat transfer regime expected. Similarly, literature correlations, which are typically derived from experimental data, have limitations in that they represent idealised geometries and conditions, and would need to be combined with each other in ways that are not part of their derivations. CFD does not have those limitations, and can also use the correlations as a way to check the accuracy and geometrical dependence of the results, and guide what physics and flow regime to expect. Also, at the concept design stage, the geometry, materials and conditions expected are subject to change, so modelling by CFD provides a more rapid and economical tool for design development and assessment than physical testing.

For the SBO transient, the correlations are required to be sufficiently accurate and applicable during three reactor conditions:

1. Forced convection when the pumps are active.
2. Mixed convection after natural circulation passive cooling is established.
3. A combination of forced, natural and mixed convection in the cross-over period following the reactor trip as the coolant pumps coast-down to a stop and the flow direction reverses. The flow in the fuel assemblies will undergo a complex transient, during which there may be non-uniform or counter-flow in the coolant spaces of the fuel assembly cross-section.

This case study will assess the first two of these (the third is possible as an extension of the techniques described, but involves significant additional modelling challenges).

The objective of this case study is to demonstrate how to start from a position of knowledge of the geometry, materials, and initial expected reactor conditions that is relevant to the concept design stage, and use CFD to create correlations suitable for a porous medium representation. This includes giving some consideration to their accuracy, and to the effect of uncertainty in material properties.

¹ The long duration of the transients involved, and the need to include the whole primary and secondary secondary circuits means that a high-fidelity CFD model is prohibitively expensive.

Planning the Analysis

2.2 Planning the Analysis

As part of the ongoing reactor design process, the SBO event has been simulated using a system code model of the reactor, where 1/3rd of the reactor (one of the three primary heat exchanger circuits) has been modelled with the core represented by four parallel channels (Figure 2.2), representing four regions from the inner (CH1) with the highest power assemblies to the outer region (CH4) with the least. The flow losses have been chosen such that the flow distribution among channels closely matches the power distribution, giving a uniform exit temperature to the lower plenum.

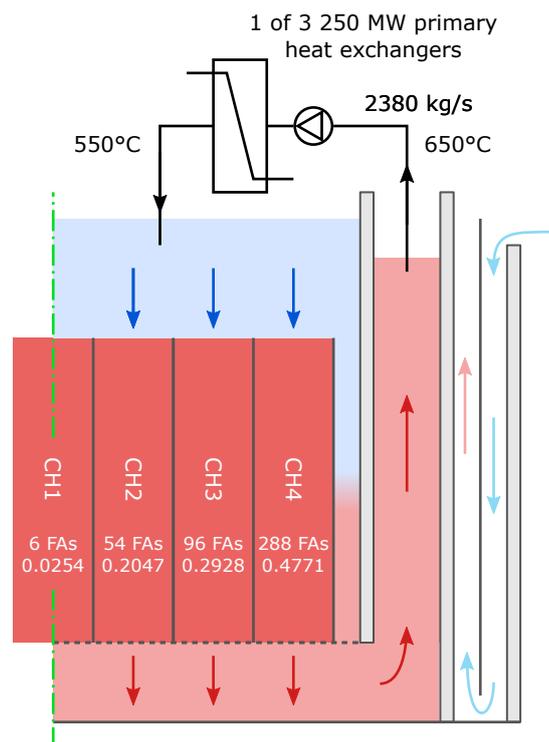


Figure 2.2: Core split into four concentric regions of channels, showing the number of Fuel Assemblies (FAs) in each for the whole core, as well as the split of core power and flow. Nominal core inlet and outlet temperatures are 550 °C (823.15 K) and 650 °C (923.15 K) respectively.

The SBO transient is initiated at time $t = 0$, where the primary pumps are 'tripped' due to loss of power; they coast-down to stop over 100 s. The nuclear heating from fission is terminated immediately (by the insertion of control rods) and a reducing decay heat is applied. This transient is shown for Region 2 (CH2) in Figure 2.3. The flow reverses direction at approximately 110 s, and the upper plenum temperature exceeds the lower plenum temperature at approximately 130 s. This case study only considers the core after this point, where natural circulation conditions have been established – three instants in the transient were chosen to analyse: 150 s, 250 s and 500 s.

Using these system code predictions as inputs, simulations can be performed of a fuel assembly for two types of reactor condition:

1. Steady-state calculations for normal operation for the four core regions, applying constant boundary conditions.

Planning the Analysis

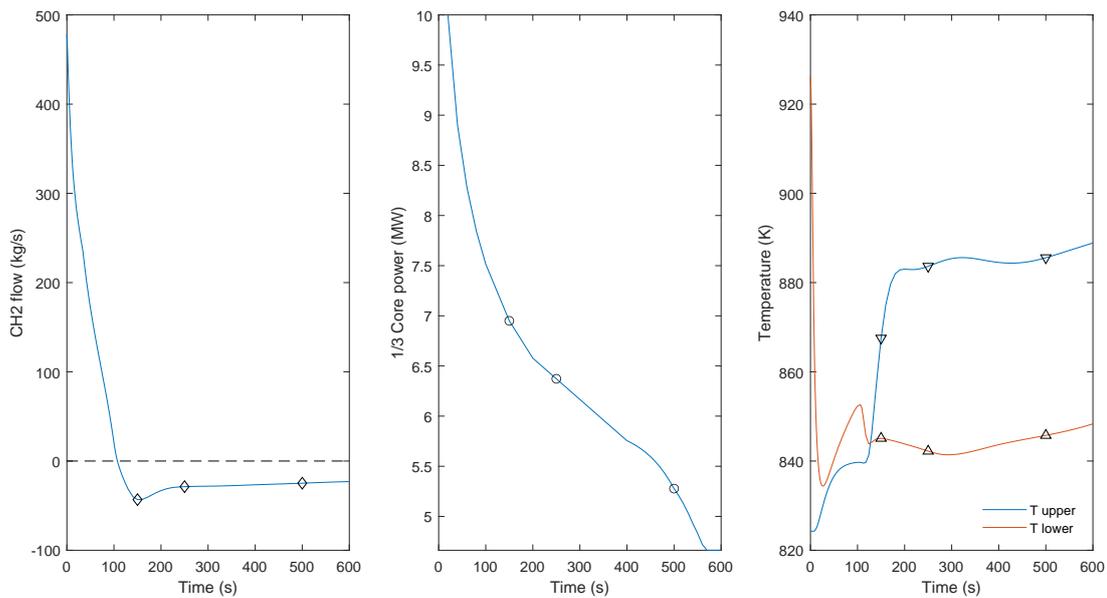


Figure 2.3: Flow in CH2, core power and upper and lower plenum temperatures during the SBO transient. Chosen times for simulation marked.

2. Quasi steady-state calculations of natural circulation in CH2 at the three chosen times, applying boundary conditions that are assumed to not change significantly over the timescales of the flow (hence ‘quasi’ steady).

These seven conditions are intended to capture the response of the fuel assemblies to variations in power and flow. The performance of the molten fuel salt is also expected to be sensitive to variations in:

- Thermophysical and radiative heat transfer (optical) properties of the fuel and coolant salt and of the cladding.
- Fuel and core geometry.

The effect of the first of these will be explored by applying Sensitivity Analysis (SA) and Uncertainty Quantification (UQ) techniques. However, these techniques could be readily extended to encompass geometry variations and applied to system code simulations too.

2.2.1 Moltex Fuel Assembly

The fuel assembly design under consideration is shown in Figure 2.4 with the expected natural convection circulating flow pattern of molten fuel salt inside the fuel pin shown in Figure 2.5.

The fuel pins are arranged in a triangular array within a hexagonal fuel assembly. Kakaç *et al.* (1987, Chapter 7) provides a description of the geometry, frictional losses and heat transfer in this kind of array. The main dimensions are the pitch of the spacing of the pins, p , the pin outer diameter, d and the cladding thickness, t , as shown in Figure 2.5.

The flow in the spaces between pins in the array is characterised largely by, p/d , the pitch to diameter ratio. The hydraulic diameter, D_h , of the coolant flow path between the pin array is calculated

Planning the Analysis

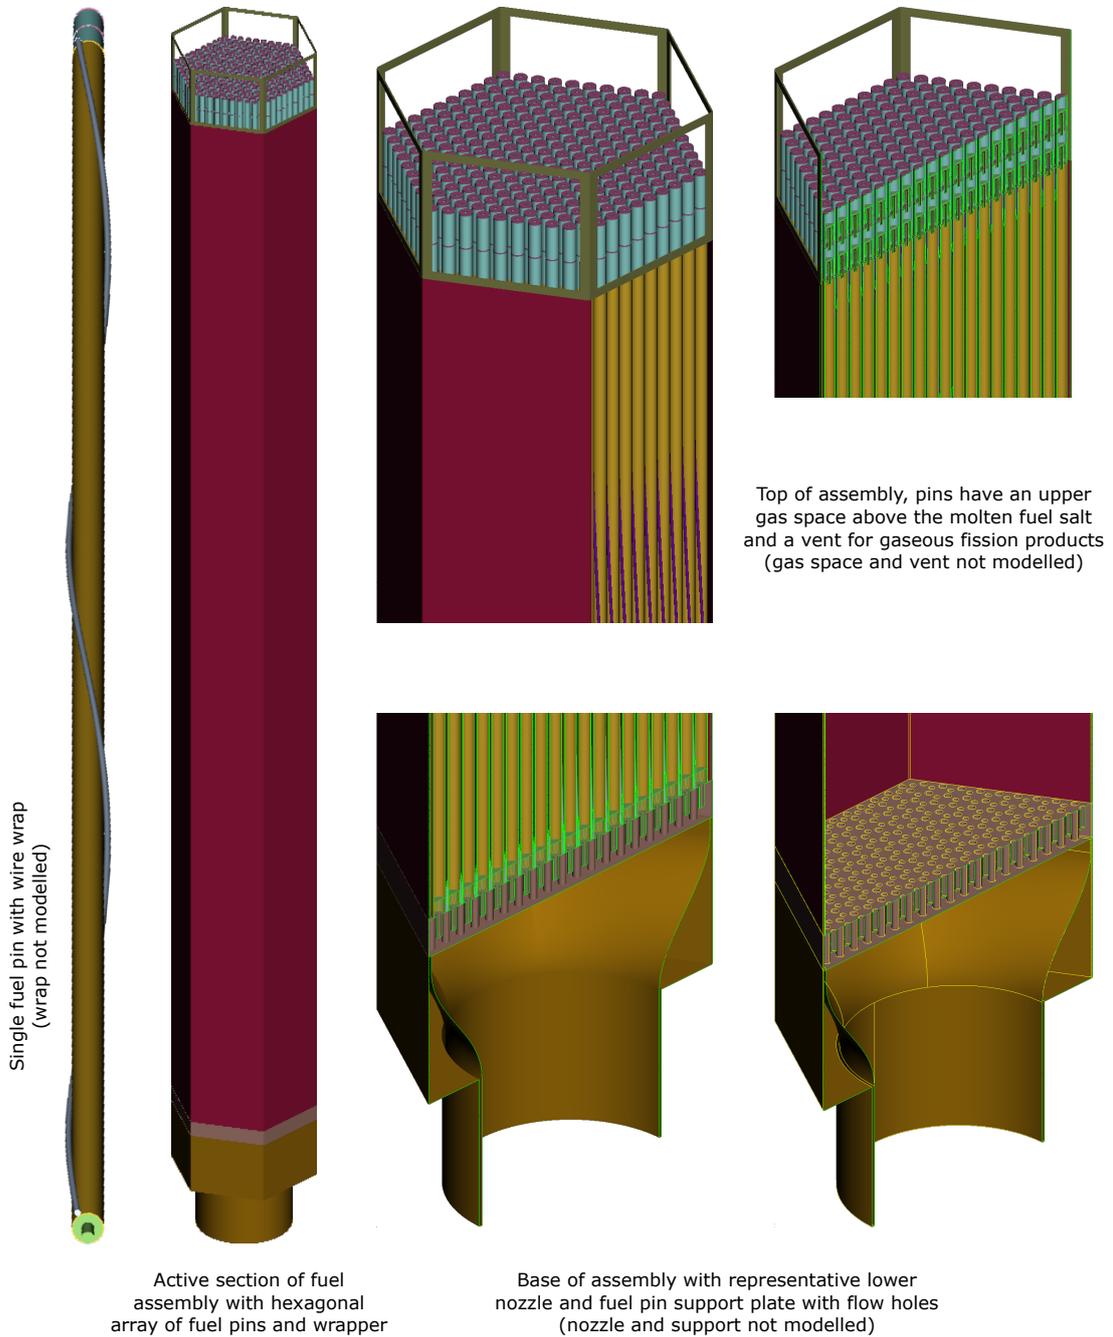


Figure 2.4: CAD model of the fuel assembly concept design, used as a source of dimensions and to visualise included and omitted features.

Planning the Analysis

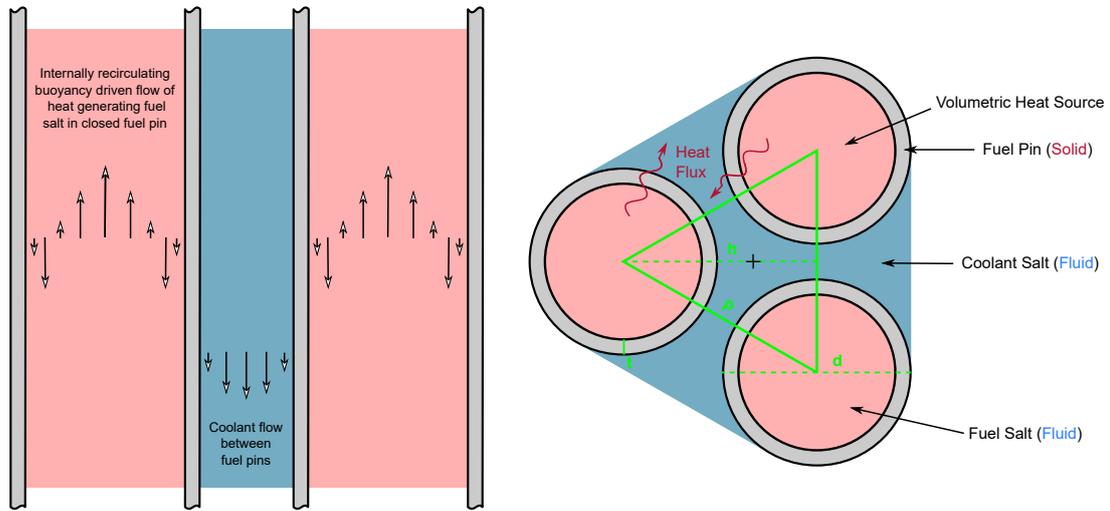


Figure 2.5: Fuel assembly pin array components and dimensions.

Array dimensions (m)	
p	0.012
d	0.01
t	0.0003
h	0.010392
L	1.6
D_h	0.0058783
h_{duct}	0.10046
Array properties	
p/d	1.2
N_{ring}	9
N_{pin}	271
Duct and gap dimensions (m)	
t_{duct}	0.001
w_{gap}	0.0031
Optical path lengths (m)	
$2h - d$	0.010785
$p - d$	0.002
$d_{fuel}/2$	0.0047

Table 2.1: Fuel assembly geometric properties.

Planning the Analysis

from the subchannel bounded by the triangle in Figure 2.5, where the coolant flow area is given by

$$A = \frac{\sqrt{3}}{4}p^2 - \frac{\pi}{8}d^2$$

The height, h , of the subchannel is

$$h = \frac{\sqrt{3}}{2}p$$

and the centre of the subchannel is at $(2/3)h$ from the fuel pin centre (marked with a cross). The fuel assembly has 9 rings of pins around the centre pin, giving a total of

$$N_{pin} = 3N_{ring}(N_{ring} + 1) + 1 = 271$$

pins in each fuel assembly cross-section. The inner wall of the duct² is placed at the centre of the outer subchannel, so is at a distance of

$$h_{duct} = (N_{ring} + 2/3)h$$

from the array centre. The total cross-section of the assembly is bounded by the duct:

$$A_{assembly} = \frac{6}{\sqrt{3}}h_{duct}^2$$

The flow area available to coolant flow is given by

$$A_{coolant} = A_{assembly} - N_{pin}\frac{\pi}{4}d^2$$

The active length (containing fuel salt) of the fuel pins is $L = 1.6$ m. The dimensions of the array are given in Table 2.1, along with other derived dimensions that are required later.

Results of Importance: The primary considerations for analysing the fuel assembly are to predict the temperature of the fuel salt and cladding material under all credible flow and heat flux conditions.

- The cladding is a barrier preventing the fuel salt from mixing with the coolant which would contaminate the reactor, and so ensuring its integrity is a primary consideration. Damage and failure mechanisms can be non-linearly sensitive at elevated temperatures, and so knowing the maximum temperature of the cladding is typically required.
- The duct also requires its structural integrity to be assured, and could be susceptible to thermal stresses, so the distribution of heat transfer and temperature around its circumference and along its length will allow this to be assessed.
- Having a negative reactivity temperature feedback coefficient is an important operational and safety feature for MSRs, so knowledge of the temperature of the fuel salt is important from a criticality perspective.
- The fuel temperature also affects the retention of gaseous fission products, and, in the most extreme cases, boiling of the salt must be avoided.

² The words 'duct' and 'wrapper' will be used interchangeably. The wrapper surrounds the whole fuel assembly, and is not to be confused with a 'wire wrap' around each pin.

Planning the Analysis

The range of credible conditions depends on the resistance to flow in the assemblies (as well as the driving forces and losses in the rest of the core) and the local heat transfer between the coolant and the cladding. Therefore, characterising these conditions provides a direct route to assessing the fuel and cladding temperatures that have safety implications.

2.2.2 Extent and Boundary Conditions for Fuel Assembly CFD Model

To develop a CFD model of the fuel assembly, it is necessary to consider how the outputs from the simulations will be used, and what is the appropriate level of simplification that should be adopted in the geometry and boundary conditions. Here the target application for the results is a reactor scale simulation with the fuel assemblies represented by porous medium, but the results could also be easily converted into inputs for system code models.

Several choices for modelling an overall core layout are described in Section 4.1 of Study C (Reactor Scale CFD for Decay Heat Removal in a Lead-cooled Fast Reactor). The approach of resolving each assembly individually, with their wrappers represented as zero thickness walls, and the inter-wrapper gaps also represented by porous media, has been adopted in Study C to give better fidelity for the prediction of core flows. The porous MSR fuel assembly model developed here demonstrates a compatible approach. The use of per-assembly thin walls means that the porous models of the fuel and inter-wrapper gap do not need to apply any specific modelling to inhibit flow in the across-core direction, which would otherwise be needed to account for the effect of the ducts.

The inclusion of both fuel and coolant requires a Conjugate Heat Transfer (CHT) approach to be adopted for the fuel pin and duct, allowing the temperature distribution to be simultaneously computed in the coolant salt, fuel pin and fuel salt. Beyond this, there are a number of considerations involved in choosing what to model, some of which require their implications to be clearly anticipated before starting the modelling process. For others, the impact will only emerge from the simulation results, or exploratory precursor simulations.

- Should the model include inlet and outlet nozzles of the assembly?
- Does the full active length need to be modelled or would a shorter section be sufficient?
- Should details of the fuel pin at each end be included?
- Is modelling a whole assembly cross-section necessary, would a 1/6th sector of the hexagonal fuel assembly, or even a single subchannel be sufficient?
- Would a 2D slice be useful in determining transfer properties?
- How to account for wire wraps or spacer grids?

The design of a core and fuel to produce a desired flow distribution involves many considerations that combine the effects of the design of inlet and outlet nozzles to the fuel and above and below core structures. The flow openings and restrictions at the side and base of the core can, for example, be intended to control inter-wrapper flow and provide routes for flow under decay heat removal conditions. Because of this, a CFD model of a single fuel assembly cannot, on its own, be expected to be predictive of the core flow rates based on above and below core pressures, unless all these additional features are accounted for (these features can be represented in a system code or reactor scale CFD model, however). Therefore, to focus this work on the characterisation of the

Planning the Analysis

fuel assembly, the flow rates for both forced and natural circulation conditions will be prescribed for the CFD models considered here. As long as the flow rates, temperatures and heat fluxes are sufficiently representative of the range of core conditions that will be encountered, the correlations derived will be applicable to the fuel assembly flows encountered in a reactor scale model.

A significant simplification in this case study is the omission of the wire wraps from the fuel and the treatment of the pins as a bare triangular rod array. This simplifies the meshing and analysis of the pin array for demonstration purposes within the case study. The inclusion of wire wraps in a more realistic case would make some of the exploitation of simple symmetry or single subchannel models shown here less straightforward or not possible. Volume 5 (Section 2.3.1) discusses the modelling of wire wrapped fuel in the context of Liquid Metal-cooled Fast Reactors (LMFRs). Fuel pins without wire wraps typically have spacer grids, which can also be assessed with CFD (CSNI, 2013).

Omitting the wire wraps impacts on the losses and heat transfer to the coolant, but should not significantly affect the fuel salt modelling inside the fuel pins. The inlet and outlet nozzle structures will also be omitted, focusing this work on the performance of the coolant and fuel salt flow and heat transfer in the active length. When using the correlations developed in this study for a reactor scale simulation in a system code or porous medium CFD model, the effect of the inlet and outlet regions would need to be added on as additional losses. Separate or combined CFD model studies, similar to those described in this study could also be used to characterise those components.

It would only be necessary to resolve the whole cross-section of an assembly if any asymmetry was to be represented, potentially arising from partial blockage or non uniform conditions for inter-wrapper flow, temperature or heat load in adjacent assemblies. Otherwise, a $1/6^{\text{th}}$ (or $1/12^{\text{th}}$) symmetric sector can be used. As a further simplification, it is expected (but not guaranteed) that, except in the near-duct region, the behaviour of each fuel pin and coolant subchannel will be similar, so a model of a single coolant subchannel should be able to determine the axial behaviour. The use of symmetry planes for the fuel pins in a single subchannel model will force a symmetric recirculation pattern within them, which may or may not be physical. A model with a whole $1/6^{\text{th}}$ sector does not force this constraint on most of the fuel pins, and so if this behaviour is present (or more precisely, if the CFD model can predict it), it would be visible.

A $1/6^{\text{th}}$ sector model also gives the ability to assess the effect of the combination of thermal radiation and conduction on the effective value of thermal conductivity in the transverse direction. However, it is also possible to do this with a 2D slice.

For a single subchannel or $1/6^{\text{th}}$ sector model, it would be possible to use a part-length model (potentially periodic in that direction) instead of the full length of the fuel. If the fuel in the pins was solid and static, then that would be more likely to be a practical option, but because the fuel salt circulates and is bounded by a solid wall at the pin base and a free surface at the top, this process would be expected to be subject to significant distortion if not resolved over its full height. In fact, the recirculating behaviour is expected to be sensitive to the heat transfer at either end, and so an entrance length section will be included in the model in the coolant passages to move the boundary away from the top and bottom of the fuel salt.

The conclusion from these considerations is that it is worth creating a $1/6^{\text{th}}$ sector of the full active length of the fuel (plus some entry region), resolving the full 3D physics. The results from this will

Planning the Analysis

provide insight into what approximations are appropriate, and will serve as a reference to compare a 1/6th sector 3D porous model to. The porous model properties will be derived from a single subchannel full length model, with additional information on transverse properties derived from a 2D slice.

Inter-wrapper Gap: The fuel assembly duct is 1 mm thick, and to complete the definition of a 1/6th sector of a fuel assembly, an inter-wrapper gap (distance between outer flat faces of the fuel assembly ducts) needs to be defined. The system code results do not include inter-wrapper flows, and the core flows in the concept design are not finalised, so at this stage in the analysis cycle, the decision was taken to match the flow characteristics in the inter-wrapper gap to the fuel bundle (this can be updated in future analysis as the design evolves). One of the drivers for this design choice is minimising thermal stresses on the duct, and so the decision to promote equal flow on the inside and outside of the duct is intended to mitigate larger temperature differences across it.

By comparing literature correlations for turbulent and laminar frictional losses in a triangular array and a duct with infinite parallel plates, the value of $w_{gap} = 3.1$ mm was determined, where the hydraulic diameter is $2w_{gap} = 6.2$ mm, similar to the 5.88 mm hydraulic diameter of the triangular array. The same inlet velocity and temperature will be prescribed at the entrance to the gap as entering the fuel bundle.

To provide confidence that this is a reasonable choice to make for this dimension, other hexagonal fuel assembly designs were compared. This design detail is not readily available for many prototype reactors, but it was ascertained that:

- In the MYRRHA reactor design and accompanying inter-wrapper heat transfer test rig, the inter-wrapper gap is 3 mm (Uitslag-Doolaard *et al.*, 2019).
- In the CEFR reactor design, the inter-wrapper gap is 2 mm (Wang *et al.*, 2020).
- In the PRANDTL-DHX rig the gap is 7 mm, although that may be larger to facilitate experimental instrument access (Kamide *et al.*, 2001).

These references support the choice of the inter-wrapper gap as being realistic in terms of clearances for fuel handling, but it must be noted that they are for LMFRs (where the fuel pin arrays are broadly comparable, but have slightly smaller diameter pins, that are more closely spaced), which have significantly different heat transfer characteristics associated with the low Prandtl number coolant.

2.2.3 Sources of Input Data

A range of geometry, properties and operating conditions are necessary to carry out the simulations planned in this case study:

- The geometry of the fuel assembly and fuel pins.
- Material properties for the fuel salt and coolant (including optical properties) and the cladding, as well as an indication of the expected uncertainty in them.
- The temperature and flow conditions during forced and natural circulation conditions.
- The fission power (for normal operation/forced circulation conditions) or decay heat load (natural circulation conditions).

Planning the Analysis

Other than the material properties, this information is available from the CAD model of the fuel assembly (Figure 2.4) and system code predictions. The material properties for the specific salts and metals have been assembled from the open literature. This information was provided by Moltex Energy as representative of an SSR–W concept design.

2.2.3.1 Flow Conditions

The flow conditions and heat source per fuel assembly in forced and natural circulation conditions, along with the coolant inlet temperature (taking into account the flow direction) from the system code results are given in Table 2.2. From this, the inlet velocity to use as a CFD boundary condition and the mean volumetric heat source in the fuel salt were calculated. The inlet velocity is a mean value to be applied over $A_{coolant}$ at the density corresponding to the inlet temperature.

	W (kg/s)	Q (kW)	Inlet T (K)	Inlet U (m/s)	q_{vol} (MW/m ³)
Forced circulation					
CH1	30.23	3175	824.3	0.7552	105.5
CH2	26.57	2843.1	824.3	0.6638	94.48
CH3	21.78	2287.5	824.3	0.5441	76.02
CH4	11.83	1242.4	824.3	0.2955	41.29
Natural circulation (CH2)					
150 s	-2.404	79.037	845.1	0.06044	2.627
250 s	-1.595	72.462	842.2	0.04006	2.408
500 s	-1.379	60.014	845.8	0.03468	1.994

Table 2.2: Per assembly flow, heat load and inlet temperature, and derived inlet velocity and mean volumetric heat source.

2.2.3.2 Nuclear Heat Source

The energy input from fission or radioactive decay will be applied as a prescribed heat source. This would typically be supplied from a neutronics/reactor physics modelling tool that would take into account the fuel temperature, composition and irradiation (burnup). The energy source will be applied to the fuel in each operating condition as a volumetric heat source, uniform over each fuel pin cross-section, but able to vary along the fuel assembly length. Nuclear heating in the pin material could also be included, but has not been in this case.

Under normal operation, there is an axial power profile along the fuel caused by the distribution of neutrons, similar to solid fuelled reactors (although modified by the presence of mobile Delayed Neutron Precursors (DNPs)). The profile applied in this case is shown in Figure 2.6, and is used as a multiplier to the mean volumetric heat source (q_{vol}) for the forced cases in Table 2.2.

For solid fuelled reactors, there is a similar axial profile in the decay heat in the fuel, representing distribution and concentration of fission products. However, because the fuel salt is molten and mobile in this case, the fission products become sufficiently well mixed that a spatially uniform decay heat source can be used soon after fission ceases. Therefore, no profile is applied under natural circulation conditions, and the mean value is used.

Planning the Analysis

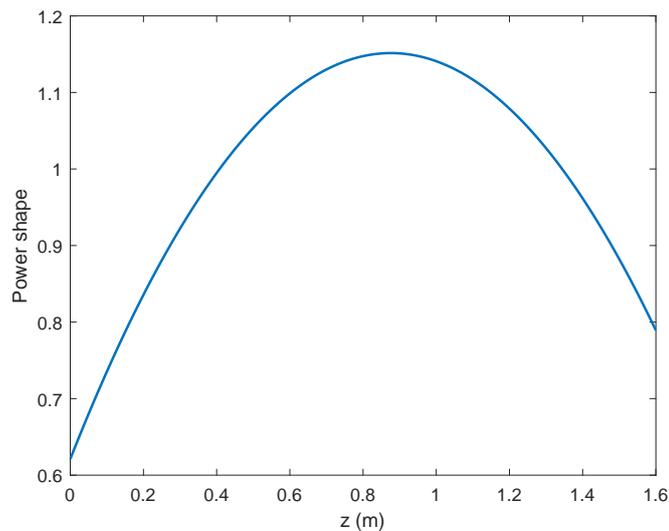


Figure 2.6: Axial power profile for normal operation (forced circulation).

2.2.3.3 Material Properties

The material properties used are defined as polynomials (or piecewise polynomials) as functions of temperature (in K) to be easily included in the chosen CFD code. The exponential functions for viscosity have fifth order polynomials fit to them. The properties are plotted over the expected range of temperatures in Figure 2.7. The salt k and c_p properties are estimated using the methods described in Volume 6 (Section 2.3.2).

The source references for experimental evaluation of salts, where available, have data that is measured at temperatures representative of the expected reactor coolant temperatures, but there are not well defined indications of the temperature range of validity or applicability of them.

Similarly, there are expected to be substantial uncertainties in the salt properties. There are uncertainty estimates available for these properties in other common salts (for example in Romatoski and Hu, 2017 and Richard *et al.*, 2014). These uncertainty estimates are used to perform UQ (Section 4.3). The cladding material is a more typical and well-studied material, so its properties have lower uncertainties and a better definition of the temperature validity range.

Fuel Salt Properties: The fuel salt used is KCl–UCl₃ (45–55 mole % eutectic)

Density (Desyatnik *et al.*, 1975):

$$\rho = 4344 - 1.015T \quad (\text{kg/m}^3)$$

Viscosity (Desyatnik *et al.*, 1975, converting to exponential form for consistency and to facilitate comparison to other salts):

$$\mu = 10^{-4.272+1644/T} = 5.3456 \times 10^{-5} e^{3785.4/T} \quad (\text{kg/ms})$$

Specific heat capacity, estimated using the Dulong and Petit method and the molecular weight of

Planning the Analysis

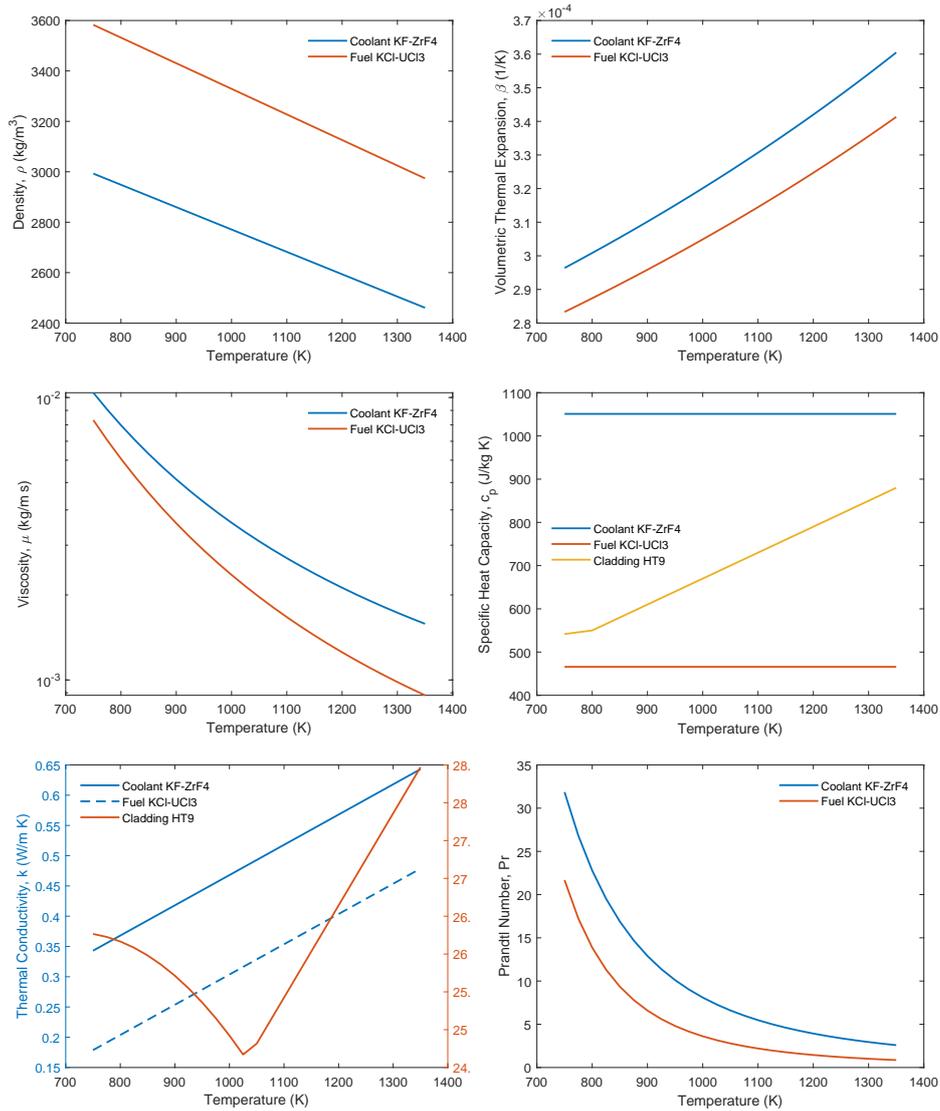


Figure 2.7: Thermophysical properties of fuel and coolant salts and fuel cladding.

Planning the Analysis

the salt constituents, which does not provide any guidance on temperature dependence:

$$c_p = 465.8 \quad (\text{J/kg K})$$

Thermal conductivity, estimated using Khoklov's equation and the molecular weight of the salt:

$$k = -0.1963 + 0.0005T \quad (\text{W/m K})$$

Coolant Salt Properties: The coolant salt used is KF-ZrF_4 (58–42 mole % eutectic)

Density (Richard *et al.*, 2014):

$$\rho = 3658 - 0.887T \quad (\text{kg/m}^3)$$

Viscosity (Williams, 2006, Richard *et al.*, 2014):

$$\mu = 1.59 \times 10^{-4} e^{3179/T} \quad (\text{kg/m s})$$

Specific heat capacity (estimated using the Dulong and Petit method, Williams, 2006):

$$c_p = 1051 \quad (\text{J/kg K})$$

Thermal conductivity (estimated using Khoklov's equation, Williams, 2006):

$$k = -0.032 + 0.0005T \quad (\text{W/m K})$$

Cladding Properties: The cladding and duct are made from HT9 martensitic stainless steel.

Density (Gelles, 1996):

$$\rho = 7800 \quad (\text{kg/m}^3)$$

Specific heat capacity (Yamanouchi *et al.*, 1992):

$$c_p = \begin{cases} \frac{T-500}{6} + 500 & T \leq 800 \\ \frac{3(T-800)}{5} + 550 & T > 800 \end{cases} \quad (\text{J/kg K})$$

Thermal conductivity, (Leibowitz and Blomquist, 1988):

$$k = \begin{cases} 17.622 + 2.42 \times 10^{-2}T - 1.69 \times 10^{-5}T^2 & T \leq 1030 \\ 12.027 + 1.218 \times 10^{-2}T & T > 1030 \end{cases} \quad (\text{W/m K})$$

HT9 has been associated with advanced reactors for a number of years, and data for more of its properties is available in collated forms (Hales *et al.*, 2016, Hofman *et al.*, 2019).

Planning the Analysis

Thermal Radiation Properties: The optical properties for the coolant and fuel salts in question are not well known, and are expected to be subject to change with the presence of fission or corrosion products. To account for this uncertainty, the approach taken was to bracket the range of possible effects by choosing absorption coefficients, κ , that provide high and low optical thicknesses over representative radiation path lengths in the coolant and fuel salts, ranging from nearly transparent to nearly opaque, as shown in Table 2.3. The three path lengths are shown in Table 2.1, representing the longest in-plane distance, from one fuel pin across two subchannels to another pin ($2h - d$), the shortest distance between pins ($p - d$) and from the centreline of the fuel to the cladding inner surface ($d_{fuel}/2 = (d - 2t)/2$).

The chosen values of κ are consistent (the same order of magnitude) with the values shown in Volume 6 (Section 2.5.2). A nominal value of $\kappa = 300 \text{ m}^{-1}$ was chosen for the fuel and coolant, to give balanced contributions from emission/absorption and across the gaps transmission, and the effect of variation between higher and lower values will be assessed. No wavelength dependence was implemented for the absorption coefficient, given the uncertainty assumed in its magnitude. A refractive index of 1.3 was chosen for both fuel and coolant salts, with no scattering, and a surface emissivity of 0.8 on all radiating surfaces. These are initial representative values, that can be revisited once the importance of radiative heat transfer and absorption coefficient have been ascertained.

$\kappa \text{ (m}^{-1}\text{)}$	$\kappa(2h - d)$	$\kappa(p - d)$	$\kappa d_{fuel}/2$
100	1.08	0.2	0.47
300	3.24	0.6	1.41
1000	10.8	2	4.7
3000	32.4	6	14.1

Table 2.3: Optical thickness of representative radiation paths for assumed absorption coefficients.

2.2.4 Thermal Hydraulic Phenomena of Importance

A number of phenomena are expected to be present in the fuel assembly that are not necessarily typical in other flows.

Heat generating fluid: The main differentiation between modelling the SSR-W fuel assembly and one suitable for a water or liquid metal cooled reactor is the presence of the molten fuel salt; it will be represented as a heat generating fluid and its motion and heat transfer will also be predicted in the CFD model.

Participating radiation: Salts can be semi-transparent to infrared radiation, and so including radiative heat transfer in the CFD model will be necessary, treating the salts as participating media, that absorb and emit radiation, as well as transmit it.

High Prandtl number fluid: The viscosity of the salts leads to moderately high Prandtl numbers, particularly in the coolant at the temperatures expected. This will lead to thin thermal boundary layers in the channels, requiring sufficient cells to resolve them.

Variation in properties: The range of temperature that could be present and the strong dependence of properties on temperature, particularly density and viscosity, mean that the effect of temperature variation must be considered. The expected density variations mean that the

Planning the Analysis

Boussinesq approximation is not applicable.

While these phenomena need to be included, the first questions to ask about modelling flow in the fuel assembly are: are the flows laminar or turbulent? how affected by buoyancy is the flow? what does that mean for the expected pressure drop and heat transfer?

2.2.4.1 Low Reynolds Number Flow and Buoyancy Effects in Coolant

From the flow and heat source conditions in Table 2.2, it is possible to estimate the expected mean coolant outlet temperature from the fuel assembly, and, by combining this flow and thermal information with the material properties and assembly geometry, then estimate the values of characteristic non-dimensional numbers to assess the relevance of these phenomena.

