



Space to Air High-Altitude Region Adjoint
Neutron Transport

THESIS

Zachary W. LaMere, 1st Lieutenant, USAF
AFIT-ENP-MS-21-M-125

DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

DISTRIBUTION STATEMENT A
APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

The views expressed in this document are those of the author and do not reflect the official policy or position of the United States Air Force, the United States Department of Defense or the United States Government. This material is declared a work of the U.S. Government and is not subject to copyright protection in the United States.

AFIT-ENP-MS-21-M-125

SPACE TO AIR HIGH-ALTITUDE REGION ADJOINT NEUTRON
TRANSPORT

THESIS

Presented to the Faculty
Department of Engineering Physics
Graduate School of Engineering and Management
Air Force Institute of Technology
Air University
Air Education and Training Command
in Partial Fulfillment of the Requirements for the
Degree of Master of Science in Nuclear Engineering

Zachary W. LaMere, B.S.

1st Lieutenant, USAF

March 2021

DISTRIBUTION STATEMENT A
APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

AFIT-ENP-MS-21-M-125

SPACE TO AIR HIGH-ALTITUDE REGION ADJOINT NEUTRON
TRANSPORT

THESIS

Zachary W. LaMere, B.S.
1st Lieutenant, USAF

Committee Membership:

Darren Holland, PhD
Chair

Lt Col Whitman Dailey, PhD
Member

John McClory, PhD
Member

Abstract

The goal of this work was to use the adjoint transport equation to characterize the energy spectrum of a neutron source located in the upper atmosphere using the time-energy fluence at the satellite. This adjoint approach directly solves for the source energy spectrum. The adjoint method also could improve computational efficiency over the forward method. The adjoint transport equation was solved via Monte Carlo methods in a new program called Space to Air High-Altitude Region Adjoint (SAHARA) written in Python 3.7. A new adjoint source event estimator was developed to improve the computational efficiency. This work explores SAHARA's development, and its performance for mono-energetic and continuous energy sources. In general, the identified source spectra were shifted towards lower energies approximately five percent, but were able to capture the source spectrum shapes. Additionally, continuous energy sources still passed a 2-D Kolmogorov–Smirnov (K-S) test. Lastly, SAHARA was applied to the real-world neutron energy spectra of Fat Man and Little Boy. Although these spectra were also noticeably shifted towards lower energies, the spectral features are still recognizable and the spectra passed a 2-D K-S test. SAHARA provides a new tool for estimating the source spectrum from space-based neutron measurements.

AFIT-ENP-MS-21-M-125

For her

Acknowledgements

I would like to thank the constant advice and support of Dr. Darren Holland and for putting up with my many random questions. This work would also not have been possible without the help of Lt Col Whitman Dailey. Lt Col Dailey managed to find time to discuss problems I had with SAHARA while being deployed. Lastly, thank you to Dr. John McClory for setting me on the path for success. It has been a pleasure working with you all.

Zachary W. LaMere

Table of Contents

	Page
Abstract	iv
Dedication	v
Acknowledgements	vi
List of Figures	x
List of Tables	xv
List of Acronyms	xvi
I. Introduction	1
1.1 Background	1
1.2 Why adjoint method?	3
1.3 Problem Statement	3
1.4 Problem Scope Assumptions	4
II. Background and Related Work	6
2.1 Adjoint Transport Equation Development	6
2.2 Methods of Solving the Neutron Transport Equation	8
2.3 Common Uses of the Adjoint Transport Equation	9
2.4 Assumptions of the Air-to-Space Neutron Transport Problem	11
2.4.1 Relative Motions	12
2.4.2 Gravity	13
2.4.3 Atmosphere Fidelity	14
2.4.4 Radioactivity Decay of Neutrons	17
III. Methodology	18
3.1 Methodology of SAHARA	18
3.2 Start Pseudo-Neutron	19
3.2.1 Sampling the Distribution	19
3.2.2 Direct Contribution Check	23
3.2.3 Sampling Angle	23
3.2.4 Move Pseudo-Neutron to the Edge of the Atmosphere	24
3.3 Move Pseudo-Neutron	24
3.3.1 Calculating Optical Thickness and Effective Path Length	24

	Page
3.3.2 Finding Geometric Distance	26
3.3.3 Verifying Move Pseudo-Neutron	27
3.4 Source Event Pseudo-Neutron	27
3.4.1 Source Event Estimator	28
3.4.2 Calculation of Required Energies	30
3.4.3 Probability of Non-Interaction Along Flight Path	31
3.4.4 Probability of Scattering	32
3.5 Interact Pseudo-Neutron	32
3.6 Kill Pseudo-Neutron	33
3.7 Tally Pseudo-Neutron	34
3.8 Atmosphere Representation	35
3.8.1 Sea Level to 86 kilometers	35
3.8.2 Temperature	35
3.8.3 Density	39
3.8.4 Cross Section Data	43
3.9 Overview of SAHARA	43
3.9.1 Commonality with High-Altitude Transport to Space for Neutrons (HATS-n)	44
3.9.2 Target Architecture	44
3.9.3 SAHARA Inputs	44
3.9.4 Modules and Descriptions	45
3.10 Summary	47
IV. Once-Scattered Scenarios	48
4.1 Mono-energetic Source Results	49
4.2 Examining Once-Scattered Peak Energy Shift	52
4.3 Mono-energetic Importance Results	53
4.4 Continuous Energy Sources	57
4.4.1 Normalized Counts vs Energy	57
4.4.2 Importance vs Energy	59
4.5 Summary	60
V. Multi-Scattered Scenarios	61
5.1 Multi-scattering Scenarios vs Once-scattered Scenarios	61
5.1.1 Multi-scattering Results and Analysis	61
5.1.2 Watt Fission Spectrum Results and Analysis	65
5.2 Approaches taken to the Direct Contributions Issue	66
5.2.1 Changing Time-Energy Grid of Input	67
5.2.2 Modifying Time Energy Bins when Selected	69
5.2.3 Directionally Limiting Direct Contributions	70
5.2.4 Modifying Pseudo-Neutron Kill Criteria if Marked Direct	70

	Page
5.2.5 Post Processing Results Option.....	71
5.3 Influence Of Tally Energy Bin Sizes.....	72
5.4 Summary	74
VI. Five-Scattered Scenarios.....	75
6.1 Five-Scatter vs Once-Scatter	75
6.1.1 Mono-energetic Sources	75
6.1.2 Watt Fission Source	78
6.2 Three Isotope Atmosphere Results.....	79
6.2.1 Varying Source Altitude	82
6.2.2 Varying Source Location	87
6.2.3 Little Boy and Fat Man Sources	91
6.3 Summary	95
VII. Conclusions	97
7.1 Summary	97
7.2 Future Work.....	99
Bibliography	102

List of Figures

Figure		Page
1	The adjoint air-to-space neutron transport problem. Pseudo-neutrons travel back towards the source. Possible source events, similar to a next event estimator, are also calculated if a scatter could reach the source.	2
2	Arrival energy as a function of time of flight from an equatorial source at $\Delta\alpha = 0^\circ$ at time of emission with and without gravity showing negligible gravitational effects below 1,000 seconds. Figure used with permission from [1]	14
3	Example normalized time-energy neutron distribution at the satellite generated from HATS-n using a Watt fission spectrum.	20
4	Sampled time-energy neutron distribution from nominal input distribution sampling time first, then energy.	21
5	Sampled time-energy neutron distribution from nominal input distribution sampling energy first, then time.	21
6	Time and energy collapsed Cumulative Distribution Function (CDF)s for a Watt fission spectrum used for testing equal distributions	22
7	First sampled interaction altitude of SAHARA verifying move pseudo-neutron step using a N-14 atmosphere	27
8	Two examples of solutions to the source event estimator. (a) shows where the solution is above the sampled interaction point. (b) shows where the solution is below the sampled interaction point.	29
9	Temperature as a function of geometric altitude below 86 km verifying correct implementation of the United States Standard Atmosphere 1976 (USSA76).	39
10	Atmospheric pressure as a function of geometric altitude below 86 km verifying correct implementation of the USSA76.	40

Figure		Page
11	Atmospheric density as a function of geometric altitude below 86 km verifying correct implementation of the USSA76.	41
12	Once-scattered Probability Distribution Function (PDF)s for mono-energetic spectra. All peaks are lower than the true source values. Energy bins with counts greater than 500 are plotted.	50
13	Once-scattered PDFs for various mono-energetic spectra with weighted Gaussian best fits applied. Gaussian means are below actual source energies. Energy bins with counts greater than 500 are plotted.	51
14	Once-scattered normalized average importance for a 100 keV, 1 MeV, and 14.056 MeV source and their relative error within each bin. Energy bins with counts greater than 500 are plotted.	53
15	Once-scattered normalized integral of importance for a 100 keV, 1 MeV, and 14.056 MeV source. Results match closely to their respective PDF plots. The peak values are still below actual source energy. Energy bins with counts greater than 500 are plotted.	56
16	Once-scattered 100 keV to 1 MeV and Watt fission sources PDFs. Energy bins with counts greater than 100 are plotted.	57
17	Once-scattered 100 keV to 1 MeV uniform source PDF with the non-uniformly spaced PDF overlaid. Energy bins with counts greater than 100 are plotted. The bin width changes the shape of the expected uniform PDF. This passes a 2-D Kolmogorov-Smirnov Test (K-S Test) with a p-value of 0.8636.	58
18	Watt fission source once scatter normalized counts with Watt fission spectrum (37) overlayed. Energy bins with counts greater than 100 are plotted. The bin widths change the expected shape of the Watt fission spectrum PDF. This passes a 2-D K-S Test with a p-value of 0.8980.	59

Figure		Page
19	Once-scattered average importance for 100 keV to 1 MeV uniform source, and Watt fission spectrum source. Although not useful for characterizing a source, all pseudo-neutrons reaching the source are important to the source spectrum. Energy bins with counts greater than 100 are plotted.	60
20	Multi-scattered, direct and scattered contributions for 100 keV and 1 MeV sources. Energy bins with counts greater than 500 are plotted.	62
21	Multi-scattered, 1 MeV source energy PDF at the satellite showing that the direct particles comprise 40% of the total particles at the satellite.	63
22	Normalized average scatter importance for a 1 MeV source. The higher energies have higher importance indicating that they are less likely to occur.	64
23	Multi-scattered, scattered contributions PDF from a 1 MeV source and a energy cut-off of 1 MeV which removes the higher energy peak seen in Figure 20d. Energy bins with counts greater than 100 are plotted.	65
24	Multi-scattering, direct and scattered contribution PDFs for a watt fission spectrum source.	66
25	Multi-scattering, direct and scattered contribution PDFs for a 1 MeV mono-energetic source using higher resolution input.	68
26	Subdividing time energy bins which are sampled uniformly	70
27	Scattered contribution PDF for a 1 MeV source allowing direct contributions to scatter once. Non-true higher energy peak is still present. Energy bins with counts greater than 500 are plotted.	71
28	Direct contribution PDF for a Watt fission source using 100 logarithmically spaced bins per decade displaying a bi-modal distribution. Plotting energy bins with counts greater than 100.	72

Figure		Page
29	Normalized time-energy fluence input produced by HATS-n for a Watt fission source displaying multiple distributions.	73
30	Five-scattered PDFs for mono-energetic sources. Peaks are significantly wider compared to the once-scattered cases. Energy bins with counts greater than 500 are plotted.	76
31	Once-scattered and five-scattered PDF for a 1 MeV source. Allowing more scatters widens the peak.	77
32	Five-scattered PDF from a Watt fission source with a Watt fission spectrum corrected for bin width overlaid. This passes a 2-D K-S Test with a p-value of 0.8298. Energy bins with counts greater than 500 are plotted.	78
33	Once-scattered and five-scattered PDFs from a Watt fission source. Allowing more scatters results in a shift towards lower energies and a depreciated peak.	79
34	Five-scattered PDFs for sources in a 3 isotope atmosphere. Energy bins with counts greater than 500 are plotted for mono-energetic sources.	81
35	Five-scatter PDFs for a 1 MeV source for both a N_{14} and three isotope atmosphere. Moving to a three isotope shifts the spectrum towards higher energies.	82
36	Five-scattered PDFs 1 MeV source in a 3 isotope atmosphere, varying source altitude between 30, 50, and 80 km. Energy bins with counts greater than 500 are plotted.	84
37	PDFs for a 1 MeV source at various altitudes. The lower in the atmosphere the source, the larger the standard deviation of the peak. The higher in the atmosphere, the larger the error in the Gaussian peak.	85
38	Five-scattered PDFs for a Watt fission source in a 3 isotope atmosphere, varying source altitude between 30, 50, and 80 km.	86

Figure		Page
39	PDFs for a Watt fission spectrum at various altitudes. The effect of the atmosphere's cross sections are more evident at lower altitudes.....	87
40	Five-scattered PDFs for a 1 MeV and Watt fission source located at 80 km altitude, and 45N, 0E in a N_{14} atmosphere.....	89
41	Five-scattered PDFs for a 1 MeV and Watt fission source located at 50 km altitude, and 45N, 45E in a N_{14} atmosphere.....	90
42	Five-scattered, 1 MeV and Watt fission source with the source located at 50 km altitude, 0N, 0E and 45N, 45E.	91
43	Neutron spectra for Fat Man and Little Boy. Plot generated from data provided by [2].	92
44	PDF from the Little Boy source with the true spectrum overlaid [2]. This passes a 2-D K-S Test with a p-value of 0.7583. Energy bins with counts greater than 100 are plotted.	93
45	PDF from a Fat Man source with the true spectrum overlaid [2]. This passes a 2-D K-S Test with a p-value of 0.8063. Energy bins with counts greater than 100 are plotted.	94
46	PDF from a Fat Man source above 10 eV with the true spectrum overlaid [2]. This passes a 2-D K-S Test with a p-value of 0.9690. Energy bins with counts greater than 100 are plotted.	95

List of Tables

Table		Page
1	Physical constants for 1976 U.S. Standard Atmosphere [3]	35
2	Reference levels, and temperature lapse rates and computed base temperatures and pressures from the surface to 86 geometric kilometers [3, 1].	37
3	Molecular weight ratios for various geometric altitudes and geopotential heights [3, 1].	38
4	Molecular weights and fractional values for constituents for sea level atmosphere [3].	42
5	Molecular volumetric fractions for a atmosphere containing all isotopes in SAHARA.	43
6	Atmosphere constituents options in SAHARA.	45
7	SAHARA settings for once-scattered scenarios.	49
8	Fitting parameters to once-scattered Gaussian best fits from Origin Pro.	52
9	Gaussian best fit parameters for both once-scattered (OS) and five-scattered (5S).	77
10	SAHARA settings for five scattering, three isotope atmosphere scenarios.	80
11	SAHARA settings for five-scattered, varied altitudes scenarios.	83
12	Gaussian best fit parameters for a 1 MeV source at altitudes of 30, 50, and 80 km.	86
13	SAHARA settings for five-scattered, varied latitude and longitude scenarios. Both the 1 MeV and Watt fission source were run at all locations.	88
14	SAHARA settings for Little Boy and Fat Man scenarios.	92

List of Acronyms

ADVANTG	Automated VARIance reducTion Generator
CADIS	Consistent Adjoint-Driven Importance Sampling
CDF	Cumulative Distribution Function
DIORAMA	Distributed Infrastructure Offering Real-time Access to Modeling and Analysis
ECI	Earth-Centered Inertial
EPL	Effective Path Length
GEANT	GEometry ANd Tracking
HATS-n	High-Altitude Transport to Space for Neutrons
K-S Test	Kolmogorov-Smirnov Test
MC	Monte Carlo
MCNA	Monte Carlo N-particle Adjoint
MCNP	Monte Carlo N-Particle Transport Code
NUDET	Nuclear Detonation
PDF	Probability Distribution Function
SAHARA	Space to Air High-Altitude Region Adjoint
ToF	time-of-flight
USNDS	United States NUDET Detection System
USSA76	United States Standard Atmosphere 1976

SPACE TO AIR HIGH-ALTITUDE REGION ADJOINT NEUTRON TRANSPORT

I. Introduction

1.1 Background

Imagine a scenario where there is a Nuclear Detonation (NUDET) at a certain altitude, and location where the neutron, gamma, and x-rays sensors of a space-based detection system. Neutrons from this event have some probability to reach a satellite's position in orbit. They could travel to a satellite directly from the source, or after an interaction in the atmosphere. From the time-energy distribution of neutrons that reach the satellite it may be possible to characterize the energy distribution of the source.

