ABSTRACT

SEDOTA, CHRISTOPHER STEPHEN. Uncertainty Quantification in Nuclear Reactor Modeling using Stochastic Sampling with the Virtual Environment for Reactor Applications. (Under the direction of Scott Palmtag.)

Uncertainty quantification is an important component of model development. In uncertainty quantification, uncertainties in the inputs to a model are propagated through the model, and the resulting uncertainty in the model output is analyzed. The various methods for uncertainty quantification can be divided into two classes: statistical methods and deterministic methods. Stochastic sampling, a statistical method, is exclusively used in this study. The model is the Virtual Environment for Reactor Applications (VERA). The uncertain inputs to the model are the multigroup cross sections and fuel rod manufacturing parameters. One thousand perturbed multigroup cross section libraries were generated. These perturbed cross section libraries were used with a variety of benchmark problems based on the Three Mile Island Unit 1 Pressurized Water Reactor, the Peach Bottom Unit 2 Boiling Water Reactor, and the Watts Bar Unit 1 Pressurized Water Reactor. When only considering the uncertainty in the effective neutron multiplication factor due to cross section uncertainty, the relative standard deviation for all cases was determined to be approximately 0.5 – 0.6 %.

In this study, one thousand perturbed cross section libraries were used for stochastic sampling. A major question that this study aims to address is whether one thousand cross section libraries is sufficient for quantification of cross section uncertainty using stochastic sampling, or whether fewer libraries could be used. The use of more or fewer cross section libraries has major implications on computational cost, particularly for very large cases. For all cases that were examined, the estimated relative standard error of the standard deviation estimate was approximately 2 % when one thousand cross section libraries were used. The use of 500 cross section libraries, which is a reduction of computational cost of 50 %, increased the relative standard error to approximately 3 % for all cases. Even when only 250 cross section libraries were used, which is a 75 % reduction in computational cost from the full one thousand libraries, the relative standard error of the standard deviation estimate remained below 5 % for all cases. Therefore it was determined that with well-sampled cross sections, fewer than the full one thousand cross section libraries could be used to generate valid uncertainty results.

Uncertainty in the dimensions and composition of the fuel rods was also quantified. The VERA format of individual pin specification facilitated this analysis, and a tool was developed to perform material parameter perturbations on a base VERA input deck. It was found that the size of the problem had a significant impact on the uncertainty due to material parameter uncertainties. For instance, while the estimated relative standard deviation of the effective neutron multiplication factor for the Three Mile Island Unit 1 pincell was 0.07 %, the relative standard deviation for the Three Mile Island Unit 1 lattice was 0.007 %, a factor of 10 lower. The effect of the exact application of material parameter uncertainties was also examined.

Uncertainty Quantification in Nuclear Reactor Modeling using Stochastic Sampling with the Virtual Environment for Reactor Applications

by
Christopher Stephen Sedota

A thesis submitted to the Graduate Faculty of
North Carolina State University
in partial fulfillment of the
requirements for the Degree of
Master of Science

Nuclear Engineering

Raleigh, North Carolina

2020

APPROVED BY:

Jason Hou	Ralph Smith

Scott Palmtag
Chair of Advisory Committee

DEDICATION

To my parents, who supported me every step of the way.

BIOGRAPHY

Chris Sedota was born in Rochester, NY and grew up in Mooresville, NC. He attended Lake Norman High School, where a number of passionate science and mathematics teachers influenced him to pursue engineering in higher education. He attended North Carolina State University and obtained a Bachelor of Science in Nuclear Engineering in 2018. He then continued his education at North Carolina State University and is pursuing a Master of Science in Nuclear Engineering.

ACKNOWLEDGEMENTS

I would like to thank my advisor, Dr. Scott Palmtag, for his immense help in the completion of this thesis. I would also like to thank Dr. Jason Hou and Dr. Ralph Smith, who agreed to serve on my graduate committee, and all the great educators who helped me along the way. Finally, I would like to thank Dr. Matt Jessee and Dr. Kang Seog Kim for generating the perturbed cross section libraries used in this study.

This research was supported by the Consortium for Advanced Simulation of Light Water Reactors (http://www.casl.gov), an Energy Innovation Hub (http://www.energy.gov/hubs) for Modeling and Simulation of Nuclear Reactors under U.S. Department of Energy Contract No. DE-AC05-00OR22725.

This research also made use of the resources of the High Performance Computing Center at Idaho National Laboratory, which is supported by the Office of Nuclear Energy of the U.S. Department of Energy and the Nuclear Science User Facilities under Contract No. DE-AC07-05ID14517.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER

1

INTRODUCTION

Uncertainty quantification is a critical component of the analysis of any system. If there are uncertainties in the inputs to a system, there will necessarily be uncertainty in the response of the system. Knowledge of the uncertainty in a system can help establish confidence in the system and allow designers to make more informed decisions on the edge of the design space, which is where efficient designs exist.

When dealing with any complex system, modeling and simulation is used to gain information of the system in place of physical experimentation whenever possible. Certainly nuclear systems can be characterized as complex systems. Consider the cost associated with physically testing a new fuel assembly design in an active power reactor. Neglecting even the initial cost to manufacture the novel fuel assembly, which would likely exceed the estimated 1 million USD cost of manufacturing a typical fuel assembly [Kok09], the cost associated with potentially reducing the power production of a commercial reactor for experimentation is very large. If the reactor needed to be taken off-line for experimentation, it could cost the operators of the reactor upwards of 500,000 USD per day [Str95], which is the approximate cost typically associated with refueling outages.

The other alternative for physical experimentation in a power-producing reactor is the construction of dedicated experimental facilities. This is another extremely expensive option. For instance, the Versatile Test Reactor (VTR) project, which is slated to begin construction at Idaho National Laboratory in 2022, has an initial cost estimate of 3-6 billion USD [Ida]. While physical experimentation is an unavoidable and necessary component of gaining knowledge of a system, in many cases it is avoided as much as possible due to its extreme cost. Clearly it is in the interest of those associated with nuclear power to use modeling and simulation in place of physical experimentation wherever possible.

The high cost associated with physical experimentation has prompted the development of very sophisticated models to model and predict reactor behavior. The modeling of light water reactors (LWRs) requires a wide range of physics to be considered, from the interaction of neutrons with matter to the complex multi-phase flow of the water coolant through the core. The Consortium for the Advanced Simulation of Light Water Reactors (CASL) was established by the US Department of Energy as a collaborative effort between universities, national laboratories, and industry participants with the goal of improving the modeling of light water reactor modeling. To achieve this goal, a high-fidelity reactor simulation suite called the Virtual Environment for Reactor Applications (VERA) was developed.

As VERA has matured into a production-level code system, increased interest has been placed on using it as a model for uncertainty quantification. Uncertainty quantification is an activity that can only be undertaken when there is sufficient confidence in the underlying model. This confidence is gained through rigorous solution and code verification, which ensures that errors are not present in the coding implementation of the underlying mathematical model [Roy05], and validation exercises, where the predictive capability of the model is assessed against experimental data. The predictive capability of VERA tools were largely tested using a series of progression problems based on the Watts Bar Unit 1 reactor [God14], a few of which have associated physical data.

The governing equations of the neutronics component of nuclear systems have been relatively well-understood from the inception of the field of nuclear engineering. The neutron transport equation, for instance, which accounts for the movement of neutrons about a system, was developed in the 1940s [Dud76]. With the rapid advancement of computing since then, the neutron transport equation has been widely implemented in computer codes to model nuclear systems, and still provides the basis for neutronics codes today.

The neutronics component of VERA is the Michigan PArallel Characteristics based Transport (MPACT) code [Koc17]. MPACT serves as the model for most of the uncertainty analysis in this study. MPACT coupled with the thermal hydraulic code CTF [Sal16] serves as the model for uncertainty analysis in a 3-dimensional Watts Bar Unit 1 system with thermal hydraulic feedback.

1.1 Motivation

With computational uncertainty quantification, the ultimate goal is to obtain high-quality uncertainty information at a low computational cost. Uncertainty quantification can be used to reduce conservative margins that have been built into reactor design. With little or no knowledge of the uncertainty of a system, generous margin must be included to ensure safety. If high-quality uncertainty information is known, then this generous margin can be reduced while still ensuring safety. Reduction in margin often means more economical design and operation can be achieved.

As will be discussed in the following sections, there are significant computational cost differences between the two main classes of uncertainty quantification. While this study is focused on stochastic sampling, which is the more computationally expensive technique, considerable attention will

be given to two concepts that could reduce computational expense: the potential use of simpler problems to serve as surrogate problems for more complex problems, and the number of samples necessary to produce reliable uncertainty information for different cases. Chapters 3 and 4 will show how simple cases compare to more complex, and thus more computationally expensive, cases. Chapter 5 deals with the latter consideration – the sample size required to generate high-quality uncertainty information for a wide range of cases.

1.2 Classes of Uncertainty Quantification

Uncertainty analysis can be broadly divided into two classes: stochastic methods and deterministic methods. Although stochastic methods are exclusively used in this study, it is beneficial to understand both classes, and the advantages and disadvantages that each method brings.

1.2.1 Uncertainty Quantification using Deterministic Methods

The k-eigenvalue neutron transport problem [Lew93] is

$$\left[\hat{\Omega}\cdot\vec{\nabla}+\Sigma(\vec{r},E)\right]\psi\left(\vec{r},\hat{\Omega},E\right)-\int dE'\int d\Omega'\,\Sigma_s\left(\vec{r},E'\rightarrow E,\hat{\Omega}'\cdot\hat{\Omega}\right)\psi\left(\vec{r},\hat{\Omega}',E'\right) = \frac{\chi(E)}{k}\int dE'\,\nu\Sigma_f\left(\vec{r},E'\right)\int d\Omega'\,\psi\left(\vec{r},\hat{\Omega}',E'\right), \quad (1.1)$$

where

$\Sigma(\vec{r},E)$ = total macroscopic cross section at position $\vec{r}$ at energy E,
$\psi(\vec{r},\hat{\Omega},E)$ = angular neutron flux at position $\vec{r}$ in direction $\hat{\Omega}$ at energy E,
$\Sigma_s\left(\vec{r},E'\rightarrow E,\hat{\Omega}'\cdot\hat{\Omega}\right)$ = macroscopic scattering cross section at position $\vec{r}$ from energy E' to E and from direction $\hat{\Omega}$' to $\hat{\Omega}$,
$\chi(E)$ = fission spectrum at energy E,
k = neutron multiplication factor,
$\nu\Sigma_f\left(\vec{r},E'\right)$ = mean number of fission neutrons produced in a fission caused by a neutron with energy E' multiplied by the macroscopic fission cross section at position $\vec{r}$ at energy E'.

In operator k-eigenvalue form, the neutron transport equation can be written as

$$H\psi=\frac{1}{k}G\psi, \quad (1.2)$$

where ψ is the angular neutron flux, H is the transport operator, and G is the fission operator. The transport operator H and the fission operator G are defined in Equations 1.3a and 1.3b respectively.

$$H\psi = \left[\hat{\Omega}\cdot\vec{\nabla} + \Sigma(\vec{r},E)\right]\psi\left(\vec{r},\hat{\Omega},E\right) - \int dE' \int d\Omega'\, \Sigma_s\left(\vec{r},E'\to E,\hat{\Omega}'\cdot\hat{\Omega}\right)\psi\left(\vec{r},\hat{\Omega}',E'\right), \tag{1.3a}$$

$$G\psi = \chi(E)\int dE'\, \nu\Sigma_f\left(\vec{r},E'\right)\int d\Omega'\, \psi\left(\vec{r},\hat{\Omega}',E'\right). \tag{1.3b}$$

The adjoint of Equation 1.2 is

$$H^+\psi^+ = \frac{1}{k^+}G^+\psi^+, \tag{1.4}$$

with the transport operator H^+ and the adjoint fission operator G^+ defined in Equations 1.5a and 1.5b as [Lew93]

$$H^+\psi^+ = \left[-\hat{\Omega}\cdot\vec{\nabla} + \Sigma(\vec{r},E)\right]\psi^+\left(\vec{r},\hat{\Omega},E\right) - \int dE' \int d\Omega'\, \Sigma_s\left(\vec{r},E\to E',\hat{\Omega}\cdot\hat{\Omega}'\right)\psi^+\left(\vec{r},\hat{\Omega},E\right), \tag{1.5a}$$

$$G^+\psi^+ = \nu\Sigma_f(\vec{r},E)\int dE'\, \chi\left(E'\right)\int d\Omega'\, \psi^+\left(\vec{r},\hat{\Omega}',E'\right). \tag{1.5b}$$

It can be shown by multiplying both the forward and adjoint equations by the opposing neutron flux, integrating over independent variables, and taking the difference of the two equations yields

$$\left(\frac{1}{k} - \frac{1}{k^+}\right)\langle\psi^+ G\psi\rangle = 0, \tag{1.6}$$

where $\langle\rangle$ denotes integration over all independent variables, which in this case is the volume V, the solid angle Ω, and the energy E.

$$\langle\cdot\rangle = \int dV \int d\Omega \int dE. \tag{1.7}$$

If the forward and adjoint neutron fluxes are both positive, then the only way for Equation 1.6 to hold is if the k-eigenvalues of the neutron transport equation and its adjoint are equal,

$$k = k^+. \tag{1.8}$$

Now assume that Equation 1.2 describes the perturbed system, and that the unperturbed system is described by Equation 1.9 as

$$H_0\psi = \frac{1}{k_0}G_0\psi. \tag{1.9}$$

The perturbed quantities are related to the unperturbed quantities by Equations 1.10a - 1.10d,

$$H = H_0 + \delta H, \tag{1.10a}$$

$$G = G_0 + \delta G, \tag{1.10b}$$

$$k = k_0 + \delta k, \tag{1.10c}$$

$$\psi = \psi_0 + \delta \psi, \tag{1.10d}$$

where δH, δG, δk, and $\delta \psi$ are small perturbations in the transport operator, the fission operator, the k-eigenvalue, and the neutron angular flux, respectively. Substituting these perturbed quantities into integral form of the neutron transport equation in Equation 1.2 yields

$$(H_0 + \delta H)(\psi_0 + \delta \psi) - \frac{1}{k_0 + \delta k}(G_0 + \delta G)(\psi_0 + \delta \psi) = 0. \tag{1.11}$$

If we assume that the perturbation is small enough that terms of order δ^2 can be neglected, the perturbed eigenvalue term can be approximated linearly by

$$\frac{1}{k_0 + \delta k} \approx \frac{1}{k_0} - \frac{\delta k}{k_0^2}. \tag{1.12}$$

Using this approximation in Equation 1.11, multiplying by the adjoint neutron flux, and integrating over the independent variables, Equation 1.13 is obtained.

$$\frac{\delta k}{k_0^2} = \frac{\langle \psi_0^+ \left(k_0^{-1} \delta G - \delta H \right) \psi_0 \rangle}{\langle \psi_0^+ G_0 \psi_0 \rangle}. \tag{1.13}$$

Equation 1.13 can be used to approximate the reactivity effect of perturbations of the cross sections or material parameters, which are accounted for in the δG and δH operators. Since the response of interest is the k-eigenvalue, Equation 1.13 can be used to calculate sensitivity indices for cross section and material parameters. If the perturbations δG and δH are replaced by partial derivatives of the transport operators with respect to a macroscopic cross section Σ, the sensitivity of k due to changes in the macroscopic cross section is, by the chain rule, [Rea16]

$$s_{k,\Sigma} \equiv \frac{\Sigma}{k_0} \frac{\partial k}{\partial \Sigma} = -\frac{\Sigma}{k_0} \frac{\langle \psi_0^+ \left(\frac{\partial H}{\partial \Sigma} - \frac{1}{k_0} \frac{\partial G}{\partial \Sigma} \right) \psi_0 \rangle}{\langle \psi_0^+ \frac{1}{k_0^2} G \psi_0 \rangle}. \tag{1.14}$$

In practice, these would be discretized and solved for the unperturbed neutron flux ψ and the unperturbed adjoint neutron flux ψ^+ iteratively.

With sensitivity indices calculated and stored in the sensitivity vector $S^T = \left[s_{k,\Sigma_1}, s_{k,\Sigma_2}, ... \right]$, and the covariance matrix V of cross sections from a library's covariance data, the variance of the k-eigenvalue can be directly calculated using Equation 1.15.

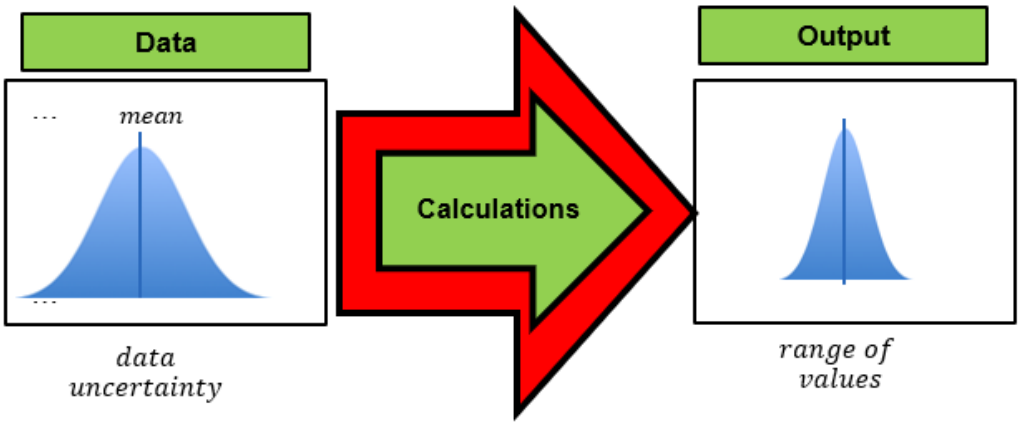

Figure 1.1 Stochastic sampling flow chart.

$$\mathrm{Var}(k) = S^T V S \tag{1.15}$$

This is an attractive method for characterizing the uncertainty in k when only neutronics are considered. In applications where thermal hydraulic feedback to the neutronics are considered, it is not possible to form the adjoint problem necessary for the calculation of the sensitivity coefficients. Even in solely neutronics codes, the adjoint solution may not be available. Therefore deterministic methods are only feasible for some stand-alone neutronics problems.

1.2.2 Uncertainty Quantification using Stochastic Sampling

Stochastic sampling as a technique for uncertainty quantification is a statistical method, as opposed to the purely deterministic method. In stochastic sampling, the parameter uncertainty information in the covariance matrix V is used as a sampling distribution. A certain number of parameter realizations are generated and used as inputs to generate responses. A one-to-one mapping of parameter realizations to responses is generated, and the resulting response distribution can be analyzed to quantify the effect of parameter uncertainty on the response. Figure 1.1 shows the flow of uncertainty information from the inputs through the model and to the response.

No direct knowledge of the model used to generate the response values is needed beyond the knowledge that the model is sufficiently representative of the real system. The model can, in this case, be treated as a black box, and the response values can be strictly analyzed from a statistical perspective. For the characterization of near-normal distributions, typical statistics used to quantify uncertainty are the mean, variance, and standard deviation.

Assume we have a model with m input parameters, $\vec{q} = [q_1, q_2, ..., q_m]$, and the m–dimensional parameter space has a fully-defined m-dimensional probability distribution. One sample from this m-dimensional probability distribution defines all m values of $\vec{q}$ for that realization. When the model is evaluated with the sampled parameter values, one model response is obtained. If the

model evaluated at the parameter realization $\vec{q}^1$ is denoted $f(\vec{q}^1)$, the response generation can be described by Equation 1.16, where k_1 is the response of the model due to $\vec{q}^1$.

$$f(\vec{q}^1) = k_1 \tag{1.16}$$

If n input parameter realizations are generated, n responses will be generated from n model evaluations, as shown in Equation 1.17 where $\vec{Q} = \left[\vec{q}^1, \vec{q}^2, ..., \vec{q}^n\right]$ and $\vec{k} = [k_1, k_2, ..., k_n]$.

$$f(\vec{Q}) = \vec{k} \tag{1.17}$$

The mean $\bar{k}$ of the response vector $\vec{k}$ can be calculated using Equation 1.18.

$$\bar{k} = \frac{1}{n}\sum_{i=1}^{n} k_i \tag{1.18}$$

The variance s_k^2 of the response vector $\vec{k}$ can be calculated using Equation 1.19, and the standard deviation s_k of the response vector $\vec{k}$ is the square root of the variance s_k^2, as shown in Equation 1.20.

$$s_k^2 = \frac{1}{n-1}\sum_{i=1}^{n}\left(k_i - \bar{k}\right)^2 \tag{1.19}$$

$$s_k = \sqrt{s_k^2} = \sqrt{\frac{1}{n-1}\sum_{i=1}^{n}\left(k_i - \bar{k}\right)^2} \tag{1.20}$$

The variance s_k^2 of the response vector $\vec{k}$ obtained through stochastic sampling can be compared to the variance obtained using Equation 1.15 through the deterministic method. It should be noted that while the variance obtained using Equation 1.15 is a direct calculation, the variance obtained through stochastic sampling is only an estimate of the true variance of k. Therefore it is necessary to evaluate the convergence behavior of the variance and other statistics obtained through stochastic sampling. This is a point that will be particularly emphasized in Chapter 5.

1.2.3 Comparison of Methods

The one-to-one mapping of the parameter realizations to the response vector means that many more model evaluations need to be performed in the stochastic sampling to obtain valid statistical estimates. For the deterministic method, each sensitivity coefficient can be calculated from a single forward and adjoint model evaluation. For m parameters, $2m$ total calculations of Equation 1.14 must be made to obtain the variance of the response. For stochastic sampling, the forward neutron flux calculation must be carried out through forward model evaluations until the response vector is large enough to calculate statistics with reasonable levels of confidence.

In addition, a useful byproduct of the deterministic method is the calculation of sensitivity coefficients for each parameter. A similar amount of sensitivity data from stochastic sampling could only be obtained by constructing a response distribution for each parameter sampled individually.

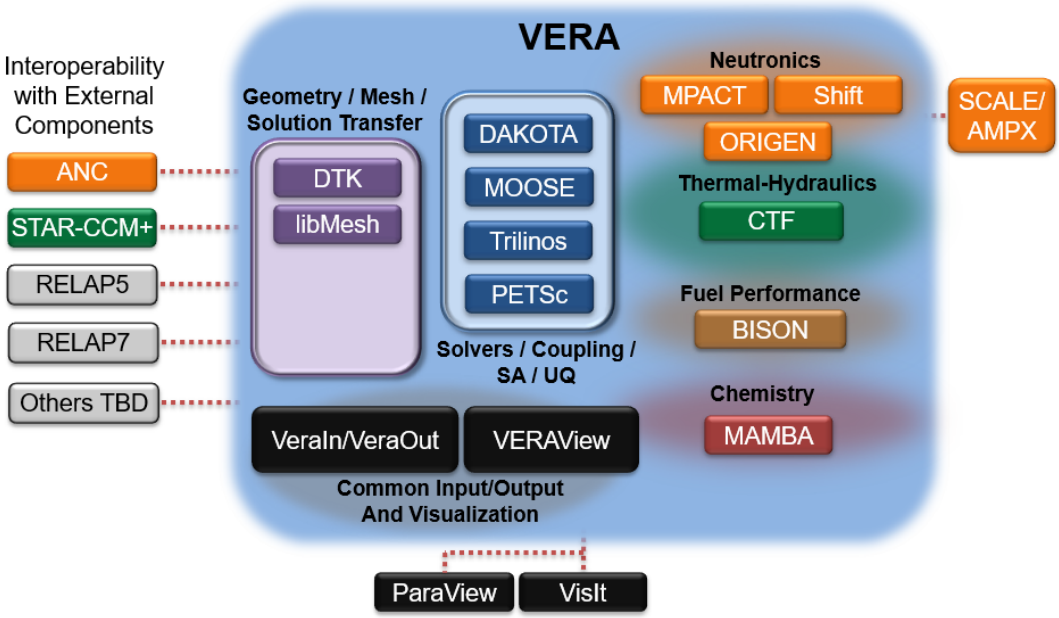

Figure 1.2 VERA code structure [CAS18].

This may be feasible for problems with low dimensionality, but is virtually impossible for analyzing cross section uncertainty because of the extremely high number of input parameters.

While the deterministic method has the advantage in terms of pure computational cost, stochastic sampling has certain important advantages. Perhaps the most important of these advantages is the fact that the model can serve as a black-box for the generation of responses. While in the deterministic method knowledge of the fundamental equations was required to calculate the adjoint neutron flux, calculation of the adjoint is unnecessary in stochastic sampling. This is important in cases where an adjoint is difficult or even impossible to calculate, such as when secondary effects are added. The deterministic method is only valid for purely neutronic applications. If thermal hydraulic feedback were added, which is certainly present in any real nuclear system, the deterministic method would not be able to capture these effects. Figure 1.2 depicts a diagram of the VERA code system as a whole. MPACT, the neutronics solver, is only one component of the whole environment. When thermal-hydraulics using CTF, fuel performance using BISON, and chemistry using MAMBA are included in the calculation, stochastic sampling is the only way to obtain uncertainty data. Throughout this thesis, stochastic sampling allowing for a code system to serve as a model for uncertainty quantification with no modification necessary is described as a “non-invasive” method.

As computational costs continue to decline with the development of high performance computing, the advantage presented by deterministic methods of limiting computational cost also declines. For simple problems with only neutronics, deterministic methods may be superior to stochastic sampling methods due to the individual sensitivity coefficients generated for each input parameter. For complex simulations, however, stochastic sampling is the obvious choice. Because nuclear systems are invariably extremely complex multi-physics systems, stochastic sampling is

Table 1.1 Pros and cons of deterministic and stochastic sampling methods for uncertainty quantification.

Method	Pros	Cons
Deterministic	• Identification of individual contributions to total uncertainty. • Lower computational cost.	• Adjoint solver required. • Very difficult for multiphysics applications.
Sampling	• No adjoint solution needed. • Easily implemented for multiphysics problems.	• Large computational cost to obtain high-quality statistics. • Identification of individual contributions to total uncertainty limited by parameter dimensionality.

the natural choice for uncertainty quantification in this study. A summary of the pros and cons of the deterministic method and stochastic sampling for uncertainty quantification is presented in Table 1.1.

