

Stand-off Nuclear Reactor Monitoring with Neutron Detectors at the McMaster
Nuclear Reactor

STAND-OFF NUCLEAR REACTOR MONITORING WITH NEUTRON
DETECTORS AT THE MCMaster NUCLEAR REACTOR

By Philip James BARRON, BAsC

*A Thesis Submitted to the School of Graduate Studies in the Partial Fulfillment
of the Requirements for the Degree Master of Applied Science*
McMaster University © Copyright by Philip James BARRON December 7, 2022

McMaster University

Master of Applied Science (2022)

Hamilton, Ontario (Department of Engineering Physics)

TITLE: Stand-off Nuclear Reactor Monitoring with Neutron Detectors at the McMaster Nuclear Reactor

AUTHOR: Philip James BARRON (McMaster University)

SUPERVISOR: Dr. Adriaan BUIJS

NUMBER OF PAGES: xiv, 94

Abstract

Nuclear reactor safeguards are how the peaceful use of nuclear material is ensured. Safeguards consist of a broad array of techniques, such as video surveillance and tamper-proof seals, to ensure that nuclear material is not diverted from declared activities. Safeguards research is conducted to ensure that safeguards techniques are capable of meeting the challenges posed by future reactor designs and operating conditions. One such technique that has broad applicability to novel reactor designs, including small modular reactors, is the method of standoff neutron detection using large area neutron detectors. In this method, neutron detectors are employed to detect neutrons which have escaped from the core, which are representative of the flux inside the core. Because the flux required to achieve a given power is dependent on the isotopes being fissioned, due to their different cross sections and fission energies, the state of the core can be assessed using the neutron detectors. Prior research has demonstrated that it is possible to correlate kilogram changes in fissile inventory using neutron detectors by employing the standoff neutron detector method. This work at the McMaster nuclear reactor details additional experiments to support prior work.

First, the apparatus and procedure to collect neutron detector data are detailed, along with persistent challenges to the collection. Next, simulations using the OSCAR-5 code to determine the fissile inventory are described. These two sections are subsequently combined, to compare changes in detector signal to the simulated core inventory. It was found that the uncertainty was too large to correlate changes in detector signal with changes in core inventory. Lastly, a method of detecting malicious interference is derived and tested.

Acknowledgements

I tend to go long, so I'll try to keep this short: I owe a lot to a lot of people. First and foremost, to Dr. Buijs and Dr. van der Ende for their advice and patience with me. To Dr. Simon Day for his advice and help with all things MNR. To Jay Grigg-Tait and the MNR operations staff for their help inside MNR, and for answering each of my one hundred emails asking for data, time after time. To Dr. Mohammed Alqahtani for his help early in this work when I was still learning the ropes. To Dr. Elizabeth MacConnachie for her help concerning uncertainties and flux at MNR.

To Cahit, PJ, Stephen, Sean, Elliot, Mike and Nik for the good conversation about nuclear engineering and everything else under the sun, and their friendship during Covid. To Zahra, and her family for their company (and food!) throughout and particularly at the end of this journey. Lastly, to my own family, for their love and guidance for the last 20-something years. As the saying goes: You go along and there you are.

Contents

Abstract	iii
Acknowledgements	iv
Acronyms	xiii
Declaration of Authorship	xiv
1 Introduction	1
1.1 Nuclear Safeguards and the Standoff Neutron Detection Technique	1
1.2 Standoff Neutron Detection Theory	2
1.3 McMaster Nuclear Reactor	4
1.4 Malicious Interference	6
2 MNR Experiments	8
2.1 Introduction	8
2.2 Experimental Apparatus	8
2.3 Timeline of Locations at MNR	12
2.4 Data Processing	17
2.4.1 Data Collection and Processing	17
2.4.2 Beam Port Interference Removal	21
3 Core Following and Weighted Isotopic Composition Generation	26
3.1 Introduction	26
3.2 Core-following of MNR using OSCAR-5	26
3.2.1 Methodology	26
3.2.2 Results	27
3.3 WIC changes at MNR	28
3.3.1 Change in Fissile Inventory and WIC During Fuel Shuffles	29
3.3.2 Change in Fissile Inventory and WIC During Cores	33
3.3.3 Change in WIC During Possible Clandestine Activity	36
4 Results	39
4.1 Introduction	39
4.2 Weekly Periodicity of Detector Data	41
4.3 Comparison Between WIC and Uncorrected Normalized Counts	45
4.3.1 Overview	45

4.3.2	Comparison During Cores	47
4.3.3	Comparison During Fuel Shuffles	51
4.3.4	63E to 63H Aggregated Comparison	55
4.4	Comparison Between WIC and Corrected Neutron Counts	56
4.4.1	Overview	56
4.4.2	Comparison During Cores	58
4.4.3	Comparison During Fuel Shuffles	62
4.4.4	63E to 63H Aggregated Comparison	66
4.5	Discussion	66
4.5.1	Cores 63E to 63H Aggregation	66
4.5.2	Aggregating with Smaller Groups	69
4.5.3	Standoff Neutron Detection and SMRs	71
5	Malicious Interference	72
5.1	Overview	72
5.2	Background	73
5.3	Theory	73
5.4	Simulation Tool	79
5.4.1	Am-Be Source Measurements	79
5.4.2	Simulation Tool	81
5.5	Low Counts	85
5.6	True FAP and Detection Probability	85
5.6.1	True FAP	85
5.6.2	Detection Probability	87
5.6.3	Malicious Interference Conclusions	88
6	Conclusion	90
6.1	Summary and Conclusion	90
6.2	Future Work	91
	Bibliography	92

List of Figures

1.1	A cutaway illustration of an MNR FA [8].	5
1.2	Cross section of MNR core loading during the relevant period, generated using OpenMC, described in Chapter 3.	6
2.1	Architecture of the NPM device. Components in blue are analog electronics while the components in green are digital. Taken from [13].	9
2.2	Photo of the detector apparatus. Seven detector tubes are mounted vertically in aluminium T-slot tubing frames. The stand and horizontal bar to keep the tubes in place are also shown. The NPM device is hidden behind the black plate on the upper right corner of the detector.	10
2.3	Photo of the detector setup, including the UPS, PoE injector, NPM, and laptop connection. The white ethernet cable runs from the PoE injector to the NPM, and the grey cable runs from the PoE injector to the computer.	11
2.4	Timeline of detector locations.	13
2.5	Floor plan showing Location 1, indicated by a star, in reference to the NE stairwell and the reactor core.	13
2.6	Floor plan showing Location 2, indicated by a star, in reference to the NE stairwell and the reactor core.	14
2.7	Floor plan showing Location 3, indicated by a star, in reference to the NE stairwell and the reactor core.	15
2.8	Floor plan showing Location 4, indicated by a star, in reference to the NE stairwell and the reactor core.	15
2.9	Floor plan showing Location 5, indicated by a star, in reference to the NW stairwell and the reactor core.	16
2.10	Floor plan showing Location 6, indicated by a star, in reference to the NW stairwell and the reactor core.	16
2.11	Time series plot at shutdown showing power and counts. By referring to this plot the count series was shifted to align with the power series.	19
2.12	Time series plot showing typical daily operation. The neutron count data in orange, and the power data in black.	20
2.13	Time series plot showing a zoomed-in segment of Figure 2.12	20

2.14	Collection of plots showing the results of the filtering function. Top: time series plot of whole day operation. Orange points indicate the results of the filtering function. Horizontal orange dashed bars indicate the final thresholds. Below, left: Successive plots of the individual time segments that were filtered. Below, right: Bin histogram of counts during the corresponding time segment on the left. The normal fit σ and $\sqrt{\mu}$ are listed for comparison.	23
2.15	Filtering function applied to the segment of lower signal between 12:00 and 16:00 when the beam port was in use. Top: Results with <code>std_cutoff</code> set to 5, showing a significantly larger standard deviation than expected. Bottom: Results with <code>std_cutoff</code> set to 2, showing that points below the mean that should not be excluded are being excluded.	24
3.1	Upper Left: Mass of U235 and U238 in the MNR over the entire period of observation. Upper Right: Mass of Pu239 and Pu241. Lower Left: The normalized WIC. Lower Right: The normalized core-average flux and Xe125 concentrations.	28
3.2	Change in U235 mass in the core, during each fuel shuffle. The core labels on the x-axis refer to the fuel shuffles at the beginning of that core cycle. The length of the red bar indicates the mass of U235 in removed FAs, and the length of the green bar indicates the mass of U235 added in new FAs.	30
3.3	Change in U238 mass in the core, during each fuel shuffle. The length of the red bar indicates the mass of U238 in removed FAs, and the length of the green bar indicates the mass of U238 added in new FAs.	30
3.4	Change in Pu239 mass in the core, during each fuel shuffle. The length of the red bar indicates the mass of Pu239 in removed FAs, and the length of the green bar indicates the mass of Pu239 added in new FAs.	31
3.5	Change in Pu241 mass in the core, during each fuel shuffle. The length of the red bar indicates the mass of Pu241 in removed FAs, and the length of the green bar indicates the mass of Pu241 added in new FAs.	31
3.6	The effects of the constituent isotopes on the WIC during fuel shuffles. Each fuel shuffle is normalized such that 0% change is equal to the WIC of the EOC core without any removals or additions.	33
3.7	The change in WIC resulting from the removal of all plutonium from the core at EOC and BOC for each core.	37
3.8	The change in WIC resulting from removing one FA with nearest burnup to 0%/10%/20%/30%/40%/50%.	38
4.1	Flow chart showing “data cleaning” procedure applied in Chapter 2.	39
4.2	Time series plot of a day of filtered count, power, and counts per unit power data.	40
4.3	Time series plot of all WIC data.	41

4.4	Plot of all reactor weeks normalized to their respective weekly mean counts. All days are aligned so that startup is at 00:00, and averaged hourly. The black line shows the average of the means for that hour. . . .	42
4.5	Flow chart showing procedure to produce Figure 4.4.	43
4.6	Plot of a sample xenon transient created using the RK4 model from [21]. The final decay is approximately 56 hours, which is equal to a typical two day weekend shutdown.	44
4.7	Flow chart showing procedure to generate xenon coefficients.	44
4.8	Time series of normalized counts and WIC data on the same axis. NE0043 data are marked by circles, while NE02A data are marked by crosses. WIC data are marked by stars.	45
4.9	Flow chart showing procedure to produce Figure 4.8.	45
4.10	Plot of residuals between the daily mean of the two detectors over the period when both detectors were in their long-term observation locations.	46
4.11	Normalized counts per unit power against WIC for cores 63C to 63H.	47
4.12	Left: Daily averaged time series plot of normalized counts per unit power, and WIC during the 63C core. Right: Daily averaged normalized counts per unit power against WIC during the 63C core.	48
4.13	Left: Daily averaged time series plot of normalized counts per unit power, and WIC during the 63D core. Right: Daily averaged normalized counts per unit power against WIC during the 63D core.	48
4.14	Left: Daily averaged time series plot of normalized counts per unit power, and WIC during the 63E core. Right: Daily averaged normalized counts per unit power against WIC during the 63E fuel core.	49
4.15	Left: Daily averaged time series plot of normalized counts per unit power, and WIC during the 63F_p1 core. Right: Daily averaged normalized counts per unit power against WIC during the 63F_p1 core.	49
4.16	Left: Reactor week averaged time series plot of normalized counts per unit power, and WIC during the 63F_p2 core. Right: Reactor week averaged normalized counts per unit power against WIC during the 63F_p2 core.	50
4.17	Left: Daily averaged time series plot of normalized counts per unit power, and WIC during the 63G core. Right: Daily averaged normalized counts per unit power against WIC during the 63G core.	50
4.18	Left: Daily averaged time series plot of normalized counts per unit power, and WIC during the 63H core. Right: Daily averaged normalized counts per unit power against WIC during the 63H core.	51
4.19	Normalized counts per unit power against WIC for the three weeks before and after fuel shuffle 63C to 63D.	52
4.20	Normalized counts per unit power against WIC for the three weeks before and after fuel shuffle 63D to 63E.	52
4.21	Normalized counts per unit power against WIC for the three weeks before and after fuel shuffle 63E to 63F_p1.	53
4.22	Normalized counts per unit power against WIC for the three weeks before and after fuel shuffle 63F_p1 to 63F_p2.	53

4.23	Normalized counts per unit power against WIC for the three weeks before and after fuel shuffle 63F_p2 to 63G.	54
4.24	Normalized counts per unit power against WIC for the three weeks before and after fuel shuffle 63G to 63H.	54
4.25	Normalized counts per unit power against WIC for cores 63E to 63H, with data averaged over three or four weeks, if available.	55
4.26	Flow chart of aggregation plotting procedure.	55
4.27	Time series of xenon corrected normalized counts and WIC on the same axis. NE0043 data are marked by circles, while NE02A data are marked by crosses. WIC data are marked by stars.	56
4.28	Time series data of the 63C and 63D cores showing both corrected and uncorrected data for both detectors.	57
4.29	Time series data of cores 63F_p1, 63F_2, and 63G showing both corrected and uncorrected data for both detectors.	57
4.30	Xenon corrected normalized counts per unit power against WIC for cores 63C to 63H.	58
4.31	Left: Reactor week averaged time series plot of xenon corrected normalized counts per unit power, and WIC during the 63C core. Right: Reactor week averaged xenon corrected normalized counts per unit power against WIC during the 63C core.	59
4.32	Left: Reactor week averaged time series plot of xenon corrected normalized counts per unit power, and WIC during the 63C core. Right: Reactor week averaged xenon corrected normalized counts per unit power against WIC during the 63D core.	59
4.33	Left: Daily averaged time series plot xenon corrected of normalized counts per unit power, and WIC during the 63E core. Right: Daily averaged xenon corrected normalized counts per unit power against WIC during the 63E core.	60
4.34	Left: Daily averaged time series plot of xenon corrected normalized counts per unit power, and WIC during the 63F_p1 core. Right: Daily averaged xenon corrected normalized counts per unit power against WIC during the 63F_p1 core.	60
4.35	Left: Daily averaged time series plot of xenon corrected normalized counts per unit power, and WIC during the 63F_p2 core. Right: Daily averaged xenon corrected normalized counts per unit power against WIC during the 63F_p2 core.	61
4.36	Left: Daily averaged time series plot of xenon corrected normalized counts per unit power, and WIC during the 63G core. Right: Daily averaged xenon corrected normalized counts per unit power against WIC during the 63G core.	61
4.37	Left: Daily averaged time series plot of xenon corrected normalized counts per unit power, and WIC during the 63H core. Right: Daily averaged xenon corrected normalized counts per unit power against WIC during the 63H core.	62

4.38	Normalized counts per unit power against WIC for the three weeks before and after fuel shuffle 63C to 63D.	63
4.39	Normalized counts per unit power against WIC for the three weeks before and after fuel shuffle 63D to 63E.	63
4.40	Normalized counts per unit power against WIC for the three weeks before and after fuel shuffle 63E to 63F_p1.	64
4.41	Normalized counts per unit power against WIC for the three weeks before and after fuel shuffle 63F_p1 to 63F_p2.	64
4.42	Normalized counts per unit power against WIC for the three weeks before and after fuel shuffle 63F_p2 to 63G.	65
4.43	Normalized counts per unit power against WIC for the three weeks before and after fuel shuffle 63G to 63H.	65
4.44	Xenon corrected normalized counts per unit power against WIC for cores 63E to 63H, with data averaged over three or four weeks, if available. . . .	66
4.45	Normalized counts per unit power against WIC, data from NRU experiment [5].	67
4.46	WIC and core average thermal flux from OSCAR.	68
4.47	Normalized counts per unit power against WIC for cores 63E to 63H, with data averaged over two to three weeks, if available.	70
4.48	Xenon corrected normalized counts per unit power against WIC for cores 63E to 63H, with data averaged over two to three weeks, if available. . . .	70
5.1	Poisson distributions ($\mu = 4.5$) with thresholds $k_L = 2$ (top) and $k_U = 8$ (bottom).	75
5.2	Poisson distribution with $\mu = 1.5$, although k_L is set to 0, the probability of observing 0 counts is still too high to set a reasonable FAP_L	77
5.3	Simulated background with transient added to show the points used in the background and decision time samples.	78
5.4	A sample of the counts observed with the AmBe source present, in the McMaster Tandem Accelerator Building. The local environment was not changed during the course of the sample being collected.	80
5.5	A sample of the counts with the AmBe source placed at 2.8 m. “3-ply” refers to the wax, rubber and cork slab. The AmBe source was not present during the “background” sample.	81
5.6	Multiple different shape function transients with a maximum height of 10 counts, on top of a background count rate of 10 counts, shown are the Gaussian pulse, square, triangle, and sawtooth functions.	82
5.7	Left: Complete time series of data taken at MNR, Right: Section of time series over which Cd foil was placed in front of the detector, between the detector and the reactor.	83
5.8	Time series taken from MNR with simulated data overlaid.	84

List of Tables

2.1	Approximate count rates per 10 s at each of the locations.	17
3.1	Change in mass of the four isotopes that comprise the WIC over the duration of the core.	34
3.2	Mass rate of change of the four isotopes that comprise the WIC, in units of [g/h], for each core cycle.	34
3.3	Mass rate of change of the four isotopes that comprise the WIC, in units of [g/MWh], for each core cycle.	35
3.4	Absolute change and rate of change of the WIC in [g/h] and [g/MWh], for each core cycle	35
4.1	FAPs used to test MNR experiment.	71
5.1	FAPs used to test MNR experiment.	84
5.2	Results of modified MAL testing. Each trial consisted of N=100,000 samples, with a mean of 30, and each summed sample consisted of the prior 100 samples.	86
5.3	Results of modified MAL testing. Each trial consisted of N=100,000 samples, with a mean of 10, and each summed sample consisted of the prior 15 samples.	86
5.4	Results of modified MAL testing. Each trial consisted of N=100,000 samples, with a mean of 1, and each summed sample consisted of the prior 20 samples.	86
5.5	Results of the five different transient magnitudes applied to the background with a count rate of 150 cps, with a mean of 30 counts, using a decision time of 20 s.	87
5.6	Results of the five different transient magnitudes applied to the background with a count rate of $6.\bar{6}$ cps, with a mean of 10 counts, using a decision time of 20 s.	88
5.7	Results of the five different transient magnitudes applied to the background with a count rate of 1 cps, with a mean of 1 counts, using a decision time of 20 s.	88

Acronyms

CSA	Comprehensive Safeguards Agreement
SMR	Small Modular Reactor
CNL	Canadian Nuclear Laboratories
SND	Standoff Neutron Detection
NRU	National Research Universal
MNR	McMaster Nuclear Reactor
WIC	Weighted Isotopic Composition
MTR	Materials Testing Reactor
SFA	Standard Fuel Assembly
CIF	Central Irradiation Facility
RPM	Radiation Portal Monitor
NPM	Neutron Pulse Monitor
UPS	Uninterrupted Power Supply
PoE	Power over Ethernet
BOC	Beginning of Core
EOC	End of Core
MGRAC	Multi-Group Reactor Analysis Code
NORM	Naturally Occurring Radioactive Material
FAP	False Alarm Probability
TM	Transient Magnitude

Declaration of Authorship

I, Philip James BARRON, declare that this thesis titled, “Stand-off Nuclear Reactor Monitoring with Neutron Detectors at the McMaster Nuclear Reactor” and the work presented in it are my own. I confirm that, under the supervision of Dr. Adriaan Buijs and Dr. Bryan van der Ende, I performed all of the work presented within including:

- **Chapter 1:** Background research, writing.
- **Chapter 2:** Apparatus setup, data acquisition, analysis, writing.
- **Chapter 3:** Analysis, Writing. OSCAR-5 burnup simulations were performed by Dr. Simon Day as part of his responsibilities as MNR Reactor Physicist.
- **Chapter 4:** Analysis, writing.
- **Chapter 5:** Apparatus setup, data acquisition, analysis, writing.
- **Chapter 6:** Writing.