	Inlet Re	Outlet Re	Bo	Ra^*	Gr^*/Re	S
Forced circulation						
CH1	1831	2875	0.0157			50867
CH2	1609	2546	0.0218			45549
CH3	1319	2071	0.0347			36649
CH4	716	1125	0.152			19906
Natural circulation (CH2)						
150 s	160.1	185.3		282	70.8	1266
250 s	104.9	128.1		394	99.2	1161
500 s	92.1	111.5		380	95.6	961

Table 2.4: Flow regime non-dimensional numbers. Bo is only applicable to turbulent flows; Ra^* and Gr^*/Re in this context are only assessed for laminar flows.

Kakaç *et al.* (1987, Chapter 7) describe that there is no critical Reynolds number for laminar to turbulent transition in rod bundles (compared to the usually clearer transition behaviour in circular tubes). It provides data showing that, for a triangular array of $p/d = 1.2$, transition begins at approximately $Re = 1000$ and fully turbulent flow is expected for $Re \gtrsim 4000$. The effect of buoyancy on transition is not considered in this data, and in a real fuel assembly, turbulence at inlets from nozzles, or generated by wire wraps or spacer grids will be significant. Therefore, from Table 2.4 the expected Reynolds numbers of the coolant under forced circulation conditions in Regions 1 to 3 (CH1, CH2 and CH3) means that it is predicted to be weakly turbulent.

The difference between inlet and outlet Reynolds number indicates the effect of the variation in viscosity³. For CH4, the flow is in the transition region, which complicates its assessment. Under natural circulation conditions in CH2, the coolant channel flow is expected to be laminar.

The effects of buoyancy on heat transfer in turbulent vertical flow in pipes have been studied in detail, and are summarised by Jackson *et al.* (1989). For buoyancy assisted flows (upwards flows where a fluid is heated by channel walls) there is expected to be impairment of heat transfer at intermediate heat fluxes. Buoyancy opposed flows (downwards flows of a heated fluid) lead to heat transfer enhancement.

The Buoyancy parameter Bo can be used to assess the extent of the effect of buoyancy for vertical

³ For a fixed mass flow rate in a channel, Reynolds number only depends on the channel geometry and the fluid viscosity: $Re = WD_h/A\mu$.

Planning the Analysis

turbulent mixed convection flows (Keshmiri *et al.*, 2012)

$$Bo = 8 \times 10^4 \frac{Gr^*}{Re^{3.425} Pr^{0.8}}$$

where a heat flux form of Grashof number⁴ is used:

$$Gr^* = \frac{qg\beta D^4}{k\nu^2}$$

For vertical laminar mixed convection flows⁵, the parameter Gr^*/Re (Jackson *et al.*, 1989) or a form of Rayleigh number based on the vertical temperature gradient (analogous to uniform heat flux conditions, Kakaç *et al.* 1987, Chapter 15)

$$Ra^* = \frac{\rho^2 \beta g c_p D^4 (dT/dz)}{\mu k}$$

describe the extent of the effect of buoyancy. These parameters are shown in Table 2.4, evaluated at the mean of the inlet and outlet conditions.

- The values of Bo for the forced circulation conditions in CH1 to CH3, where there is buoyancy opposed (downward) flow indicate that there is not expected to be a significant buoyancy effect. The Reynolds number in CH4 are too small to use this (turbulent) correlation reliably, and so the higher value of Bo is harder to interpret.
- For laminar, buoyancy assisted natural circulation (upward) flow conditions, comparing the values of Gr^*/Re and Ra^* to the data in their source references indicates that only a small enhancement to the heat transfer is expected to be added to the values of Nusselt number (for fully developed flow) that would be present without buoyancy.

The effect of buoyancy on these flows is to introduce driving forces for the flow at the fuel pin surface due to body forces, and as such alter the velocity profile. This affects both heat transfer, but also friction factor. The extent to which flow losses are influenced by buoyancy is expected to be small, similarly to heat transfer.

The assessments described here use correlations and results derived for flows in vertical tubes, so applying them to the subchannel spaces in a vertical rod array is an approximation.

2.2.4.2 Recirculating Buoyancy Driven Low Reynolds Number Flow in Fuel Salt

It is not as straightforward to assess the nature of the expected self-generated recirculating flow behaviour in the heat generating salt in the closed fuel pin, however some guidance was found in the literature, dating from nuclear reactor development programmes of the 1950s and 1960s⁶. Murgatroyd and Watson (1970) describe experimental and theoretical studies of the same geometry and configuration, using water with uniform internal heat generation⁷ and a constant tube

⁴ The parameter of $\overline{Gr_b}/Re^{2.7}$, can also be used (Jackson *et al.*, 1989, Kakaç *et al.* 1987, Chapter 15), but applies a different form of Grashof number, $\overline{Gr_b} = g(\rho_b - \bar{\rho})D^3/(\bar{\rho}\nu^2)$, based on integrated density, $\bar{\rho}$ and bulk density, ρ_b .

⁵ The comparisons are made to data for experiments with uniform heat fluxes, which is a better representation of the conditions in a fuel assembly than the alternative idealisation of a uniform temperature. It is always necessary to check which thermal boundary condition applies to a heat transfer scenario, and to which condition a correlation or theory applies.

⁶ Such as the The Los Alamos Molten Plutonium Reactor Experiment (LAMPRE, Los Alamos, 1962)

⁷ Achieved by passing an electrical current through 0.4% HCl as the working fluid.

Planning the Analysis

outer temperature. The experimental data includes resolved axial and radial profiles of velocity and temperature, suitable for comparison to CFD for validation.

The flow is non-dimensionalised by the S parameter:

$$S = \frac{g\beta r^3}{\nu^2} \frac{q_{vol} r^2}{k} Pr \frac{r}{L} = Ra \frac{r}{L}$$

where r , is the tube radius, and q_{vol} is the uniform heat source per unit volume. This is equivalent to a Rayleigh number multiplied by the radius to length ratio.

Watson (1971) provides further analysis and plots of the experimental data, but also shows that the agreement of theoretical solutions with some the data should not be relied on over a wide range of S . Investigations of this type of configuration and flow do not seem to have been continued in the open literature beyond this point, possibly because the design of nuclear reactors based on this configuration were not pursued. The estimated value of S for the forced and natural circulation conditions is shown in Table 2.4.

- For the natural circulation conditions, the values of $S \approx 1000$ are similar to the lowest values used by Murgatroyd and Watson (1970) and are expected to result in stable laminar flow.
- For the forced circulation conditions, the values of $S \approx 2 \times 10^4$ to 5×10^4 coincide with the range where much of the experimental data is present. At these higher values of S , some periodic instabilities or spiral motions around the tube axis were observed, and at $S = 1.7 \times 10^6$, the flow appeared to be turbulent.

From this assessment, it is expected that the fuel salt flow will be laminar, but that the full active length of the fuel pin is needed to correctly simulate the recirculation, and that a larger cross-section of assembly than a single coolant subchannel is necessary to allow any instabilities and secondary motions, such as a spiral to develop. It should be noted that the experimental data was collected using an aspect ratio of $r/L = 1/50$, whereas for the fuel pin $(d_{fuel}/2)/L = 1/340.4$. It is also the case that the temperature on the pin outer surface is not uniform, and in forced circulation conditions, the heat generation rate is not uniform with length.

2.2.5 Modelling Strategy

Based on the considerations discussed so far, an analysis scheme as shown in Figure 2.8 has been devised to demonstrate a process of deriving outputs that would contribute to enabling reactor scale simulations with more confidence.

This involves building parts of the solution approach and modelling separately and sequentially, allowing a fine-grained and incremental assessment of what works well and what each part adds. This is possible in the scope of this case study because the fuel assembly (considered in isolation) is geometrically simple with well-defined flows. For larger and more complex models (such as the reactor-scale model in Study C) attempting this approach from the outset would be substantially more costly and time consuming, with less guarantee of success, because it would be harder to anticipate which phenomena will matter most and will present the most difficult aspects for the solution. For such larger cases, the approach shown here could be applied to a number of selected components, once their importance was better established.

Planning the Analysis

Figure 2.8 shows the flow of information, understanding and decision making between the various models:

Precursor and auxiliary models which test the applicability of the available Reynolds-Averaged Navier-Stokes (RANS) models to vertical mixed convection flow (described in Section 3.1); provide validation and mesh resolution optimisation for an internally circulating heat generating fluid (Section 3.2), and explore the transition to turbulence in such circulating flows (Section 3.4). The first two were planned in advance of the simulation starting, and the last one was performed part-way through the process to understand the behaviour seen.

Sixth sector assembly resolving each fuel pin and including all of the physical processes simultaneously (Section 3.3).

Single coolant subchannel modelled over the full height of the assembly to evaluate the pressure loss and heat transfer characteristics of the bundle and the maximum fuel and cladding temperature (Section 3.5).

2D slice with no flow to evaluate the effective transverse conductivity of the assembly, including conduction in the coolant, cladding and fuel, and thermal radiation in the salts (Section 3.6).

Porous sixth sector assembly implementing the information derived in the subchannel and 2D slice models, and comparing the results of the porous model to the resolved model to assess the accuracy of the approximation (Section 4.1) .

Automation of the subchannel model to build surrogate models for the figures of merit that are not directly available from a porous model solution (Section 4.2). This technique is then extended to perform SA and UQ of the single subchannel to assess the sensitivity of chosen figures of merit to uncertainties in the salt material properties (Section 4.3).

The reasons why the models form this particular sequence depends partly on the learning gained and choices made from each, and is elaborated in detail in each section. It should be noted that they are not simply a sequential increase in complexity – the most complex model (resolved 1/6th assembly) is built relatively early in the process.

Turbulent flow is expected in the coolant, and a RANS turbulence modelling approach has been selected because it is a pragmatic and cost effective option, but some inaccuracies are expected. To provide a significantly more accurate answer, a resolved turbulence approach such as Large Eddy Simulation (LES) would be required, but could be orders of magnitude more costly to solve per model, given the size of the domain for the resolved 1/6th assembly. Also the option to use a single subchannel model would not be available because symmetry boundary conditions are not usable, and there is no natural periodicity available. The Reynolds numbers present in CH₄ mean that it is likely to be in laminar-turbulent transition, and so for simplicity, it will not be considered as part of this case study.

When implementing the porous medium models, it is expected that a non-equilibrium approach will be necessary, where the solid part of the porous section exists as an object in its own right, with thermal mass and a separate temperature field. This means that the coolant and fuel do not need to be assumed or forced to locally be at the same temperature, but it also requires the heat transfer characteristics between the coolant and cladding to be specified.

Planning the Analysis

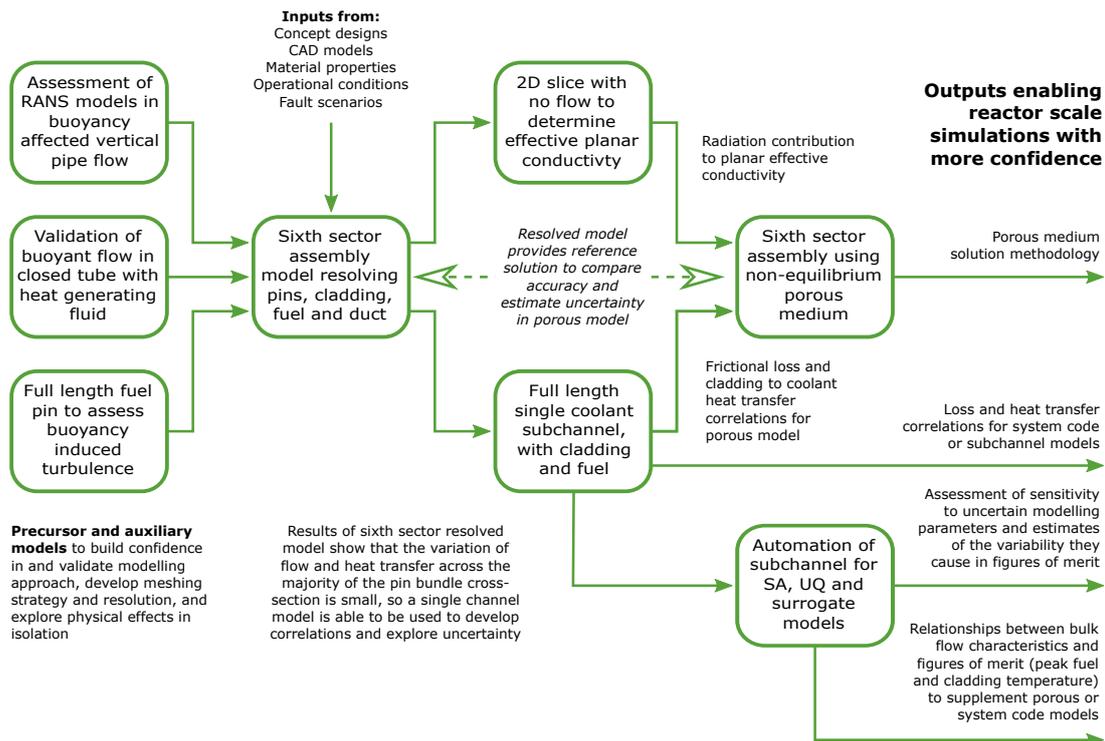


Figure 2.8: The sequence and flow of insight, decision making and information between the various models created as part of this case study, which together produce an overall modelling strategy.

2.2.5.1 Modelling Tool Selection

The ANSYS Fluent CFD code was chosen because it is available, and familiar, to the engineer performing the analysis, and has the key required functionality, that is able to be easily and reliably combined:

- The ability to obtain steady-state solutions in closed domains without open boundaries via SIMPLE based pressure-velocity solvers.
- Easy and reliable handling of CHT solids.
- Non-equilibrium porous medium models.
- Thermal radiation modelling of participating media.
- A text-based run configuration and solution interface that is well suited to scripting a staged convergence strategy and to automating solutions.

A RANS turbulence modelling approach has been chosen, but a remaining unknown at the planning stage is whether there are adequate turbulence models available in the selected CFD code. This will be assessed by a precursor model.

The open source Dakota toolkit, developed by Sandia National Laboratories will be used to implement the SA, UQ and surrogate model aspects. It can be interfaced to Fluent to automate runs without significant effort and contains all of the sampling and statistical functionality required.

Planning the Analysis

2.2.5.2 Quality Assurance and Validation

This study is intended to demonstrate the decision making and discovery process of starting from a design and choosing to applying CFD to characterise aspects of its behaviour. There are no sources of 'full-effect' validation data, but it is prudent to validate the key phenomena where possible, and perform studies to assess whether the results are sensitive to the effects of varying numerical modelling options.

In addition to this, the process involves making decisions and choices about what modelling approaches work well and should be pursued, and which should not. In a design and assessment setting, these choices may persist for a prolonged period and become part of the analysis 'custom and practice'. This introduces two possible latent conditions that may lead to inefficient or erroneous future work:

- The decisions apply to the fuel design and reactor conditions assessed at this initial stage, and to the capabilities of the computational tools available, all of which may change, altering the best options for modelling.
- Undetected mistakes may have been made at this stage because it was exploratory work that was not directly nuclear safety related, and these may lead to an incorrect or unsubstantiated conclusion being drawn.

Therefore, the insight and outputs derived from this work should be created on the basis of a rigorous and robust selection process, that can be revisited in future as necessary. To assure this, the geometry, mesh, CFD setup, solutions, post-processing and descriptions in this case study report have been subject to independent quality assurance verification. Further, the simulations were all performed in a scripted and automated manner, so that they can be easily re-evaluated or extended and compared to the existing results. This is particularly pertinent in this study because it includes results from many hundreds of individual CFD simulations, of a number of different types, which build on each other sequentially.

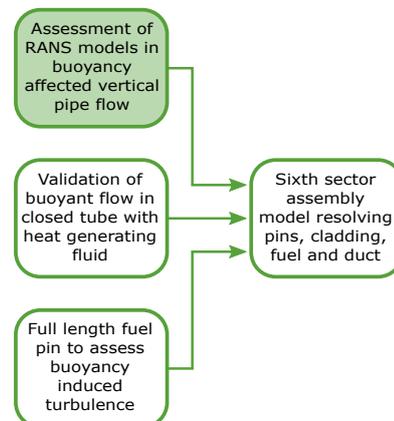
3 Performing the Analysis

The individual elements of the modelling strategy are described in the following sections, approximately following the sequence that they were planned and executed in (although there are several feedback and iteration processes that have taken part along the way). The conclusions from each step and the decision processes for what it meant for the subsequent stages are described in each step, with the modelling used to determine overall comparisons, conclusions and recommendations discussed in Section 4.

The precursor and auxiliary models are not described in the same level of detail as the later fuel assembly models, but the details of their configuration can be found in the references discussed.

3.1 Buoyancy Affected Vertical Channel Flow

The presence of relatively low Reynolds number flows and buoyancy forces in the coolant channels raised the question about whether the turbulence models available in Fluent are able to accurately represent these effects. The work of Keshmiri *et al.* (2012) suggests that some RANS models are able to predict the reduction in Nusselt number seen at intermediate values of Bo for buoyancy assisted flow, because they capture boundary layer relaminarisation, and that some models do not predict this. The models shown to work are not available in Fluent, but in discussion with ANSYS, it was suggested that their $k-\omega$ SST γ (intermittency) transition model may be able to.



This model was not designed for internal flows, however, instead it was intended to capture transition and relaminarisation in aerodynamic applications (Menter *et al.*, 2015). This section shows the results of testing this model and also assessing the performance of more common models.

A 0.01 m long axially periodic section of a 0.1 m diameter pipe was represented as a 10° wedge with symmetry boundary conditions, and meshed to achieve $y^+ < 1$ in the first cell (Figure 3.1). The flow rate was prescribed to give a Reynolds number of 5300 and the wall heat flux is modified to create a range of values of Bo . The fluid used was air, applying the Boussinesq approximation with a fixed density of 1.205 kg/m^3 and constant properties giving $Pr = 0.71$. The details of the setup are given by Keshmiri (2010), along with comparison data from experiments, Direct Numerical Simulation (DNS), LES and successful RANS results from other CFD codes.

Performing the Analysis

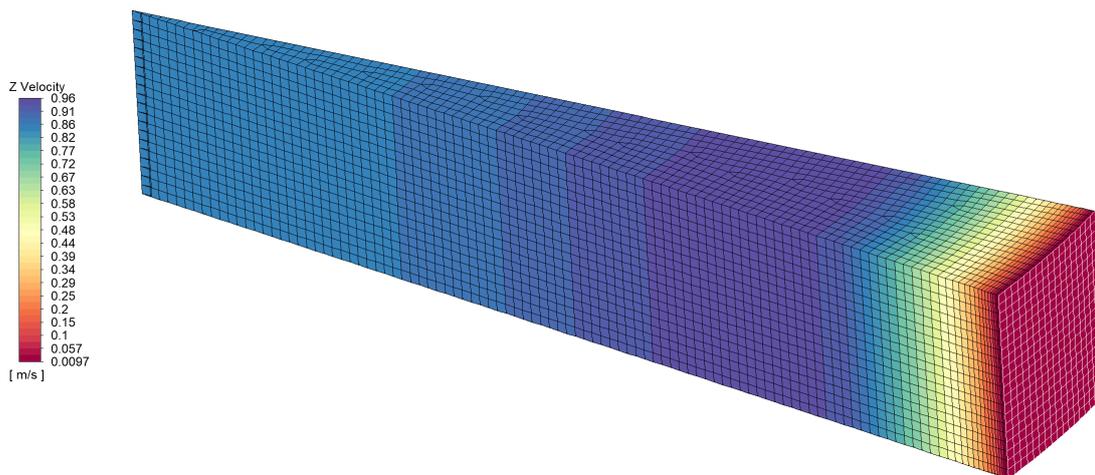


Figure 3.1: Mesh coloured by axial (z) velocity for $Bo = 0.25$, for buoyancy assisted flow using the intermittency model.

Cases where the flow is assisted by buoyancy ('ascending' buoyancy in Keshmiri *et al.*, 2012) show an impairment to heat transfer at moderate values of buoyancy (approximately $Bo = 0.2$). Buoyancy opposed flow ('descending') always show enhancement to heat transfer from the action of buoyancy forces.

A large number of cases were used to assess the performance of the models and to evaluate the influence of initial conditions, only a subset of which are reported here in Figure 3.2. Solutions were converged for values of Bo from 0.02 to 100¹, with values clustered around $Bo = 0.2$. The Nusselt number, Nu , and Darcy friction factor, f , for the flow was calculated for each result. An additional simulation was run first for each setup configuration without gravity to provide a comparison without the effect of buoyancy, giving the prediction in the $Bo = 0$ limit (forced convection). These values are referred to as Nu_0 and f_0 , and are used to normalise the results (Nu/Nu_0 and f/f_0). The results are also compared to the Gnielinski forced convection correlation and the Moody chart, via the Haaland correlation (Volume 6, Section 3.1) – these values are shown in the figure titles. The cases reported are:

k- ω SST assisted/opposed: the *k*- ω SST model was used, with results for buoyancy assisted and buoyancy opposed conditions.

k- ω SST all options assisted/opposed: the *k*- ω SST model with the 'low Reynolds number corrections', 'full' buoyancy terms and γ intermittency enabled.

k- ω SST γ assisted/opposed: adding only the intermittency option, using a converged *k*- ω SST solution first for each case.

k- ω SST γ (0 initial) assisted: the intermittency option added at the start of the solution with the γ field set to zero (laminar) everywhere.

k- ω SST low *Re* correction opposed: the *k*- ω SST model with the 'low Reynolds number corrections' enabled.

realizable *k*- ϵ buoyancy assisted/opposed: the realizable *k*- ϵ with 'full' buoyancy terms enabled.

The '*k*- ω SST γ (0 initial) assisted' simulation produced the same result as the γ model results

¹ Only from 0.02 to 2 for some configurations, because the very high heat flux cases were of less relevance.

Performing the Analysis

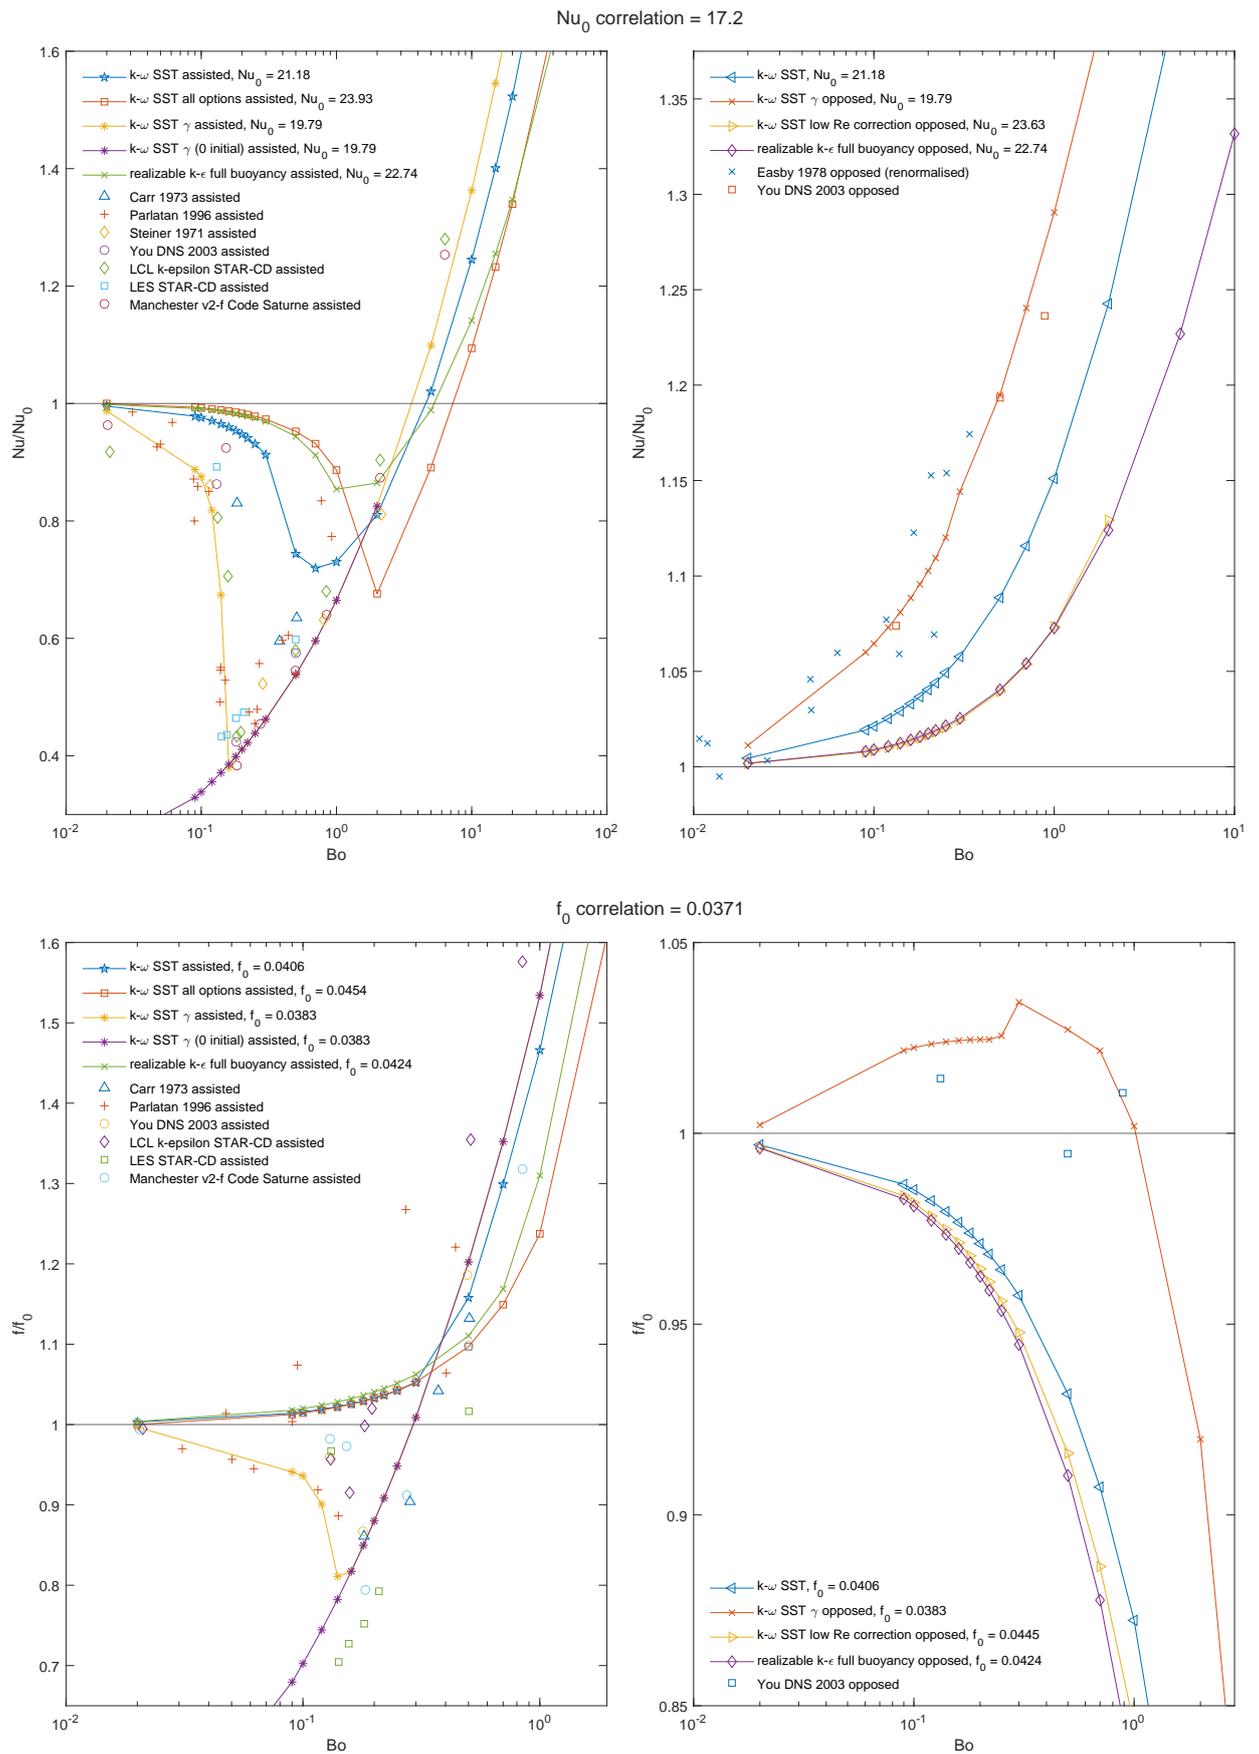
that are started from a converged $k-\omega$ SST result for $Bo \geq 0.16$. However, at lower values of Bo , the simulations that were initialised with zero γ produce values of Nu and f significantly lower than Nu_0 and f_0 because they erroneously predict laminar behaviour across a large portion of the cross-section of the channel. This indicates that the γ model is unable to correctly generate its own turbulence from some starting points. The same effect is seen if the solution is initialised with a uniform $\gamma = 1$, and is also observed when using an approach where the solution at one value of Bo is started from the converged result from the next higher value of Bo (the solutions being run in sequence without being reinitialised each time). This is interpreted as indicating that the solution produced by the γ model includes an element of hysteresis, or solution trajectory dependence, for whether or not it predicts the correct boundary region laminarisation.

From the results of all of the cases, it can be seen that.

- All of the RANS models tested over-predict Nu_0 and f_0 compared to the correlations for forced convection, although the relatively low Re flow may make this comparison more challenging. The γ model results are close to the experimental, DNS and LES results reported by Keshmiri (2010).
- The $k-\omega$ SST γ intermittency model (with default parameters) is the only case plotted that performs well for the buoyancy assisted and opposed flows for Nu and f . For buoyancy assisted flows it predicts the relaminarisation and heat transfer impairment behaviour (reduction in Nu/Nu_0 to ≈ 0.4 at around $Bo = 0.2$, with corresponding reduction in f/f_0) to the same extent as the other successful RANS and LES results (from Keshmiri, 2010) that are plotted for comparison with the experimental data. It does appear to be sensitive to solution strategy, however. It also produces good predictions of the enhancement to Nu and the behaviour of f for buoyancy opposed flow, compared to the available experimental and DNS results.
- The performance of the $k-\omega$ SST model is in-line with the behaviour found by Keshmiri *et al.* (2012), where heat transfer impairment was also found to be not well predicted.
- The complexity of the numerical implementation of the models, and of the physics of buoyancy generated relaminarisation behaviour makes it hard to tell (without significantly more detailed investigations) if the ‘hysteresis’ effect is present for true physical reasons, or is purely a modelling artefact. For these relatively low Reynolds numbers, delayed transition is known to occur in the absence of disturbances (Eckhardt *et al.*, 2007, Mullin and Peixinho, 2006) and has been shown to be reluctant to transition in DNS simulations (Wu *et al.*, 2015).
- Starting from a converged $k-\omega$ SST solution before enabling γ appears to be the most reliable and general approach for using the intermittency model in this case.
- The ‘low Reynolds number corrections’ to the $k-\omega$ SST model worsen its performance at resolving the heat transfer impairment and increase the values of Nu_0 . When combining this option with the γ model, it is predominately the ‘low Reynolds number corrections’ behaviour that is observed.
- Adding the ‘full’ buoyancy terms to the $k-\omega$ SST or realizable $k-\epsilon$ models has little effect.

From Table 2.4, buoyancy opposed flow in the coolant channel at $Bo < 0.035$ is of most relevance to this case study, although for flow at lower Reynolds numbers ($Re < 3000$). Details of the propensity for and extent of laminarisation in the rod array under these conditions is not well known, and

Performing the Analysis



Performing the Analysis

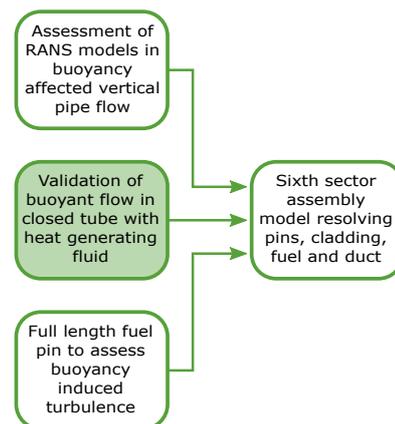
cannot be inferred reliably from these vertical tube results. The turbulent coolant flow in the fuel assembly that is being considered is all buoyancy opposed, which should promote turbulence, and so some enhancement of Nu is expected, but only by a small amount. The expected effect on f is not as clear, but also not substantial.

Despite the promising performance of the γ intermittency model, it will not be used in the following simulations. This is because, there is a risk of the solutions exhibiting the ‘hysteresis’ effect found, because low levels of intermittency may be transported from laminarised regions into higher Re sections, where they could cause problems with creating or sustaining the correct state of turbulence. This would be possible to examine and explore for a small number of simulations, but because these models are intended to be run hundreds of times, being sure of the correctness and robustness of this behaviour in each would require assessment of each, adding a potentially significant burden.

3.2 Closed Vertical Tubes With a Heat Generating Fluid

One of the aspects of the fuel assembly that is unfamiliar and uncertain is the closed recirculating fuel salt with internal heat generation. It is not clear whether there will be a stable flow within the tube and whether there will be one full-length recirculation, or a number of convection cells, as found in tall parallel plate cavities (Jin and Chen, 1996). Bergholz (1980) suggests that there are not expected to be instabilities that cause a cellular structure to occur, and a stable upward flow at the centre of each pin is likely.

The experiments and analysis of Murgatroyd and Watson (1970) and Watson (1971) provide spatially resolved validation data, insight into the expected phenomena, and a straightforward case to test solution methods against.



While the design feature of molten fuel in long closed tubes is specific to the Moltex design, the requirement to simulate buoyancy driven recirculation of variable density fluids in a closed space, or regions of high heat load but low circulation, is a commonly encountered situation. The process shown here of finding validation data, accompanied by dimensional analysis, developing meshing approaches and testing the necessary resolution, and proving the numerical solution route are a transferable pattern and thought process.

Another aspect that required testing in a simple case is the simulation of a closed domain containing a fluid with variable density. In reality the fuel has a gas space above it, but it would be unnecessarily complex to try to simulate this gas space explicitly using a multiphase approach. Instead, a pragmatic approximation is to apply a zero-shear wall (slip) velocity field boundary condition to the top of the domain to represent the free surface. Doing this creates a closed domain. If the Boussinesq approximation is used then the density of the fluid is fixed for the purposes of the continuity equation, and the mass in the domain is fixed. However, the density of fuel salt depends on temperature, and the Boussinesq approximation is not applicable to these cases with large

Performing the Analysis

changes in temperature over the tube cross-section, and with axial distance. In this situation, the mass of salt in the fluid must change as the temperature solution converges, and there is no link to the pressure of the domain in the equation of state (for liquids which are nearly incompressible, this is generally a good approximation). For closed spaces containing gases, this could be accounted for by conserving mass and allowing the pressure in the space to solve to the physically correct value.

The geometry and mesh is shown in Figure 3.3. The mesh shown contains approximately 100,000 cells, and was optimised in wall normal and axial dimensions over four loops of iterative testing to give results almost indistinguishable from a 770,000 cell mesh (mainly achieved by stretching the axial cells to 20 mm in length and grading them towards the top and bottom). This mesh contained a small 'pin-hole' aperture/vent in the top surface that could be made either a zero-shear wall (to close the domain) or a pressure boundary, able to allow mass flow in either direction.

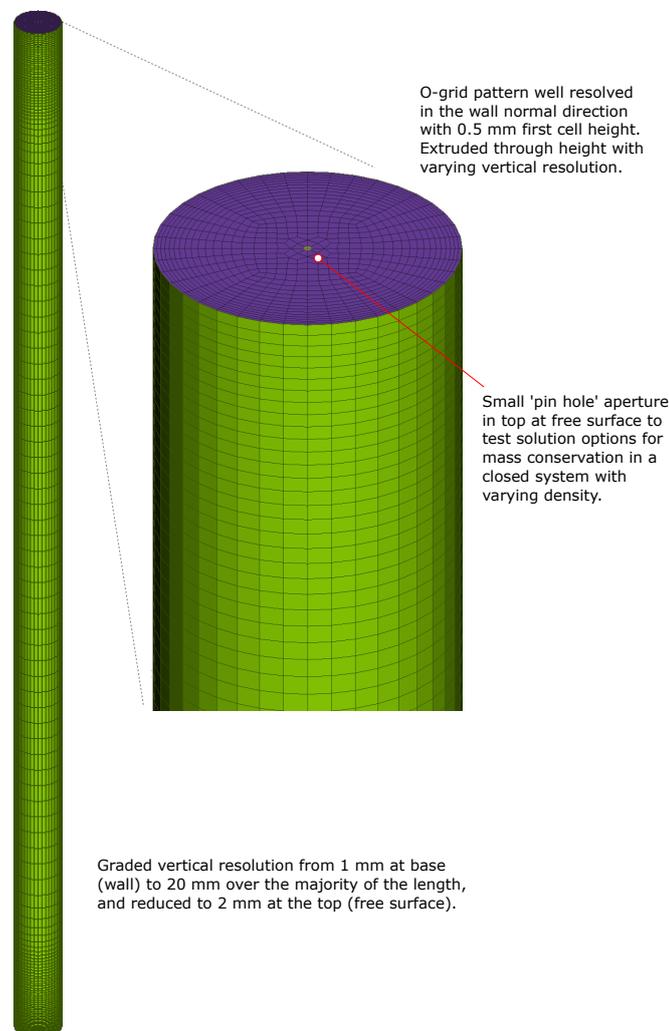


Figure 3.3: Geometry to compare to experimental data from Murgatroyd and Watson (1970) and Watson (1971) (length $L = 1270$ mm, radius $r = 25.4$ mm, $r/L = 1/50$).

The details of the experimental external boundary conditions, heat sources or temperatures used are not given precisely by Murgatroyd and Watson (1970) or Watson (1971). They are either omitted or presented in non-dimensional form only. For a given reference temperature, the volumet-

Performing the Analysis

ric heat source, q_{vol} required for a given value of S can be found using the expression in Section 2.2.4.2 by evaluating the material properties. However, the temperature to evaluate them at is not stated.

The best information is that the sources refer to the Pr value for water in the range of temperatures used as being $3 < Pr < 8$ which gives a range of temperatures of approximately 15 to 50 °C. It is also not clear whether it is the fluid temperature or the reference (wall) temperature that is used to evaluate the properties, and the source of the water properties is not stated, and hence no assessment can be made of their accuracy.

The experiment is designed to provide a close approximation of a constant temperature on the outer surface of the perspex tube containing the fluid, and so an external convection boundary condition was applied using a Heat Transfer Coefficient (HTC) of 100 W/m² K and an external temperature of 300 K. 305 K was used as a representative fluid temperature to determine the fluid properties and heat source. These were sourced from the IAPWS-IF97² (Wagner *et al.*, 2000) and are shown in Table 3.1.

	293.15 K	305 K
ρ (kg/m ³)	998.21	995.08
k (W/m K)	0.5995	0.6177
μ (kg/m s)	0.001002	0.0007668
c_p (J/kg K)	4184.8	4179.5
Pr	6.991	5.189
β (1/K)	0.0002069	0.0003191
q_{vol} (W/m ³) for $S = 8900$	1790.8	950.3
q_{vol} (W/m ³) for $S = 84300$	16963	9001.3

Table 3.1: Water properties and required volumetric heat sources for two values of S .

The necessary values of q_{vol} for two example values of S used for the experiments are also included. It can be seen that reducing the reference temperature to 293.15 K results in a difference in the heat source terms of a factor of 1.9, driven mainly by the changes in $1/\nu^2$ and β . This significant sensitivity to operating temperature introduces a large uncertainty in the comparison of the CFD results to the experiment for a given value of S .