1.3 Neutron Multiplication Factor

Perhaps the most important quantity in reactor analysis is the effective neutron multiplication factor. The effective neutron multiplication factor (k-eff) can be thought of as the rate at which neutrons are produced in a system normalized by the rate at which neutrons are destroyed or lost. Neutrons can only be produced as a result of a fission reaction within the system, but can be lost either by being absorbed by an atom within the system or by streaming out of the system. The rate of neutron production is simply the rate of fission reactions multiplied by the number of neutrons produced by one of these reactions.

k-eff is a commonly used quantity in reactor analysis because it characterizes the criticality of the system, as shown in Equation 1.21.

$$\text{k-eff} \rightarrow \begin{cases} > 1; \text{ supercritical} \\ = 1; \text{ critical} \\ < 1; \text{ subcritical} \end{cases} \tag{1.21}$$

k-eff is dependent on the geometry of the system, the neutron flux in the system, and the

microscopic cross sections of the material of the system. The uncertainty in the microscopic cross sections can be propagated through the calculation of k-eff to obtain the uncertainty in k-eff. k-eff is therefore used as the sole response for this study. It should be noted that k-eff and the k introduced in Equation 1.1 refer to the same quantity, k-eff $= k$.

1.4 Cross Section Uncertainty

Nuclear cross sections have been calculated using experimental data since at least the early 1940s [Seg53]. However, due to the sensitive nature of nuclear research in the 1940s, international organizations independently generated nuclear data for internal use. Differences in nuclear data storage and communication between national laboratories in the United States prompted the creation of the standard Evaluated Nuclear Data File (ENDF) in the early 1960s.

Evaluation of nuclear data involves combining sufficiently high-quality experimental data with sufficiently high-quality model calculations to approximate the true value of cross sections [Trk18]. As with any experimental procedure and model evaluation, there is uncertainty in evaluated cross section data. While some early (1950s) nuclear data was reported with no experimental uncertainty [Hol01], it has been a focus since the creation of ENDF to compile uncertainty associated with cross section data.

While the most recent ENDF release was ENDF/B-VIII.0 in 2018, it has not yet been widely adopted. ENDF/B-VII.1 [Cha11], released in 2011, is more commonly used. One of the most important advances from ENDF/B-VII.0 to ENDF/B-VII.1 is the inclusion of much more covariance data. While ENDF/B-VII.1 included covariances for only 27 materials, ENDF/B-VII.1 includes covariances for 190 materials. This increase in covariance data over the span of only five years illustrates the rapidly growing interest in high-quality nuclear data uncertainty information.

Unfortunately, while ENDF is largely used as the standard source of cross section data in the United States, other cross section libraries are used around the world. For instance, another commonly used cross section library is the Joint Evaluated Fission and Fusion File (JEFF) [OEC09], produced by the Organization for Economic Co-Operation and Development (OECD) Nuclear Energy Agency (NEA). The lack of a universally accepted nuclear data source somewhat complicates the collection of benchmark data [Hou19].

This study uses cross section uncertainty data included in SCALE [Rea16], which uses covariance data included in ENDF/B-VII.1 for 187 materials and supplements with additional covariance data generated by Oak Ridge National Laboratory (ORNL). Considerable work goes into converting the point-wise nuclear data in the ENDF format into a multigroup library that can be used for specific simulations. The details of the data processing are outside the scope of this study, but further discussion of cross section processing can be found in Chapter 2.

Table 1.2 UAM LWR benchmark problem naming convention.

Case	Abbreviated Name
PWR **P**incell at H**Z**P	**PPZ**
PWR **P**incell at H**F**P	**PPF**
BWR **P**incell at H**Z**P	**BPZ**
BWR **P**incell at H**F**P	**BPF**
PWR **L**attice **U**nrodded at H**Z**P	**PLUZ**
PWR **L**attice **U**nrodded at H**F**P	**PLUF**
PWR **L**attice **R**odded at H**Z**P	**PLRZ**
PWR **L**attice **R**odded at H**F**P	**PLRF**
BWR **L**attice **U**nrodded at H**Z**P	**BLUZ**
BWR **L**attice **U**nrodded at H**F**P	**BLUF**
BWR **L**attice **R**odded at H**Z**P	**BLRZ**
BWR **L**attice **R**odded at H**F**P	**BLRF**
PWR **C**olorset at H**Z**P	**PCZ**
PWR **C**olorset at H**F**P	**PCF**

1.5 Material Uncertainty

While most of the focus of uncertainty quantification of nuclear systems is focused on the uncertainty of nuclear data, uncertainties in material parameters due to imperfect manufacturing processes must be considered in order to get a full picture of the total uncertainty of the system. Chapter 6 contains analysis of the effect of certain material parameters with uncertainties on the neutron multiplication factor. The uncertainty in the neutron multiplication factor stemming solely from cross section uncertainty and the uncertainty from solely the material uncertainties is compared, as well as uncertainty from both of these sources combined.

1.6 Benchmark Problem Naming Convention

For the UAM LWR benchmark [Iva13] problems, a naming convention was developed to more easily refer to specific problem configurations. Details of the UAM LWR benchmark can be found in Chapter 3. The naming conventions, which are used throughout this thesis, are shown in Table 1.2. In general, the first letter of the abbreviation corresponds to the type of reactor and the second letter corresponds to the size or complexity of the model. For lattice cases, the third letter in the abbreviation indicates the rod position, but this does not apply for the other cases. The final letter corresponds to the state of the reactor. More details can be found in Chapter 3.

1.7 Research Objectives

The purpose of this study is first and foremost to establish stochastic sampling with VERA tools as a valid method for uncertainty quantification. To accomplish this, a number of benchmark test cases were examined, and the results from this research were compared to results submitted by other benchmark participants. Some of the results were also submitted to the UAM LWR benchmark team for inclusion in the benchmark.

While the focus of the benchmark cases is largely on the quantification of cross section uncertainties, another research objective was to also quantify uncertainty due to material parameters. The VERA common input format, with individual pin dimension specification, can easily accommodate stochastic sampling to quantify material parameter uncertainties. The uncertainty in k-eff due to material uncertainties is compared with the uncertainty due to cross section uncertainties.

Since stochastic sampling is a computationally expensive uncertainty quantification method, it is beneficial to limit the number of samples used. Therefore jackknife resampling was used to estimate the standard error of the estimates of the standard deviation at different sample sizes for all cases. Also in the interest of limiting computational expense, the possibility of using a simplified, less expensive model for uncertainty quantification is discussed.

1.8 Thesis Layout

Chapter 2 describes the methodology used for uncertainty quantification in this study. Details of the tools used for cross section library generation and neutronics calculations are presented. A description of the new material perturbation tool developed during the course of this research is provided. A comparison of the computational cost of all of the models used in this study is also provided.

In Chapter 3, results for cross section uncertainty quantification for the UAM LWR benchmark problems are presented and compared. The different models, ranging from a pincell to a colorset for the PWR and from a pincell to a rodded lattice for the BWR are described in detail. The results generated from this study are compared to the published average benchmark results. Chapter 4 also presents cross section uncertainty quantification results, but for the Watts Bar Unit 1 full-core. Again a range of cases, from a two-dimensional full-core to a three-dimensional full-core with thermal hydraulic feedback, are examined and compared.

Chapter 5 addresses the question of a sufficient sample size for cross section uncertainty quantification using stochastic sampling. The theory behind the standard error of estimators is described, and jackknife resampling is introduced as a method to approximate the standard error of estimators. The estimated standard errors of the standard deviation estimates for k-eff due to cross section uncertainty are tabulated. The standard error estimates are presented for a select number of sample sizes for all cases described in Chapters 3 and 4.

Chapter 6 describes the uncertainty quantification of k-eff due to material parameter uncer-

tainties. The two cases examined in detail are the PWR pincell and the PWR lattice. The relative contribution to the uncertainty in k-eff from each material parameter is presented for the pincell model in a sensitivity study using stochastic sampling. The overall uncertainty in k-eff due to different combinations of material parameter uncertainties is then calculated for the PWR lattice, and the results are compared to the contributions to k-eff uncertainty from the cross sections.

Finally, conclusions of the study are presented in Chapter 7. The research objectives presented in the previous section are addressed, and recommendations are made for future work.

CHAPTER

2

METHODOLOGY

A general description of the techniques and tools used for stochastic sampling will follow. In short, cross section uncertainty data in SCALE [Rea16] was used by the Medusa module in XSUSA [Wil13] to calculate multiplicative cross section perturbation factors which were in turn used by Sampler to generate 1000 perturbed cross section libraries. These perturbed cross section libraries were then converted from their default AMPX format into a format usable by the neutronics code MPACT [Koc17], and served as inputs for the uncertainty analysis. The resulting values of k-eff calculated in MPACT for each of the 1000 perturbed cross section libraries served as the model response. Additionally, a perturbation method was developed specifically for MPACT to account for material parameter distributions. Python was used to sample from material parameter distributions and generate perturbed MPACT input files for stochastic sampling.

2.1 SCALE

SCALE [Rea16] was used to generate the perturbed cross section libraries used in this study. SCALE is a modeling and simulation code suite that, in addition to a large number of criticality safety and reactor physics tools, has uncertainty and sensitivity analysis capabilities. The Sampler tool in SCALE is one of these uncertainty analysis tools. Sampler allows for a streamlined approach to stochastic-sampling-based uncertainty quantification, generating perturbed input parameters and analyzing responses from code output. In this study, Sampler was only used to generate the perturbed cross section libraries. While SCALE has the capability to sample input parameters and analyze responses on-the-fly, significantly less computational overhead is incurred if Sampler is used to simply generate perturbed cross section data files and use the perturbed data case-by-case.

That is the approach used in this study – perturbed cross section libraries were generated from Sampler, and these cross section files were converted to MPACT format and each used as cross section data for the benchmark problems. The perturbed cross section libraries were generated at Oak Ridge National Laboratory.

The Medusa module in XSUSA [Wil13] was used to generate multiplicative perturbation factors for microscopic cross sections. Cross section data from the ENDF/B VII.1 nuclear data library are included in SCALE. The covariance data is stored in the form of relative values of the infinitely dilute, or unshielded, cross sections. Shielding is an important phenomenon in which large resonance peaks for certain energy levels will greatly increase the probability of interaction for one type of reaction at that energy level, decreasing the neutron flux at that energy level. Shielding, or self-shielding, must be considered in both energy and space, or simulated reaction rates will not be reflective of reality. Unfortunately, self-shielding is highly dependent on the local geometry and the nuclide mixture, so it is problem-dependent. For this reason, the effects of self-shielding are removed from cross section data, thereby creating unshielded cross sections which are problem-independent. This means that the cross sections generated by SCALE can be used for any application, but multi-group cross sections must be independently processed for each problem. The independent cross section generation is handled by MPACT, and will be discussed in the following section.

Because cross section covariances in SCALE are relative values, a random sample in XSUSA for a microscopic cross section $\sigma_{x,g}$ for reaction type x in energy group g takes the form $\Delta\sigma_{x,g}/\sigma_{x,g}$, and is converted to a multiplicative perturbation factor $Q_{x,g}$,

$$Q_{x,g} = 1 + \frac{\Delta\sigma_{x,g}}{\sigma_{x,g}}. \tag{2.1}$$

The perturbed cross section $\sigma'_{x,g}$ for reaction type x in energy group g is then calculated using the associated perturbation factor $Q_{x,g}$,

$$\sigma'_{x,g} = Q_{x,g}\sigma_{x,g}. \tag{2.2}$$

After all cross sections have been perturbed according to Equation 2.2, the perturbed cross sections are stored in a perturbed cross section library. 1000 uniquely perturbed cross section libraries were used in this study.

2.2 MPACT

The neutronics solver in VERA is the Michigan PArallel Characteristics based Transport (MPACT) code [Koc17]. MPACT is a steady-state eigenvalue problem solver that uses the Method of Characteristics (MOC) [Col16]. In Method of Characteristics solvers, the steady-state neutron transport equation is solved along characteristic direction imposed onto the problem domain, as shown in Figure 2.1 for a pincell. The characteristic form of the neutron transport equation is shown in Equation 2.3,

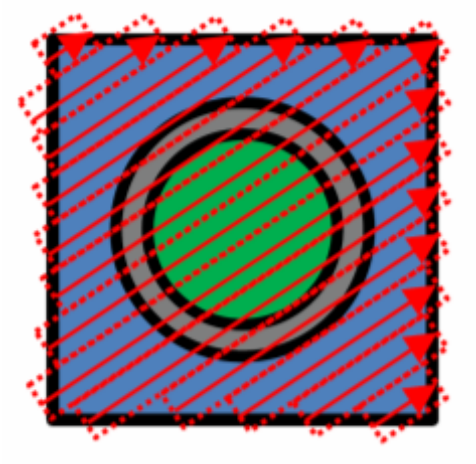

Figure 2.1 Characteristics imposed on a pincell [Col16].

$$\frac{d\psi}{ds}\left(\vec{r}_0+s\hat{\Omega},\hat{\Omega},E\right)+\Sigma\left(\vec{r}_0+s\hat{\Omega},E\right)\psi\left(\vec{r}_0+s\hat{\Omega},\hat{\Omega},E\right)= \\ \frac{\chi\left(\vec{r}_0+s\hat{\Omega},E\right)}{k}\int_0^\infty dE'\int_0^{4\pi}d\Omega'\;\nu\Sigma_f\left(\vec{r}_0+s\hat{\Omega}',E'\right)\psi\left(\vec{r}_0+s\hat{\Omega}',\hat{\Omega}',E'\right)+ \\ \int_0^\infty dE'\int_0^{4\pi}d\Omega'\,\Sigma_s\left(\vec{r}_0+s\hat{\Omega}',\hat{\Omega}'\cdot\hat{\Omega},E'\rightarrow E\right)\psi\left(\vec{r}_0+s\hat{\Omega}',\hat{\Omega}',E'\right), \quad (2.3)$$

where s is the distance along the characteristic direction $\hat{\Omega}$. In MPACT, rather than purely solving the 3-dimensional neutron transport equation in the entire discretized domain, which would be very computationally expensive, the solution is approximated by a "2D/1D" method. In the 2D/1D method, the transport equation is solved radially (in two dimensions) using the MOC method, and axial transport is instead approximated by one-dimensional transport solvers. The 2-dimensional $\vec{r}=(x,y)$ radial transport equation is modified to include an axial z leakage term, as shown in Equation 2.4 for one region of the domain with constant material properties.

$$\Omega_x\frac{\partial\psi}{\partial x}\left(\vec{r},\hat{\Omega}\right)+\Omega_y\frac{\partial\psi}{\partial y}\left(\vec{r},\hat{\Omega}\right)+\Sigma\psi\left(\vec{r},\hat{\Omega}\right)= \\ \Sigma_s\int_{4\pi}d\Omega'\,\psi\left(\vec{r},\hat{\Omega}'\right)+\frac{\nu\Sigma_f}{k}\int_{4\pi}d\Omega'\,\psi\left(\vec{r},\hat{\Omega}'\right)-\left[\frac{\partial J_z}{\partial z}\left(\vec{r}\right)\right] \quad (2.4)$$

In MPACT, either a P_1 or P_3 Legendre polynomial (one-dimensional spherical harmonics) approximation is used for the axial current J_z. For the simpler P_1 approximation, the axial current J_z is obtained by solving for the scalar flux ϕ in Equation 2.5 with appropriate boundary conditions at the top and bottom of the axial mesh interval, and using the scalar flux in the definition of the axial

current in Equation 2.6,

$$-\frac{\partial}{\partial z}\frac{1}{3\Sigma}\frac{\partial \phi}{\partial z}(\vec{r})+\sigma_a\phi(\vec{r})=\frac{\nu\Sigma_f}{k}\phi(\vec{r})-\left[\frac{\partial J_x}{\partial x}(\vec{r})+\frac{\partial J_y}{\partial y}(\vec{r})\right], \tag{2.5}$$

$$J_z(\vec{r})=-\frac{1}{3\Sigma}\frac{\partial \phi}{\partial z}(\vec{r}), \tag{2.6}$$

with the radial currents J_x and J_y defined by Equations 2.7a and 2.7b, respectively.

$$J_x(\vec{r})=\int_{4\pi}\Omega_x\psi\left(\vec{r},\hat{\Omega}\right)d\Omega \tag{2.7a}$$

$$J_y(\vec{r})=\int_{4\pi}\Omega_y\psi\left(\vec{r},\hat{\Omega}\right)d\Omega \tag{2.7b}$$

After the 2D/1D equations are solved over the discretized domain, a total effective neutron multiplication factor k-eff can be calculated for the entire domain. As mentioned in Chapter 1, this effective neutron multiplication factor serves as the response for our stochastic sampling-based uncertainty quantification.

2.2.1 Resonance Self-Shielding

As mentioned Chapter 1, the nuclear data in SCALE is problem-independent, and does not account for the effects of self-shielding in the resonance regions of the cross sections. Self-shielding is handled in MPACT by the subgroup method, in which pre-computed resonance integral tables are converted to a quadrature set and effective cross sections are computed. This effective cross section calculation is performed using Equation 2.8, where the unshielded cross section $\sigma_{x,g,n}$ of reaction type x in energy group g in resonance-table derived subgroup n is converted to an effective group cross section. $\sigma_{a,g,n}$ is the unshielded microscopic absorption cross section for energy group g and subgroup n and $\sigma_{b,g,n}$ is the unshielded background cross section for energy group g and subgroup n.

$$\sigma_{x,g}\approx\frac{\sum_n\sigma_{x,g,n}\frac{\sigma_{b,g,n}}{\sigma_{a,g,n}+\sigma_{b,g,n}}w_{x,g,n}}{\sum_n\frac{\sigma_{b,g,n}}{\sigma_{a,g,n}\sigma_{b,g,n}}w_{x,g,n}} \tag{2.8}$$

More details on the subgroup treatment of cross section resonances in multigroup calculations in MPACT can be found in the MPACT theory manual [Koc17].

2.3 Computational Cost

All cases involving the UAM LWR benchmark problems were run on the NC State High Performance Computing cluster. The Watts Bar cases, which are much more computationally expensive, were

Table 2.1 Profiling data for UAM LWR benchmark cases.

Case	Single Case Time	Est. 1000 Case Time	Single Case Memory Usage [MB]
PPZ	3.05 s	50 m, 50 s	107.6
PPF	3.11 s	51 m, 50 s	107.6
PLUZ	22.04 s	6 h, 07m, 20 s	214.2
PLUF	19.80 s	5 h, 30 m, 00 s	214.4
PLRZ	26.28 s	7 h, 18 m, 00 s	213.8
PLRF	22.71 s	6 h, 18 m, 30 s	215.7
PCZ	3 m, 04.87 s	51 h, 21 m, 10 s	640.7
PCF	3 m, 10.17 s	52 h, 49 m, 30 s	640.7
BPZ	3.02 s	50 m, 20 s	109.2
BPF	2.46 s	41 m, 00 s	111.2
BLUZ	50.28 s	13 h, 58 m, 00 s	261.1
BLUF	53.78 s	14 h, 56 m, 20 s	259.0
BLRZ	1 m, 16.47 s	21 h, 14 m, 30 s	266.6
BLRF	57.77 s	16 h, 02m, 50 s	266.4

run at Idaho National Labs. Since computational cost is one of the key limiting factors with the implementation of stochastic sampling, it is beneficial to understand the computational power that was required to generate the data in this study. Table 2.1 shows the computation times for each UAM LWR benchmark problem.

Each case was run on one core so that the times could be compared between cases, and the wall time in Table 2.1 is approximately equivalent to the CPU time. As expected, as the complexity and size of the cases increase, the computational cost increases. This reinforces the desire to use simpler cases to estimate uncertainty in larger problems. For instance, if we wish to quantify the uncertainty in k-eff due to cross section uncertainty for a PWR, it would obviously be desirable to use pincell cases, which take approximately three seconds to run, instead of larger colorset problems that take three minutes to run. Even if the colorset problem is more similar in size to the real configuration of the problem, if there is little difference between the uncertainty from the pincell, lattice, and colorset cases, it would be much more efficient to use the pincell cases to generate data.

In addition to the UAM LWR benchmark cases, larger cases based on the CASL Watts Bar Unit 1 benchmark problems were examined. These larger problems include a 2-dimensional full-core, a 3-dimensional full-core, a full-core with thermal hydraulic feedback, and a full-core depletion study. These problems are much larger and more complex than the UAM LWR benchmark problems, and therefore required much more computational power. For instance, the full-core Watts Bar depletion study, which was only run with 29 perturbed cross section libraries, required 565,152 core-hours on the Lemhi supercomputer at Idaho National Laboratory. Some profiling data for the Watts Bar Unit 1 full-core cases are shown in Table 2.2.

Table 2.2 Profiling data for Watts Bar Unit 1 full-core cases.

Case	XS Libraries Used	Total Core-Hours
2D	1000	2,667
3D	1000	342,000
3D + Feedback	1000	552,462
3D + Feedback + Depletion	29	565,152

The profiling statistics associated with these high-fidelity simulations further emphasizes the large computational cost of stochastic simulation as an uncertainty quantification technique. However, even with the high computational cost, the fact that these studies could be performed at all lends credence to the idea that the rapid increase in raw computational power over the past decades makes stochastic sampling perhaps the best technique for uncertainty quantification in nuclear physics simulation.

2.4 Material Perturbation Mechanism

To aid in the material uncertainty analysis, a tool was developed to easily perform pin-by-pin material perturbations. The fact that MPACT, and the VERA environment as a whole, allows for each pin in an assembly to be individually defined made a material uncertainty analysis to be performed without invasiveness. In MPACT, the geometry of each individual fuel pin is defined radially by specifying the outer radius of the fuel pellet, the inner radius of the cladding, and the outer radius of the cladding. With these three parameters specified, three regions of the fuel pin are defined: the center of the fuel pellet to the edge of the fuel pellet, the edge of the fuel pellet to the inner surface of the cladding, and the inner surface of the cladding to the outer surface of the cladding. The dimensions of these three regions can be described by the radius of the fuel pellet, the gap width, and the cladding thickness respectively.

As will be discussed in Chapter 6, we are assuming that there is uncertainty in each of the fuel pellet radius, gap width, and cladding thickness. Since MPACT allows for the specification of each of these parameters for each fuel pin in an assembly, one could simply sample out of the respective uncertainty distributions for each pin, combine these pin realizations into an assembly, and run the realized assembly in MPACT to observe the uncertainty due to the fuel pin geometric parameters.

In addition to the fuel pin geometry, the UAM benchmark also provides uncertainty information about the fuel enrichment and the fuel pellet density. Both of these parameters can also easily be specified in MPACT using an individual fuel definition for each realization.

In a typical MPACT input deck, each pin type is only defined once, and then is repeated in an assembly map. This assumes that each fuel pin of a certain type is identical to every fuel pin of that type. With manufacturing uncertainty, each pin is not identical to each pin of the same type.

Every pin is uniquely generated by sampling out of the material uncertainty distributions. This complicates the construction of an input deck for material uncertainty analysis. Instead of, for instance, five pin types being defined in a lattice, all pins (typically more than two-hundred per lattice for PWRs) would need unique geometric and compositional definitions.

A script was developed to modify a base input deck to account for material parameter uncertainty. The base input deck is identical to any other VERA input deck, but with an additional material uncertainty block to define how the material parameters should be perturbed. The current implementation assumes the parameters in the normal input deck are the mean values, and the uncertainty data is included in the additional uncertainty block. This should theoretically make uncertainty analysis much easier on commonly examined problems, as the input decks are already completed, and the only modification a user would need to make is to add the material uncertainty data in the material uncertainty block.

Within the material uncertainty block, which can be placed at any location in the base input deck, the material uncertainty data is defined. Currently, it is assumed that all material parameters have normal distributions, but that could easily be extended to other distributions. The only issue with the implementation of more exotic probability distributions is that they are slightly more difficult and less intuitive to define with a small number of quantities. The normal distribution is quite simple and intuitive – the entire distribution can be defined by the mean and standard deviation. If another probability distribution was nonetheless desired, only the material uncertainty block interpreter and sampling kernel would need to be modified. Since the benchmark data used in the material uncertainty quantification in this study was a normal distribution, that was the only probability distribution included.

Assuming the material parameters are normally distributed and the mean values are included in the non-material-uncertainty blocks, only the standard deviations for each material parameter need to be defined in the material uncertainty block to fully define the probability distributions. The five material parameters that were considered in this study are the pellet radius, gap width, clad thickness, fuel density, and fuel enrichment. More information on these quantities is included with the results of the study, which are reported in Chapter 6. The five material parameters are defined by keywords in the input. These keywords are listed below.

- pellet radius – `pellet_radius`
- gap width (gap thickness) – `gap_thickness`
- clad thickness – `clad_thickness`
- fuel density – `fuel_density`
- fuel enrichment – `fuel_enrichment`

The keywords can be included in the material uncertainty block in any order. Not all keywords need to be defined. If a keyword is not present in the material uncertainty block, the parameter

will default to the mean value. As mentioned previously, the mean values for these five parameters are obtained from the non-material-parameter blocks of the base input deck. Now that the five parameters have been defined, it should be noted specifically where the mean values are defined.

Since all five parameters are focused on the fuel rods, the mean values are defined near each other. They can be split into two groups: fuel composition parameters (fuel density and fuel enrichment) and geometric parameters (fuel pellet radius, gap width, and clad thickness). All of these parameters are defined in the assembly block in a VERA input deck. The fuel parameters are defined with a `fuel` keyword, and the geometric parameters are defined with a `cell` keyword in this block. For example, to define a fuel with a density of 10.283 g/cm^3 and an enrichment of 4.85 w/o, one would include the following line in the assembly block:

```
fuel f1 10.283 95.0 / 4.85
```

In above line, the fuel is given the name `f1`, the first quantity after the fuel name is the fuel stack density in units of g/cm^3 (10.283), the second quantity is the theoretical pellet density in % used in fuel performance calculations (95.0), and the quantity after the forward slash is the fuel enrichment in weight percent (4.85). In any fuel definition encountered by the material perturbation script, the quantity after the fuel name is stored as the mean fuel density for fuel `f1`, and the quantity after the forward slash is stored as the mean fuel enrichment for fuel `f1`. The gadolinia concentration, if there is any, can also be included after the fuel enrichment, and is read into the script, but is not perturbed.