Chapter 1

Introduction

1.1 Nuclear Safeguards and the Standoff Neutron Detection Technique

Nuclear reactor safeguards are how the peaceful use of nuclear material is ensured [1]. Specifically, safeguards ensure that states which operate nuclear facilities do not themselves divert nuclear material from declared nuclear activities. The diversion of nuclear material may take many forms, and may occur at different points in the nuclear fuel cycle; either prior, during or after irradiation inside of the reactor core. Consequently, there is a broad range of safeguards techniques, from tamper-proof seals to security cameras and handheld particle detectors that are employed at different parts of the nuclear fuel cycle [2]. The IAEA, as the international body responsible for promoting the peaceful use of nuclear material, creates and enforces Comprehensive Safeguards Agreements (CSAs) with member states, which detail the state's legal requirements to submit themselves to a safeguards regime, and regular inspection [3]. The IAEA Department of Safeguards is also responsible for pursuing safeguards research, by creating research plans detailing its priorities for future research and directing safeguards research among member states.

Safeguards research is important to ensure that safeguards methods are capable of meeting the challenges posed by future reactor designs and deployment scenarios. There has been significant interest in Small Modular Reactors (SMRs) [4]. SMRs may pose problems for some existing safeguards techniques: Many SMRs are sealed-core designs, meaning that the core cannot be opened to verify the fissile inventory, which is may be comprised of higher enrichment fuel than typical power reactors to offset their smaller size. The mass deployment of SMRs in remote locations may also make timely safeguards inspections difficult. To meet these challenges, one such novel safeguards technique developed by Canadian Nuclear Laboratories (CNL) is Standoff Neutron Detection (SND). In brief, it is possible to correlate emanations from a nuclear reactor, in this case, neutrons, with the internal state of the reactor, namely the quantity of fissile material in the core. Dr. Bryan van der Ende and colleagues at CNL have demonstrated that it is possible to correlate kilogram quantity changes in reactor fissile isotope inventory with changes in neutron count rate collected by strategically positioned neutron detectors

surrounding the core [5]. Since the decommissioning of the National Research Universal (NRU) reactor in 2018 (the reactor studied in the aforementioned work), the McMaster Nuclear Reactor (MNR) is the largest operational research reactor in Canada, and therefore a prime candidate for additional experimentation into the SND method. This work details those experiments. In this introductory chapter, background concerning the SND method, a description of the MNR and so-called “malicious interference” are provided. Chapter 2 describes the neutron detector data collection and analysis procedure. Chapter 3 details computer simulations employed to estimate the fissile inventory of the MNR for comparison to the neutron detector data. Chapter 4 provides the results of comparisons between the neutron detector data, and the estimates of fissile inventory. Chapter 5 considers the resilience of the SND method to “malicious interference”, which refers to purposeful attempts to manipulate the local neutron environment around the detector to mask clandestine activity. Conclusions are provided in Chapter 6.

1.2 Standoff Neutron Detection Theory

The primary particle of interest in nuclear engineering is the neutron, as the neutron is the particle that initiates nuclear fission. For this reason, quantities related to the number of neutrons inside a reactor are very important for determining the state and dynamics of the reactor. The most common quantity related to the number of neutrons is the neutron flux. Like other mathematical fluxes, the neutron flux is equal to the number of neutrons passing through some infinitesimal surface area per time, consequently, it has units of [neutrons/cm²s]. Often the central goal of reactor physics is to determine the spatial and energy distribution of the flux inside a reactor. The fission power is related to the neutron flux by Equation (1.1), where the index i ranges over all fissile isotopes, $\phi(r, E)$ is the flux at position \vec{r} and energy E , $\sigma_{f,i}$, $E_{f,i}$, N_i are the fission cross-section, fission energy, and atom density of isotope i , respectively.

$$P = \sum_i \int dE \int d^3\vec{r} N_i \phi(\vec{r}, E) \sigma_{f,i}(\vec{r}, E) E_{f,i} \quad (1.1)$$

Calculating the integrals over space and energy necessary to solve this equation is very complex, as the reactor geometry itself is composed of countless parts made of many different materials, each with its own cross sections that influence the spatial and energy distribution of the flux. Consequently, many simplifications are often used to find analytic solutions to the neutron flux, or more complex computer-based approaches are employed.

Perhaps the most extreme simplification is to determine representative spatial and energy averages of the flux, fission cross-sections and fission energy of each isotope, and to replace the integrals with those appropriately averaged quantities. These simplifications could be conceptualized as a 1-group, homogeneous slab reactor model. The resulting

analogue to Equation (1.1) after averaging is Equation (1.2), where $\langle\phi\rangle$ and $\langle\sigma_f\rangle$ are spatial and energy averaged quantities, and E_f has no energy or spatial dependence.

$$P_{tot} = V \langle\phi\rangle \sum_i N_i \langle\sigma_{f,i}\rangle E_{f,i} \quad (1.2)$$

This in turn can be rearranged to obtain Equation 1.4, as shown below, where the volume and sum over the atomic densities are re-written terms of Avogadro's number, N_A and the mass inventory and atomic weight of the i th fissile isotope, m_i and w_i .

$$\frac{\langle\phi\rangle}{P_{tot}} = \left[V \sum_i N_i \langle\sigma_{f,i}\rangle E_{f,i} \right]^{-1} \quad (1.3)$$

$$\implies \frac{\langle\phi\rangle}{P_{tot}} = \left[N_A \sum_i \frac{m_i \langle\sigma_{f,i}\rangle E_{f,i}}{w_i} \right]^{-1} \quad (1.4)$$

Before remarking on Equation 1.4, consider Equation 1.5, which states that the number of neutrons detected, n_{det} , is proportional to the total neutron population in the core n_{pop} , which can be represented in terms of the flux, average neutron velocity $\langle v \rangle$ and volume V . Note that the average velocity and volume of the reactor do not change during operation and that therefore, the number of neutrons detected is proportional to the neutron flux.

$$n_{det} \propto n_{pop} \propto \frac{\langle\phi\rangle}{\langle v \rangle} V \quad (1.5)$$

Returning to Equation 1.4, in this work, the quantity on the right-hand side will be referred to as the Weighted Isotopic Composition (WIC). As the name suggests, it is equal to the mass of the respective fissile isotope weighted by the likelihood of fission and energy released per fission. Note that the right-hand side represents a model for the state of the reactor core, and the left-hand side involves two measurable quantities, the flux by Equation 1.5, and the thermal power by conventional means using a flow meter and thermocouples (at least in this work).

Isotopic composition of the core changes as U235 is depleted and as plutonium isotopes are created, resulting in subtle changes in flux per unit power caused by the difference in cross-sections and fission energies. This is what makes this technique possible to employ as a safeguards approach; the equation is separated between observables on the LHS, by which the state of the reactor, as represented by the RHS, can be inferred. For the purpose of this work, because U235, U238, Pu239 and Pu241 comprise 99.97% of the total fission power, the index i in Equation 1.4 spans those four isotopes (which are referred to colloquially as the four fissile isotopes) [6]. The objective of this work is therefore to establish this correlation at MNR, as Dr. van der Ende et al. were able to

establish at NRU, to show that changes in fissile inventory can be correlated to changes in neutron counts per unit power. To better understand this work, it can be broken into two major parts: the first part being the experimental work of obtaining neutron count rate measurements at MNR over time (in essence, the LHS of Equation 1.4), and the second part being the generation of an estimate of the fissile inventory, and by extension, the WIC by computational means.

1.3 McMaster Nuclear Reactor

The McMaster Nuclear reactor is an open pool type research reactor (known as a Materials Testing Reactor (MTR)) with a maximum power of 5 MW_{th} , originally built in the mid-1950s. The reactor is currently operated at a nominal 3 MW_{th} . Before 2008, the reactor was fuelled with 93% enriched uranium, which was repatriated to the United States as part of the Reduced Enrichment for Research and Test Reactors program [7]. Since 2008, the reactor has used high assay low enriched uranium with a nominal enrichment of 19.75%. Like other MTRs, the MNR fuel is composed of U_2Si_3 in an aluminium matrix, which is then formed into curved fuel plates. The Standard Fuel Assemblies (SFAs) are composed of 16 curved plates welded to aluminium side plates, with one fuel-less curved endplate on either side, forming a rectangular tube through which coolant travels. The bottom of each assembly has a snout, which slots into the grid plate, below the core (described below). The Control Fuel Assemblies are modified FAs with seven central fuel plates removed to make room for the control rod. There are five AgInCd shim rods that are operated in unison, which are used for coarse reactivity control and shutdown, and one stainless-steel regulating rod, used for fine reactivity control. A cutaway illustration of an MNR FA is shown below in Figure 1.1:




FIGURE 1.1: A cutaway illustration of an MNR FA [8].

The fuel assemblies are loaded into a 9 by 6 grid, labelled “1” through “9” and “A” through “F”. Row 8, on the south side of the core, is a wall of graphite reflector assemblies, the west wall is lined with a lead block, and the north and east walls are left open to the six beam port tubes, operated by a variety of research and commercial users. Fuelling is performed on a burn-up basis, meaning that FAs are removed when they reach the target exit burnup [9]. Burnup is estimated using CuMn flux wire measurements and inferring integrated power produced per element. Flux wire measurements are performed during each fuel shuffle. MacConnachie provides a more detailed description of the apparatus and procedure used in flux wire measurements in [10]. During the period relevant to this work, the MNR was loaded with 37 or 38 fuel assemblies. A typical core loading during the relevant period is shown below in Figure 1.2:

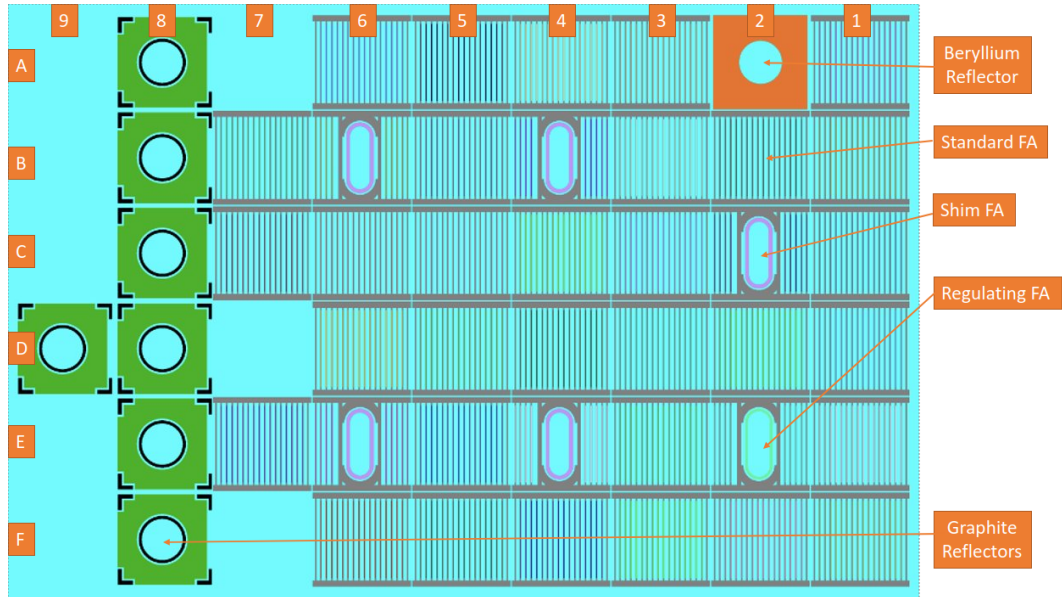


FIGURE 1.2: Cross section of MNR core loading during the relevant period, generated using OpenMC, described in Chapter 3.

An additional assembly, called the Central Irradiation Facility (CIF) is occasionally inserted into the core for research purposes at site 5C. Because the CIF does not contain any fissile material, it is often inserted with fresh FAs surrounding it to compensate for the negative reactivity of the CIF itself.

The MNR building consists of three floors, in order from top to bottom: the experimental floor, the mechanical floor, and the beam port floor. The control room and reactor pool are on the experimental floor. The mechanical floor houses the hot cell. The beam port floor is level with the reactor core itself, and the working area for the beam ports.

1.4 Malicious Interference

A secondary focus of this work is “malicious interference” attacks on the SND method. In essence, the objective is to determine the resilience of the SND method itself to cheating by the addition of exogenous neutron sources or shielding to alter the local neutron field surrounding the detector, in order to alter the neutron count rate in the detector. Because of the novelty of the SND method, the closest analogue to the malicious interference scenario is Radiation Portal Monitors (RPMs), which are radiation detectors (usually coupled gamma and neutron detectors) deployed at checkpoints to search for radioactive material [11]. RPMs employ detection algorithms that interpret detector data and sense when excess radioactivity produced by exogenous radioactivity sources in cargo, for example, is present. Similarly, a modified RPM algorithm has been implemented in order to sense the presence of exogenous sources or shielding surrounding the neutron detector. Such an algorithm is not to be confused with the SND method

itself, which provides a means of estimating the state of the core by inferring the flux using neutron detectors; rather, this algorithm is responsible for ensuring that the local neutron environment surrounding the detector has not been tampered or cheated and that the results of the SND method can be trusted.

Chapter 2

MNR Experiments

2.1 Introduction

This chapter pertains to determining the neutron count rate per unit power at MNR as a function of time, which is the left-hand side of Equation 1.4. This section details the detector and equipment used in the experiment, the locations at which measurements were performed, and the data processing and analysis performed in part to overcome certain limitations associated with counting neutrons at MNR.

2.2 Experimental Apparatus

The two neutron detectors used over the course of this project were composed of seven He3 proportional thermal neutron detectors from General Electric Reuter Stokes (model RS-P4-0840-223, similar to those shown under the “Resources” section in [12]). These detectors are stainless steel tubes filled with He3 gas to 4 atm. The active length of the He3 tubes measured 2.54 cm in diameter and 100.0 cm in length. The detectors were biased using an NPM3100E “all-in-one” Neuchrometer Neutron Pulse Monitor (NPM) by Quaesta Instruments [13]. Figure 2.1 shows the architecture of the NPM device. In brief, it consists of a high voltage source, a series of analog amplifiers, an analog to digital converter, and a microcontroller to tally the counts and store the data in text or binary files. Finally, a frame of aluminum T-slot framing rails with a stand for stability, and a horizontal bar to keep the He3 tubes in place was custom-built for these detectors. The complete detector apparatus is shown in Figure 2.2.

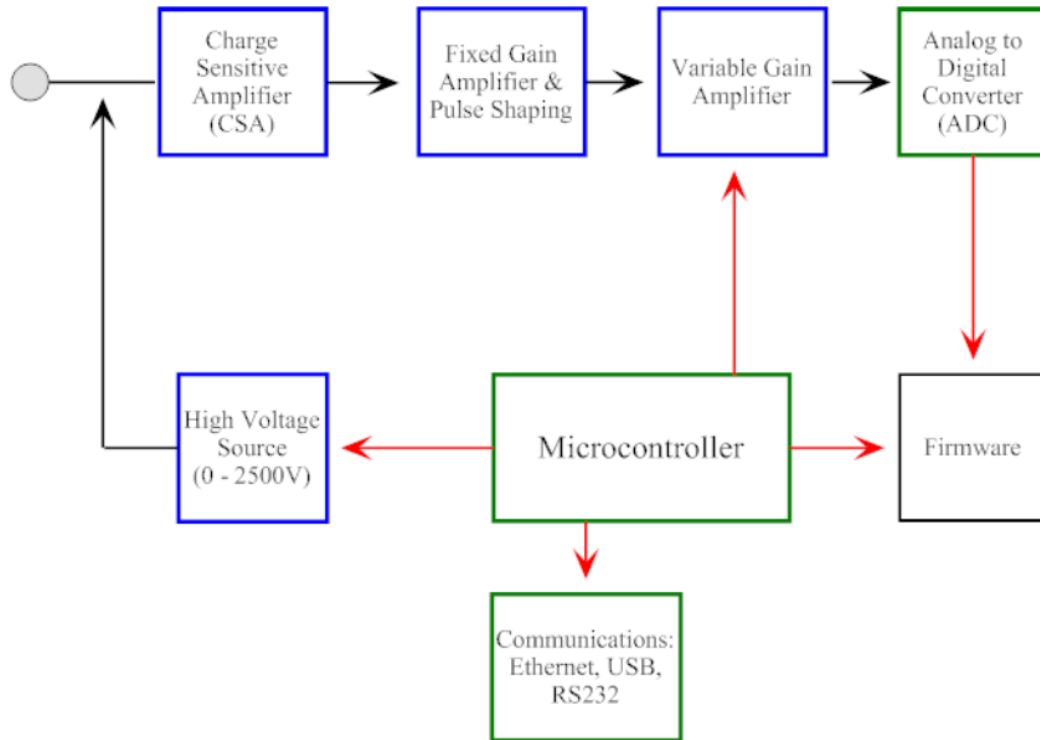


FIGURE 2.1: Architecture of the NPM device. Components in blue are analog electronics while the components in green are digital. Taken from [13].




FIGURE 2.2: Photo of the detector apparatus. Seven detector tubes are mounted vertically in each aluminium T-slot tubing frame. The stand and horizontal bar to keep the tubes in place are also shown. The NPM device is hidden behind the black plate on the upper right corner of the detector.

Figure 2.3, below, shows how the detector apparatus was powered. The setup consists of an uninterruptible power supply (UPS), plugged into an electrical outlet, and a power over ethernet (PoE) injector, plugged into the UPS and the NPM (housed inside the detector framing). The PoE injector was also where a computer was connected (via ethernet cable) to control and retrieve data from the NPM via proprietary software.

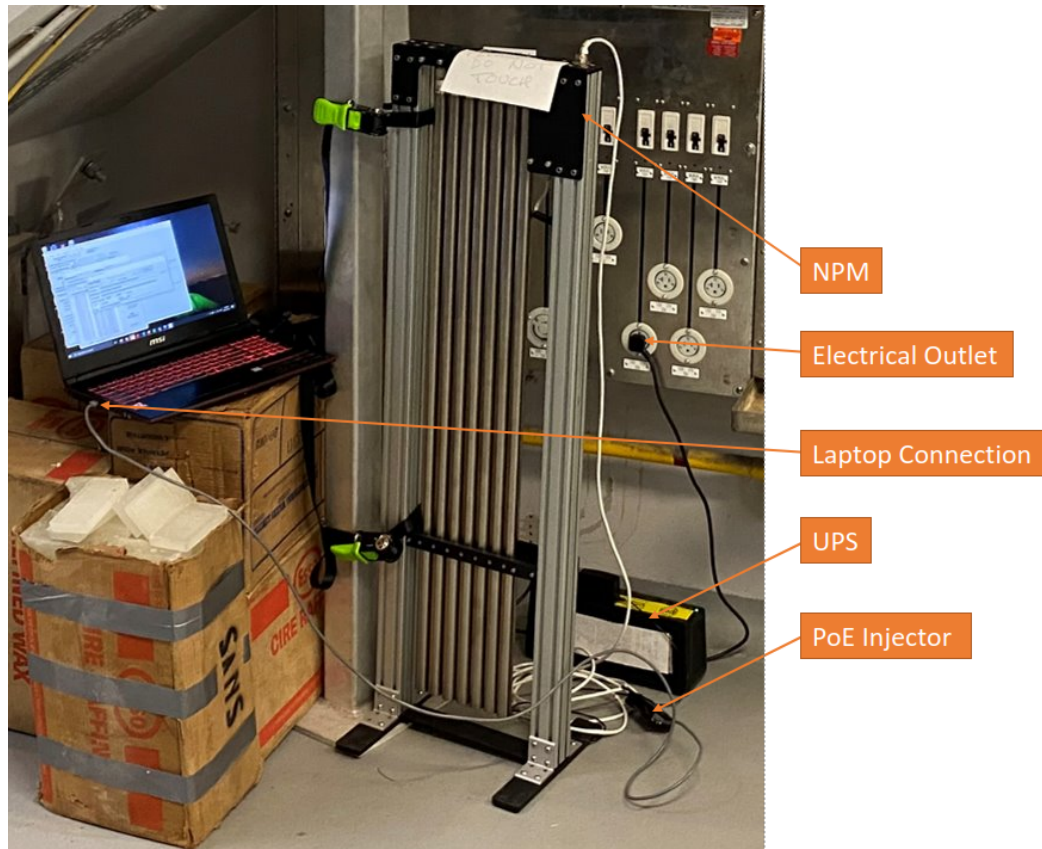


FIGURE 2.3: Photo of the detector setup, including the UPS, PoE injector, NPM, and laptop connection. The white ethernet cable runs from the PoE injector to the NPM, and the grey cable runs from the PoE injector to the computer.