An estimate of the time-energy distribution of neutrons may be calculated via Monte Carlo (MC) methods by following the neutrons from the source to the satellite location and then determining if the modeled satellite response matches real world data. Characterizing the source using this process is known as the *forward* air-to-space neutron transport problem. Alternatively, given a time-energy distribution at the satellite, the energy spectrum at the source location can be determined via a similar MC method. This approach is known as the *adjoint* air-to-space neutron transport problem. The general adjoint transport problem is illustrated in Figure 1. The forward air-to-space neutron transport problem, although not exhausted, has been thoroughly studied by Lt Col Whitman Dailey, USAF [1], and others. There has also been production-level programs such as Distributed Infrastructure Offering Real-

time Access to Modeling and Analysis (DIORAMA), which is the current program used by mission planners, and operators of United States NUDET Detection System (USNDS). However, the adjoint problem for the air-to-space problem warrants further investigation.




Figure 1. The adjoint air-to-space neutron transport problem. Pseudo-neutrons travel back towards the source. Possible source events, similar to a next event estimator, are also calculated if a scatter could reach the source.

Neutral particle transport, which includes neutron and photon transport, is typically taught and implemented in the forward direction. For example, in *Computational Methods of Nuclear Particle Transport* by Lewis and Miller, a common text for this problem, all methods presented in detail are forward methods. Additionally production Monte Carlo codes such as Monte Carlo N-Particle Transport Code (MCNP), and GEometry ANd Tracking (GEANT) are also generally forward methods, though both have some adjoint method capabilities. However, the adjoint neutron transport method provides an equally valid solution [4] and corresponds to the realistic scenario where the response is known, but the source is not.

The adjoint method is not a new development in nuclear engineering, and has been

around since the 1950's and has many common uses within field. This includes, but is not limited to, low probability detector problems, reactor criticality perturbation theory, and initial seeds for criticality calculations.

1.2 Why adjoint method?

Although solving a neutral particle transport problem is equally valid in both the forward and adjoint directions, finding a adjoint solution provides the unique benefit of directly solving for what is of interest: the energy spectrum of the source. In the forward problem, the neutron time-energy distribution at the satellite is determined, and the neutron energy distribution of the source is the input. However, for the adjoint approach, the source neutron energy distribution is determined, and the time-energy distribution of the satellite detected response is the input.

Using the forward method, one would have to make an educated guess regarding the source, and then assess how well the result matches the time-energy neutron distribution observed at the satellite. This process is often repeated with multiple source approximations and iterations as the simulated and measured responses converge. This becomes especially burdensome where there are multiple satellite responses to consider. However, the adjoint approach yields a source estimate which produces a simulated source correlating to the measured response (within the MC limitations and modeling assumptions) after only one simulation.

1.3 Problem Statement

The primary goal of the work is to *characterize* the energy spectrum of a neutron source from a time-energy neutron distribution at a single satellite location by using the adjoint neutron transport method via Monte Carlo simulation.

1.4 Problem Scope Assumptions

The following assumptions provide for an initial scope of the problem for the adjoint method for characterizing the energy spectrum of a point neutron source within the earth's atmosphere. The most important assumption, one that enables the primary goal to be achieved, is that the source location is known. For a single satellite study such as this, if the location of the source is not known then solving the adjoint problem will yield a most probable location of the source, but not source energy spectrum characterization. However, assuming the location of the source is not unwarranted.

Further assumptions are:

1. Due to the distances involved, the satellite and source are treated as points.
2. The 1,000 second maximum detector measurement time begins at the instantaneous detonation.
3. The 1976 US Standard Atmosphere model is representative of the true atmospheric conditions at all latitudes and longitudes, and ends at 86 km.
4. Motion of the atmosphere, source and satellite is assumed negligible.
5. Gravitational effects are negligible.
6. The earth is spherical.
7. The decay of neutrons is negligible.
8. Source and satellite position is known.
9. Scattering is elastic and isotropic.
10. Detector characteristics can be ignored.

The rationale behind these problem scope assumptions will be discussed as possible future expansions of SAHARA.

II. Background and Related Work

The concept of the adjoint neutron transport equation was a key landmark in the evolution of the understanding of nuclear particle transport [5]. The basic concept was first developed by Weinberg and Wigner in 1958. The solutions to the adjoint neutron transport equation are interpreted as a neutron importance and can be used to expedite different classes of problems [4]. One common use of the adjoint method is the reduction of uncertainty in problems where there is a small detection volume, and a large source volume where the probability of detection is low [6, 7]. Alternatively, coupling adjoint and forward transport methods are useful for small source, detector problems where next event estimation is used [8, 9]. It is also used to improve convergence speeds for reactor criticality calculations as the adjoint equation can solve for better starting locations for neutron seeds [10]. Other common uses are in various perturbation methods. Underlying all these methods is the fact that for every form of the neutron transport equation, a related adjoint equation can be developed [4].

2.1 Adjoint Transport Equation Development

The general form of the non-multiplying forward neutron transport equation under isotropic scattering is shown in Equation 1. From this equation an equivalent adjoint transport formulation can be developed. In forward transport, the particle in question is the neutron, whereas in the adjoint equation the particle is referred to as a pseudo-neutron or adjunction [5, 6]. This work will refer to them as pseudo-neutrons.

$$\left(\frac{1}{v} \frac{\partial}{\partial t} + \hat{\Omega} \cdot \vec{\nabla} + \Sigma_t(\vec{r}, E) \right) \phi(\vec{r}, E, \hat{\Omega}, t) = S_{ext}(\vec{r}, E, \hat{\Omega}, t) + \int_{4\pi} d\hat{\Omega}' \int_0^\infty dE' \Sigma_s(\vec{r}, E' \rightarrow E, \hat{\Omega}' \cdot \hat{\Omega}) \phi(\vec{r}, E', \hat{\Omega}', t) \quad (1)$$

Following the derivation presented in Lewis and Miller [4] pages 46-50, the corresponding adjoint transport equation is

$$\left(\frac{1}{v}\frac{\partial}{\partial t} - \hat{\Omega} \cdot \vec{\nabla} + \Sigma_t(\vec{r}, E)\right)\phi^\dagger(\vec{r}, E, \hat{\Omega}, t) = S_{ext}^\dagger(\vec{r}, E, \hat{\Omega}, t) + \int_{4\pi} d\hat{\Omega}' \int_0^\infty dE' \Sigma_s(\vec{r}, E \longrightarrow E', \hat{\Omega} \cdot \hat{\Omega}')\phi^\dagger(\vec{r}, E, \hat{\Omega}', t). \quad (2)$$

A few comments are warranted on the physical meaning of Equation 2. The first is although $\phi^\dagger$ is usually referred to as the adjoint flux, physically this is not the case. Instead $\phi^\dagger$ is neutron importance which is a dimensionless quantity [5]. However, the lack of units does not mean that the value is meaningless. The solution to the adjoint equation is adjoint to the solution of the forward equation. This means that even though the solutions are different in units, the values numerically coincide. Although neutron importance is dimensionless, it can be interpreted as a flux [5, 6]. Additionally the superscript $\dagger$ does not refer to a transpose or a complex conjugate, but rather the parameters defining the transport of pseudo-neutrons. $S^\dagger$ represents a source density of the pseudo-neutrons [6].

Further examining Equations 1 and 2, reveals that the collision kernels are reversed. In the forward transport equation, neutrons are moving from energy E' , and down-scattered to energy E which corresponds to direction $\hat{\Omega}$. In the adjoint transport equation, the pseudo-neutrons are up-scattered from energy E to energy E' corresponding to direction $\hat{\Omega}'$. This fact presents an issue unique to the adjoint equation when implementing a Monte Carlo solution. The scattering cross section value used will be at E , but because the pseudo-neutron is moving backwards the real cross section value needed is at E' . It also means the scattering kernel may be greater than unity which can cause poor convergence properties. There are two different approaches to solving this issue. One method involves normalizing the scattering kernel,

a non-trivial task and changes the neutron importance to a form similar to a collision density [11]. Alternatively particle splitting can be used to reduce pseudo-neutron weights if they become much greater than unity, thus improving convergence [5]. The latter approach will be used in this work, and more discussion on this can be found in section 3.5.

2.2 Methods of Solving the Neutron Transport Equation

The two primary approaches to solving the neutron transport equation are deterministic and stochastic. Deterministic methods aim to solve the neutron transport equation directly, a common deterministic method is discrete ordinates method. Stochastic methods, which includes the MC method, can be described as a probabilistic methods [4].

The deterministic method is a common way to solve simple neutron transport problems. When using this method all independent variables are discretized. Energies are broken into bins creating a multi-group approximation. For reactors, as few as two groups can be used with fairly accurate results, but for the air-to-space problem many more would be needed. Spatially, discretizing breaks down the geometry of the problem space into smaller regions which can be solved via finite difference methods. Angular variables are also discretized which can be solved via S_n or P_n methods [4]. The time component is also discretized, and can be solved via finite difference methods. Discretization creates various systems of equations that can be solved directly, or iteratively depending on the approach. For problems like in this work, the discrete ordinate method becomes unwieldy, and computationally expensive. Additionally any errors introduced by discretization are systematic. As a result, the preferred methods for complex problems like the air-to-space neutron transport problem are stochastic methods.

Stochastic methods like MC is typically summed up as a *random walk* process. A more precise definition is that the MC method uses a stochastic model where the expected value of random variables is equivalent to the physical quantity being determined [7]. An example of a stochastic process neutron transport is whether a neutron that has interacted is absorbed or scattered. Whether a neutron is scattered or absorbed has a predefined probability based off of interaction cross sections. When randomly sampled many times, the expected value will converge to the probabilities in the actual physical system. Using the MC method has the main advantage is discretization is not needed: time, spatial, angle, and energy can all be treated continuously [4]. However, this comes at the cost of requiring enough histories, or particles followed, to converge to the correct solution. Because of the geometric and atmospheric complexity of this problem though, a MC approach was used. It is far less complex an approach to follow simulated pseudo-neutrons as they *random walk* back to the source, then to solve an approximated neutron transport equation directly. A thorough explanation of how this was implemented in SAHARA is found in Chapter 3.

2.3 Common Uses of the Adjoint Transport Equation

The adjoint transport equation is powerful, especially in certain classes of problems. One such example is the case where a detector region is small and the source region is large, therefore the probability of any one neutron reaching the detector is small [4]. This is a common enough technique that it is even a built in function within MCNP [12]. The adjoint is preferred here because the probability of a pseudo-neutron reaching back to the source volume is much greater than the forward case. It solves for the source characteristics, such as the strength of source, directly from a given detector response. A similar problem exists for radiation shielding problems. Radia-

tion shielding aims to reduce the exposure of any personal to extremely low values. Like the first problem, the probability of any one neutron, or other radiation particle, reaching the region of interest might be extremely low. To overcome this challenge a hybrid approach has been developed to reduce the variance in these calculations by using the adjoint flux to generate appropriate variance reduction parameters for a forward MC solution [13]. This method, called Consistent Adjoint-Driven Importance Sampling (CADIS), is now considered a standard in deep-penetrating radiation applications. This method also significantly reduces necessary computation time.

Other common uses for the adjoint method are included in a wide variety of nuclear reactor calculations. One such method is used to find the adjoint flux within a reactor [5]. The adjoint flux itself can then be used to perform perturbation analysis on changes to the estimated neutron multiplication caused by changes in material properties or time [4, 14]. The adjoint flux allows one to estimate the perturbed reactivity change without actually having to solve independently the complex perturbed system [15]. The adjoint method is also used to speed up convergence of reactor criticality calculations. Using a discrete ordinate approximation, the initial neutron distribution within the reactor for a forward MC calculation can be found [10]. This can increase the rate of convergence significantly over that of estimating the starting distribution. This is the same methodology used in shielding problems. There are also current production codes such as Automated VARIance reduction Generator (ADVANTG) dedicated to performing such initial adjoint calculations.

Most applications of the adjoint problem involve reducing the uncertainty and reducing computational time. For adjoint detector problems another added benefit is the inclusion source characteristics. The adjoint air-to-space problem is similar to low probability detector problems, but differs because the source is also considered a point. While the source is solved for using the adjoint method, the efficiency and

variance issues will still be present as they are in the forward problem.

Lastly, a critical feature present in the adjoint air-to-space problem not seen in the forward is the introduction of time as a constraint. As stated above, adjoint approaches have been used in perturbation theory to estimate the time dependence of reactors, or the time eigenvalue α [4, 14]. Essentially, these works use the set of forward and adjoint eigenfunctions to approximate the time-dependent response for small changes in time [16]. In the air-to-space problem the time component appears differently. Not only do sampled pseudo-neutrons have to make it back to the source location, but must do so with at time equals zero. The approach SAHARA uses to solve this problem is explained in section 3.4.

2.4 Assumptions of the Air-to-Space Neutron Transport Problem

In order to derive the neutron transport equation, and the subsequent adjoint neutron transport equation, the following assumptions are made in Lewis and Miller [4]:

1. *Particles may be considered as points.*
2. *Particles travel in straight lines between point collisions.*
3. *Particle-particle interactions may be neglected.*
4. *Collisions may be considered instantaneous.*
5. *The material properties are assumed to be isotropic.*
6. *The properties of nuclei and the compositions of materials under consideration are assumed to be known and time-independent unless explicitly stated otherwise.*
7. *Only the expected or mean value of the particle density distribution is considered.*

Because these are assumed in the derivation of the neutron transport equation there are also assumptions in both the deterministic and stochastic methods to solve it. However, most solution methods also assume the following assumptions which were examined in Lt Col Whitman Dailey's dissertation, *Special Features of the Air-to-Space Neutron Transport Problem* [1]. These assumptions are:

8. The source, scattering medium, and satellite are stationary in the same reference frame.
9. The composition, density, and temperature are uniform within regions that define the problem.
10. Radioactive decay of neutrons is not significant in the time scale relevant to the problem.

Within both the forward and adjoint air-to-space neutron transport problem, assumptions 2, 8, 9, and 10 are not necessarily valid [1]. One of the problems is simply the geometric scale. Most neutron transport codes were designed to be used for reactors and experimental designs whose geometric sizes are on the order of meters. The air-to-space neutron transport problem is on the scale of tens of thousands of kilometers. Neutron mean free paths also vary by six orders of magnitude from the bottom to the top of the atmosphere. The mechanisms by which these special features may influence the problem can be grouped into the topics of relative motions, gravity, atmospheric fidelity, and radioactive decay.

2.4.1 Relative Motions

Relative motions refers to the motion of the source, the scattering medium itself, and the point at which the flux is estimated. This topic violates assumption 8 of traditional neutron transport codes. The relative motions impact the rendezvous

problem, the divergence factor, the neutron/pseudo-neutron optical thickness in the atmosphere, and the pseudo-neutron velocity.

Since both the source and the satellite are moving in the Earth-Centered Inertial (ECI) frame of reference, the rendezvous calculations change. However, it has been shown that this problem has a minimal impact on the solution for measurement times less than 1 hour [1]. The problem scope states only a 1000 second time-energy neutron distribution will be used, and as such the satellite and source motion will not be considered. This assumption was made not only to simplify the problem, but the measurement time frame is representative of real world detectors.

Relative motion also affects the divergence factor of the problem. With gravity neglected (see following section), and relative motions neglected this factor is $1/r^2$. However, for short time problems the difference when including relative motions is negligible [1].

The optical thickness through the atmosphere changes when the rotation of the earth within the ECI is considered. Additionally the thermal motion, and bulk motion of the atmosphere will influence the interaction cross sections. However, both of these impacts were found to be negligible [1] and will be ignored. These results indicate that assumption 8 is valid for this problem.

2.4.2 Gravity

When gravity is included, neutrons/pseudo-neutrons no longer travel in straight lines, but in orbital trajectories. This violates assumption 2 of traditional neutral particle transport. Introducing gravity has a large effect on the rendezvous problem, the divergence factor, and the optical thickness through the atmosphere. For problems where transport occurs for longer than an hour, the effect of gravity is quite dramatic and can even result in over the horizon fluxes at the point of interest. However,

for this work's 1000 second time-energy neutron distribution the effect of gravity is negligible and is not modeled as shown in Figure 2. Thus, assumption 2 is valid.




Figure 2. Arrival energy as a function of time of flight from an equatorial source at $\Delta\alpha = 0^\circ$ at time of emission with and without gravity showing negligible gravitational effects below 1,000 seconds. Figure used with permission from [1]

2.4.3 Atmosphere Fidelity

As altitude increases, the composition, density, and temperature of the atmosphere varies continuously which violates assumption 9. The manner in which the atmosphere is treated can impact both calculation time and accuracy. This includes not only the model chosen, but also how the model is implemented [1]. Currently the standard atmosphere model for production level programs is NRLMSISE-00 [17] and is used in the DoD/DoE code DIORAMA[18]. DIORAMA is the current source-to-sensor production code for USNDS programmers, planners, and operational users [18]. However, for this work the US 1976 Standard Atmosphere is used as the NRLMSISE-00 adds additional complexities that is not needed for a research code such as this one. Additionally, the effect of including an atmosphere above 86 km is of little consequence and will be disregarded [1]. The fidelity of the atmospheric composition will

also influence the problem. However it is primarily composed of nitrogen, oxygen, and argon. Statistically other elements are of little significance. More information on the atmospheric model is found in section 3.8.