The `cell` keyword then defines the radial geometry of the pin. For a fuel pin composed of the already-defined fuel `f1`, with a fuel pellet radius of 0.4695 cm, a cladding inner radius of 0.4788 cm, and a cladding outer radius of 0.5461 cm, the cell could be defined as follows:

```
cell c1 0.4695 0.4788 0.5461 / f1 he zirc4
```

The cell is given the name `c1`, and the radii of the fuel pin components are defined between the cell name and forward slash. The terms after the forward slash define the materials of the fuel pin, and correspond to the dimensions listed before the forward slash. In this case, the fuel pin is composed of fuel f1 from the center of the pin to 0.4695 cm from the center, helium from the edge of the fuel pellet to the inner surface of the cladding, and zirc4 in the cladding. Note that the geometry of the pin is defined in a radial fashion, but the uncertainty in the gap width and cladding thickness are defined in absolute terms. The cladding thickness is assumed to be the difference between the last quantity before the forward slash and the second-to-last quantity, and the gap thickness is the difference between the second-to-last quantity and the third-to-last quantity.

In any realistic fuel lattice, there will be multiple fuel types, particularly with regard to enrichment. In the UAM BWR lattice benchmark, for instance, there are six unique rod types with varying levels of enrichment and fuel density. Therefore the capability to handle multiple fuel definitions was included. The way this is handled in the script is to default to a perturbation for all fuel rods, but allow the user to specify an uncertainty for a certain fuel or cell type. To specify an uncertainty for a

certain fuel type, the user would type the keyword `fuel` and the desired fuel name after the keyword and standard deviation.

For example, say we wish to model a lattice with two fuel types, which we have named `f1` and `f2` in the fuel definitions in the assembly block of the input deck. The fuel enrichment for fuel type `f1` has a lower standard deviation (0.3 w/o) than fuel `f2` (0.6 w/o). Each of the uncertainties for the fuel types must be defined individually. The correct input in the material uncertainty block for this example would be the following:

```
fuel_enrichment 0.3 fuel f1
fuel_enrichment 0.6 fuel f2
```

If one fuel type is defined with a distinct uncertainty, all fuel types with uncertainty must be defined individually. The same syntax is used when describing uncertainty in the geometric parameters, but the `cell` keyword and cell name are used instead of the `fuel` keyword and fuel name. It is not valid to apply a fuel type parameter to a cell. `fuel_density` and `fuel_enrichment` must be applied to a fuel name, while `pellet_radius`, `gap_thickness`, and `clad_thickness` must be applied to a cell name.

The other main component of the material uncertainty block in the base input deck is on which level the user would like to perturb the parameters. There are currently two perturbation modes: pin-level and assembly-level. These two modes are triggered by the keywords `pin_level` and `assembly_level`, with each keyword followed by a list of the associated material parameters.

A pin-level perturbation means that each pin in the input deck is realized independently by sampling out of the distribution defined by the pin's mean value and uncertainty. This would be a good choice for a parameter that is associated with independent manufacturing processes. For instance, the fuel cladding is formed for each fuel pin independently of the other fuel pins, so each fuel pin will have a unique cladding thickness, within the manufacturing tolerance. In five assemblies with two-hundred pins each, every pin-level parameter would be realized 1000 times.

An assembly-level perturbation, on the other hand, means that each pin in an assembly is identical with regard to a certain parameter. This would be a good choice for a parameter that is associated with a manufacturing process in which materials are produced simultaneously. Since fuel pellets are manufactured in batches, the fuel pellet composition parameters of fuel density and fuel enrichment would likely not vary within a batch, but may vary when comparing different batches. In five assemblies with two-hundred pins each, every assembly-level parameter would only be realized five times, with all two-hundred pins in each assembly sharing the parameter realization for that assembly.

As will be discussed in Chapter 6, the most realistic material parameter configuration is to vary the fuel density and fuel enrichment on an assembly-level, and to vary the pellet radius, gap thickness, and cladding thickness on a pin-level. To specify the material parameter perturbations taking place on the pin-level and those taking place on the assembly-level. the additional lines that would need to be added to the material uncertainty block in the base input deck are:

```
pin_level pellet_radius gap_thickness clad_thickness
assembly_level fuel_density fuel_enrichment
```

The order of the parameters following each of the keywords is not important. If these keywords are not included, all parameters default to a pin-level perturbation.

The ability to simultaneously use perturbed cross section libraries was added to the script. The user simply needs to specify the path to one cross section library in its normal position in the base input deck, and the script will search within that directory for other files with the cross section library extension `.fmt`. The script then randomly chooses cross section libraries to use in each realization. If it is not desired to use perturbed cross section libraries, the script will default to the cross section library included in the base input deck.

Once the base input deck is completed, which should simply entail the addition of the material uncertainty block in the input deck, the user specifies the number of realizations to produce, and the desired destination for these files. Each realization takes the form of one input deck. Further automation was developed to submit these individual input decks as a batch to the HPC for execution, but obviously this would need to be modified to submit jobs to other clusters.

```
[CASEID]
    title 'TMI1PWR Lattice Benchmark HZP Unrodded with Library 0'

[STATE]
    power    0.0
    tinlet   551 K
    tfuel    551 K
    modden   0.766
    boron    0
    sym      qtr
    feedback off

[CORE]
    size   1
    apitch 21.8110
    height 1.0
    rated  0.04832 0.01

    core_shape
        1

    assm_map
        ASSY

    bc_rad reflecting
    bc_top reflecting
    bc_bot reflecting

[ASSEMBLY]
    npin 15
    ppitch 1.4427

    fuel U31 10.283 95.0 / 4.85

    fuel Ugd 10.144 95.0 / 4.12 / gad=2.0

    cell 1 0.4695 0.4788 0.5461 / U31 he zirc4
    cell 2 0.4695 0.4788 0.5461 / Ugd he zirc4
    cell 3        0.63245 0.6731 /    mod zirc4
    cell 4        0.56005 0.6261 /    mod zirc4
```

```
    lattice LAT
        1 1 1 1 1 1 1 1 1 1 1 1 1 1 1
        1 2 1 1 1 1 1 1 1 1 1 1 1 2 1
        1 1 1 1 1 3 1 1 1 3 1 1 1 1 1
        1 1 1 3 1 1 1 1 1 1 1 3 1 1 1
        1 1 1 1 1 1 1 1 1 1 1 1 1 1 1
        1 1 3 1 1 3 1 1 1 3 1 1 3 1 1
        1 1 1 1 1 1 1 1 1 1 1 1 1 1 1
        1 1 1 1 1 1 1 4 1 1 1 1 1 1 1
        1 1 1 1 1 1 1 1 1 1 1 1 1 1 1
        1 1 3 1 1 3 1 1 1 3 1 1 3 1 1
        1 1 1 1 1 1 1 1 1 1 1 1 1 1 1
        1 1 1 3 1 1 1 1 1 1 1 3 1 1 1
        1 1 1 1 1 3 1 1 1 3 1 1 1 1 1
        1 2 1 1 1 1 1 1 1 1 1 1 1 2 1
        1 1 1 1 1 1 1 1 1 1 1 1 1 1 1

    axial ASSY 0.0 LAT 1.0

[MPACT]
    k_shift 2.0
    res_up_scatter  false
    xs_filename  /pathtofile/xsfile0000.fmt

[SHIFT]
    num_cycles           1100
    num_inactive_cycles  100
    Np                   1000000

[MAT_UNC]
    fuel_density 0.0032111
    pellet_radius 4.6944e-8
    gap_thickness 6.4e-7
    clad_thickness 6.9444e-7
    fuel_enrichment 5.5751e-7

    pin_level gap_thickness clad_thickness pellet_radius
    assm_level fuel_density fuel_enrichment
```

Figure 2.2 Sample input deck for PWR Lattice with material parameter uncertainty.

CHAPTER

3

BENCHMARK CROSS SECTION UNCERTAINTY

3.1 Benchmark Configuration

The 1000 perturbed cross section libraries have been used to calculate the uncertainty in k-eff for the problems specified in the Uncertainty Analysis in Modeling (UAM) Volume I Benchmark [Iva13] for the Pressurized Water Reactor (PWR) and Boiling Water Reactor (BWR). The PWR is based on the Three Mile Island Unit 1 (TMI-1) design, and the BWR is based on Peach Bottom Unit 2 (PB-2). Generally, the progression of the individual benchmark tests increases in size and complexity, from a single pincell to a two-dimensional lattice to a colorset of two-dimensional lattices. The details of the specific test cases are provided in the following sections. An abbreviated naming convention, shown in Table 1.2, is used throughout this chapter.

For all models in this chapter, cases were run at both Hot Zero Power (HZP) and Hot Full Power (HFP). Hot Zero Power conditions are reached when the core is heated to operating temperature, but no thermal power is produced. Hot Full Power is representative of the typical operating conditions in the reactor. The state parameters associated with HZP and HFP conditions for the PWR and BWR models are shown in Table 3.1.

Table 3.1 State parameters for the UAM LWRs.

	Hot Zero Power		Hot Full Power	
Parameter	PWR	BWR	PWR	BWR
Fuel temp.	551 K	552.833 K	900 K	900 K
Moderator temp.	551 K	552.833 K	562 K	557 K
Moderator density	766 kg/m^3	753.978 kg/m^3	748.4 kg/m^3	460.72 kg/m^3
Void fraction	-	0 %	-	40 %

3.1.1 PWR Pincell

The pincell model is commonly used in benchmarking due to its small size. A two-dimensional pincell model consists of a single fuel pin surrounded by moderator, as shown in Figure 3.1. When an individual pincell is modeled with reflective boundary conditions, the neutron multiplication factor (k-eff) of the pincell is the same as for an infinite array of pincells. This is often referred to as k-inf.

An infinite array of pincells in some cases is a reasonable approximation to an actual core. The TMI-1 reactor consists of 177 fuel assemblies [Wim01], each with 225 rod locations. Fuel rods are not placed in every rod location, but there are still more than 35,000 fuel rods in the core. The k-eff of a pincell with reflective boundary conditions should be an overestimation of the k-eff of a full-core due to the lack of neutron leakage. Key parameters for the configuration of the PWR pincell test problem, based on the TMI-1 design, are reproduced in Table 3.2.

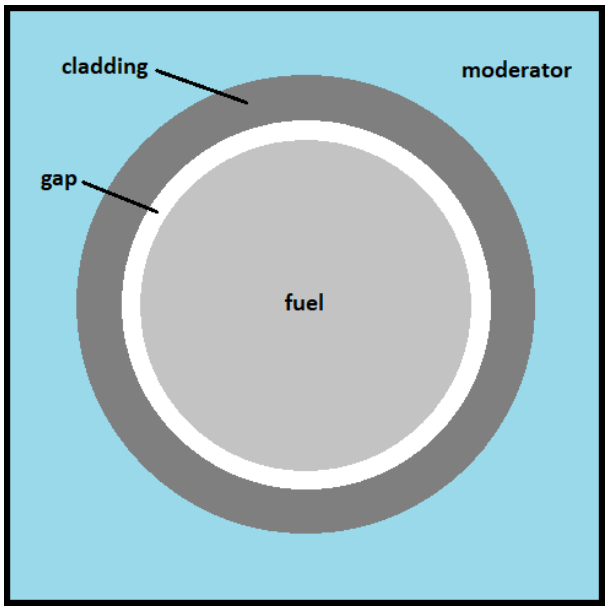

Figure 3.1 Two-dimensional pincell geometry. The pitch in Table 3.2 corresponds to the side length of the black boundary.

Table 3.2 PWR pincell parameters [Iva13].

Parameter	Value
Pitch	1.4427 cm
Pellet radius	0.46955 cm
Pellet density	10.283 g/cm^3
Fuel enrichment	4.85 w/o
Gap thickness	0.00955 cm
Cladding thickness	0.0673 cm
Pellet material	UO_2
Gap material	He
Cladding material	Zirc-4
Moderator material	H_2O

3.1.2 PWR Lattice

Parameters for the configuration of the two-dimensional PWR lattice (assembly) problem were provided in the benchmark document [Iva13] and are reproduced in Table 3.3. A diagram of the two-dimensional lattice is shown in Figure 3.2. The dimensions of the fuel pins are the same as for the pincell, so they are omitted from this section. The PWR lattice was also run at HZP and HFP. In addition, the control rods were either fully inserted (ARI) or fully withdrawn (ARO). This specifies four unique problems for the two-dimensional PWR lattice: HZP ARO, HFP ARO, HZP ARI, and HFP ARI. Information on the control rods for the PWR lattice is also included in Table 3.3. The state information for the HZP and HFP cases were the same as the PWR pincell case.

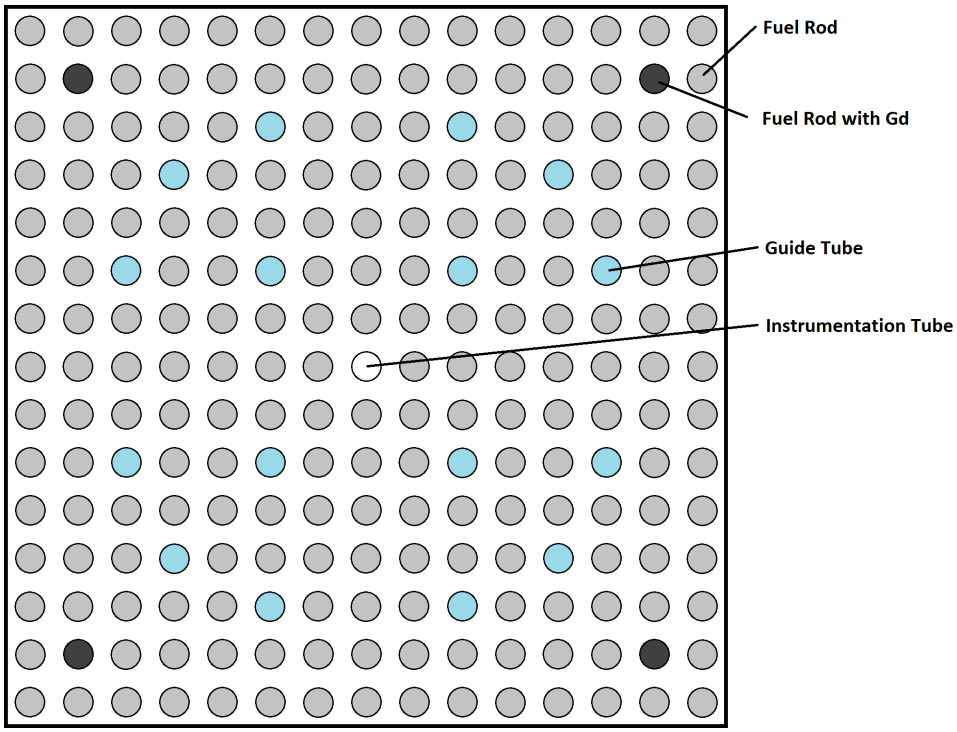

Figure 3.2 PWR lattice configuration.

3.1.3 PWR Colorset

The configuration for the two-dimensional PWR colorset problem is shown in Figure 3.3. The colorset consists of a three-by-three combination of the ARO and ARI assemblies described in the PWR Lattice section – with the ARI assembly surrounded by ARO assemblies. This configuration was run at HZP and HFP.

Table 3.3 PWR lattice parameters [Iva13].

Parameter	Value
Assembly pitch	21.811 cm
Dimension	15 x 15
Fuel rods	208
Guide tubes	16
Instrumentation tubes	1
Gadolinia pins	4
Guide tube inner radius	0.63245 cm
Guide tube thickness	0.04065 cm
Guide tube material	Zirc-4
Inst. tube inner radius	0.56005 cm
Inst. tube thickness	0.06605 cm
Inst. tube material	Zirc-4
Control rods	16
Control rod inner radius	0.50292 cm
Control rod absorber material	Ag-In-Cd
Control rod cladding thickness	0.05715 cm
Control rod cladding material	Inconel 625

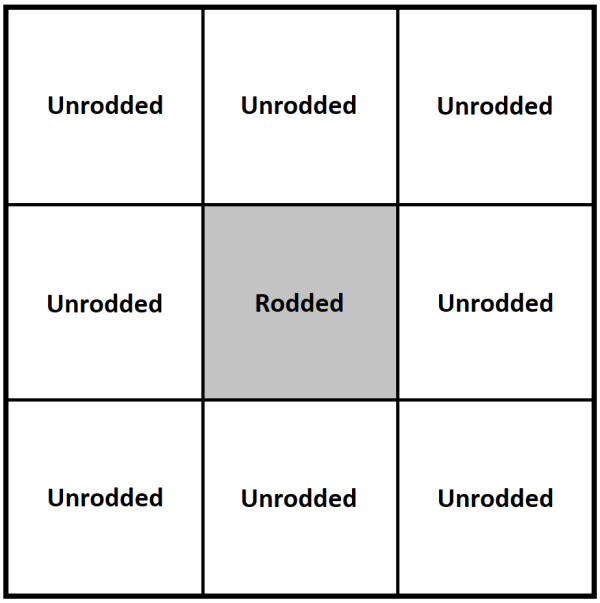

Figure 3.3 PWR colorset configuration. Each square corresponds to one assembly.

3.1.4 BWR Pincell

The parameters of the BWR pincell model are reproduced in Table 3.4, which comes from the UAM LWR Benchmark specification [Iva13]. The geometry of the BWR pincell is the same as the PWR pincell shown in Figure 3.1, with slightly different dimensions. Like the PWR test problems, the BWR problems were examined at two states: HZP and HFP, with state parameters given in Table 3.1.

Table 3.4 BWR pincell parameters [Iva13].

Parameter	Value
Pitch	1.875 cm
Pellet radius	0.60579 cm
Pellet density	10.42 g/cm^3
Fuel enrichment	2.93 w/o
Gap thickness	0.01524 cm
Cladding thickness	0.09398 cm
Pellet material	UO_2
Gap material	He
Cladding material	Zirc-2
Moderator material	H_2O

3.1.5 BWR Lattice

For the BWR two-dimensional lattice problems, the fuel-pin-level parameters and state parameters are maintained from the BWR pincell case shown in Table 3.4. The assembly geometry parameters needed to model the problem in MPACT are given in Table 3.6. The wide gap and narrow gap lengths are defined as the distance between the exterior of the channel box and the midpoint of the channel. This is equivalent to half of the distance between exterior channel boxes of neighboring assemblies in locations with control blades (wide gap) and without control blades (narrow gap). The assembly is a 7 x 7 geometry with six fuel types, as shown in Figure 3.4. Specific information on the different rod types is reproduced from the report in Table 3.5.

The 'A' and 'B' in rod types 5A and 6B correspond to the different axial placement of Gadolinium in the fuel rods. Type 'A' rods are completely composed of $UO_2 + Gd_2O_3$ pellets, while type 'B' rods have both regular UO_2 pellets and $UO_2 + Gd_2O_3$ in certain axial levels. The difference in the 'A' and 'B' rod types is illustrated in Figure 3.5. Since we are only considering the two-dimensional lattice, it is sufficient to use the stack density listed in Table 3.5 in conjunction with the Gd_2O_3 concentration to model the different rod types, homogenizing the rods axially. Figure 3.5 shows the axial BWR diagram.

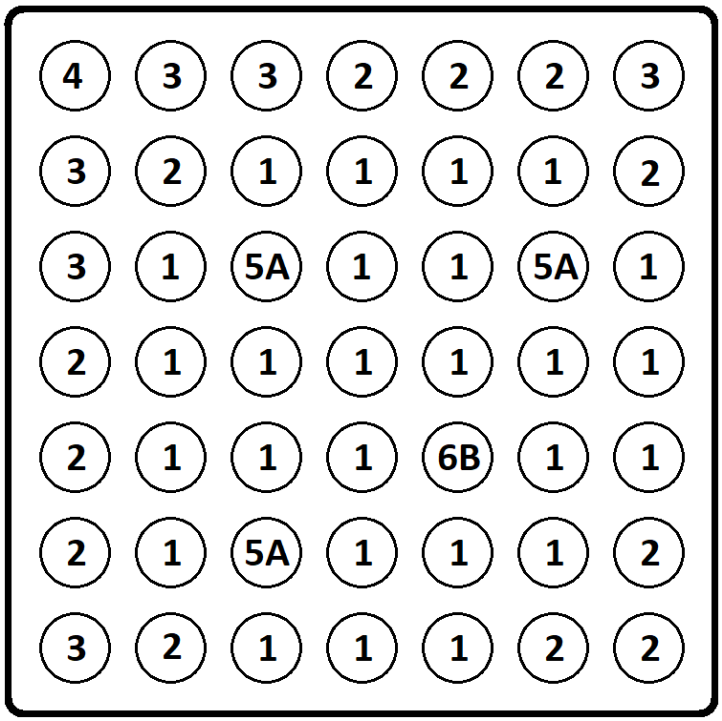

Figure 3.4 BWR lattice configuration with rod type labels.

Table 3.5 BWR rod types [Iva13].

Rod Type	Stack Density [g/cm^3]	Enrichment [w/o]	Gd_2O_3 [w/o]
1	10.32	2.93	0
2	10.32	1.94	0
3	10.32	1.69	0
4	10.32	1.33	0
5A	10.19	2.93	3.0
6B	10.27	2.93	3.0

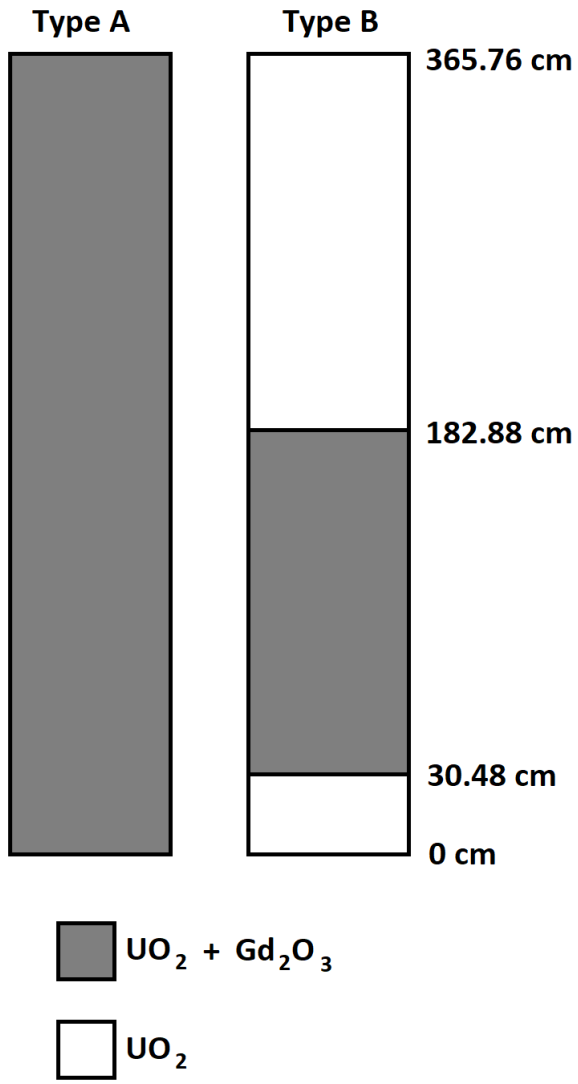

Figure 3.5 BWR axial configuration.

Table 3.6 BWR lattice parameters [Iva13].

Parameter	Label (Figure 3.6)	Value
Assembly pitch	A	15.24 cm
Wide gap length	B	0.9525 cm
Narrow gap length	C	0.47498 cm
Channel box thickness	D	0.2032 cm
Channel box inside corner radius	E	0.9652 cm

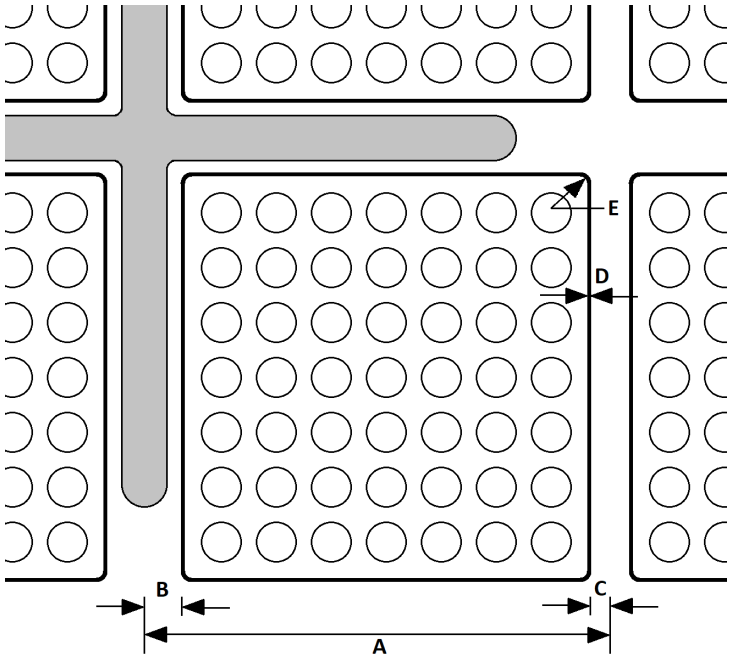

Figure 3.6 BWR assembly diagram. Labels correspond to Table 3.6.

Table 3.6 lists the lattice parameters for the BWR. A diagram showing the quantities in Table 3.6 is shown in Figure 3.6. The channel box, which forms a solid barrier around the fuel lattice, is made of Zirc-4. The central assembly in Figure 3.6 is the single lattice modeled in this study, with the wide-wide gap corresponding to the northwest corner.

Parameters that specify the BWR lattice control blades are provided in Table 3.7 and correspond to the labels in Figure 3.7. Control blade span is defined as the length of the control blade from the center of the blade structure to the tip of one wing. The control blade thickness is the width of one wing of the control blade – perpendicular to the span when viewing the blade from the top-down. The central structure wing length is the distance between the center of the entire control blade and the nearest point on the exterior of the nearest absorber tube. This is also referred to as the control blade dead length. Note that for the single assembly case, the control blade is placed outside of the northwest corner of the assembly – in the corner with two wide gaps. Due to limitations in

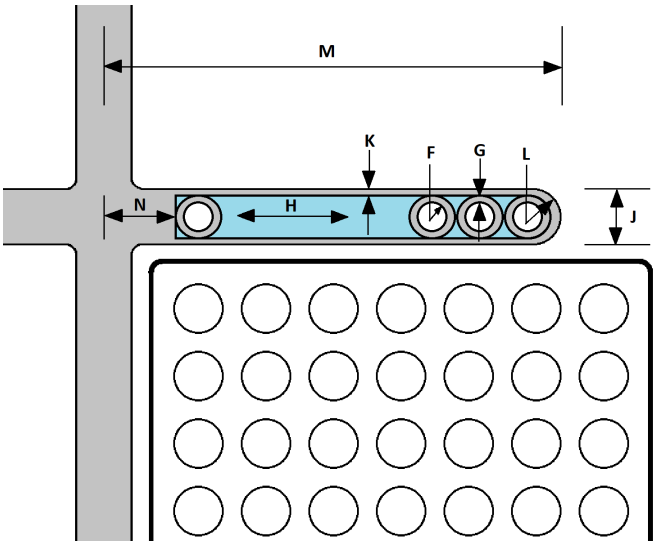

Figure 3.7 BWR control blade structure. Labels correspond to Table 3.7.

the placement of the control blades in MPACT, the colorset cases could not be run for the BWR. The BWR colorset is a 4 x 4 assembly with control blades on the northwest corner of the northwest assembly and southeast corner of the southeast assembly. The absorber material is B_4C, and all other structural material in the control blade is Type-304 stainless steel.