The two detectors were given the IDs “NE02A and “NE0043”. Using the proprietary software provided by Quaesta Instruments, the following settings were applied to the NPM devices. For data acquisition, the NPM was set to the “DAT” mode, which produces binned, time series count data over a user-specified time-bin width. The voltage of both detectors was set to 1600 V, and the gain on detectors NE02A and NE0043 were set to 10 and 2, respectively. The lower discriminator was adjusted such that it was set at the lower edge of the wall effect continuum of the He3 pulse spectrum. These values were determined from prior testing of the detector apparatus by van der Ende et al. at CNL.

2.3 Timeline of Locations at MNR

Neutron counting at MNR was complicated by other research and commercial activities that generated a significant and varied neutron background. In particular, the beam ports, operated by a variety of users, would generate a neutron signal that could be many times larger than the neutron signal produced by leakage from the reactor. Measurements were taken at several different locations, all inside of the MNR facility, before locations for long-term placement were determined. Detector locations were selected to maximize the signal from the reactor, minimize the noise from beam port use, and ensure that the detectors would not be damaged during the counting campaign (that is, that there are adequate railings or other objects to strap down the detectors, and that the detectors were not placed in high-traffic areas), as well as the availability of electrical outlets. Below is a list of locations tested and the rationale for changing locations before determining the final locations. Figure 2.4 shows a timeline of detector locations throughout the detection campaign. Each location has a unique local neutron environment, which is a product of the distance between the reactor core and the detector, and the shielding between the core and the detector. Dr. van der Ende and colleagues have shown that the quantity of shielding (namely water, steel and high-density concrete at NRU) can drastically impact the local neutron environment to the degree that locations with less shielding, at a much farther distance (67 m versus 17 m from the reactor core), experienced approximately three to four times greater neutron count rate, adjusted for detector efficiency [5]. Because each location has a unique local neutron environment, to detect small changes in count rate that can be associated with changes in WIC, the detectors must be left in long-term locations to maximize the data available for comparison.

Finally, Dr. van der Ende and colleagues note that, at one of their detector locations, they were able to observe short changes in neutron count rate that they were able to ascertain were caused by online refuelling activities at NRU. Because the MNR cannot perform online refuelling, all fuel shuffle activities were performed when the core was shut down, and no variance from the typical background count rate was observed when the reactor was being refuelled, although an attempt was made to position detectors on the experimental floor (at the same elevation as the top of the reactor pool), where they might be able to better observe some deviation.

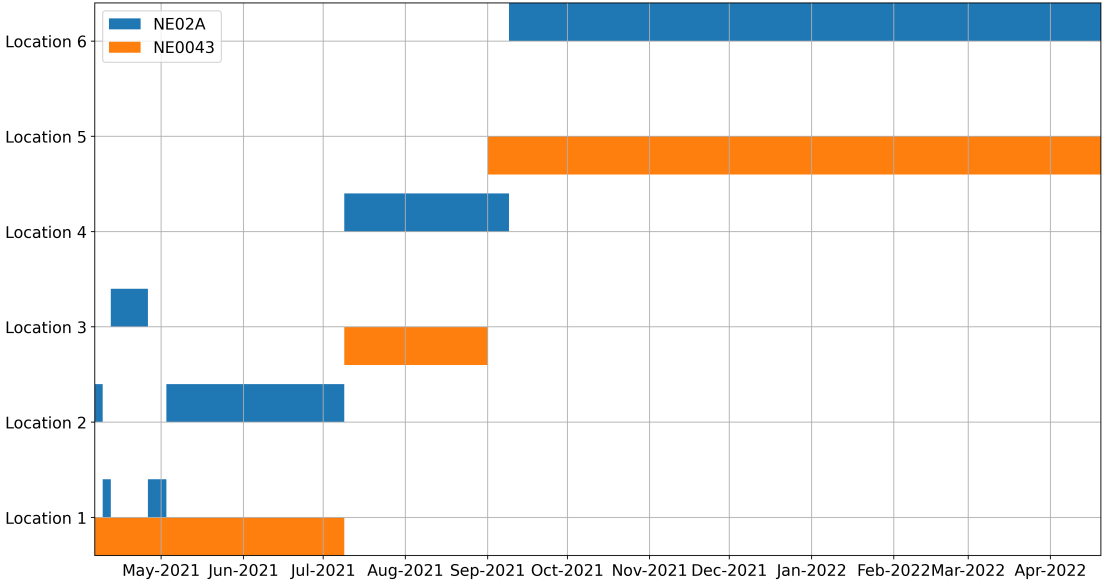


FIGURE 2.4: Timeline of detector locations.

The locations are listed below, in chronological order:

- Location 1:** Location 1 is at the end of the north corridor on the mechanical floor, to the right of a small stairwell to a storage closet. The location is marked on the floor plan, shown below in Figure 2.5.




FIGURE 2.5: Floor plan showing Location 1, indicated by a star, in reference to the NE stairwell and the reactor core.

- **Location 2:** Location 2 is directly below the north-east stairwell on the beam port floor, in front of the electrical panel. The location is in front of beam ports 4 and 5. The location is marked on the floor plan, shown below in Figure 2.6.




FIGURE 2.6: Floor plan showing Location 2, indicated by a star, in reference to the NE stairwell and the reactor core.

- **Location 3:** Location 3 is in the east corner of the control room floor, placed directly below the electrical panel. The location is marked on the floor plan, shown below in Figure 2.7.




FIGURE 2.7: Floor plan showing Location 3, indicated by a star, in reference to the NE stairwell and the reactor core.

- **Location 4:** Location 4 is along the back wall of the beam port floor, in front of beam port 6, on the east side of the building. The location is in front of beam ports 4 and 5. The location is marked on the floor plan, shown below in Figure 2.8.




FIGURE 2.8: Floor plan showing Location 4, indicated by a star, in reference to the NE stairwell and the reactor core.

- **Location 5:** Location 5 on top of the enclosure on the west side of the experimental floor. The location is marked on the floor plan, shown below in Figure 2.7.




FIGURE 2.9: Floor plan showing Location 5, indicated by a star, in reference to the NW stairwell and the reactor core.

- **Location 6:** Location 6 is in front of the hot cell, on the mechanical floor, it is marked on the floor plan, shown below in Figure 2.10.




FIGURE 2.10: Floor plan showing Location 6, indicated by a star, in reference to the NW stairwell and the reactor core.

Table 2.1 shows typical neutron count rates at each location.

Location	Approximate Counts per 10 s
Location 1	150
Location 2	2200
Location 3	350
Location 4	1200
Location 5	1100
Location 6	1200

TABLE 2.1: Approximate count rates per 10 s at each of the locations.

Location 1, the original detector location, was suggested by MNR staff as a suitable location to set up the detectors due to it being in a low-traffic area. Location 2 was tested to determine if a location on the beam port floor (that is, at the same elevation as the core itself) would have a higher local count rate. Location 3 was tested to determine the count rate on the experimental floor and to attempt to observe a change in count rate that could be associated with fuel shuffle activities, however, none was observed. The NE0043 detector was moved from location 3 because large hard plastic (likely polyethylene) bins were occasionally placed in front of the detector. Location 4 was tested as a potential alternative to location 2 for long-term observation because location 2 was no longer available. Location 5 was employed as an alternative to location 3 because it was in a low-traffic area with a consistent local environment (meaning the local shielding would not change over time). During the counting period when NE02A was positioned at location 4, there was beam port interference that presented a constant signal, as opposed to the square wave interference that was typical for beam port interference. Location 6 was tested to determine if it was possible to “block” or mitigate the constant beam port interference by placing the detector on the other side of the core from the beam port producing the constant signal. The activities that produced the constant beam port signal were ended shortly thereafter, however, NE02A remained at location 6 because the signal was comparably strong to location 4, and it was in a lower-traffic area.

2.4 Data Processing

2.4.1 Data Collection and Processing

Data is collected via ethernet cable connection using the NPM software. With the NPM in “DAT” mode, the data is gathered at 10 s bins and written in text files, that are timestamped at the beginning of the bin. At the end of a day, the current data file is terminated and a new file is written. Data analysis was performed in Python using the Pandas package, which provides tabular data structures to facilitate data manipulation and analysis [14]. Pandas also has extensive features for use with time series data, which were used in this project. To process the data in Python, the entire data set was added to a folder, and the individual files were sequentially read and concatenated to

create one large variable corresponding to all of the files in the folder. For ease of use, the timestamps were rounded to the nearest even 10 s interval (data with time-stamps ending in 0 through 4 would be rounded down to 0, and ending in 5 through 9 would be rounded up to 10). This rounding procedure was performed to merge the detector count data with thermal power data provided by MNR operations. The power is measured by a flow meter, and four thermocouples (two on the inlet and outlet), the power is calculated using Equation 2.1:

$$P = qC_p\Delta T \quad (2.1)$$

Where q is the mass flow rate, C_p is the specific heat, and ΔT is the change in temperature from the outlet to the inlet. A discussion of uncertainty in power measurements at MNR is found in MacConnachie et al. [15]. Power measurements were logged at 10 second intervals starting on the minute (such that the data was timestamped at 0 s, 10 s, 20 s and so on). The power and rounded count data were then merged, and only the timestamps for which power and count data both existed were kept.

The digital recorders in the MNR control room were not synchronized to any time, that is, there was an offset between both the time-stamp of the power data and the time on any internet-connected devices, and between the NPM count data and the power data time-stamps. For this reason, before the two time series could be compared, they first had to be synchronized. The simplest method of synchronizing was to compare the shutdown time of the power and count data and to shift that day of count data forward or backward to be aligned with the power data. For each day, before the filtering operation discussed below in Section 2.4.2 was performed, a plot of the two time series at shutdown was produced, and the count data was shifted until the first decrease in detector count rate occurred at the same time as the first decrease in thermal power. Because the neutron detector count rates are proportional to the flux in the core, which is in turn proportional to the neutron power (equivalent to the fission rate), the count signal decays more quickly than the thermal power. This procedure was repeated for each day during the counting campaign, to ensure that all days were synchronized.




FIGURE 2.11: Time series plot at shutdown showing power and counts. By referring to this plot the count series was shifted to align with the power series.

Figure 2.12 shows a time series plot of typical daily operations, with the neutron count rate in orange, and the power in black. The neutron count rate data shows a consistent signal at approximately 1200 counts per 10 s, and a fluctuating higher signal at approximately 12000 counts per 10 s. For example, in Figure 2.12, the higher signal is visible from approximately 9:00 to 16:00 with a brief pause around noon.




FIGURE 2.12: Time series plot showing typical daily operation. The neutron count data in orange, and the power data in black.

By comparing daily plots of the detector signal to the beam port use schedule, it was clear that the elevated signal to the detector (the signal above the consistent signal at 1200 counts per 10 s) was caused by beam port use. Figure 2.13 shows a smaller segment of time taken from the same day of operation as Figure 2.12.




FIGURE 2.13: Time series plot showing a zoomed-in segment of Figure 2.12

As can be seen in 2.13, although the consistent “lower” signal and higher signal appear continuous in 2.12, there are minute-long periods of beam port use followed by

minute-long periods free of the beam port signal. Points in between the “lower” and “higher” signal are caused by the time-binning, meaning that if the beam port shutters were opened or closed halfway through a time-bin, the count rate will be a combination of the two signals, resulting in an averaged count rate between the two signals. In the following section, the technique employed to remove the beam port signal is described.

2.4.2 Beam Port Interference Removal

For later comparison, the signal to the neutron detectors must be representative of the flux in the core. That is, if the signal without any filtering of the beam port signal were to be used, the result would be that the count signal would be more representative of the fraction of the day that the beam port shutters were open. As stated above, the “true” signal that represents the state of the core is the persistent signal around 1200 counts per 10 s in Figure 2.12. Therefore, instead of using the entire signal, only the low signal should be used for later comparison. To that end, a method of filtering the time series to only select the low signal was developed. The pseudocode corresponding to the function is shown below:

```
1 Function lowband_filter
2   Pass In: start_time, stop_time, initial_guess, std_cutoff
3
4   Set data to data between start_time and stop_time
5   Set mean and std to initial_guess
6
7   WHILE mean and std have not converged
8     Set data to data between mean - std_cutoff*std and mean +
9     ↪ std_cutoff*std
10    Perform normal distribution fit on data
11    Set new mean and std to results from normal distribution fit
12    Compare new mean and std to previous iteration
13
14   Set final mean and std to final interation mean and std
15   Set data to data between mean - std_cutoff*std and mean +
    ↪ std_cutoff*std
    Pass Out: mean, std, data
```

In this case, “convergence” means that the data points used in one iteration do not change from one iteration to the next. Recall that particle counting is a Poisson process, and thus the resulting distribution of particles counted per unit time will be the Poisson distribution [16]. Because the Poisson distribution can be approximated by a normal distribution for large λ , a normal distribution fit was applied to the data to determine the mean and standard distribution of the normal distribution representing the data.

In effect, we are trimming our data to be between `start_time` and `stop_time`, and creating a band-pass filter to select data around the mean. The while loop shows one of many possible ways of making the band-pass filter “dynamic”, meaning instead of simply

setting the bounds of the filter manually, the bounds are set automatically by iteratively applying a normal fit and selecting the data within `std_cutoff` number of standard deviations of the mean. `std_cutoff` was usually set to 5, and the filter usually captured all of the data without excluding any points during regular operation, as shown below. Finally, as a check, the standard deviation of the final normal fit can be compared to $\sqrt{\mu}$, as they should be approximately equal (as counting particles is a poisson process), with some additional variance caused by limited sample sizes and small changes in the underlying mean count rate.

Figure 2.14, shows the output of the filtering operation on the same day of data shown in Figure 2.12. The top plot is identical to Figure 2.12, with the orange points indicating data resulting from the filtering function. The plots below and on the left are time series plots of individual segments that were filtered. The red points are points that were between `start_time` and `stop_time`, but did not pass through the filter. The plots below and on the right are bin histograms of the corresponding time series on the left. Note that the first segment, between 10:00 and 10:20, has large variance due to change in reactor power during the 20 minute segment.

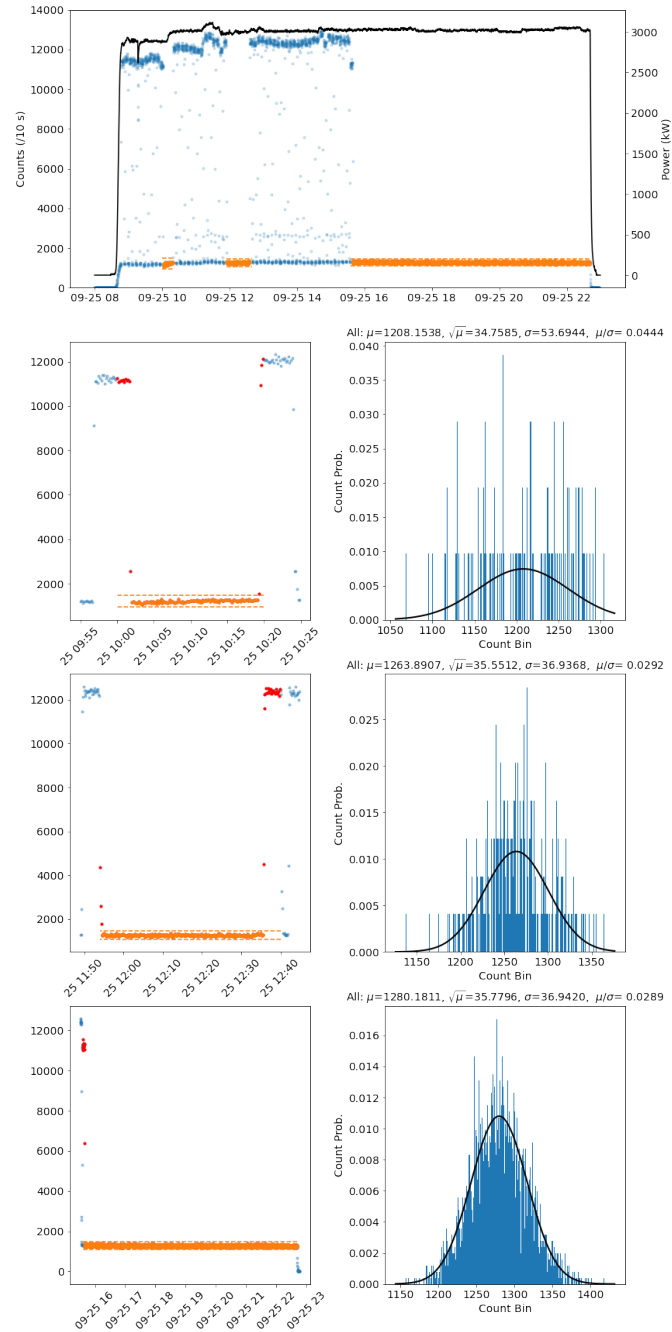


FIGURE 2.14: Collection of plots showing the results of the filtering function. Top: time series plot of whole day operation. Orange points indicate the results of the filtering function. Horizontal orange dashed bars indicate the final thresholds. Below, left: Successive plots of the individual time segments that were filtered. Below, right: Bin histogram of counts during the corresponding time segment on the left. The normal fit σ and $\sqrt{\mu}$ are listed for comparison.

One additional alternative is to use *all* of the “lower” signal data, including the minute-long segments in between beam port use, regardless of duration. Figure 2.15 shows this result. The standard deviation of the resulting normal fit is significantly larger (in the case shown below, approximately 1.5 times larger), and decreasing the bandwidth to decrease the standard deviation begins to exclude data too far below the mean, meaning that the filter is not working as intended (that is, there is no situation where the beam port use reduces the number of neutrons being counted, and therefore the filtering operation should not exclude points too far below the mean).

