

**AN IMPROVED MODEL OF NUCLEAR DEBRIS ANALYSIS SYSTEMS IN
ATMOSPHERIC COLLECTION AIRCRAFT**

A Dissertation
Presented to
The Academic Faculty

By

Jacob W. Inman

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in Nuclear Engineering in the
School of Mechanical Engineering

Georgia Institute of Technology

May 2021

Copyright © Jacob W. Inman 2021

**AN IMPROVED MODEL OF NUCLEAR DEBRIS ANALYSIS SYSTEMS IN
ATMOSPHERIC COLLECTION AIRCRAFT**

Approved by:

Dr. Nolan E. Hertel, Advisor
Nuclear and Radiological Engineering
Georgia Institute of Technology

Dr. Vincent J. Jodoin
Detonation Forensics and Response
Oak Ridge National Laboratory

Dr. Steven R. Biegalski
Nuclear and Radiological Engineering
Georgia Institute of Technology

Dr. Michael P. Shannon
Nuclear and Radiological Engineering
Georgia Institute of Technology

Dr. Sotira Z. Yiacoumi
School of Civil and Environmental
Engineering
Georgia Institute of Technology

Date Approved: April 29, 2021

TABLE OF CONTENTS

List of Tables	iii
List of Figures	iv
Acronyms, Abbreviations, and Initialisms	vii
Summary	ix
Chapter 1: Introduction	1
Chapter 2: Background and Theory	4
2.1 DELFIC	4
2.1.1 Nuclear Cloud Rise in DELFIC	5
2.1.1.1 Entrainment of Ambient Material in DELFIC	5
2.1.1.2 DELFIC Elevation Sensitivity	13
2.1.2 Output Processing in DELFIC	17
2.1.2.1 Fission Product Mass Chains and Fractionation	19
2.2 HYSPLIT	23
2.2.1 Spatial Resolution in HYSPLIT Concentration Grid	25
2.2.2 HYSPLIT Elevation Sensitivity	28
2.3 APTool	35
2.3.1 Radiological Concentration Grid using DELFIC's PAM	37
2.4 WC-135 Analysis Systems	40

2.4.1	ERD System	41
2.4.2	DGS System	43
2.5	Gamma-ray Transport and Detection	45
2.5.1	Deterministic Radiation Transport	45
2.5.1.1	Buildup Factor	47
2.5.2	Radiation Detection Model	48
2.5.2.1	Scintillation Detection Model	48
2.5.2.1.1	Photoelectric Absorption	49
2.5.2.1.2	Compton Scattering	50
2.5.2.1.3	Pair Production	51
2.6	Model Nuclear Datasets	52
2.6.1	Gamma-ray Emission Dataset Quality Analysis	52
Chapter 3: Methodology		58
3.1	Nuclear Data Implementation	59
3.1.1	Isotopic Composition of HYSPLIT Dispersed Cloud	59
3.1.2	External Dataset Integration	60
3.2	APTool Model Additions	61
3.2.1	Nuclear Cloud Particulate Model	61
3.2.1.1	Concentration Grid Particulate Representation	64
3.2.1.1.1	Collection Filter Particulates	66
3.2.1.1.2	External Cloud Particulates	68
3.2.2	WC-135 Model	72

3.2.2.1	WC-135 Grid Traversal	72
3.2.2.2	Cloud Sample Collection	76
3.2.3	Radiation Transport and Detection	79
3.2.3.1	Source-Detector Radiation Transport	79
3.2.3.2	Radiation Detection Algorithm	82
3.2.3.2.1	Csl(Tl) Detection Algorithm	86
3.2.3.2.2	Sampler Dose and Activity Calculation	91
3.3	MCNP Verification Technique	94
3.4	Sample Case Parameters	97
Chapter 4: Results	100
4.1	Legacy APTool Capability Improvements	100
4.2	WC-135 Onboard Detector Responses	104
4.2.1	Sample Case Verification	107
4.3	Model Computational Efficiency	115
Chapter 5: Conclusions	118
Appendices	121
A.	Nuclear Test Data	122
References	129

LIST OF TABLES

2.1	Values of μ for selected event yields	6
2.2	Agreement statistics for the data shown in figure 2.3	10
2.3	ICRP recommendations for tissue weighting factors in <i>Publication 26</i> and <i>Publication 103</i>	38
2.4	Relevant properties of detectors used by the ERD and DGS	41
2.5	Summary of nuclide data present in ORIGEN resources	53
2.6	Energy-isotope abbreviations and number of time intervals used in Foster <i>et al.</i> compared to this work	54
2.7	Comparison of ENDF/B-IV data used by Foster <i>et al.</i> and ENDF/B-VIII data used in this work	56
3.1	US surface and atmospheric test events selected to assess the capabilities of the improved route planning model	98
4.1	Legacy APTool route planning model, updated model, and MCNP verification computation time for sample cases	116
A.1	US surface and atmospheric test events with sufficient observed cloud rise data for the analysis described in section 2.1.1.1	122
A.2	US and UK surface and atmospheric test events with sufficient observed stabilized cloud top data for the analysis described in section 2.1.1.2	124
A.3	US surface and atmospheric test events selected for the analysis described in section 2.2.2	127

LIST OF FIGURES

2.1	DELFIG apparent cloud center height change due to asymmetric entrainment	7
2.2	Cloud top height evolution using entrainment parameter formulation by Knowles and Jodoin	9
2.3	Stabilized cloud top height predictions, legacy DELFIG vs. morphological entrainment	11
2.4	Variation of morphological entrainment parameter with cloud rise time .	12
2.5	Cloud top height evolution using two different entrainment parameter calculation methods	13
2.6	DELFIG stabilized cloud top height predictions using NFDB inputs . . .	16
2.7	DELFIG stabilized cloud top height predictions using ASD layers with 250 m maximum thicknesses	17
2.8	DELFIG stabilized cloud top height predictions using wind data with 1,000 m minimum thicknesses	18
2.9	DELFIG stabilized cloud top height predictions using meteorology data shifted +500 m	19
2.10	DELFIG stabilized cloud top height predictions using changed burst classifications	20
2.11	F71 file structure in the context of a DELFIG run	21
2.12	Fast neutron-induced ^{235}U fission product yield	22
2.13	HYSPLIT concentration grid indexing hierarchy	26
2.14	Route latitude-longitude waypoint selection	29
2.15	Total doses for Buster-Jangle Easy	32
2.16	Average activities for Redwing Dakota	32
2.17	Maximum activities for Redwing Blackfoot	33
2.18	Maximum dose rates for Plumbbob Coulomb-B	34
2.19	Typical workflow for a simulated problem in APTool	36
2.20	Rate of decay of fallout dose rate over time post-event	39
2.21	WC-135 aircraft on tarmac at Patrick AFB	40
2.22	WC-135 geometric model and layout	42
2.23	ERD system design	43
2.24	DGS system design	44
2.25	DGS directionality quadrant boundaries	45
2.26	Gamma-ray interaction probabilities by medium Z	49
2.27	Example detector response in a 1×1 in. CsI(Tl) scintillator using a monoenergetic ^{137}Cs source	50
2.28	Relative gamma-ray emission rate per fission using fission-spectrum high-energy neutrons for three fissionable isotopes	55
2.29	Relative gamma-ray emission rate per fission for three isotopes compared to Way-Wigner decay estimate	57