How the atmosphere is modeled also impacts the computation of the Effective Path Length (EPL) which is defined as the path length through a medium of uniform density, $\bar{\rho}$, having the same optical thickness through the actual medium [1]. This means that EPL and optical thickness in effect represent the same value. Optical thickness refers to the number of mean free paths through some medium whereas EPL is the geometric optical path length if the pseudo-neutron was traveling through air at sea level. EPL is expressed in *km*.

In a simplified model, one would only need two concentric spheres, one for the earth and one for the atmosphere of constant properties. However, this would negate the changing properties of the atmosphere with altitude. A more accurate approach breaks the atmosphere into thousands or more concentric spheres. Each layer has its own constant composition, density and temperature. This complicates the geometry in any ray-tracing procedure, increasing the computational cost of the problem. It also can cause a systematic accumulation of error [1]. Previous layer models used this piecewise-constant-density model, where each concentric spherical annuli has constant properties. This approach simplifies the calculations in any one region, but drastically increases number of boundary crossings. In the worst case, 2 million layers is needed for 6 digits of precision for the EPL[1].

Alternatively, a continuous representation could be used within a fewer layers. This would minimize the number of layers needed, but increases the layer model complexity.

Consider the simplest case, calculating the EPL from the surface of the earth to 86 km in the surface normal direction. For the continuous representation, the EPL is

$$L = \int_0^{86} \frac{\rho(z)}{\bar{\rho}} dz, \quad (3)$$

which can be computed using a n -point Gauss-Legendre quadrature. Alternatively a discrete representation can be used, where the atmosphere is a set of N concentric annuli each with n uniformly spaced shells within each atmospheric layer; each with its own uniform density. The uniform density in each shell is chosen to be the density at the center of the shell. The EPL is essentially a composite midpoint approximation to Equation 3. The EPL for this approximation is

$$L \approx \sum_{i=1}^N \frac{\Delta Z_{b,i}}{n} \sum_{j=1}^n \frac{\rho(z_j)}{\bar{\rho}_i}, \quad (4)$$

where ΔZ_b is the geometric thickness of the atmospheric layer, n is the number of shells in each layer (a constant), and z_j is the altitude at each shell. For the continuous representation, 6 digits of precision is obtained using only a 5 point quadrature whereas the discrete representation needs 3584 shells (512 shells per layer, 7 layer atmosphere model) to get the same precision for a vertical path [1]. The number of shells per layer would need to significantly increase for non-vertical paths. Additionally, the continuous representation of the atmosphere can be applied to a layered approach without degrading precision, but does add additional computational cost.

In practice the computation of the EPL is not the full path from the earth to the edge of the atmosphere, nor is it completely vertical. Both assumptions degrade the calculation precision. For any path not normal, the path curvature puts more of the path below the center altitude than above. Ideally the density would be evaluated at the center of the path s , but this is impractical and the loss of precision is low for a path angle $\cos(\theta) = \zeta > 0.07$ [1]. For a path other than vertical using a continuous representation with the path angle, $\zeta_o > 0.07$ can be formulated as

$$L = \int_0^s \frac{\rho(z(s))}{\bar{\rho}} ds = \int_0^{86} \frac{\rho(z)}{\bar{\rho}} \frac{R_e + z}{\sqrt{R_e^2 \zeta_o^2 + 2R_e z + z^2}}, \quad (5)$$

where R_e is the radius of the earth. For path angle, $\zeta_o < 0.07$, Equation 6 should be used [1].

$$L = \int_0^{\Delta s} \frac{\rho(z_o + \Delta z(s))}{\bar{\rho}} ds \quad (6)$$

Although using Equation 6 comes at a higher computation cost, it preserves the precision [1]. As a result, assumption 9 is valid for each atmospheric sub-layer. For further background on how the choice of atmosphere treatment affects the EPL please see reference [1].

2.4.4 Radioactivity Decay of Neutrons

A free neutron is not a stable particle, and has a half-life of approximately 10 minutes [19]. In most neutron codes, this is of little importance as a neutron lifetime is on the order of milliseconds [4]. However, in the air-to-space neutron transport problem, this is not the case as the problem can span several hours. It follows that for longer time, neutron decay impacts the resultant solution more. Although this work considers a 1000 second time-energy neutron distribution at the satellite, the fluence is low above 100 seconds for most sources, which is well below the decay half-life. Thus, neutron decay is considered negligible for the low fluence, 100-1,000 second lifetime pseudo-neutrons and assumption 10 remains valid.

III. Methodology

3.1 Methodology of SAHARA

A Python 3.7 code called SAHARA was developed and implemented to solve the adjoint neutron equation for the air-to-space problem. The code can be broken down into the six processes outlined below:

1. *Start Pseudo-Neutron:* Start a pseudo-neutron whose time-energy is sampled from the time-energy neutron distribution at the satellite location. At this point a check is done to see if it was possible it could have been a direct contribution. The pseudo-neutron is then moved to the edge of the atmosphere.
2. *Move pseudo-neutron:* The pseudo-neutron is transported to its next interaction point. This involves sampling the optical path length and transforming this to a geometric distance. The pseudo-neutron is then moved to the sampled location and its time is updated.
3. *Generate Source Event Pseudo-Neutron:* Along the path of the pseudo-neutron, check if a interaction point exists such that the energy from upscattering back to the source and the energy required from the time-of-flight (ToF) are equal. If the required criteria is met, adjust the source event pseudo-neutron weight due to probability of scattering, and the probability of non-interaction along the paths.
4. *Interact Pseudo-Neutron:* Sample the interaction nuclei and type. If absorption is sampled, scattering is forced since the particle had to come from the source. Use the properties of the pseudo-neutron energy after upscattering to correct the interaction probabilities.

5. *Kill Pseudo-Neutron:* Kill the pseudo-neutron if it has moved outside the range of the problem of interest (leakage, time or energy out-of-range or other kill criteria). If kill criteria is not met, repeat steps 2-4.
6. *Tally Pseudo-Neutron:* The contributions from all the source event pseudo-neutrons are tallied under one, pseudo-neutron energy tally.

The following sections provide details regarding each step's implementation.

3.2 Start Pseudo-Neutron

The start pseudo-neutron module selects the pseudo-neutron's initial time, energy, and direction. The time and energy is selected from the inputted time-energy neutron distribution at the satellite location. The direction of the pseudo-neutron is sampled uniformly from the solid angle subtended upon the earth's atmosphere.

3.2.1 Sampling the Distribution

An exemplar time-energy neutron distribution is shown in Figure 3. The general process for sampling any two dimensional distribution is:

1. Collapse distribution to one dimension by summing along either columns or rows
2. Construct a CDF from the one dimensional distribution
3. Sample one dimension's CDF and identify the sampled row or column bin_i
4. Construct a CDF from in the second dimension based within first dimension's sampled bin_i
5. Sample second dimension's CDF and identify the sampled row or column bin_j

6. Sample uniformly both dimensions within the selected $\text{bin}_{i,j}$




Figure 3. Example normalized time-energy neutron distribution at the satellite generated from HATS-n using a Watt fission spectrum.

This general process is used for the sampling method. However, different sampling results depending on the order of the dimension collapse. Figure 4 displays the sampled time-energy distribution when time was sampled first, followed by energy. The change in flux magnitude, and shape indicate that the distribution was not sampled correctly. When energy is sampled first the sampled time-energy distribution, shown in Figure 5, produces a distribution more similar to the nominal input distribution.




Figure 4. Sampled time-energy neutron distribution from nominal input distribution sampling time first, then energy.




Figure 5. Sampled time-energy neutron distribution from nominal input distribution sampling energy first, then time.


Although the distribution sampled in Figure 5 appears correct, a statistical test needs to be done to ensure they are the same distribution. A non-parametric Kolmogorov–Smirnov test K-S Test is preferred as it makes no assumptions about the underlying distributions [20]. The K-S Test quantifies the probability that two samples were created by the same underlying probability distribution. This was expanded to two dimensions by Peacock [21].

$$H_o : F(x) = G(x)$$

$$H_A : F(x) \neq G(x)$$

Kolmogorov-Smirnov Test of equal distribution

The K-S Test is used to test if the sampled time and energy CDFs come from the same distribution as the input time and energy CDFs. In both time and energy the null hypothesis is failed to be rejected indicating they are the same distribution. The time and energy collapsed sampled CDFs are shown in Figure 6.



(a) Sampled energy CDF with nominal spectrum energy CDF overlayed (b) Sampled time CDF with nominal spectrum time CDF overlayed

Figure 6. Time and energy collapsed CDFs for a Watt fission spectrum used for testing equal distributions

The distribution sampled from is binned data, and multiple options exist for deciding the pseudo-neutron time and energy once a bin is selected. One could chose

the midpoint value, bin edge values, or sample uniformly or by some other distribution. For this work, the time and energy value is sampled logarithmically uniformly between the end points of that bin. This was decided as it allows for a full range of values within the bin. The logarithmic sampling is specific to using HATS-n as the program generates input files with logarithmically spaced bins. If other programs are used to generate the input file this sampling methodology may need to be modified.

3.2.2 Direct Contribution Check

At this point in the algorithm, a time and energy has been selected for the pseudo-neutron. Additionally, the lower and upper values of the time bin, and the distance to the source location are known. With these values a check is done to see if the selected pseudo-neutron could have originated directly from the source. This step is done by computing the ToF to the source at that energy, and then checking if that time falls within the bounds of the time bin. If so, it is tallied as a separate, direct tally. The pseudo-neutron is still allowed to scatter even if tallied. It should be noted that the bin structure of the input file greatly affects the results of this direct contribution check, which will be discussed in chapter 5.

3.2.3 Sampling Angle

The pseudo-neutron's initial angle is sampled isotropically from the solid angle subtended from the satellite location to the Earth's atmosphere as pseudo-neutrons released outside of this angle range would not interact in the assumed vacuum resulting in impossible scatter to the source. Once the angle and energy are sampled, the initial velocity of the pseudo-neutron is calculated.

3.2.4 Move Pseudo-Neutron to the Edge of the Atmosphere

Once the initial velocity for the pseudo-neutron is found, it is then moved to the edge of the atmosphere and the particle's time is updated. A check is done to determine if the pseudo-neutron time has become negative demonstrating that the particle could not have traveled that distance with the current velocity and come from the source. If so, a new pseudo-neutron is sampled in order to find a velocity such that the particle could have come from the source. This check concludes the start pseudo-neutron step. Of note, for a Watt fission spectrum source approximately 80 percent of sampled pseudo-neutrons are rejected using this method. More effective direction sampling techniques certainly exist, but the computational cost of rejecting and sampling a new pseudo-neutron is low.

3.3 Move Pseudo-Neutron

The move pseudo-neutron step finds the location of the next interaction and moves the particle to that location. This step involves calculating and sampling the atmosphere's optical thickness, transforming the sampled optical thickness into an EPL, and root-solving the EPL to find the geometric distance traversed. The methodology of this process is reused from HATS-n [1].

3.3.1 Calculating Optical Thickness and Effective Path Length

Dailey states, "The effective path length is defined as the path length through a medium of uniform density, $\bar{\rho}$, having the same optical thickness through the actual medium" [1]. For a general path, EPL, L , can be defined as:

$$L = \int_0^{\Delta s} \frac{\rho(z_o + \Delta z(s))}{\bar{\rho}} ds, \quad (7)$$

where z_o is the starting altitude, Δs is the path length, and $\Delta z(s)$ is the changed in altitude with position s defined in Equation 8. r_o is the distance from the center of the Earth to the start of the path and ζ_o is the cosine of the zenith angle. The zenith angle is the angle of the direction of the pseudo-neutron perpendicular to the surface of the earth.

$$\Delta z(s) = \frac{s^2 + 2r_o\zeta_o s}{r_o + \sqrt{r_o^2 + s^2 + 2r_o s\zeta_o}} \quad (8)$$

In SAHARA, the atmosphere is broken down into multiple layers, and approximated via a Gaussian quadrature. The EPL through any layer b is

$$\tilde{L}_b = \frac{\Delta s_b}{2\bar{\rho}} \sum_{i=1}^n w_i \rho(Z_{b-1} + \Delta z(s_i)) \quad (9)$$

This formulation is used if the zenith angle, ζ_o , is less than 0.07, but this comes at a significant computational cost [1]. For $\zeta > 0.07$, Equation 7's $z(s)$ can be changed to z which improves computational efficiency. The solution is also numerically better conditioned if the integral is taken in upward paths. If a path is downwards, the bounds of the integral is switched which has no effect on the value of the EPL. This formulation is shown in Equation 10.

$$L = \int_{z_o}^{z_o + \Delta z} \frac{\rho(z)}{\bar{\rho}} \frac{r_o + z}{\sqrt{r_o^2 \zeta_o^2 + 2r_o z + z^2}} \quad (10)$$

Equation 10 can also be approximated by Gauss-Legendre quadratures.

The path length through any given layer is

$$\tilde{L}_b = \frac{\Delta Z_b}{2\bar{\rho}} \frac{w_i \rho(Z_{b-1} + z_i) ((R_e + Z_{b-1}) + (Z_{b-1} + z_i))}{\sqrt{(R_e + Z_{b-1})^2 \zeta_o^2 + 2(R_e + Z_{b-1})(Z_{b-1} + z_i) + (Z_{b-1} + z_i)^2}}, \quad (11)$$

where ΔZ_b is the change in altitude in a given layer b and z_i corresponds to the

Gauss-Legendre abscissa a_i ,

$$z_i = \frac{\Delta Z_b}{2}(a_i + 1). \quad (12)$$

Because the full layers in Equation 11 only depend on ζ_o , the quadrature points can be pre-computed to decrease the computational cost.

With the EPL calculated, the probability of not leaking to the edge of atmosphere is

$$P_{NL} = 1 - e^{-\Sigma_T L_{total}}, \quad (13)$$

where Σ_T is the total cross section of the atmosphere at sea level. The P_{NL} can be used to suppress the pseudo-neutrons ability to leak out of the system. Because this is a forced result, the weight would need to be updated accordingly. If leakage is suppressed, the optical thickness, τ , is then sampled by:

$$\tau = \ln(1 - \xi P_{NL}), \quad (14)$$

where ξ is a random number between 0 and 1. This randomly sampled optical thickness can then be converted back to EPL:

$$L = \frac{\tau}{\Sigma_t}. \quad (15)$$

3.3.2 Finding Geometric Distance

The sampled EPL is converted to a geometric distance through a root-solving method. Equation 7 or 10 needs to be solved for s or z . In the case of the latter, z needs to be converted back to s . This portion of SAHARA was ported from HATS-n. Once the distance traveled Δs is found, the pseudo-neutron's location and the time is updated.

3.3.3 Verifying Move Pseudo-Neutron

The move pseudo-neutron step was verified by running 10,000 histories to observe how the location of sampled first interactions is distributed. The histogram for this distribution is shown in Figure 7. Since the pseudo-neutron's flight path is downwards towards the surface of the earth, it follows that most interactions occur at lower altitudes because as the pseudo-neutron travels to lower altitude, density increases, shortening path lengths. As the atmosphere gets more dense, at around 40 km, more and more pseudo-neutrons interact. The histogram should then decrease after the peak as there is a reduced probability that pseudo-neutrons reach that altitude. Figure 7 meets all of the physics-based expectations for correct implementation of the interaction location.




Figure 7. First sampled interaction altitude of SAHARA verifying move pseudo-neutron step using a N-14 atmosphere

3.4 Source Event Pseudo-Neutron

In order to improve efficiency for similar problems, a next event estimator is commonly used. For the forward problem, this means at every *sampled* interaction

point, a pseudo-neutron is created and forced to the region of interest. Its weight is then corrected by the probability of non-interaction along the path, and probability of scattering in that direction. This method works well in the forward problem as there are no constraints on the generated pseudo-neutron. However, this version of the next event estimator fails in the adjoint problem.

Every pseudo-neutron has both a time and energy. If the same next event estimator technique was used in this problem and a “next event pseudo-neutron” was forced to the source location at each sampled interaction point, a solution would not be possible due to the over constrained nature of the adjoint problem. Because the sampled pseudo-neutron has both energy and time, there are two ways to determine the energy at the source of the “next event pseudo-neutron”. Energy can be calculated either from the scattering angle in the lab frame to scatter to the source, or the ToF remaining to reach the source at time equal to zero. Because of this, the typical next event estimator method is invalid as the probability of both definitions of energy being equal at a sampled interaction location is near zero. Thus a new method for a next event estimator was developed called the source event estimator.

3.4.1 Source Event Estimator

Due to the issues highlighted in the previous section the source event estimator was developed to improve the efficiency of SAHARA. Instead of forcing a pseudo-neutron to be generated at the sampled interaction point, as the required energies from both the scattering angle and ToF are unlikely to match, the point along the flight path is searched to see if such a point exists where the required energies do match. It is important to recognize this solution is independent of the sampled next event interaction location as shown in Figure 8.