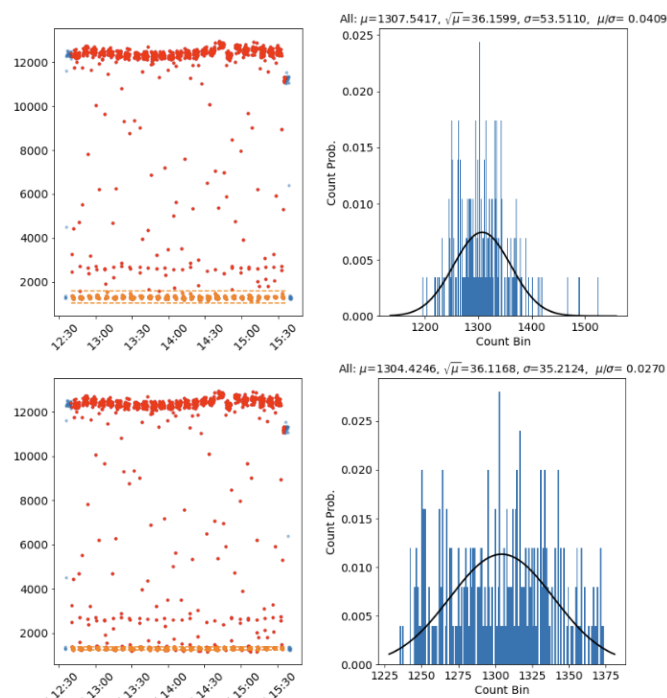


FIGURE 2.15: Filtering function applied to the segment of lower signal between 12:00 and 16:00 when the beam port was in use. Top: Results with `std_cutoff` set to 5, showing a significantly larger standard deviation than expected. Bottom: Results with `std_cutoff` set to 2, showing that points below the mean that should not be excluded are being excluded.

Moreover, if the “lower” signal when the beam port is in use tends to have a larger mean than the “lower” signal when the beam port was not in use, then the same problem would arise as simply using all of the data, that is, the mean would become a function of the fraction of the day that the beam port was in use.

As noted above, the frequency of the data logged in the power recorder and the frequency of the binning performed by the NPM system was the same, allowing for each pair of count and power values to be divided to obtain a normalized count per unit power value. This filtered count per unit power data is ultimately what is used in

comparison with the WIC values, discussed next in Chapter 3, to attempt to correlate neutron counts per unit power with the WIC.

Chapter 3

Core Following and Weighted Isotopic Composition Generation

3.1 Introduction

This chapter pertains to the weighted isotopic coefficient, and core following simulations of the MNR to compute the fissile inventory, which is used to generate the timestamped weighted isotopic coefficients that are compared to the detector data. The first section describes the WIC, and particularly how the WIC changes during and between fuel shuffles at the MNR. The following section contains a description of the core-following methodology of Dr. Simon Day, using the OSCAR-5 code suite.

3.2 Core-following of MNR using OSCAR-5

3.2.1 Methodology

As part of regular operations at MNR, the Reactor Analysis group, led by Dr. Simon Day is responsible for performing core-following simulations to track the MNR for safety and optimization of commercial radioisotope activities. For this work, it was valuable that MNR staff perform in-house core-following, as it meant that an independent core-following methodology did not need to be created and applied over the entire observation period to generate the WIC values.

The code used in core-following is the MGRAC (Multi-Group Reactor Analysis Code) nodal diffusion solver in the OSCAR-5 code package. OSCAR-5 is a code suite developed by the Radiation and Reactor Theory Section of the South African Nuclear Energy Corporation, designed for use in reactor analysis and core-following [17]. OSCAR-5 stores the model, including the geometry, materials, and cross-section (both standard ENDF/B-VII.1 .ACE files, and multi-group cross-section data generated by the SERPENT2 code) inside the code suite, and automates the creation of simulations using different codes (namely MCNP, SERPENT and MGRAC) and model inputs [18][19].

As noted, one aspect that sets MTRs apart from other reactor types is the day/night operation. The day/night operation has significant consequences for the core-following

procedure: It means that accurate modelling must use enough steps to track the xenon oscillations through regular operation. To that end, Dr. Day’s method creates time steps for use in the depletion by first noting the startup and shutdown times recorded by MNR operations staff, as well as the control rod positions (often described in terms of percent control rod insertion) when the reactor is first critical at 3 MW. The method involves creating time steps by working backwards in four-hour increments until the startup time is reached. Because the reactor never truly operates for a multiple of 4 hours, the earliest time step is of variable length. For example, say the reactor starts at 915h and shuts down at 2245h. The code would split the time into time stamps at 2245h, 1845h, 1445h, 1045h, and 915h (all 4-hour segments, except for the first, which is only 1.5 hours). The reason behind this operation is so that the earliest, and most precise time step occurs at the beginning of the day when the xenon concentration is changing the fastest. The method assumes that the reactor is operating at the nominal 3 MW at all times.

Using those time steps and the nominal power as inputs, the depletion stepping is performed. For daily operation, the next step is to run a ”flux step“ using MGRAC, with the CRs in their initial positions as recorded by MNR operations staff, these initial daily steps are referred to as ”criticality estimates”. The position of the CRs is essential to ensuring that the reactor maintains the correct flux shape, such that the fissile inventory depletes correctly. To that end, before every depletion step, the regulating rod is placed at 50% insertion, and a criticality search is performed (with a target k_{eff} of 1 minus a systematic offset, described below), with all shim rods set to the same position. The depletion is performed using a predictor-corrector integrator method. The criticality search systematic offset is computed by taking the average of the criticality estimates over the previous fuelling cycles. This procedure also provides a convenient way of ensuring that the core-follow is accurately depleting the fissile inventory; if the criticality estimates do not diverge from the history of criticality estimates (to within the model uncertainty), then the core-follow is accurately tracking the true operation of the core¹. Finally, at the end of the day, a decay step is performed between the shutdown time and the startup time the next morning.

Over fuel shuffling periods, the geometry and materials in the model are changed to match the inventory at BOC of the following core, and a “core-snapshot” simulation is computed within the OSCAR-5 suite using SERPENT2. The core snap-shots in SERPENT2 serve as another means of verifying the criticality estimates from OSCAR.

3.2.2 Results

The MGRAC depletion results can be written to text files, and a text parser was coded in Python to extract the desired quantities at each step: the time stamp, the masses of the four main fissile isotopes of interest, and additional quantities (a spatial average flux and average xenon concentration, k_{eff} and the reg and shim rod positions). These quantities were imported into a Pandas DataFrame to calculate the WIC for each step,

¹Indeed, during the course of this work, a bug was discovered and corrected by examining the criticality estimates over time.

and to perform additional analysis. Figure 3.1, below contains a series of plots that show the quantities extracted from the MGRAC output files over the entire period of observation, from May 18th 2021, until April 20th 2022: The upper left plot shows both the mass of U235 and U238 on separate axes, the upper right plot shows the quantity of Pu239 and Pu241 on separate axes, the bottom left plot shows the normalized WIC, and the normalized core-average flux and Xe125 concentrations. Further analysis of these results is found below and in Chapter 4.

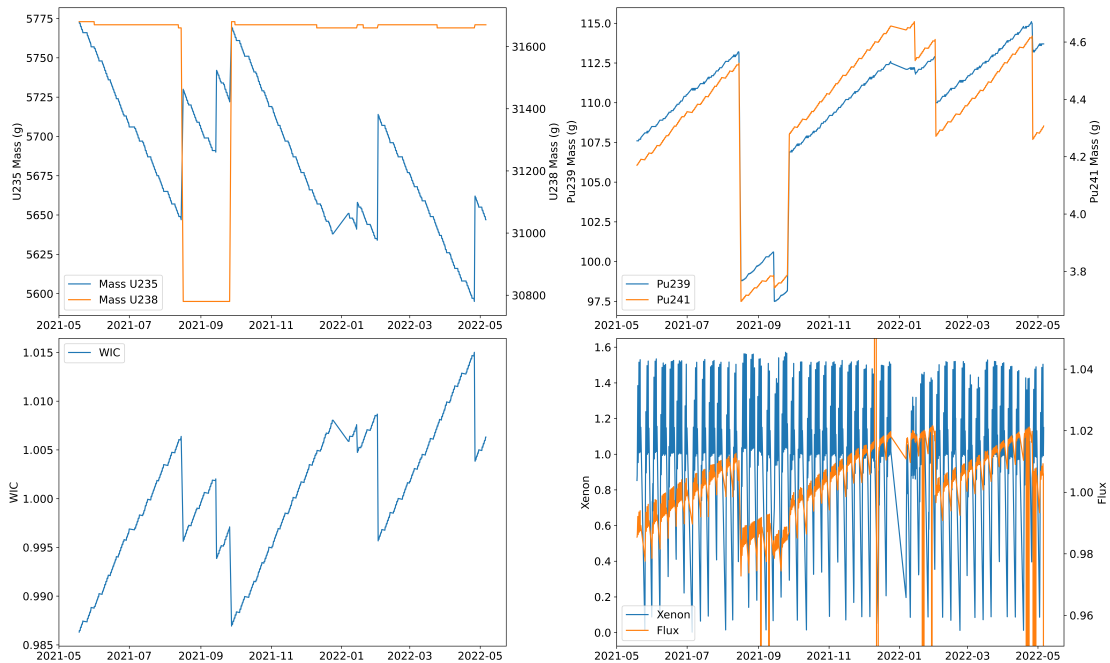


FIGURE 3.1: Upper Left: Mass of U235 and U238 in the MNR over the entire period of observation. Upper Right: Mass of Pu239 and Pu241. Lower Left: The normalized WIC. Lower Right: The normalized core-average flux and Xe125 concentrations.

3.3 WIC changes at MNR

Before comparing the WIC to the detector data in the following chapter, it is important to understand the changes in the WIC during fuel shuffles and by depletion during the individual cores at MNR. This section considers several aspects of the change in fissile inventory and WIC over time, as well as calculating the change in WIC from the removal of fissile inventory that might be removed in the event of clandestine activity. To generate the data used in this section, the OSCAR-5 core-follow model was used, in particular, the capability of the OSCAR-5 code to take the MNR model at a moment in time and generate a SERPENT model, with the geometry and materials corresponding to that time. The materials from the SERPENT models created in OSCAR-5 were then stripped from the input file and converted into the OpenMC material input file format.

OpenMC is a modern, open-source Monte Carlo code, complete with a Python API that allows for easier and more sophisticated model creation than in SERPENT was used for this limited purpose [20]. The OpenMC Python API allowed for the creation of arrays of materials containing the isotopic composition of the fuel to be stored and manipulated with ease, which was the purpose of converting from OSCAR to SERPENT to OpenMC. The materials variables could be altered and summed over to easily calculate the total mass and WIC. For each core, this process of generating OpenMC materials data was performed for the BOC and EOC materials, such that the change in mass and WIC could be calculated during individual cores (depletion during regular reactor operation) and during fuel shuffles (as FAs were added and removed).

3.3.1 Change in Fissile Inventory and WIC During Fuel Shuffles

The total change in fissile inventory during a fuel shuffle can be thought of as the sum of two terms, the first is the increase in fissile inventory as new FAs are added, and the second term is the reduction in fissile inventory as old FAs are removed. Consider the first term: Often, the only fissile inventory that is added is the nominal 225 g of U235, as well as approximately 915 g of U238 in fresh FAs (sometimes FAs that have previously been added to the core are re-added to the core). The second term is conversely the sum of all of the U235, U238 and plutonium isotopes in FAs that are removed. In this section, when considering the effects of Pu239 and Pu241 the WIC, the contributions of Pu239 and Pu241 are grouped for simplicity. Figures 3.2-3.5, below, show the change in the four isotopes during fuel shuffle activities. The fuel shuffles are labelled by the beginning of the new core, for example, on the x-axis labels of the plots below, the label “63B” corresponds to the 63A to 63B fuel shuffle. For each fuel shuffle, the red bar indicates the decrease in the respective isotope mass following the removal of the FAs, and the green bar indicates the increase in the respective isotope mass following the addition of the new FAs. These plots were created by comparing the EOC and subsequent BOC OpenMC materials files, selectively removing the isotope masses (corresponding to the FAs being removed) and adding the isotope masses (corresponding to the FAs being added).

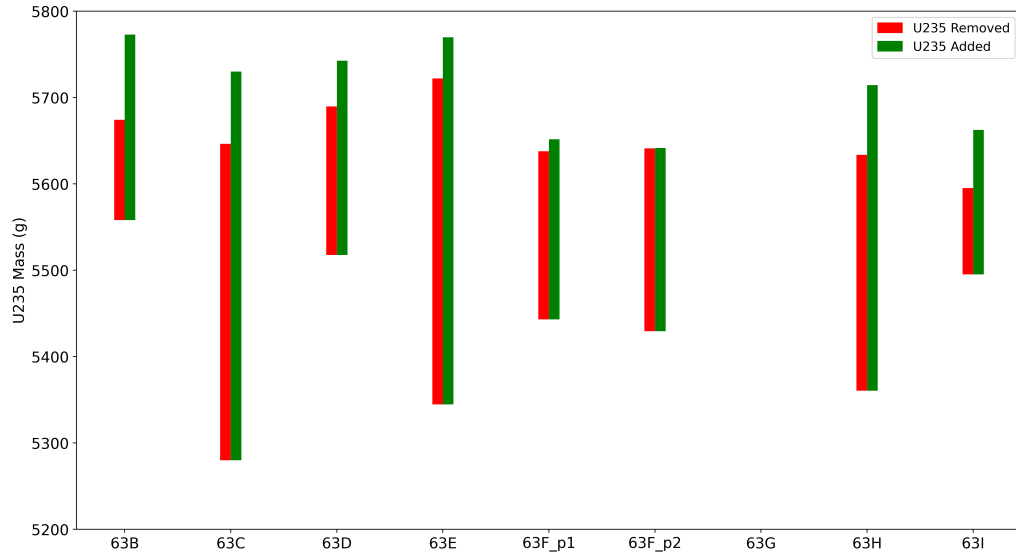


FIGURE 3.2: Change in U235 mass in the core, during each fuel shuffle. The core labels on the x-axis refer to the fuel shuffles at the beginning of that core cycle. The length of the red bar indicates the mass of U235 in removed FAs, and the length of the green bar indicates the mass of U235 added in new FAs.




FIGURE 3.3: Change in U238 mass in the core, during each fuel shuffle. The length of the red bar indicates the mass of U238 in removed FAs, and the length of the green bar indicates the mass of U238 added in new FAs.

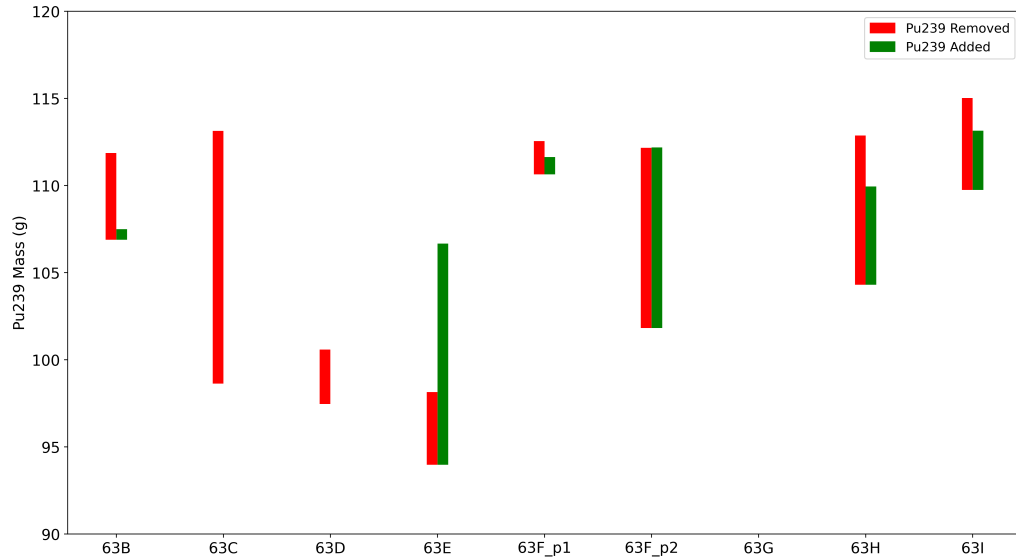


FIGURE 3.4: Change in Pu239 mass in the core, during each fuel shuffle. The length of the red bar indicates the mass of Pu239 in removed FAs, and the length of the green bar indicates the mass of Pu239 added in new FAs.




FIGURE 3.5: Change in Pu241 mass in the core, during each fuel shuffle. The length of the red bar indicates the mass of Pu241 in removed FAs, and the length of the green bar indicates the mass of Pu241 added in new FAs.

The overall change in mass is therefore the difference between the top of the red and

green bars. Note that the 63F_p2 to 63G fuel shuffle did not involve any removal or addition of fissile inventory. Figures 3.4 and 3.5 are quite similar, both show that in the majority of fuel shuffles, more plutonium is removed than is added, and the addition of plutonium indicates that a FA that was previously in the core has been re-added.

Figure 3.3 shows very discrete quantities of U238, that is, if the quantity of U238 added and removed in the 63A to 63B shuffle is a bar of relative length one, then the mass of U238 added in the 63B to 63C shuffle is of length two, and the mass of U238 removed in the 63B to 63C shuffle is of length three. One “relative length” in actual terms is the mass of U238 in one FA, and consequently, two or three “relative lengths” of U238 being added or removed indicates that two or three FAs were added or removed. Because a relatively small fraction of U238 in each FA undergoes neutron capture to eventually become plutonium isotopes, the quantities of U238 in fresh and highly depleted FAs are quite similar, which gives rise to the discretization described above. The figure also clearly shows the addition of the CIF (discussed in Section 1.3) in the 63B to 63C shuffle (one less FA is added than removed) and its removal in the 63D to 63E shuffle.

Because of the mass of U235 present in the core, and its comparable fission energies and cross sections to those of the plutonium isotopes, Figure 3.2 shows the clearest relationship to the overall change in WIC during fuel shuffles, as discussed below. Recall that burnup in the MNR context generally refers to the percent U235 depleted, “exit target burnup” (that is, the desired U235 depletion of a fully depleted FA when removed) is around 50%-60%. Consequently, fuel shuffles like 63A to 63B, where a highly-depleted FA is removed and a (nearly) fresh FA is added constitute the largest changes in WIC.

To determine the overall change in WIC during a fuel shuffle, Figures 3.2-3.5 can be effectively scaled by their respective fission energies and cross-sections and summed together, as shown in Figure 3.6. Note that the contribution of U238, despite its relative abundance in the core, to the WIC is minimal– generally 6 orders of magnitude smaller than the effects of plutonium isotopes or U235.




FIGURE 3.6: The effects of the constituent isotopes on the WIC during fuel shuffles. Each fuel shuffle is normalized such that 0% change is equal to the WIC of the EOC core without any removals or additions.

For each core, a 0% change in WIC refers to the state of the core before removals or additions of fissile inventory. Because the WICs of each core before the fuel shuffle is not identical, the exact magnitude of the change in magnitude cannot be compared to plots (for example, Figure 3.1, lower left) where multiple cores of reactor history are normalized to the same value.

Figure 3.6 shows that the total change in WIC is dominated by changes in U235 and that changes in U238 have a negligible effect on the WIC. Changes in plutonium tend to have a small but not negligible effect on the WIC. These results, namely the change in magnitude during refuelling, provide some perspective on the approximate change in WIC that would be incurred during clandestine activity, discussed in Section 3.3.3.

3.3.2 Change in Fissile Inventory and WIC During Cores

After considering the change in WIC during fuel shuffles, the second period over which the WIC changes is during cores, that is, depletion during regular operation. In the previous section, to calculate the change in isotope masses and the WIC, the EOC core from one core cycle was compared to the BOC core from the subsequent core cycle. Conversely, to generate these results the BOC and EOC of the same core cycle were compared to determine the change in mass and WIC during operation. The absolute change in fissile inventory is shown in Table 3.1.