3.1	Schematic of filter collection scheme used in route planning model. . .	67
3.2	Increase in size of the working system within the concentration grid due to directional detection algorithm	70
3.3	Legacy APTool simple grid traversal using ray tracing	73
3.4	Grid traversal using sampler with physical model	74
3.5	Loss of incident detector gamma-ray flux using a single voxel system .	76
3.6	Schematic of radiation transport from external particulate to DGS system detector	81
3.7	Gamma-ray absorption efficiencies for CsI(Tl) and G-M detectors . . .	84
3.8	CsI(Tl) interaction probabilities as a function of gamma-ray energy . . .	87
3.9	Factors contributing to detector energy resolution	90
3.10	CsI(Tl) detector response to fission product gamma-rays with some prominent E_γ included	92
4.1	Sampler route and mission accumulated dose comparison between legacy and updated models for Redwing Tewa	101
4.2	Dose rate time dependence comparison between legacy and updated models for Hardtack I Elder	102
4.3	Sampler route and mission accumulated activity comparison between legacy and updated models for Greenhouse George	103
4.4	WC-135 left wing ERD responses over time for the three route mission flown for the Tumbler-Snapper George event	105
4.5	WC-135 upper quadrant DGS detector responses over time for the third route flown for the Tumbler-Snapper George event	106
4.6	Sample raw output of WC-135 right wing ERD detector responses by isotope at 24 h following the Redwing Tewa event	108
4.7	Route planning model left wing ERD detector response compared to MCNP pulse-height tally for Tumbler-Snapper George	109
4.8	Route planning model left wing ERD detector response compared to MCNP pulse-height tally for Greenhouse George	110
4.9	Route planning model left wing ERD detector response compared to MCNP pulse-height tally for Hardtack I Elder	111
4.10	Route planning model left wing ERD detector response compared to MCNP pulse-height tally for Redwing Tewa	112
4.11	Agreement statistics between route planning model detector responses and MCNP pulse-height tallies for two lower yield events	113
4.12	Agreement statistics between route planning model detector responses and MCNP pulse-height tallies for two higher yield events	114

ACRONYMS, ABBREVIATIONS, AND INITIALISMS

AFB	Air Force Base
AFTAC	Air Force Technical Applications Center
AGL	above ground level
APTDriver	APTool Driver
APTool	Airborne Debris Collection Planning Tool
ARL	Air Resources Laboratory
ASD	atmospheric stability data
CRM	Cloud Rise Module
CsI(Tl)	thallium-activated cesium iodide
CSV	comma-separated value
DAEM	DELFIc Adapter and Executor Model
DASA	Defense Atomic Support Agency
DCF	dose conversion factor
DELFIc	Defense Land Fallout Interpretive Code
DFR	Detonation Forensics and Response
DGS	Directional Gamma Sensor
DOD	Department of Defense
DOE	Department of Energy
DOE/NV	DOE Nevada Operations Office
DTM	Diffusive Transport Module
ENDF/B	Evaluated Nuclear Data File, Part B
ENSDF	Evaluated Nuclear Structure Data File
EPA	Environmental Protection Agency
ERD	External Radiation Detector
F71	ORIGEN binary concentration
FGR	Federal Guidance Report
FMD	fractional mean deviation
FRMS	fractional root mean square
FWHM	full width at half-maximum
GCS	geographic coordinate system
G-M	Geiger-Mueller
GUI	graphical user interface
GZ	ground zero
HAEM	HYSPLIT Adapter and Executor Model
HYSPLIT	Hybrid Single-Particle Lagrangian Integrated Trajectory

ICRP	International Commission on Radiological Protection
JSON	JavaScript Object Notation
LUT	look-up table
MCNP	Monte Carlo N-Particle
MSL	mean sea level
NCAR	National Center for Atmospheric Research
NCEP	National Centers for Environmental Prediction
NFDB	Nuclear Fallout Data Backbone
NNSS	Nevada National Security Site
NOAA	National Oceanic and Atmospheric Administration
ODE	ordinary differential equation
OPM	Output Processor Module
ORIGEN	Oak Ridge Isotope Generation
ORNL	Oak Ridge National Laboratory
PAM	Particle Activity Module
SCALE	Standardized Computer Analyses for Licensing Evaluation

SUMMARY

Nuclear treaty verification via the use of nuclear cloud sampling aircraft, such as the US Air Force's WC-135, has been common practice since 1965. But while other verification methods have seen extensive computational models developed to support their mission, a detailed model of airborne cloud sampling has yet to be fully realized. In support of an effort to enhance and optimize the predictive capabilities of nuclear fallout software at Oak Ridge National Laboratory, a software tool capable of simulating the collection and characterization of airborne radioactive debris has been developed. In combination with the nuclear event simulation software DELFIC and the atmospheric transport and dispersion code HYSPLIT, this model's functionality includes every aspect of a sample collection mission, from the movement of the aircraft through the nuclear cloud to explicit radiation transport and detection from cloud material to onboard analysis systems. Designed with both end users and future researchers in mind, this software offers the automatic production of MCNP input files for model validation and a variety of user options to control data output and model operation. With the Air Force on the cusp of rolling out its newest line of WC-135R aircraft, this model is well-suited to play an important mission planning role in the modern age of treaty verification.