Table 3.7 BWR control blade parameters [Iva13].

Parameter	Label (Figure 3.7)	Value
Absorber tube inner radius	F	0.17526 cm
Absorber tube sheath thickness	G	0.0635 cm
Absorber tubes per wing	H	21
Control blade thickness	J	0.7925 cm
Control blade sheath thickness	K	0.14224 cm
Control blade tip radius	L	0.39624 cm
Control blade span	M	12.3825 cm
Central structure wing length	N	1.98501 cm

3.2 Benchmark Results

A complete list of the resulting estimated mean k-eff (k-eff), relative standard deviation (RSD), and standard deviation (SD) for each of the fourteen test problems is shown in Table 3.8. Note that the abbreviated case names are used. Refer to Table 1.2 for the case naming convention.

Table 3.8 Mean and standard deviation of k-eff.

Case	k-eff	RSD [%]	SD [pcm]
PPZ	1.427559	0.513	732
PPF	1.409660	0.517	729
PLUZ	1.414477	0.502	710
PLUF	1.398791	0.505	706
PLRZ	1.077715	0.529	570
PLRF	1.062718	0.534	567
PCZ	1.384679	0.500	693
PCF	1.368892	0.504	690
BPZ	1.3444756	0.553	744
BPF	1.229762	0.607	747
BLUZ	1.107253	0.559	619
BLUF	1.078942	0.576	622
BLRZ	0.863599	0.550	475
BLRF	0.794981	0.591	470

The relative standard deviation for all of the benchmark cases falls within the range 0.500 % to 0.607 %. The two extreme values of the relative standard deviation come from arguably the most disparate cases – the largest PWR problem at Hot Zero Power has the smallest relative standard deviation, and the smallest BWR problem at Hot Full Power has the largest relative standard deviation. Still, the difference in relative standard deviation of 0.107 % is small when considering the fact that the uncertainty was calculated using two very different problems. A difference in relative standard deviation of 0.107 % corresponds to a difference of approximately 140 percent mille (pcm) when adjusting by the average k-eff of the two cases. Percent mille is equivalent to $10^{-}5\Delta k$.

This difference is reduced when considering two problems that are similar in their configuration. For instance, consider the PWR Pincell at HFP (case PPF) and the PWR colorset at HFP (case PCF). Before comparing the cases, it can be observed that these two cases are similar in nature. They are both derived from the same PWR, so that the colorset case is essentially composed of a series of pincell cases, and they have the same state parameters. The main difference between the two cases is their complexity, or the size of the problem. As a result, each calculation of k-eff, which is performed one thousand times, takes less than one minute for the pincell case but takes approximately five minutes for the colorset case. The relative standard deviations are 0.517% for the pincell and 0.504% for the colorset, which is a difference of only 0.013%. In terms of pcm, this difference is only approximately 18 pcm.

Assume the PWR colorset at HFP is our "real world" case, or the case in which we are interested in quantifying the uncertainty of k-eff due to cross section uncertainty. Rather than running the calculation on the full "real world" case one thousand times, at a relatively large computational

Table 3.9 Comparison of results with benchmark average (B.A.) results.

Case	k-eff	B.A. k-eff	RSD [%]	B.A. RSD [%]
PPZ	1.4276	1.4235	0.513	0.586
PPF	1.4097	1.4075	0.517	0.293
PLUZ	1.4145	1.4123	0.502	0.458
PLUF	1.3988	1.3979	0.505	0.471
BPZ	1.3445	1.3408	0.553	0.603
BPF	1.2298	1.2257	0.607	0.680
BLUZ	1.1073	1.1067	0.559	0.519
BLUF	1.0789	N/A	0.576	0.471

cost, we could use the simpler PWR pincell at HFP case and obtain an uncertainty in k-eff due to cross section uncertainty that differs from the results from the colorset case by only 18 pcm. This is a negligible difference when considering the fact that the standard deviation of k-eff is on the order of 700 pcm for these cases, and that the error incurred by limiting the sample to one thousand cross section libraries is on the order of tens of pcm. Therefore, for this case at least, the uncertainty in k-eff due to cross section uncertainty can be approximated quite accurately by a simplified version of the problem at a reduced computational cost.

This trend of problem complexity not affecting the relative standard deviation is also present in the BWR benchmark problems, but to a lesser degree, as the largest BWR problem run was only a single lattice. One trend that was noticed in the BWR data is that the difference in relative standard deviation between cases at HZP and HFP was larger than for the PWR cases.

3.2.1 Comparison to Benchmark Averages

To establish confidence in the validity of the obtained uncertainty data, it is helpful to compare the k-eff uncertainty data to results others independently obtained, usually by different methods. To alleviate any issues in the potential use of different cross section data, a comparison will only be made with the reported benchmark average quantities [Iva13]. In the future, it would be beneficial to perform a direct comparison with only the results generated using the same cross section data. Due to limited availability of benchmark average data, only eight of the fourteen cases could be compared. The comparison for these eight cases is shown in Table 3.9.

Two notable deviations from the reported benchmark average results [Bra14] are present. One is the extremely low benchmark average RSD for the PWR Pincell at HFP case (PPF). Our study produced a relative standard deviation of 0.517 %, while the reported benchmark average was 0.293 %. Analysis of the submitted results [Hou19] revealed that one benchmark participant submitted results with extremely low uncertainty, and both benchmark results summaries note that benchmark participation was limited for many models. It is likely that the benchmark participant with low

Table 3.10 Absolute differences in obtained results with respect to benchmark average results.

Case	k-eff [%]	k-eff [pcm]	RSD [%]
PPZ	+0.41	+410	-0.073
PPF	+0.22	+220	N/A
PLUZ	+0.22	+220	+0.044
PLUF	+0.09	+90	+0.034
BPZ	+0.40	+400	-0.050
BPF	+0.41	+410	-0.073
BLUZ	+0.06	+60	+0.040
BLUF	N/A	N/A	+0.105

uncertainty severely reduced the benchmark average uncertainty for this case. The benchmark average uncertainty for the PPF case is a notable outlier from all other benchmark averages, which are mostly consistent across cases.

The other notable discrepancy between our results and the benchmark average results is for the BWR Lattice Unrodded at HFP (BLUF). In the results summary [Bra14], the reported data for the BLUF case is identical to the data for the PWR Lattice Unrodded at HFP case (PLUF). It is therefore likely that this is simply a transcription error in the results summary, and is omitted from Table 3.9.

A quantified comparison of the obtained results to the benchmark average quantities is shown in Table 3.10. Results for the difference in k-eff are shown in units of percent and percent mille (pcm), which is a common unit of reactivity comparison, and is simply one ten-thousandth of a percent. The largest discrepancy between our mean k-eff values and the benchmark average values was for the PPZ and BPF cases, with a difference of +0.41 % (+410 pcm) for both cases. The largest difference between our results for the relative standard deviation of k-eff and the benchmark average results was for the BLUF case, with a difference in RSD of +0.105 %.

It can be beneficial to represent the standard deviation in terms of pcm. The obtained standard deviations are compared with the benchmark quantities in Table 3.11 in units of pcm. The largest underestimation of the standard deviation with respect to the benchmark average was for the PPZ case, with an underestimation of 102 pcm. However, it should be noted that the values in Table 3.11 are not relative, while those in Table 3.10 are. In fact, from Table 3.10, the underestimation of the BPF case is the same as the PPZ case in relative terms. The largest overestimation with respect to the benchmark average of the standard deviation was for the PLUZ case, both in absolute (+63 pcm) and relative (+0.044 %) terms.

For the cases with realistic results for the benchmark average values, the k-eff and k-eff uncertainty results generated using our 1000 perturbed cross section libraries with VERA are comparable to the benchmark average results. For the six cases that could clearly be compared, the differences in relative standard deviation ranged from -0.073 % to +0.040 %. The results used to generate the

Table 3.11 Comparison of standard deviation (SD) with benchmark average (B.A.) standard deviation.

Case	SD [pcm]	B.A. SD [pcm]	Diff. [pcm]
PPZ	732	834	-102
PPF	729	N/A	N/A
PLUZ	710	647	+63
PLUF	706	658	+48
BPZ	744	809	-65
BPF	747	833	-86
BLUZ	619	574	+45
BLUF	622	N/A	N/A

benchmark average values were generated from many different benchmark participants using different nuclear data libraries, transport codes, and uncertainty quantification methods [Hou19]. The fact that our uncertainty results are close to the benchmark average uncertainty results provides evidence not only that our method of stochastic sampling using perturbed cross section libraries is sound, but that VERA/MPACT is a valid model for uncertainty propagation. The confidence in our uncertainty quantification method gained by the comparison of our results to the UAM LWR benchmark average results is particularly important when considering its application to non-benchmark cases, such as in Chapter 4.

CHAPTER

4

WATTS BAR CROSS SECTION UNCERTAINTY

One of the main benchmarks used for the evaluation of modeling and simulation tools in CASL is the Watts Bar reactor. Just as in the UAM LWR benchmark problems, where problem complexity was increased, the complexity of the Watts Bar test cases will be increased. The simplest Watts Bar case is a two-dimensional full-core calculation, followed by a three-dimensional full-core calculation. After these steady-state cases with no thermal hydraulic feedback are analyzed, two full-core cases with feedback will be introduced – a three-dimensional full-core with thermal hydraulic feedback, and finally a three-dimensional full-core depletion calculation. The computational power associated with these cases can be compared with the much simpler UAM LWR benchmark problems, as shown in the profiling section in Chapter 2.

4.1 Watts Bar Configuration

The fuel assemblies in the Watts Bar core model are based on the actual Westinghouse assemblies loaded into the Watts Bar Unit 1 reactor [God14]. The dimensions of the fuel pins compared to the dimensions of the PWR benchmark described in Chapter 3 are shown in Table 4.1. While the fuel pins are composed of the same material, the dimensions of the Watts Bar pins are slightly smaller than the UAM PWR dimensions. The fuel pellet density of the Watts Bar fuel pins is also slightly lower than the fuel density of the UAM PWR pins.

The assembly-level parameters of the Watts Bar fuel assemblies are compared to the UAM PWR parameters in Table 4.2. The main difference between the two models is that the Watts Bar

Table 4.1 Watts Bar pin parameters [God14].

Parameter	Watts Bar	UAM PWR
Pin Pitch [cm]	1.26	1.4427
Pellet radius [cm]	0.4096	0.46955
Pellet density [g/cm^3]	10.257	10.283
Gap thickness [cm]	0.0084	0.00955
Cladding thickness [cm]	0.057	0.0673
Pellet material	UO_2	UO_2
Gap material	He	He
Cladding material	Zirc-4	Zirc-4

Table 4.2 Watts Bar and UAM LWR lattice parameters [God14] [Iva13].

Parameter	Watts Bar	UAM PWR
Assembly pitch [cm]	21.5	21.811
Dimension	17 x 17	15 x 15
Guide tubes	24	16
Instrumentation tubes	1	1
Guide tube inner radius [cm]	0.561	0.63245
Guide tube thickness [cm]	0.041	0.04065
Guide tube material	Zirc-4	Zirc-4
Inst. tube inner radius [cm]	0.559	0.56005
Inst. tube thickness [cm]	.046	0.06605
Inst. tube material	Zirc-4	Zirc-4

assemblies have 17 pins across each assembly, while the UAM LWR only has 15. This is compensated for by the smaller pin dimensions and pin pitch in the Watts Bar assemblies so the assembly pitches are comparable between the two models.

A loading pattern for the full-core in quarter-core symmetry for the Watts Bar reactor with fuel enrichment and pyrex absorber rods is shown in Figure 4.1. This is the loading pattern of the first cycle of the Watts Bar Unit 1 reactor. The levels of enrichment are also in line with the enrichment in the UAM LWR assembly, which was 2.93 w/o U-235. The Watts Bar core consists of assemblies with three enrichment levels: 2.11 w/o, 2.619 w/o, and 3.1 w/o. While the Watts Bar model does not contain any integral burnable absorbers, such as gadolinia in the UAM LWR, it does contain a number of discrete burnable absorbers in the form of Pyrex glass rods. The number of Pyrex rods per assembly is shown in Figure 4.1 below the fuel enrichment.

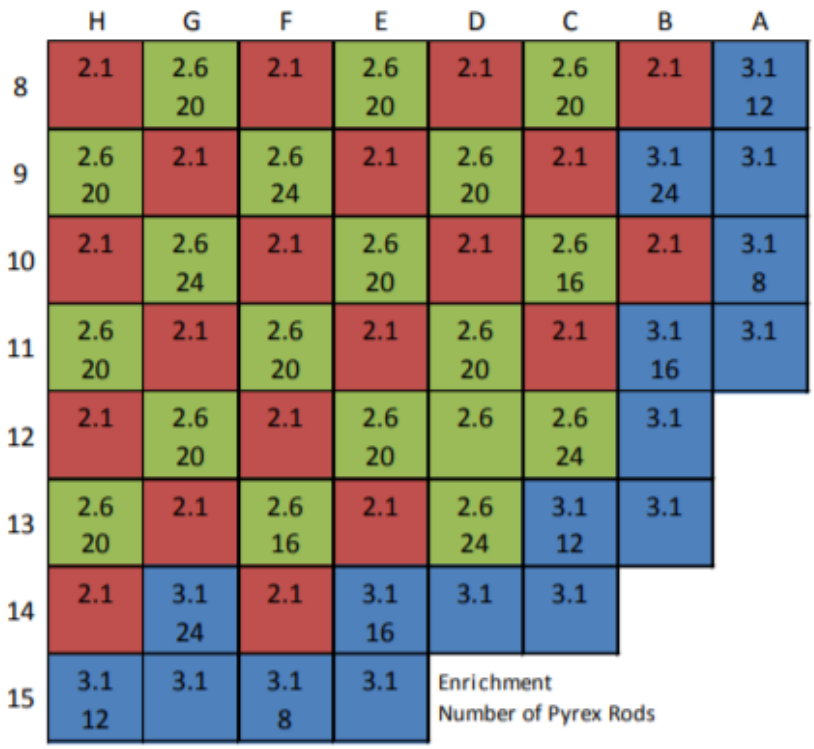

Figure 4.1 Watts Bar loading pattern [God14].

4.2 2D Watts Bar

The two-dimensional full-core Watts Bar case is similar in complexity to the colorset cases from the UAM LWR benchmark, but with more fuel lattices in a more realistic arrangement. The size of the full-core lattice problem is obviously much larger than the UAM LWR colorset benchmark problems. While the colorset problems were an arrangement of nine fuel lattices, the Watts Bar full-core consists of 193 fuel lattices. The two-dimensional case corresponds to a two-dimensional version of the fifth VERA progression problem.

The reference k-eff value calculated by the continuous-energy Monte Carlo code KENO-VI [Gol11] agrees well with the average k-eff calculated with 1000 samples. The reference solution [God14] for this problem is given as 1.004085 ± 0.000008, which was obtained with 2.5×10^{10} total particles over 5000 total generations. The estimated mean obtained through stochastic sampling with 1000 cross section libraries was 1.004771. The absolute difference between the estimated mean and the reference solution is approximately 70 pcm. This small difference provides evidence that attests both to the accuracy of this model's simulation in MPACT and that the cross section libraries are well–sampled statistically.

Shown in Figure 4.2 is the k-eff distribution for the two-dimensional Watts Bar full-core case. The distribution appears approximately normal, with a mean of 1.004771 and standard deviation of 556 pcm. Comparison with other Watts Bar cases and the UAM LWR benchmark problems is shown

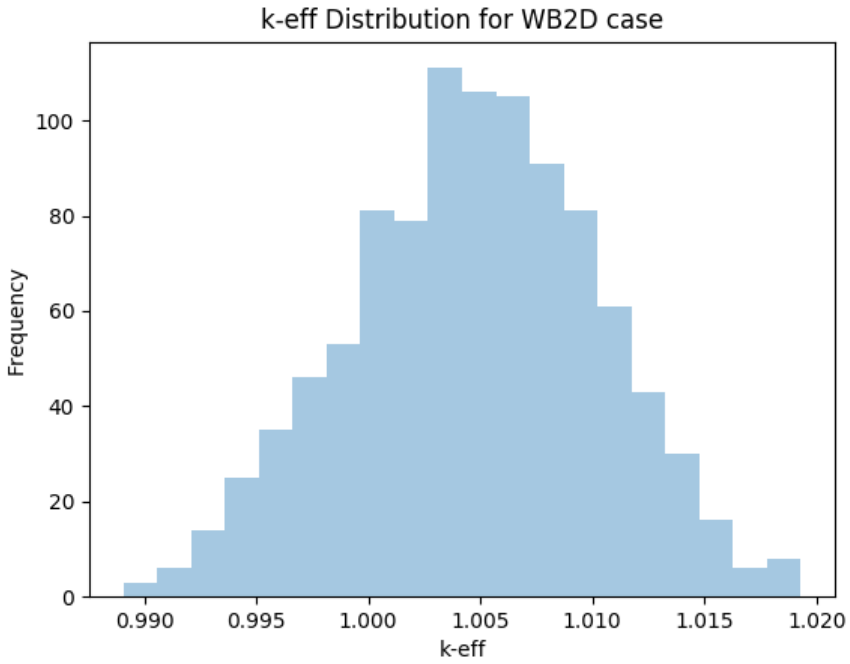

Figure 4.2 Two-dimensional Watts Bar k-eff distribution.

later in this chapter.

4.3 3D Watts Bar

The three-dimensional Watts Bar configuration is a direct model of the Watts Bar reactor at conditions seen in the actual reactor during zero-power physics tests at the start of each cycle. Because this is a simulation of a real reactor configuration, this case represents an important basis for validation of the modeling tools in VERA. Since the simulation tools in VERA were in part evaluated and validated using this case, we should have very high confidence in the validity of the obtained data for this case.

This case is of course very similar to the preceding case – the two-dimensional Watts Bar full-core. The fuel composition, loading pattern, and state parameters are identical between the two cases, with the only difference being that axial features such as the fuel rod stack, upper and lower plenums, and end plugs are explicitly modeled. Spacer grids and other core structural materials are semi-explicitly modeled. Specific details of the model, which corresponds to the fifth progression problem in the VERA benchmark, can be found in the VERA benchmark document [God14].

The simulated case was evaluated by the continuous-energy Monte Carlo code in SCALE, KENO-VI [Gol11], and deviates from the measured value by only 10 pcm, confirming the validity of KENO-VI as a reference. To obtain an error of 1 pcm in k-eff from KENO-VI, 7.9×10^9 particles were used over 1,500 generations. Note that because fewer particles were used in the three-dimensional case

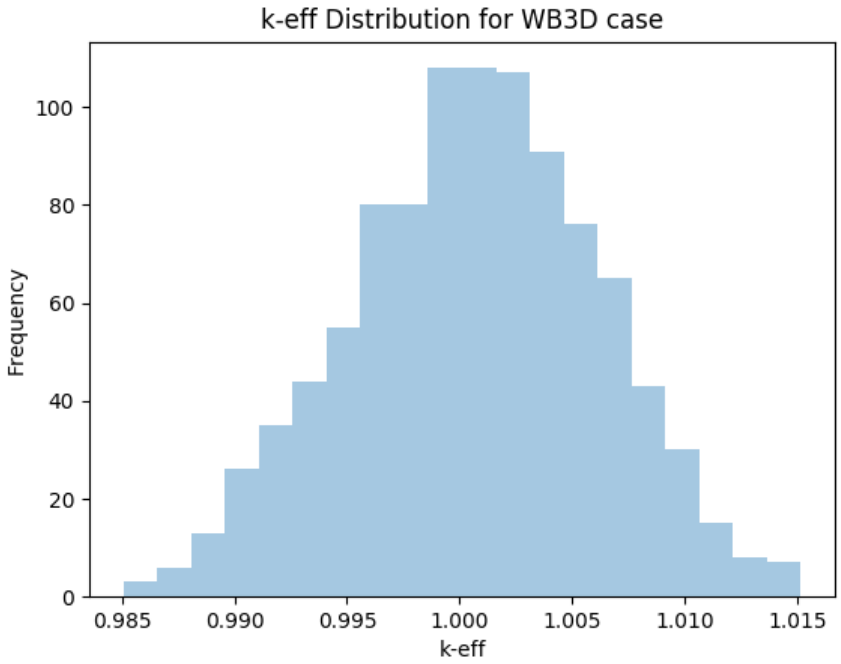

Figure 4.3 Three-dimensional Watts Bar k-eff distribution.

than in the two-dimensional case, the uncertainty in the two-dimensional case is lower than in the three-dimensional case. The reference solution [God14] is 0.999899 ± 0.000010, compared to a measured value of 1.00000. Our estimate of the mean k-eff for this case was 1.000704. The mean k-eff from 1000 samples of this case is again close to the reference value, overestimating k-eff by approximately 80 pcm.

Shown in Figure 4.3 is the k-eff distribution for the three-dimensional Watts Bar full-core case. Just as in the two-dimensional case, k-eff appears to be normally distributed about the mean of 1.000704, with a standard deviation of 552 pcm. The estimated relative standard deviation for the two-dimensional model and three-dimensional model differed by only 0.001%, or 4 pcm in absolute terms.

4.4 3D Watts Bar with Thermal Hydraulic Feedback

The addition of thermal-hydraulic feedback to the previous case represents a model which is closer to the typical operation conditions of the reactor. While the three-dimensional case with no feedback models a reactor state before the beginning of each cycle, the three-dimensional case with feedback models a reactor state during operation. Since this is a case in which neutronics and thermal hydraulics are coupled, and is not a pure neutronics exercise, no Monte Carlo reference solution exists for this case. This case reaches the level of complexity where a coupled core simulator such as VERA is the only real option for modeling reactor behavior. This is the seventh progression problem

in the CASL benchmark [God14].

The geometry of this case is identical to the previous case, the three-dimensional full-core at zero power. The difference is that this case models the reactor at power at the beginning of life (BOL). The idea of thermal hydraulic feedback, or coupling between neutronics and thermal hydraulics, is important because of the effect of temperature on the behavior of the moderator. For instance, the neutron distribution within the reactor affects the power distribution within the reactor. While it is desirable to have the neutron, and power, distribution as uniform as possible throughout the core, local hot-spots of high power will occur in high-reactivity regions. The increased power production increases temperature in these locations. Moderator, which is liquid water, flowing in these high-temperature locations subsequently increases in temperature and decreases in density. The decreased moderator density reduces moderation in these locations, reducing reactivity, and reducing the neutron density.

In VERA, thermal hydraulic behavior is modeled using the subchannel code CTF [Sal16]. In CTF, mass, momentum, and energy equations in addition to a large number of closure equations are solved for each control volume in the flow region. This makes the deterministic uncertainty analysis described in Chapter 1 virtually impossible, so a case of this complexity can best be studied using the stochastic sampling method.

The values of k-eff were obtained in a slightly different method for this case. Rather than directly calculating k-eff, a critical boron search was performed. This is a common calculation performed on PWRs. The critical boron concentration is the concentration of soluble boron, which is a relatively strong absorber, in the moderator that would bring the system into a critical configuration ($k = 1$). So for each of the 1000 samples, a critical boron concentration in parts per million (ppm) was obtained. The conversion from critical boron concentration to k-eff was accomplished using a quadratic fit on a range of critical boron concentrations. Eighteen cases were run, and values of k-eff obtained, with boron concentrations ranging from 825 ppm to 945 ppm for the same system configuration. A quadratic fit of the form shown in Equation 4.1 was performed on these eighteen data points, and the coefficients A, B, and C were obtained. In Equation 4.1, k represents k-eff and b_c represents the critical boron concentration.

$$k - 1 = A b_c^2 + B b_c + C \tag{4.1}$$

This quadratic fit was applied to the 1000 values of critical boron concentration to convert the critical boron concentration to k-eff. Shown in Figure 4.4 is the k-eff distribution for the Watts Bar Unit 1 full-core with thermal hydraulic feedback. The estimated mean k-eff was 1.000072, with a standard deviation of 556 pcm using 1000 samples. Although there is neither measured data nor a reference solution for comparison, both of these estimated parameters agree well with the three-dimensional case with no thermal hydraulic feedback.

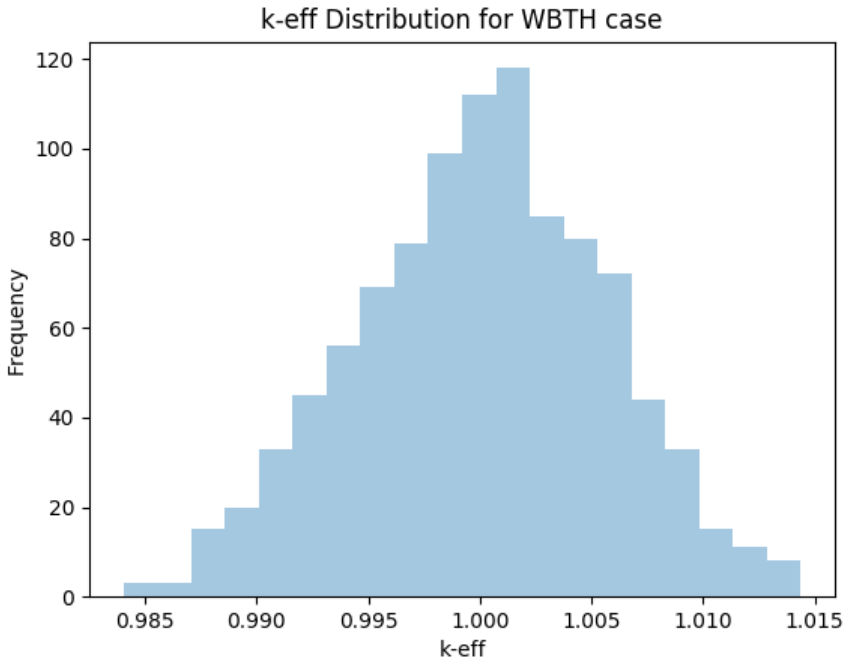

Figure 4.4 Watts Bar with thermal hydraulic feedback k-eff distribution.