The case was setup in Fluent, as the intended CFD code, but also in CFX because it uses a different numerical approach to the solution, and also to test the use of closed domains in CFX, which is relevant to Study C. Three setup options were explored:

- The water density represented by the Boussinesq approximation, with a fixed value of density evaluated at 305 K, or a variable density implemented as a polynomial with a dependence on temperature only.
- The small aperture was either open (able to exchange mass in either direction) or closed (a zero-shear wall as part of the free surface).
- Pressure-velocity coupling using either an algorithm based on a segregated solution using the SIMPLE algorithm, or a coupled solution. The SIMPLE algorithm is not available in CFX.

² Using an open source MATLAB implementation available from github.com/mikofski/IAPWS_IF97.

Performing the Analysis

All solutions were configured to run in steady-state mode – for SIMPLE this is a true steady-state solution, for coupled solutions, the numerical time step is manipulated to reach steady-state.

The solutions attempted are shown in Table 3.2, which were all run with laminar settings.

- Using the Boussinesq approximation, the cases solve well in both codes, and no opening is needed. These also form good starting conditions for variable density cases. In Fluent when allowing a variable density, running with a SIMPLE based solver was successful with and without using an opening – the changes in system mass were absorbed by the algorithm when enforcing continuity. In Fluent when using variable density, a converged steady-state solution was not obtained using the coupled solver with the settings tested³, whether including an opening or not.
- In CFX, for variable density without an opening, the solution sometimes converged, but the velocity field was occasionally subject to sudden large excursions.
- In CFX, for variable density with an opening, the solution converged, but chose the wrong numerical timescales by default, so converged very slowly. A manual timescale estimate was necessary (10 to 100 times bigger) to obtain rapid convergence. The Boussinesq solution estimated the timescale correctly, and so was a good guide for what to apply.

Density	Aperture	$P-U$ coupling	Fluent	CFX
Boussinesq	Closed	SIMPLE	Stable	
Boussinesq	Closed	Coupled		Stable
Variable	Closed	SIMPLE	Stable	
Variable	Open	SIMPLE	Stable	
Variable	Closed	Coupled	Unstable	Unstable
Variable	Open	Coupled	Unstable	Stable

Table 3.2: Solution attempts for closed domain.

Monitoring the mass of fluid in the domain (and the mass flow rates across the opening, where present), the minimum and maximum vertical velocities, maximum temperature, and wall heat flux were used as guides to assess convergence and energy conservation, in addition to assessing the solution residuals.

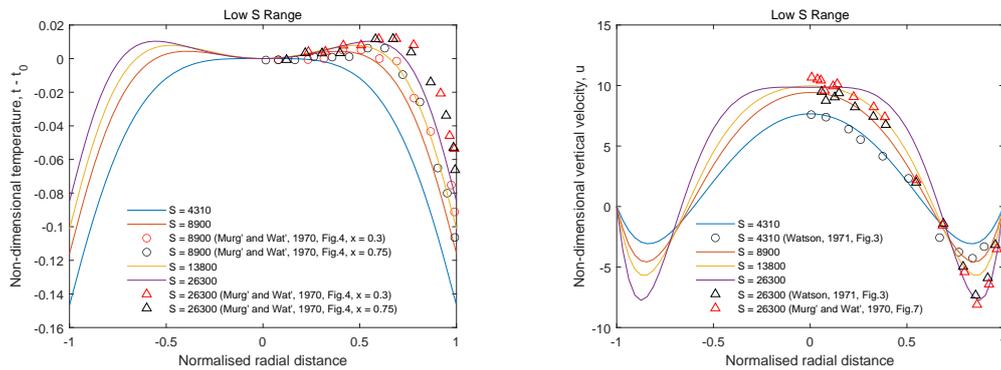
To obtain any converged solutions with a variable density, it was found that a good estimate of operating density was necessary. The operating density, ρ_{op} redefines the pressure field for the solution to be $P' = P - \rho_{op}gz$, where z is the co-ordinate in the direction of gravity. This removes a reference hydrostatic contribution⁴ from the pressure field for numerical reasons. It improves convergence by separating the small pressure differences generated by buoyancy forces and flow from the large hydrostatic variation⁵. A value of $\rho_{op} = 995.08 \text{ kg/m}^3$ was used.

³ Including with or without the 'pseudo-transient' setting.

⁴ It cannot be said to remove *the* hydrostatic contribution unless the density is constant. When density varies, the hydrostatic pressure is the result of the integral of the density field above a point.

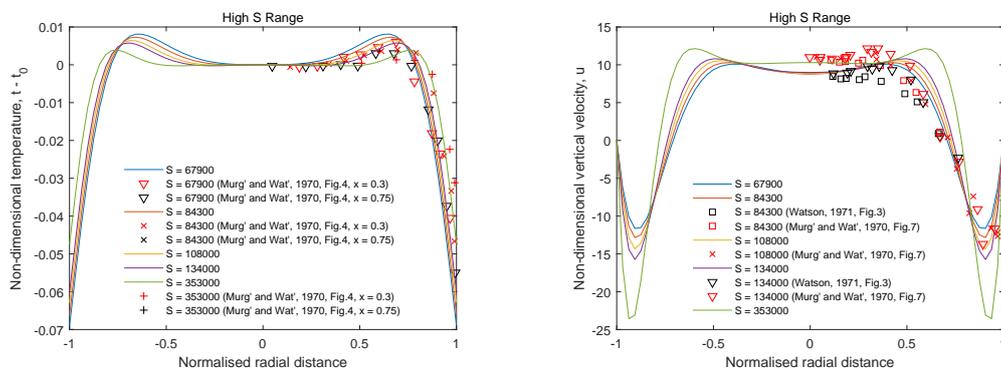
⁵ This is true even when using a double precision solution. All simulations reported in this case study use double precision, and it should be the default for all CFD solutions, except in specific situations of relatively simple physics, such as incompressible external aerodynamics, and where the usable mesh sizes are limited by available memory.

Performing the Analysis



a: Non-dimensional temperature, low S range.

b: Non-dimensional vertical velocity, low S range.



c: Non-dimensional temperature, high S range.

d: Non-dimensional vertical velocity, high S range.

Figure 3.4: Comparisons of CFD predictions of non-dimensionalised vertical velocity and temperature to the available experimental data (symbols).

The solutions produced by the two CFD codes were very similar, where they could both be obtained. The comparisons between codes and the assessments of the iteration of mesh resolution were made by comparing the details of the axial variation of centreline temperature and vertical velocity.

This exploration shows that using Fluent with a steady-state SIMPLE solver, using an operating density and starting a simulation from a Boussinesq condition before switching to a variable density fluid provides a reliable solution, without needing an open aperture. CFD results produced using this approach were compared to the experimental results combined from both Murgatroyd and Watson (1970) and Watson (1971). They were compared using radial profiles of non-dimensionalised axial (vertical) velocity and temperature at mid-height ($x = 0.5$, where x is the non-dimensional axial distance). The non-dimensional quantities are defined as

$$u = \frac{Ur^2}{\alpha L}$$

and

$$t = \frac{k(T - T_r)}{q_{vol}r^2}$$

where T_r is the local inner wall temperature at that height. The results are shown in Figure 3.4,

Performing the Analysis

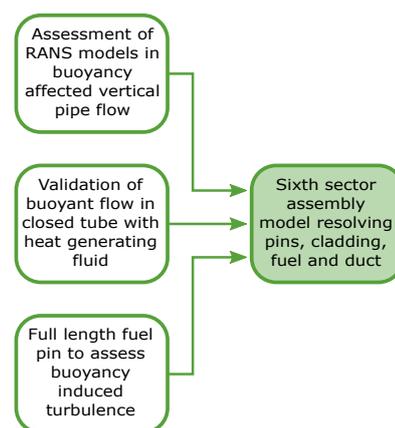
where the comparisons are split into those at high and low S for clarity. The t values are plotted relative to their centreline value to make the result coincide at $r = 0$ to emphasise the differences in shape. The agreement is convincing, showing the peaking effect in the temperature field that coincides with the radius where zero velocity occurs. More detailed numerical comparison of the level of agreement has not been pursued because of the sensitivity of S for a given heat load to the operating temperature.

Aside from demonstrating an initial validation comparison, giving confidence that CFD is able to predict these laminar flows well, performing this precursor simulation was a useful step for a number of reasons. It allowed a solution and monitoring strategy to be developed, and practised with a simple and small model, which could then be applied to the resolved fuel assembly case. It also provided guidance on the appropriate mesh resolution for the resolved model, and gave confidence that an aperture does not need to be included, simplifying the meshing and setup process of the larger model.

3.3 Resolved Rod Assembly

The largest and most complex model solved in the case study was the 1/6th sector resolved geometry, containing all of the components and physics simultaneously.

For brevity, only one result for forced circulation conditions (in CH2) and one for the natural circulation conditions (at 250 s in CH2) were simulated (these will be referred to as the 'baseline' conditions), although the methods described below would allow any or all of the flow conditions to be evaluated.



3.3.1 Geometry and Mesh

The mesh for the resolved mesh was created using ANSA. This tool was selected because it is powerful and flexible for producing quad dominated patterns on surfaces, and then using them to create extruded meshes with axial resolution control. It also has the facility to 'link' faces in the geometry, so that the cross-section for a single fuel pin, coolant subchannel and piece of side symmetry could be created and then replicated. Changing the mesh on any one of these linked faces changed it on all of them automatically, allowing modifications to the mesh to be made rapidly and reliably.

The mesh was built from an analytical description of the geometry, rather than retaining any of the CAD model geometry itself. The CAD model was imported and overlaid to check that the major features were the correct size and in the correct place. The locations for the repeating units for linked faces are shown in Figure 3.5. These points were generated analytically and imported into ANSA as a text file of positions to minimise manual transcription errors.

The mesh pattern chosen for the cross-section is shown in Figure 3.6. The fuel pins were meshed with an adapted O-grid pattern based on six segments (to allow the centre and symmetry plane pins to be identical to those in the main assembly). The mainly quad face pattern was allowed to

Performing the Analysis

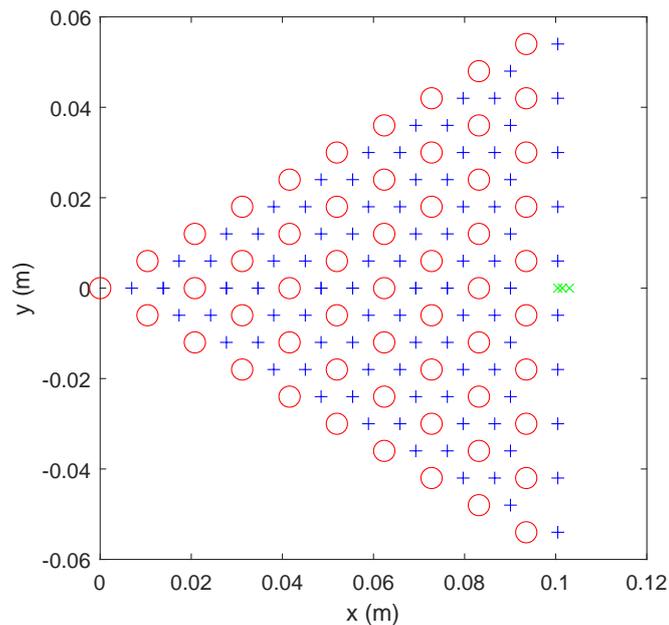


Figure 3.5: Locations of the centres of fuel pins (red) and subchannels (blue) to provide replication points for linked units of mesh, and the location of the duct inner and outer surface, and centre of the inter-wrapper gap (green).

introduce triangle faces to reduce the circumferential number of cells with reducing radius. This approach is also used in the cladding and boundary layer region of the coolant. Without this approach, and using quad faces only, an excessive number of converging radial lines would be present at the centre of each fuel pin. The cladding and duct are included explicitly as solid regions and so will be solved for using a CHT approach.

This profile was extruded over the 1600 mm active length of the fuel. 100 mm was added at the top to allow an entrance region for the coolant, and 10 mm was added at the bottom, representing the gap between the bottom of the fuel salt and the pin support plate (Figure 2.4). The resulting mesh is shown in Figure 3.7, where the axial and planar resolutions were guided by the optimised mesh from Section 3.2. The axial resolution expands to 4 mm over the first 100 mm from each end of the active length, and then to 20 mm over the next 100 mm, so is constant at 20 mm over the middle 1200 mm of active length.

Other key cell size are:

- 0.03 mm first cell height (normal to surface) in the coolant at the duct and cladding surface. The cells expand in layers, and have a target size of 0.16 mm in the middle parts of the coolant gap between fuel pins. An estimate from the expected flow rates indicated that a first cell height of 0.06 mm would be necessary for $y^+ = 1$ in the coolant, so the near-wall flow is expected to be well resolved.
- 0.09 mm first cell height in the fuel salt normal to surface.
- 0.3 mm circumferential cell size on fuel pin surface.
- 4 cells were used through the cladding thickness.

The resulting mesh contained 44 million cells. This approach to meshing results in cells that are

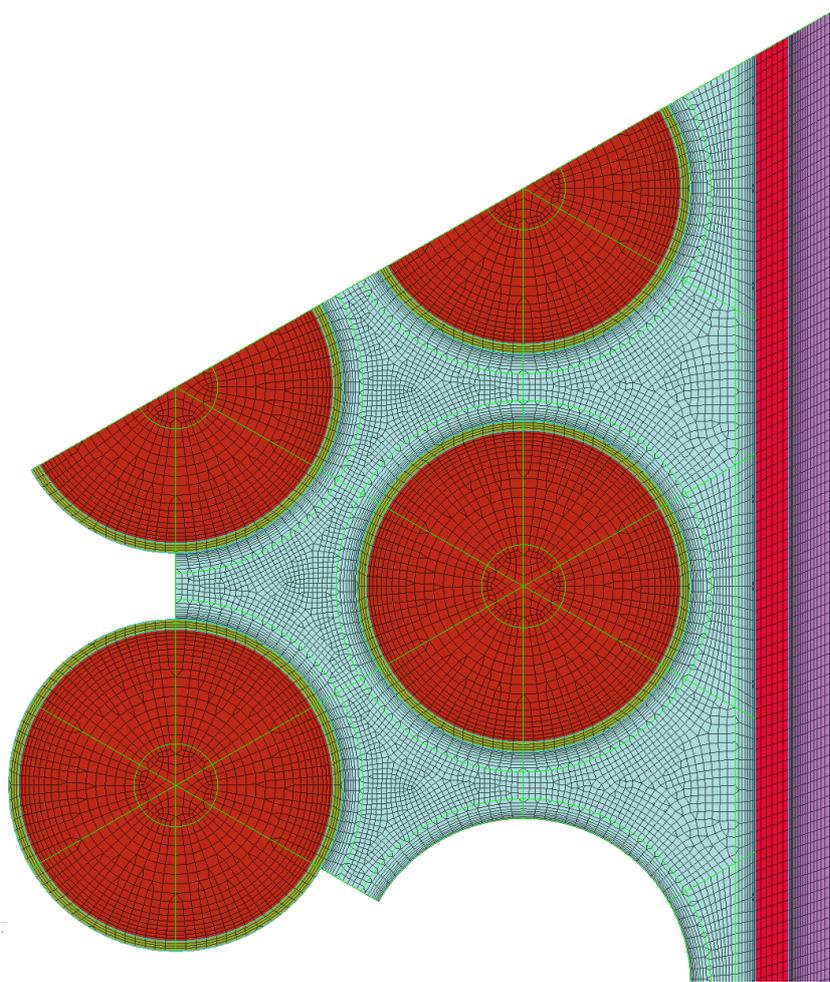


Figure 3.6: Mesh cross-sections of fuel pin, coolant subchannels and duct region for the 1/6th resolved fuel assembly. Only showing a subset of the mesh at the corner of the domain.

almost entirely hexahedra, with a small number of triangular prisms and no pyramids, tetrahedra or polyhedra. The resulting mesh had a worst cell quality (using Fluent's metrics, ANSYS, 2020b) of 0.846 equivolume skew (where values > 0.95 indicate poor quality, that would badly compromise local accuracy or overall solution stability) and 0.231 orthogonal quality (where numbers close to zero indicate poor quality, and < 0.01 would lead to solution stability difficulties). Because they are mainly flow-aligned hexahedra, most cells in the mesh are of significantly higher quality than these worst values.

The high quality cell shapes give confidence that the solution should not exhibit divergences for mesh quality reasons. High aspect ratio cells are present, and not just in boundary layers. This is due to the stretched 20 mm axial resolution necessary to limit the cell count. Because the cells are stretched in the flow direction, and the axial variation in velocity and temperature is slow and smooth, this does not cause problems for converged, stable, unidirectional flow solutions. If the flow was subject to recirculation in these regions, these cells could be problematic for the solution.

A non-conformal interface was applied between the cladding and the fuel salt over 100 mm at either end of the fuel pin – without it, the fine axial resolution of the fuel would need to be present

Performing the Analysis

in the mesh for the coolant channel. Using the non-conformal interface in this way reduces cell count and avoids unnecessary changes in refinement in the flow direction in the coolant, whilst allowing the flows at the ends of the fuel salt to be sufficiently resolved. The local axial temperature gradients present in the fuel salt at the ends of the pins, in the sections covered by the non-conformal interface, are not sufficiently high that the coarser axial solid (and coolant) would be unable to resolve them.

3.3.2 Case Setup

Materials were created with polynomial dependences on temperature as described in Section 2.2.3.3 and applied to the fuel salt, coolant salt, cladding and duct.

Inlet and Outlet: For the forced circulation cases, the coolant inlet is at the top of the domain; for the natural circulation cases the inlet is at the base. The flow was introduced with a uniform coolant velocity and temperature at the values given in Table 2.2. Uniform inlet conditions are an approximation:

- For the forced circulation flow, the inlet is $100/5.88 = 17$ hydraulic diameters from the top of the active section containing fuel, so a substantially well developed velocity profile is expected to exist⁶.
- For the natural circulation flow, the coolant must pass through the holes in the support plate only 10 mm below the active section, and so will not be a developed profile⁷.

The inlet turbulence values were obtained from a pipe flow correlation based on the inlet Reynolds number and hydraulic diameter (ANSYS, 2020b). The outlet for the coolant is set to zero gauge pressure. For simplicity, the same inlet and outlet boundary conditions were applied to the inter-wrapper gap as to the coolant in the fuel pin bundle. The bundle and gap fluid domains are separated by the duct and do not exchange flow or communicate pressure.

Outlet boundary conditions usually require temperature and turbulence values to be specified for flow introduced by any backflow. However, this case is a unidirectional channel flow in the coolant, so no backflow could occur physically. As a solution convergence improvement, the outlets were set to create wall faces to prevent backflow if, at an intermediate point during the solution process, the pressure field was such as to generate any reversal.

Heat Source: The heat source given in Table 2.2 was applied to the fuel salt regions of the domain, with the spatial distribution shown in Figure 2.6 for the forced circulation conditions implemented as an axial profile.

⁶ Longer entrance lengths would be needed to produce fully developed velocity profiles, but these would be longer than the available upstream distances in the actual fuel assembly. The cladding outer surface over the entrance length is also approximated as being adiabatic, but the coolant environment above the top of the fuel salt would create some heat transfer, so there would also be a partially developing thermal profile present in reality.

⁷ The design of the holes in the support plate is schematic only at this stage in the reactor development, and so detailed consideration of flow through them is not warranted.

Performing the Analysis

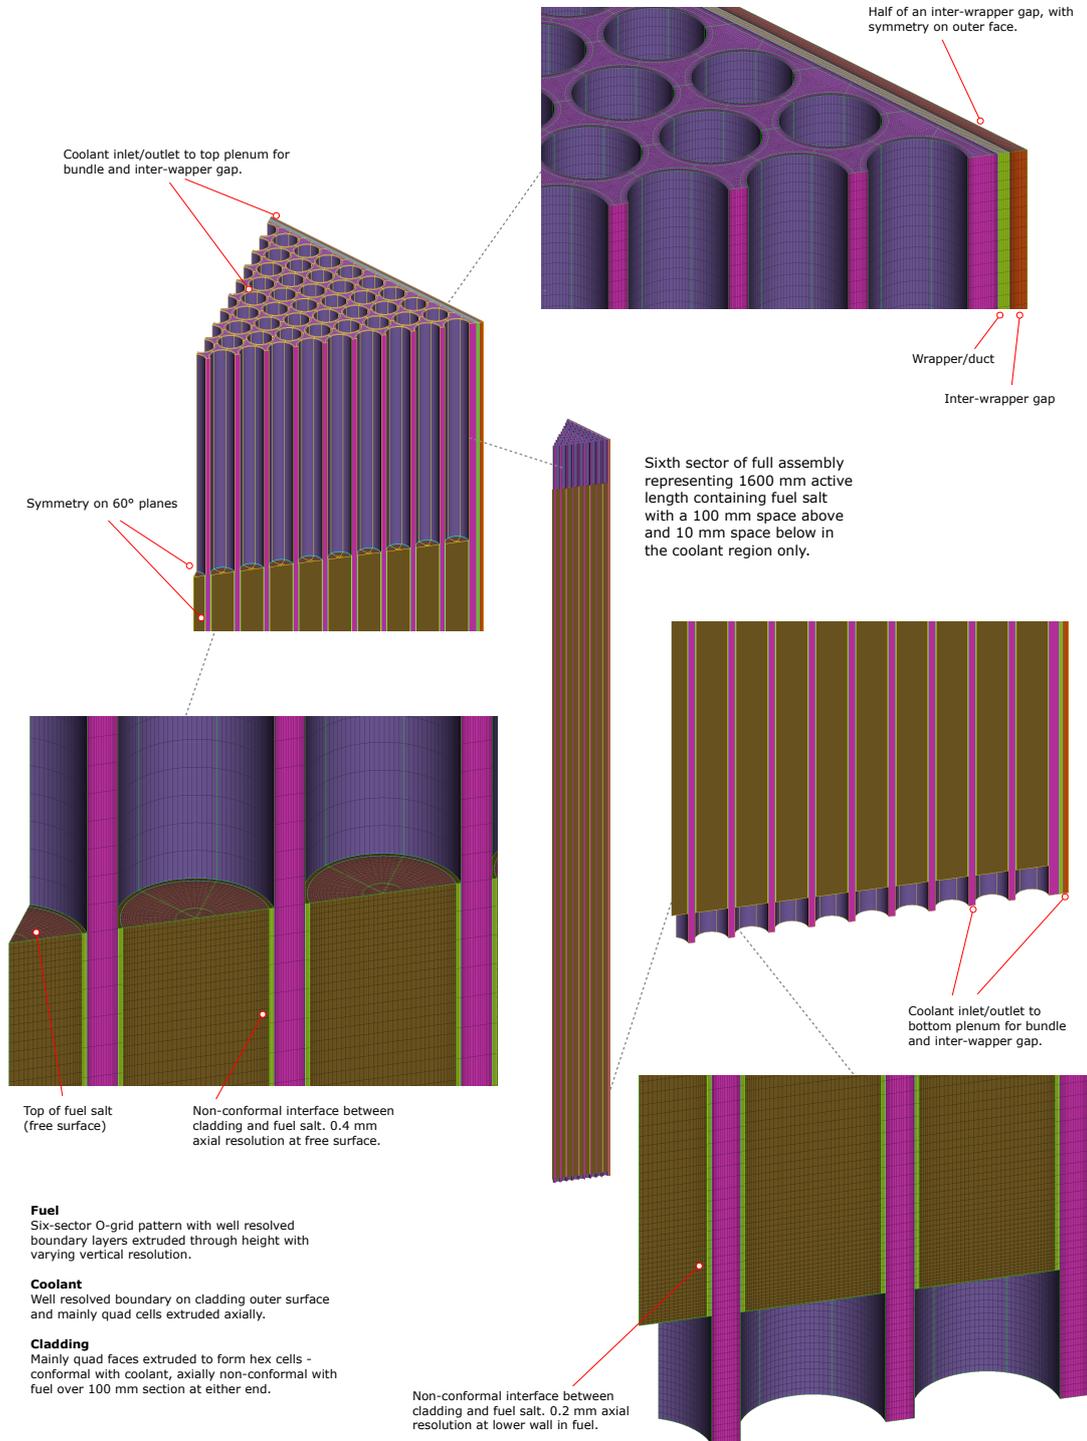


Figure 3.7: Geometry and mesh of the 1/6th resolved fuel assembly.

Performing the Analysis

Operating Conditions: The coolant outlet pressures are relative to a reference pressure of 2 bar applicable in the domain coolant. The fuel pins are closed, and have their pressure referenced to the same value of reference pressure at a fixed location at the top of the fuel pins. However, the absolute pressure in the domain has no effect on the solution, because the density of the fluid depends only on temperature, and none of the other material properties depend on it.

The operating density was set to be the fuel salt density at a temperature estimated to be representative of the final solved value. For the forced circulation case, this was 3126 kg/m^3 representing 1200 K, and for the natural circulation case it was 3430.5 kg/m^3 representing 900 K.

Turbulence: The $k-\omega$ SST RANS model was applied, and the low Reynolds number corrections were activated. This was done, despite the findings of Section 3.1, because a stable solution proved difficult to obtain without them. The choice of this model is discussed further below.

Participating Radiation: The Discrete Ordinates (DO) radiation model was chosen to represent the salts as participating media, and the nominal material and surface properties described in Section 2.2.3.3 were applied. The radiation models typically available in CFD codes are described in Volume 2 (Section 3.4.3.2), and the considerations necessary for participating media are described in Volume 6. DO is the most suitable model because it can accommodate any optical thickness, and can be easily extended to represent a banded spectral (wavelength dependent) absorption coefficient.

It was found to be necessary to have a minimum level of angular discretisation and pixelation (ANSYS, 2020a, Murthy and Mathur, 1998) to avoid spurious results being caused in the temperature field at the 60° symmetry planes. It was necessary to use 3×3 angular discretisation per octant of angular space, and 5×5 pixelation of the control angles.

3.3.3 Obtaining a Converged Solution

Obtaining a stable solution for this case was not straightforward. It was found to be necessary to carefully control the numerical settings, turbulence modelling and solution initialisation, as well as adding some aspects of the solution incrementally. At each stage, if any of these conditions were not appropriate, the solution residuals would increase, the coolant outlets would attempt to backflow over a significant fraction ($> 50\%$) of their area, introducing wall faces, and eventually the solution would diverge, resulting in a floating point error. The stages of the solution strategy described below are summarised in Tables 3.4 and 3.5.

Initialisation: The solution initialisation is summarised in Table 3.3. In the forced circulation case, it was found to be advantageous to provide an initial vertical velocity, and in both forced and natural circulation cases it was found to be beneficial to provide domain specific initial guesses for temperatures to the cladding and fuel salt. For the turbulence fields in the coolant in both cases, the inlet value estimates were applied for the k and ω fields. For the fuel salt in the forced circulation case, the turbulence in the coolant was too high, and the same low turbulence values as the natural circulation case were applied. Zero gauge pressure was applied everywhere.

Performing the Analysis

	Forced circulation		Natural circulation	
	Coolant	Fuel	Coolant	Fuel
P (N/m ²)		0		0
U (m/s)	(0 0 -0.05)	(0 0 0)	(0 0 0)	
k and ω	As inlet	As natural	As inlet	
T (K)	Coolant As inlet	Fuel/Clad 1200	Coolant As inlet	Fuel/Clad 890/850

Table 3.3: Initialisation settings for forced and natural cases.

Source and Buoyancy: It was found necessary to obtain a relatively well converged solution for the flow and temperature fields in the coolant, to provide a stable boundary condition to the cladding outer surface, before attempting to solve for the flow in the fuel salt. This is achieved by starting the case with gravity off, and adding the heat source to the fuel incrementally, starting at 25% for an initial period to allow the flow to develop before stepping up to 100% and converging the temperatures. While gravity is off, there is no circulation in the fuel salt, so energy can only be transported by conduction. To limit the rise in centreline temperature, the fuel salt conductivity was artificially increased by a factor of 10. Gravity was then introduced in three steps, with the number of iterations used between each increment chosen by monitoring the behaviour of the solution to observe a tendency towards convergence before each change.

A key method for obtaining a solution for the salt flow is running with the Boussinesq approximation in the fuel salt to establish the flow and temperature fields before allowing density to vary. Because it is an open channel and there was less local variation in temperature, the coolant density was allowed to vary with temperature throughout. While running with Boussinesq, the fixed fuel salt density was set to be the same as the operating density for each case (3126 kg/m³ or 3430.5 kg/m³).

This solution is harder to run because there are disconnected fluid regions that are strongly affected by buoyancy forces, with substantially different fluid densities (e.g. 3430.5 kg/m³ in the fuel vs. 2860 kg/m³ in the coolant at 900 K), but only a single operating density can be specified. The operating density was chosen to suit the fuel because that domain is more challenging to converge. Similarly, it is also only possible to run a single RANS model in both domains (although it is possible to solve for laminar flow in a subset of domains), which may not be ideally suited to both.

Numerical Settings: Based on the experience of the cases described in Section 3.2, the SIMPLE segregated pressure-velocity coupling scheme was chosen, with body-force weighted pressure discretisation applied, as recommended for buoyancy driven flow (ANSYS, 2020b). By observing the solver messages generated during the early solution process, it was also found that applying the BCGSTAB pressure stabilisation method was beneficial.

The solutions were started with first-order upwind discretisation for the momentum and energy field, and were changed to second-order upwind discretisation when the solution was well enough converged to allow 'full' buoyancy (variable density) in the fuel salt, rather than the Boussinesq approximation. The fixed densities for the Boussinesq representation were chosen to be a good estimate of the mean density that would exist in the fuel at the temperature reached at this switchover point, so the jump in system mass is minimised, aiding the convergence of mass continuity.

Performing the Analysis

The Under-Relaxation Factors (URFs) for the density, turbulent viscosity, temperature and body-force terms were reduced from their default values of 1.0 to 0.95 for the start of the solution and increased to 0.98 later. The default values for other URFs (particularly 0.3 for pressure and 0.7 for momentum) were retained throughout.

The thermal radiation solution is computationally expensive and only changes as a result of the evolution of the temperature field (which depends on the flow), so the radiation solution was only turned on after the temperature and flow fields had converged. In addition, the radiation fields were only solved every 4 flow iterations for the forced circulation case, and every 10 iterations for the natural circulation case.

At the end of the natural circulation case, the turbulence modelling was turned off, and the solution only solved for laminar flow. Despite the flow being laminar in its final solution, it was not found to be possible to run the case without a RANS model active in the early stages, to damp out the 'spikes' in the velocity fields that occur, particularly with the addition of gravity.

Iteration	URFs	Source	Gravity	Buoyancy	U, T order	Radiation
start	0.95	25%	0	Boussinesq	1 st	off
50	0.98 T	100%				
100						
300			-1			
325			-5			
400			-9.81			
600	0.98			full	2 nd	
982*						on
1769*	0.98	100%	-9.81	full	2 nd	on

Table 3.4: Sequential steps of the solution strategy for the forced resolved rod case.

Iteration	URFs	Source	Gravity	Buoyancy	U, T order	Radiation	Turbulence
start	0.95	25%	0	Boussinesq	1 st	off	RANS
25	0.98 T	100%					
250			-1				
275			-5				
300	0.98		-9.81				
400				full	2 nd		
1400						on	
2824*							laminar
2875*	0.98	100%	-9.81	full	2 nd	on	laminar

Table 3.5: Sequential steps of the solution strategy for the natural resolved rod case.

Monitors and Convergence: The convergence of the solutions was assessed by monitoring the progress of the solution residuals, and also key quantities providing insight into the solution state and stability, and that also provide criteria for judging when a steady-state had been achieved. Because of the lack of features such as recirculation and separated flow, all of the solutions in this case study could be converged to true steady-states. This would not be the case with more complex geometries, where small regions of unsteadiness could prevent it.

The following quantities were recorded at each iteration:

Performing the Analysis

Maximum temperatures: The maximum cell temperature in the bundle coolant, fuel, cladding, duct and gap indicate if the solution remains bounded between acceptable values, and when the flow and heat transfer processes have stabilised. Jumps in the value or gradient of these monitors were seen as the steps of the solution strategy were applied. The application of thermal radiation results in a substantial reduction in maximum fuel temperature, for example.

Fuel salt velocities: The minimum and maximum vertical (z) velocities as well as the average velocity magnitude in the salt provided an indication of the flow stability in the fuel salt, as well as when it became steady. The maximum and minimum values were very 'spiky' when gravity was applied. Cases that diverged during the development of the solution strategy were seen to produce sustained unphysically large velocities.

Heat transfer: The integral of the heat flux across the cladding outer surface and to the inner surface of the duct provided an assessment of the convergence of the heat transfer solution, and the former also allowed energy conservation to be judged, because all of the heat source must cross the cladding to the coolant. The expected value is known from the fuel salt volume and volumetric source term, and so can be compared to an absolute value, and not just a measure of change in the solution.

Salt mass: The effect of changing from Boussinesq to 'full' buoyancy can be seen in changes in the mass of fuel salt, giving a positive indication that the change has worked as expected. The stabilisation of this value reflects the stabilisation of the mean fuel temperature.

A selection of these monitors are plotted in Figure 3.8 for the forced circulation conditions.

The solution was configured to be considered converged and automatically stopped (further iterations terminated) when the residuals for all fields fell below 10^{-5} (10^{-7} for energy) and the relative change in all of the four monitors underlined above was less than 10^{-5} over the preceding 5 iterations. Iteration numbers in Tables 3.4 and 3.5 marked with an asterisk* stopped automatically at these points when their residual and monitor convergence criteria were all met. The solutions were enabled to continue to solve for several thousand more iterations, if convergence had been not obtained.

The attention paid to the numerical determination of convergence criteria in these baseline models, and to the methods for ensuring solution robustness will be seen later to be key in enabling the minimum computation cost and manual intervention to be applied when running a large number of automated simulations.

3.3.4 Interpretation of Results

The baseline solutions for the forced and natural circulation conditions contain a large quantity of data, and despite the apparent good convergence, it is necessary to assess the result for the presence of expected features, and for the the absence of artefacts.

Solid surface fields: The temperature, total heat flux, radiation heat flux and shear stress distributions on wall surfaces were assessed, looking for discontinuities, patterns or manifestations of the mesh in the results. These patterns could indicate that the resolution is inadequate, or the fluid-solid interface is not working correctly, or another solution defect. The y^+ field was also inspected to confirm that it was less than one.

Performing the Analysis

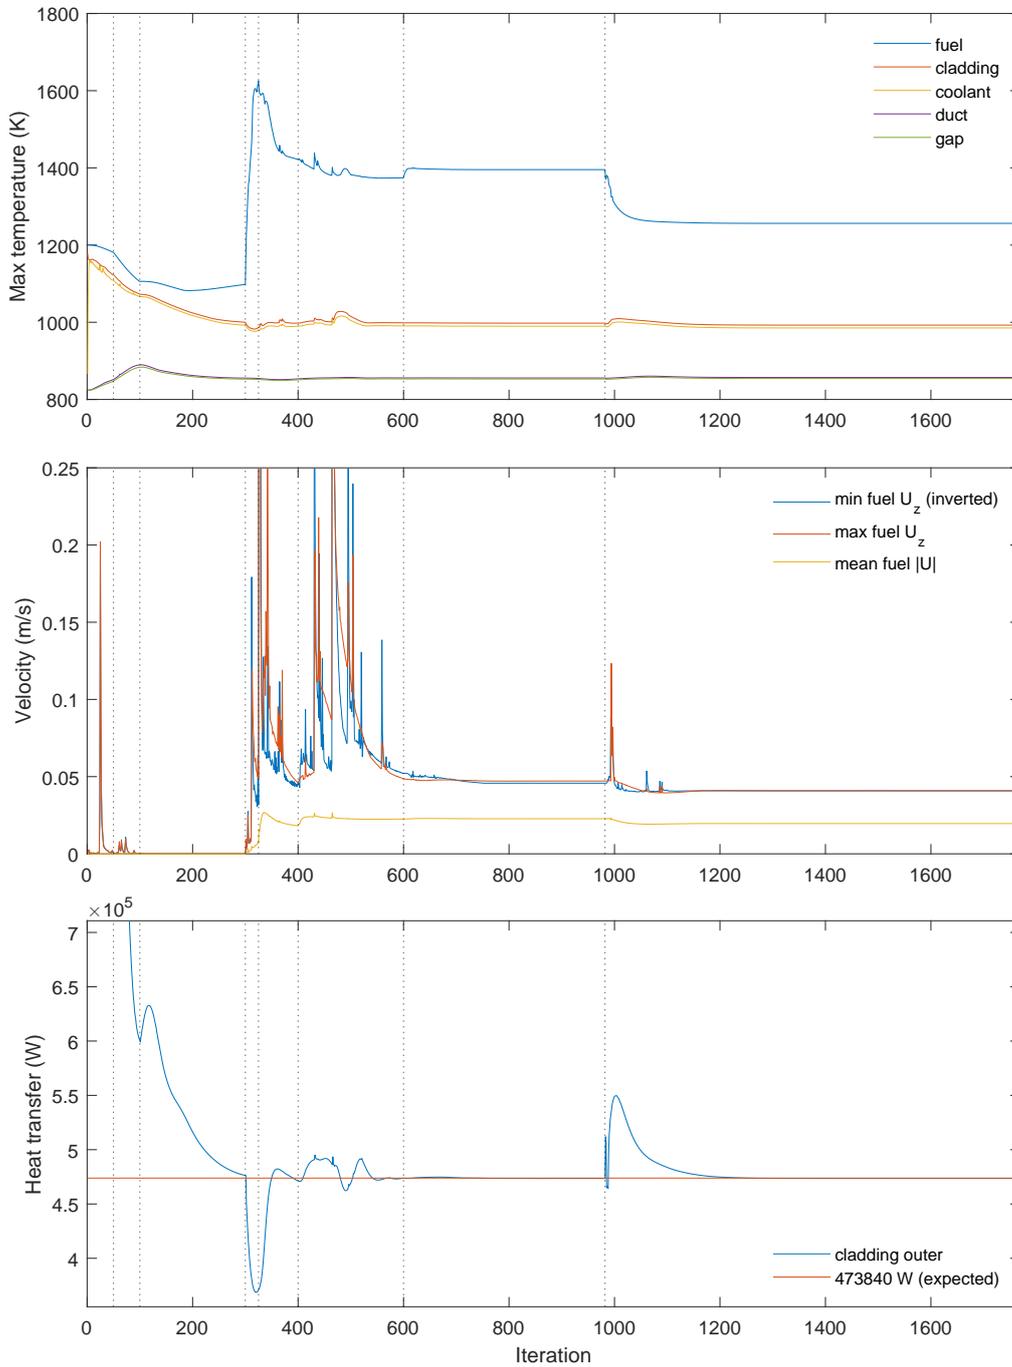


Figure 3.8: Convergence monitors from the forced baseline resolved assembly case solution. Vertical iteration markers correspond to events in Table 3.4. The minimum vertical fuel salt velocity has been inverted to aid comparison.

Performing the Analysis

Conservation: The mass and energy conservation was assessed by checking the inlet and outlet mass flows and the energy transfers across interfaces, compared to the energy source.

Cross-section plots: Planes through the fuel, coolant, cladding and duct regions were used to plot key quantities as contour plots at a range of heights. These plots allow the results to be interpreted for their meaning as a solution result, but also allow solution artefacts related to mesh or poor convergence to be seen. It is also possible to check overall features, such as whether the solution is symmetric as expected, that the effect of gravity is in the correct direction, or that fluid properties or turbulence fields have values in the expected ranges.