CHAPTER 1

INTRODUCTION

With the signing of the Limited Nuclear Test Ban Treaty in 1963, scientific investigations of above-ground nuclear weapons tests in the United States transitioned from physical measurements and test planning to developing treaty compliance methods and further advancing theoretical frameworks capable of predicting the effects of nuclear events. These efforts have culminated in the creation of sophisticated aircraft capable of characterizing a nuclear event based solely on the collection of radioactive particulates it disperses, as well as the development of several computational models that can accurately simulate these nuclear events, notably the Department of Defense (DOD)-developed Defense Land Fallout Interpretive Code (DELFI) [1]. Until recently, however, practical applications of these predictive models have not expanded their scope into simulating nuclear event sampling missions: though software packages such as DELFI are effective in the prediction of nuclear fallout deposition in the immediate aftermath of an event, the persistence of radioactive particulates in the atmosphere in the days and weeks following an event requires the use of additional tools not typically employed in nuclear forensics contexts.

The Oak Ridge National Laboratory (ORNL)-developed Airborne Debris Collection Planning Tool (APTool) represents one such effort to create a cohesive model that leverages all of these capabilities. By modeling the sampling tools used to collect and analyze debris produced in a nuclear event, APTool provides a method of rigorously testing existing sampling plans for nuclear events as well as the development of more effective procedures for doing so. In its current iteration, APTool allows a user to model both ground-based air debris samplers placed near the site of a nuclear event and samplers placed on aircraft flown through the nuclear cloud in the hours and days after

an event has occurred. The airborne sampler model utilized a voxelized atmospheric concentration grid to estimate dose rate and activity incident on a sampler as it traveled on a direct path between user-defined waypoints [2]. But while the ground-based sampling model has proven robust in past studies [3], the airborne sampler implementation provides only rough estimates of the quantity and species of radioactive material such a sampler would encounter during a mission. These limitations in the airborne sampler model, a consequence of the integration of software within APTool that is atypical in nuclear forensics analyses, are the gap that this work sought to fill.

To this end, three primary objectives were established to guide this effort:

1. To develop, test, and demonstrate the basic functionality of an airborne nuclear cloud sampling aircraft and its onboard analysis systems,
2. To investigate, make recommendations, and edit software tools underlying APTool in support of this effort, and
3. To establish an automated validation framework around the model to guarantee the model's integrity when in the hands of a user

As this work involved the ground-up development of many software capabilities without the influence of sponsor-specific application requests, model robustness and flexibility were the guiding principles underpinning development. To this end, this study's primary focus is to detail the incorporation of radiation transport, atmospheric modeling, computer science, and health physics concepts into the novel airborne sampler collection model; demonstration of model functionality is provided via the study of four sample cases, comparing the functionality of the legacy route planning model to current capabilities, and validation of model detector response results. As experimental WC-135 sample collection data was not made available for this work, assessing the quality of model results could proceed only via comparison to data obtained using independently-verified software.

Therefore, this study is separated into three broadly-defined chapters: chapter 2 establishes the bases on which this work builds, providing the theoretical foundation, a description of the extant software informing the route planning model, the nuclear data implemented as part of the model, and a series of short studies detailing how these data and software packages were validated for the specific needs of this work. This is followed in chapter 3 by a detailed explanation of the model itself, including the development and implementation of new model functional parts, the grid traversal and radiation transport and detection algorithms, the automated detector response validation capabilities, and how underlying software is leveraged to provide necessary model inputs. Finally, chapter 4 will detail the results of the four sample cases chosen for this work, comparing and contrasting the performance of the model with its legacy counterpart, analyzing the radiation transport and detection capabilities by a validation study using a stochastic modeling approach, and an overview of model computational efficiency and potential future improvements to be implemented in the model to augment its functionality.

CHAPTER 2

BACKGROUND AND THEORY

As APTool synthesizes the capabilities of both DELFIC and the National Oceanic and Atmospheric Administration (NOAA)'s Hybrid Single-Particle Lagrangian Integrated Trajectory (HYSPLIT), it will be necessary to first describe the operational characteristics of these models, providing the basis for further details regarding their integration. Supporting work performed using DELFIC or HYSPLIT (i.e., outside of APTool) will also be presented. This will then be followed by a brief overview of the WC-135 sampler aircraft and its onboard analysis systems, which will serve as the basis for this work. Finally, a description of the radiation transport and detection fundamentals that underpin this modeling effort will be provided, alongside a description of the decay characteristics of nuclear fallout that will serve as the basis for the methodology detailed in chapter 3.

2.1 DELFIC

DELFIC is a phenomenological code designed primarily to simulate local fallout deposition from a user-defined nuclear event, though its scope has expanded somewhat since it was first released in the 1960s. The code structure is divided into three overarching modules: the Cloud Rise Module (CRM), the Diffusive Transport Module (DTM), and the Output Processor Module (OPM). The first two of these are devoted to establishing, transporting, and depositing the cloud's particulates¹ in the immediate aftermath of a nuclear event, while the OPM handles the logistics behind formatting the resulting data. For the purposes of this work, the CRM will be of primary importance, while only some components of the OPM are heavily used and the DTM is largely ignored (though

¹DELFIc terminology defines its cloud in terms of "particles"; instead, the term "particulate" will be used in this work to differentiate the fundamental unit of the nuclear cloud from the generic term "particle" used to describe discrete units of ionizing radiation.

the DTM is used as part of DELFIC, the time scale of transport in DELFIC is insufficient for long-range sampler planning, and so this work will focus almost exclusively on HYSPLIT's transport).

In a general sense, DELFIC's CRM is used to define the size, composition, and location of a nuclear cloud at stabilization time in a format digestible by HYSPLIT. The OPM is used not to display output directly to the user, but to leverage its fallout mapping tools in a variety of ways to inform both HYSPLIT and APTool as a whole (additionally, the OPM houses the Particle Activity Module (PAM), which will be used extensively throughout this work). More details regarding the operation of these modules is provided in the following sections.

2.1.1 Nuclear Cloud Rise in DELFIC

As with any phenomenological code, the CRM applies as much of the true physics of the problem as possible, with relatively little in the way of empiricism. Thus, the bulk of the CRM consists of modeling the growth and transport of an entraining one-dimensional bubble of hot gases. Beginning the moment the fireball reaches pressure equilibrium with the surrounding atmosphere, the bubble's motion is governed by a set of coupled ordinary differential equations (ODEs) representing conservation of momentum, mass, heat, and turbulent kinetic energy [1, vol. 3], solved simultaneously using numerical methods. Stabilization of the cloud (the time at which the CRM calculations are complete) can occur either due to insufficient cloud internal energy or when the cloud is no longer rising or expanding horizontally.