Core	U235 (g)	U238 (g)	Pu239 (g)	Pu241 (g)
63B	-1.266e+02	-1.192e+01	5.573e+00	3.592e-01
63C	-4.036e+01	-3.621e+00	1.796e+00	9.309e-02
63D	-2.067e+01	-1.851e+00	6.386e-01	4.794e-02
63E	-1.322e+02	-1.244e+01	5.671e+00	3.782e-01
63Fp1	-1.078e+01	-1.020e+00	1.049e-02	3.180e-02
63Fp2	-1.470e+01	-1.387e+00	4.178e-01	4.468e-02
63G	-1.041e+01	-9.841e-01	2.611e-01	3.173e-02
63H	-1.191e+02	-1.127e+01	5.031e+00	3.514e-01

TABLE 3.1: Change in mass of the four isotopes that comprise the WIC over the duration of the core.

Because each core cycle spans a different amount of time, to determine the rate of change of the masses, these values must be normalized to some measure of the duration of the core cycle. Using operational data from MNR, the core duration (in hours and MWh) of each cycle was determined. To determine the duration of the reactor in units of MWh, a simple trapezoidal integration was performed over the reactor power time series. Tables 3.2 and 3.3 show the rates in units of [g/h] and [g/MWh].

Core	hours	U235 (g/h)	U238 (g/h)	Pu239 (g/h)	Pu241 (g/h)
63B	8.560e+02	-1.479e-01	-1.393e-02	6.510e-03	4.195e-04
63C	2.682e+02	-1.505e-01	-1.350e-02	6.696e-03	3.471e-04
63D	1.384e+02	-1.494e-01	-1.337e-02	4.614e-03	3.464e-04
63E	8.958e+02	-1.476e-01	-1.389e-02	6.331e-03	4.221e-04
63Fp1	7.343e+01	-1.468e-01	-1.389e-02	1.428e-04	4.330e-04
63Fp2	1.009e+02	-1.457e-01	-1.375e-02	4.141e-03	4.428e-04
63G	7.136e+01	-1.458e-01	-1.379e-02	3.659e-03	4.447e-04
63H	7.214e+02	-1.651e-01	-1.562e-02	6.975e-03	4.871e-04

TABLE 3.2: Mass rate of change of the four isotopes that comprise the WIC, in units of [g/h], for each core cycle.

Core	mwh	U235 (g/MWh)	U238 (g/MWh)	Pu239 (g/MWh)	Pu241 (g/MWh)
63B	2.566e+03	-4.933e-02	-4.646e-03	2.172e-03	1.400e-04
63C	7.966e+02	-5.066e-02	-4.546e-03	2.255e-03	1.169e-04
63D	4.139e+02	-4.994e-02	-4.471e-03	1.543e-03	1.158e-04
63E	2.689e+03	-4.916e-02	-4.628e-03	2.109e-03	1.406e-04
63Fp1	2.193e+02	-4.916e-02	-4.649e-03	4.782e-05	1.450e-04
63Fp2	3.004e+02	-4.891e-02	-4.618e-03	1.391e-03	1.487e-04
63G	2.118e+02	-4.912e-02	-4.646e-03	1.233e-03	1.498e-04
63H	2.158e+03	-5.520e-02	-5.223e-03	2.331e-03	1.628e-04

TABLE 3.3: Mass rate of change of the four isotopes that comprise the WIC, in units of [g/MWh], for each core cycle.

The agreement between individual rows of Tables 3.2-3.3 reflects that MNR is reliably operating at a nominal 3 MW power. Finally, using the information from the above tables, we can compute the absolute change and rate of change of the WIC over the different core cycles, as seen in Table 3.4.

Core	hours	mwh	WIC Change	WIC (/h)	WIC (/MWh)
63B	8.560e+02	2.566e+03	-1.146e-32	-1.339e-35	-4.468e-36
63C	2.682e+02	7.966e+02	-3.675e-33	-1.370e-35	-4.613e-36
63D	1.384e+02	4.139e+02	-1.905e-33	-1.377e-35	-4.602e-36
63E	8.958e+02	2.689e+03	-1.202e-32	-1.342e-35	-4.470e-36
63Fp1	7.343e+01	2.193e+02	-1.057e-33	-1.439e-35	-4.819e-36
63Fp2	1.009e+02	3.004e+02	-1.386e-33	-1.373e-35	-4.612e-36
63G	7.136e+01	2.118e+02	-9.896e-34	-1.387e-35	-4.672e-36
63H	7.214e+02	2.158e+03	-1.101e-32	-1.526e-35	-5.100e-36

TABLE 3.4: Absolute change and rate of change of the WIC in [g/h] and [g/MWh], for each core cycle

There is good agreement in the rate of change of the WIC between the various cores. An added benefit of this procedure is that it provides another means of confirming that the results of the OSCAR-5 core-follow can be trusted. Because MNR operates at a steady nominal power, MNR operations staff desire to have a consistent fissile inventory in the core, and therefore that the effects of fuel shuffles (adding new fissile inventory) and depletion should balance each other. And consequently changes in fissile inventory during fuel shuffles and during cores should be quite similar, as they are. Lastly, this method of splitting changes in WIC into changes during and between fuel shuffles can also be applied to determining if a change in detector count rate resulted from a change in fissile inventory; changes can be observed during, and in between fuel shuffles.

3.3.3 Change in WIC During Possible Clandestine Activity

After considering the changes in WIC that arise during normal operation, it is valuable to consider the possible changes to WIC that would be incurred during possible undeclared diversion of material. Two possible removals are postulated:

- The removal of Pu239 and Pu241 from the core.
- The removal of one FA of variable burnup.

In the first case, the removal of all Pu239 and Pu241 by some hypothetical means is considered. This scenario is postulated as plutonium presents the ideal material for the construction of a nuclear weapon, due to its large fission cross section and high fission energy. For each BOC and EOC core, the entirety of both plutonium isotopes was removed and the resultant change in WIC was compared to the WIC before removal. Figure 3.7 shows the change in WIC for each core. The mean plutonium removed was 113.45 g, and the mean change in WIC was $2.66\% \pm 0.16\%$. As will be shown below, this is equivalent to the removal of a partially depleted FA. The change in WIC is roughly the same magnitude (though in the opposite direction) as the change in WIC from refuelling, meaning that if the neutron detectors are sensitive enough to observe a fuel shuffle, then they will be sensitive enough to observe the removal of all plutonium. Additionally, it is worth noting that for every core, the change in WIC from removing the plutonium at EOC is larger than removing the plutonium at BOC, as there is always more plutonium after depletion than before (although for some shorter-duration cores is it difficult to see).




FIGURE 3.7: The change in WIC resulting from the removal of all plutonium from the core at EOC and BOC for each core.

In the second case, the removal of a single FA is considered. This removal is more consistent with possible clandestine activity, as FAs are added and removed from the core during fuel shuffles. In this case, for each BOC and EOC core, six tests were performed in which one specific FA was removed. In the first test, the FA with the burnup (measured by percent U235 depletion) with the closest to 0% burnup was removed. In the second test, the FA with the burnup closest to 10% was removed, and so forth in increments of 10% burnup until 50% burnup. The results of repeating this process for all cores are shown below in Figure 3.8.

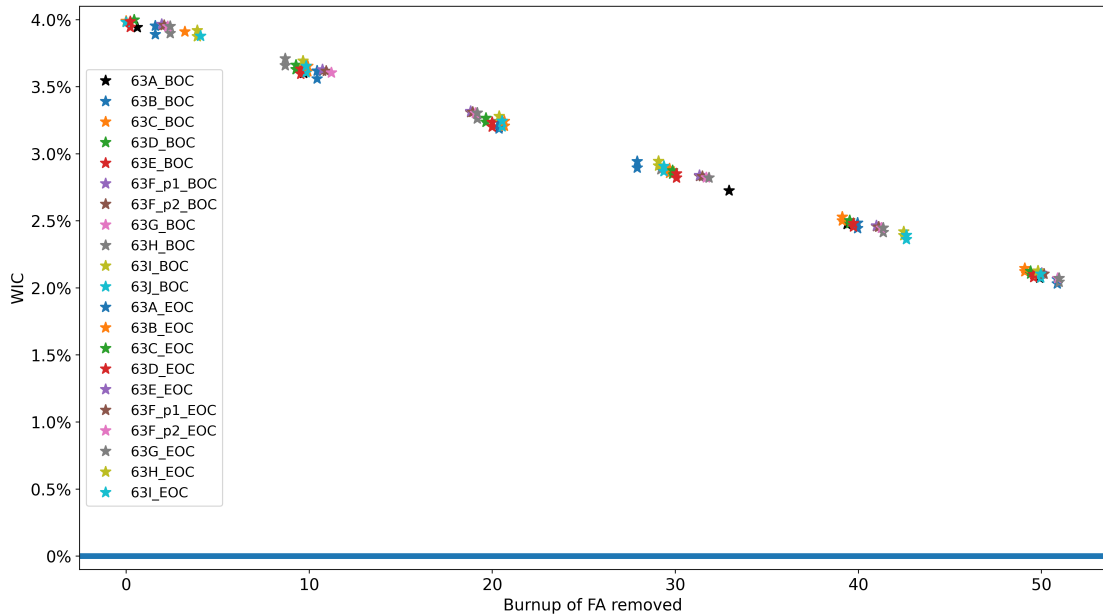


FIGURE 3.8: The change in WIC resulting from removing one FA with nearest burnup to 0%/10%/20%/30%/40%/50%.

Note that the burnup of the FA that was selected to be removed is never exactly the desired 0%/10%/20%/30%/40%/50% burnup, instead, it is the nearest burnup to that desired value. In the 0% case, there is no fresh FA in the core if none was added in the most recent fueling. Additionally, the EOC and BOC points of each core are coloured the same. Like in the plutonium removal scenario where the EOC change in WIC from withdrawing all plutonium was larger in magnitude than the BOC change, in cases where the same FA was selected in the EOC and BOC case, the stars of the same colour appear “stacked” one atop the other, however, the change in WIC is not the same; this is because the WIC at EOC is always less than at BOC, and therefore the change in removing an FA at BOC will always have a smaller relative effect. The changes in WIC were $3.94\% \pm 0.04\%$ / $3.63\% \pm 0.03\%$ / $3.26\% \pm 0.04\%$ / $2.86\% \pm 0.05\%$ / $2.45\% \pm 0.04\%$ / $2.09\% \pm 0.03\%$, respectively.

As in the plutonium removal case, removing one FA results in a change in WIC comparable to the change incurred during refuelling, but in the opposite direction, indicating that if the neutron detection system is sensitive enough to detect fuel shuffling events, it will be able to detect the removal of a FA.

Chapter 4

Results

4.1 Introduction

In this chapter, the results from the prior two chapters are combined to determine if it is possible to correlate changes in WIC and normalized neutron counts. It is worth reviewing the results from the previous two chapters, such that the exact nature of the data being compared is better understood. A flow chart showing the procedure applied to the counts data in Chapter 2 is shown below in Figure 4.1, and a sample of the resulting data is shown in Figure 4.2.

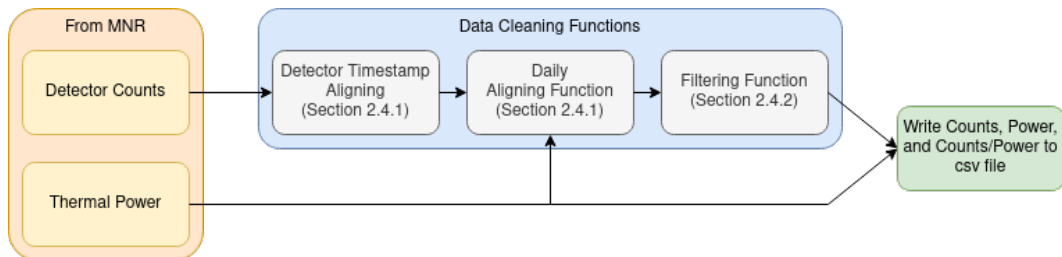


FIGURE 4.1: Flow chart showing “data cleaning” procedure applied in Chapter 2.

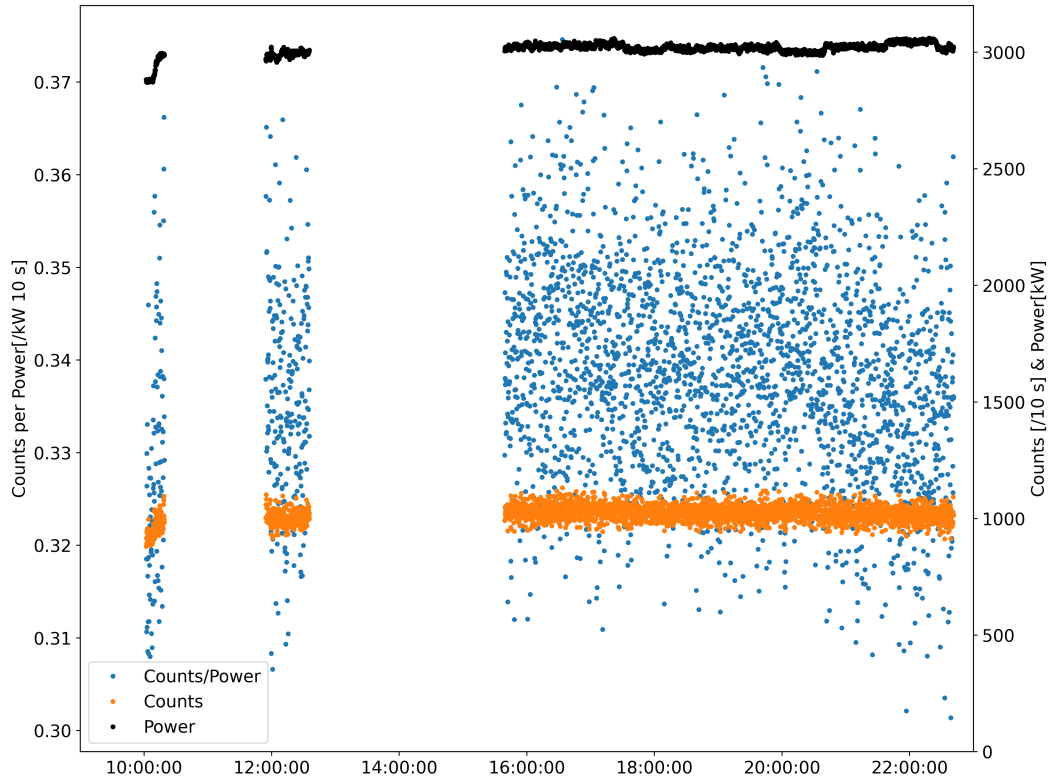


FIGURE 4.2: Time series plot of a day of filtered count, power, and counts per unit power data.

In Chapter 3, a time series of WIC values (typically four or five per day) was created from the results of OSCAR5 depletion simulations. The entire time series is shown in Figure 4.3.




FIGURE 4.3: Time series plot of all WIC data.

From 1.4, recall that counts per unit power and the WIC are proportional, and therefore for clarity they can be normalized. Although they can be normalized by any arbitrary scaling, because this work employed two neutron detectors, they should be normalized in such a way that the results from the two detectors can be compared. In this case, the mean counts per unit power of the respective detector, over the period when both detectors were in their long-term observation locations (locations 5 and 6) were used to normalize the data. Detector NE02A was placed at Location 6 on September 9th, 2021, and removed January 25th, 2022. Therefore both detectors were normalized to their respective means between September 9th and January 25th. Likewise, the WIC was normalized to the mean of the series between September 1st and April 15th. The period over which the detectors were in their long-term locations spanned six cores, 63C to 63H, over approximately eight months.

4.2 Weekly Periodicity of Detector Data

Because MNR does not always operate on a consistent Monday to Friday schedule, the definition of a “week” for this work refers to a continuous series of days without a shutdown day. As an example, if (instead of shutting down on Saturday and Sunday) the reactor were operating on Sunday, it would become day 1 of the “week” and Monday

would become day two. The term “reactor week” is used in the subsequent analysis to refer to the aforementioned “weeks”. The use of reactor weeks instead of normal weeks is important to the analysis below.

Before comparing detector data to the WIC, there are trends in the detector data that must be addressed. By examining the daily detector data, it can be seen that there is a clear weekly periodicity to the detector data. Data from earlier in the week tend to have an elevated count rate, on the order of 5% above the count rate later in the week. Moreover, it appears as though data from midday tends to have a higher count rate than in the morning or the evening. Figure 4.4 shows this weekly periodicity for all reactor weeks. That is, this plot features all reactor weeks on the same plot, where every week is normalized to the mean count of that week. Because the reactor does not startup at the same time every day, every day of data was aligned such that the startup is at 00:00. Finally, all the individual days of data were averaged into hours, and the mean of that hour is plotted below. The black line shows the average of the hourly means for all data with any counts within that hour range. Figure 4.5 shows a flow chart of the procedure to produce Figure 4.4

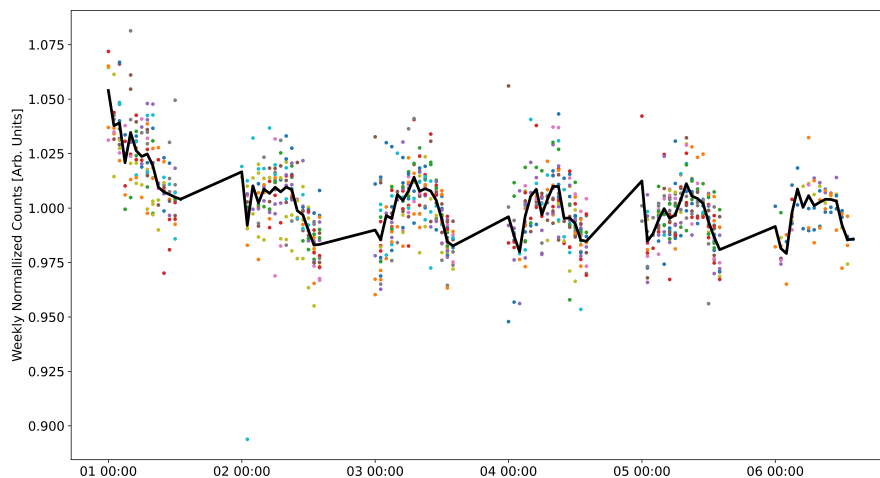


FIGURE 4.4: Plot of all reactor weeks normalized to their respective weekly mean counts. All days are aligned so that startup is at 00:00, and averaged hourly. The black line shows the average of the means for that hour.



FIGURE 4.5: Flow chart showing procedure to produce Figure 4.4.

Before discussing Figure 4.4 further, it is important to understand why the weekly periodicity must be accounted for: Consider a hypothetical two weeks of data, where in the first week, there are many data points at the beginning of the week, and fewer data points at the end of the week (recall that the availability of the detector data is limited to whenever the beam ports are not in use), and in the second week, there are fewer data points at the beginning of the week, and more at the end of the week. Based on Figure 4.4, if the data were to be averaged weekly, then the mean of the first week of data would be significantly higher than the second week, simply because of the limited availability of the detector data.