Line plots: The same quantities as investigated with cross-section plots were sampled on lines drawn from the centre to the inter-wrapper gap, through the fuel, coolant, cladding and duct, at a range of heights. Similar assessments were made as for the contour plots by assessing line plots of these quantities – the visualisation as a line plot identifies features that are harder to notice on a coloured contour plot, and vice-versa.

Examples of the cross-section contour plots at the mid-height of the fuel pin for a range of variables are shown in Figure 3.9 for the forced circulation baseline case and in Figure 3.10 for the natural circulation case. For the forced circulation case, it is not useful to visualise the temperature and velocity fields on the same contour plot for fuel and coolant simultaneously. The fuel is too hot and slow, and the scale of the colour bar means that both cannot be distinguished. Line plots allow that comparison to be made, however, and further assessment of these results using them are shown in Section 3.5 in comparison to a single subchannel model.

The most obvious result visible in both cases is that there is not a significant variation over the fuel pins except at the outer row adjacent to the duct, where the velocities and temperatures change rapidly. This result can also be seen in the line plots in Section 4.1 when they are compared to an equivalent porous simulation. No variation was seen in the fuel flow pattern that indicated widespread asymmetry in the natural circulation. Visualising temperature in the coolant for the forced circulation conditions is difficult because the thermal boundary layer is thin compared to the velocity boundary layer, as a result of the $Pr > 1$ coolant salt.

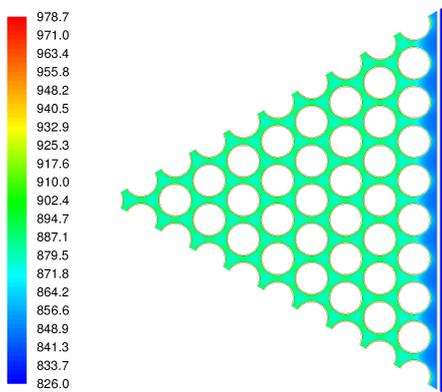
3.3.5 Turbulence Modelling in the Fuel Pin

For the natural circulation cases, the models become effectively laminar when converged in both the coolant and fuel – the turbulence decayed to almost nothing. This confirms that the damping terms in the RANS model used are able to represent laminar conditions. The switch to a truly laminar solution changes very little in the solution (and hence the model converges after only 51 more iterations, Table 3.5).

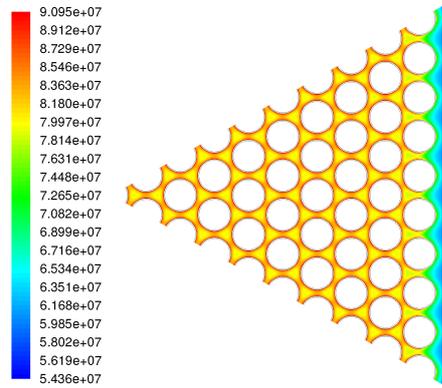
It was expected that, once convergence had been achieved, the fuel salt flow in the forced circulation case could also be made laminar by using Fluent's 'laminar zone' option for the fuel tube regions only. However, it was found that it was necessary to keep the RANS model active in the fuel; running with the fuel as a laminar region produced erratic maximum salt velocities observable in the monitors, approximately 10 times higher than with a turbulence model active, and high solution residuals were seen.

From a simple pipe flow Reynolds number estimate using the pin dimensions and the expected velocities (0.05 m/s), and from the comparison in Section 3.2, it was expected that the fuel circula-

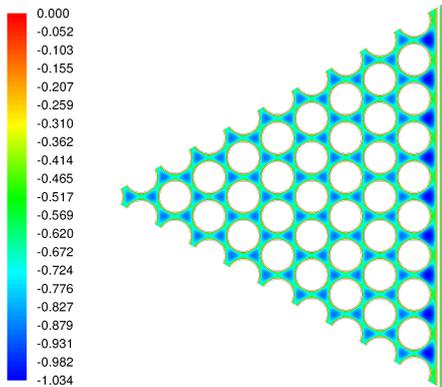
Performing the Analysis



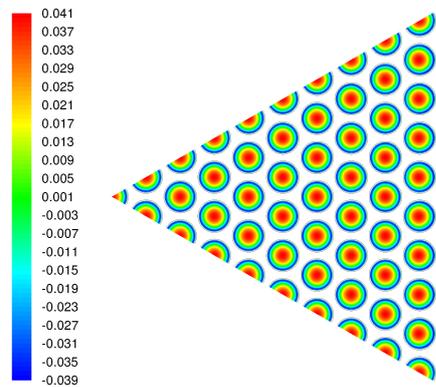
a: Coolant temperature (K)



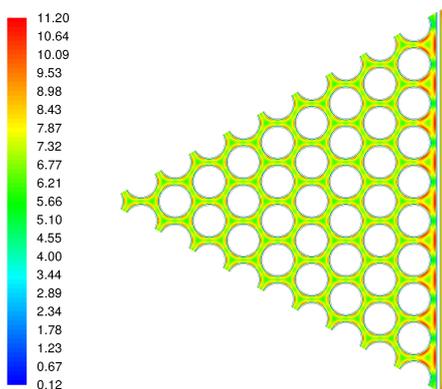
b: Volumetric absorption of radiation in coolant (W/m^3)



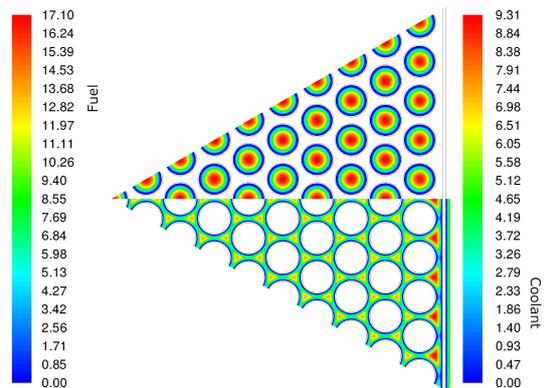
c: Coolant vertical velocity (m/s)



d: Fuel vertical velocity (m/s)



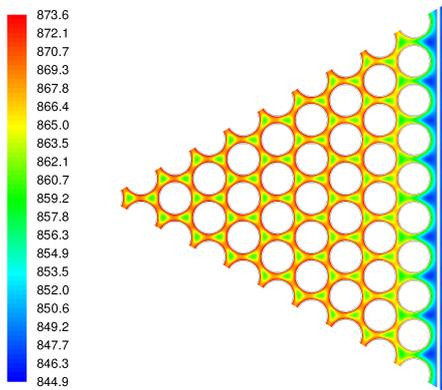
e: Coolant turbulent intensity (%)



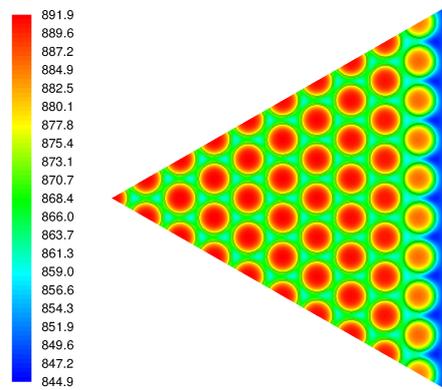
f: Fuel (upper, left hand scale) and coolant (lower, right hand scale) turbulent viscosity ratio

Figure 3.9: Cross-sections of the resolved assembly model under forced circulation conditions at $z = 0.8 \text{ m}$ (mid-height).

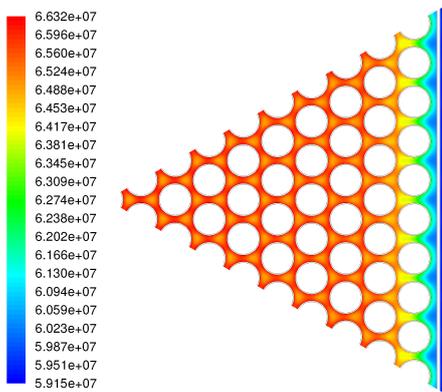
Performing the Analysis



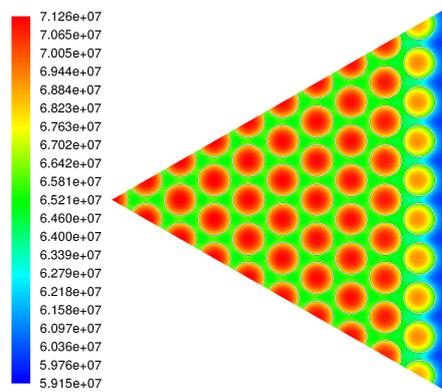
a: Coolant temperature (K)



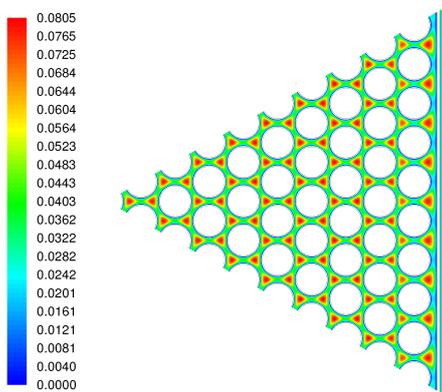
b: Coolant, fuel, cladding and duct temperature (K)



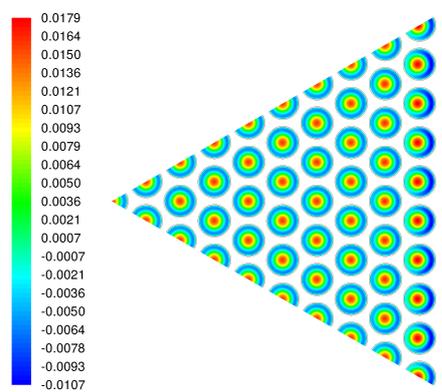
c: Volumetric absorption of radiation in coolant (W/m^3)



d: Volumetric absorption of radiation in coolant and fuel (W/m^3)



e: Coolant vertical velocity (m/s)



f: Fuel vertical velocity (m/s)

Figure 3.10: Cross-sections of the resolved assembly model under natural circulation conditions at $z = 0.8$ m (mid-height).

Performing the Analysis

tion would be laminar. Further investigation described in Section 3.4 showed that the need to use a turbulence model was not just a numerical solution artefact, however. The following facts indicate that the solution was legitimately picking up on the need to model the presence of turbulence in the fuel:

- The natural circulation cases were able to return laminar flow even with the RANS model active means that it would not spuriously ‘insist’ on turbulence being present.
- The poor behaviour when attempting to run with laminar flow in the salt for the forced circulation case.
- When solving with RANS turbulence modelling active in the fuel, the turbulent viscosity ratios for the forced circulation case are predicted to be >10 (Figure 3.9f) which indicates non-negligible turbulent effects. These ratios are actually higher than in the coolant for the same case, although the turbulence intensity in the fuel is only $\approx 1\%$, whereas in the coolant it is 5 to 10%.

3.4 Single Fuel Pin

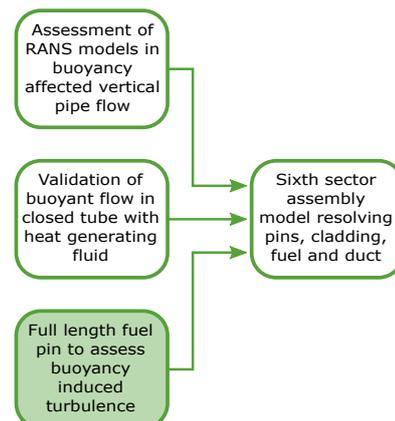
The validation comparisons described in Section 3.2 ran successfully with laminar settings in the closed tube domain, as did the low heat load natural circulation cases above for the resolved fuel assembly, but the high power forced circulation cases did not.

To understand the difference, and determine if it was present for numerical or physical reasons, a simplified version of a whole single fuel pin (with the same dimensions as the pins in the fuel assembly) was created. The fluid properties were made constant (sampled at 1200 K) the Boussinesq approximation was applied to maintain constant density, a uniform heat source of 100 MW/m^3

was used and a constant temperature external boundary condition of 900 K was applied on the outside of a cladding. This results in a value of $S = 4.78 \times 10^4$. Turning the RANS model ($k - \omega$ SST with low Re corrections) on and off in Fluent repeatedly and reversibly generated erratic high velocities and high residuals when laminar, which disappeared and converged and a steady, symmetric convection pattern appeared with the turbulence model active. Running the same setup in CFX resulted in the same behaviour, which suggested that the behaviour was more likely to be the result of a physical turbulent process, not a numerical quirk.

This single fuel pin model was investigated in detail, and it allowed the conclusion to be drawn that the flow of fuel salt for the forced circulation case is expected to be turbulent, and is able to be represented by the chosen RANS models. The steps used, and the process for understanding this buoyancy driven effect, arriving at this conclusion and creating robust supporting justifications are discussed in Appendix A.

The experiments discussed in Section 3.2 were performed in wider tubes of similar length, with lower heat sources, but these changes substantially cancelled out to produce a similar value for the



Performing the Analysis

non-dimensional parameter S . S is a Rayleigh number multiplied by r/L , the radius to length ratio of the tube, which is $1/50$ in the experiments, and $(d_{fuel}/2)/L = 1/340.4$ for the fuel pin. Therefore, for a condition of a given S , the value of Ra in the fuel pin must be higher. The equivalent value of S in the experiments by Murgatroyd and Watson (1970) at the Ra used in the single fuel pin is

$$4.78 \times 10^4 \left(\frac{(d_{fuel}/2)/L}{1/50} \right) = 3.25 \times 10^5$$

which is in the range where instabilities, and towards the transition to turbulence, were observed in the experiment. In strongly buoyancy affected flows, it is often the case that a critical Rayleigh number determines the transition to turbulence, and so the higher Ra in the fuel pin is likely generate transition at lower values of S .

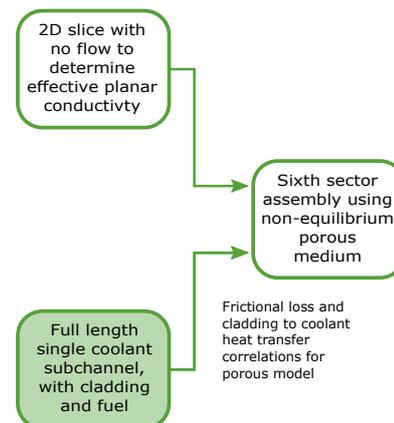
This is an example of how important it is to understand the meaning and validity of the parameter range for any validation data. While the fuel pin is within the tested range of S values, it is not in the experimental r/L , hence Ra range, and there is not a parameter elucidated by Watson (1971) that characterises laminar-turbulent transition explicitly.

3.5 Single Coolant Subchannel

The resolved assembly model demonstrated that the behaviour in the fuel and coolant can be well represented by a single subchannel and surrounding three fuel pins. A mesh was created to represent this subset of the cross-section, but with the same axial extent and resolution as the $1/6^{\text{th}}$ assembly (Figure 3.11). The mesh contained 460,000 cells – a factor of approximately 100 less. Symmetry conditions were applied on the exposed sides of the fuel, cladding and coolant.

The intention of this model is to assess several different flow conditions, and then derive friction factor and heat transfer correlations from the results. It will also be used as the basis for the SA and UQ assessment described in Section 4.3.

The same model setup (with the duct and gap settings removed) and solution strategy as applied to the resolved $1/6^{\text{th}}$ assembly model (referred to as the ‘full’ model below for brevity) was able to be directly applied to the single subchannel model, and was still found to be necessary to successfully achieve a converged solution.



3.5.1 Comparison to Full Assembly Model

To assess the agreement between this subchannel model and the full assembly, the same baseline forced and natural circulation conditions were applied. The results of this can be seen in Figures 3.12 and 3.13 for the forced and natural circulation conditions respectively. The details of the fields are easier to see and compare in the more compact image, compared to the contour plots of the whole assembly cross-section. There is a horizontal line drawn on Figures 3.12a and 3.13a –

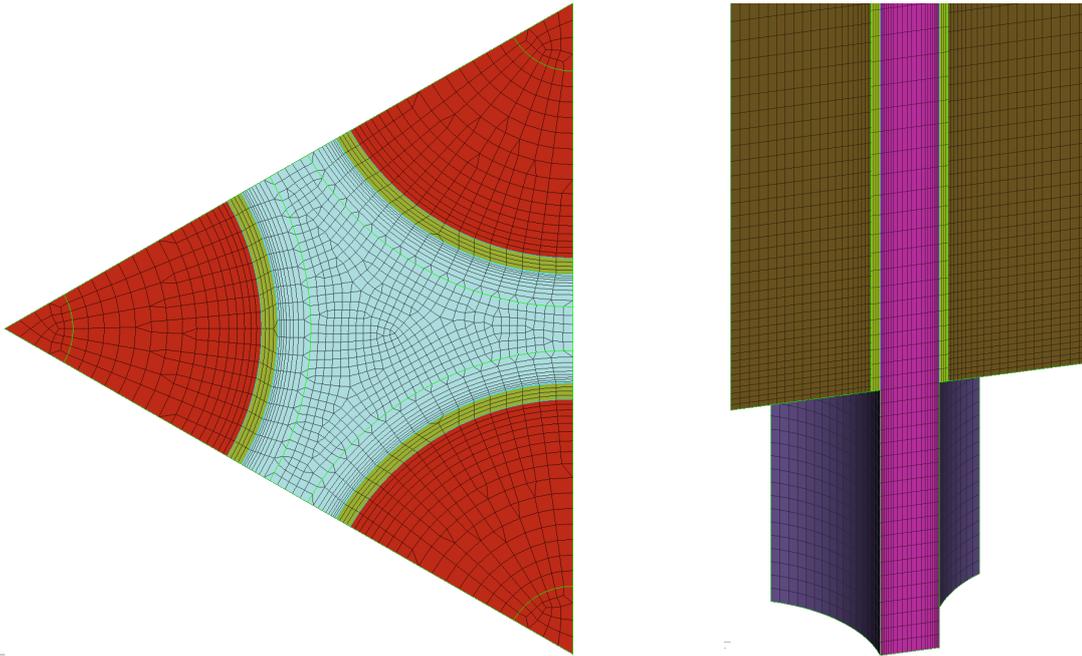


Figure 3.11: Single subchannel mesh with surrounding fuel pins. The extended sections of coolant at top and bottom (shown) were retained from the 1/6th assembly model.

Figures 3.14 and 3.15 show the temperature and axial velocity fields along these lines, compared to the equivalent results from the centre channel of the full assembly model. The results show the key features of the flow:

- The temperature in the fuel varies significantly across the distance from the cladding to the centreline of the fuel pin.
- In the fuel, the velocities are upwards in the centre of the pin, and downwards at the edge.
- In the forced circulation case, the fuel temperatures are much higher than in the coolant, and the fuel velocities are much lower. In the natural circulation case they are both closer in scale.
- In the forced circulation case, the coolant flows downwards, with steep momentum and thermal boundary layers. In the natural circulation case, the coolant flow is upwards and shows a laminar velocity profile.
- The 'dips' in the middle of Figures 3.14 and 3.15 occur where the flow passes through the closest gap between fuel pins.
- The absorption and emission fields for participating thermal radiation are closely related to the temperature fields, and have volumetric power densities of 60 to 100 MW/m³, comparable to the nuclear heat source.

There are small differences between the profiles along these radial lines, but Figures 3.16a to 3.16c show that this is because of accumulated differences in the axial profiles. The radial comparison is made to the centre of the full assembly, which will have some variation in temperature and flow distribution across its cross-section with height.

Performing the Analysis

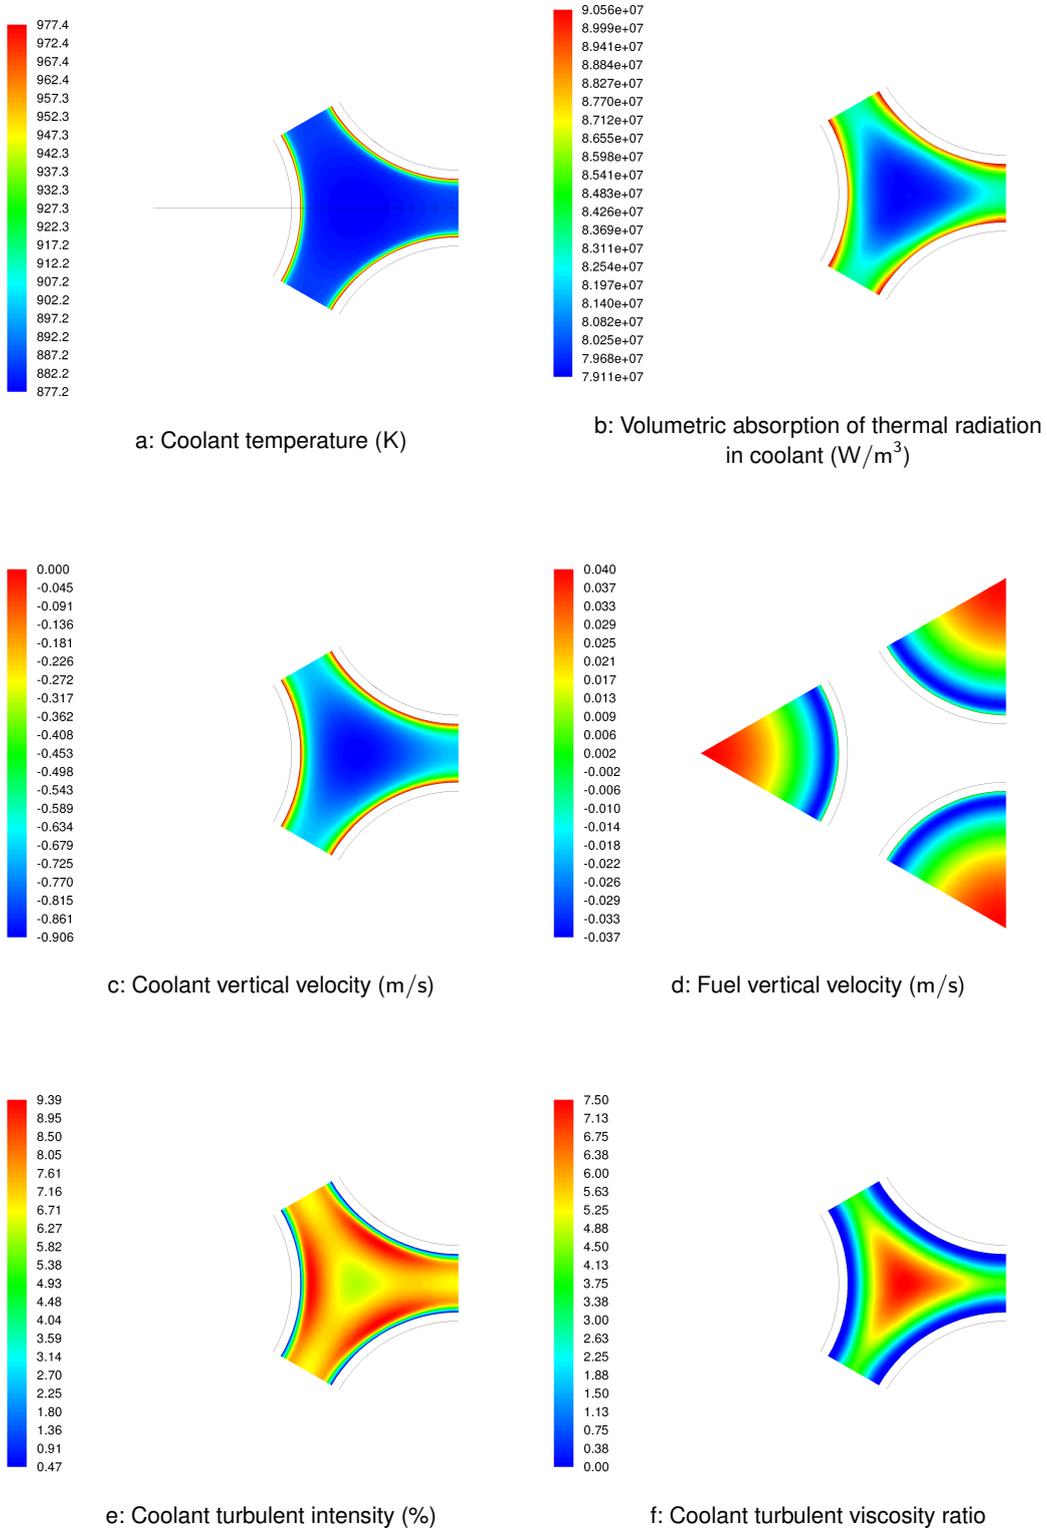
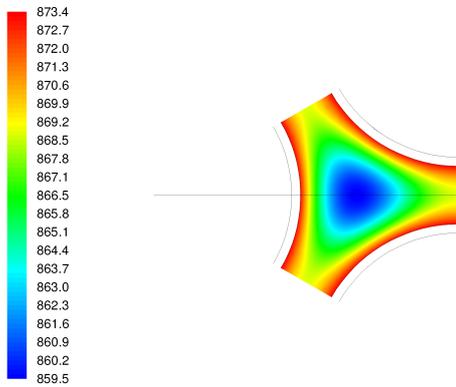
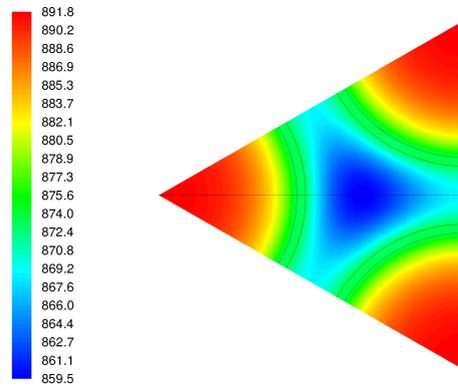


Figure 3.12: Cross-sections of the single channel model under forced circulation conditions at $z = 0.8$ m (mid-height).

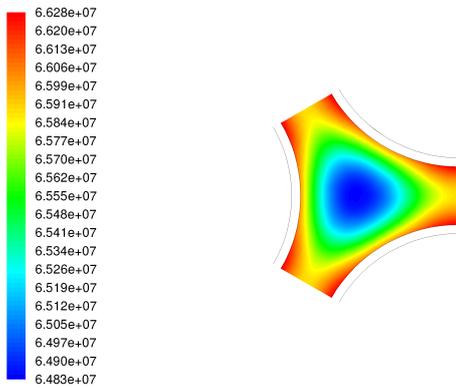
Performing the Analysis



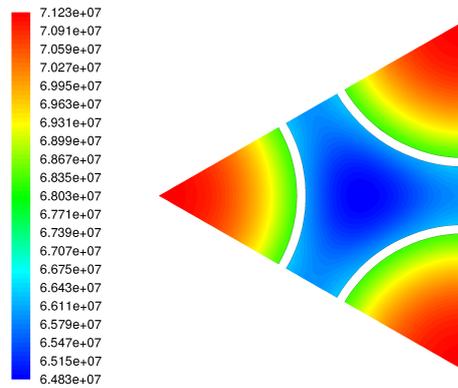
a: Coolant temperature (K)



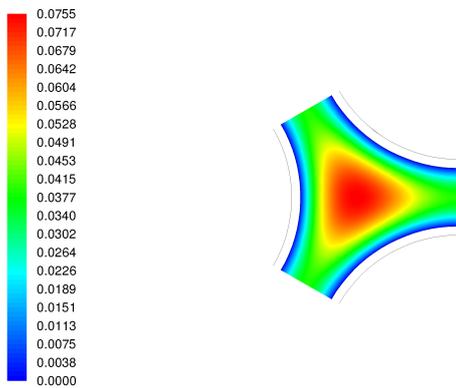
b: Coolant, fuel, cladding and duct temperature (K)



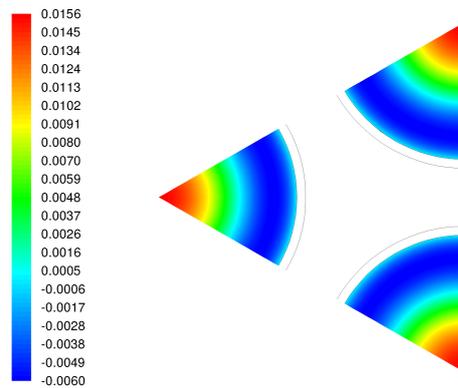
c: Volumetric absorption of thermal radiation in coolant (W/m^3)



d: Volumetric absorption of thermal radiation in coolant and fuel (W/m^3)



e: Coolant vertical velocity (m/s)



f: Fuel vertical velocity (m/s)

Figure 3.13: Cross-sections of the single channel model under natural circulation conditions at $z = 0.8$ m (mid-height).

Performing the Analysis

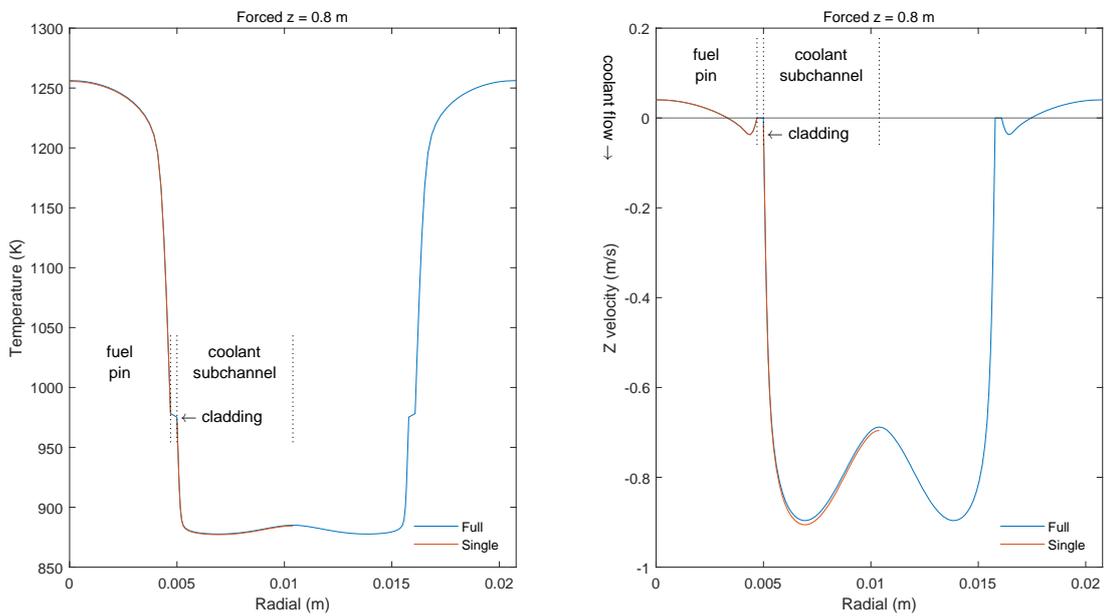


Figure 3.14: Temperature and axial velocity in single subchannel model compared to the centre of the full assembly model for baseline forced circulation conditions.

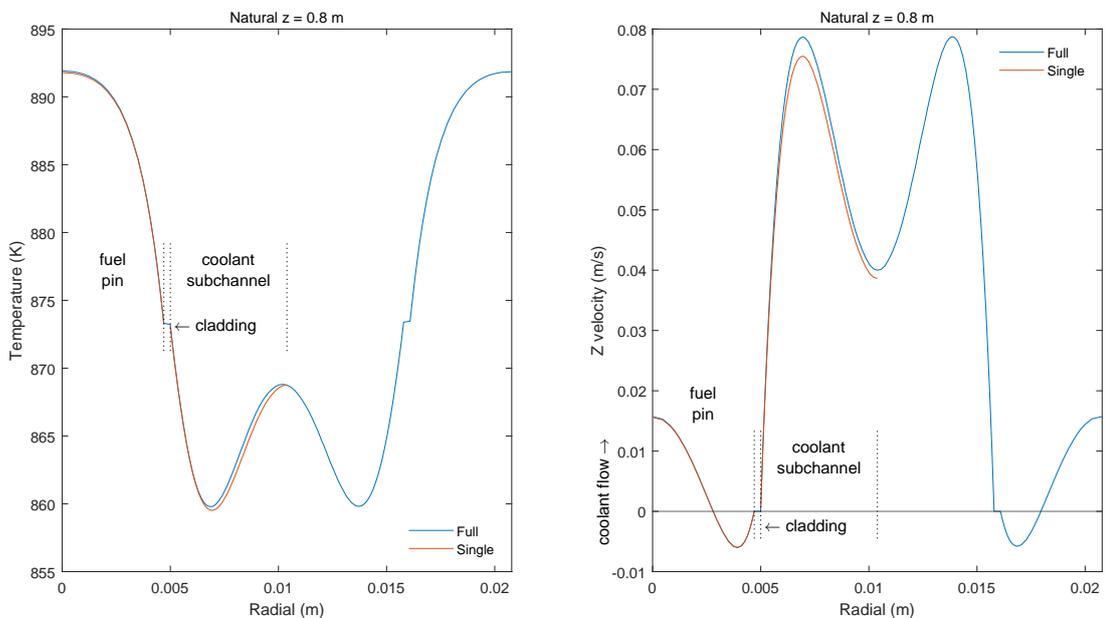
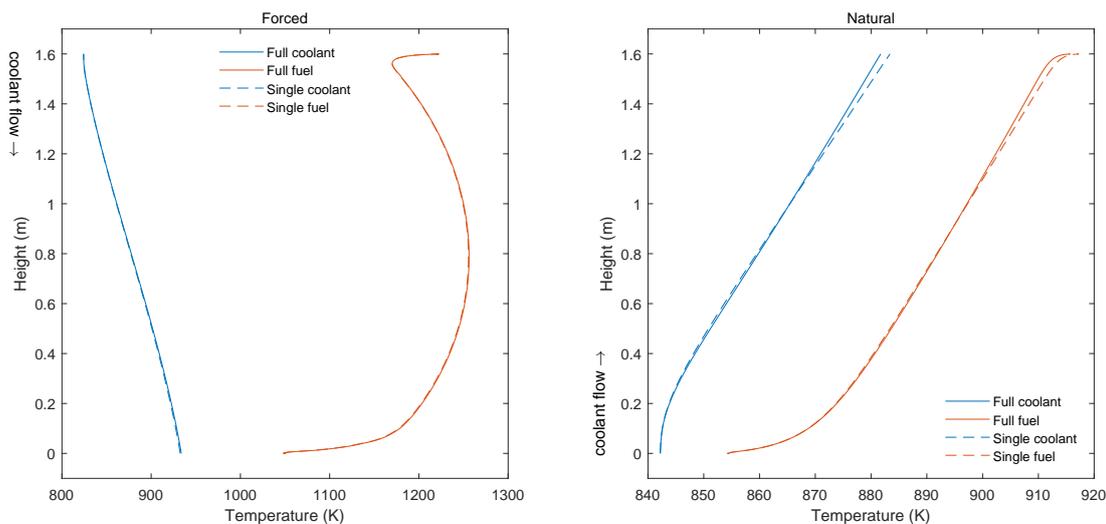
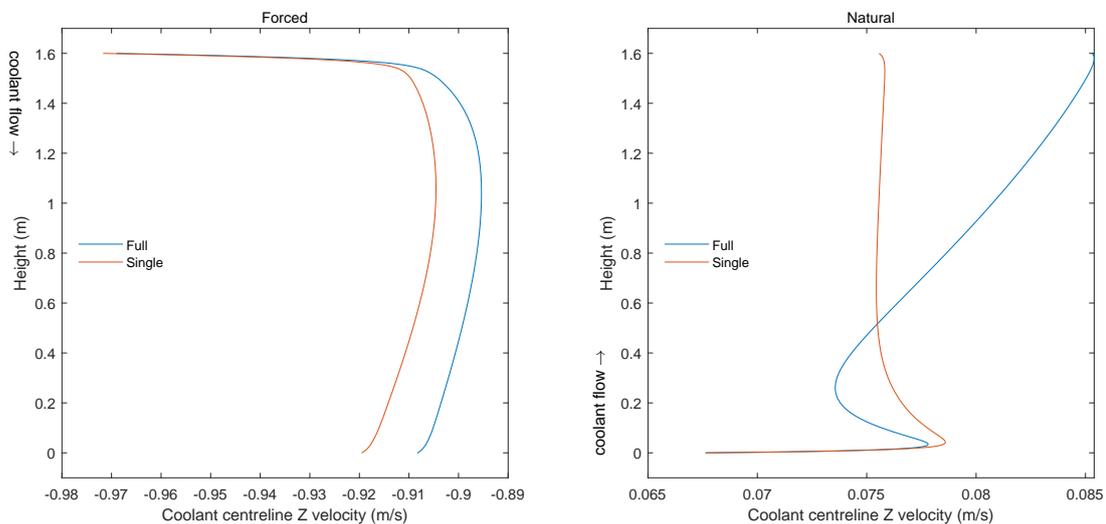


Figure 3.15: Temperature and axial velocity in single subchannel model compared to the centre of the full assembly model for baseline natural circulation conditions.

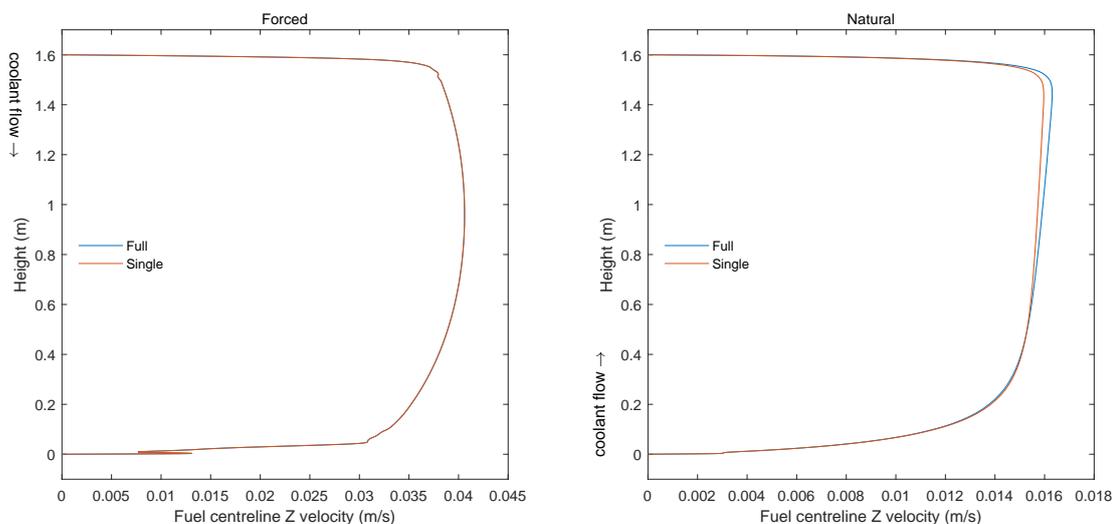
Performing the Analysis



a: Temperature at the fuel pin centreline and centre of coolant subchannel.



b: Coolant axial velocity at the centre of coolant subchannel.



c: Fuel axial velocity at the fuel pin centreline.

Figure 3.16: Comparisons of vertical profiles for the full assembly model and equivalents in the single channel model for the forced (left) and natural (right) circulation cases. Central fuel pin and nearest coolant subchannel used for full assembly. Note the direction of coolant flow.