2.1.1.1 Entrainment of Ambient Material in DELFIC

If the cloud were uniform in composition throughout cloud rise, these ODEs would likely suffice for prediction purposes. The consideration of water, soil, and air external to the cloud (collectively described as *entrained material*) introduces significant complexity

Table 2.1: Values of μ for selected event yields. Source: Adapted from [1, vol. 3, table B.1.1].

W (kt)	μ
0.01	0.12
0.1	0.12
1	0.12
10	0.124
100	0.167
1,000	0.226
10,000	0.305
100,000	0.411

to the problem, in some cases requiring entirely separate ODEs for each constituent material. In the case of the conservation of momentum, the change in cloud center height at a given time as a result of the uptake of this entrained material could not be adequately captured in the momentum ODEs, necessitating the introduction of the *entrainment parameter* μ :

$$\mu = 1 - \frac{u}{u + \frac{dH_C}{dt}} \quad (2.1)$$

where u is the cloud's rise velocity due to momentum alone and $\frac{dH_C}{dt}$ is the difference between the cloud's actual center height and what its center height would be using only momentum conservation (see figure 2.1). The entrainment parameter varies between zero and one, measuring the cloud center height change relative to the momentum of the cloud, and becoming significant at high yields (see table 2.1). Notably, this cloud center height change is always positive in value: entrained material is generally denser than the cloud itself, and so must congregate above the center of the cloud in order to rise quickly enough to avoid falling out [1, vol. 3, p. 31].

Introducing this parameter in the form of equation 2.1 would quickly result in a feedback loop, as the cloud's rise velocity itself is a function of the amount of entrained material already present. Instead, the entrainment parameter is set based on a pa-




Figure 2.1: DELFIC apparent cloud center height change due to asymmetric entrainment of external material. Source: Reproduced from [1, vol. 3, figure B.1.1].

parameterized fit to observed nuclear test data and scaled by event yield, resulting in the form

$$\mu = \max \begin{cases} 0.12 \\ 0.1W^{0.1} \\ 0.01W^{1/3} \end{cases} \quad (2.2)$$

where W is the event yield in kt. The case of the entrainment parameter is only one such example of the use of observed data to inform the CRM; other such simplifications can also be found in calculating the eddy-viscous drag force on the rising cloud.

This yield-dependent formulation suggests that entrainment does not vary with time or event conditions, but instead remains constant throughout cloud rise. This in turn implies that although the uptake of external material into the cloud continues throughout the cloud rise process, the increase in cloud top height due to entrainment of this material is a constant fraction of the cloud's vertical size. Thus, although the cloud top height change due to entrainment varies over the cloud rise process, it is inherently tied to the cloud's momentum. An entrainment parameter capable of responding to

changing cloud conditions could potentially provide better performance than the use of a constant parameter.

One attempt to develop an improved entrainment parameter from Knowles and Jodoin [4] took advantage of the cloud's morphology (encompassing its rise and radial growth rate) to inform the parameters in equation 2.1. Within DELFIC, the cloud is defined as an oblate spheroid of eccentricity 0.75 until the cloud's rise velocity u reaches zero (though the cloud rise calculation may continue as the cloud grows horizontally) [1, vol. 1]. Thus, while the cloud is rising, any horizontal cloud growth must be matched by commensurate vertical cloud growth by a factor of

$$\frac{a}{b} = \sqrt{1 - e^2} = \sqrt{1 - 0.75^2} = 0.661$$

where a is the cloud vertical radius, b is the cloud horizontal radius, and e is the cloud eccentricity. Thus, the change in cloud center height due only to entrainment can be expressed as

$$\frac{dH_C}{dt} = \frac{dr}{dt} \frac{a}{b}$$

where $\frac{dr}{dt}$ is the cloud horizontal growth rate, expressed as r' hereafter. Substituting this definition into equation 2.1 yields

$$\mu = 1 - \frac{u}{r' \frac{a}{b} + u} \quad (2.3)$$

The study by Knowles and Jodoin used observed cloud rise velocity and horizontal growth data from past test events to inform this calculation, implementing this entrainment parameter formulation in DELFIC as a replacement for the yield-dependent method described above in equation 2.2 rather than changing its implementation to allow the parameter to update during the calculation. Results from this study indicate




Figure 2.2: Cloud top height change over time, simulation vs. observation using the DELFIC entrainment parameter formulation described by Knowles and Jodoin [4]. All heights given relative to mean sea level (MSL).

that this morphological entrainment parameter formulation provided a more reliable approximation of stabilized cloud top height than the current DELFIC method.

As the accuracy of cloud rise calculations is central to the quality of the route planning model, one thrust of this work was to evaluate the impact of a dynamic, morphological entrainment parameter on the predictive capabilities of the CRM throughout the process of cloud rise. While use of equation 2.3 provides a reasonable approximation of stabilized cloud top height, its predictions during the process of cloud rise diverge significantly from observation (see figure 2.2).

In an attempt to improve these predictions, DELFIC was modified in order to implement the entrainment parameter formulation in equation 2.3 with the ability to update the entrainment parameter value during computation. Because the data informing equation 2.3 cannot come from DELFIC itself (in order to avoid the feedback loop discussed previously), this also required exposing the model to observed nuclear test data from which rise velocity and radial growth rates could be interpolated and drawn. Sufficient observed morphology data was found for 32 historical test events (digitized

Table 2.2: Agreement statistics for the data shown in figure 2.3. $N=32$ events.

Entrainment Method	FRMS	FMD
Legacy DELFIC	0.268	0.143
Morph. DELFIC	0.264	0.127

from DASA-1251 [5] and provided by the Detonation Forensics and Response (DFR) Group at ORNL); the list of included events is provided in table A.1. This data was implemented in DELFIC as a set of look-up tables (LUTs).

Implementing equation 2.3 with an update capability required significant code structure changes within the CRM. In the most general sense, this involved inserting an update step for the entrainment parameter that preceded the model ODE solution and update step. In practice, this change necessitated a cascade of additional changes across several modules due to widespread entrainment dependence among CRM variables. Additional changes to the OPM added the entrainment parameter to the cloud history table in the DELFIC output.

In evaluating this new model, it was first necessary to ensure that this morphological entrainment formulation predicted stabilized cloud top heights as effectively as the current DELFIC method; even a model that improves on the dynamic cloud top prediction during rise would be rendered useless if the final stabilized cloud top height is incorrect. To compare the results, each simulated test event was run on both a vanilla build of DELFIC and a DELFIC build implementing the morphological entrainment method. Stabilized cloud top heights were extracted from cloud history tables in each output and compared against observed stabilization heights, as shown in figure 2.3.