Figure 8. Two examples of solutions to the source event estimator. (a) shows where the solution is above the sampled interaction point. (b) shows where the solution is below the sampled interaction point.

To improve the efficiency of this technique, the path is first checked to determine if a solution exists. This is done by checking the atmospheric boundary layer intersection points across all atmospheric constituents, and computing the required energies. If such a point exists, then a root-solving technique is used to find the solution point within the corresponding layer. It is important to note that there are 2 convergence criteria, one for energy and one for location.

Because this is a forced interaction, the weight of the source event pseudo-neutron must be updated according to the probability of non-interaction along the path and the scattering probability.

3.4.2 Calculation of Required Energies

As previously stated the required upscattered energy can be computed both from the ToF remaining to reach the source, and from the required scattering angle in the lab frame. Calculating the required energy, E_{tof} , in keV from the remaining ToF is a straight forward calculation shown in Equation 16 where s is the distance to the source in km and κ is a conversion constant.

$$E_{tof} = \left(\frac{s\kappa}{t_{tof}} \right)^2 \quad (16)$$

Calculating the required upscattered energy from the scattering angle in the lab frame is a trivial. Finding the scattering angle in the lab frame is straight forward by taking the dot product of the unit direction vector of the source event pseudo-neutron $\hat{\Omega}$ and the unit direction vector of the path to the source $\hat{\Omega}_s$. The relationship between scattering angle in the lab (θ_{lab}) and center-of-mass (θ_{cm}) frame is shown in Equation 17 where A is the mass number of the scattering isotope.

$$\cos(\theta_{lab}) = \frac{1 + A \cos \theta_{cm}}{\sqrt{A^2 + 2A \cos \theta_{cm} + 1}} \quad (17)$$

Solving Equation 17 for θ_{cm} yields two possibilities depending on θ_{lab} shown in equation 18 [7, 22, 23].

$$\cos \theta_{cm} = \begin{cases} \frac{\sqrt{2}A \cos \theta_{lab} \sqrt{2A^2 - 1 + \cos 2\theta_{lab} + 2A(\cos^2 \theta_{lab} - 1)}}{2A^2}, & \text{if } \theta_{lab} < \frac{\pi}{2} \\ -\frac{\sqrt{2}A \cos \theta_{lab} \sqrt{2A^2 - 1 + \cos 2\theta_{lab} + 2A(\cos^2 \theta_{lab} - 1)}}{2A^2}, & \text{if } \theta_{lab} > \frac{\pi}{2} \end{cases} \quad (18)$$

Once the scattering angle in the center-of-mass frame is found, the upscattered

energy, E' , can then be calculated via Equation 19 [7, 24].

$$E' = \frac{(A + 1)^2}{1 + 2A \cos(\theta_{cm}) + A^2} E \quad (19)$$

3.4.3 Probability of Non-Interaction Along Flight Path

Because the source event estimator is a forced interaction, the pseudo-neutron weight needs to be adjusted. In the next event estimator case, the weight is equal to the probability of non-interaction, P_{NI} along the flight path to the region of interest. However, in the source event estimator case it also needs to account for the flight path to the required interaction point. Referring back to Figure 8, this is the flight path to the required source event interaction point, L_1 , and the source event flight path, L_2 . The total probability of non-interaction is given in Equation 20, keeping in mind the cross sections are taken at the original energy, E , for L_1 , and the upscattered energy, E' , for L_2 .

$$P_{NI} = e^{-L_1 \Sigma_t^*(E)} e^{-L_2 \Sigma_t^*(E')} \quad (20)$$

The weight of the source event pseudo-neutron is increased according to the probability of non-interaction, Equation 21. At this point it is worth noting that this source event estimator method has not been done in this way before. A version of the source event estimator was developed for the computer program Monte Carlo N-particle Adjoint (MCNA) [12], but it lacks the root solving component done within SAHARA. This causes interesting results as the weights of the source event estimator gets extremely large. This will be discussed further in chapter 4.

$$w_{new} = \frac{w_{old}}{P_{NI}} \quad (21)$$

3.4.4 Probability of Scattering

The last update to weight needed is the probability of scattering. This probability is dependent on the probability of scattering off a given isotope, P_{iso} in the atmosphere and the probability of scattering for that isotope at the upscattered energy. Depending on the atmosphere composition, the source event pseudo-neutron loops through all isotopes until a possible source event interaction is identified. Since the isotope selection is not a random process, it must be corrected for. Thus, the final weight is increased according to Equation 22.

$$w_{new} = \frac{w_{old}}{\frac{\sigma_s(E')}{\sigma_t(E')} P_{iso}} \quad (22)$$

3.5 Interact Pseudo-Neutron

Returning to the original pseudo-neutron, the location point is known and thus sampling which nuclei it interacts with is straight forward. Once the nuclei is chosen, the pseudo-neutron can then either scatter or be absorbed, although it is recommended that scattering is forced. Isotropic scattering in the center-of-mass frame was assumed and sampled angles are chosen according to Equations 23 and 24 where ξ_1 and ξ_2 are a random numbers between 0 and 1.

$$\cos \theta = 1 - 2\xi_1 \quad (23)$$

$$\phi = 2\pi\xi_2 \quad (24)$$

Upscattered energy is then found via Equation 19. The new direction is updated from Equation 25 [25]. If $1 - \gamma^2$ approaches zero, Equation 26 should be used.

$$\begin{aligned}
\alpha' &= \alpha \cos \theta + \gamma \alpha \frac{\sin \theta \cos \phi}{\sqrt{1 - \gamma^2}} - \beta \frac{\sin \theta \sin \phi}{\sqrt{1 - \gamma^2}} \\
\beta' &= \beta \cos \theta + \gamma \beta \frac{\sin \theta \cos \phi}{\sqrt{1 - \gamma^2}} + \alpha \frac{\sin \theta \sin \phi}{\sqrt{1 - \gamma^2}} \\
\gamma' &= \gamma \cos \theta - \sqrt{1 - \gamma^2} \sin \theta \cos \phi
\end{aligned} \tag{25}$$

$$\begin{aligned}
\alpha' &= \alpha \sin \theta \cos \phi \\
\beta' &= \beta \sin \theta \sin \phi \\
\gamma' &= \gamma \cos \theta
\end{aligned} \tag{26}$$

The weight of the pseudo-neutron is updated according to the cross sections at upscattered energy by Equation 27. Again weight is increasing as occurs in other adjoint transport methods [5]. If the weight gets too large, (recommended greater than 3 [5]), the particle could be split, however this is not currently available in SAHARA.

$$w_{new} = \frac{w_{old}}{\frac{\sigma_s(E)}{\sigma_t(E)}} \tag{27}$$

3.6 Kill Pseudo-Neutron

There are four criteria for killing a pseudo-neutron history in SAHARA, two of which are user defined. The first occurs if the energy of the pseudo-neutron upscatters above a given level. This value is outside the expected source range for this problem as the probability of a neutron emitted from nuclear fission greater than 20 MeV is approximately 3.3E-7. Another kill criteria option occurs when the pseudo-neutron leaks out of the system if leakage is not suppressed. Thirdly, if the ToF of the pseudo-neutron becomes negative, then the pseudo-neutron can no longer reach the source

at detonation time. Finally, the number of allowed scatters can be adjusted. If none of the kill criteria is met the steps described in sections 3.3-3.5 are repeated.

3.7 Tally Pseudo-Neutron

All of the source event pseudo-neutrons are tallied under one event history and are added to the total of the simulation. They are not scored separately to preserve appropriate figure of merit statistics. Straight counts are tallied along with a running average importance value to reduce computational overhead. A user input to the tally eliminates scoring pseudo-neutrons above a certain weight. In order to limit the memory requirements, an approach for computing variance described in *Monte Carlo Particle Transport Methods: Neutron and Photon Calculations* is used [7]. The average importance is calculated via Equation 28. This approach allows just the sum of the importances to be saved.

$$\bar{\mu} = \frac{1}{n} \sum_{i=1}^n \mu_i \quad (28)$$

The standard deviation of the mean and relative standard deviation of the mean in each bin is computed using Equations 29 and 30. This again eliminates the need to store all the values, but rather just the squared sum.

$$s_{\bar{\mu}} = \sqrt{\frac{1}{n(n-1)} \sum_{i=1}^n \mu_i^2 - \frac{1}{n} \left(\sum_{i=1}^n \mu_i \right)^2} \quad (29)$$

$$s_r = \sqrt{\frac{n}{n-1} \left[\frac{\sum \mu_i^2}{(\sum \mu_i)^2} - \frac{1}{n} \right]} \quad (30)$$

3.8 Atmosphere Representation

The atmosphere in SAHARA can be run with either a isothermal representation or the USSA76 representation [3]. Although both are functional in SAHARA, all results presented in future sections use the USSA76 implemented using the following equations and constants within the code. To maintain consistency with the published atmospheric model the constants in Table 1 are used.

Table 1. Physical constants for 1976 U.S. Standard Atmosphere [3]

Constant	Symbol	Value
Earth Radius	$R_{\oplus}$	6356.766 <i>km</i>
Acceleration due to gravity at sea level	g_o	6.80665 <i>m/s²</i>
Gas Constant	R^*	8314.32 <i>N · m/(kmol · K)</i>
Mean molecular weight at sea level air	M_o	28.9644 <i>kg/kmol</i>
Avogadro constant	N_A	6.022169x10 ²⁶ <i>kmol⁻¹</i>

3.8.1 Sea Level to 86 kilometers

The atmosphere above 86 km is sparse, and was shown by Dailey [1] to have little effect on the air-to-space neutron transport problem. Thus only the atmosphere representation in USSA76 from sea level to 86 km is used in SAHARA. Below 86 km, the atmosphere is divided into seven layers as defined by the geopotential height, H , which is related to geometric altitude, Z , by Equation 31. The units of geopotential height is km'.

$$H = \frac{R_{\oplus} Z}{R_{\oplus} + Z} \quad (31)$$

3.8.2 Temperature

The molecular temperature of the atmosphere is a function of geopotential height and is given by

$$T(H) = f_M(Z)(T_b + L_b(H - H_b)), \quad (32)$$

where the subscript b represents the value at base of the layer. L is the lapse rate, and f_M is the function used to interpolated the ratio of the mean molecular weight at sea level to the mean molecular weight at the geometric altitude Z corresponding to geopotential height H [1, 3]. This is given by Equation 33. i in Equation 33 corresponds to Table 3. The values for T_b , and L_b are listed in Table 2 along with base pressure, P_b which will be discussed in the next section. The mean molecular weight ratios are listed in Table 3. For reference, a plot of temperature vs geometric altitude is also provided in Figure 9. Although the variation in temperature of the atmosphere is not considered for Doppler broadening, which was shown to have negligible effect in the air-to-space problem [1], it's important when calculating atmospheric density.

$$f_M(Z) = (1 - 2(Z - Z_{i-1}))\left(\frac{M_o}{M}\right)_{i-1} + 2(Z - Z_{i-1})\left(\frac{M_o}{M}\right)_i \quad (33)$$

Table 2. Reference levels, and temperature lapse rates and computed base temperatures and pressures from the surface to 86 geometric kilometers [3, 1].

Subscript b	Geopotential Height [km'] H_b	Lapse Rate [K/km'] L_b	Base Temp [K] T_b	Base Pressure [N/m ²] P_b
0	0	-6.5	288.15	101325
1	11	0.0	216.65	22632.0336239
2	20	1.0	216.65	2474.87437676
3	32	2.8	228.65	868.014988511
4	47	0	270.65	110.905629144
5	51	-2.8	270.65	66.9384346264
6	71	-2.0	217.65	3.9563844998

Table 3. Molecular weight ratios for various geometric altitudes and geopotential heights [3, 1].

Geometric Altitude [km] Z_i	$(M_o/M)_i$
≤ 80	1.000000
80.5	0.999996
81.0	0.999989
81.5	0.999971
82.0	0.999941
82.5	0.999909
83.0	0.999870
83.5	0.999829
84.0	0.999786
84.5	0.999741
85.0	0.999694
85.5	0.999641
86.0	0.9995788




Figure 9. Temperature as a function of geometric altitude below 86 km verifying correct implementation of the USSA76.

3.8.3 Density

To compute atmospheric density the pressure is needed. The pressure in each layer is given by

$$P(H) = \begin{cases} P_b \left(\frac{T_b}{T(H)} \right)^{\frac{g'_o M_o}{R^* L_b}}, & \text{if } L_b \neq 0 \\ P_b e^{-\frac{g'_o M_o (H - H_b)}{R^* T_b}}, & \text{if } L_b = 0 \end{cases} \quad (34)$$

where P_b is the pressure at the base of the layer. These values are included in Table 2. Atmospheric pressure as a function of geometric altitude is shown in Figure 10. Atmospheric pressure is key to calculating atmospheric density.




Figure 10. Atmospheric pressure as a function of geometric altitude below 86 km verifying correct implementation of the USSA76.

With both pressure and temperature defined, atmospheric density as a function of geopotential height is given by

$$\rho(H) = \frac{M_o P(H)}{R^* T(H)} \quad (35)$$

[3]. The atmospheric pressure plotted as a function of geometric altitude is shown in Figure 11. Air density is key to computing macroscopic atmospheric cross sections as below 86 km density may be approximated as directly proportional to the number density of the atmosphere.




Figure 11. Atmospheric density as a function of geometric altitude below 86 km verifying correct implementation of the USSA76.

3.8.3.1 Composition

At geometric altitudes below 86 km, the atmosphere is considered homogeneously mixed. This is an assumption within USSA76 and is continued in SAHARA. The relative fractions of constituents constant from sea level are listed in Table 4.

Table 4. Molecular weights and fractional values for constituents for sea level atmosphere [3].

Species	Molecular Weight [kg/kmol]	Fractional Volume F_i
N_2	28.0134	0.78084
O_2	31.9988	0.209476
Ar	39.948	0.00934
CO_2	44.00995	0.000314
Ne	20.183	0.00001818
He	4.0026	0.00000524
Kr	83.80	0.00000114
Xe	131.30	0.000000087
CH_4	16.04303	0.000002
H_2	2.01594	0.0000005

Although more components are listed in Table 4, SAHARA only has options to use the top three: nitrogen, oxygen, and argon. The nitrogen isotopes included are 14 and 15. The isotopes of oxygen included are 16, 17, 18. In addition, argon 40 is included. When all isotopes are considered the isotopic fractions are shown in Table 5.

Table 5. Molecular volumetric fractions for a atmosphere containing all isotopes in SAHARA.

Element	Isotope	Fraction
N	14	0.77771664
	15	0.00312336
O	16	0.2089732576
	17	8.37904×10^{-5}
	18	0.000418952
Ar	40	0.00934

3.8.4 Cross Section Data

SAHARA uses ENDF/B-VII.1 cross sections [26] which are formatted according to ENDF-6 [27]. Oxygen 18 is not included in ENDF/B-VII.1, this cross section data comes from FENDL-3.1d cross section data [28]. The wrapper to retrieve cross section values is the Python package called OpenMC [29]. OpenMC contains many other features for nuclear data processing. The use of OpenMC is highly recommended for any nuclear related programming in Python. However, directly calling OpenMC for cross section values is computationally expensive. To speed up calculations, OpenMC was used to generate lookup tables which are then linearly interpolated. The lookup tables were generated using 100 log spaced bins per decade from 0 to 1 keV, and 1000 log spaced bins per decade thereafter.

3.9 Overview of SAHARA

The SAHARA code is a purpose-built adjoint Monte Carlo code written in Python 3.8 to simulate an adjoint neutron from a position, energy and time in space to a source within the atmosphere. At this time SAHARA has not been tested for sources

outside of the atmosphere. SAHARA is not a production code, but rather a code to research approaches and solutions to the air-to-space adjoint transport problem.

3.9.1 Commonality with HATS-n

SAHARA was originally envisioned to be written in the same language as HATS-n, Modern Fortran, allowing for ease of use by reusing a large portion of the code, namely the definition of the atmosphere, cross sections, and the effective and geometric path length calculations. However, the decision was made early on to switch to Python. Although no longer in Fortran, HATS-n functions and methodology, namely atmosphere and path length functions, were rewritten in Python using the same processes as HATS-n (with the creator’s permission). Overall, SAHARA contains some underlying modules and subroutines that are identical in function as HATS-n, but the functionality of SAHARA as whole is a separate entity.

3.9.2 Target Architecture

Because SAHARA utilizes a open source program in Python, any architecture can be used. SAHARA currently only runs single threaded, but parallelization is recommended for future work. The Fortran subroutines directly ported using the Python package Numpy have been completed for Mac, Windows, and Linux operating systems.

3.9.3 SAHARA Inputs

SAHARA features multiple options to constrain and test different functionalities of the problem space. The first input is the atmosphere definition. Here either the USSA76 or an isothermal atmosphere can be chosen. The constituents of the atmosphere can be picked from the options outlined in Table 6.