Performing the Analysis

- The variation of coolant and fuel temperature in the forced circulation case is almost identical (Figure 3.16a). For the natural circulation case, the conditions at the inlet (base) diverge slightly with height (the largest discrepancy is at top/outlet).
- The divergence in temperature in the natural circulation case is caused by a difference in coolant velocity (Figure 3.16b). The coolant velocity increases with height in the natural circulation case – this is caused by a subtle concentration of flow towards the full assembly centre, driven by buoyancy, and the increasing radial distance from the duct. There is a similar effect in the forced circulation case, but it is proportionally smaller.
- The steep change in velocity along the centreline of the coolant channel at a height of 1.6 m (at the top of fuel pins, Figure 3.16b) for the forced circulation case is not a result of developing flow (noting that the flow inlet is at 1.7 m). It is associated with the sudden change in wall heat flux (from zero) at the start of the solid cladding section, which results in a change in turbulence, flattening the coolant velocity profile, reducing the centreline value.
- The variation of fuel centreline velocity is driven by the local coolant temperature on the other side of the cladding, and so there is bigger difference between the full and single case for natural circulation conditions, although it is relatively small (Figure 3.16c). The presence of the inflection in the velocity field near the base of the fuel pin in the forced circulation case, and its absence for the natural circulation case is consistent with the cases described in Section 3.2 at the equivalent values of S . Similarly, the larger fuel axial velocities in the forced circulation case compared to the natural circulation case were expected, given the higher buoyant driving forces.

These differences are relatively small, and perfect agreement between the centre of the full case and the single cases is not needed to allow the latter to be used to derive friction factor and heat transfer correlations over a range of flow conditions.

3.5.2 Variation in Flow and Derivation of Correlations

The three forced circulation conditions from CH1 to CH3 and the natural circulation conditions from the three instants in CH2 defined in Table 2.2 were used to provide a range of flows and heat transfer to enable the range of applicability of friction factor and heat transfer correlations to be broader. These flow conditions were applied to the single subchannel model. The CH1 and CH3 flows were continued from a partially converged baseline (CH2) result, this provides a saving in computation cost, but was also found necessary to obtain convergence. The solution strategy developed in Section 3.3.3 is therefore not universal, and would need some modification to successfully start these cases with different flow conditions from an initialised state. The same was found with the natural circulation conditions at 150 s and 500 s – they were started from the 250 s solution.

Correlations⁸ for friction factor and Nusselt number in the coolant were required, parameterised by Prandtl and Reynolds numbers, so that the dimensionless quantities can link the flow and heat transfer processes with the material properties of the salt, giving a more transferable result. The characterisation of convection is discussed in Volume 2 (Section 2.1.2).

⁸ Because the temperature and heat flux fields vary along the length of the coolant channel, and there are coolant density and viscosity changes associated with this, the flow is never completely ‘fully developed’ with constant heat flux (as idealised in mathematical analysis). However the differences are expected to be small, given the relatively slow changes in conditions and the small hydraulic diameter relative to the length.

Performing the Analysis

Planes were created across the coolant and cladding outer surface in the model at 0.2 m intervals from 0.2 m to 1.4 m along the active length (the first and last 0.2 m were not considered, to reduce inlet and end effects). The CFD solution was sampled on these planes⁹ to return a range of quantities at each height:

- Coolant channel mass flow rate, W .
- Coolant channel area, A .
- Mass flux weighted¹⁰ average (bulk) coolant temperature, T_c .
- Area weighted coolant density, ρ .
- Area weighted cladding outer surface (forming the coolant flow path) wall temperature, T_w .
- Area weighted cladding wall shear stress, τ .
- Area weighted cladding wall total heat flux, q , (including the net radiative heat transfer from the cladding surface to the coolant).
- The volumetric heat source in the fuel, q_{vol} .
- The fuel centreline temperature at the given height, T_{cl} .
- The cladding inner surface, T_i .

The bulk temperature, T_c , was used to evaluate the coolant material properties that are used in further calculations. The HTC, h , was obtained from the cladding wall heat flux

$$h = \frac{q}{T_w - T_c}$$

and converted to Nusselt number using

$$Nu = \frac{hD_h}{k}$$

The Reynolds number at each height was calculated using the coolant mass flow and subchannel hydraulic diameter and area

$$Re = \frac{WD_h}{A\mu}$$

and the Darcy friction factor was obtained from the wall shear stress

$$f = \frac{4\tau}{W^2/(2\rho A^2)}$$

These results are plotted in Figure 3.17 and 3.18 for the forced and natural circulation cases respectively.

⁹ Results extracted on planes were effective in this case because the flow was steady and the domain was a simple 1D channel. For more complex cases it may be preferable to average fluid properties over a short section of volume, and average wall properties over the appropriate adjacent bounding surface. Care is needed in selecting the volumes and surfaces, and the mass flux weighted average temperature still needs to be calculated, not a density or volume weighted average.

¹⁰ `surface-integrals/mass-weighted-avg` in Fluent – using a different surface average will give the wrong result.

Performing the Analysis

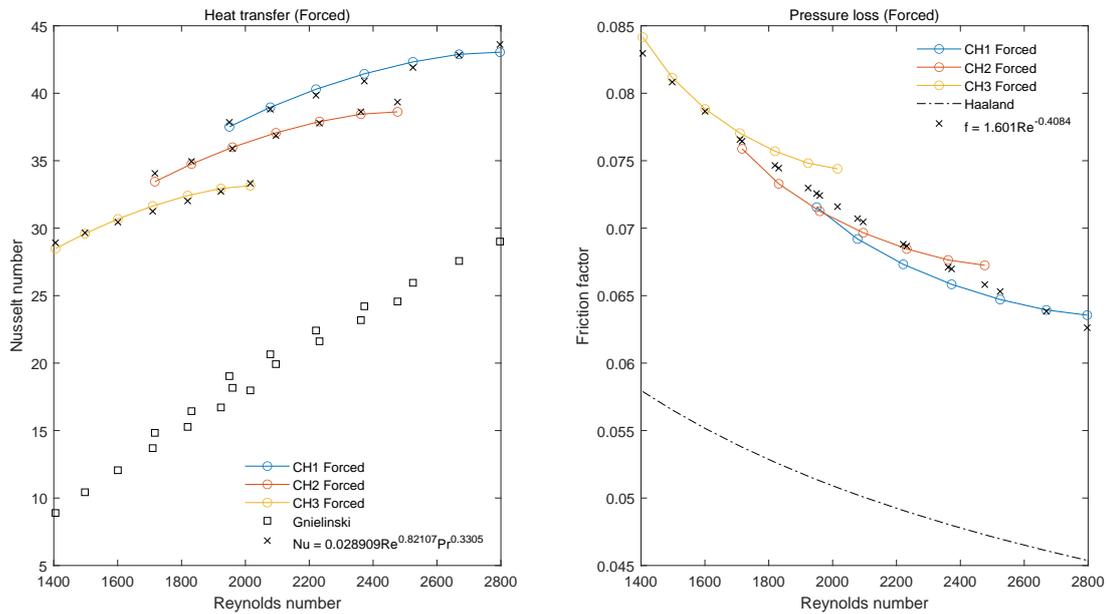


Figure 3.17: The variation of Nusselt number and friction factor with Reynolds number for the three forced circulation conditions in the single coolant subchannel case. A relevant value that would be predicted for turbulent flow from the literature is shown, along with the best fit equation delivered from the CFD results.

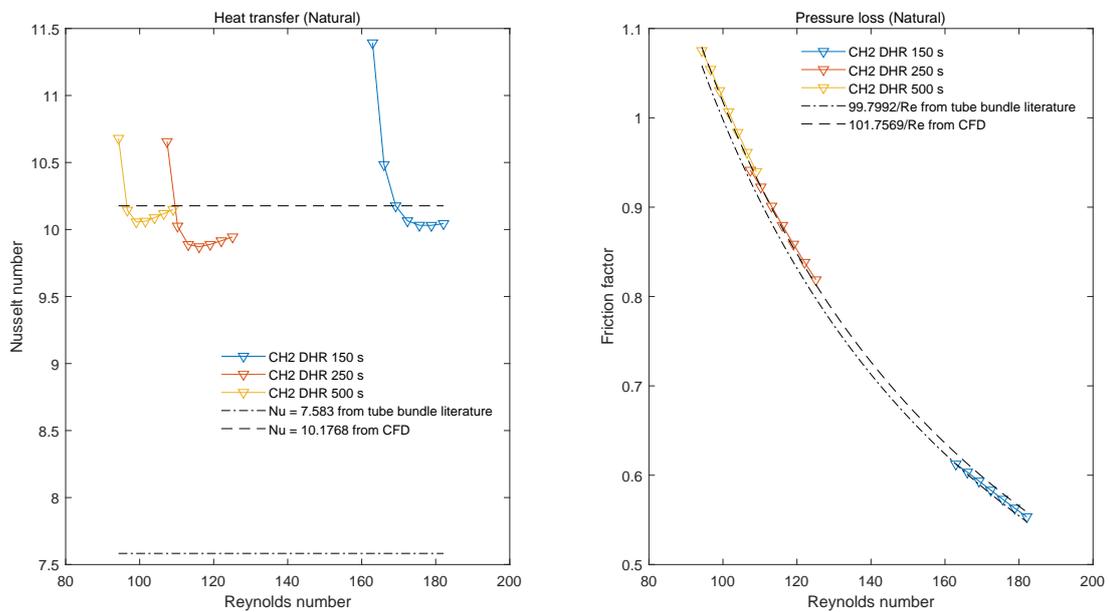


Figure 3.18: The variation of Nusselt number and friction factor with Reynolds number for the three natural circulation conditions in the single coolant subchannel case. Relevant values of Nu and f that would be predicted for laminar flow from the literature is shown, along with the best fit values delivered from the CFD results.

Performing the Analysis

Forced circulation: For turbulent flow, Nu increases with Re , and also depends on the Pr of the fluid. Friction factor decreases with Re at transition Re values, and is constant at high Re . Figure 3.17 shows that these trends are observed in the single channel results. The variation in Re within each of the three region results (CH1 to CH3) is caused by the variation in viscosity with temperature.

For Nu , Kakaç *et al.* (1987, Chapter 7, for a fluid with $Pr = 10$ and $Re \geq 10^4$) suggests that the expected value in a triangular rod array with $p/d = 1.2$ should be approximately 1.1 times greater than the Nu in a circular tube. The expected circular tube Nu is represented by the Gnielinski correlation (Volume 6, Section 3.1, also available in Incropera *et al.*, 2011 or Rohsenow *et al.*, 1998). The same reference suggests that the expected friction factor for this array would be approximately the same as would be found for a circular tube, although this guidance is for $Re = 10^4$. The Haaland correlation (Volume 6, Section 3.1, Kakaç *et al.*, 1987 or Rohsenow *et al.*, 1998) for circular pipes (applying zero roughness) is included for reference. The Gnielinski correlation requires a friction factor; this was provided by the values of f found from the CFD.

Curve fits have been applied to the CFD results, with the expressions and coefficients shown on the figure. The Nu expression has the same form as the Dittus-Boelter correlation (producing the same power of ≈ 0.8 for Re , and $\approx 1/3$ for Pr as found in several other correlations), and the friction factor expression has the same form as the Blasius correlation.

The results obtained from the CFD solution are not in particularly good agreement with the literature expectations – the Nu values follow a similar trend but are approximately 1.6 to 2 times the Gnielinski values. The friction factor behaviour is similar, being approximately 1.4 times higher than the Haaland values. The CFD results (and the resulting best fit correlations) simultaneously combine the effects of the subchannel shape, turbulence, buoyancy, thermal radiation and property variation; the contributions of each are hard to separate out without rigorous study, that is often not possible in an industrial design or assessment process. There are several reasons why the lack of agreement could occur:

- The assessments in Section 3.1 indicated that (for a circular tube at $Re = 5300$) using the low- Re corrections to the $k - \omega$ SST model were likely to over predict Nu and f , but not by as much as seen here. It is also the case that in the natural circulation results (below) the laminar Nu is substantially larger than the value predicted by the literature, and that is unaffected by the turbulence model.
- The transition Re means that the shape of the channel, and hence the velocity profile (including the narrow gap between pins), has more influence, compared to high Re flow where the increased mixing causes the velocity to be more uniform.
- The inclusion of the contribution of thermal radiation to the coolant increases the apparent convective HTC. It is appropriate to do this because the target application of the correlations are porous models and system codes, neither of which would model the radiation separately at the pin array scale.
- While the values of Bo shown in Table 2.4 mean that substantial buoyancy effects are not expected, the buoyancy opposed flow is expected to enhance heat transfer to an extent.
- Convection correlations generally assume that fluid properties are constant across a cross-

Performing the Analysis

section. The high Pr of the coolant and the strong dependence of its viscosity on temperature mean that there is approximately a 100 K drop in temperature near to the cladding surface (Figures 3.12a and 3.14). This means that there is a localised decrease in viscosity near the wall, increasing the gradients and increasing the heat transfer. This is discussed in Volume 6, Section 3.1, and that guidance is used to estimate the size of the effect. Evaluating the coolant viscosity at 850 K, giving μ_c (bulk) and at 950 K, giving μ_w (wall)

$$\left(\frac{\mu_c}{\mu_w}\right)^n = 1.4824^{0.11} = 1.044$$

it can be seen that approximately 4% of the difference between the results is predicted to be attributable to the effect of property variation (although the value of $n = 0.11$ is only valid for $Re > 10^4$, and the effect at lower values could be different). It will also be shown in Section 4.3 that the coolant viscosity has a substantial influence on the cladding temperature.

Natural circulation: For laminar flow, constant values of Nu are expected, and the friction factor is expected to vary as

$$f = \frac{fRe}{Re}$$

where $fRe = 64$ for a circular tube, for example. Figure 3.18 shows that this behaviour is observed in the CFD results. For Nu , at the lower value of Re for each of the flow conditions (corresponding to different times in the SBO transient) there is an increased value. This is at the base of the coolant channel, where there is a strong inlet effect due to the short distance from the inlet to the base of the fuel pins. The expected value of Nu for a constant heat flux boundary for laminar flow in a triangular array of $p/d = 1.2$ is given by Kakaç *et al.* (1987, Chapter 7) and plotted for comparison. The same reference provides a value of $fRe = 99.8$, which is also plotted for comparison.

Best fit values have been determined from the CFD results and are also shown. The quality of the fRe best fit and agreement between the literature value and the CFD is very good. For Nu the CFD predicts a value of Nu that is about 1.3 times higher. Compared to the discussion of forced circulation conditions above, there are no turbulence modelling effects, and the velocity profile should be similar to the literature value (the friction factor agreement supports this). The temperature variation with wall distance (Figures 3.13a and 3.15) is smaller, so the effect of varying properties will be less. Some enhancement to heat transfer for this buoyancy assisted laminar flow is expected, but the effect is expected to be small (Section 2.2.4). The contribution to cladding surface heat flux from thermal radiation is proportionally greater in the natural circulation cases (approximately 15 to 20% for the natural circulation case compared to 7 to 9% for the forced circulation case), and so this will also increase the apparent convective HTC.

While some details remain that could be pursued further, it is judged that the correlations derived are suitable to use as the basis of a porous model representation, as long as the simulations stay within or close to the Reynolds number range used to determine them.

Performing the Analysis

3.5.3 Solid Fuel Approximation

Much of the solution difficulty associated with solving these cases is related to the circulating fuel salt. It was postulated that an approximation could be made that the fuel was solid, but with an enhanced effective conductivity to replicate the effect of the circulating flow. This would allow the coolant heat transfer and cladding temperature to still be evaluated with minimal approximation by resolving the coolant flow path, and the energy distribution and centreline temperature in the fuel to be predicted by a simplified equivalent solid model.

From the results extracted at each plane, the relationship between the centreline temperature of the fuel and cladding was used with an analytical expression for conduction in a cylinder with a uniform heat source (Incropera *et al.*, 2011) to calculate an equivalent conductivity, if the fuel was a solid.

$$k_{eff} = \frac{q_{vol}(d_{fuel}/2)^2}{4(T_{cl} - T_i)}$$

This value is approximately 5 times higher than the fuel's conductivity for the forced circulation conditions and 2 times higher for the natural circulation conditions.

By making the fuel a solid with a constant conductivity of k_{eff} , evaluated at mid-height ($z = 0.8$ m) then the coolant temperature and velocity can be matched to the circulating fuel case, as can the centreline temperature in the mid-part of the fuel. However, the details of the fuel's radial temperature distribution, and the centreline temperatures at either end do not match well (they are not reported here). The treatment of the fuel as a solid also does not represent a significant computational cost saving because the coolant flow still needs to be resolved, requiring a substantial number cells.

The insight gained from this assessment is that there is little net axial heat transfer by the circulating fuel salt over the majority of the fuel pin height. The local energy source and the local cladding surface heat flux were closely balanced – the circulation acts to reduce the fuel centreline temperatures but does not redistribute the energy source. This is not true at either end of the fuel pin, where the velocities reduce and reverse, and the distribution of energy is more complex. This means that any simplified approximation of the internal fuel behaviour details (including a porous representation) is likely to contain inaccuracies at the top and bottom of the assembly. This could be overcome by creating a spatially varying effective conductivity, but deriving this, and generalising it to other flow conditions is a non-trivial task, and may not be able to be easily linked to parameters with physical significance.

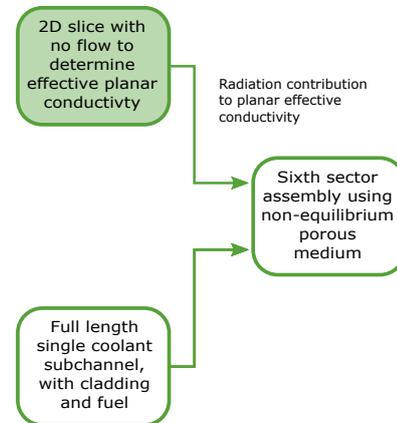
An alternative approach to using an auxiliary evaluation to predict the internal state of a fuel rod, without solving for the flow in it explicitly, is shown in Section 4.2.

Performing the Analysis

3.6 Transverse Effective Conductivity

The single subchannel models described above provide the coolant pressure loss and heat transfer characteristics from the cladding to the coolant. Beyond this, the remaining properties needed for a porous model depend on the geometry and properties of the materials only, except for the net heat transfer in the planar (transverse, or cross-bundle) direction. The effective conductivity in these directions in the pin bundle is the net result of the combination of several simultaneous heat transfer mechanisms:

- Conduction through the coolant, fuel and cladding.
- Thermal radiation transport through the fuel and coolant.



These mechanisms and their inter-relationships are too complex to be evaluated reliably by simplified calculations, and so are simulated using a 2D slice model. This allows a temperature dependent material property to be derived to represent them in a porous model.

The fuel, coolant and cladding mesh from all pins in the 1/6th assembly model at a single plane were used, bounded by the inner surface of the duct, which was held at a known temperature using a large HTC. The salt and cladding were all created as solids, with their normal material properties, except the fuel salt, which was given an effective conductivity 5 times normal, as derived above, to account for the increased cross-pin heat transfer caused by the circulation¹¹.

A heat source was applied to the fuel that was 1000 times lower than the baseline natural circulation case. The intention is to assess the shape of the temperature profile from the bundle centre to the duct and the difference between centre and duct in as close to isothermal a condition as possible. Using this, the fixed duct temperature can be varied to evaluate the dependence of the shape and temperature difference (hence effective conductivity) on array temperature, given the T^4 dependence of the thermal radiation contribution. This is an approach that has been used previously for modelling spent PWR fuel in dry storage, where thermal radiation transport plays an important cooling role (US DOE, 1996).

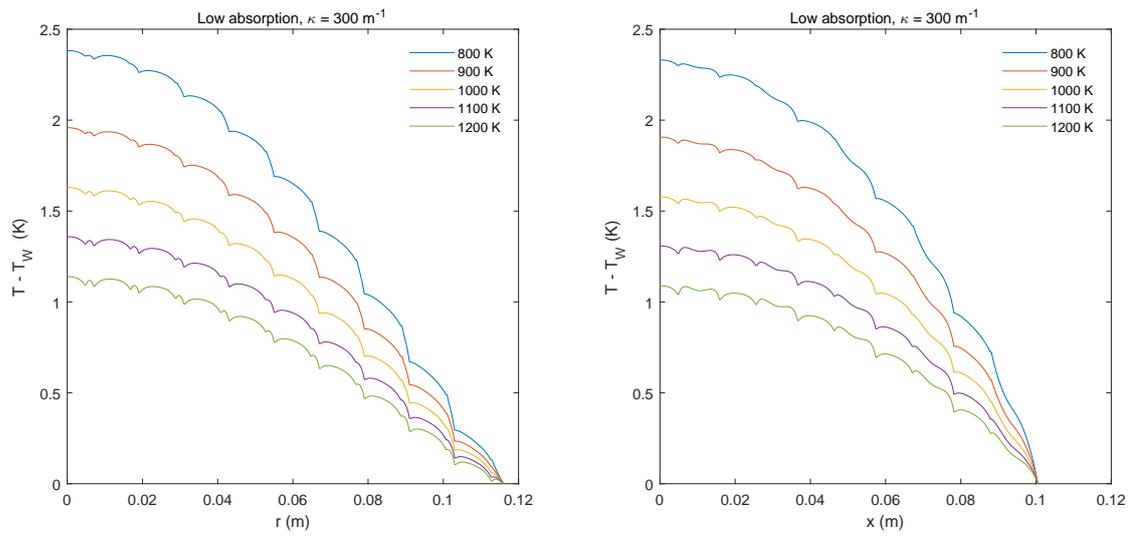
Four thermal radiation conditions were considered:

- Transparent: DO radiation, but with zero absorption.
- Low absorption: DO radiation, with an absorption coefficient of $\kappa = 300 \text{ m}^{-1}$.
- High absorption: DO radiation, with a higher absorption coefficient of $\kappa = 3000 \text{ m}^{-1}$.
- No radiation: The high absorption limit, where no radiation transport is modelled.

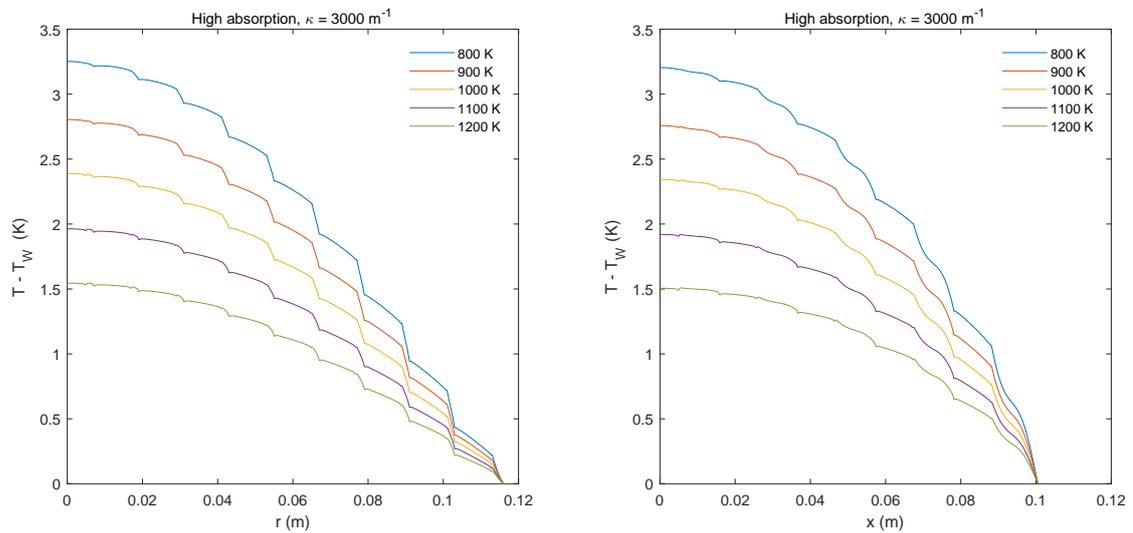
An assessment of the optical thicknesses for the high and low absorption case is given in Table 2.3. The cases with radiation used a refractive index of 1.3 and a surface emissivity of 0.8, as used in the 3D models. The wide range of conditions from transparent to opaque is assessed because

¹¹ Testing the derivation of planar conductivity without this enhancement showed that it added only a relatively small contribution to the result, so the use of a single representative value is an acceptable approximation.

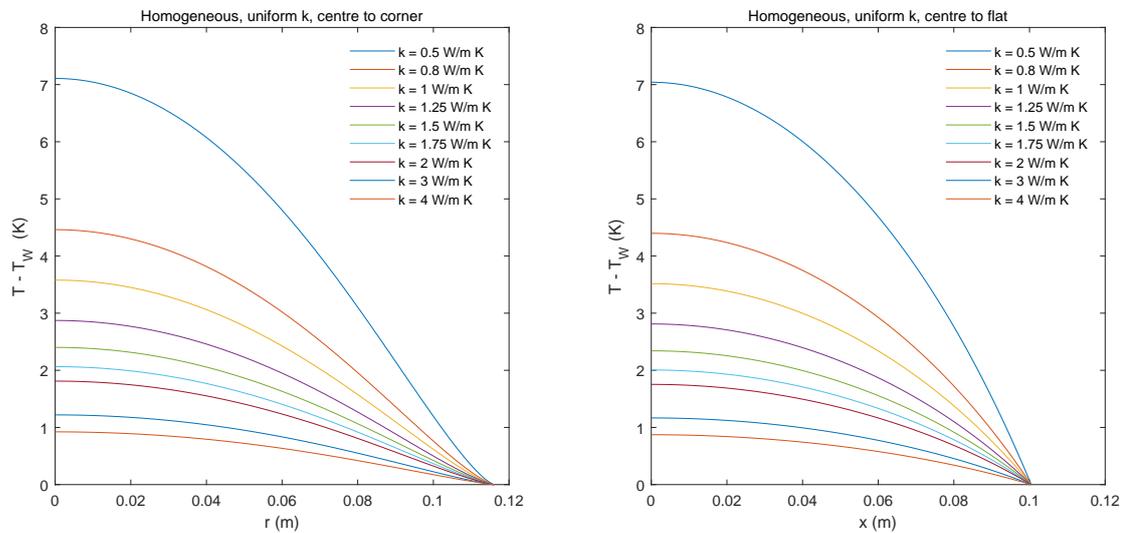
Performing the Analysis



a: Resolved pins (inhomogeneous model), $\kappa = 300 \text{ m}^{-1}$



b: Resolved pins (inhomogeneous model), $\kappa = 3000 \text{ m}^{-1}$



c: Various homogeneous conductivities.

Figure 3.19: Temperature profiles from the centre of the pin bundle slice to the corner of the duct (left) and to the middle of the duct wall (right).

Performing the Analysis

the absorption coefficient in the salts is unknown and variable. The approximation has been made for simplicity here that the fuel and coolant absorption coefficient are the same. The effect of independent variation in the absorption coefficients is assessed in Section 4.3.

Results for the low and high absorption cases are shown in Figure 3.19a and 3.19b for five duct temperatures from 800 to 1200 K each. The results are presented relative to duct wall temperature T_W for each temperature. Two results are plotted on each figure, one for the temperature profile from the centre of the bundle to the corner of the duct (along a symmetry plane) and the other in the x-direction from the centre to the middle of the flat face of the duct wall. Using these two results provides a cross-check and assurance that the representation of the effective conductivity does not exhibit a preferred direction when compared to a homogeneous conductivity (see below). The changes in temperature gradient with distance can be associated with passing through the various materials, noting that the heat source is only present in the fuel pin¹². There is a clear non-linear dependence of the centreline-to-duct temperature difference on duct temperature. These differences are substantially larger, for a given duct temperature, in the high absorption case compared to the lower absorption case, indicating that the effective conductivity is lower with more absorption.

The relationship between these simulations that contain all of the planar heat transfer mechanisms and an effective single value of conductivity at each temperature is established by creating a series of cases where the coolant, cladding, and fuel conductivities are all assigned the same constant value, and the same total heat source is distributed evenly through all three materials. Conductivity values from 0.5 W/m K to 4 W/m K were chosen and the results are shown in Figure 3.19c. The overall shape of the curves agrees with the inhomogeneous models for both the centre-to-corner and centre-to-duct results if overlaid, but the structure of the pins is not visible in the results because of the uniform conductivity and heat source. The range of centre to wall temperature differences covers the range seen in the inhomogeneous models.

The results for centreline to wall temperature differences vs homogenous conductivity are plotted in Figure 3.20 and a best fit is found with the following form

$$k_{eff} = a(T - T_W)^b + c$$

The best fit coefficients (a, b and c) and the relative fit error for the duct corner and flat results are shown in the figure legend. The corner and duct variants are effectively the same.

This relationship allows an effective conductivity to be evaluated at each duct temperature for the inhomogeneous models based on the predicted centreline to duct temperature. These relationships are plotted in Figure 3.21, where the transparent setup has the highest effective conductivity and greatest dependence on temperature, with increasing absorption producing lower and flatter curves. A third order polynomial (see Volume 2, Section 3.4.3.3) was fit to each result to allow its easy inclusion in Fluent.

¹² The differences in the details of the shapes of the temperature profiles between the $\kappa = 300 \text{ m}^{-1}$ and 3000 m^{-1} results is caused by repeated radiation absorption and (isotropic) re-emission, which allows the radiation to transport energy around the fuel pins through the coolant 'beyond line-of-sight', as well as by reflection and absorption/emission from the cladding surfaces. This is most prominent in the intermediate optical thickness $\kappa = 300 \text{ m}^{-1}$ case. This is relevant in the full assembly case near to the duct, where there are high temperature gradients over the outer row of pins.

Performing the Analysis

Although these relationships were derived using an approximately isothermal whole bundle cross-section, they are also applicable to porous model cases where there are significant localised pin-to-pin temperature gradients, caused by large planar heat fluxes – the local effective planar conductivity would depend on temperature in the vicinity of each pin.

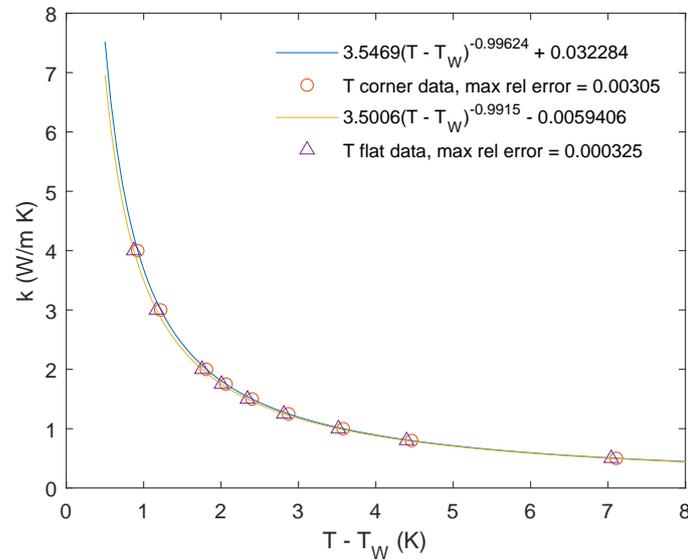


Figure 3.20: Relationship and curve fit between centreline to wall temperature difference and homogeneous conductivity.

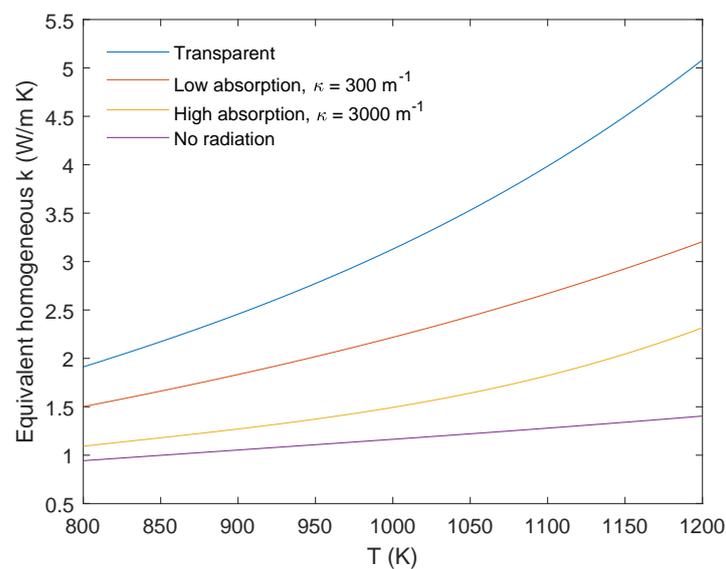


Figure 3.21: Effective planar conductivity as a function of duct temperature for the four considered radiation conditions.

4 Application of the Results

The models described in Section 3 and the information derived from them can now be used for two purposes:

- Implementing a porous representation of a fuel assembly that is suitable for inclusion in a reactor scale simulation.
- Exploring the effect of uncertainty in material properties on quantities of interest.

The rest of this section shows example of how to accomplish these applications. However, before describing these activities, which would be the point at which this modelling would start to be applied to making real decisions, it is worth recapitulating some of the caveats and limitations on what has been derived so far.

- Wire wraps have been ignored in this modelling for simplicity. Real fuel assemblies will have wire wraps or spacer grids as part of their construction, which will need to be accounted for. There will also be inlet and outlet structures (filters, nozzles and fuel supports) that will contribute to the pressure loss across the core, that would also need to be present in a whole assembly model.
- The position reached regarding the appropriate choice of RANS turbulence model could benefit from further investigation and refinement. For example, the γ intermittency model showed promise, but would need further assessment of its robustness and generality to make use of it.
- Rigorous mesh sensitivity studies have not been undertaken, and would be a necessary part of making use of the results to make decisions. Examples of mesh sensitivity studies are described in Study A and Study D.

These points notwithstanding, the purpose of the case study has not been the determination of a specific model and its numerical results. Instead, its intention has been to demonstrate the process of planning and executing the simulations, extracting the results and using them to derive characteristics of the fuel bundle, and then the implementation of those characteristics, with comparison of their predictions with a reference solution.

Application of the Results

4.1 Porous Assembly

To test the implementation of a porous assembly model, the same baseline forced and natural circulation conditions as evaluated by the resolved assembly model in Section 3.3 were assessed. By carefully aligning the mass flow rates and heat source between the two models, a detailed comparison can be made of the quality of the agreement between the approaches.

The porous 1/6th assembly model covers the same region of space as the resolved model. As shown in Figure 4.1, the domain was split into three segments axially:

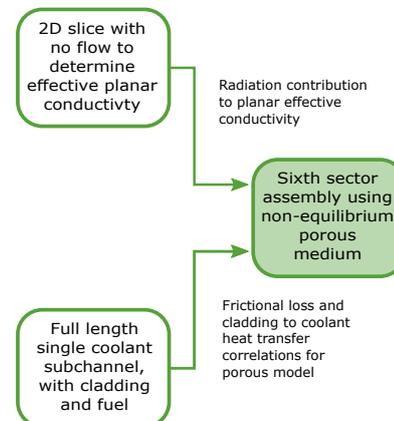
- A 10 mm lower region below the rod bundle.
- The 1600 mm rod bundle region.
- A 100 mm upper region above the rod bundle.

and into four regions radially:

- r1: The majority of the cross-section from the centre to the mid-line of the final row of pins. This region contains a heat source in the rod bundle axial region only, and has the same geometrical characteristics (porosity, hydraulic diameter and surface area density, influencing pressure loss and heat transfer) as the subchannel models.
 - r2: From the mid-line of the final row of fuel pins to the outer edge of the final row. This region also contains a heat source in the rod bundle axial region, but the change in coolant path geometry caused by being adjacent to the duct results in a different porosity, hydraulic diameter and surface area density.
 - r3: From the outer edge of the final row of fuel pins to the duct. This region contains no heat source, and has different geometrical characteristics to r1. The r2 and r3 regions form a 'wall subchannel' between the centres of adjacent fuel pins and the duct, the characteristics of which¹ are available in the literature (Kakaç *et al.*, 1987, Chapter 3).
- gap: The duct is represented by a thin two sided wall so the gap region represents the combination of the duct and coolant in the gap. There is no heat source, and the geometrical properties are those of the gap, not the bundle.

The mesh contains 46,000 cells for the fluid, with the same number of solid cells co-incident with them (an identical overlaid mesh) to represent the solid parts of the region². This is necessary for the non-equilibrium porous approach, where the solid can have its own temperature field that is different to the flow fluid temperature. This requires the surface area density (interface surface area per unit volume), A_{fs} , of the fluid-solid interface and the associated HTC, h_{fs} , to be specified.

The axial splits are required to place the heat source in the correct places, and the gap requires a separate region because it is not connected to the bundle. However, splitting the bundle into r1 to r3 is not essential – it was performed to test various options for modelling the behaviour of the



¹ Compared to the 'central subchannel' assessed in Section 3.5. There is a third 'corner subchannel' type at the duct corner that has also been analysed in the literature, but the distinction has not been considered here.

² 92,000 cells in total, compared to 44 million cells in the resolved model, a factor of approximately 480 less.

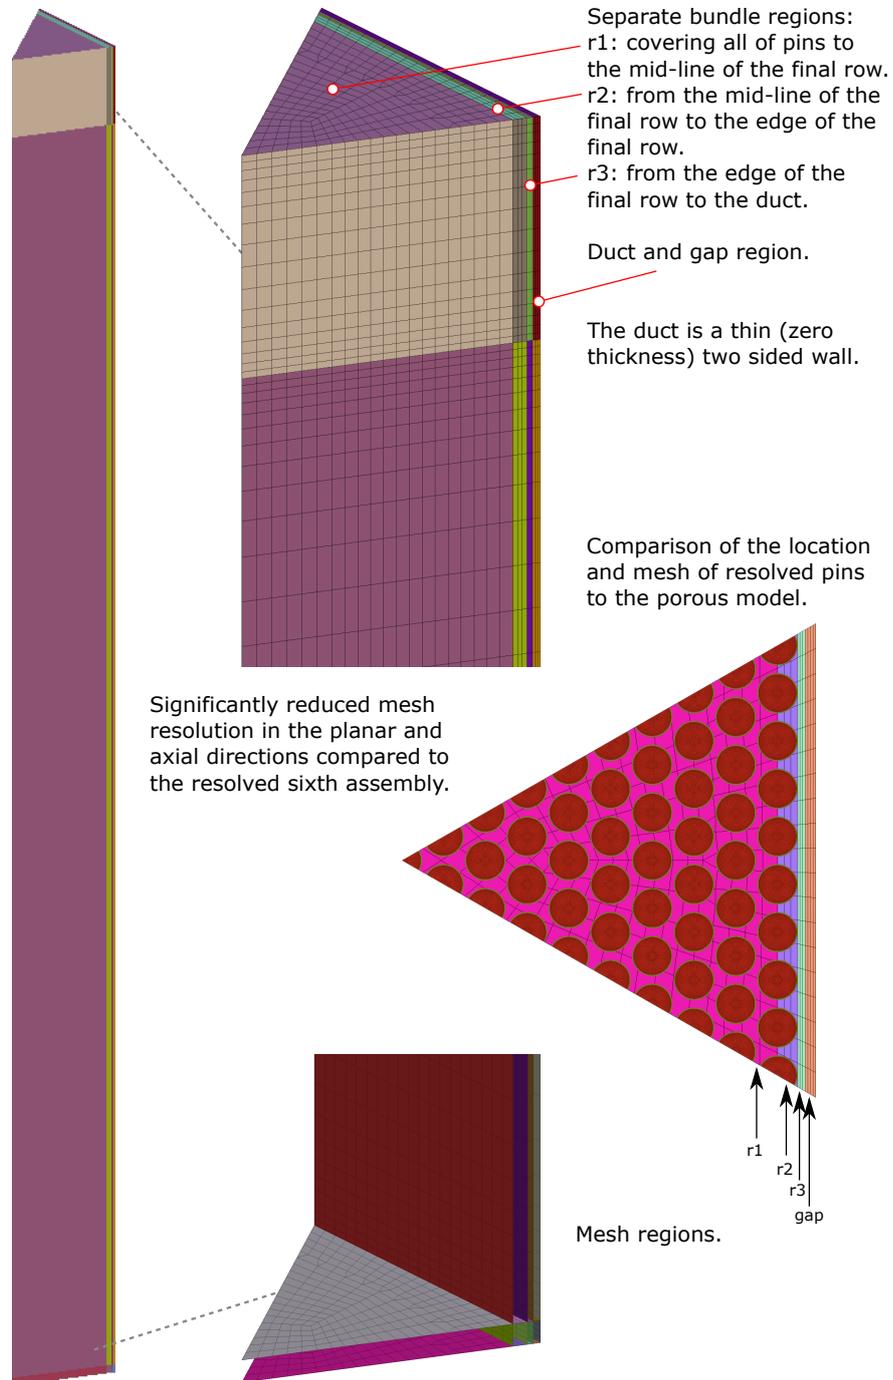


Figure 4.1: Geometry and mesh of the 1/6th porous fuel assembly.