Agreement statistics provided in table 2.2 indicate that the models are virtually indistinguishable in their cloud top height prediction capabilities. The fractional root mean square (FRMS) deviation provides a measure of the average error across the predicted dataset, while the fractional mean deviation (FMD) describes the average fractional difference between predictions and observations (with a positive FMD corresponding to




Figure 2.3: Stabilized cloud top height predictions using legacy DELFIC entrainment calculation methods and the morphological entrainment formulation. Solid line indicates 100 % prediction accuracy. All heights given relative to MSL. $N=32$ events.

underprediction and a negative FMD corresponding to overprediction):

$$FRMS = \sqrt{\frac{\sum_{i=1}^N \left(\frac{x_{obs} - x_{calc}}{x_{obs}} \right)^2}{N}} \quad (2.4a)$$

$$FMD = \frac{\sum_{i=1}^N \left(\frac{x_{obs} - x_{calc}}{x_{obs}} \right)}{N} \quad (2.4b)$$

In absolute terms, neither DELFIC as it currently exists, nor a version of DELFIC employing a more responsive entrainment parameter, offer high accuracy in stabilized cloud top height predictions, particularly at higher yields. The use of morphological entrainment performed slightly more favorably, but almost negligibly so.

The true metrics of interest for this work were the entrainment parameter values and cloud top heights during the rise phase. Whereas the magnitude of the entrainment parameter in the current DELFIC model is highly constrained (see table 2.1), the morphological formulation produces a much more volatile entrainment parameter, as shown in figure 2.4. Notably, the morphological entrainment parameter formula-




Figure 2.4: Variation of morphological entrainment parameter with cloud rise time. Entrainment parameters calculated using legacy DELFIC method provided for reference.

tion produces entrainment values reasonably similar to the current yield-dependent approach (though as with stabilization time, model divergence from observation grows with yield). In all cases, the event is initialized with the yield-dependent entrainment parameter and the morphological model only takes over at the point in time at which sufficient observed test data becomes available for interpolation (at approximately 20 s for the cases shown in figure 2.4). Because the yield-dependent and morphological models often provide drastically different values for the entrainment parameter, this results in a large correction at the point where the morphological model takes over, as shown in figure 2.4. Further work investigating the early-time nature of the entrainment parameter is warranted.

Though the entrainment parameter itself behaved predictably, cloud top height as a function of rise time did not change appreciably from what had been obtained using the yield-dependent formulation. The lack of useful early-time morphological observed data forced the use of the yield-dependent entrainment method at a time in which cloud conditions are changing rapidly, while the time-dependent technique was only applied in the latter stages of the cloud rise process. While this work raised more questions




Figure 2.5: Cloud top height change over time, simulation vs. observation using two different entrainment parameter calculation methods. All heights given relative to MSL.

than it answered, efforts to analyze nuclear test films to extract additional observation data are underway [6], paving the way for future efforts to improve the entrainment parameter calculation.

2.1.1.2 DELFIC Elevation Sensitivity

Another aspect of the CRM of interest for this work was the model's sensitivity to changes in the event's height of burst, atmospheric stability data (ASD) layer thickness, and wind data layer thickness. In effect, this analysis would quantify how dependent DELFIC is on altitude via multiple pathways and provide a baseline level of meteorological data fidelity to ensure a successful DELFIC run.

The sensitivity analysis sample consisted of 93 historic US nuclear test events, comprising historic tests for which both observed stabilized cloud top height and adequately accurate meteorology data were available, as defined by ORNL's Nuclear Fallout Data Backbone (NFDB) (the events used for the purposes of this study are listed in table A.2). NFDB-provided standard DELFIC input files were run as baseline cases and a method of input file automation was written to streamline the creation and

execution of the perturbed cases.

Perturbations specific to the meteorological data (i.e., the ASD and wind layer data) took three different forms: imposition of a maximum or minimum value on the height of each layer or translating the entire dataset upward or downward in altitude. As is shown the elevation sensitivity study of HYSPLIT (discussed in section 2.2.2), models utilizing a stratified atmosphere with sufficiently thin layers can yield inexplicable results for high-yield events, and DELFIC's susceptibility to this issue had not yet been quantified. The applied perturbations were then interpolated with the original meteorological dataset in order to produce the corresponding ASD and wind data.

Though the direct effects of changing an event's burst height are plain (even without considering the operational characteristics of DELFIC), the secondary effects that can be brought about in DELFIC by making such a change are more significant than a simple z -translation of the event point would suggest. Using initial conditions provided by the user, DELFIC classifies each event according to its *scaled height of burst* Λ [1, vol. 1]:

$$\Lambda = z_b W^{-\frac{1}{3.4}} \quad (2.5)$$

where z_b is the event height of burst measured above ground level (AGL) and W is the event yield (kt). Based on the value of the event's scaled height of burst, DELFIC characterizes the event as a

$$\begin{cases} \text{subsurface burst,} & \text{if } \Lambda < 0 \\ \text{surface burst,} & \text{if } 0 \leq \Lambda < 180 \\ \text{air burst,} & \text{if } \Lambda \geq 180 \end{cases}$$

As the focus of this work is direct sampling of nuclear clouds, only surface and air bursts were considered for this study. This classification is crucial within DELFIC

because of its effect on the initial conditions within the cloud: because surface and especially subsurface bursts will entrain significant quantities of soil, the parameters governing its cloud rise must be adjusted from those that would be expected for an event with no ground interaction whatsoever. For example, the cloud momentum and mass ODEs must be initialized with the cloud total mass, which will be much greater for a surface burst than a pure air burst.

Thus, perturbations to an event's burst height took two forms: adjusting the height of burst upward or downward by a set amount, or changing the height of burst such that events originally characterized by DELFIC as surface bursts became air bursts and vice-versa. This last perturbation was applied by adjusting the height of burst term z_b in equation 2.5 to the extent required to change the characterization.

The baseline case using DELFIC inputs with no perturbations applied is shown in figure 2.6, followed by a selected set of representative perturbation examples (figures 2.7 to 2.10). The line in each plot denotes complete agreement between DELFIC and observation. Using the agreement statistics defined in equations 2.4a and 2.4b, the reduction in DELFIC's predictive capabilities could be quantitatively assessed for each perturbation (all agreement statistics are in reference to observed cloud top heights). Notably, even the most substantial adjustments to input parameters had little absolute effect on predictions, a testament to DELFIC's robust interpolation scheme for meteorological data.