From conversations with MNR staff, supervisors, and colleagues, it was agreed that weekly periodicity appears similar to the xenon concentration caused by the day-night operation cycle at MNR. MacConnachie discusses the use of a correction factor in [10], similarly, a correction factor will be employed to attempt to allow for the comparison between days. The xenon coefficient quantifies the effect of the xenon concentration on the count rate of the detectors, relative weekly average count rate. For each 10 s data point, a xenon coefficient can be created using Equation 4.1, where C is the number of counts (per 10 s time bin), and ϕ_{Xe} is the xenon coefficient.

$$C_{corrected} = \frac{C}{\phi_{Xe}} \quad (4.1)$$

To estimate the xenon concentration at MNR, the Bateman equations for iodine and xenon were coded and time stepping was performed using the classic fourth-order Runge-Kutta method as described in [21]. To fit the xenon transient to the detector periodicity, the detector data was resampled to 15 minutes, and the number of 10 s in each 15-minute sample was calculated. The SciPy python package non-linear least squares method curve fitting function was used over the entire final observation period for both detectors, with the flux (used as input to the Bateman equation), and constants to scale the magnitude of the oscillations as fit parameters [22]. The 15-minute resampling was performed to reduce computation time, and the number of 10 s points in each 15-minute sample were used as weights in the fitting function, using $w_i = 1/\sqrt{n_i}$. The flux

was 1.71×10^{13} (in good agreement with typical fluxes from OSCAR simulations, which were approximately 1.5×10^{13} , and within the acceptable range of “high 1×10^{12} to low 1×10^{13} , no larger than 6×10^{13} ” from conversations with Dr. MacConnachie). After the fit parameters were determined, the Bateman function was used again, with 10 s time stepping instead of 15 minutes, to determine the xenon correction factor at each point.




FIGURE 4.6: Plot of a sample xenon transient created using the RK4 model from [21]. The final decay is approximately 56 hours, which is equal to a typical two day weekend shutdown.

Figure 4.7 shows a flow chart detailing the procedure to generate the xenon coefficients.

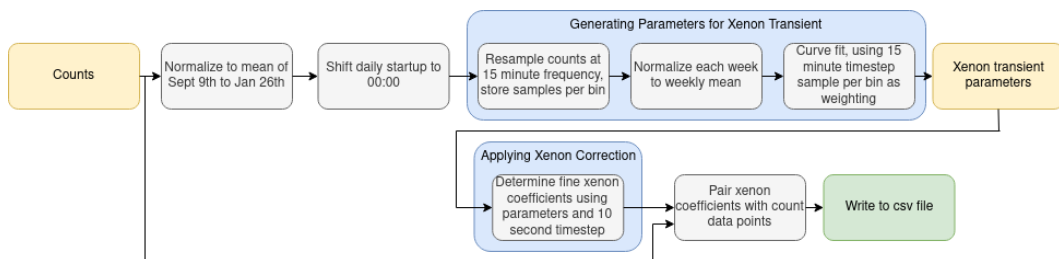


FIGURE 4.7: Flow chart showing procedure to generate xenon coefficients.

There is one complication associated with this method. First, to create Figure 4.4, each reactor week was normalized to the weekly mean. However, in doing so, the weeks with more data earlier in the week (when the signal is higher) have larger weekly means, and consequently, after being normalized, the week of data would be lower than if no normalization had been performed. In the following sections, the xenon correction is

applied after discussing the results of the uncorrected data. Despite the aforementioned complications, the xenon correction factor can reduce the variance between weekly counts significantly.

4.3 Comparison Between WIC and Uncorrected Normalized Counts

4.3.1 Overview

Figure 4.8, shows a time series, with normalized counts per unit power and WIC on the same axis. Both the count and WIC data were first grouped by the corresponding core, and then by day. Each point represents one day of data, with the standard deviation shown. NE0043 data is presented as dots, while NE02A data is presented as crosses and the WIC is shown as stars. Figure 4.9 shows the procedure to create Figure 4.8.

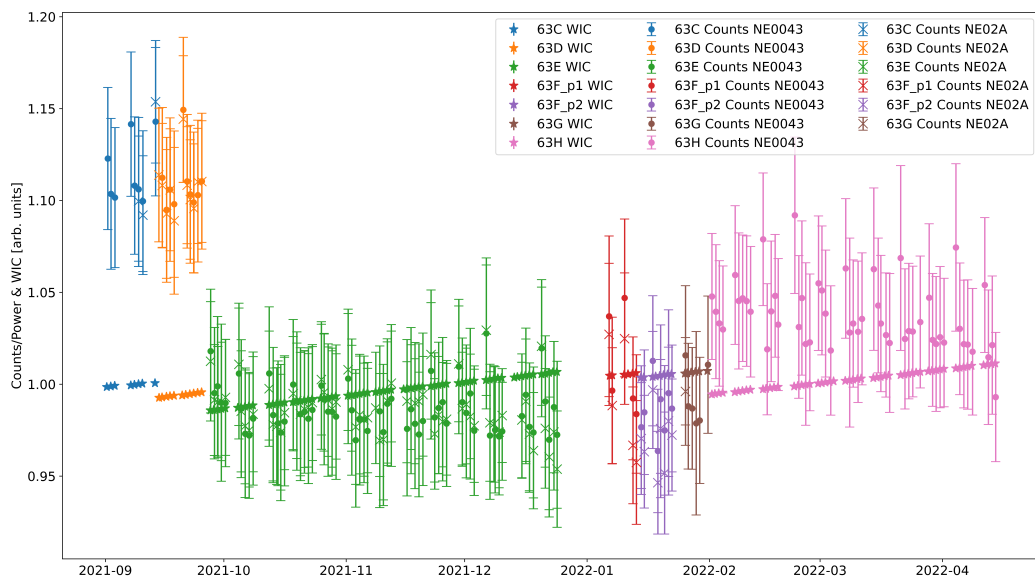


FIGURE 4.8: Time series of normalized counts per unit power and WIC data on the same axis. NE0043 data are marked by circles, while NE02A data are marked by crosses. WIC data are marked by stars.




FIGURE 4.9: Flow chart showing procedure to produce Figure 4.8.

Figure 4.10, shows the residuals of the daily means between the two detectors over the period where both detectors were in their long-term observation locations. The detectors agree with each other to within 1% initially, however, later into 2021, the NE02A detector begins to observe proportionally more counts, followed by the NE0043 detector observing proportionally more counts.




FIGURE 4.10: Plot of residuals between the daily mean of the two detectors over the period when both detectors were in their long-term observation locations.

Only the NE0043 counts data will be used in this analysis, as the detector was in its long-term location for longer. Figure 4.11 contains a plot of normalized counts against WIC, where each point corresponds to a day of WIC and normalized count data.




FIGURE 4.11: Normalized NE0043 counts per unit power against WIC for cores 63C to 63H.

After examining Figures 4.8 and 4.11, the most evident feature is the increased counts during the 63C and 63D cores. During both cores, the CIF (discussed in Section 1.3) assembly was in the core at site C5. It is possible that the change in counts during refuelling was caused by the removal of the CIF site, this possibility is discussed further below.

4.3.2 Comparison During Cores

Figures 4.12-4.18 show the counts and WIC data averaged by day. The data is presented as a time series, like 4.8, and as counts against WIC plots, like Figure 4.11, for each core period individually.




FIGURE 4.12: Left: Daily averaged time series plot of normalized counts per unit power, and WIC during the 63C core. Right: Daily averaged normalized counts per unit power against WIC during the 63C core.




FIGURE 4.13: Left: Daily averaged time series plot of normalized counts per unit power, and WIC during the 63D core. Right: Daily averaged normalized counts per unit power against WIC during the 63D core.




FIGURE 4.14: Left: Daily averaged time series plot of normalized counts per unit power, and WIC during the 63E core. Right: Daily averaged normalized counts per unit power against WIC during the 63E fuel core.




FIGURE 4.15: Left: Daily averaged time series plot of normalized counts per unit power, and WIC during the 63F_p1 core. Right: Daily averaged normalized counts per unit power against WIC during the 63F_p1 core.




FIGURE 4.16: Left: Reactor week averaged time series plot of normalized counts per unit power, and WIC during the 63F_p2 core. Right: Reactor week averaged normalized counts per unit power against WIC during the 63F_p2 core.




FIGURE 4.17: Left: Daily averaged time series plot of normalized counts per unit power, and WIC during the 63G core. Right: Daily averaged normalized counts per unit power against WIC during the 63G core.




FIGURE 4.18: Left: Daily averaged time series plot of normalized counts per unit power, and WIC during the 63H core. Right: Daily averaged normalized counts per unit power against WIC during the 63H core.

4.3.3 Comparison During Fuel Shuffles

In this section, normalized counts per unit power data are compared before and after fuel shuffles. In Figures 4.19-4.24, three reactor weeks of data before and after the fuel shuffle are used in the comparison below. If less than three reactor weeks of data were available, the entire data set was used. The error in y is the error on the daily mean, and the error in x is the standard deviation of the WIC values.




FIGURE 4.19: Normalized counts per unit power against WIC for the three weeks before and after fuel shuffle 63C to 63D.




FIGURE 4.20: Normalized counts per unit power against WIC for the three weeks before and after fuel shuffle 63D to 63E.




FIGURE 4.21: Normalized counts per unit power against WIC for the three weeks before and after fuel shuffle 63E to 63F_p1.




FIGURE 4.22: Normalized counts per unit power against WIC for the three weeks before and after fuel shuffle 63F_p1 to 63F_p2.




FIGURE 4.23: Normalized counts per unit power against WIC for the three weeks before and after fuel shuffle 63F_p2 to 63G.




FIGURE 4.24: Normalized counts per unit power against WIC for the three weeks before and after fuel shuffle 63G to 63H.

4.3.4 63E to 63H Aggregated Comparison

Finally, in Figure 4.25, longer averaging windows are applied. Cores 63C and 63D were excluded because of the significantly larger normalized counts than subsequent cores, possibly a product of the CIF being in the core (discussed below in Section 4.5). For this reason, the counts and WIC were renormalized to the mean of the data shown on the plot. Core 63E was averaged into 4 groups, with the first group composed of the first four weeks, and each subsequent group being three reactor weeks. Core 63F_p1, 63F_p2, and 63G were all grouped individually. Core 63H was averaged into three groups, the first and last groups were composed of four reactor weeks, while the middle group was 3 reactor weeks. The error in y is the error on the daily mean, and the error in x is this standard deviation of the WIC values. A bivariate linear regression was performed using the York method from [23]. The slope of the resulting fit was 0.80 ± 0.48 , and the intercept was 0.21 ± 0.48 . These results are discussed below in Section 4.5, after presenting the xenon corrected results. A flow chart of the aforementioned procedure is shown below in Figure 4.26.




FIGURE 4.25: Normalized counts per unit power against WIC for cores 63E to 63H, with data averaged over three or four weeks, if available.




FIGURE 4.26: Flow chart of aggregation plotting procedure.

4.4 Comparison Between WIC and Corrected Neutron Counts

4.4.1 Overview

This section, which has the same structure as the previous section, presents the same results as the previous section, but with the xenon correction applied to the normalized counts per unit power data, as described in Section 4.2. Figure 4.27 shows the same time series data, with the aforementioned correction applied. To produce this plot, the procedure shown in Figure 4.7, followed by Figure 4.9 were applied.




FIGURE 4.27: Time series of xenon corrected normalized counts and WIC on the same axis. NE0043 data are marked by circles, while NE02A data are marked by crosses. WIC data are marked by stars.

The most prominent feature is that there is no longer a higher mean on the first day of the week every week (by chance, some weeks do have a higher daily mean on the first day, however, the persistent periodicity is no longer present). To compare the effect of the xenon correction, select parts of the time series are presented (as opposed to the entire series, for readability). Figures 4.28 and 4.29 show time series plots of cores 63C and 63D, and 63F_p1 to 63G, respectively. The plots show the data from the NE0043 detector, both with and without the xenon correction. The uncorrected data is shown as faded data points, with smaller end caps, while the xenon corrected data are solid points with larger end caps.




FIGURE 4.28: Time series data of the 63C and 63D cores showing both corrected and uncorrected data for both detectors.




FIGURE 4.29: Time series data of cores 63F_p1, 63F_2, and 63G showing both corrected and uncorrected data for both detectors.

From the plots, it is clear that the variance between reactor weeks during the same core decreases significantly (this decrease is quantified below). Trying to attribute the changes in count rate between weeks to changes in fissile inventory is difficult because

the changes in fissile inventory (and therefore in WIC) between weeks are minute (see Chapter 3).

Figure 4.30, shows the same plot as 4.11, but with xenon corrected count data.




FIGURE 4.30: Xenon corrected normalized counts per unit power against WIC for cores 63C to 63H.

4.4.2 Comparison During Cores

Figures 4.31-4.37 show xenon corrected time series and counts against WIC plots for the individual cores. The error in y is the standard deviation of the daily count rate, and the error in x is the standard deviation of the WIC.




FIGURE 4.31: Left: Reactor week averaged time series plot of xenon corrected normalized counts per unit power, and WIC during the 63C core. Right: Reactor week averaged xenon corrected normalized counts per unit power against WIC during the 63C core.




FIGURE 4.32: Left: Reactor week averaged time series plot of xenon corrected normalized counts per unit power, and WIC during the 63D core. Right: Reactor week averaged xenon corrected normalized counts per unit power against WIC during the 63D core.




FIGURE 4.33: Left: Daily averaged time series plot xenon corrected of normalized counts per unit power, and WIC during the 63E core. Right: Daily averaged xenon corrected normalized counts per unit power against WIC during the 63E core.




FIGURE 4.34: Left: Daily averaged time series plot of xenon corrected normalized counts per unit power, and WIC during the 63F_p1 core. Right: Daily averaged xenon corrected normalized counts per unit power against WIC during the 63F_p1 core.




FIGURE 4.35: Left: Daily averaged time series plot of xenon corrected normalized counts per unit power, and WIC during the 63F_p2 core. Right: Daily averaged xenon corrected normalized counts per unit power against WIC during the 63F_p2 core.




FIGURE 4.36: Left: Daily averaged time series plot of xenon corrected normalized counts per unit power, and WIC during the 63G core. Right: Daily averaged xenon corrected normalized counts per unit power against WIC during the 63G core.




FIGURE 4.37: Left: Daily averaged time series plot of xenon corrected normalized counts per unit power, and WIC during the 63H core. Right: Daily averaged xenon corrected normalized counts per unit power against WIC during the 63H core.

4.4.3 Comparison During Fuel Shuffles

This section compares, xenon corrected normalized counts per unit power and WIC data before and after fuel shuffles. In Figures 4.38-4.43, three reactor weeks of data before and after the fuel shuffle are used in the comparison below. If less than three reactor weeks of data were available, the entire data set was used. The error in y is the error on the daily mean, and the error in x is the standard deviation of the WIC values.




FIGURE 4.38: Normalized counts per unit power against WIC for the three weeks before and after fuel shuffle 63C to 63D.




FIGURE 4.39: Normalized counts per unit power against WIC for the three weeks before and after fuel shuffle 63D to 63E.




FIGURE 4.40: Normalized counts per unit power against WIC for the three weeks before and after fuel shuffle 63E to 63F_p1.




FIGURE 4.41: Normalized counts per unit power against WIC for the three weeks before and after fuel shuffle 63F_p1 to 63F_p2.




FIGURE 4.42: Normalized counts per unit power against WIC for the three weeks before and after fuel shuffle 63F_p2 to 63G.




FIGURE 4.43: Normalized counts per unit power against WIC for the three weeks before and after fuel shuffle 63G to 63H.

4.4.4 63E to 63H Aggregated Comparison

Figure 4.44 shows the same larger averaging windows as Figure 4.25, applied to the xenon corrected data. The reactor weeks were aggregated exactly as before. The slope of the linear regression is 1.43 ± 0.76 , and the intercept is -0.42 ± 0.75 . These results are discussed below in Section 4.5.




FIGURE 4.44: Xenon corrected normalized counts per unit power against WIC for cores 63E to 63H, with data averaged over three or four weeks, if available.

4.5 Discussion

4.5.1 Cores 63E to 63H Aggregation

Figures 4.25 and 4.44 show the fundamental problem with this work: The changes in fissile inventory during regular operation at the MNR are small, both in comparison to the change in detector counts and in comparison to changes in fissile inventory in other reactors. Consider the change observed at the NRU by Dr. van der Ende and colleagues, shown in Figure 4.45:

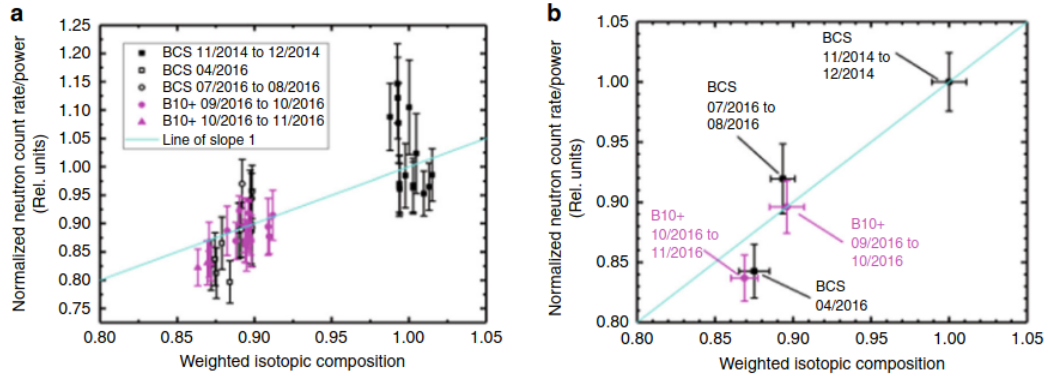


FIGURE 4.45: Normalized counts per unit power against WIC, data from NRU experiment [5].

The change in WIC observed at NRU is on the order of 10%, whereas at MNR it is on the order of 2%. Quantitatively, this effect is reflected in the uncertainty in the linear regression applied to both data sets. The linear regression with standard error at NRU was 1.08 ± 0.13 and -0.08 ± 0.12 for the slope and intercept respectively, whereas the linear regression with standard error for the uncorrected data was 0.80 ± 0.48 and 0.21 ± 0.48 for the slope and intercept respectively, and 1.43 ± 0.76 and -0.42 ± 0.75 for the slope and intercept for the corrected data. With these large uncertainties, it is not possible to correlate changes in fissile inventory. For example, to a 95% confidence interval, the slope estimators are -0.31 to 1.90 for the uncorrected case, and -0.32 to 3.17 for the corrected case.

After examining the standard error on the uncorrected and corrected regressions, it might seem as though the xenon correction did not have the desired effect of reducing the variance in count rate means. However, the average standard error on the daily mean in the uncorrected aggregate groups was 0.016, and the average standard error in the corrected aggregate groups was 0.09, an almost two-fold reduction in average standard error. The larger uncertainty can only be attributed to the means themselves, after xenon correction.

In particular, the cores 63F_p1, 63F_p2 and 63G all appear lower in the corrected plot than the uncorrected plot, likely because those shorter cores had more days earlier in the week, which lead to those points being corrected downwards. In the regression, this likely had the effect of lowering the y-intercept estimator and increasing the slope estimator to compensate for more data being clustered in the lower-right quadrant of the plot.

The large variance in counts occurs both during fuel shuffles and cores. The average change in WIC during cores 63E–63H was 0.35% when averaged over the first and last three weeks (as performed in Section 4.4.3), and the change in counts was 2.99% and 2.31% for the uncorrected and corrected data. The average reactor week over reactor week change in WIC during cores 63E–63H was 0.15%. The average reactor week over

reactor week change in counts was 0.85% and 0.61% and the uncorrected and corrected data over the same period.

The omission of cores 63C and 63D from Sections 4.4.3 and 4.4.4 was on the basis that their count rates were significantly higher than the subsequent cores, which coincided with the CIF being removed from the core. The decrease in counts after the CIF was removed is not predicted by the OSCAR core-average thermal flux shown in Figure 4.46:




FIGURE 4.46: WIC and core average thermal flux from OSCAR.

During the 63B to 63C fuel shuffle, the flux decreased more than during other fuel shuffles as the CIF was inserted, and during the 63D to 63E fuel shuffle, the flux increased as the CIF was removed (the only fuel shuffle over which the flux increased). One possible explanation for this phenomenon is that the CIF being in-core leads to lower than typical flux in the center of the core, which leads to a broader flux shape to compensate for that effect. If the detectors are more sensitive to the flux at the side of the core that the detectors are nearest to, then the broader flux shape may increase the counts observed by the detectors when the CIF is in-core. Without more data, this explanation cannot be confirmed.