4.5 3D Watts Bar Full-Core Depletion

The three-dimensional full-core Watts Bar depletion is the most complex case examined in this study. While the three-dimensional full-core with feedback represents operating conditions at the beginning of the cycle, the depletion calculation represents operating conditions at several points in the cycle. In this case, 32 state-points within the cycle were examined, with exposures from 0 Effective Full Power Days (EFPD), which corresponds to the Beginning of the Cycle (BOC) to 441 EFPD, which is the End of the Cycle (EOC). This case is the ninth progression problem in the CASL benchmark [God14].

Like the three-dimensional case with thermal hydraulic feedback, critical boron concentration was calculated in lieu of k-eff. For this case, the critical boron concentration was converted to reactivity assuming a boron reactivity coefficient of $\alpha_B = -8$ pcm/ppm, which is typical in PWRs.

Shown in Figure 4.5 is the critical boron concentration for the 29 full-core depletion cases. Each of the 29 series in Figure 4.5 is the critical boron concentration as a function of exposure for a perturbed cross section library. With 29 cross section libraries and 32 depletion steps, a full-core calculation with thermal hydraulic feedback was performed 928 times. This case is illustrative of the large computational cost that is necessary to perform stochastic sampling for large cases.

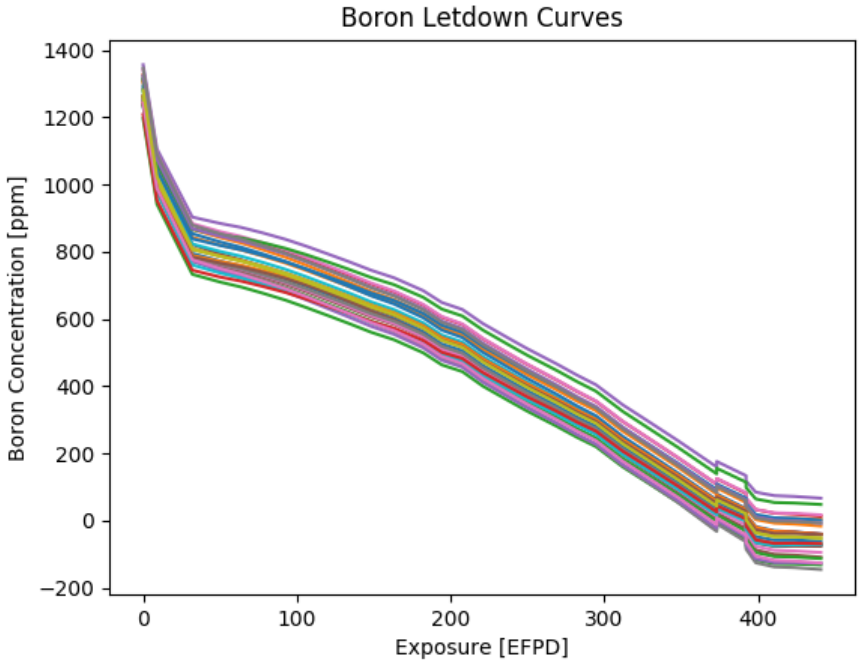

Figure 4.5 Critical boron concentration as a function of exposure, full-core Watts Bar.

4.6 Comparison of Cases

For the three full-core Watts Bar cases that used all 1000 perturbed cross section libraries, the uncertainty in k-eff due to cross section uncertainty was virtually identical. The results for these full-core Watts Bar cases are shown in Table 4.3. For the simplest full-core case, the two-dimensional model with no thermal-hydraulic feedback, the relative standard deviation of k-eff was estimated to be 0.553%, while the full three-dimensional case with thermal-hydraulic feedback was 0.556%, a difference of only 0.003%. The three-dimensional case with no thermal-hydraulic feedback, which falls between these two cases in terms of complexity, produced a similar relative standard deviation of 0.552%.

Interestingly, the introduction of thermal hydraulic feedback to the three-dimensional full-core caused almost no change in the estimated standard deviation. The fact that the k-eff uncertainty due to cross section uncertainty is not significantly dependent on thermal hydraulic feedback in this case reduces the advantage that stochastic sampling has over deterministic methods.

The limited number of cases run for each state point in the depletion study precludes any full comparison with the models for which 1000 cases were used. Therefore the uncertainty results for the depletion study are omitted from Table 4.3. However, the convergence of statistical parameters through 29 cross section libraries is possible, and is primarily examined in Chapter 5.

Table 4.3 Estimated Mean and Standard Deviation for Watts Bar Cases.

Case	k-eff	RSD [%]	SD [pcm]
2D	1.004771	0.553	556
3D	1.000704	0.552	552
3D-TH	1.000072	0.556	556

CHAPTER

5

SAMPLE SIZE

One of the most important considerations when using stochastic sampling for uncertainty quantification is determination of sample size. For all cases examined to this point, with the exception of the Watts Bar full-core depletion, one thousand samples in the form of one thousand perturbed cross section libraries were used to form parameter estimates. However, if the use of five hundred libraries yielded results within an acceptable margin of the results from one thousand libraries, the slight reduction in accuracy of uncertainty quantification could be offset by the 50 % reduction in computational cost. This is particularly important for extremely expensive models, such as the full-core Watts Bar Unit 1 cases.

In this chapter, the impact of sample size on estimates of population parameters will be discussed, and jackknife resampling will be introduced as a technique for approximating the relative standard error of parameter estimates. Convergence behavior of the uncertainty in k-eff due to cross section uncertainty is tabulated for every case described in Chapters 3 and 4.

5.1 Standard Error of Estimators

Consider a set of n observations $(x_1, x_2, ..., x_n)$. The mean value $\bar{x}$ of this set can be calculated using Equation 5.1.

$$\bar{x} = \frac{1}{n}\sum_{i=1}^{n} x_i \tag{5.1}$$

If it were possible to obtain an infinite sample, $n \to \infty$, the mean value $\bar{x}$ of the set of observations would be the mean value μ of the population. This is shown in Equation 5.2.

$$\lim_{n\to\infty} \frac{1}{n}\sum_{i=1}^{n} x_i = \mu \tag{5.2}$$

An infinite sample is not possible, so the true population mean μ is unobtainable. It can, however, be estimated. The estimate for the population mean will be denoted by $\hat{x}$. The best estimate $\hat{x}$ for the population mean μ is simply the sample mean $\bar{x}$.

$$\hat{x} = \bar{x} \approx \mu \tag{5.3}$$

The reason the estimate of the mean $\hat{x}$ is not equivalent to the population mean μ is due to the approximation $n \approx \infty$. As the sample size n grows, the approximation $n \approx \infty$ improves, and the estimate of the mean improves. The degree to which it improves can be directly calculated.

Another statistic of interest is the variance. The variance of a population can be calculated using Equation 5.4.

$$\sigma^2 = \frac{1}{n}\sum_{i=1}^{n}\big(x_i - \mu\big)^2 \tag{5.4}$$

It has already been established that in the case of an infinite population, the population mean μ is unobtainable. Therefore the population variance is also unobtainable. Once again, an estimate for the population variance must be used. The variance estimator will be denoted by $\hat{\sigma}^2$.

A reasonable strategy for estimating the variance would be to simply use the variance of the sample. Consider again the n observations $(x_1, x_2, ..., x_n)$. The variance $\hat{\sigma}^2$ of this sample of n observations can be calculated using Equation 5.5. Note that the estimate for the mean $\hat{x}$ and the sample mean $\bar{x}$ are interchangeable.

$$\hat{\sigma}^2 = \frac{1}{n}\sum_{i=1}^{n}(x_i - \hat{x})^2 \tag{5.5}$$

This estimate of the variance is biased by a factor of $\frac{n-1}{n}$ [Weia]. An unbiased estimate of the variance is shown in Equation 5.6. This is used as the estimate for the population variance.

$$\hat{\sigma}^2 = \frac{1}{n-1}\sum_{i=1}^{n}(x_i - \hat{x})^2 \tag{5.6}$$

The standard deviation σ is a desirable alternative to the variance σ^2 because it shares units with the observations x_i and the mean μ. The standard deviation is the square root of the variance, as shown in Equation 5.7

$$\sigma = \sqrt{\sigma^2} = \left[\frac{1}{n}\sum_{i=1}^{n}\big(x_i - \mu\big)^2\right]^{1/2} \tag{5.7}$$

An estimate $\hat{\sigma}$ for the standard deviation σ is the square root of the variance estimate $\hat{\sigma}^2$, shown in Equation 5.8.

$$\hat{\sigma} = \sqrt{\hat{\sigma}^2} = \left[\frac{1}{n-1} \sum_{i=1}^{n} (x_i - \hat{x})^2 \right]^{1/2} \tag{5.8}$$

It is evident that the estimates of any of the mean, variance, or standard deviation are not perfect calculations of the true population parameters. The error introduced by using estimates to estimate population parameters is known as the "standard error" (se) of the statistic. Once the standard errors of the population parameters are computed, each of the sample parameters can be reported with their associated standard errors, as shown in Equations 5.9a, 5.9b, and 5.9c for the mean, variance, and standard deviation respectively.

$$\mu = \hat{x} \pm \text{se}_{\hat{x}} \tag{5.9a}$$

$$\sigma^2 = \hat{\sigma}^2 \pm \text{se}_{\hat{\sigma}^2} \tag{5.9b}$$

$$\sigma = \hat{\sigma} \pm \text{se}_{\hat{\sigma}} \tag{5.9c}$$

The focus now lies on the computation of the standard error $\text{se}_{\hat{x}}$ of the mean, the standard error $\text{se}_{\hat{\sigma}^2}$ of the variance, and the standard error $\text{se}_{\hat{\sigma}}$ of the standard deviation.

5.1.1 Standard Error of the Mean

The standard error of the mean is given by Equation 5.10.

$$\text{se}_{\hat{x}} = \sqrt{\text{Var}(\hat{x})} \tag{5.10}$$

Note that when the true population mean μ is used in Equation 5.10 in place of the estimate $\hat{x}$, $\text{Var}(\mu) = 0$, so $\text{se}_{\mu} = 0$. Inserting the estimate for the mean, shown in Equation 5.3, into Equation 5.10,

$$\text{se}_{\hat{x}} = \sqrt{\text{Var}\left(\frac{1}{n} \sum_{i=1}^{n} x_i \right)} \tag{5.11}$$

Using the identity $\text{Var}(a\,x) = a^2 \text{Var}(x)$, where a is a constant and x is a random variable,

$$\text{se}_{\hat{x}} = \frac{1}{n} \sqrt{\text{Var}\left(\sum_{i=1}^{n} x_i \right)} \tag{5.12}$$

The variance of the sum of i random variables x_i is simply the sum of the covariances between each of the random variables.

$$\text{Var}\left(\sum_{i=1}^{n} x_i \right) = \sum_{i,j=1}^{n} \text{Cov}(x_i, x_j) \tag{5.13}$$

Under the assumption that the observations are independently generated, the covariances between different observations are zero, and the sum of the covariances simplifies to the sum of the variances.

$$\text{Var}\left(\sum_{i=1}^{n} x_i\right) = \sum_{i=1}^{n} \text{Var}(x_i) \tag{5.14}$$

Recall that the variance of the observations is denoted by $\hat{\sigma}^2$ and is constant. Therefore the variance of the sum of the n observations reduces to

$$\text{Var}\left(\sum_{i=1}^{n} x_i\right) = \sum_{i=1}^{n} \hat{\sigma}^2 = n\hat{\sigma}^2 \tag{5.15}$$

Inserting this into Equation 5.12,

$$\text{se}_{\hat{x}} = \frac{1}{n}\sqrt{n\hat{\sigma}^2} \tag{5.16}$$

which simplifies to

$$\text{se}_{\hat{x}} = \frac{\hat{\sigma}}{\sqrt{n}} \tag{5.17}$$

5.1.2 Standard Error of the Variance

The standard error of the variance is given by Equation 5.18.

$$\text{se}_{\hat{\sigma}^2} = \sqrt{\text{Var}(\hat{\sigma}^2)} \tag{5.18}$$

The derivation of $\text{Var}(\hat{\sigma}^2)$ for any distribution is illustrated in Rao's work [Rao73]. The variance of the unbiased variance estimate is

$$\text{Var}(\hat{\sigma}^2) = \frac{1}{n}\left(\mu_4 - \sigma^4\right) + \frac{2}{n(n-1)}\sigma^4, \tag{5.19}$$

and the standard error of the sample variance is

$$\text{se}_{\hat{\sigma}^2} = \left[\frac{1}{n}\left(\mu_4 - \sigma^4\right) + \frac{2}{n(n-1)}\sigma^4\right]^{1/2} \tag{5.20}$$

As established previously, the population standard deviation σ^2 is unobtainable with a finite sample. The exact fourth central moment μ_4 of the population distribution is also unobtainable without the population distribution, so the true standard error of the sample variance is incalculable without full knowledge of the distribution. The only remedy would be to calculate an estimate of the fourth central moment $\hat{\mu_4}$ of the distribution. Since exact knowledge of the population distribution is seldom known, other standard error estimation techniques are often used, such as jackknife resampling [Efr81a].

5.2 Jackknife Resampling as a Technique to Estimate Standard Errors of Parameters

It was mentioned in the previous section that while the standard error $\text{se}_{\hat{x}}$ of the sample mean could be calculated using the sample variance, the standard error $\text{se}_{\hat{\sigma}^2}$ of the sample variance and the standard error $\text{se}_{\hat{\sigma}}$ of the sample standard deviation could not be calculated without knowing the population distribution. There are however techniques for estimating the standard errors of the variance and standard deviation. Equations 5.9a, 5.9b, and 5.9c can be re-written in terms of obtainable quantities.

$$\mu \approx \hat{x} \pm \text{se}_{\hat{x}} \tag{5.21a}$$

$$\sigma^2 \approx \hat{\sigma}^2 \pm \hat{\text{se}}_{\hat{\sigma}^2} \tag{5.21b}$$

$$\sigma \approx \hat{\sigma} \pm \hat{\text{se}}_{\hat{\sigma}} \tag{5.21c}$$

One technique for estimating the distribution of a population parameter is jackknife resampling [Efr81a]. In jackknife resampling, a sample of size n is decomposed into n samples by leaving one datapoint out of the dataset. Consider a population parameter ρ, with an estimator $\hat{\rho}(x_1, x_2, ..., x_n)$ that generates estimates of the population parameter based on a sample $(x_1, x_2, ..., x_n)$ of size n. In the jackknife method, the parameter is instead estimated n times. The jackknife estimators are generated using Equation 5.22.

$$\hat{\rho}_i = \hat{\rho}(x_1, x_2, ..., x_{i-1}, x_{i+1}, ..., x_n)\ ,\ i = 1, ..., n \tag{5.22}$$

We now have a set of n jackknife estimates $(\hat{\rho}_1, \hat{\rho}_2, ..., \hat{\rho}_n)$ of the parameter ρ. The mean of the jackknife estimates can be calculated, as shown in Equation 5.23

$$\hat{\rho}_{(.)} = \frac{1}{n}\sum_{i=1}^{n}\hat{\rho}_i \tag{5.23}$$

The jackknife estimate of the standard error $\hat{\text{se}}_{\rho}$ of the parameter ρ is calculated using Equation 5.24 [Efr81b].

$$\hat{\text{se}}_{\rho} = \sqrt{\frac{n-1}{n}\sum_{i=1}^{n}\left(\hat{\rho}_i - \hat{\rho}_{(.)}\right)^2} \tag{5.24}$$

5.3 Jackknife Estimates of Standard Error of Standard Deviation Estimate for Benchmark Cases

Of particular interest in this case is the standard error $\text{se}_{\hat{\sigma}}$ of the standard deviation estimate $\hat{\sigma}$. Tables 5.1 – 5.14 show the reduction in the estimated standard error of the estimated standard

Table 5.1 Reduction of standard error of standard deviation estimate with increasing sample size, PWR Pincell at HZP (PPZ).

n	$\hat{\sigma}$ [pcm]	$\hat{\text{se}}_{\hat{\sigma}}$ [pcm]	RSE [%]
5	700	301	43.0
10	566	141	24.9
30	612	80.0	13.7
50	700	72.5	10.4
100	722	52.6	7.29
250	712	33.0	4.63
500	722	22.1	3.06
1000	732	15.3	2.09

Table 5.2 Reduction of standard error of standard deviation estimate with increasing sample size, PWR Pincell at HFP (PPF).

n	$\hat{\sigma}$ [pcm]	$\hat{\text{se}}_{\hat{\sigma}}$ [pcm]	RSE [%]
5	702	295	42.0
10	572	138	24.1
30	615	82.8	13.5
50	697	72.4	10.4
100	719	52.3	7.27
250	710	32.9	4.63
500	719	22.0	3.06
1000	729	15.3	2.10

deviation as sample size n grows. Also shown is the relative standard error, which is defined by Equation 5.25. The full data used to generate the tables can be used to construct confidence intervals. This is shown for all cases in Appendix A.

$$\text{RSE} = \frac{\hat{\text{se}}_{\hat{\sigma}}}{\hat{\sigma}} \tag{5.25}$$

For all fourteen UAM LWR benchmark cases, the estimated standard error of our standard deviation estimate fell within the range of 9.57 pcm to 15.8 pcm when using all 1000 samples. When adjusted for the magnitude of the standard deviation estimate, the relative standard error estimates ranged from 2.04 % to 2.12 %. Little variation is seen between the relative standard error estimates for the different cases. Interestingly, the estimated relative standard error remains below 5 % when only using 250 samples for all cases, ranging from 4.32 % to 4.76 %. The estimated relative standard error for most cases reaches 10 % at approximately 50 samples, so significant error would be incurred by using fewer than 50 samples.

Table 5.3 Reduction of standard error of standard deviation estimate with increasing sample size, BWR Pincell at HZP (BPZ).

n	$\hat{\sigma}$ [pcm]	$\hat{se}_{\hat{\sigma}}$ [pcm]	RSE [%]
5	712	375	52.7
10	548	178	32.5
30	599	95.8	16.0
50	696	76.5	11.0
100	733	55.4	7.56
250	716	34.1	4.76
500	733	22.4	3.06
1000	744	15.6	2.10

Table 5.4 Reduction of standard error of standard deviation estimate with increasing sample size, BWR Pincell at HFP (BPF).

n	$\hat{\sigma}$ [pcm]	$\hat{se}_{\hat{\sigma}}$ [pcm]	RSE [%]
5	727	321	44.2
10	589	141	23.9
30	625	86.7	13.9
50	711	77.5	10.9
100	749	55.5	7.41
250	732	34.6	4.73
500	745	22.9	3.07
1000	747	15.8	2.12

Table 5.5 Reduction of standard error of standard deviation estimate with increasing sample size, Unrodded PWR Lattice at HZP (PLUZ).

n	$\hat{\sigma}$ [pcm]	$\hat{se}_{\hat{\sigma}}$ [pcm]	RSE [%]
5	678	297	43.8
10	547	138	25.2
30	583	79.7	13.7
50	667	67.9	10.2
100	688	49.7	7.22
250	685	31.5	4.60
500	696	21.1	3.03
1000	710	14.8	2.08

Table 5.6 Reduction of standard error of standard deviation estimate with increasing sample size, Unrodded PWR Lattice at HFP (PLUF).

n	$\hat{\sigma}$ [pcm]	$\hat{se}_{\hat{\sigma}}$ [pcm]	RSE [%]
5	673	291	43.2
10	548	134	24.5
30	583	78.6	13.5
50	663	67.5	10.2
100	684	49.3	7.21
250	682	31.4	4.60
500	693	21.0	3.03
1000	706	14.7	2.08

Table 5.7 Reduction of standard error of standard deviation estimate with increasing sample size, Rodded PWR Lattice at HZP (PLRZ).

n	$\hat{\sigma}$ [pcm]	$\hat{se}_{\hat{\sigma}}$ [pcm]	RSE [%]
5	596	254	42.6
10	486	131	27.0
30	497	64.6	13.0
50	547	53.6	9.80
100	560	36.9	6.59
250	556	24.1	4.33
500	557	17.0	3.05
1000	570	11.7	2.05

Table 5.8 Reduction of standard error of standard deviation estimate with increasing sample size, Rodded PWR Lattice at HFP (PLRF).

n	$\hat{\sigma}$ [pcm]	$\hat{se}_{\hat{\sigma}}$ [pcm]	RSE [%]
5	591	250	42.3
10	485	128	26.4
30	498	64.0	12.9
50	545	53.3	9.78
100	557	36.7	6.59
250	553	23.9	4.32
500	555	16.9	3.05
1000	567	11.7	2.06

Table 5.9 Reduction of standard error of standard deviation estimate with increasing sample size, Unrodded BWR Lattice at HZP (BLUZ).

n	$\hat{\sigma}$ [pcm]	$\hat{se}_{\hat{\sigma}}$ [pcm]	RSE [%]
5	626	322	51.4
10	481	140	29.1
30	477	73.7	15.5
50	550	57.1	10.4
100	571	41.1	7.20
250	574	26.6	4.63
500	598	17.9	2.99
1000	619	12.9	2.09

Table 5.10 Reduction of standard error of standard deviation estimate with increasing sample size, Unrodded BWR Lattice at HFP (BLUF).

n	$\hat{\sigma}$ [pcm]	$\hat{se}_{\hat{\sigma}}$ [pcm]	RSE [%]
5	628	306	48.7
10	488	131	26.8
30	487	70.8	14.5
50	555	57.1	10.3
100	578	40.9	7.08
250	580	26.9	4.64
500	605	18.1	2.99
1000	622	12.9	2.07

Table 5.11 Reduction of standard error of standard deviation estimate with increasing sample size, Rodded BWR Lattice at HZP (BLRZ).

n	$\hat{\sigma}$ [pcm]	$\hat{se}_{\hat{\sigma}}$ [pcm]	RSE [%]
5	515	247	48.0
10	378	125	33.1
30	366	54.7	14.9
50	421	44.2	10.5
100	437	30.4	6.96
250	443	20.1	4.54
500	460	13.7	2.98
1000	475	9.76	2.05

Table 5.12 Reduction of standard error of standard deviation estimate with increasing sample size, Rodded BWR Lattice at HFP (BLRF).

n	$\hat{\sigma}$ [pcm]	$\hat{se}_{\hat{\sigma}}$ [pcm]	RSE [%]
5	507	235	46.4
10	375	116	30.9
30	374	49.4	13.2
50	427	42.9	10.0
100	439	29.7	6.77
250	446	19.7	4.42
500	460	13.5	2.93
1000	470	9.57	2.04

Table 5.13 Reduction of standard error of standard deviation estimate with increasing sample size, PWR Colorset at HZP (PCZ).

n	$\hat{\sigma}$ [pcm]	$\hat{se}_{\hat{\sigma}}$ [pcm]	RSE [%]
5	668	292	43.7
10	538	137	25.5
30	569	78.1	13.7
50	652	66.6	10.2
100	673	48.5	7.21
250	669	30.7	4.59
500	680	20.6	3.03
1000	693	14.5	2.09

Table 5.14 Reduction of standard error of standard deviation estimate with increasing sample size, PWR Colorset at HFP (PCF).

n	$\hat{\sigma}$ [pcm]	$\hat{se}_{\hat{\sigma}}$ [pcm]	RSE [%]
5	598	349	58.4
10	538	133	24.7
30	569	77.0	13.5
50	648	66.2	10.2
100	668	48.2	7.22
250	665	30.6	4.60
500	676	20.5	3.03
1000	690	14.4	2.09

5.4 Jackknife Estimates of Standard Error of Standard Deviation Estimate for Watts Bar Cases

Tables 5.15, 5.16, and 5.17 show the convergence of the standard deviation for the two-dimensional, three-dimensional, and three-dimensional with thermal hydraulic feedback models of the full-core Watts Bar Unit 1 reactor. The Watts Bar Unit 1 cases exhibit behavior very similar to the UAM LWR cases. The estimated relative standard error of our standard deviation estimate with 1000 samples ranges from 2.05 % to 2.07 %, which falls within the range of 2.04 % to 2.12 % observed in the UAM LWR cases. The reduction of the estimated relative standard error of our standard deviation estimate below 5 % also occurs at approximately $n = 250$, just as in the UAM LWR benchmark cases.

Table 5.15 Reduction of standard error of standard deviation estimate with increasing sample size, 2D Watts Bar full-core.

n	$\hat{\sigma}$ [pcm]	$\hat{\text{se}}_{\hat{\sigma}}$ [pcm]	RSE [%]
5	599	320	53.4
10	435	172	39.5
30	465	72.4	15.6
50	520	57.3	11.0
100	549	39.6	7.21
250	536	24.4	4.55
500	548	16.3	2.97
1000	556	11.4	2.05

Table 5.16 Reduction of standard error of standard deviation estimate with increasing sample size, 3D Watts Bar full-core.

n	$\hat{\sigma}$ [pcm]	$\hat{se}_{\hat{\sigma}}$ [pcm]	RSE [%]
5	598	315	52.7
10	434	180	41.5
30	461	72.1	15.6
50	517	57.3	11.1
100	546	39.5	7.23
250	532	24.3	4.57
500	544	16.3	3.00
1000	552	11.4	2.07

Table 5.17 Reduction of standard error of standard deviation estimate with increasing sample size, 3D Watts Bar full-core with thermal hydraulic feedback.

n	$\hat{\sigma}$ [pcm]	$\hat{se}_{\hat{\sigma}}$ [pcm]	RSE [%]
5	598	329	55.0
10	429	180	42.0
30	461	71.7	15.6
50	520	57.6	11.1
100	550	40.2	7.31
250	537	24.6	4.58
500	549	16.5	3.01
1000	556	11.5	2.07

5.5 Standard Error of Standard Deviation Estimate for Watts Bar Unit 1 Full-Core Depletion

For full-core depletion simulations, it is unrealistic to run one thousand cases with 1000 perturbed cross section libraries. Data was obtained using 29 libraries with 32 depletion steps. Critical boron concentration was calculated at each depletion step using each of the 29 cross section libraries, and converted to reactivity assuming a boron reactivity coefficient of -8 pcm/ppm. The standard error of the standard deviation estimate will be analyzed at three exposures: Beginning of Cycle (BOC - 0 MWD/MTU), Middle of Cycle (MOC - 221 MWD/MTU), and End of Cycle (EOC - 441 MWD/MTU). BOC data is shown in Table 5.18, MOC data is shown in Table 5.19, and EOC data is shown in Table 5.20.