Table 6. Atmosphere constituents options in SAHARA.

Single Isotope Atm	Mixed Isotope Atm
N_{14}	$N_{14}O_{16}$
N_{15}	$N_{14}O_{16}Ar_{40}$
O_{16}	$N_{14}N_{15}O_{16}O_{17}O_{18}Ar_{40}$
O_{17}	
O_{18}	
Ar_{40}	

Scatter and leakage criteria is specified next with a flags to force scatter and suppress leakage. All results in this work force scatter and suppress leakage. Additionally there are flags built in to include inelastic and anisotropic scattering, however these features are currently not available. It is recommended these features be added in future work. Kill criteria options include the maximum number of scatters allowed, energy cut-off and weight cut-off. The number of histories is specified along with the satellite measurement input file.

3.9.4 Modules and Descriptions

The list below outlines all the modules and their purpose within SAHARA. In general, modules were broken down according to their purpose.

1. SAHARA Driver: Main program that manages the setup, runs histories and generates results files and plots.
2. Data Analysis: Contains the routines for parsing through tally classes, saving data to file, and generating plots.
3. MC Pseudo-Neutron: Contains the subroutines for executing a single and loop

pseudo-neutron history. Creation, raytracing, scattering, source event estimation, tallying and killing are included.

4. Tally Module: Contains the tally class, and initializes tally and update tally subroutines.
5. Atmosphere Module: Contains variables and functions needed to build the atmosphere.
6. Build Atmosphere: Contains the atmosphere class, EPL class; builds the atmosphere class that saves all relevant information of the atmosphere including cross section data.
7. Cross Sections: Module that loads the cross section look-up tables and generates interpolation functions to be stored in the atmosphere class. Also contains functions to build new lookup tables if needed.
8. Gauss Legendre: Contains lists of Gauss-Legendre weights and abscissas. Directly converted from HATS-n using `numpy.f2py`.
9. Global Variables: Contains variables commonly used across all modules.
10. HASTE Functions: Contain simple subroutines directly ported from HATS-n using `numpy.f2py`.
11. Initialize Pseudo-Neutron: Contains the definition of the pseudo-neutron class and functions to move it to the edge of the atmosphere.
12. Neutron Utilities: Contains subroutines for converting between pseudo-neutron energy, speed, and velocity.
13. Pathlengths: Contains subroutines necessary for computing EPL and converting EPL to geometric distance. Converted from HATS-n subroutines into Python.

14. Sample Pseudo-Neutron Time Energy: This module reads the input file and builds the input distribution for sampling. Also contains functions to sample time, energy and direction of the pseudo-neutron.
15. Scatter Model Module: Contains the scatter model class and function to initialize it.
16. Source Information: Contains the source class and function to initialize it.
17. Utilities: Contains many different subroutines such as calculating required energies and bisection searches.

3.10 Summary

This chapter laid out the methodology behind SAHARA and the features of the code. The main steps to find the solution to the adjoint transport equation in SAHARA are:

1. *Start Pseudo-Neutron,*
2. *Move Pseudo-Neutron,*
3. *Source Event Pseudo-Neutron,*
4. *Interact Pseudo-Neutron,*
5. *Kill Pseudo-Neutron,* and
6. *Tally Pseudo-Neutron.*

IV. Once-Scattered Scenarios

The SAHARA approach to the space-to-air neutron transport problem is a new method which characterizes the neutron energy spectrum of a source within the atmosphere directly from a time-energy spectrum at the satellite. Additionally, the formulation of the source event estimator within SAHARA is a novel Monte Carlo approach to solving the adjoint neutron transport problem. The simulated scenarios in this work are primarily used to demonstrate the feasibility of this approach, and not a full validation study. The simulations described in this section use once-scattered input spectra from various mono-energetic and continuous energy sources at a single source position in the atmosphere. Later, the results will be compared to the same spectra, but using a multi-scattering approach.

The input is generated from the forward transport code HATS-n developed by Lt Col Whitman Dailey [1]. A key fact about the input from HATS-n is that the time component is the time from particle creation. This means that the distribution does not start at the time equals zero, but rather the ToF of the most energetic particle and does not have to be corrected for.

The following results use time-energy distributions generated from HATS-n, which has the ability to turn off all contributions to the satellite except once-scattered neutrons. To duplicate the results, SAHARA's direct contribution portion is turned off. The settings for these scenarios are shown in Table 7. Proper convergence is commonly listed as 0.1% or less relative variance [7]. However, only five to ten million histories are run depending on the source; thus proper convergence is not always achieved. Also, the starting weight of each particle is $1/n_{hist}$. This step was performed to both increase the number of particles that reach the source location (due to the weight cut-off), and allow for easy normalization.

Table 7. SAHARA settings for once-scattered scenarios.


Setting	Value
Source Position (ECI)	6421, 0, 0
Source Altitude [km]	50
Satellite Position (ECI)	42164, 0, 0
Satellite Altitude [km]	35,793
Weight Cutoff	1×10^{23}
Energy Cutoff [MeV]	20
Scatters Allowed	1
Energy Bins per Decade	50
Atmosphere Constituents	N_{14}
Mono-energetic Source	100 keV 1 MeV 14.056 MeV
Histories	5×10^6
Initial Weight	5×10^{-6}
Source Spectrum	Uniform 100 keV - 1 MeV Watt Fission Spectrum
Histories	1×10^7
Initial Weight	1×10^{-7}

4.1 Mono-energetic Source Results


To remove the result's dependence on the number of source particles, typically the number of counts are normalized by the total number of source particles counts. However, SAHARA does not know how many source particles (aka the yield) produced the measured spectrum. Thus, the counts are normalized by the total number of

detected particles. If SAHARA is working correctly, the true source distribution should be reconstructed accurately. Figure 12 displays results for the three mono-energetic sources. For all energies, the peak does not correspond to the source energy's correct energy bin warranting further investigation. In addition, a mono-energetic source input creates a broadened adjoint source. The relative error on the second axis corresponds to the relative error due to counting statistics within the bin, Equation 36.

$$\sigma_{rel} = \frac{\sqrt{counts_i}}{counts_i} \quad (36)$$




(a) Once-scattered PDF for 100 keV source. Peak bin is below 100 keV. (b) Once-scattered PDF for a 1 MeV source. Peak bin is below 1 MeV.




(c) Once-scattered PDF for a 14.056 MeV source. Peak bin is below 14.056 MeV.

Figure 12. Once-scattered PDFs for mono-energetic spectra. All peaks are lower than the true source values. Energy bins with counts greater than 500 are plotted.


To perform a Gaussian fit to the data, the program *Origin Pro 2020b* was used [30]. The Gaussian line of best fit for each energy are shown in Figure 13. The summary of each best fit peak energy, σ , and FWHM is found in Table 8. With only three data points trends are not definitive, however peak energies found by SAHARA are less than the mono-energetic source by an average of 4.8% percent, while the standard deviation of the Gaussian curve is between 7-8%.



(a) Once-scattered PDF for 100 keV source spectrum with weighted Gaussian best fit. Gaussian mean is below 100 keV.



(b) Once-scattered PDF for a 1 MeV source spectrum with weighted Gaussian best fit. Gaussian mean is below 1 MeV.



(c) Once-scattered PDF for a 14.056 MeV source spectrum with weighted Gaussian best fit. Gaussian mean is below 14.056 MeV

Figure 13. Once-scattered PDFs for various mono-energetic spectra with weighted Gaussian best fits applied. Gaussian means are below actual source energies. Energy bins with counts greater than 500 are plotted.

Table 8. Fitting parameters to once-scattered Gaussian best fits from Origin Pro.

Source	Peak Energy [keV]	Relative Error %	Std Dev % of Peak	FWHM % of peak
100 keV	96.416	3.854	8.191	19.288
1 MeV	941.146	5.885	7.385	17.391
14.056 MeV	13375.520	4.841	7.334	17.270

4.2 Examining Once-Scattered Peak Energy Shift

There are a few possibilities as to why the peak energy found is consistently five percent lower than the actual source energy. First, the data points used to perform the line of best fit comes from binned data. The x-value (energy) of the points plotted in Figure 13 are the logarithmic centers, which are lower than if the geometric center was used. For example, the geometric center between 10 and 20 is 15, but the logarithmic center is 14.14. Theoretically more pseudo-neutrons could also have energies at the higher end of the bin. This being said, the peak bin still is not at the correct energy. For example, in the 14.056 MeV source's case it still fails to have the largest peak in the correct energy bin. Given that this is an inverse problem there is an infinite number of possibilities of source spectra. The hope is that this the average of all these solutions converge to the actual source. As a result, the source's identified energy is within the standard deviation of the Gaussian mean demonstrating that SAHARA can achieve an accurate prediction of a mono-energetic source's energy.

Examining counts for mono-energetic, once-scattered sources provided a baseline of how SAHARA was performing, as well as what to expect when moving towards continuous sources. However, this was only particle counts and did not take into account the importance weighting.

4.3 Mono-energetic Importance Results

Recall that neutron importance is a dimensionless quantity, but because its solution is numerically coincident with the forward solution, it can be thought of as a fluence. Figure 14 shows the average importance per bin for the three mono-energetic sources.




Figure 14. Once-scattered normalized average importance for a 100 keV, 1 MeV, and 14.056 MeV source and their relative error within each bin. Energy bins with counts greater than 500 are plotted.

Figure 14 does not match expectations. The initial expectation is that there would be a similar result to what was seen in examining normalized counts, a Gaussian peak around 1 MeV for Figure 14b, as an example. This difference is caused by both the

construction of the adjoint problem, and the how the source event estimator was developed. When doing an adjoint problem, the pseudo-neutron weight, or importance is increasing as its transported backwards in time, and upwards in energy if scatter and non-leakage is forced.


Examining the forcing of scatter, assume the probability of scattering at a given energy is 50%. In the forward problem, if scattering was forced the weight of that particle would be decreased by half, i.e. that neutron is half as important after the forced scatter. It follows that, if working backwards, if a pseudo-neutron upscattered to that energy and scattering is forced, the weight should double. However, this can create large weights, impacting convergence. In these cases particle histories can be split which allows for better convergence [5], but with the isotopes in the atmosphere weights increase slowly because their absorption cross sections are low. The problem of the increase in weight most likely does not lie in the original pseudo-neutron, but rather the use of the source event estimator.

Re-examining SAHARA's approach and the source event estimator, the problem becomes more clear. In general Monte Carlo approaches, a forced particle action results in a need to modify the weight accordingly to preserve probabilities. For the source event estimator there are four things being forced; isotope of interaction, scatter, interaction location and termination location. Further more, because this is an adjoint approach the weight should increase according to those parameters. Since the probability of non-interactions along the paths (see figure 8) dominates, the weights increase rapidly and can even cause numerical overflows if no weight cutoff is employed. As a result, tallied importances have can have large variances, especially when there are not sufficient counts in a bin. In these single scatter cases the SAHARA weight cutoff is set to 1×10^{23} . This creates a large span in value weights and thus allows outliers to heavily impact the resulting average importance.


This approach also creates an interesting dichotomy when examining the mono-energetic sources. Looking at 1 MeV source, the lower average importance's are around the actual source energy with peaks at the lower and upper ends in energy. For the lower energies, this could be due to increases in the total cross section, which would then increase weight. At higher energies, in order to successfully match energies, the location of these required scattering points could require longer paths and thus increases the weight. A recommended further study would be to investigate the required scattering locations as a function of energy and time of the pseudo-neutrons.

The outliers sensitivity and poor convergence is a problem for mono-energetic sources, and makes gaining any useful information out of continuous source spectra difficult. Consider that *importance* in this sense could be thought of as estimating how important pseudo-neutron energies are to generating the source spectrum, and not the source spectrum itself. So when comparing Figure 12 to Figure 14, one could possibly use these results to modify Figure 12 to account for the average importance. Additionally one could use Figure 14 to determine proper a proper energy cut-off, maximum weight allowed, or when to split particles. This would be worthwhile to examine in future work.


Another possible option to handle the outliers and poor convergence is to use the integral value of the importance, i.e. average importance multiplied by the counts in each bin. These results are shown in Figure 15.



(a) Once-scattered normalized integral of importance for a 100 keV source.



(b) Once-scattered normalized integral of importance for a 1 MeV source.



(c) Once-scattered normalized integral of importance for a 14.056 MeV source.


Figure 15. Once-scattered normalized integral of importance for a 100 keV, 1 MeV, and 14.056 MeV source. Results match closely to their respective PDF plots. The peak values are still below actual source energy. Energy bins with counts greater than 500 are plotted.

As seen, the integral of importance plots more accurately reflects the mono-energetic sources than the average importance does. The product of average importance and counts suppresses the outliers in bins with lower total counts. It also displays that high average importance, but low counts does not contribute to a source spectrum compared to low importance and high counts. Additionally it shows that if the energy cut-off was changed to closer to the source energy, it would have little effect on the source spectrum as those bins are less important. Ultimately, though Figure 15 does not reveal new information towards the goal of this work, but it could help inform future development of SAHARA and other approaches to this problem.

4.4 Continuous Energy Sources

Two continuous, once-scattered energy sources were chosen to test the initial performance of SAHARA: a uniform 100 keV to 1 MeV source, and the U_{235} thermal Watt fission spectrum (Equation 37). The SAHARA settings for these runs are found in Table 7. Figure 16 shows the normalized counts results.

$$P(E) = 0.4865 \sinh \sqrt{2E} e^{-E} \quad [MeV]^{-1} \quad (37)$$



(a) Once-scattered PDF for a 100 keV to 1 MeV uniform source. PDF skews to the right due to increasing bin widths. (b) Once-scattered PDF for a Watt fission source. PDF depreciated due to bin width at lower energies.

Figure 16. Once-scattered 100 keV to 1 MeV and Watt fission sources PDFs. Energy bins with counts greater than 100 are plotted.

4.4.1 Normalized Counts vs Energy

Examining the 100 keV to 1 MeV uniform source spectrum of Figure 16a, it would be expected that the resulting plot should be a flat line between 100 keV and 1 MeV. However, the results display a exponential-like increase to approximately 0.8 MeV.

One aspect not captured when examining strictly normalized counts is the non-uniform binned nature of the data. The equation for a uniformly spaced PDF of constant value per interval is shown in Equation 38 where b and a are the bounds of the distribution, in this case 100 and 1000 keV.

$$P(E) = \frac{1}{b-a} [keV]^{-1} \quad (38)$$

The units of the uniform PDF is per keV, this means if a comparison is to be made, the PDF must be multiplied by the bin width. This result is shown in Figure 17. The results match closely, and passes a K-S Test with a p value of 0.0613. A significance level, or α , is used for these tests.




Figure 17. Once-scattered 100 keV to 1 MeV uniform source PDF with the non-uniformly spaced PDF overlaid. Energy bins with counts greater than 100 are plotted. The bin width changes the shape of the expected uniform PDF. This passes a 2-D K-S Test with a p-value of 0.8636.

Looking now at the Watt spectrum normalized counts, many of the same trends can be observed. There is a decrease in counts, most notably at lower energies. This can still be compared to the Watt spectrum, correcting for bin width, as shown in Figure 18. The same trend of seeing the spectrum shifted towards lower energies, like the mono-energetic sources, is observed. This passes the K-S Test with a p-value of 0.8980 indicating that the two observations are likely from the same underlying distribution.




Figure 18. Watt fission source once scatter normalized counts with Watt fission spectrum (37) overlaid. Energy bins with counts greater than 100 are plotted. The bin widths change the expected shape of the Watt fission spectrum PDF. This passes a 2-D K-S Test with a p-value of 0.8980.

4.4.2 Importance vs Energy

Figure 19 shows the average normalized importance and relative error. Once again these plots appear to not be of much use other than to say all bins are important to the source spectrum. Utility is degraded further by the obvious lack of convergence. Even with ten million histories, some bins have a relative error approaching 100%. For future work, SAHARA needs to be parallelized to allow quicker runs of higher histories for better convergence, and an investigation needs to be conducted exploring how to handle the importance while using a source event estimator.




Figure 19. Once-scattered average importance for 100 keV to 1 MeV uniform source, and Watt fission spectrum source. Although not useful for characterizing a source, all pseudo-neutrons reaching the source are important to the source spectrum. Energy bins with counts greater than 100 are plotted.

4.5 Summary

Once-scattered results for both mono-energetic spectra and continuous spectra proved to be informative regarding SAHARA's limits and result interpretations. One limit is that SAHARA under-reports true source energies by about five percent. This result was explored in detail to ensure that the pseudo-neutrons were being tallied in the correct bin, and that where bin edges were located played no role in the results. Unfortunately, a conclusive reason for this behavior was not found. Currently, importance has no clear value towards characterizing the source spectrum, especially for continuous spectra, but does provide insight into where source spectrum is most important given the number of counts. Additionally it could be used in future work to enable particle splitting, and proper energy and weight cut-offs. SAHARA successfully matched both the uniform and Watt fission spectra.