Using the meteorological data the user provides as part of an input file, DELFIC interpolates it at 200 m intervals from -1 to 50 km relative to MSL, supplying data at these two extremes if the user has not already done so [1, vol. 2]. Thus, the meteorological data grid the user provides to DELFIC is only used to inform the auto-generated grid DELFIC will actually use (note that this interpolation scheme also produces arrays of size 256, the maximum array size currently allowed in DELFIC). The effectiveness of this interpolation scheme is evident when the user provides a sparse meteorological




Figure 2.6: DELFIC stabilized cloud top height prediction capability using unmodified input files from ORNL's NFDB. Solid line indicates 100 % prediction accuracy. $N=93$ events.

dataset (as in figure 2.8), but the more data the user provides, the more DELFIC's predictive capabilities become dependent on its quality (as in figure 2.7). Merely shifting the already-defined meteorological grid upward or downward (as in figure 2.9) had little effect, as the total number of interpolation points did not change.

DELFIC's capabilities also showed resilience when exposed to height of burst changes, as shown in figure 2.10. The accumulation of the many under-the-hood model adjustments associated with scaled height of burst produced slightly different DELFIC results, though these differences are on the scale of those produced by meteorological perturbations. High-yield events (those producing the highest cloud tops) appear relatively unaffected by this classification change, while low- to mid-yield event predictions are far less accurate, suggesting that these classifications become irrelevant for events of sufficient size. On the whole, however, it is clear that DELFIC is highly resistant to height of burst changes, and in general DELFIC is significantly less susceptible to elevation changes than HYSPLIT (as will be detailed in section 2.2.2).




Figure 2.7: DELFIC stabilized cloud top height prediction capability with ASD layer thicknesses constrained to a maximum of 250 m. Solid line indicates 100 % prediction accuracy. $N=93$ events.

2.1.2 Output Processing in DELFIC

Though the user rarely interacts directly with DELFIC output while using APTool, the OPM facilitates the hand-off between DELFIC and HYSPLIT in its workflow (discussed in section 2.3), while also housing the submodules that track the cloud's isotopic composition (heavily used in this work). The OPM operates on data produced from the CRM and DTM according to specifications provided by the user in the DELFIC input file.

DELFIC's output capabilities center around mapping quantities of interest onto a user-defined spatial grid (though some tools use these maps only superficially), a product of DELFIC's core mission as fallout mapping software. The user could, in theory, map any quantity associated with the cloud's mass, but DELFIC also provides several built-in options for the user to choose from, such as dose rate, deposition by particulate size, and, because DELFIC tracks the isotopic composition of the cloud, activity concentration by fission product mass chain (described in detail in section 2.1.2.1).

Maps in DELFIC are defined as two-dimensional grids at ground level, with the




Figure 2.8: DELFIC stabilized cloud top height prediction capability with wind layer thicknesses of 1,000 m. Solid line indicates 100 % prediction accuracy. $N=93$ events.

user specifying the grid domain and spacing, among other components. Users then select data of interest for mapping, which is then included as part of the DELFIC output file [1, vol. 2]. However, users can also leverage the OPM's submodules to produce radiological data of interest by providing special entries to the map request portion of the DELFIC input file. Among these is an option to produce binary concentration files compatible with the Oak Ridge Isotope Generation (ORIGEN) software [7], referred to as F71 files due to their file extension, containing concentration data for every nuclide present in each particulate size bin in the DELFIC cloud at a given time (see figure 2.11; all concentration data from F71 files used in this work is given in mol). The user can provide a list of desired transport time points for which to produce F71 files, even time points typically outside the scope of DELFIC, up to months post-event. Because the production of F71 files at a given time point requires only the application of the Bateman equation with no external source term [8]




Figure 2.9: DELFIC stabilized cloud top height prediction capability with ASD and wind data layers translated +500 m. Solid line indicates 100 % prediction accuracy. $N=93$ events.

$$N_n(t) = \sum_{i=1}^n \left[N_i(0) \prod_{j=i}^{n-1} \lambda_j \times \sum_{j=i}^n \frac{e^{-\lambda_j t}}{\prod_{\substack{k=i \\ k \neq j}}^n (\lambda_k - \lambda_j)} \right] \quad (2.6)$$

where $N_n(t)$ is the concentration of nuclide n at time t and

$$\lambda_x = \frac{\ln 2}{t_{1/2,x}} \quad (2.7)$$

is the decay constant for any nuclide x with a half-life $t_{1/2,x}$, DELFIC could effectively follow the decay of an event's debris indefinitely without introducing significant sources of uncertainty. It is the isotopic data contained in these F71 files that served as the basis for the route planning model described in this work.

2.1.2.1 Fission Product Mass Chains and Fractionation

Although measuring the yield of individual isotopes within fission products over time would be the ideal method of calculating the activity of those fission products, in prac-





Figure 2.10: DELFIC stabilized cloud top height prediction capability with event heights of burst necessary to change burst classification. Solid line indicates 100 % prediction accuracy. $N=93$ events.

tice, this type of direct measurement is not feasible. Individual isotopic concentrations are dependent on the device's target nucleus, the energy spectrum of bombarding neutrons, the time elapsed between the event and radiochemical sample analysis, and the rate of decay of fission product isotopes, resulting in a calculation that quickly becomes computationally intensive for even the simplest cases. To reduce the computational burden, and because many unstable fission product isotopes are β^- emitters, fission yields are often determined by fission product *mass chain* instead of by isotope [9].

A fission product mass chain is the series of isobars that a fission product will decay through before it reaches a stable daughter. For example, fission product mass chain 140 contains five isobars from fission product creation until stability:



For yield calculation purposes, quantities of these isotopes will be grouped together,




Figure 2.11: Breakdown of the F71 file structure in the context of a DELFIC run using N particulate size bins for the cloud.

avoiding the need to account for half-lives in time-dependent computations. The fission process creates mass chains for isobars ranging from $A \approx 60$ to 170 u, with the most prominent mass chains clustering between $A \approx 90$ to 100 u and $SA \approx 130$ to 140 u (an example of this so-called “double-hump” fission product yield curve is shown in section 2.1.2.1)

Using mass chains simplifies fission product yield calculations, but measuring their corresponding activities introduces additional complications. Because each constituent isotope of a mass chain emits radiation of differing types and energies, defining the activity of a particular mass chain in terms of *equivalent fissions* is useful. A mass chain’s number of equivalent fissions is the number of device fissions necessary to produce the number of mass chain particles that are present in the sample. Because fission product isotopic yields are well known for many device types, determining the number of equivalent fissions for a mass chain requires only a radiochemical sample measurement and a simple calculation to produce a mass chain activity concentration value that can be compared to other mass chains to determine relative yields from the




Figure 2.12: Fast neutron-induced fission product yield for ^{235}U arranged by product mass number. Data obtained from ENDF/B, version VIII.0 [10].

device in question.