Application of the Results

coolant in the near duct region. This followed the descriptions of Wang *et al.* (2020), who were concerned with duct temperatures and inter-wrapper heat transfer in a LMFR application. They used a separate zone between the rod bundle and the duct too, that treated the equivalent of r3 as a pure fluid layer (e.g. without any porous modelling).

The details of the chosen setup for the porous model were determined by a series of precursor simulations with representative losses and heat transfer. The modelling choices and findings were:

- The physical velocity formulation was assessed as well as the default superficial velocity formulation. The guidance in the solver documentation (ANSYS, 2020b), suggests that the physical formulation is able to resolve velocity gradients more accurately. The physical velocity, U_p , in a porous region is related to the superficial velocity, U_s , by the porosity, ϕ , (the fraction of the volume that the fluid occupies) by

$$U_s = \phi U_p$$

Regardless of which formulation is used, however, Fluent requires that inlet velocities are superficial, and that the porous loss coefficients are specified as being relative to the superficial velocity.

- Attempting to run with r3 represented by a clear fluid (no solid) or using the physically realistic higher porosity of r2 and r3 was problematic when using the physical velocity formulation. The flow tended to traverse across the bundle at the inlet and concentrate in r3. The topic of the physics and modelling of flow parallel to a porous/non-porous interface is well described in the literature (for example Ochoa-Tapia and Whitaker, 1995, de Lemos, 2012 or Al-Aabidy *et al.*, 2020) but involve a relatively complex implementation. It is possible that using these methods and resolving the mesh in the radial direction in r2 and r3, to properly capture the shear and turbulence would be successful, but this would be counterproductive to the aim of minimising cell count and model complexity, so they have not been employed here.
- A slip (zero-shear) condition was applied to both sides of the duct wall. The friction associated with the duct is included in the porous resistance term already, and there are not enough cells to resolve the near-wall gradients.
- Applying a viscosity scaling function (such as the Brinkman correction) was found to make little difference, and was omitted.

From these initial investigations, and the fact that each additional per-region setting adds to the complexity of including them at large scale in a reactor model, a relatively simple approach was taken. For the forced circulation case:

- The r1 region was represented using the subchannel values of ϕ , D_h and A_{fs} , and the f and Nu correlations derived from the subchannel model in Section 3.5.
- Two successful configurations for modelling r2 and r3 were found:
 1. Using the physical velocity formulation for all regions and with r2 and r3 represented using their own values of D_h and A_{fs} , but the same value of ϕ as r1.
 2. An alternative configuration, using the superficial velocity formulation for all regions and the physically accurate value of ϕ for r2 and r3 (as well as their own values of D_h and A_{fs}).

Application of the Results

The same f and Nu correlation as r1 was used in each configuration. Using the correct value of ϕ in r2 and r3 with the physical velocity formulation resulted in unphysical discontinuities in the turbulence and velocity fields at the r1 to r2 interface. This is largely caused by the fact that by default Fluent does not account for the presence of the porous medium in the modelling of turbulence at all. An option to avoid this is to run the model with a laminar setup (removing all turbulence); this was successful in removing the discontinuities, and the results were similar to, but not quite as good as, the 'alternative' configuration above, which used the superficial velocity formulation.

- The gap region used its own value of ϕ , D_h , A_{fs} and the same f and Nu correlation as r1 to r3 for the turbulent flow conditions.

For the natural circulation case:

- The r1 region was represented using the subchannel values of ϕ , D_h and A_{fs} , and the fRe and Nu value derived from the subchannel model in Section 3.5.
- The r2 and r3 regions were represented using their own values of ϕ , D_h and A_{fs} , and the same fRe and Nu values as r1.
- The gap region used its own value of ϕ , D_h , A_{fs} and $fRe = 96$ and $Nu = 140/17$, which are analytical results for laminar flow in an infinite parallel plate channel with constant heat flux (Kakaç *et al.*, 1987, Chapter 3).

Some of the approximations made are likely to lead to some reduced accuracy of the flow and heat transfer in the near-duct region. If the details of the near-duct region were to become of increased importance and interest, the modelling details in r2, r3 and the gap could be revisited.

4.1.1 Porous Model Implementation

More precursor calculations and modelling parameter evaluations were necessary to create a porous model in CFD than explicitly resolving the geometry. These are areas where both conceptual and mathematical errors are likely, and so require particular attention in verification.

Geometrical Characteristics: The expressions and values for ϕ , D_h and A_{fs} for the four regions are given in Table 4.1. The wall subchannel expressions for r2 and r3 use $w = (2/3)h$, which is the distance covered by r2 and r3 (centre of the last pin to the duct, Section 2.2.1). The gap geometry has been calculated by treating it as an infinite parallel plate channel, which ignores corner effects.

Boundary Conditions and Source Term: The inlet and outlet conditions for the porous model are the same as the resolved model, except that the velocity specified is the superficial velocity across the entire inlet face, giving for example 0.2597 m/s compared to 0.6638 m/s (Table 2.2). Care is needed with the flow areas to ensure that exactly the same inlet mass flow rate as the resolved model is introduced.

The heat source was distributed evenly across the solid regions of r1 and r2, using the same axial power profile for scaling with height as the resolved case. Care was also needed with the details of the cross-sectional areas of the regions to introduce exactly the same power input as the resolved case. It should be noted that the source terms are related to the geometrical volume of the region, and are not scaled by $1 - \phi$ in the solid.

Application of the Results

	r1	r2 and r3	gap
A_f	$\sqrt{3}p^2/4 - \pi d^2/8$	$(w - d/2)p - \pi d^2/8$	
A_o	$\sqrt{3}p^2/4$	$(w - d/2)p$	
P_f	$\pi d/2$	$\pi d/2 + p$	
ϕ	A_f/A_o	A_f/A_o	$w_{gap}/(w_{gap} + 2t_{duct})$
D_h	$4A_f/P_f$	$4A_f/P_f$	$2w_{gap}$
A_{fs}	P_f/A_o	P_f/A_o	$2/(w_{gap} + 2t_{duct})$
ϕ	0.37021	0.52766	0.60784
D_h (m)	0.0058783	0.0063330	0.0062
A_{fs} (m ² /m ³)	251.92	333.28	392.16

Table 4.1: Expressions for the geometrical characteristics of regions and their resulting numerical values. A_f is the cross-sectional area available for coolant flow in a subchannel and A_o is the overall cross-sectional area of the subchannel, including the region occupied by cladding or fuel.

Equivalent Material Properties: The coolant properties can be used unmodified in the porous model and in the gap, as Fluent will correctly weight the contributions to density, specific heat capacity and thermal conductivity from the coolant and the HT9 material of the duct by ϕ and $1 - \phi$ respectively. If modelling the heat transfer across the inter-wrapper gap becomes more of a focus, then the effective conductivity in the planar direction would need to be modified compared to the axial conductivity to represent the layout of the duct and coolant gap (as shown in Volume 2, Figure 3.5).

The solid in the bundle region is intended to represent the combined effect of the fuel salt and cladding as an equivalent solid material. The conductivity is split into an axial component and a planar component represented using an orthotropic material.

The effective axial conductivity is evaluated from material properties only, so, as a simplification, no contribution from the natural circulation of the fuel is included.

$$k_{equiv,axial} = \frac{k_{clad}A_{clad} + k_{fuel}A_{fuel}}{A_{clad} + A_{fuel}}$$

It should be noted that the k properties here are temperature dependent polynomials, and k_{clad} is a piecewise function, so the resulting equivalent property must be algebraically or numerically derived to be a piecewise polynomial in T too.

The effective planar conductivity, including thermal radiation, k_{eff} , has already been derived in Section 3.6, and the function derived for $\kappa = 300 \text{ m}^{-1}$ was used to align with the baseline conditions. This function already accounts for the fact that the cladding and fuel only occupy $1 - \phi$ of the cross-section, and contains the contribution from the molecular conductivity of the coolant. However, Fluent will perform a porosity weighting of the solid property that would 'double count' the fluid contribution. This is anticipated by specifying the conductivity of the solid in the planar directions to be

$$k_{equiv,planar} = \frac{k_{eff} - \phi k_{coolant}}{1 - \phi}$$

It should be noted that this expression is strictly only valid for an equilibrium porous model (where there is no difference in local fluid and solid temperature) and it is also only ever an approximation to segregate the fluid and solid effects; conduction and thermal radiation in the coolant will be altered

Application of the Results

by the velocity and temperature profiles when there is flow, including by the additional transport created by turbulence.

Calculating effective values for density and specific heat capacity have a further subtlety – it is not possible to specify a temperature dependent density for a solid in Fluent (because no thermal expansion is modelled), however, it is only the product ρc_p that matters (and only during a transient simulation). The equivalent density is evaluated as

$$\rho_{equiv} = \frac{\rho_{clad} A_{clad} + \rho_{fuel} A_{fuel}}{A_{clad} + A_{fuel}}$$

and a nominal value of this, ρ_n , is established at a representative temperature (say, 1000 K). This is used as the material property in Fluent, and also to derive a modified equivalent specific heat capacity, which can vary with temperature when implemented

$$c_{p,equiv} = \frac{\rho_{clad} c_{p,clad} A_{clad} + \rho_{fuel} c_{p,fuel} A_{fuel}}{\rho_n A_{clad} + \rho_n A_{fuel}}$$

This evaluates to the correct temperature varying function as $\rho_n c_{p,equiv}$ because ρ_n is cancelled out. This is also needs to be derived as a piecewise polynomial because of the piecewise expression for the HT9 specific heat capacity.

Pressure Loss: The resistance (pressure loss) caused by the porous medium is implemented as a momentum source term in Fluent in terms of superficial velocity U_s :

$$S_i = - \left(\sum_{j=1}^3 D_{ij} \mu U_{si} + \sum_{j=1}^3 C_{ij} \frac{1}{2} \rho |U_s| U_{si} \right)$$

Where S_i is the momentum source (N/m^3) in direction i , D_{ij} is the viscous loss tensor (m^{-2}), and C_{ij} is the inertial loss tensor (m^{-1}). The directions used are the Cartesian axes, with the axial loss in the z direction, and the planar loss coefficients in the x and y directions, so there are no off-diagonal terms (D and C are diagonal). Only the axial losses are of importance in the fuel assembly (there is negligible cross-flow), and the transverse direction losses would not be substantially different, so only a single value of either D or C needs to be specified for all three directions (making the approximation that the losses are isotropic). These assumptions and approximations will not always be valid, depending on the geometry and application/flow.

Comparing this equation to the expression for pressure loss per unit length in a channel flow (also with units N/m^3)

$$\frac{dP}{dx} = \frac{f(Re_p)}{D_h} \frac{1}{2} \rho |U_p| U_p$$

This equation is framed in terms of the physical velocity, which is also used to evaluate the Reynolds number dependent friction factor, as derived in Section 3.5. For the turbulent forced circulation cases, only the inertial term will be used, so an expression for C is required; for the laminar natural circulation cases, only the viscous term will be used, so an expression for D is required. More complex loss characteristics could use both viscous and inertial terms simultaneously.

Equating the expressions for the inertial term and relating the superficial and physical velocities

Application of the Results

using $U_s = \phi U_p$ gives

$$C \frac{1}{2} \rho |U_s| U_s = \frac{f(Re_p)}{D_h} \frac{1}{2} \rho |U_p| U_p$$

$$C = \frac{f(Re_p)}{D_h} \frac{|U_p| U_p}{|U_s| U_s}$$

$$C = \frac{f(Re_p)}{D_h \phi^2}$$

Algebraically expanding the form of $f(Re_p)$ further is not helpful because the expression cannot be directly used in Fluent, and to retain the Reynolds number dependence derived in Section 3.5, it must be implemented as a User Defined Function (UDF). Similarly equating the viscous term

$$D \mu U_s = \frac{f(Re_p)}{D_h} \frac{1}{2} \rho |U_p| U_p$$

expanding the laminar friction factor expression $f = fRe/Re$, and expanding $Re = \rho |U_p| D_h / \mu$

$$D \mu U_s = \frac{fRe}{Re D_h} \frac{1}{2} \rho |U_p| U_p$$

$$D \mu = \frac{fRe \mu}{2 \rho |U_p| D_h^2} \rho |U_p| \frac{U_p}{U_s}$$

$$D = \frac{fRe}{2 D_h^2 \phi}$$

This is a single number, where the viscosity term has been cancelled, so can be implemented in the Fluent interface without a UDF.

Heat Transfer: The heat transfer between the fluid and solid at any point in the porous region is a volumetric energy source term

$$q_{vol} = A_{fs} h_{fs} (T_f - T_s)$$

and

$$h_{fs} = \frac{Nu k}{D_h}$$

Because k varies with temperature, h_{fs} cannot be specified as a single number and requires a UDF, even when Nu is assumed constant (in the natural circulation cases). Where Nu depends on the local value of Pr and Re (in the forced circulation case), then h_{fs} must be specified as a UDF implementing the relationships derived in Section 3.5.

The solid contains the heat source, which at every location is transferred to the fluid via the HTC and surface area density, by a temperature difference. From this equation, it can be seen that the interpretation of the solid temperature, T_s , in the non-equilibrium porous formulation is that it is the cladding outer surface temperature necessary to dissipate the internal source. There is no representation of the 'inside' the solids. This is only conceptually accurate at steady-state, and could potentially be significantly inaccurate in a transient analysis for two reasons:

- The porous solid temperature is assumed to be representative of all of the material in the cross-section of a fuel pin with an equivalent heat capacity, ρc_p . From the results of the resolved rod cases (compared in more detail below) the internals of the fuel are significantly

Application of the Results

hotter than the surface because they contain the heat source. There is a quantity of energy that would need to be accumulated or lost in a transient that represents that elevated temperature, and this would affect the timescale of the transient response. The presence of this 'stored' energy would not be accounted for using the setup described here.

- The uniform temperature of the solid implies that it has an infinitely high conductivity (zero Biot number) at the scale of individual pins, because there are no local/internal temperature gradients represented. In a resolved model, if the coolant flow was suddenly increased, then the fuel pin surface would drop rapidly, but the internal temperature profile of the fuel would not respond instantaneously. In a porous model, the thermal mass of the whole pin cross-section would notionally need to be cooled to reduce its temperature, which would happen more slowly. Therefore the transient solid temperature would no-longer represent the local surface condition, and the predicted rate of heat transfer to the coolant would be inaccurate.

Modelling this more accurately within the porous model framework would require a combination of modified equivalent material properties for ρc_p and tracking an additional variable (or variables) representing the internal state (temperature profile) of the pins.

Porous Model Solution: Similar initialisation and numerical settings were applied to the porous model as to the resolved model, and some of the same incremental solution strategy steps were used, but less caution was needed. It was found that in the natural circulation case, starting the solution with turbulence active and only making it laminar after the initial solution stage was a reliable approach. Conversely, it was found that when using the superficial velocity formulation (in the 'alternative' forced circulation configuration), this condition was harder to converge successfully and a period of laminar solution at the start of the simulation before activating turbulence modelling was successful in this case.

4.1.2 Comparing Porous and Resolved Assembly Models

The porous and resolved full assembly model have deliberately been configured to have the same mass flow rate and heat source. Therefore, assessing the agreement between the models in coolant velocity and temperature does not provide significant insight into the accuracy of the porous modelling approach, but it does allow a verification cross-check that the porous implementation is correct. Comparisons of temperature prediction of the solid in the porous medium to the cladding outer surface in the full model, and an assessment of the overall pressure drop do test the accuracy of the approach.

Correctly predicting overall pressure drop is important because it will determine the flow in each fuel assembly in a reactor scale model in response to the above and below core conditions, which is likely to be the most important result. The static pressure difference, dP in the coolant from top to bottom of the fuel pin section has two contributions:

$$dP = dP_{hydro} + dP_{friction}$$

The hydrostatic pressure gradient increases the pressure with height (and so dP_{hydro} is negative). Flow from top to bottom in the forced circulation case leads to a positive pressure difference for the frictional loss term, $dP_{friction}$. The natural circulation case, with flow from bottom to top, gives a neg-

Application of the Results

ative $dP_{friction}$. The pressure differences and their contributions are shown in Table 4.2, comparing the full resolved assembly, single subchannel and porous model.

	Full	Single	Porous	Porous (alt)
Forced circulation				
dP_{hydro}	-45197	-45204	-45176	-45132
$dP_{friction}$	12574	12814	13751	12479
dP	-32623	-32390	-31425	-32651
Natural circulation				
dP_{hydro}	-45442	-45440	-45368	
$dP_{friction}$	-490	-468	-590	
dP	-45931	-45908	-45958	

Table 4.2: Comparison of pressure difference (N/m^2) contributions for different model implementations.

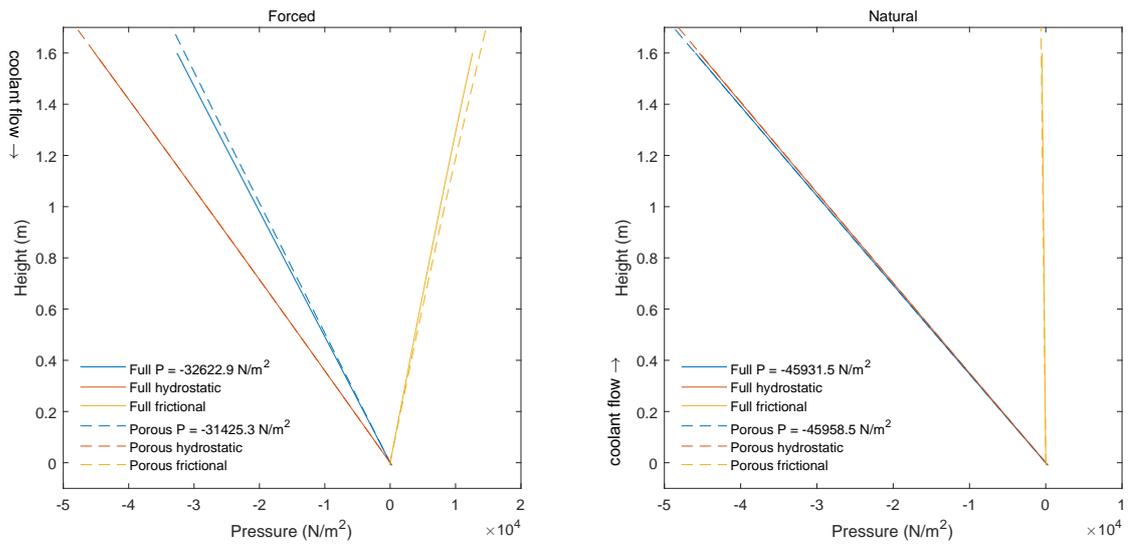
These results are visualised in Figure 4.2a, along with the vertical profile of coolant density and temperature, comparing the porous to the full assembly.

- The hydrostatic part of the pressure difference compares well between the full and porous models. It is the integral of the density field with height (Figure 4.2b), and this density field arises from the coolant temperature field (Figure 4.2c).
- The coolant temperature field for the forced circulation case agrees well over the fuel pin height. For the natural circulation case, the gradient with height is similar, but has an offset that appears at the base where the coolant enters. The porous model does not represent the entrance effects in the coolant flow paths, nor the fuel pin recirculation end effects that occur in the full model, leading to a different local heat transfer.
- The natural circulation case has a very low frictional loss – this loss balances the differences in hydrostatic driving force produced by temperature differences in the rest of the reactor to produce natural circulation.
- The dP for the ‘alternative’ configuration for the forced circulation case (using the superficial velocity formulation and accurate value of ϕ in r2 and r3) is also shown in Table 4.2 where it can be seen that it produces good agreement with the full assembly. However, it produces worse agreement with the flow distribution across the cross-section and with the coolant and solid temperatures (these are not reported).

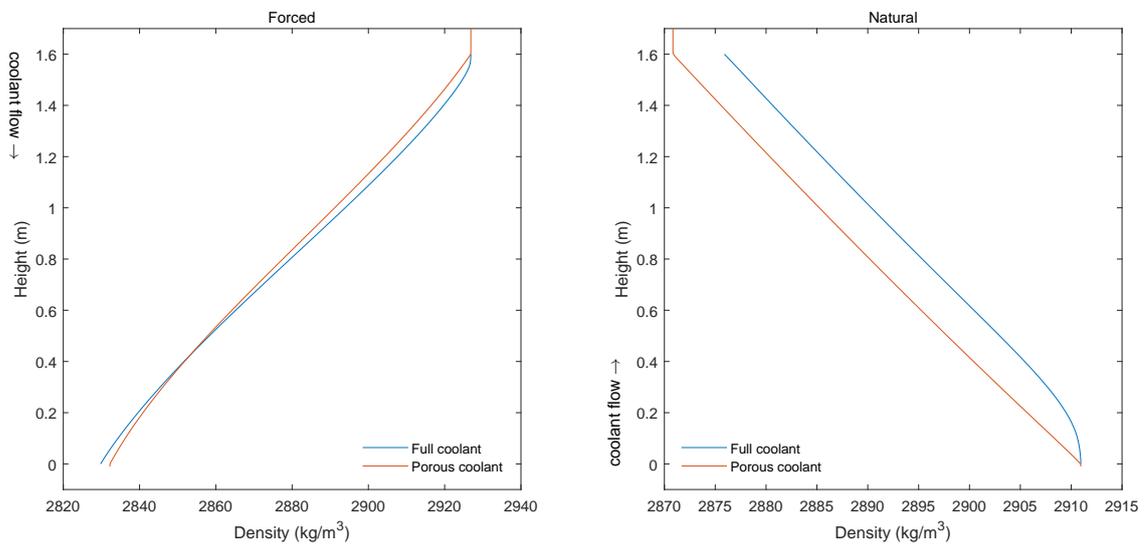
Cross-sections of the temperature and velocity fields are shown in Figure 4.3, which can be compared their equivalents in Section 3.3.4, but the porous results produce little visible structure to the flow. The per-cell solution values are shown instead of smooth fields (which are interpolated by the post-processor) to emphasise the coarseness of the mesh, and also because interpolation of such coarse results can introduce misleading visual artefacts.

Radial line profiles of the velocity and temperature fields produce more meaningful results, and results at three heights are shown in Figures 4.4 and 4.5 for the forced and natural circulation cases respectively.

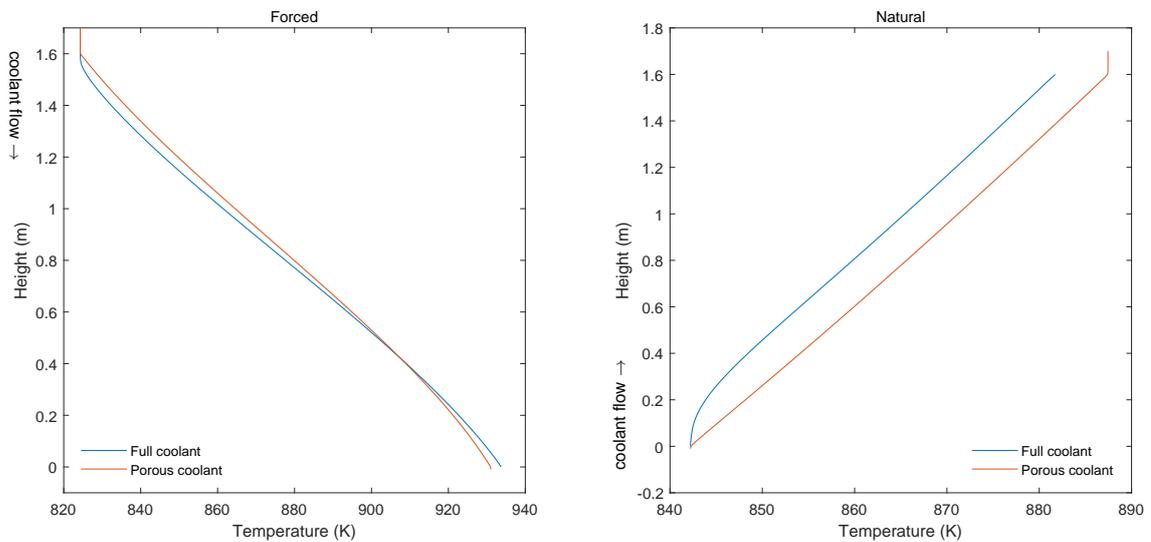
Application of the Results



a: Comparison of overall dP and hydrostatic and frictional components.



b: Comparison of coolant density.



c: Comparison of coolant temperature.

Figure 4.2: Comparisons of centre subchannel of the full assembly model to the centre of the porous model for the forced (left) and natural (right) circulation cases.

Application of the Results

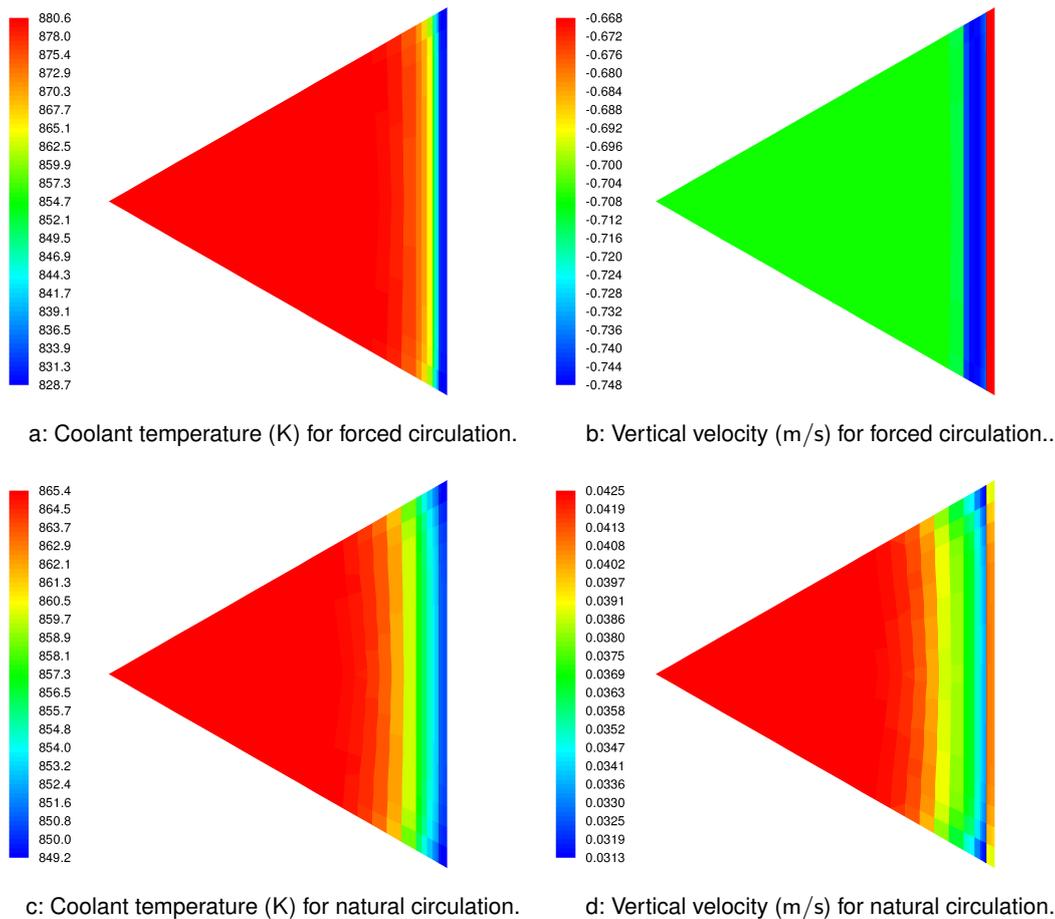


Figure 4.3: Cross-sections of coolant temperature and vertical velocity in the porous assembly model at $z = 0.8$ m (mid-height).

- The agreement with the coolant velocity is relatively good across the cross-section, and the locations near to the duct where it is either increased or reduced compared to the main bundle in the full model are captured, although a qualitative comparison has not been made. The velocities are comparable because the physical velocity formulation is used; for the superficial formulation, the porous model results would need to be divided by their local value of ϕ .
- The agreement with coolant temperature in the main part of the bundle and the reduction toward the duct is well captured, but the extent of the reduction at the duct is under-estimated in the porous fluid.
- For the natural circulation case, the agreement between the solid temperature in the porous model and the full model is good, and the reduction in pin temperature towards the duct is represented. The interface between the coolant and the outer surface of the cladding in the full model is marked with crosses for each pin. The effect of the combination of conduction in the fuel cladding and radiation can be seen in the full model, where there is a small temperature offset between the inward facing side of each pin compared to the outward facing side. The difference at each height between the porous fluid and solid temperature indicates that the choice of the non-equilibrium approach was successful.

Application of the Results

- For the forced circulation case, the porous solid temperature is generally in good agreement with the full model – the differences are larger at higher values of z in the bundle (closer to the coolant inlet). Contributions to the transverse thermal transport from turbulence are lower in the porous model and likely to introduce some inaccuracy. Turbulence is not correctly affected by the presence of the solid, and is under-estimated, resulting in lower thermal diffusivity, including in the near duct region.
- The ‘alternative’ forced circulation configuration (not shown), exhibits lower turbulence in the bundle. Its coolant velocity prediction in the main bundle section is lower and does not overlap with the full model in each coolant passage, and the solid temperatures are higher. It does, however, produce better dP agreement.

The relatively good comparison of cladding outer temperature is reassuring that the heat transfer processes have been characterised correctly, but they do not allow any prediction of the internal state of the fuel pin. An approach to obtaining such a prediction is described in Section 4.2 below.

These comparisons demonstrate the process of ‘closing the loop’, where the resolved model has been evaluated, a porous model created, and then compared back to the resolved model as a reference. More improvements to the porous model could be made to improve the agreement, especially in the duct region if local temperatures are of interest, or to tune the loss representation to produce closely matched pressure drops.

This process could be done in an ad-hoc manner, or could be automated using a tool such as Dakota to perform an optimised calibration. However, at each possible iteration of model refinement, the question should be asked about how much additional effort is justified? How much more time and cost in porous model development, with potentially increased complexity in setup of a reactor scale model, is worth how much extra accuracy? Particularly when weighed against other sources of inaccuracy or uncertainty.

Application of the Results

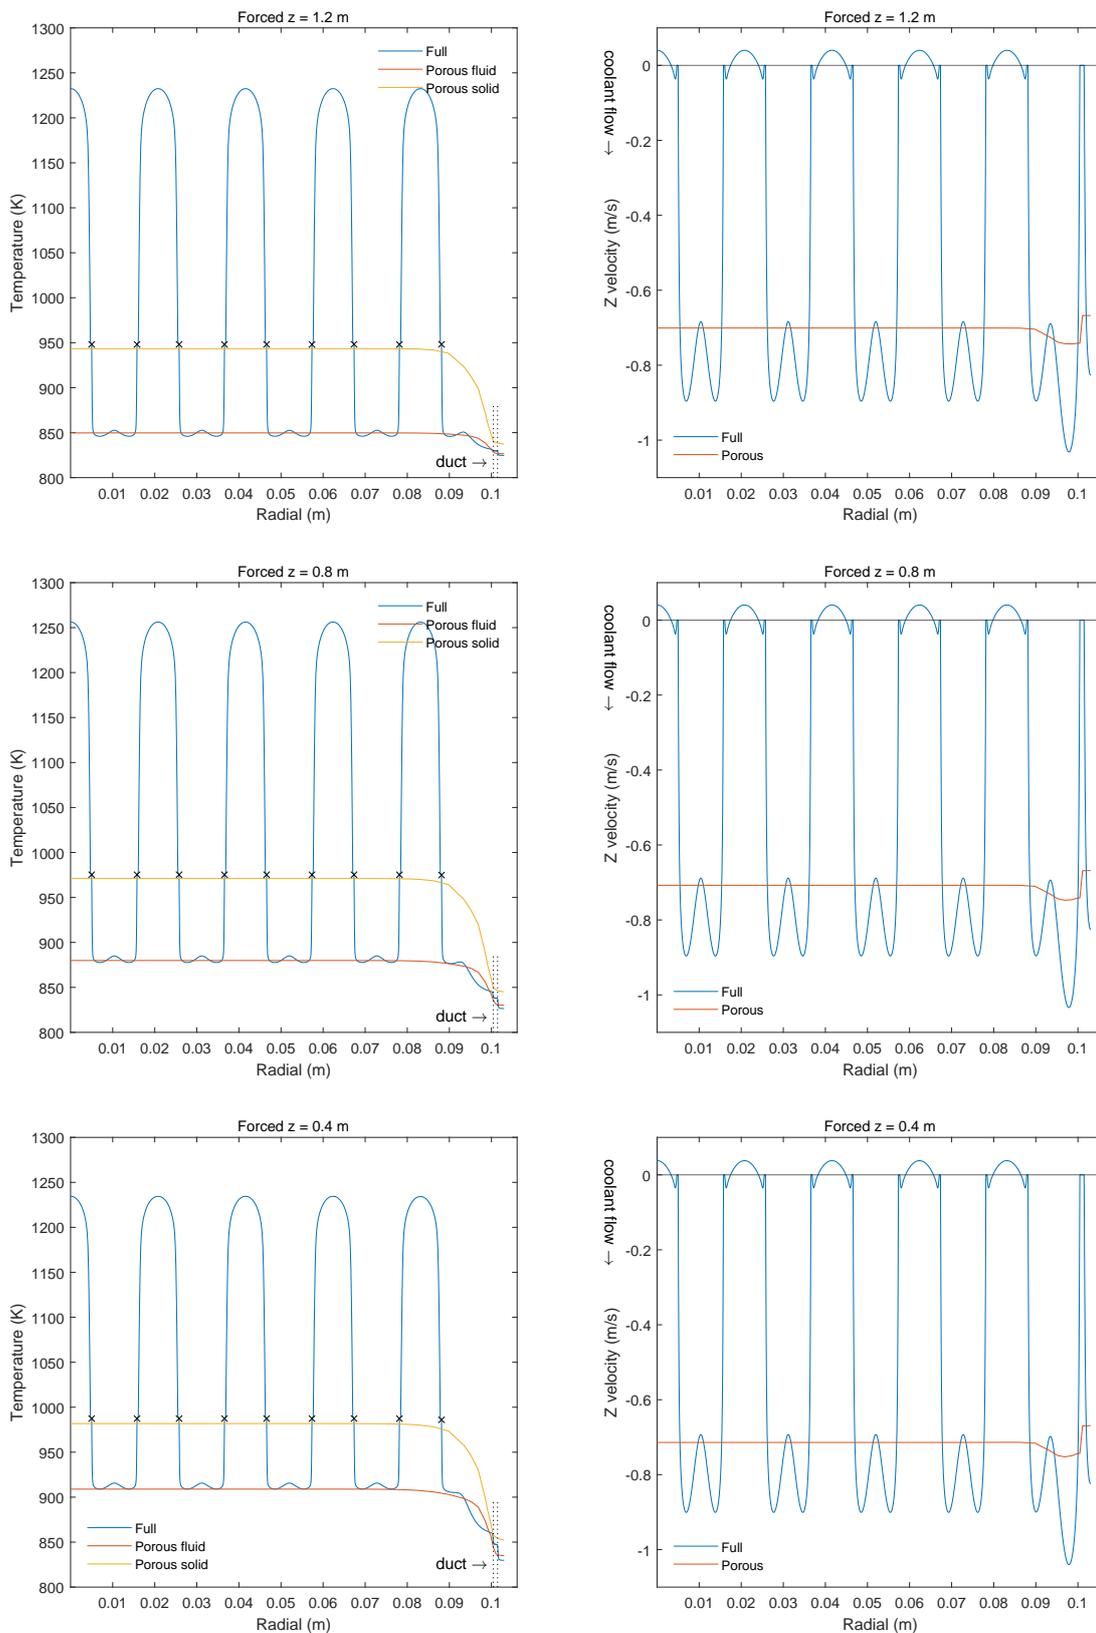


Figure 4.4: Comparison of resolved and porous assembly radial profiles of temperature and axial velocity at three heights for forced circulation case.