V. Multi-Scattered Scenarios

Multi-scattering in this work is defined as allowing up to five scatters plus direct contributions for both HATS-n and SAHARA computations. The most complex aspect is the handling of direct contributions. Five scatters was chosen after a parametric study examined the impact of varying the number of allowed scatters on the total count tallies. It was found there was no significant increase in the tally values for more than five scatters, but a significant increase in computation time. Similar to the once-scattered scenarios, the approach for multi-scattering starts with mono-energetic sources, and then moves to continuous sources.

5.1 Multi-scattering Scenarios vs Once-scattered Scenarios

To compare multi-scattering to once-scattering, the same inputs are used as the once-scattered as cases shown in table 7, with the exception of the number of scatters allowed which is now five.

5.1.1 Multi-scattering Results and Analysis

When performing multi-scatter calculations in SAHARA, a check is performed when the pseudo-neutrons are sampled to determine if the particle could have traveled directly from source to the satellite. This is done by calculating the ToF of the pseudo-neutron directly to the source from the sampled energy and determining if it falls within the upper and lower limits of the sampled time bin. If so, the pseudo-neutron is tallied separately as a potential direct contribution, but still allowed to continue to scatter for further contributions. The scatter continues, because even though a time-energy bin could have a direct component, scatters could also have contributed to the measurements. As a result, for each source, there are two separate tallies: a

direct contribution tally, and a scattered tally. The mono-energetic source normalized counts are shown in Figure 20.




Figure 20. Multi-scattered, direct and scattered contributions for 100 keV and 1 MeV sources. Energy bins with counts greater than 500 are plotted.

The issues with both the direct, and scatter tallies are apparent in Figure 20. These are mono-energetic sources and there should only be one energy bin with direct contribution tallies in it. Thus, the direct contribution tally does capture the direct contribution energy correctly, but it also tallies energies that are clearly not directly from the source. This result is a product of time being discrete with bins large enough to allow for a scattered particle to retain enough energy to arrive during the same time bin as a non-scattered particle. The other obvious issue is seen in the scatter tallies.

These results show a much higher peak at an energy above the true source energy. This peak happens because even though a sampled pseudo-neutron could have been direct, it was still allowed to scatter because that also could have been a valid path. Figure 21 contains the 1 MeV source energy PDF at the satellite. Approximately 40% of the spectrum at the satellite is direct, and by allowing those to scatter there is a large peak above the actual source energy. This is reinforced by examining the average normalized importance of the scatter tally in Figure 22. The higher energies have higher average importance meaning they are less likely to occur. This is because a true direct pseudo-neutron would have to traverse more atmosphere to scatter and reach back to the source compared to true scattered pseudo-neutrons. The lower energy peak is similar to the once-scattered peaks previously seen in Figure 12. It is also notable that the second peak is also lower than the true source energy, as was seen in the once-scattered results. Approaches to handle the issues introduced by direct contribution are discussed in Section 5.2.




Figure 21. Multi-scattered, 1 MeV source energy PDF at the satellite showing that the direct particles comprise 40% of the total particles at the satellite.




Figure 22. Normalized average scatter importance for a 1 MeV source. The higher energies have higher importance indicating that they are less likely to occur.

For mono-energetic sources, the information provided by the direct contribution is of use. It shows that there are no source particles above 1 MeV, so scatters above this energy are non-real. This could then be used as the energy cut-off parameter. The PDF using this technique is shown in Figure 23. The effective result is essentially Figure 20d up to 1 MeV. Although a very appropriate technique for mono-energetic sources, this is less useful for continuous sources as the energy cut-off would only affect the highest energy pseudo-neutrons.





Figure 23. Multi-scattered, scattered contributions PDF from a 1 MeV source and a energy cut-off of 1 MeV which removes the higher energy peak seen in Figure 20d. Energy bins with counts greater than 100 are plotted.

5.1.2 Watt Fission Spectrum Results and Analysis

Using the same approach as for mono-energetic sources, both the normalized counts for direct and scattered tallies are shown in Figure 24. Additionally, the results are compared to a Watt spectrum for each tally. Although less obvious than the mono-energetic scenarios, the results for the Watt fission spectrum do not compare favorably. Upon examining the direct contribution, one would expect the values to match closely in the higher energy regime. The match should not be perfect due to the spectrum at the satellite location having already been transported through the atmosphere, but cross sections are small at higher energies resulting in fewer scatter energy changes. Unfortunately, there is significant deviation from this expectation. Interestingly the lower energies are greater than the true spectrum. This behavior is most likely due to the direct scatter binning issues discussed in the multi-scatter mono-energetic source results. At lower energies, the bin widths get larger in time

which allows creates a higher probability that pseudo-neutrons sampled from that bin could have been direct. Examining the scattered contribution, the probability is shifted towards the spectrum tails. This is most notable above approximately 8 MeV and below 400-500 keV. This, coupled with the difference in the peak, results in failure of a 2-D K-S Test with a p-value of 2.069×10^{-3} indicating that the underlying distributions are not exactly the same.



(a) Direct contribution PDF fails a 2-D K-S Test with a p-value of 1.920×10^{-9} . Plotting energy bins with counts greater than 200. (b) Scatter contribution PDF fails a 2-D K-S Test with a p-value of 2.069×10^{-3} . Plotting energy bins with counts greater than 500.

Figure 24. Multi-scattering, direct and scattered contribution PDFs for a watt fission spectrum source.

5.2 Approaches taken to the Direct Contributions Issue

The previous approach highlighted some of the issues with sampling from a binned spectrum. One quickly rejected alternative was not allowing pseudo-neutrons labeled as a possible direct contribution to scatter. This methodology makes complete sense for a single source energy, but is not necessarily valid for continuous energy spectra. For any given bin it is impossible to know if a sampled pseudo-neutron was a direct or scattered particle.


To determine if a sampled pseudo-neutron could have been direct, SAHARA calculates the ToF of the pseudo-neutron at the sampled energy directly from the source to the satellite. This ToF is then compared to the time bin edges. If it falls within

that range, the pseudo-neutron could have been direct. One solution to the binning issue is to increase the probability that the particle was direct by decreasing the time bin size. This in theory would alleviate the binning issues.

5.2.1 Changing Time-Energy Grid of Input


The standard settings for HATS-n are 100 logarithmically spaced bins per decade in both energy and time. To test the theory that this issue was the result of bin size, the HATS-n resolution was increased to 1000 logarithmically spaced bins per decade in both energy and time for a 1 MeV mono-energetic source. The results of both the direct contribution tally and scatter tally are shown in Figure 25 for the 1 MeV source.

Looking at the direct contribution plot, it appears similar to the previous results in Figure 20 at least in shape. However, the magnitudes are different. In fact, when examining the direct contribution PDF shown in Figures 20c and 25a, the fine input peak at 1 MeV is higher by a factor of 2 indicating more true direct contributions. However, the ratio of true direct contributions to false direct contributions, energies lower than 1 MeV, remains mostly constant. Examining the scattered tally PDF shows that the second peak seen in Figure 20d has largely disappeared. Smaller time bins has increased the fidelity of direct versus scatter categorization. With smaller time bins, if a higher energy is selected and the time bin is smaller, there is less chance for that pseudo-neutron to scatter and reach the source at time equal to zero. This is true for all sampled pseudo-neutrons. Since five scatters are allowed, the Gaussian peak around the source energy is bound to widen, and once again the found peak energy is lower than the actual source, but it has increased to approximately 10%. The σ increased by approximately 300% compared to the once-scattered case. This is further studied in chapter 6. Overall, use of a higher resolution time and energy



(a) Multi-scattering, direction contribution PDF from 1 MeV source. Non-physical direct contributions are continued to be found despite greater input resolution. Energy bins with counts greater than 200 are plotted.

(b) Multi-scattering, scattered contribution PDF from a 1 MeV mono-energetic source. The higher energy peak has been removed. Energy bins with counts greater than 500 are plotted.



(c) Multi-scattering, scattered contributions from a 1 MeV mono-energetic source with a Gaussian best fit applied.

Figure 25. Multi-scattering, direct and scattered contribution PDFs for a 1 MeV mono-energetic source using higher resolution input.

bins has resolves the second peak seen in the scattered tally, but the fact there is still non-direct contributions still being tallied as direct is still cause for concern. Additionally for real-world detectors it may not be feasible to obtain this resolution the time component to this fine of a level. As a result this method was rejected in this work, but may an avenue worth pursuing in future work.

5.2.2 Modifying Time Energy Bins when Selected

Another approach was subdividing the sampled time energy bins into smaller bins themselves. An example of this process is shown in Figure 26 where a selected bin would be subdivided into 16 smaller bins. In theory, if a sampled pseudo-neutron could have been direct, further differentiation can be obtained by subdividing the bin and re-sampling uniformly for a new time and energy. This method has the additional benefit of not requiring an increase in the measurement resolution. The new ToF associated with the newly sampled energy is compared to the selected subdivided time bin lower and upper values to determine if it could have been direct. Like modifying the actual input file from HATS-n, breaking the larger initial bin into smaller bins should increase the categorization fidelity for true direct contributions. One necessary assumption is that the true, finer time-energy spectrum is uniform. Although a more accurate normalized direct contribution is found using this method, the trade-off is an increase in the computational cost equal to the number of subdivisions in order to achieve the same number of total counts. Additionally, because the sampled pseudo-neutrons are still allowed to scatter, the peak above the source energy was still seen as in Figure 20. Because of these reasons this method has limited applications.




Figure 26. Subdividing time energy bins which are sampled uniformly

5.2.3 Directionally Limiting Direct Contributions

The original direct contribution check ignores the sampled pseudo-neutron direction, but this information may be useful. In reality, only pseudo-neutrons with a sampled direction within some solid angle towards the source could truly be direct. However, this approach leads to the same problem as the subdividing bins approach. The initial direction of the pseudo-neutron in SAHARA is sampled uniformly as the HATS-n input only provides a time-energy fluence distribution, and not direction consistent with real-world detection systems. If it is assumed that this information is available, applying a direction check will reduce the likelihood that a sampled pseudo-neutron is incorrectly labeled as direct. Currently USNDS neutron detectors have no directional capability. Is it then worthwhile to develop a solution which needs inputs that are currently not available? One could easily argue no. However, if directional sensing is added in the future, this approach should be explored further.

5.2.4 Modifying Pseudo-Neutron Kill Criteria if Marked Direct

Another option explored to solve the direct contribution problem was changing the kill criteria of pseudo-neutrons marked as direct; specifically, how many scatters should be allowed. The theory is that if a pseudo-neutron could have been direct, then the likelihood that it scattered more than once, or at all is low. The direct

results remain unaffected by this modification. The scattered tally results are shown in Figure 27. Although the results have not resolved the previously noted issues, this result gives insight into what is occurring to pseudo-neutrons categorized as direct. Pseudo-neutrons that could have been direct, but are allowed to scatter, largely only have enough time remaining to scatter once. In effect, changing the kill criteria had no impact on the results because the time kill criteria was already being met.




Figure 27. Scattered contribution PDF for a 1 MeV source allowing direct contributions to scatter once. Non-true higher energy peak is still present. Energy bins with counts greater than 500 are plotted.

5.2.5 Post Processing Results Option

One additional possible solution to handling the direct contribution is to do so as a post processing step. One could assume some apriori knowledge of the source, and then shift some percentage of the results accordingly. Although more data points are needed, Figure 20 indicates that the second peak is approximately 15% higher than the actual source energy. The amount that the SAHARA spectrum should be shifted would then be based off the fraction of scattered vs unscattered neutrons. For instance, for a 1 MeV source at 50 km directly below the satellite, approximately

40% of neutrons are not scattered, so 40% of 1.15 MeV counts at the source should be shifted back to 1 MeV. The approach outlined here is rather simplistic and would require more investigation on how to handle the Gaussian nature of the upscattered direct pseudo-neutrons, but nevertheless could still be explored in future work.

5.3 Influence Of Tally Energy Bin Sizes

While exploring the direct contribution issue, an interesting result regarding the energy bin size was discovered. Originally the energy bin tallies in SAHARA were set to match the ones used in the input: 100 logarithmically spaced bins per decade. This was chosen for no specific reason other than convenience. When a multi-scattering Watt fission source, was run the direct contribution tally seemed to display two nested distributions as shown in Figure 28.




Figure 28. Direct contribution PDF for a Watt fission source using 100 logarithmically spaced bins per decade displaying a bi-modal distribution. Plotting energy bins with counts greater than 100.

A parametric study investigated the effects of sampling with 10, 20, 30, 50, 60, and 120 logarithmically spaced bins per decade. Only in the 50 bins per decade results did this nested distribution behavior fail to appear. This led to a re-investigation of

the input time-energy spectrum shown in Figure 29. Unlike in the surface plot in Figure 3, it is possible that the input sampling causes more than one distribution to be present. Specifically there is one main peak existing along the diagonal of the time-energy fluence matrix, and smaller magnitude, off-diagonal distributions. For a given energy, there typically exists only 2 to 5 time bins with non-zero fluences. As a result of this binning in the input, changing the output binning can cause the appearance of nested distributions. Just to reinforce, this is only the case for the direct contribution tally, as the scattered tally is smoothed due to scatter “blurring” the spectrum. Even though the number of bins per decade was reduced from 100 to 50, it does provide better statistics as there are more tallies per bin. This is helpful for this unparallelized code.




Figure 29. Normalized time-energy fluence input produced by HATS-n for a Watt fission source displaying multiple distributions.

5.4 Summary

When moving to multi-scattering scenarios, the most important consideration was how the direct contributions are handled. Without this, false peaks can appear as in Figure 20, and continuous source spectra fail to match. A direct contribution check was developed to determine if a sampled pseudo-neutron could have been direct. This was done by checking if the ToF for the sampled pseudo-neutron falls within the time ranges of the time bin from which it was sampled. The challenge is that even if the pseudo-neutron could have been direct, it cannot be ruled out that it could have scattered as well. The other issue is the bin size. If the bins are too large, then true non-direct sampling could be categorized as being direct.

To address these issues there were several different approaches explored. The first approach involved obtaining higher resolution input time-energy spectra. In this case the scattered tally matched closely to the once-scattered results. However, there were still sampled pseudo-neutrons tallied as direct that were not valid and high resolution input spectra may not be available. Nevertheless this should still be explored in future work. Two other options explored were subdividing the selected time-energy bin and checking the sampled direction. Both options reduced the likelihood a sampled pseudo-neutron would be selected as direct according to a uniform probability. Thus, better categorization of direct contributions is achieved, but at an increase in computational cost. The latter method would be promising if directional information is available. Finally, post-processing scattered tallies to shift the spectrum may be a valid approach, but was not further developed in this work.

VI. Five-Scattered Scenarios


With the problems handling direct contributions still persisting, five-scattered scenarios that did not include direct contributions were run using HATS-n. For terminology sake, "multi-scattering" refers to five scatters allowed plus direct contribution, and "five-scattered" refers to five scatters allowed and no direct contributions. A further comparison is made with the once-scattered results. Additionally, the atmospheric composition will be changed, as well as the source's location within the atmosphere. Finally, 1-D Little Boy and Fatman input spectra will be used to determine if SAHARA can match to them successfully.

6.1 Five-Scatter vs Once-Scatter


The inputs to SAHARA for the five-scatter scenarios are consistent with the once-scattered cases (see Table 7), with the exception of the number of scatters allowed.

6.1.1 Mono-energetic Sources


Mono-energetic source of 100 keV and 1 MeV sources were used for comparison to the once-scattered scenarios. Their normalized counts and Gaussian fits are shown in Figure 30.




(a) Five-scattered PDF for a 100 keV source.



(b) Five-scattered PDF for a 100 keV source with a Gaussian best fit applied.



(c) Five-scattered PDF for a 1 MeV source.



(d) Five-scattered PDF for a 1 MeV source with a Gaussian best fit applied.

Figure 30. Five-scattered PDFs for mono-energetic sources. Peaks are significantly wider compared to the once-scattered cases. Energy bins with counts greater than 500 are plotted.

Comparing the once-scattered and five-scatter scenarios results, one feature remains constant: the peak energy found by SAHARA continues to be less than the actual source. As expected, allowing more scatters increased the higher energy contributions. Because this is the adjoint problem, pseudo-neutrons are up-scattered in energy. The more scattering allowed, the more high energy contributions are expected. As a result a higher energy shoulder appears to the right of the peak that was not present in the once-scattered results. Another large difference between the once-scattered and five-scattered results is the significant broadening of the peak. A full comparison between once-scattered (OS) and five-scattered (5S) Gaussian parameters

are shown in Table 9. The standard deviation around the peak energy increased by approximately 150%. Although the broadening was expected due to allowing more scattering to occur, this is a significant increase. An interesting case study for future work would be to see if the trend between the number of allowed scatters and the peak broadening continues. A visualization of how much the peak has broadened is shown in Figure 31.