Of particular interest when considering fission chain yields is the issue of *fractionation*. Fractionation is the change in isotopic makeup of a fallout sample from the event time until the sample is analyzed radiochemically [11]. Because mass chain analysis does not capture the effects of decay between isobars, however, additional analysis is necessary in order to quantifiably account for the isotope-by-isotope change in a fallout sample's composition. One method by which the degree of fractionation can be determined is through the use of r -values, defined in equation 2.8:

$$r = \frac{EF_{\text{volatile}}}{EF_{\text{refractory}}}, \quad (2.8)$$

where r describes the degree of fractionation that has occurred, EF_{volatile} is the number of equivalent fissions from a volatile mass chain, and $EF_{\text{refractory}}$ is the number of equivalent fissions from a refractory mass chain, such that the fractionation taxonomy

takes the following form:

$$\begin{cases} \text{refractory,} & \text{if } r < 1.0 \\ \text{unfractionated,} & \text{if } r = 1.0 \\ \text{volatile,} & \text{if } r > 1.0 \end{cases}$$

r -values involving mass chains 99 and 140 are often reported in literature both for their high fission yields as well as the highly refractory nature of mass chain 99 and the volatility of mass chain 140; these values are utilized as defaults for fission product calculations within APTool. Users can select mass chains of interest for a given problem and auto-generate the DELFIC map request needed to output the mass chain-specific data².

2.2 HYSPLIT

Where DELFIC uses a combination of event physics and empirical parameterizations to inform a numerical model, HYSPLIT is based entirely around numerical solutions to gridded advection-diffusion problems [12]. The core of the model centers around computing the time history of particulate advection and dispersion using a Lagrangian numerical scheme (i.e., using the moving particulate's reference frame to step through the grid). These calculations serve to inform an implicit Eulerian model (using a static reference frame) that computes grid voxel concentration data. This combined approach maximizes the allowable time step between grid iterations, improving computational efficiency with negligible impacts on accuracy and avoiding model instability (i.e., uncontrolled deviations from typical model behavior due to accumulated round-off and truncation errors) [13].

Within the context of APTool, HYSPLIT's capabilities are flexible enough to simu-

²though the detector response capability provides more details about cloud composition than the mass chain output is capable of; as a result, these results have been excluded from this work

late both the trajectory of an individual particulate within the DELFIC cloud and the transport and dispersion of the cloud as a whole. In the former, a particulate traverses the grid by atmospheric advection informed by meteorological data over a time period defined by the user, producing a time history of particulate locations starting with the stabilized cloud position. In the latter, the cloud is split into *puffs* that follow a similar advection scheme, but which can also fluctuate in size according to the dispersive nature of atmospheric turbulence [12]. Of the two HYSPLIT methods, this *dispersion model* is used exclusively in this work, and so it will be the focus of the HYSPLIT discussion to follow.

While HYSPLIT's trajectory model is narrowly focused on the path of individual particulates, the task of the dispersion model has a much broader scope. The model does not seek to follow the paths of each individual cloud particulate, but to track the general distribution of cloud material over a user-defined three-dimensional map. To do this, the program requests the cloud definition and the desired map extent (i.e., its span in latitude, longitude, and altitude) and grid spacing, as well as the dispersion time intervals of interest. The cloud is then dispersed over the area for the specified length of time, with the mass concentration of cloud material within each grid voxel recorded in units of g cm^{-3} throughout the computation. DELFIC particulate size classes (characterized as distinct air pollutants in HYSPLIT) are transported separately during the calculation, resulting in a "concentration grid" that is, in fact, multiple time-dependent concentration grids superimposed on one another; the indexing order to access these superimposed grids is shown graphically in figure 2.13. These concentration grids are stored in memory under this hierarchy with one addition: all voxels with a non-zero concentration of particulates of a given size class at a given time are flagged in the model; this permits the model to skip over these voxels when traversing the grid in order to improve computational efficiency. Thus, HYSPLIT can provide a user-customizable time history of each particulate size class's concentration grid, and also serves as the working system

for the route planning model developed for this work.

In some ways, HYSPLIT's dispersion model behaves similarly to DELFIC's atmospheric transport and fallout deposition mapping capabilities described in section 2.1: both map out mass concentration data on a user-defined grid in geographic coordinates. For the purposes of this work, however, HYSPLIT has several advantages over DELFIC's DTM. Simplest and most significant of these is HYSPLIT's use of a three-dimensional concentration grid, where DELFIC is concerned primarily with deposited debris and thus requires only two-dimensional grids (though DELFIC uses a three-dimensional grid when modeling particulate dispersion in the atmosphere). HYSPLIT is also designed to provide realistic advection over thousands of miles, while DELFIC describes itself as "intended for research in local nuclear fallout prediction" [1, vol. 1]; DELFIC's atmosphere, once interpolated, is constant for the length of the problem, and thus will quickly lose accuracy as the cloud moves away from ground zero (GZ), but HYSPLIT requires the user to provide packed meteorological data files containing continent-level and often global information to ensure accuracy over long trajectories. Finally, DELFIC's DTM and HYSPLIT's dispersion model are, in sum, built for entirely different purposes: DELFIC's atmospheric transport is designed to accurately model particulates as they settle to the ground, while HYSPLIT is concerned with modeling particulates that remain airborne for long periods of time.

2.2.1 Spatial Resolution in HYSPLIT Concentration Grid

Though the geographic coordinate system (GCS) used by both the DELFIC and HYSPLIT spatial grids is convenient for mapping purposes, a polar coordinate system is not ideal for grid traversal algorithms, such as those required for this work. Users unaware of this characteristic of the GCS are likely to request grids with non-cubic voxels, slowing down not only the grid traversal algorithm but any methods that require position sampling (such as those discussed in section 3.2.1.1). DELFIC avoids this potentiality




Figure 2.13: HYSPLIT concentration grid indexing hierarchy, where t_i is the time interval, d_i is the particulate size class, z_i is the altitude layer, ϕ_i is voxel latitude, and $\lambda_{\ell,i}$ is the voxel longitude. Non-zero concentration indexing not included.

by requiring users to enter grid voxel spacing in Cartesian coordinates (and because the user-defined grid in DELFIC is two-dimensional), while HYSPLIT requests grid data in polar coordinates: latitude and longitude with grid domain and spacing in $^\circ$, altitude as a list of layer heights in m.