Similarly, it can be seen in Figure 4.46 that there is a slight variation in the ratio of OSCAR predicted flux to WIC when comparing the last week of core 63B to the first

week of core 63E, which is not predicted by the SND model. In the last week of core 63B, the ratio of flux to WIC was 0.9965, whereas in the first week of 63E it was 0.9920, a difference of 0.45%. This difference may be due to OSCAR, or it could be an indication that changes in core configuration affect the ratio of flux to WIC that is not predicted by the SND model. This effect merits further investigation in subsequent work, both in simulation and experiment.

Similar to how fuel shuffling operations at NRU were able to be detected by neutron detectors due to their location relative to the core and the online refuelling capability of NRU, the presence of the CIF is an event with a neutron signature that can be observed in the detector. Because only one removal of the CIF was observed, this theory could not be confirmed, however, it does speak to the need for long-term placement of the neutron detectors in future research, such that expertise working with that particular detector, and enough data to search for features in the particular neutron count signal from that reactor can be obtained.

It may have been possible to correlate comparatively small changes in WIC (like those at MNR) with normalized counts per unit power, without the problems with neutron counting caused by beam port use, and the weekly periodicity of neutron counting. Though it was not possible to correlate changes in WIC with changes in neutron counts at MNR over this period, other experiments at more suitable reactors may well be able to establish the correlation. Some reasons why the SND method would be more suitable for SMRs are given in Section 4.5.3.

4.5.2 Aggregating with Smaller Groups

This section pertains to repeating the results of the previous section, this time with smaller aggregation groups. In the original aggregation, the longer cores were split into groups of 3-4 weeks, while the shorter cores contained only one group of 1-2 weeks. This has the effect of biasing the results towards the smaller groups, as the groups themselves are weighted equally (provided their errors are equal). In this section, the 63E core was split into 6 groups of 2-3 weeks instead of 4 groups, and the 63H core was split into 5 groups of 2-3 weeks, instead of 3. Figures 4.47 and 4.48 show the results applying the York regression to the the data with finer groups.




FIGURE 4.47: Normalized counts per unit power against WIC for cores 63E to 63H, with data averaged over two to three weeks, if available.




FIGURE 4.48: Xenon corrected normalized counts per unit power against WIC for cores 63E to 63H, with data averaged over two to three weeks, if available.

Most notably, the fact that adjusting the size of aggregation groups has such a

substantial effect on the statistics is indicative that the correlation between the two variables is weak in the first place. The 95% confidence estimators for the xenon corrected data are 0.43 to 1.24 and -0.24 to 0.57 for the slope and intercept, respectively. Table 4.1 shows the coefficient of determination for the data without errors, for the uncorrected and corrected data, with and without the finer groups. Although this statistic does not consider the error in the measurements, it does support that the correlation between the two variables is weak.

	Uncorrected	Corrected
Original Groups	0.043	0.019
Finer Groups	0.065	0.039

TABLE 4.1: FAPs used to test MNR experiment.

4.5.3 Standoff Neutron Detection and SMRs

In principle the MNR was a suitable choice for this work because of its “small” thermal power, however, the features that make neutron counting at MNR difficult would likely not exist at conventional SMR power plants (namely, interference from research and commercial activities and day/night operation). Below, the desirable qualities of SMRs are listed.

- **Neutron background and interference:** SMR designs intended to be deployed in remote sites would likely not experience the sort of neutron background that has been problematic in this research [24]. Similarly, power reactors in general do not have beam ports or other research activities that produce a significant neutron background.
- **Fuel shuffles:** The SMRs that are the furthest along in the licensing process in Canada and the US have more regular refuelling schedules, long enough to observe larger changes in WIC. For the Nuscale SMR, refuelling is every two years, and for the GE-Hitachi BWRX-300, refuelling is every one to two years [25][26].
- **Sealed cores:** Many SMRs, particularly smaller SMRs intended for remote applications, have sealed cores. Over years a much broader range of WIC values could be observed (see the MMR by Ultra Safe Nuclear Corporation in [27]).
- **Operation schedule:** Although some SMRs are often touted as being able to load follow (to adjust the operating power to balance grid demand), all SMRs would likely be operated continuously, instead of the day/night schedule at MNR.

For these reasons, the SND method still has potential applicability to SMR designs, despite the results of this work.

Chapter 5

Malicious Interference

5.1 Overview

This chapter concerns “malicious interference” attacks on the standoff neutron detector system. A “malicious interference” attempt, in this context, is any attempt to mask clandestine activity by altering the neutron signal to the detector, by physically adding shielding or exogenous neutron sources. Two example scenarios are:

- Attempts to spoof the neutron detectors while the reactor is operated at a higher power level (to increase the production of plutonium, for example), by placing shielding in front of the neutron detectors.
- Attempts to spoof the neutron detectors while the reactor is shut down to unload fuel, by placing exogenous neutron sources near the detectors to emulate the neutron signature of the reactor.

One of the objectives of this thesis is to determine the resilience of the standoff neutron detection system against these malicious interference attacks. It is assumed that during the act of adding or removing exogenous material, some deviation in neutron count rate will be incurred and that such a deviation can be detected by a suitably designed algorithm. Such an algorithm should only be triggered in the event of a malicious interference attack, and not as a consequence of normal background fluctuations in neutron count rate.

An algorithm to prevent malicious interference is not the same algorithm that is used to compare the neutron count rate to the WIC to determine if a deviation from the expected count rate per unit power is observed, which would indicate that the fissile inventory of the reactor has been altered. The malicious interference detection algorithm would attempt to ensure that the neutron count rate has not been altered and that the conclusions of the standoff neutron detector method can be trusted.

This section describes the background and research into the development of the malicious interference detection algorithm, before describing the algorithm itself, the experiments to test experiment, and the results of those experiments.

5.2 Background

Because this technique of using neutron detectors at standoff distances for safeguards purposes is new, there exists no research directly on this question. Therefore, the first objective was to gather pertinent research that might inform the development of the malicious interference detection algorithm. The area of research that informed this work was radiation portal monitors (RPMs) found at border crossings to search for illicit radioactive material. RPMs usually consist of combined gamma and neutron detector systems. These systems must be able to discriminate between background radiation (primarily from cosmic rays), and nuclear material, and should have some capability of distinguishing between illicit radioactive materials and so-called NORMs (Naturally Occurring Radioactive Materials, like radioactive materials found in industrial waste). Such systems generally employ dynamic thresholding algorithms like the one shown in Equation (5.1) [28]:

$$T = B + m\sqrt{B} + C \quad (5.1)$$

Where T is the threshold counts, B is the background counts, m is a variable to tune the sensitivity of the threshold and C is some constant (like one or two) to prevent the threshold from being triggered if $B = 0$. When the number of counts is sufficiently high the Poisson distribution can be approximated as a Gaussian distribution, the term $m\sqrt{B}$ corresponds to m standard deviations (the standard deviation of the Poisson distribution is $\sigma = \sqrt{B}$) above the mean count, B , and hence this method of thresholding is only applicable for situations in which enough counts are detected. This is not true in general for standoff neutron detector systems, for example, B. van der Ende et al. observed 2-10 cps at the NRU reactor. With such low counts, an alternative method of detection must be used in order to permit low numbers of counts while still allowing for sufficient time resolution to be able to sense the deviation in counts incurred while adding exogenous material. A low count method relying on Poisson statistics, gross counting using Poisson statistics, was developed for radiation portal monitors by Blessinger and York, and is described below before being adapted for use with the standoff neutron detection system [29].

5.3 Theory

It is well understood that counting particles is a Poisson process and, as noted above, that the resulting Poisson distribution can be approximated as a Gaussian distribution at high counts. The Poisson distribution is shown below in Equation (5.2), where μ is the mean number of counts observed over some sample time t , and $P(\mu, x)$ is the probability of observing x counts.

$$P(\mu, x) = \frac{\mu^x e^{-\mu}}{x!} \quad (5.2)$$

Because the Poisson distribution is only defined for integer counts (because it is impossible to observe, say, 2.5 counts over a given period), the probability of observing more, or less than a given number of counts is computed by taking a sum of the counts above and below the given counts, as shown in Equation (5.3) for counts below some arbitrary threshold k_L and Equation (5.4) for counts above some arbitrary threshold k_U .

$$P(\mu, x) = \sum_{x=0}^{k_L} \frac{\mu^x e^{-\mu}}{x!} \quad (5.3)$$

$$P(\mu, x) = \sum_{x=k_U}^{\infty} \frac{\mu^x e^{-\mu}}{x!} \quad (5.4)$$

These thresholds are shown visually in Figure 5.1:




FIGURE 5.1: Poisson distributions ($\mu = 4.5$) with thresholds $k_L = 2$ (top) and $k_U = 8$ (bottom).

One consideration specific to the malicious interference detection algorithm that is not found in the RPM context is the need for a low count rate alarm. In the RPM context, because all that is required is to detect counts above background, there is no

need to set a low count threshold, as it would not indicate the presence of NORMs or other nuclear materials. In the context of malicious interference, the low count alarm could correspond to adding exogenous shielding, which would decrease the efficiency of the detector system and reduce the counts observed.

With these sums, the False Alarm Probability (FAP) can be defined. The FAP is the probability that, absent malicious interference, an anomalously high or low number of counts are detected. Similar to m from Equation (5.1), the FAP is a variable that will adjust the sensitivity of the detection algorithm.

Because the algorithm must be sensitive to both low and high counts, it must have low and high count FAPs, and because the Poisson distribution is discrete, in practice the FAPs will never correspond exactly to the desired values. Instead for FAP_L , the maximum k_L that satisfies Equation (5.5) is selected, and for FAP_U , the minimum k_U that satisfied Equation (5.6) is selected.

$$FAP_L \geq \sum_{x=0}^{k_L} \frac{\mu^x e^{-\mu}}{x!} \quad (5.5)$$

$$\begin{aligned} P_{total} &= \sum_{x=0}^{k_U-1} \frac{\mu^x e^{-\mu}}{x!} + \sum_{x=k_U}^{\infty} \frac{\mu^x e^{-\mu}}{x!} = 1 \\ \implies FAP_U &\geq 1 - \sum_{x=0}^{k_U-1} \frac{\mu^x e^{-\mu}}{x!} \end{aligned} \quad (5.6)$$

Note that Equation (5.5) implies that the sampling duration must be long enough to ensure that the probability of observing low counts is sufficiently low such that a reasonable FAP_L that satisfies Equation (5.5) can be defined. Graphically this corresponds to Figure 5.2, where μ is too low to define a reasonable FAP_L




FIGURE 5.2: Poisson distribution with $\mu = 1.5$, although k_L is set to 0, the probability of observing 0 counts is still too high to set a reasonable FAP_L .

To summarize, there are several constraints and variables: The FAPs, and their corresponding thresholds, k_U and k_L , and the sample time t . Additionally, the algorithm should be equally sensitive to low and high count rates, which requires $FAP_L \approx FAP_U$ (for the Gaussian threshold shown in 5.1, the term $m\sqrt{B}$ could be simply modified to $\pm m\sqrt{B}$ to incorporate a low threshold alarm).

Blessinger and York also suggest the use of a decision time T , that is, a characteristic time over which a decision needs to be made as to whether or not the alarm will trigger, for use in a running sum. That is if the decision time is T , and the sampling time t , then T/t (rounded down) samples can be used in the running sum. The sum is shown in Equation 5.7:

$$S_i = \sum_{k=-(T/t-1)}^0 s_{i+k} \quad (5.7)$$

In other words, all of the samples prior to i within the decision time T are included. The background (B in Equation (5.1)) is taken as the mean of the samples prior the samples in the decision time, the length samples included in the background is dependent on how rapidly the neutron background changes.

This method is shown in Figure 5.3, using simulated data (discussed in section 5.4). A pulse with Poisson distributed noise lasting 10 minutes was added starting at 45 minutes on top of a Poisson background with a mean of 50 counts. The decision time is set to 3 minutes, centred at an arbitrary point i , and the background sample time is 30 minutes.




FIGURE 5.3: Simulated background with transient added to show the points used in the background and decision time samples.

The pseudocode for the thresholding algorithm is shown below:

```

16 Function MAL_threshold
17     Pass In: count_series, background_time, decision_time, FAP
18
19     Set t to sample time count_series
20     Set N to decision_time/t rounded down
21     Set parser_time to the end of background_time
22
23     WHILE parser_time <= end_time
24         Set B to mean of count_series from (parser_time -
        ↪ decision_time - background_time to (parser_time -
        ↪ decision_time)
    
```

```
25         Set S to mean of count_series from (parser_time) to  
        ↪ (parser_time - decision_time)  
26  
27         Set K_L to Poisson ppf, p=FAP  
28         Set K_H to Poisson ppf, p=1-FAP  
29  
30         IF S <= K_L or S >= K_U  
31             Append time, S, B, K_L, K_U
```

The simulated plot also makes clear that for this method to be as sensitive as possible, the decision time should be exactly the time span of the deviation in the count rate. If the decision time is too short, it is possible that too few points will be sampled to observe the deviation in the count rate incurred by the malicious interference. Conversely, if the decision time is too long, samples that do not correspond to interference will be included in the running sum and decrease the signal from the interference.

5.4 Simulation Tool

5.4.1 Am-Be Source Measurements

Originally, exogenous neutron sources were used to attempt to validate the thresholding algorithm. An Am-Be neutron source with an activity of 1 mCi was used to replicate leakage from the reactor, and the distance between the source and the detector was varied to adjust the number of counts observed. However, it was found that the testing environment (the McMaster Tandem Accelerator Building) was prone to abrupt changes in neutron flux, likely from other sources, or experiments in proximity to the testing area, that made testing the algorithm more complicated. Consequently, a simulation tool was created that allows for testing of the thresholding algorithm without the need to obtain actual data. A brief description of the simulation tool is provided, along with a simple experiment to validate the results of the simulation tool. Figure 5.4 shows a sample of the count rate observed with the Am-Be source present. During the period shown, there were no changes to the location of either the detector or neutron source, and the local environment was not changed.




FIGURE 5.4: A sample of the counts observed with the AmBe source present, in the McMaster Tandem Accelerator Building. The local environment was not changed during the course of the sample being collected.

Figure 5.5 shows a comparison between the background count rate, and the count rate with the AmBe source at approximately 2.8 m with various moderators and absorbers placed in front of the detector. “3-ply” refers to a large slab of wax blocks, boronated rubber, and cork, approximately 8 inches thick. Two sheets of $1/16$ inch thick cadmium foil were also used in this test. The figure shows that the tests using two sheets of Cd foil produce approximately a 10% change in detector counts, compared to the one sheet. The background sample taken during the same day (without the AmBe source present), shows the count rate more than doubling without an apparent cause.




FIGURE 5.5: A sample of the counts with the AmBe source placed at 2.8 m. “3-ply” refers to the wax, rubber and cork slab. The AmBe source was not present during the “background” sample.

5.4.2 Simulation Tool

The simulation tool is a series of python functions that allow for the creation of a background and a transient that can be added to the background. The background function takes as parameters the desired count rate (in counts/s), the desired sample mean, and the length of the background sample, and creates a series with Poisson distributed noise.

The tool allowed for the creation of a transient to represent the deviation from the observed count rate incurred in the process of adding interfering material. The transient generating function takes as parameters the background time series, the time of beginning and end of the transient, the shape of the transient, and the maximum height of the transient. The transient counts for each sample during the transient are computed (shown in orange below) and Poisson noise is applied to the function. The transient is then added or subtracted to the background data element-wise (that is, the Poisson noise is applied to both the background and the transient before they are added). Figure 5.6 shows a number of transients that were created.

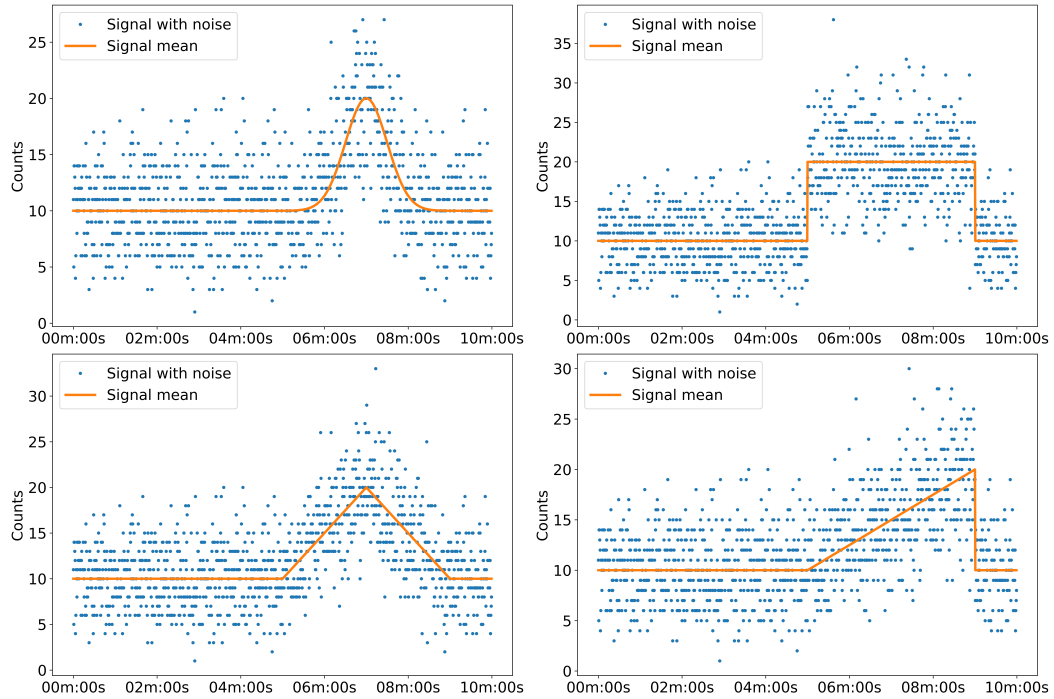


FIGURE 5.6: Multiple different shape function transients with a maximum height of 10 counts, on top of a background count rate of 10 counts, shown are the Gaussian pulse, square, triangle, and sawtooth functions.

In order to validate the simulation tool, a sample set of data was taken from the reactor when no beam ports were in use (when the signal to the detector was just leakage). The detector was left to generate a background for about 20 minutes before 1/8" of cadmium foil was used to shield the detector for approximately 30 seconds to create the transient. The signal is shown below, the complete series is shown on top, and just the transient is shown on the bottom. A longer background period was not possible due to beam port use.




FIGURE 5.7: Left: Complete time series of data taken at MNR, Right: Section of time series over which Cd foil was placed in front of the detector, between the detector and the reactor.

By examination, it appears as though the transient is best represented by the triangular function. If the transient can be represented by a triangular wave, the maximum height of the triangular wave is twice the mean of the wave. Therefore a simulated transient with a height of twice the mean of the experimental transient was created. The experimental transient and simulated transient are shown below.




FIGURE 5.8: Time series taken from MNR with simulated data overlaid.

Next, a variety of FAPs was tested against 20 hours of simulated background data, to observe if the alarm thresholds were triggered by background operation with the same count rate and mean of the MNR data. The four FAPs are shown in Table 5.1, they are the probabilities equivalent to 2σ , 3σ , 4σ and 5σ confidence. These FAPs are used below in further tests.

σ Confidence	FAP
2	0.954 499 736
3	0.997 300 203
4	0.999 936 657
5	0.999 999 426

TABLE 5.1: FAPs used to test MNR experiment.