Table 9. Gaussian best fit parameters for both once-scattered (OS) and five-scattered (5S).

Source	Peak Energy [keV]	Relative Error %	Std Dev % of Peak	FWHM % of peak
100 keV OS	96.353	3.649	8.292	19.526
100 keV 5S	93.813	6.184	15.033	35.399
Difference	-2.54 keV	169.5%	181.3%	181.3%
1 MeV OS	941.146	5.885	7.385	17.391
1 MeV 5S	922.681	7.373	9.953	23.467
Difference	-18.465 keV	125.3%	134.8%	134.8%




Figure 31. Once-scattered and five-scattered PDF for a 1 MeV source. Allowing more scatters widens the peak.

6.1.2 Watt Fission Source

For the mono-energetic sources, increasing the number of allowed scatters resulted in a further lowering of peak energies, and a broadening of the peak. These effects will undoubtedly affect a continuous spectrum as well. Watt fission spectrum results are shown in Figure 32 and with the once-scattered in Figure 33.




Figure 32. Five-scattered PDF from a Watt fission source with a Watt fission spectrum corrected for bin width overlaid. This passes a 2-D K-S Test with a p-value of 0.8298. Energy bins with counts greater than 500 are plotted.




Figure 33. Once-scattered and five-scattered PDFs from a Watt fission source. Allowing more scatters results in a shift towards lower energies and a depreciated peak.


The five-scatter results retain the overall Watt shape. Additionally, performing a 2-D K-S Test still results in failing to reject the null hypothesis with a p-value of 0.8298, indicating that they cannot be distinguished as different distributions. Between the once-scattered and five-scattered results, the five-scattered is shifted towards lower energies, which was also seen in the mono-energetic sources. The exception is at the tails where the number of counts is actually higher. This is most likely due to the higher energy shoulder seen in the monoenergetic sources.

6.2 Three Isotope Atmosphere Results


To examine the influence of other isotopes in the atmosphere, SAHARA was run using an atmosphere consisting of N_{14} , O_{16} , and Ar_{40} ; hereafter referred to as the 3-isotope atmosphere. For these scenarios, the inputs are shown in Table 10. The results from these scenarios are shown in Figure 34.

Table 10. SAHARA settings for five scattering, three isotope atmosphere scenarios.


Setting	Value
Source Position (ECI)	6421, 0, 0
Source Altitude [km]	50
Satellite Position (ECI)	(42164, 0, 0)
Satellite Altitude [km]	35,793
Histories	5×10^6
Initial Weight	5×10^{-6}
Weight Cutoff	1×10^{23}
Energy Cutoff [MeV]	20
Scatters Allowed	5
Energy Bins per Decade	50
Atmosphere Constituents	N_{14}, O_{16}, Ar_{40}
Sources	100 keV MeV 1 MeV MeV Watt Fission Spectrum




(a) Five-scattered PDF for a 100 keV source.




(b) Gaussian best fit applied to the 100 keV PDF.



(c) Five-scattered PDF for a 1 MeV source.



(d) Gaussian best fit applied to the 1 MeV PDF.



(e) Five-scattered PDF for a Watt fission source to compare to a Watt fission spectrum corrected for bin width. This passes a 2-D K-S Test with a p-value of 0.7845. Energy bins with counts greater than 100 are plotted.

Figure 34. Five-scattered PDFs for sources in a 3 isotope atmosphere. Energy bins with counts greater than 500 are plotted for mono-energetic sources.

The mono-energetic source results appear much the same as what was seen in the N_{14} atmosphere, five-scattered results. The peak energies once again fall below the

true source energy although slightly closer. The shoulder to the right of the peak is also still evident, and higher in magnitude. Examining Figure 35 reveals that the whole spectrum is shifted slightly up in energy compared to the N_{14} atmosphere, a result that is somewhat surprising. The other isotopes added were O_{16} and Ar_{40} which are heavier than N_{14} meaning possible up-scattered energies are lower. A further study of this effect is recommended for future work.





Figure 35. Five-scatter PDFs for a 1 MeV source for both a N_{14} and three isotope atmosphere. Moving to a three isotope shifts the spectrum towards higher energies.

6.2.1 Varying Source Altitude


SAHARA's performance is evaluated for various altitudes at the same latitude and longitude, in this case the satellite position is directly overhead the source position. The inputs for these cases are shown in Table 11. Each source is simulated at each altitude. The results are shown in Figures 36 and 38.

Table 11. SAHARA settings for five-scattered, varied altitudes scenarios.


Setting	Value
Source Position (ECI)	6401, 0, 0
	6421, 0, 0
	6451, 0, 0
Source Altitude [km]	30
	50
	80
Satellite Position (ECI)	42164, 0, 0
Satellite Altitude [km]	35,793
Histories	5×10^6
Initial Weight	5×10^{-6}
Weight Cutoff	1×10^{23}
Energy Cutoff [MeV]	20
Scatters Allowed	5
Energy Bins per Decade	50
Atmosphere Constituents	N_{14}, O_{16}, Ar_{40}
Sources	1 MeV Watt Fission Spectrum




(a) Five-scattered PDF for a 1 MeV source at 30 km. Peak is broader at lower altitudes.




(b) Gaussian best fit applied to the 1 MeV PDF at 30 km.




(c) Five-scattered PDF for a 1 MeV source at 50 km.



(d) Gaussian best fit applied to the 1 MeV PDF at 50 km.



(e) Five-scattered PDF for a 1 MeV source at 80 km. The right shoulder is more prominent than at 50 km.



(f) Gaussian best fit applied to the 1 MeV PDF at 80 km.

Figure 36. Five-scattered PDFs 1 MeV source in a 3 isotope atmosphere, varying source altitude between 30, 50, and 80 km. Energy bins with counts greater than 500 are plotted.

The overlay for the 30, 50, and 80 km source altitudes is shown in Figure 37. Looking at the plot, it is clear that as the source moves deeper into the atmosphere,

the wider the peak becomes. However, the more allowed scatters, the greater the chance for a particles path to diverge from the true path. Combining this effect with a decrease in the number of counts due to the various kill criteria explains why the uncertainty increases as the altitude decreases and the peak widens. This is further backed up by examining the Gaussian fitting parameters generated by OriginPro shown in Table 12. Interestingly though, the error in SAHARA peak location approaches the true mono-energetic source value as altitude decreases. Throughout all altitudes the right shoulder is still present, and becomes more prominent at lower altitudes.




Figure 37. PDFs for a 1 MeV source at various altitudes. The lower in the atmosphere the source, the larger the standard deviation of the peak. The higher in the atmosphere, the larger the error in the Gaussian peak.

Source Altitude	Peak Energy [keV]	Relative Error %	Std Dev % of Peak	FWHM % of peak
30 km	976.588	2.341	21.726	51.161
50 km	942.334	5.767	10.395	24.478
80 km	911.768	8.823	8.241	19.405

Table 12. Gaussian best fit parameters for a 1 MeV source at altitudes of 30, 50, and 80 km.




Figure 38. Five-scattered PDFs for a Watt fission source in a 3 isotope atmosphere, varying source altitude between 30, 50, and 80 km.

Figure 38 shows the individual SAHARA responses due to a Watt fission spectrum input at 30, 50, and 80 km source altitudes. Figure 39 shows all three overlaid together. The 30 km source altitude is the most different from the actual Watt spectrum. However, all three altitudes still pass the K-S Test indicating that cannot be differentiated from the underlying Watt spectrum. The 50 km and 80 km source locations are extremely similar to one another, not a unexpected result due to the low density of the atmosphere above 40 km. The 30 km has larger discrepancies at lower energies due the increased atmospheric density and resulting higher uncertainties. The influence of the atmospheric cross section resonances are also more pronounced at this altitude. Additionally, the overall efficiency of pseudo-neutrons reaching the source location is much lower for the 30 km source location versus the 50 and 80 km.




Figure 39. PDFs for a Watt fission spectrum at various altitudes. The effect of the atmosphere's cross sections are more evident at lower altitudes.

6.2.2 Varying Source Location

The source's latitude and longitudinal position were also varied to see the effect. Changing the location from directly below the satellite, to some other position increases the path length through the atmosphere, decreasing the efficiency, and in-


creasing atmospheric effects. Inputs to SAHARA for these cases are shown in Table 13. The results are shown in Figures 40 and 41.

Table 13. SAHARA settings for five-scattered, varied latitude and longitude scenarios. Both the 1 MeV and Watt fission source were run at all locations.

Setting	Value
Latitude, Longitude	45N, 0E
Source Position (ECI)	4561.55, 0, 4561.55
Source Altitude	80 km
Latitude, Longitude	45N, 45E
Source Position (ECI)	3210.5, 3210.5, 4540.33
Source Altitude	50 km
Satellite Position (ECI)	42164, 0, 0
Satellite Altitude [km]	35,793
Histories	5×10^6
Initial Weight	5×10^{-6}
Weight Cutoff	1×10^{23}
Energy Cutoff [MeV]	20
Scatters Allowed	5
Energy Bins per Decade	50
Atmosphere Constituents	N_{14}
Sources	1 MeV Watt Fission Spectrum


As expected, the overall counts decreased which increases the uncertainty in any given bin. Where peaks in the previous directly overhead scenarios may have around 30,000 tallies these cases received under 1,000. Thus, without additional computation times making meaningful, quantitative conclusions is difficult. Looking at the 1 MeV

source, the broadening is still present, along with the shoulder to the right of the peak. Due to the efficiency decrease, the maximum relative error shown in Figures 40a and 41a increased from 4% to 10% and 16% and required lowering the plotted counts limit from 500 to 100 and 25 particles respectively. Additionally, the Watt shape shown in Figure 41 is still present and clearly has more variance and corresponding relative error due to the increased atmospheric influence than the directly overhead cases.




(a) Five-scattered PDF for a 1 MeV source. Plotting energy bins with counts greater than 100.

(b) Gaussian best fit applied to the 1 MeV PDF.




(c) Five-scattered PDF for a Watt fission source at 50 km to compare to a Watt fission spectrum accounting for bin width. This passes a 2-D K-S Test with a p-value of 0.9991. Plotting energy bins with counts greater than 10.

Figure 40. Five-scattered PDFs for a 1 MeV and Watt fission source located at 80 km altitude, and 45N, 0E in a N_{14} atmosphere.



(a) Five-scattered PDF for a 1 MeV source. Plotting energy bins with counts greater than 25.


(b) Gaussian best fit applied to the 1 MeV PDF.



(c) Five-scattered PDF for a Watt fission source at 50 km to compare to a Watt fission spectrum accounting for bin width. This passes a 2-D K-S Test with a p-value of 0.96923. Plotting energy bins with counts greater than 10.

Figure 41. Five-scattered PDFs for a 1 MeV and Watt fission source located at 50 km altitude, and 45N, 45E in a N_{14} atmosphere.

Figure 42 overlays the previous 50 km directly overhead scenario and 50 km 45N, 45E scenarios. Due to the increased uncertainty, drawing meaningful shape characteristics for the Watt source is impossible. In contrast, the 1 MeV source result appears to broaden slightly which is expected due to the increase in atmospheric scatter. Due to the relative increased uncertainty this conclusion is hesitantly made.



(a) Five-scattered PDFs for a 1 MeV source. The peak broadens with an increased path through the atmosphere. The uncertainty in this result is high. (b) Five-scattered PDFs for a Watt fission source and a Watt fission spectrum accounting for bin width.

Figure 42. Five-scattered, 1 MeV and Watt fission source with the source located at 50 km altitude, 0N, 0E and 45N, 45E.

6.2.3 Little Boy and Fat Man Sources

The Little Boy and Fat Man energy spectra [2] are used as the input for HATS-n. These spectra are shown in Figure 43. HATS-n sampled from the normalized version of these spectra. This case is the true test of SAHARA's ability to reconstruct realistic source spectra. As such, a full atmosphere representation was modeled. The other inputs are shown in Table 14. The energy bin structure of the resulting tally was matched to the real spectrum to enable an accurate comparison.




Figure 43. Neutron spectra for Fat Man and Little Boy. Plot generated from data provided by [2].

Setting	Value
Source Position (ECI)	6421, 0, 0
Source Altitude [km]	50
Satellite Position (ECI)	42,164, 0, 0
Satellite Altitude [km]	35,793
Histories	1×10^7
Initial Weight	1×10^{-7}
Weight Cutoff	1×10^{23}
Energy Cutoff [MeV]	20
Scatters Allowed	4
Energy Bins Structure	Same as [2]
Atmosphere Constituents	N_{14} , N_{15} , O_{16} , O_{17} , O_{18} , Ar_{40}
Sources	Little Boy Fat Man

Table 14. SAHARA settings for Little Boy and Fat Man scenarios.

The PDF from the Little Boy source is shown in Figure 44. This passes a 2-D K-S Test with a p-value of 0.7583. The same trend of the energy spectrum being shifted towards lower energies can still be seen which is most notable at the peak around 25 keV. The PDF from a Fat Man source is shown in Figure 45. Like Little Boy, this passes a 2-D K-S Test with a p-value of 0.8063.




Figure 44. PDF from the Little Boy source with the true spectrum overlaid [2]. This passes a 2-D K-S Test with a p-value of 0.7583. Energy bins with counts greater than 100 are plotted.




Figure 45. PDF from a Fat Man source with the true spectrum overlaid [2]. This passes a 2-D K-S Test with a p-value of 0.8063. Energy bins with counts greater than 100 are plotted.

Where both the Fat Man and Little Boy sources differ from the true sources are at lower energies. This behavior is connected to low energy particles in the true spectrum being absorbed by the atmosphere, and thus not appearing in the satellite's measurements. Note that higher energy scatter can contribute to the detected lower energies, however some of the lower energy counts remain missing. As a result, SAHARA is unable to recreate the missing information. In addition, the total number of spectrum particles is underestimated resulting in a skewed normalization. As a result, SAHARA overestimates the spectrum below approximately 80 to 100 eV for both Fat Man and Little Boy. The significant $P(E)$ decrease at energies at and below 10 eV for both spectra is an atmospheric cross section effect, where the weight of the source event pseudo-neutron increases past the weight cutoff resulting in the particles being killed. Additionally, because only a 1000 second time-energy distribution is used, this means there is no possible way for SAHARA pseudo-neutrons below approximately 5 eV to even be within the reconstructed spectrum. If the portion below 10 eV is removed from the Fat Man results, the match becomes better as shown in Figure 46.

Another option may be to run SAHARA with a higher weight cutoff.




Figure 46. PDF from a Fat Man source above 10 eV with the true spectrum overlaid [2]. This passes a 2-D K-S Test with a p-value of 0.9690. Energy bins with counts greater than 100 are plotted.

6.3 Summary

Moving to multi-scattered scenarios resulted in a widening of the peak for mono-energetic sources. Additionally, for mono-energetic sources a higher energy shoulder is present to the right of the peak which was not seen in the once-scattered scenarios. When moving to a Watt fission source spectrum, the spectrum is noticeably shifted towards lower energies compared to the once-scattered results, but still was found to match a Watt spectrum.

Moving to a three isotope atmosphere, the Gaussian peak widened slightly when compared to the N_{14} atmosphere results. The higher energy shoulder also increased in magnitude by a small margin. For a Watt fission source, a match was still achieved.

The altitude and location of the source also effected the SAHARA results. In general, the more atmosphere there is to travel through, the wider the peak of a mono-energetic source. As expected, SAHARA's efficiency decreases significantly

due to the increase of the path through the atmosphere. Running SAHARA with higher weight cutoff may mitigate some of this effect.

SAHARA was then tested against real world sources within a six isotope atmosphere, and demonstrated reconstruction of sources energy spectra well. Both Fat Man and Little Boy sources generated by SAHARA passed a 2-D K-S Test.

VII. Conclusions

The primary goal of this work was to characterize the energy spectrum of a point neutron source from a time-energy neutron distribution at the satellite location by using the adjoint neutron transport method via Monte Carlo simulation. In this work, the Python code SAHARA was developed to both accomplish this goal and study various effects of the air-to-space adjoint problem.

7.1 Summary

The methodology behind SAHARA is broken down into the following steps:

1. *Start Pseudo-Neutron,*
2. *Move pseudo-neutron,*
3. *Source Event Pseudo-Neutron,*
4. *Interact Pseudo-Neutron,*
5. *Kill Pseudo-Neutron,* and
6. *Tally Pseudo-Neutron.*

In order to improve the efficiency of this point detector to point source problem, a new version of the source event was developed that incorporated both the scattering angle in the lab frame and the time left to reach the source. Although this method successfully root-solves to find the required scattered locations, it led to pseudo-neutrons with extremely high weights. This results in importances that never converge, but are useful for future particle splitting and weight reduction approaches. Using tallies of pseudo-neutrons counts within energy bins provided the basis for various proof-of-concept studies.