A slight dependence of the estimated standard deviation and its estimated error on the exposure level is observed. The standard deviation and its estimated error is slightly lower at low exposures than towards the end of the cycle. The estimated standard deviation using all 29 libraries at BOC was 333 pcm, while the estimated standard deviation at EOC was 433 pcm, a 100 pcm increase. The estimated relative standard error at BOC was 11 %, while at EOC it was 12 %.

The effect of a very small sample size $n = 29$ on our estimate of the standard deviation is evident in these three cases. At all three exposure levels, the estimated relative standard error for the standard deviation estimate is above 10 %. From the data presented for the cases with 1000 samples, the general trend of our standard deviation estimate was an underestimation when fewer than 100 libraries were used, followed by a gradual increase in standard deviation with increasing sample size. It is therefore very likely that our standard deviation estimates for the depletion study are underestimations at $n = 29$, as the trend for these cases had not yet entered the gradual increase regime seen in the other cases for $n > 100$. The underestimation of the standard deviation estimate is not a surprising result. With only 29 libraries, it is unlikely that any significant outliers are included. The chance of outliers appearing in the dataset, which would cause the standard deviation estimate to increase, increases with increasing sample size.

Table 5.18 Reduction of standard error of standard deviation estimate with increasing sample size, Watts Bar full-core at BOC.

n	$\hat{\sigma}$ [pcm]	$\hat{se}_{\hat{\sigma}}$ [pcm]	RSE [%]
5	482	190	39.4
10	381	90.6	23.8
15	352	59.5	16.9
20	323	50.2	15.5
25	330	40.3	12.2
29	333	36.6	11.0

Table 5.19 Reduction of standard error of standard deviation estimate with increasing sample size, Watts Bar full-core at MOC.

n	$\hat{\sigma}$ [pcm]	$\hat{se}_{\hat{\sigma}}$ [pcm]	RSE [%]
5	570	190	33.3
10	472	91.6	19.4
15	429	70.1	16.3
20	386	64.6	16.7
25	394	49.0	12.4
29	381	44.7	11.7

Table 5.20 Reduction of standard error of standard deviation estimate with increasing sample size, Watts Bar full-core at EOC.

n	$\hat{\sigma}$ [pcm]	$\hat{se}_{\hat{\sigma}}$ [pcm]	RSE [%]
5	612	216	35.3
10	564	98.0	17.4
15	506	76.2	15.1
20	444	70.8	15.9
25	445	56.9	12.8
29	433	52.0	12.0

CHAPTER

6

MATERIAL UNCERTAINTY

To this point, only uncertainty in the microscopic cross sections has been considered. To illustrate the dependencies of k-eff, it is helpful to consider the value of k-eff in terms of reaction rates. The dependence of reaction rates on neutron energy, which appears in the energy dependence of the cross sections, will be suppressed.

$$\text{k-eff} \equiv \frac{\text{Rate of neutron production}}{\text{Rate of neutron loss}} \tag{6.1}$$

The neutron production in the system comes exclusively from fission. The fission reaction rate for isotope i in volume V is shown in Equation 6.2.

$$R_{f,i,V} = N_{i,V} \int_V dV' \, \sigma_{f,i} \phi(\vec{r}) \tag{6.2}$$

Similarly, the absorption reaction rate for isotope i in the same volume V is

$$R_{a,i,V} = N_{i,V} \int_V dV' \, \sigma_{a,i} \phi(\vec{r}). \tag{6.3}$$

If the neutron leakage out of the volume V is represented by the leakage term L, then k-eff for a system composed solely of isotope i can be expressed in terms of the fission and absorption reaction rates as in Equation 6.4

$$\text{k-eff} \equiv \frac{R_{f,i,V}}{R_{a,i,V} + L} \tag{6.4}$$

It is clear that k-eff depends on the same quantities as the fission reaction rate, the absorption

reaction rate, and the neutron leakage.

Let us now take a closer look at the dependencies of the reaction rate of reaction type x for isotope i in volume V, $R_{x,i,V}$, shown in Equation 6.5. The number density $N_{i,V}$ is the number of atoms per unit volume in the volume of interest V. The number density can be expressed in terms of the total number of atoms $n_{i,V}$ of isotope i in volume V and the volume of interest V, as shown in Equation 6.6.

$$R_{x,i,V} = N_{i,V} \int_V dV' \sigma_{x,i} \phi(\vec{r}) \tag{6.5}$$

$$N_{i,V} = \frac{n_{i,V}}{V} \tag{6.6}$$

Inserting Equation 6.6 into the reaction rate in Equation 6.5 yields

$$R_{x,i,V} = \frac{n_{i,V}}{V} \int_V dV' \sigma_{x,i} \phi(\vec{r}) \tag{6.7}$$

From Equation 6.7, it is clear that the reaction rate $R_{x,i}$ is dependent on the volume V, the number of atoms $n_{i,V}$ of isotope i in volume V, the microscopic cross section $\sigma_{x,i}$, and the neutron flux $\phi(V,...)$ integrated over volume V. Therefore the reaction rate can be specified with the dependencies shown in Equation 6.8.

$$R_{x,i,V}\left(n_{i,V}, V, \sigma_{x,i}\right) = \frac{n_{i,V}}{V} \int_V dV' \sigma_{x,i} \phi(\vec{r}) \tag{6.8}$$

In addition to uncertainty in $\sigma_{x,i}$, uncertainty in the volume V and number of atoms of isotope i in volume V affects the calculation of the reaction rate $R_{x,i}$ and therefore the neutron multiplication factor. These two sources of uncertainty can be separated into two distinct classes: uncertainty due to cross section uncertainty, and uncertainty due to material uncertainty.

6.1 Sources of Material Uncertainty

In this study, the material uncertainty is characterized by uncertainty in the values of five quantities:

1. Fuel density
2. Fuel enrichment
3. Pellet radius
4. Gap thickness
5. Clad thickness

If these five quantities are all specified exactly and with no uncertainty, two independent parties could run a pincell calculation and would obtain the exact same neutron multiplication factor

Table 6.1 Material Parameter Uncertainties for the PWR [Iva13].

Parameter	Mean Value	Standard Deviation
Fuel Density	10.283 g/cm^3	± 0.0567 g/cm^3
Pellet Diameter	9.391 mm	± 0.00433 mm
Gap Thickness	0.0955 mm	± 0.008 mm
Clad Thickness	0.673 mm	± 0.00833 mm
Fuel Enrichment	4.85 w/o	± 0.000747 w/o

with no uncertainty. The five quantities listed can also be separated into distinct categories; for instance, items 1 and 2 correspond to material composition, while items 3, 4, and 5 correspond to the dimensions of the material. Another classification set could be that items 1, 2, and 3 correspond to the pellet, while items 4 and 5 correspond to the structure surrounding the pellet. Therefore the same analysis that was performed with regard to analyzing uncertainty due to cross sections can be performed here. To analyze uncertainty due to cross section uncertainty, the five material parameters were fixed to their mean values and the cross sections were allowed to vary. To analyze uncertainty due to uncertainty in the pellet manufacturing process, the cross sections and clad thickness could be fixed to their mean values while varying the fuel density, fuel enrichment, and pellet radius.

Table 6.1 shows the mean values and standard deviations of the five material parameters described above for the LWR. This data was taken from the UAM benchmark document [Iva13]. All five parameters were assumed to be normally distributed with the standard deviations listed in Table 6.1.

To this point, a large number of problem configurations have been introduced. In addition to the fourteen benchmark problem variants described in Chapter 3, a number of Watts Bar cases were analyzed in Chapter 4. Analyzing the uncertainty due to material parameter uncertainty for every case introduced would be extremely difficult, especially considering the many different possibilities for implementing the material uncertainty described earlier. For this reason, a few representative cases were analyzed for the uncertainty due to material uncertainty. These cases are primarily the PWR pincell at Hot Zero Power (case PPZ) and the Unrodded PWR lattice at Hot Zero Power (case PLUZ).

6.2 PWR Pincell

One of the first activities to quantify the uncertainty in k-eff due to these material parameters was to observe the uncertainty produced by varying one parameter individually, and fixing the others, including cross sections, to their mean values. Table 6.2 shows the standard deviation in k-eff observed when varying the individual parameter individually for the simplest case, the PWR

Table 6.2 Material Parameter Uncertainties for the PWR Pincell at HZP.

Varied Parameter	N	SD [pcm]	RSD [%]
Fuel Density	500	59	0.0413
Pellet Diameter	500	29	0.0203
Gap Thickness	500	44	0.0308
Clad Thickness	500	67	0.0469
Fuel Enrichment	500	2	0.00140
All Materials	500	101	0.0708
Cross Sections	500	722	0.506

pincell at Hot Zero Power.

The standard deviation for each of the parameters in Table 6.2 provides an indication of the amount of uncertainty introduced by each parameter individually. In other words, it is a useful metric for identifying the primary contributors to the overall uncertainty in k-eff. For instance, according to this study, the uncertainty in the clad thickness contributes much more to the overall uncertainty in k-eff than the uncertainty in the fuel enrichment does. The uncertainty in the clad thickness produced a k-eff distribution with a standard deviation of 67 pcm, while the uncertainty in the fuel enrichment produced a k-eff distribution with a standard deviation of only 2 pcm. It is important to note that this is entirely dependent on the initial material parameter distributions. It is not valid to conclude that uncertainty in the clad thickness is much more impactful than the uncertainty in the fuel enrichment. With a larger uncertainty in fuel enrichment, it is entirely possible for the uncertainty in fuel enrichment to contribute more to the overall k-eff uncertainty than the cladding thickness. This study is only valid for the material uncertainty parameters listed in Table 6.1. In the fuel manufacturing process, the uncertainty in material parameters is influenced by the manufacturing tolerances which are specified by each individual fuel manufacturer.

The study that produced the results in Table 6.2 could be performed by fuel manufacturers to identify the parameters that produce the highest uncertainty. If large uncertainty in a certain parameter is a result of a poor manufacturing process rather than a poor measurement, then that particular process could be optimized to lower the overall uncertainty.

For this PWR pincell case at HZP, the uncertainty in k-eff due to uncertainty in all of the material parameters was 101 pcm. This was obtained by sampling out of every material parameter distribution simultaneously, so this is the combined effect of the uncertainty of the material parameters. While this 101 pcm standard deviation is much lower than the 722 pcm standard deviation obtained by accounting for only cross section uncertainty, neglecting the material uncertainty may be an unacceptable simplification for applications in which a very high degree of precision is desired.

Figure 6.1 shows the kernel density estimate for the k-eff distribution obtained when varying each of the five material parameters individually. This is simply a visual representation of the data

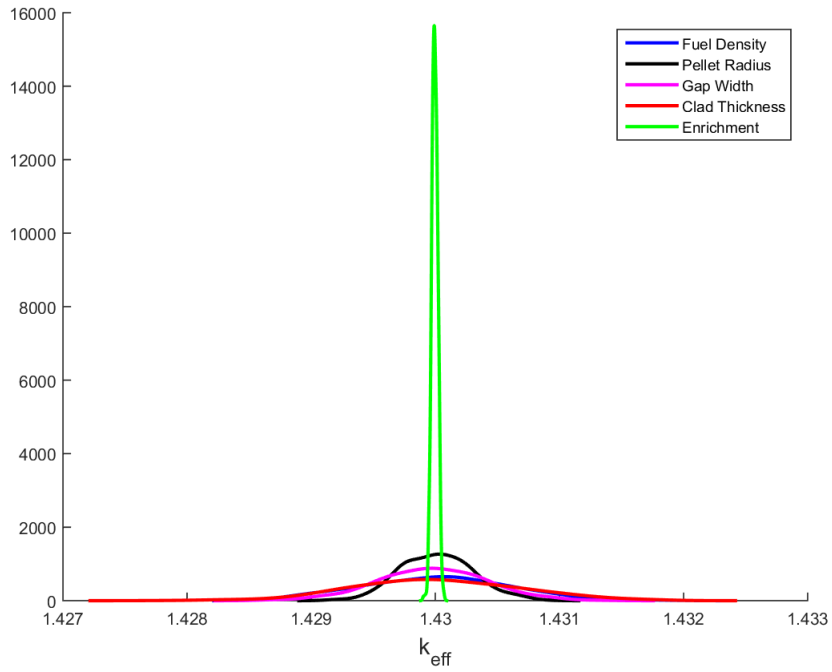

Figure 6.1 Kernel density estimate of the k-eff distributions when varying each material parameter individually.

shown in Table 6.2. The low uncertainty in the fuel enrichment is captured by the low uncertainty in the resulting k-eff distribution, while the uncertainty of the other four parameters is very similar in magnitude. To illustrate the earlier point about these results being highly dependent on the uncertainty of the individual parameters, consider what the k-eff distributions would look like for a fuel manufacturer who had a less-strict manufacturing tolerance on the fuel enrichment. From the UAM LWR benchmark, the standard deviation of the fuel enrichment is 0.000747 w/o, as shown in Table 6.1. If the fuel enrichment uncertainty were ten times higher, with a standard deviation of 0.00747 w/o, the resulting k-eff standard deviation would be approximately 24 pcm. Using this uncertainty, the k-eff distributions for all five parameters are much more similar in magnitude, as shown in Figure 6.2.

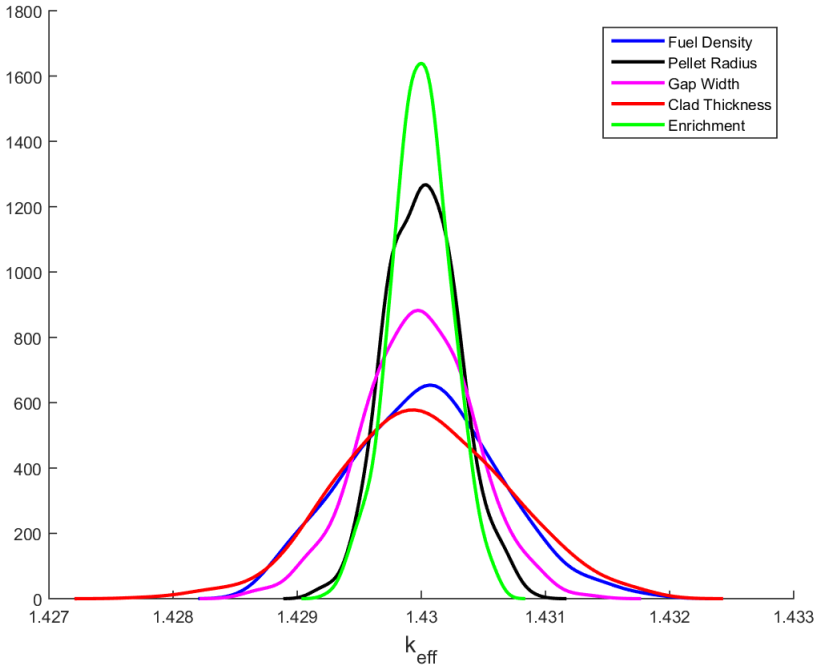

Figure 6.2 Kernel density estimate of the k-eff distributions when varying each material parameter individually, 10x fuel enrichment uncertainty.

6.3 PWR Lattice

More complexity is introduced when analyzing the effect of material uncertainty on lattices. Recall that fuel lattices are essentially a series of pincells arranged in a square, and that all of the material uncertainties given in the benchmark are pin-level uncertainties. As a result, there are a number of ways to model material uncertainty in fuel lattices. Consider three potential options:

1. Construct n lattices by sampling out of each material parameter distribution n times and using the sampled quantities for every pin in a lattice.

2. Construct n lattices by sampling out of each material parameter distribution $n * n_{\text{pin}}$ times, and using the unique pin samples in each lattice.

3. Construct n lattices by sampling out of the distributions of a subset of the material parameters p_p n times, and sampling out of the remaining $p - p_p$ parameter distributions $n * n_{\text{pin}}$ times.

The first method can be considered an assembly-level variation. If the first method is employed, we are assuming that there is no variation in any material parameters in a given assembly, and the only variation in material parameters occurs in different assemblies. In other words, it is assumed that, in a lattice with 208 fuel pins, the fuel pellet radius, the gap width, the clad thickness, the fuel density, and the fuel enrichment are all identical. The second method can be considered a

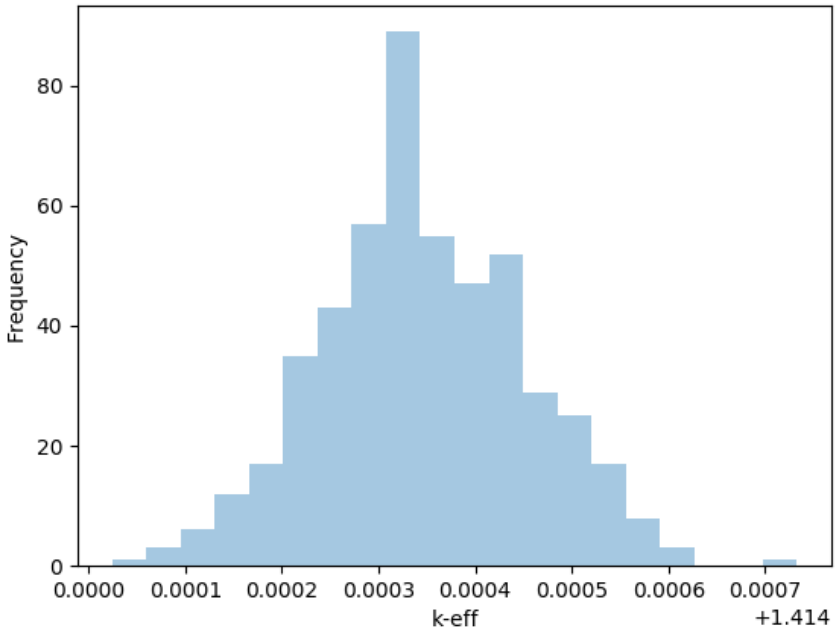

Figure 6.3 k-eff distribution for PWR Lattice Unrodded at HZP, varying all material parameters on a pin-level.

pin-level variation. In the second method, it is assumed that all of the materials are manufactured independently. The third method is a combination of the two methods, where some parameters are varied on a pin-level and others are varied only on an assembly-level.

Shown in Figure 6.3 is the distribution of k-eff for the unrodded PWR lattice at HZP when applying the pin-level variation strategy, the second option from the list, for all five material parameters. In reality, both the material parameters and cross sections have uncertainty. Figure 6.4 shows the k-eff distribution for the unrodded PWR lattice at HZP when varying both the material parameters on a pin-level and the cross sections simultaneously.

Of the three scenarios presented earlier, the one most likely to be encountered in reality is the third option – that some material parameters will vary on a pin-by-pin basis, and some will vary on an assembly-by-assembly basis. The material parameter perturbation tool that was developed and described in Chapter 2 has the capability to handle both simultaneously. Consider the scenario described earlier in this chapter, where fuel pellets are manufactured in batches. Assuming the fuel is well-mixed in a batch before it is pressed into individual pellets, the density and enrichment of the fuel will be consistent throughout a batch. Figure 6.5 shows the k-eff distribution when the fuel pellet radius, gap thickness, and clad thickness are varied on a pin-level basis, and the fuel density and enrichment are varied on an assembly-level basis. Figure 6.6 shows the same scenario, where the fuel composition parameters are varied on an assembly-level basis, but simultaneously varying the cross sections.

A summary of the results from the lattice study is shown in Table 6.3. For the first case, labeled "XS", only cross section uncertainties are considered, and material uncertainties in the lattice are

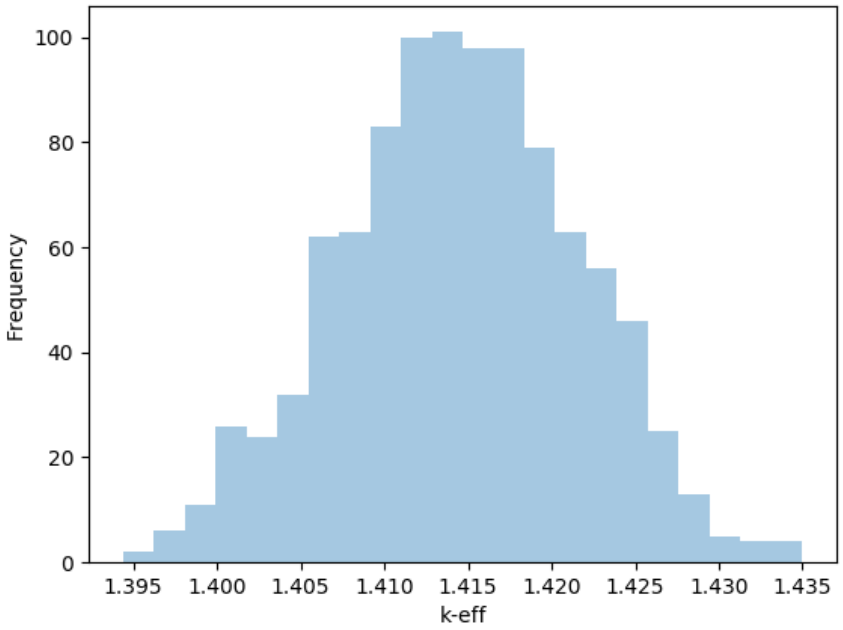

Figure 6.4 k-eff distribution for PWR Lattice Unrodded at HZP, varying all material parameters on a pin-level and varying cross sections.

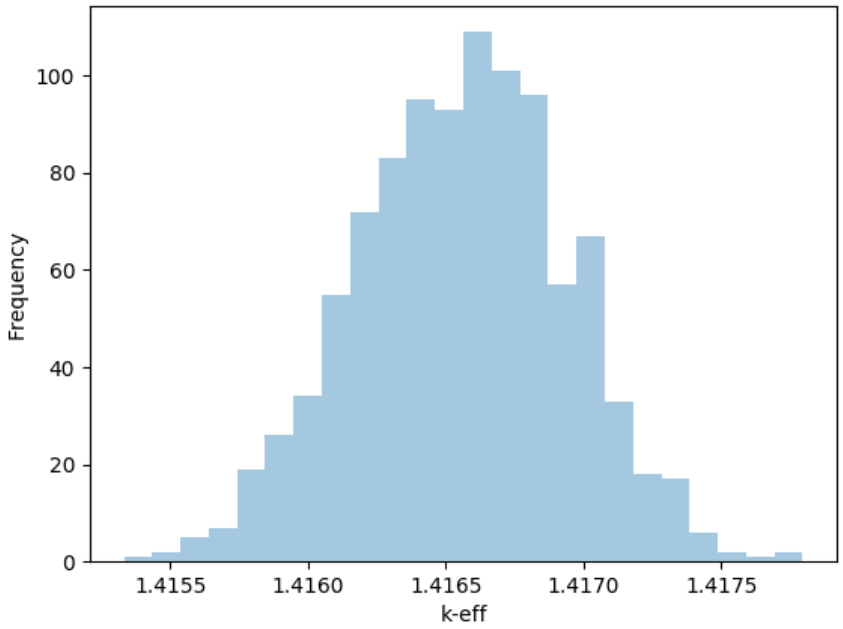

Figure 6.5 k-eff distribution for PWR Lattice Unrodded at HZP, varying fuel composition parameters on an assembly-level.

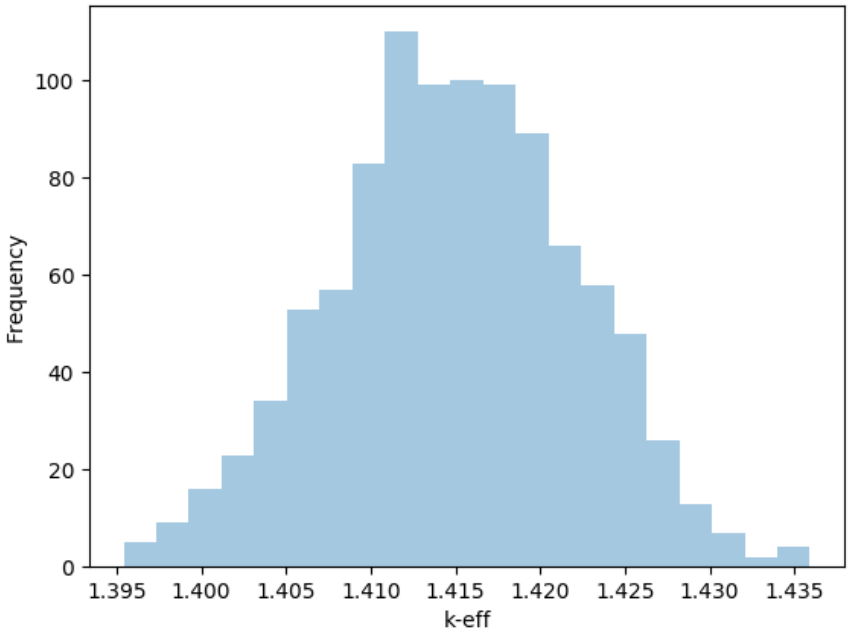

Figure 6.6 k-eff distribution for PWR Lattice Unrodded at HZP, varying fuel composition parameters on an assembly-level with cross sections.

completely neglected. In second case, labeled “Pin-level mat.”, it is assumed that all material parameters are varied on a pin-level, and there is no uncertainty in the cross sections. The third option is the same as the second, but also accounts for the cross section uncertainties. Note that the second and third options correspond to the distributions shown in Figures 6.3 and 6.4, respectively. The fourth case, named “Mixed mat.”, corresponds to the assumption that the fuel composition parameters are assembly-level variations, while the pellet dimension parameters are pin-level variations. The final case is the same as the fourth, but also accounts for cross section uncertainties. These final two cases correspond to Figures 6.5 and 6.6, respectively.

When all material parameters are varied at a pin-level, the relative uncertainty in k-eff due to material parameter uncertainty of 0.00749 % is dwarfed by the relative cross section uncertainty of 0.502 %. The relative standard deviation from the material parameter perturbations is 67 times smaller in magnitude than the relative standard deviation from the cross section perturbations. When the fuel composition parameters are modeled as assembly-level perturbations instead of pin-level, the relative standard deviation of k-eff due to material parameter uncertainty increases to 0.0268 %. This is only approximately 18 times smaller than the relative standard deviation from the cross section perturbations. Still, in both cases, the total relative uncertainty is dominated by the cross section uncertainty.

Table 6.3 Material Parameter Uncertainties for the unrodded PWR lattice at HZP.