It was found that for the 20-hour background sample, all FAPs except for the 5σ triggered the alarm, and all FAPs were able to sense the deviation from the background of both the simulated and experimental samples. In the following section, some more results are presented concerning the FAP resulting from correlated samples, and the detection probability for various transients and FAPs is presented.

5.5 Low Counts

The primary purpose of the development of this method is its applications to low count rate regimes. Recall that the thresholds k_L and k_U are set based on the mean of the background count rate multiplied by the number of samples inside the decision time. As described in Section 5.4, if the mean of the Poisson distribution used to set the thresholds is too low (either because the decision time is too short, or the count rate is too low), the probability of detecting no counts will be too large to set a reasonable bound. Because the probability of detecting no counts can be calculated simply as the probability mass function of the Poisson distribution at $k = 0$, noting that $\mu^k/k! = 1$ the required mean counts to set a low count threshold is shown in Equation (5.8):

$$\exp(-\mu) = FAP \implies \mu \geq -\ln(FAP) \quad (5.8)$$

Table 5.5 shows the required means for common FAPs:

σ Confidence	FAP	μ
2	0.954 499 736	3.0900372
3	0.997 300 203	5.9145776
4	0.999 936 657	9.6669462
5	0.999 999 426	14.370636

These results show that, in theory, only a few counts per second are adequate to set the FAP at 5σ with a decision time on the order of a few seconds. In practice, because of fluctuations in neutron count rate, it is possible that an fewer counts are detected over some time period, such that for a moment, the background count rate is below these theoretical limits, and consequently a lower threshold can not be set. Consequently, some counts above these minimum thresholds would be required to feasibly set the low threshold, and more counts increase the detection probability significantly.

5.6 True FAP and Detection Probability

5.6.1 True FAP

One interest in this work was determining the "true" FAP in the case where many samples were summed using the decision time methodology described above to test the alarm. In the case where a single sample is used to test the alarm, and consequently, all alarm tests are independent, the probability of any number of samples being above some threshold is given simply by the binomial distribution. In the case where the samples are summed over, like in this work, the summed probabilities calculated using Equation (5.7) are no longer independent and consequently, the distribution of points can no longer simply be calculated using the binomial distribution. This is of interest because it should be the case that, regardless of count rate or sample time, the FAP should remain consistent. To determine this, a modified algorithm was created to run thousands of trials using different FAPs, means, and decision times. Instead of inferring the background from

the mean of the prior sample, the background was fixed at the desired mean (Blessinger and York note on pg. 4 of [29] that the contribution of the error on the mean of the background to the error of future samples becomes small as the number of samples used to compute the background becomes large). A constant background was used to speed up computation. Tables 5.2-5.4 show the results of these tests, on three pairs of sample means and number of samples included in the running sum. Each trial consisted of 100,000 sample points, and 25000 trials were computed, the quantities presented are the mean number of counts that triggered the alarm, divided by the total number of samples, with standard error. The three tests used 100 samples per sum with a mean of 30, 15 samples with a mean of 10, and 20 samples with a mean of 1. The results show that the FAP converges to the desired FAP, with some systematic uncertainty created by the discrete nature of Poisson statistics, for the samples with more counts, the agreement is substantially better.

σ Confidence	FAP	True FAP
2	4.550 026 4%	$0.04552821 \pm 2.784e-5$
3	0.269 979 7%	$0.002669570 \pm 5.607e-6$
4	0.063 342 5%	$0.0006380952 \pm 2.506e-6$
5	0.000 057 3%	$5.648e-7 \pm 5.61e-8$

TABLE 5.2: Results of modified MAL testing. Each trial consisted of N=100,000 samples, with a mean of 30, and each summed sample consisted of the prior 100 samples.

σ Confidence	FAP	True FAP
2	4.550 026 4%	$0.04526410 \pm 1.075e-5$
3	0.269 979 7%	$0.002841065 \pm 2.326e-6$
4	0.063 342 5%	$0.00059302 \pm 9.694e-7$
5	0.000 057 3%	$5.936e-7 \pm 2.39e-8$

TABLE 5.3: Results of modified MAL testing. Each trial consisted of N=100,000 samples, with a mean of 10, and each summed sample consisted of the prior 15 samples.

σ Confidence	FAP	True FAP
2	4.550 026 4%	$0.05570155 \pm 1.407e-5$
3	0.269 979 7%	$0.002264964 \pm 2.363e-6$
4	0.063 342 5%	$0.0006773136 \pm 1.212e-6$
5	0.000 057 3%	$5.14e-7 \pm 2.72e-8$

TABLE 5.4: Results of modified MAL testing. Each trial consisted of N=100,000 samples, with a mean of 1, and each summed sample consisted of the prior 20 samples.

Below, transients are applied to these FAPs and count rates to determine the probability of detecting the transient.

5.6.2 Detection Probability

As Blessinger and York note on page 7 of [29], the detection probability for any given sample can be determined by calculating the area of the Poisson distribution during malicious interference with the thresholds set using the background count rate (in practice, the magnitude of the transient created during malicious interference would not be known, and therefore the detection probability could not be calculated in this way). Like the above case, if all samples are independent, then the mean number of alarms during a malicious interference event can be calculated by taking the sum of individual detection probabilities, as described above, when the summed samples are correlated, such a procedure is no longer possible. Therefore, thousands of transients were simulated to determine the probability of at least one alarm being triggered during the transient. In this case, a square wave transient was used, using 5 different magnitudes (abbreviated as transient magnitude (TM) in the tables below): 5%, 10%, 15%, 20%, and 25% of the background mean. The transient used was a square wave with a duration of 20 s.

100 batches of 100 simulations were performed for each test, with a background of 30 minutes, and a decision time of 20 s (Because each test either triggered or did not trigger the alarm, each batch represents the fraction of simulations in which the alarm was triggered). The mean of the 100 batches and the standard error of the batches are shown below. The background count rate and sampling time were set such that they were equivalent to the FAP tests from the previous section. The results are shown in Tables 5.5-5.7:

T.M.	2σ	3σ	4σ	5σ
5%	0.9654 +/- 0.0017	0.6692 +/- 0.0048	0.4836 +/- 0.0052	0.0373 +/- 0.0017
10%	1.0 +/- 0.0	0.9989 +/- 0.0003	0.9929 +/- 0.0009	0.7764 +/- 0.0049
15%	1.0 +/- 0.0	1.0 +/- 0.0	1.0 +/- 0.0	0.9996 +/- 0.0002
20%	1.0 +/- 0.0	1.0 +/- 0.0	1.0 +/- 0.0	1.0 +/- 0.0
25%	1.0 +/- 0.0	1.0 +/- 0.0	1.0 +/- 0.0	1.0 +/- 0.0

TABLE 5.5: Results of the five different transient magnitudes applied to the background with a count rate of 150 cps, with a mean of 30 counts, using a decision time of 20 s.

T.M.	2σ	3σ	4σ	5σ
5%	0.6698 +/- 0.0046	0.1289 +/- 0.0033	0.0472 +/- 0.0021	0.0 +/- 0.0
10%	0.6493 +/- 0.0046	0.1110 +/- 0.0028	0.0350 +/- 0.0019	0.0002 +/- 0.0001
15%	0.6765 +/- 0.0051	0.1225 +/- 0.0031	0.0394 +/- 0.0020	0.0 +/- 0.0
20%	0.7565 +/- 0.0045	0.1779 +/- 0.0039	0.0697 +/- 0.0025	0.0005 +/- 0.0002
25%	0.8287 +/- 0.0036	0.2939 +/- 0.0044	0.1465 +/- 0.0039	0.0021 +/- 0.0005

TABLE 5.6: Results of the five different transient magnitudes applied to the background with a count rate of $6.\bar{6}$ cps, with a mean of 10 counts, using a decision time of 20 s.

T.M.	2σ	3σ	4σ	5σ
5%	0.5513 +/- 0.0044	0.0606 +/- 0.0023	0.0179 +/- 0.0014	0.0 +/- 0.0
10%	0.5656 +/- 0.0053	0.0697 +/- 0.0025	0.0201 +/- 0.0015	0.0 +/- 0.0
15%	0.5870 +/- 0.0044	0.0794 +/- 0.0025	0.0255 +/- 0.0015	0.0004 +/- 0.0002
20%	0.6130 +/- 0.0055	0.0945 +/- 0.0030	0.0325 +/- 0.0017	0.0 +/- 0.0
25%	0.6389 +/- 0.0045	0.1197 +/- 0.0031	0.0443 +/- 0.0020	0.0001 +/- 0.0001

TABLE 5.7: Results of the five different transient magnitudes applied to the background with a count rate of 1 cps, with a mean of 1 counts, using a decision time of 20 s.

For the high count rate case shown in Table 5.5, all FAPs can detect transients above 20% of the background mean. For the medium and low count rate cases shown in Table 5.6 and 5.7, the probability of detection is low for the 4σ and 5σ FAPs, and uncertain for less sensitive FAPs, regardless of transient magnitude. These results show that significantly more counts than are required to set the lower threshold are required to increase the detection probability to an acceptable level.

These results are not meant to be interpreted as precise detection probabilities. Instead, they provide a sense of scale about the size of the transients that can be observed using this algorithm at various FAPs.

5.6.3 Malicious Interference Conclusions

The results of the prior two subsections, when interpreted together, show that a FAP of 4-5 σ is appropriate to avoid regular false alarms, and the detection probability of malicious interference transients with the FAP set to 4-5 σ is too low to be useful in low count rate scenarios. Typical count rates observed at MNR were comparable to the 150 cps case, and therefore adequate to set sensitive FAPs, and have good detection probability. It bears emphasizing that these results only bound reasonable FAPs and transient magnitudes, and that appropriate parameters would have to be established by trial and error for each neutron detector and local environment.

In the context of RPMs, from which this algorithm was derived, a false alarm would likely lead to an unnecessary search of the vehicle in question. Similarly in the SND context, a procedure should be developed to verify the integrity of the neutron monitoring system. Based on the frequency of false alarms and the burden of verifying false alarms, that verification could involve reviewing video records of the detectors during the relevant period or physical inspection of the neutron detectors and their local environment for signs of malicious interference. The burden associated with verifying false alarms would be an important consideration when determining an appropriate FAP.

Finally, the triggering of the alarm would likely occur during several possible operations, primarily refuelling, as the reactor changes power. These alarms could either be discarded or used to corroborate that the stated operation has occurred.

Chapter 6

Conclusion

6.1 Summary and Conclusion

This work describes experiments into the use of standoff neutron detectors as a safeguards method at the McMaster Nuclear Reactor. The objective of which was to show that the weight isotopic composition, a measure of the quantity of fissile material in the core, scaled by the fission likelihood and fission energy, could be correlated with the neutron counts per unit power observed by the detector. This work broadly consists of four parts; in the first part (Chapter 2), the procedure and apparatus for neutron data collection are described. This section also introduces two of the primary difficulties with neutron counting at MNR as compared to other, more conventional facilities, those being the beam port use, and day/night operation. The second part (Chapter 3) details how the isotopic inventory was determined from the results of OSCAR5 depletion simulations, and provides analysis that decomposes changes in WIC into its constituent isotopes, and analyzes those changes over time. The third part (Chapter 4) combines the previous two parts to determine if it is possible to correlate changes in counts per power with change in WIC. Ultimately, the changes in neutron counts per unit power were significantly larger than the changes in WIC over the same period, in spite of attempts to correct for the influence of the day/night periodicity on the neutron count data. This resulted in significant error in regressions and made it impossible to provide reasonable agreement on thresholds on the neutron count rate per unit power to determine if the removal of fissile material has occurred.

The final part (Chapter 5) concerns malicious interference. Malicious interference refers to attempts to add exogenous material to alter the neutron count rate to defeat the standoff neutron method. It is postulated that during that addition of exogenous material a deviation in count rate will be incurred and that this deviation can be detected using an algorithm. The design of one such algorithm is detailed. It was determined that, for sufficiently high count rates (approximately 100 counts per second), both low false alarm rates and high detection probabilities can be achieved.

6.2 Future Work

This work has suffered from a variety of difficulties specific to neutron counting at MNR (beam port use, day/night operation, small changes in WIC) that have made drawing clear conclusions from this work difficult. Although the results of this work have been inconclusive, that does not mean that the results from attempting the SND at a more suitable site would not be more conclusive, as they were in the case of the NRU experiments before this work. As mentioned in the discussion in Chapter 4, these experiments can also be used to inform the selection of new facilities to attempt the SND method. Aside from the selection of new facilities, other possibilities for future work are discussed below.

First, future changes to the MNR operations schedule and activities may allow for the SND method to be attempted at a later date without the difficulties that hampered this work. It is also possible that a more sophisticated time series analysis could be done to attempt to recover the portions of the daily signal during beam port use.

One possible way to acquire more data to support the SND method would be to use existing neutron detectors installed for neutron monitoring as detectors in the SND method. The use of such detectors would be only possible with the support of the operator. Additionally, the use of existing detectors in the SND method would be limited by the detector efficiency and size, as well as location relative to the reactor and ability to control the local neutron environment. Validation work would need to be performed to interpret the neutron signal during regular operation.

Bibliography

- [1] *IAEA safeguards overview*, Jul. 2014. [Online]. Available: <https://www.iaea.org/publications/factsheets/iaea-safeguards-overview>.
- [2] International Atomic Energy Agency, *Safeguards techniques and equipment*, Vienna: IAEA, 2011. [Online]. Available: <https://www.iaea.org/publications/8695/safeguards-techniques-and-equipment>.
- [3] *Safeguards agreements*, Jun. 2016. [Online]. Available: <https://www.iaea.org/topics/safeguards-agreements>.
- [4] *Canada's small modular reactor action plan*, May 2021. [Online]. Available: <https://smractionplan.ca/>.
- [5] B. van der Ende, L. Li, D. Godin, and B. Sur, Stand-off nuclear reactor monitoring with neutron detectors for safeguards and non-proliferation applications, *Nature Communications*, vol. 10, Apr. 2019.
- [6] M. Alqahtani, A. Buijs, and S. Day, Serpent-2 and OSCAR-4 computational tools compared against mcmaster nuclear reactor improved operational data history for U-235 fuel inventory tracking, local power tracking and validation of multiplication factor, *Annals of Nuclear Energy*, vol. 145, May 2020.
- [7] *Canada's small modular reactor action plan*, May 2021. [Online]. Available: <https://smractionplan.ca/>.
- [8] S. Day, The use of experimental data in an mtr-type nuclear reactor safety analysis, Ph.D. dissertation, 2006. [Online]. Available: https://www.nuceng.ca/grads/sed-phd_use-of-expt-data-in-MTR-SA_r0.pdf.
- [9] *Benchmarking Against Experimental Data of Neutronics and Thermohydraulic Computational Methods and Tools for Operation and Safety Analysis of Research Reactors* (TECDOC Series 1879). Vienna: INTERNATIONAL ATOMIC ENERGY AGENCY, 2019, ISBN: 978-92-0-109619-7. [Online]. Available: <https://www.iaea.org/publications/13547/benchmarking-against-experimental-data-of-neutronics-and-thermohydraulic-computational-methods-and-tools-for-operation-and-safety-analysis-of-research-reactors>.
- [10] E. MacConnachie, Neutron flux measurements for uncertainty quantification at the mcmaster nuclear reactor, Ph.D. dissertation, 2021. [Online]. Available: <https://macsphere.mcmaster.ca/handle/11375/26906>.

Bibliography

- [11] R. T. Kouzes, Detecting illicit nuclear materials: The installation of radiological monitoring equipment in the united states and overseas is helping thwart nuclear terrorism, *American Scientist*, vol. 93(5), 422–427, 2005, ISSN: 00030996. [Online]. Available: <http://www.jstor.org/stable/27858641> (visited on 09/12/2022).
- [12] *Radiation detection equipment for homeland security: Reuter stokes*. [Online]. Available: <https://www.bakerhughes.com/reuter-stokes/radiation-measurement-monitoring/homeland-security-4>.
- [13] *NPM3100 operation manual and specifications*, version OM_9-2015_v2, Questa Instruments, Tuscon, Arizona, 2015. [Online]. Available: <https://www.quaestainstruments.com/neuchrometer>.
- [14] W. McKinney, Data Structures for Statistical Computing in Python, in *Proceedings of the 9th Python in Science Conference*, S. van der Walt and J. Millman, Eds., 2010, 56–61.
- [15] E. MacConnachie, D. Novog, and S. Day, Quantification of system uncertainties in activation experiments at nuclear research reactors, *Annals of Nuclear Energy*, vol. 134, 432–440, Dec. 2019.
- [16] *28.2 - normal approximation to poisson: Stat 414*. [Online]. Available: <https://online.stat.psu.edu/stat414/lesson/28/28.2>.
- [17] R. Prinsloo, F. Van Heerden, D. Botes, and R. Mudau, Recent developments of the oscar calculational system, as applied to selected examples from iaea research reactor benchmarks, 2017.
- [18] J. Leppänen, M. Pusa, T. Viitanen, V. Valtavirta, and T. Kaltiaisenaho, The serpent monte carlo code: Status, development and applications in 2013, *Annals of Nuclear Energy*, vol. 82, 142–150, 2015, Joint International Conference on Supercomputing in Nuclear Applications and Monte Carlo 2013, SNA + MC 2013. Pluri- and Trans-disciplinarity, Towards New Modeling and Numerical Simulation Paradigms, ISSN: 0306-4549. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0306454914004095>.
- [19] E. Fridman and J. Leppänen, Revised methods for few-group cross section generation in the serpent monte carlo code, vol. 1, Apr. 2012.
- [20] P. K. Romano, N. E. Horelik, B. R. Herman, A. G. Nelson, B. Forget, and K. Smith, Openmc: A state-of-the-art monte carlo code for research and development, *Annals of Nuclear Energy*, vol. 82, 90–97, 2015, Joint International Conference on Supercomputing in Nuclear Applications and Monte Carlo 2013, SNA + MC 2013. Pluri- and Trans-disciplinarity, Towards New Modeling and Numerical Simulation Paradigms, ISSN: 0306-4549. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S030645491400379X>.
- [21] Ding, Zechuan, Solving bateman equation for xenon transient analysis using numerical methods, *MATEC Web Conf.*, vol. 186, 01004, 2018. [Online]. Available: <https://doi.org/10.1051/mateconf/201818601004>.

Bibliography

- [22] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, Í. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors, SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python, *Nature Methods*, vol. 17, 261–272, 2020.
- [23] D. York, N. Evensen, M. Ló, M. Nez, J. De, and B. Delgado, Unified equations for the slope, intercept, and standard errors of the best straight line, *American Journal of Physics - AMER J PHYS*, vol. 72, Mar. 2004.
- [24] *Evinco micro reactor*. [Online]. Available: <https://www.westinghousenuclear.com/energy-systems/evinco-micro-reactor>.
- [25] *Status report - NuScale smr*. [Online]. Available: https://aris.iaea.org/PDF/NuScale-NPM200_2020.pdf.
- [26] *Status report - BWRX-300*. [Online]. Available: https://aris.iaea.org/PDF/BWRX-300_2020.pdf.
- [27] *Advances in small modular reactor technology developments*. [Online]. Available: https://aris.iaea.org/Publications/SMR-Book_2018.pdf.
- [28] R. T. Kouzes, E. R. Siciliano, J. H. Ely, P. E. Keller, and R. J. McConn, Passive neutron detection for interdiction of nuclear material at borders, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 584(2), 383–400, 2008, ISSN: 0168-9002. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0168900207022085>.
- [29] C. Blessinger and R. York, Neutron detection algorithm development, *Oak Ridge National Laboratory ORNL/TM-2006/157*, 2006.