To provide baseline tests of SAHARA, mono-energetic and continuous energy once-scattered input spectra were used. For mono-energetic sources, SAHARA reports source energies approximately five percent lower than the true source energy, and Gaussian peaks with a standard deviation of approximately seven percent of the peak value. The consistent under-reporting of energy was investigated. For continuous energy sources, SAHARA performed well and 2-D K-S tests showed it was able to match to a Watt fission source. However, the energy spectrum is noticeably shifted towards lower energies.

The next case studies investigated moving to using five-scattered plus direct contribution input spectra. The direct contribution method was developed to examine if the sampled pseudo-neutron time-energy could have come directly from the source. This process was completed by checking the time of flight back to the pseudo-neutron source given the particle sampled energy against the sampled time bin boundaries. This study revealed the challenges associated with binned input spectra. When working with mono-energetic sources, SAHARA was finding direct contributions where none should be present. However, interpreting the results with the understanding that the source is mono-energetic revealed that SAHARA correctly assigns the mono-energetic energy bin a high probability of occurrence. Several approaches to solving this issue were attempted, with varying success. It is left to future work to find if direct contributions can effectively be filtered.

Next, a variety of five scattered, no direct contribution scenarios were presented for comparison to the once-scattered results and for examining the influence of source position and atmospheric composition. Finally, SAHARA was used to reconstruct the energy spectra of Little Boy and Fat Man sources. When moving from once scattered to five scattered mono-energetic sources, the peak location error increased from approximately five to seven percent toward lower energies. The standard deviation

of the Gaussian peak also increased as expected by approximately 150% from the once-scattered cases. The Watt fission spectrum also shifted towards lower energies, however, a 2-D K-S test showed that

SAHARA was still able to match the Watt fission spectrum.

When moving from a ^{14}N atmosphere to a ^{14}N , ^{16}O , ^{40}Ar atmosphere, a slight broadening of the Gaussian peak was seen for mono-energetic sources. More source scenarios are recommended to fully examine this influence. Again, a 2-D K-S test demonstrated that SAHARA was still able to match the Watt fission spectrum.

Changing the altitude of the source was examined next. In general, the lower the source altitude, the wider the Gaussian peak for a mono-energetic source. However, the corresponding 2-D K-S test showed that SAHARA was still able to match a Watt fission spectrum. Then, the source's location within the atmosphere was varied. For these problems the decreased efficiency led to large uncertainties. For sources at the same altitude, the Gaussian peak standard deviation increased slightly, but this result may be an artifact of poor statistics .

Lastly, the real world sources, Little Boy and Fat Man, were analyzed using SAHARA. Although like the Watt spectrum they are noticeably shifted towards lower energies the spectral features are still recognizable, a 2-D K-S Tests proved they matched to the provided spectra.

Although many questions still persists, SAHARA has successfully met the research goals posed for this work.

7.2 Future Work

As with many research problems, the attempt to answer one question has led to many others. This portion will highlight some of the most pertinent problems to be explored in future work.

The once-scattered scenarios were used as an initial benchmark to test SAHARA. For mono-energetic sources the peaks were expected to be Gaussian in nature, but the peaks center should still be extremely close. However, the peak centers were consistently five percent lower than the actual values. One method to correct this was attempted, but failed to address the issue. Although a match was still obtained for the Watt fission spectrum, it is also shifted towards lower energies. This effect also worsens as the number of scatters increases. As such it is important to further investigate this behavior.

The lack of direct relevance to the source spectrum for the importance was also observed. This was due to the characteristics of the source event estimator. Although using just counts is a valid method of examining source characteristics, it provides no information on source intensity. Including this estimate may invoke a redefinition of the cross sections to conform to other adjoint solution methodology such as described in Hoogenboom's dissertation [6]. Another possible method may be to use one similar to the DXTRAN spheres used within MCNP, and split particles with high weight that reach a given volume using nested spheres [31].

Identifying and filtering direct contribution pseudo-neutrons is another key issue for future work. There were several attempts presented in this work, but adding an angular dimension is promising assuming a segmented neutron detector is ever fielded. The issue in identifying direct contributions is inherent in using binned data. Ultimately adding a directional variable to the input makes the sampled spectrum a time-energy-direction fluence. All direct contributions should be in the same direction defined by the satellite and source positions and thus, the same directional bin. Another promising solution is to assume and incorporate some apriori source information within a post-processing step.

Another improvement to SAHARA would be parallel computing. The run times

for the scenarios presented in this work range from 12 hours for the once-scattered scenarios to 4 days for the Fat Man and Little Boy scenarios. The computational run time will be decreased if parallelization is implemented. Also, decreased run times will enable or increase the fidelity of additional case studies examining the impacts of source location, and other low efficiency scenarios.

Lastly, the other problem scope assumptions should be lessened. Some of the more important assumptions were isotropic, and elastic only scattering and ignoring detector characteristics. Although these were important to make to scope the project for thesis work, their inclusion would have an impact on the results. It is well known that high energy neutrons do not scatter isotropically, but rather are forward peaked. Neglecting this may have resulted in the high energy shoulder seen in the five-scatter scenarios. Additionally, inelastic scattering adds further possibilities for a pseudo-neutron to make it back to the source. Including detector characteristics would also have a large effect on the solution. This would introduce uncertainties in the time-energy spectrum which would follow into the determined energy spectrum. Further work should relax these assumptions.

Although there are additional possible improvements within SAHARA, this work provides a foundation for continuing work.

Bibliography

1. W. T. Dailey, “Special Features of the Air-to-Space Neutron Transport Problem,” PhD Dissertation, Air Force Institute of Technology, 2017. [Online]. Available: <https://scholar.afit.edu/cgi/viewcontent.cgi?article=1771&context=etd>
2. R. L. Holmes and S. W. White, “Standardized Unclassified Little Boy and Fat Man Outputs,” Defense Threat Reduction Agency, Fort Belvoir, VA, USA, Tech. Rep., 2013. [Online]. Available: <http://www.dtic.mil/dtic/tr/fulltext/u2/a590626.pdf>
3. “U.S. Standard Atmosphere 1976,” NOAA, NASA, USAF, Washington D.C., USA, Tech. Rep., 1976. [Online]. Available: <https://strives-uploads-prod.s3.us-gov-west-1.amazonaws.com/19770009539/19770009539.pdf?AWSAccessKeyId=AKIASEVSKC45ZTTM42XZ&Expires=1599328433&Signature=SJfdOIS4etD3LuqgHfSFLcstRPw%3D>
4. E. Lewis and W. Miller, *Computational Methods of Neutral Particle Transport*, 1st ed. American Nuclear Society, 1984.
5. P. Saracco, S. Dulla, and P. Ravetto, “The adjoint neutron transport equation and the statistical approach for its solution,” *European Physical Journal Plus*, vol. 131, no. 412, 2016.
6. J. E. Hoogenboom, “Adjoint Monte Carlo Methods in Neutron Transport Calculations,” PhD Dissertation, Delft University of Technology, 1977. [Online]. Available: https://d1rkab7tlqy5f1.cloudfront.net/TNW/Afdelingen/RadiationScienceandTechnology/ResearchGroups/RPNM/Publications/PhD_{_}Hoogenboom.pdf

7. I. Lux and L. Koblinger, *Monte Carlo Particle Transport Methods : Neutron and Photon Calculations*, 1st ed. Boca Raton, FL, USA: CRC Press, 1991. [Online]. Available: <https://gnssn.iaea.org/NSNI/SharedDocuments/OPENSharedFiles/MonteCarloParticleTransportMethodsNeutronAndPhotonCalculations.pdf>
8. S. N. Cramer, “Bounded next-event estimation and Monte Carlo forward-adjoint correlated coupling,” *Nuclear Science and Engineering*, vol. 149, no. 3, pp. 247–258, 2005.
9. T. Ueki, J. Hoogenboom, and J. Kloosterman, “Monte Carlo forward-adjoint coupling by next event estimation,” p. 241, 1999. [Online]. Available: <http://www.janleenkloosterman.nl/papers/ueki9901.pdf>
10. A. Jinaphanh, Y. Richet, and N. Martin, “Exploring the Use of a Deterministic Adjoint Flux Calculation in Criticality Monte Carlo Simulations,” in *International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering*, no. 1, 2011, pp. 1–9.
11. D. C. Irving, “The Adjoint Boltzmann Equation and its simulation by Monte Carlo,” Oak Ridge National Laboratory, Oak Ridge, TN, USA, Tech. Rep., 1970.
12. L. L. Carter, “MCNA: A Computer Program to Solve the Adjoint Neutron Transport Equation by Coupled Sampling with the Monte Carlo Method,” Los Alamos Scientific Laboratory, Los Alamos, NM, USA, Tech. Rep. August, 1970.
13. M. Munk and R. N. Slaybaugh, “Nuclear Science and Engineering Review of Hybrid Methods for Deep-Penetration Neutron Transport Review of Hybrid Methods for Deep-Penetration Neutron Transport,” 2019. [Online]. Available: <https://www.tandfonline.com/action/journalInformation?journalCode=unse20>

14. A. Zhu, “Transient Methods for Pin-Resolved Whole-Core Neutron Transport,” Ph.D. dissertation, University of Michigan, 2016.
15. K. Ott and R. Neuhold, *Introductory Nuclear Reactor Dynamics*. American Nuclear Society, 1985.
16. B. R. Betzler, W. R. Martin, B. C. Kiedrowski, and F. B. Brown, “Journal of Computational and Theoretical Transport Calculating α Eigenvalues of One-Dimensional Media with Monte Carlo Calculating α Eigenvalues of One-Dimensional Media with Monte Carlo,” *Journal of Computational and Theoretical Transport*, vol. 43, pp. 38–49, 2014. [Online]. Available: <https://www.tandfonline.com/action/journalInformation?journalCode=ltty21>
17. J. M. Picone, A. E. Hedin, D. P. Drob, and A. C. Aikin, “NRLMSISE-00 EMPIRICAL MODEL OF THE ATMOSPHERE: STATISTICAL COMPARISONS AND SCIENTIFIC ISSUES,” *Journal of Geophysical Research*, vol. 107, no. A12, 2001. [Online]. Available: <https://agupubs.onlinelibrary.wiley.com/doi/epdf/10.1029/2002JA009430>
18. S. L. Tornga and W. L. Wakeford, “The DIORAMA Approach to Modeling Satellite Detection of Nuclear Detonations,” *Journal of Radiation Effects Research and Engineering*, vol. 37, no. 1, pp. 208–212, 2019.
19. A. T. Yue, M. S. Dewey, D. M. Gilliam, G. L. Greene, A. B. Laptev, J. S. Nico, W. M. Snow, and F. E. Wietfeldt, “Improved determination of the neutron lifetime,” *Physical Review Letters*, vol. 111, no. 22, p. 222501, nov 2013. [Online]. Available: <https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.111.222501>

20. M. Hollander, D. Wolfe, and E. Chicken, *Nonparametric Statistical Methods*, 3rd ed. Wiley, 2013.
21. J. Peacock, “Two-dimensional goodness-of-fit testing in astronomy,” *Astronomical Society*, vol. 202, no. 1, pp. 615–627, 1983. [Online]. Available: <http://adsabs.harvard.edu/pdf/1983MNRAS.202..615P>
22. J. Bernard, “Reactor Physics-Part Fourteen Kinematics of Elastic Neutron Scattering,” 2006. [Online]. Available: <https://dspace.mit.edu/bitstream/handle/1721.1/74136/22-05-fall-2006/contents/lecture-notes/lecture14.pdf>
23. M. Ragheb, “NEUTRON COLLISION THEORY,” Tech. Rep., 2006.
24. K. Krane, *Introductory Nuclear Physics*, 2nd ed. Wiley, 1988.
25. R. Pevey, “Nuclear Criticality Safety Engineer Training Module 7,” 2010. [Online]. Available: https://ncsp.llnl.gov/docs/training/ncset/Module7{_}r1.pdf
26. M. B. Chadwick, M. Herman, P. Obložinský, M. E. Dunn, Y. Danon, A. C. Kahler, D. L. Smith, B. Pritychenko, G. Arbanas, R. Arcilla, R. Brewer, D. A. Brown, R. Capote, A. D. Carlson, Y. S. Cho, H. Derrien, K. Guber, G. M. Hale, S. Hoblit, S. Holloway, T. D. Johnson, T. Kawano, B. C. Kiedrowski, H. Kim, S. Kunieda, N. M. Larson, L. Leal, J. P. Lestone, R. C. Little, E. A. McCutchan, R. E. MacFarlane, M. MacInnes, C. M. Mattoon, R. D. McKnight, S. F. Mughabghab, G. P. Nobre, G. Palmiotti, A. Palumbo, M. T. Pigni, V. G. Pronyaev, R. O. Sayer, A. A. Sonzogni, N. C. Summers, P. Talou, I. J. Thompson, A. Trkov, R. L. Vogt, S. C. van der Marck, A. Wallner, M. C. White, D. Wiarda, and P. G. Young, “ENDF/B-VII.1 nuclear data for science and technology: Cross

sections, covariances, fission product yields and decay data,” *Nuclear Data Sheets*, vol. 112, no. 12, pp. 2887–2996, dec 2011.

27. A. Trkov, M. Herman, D. A. Brown, N. Holden, and G. Hedstrom, “ENDF-6 Formats Manual With contributions from,” Brookhaven National Laboratory, Upton, NY, USA, Tech. Rep., 2018. [Online]. Available: www.nndc.bnl.gov
28. R. Forrest, R. Capote, N. Otsuka, T. Kawano, A. Koning, S. Kunieda, J.-C. Sublet, and Y. Watanabe, “FENDL-3 Library–Summary documentation,” International Atomic Energy Commission, Vienna, Austria, Tech. Rep., 2012. [Online]. Available: <https://www-nds.iaea.org/fendl/>
29. P. Romano, N. Horelik, B. Herman, A. Nelson, B. Forget, and K. Smith, “OpenMC: A State of the Art Monte Carlo Code for Research and Development,” 2015.
30. “Origin(Pro),” Northhampton, MA, USA, 2020.
31. T. E. Booth, K. C. Kelley, and S. S. McCreedy, “Monte Carlo variance reduction using nested dxtran spheres,” *Nuclear Technology*, vol. 168, no. 3, pp. 765–767, 2009.

REPORT DOCUMENTATION PAGE					Form Approved OMB No. 0704-0188	
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.						
1. REPORT DATE (DD-MM-YYYY)		2. REPORT TYPE		3. DATES COVERED (From — To)		
21-03-2019		Master's Thesis		June 2019 — Mar 2021		
4. TITLE AND SUBTITLE Space to Air High-Altitude Region Adjoint Neutron Transport				5a. CONTRACT NUMBER		
				5b. GRANT NUMBER		
				5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S) LaMere, Zachary W., 1Lt, USAF				5d. PROJECT NUMBER		
				5e. TASK NUMBER		
				5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Air Force Institute of Technology Graduate School of Engineering and Management (AFIT/EN) 2950 Hobson Way WPAFB OH 45433-7765				8. PERFORMING ORGANIZATION REPORT NUMBER AFIT-ENP-MS-21-M-125		
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Technical Applications Center (AFTAC) ATTN: Dr. Jeremy Kephart 10989 South Patrick Drive Patrick SFB, FL 32935				10. SPONSOR/MONITOR'S ACRONYM(S)		
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION / AVAILABILITY STATEMENT DISTRIBUTION STATEMENT A: APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.						
13. SUPPLEMENTARY NOTES This material is declared work of the U.S. Government and is not subject to copyright protection in the United States.						
14. ABSTRACT The goal of this work was to use the adjoint transport equation to characterize the energy spectrum of a neutron source located in the upper atmosphere using the time-energy fluence at the satellite. This adjoint approach directly solves for the source spectrum thereby potentially enabling device reconstruction. The adjoint method also could improve computational efficiency over the forward method. The adjoint transport equation was solved via Monte Carlo methods in a new program called SAHARA written in Python 3.7. A new adjoint source event estimator was developed to improve the computational efficiency. This work explores SAHARA's development, and its performance for mono-energetic and continuous energy sources. In general, the identified source spectra were shifted towards lower energies approximately five percent, but were able to capture the source spectrum shapes. Additionally, continuous energy sources still passed a 2-D Kolmogorov-Smirnov (K-S) test. Lastly, SAHARA was applied to the real-world neutron energy spectra of Fat Man and Little Boy. Although these spectra were also noticeably shifted towards lower energies, the spectral features are still recognizable and the spectra passed a 2-D K-S test. SAHARA provides a new tool for estimating the source spectrum from space-based neutron measurements.						
15. SUBJECT TERMS Adjoint Transport, Monte Carlo, Source Characterization						
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON	
a. REPORT	b. ABSTRACT	c. THIS PAGE			Dr. Darren Holland, AFIT/ENP	
U	U	U	UU	124	19b. TELEPHONE NUMBER (include area code) (937) 255-6565 x4697; darren.holland@afit.edu	