Application of the Results

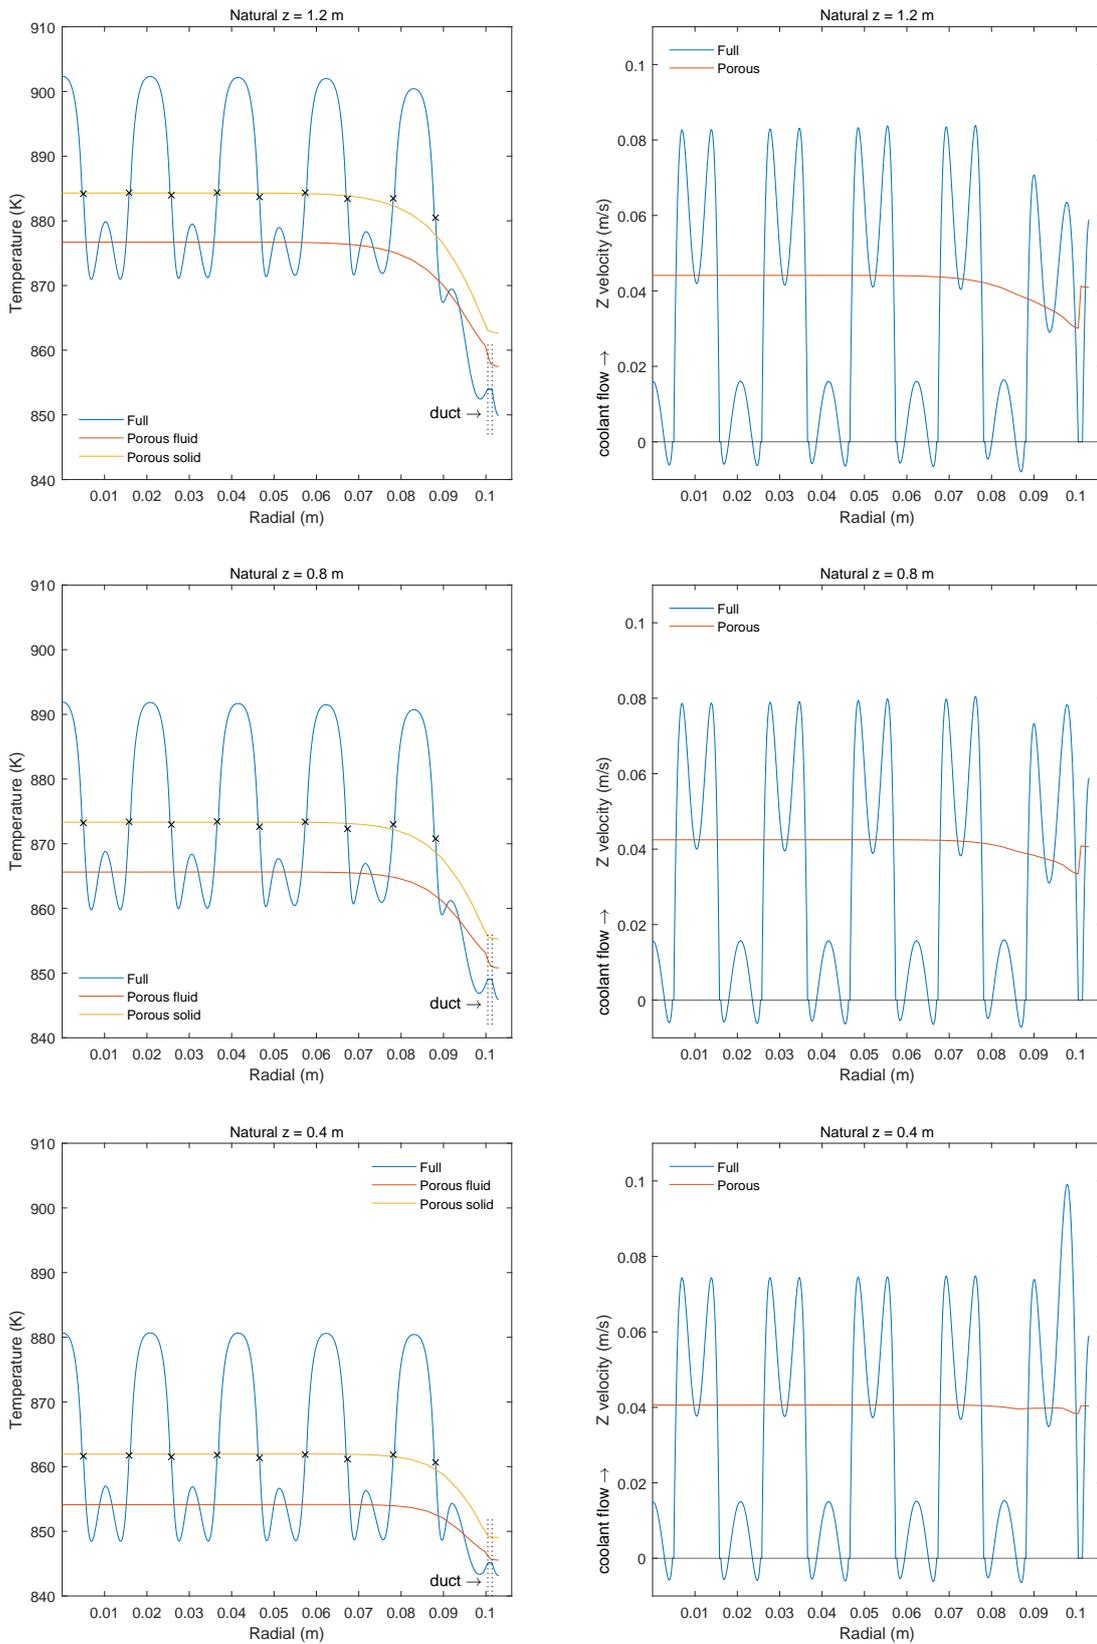
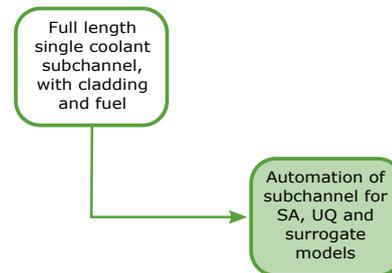


Figure 4.5: Comparison of resolved and porous assembly radial profiles of temperature and axial velocity at three heights for natural circulation case.

Application of the Results

4.2 Surrogate Model of Fuel Conditions

The porous model described above gives a good prediction of cladding temperatures over the whole assembly, but no internal details of the fuel. It is often the case that only the maximum values are of interest. This section demonstrates the construction of a surrogate model where, using the single subchannel case (Section 3.5), the flow rate, decay heat and inlet temperature will be varied over values relevant to the SBO conditions, and a model representing the maximum fuel salt temperature will be derived.



Given the assumption of quasi-steady conditions³ then the surrogate model could be used as an auxiliary calculation to a porous or system code model to post-process the expected maximum temperature at any instant, or it could be used on its own in a design optimisation.

The Dakota toolkit was used to automate the selection of the cases to be run and also generated the surrogate model – a quadratic polynomial response surface model was chosen.

- A polynomial response surface model is expected to perform well with the smooth variation of output quantity, but may give inaccurate results at extremes of the variable range, and is liable to be invalid for extrapolation outside of them.
- A Gaussian Process model could also have been chosen, and would have the advantage of providing an estimate of the uncertainty at each point. However, it would be more complex to evaluate, and the polynomial model performance is sufficiently accurate to make the per-location uncertainty estimate unnecessary.

More details on the characteristics of surrogate models can be found in Volume 4 (Section 4.4).

The natural circulation case setup was used, and three input variables were chosen to be varied:

- The volumetric heat source in the fuel, varying uniformly from 1.8 to 3 MW/m³.
- The inlet velocity to the coolant, varying uniformly from 0.03 to 0.065 m/s.
- The coolant inlet temperature, varying uniformly from 830 to 860 K.

The definition of these inputs was provided to Dakota, and 64 evaluations requested, chosen by Latin Hypercube Sampling (LHS). It was necessary to write some interface functionality to supply these variables to placeholder locations in a Fluent input journal, launch the Fluent run, and then extract and return the quantity of interest (fuel salt maximum temperature) to Dakota. Python scripting was used for the interface.

The subchannel model is well suited to this kind of automation because, based on the solution approach developed for it in Section 3.5, it already had a suitably parametrised run journal, and solved robustly, stopping with a self identified definition of convergence. This is beneficial because it allows solutions that need more iterations to converge to have them, but also to not waste computing

³ i.e. the prediction remains valid as long as none of the inputs change too rapidly, although the definition of what 'too rapidly' is, is not addressed here.

Application of the Results

time for those that do not. In addition, it means that all cases are converged to an equivalent state of accuracy. The model is also fast to run, taking 20 to 30 minutes to converge on 16 cores. This level of automation could equally well be applied to the porous model, but could be prohibitively expensive to apply to the resolved assembly case.

The presence or absence of known ‘sentinels’ (specific phrases in run log file) was used to detect run failure and non convergence. Solution robustness is necessary for the combinations of inputs provided to Fluent by Dakota, but is hard to ensure *a priori* across the whole parameter space. Therefore, if a solution fails to pass its automated converge checks, it can be manually inspected to assess if the solution is still acceptable. In some instances, a CFD solution may diverge or hang, and will only run successfully with some changes to the solution strategy, or computer failures may interrupt the solution. In these cases, Dakota can be restarted part-way through an interrupted assessment, without needing to repeat the already completed CFD evaluations.

The form of the polynomial model returned by Dakota is a sum of the coefficients, C , each multiplying the product of the input variables, x , raised to various powers, p

$$f(x) = \sum_{k=1}^{N_k} C_k \prod_{i=1}^{N_i} x_i^{p_{k,i}}$$

In the surrogate model derived for this example $N_k = 10$, for the various combinations of powers of the $N_i = 3$ input variables. The coefficients and corresponding powers are shown in Table 4.3. Dakota also returns the quality metrics for the surrogate model:

$$\begin{aligned} \text{RMS} &= 0.3388 \text{ K} \\ \text{mean absolute error} &= 0.2723 \text{ K} \\ \text{max absolute error} &= 1.151 \text{ K} \\ R^2 &= 0.9995 \end{aligned}$$

This indicates that the fit quality was very good, and that the surrogate can be used, within the range of the input variables, to provide a rapid evaluation of what the single coolant subchannel CFD model would have predicted, with an error typically less than ± 1 K.

C_k	$p_{k,1}$	$p_{k,2}$	$p_{k,3}$
2.925205×10^2	0	0	0
9.152148×10^{-5}	1	0	0
$-1.299638 \times 10^{-12}$	2	0	0
-3.188494×10^{-4}	1	1	0
-5.110499×10^{-8}	1	0	1
-2.110934×10^3	0	1	0
1.829253×10^4	0	2	0
3.868871×10^{-1}	0	1	1
3.630437×10^{-1}	0	0	1
3.995353×10^{-4}	0	0	2

Table 4.3: Coefficients, C , and p (powers of x) for the example maximum fuel temperature surrogate model.

The order of the input variables corresponds to the order that they were supplied to Dakota, so as

Application of the Results

an example using the 250 s (baseline) conditions, for an input vector of

$$\begin{aligned}
 x_1 &= 2.40812 \times 10^6 \text{ W/m}^3 \\
 x_2 &= 0.040058 \text{ m/s} \\
 x_3 &= 842.22 \text{ K} \\
 f(x) &= 917.98 \text{ K}
 \end{aligned}$$

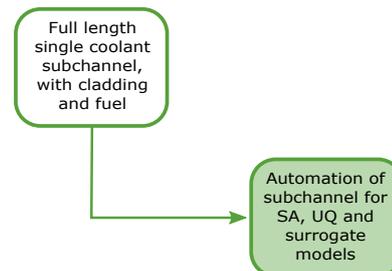
The surrogate prediction compares favourably with the CFD result of 917.53 K.

This process produces the example maximum fuel temperature surrogate model, but it also generated 64 CFD solutions, sampling a wide range of the relevant parameter space. Many quantities of interest can be extracted and surrogate models fit for them simultaneously, maximising the benefit derived from the CFD runs performed. The data generated could also be used for a range of other purposes, either by Dakota, or by further manual processing. For example, it could be used to revisit and improve the range of applicability of the heat transfer and friction factor correlations.

4.3 Uncertainty Quantification of Single Coolant Channel

The process of arriving at the comparison between the porous and resolved models identified potential significant sources of uncertainty:

- If simple correlations from the literature had been used to create the porous model inputs, then significant differences to the CFD prediction would have arisen.
- However, there are uncertainties in the CFD results, particularly related to the ability of RANS turbulence models to represent the fuel and coolant flow.
- The porous implementation of the same conditions as the resolved case does not produce perfect agreement because of the limitations and simplifications inherent in the porous approximation.



The importance of these is hard to judge until they can be compared to the dependence of the results of interest to safety or operational limits, and to the differences caused by other sources of uncertainty. There is potentially substantial uncertainty present in the material properties. The final application of the models derived in this case study is to use the Dakota toolkit and the solution automation mechanisms described above to evaluate the effect of variability in salt properties.

The type of UQ that Dakota performs is ‘forward propagation of parametric uncertainties’, where individual high-fidelity model evaluations are based on sampled distributions of input values, and the outputs are assembled using statical methods to provide meaningful insight. The Dakota training materials⁴ describe *A Practical Process for UQ* comprising seven points:

⁴ dakota.sandia.gov/training/dakota-training-materials

Application of the Results

1. Determine your UQ analysis goal:
 - What are the key model responses (quantities of interest)?
 - What kinds of statistics or metrics do you want on them?
2. Identify potentially influential uncertain input parameters, which includes parameters that influence the trend in a response, as well as those that influence variability in a response.
3. Characterise input uncertainties and map them into Dakota variable specifications.
4. What are the model characteristics/behaviours? Simulation cost, model robustness, input/output properties such as kinks, discontinuities, multi-modal, noise, disparate regimes.
5. Select a method appropriate to the variables, goal, and problem.
6. Set up Dakota input file and interface to simulation.
7. Run study and interpret the results.

The goals for the MSR fuel assembly UQ (Point 1) are to determine the influence of uncertainty and variability in salt material properties on three quantities of interest:

Max clad T : the maximum temperature in the fuel pin cladding solid⁵.

Max fuel T : the maximum temperature in the fuel salt.

Max U_z : the maximum axial velocity in the circulating fuel salt.

The probability distribution of the range of predicted values that are possible for these quantities is sought, as well as the contribution to variability in each from the chosen uncertain properties. This will provide the contribution of these variables to the expected uncertainty band that should be applied to predictions of these quantities of interest. It will also allow a judgement to be made about which properties should be subject to further modelling or measurement investigation. These three quantities were already used as solution convergence monitors, so were readily available from each CFD solution, and their expected behaviour was already well understood.

The modelling performed in the proceeding sections addresses Point 4 – the cost and robustness of the model is well understood, and kinks and discontinuities are not expected, except at laminar-turbulent transition. For this reason, the UQ study is restricted to turbulent flow conditions, and the forced circulation case setup for the single coolant subchannel model with baseline (CH2) flow conditions was used.

Points 2 and 3 are discussed in Section 4.3.1 below, Points 5 and 6 are described in Section 4.3.2, and Point 7 in Section 4.3.3.

4.3.1 Selection of Uncertain Input Parameters

The properties of the intended fuel and coolant salts contain significant uncertainty, and are assumed to represent a substantial contribution to the overall uncertainty in any assessment. This assumption is tested here by performing what would be best considered to be a screening study, where relatively wide uncertainties are assumed and both SA and UQ are performed at the same time on the same set of CFD simulations.

⁵ Abbreviating 'cladding' to 'clad' for compactness in figures and tables.

Application of the Results

Not all material properties have been chosen to be varied, and there is no simultaneous varying of flow conditions and material properties. Studying a wide range of variables at once is possible using the methods described here, but some caution and thought is warranted, rather than attempting a brute force approach and varying everything. This is because the complexity of the responses is unknown at this initial stage, and the number of runs necessary, as well as the task of assessing the adequacy of the outcome and methods, can easily grow to be become difficult to manage and interpret.

The chosen properties, the form of their Probability Density Function (PDF) (normally or uniformly distributed) and the assumed uncertainty in them is specified in Table 4.4. The uncertainty in the viscosity and thermal conductivity was chosen as the largest values for these properties from any of the salts described by Romatoski and Hu (2017). That reference assumes that these uncertainties are 2σ values (± 2 standard deviations, 95 % coverage), but acknowledges that there is ambiguity about whether the source data is a 1σ or a 2σ value. Because the salts used in this case study are less well known, with some of their properties estimated, then the stated uncertainties represent $\pm 1\sigma$, representing one standard deviation (68 % coverage) as a ‘standard uncertainty’ (JCGM, 2008). This means that, for example, if 15 % represents 1σ , the uncertainty in the value is twice as much as if 15 % was assumed to represent 2σ , i.e. that the PDF is twice the width.

Normal	standard uncertainty	
fuel $\partial\rho/\partial T$	$\pm 15\%$	
fuel k	$\pm 15\%$	
fuel μ	$\pm 20\%$	
coolant k	$\pm 15\%$	
coolant μ	$\pm 20\%$	
Uniform	min	max
fuel κ (m^{-1})	100	1000
coolant κ (m^{-1})	100	1000

Table 4.4: Chosen uncertain variables and their estimated uncertainty.

When varying a property it was not considered adequate to vary it as a uniform value applied everywhere in the fluid in question. The local variation in temperature near to surfaces, and the local variation in μ and k that this produces is expected to lead to differences in heat transfer. This is a reason that these properties were selected as a subset of properties, in preference to specific heat capacity, because there is no guidance available on the variation of c_p with temperature. A given perturbation to the μ and k properties was implemented by multiplying the entire polynomial by a factor and specifying these new polynomial coefficients in Fluent. This altered the way that the interface to Dakota was implemented. Dakota specifies a property to be varied as a single number, and therefore, for each normally distributed variable, Dakota was requested to select from a standard normal distribution (mean = 0, variance = 1), and the corresponding uncertainty value and polynomial manipulation was implemented in the python interface, which was developed as a ‘driver’ for Fluent.

The error associated with fitting the fifth order polynomial to the exponential viscosity function is not accounted for. For the nominal (zero uncertainty) functions the error in this polynomial fit is 0.43 % for the coolant and 0.89 % for the fuel ($< 1\%$ error being the criterion used to select the

Application of the Results

necessary polynomial order) and while this is small, it is an example of a range of small sources of uncertainties that can be introduced as part of the practical implementation of the modelling process that are often overlooked, or not well captured and dealt with rigorously. Completeness and comprehensiveness in the treatment of all sources of uncertainty in a complex model is likely to be either impossible, or disproportionate when considered as part of a graded approach (Volume 1, Section 2.2.5).

In addition to μ and k in both coolant and fuel, uncertainty in the temperature coefficient of the density of the fuel salt was included. Contradictory evidence exists from some historical references (such as that reported in Janz *et al.*, 1975), and so the value of $\partial\rho/\partial T$ ($-1.015\text{ kg/m}^3\text{ K}$ as the baseline value in Section 2.2.3.3) was varied. This is a different type of polynomial manipulation to that applied to μ and k because instead of shifting a whole curve to be greater or smaller value, this changed the slope only. The value of fuel density at 1200 K was used as the reference point, with the gradient ‘pivoting’ around this.

Dakota allows a mixture of forms of PDF for uncertain variables to be assessed, and the normal distributions for thermophysical properties were able to be combined with uniform distributions for absorption coefficients. The range of absorption coefficients was chosen to span an optically thick to an optically thin situation (Table 2.3), and the independent effects of the fuel and coolant salt absorption can be established.

The potential variation in material properties is assumed to have no correlation between them, so they can be treated as independent uncertain variables.

4.3.2 Sensitivity Analysis and Uncertainty Quantification Method

The chosen approach for SA and UQ of the uncertainty in material properties is to apply a Polynomial Chaos Expansion (PCE) surrogate model to a sampled set of CFD solutions and to evaluate the desired statistical quantities from it. Once the PCE surrogate model has been created, its properties allow the statistical moments of the output distribution (mean, variance, skew and kurtosis), the linear sensitivity coefficients and Sobol indices to be evaluated analytically.

Two different methods for creating the PCE were evaluated, each being a cross-check of the other.

LHS: 128 samples were drawn from the PDFs of the 7 uncertain input variables and a CFD model evaluation used for each. The baseline single subchannel converged solution was used as a starting point for each evaluation to reduce run-time and reduce the risk of the solution not converging in the early phases of establishing the flow. The LHS method is the most common sampling method and typically used as the starting point for an SA or UQ study. The samples and their solved evaluations can be used for a number of purposes after they have been generated. More samples can be added to an existing set if more data is found to be beneficial, but to retain the ‘latin’ property of covering the parameter space with one sample per ‘bin’, the number of samples must double each time.

The results from the LHS samples themselves can be used directly without applying the PCE, as will be shown in Section 4.3.3, and this kind of sampling combined with a straightforward evaluation of the results would be the only option if some of the uncertain variables were not smooth.

Application of the Results

For example, an LHS sampling approach could vary the RANS turbulence models as one of its parameters, but PCE could not, because its inputs must be smooth.

When building the PCE from the LHS samples, Dakota can perform its own assessment to find the best polynomial order for each output variable polynomial, and third to fifth order responses were generally selected. The advantage of building the PCE is mainly the additional quantities that can be derived from it; to obtain the Sobol indices by LHS sampling alone with 7 input variables would require $7 + 2 = 9$ times as many CFD evaluations.

Sparse Grid: The PCE model applied to the LHS samples is a regression fit to that data. An alternative method directly uses the orthogonality in relationships between the inputs by creating a grid of samples and building polynomials along the various directions. To use the full grid of points would need N^p samples for N uncertain variables with a polynomial order of p . This quickly results in a large number of evaluations being necessary.

An alternative is to apply a 'sparse grid' approach, where only a subset of the grid is evaluated. Dakota selects the number of evaluations needed based on the 'level' of occupancy of the grid. For the 7 uncertain variables used here, 137 CFD evaluations were created for sparse grid Level 2. Level 3 would need 889 evaluations. Similarly to the LHS PCE, the sparse grid output contained up to 5th order polynomial responses.

Both types of PCE were sampled 10^5 times after being built to provide smooth Cumulative Distribution Functions (CDFs) with well resolved 'tails', in addition to their analytically derived quantities. Once built the PCEs could also be potentially re-used as fast running surrogate models for the quantities of interest in their own right.

4.3.3 Uncertainty Quantification Results

The LHS results were assessed first by using a scatter plot of each of the three quantities of interest against the seven uncertain inputs, and by drawing histograms of the outputs – these are shown together in Figure 4.6.

It is good practice to always visualise the results like this to look for features and anomalies in the results, rather than relying purely on statistical measures. In particular, the presence of fine structure in the response of a variable is significant because it would make the application of a method like PCE less likely to be successful.

For each scatter plot, the partial correlation coefficients and partial rank correlation coefficients are shown for the quantity of interest vs. input variable (the rank coefficient is the number shown underneath). The normally distributed variables are plotted against the value of σ chosen by Dakota (noting that these are then scaled by the uncertainties in Table 4.4), showing that a span of approximately $\pm 2\sigma$ was sampled for each.

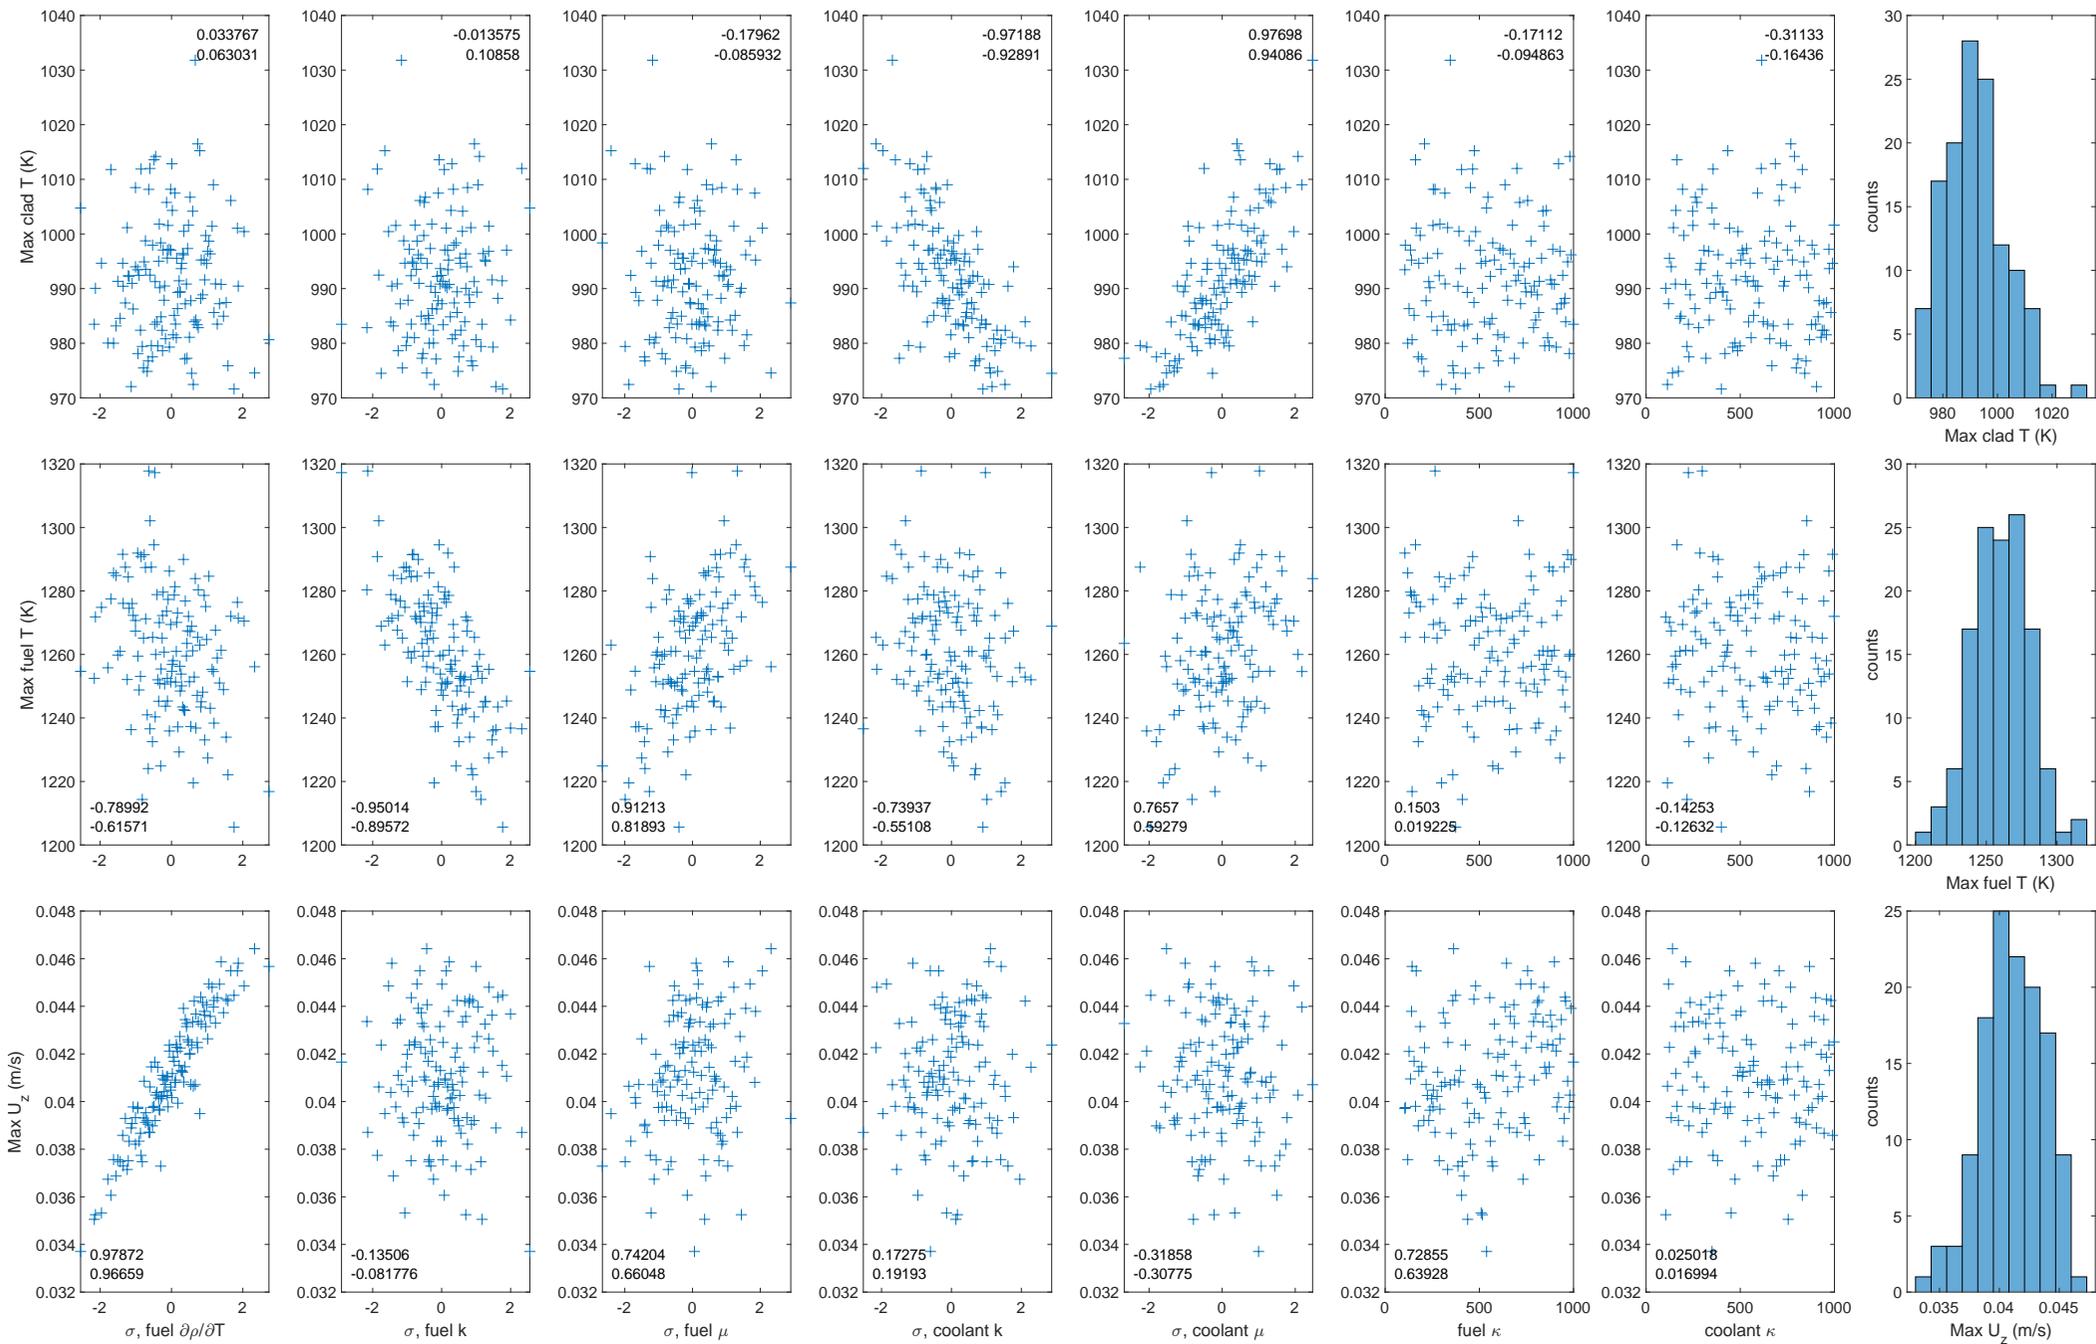


Figure 4.6: Scatter plots and histograms for all three response variables against the seven uncertain variables.

Application of the Results

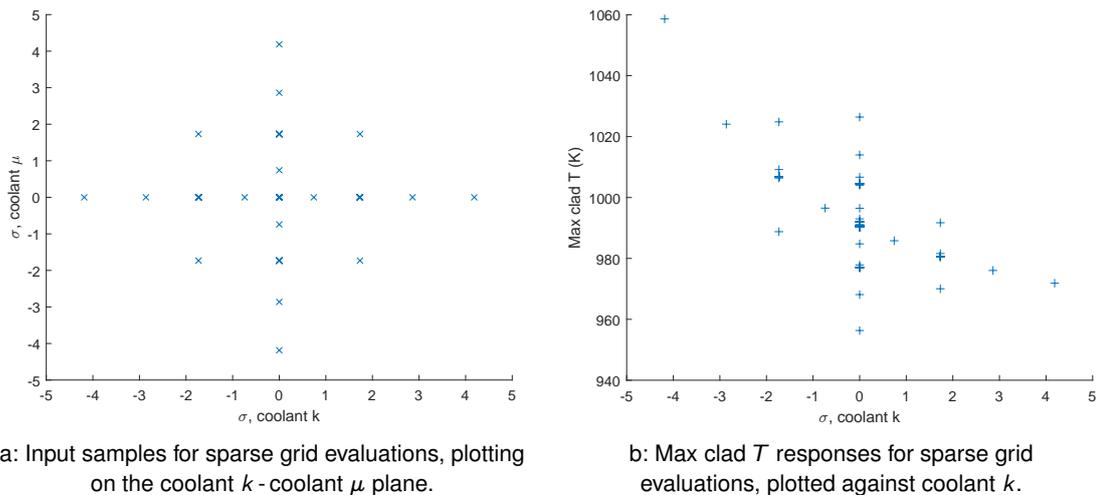


Figure 4.7: Example of sparse grid sampling.

The difference between the random layout of the LHS samples and the sparse grid sampling can be seen in Figure 4.7. All 137 samples are plotted on Figure 4.7a (this is a projection onto 2 dimensions of the 7 dimensional sample space) so the variation of the other parameters occurs along lines orthogonal to this plane.

The standard normal distribution values chosen by Dakota for these variables mainly lie along the 0σ axis for the other variables (a higher sparse grid level than 2 would place more samples towards the corners, to capture more coupled effects, but would require more evaluations), and the maximum range is $\pm 4.2\sigma$.

This is a wider range of variation than the LHS samples, so it was necessary to check that these greater variations in properties still produced valid CFD solutions. An example of the response of the CFD evaluations on the sparse grid parameter space is shown in Figure 4.7b, where the maximum cladding temperature variation is shown against the variation in coolant thermal conductivity. The clear trend of reducing cladding temperature with increasing conductivity can be seen, which agrees with the correlation coefficients with values that are nearly -1 in the corresponding scatter plot.

The CDFs of the LHS samples are plotted for the resulting variation in the three quantities of interest as a result of the variation in all inputs in Figure 4.8. The CDFs from both methods of PCE are also plotted, and show effectively identical results as each other, and as the LHS results. The advantage of the PCE curves is that they do not contain the same ‘noise’ (variation in gradient from point to point) as the LHS, and they can also extend into the low probability high and low response ‘tails’ of the distribution.

The mean and standard distribution of the three output responses (which are not normally distributed, and have no requirement to be to apply these measures) from the LHS samples corresponding to the CDF figures are shown in Table 4.5, including the upper and lower 95% Confidence Interval (CI) estimated by Dakota for each. The mean and σ from the sparse grid PCE implementation is also shown in this table – the nominal values agree well with the LHS values and lie comfortably within the CIs.

Application of the Results



	Max clad T (K)	Max fuel T (K)	Max U_z (m/s)
LHS samples			
lower 95% CI of mean	990.01	1257.05	0.04060
mean	991.95	1260.61	0.04106
upper 95% CI of mean	993.90	1264.18	0.04151
lower 95% CI of σ	9.91	18.14	0.00230
σ	11.13	20.37	0.00259
upper 95% CI of σ	12.67	23.22	0.00295
Sparse grid PCE			
mean	992.06	1260.61	0.04103
σ	11.39	20.62	0.00257

Table 4.5: Statistics for UQ for forced circulation conditions.

	Max clad T	Max fuel T	Max U_z
fuel $\partial\rho/\partial T$		-5.81	2.375×10^{-3}
fuel k		-13.8	
fuel μ		9.45	4.061×10^{-4}
coolant k	-7.16	-4.95	9.800×10^{-5}
coolant μ	7.89	5.56	-1.122×10^{-4}
fuel κ		0.0143	3.170×10^{-6}
coolant κ	-0.00173		

Table 4.6: Linear sensitivity coefficients for SA for forced circulation conditions (omitting small values).

	Max clad T	Max fuel T	Max U_z
fuel $\partial\rho/\partial T$		0.086	0.878
fuel k		0.492	
fuel μ		0.234	0.057
coolant k	0.471	0.067	0.002
coolant μ	0.512	0.078	0.002
fuel κ		0.040	0.056
coolant κ	0.005		

Table 4.7: Sobol indices for SA for forced circulation conditions, corresponding to Figure 4.9 (omitting $< 10^{-3}$).

Application of the Results

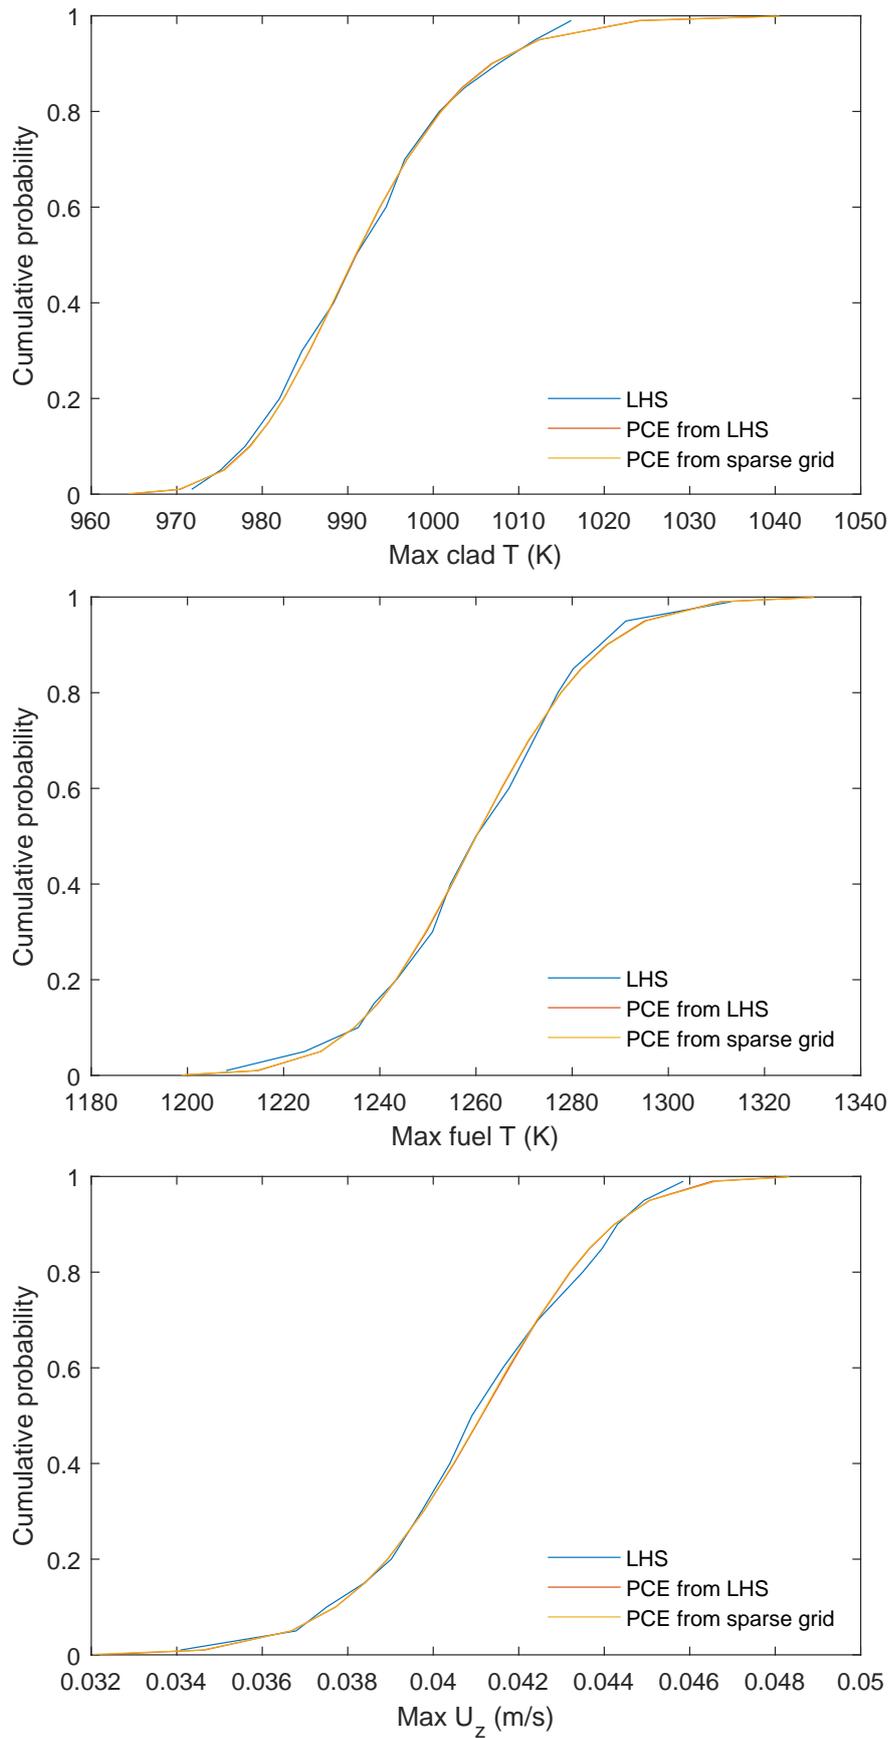


Figure 4.8: CDFs for the quantities of interest.

Application of the Results

Comparing several methods of calculating the same quantities is recommended, and similar to the process of building confidence in a CFD model. Applying SA and UQ successfully also involves a gradual accumulation of evidence, insight and decisions, adding understanding of the methods and system incrementally – at the early stages of an assessment there is rarely one complete method or setup that can be applied without checking its performance and questioning its appropriateness.

This means that SA and UQ do not ‘come for free’ (even with convenient tools like Dakota). They add an additional burden to asking ‘*is my physical simulation adequate?*’ of also asking:

- *Is my UQ adequate?*
- *Have I chosen suitable uncertain variables and distributions?*
- *Are the surrogate model polynomials high enough order?*
- *Have enough samples been taken?*
- *Did the best fits converge well enough? etc.*

If SA and UQ are to be used and relied on to make decisions or justifications, then they need to be of high enough quality and sufficiently trustworthy to justify that reliance.