Case	N	k-eff	SD [pcm]	RSD [%]
XS	1000	1.414477	710	0.502
Pin-level mat.	1000	1.414347	10.6	0.00749
Pin-level mat. & XS	1000	1.414473	709	0.501
Mixed mat.	1000	1.416559	38.0	0.0268
Mixed mat. & XS	1000	1.414924	711	0.503

CHAPTER

7

CONCLUSIONS

From the results reported in Chapters 3-5, it was confirmed that stochastic sampling is a valid and feasible method for determining the uncertainty in k-eff due to uncertainty in cross sections. The close agreement of the results obtained for the UAM LWR benchmark with the benchmark average results, and the agreement between the UAM LWR benchmark results and the more complex Watts Bar models provides strong evidence to the effectiveness of the stochastic sampling method in this application.

For the UAM LWR benchmark PWR models, based on the dimensions of the TMI Unit 1 plant, the relative standard deviation of k-eff due to cross section uncertainty ranged from 0.500 % for the colorset at Hot Full Power to 0.534 % for the rodded lattice at Hot Full Power. All eight PWR cases fell within this range of 0.034 %, which corresponds to a range of approximately 40 pcm when adjusting by the average k-eff of the disparate cases. For the UAM LWR benchmark BWR models, which were based on the Peach Bottom Unit 2 plant, the relative standard deviation of k-eff due to cross section uncertainty ranged from 0.550 % for the rodded lattice at Hot Zero Power to 0.607 % for the pincell at Hot Full Power. The six cases all fell within the range of 0.057 %, which corresponds to an adjusted range of approximately 60 pcm.

While the relative standard deviations of the PWR and BWR models were similar in magnitude, one trend stood out between the two datasets. The relative standard deviation estimates for PWR cases at HFP were larger by approximately 0.004 % than the estimates for the same cases at HZP. However, the relative standard deviation estimates for the BWR cases at HFP were larger by approximately 0.04 %. If this trend held for other PWR and BWR models, it could be concluded that the state parameters of the model is more important in cross section uncertainty quantification in BWRs than PWRs. Cross section uncertainty results for more PWR and BWR models would needed

to definitively make this claim. It is likely that this dependence on the state of the reactor is due to the hardening of the neutron spectrum from the decreased moderator density in the BWR.

The Watts Bar Unit 1 full-core models produced relative standard deviations of 0.553 %, 0.552 %, and 0.556 % for the two-dimensional, three-dimensional, and three-dimensional with thermal hydraulic feedback cases respectively. These values are consistent with the range of relative standard deviations obtained from the UAM LWR models. There is remarkable consistency between the relative standard deviations for these models, as the two-dimensional full-core model and three-dimensional full-core model only differ by 0.003 %. The small difference in the results compared with the large difference in computational cost between the cases provides evidence to the possibility of using surrogate models to model more complex cases, as will be addressed in the following section.

A major topic of investigation in this study was whether 1000 samples was sufficient to generate high-quality statistics for uncertainty quantification. The estimation of the standard error of the standard deviation estimate, which was performed by jackknife re-sampling of the datasets in Chapter 5, showed that for every case that was examined, the relative standard error of the standard deviation estimate was approximately 2 %. The low relative standard error and the consistency of the relative standard error across different cases indicates that 1000 libraries should be more than sufficient to estimate uncertainty due to cross section uncertainty for any LWR model. In fact, the relative standard error of the standard deviation estimate was less than 5% for all models with only 250 cross section libraries. For applications where the use of a surrogate model is non-viable, and a computationally expensive model must be used, 250 perturbed cross section models should produce a standard deviation estimate with a relative standard error less than 5%, which may be an acceptable level of error in certain applications. In any application of the results of this study, an acceptable level of error must be determined and weighed against the computational cost that would be incurred by reducing this level of error.

Uncertainty in k-eff due to material uncertainty was also studied. Using the UAM PWR as a representative case, and material uncertainties given in the UAM LWR benchmark as representative parameter uncertainties, it was determined that uncertainty in k-eff due to material parameter uncertainties is significantly lower than due to cross section uncertainty. From the cross section uncertainty results reported in Chapter 3 and the material uncertainty results reported in Chapter 6, the relative standard deviation for cross section uncertainty for the PWR pincell at HZP was 0.511 %, and the relative standard deviation for the material uncertainties for the same model was 0.0708 %. It is likely that even this is an overestimation of the uncertainty introduced by material uncertainties, as the relative standard deviation for the "realistic" material uncertainty model on the assembly-level, in which fuel parameters were fixed across the assembly and the other parameters were varied pin-by-pin, was only 0.0268 %.

In applications where a very high degree of statistical confidence is necessary, the material parameter uncertainties should be accounted for. However, in less strict applications, and particularly in large-scale applications where a large amount of pin-by-pin averaging occurs naturally, the vast majority of the uncertainty in the system could be accounted for by only considering the cross

section uncertainties.

One conclusion that can certainly be drawn from this study is the validity of VERA as a model for propagation of uncertainty using stochastic sampling. From the comparison of the VERA stochastic sampling results to the benchmark average results in Chapter 3, our results are very similar to results generated by other benchmark participants using different uncertainty quantification techniques. The application of VERA stochastic sampling to the full-core Watts Bar Unit 1 cases in Chapter 4 demonstrated that this method could be scaled up to very large problems, but not without considerable computational expense.

7.1 Suggestion for Future Work

One potential use of the data presented in this study could be to determine whether a simplified model could serve as a surrogate for quantifying uncertainty due to cross section uncertainty for a more complex model. The motivation for this can clearly be seen from the profiling data in Tables 2.1 and 2.2. While a 1000-sample study of a PWR pincell took less than one core-hour, a 1000-sample study of a full-core PWR with no thermal hydraulic feedback took over 340,000 core-hours. A relative standard deviation of 0.511 % was obtained for the pincell, and a relative standard deviation of 0.552 % was obtained for the full-core. Is a difference of 0.041 % in the quantification of cross section uncertainty worth a 340,000-fold increase in computational cost?

The pincell and full-core examples are based on two different PWR configurations. The pincell model is taken from the TMI-1 specification in the UAM LWR benchmark [Iva13], while the full-core is a model of the Watts Bar plant [God14]. An argument could be made that the comparison of different models, even though the results are very similar, is a bridge too far for simplification. In this case, comparisons can simply be made between the simplest and most complex cases in a certain model. For this comparison, no heuristic knowledge will be known about the individual cases: the simplest case will be defined as the case with the smallest computational cost, and the most complex will be the case with the largest. This means that no distinction will be made between models run at HZP or HFP. The three models that were examined in this study were the TMI-1 PWR [Iva13], the PB-2 BWR [Iva13], and the Watts Bar PWR [God14].

The percent difference in relative standard deviation compared to the percent difference in computational cost is a useful metric in this case. For the TMI-1 case from the UAM LWR benchmark, increasing the complexity from a pincell at Hot Zero Power to a colorset at Hot Full Power resulted in a change in the relative standard deviation of -1.76 %, and a change in the computational cost of +10,200 %. For the PB-2 BWR case from the UAM LWR benchmark, increasing the complexity from a pincell at Hot Full Power to a rodded lattice at Hot Zero Power resulted in a change in the relative standard deviation of -9.57 %, and a change in the computational cost of +3,390 %.

One strategy could be to run a limited number of samples for a complex and simplified model, establish the similarity in convergence between the models for the limited number of samples, and then run the large number of samples on the simplified model. The limiting factor of this strategy

Table 7.1 Comparison between simplest and most complex case for each model.

Case	RSD [%]	Comp. Cost [CPU-hr]
PPZ	0.511	0.848
PCF	0.502	52.8
BPF	0.606	0.683
BLRZ	0.548	21.2
WB2D	0.553	2,667
WB3D-TH	0.556	552,462

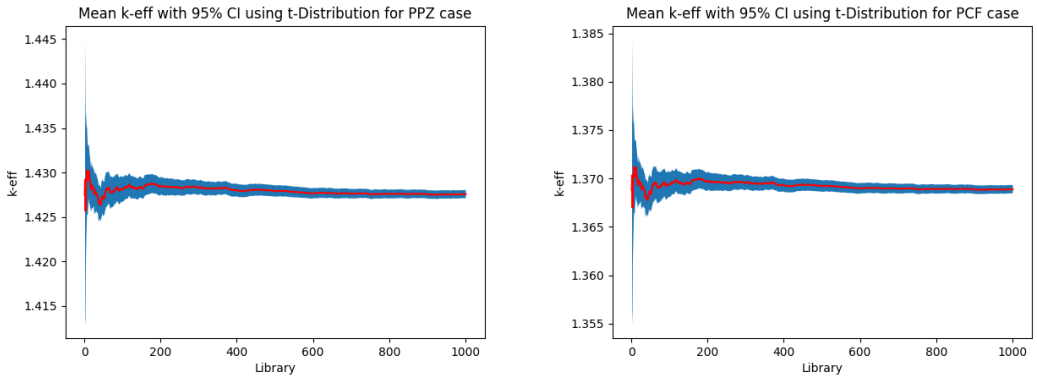

Figure 7.1 Confidence intervals for PWR Pincell and PWR Colorset models.

would be determining the number of samples necessary to establish similarity between the models. An initial check for similarity could be a simple visual comparison of the confidence intervals for k-eff for each of the models.

For instance, consider the confidence intervals for the simplest UAM PWR case (PPZ) and the largest (PCF). From a purely visual perspective, there is no discernible difference in the running average k-eff, shown in red in Figures 7.1 and 7.2 and in the 95 % confidence interval, shown in blue in the same figures, between the two cases.

A numerically-based similarity analysis could also easily be carried out. Two statistics that are independent of the model are the relative standard deviation and the relative standard error of the standard deviation estimate. The relative standard deviation is the standard deviation normalized by the mean k-eff, so it is independent of the magnitude of k-eff, and the relative standard error of the standard deviation estimate is normalized by the standard deviation estimate, so it is independent of the magnitude of the standard deviation estimate. The relative standard deviation provides an estimate of the total uncertainty, and the relative standard error of the standard deviation provides an estimate of the convergence of this total uncertainty to the actual uncertainty.

Table 7.2 shows the relative standard deviation and relative standard error of the standard

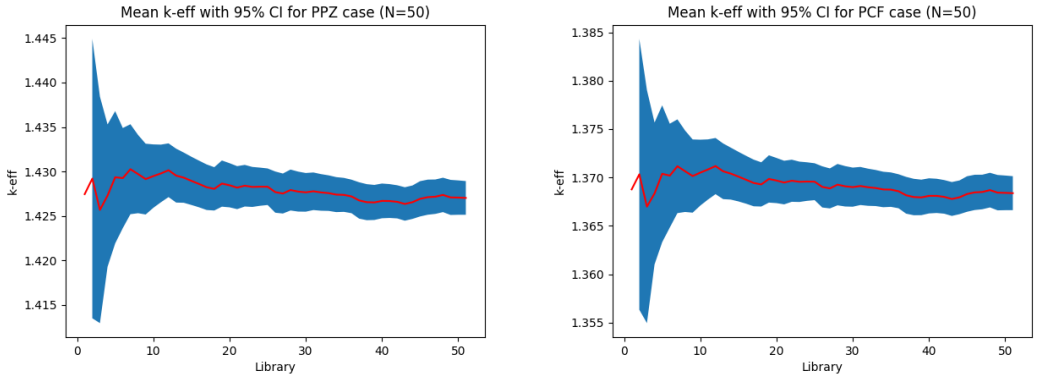

Figure 7.2 Confidence intervals through 50 samples for PWR Pincell and PWR Colorset models.

Table 7.2 Comparison between simplest and most complex case for each model.

	PPZ		PCF		Difference	
n	RSD [%]	RSE [%]	RSD [%]	RSE [%]	RSD [%]	RSE [%]
5	0.490	43.0	0.485	43.1	-0.005	+0.1
10	0.396	25.0	0.393	24.8	-0.002	-0.2
15	0.378	15.9	0.373	16.3	-0.005	+0.4
20	0.400	13.0	0.392	13.5	-0.008	+0.5
25	0.373	12.4	0.362	13.1	-0.011	+0.7
30	0.429	13.7	0.416	13.5	-0.013	-0.2
35	0.408	13.0	0.397	12.9	-0.011	-0.1
40	0.444	12.6	0.431	12.5	-0.013	-0.1
45	0.478	11.6	0.461	11.4	-0.017	-0.2
50	0.490	10.4	0.474	10.2	-0.016	-0.2

deviation for the PWR pincell at HZP and the PWR colorset at HFP. Even with only fifty samples for comparison, the difference between the relative standard deviations for the two cases is minimal, with a maximum difference of 0.017 % at 45 samples. The small difference between the relative standard errors shows the similar error reduction behavior between the cases.

Figure 7.3 shows the relative standard deviation as a function of sample size for the PWR pincell at HZP and the PWR colorset at HFP, and Figure 7.4 shows the difference between the two values, defined by Equation 7.1, at each sample size.

$$\text{Diff.} = \text{RSD}_{\text{PCF}} - \text{RSD}_{\text{PPZ}} \tag{7.1}$$

From Figure 7.4, information about the maximum difference between the cross section uncertainties and the behavior of this difference can be studied. The maximum absolute difference between the two cases is -0.0200 at $n = 71$ samples. After the maximum difference at $n = 71$ samples,

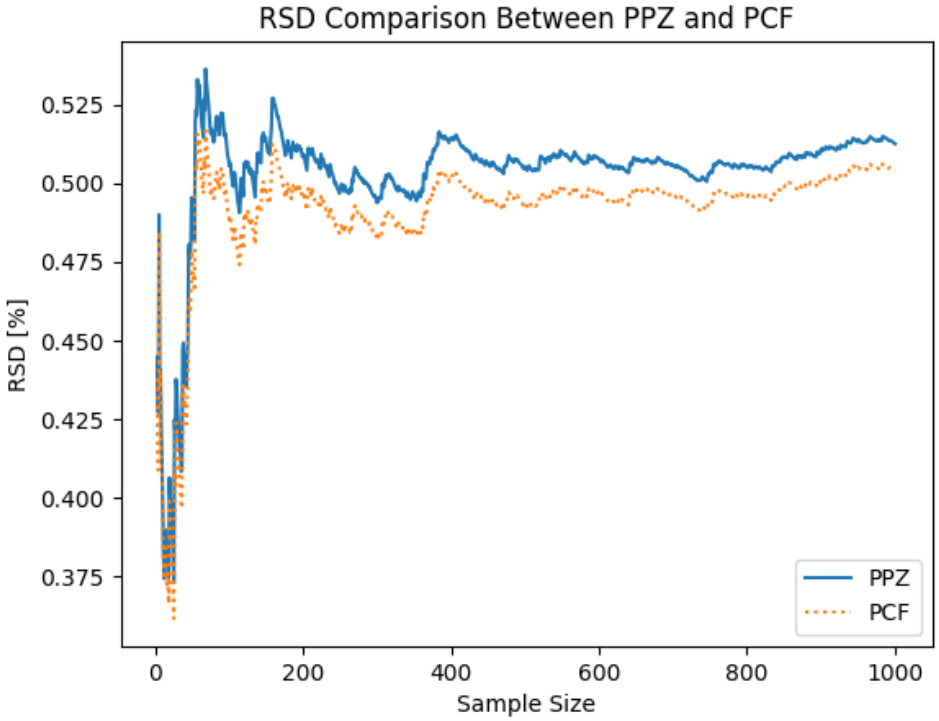

Figure 7.3 Relative standard deviations for the PWR pincell at HZP and the PWR colorset at HFP.

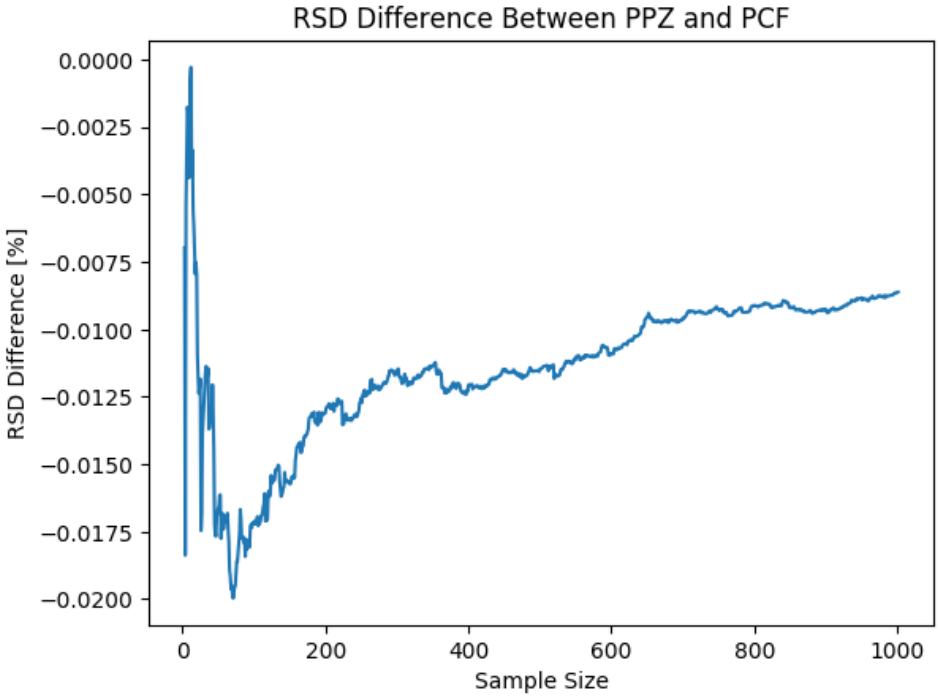

Figure 7.4 Relative standard deviation difference between the PWR pincell at HZP and the PWR colorset at HFP.

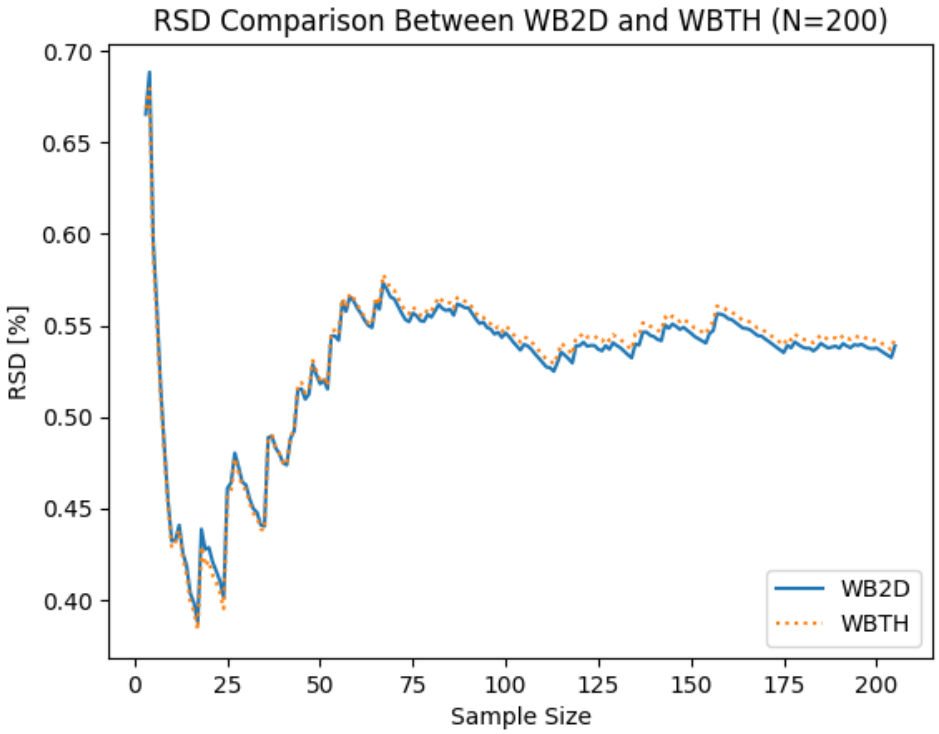

Figure 7.5 Relative standard deviations for the 2D Watts Bar model and the 3D Watts Bar full- core with thermal hydraulic feedback, N=200.

the absolute difference between the relative standard deviations for the simple and complex case gradually reduces.

Using the results of this study to develop a strategy to deal with the Watts Bar case, one may decide to use 200 samples to establish similarity between a simple and complex case, and then use 1000 samples with the simple case as a surrogate for the complex case. The choice of 200 samples is solely based on Figure 7.4, as the maximum absolute difference is reached within 100 samples and a gradual decline in maximum absolute difference is reached by 200 samples.

Figure 7.5 shows the relative standard deviations for the 2D model and the 3D model with thermal hydraulic feedback. Only the first 200 samples are shown for each model. The two models produced nearly identical relative standard deviations through 200 samples, so in this case it would be assumed that the 2D model is a valid surrogate for the 3D model with thermal hydraulic feedback. Figure 7.6 shows the difference between the relative standard deviations for the two models at each sample size. Just as in the PWR benchmark, the initial instability in the difference is diminished by approximately 200 libraries, and the difference gradually tends towards zero.

If we had accepted Figure 7.5 as evidence that the 2D model is a valid surrogate for uncertainty quantification, computational cost could greatly be reduced. The cost of 200 cases for the 3D model with thermal hydraulic feedback is approximately 110,492 CPU-hr. The cost of 200 cases for the simplified 2D model is approximately 533 CPU-hr, bringing the total computational cost for

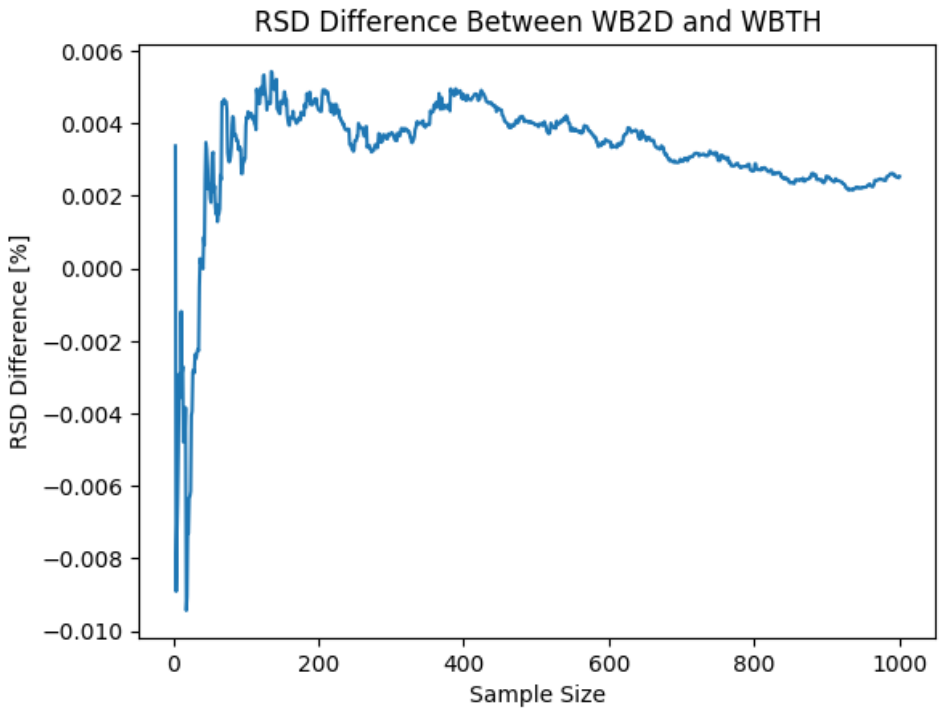

Figure 7.6 Relative standard deviations for the 2D Watts Bar model and the 3D Watts Bar full- core with thermal hydraulic feedback.

establishment of the 2D model as a surrogate to 111,025 CPU-hr. The cost of running the additional 800 cases with the 2D model is only 2,134 CPU-hr, so the total computational cost of this study would be 113,159 CPU-hr, compared to the 552,462 CPU-hr used for 1000 cases of the 3D model with thermal hydraulic feedback.

This surrogate method could be used to reduce computational cost and to increase the quality of the statistical estimates. The computational cost of 800 3D cases with thermal hydraulic feedback is approximately 441,970 CPU-hr. If the computational expense used to run 800 additional cases of the 3D model with thermal hydraulic feedback had instead been used for the 2D model, which only uses 2.7 CPU-hr per evaluation, approximately 165,000 additional 2D cases could have been evaluated.

BIBLIOGRAPHY

[Bra14] Bratton, R. N. et al. "OECD/NEA Benchmark for Uncertainty Analysis in Modeling (UAM) for LWRs - Summary and Discussion of Neutronics Cases (Phase I)". *Nuclear Engineering and Technology* **46**.3 (2014), pp. 313–342.

[CAS18] CASL. *CASL Annual Report*. Tech. rep. Consortium for the Advanced Simulation of Light Water Reactors, 2018.

[Cha11] Chadwick, M. B. et al. "ENDF/B-VII.1 Nuclear Data for Science and Technology: Cross Sections, Covariances, Fission Product Yields and Decay Data". *Nuclear Data Sheets* **112**.12 (2011), pp. 2887–2996.

[Col16] Collins, B. et al. *MPACT Theory Manual*. 2.2.0. Oak Ridge National Laboratory. Oak Ridge, TN, 2016.

[Dud76] Duderstadt, J. J. & Hamilton, L. J. *Nuclear Reactor Analysis*. John Wiley & Sons, 1976.

[Efr81a] Efron, B. & Stein, C. "The Jackknife Estimate of Variance". *The Annals of Statistics* **9**.3 (1981), pp. 586–596.

[Efr81b] Efron, B. "Nonparametric Estimates of Standard Error: The Jackknife, the Bootstrap, and Other Methods". *Biometrika* **68**.3 (1981), pp. 589–599.

[God14] Godfrey, A. T. *VERA Core Physics Benchmark Progression Problem Specifications*. 4th ed. Oak Ridge National Laboratory. 2014.

[Gol11] Goluoglu, S. et al. "Monte Carlo Criticality Methods and Analysis Capabilities in SCALE". *Nuclear Technology* **174**.2 (2011), pp. 214–235.

[Hol01] Holden, N. E. "A Short History of Nuclear Data and its Evaluation". *A CSEWG Retrospective*. Cross Section Evaluation Working Group, Brookhaven National Laboratory. 2001.

[Hou19] Hou, J. et al. "Comparative Analysis of Solutions of Neutronics Exercises of the LWR UAM Benchmark". 2019.

[Ida] Idaho National Laboratory. "Versatile Test Reactor, Frequently Asked Questions". Accessed online, April, 2020.