4.3.4 Sensitivity Analysis Results

The CDF results discussed above focus on the distribution of the output quantities of interest. The same PCE can be used to look at the significance of the uncertainty in each input (distinguishing SA from UQ, as discussed in Volume 4). Two results are shown here that allow this:

Linear sensitivities: The gradient of the response of each quantity of interest to changes in each input at the centre of its uncertainty distribution is shown in Table 4.6. Small values have been omitted for clarity. These values are from the sparse grid PCE, although the values determined for the PCE fit to the LHS samples were very similar. For the normally distributed variables, these values are sensitivity per σ , for the uniformly distributed radiation absorption coefficients, they are per unit κ .

As an example of their application, Figure 4.7b shows that for an increase in coolant k from 0σ to $+1.73\sigma \equiv 26\%$ there is a 10.25 K reduction in cladding temperature in the CFD model result. Using the sensitivity coefficient predicts $-7.16 \text{ K}/\sigma \times 1.73 \sigma = -12.4 \text{ K}$. The agreement is reasonable, with the difference arising from the obvious curvature (non-linearity) in the actual response. This shows that in this case, projecting the first-order gradient is only accurate for small variations.

Sobol indices: While the sensitivity coefficients can quantify the expected change in an output with an input, they are less intuitive to interpret and compare how important the uncertainty in each inputs is. The Sobol indices for each input sum to 1, and show the relative contribution to variation in the output of each input (although they do not give the sign information that sensitivity coefficients do). The Sobol indices⁶ are shown in Table 4.7 (again omitting small values) and are plotted in Figure 4.9.

The Sobol indices give a clear result for each quantity of interest.

⁶ These are the ‘main effect’ indices (see Volume 4, Section 4.5); Dakota also produces the ‘total effect’ indices, which have similar values in this case.

Application of the Results

Max clad T : Only the coolant thermophysical properties have any significant influence. Variation in the fuel properties have no effect, and the coolant absorption coefficient had almost no influence. The dependence of the cladding temperature on coolant k was expected, but the equal significance of coolant μ was not; this is interpreted as supporting the point discussed in Section 3.5.2, where, for this low Re turbulent flow, the effect of the variation of viscosity in the near-wall momentum boundary layer is important for heat transfer.

Max fuel T : The maximum fuel temperature is affected mainly by the fuel k , and to a lesser extent by its viscosity. The influence on the cladding maximum temperature of the coolant μ and k is also seen, because this is passed on as an offset to the fuel temperature. The fuel radiation absorption coefficient is seen to have only a small effect.

Max U_z : The maximum fuel circulating velocity is sensitive to $\partial\rho/\partial T$; this is expected because this is effectively the same as β that appears in S (Section 2.2.4). Of the other variables, only a small sensitivity to fuel μ and absorption coefficient was observed.

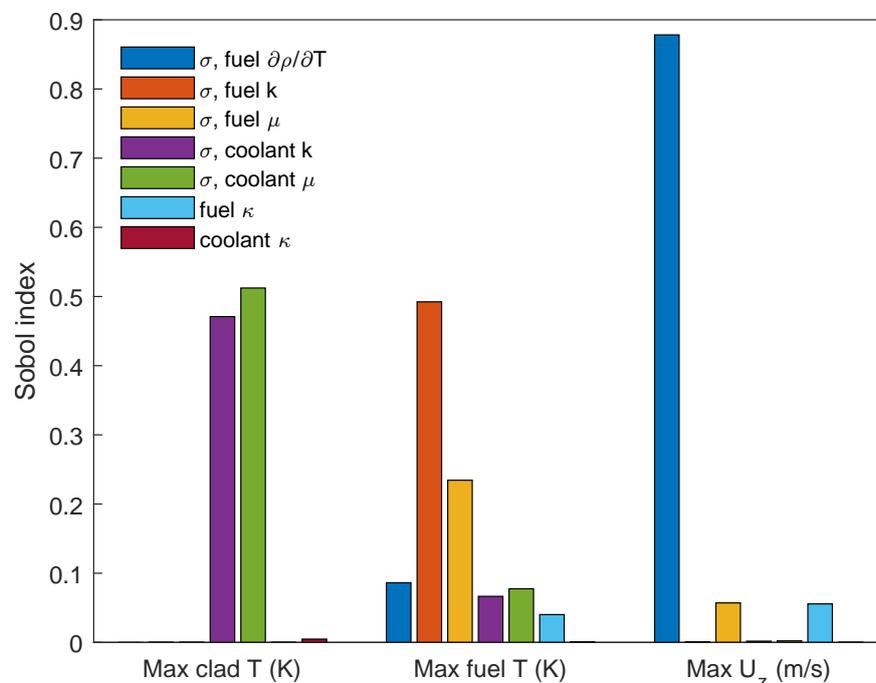


Figure 4.9: Sobol indices for SA for forced circulation conditions, corresponding to Table 4.7.

The conclusions that can be drawn from the Sobol indices are supported by comparing the relative magnitude of the sensitivity coefficients, and also by the correlation coefficients shown on Figure 4.6. Both types of correlation coefficient generally agree with each other, and where there is a large sensitivity contribution, there is a correlation that is close to ± 1 . Where there is less sensitivity, there is less correlation.

The lack of sensitivity to absorption coefficient is a significant result, because knowing the value of this property through all stages of reactor life would be challenging, and modelling its effect in detail can be computationally costly and complex.

The comparative lack of sensitivity of the fuel maximum temperature to $\partial\rho/\partial T$ is notable, given its influence on the maximum fuel velocity. This is consistent with the conclusion reached earlier that,

Application of the Results

at the mid-length of the pin (where the maximum temperature occurs in the forced circulation case, because of the power profile), the heat transfer is mainly in the radial direction in the pin, and so the fuel conductivity is the dominant effect, as shown. It is expected that the profile of fuel temperature at either end of the pin would be more sensitive to $\partial\rho/\partial T$. The interaction of this parameter with the turbulence modelling in the fuel salt may also be an area worth investigating further, if desirable.

This is an initial screening assessment, and it should be noted that although an uncertain input may not exhibit sensitivity here, it cannot be excluded from all future consideration based on this assessment alone. A different flow condition (such as natural circulation) or a different quantity of interest (such as duct temperature or coolant pressure drop) may be significantly more sensitive to an input than found here. To be confident in a conclusion, the process of assessing the importance of uncertain variables, and the conditions that they are applied in, should be sufficiently extensive, and will likely need to be iterative.

5 References

- Al-Aabidy Q, Craft T J, Iacovides H** (2020) Improved Eddy-Viscosity Modelling of Turbulent Flow around Porous–Fluid Interface Regions. *Transport in Porous Media*, 131(2), 569–594, dx.doi.org/10.1007/s11242-019-01357-0.
- ANSYS** (2020a) Fluent Theory Guide, Release 2020 R1.
- ANSYS** (2020b) Fluent Users’s Guide, Release 2020 R1.
- Bergholz R F** (1980) Natural Convection of a Heat Generating Fluid in a Closed Cavity. *Journal of Heat Transfer*, 102(2), 242–247, dx.doi.org/10.1115/1.3244267.
- CSNI** (2013) Report of the OECD/NEA-KAERI Rod Bundle CFD Benchmark Exercise. NEA/CSNI/R(2013)5, OECD NEA Committee on the Safety of Nuclear Installations.
- de Lemos M J S** (2012) Turbulence in Porous Media: Modeling and Applications. Second edition, Elsevier.
- Desyatnik V N, Katyshev S F, Raspopin S P, Chervinskii Y F** (1975) Density, Surface Tension, and Viscosity of Uranium Trichloride-Sodium Chloride Melts. *Soviet Atomic Energy*, 39(1), 649–651, dx.doi.org/10.1007/BF01121527.
- Eckhardt B, Schneider T M, Hof B, Westerweel J** (2007) Turbulence Transition in Pipe Flow. *Annual Review of Fluid Mechanics*, 39(1), 447–468, dx.doi.org/10.1146/annurev.fluid.39.050905.110308.
- Gelles D** (1996) Microstructural Examination of Commercial Ferritic Alloys at 200 Dpa. *Journal of Nuclear Materials*, 233–237, 293–298, dx.doi.org/10.1016/S0022-3115(96)00222-X.
- Hales J D, Williamson R L, Novascone S R, Pastore G, Spencer B W, Stafford D S, Gamble K A, Perez D M, et al.** (2016) BISON Theory Manual The Equations behind Nuclear Fuel Analysis. BISON Release 1.3. INL/EXT-13-29930 Rev.3, Idaho National Laboratory, dx.doi.org/10.2172/1374503.
- Hofman G L, Billone M C, Koenig J F, Kramer J M, Lambert J D B, Leibowitz L, Orechwa Y, Pedersen D R, et al.** (2019) Metallic Fuels Handbook. ANL-NSE-3, Argonne National Laboratory, dx.doi.org/10.2172/1506477.
- Incropera F P, DeWitt D P, Bergman T L, Lavine A S** (2011) Fundamentals of Heat and Mass Transfer. Seventh edition, John Wiley & Sons.
- Jackson J D, Cotton M A, Axcell B P** (1989) Studies of Mixed Convection in Vertical Tubes. *International Journal of Heat and Fluid Flow*, 10(1), 2–15, dx.doi.org/10.1016/0142-727X(89)90049-0.
- Janz G J, Tomkins R P T, Allen C B, Downey J R, Garner G L, Krebs U, Singer S K** (1975) Molten Salts: Volume 4, Part 2, Chlorides and Mixtures—Electrical Conductance, Density, Viscosity, and Surface Tension Data. *Journal of Physical and Chemical Reference Data*, 4(4), 871–1178, dx.doi.org/10.1063/1.555527.

References

- JCGM** (2008) Evaluation of Measurement Data - Guide to the Expression of Uncertainty in Measurement. JCGM 100:2008 (GUM 1995), Joint Committee for Guides in Metrology.
- Jin Y, Chen C** (1996) Natural Convection of High Prandtl Number Fluids with Variable Viscosity in a Vertical Slot. *International Journal of Heat and Mass Transfer*, 39(13), 2663–2670, dx.doi.org/10.1016/0017-9310(95)00366-5.
- Kakaç S, Shah R K, Aung W** (1987) Handbook of Single-Phase Convective Heat Transfer. Wiley.
- Kamide H, Hayashi K, Isozaki T, Nishimura M** (2001) Investigation of Core Thermohydraulics in Fast Reactors—Interwrapper Flow during Natural Circulation. *Nuclear Technology*, 133(1), 77–91, dx.doi.org/10.13182/NT01-A3160.
- Keshmiri A** (2010) Thermal-Hydraulic Analysis of Gas- Cooled Reactor Core Flows. Ph.D. thesis, University of Manchester.
- Keshmiri A, Cotton M A, Addad Y, Laurence D** (2012) Turbulence Models and Large Eddy Simulations Applied to Ascending Mixed Convection Flows. *Flow, Turbulence and Combustion*, 89(3), 407–434, dx.doi.org/10.1007/s10494-012-9401-4.
- Leibowitz L, Blomquist R A** (1988) Thermal Conductivity and Thermal Expansion of Stainless Steels D9 and HT9. *International Journal of Thermophysics*, 9(5), 873–883, dx.doi.org/10.1007/BF00503252.
- Los Alamos** (1962) LAMPRE I Final Design Status Report. LA-2833, Los Alamos Scientific Laboratory, dx.doi.org/10.2172/4748402.
- Menter F R, Smirnov P E, Liu T, Avancha R** (2015) A One-Equation Local Correlation-Based Transition Model. *Flow, Turbulence and Combustion*, 95(4), 583–619, dx.doi.org/10.1007/s10494-015-9622-4.
- Mullin T, Peixinho J** (2006) Transition to Turbulence in Pipe Flow. *Journal of Low Temperature Physics*, 145(1-4), 75–88, dx.doi.org/10.1007/s10909-006-9242-4.
- Murgatroyd W, Watson A** (1970) An Experimental Investigation of the Natural Convection of a Heat Generating Fluid within a Closed Vertical Cylinder. *Journal of Mechanical Engineering Science*, 12(5), 354–363, dx.doi.org/10.1243/JMES_JOUR_1970_012_059_02.
- Murthy J Y, Mathur S R** (1998) Finite Volume Method for Radiative Heat Transfer Using Unstructured Meshes. *Journal of Thermophysics and Heat Transfer*, 12(3), 313–321, dx.doi.org/10.2514/2.6363.
- Ochoa-Tapia J, Whitaker S** (1995) Momentum Transfer at the Boundary between a Porous Medium and a Homogeneous Fluid—II. Comparison with Experiment. *International Journal of Heat and Mass Transfer*, 38(14), 2647–2655, dx.doi.org/10.1016/0017-9310(94)00347-X.
- Richard J, Wang D, Yoder G, Carbajo J, Williams D, Forget B, Forsberg C** (2014) Implementation of Liquid Salt Working Fluids into TRACE. In *Proceedings of ICAPP 2014*.
- Rohsenow W M, Hartnett J P, Cho Y I** (1998) Handbook of Heat Transfer. Third edition, McGraw-Hill.

References

- Romatoski R, Hu L** (2017) Fluoride Salt Coolant Properties for Nuclear Reactor Applications: A Review. *Annals of Nuclear Energy*, 109, 635–647, dx.doi.org/10.1016/j.anucene.2017.05.036.
- Uitslag-Doolaard H, Roelofs F, Pacio J, Batta A** (2019) Experiment Design to Assess the Inter-Wrapper Heat Transfer in LMFR. *Nuclear Engineering and Design*, 341, 297–305, dx.doi.org/10.1016/j.nucengdes.2018.11.019.
- US DOE** (1996) Spent Nuclear Fuel Effective Thermal Conductivity Report. MOL.19961202.0030, U.S. Department of Energy, dx.doi.org/10.2172/778872.
- Wagner W, Cooper J R, Dittmann A, Kijima J, Kretzschmar H J, Kruse A, Mareš R, Oguchi K, et al.** (2000) The IAPWS Industrial Formulation 1997 for the Thermodynamic Properties of Water and Steam. *Journal of Engineering for Gas Turbines and Power*, 122(1), 150–184, dx.doi.org/10.1115/1.483186.
- Wang X, Zhang D, Wang M, Song P, Wang S, Liang Y, Zhang Y, Tian W, et al.** (2020) Hybrid Medium Model for Conjugate Heat Transfer Modeling in the Core of Sodium-Cooled Fast Reactor. *Nuclear Engineering and Technology*, 52(4), 708–720, dx.doi.org/10.1016/j.net.2019.09.009.
- Watson A** (1971) Natural Convection of a Heat Generating Fluid in a Closed Vertical Cylinder: An Examination of Theoretical Predictions. *Journal of Mechanical Engineering Science*, 13(3), 151–156, dx.doi.org/10.1243/JMES_JOUR_1971_013_027_02.
- Williams D** (2006) Assessment of Candidate Molten Salt Coolants for the Advanced High Temperature Reactor (AHTR). ORNL/TM-2006/12, Oak Ridge National Laboratory, dx.doi.org/10.2172/885975.
- Wu X, Moin P, Adrian R J, Baltzer J R** (2015) Osborne Reynolds Pipe Flow: Direct Simulation from Laminar through Gradual Transition to Fully Developed Turbulence. *Proceedings of the National Academy of Sciences*, 112(26), 7920–7924, dx.doi.org/10.1073/pnas.1509451112.
- Yamanouchi N, Tamura M, Hayakawa H, Hishinuma A, Kondo T** (1992) Accumulation of Engineering Data for Practical Use of Reduced Activation Ferritic Steel: 8%Cr-2%W-0.2%V-0.04%Ta-Fe. *Journal of Nuclear Materials*, 191–194, 822–826, dx.doi.org/10.1016/0022-3115(92)90587-B.

6 Nomenclature

Latin Symbols

A Area, m^2

At Atwood number ($At = (\rho_1 - \rho_2)/(\rho_1 + \rho_2)$)

Bi Biot number ($Bi = hL/k_s$)

c_p, c_v Specific heat at constant pressure or volume, $\text{J kg}^{-1} \text{K}^{-1}$

d or D Diameter ($D_h = 4A_{cs}/p_{cs}$ for hydraulic diameter), m

f Darcy-Weisbach friction factor

Fo Fourier number ($Fo = \alpha_s t/L^2$)

Gr Grashof number ($Gr = gL^3 \Delta\rho/\nu^2 \rho = gL^3 \beta \Delta T/\nu^2$, using the Boussinesq approximation $\Delta\rho/\rho \approx -\beta \Delta T$, where ΔT is often taken as $T_w - T_{s,\infty}$)

g Acceleration due to gravity, m s^{-2}

h Specific enthalpy, J kg^{-1} , Heat Transfer Coefficient (HTC), $\text{W m}^{-2} \text{K}^{-1}$ or height, m

I Radiative intensity, $\text{W m}^{-2} \text{sr}^{-1}$ or $\text{W m}^{-2} \text{sr}^{-1} \mu\text{m}^{-1}$ for a spectral density, where sr (steradian) is solid angle

J Radiosity, W m^{-2}

k Thermal conductivity, $\text{W m}^{-1} \text{K}^{-1}$

L Length or wall thickness, m

M Molar mass of a species, kg kmol^{-1}

Ma Mach number ($Ma = U/a$, where a is the speed of sound)

n Refractive index

Nu Nusselt Number ($Nu = hL/k_f$)

p Perimeter, m

P Pressure ($P_s =$ static pressure, $P_T =$ total pressure), N m^{-2} or Pa

Pe Péclet number ($Pe = RePr$)

Pr Prandtl number ($Pr = c_p \mu/k_f$)

q Heat flux (rate of heat transfer per unit area, $q = Q/A$), W m^{-2}

Q Rate of heat transfer, W

r Radius, m

R Gas constant (for a particular gas, $R = \tilde{R}/M$), $\text{J kg}^{-1} \text{K}^{-1}$

\tilde{R} Universal gas constant, $8314.5 \text{ J kmol}^{-1} \text{K}^{-1}$

R_{th} Thermal resistance, K W^{-1}

Ra Rayleigh number ($Ra = GrPr$)

Re Reynolds number ($Re = \rho UL/\mu$, or for an internal flow $Re = WD_h/A_{cs}\mu$)

Ri Richardson number ($Ri = Gr/Re^2$)

Sr Strouhal number ($Sr = fL/U$, where f is frequency)

Nomenclature

St	Stanton number ($St = Nu/RePr$)
t	Time, s
T	Temperature ($T_s =$ static temperature, $T_T =$ total temperature), K
u_τ	Wall friction velocity ($u_\tau = \sqrt{\tau_w/\rho}$), m s^{-1}
U	Velocity, m s^{-1} or thermal transmittance, $\text{W m}^{-2} \text{K}^{-1}$
v	Specific volume, $\text{m}^3 \text{kg}^{-1}$
V	Volume, m^3
W	Mass flow rate, kg s^{-1}
y	Wall distance, m
y^+	Non-dimensional wall distance ($y^+ = yu_\tau/\nu$)

Greek Symbols

α	Thermal diffusivity ($\alpha = k/\rho c_p$), $\text{m}^2 \text{s}^{-1}$
β	Volumetric thermal expansion coefficient ($\beta = -(1/\rho)(\partial\rho/\partial T)$), K^{-1}
γ	Ratio of specific heats ($\gamma = c_p/c_v$)
ϵ	Emissivity or surface roughness height, m
κ	Absorption coefficient, m^{-1}
λ	Wavelength, m
μ	Viscosity, $\text{kg m}^{-1} \text{s}^{-1}$
ν	Kinematic viscosity and momentum diffusivity ($\nu = \mu/\rho$), $\text{m}^2 \text{s}^{-1}$
ρ	Density, kg m^{-3}
σ	Stefan Boltzmann constant, $5.67 \times 10^{-8} \text{W m}^{-2} \text{K}^{-4}$
τ	Shear stress, N m^{-2}
ϕ	Porosity or void fraction

Subscripts and Modifications

b	Bulk (mass-averaged) quantity
cs	Cross-sectional quantity
f	Quantity relating to a fluid
i	Quantity relating to a particular species
T	Total (stagnation) quantity
t	Turbulent quantity
s	Static quantity or quantity relating to a solid
w	Quantity relating to a wall or surface
∞	Quantity far from a wall or in free-stream
\square	Average quantity
\checkmark	Molar quantity
\square'	Varying quantity

7 Abbreviations

CAD	Computer Aided Design
CDF	Cumulative Distribution Function
CFD	Computational Fluid Dynamics
CHT	Conjugate Heat Transfer
CI	Confidence Interval
DNP	Delayed Neutron Precursor
DNS	Direct Numerical Simulation
DO	Discrete Ordinates
EHRS	Emergency Heat Removal System
FA	Fuel Assembly
HTC	Heat Transfer Coefficient
LES	Large Eddy Simulation
LHS	Latin Hypercube Sampling
LMFR	Liquid Metal-cooled Fast Reactor
MSR	Molten Salt Reactor
PCE	Polynomial Chaos Expansion
PDF	Probability Density Function
PWR	Pressurised Water Reactor
RANS	Reynolds-Averaged Navier-Stokes
RMS	Root Mean Square
SA	Sensitivity Analysis
SBO	Station Blackout
SSR-W	Stable Salt Reactor – Wasteburner
UDF	User Defined Function
UQ	Uncertainty Quantification
URF	Under-Relaxation Factor

A Single Fuel Pin Turbulence Assessment

In the 1/6th fuel assembly model, the fuel pins have high aspect ratio 20 mm long cells for most of their length. The single fuel pin mesh used the same resolution as a starting point, but these cells were suspected of being problematic, so an additional revised mesh was created, limiting the maximum axial cell length to 1 mm, so that the cells were not high aspect ratio. This mesh contained 3 million cells for the single pin instead of 375,000.

With this axially resolved mesh, using the simplified conditions described in Section 3.4, it was found that when running laminar with a steady-state solver, the maximum velocities were much reduced, but the solution residuals would still not converge.

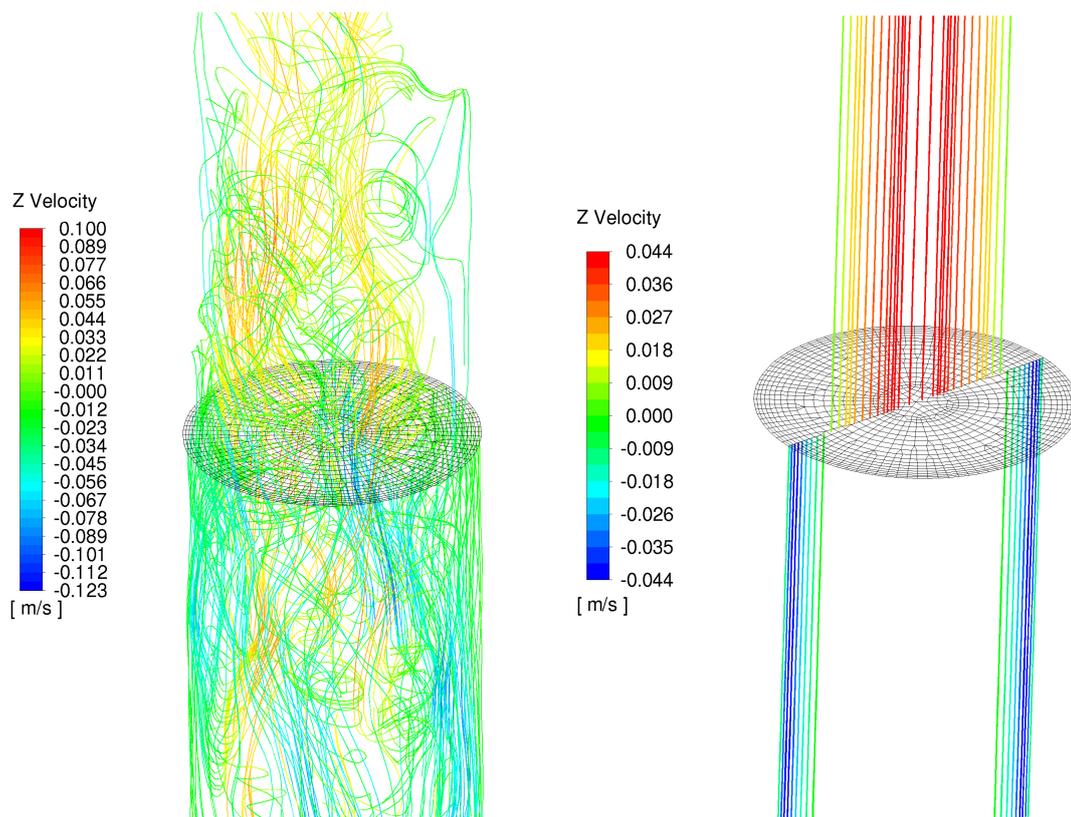


Figure A.1: Streamlines for the transient laminar axially resolved single pin case (left) where the streamlines are released from the whole cross-section at mid-height and indicate chaotic flow. For the equivalent case with a RANS model active (right), the streamlines released along a line across the profile show steady, symmetric flow upwards in the middle of the pin and downwards at the edge. There are small localised regions of significantly faster flow in the transient laminar case than in the turbulent case.

Single Fuel Pin Turbulence Assessment

Starting from a converged RANS steady-state solution, the solution was solved transiently with a 2 ms time step (keeping Courant numbers at approximately 0.5), and the RANS model was turned off. The symmetric flow pattern accelerated, keeping its shape for a period, moving several times faster than the initial condition, and then suddenly the maximum velocities dropped and became unsteady. The solution residuals were converged at every time step throughout. The interpretation is that the large buoyancy forces created the high velocities, generating an instability that eventually underwent a process that appeared like laminar-turbulent transition, resulting in chaotic flow patterns, that were well resolved by the axially refined mesh (Figure A.1).

The transient laminar case was solved for 20 s and the velocity field was averaged (also gathering the RMS values) over the last 11 s, once the flow had appeared to reach a statistically stable condition. Figures A.2 to A.5 show the centreline and cross-pin profiles of vertical velocity and temperature for the transient laminar solution¹ and several cases with RANS models:

k - ω SST: the *k* - ω SST model with low *Re* corrections, run on the coarse (20 mm) axial resolution mesh.

k - ω SST high res: the same turbulence model run on the fine (1 mm) axial resolution mesh. The results are almost indistinguishable from those for the coarse mesh. A transient version of this fine mesh with the RANS model active was solved, and no instabilities occurred – the result was identical to the steady-state RANS solutions.

k - ω SST gamma: the γ intermittency option was added and the low *Re* corrections turned off in the *k* - ω SST model. This model was started from a converged *k* - ω SST solution and used the coarse mesh

realizable *k* - ϵ : the realizable *k* - ϵ model, using the coarse mesh.

The RANS model results are in reasonable agreement with the averaged transient laminar results:

- The realizable *k* - ϵ result for centreline velocity looks convincing in Figure A.2, but Figure A.4 shows that agreement at the centre is not matched by equally good agreement over the cross-section; the shape of the velocity profile does not agree as well with the averaged laminar transient result. The velocity and temperature behaviour of this model at the top of the centreline also contains features that disagree with the laminar solution.
- The centreline velocity for the *k* - ω SST based models is higher than the laminar solution (thereby predicting a higher fuel salt circulation rate), but the shape of the velocity profile is in better agreement than for the realizable *k* - ϵ model.
- The agreement for temperature is worst for realizable *k* - ϵ , although not substantially different to the *k* - ω SST results – they predict higher temperatures than the laminar solution. The exception to this is the intermittency model, which produces a much closer centreline and cross-pin profile temperature prediction.
- The intermittency model prediction for centreline velocity at the base of the pin produces a high velocity peak that is not seen in the other models or the laminar results.

None of the RANS models tested are ideal, and which to choose depends on what other features of the analysis of the fuel performance depend on the behaviour. The effect of the difference in

¹ Averaging the flow for 11 s was considered sufficient to provide a comparison to the RANS models; a significantly longer simulation would be necessary to remove the remaining noise to produce a smooth profile.

Single Fuel Pin Turbulence Assessment

predictions is not limited to the impact on maximum fuel temperatures for thermal reasons. The temperatures also impact on neutronic thermal feedback, and hence the reactivity and heat source. The circulation rate of the fuel salt in the pin will affect the transport and distribution of DNPs, and so could have some impact on the fission power profile.

While it works well for the temperature prediction in this case, the intermittency model has not been selected because it was not designed for application to these conditions, and has a sensitivity to solution trajectory as described in Section 3.1. It is possible that it will work well at some locations or under some conditions, but be highly inaccurate at others – further work would be necessary to prove its robustness. The $k-\omega$ SST model with low Re corrections was retained in preference to the realizable $k-\epsilon$ model, primarily because of the features seen in the latter at the top of the tube, the slightly better temperature predictions of $k-\omega$ SST, and the confidence in its ability to revert to fully laminar flow where necessary.

Although the laminar case is a spatially resolved, transient flow with no turbulence modelling applied, not enough rigour has been applied to refer to it as a DNS solution. There has been no assessment of mesh and time resolution sensitivity, and the result is for a single, simplified case without varying density or viscosity. However, the result does offer reassurance that RANS models are activating and generating turbulence fields for good reasons, and justifies that it is necessary to use one in the resolved fuel assembly for physical, not purely numerical reasons.

It is surprising to an extent that the RANS models performed as well as they have, because this situation is far from the typical flows used to calibrate their parameters. The turbulence generation in the models is likely to be partly caused by temperature gradients in the buoyancy production term, as well as by shear. The details of the production and low Reynolds number damping processes are expected to be important to the behaviour and success of the models. It would be necessary to study more features of the result (such as Reynolds stresses) in detail at more locations to ensure that the apparent success of the RANS models is for the 'right' reasons, and to potentially improve or adjust the models for this situation.

This case also demonstrates a feature of the meshing approach that could cause solution difficulties in other situations – the very elongated cells in the coarse mesh cannot represent the unstable laminar flow. As an instability forms that turns the flow, it cannot be resolved and represented by the high aspect ratio cells, so the solution does not perform well.

Single Fuel Pin Turbulence Assessment

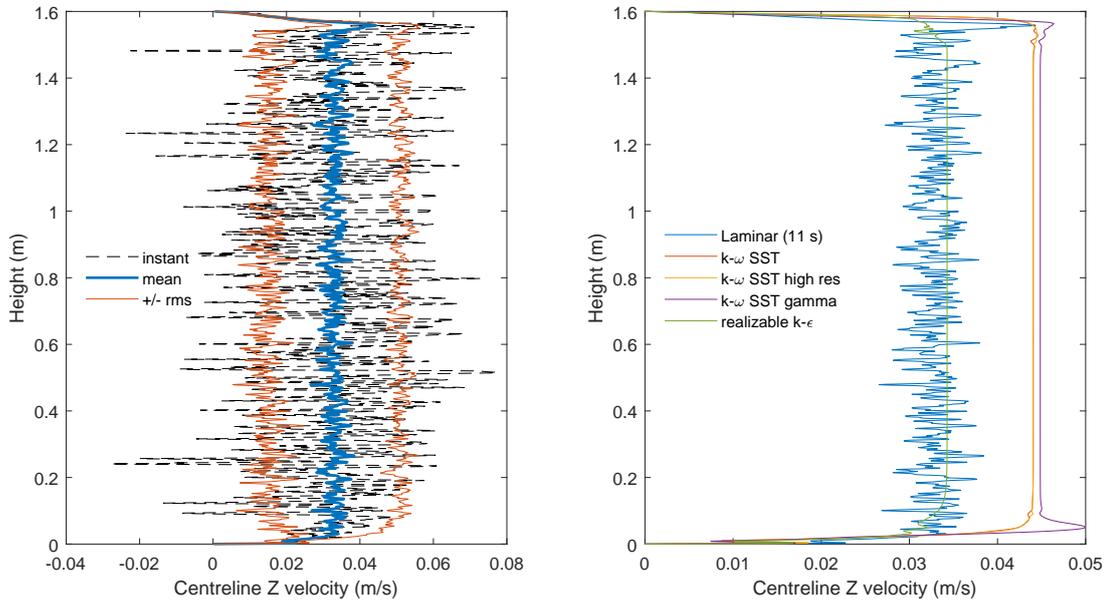


Figure A.2: Single pin centreline vertical velocity. Left: transient laminar case with mean value over 11 s compared to an instantaneous snapshot, and the envelope of the RMS value either side. Right: mean transient laminar solution compared to steady-state RANS solutions.

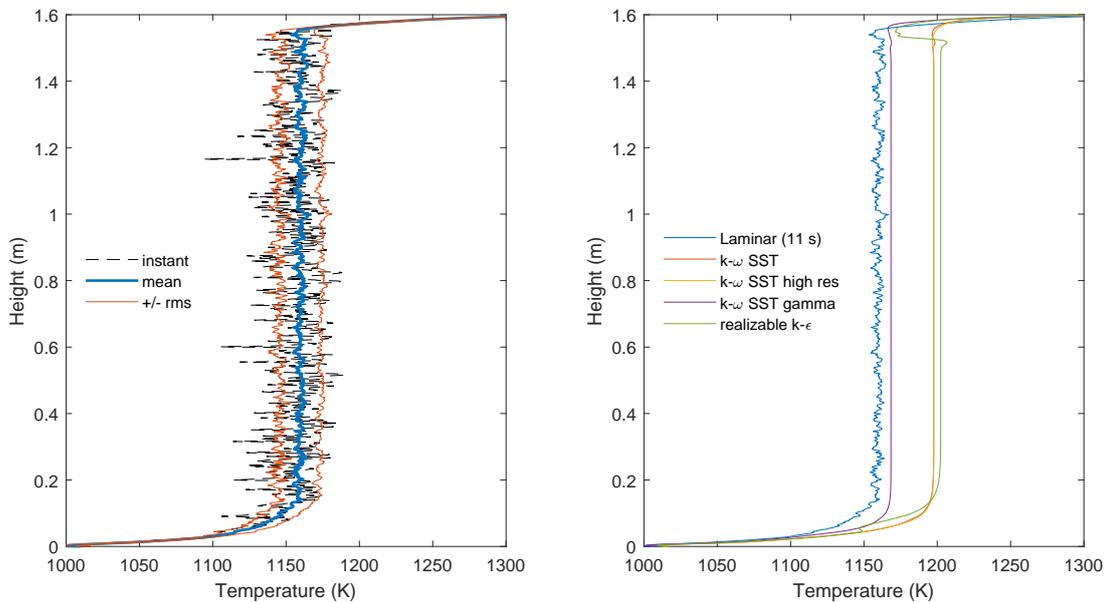


Figure A.3: Single pin centreline temperature. Left: transient laminar case with mean value over 11 s compared to an instantaneous snapshot, and the envelope of the RMS value either side. Right: mean transient laminar solution compared to steady-state RANS solutions.

Single Fuel Pin Turbulence Assessment

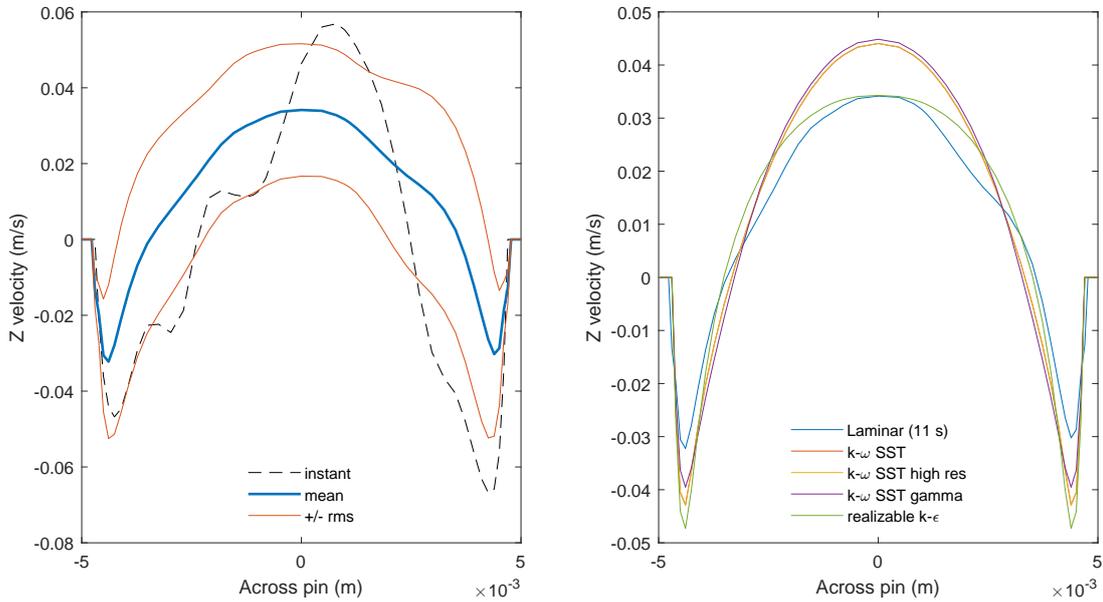


Figure A.4: Single pin cross-pin vertical velocity profile at mid-height. Left: transient laminar case with mean value over 11 s compared to an instantaneous snapshot, and the envelope of the RMS value either side. Right: mean transient laminar solution compared to steady-state RANS solutions.

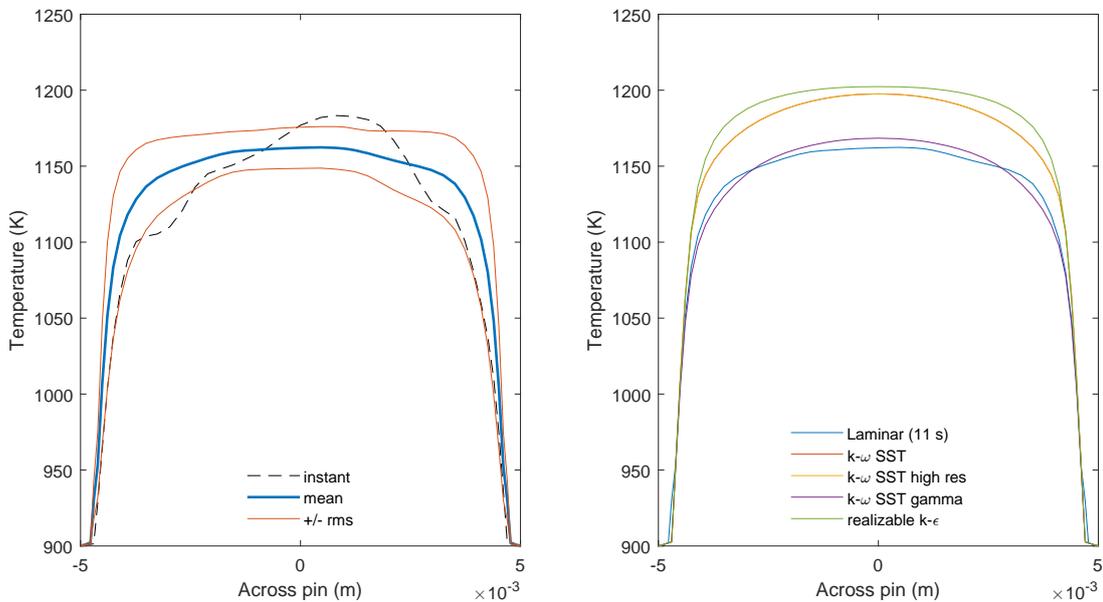


Figure A.5: Single pin cross-pin temperature profile at mid-height. Left: transient laminar case with mean value over 11 s compared to an instantaneous snapshot, and the envelope of the RMS value either side. Right: mean transient laminar solution compared to steady-state RANS solutions.



Study B