[Iva13] Ivanov, K. et al. *Benchmarks for Uncertainty Analysis in Modelling (UAM) For the Design, Operation and Safety Analysis of LWRs*. Tech. rep. OECD Nuclear Energy Agency, 2013.

[Koc17] Kochunas, B. et al. "VERA Core Simulator Methodology for Pressurized Water Reactor Cycle Depletion". *Nuclear Science and Engineering* **185**.1 (2017), pp. 217–231.

[Kok09] Kok, K. D. *Nuclear Engineering Handbook*. 1st ed. Boca Raton, FL: CRC Press, 2009.

[Lew93] Lewis, E. E. & Miller Jr., W. F. *Computational Methods of Neutron Transport*. Lagrange Park, IL: American Nuclear Society, 1993.

[OEC09] OECD/NEA Data Bank. *The JEFF-3.1.1 Nuclear Data Library*. Tech. rep. JEFF Report 22. OECD/NEA Data Bank, 2009.

[Rao73] Rao, C. R. "Linear Statistical Inference and its Applications". 2nd ed. New York City, NY: John Wiley & Sons, 1973, p. 438.

[Rea16] Rearden, B. T. et al. *SCALE Code System*. 6.2. Available from Radiation Safety Information Computational Center as CCC-834. Oak Ridge National Laboratory. Oak Ridge, Tennessee, 2016.

[Roy05] Roy, C. J. "Review of Code and Solution Verification Procedures for Computational Simulation". *Journal of Computational Physics* **205**.1 (2005), pp. 131–156.

[Sal16] Salko, R. K. & Avramova, M. *CTF Theory Manual*. 0th ed. Consortium for the Advanced Simulation of Light Water Reactors. 2016.

[Seg53] Segre, E. et al. "Experimental Nuclear Physics". 1st ed. Vol. 2. New York City, NY: John Wiley & Sons, 1953, p. 208.

[Str95] Strauss, S. D. "Nuclear Cost Control Focuses on Refueling Outages". *Power* **139**.12 (1995), pp. 39–41.

[Trk18] Trkov, A. et al. *ENDF-6 Formats Manual*. Appendix I: Historic Perspective by Norman E. Holden. Brookhaven National Laboratory. 2018.

[Weia] Weisstein, E. W. *Bessel's Correction. From MathWorld—A Wolfram Web Resource*. Accessed May 26, 2020.

[Weib] Weisstein, E. W. *Student's t-Distribution. From MathWorld—A Wolfram Web Resource*. Accessed April 23, 2020.

[Wil13] Williams, M. L. et al. "A Statistical Sampling Method for Uncertainty Analysis with SCALE and XSUSA". *Nuclear Technology* **183**.3 (2013), pp. 515–526.

[Wim01] Wimmer, L. B. *Summary Report of Commercial Reactor Criticality Data for Three Mile Island Unit 1*. Tech. rep. Bechtel SAIC Company LLC., 2001.

APPENDIX

APPENDIX

A

CONFIDENCE INTERVALS

A.1 Theory

Confidence intervals are an often-used tool for representing the error introduced to our estimators by limiting the sample size. A confidence interval can be thought of as the range of values that we are confident (with a certain confidence level) contains the true population value. If a confidence level of 95 % is used, then the probability that the true population parameter lies within the range given by the confidence interval is 0.95.

When only estimates for the expected value and standard deviation are available, the t-distribution must be used to construct confidence intervals. The derivation for the t-distribution will be omitted. The probability density function for the t-distribution is given by Equation A.1, where ν refers to the degrees of freedom with $\nu = n-1$, where n is the sample size [Weib].

$$f_\nu(t) = \frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\sqrt{\pi\nu}\Gamma\left(\frac{\nu}{2}\right)\left(1+\frac{t^2}{\nu}\right)^{(\nu+1)/2}} \tag{A.1}$$

The t-statistic is defined in Equation A.2, where n is the sample size, $\hat{x}$ is the sample mean, μ is the true mean, and $\hat{\sigma}$ is the estimated standard deviation.

$$t = \sqrt{n}\frac{\hat{x}-\mu}{\hat{\sigma}} \tag{A.2}$$

As the sample size n increases, it can be shown that the t-distribution approaches the normal distribution.

$$\lim_{n\to\infty} f_{n-1}(t) = N(\mu,\sigma) \tag{A.3}$$

The cumulative distribution function for the t-distribution can be obtained by integrating the probability density function, as shown in Equation A.4. The cumulative distribution function for the t-distribution is monotonically increasing and maps the t-statistic range $(-\infty,\infty)$ onto the full probability space $[0,1]$.

$$F_\nu(t) = \int_{-\infty}^{t} f_\nu(x)dx \tag{A.4}$$

The inverse of the cumulative distribution function, also known as the quantile function or percent point function, is an important function in the calculation of confidence intervals. While the cumulative distribution function returns the probability level at a given t-statistic value, the quantile function returns the t-statistic value at a given probability level. The quantile function $Q_\nu(p)$ for the t-distribution is defined in Equation A.5.

$$t = Q_\nu(p) = F_\nu^{-1} = Q_\nu(F_\nu(t)) \tag{A.5}$$

If for instance we were interested in constructing an interval that contained 97.5 % of the probability space of a t-distribution, we could use the quantile function $Q_\nu(0.975)$ to find the value of t, denoted t^* that would yield 0.975 when inserted into the cumulative distribution function. This relationship is demonstrated in Equations A.6 and A.7.

$$F_\nu(t^*) = \int_{-\infty}^{t^*} f_\nu(x)dx = 0.975 \tag{A.6}$$

$$Q_\nu(0.975) = t^* \tag{A.7}$$

Due to the symmetry of the t-distribution about $t = 0$, the left and right tails of the distribution are identical and the critical value of t that contains 2.5 % of the probability space is simply the negative of the critical value of t that contains 97.5 % of the probability space.

$$Q_\nu(0.025) = -Q_\nu(0.975) = -t^* \tag{A.8}$$

Therefore the interval $[-Q_\nu(0.975), Q_\nu(0.975)] = [-t^*, t^*]$ contains 95 % of the probability space of the t-distribution with ν degrees of freedom. Using the definition of the t-statistic given in Equation A.2, the critical value interval can be transformed into an interval of the sample mean $\hat{x}$. This interval can be expressed as $[\hat{x} - \frac{t^*\hat{\sigma}}{\sqrt{n}}, \hat{x} + \frac{t^*\hat{\sigma}}{\sqrt{n}}]$. This interval is commonly referred to as a confidence interval.

As the sample size increases, the critical value of the t-statistic t^* converges to the critical value of the z-statistic, which for a probability of 0.975 is $t^*_\infty = z^* = 1.960$. The sample standard deviation $\hat{\sigma}$ converges to the true standard deviation σ. The denominator of the confidence interval grows without bound, so the confidence interval converges to zero as the sample size grows.

A.2 Confidence Interval Plots

Confidence interval plots for a 95% confidence interval for the cases examined in Chapters 3 and 4 are presented in Figures A.1 – A.17. Note that while it may appear that the Watts Bar cases have a much smaller confidence interval than the UAM LWR benchmark cases, this is due to the very high uncertainty at the lowest sample size distorting the scale of the plot. From the results presented in Chapters 3 and 4, in Tables 3.8 and 4.3 respectively, the relative uncertainties with a sample size of 1000 cross section libraries for the Watts Bar cases and UAM LWR cases are comparable.

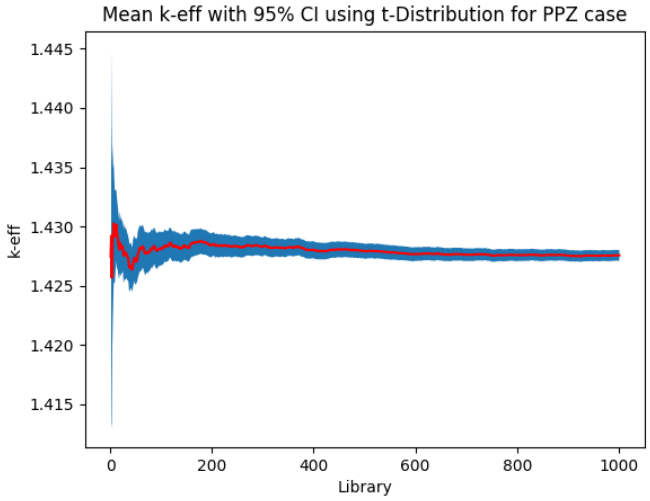

Figure A.1 PWR Pincell at HZP 95 % confidence interval.

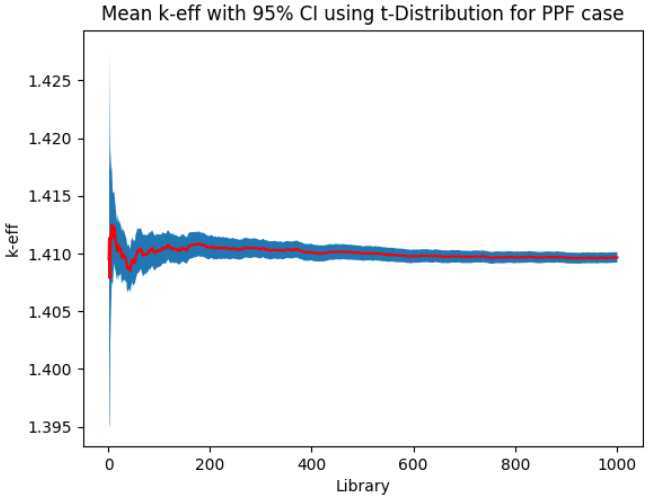

Figure A.2 PWR Pincell at HFP 95 % confidence interval.

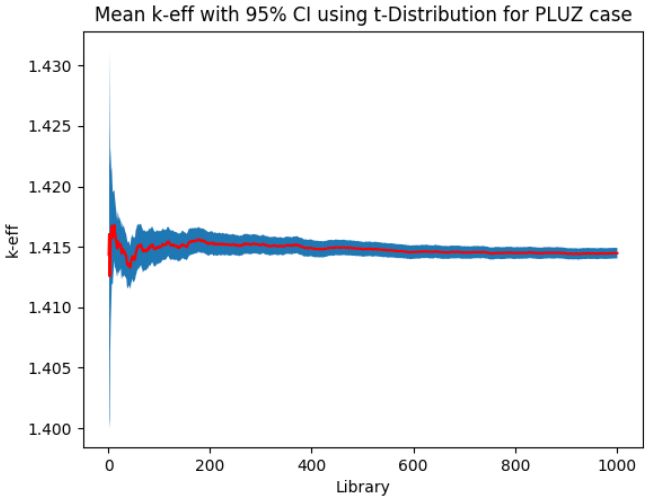

Figure A.3 PWR Lattice Unrodded at HZP 95 % confidence interval.

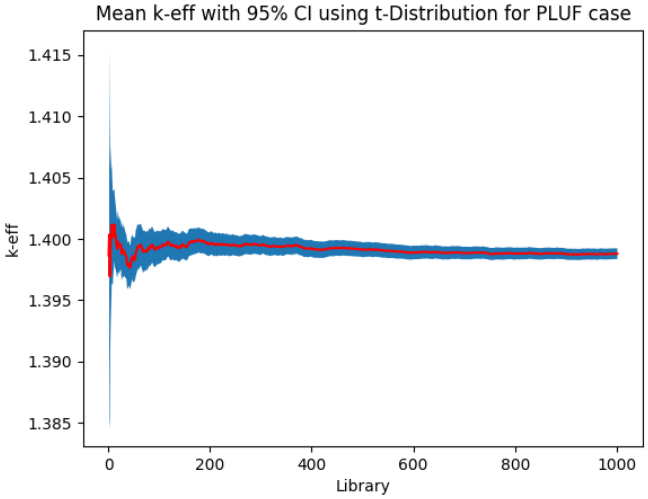

Figure A.4 PWR Lattice Unrodded at HFP 95 % confidence interval.

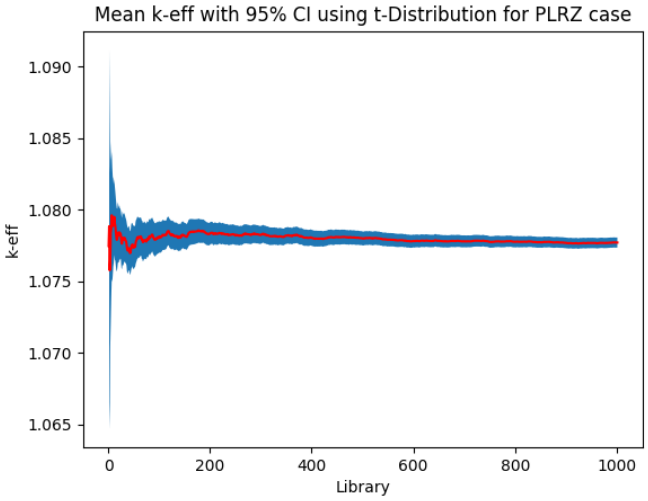

Figure A.5 PWR Lattice Rodded at HZP 95 % confidence interval.

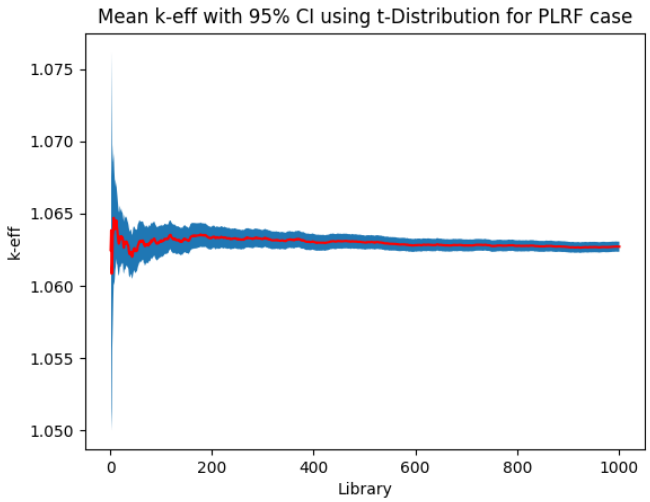

Figure A.6 PWR Lattice Rodded at HFP 95 % confidence interval.

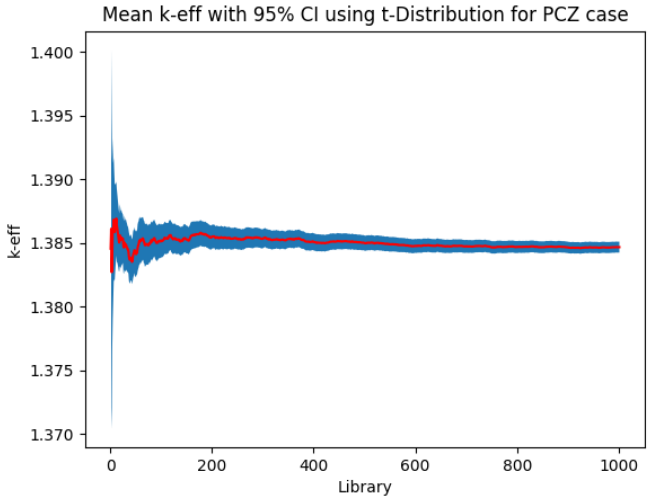

Figure A.7 PWR Colorset at HZP 95 % confidence interval.

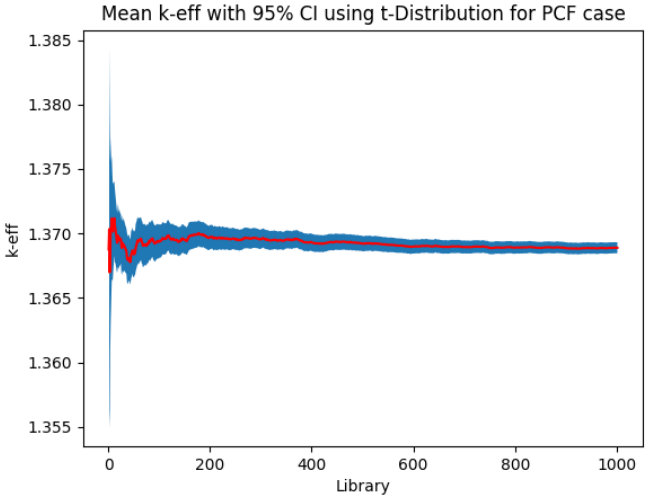

Figure A.8 PWR Colorset at HFP 95 % confidence interval.

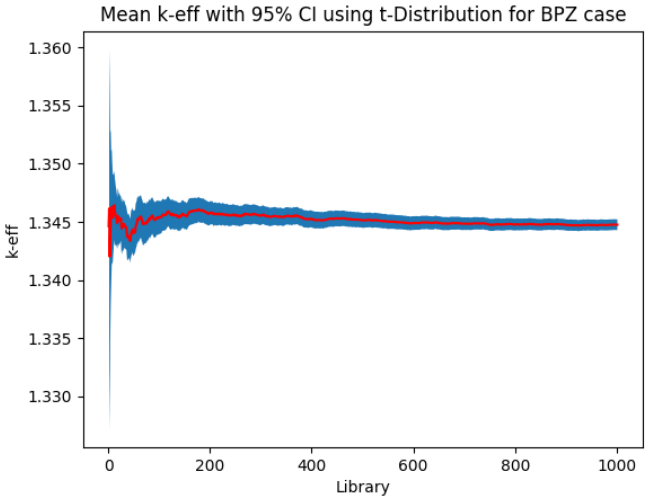

Figure A.9 BWR Pincell at HZP 95 % confidence interval.

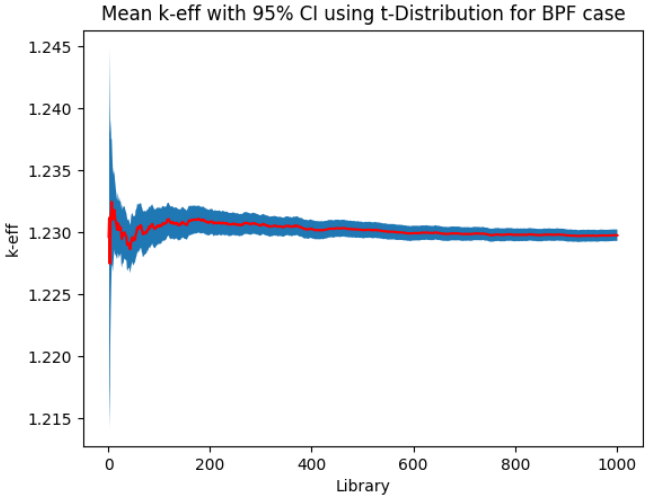

Figure A.10 BWR Pincell at HFP 95 % confidence interval.

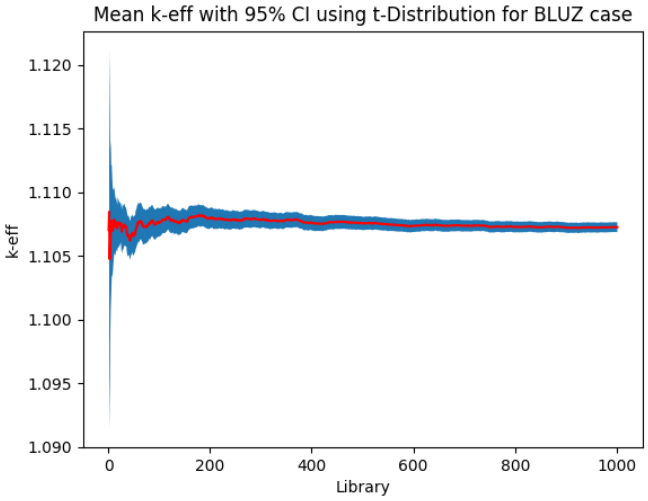

Figure A.11 BWR Lattice Unrodded at HZP 95 % confidence interval.

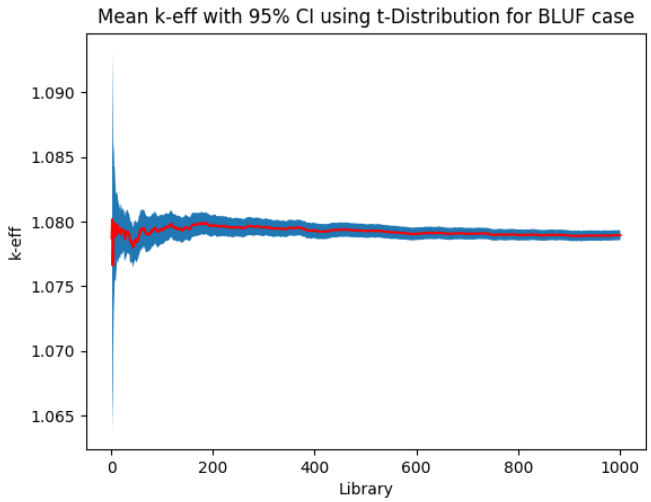

Figure A.12 BWR Lattice Unrodded at HFP 95 % confidence interval.

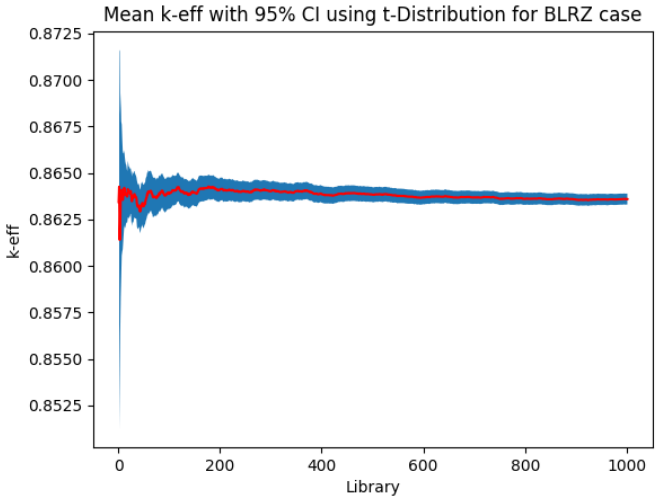

Figure A.13 BWR Lattice Rodded at HZP 95 % confidence interval.

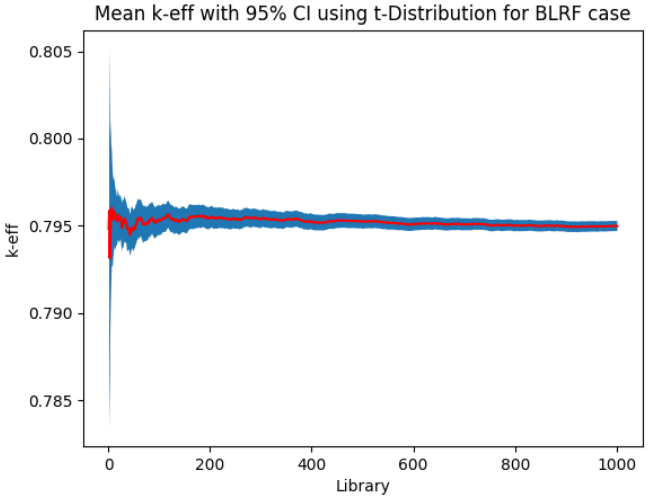

Figure A.14 BWR Lattice Rodded at HFP 95 % confidence interval.

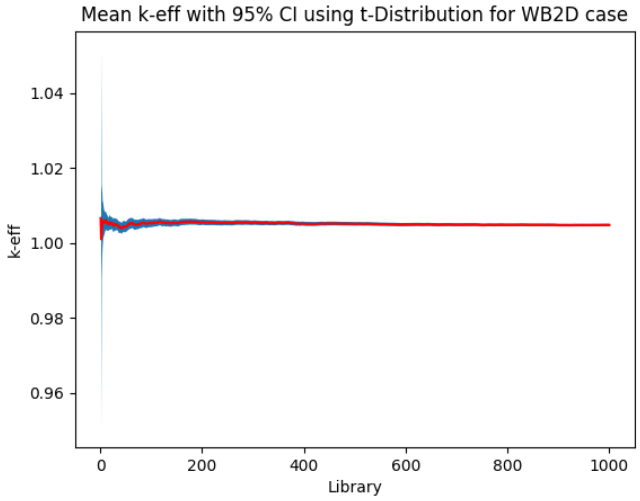

Figure A.15 Watts Bar two-dimensional full core 95 % confidence interval.

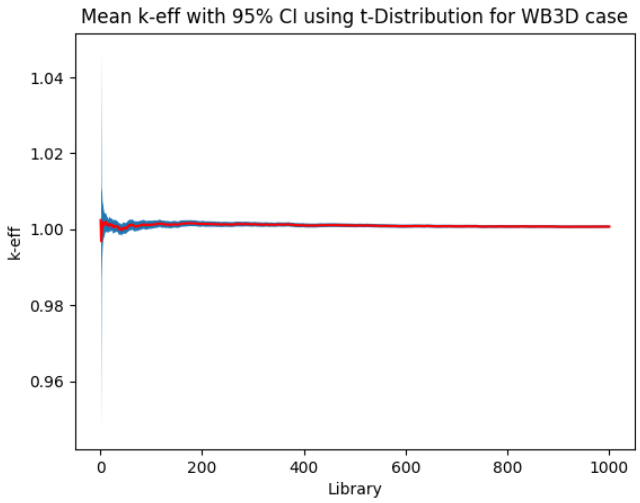

Figure A.16 Watts Bar three-dimensional full core 95 % confidence interval.

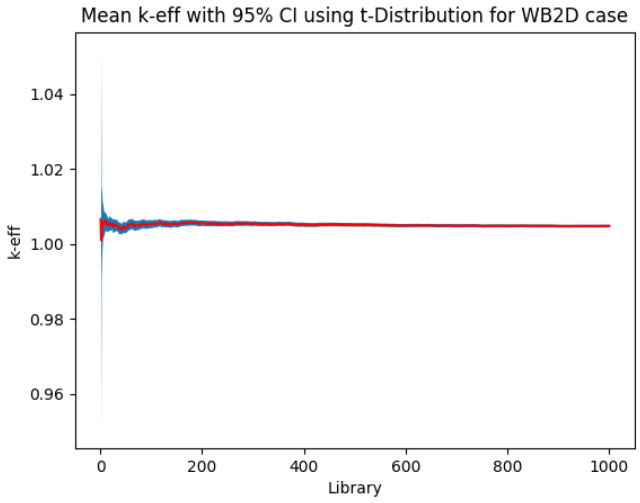

Figure A.17 Watts Bar three-dimensional full core with thermal hydraulic feedback 95 % confidence interval.