By default, HYSPLIT uses 1.0° as the latitude/longitude grid spacing along with 50 m atmospheric layer heights [14]. Converting to Cartesian coordinates using the Haversine formula

$$\text{hav}(\theta) = \sin^2\left(\frac{\theta}{2}\right)$$

$$d = 2R_E \arcsin\left(\sqrt{\text{hav}(\varphi_2 - \varphi_1) + \text{hav}(\lambda_{\ell,2} - \lambda_{\ell,1}) \cos \varphi_1 \cos \varphi_2}\right)$$

where φ_i is the latitude of point i (the abscissa of the GCS), $\lambda_{\ell,i}$ is its longitude (the ordinate), R_E is the radius of the Earth, and d is the distance between points one and two (choosing an arbitrary point as reference), grid voxels will be 111.111 km in x , $111.111 \cos \varphi$ km in y , and 50 m in z , where φ is the point's latitude. Thus, for the purposes of this work, default HYSPLIT grid spacing is set to 0.01° in latitude and longitude, while altitude layer spacing is set to 1 km; this produces grid voxels approximately 1 km per side.

This scheme provides a significant increase in horizontal grid fidelity, but at what would seem like a corresponding decrease in vertical fidelity. As with DELFIC (as described in section 2.1.1.2), the level of atmospheric stratification can occasionally have a detrimental effect on prediction quality. Note also that the level of fidelity also impacts the contours of the concentration grid as a whole: because cloud concentration in a given voxel is uniform, excessively large grid spacing will smooth out changes in concentration that would have been captured had the grid been meshed more finely.

2.2.2 HYSPLIT Elevation Sensitivity

In a manner similar to the DELFIC elevation study described in section 2.1.1.2, HYSPLIT's sensitivity to elevation change was investigated, with the ultimate goal of determining optimal atmospheric layer thicknesses for use with nuclear clouds. The motivation for this study was the concern that coarse atmospheric layers in the concentration grid would fail to capture the true distribution of radioactive material within the cloud. As this study was performed primarily to determine elevation sensitivity of airborne sampler simulations, all testing was performed using either the APTool graphical user interface (GUI) or its command-line executable APTool Driver (APTDriver), and uses sampler route data (discussed in greater detail in section 2.3) as figures of merit.

To perform this study, a set of US nuclear test events first had to be chosen to serve as a cross-section of the event parameter space. These parameters included the event yield W , event GZ, the event's height of burst, and the event classification (as discussed in section 2.1.1.2).

A total of 29 test events were chosen for the study, with relevant event parameters taken from Department of Energy (DOE) sources [15]. These events are listed in table A.3. Overall, a total of 12 of these events occurred at the Nevada National Security Site (NNSS) and 17 occurred in the Pacific Ocean; 12 events had yields $W < 10$ kt, nine had yields with $10 < W < 100$ kt, four had yields with $100 < W < 1,000$ kt, and four had yields $W > 1,000$ kt; and 18 of the events were characterized by DELFIC as surface events, with 11 characterized as air bursts. Where possible, the number of chosen events of a specific type was determined by the overall fraction of test events sharing that characteristic among the population in order to provide as representative a sample as possible.

After the representative test events were selected, APTool project files were produced for each event using default values for the APTool concentration grid altitudes (that is, with grid request levels at 0, 100, 1200, 1800, 2400, 3600, and 5400 m AGL),




Figure 2.14: An example of how route latitude-longitude waypoints were selected for a given event, such that the maximum activity concentration is encountered between the waypoints. Both waypoints were placed at an altitude of 7,500 m in this example. Each route was flown at a single altitude.

20 longitudinal grid points spaced 0.2° apart, and 15 latitudinal grid points spaced 0.2° apart. These longitudinal and latitudinal spacings deviate from the default HYSPLIT settings, as discussed in section 2.2.1. This reduction in the number of latitude-longitude grid points was necessary due to the fact that, in later cases in which the number of defined altitude grid points was significantly higher than the default, the total number of grid points would exceed HYSPLIT's defined maximum (it was also necessary to reduce the number of particulate size classes in the initial DELFIC cloud to avoid exceeding memory capacity). These base cases were also defined with sampling intervals of 24 h and a total of 96 h of dispersion time.

The HYSPLIT dispersion model was then executed for these base cases. This produced a base concentration grid onto which a user could plot waypoints for an airborne sampler to pass through in a hypothetical mission scenario. It is due to this step that the APTool GUI was used for these base cases: determining adequate sampler latitude-longitude waypoints without the visual aid of the plotted concentration grid would have required considerably more effort. Only two latitude-longitude locations were chosen for these waypoints, at each extreme end of the portion of the concentration grid where concentration was present at the time of the event (an example of the positions at which waypoints would be placed is shown in figure 2.14).

After selecting the route waypoints, the APTool project file for the base case was saved in JavaScript Object Notation (JSON) format. From this point, the APTool GUI was no longer used and all project file changes were made to these JSON files directly. For each event, two parameters were adjusted to fit the needs of this study: the concentration grid altitudes for the HYSPLIT dispersion model, and the altitude component of the route waypoints plotted on the resulting concentration grid. Across all cases, the changes to the route waypoints were the same: separate routes were defined every 500 m from ground level to the top of the HYSPLIT domain (defaulted to 25,000 m) with the same latitude-longitude positions as were selected in the base case. This produced a set of 49 routes for each case that would, ideally, encompass the maximum activity concentration from a given event throughout the altitude domain of the problem and would remain consistent even with changes in the concentration altitude grid points (note that each route was also defined at sampling time index zero, such that the activity concentration throughout the cloud would be at its maximum).

Because the route positions would stay consistent across the cases for a given event, the only difference in project files between each case was the definition of HYSPLIT concentration grid altitude points. For each case, new altitude levels would be defined based on thickness from ground level to the top of the domain. New cases were produced with altitude layer thicknesses from 50 m to 10 km with a step of 50 m between cases (a total of 200 separate cases per test event, resulting in a total of 5,800 cases). Because HYSPLIT limits the number of altitude layers that can be defined for a given concentration grid (30 per grid), cases that required additional layers had new concentration grids added in, with the results combined after the cases had been run.

These cases were fed into APTool's APTDriver command-line utility, which was partially modified for the purposes of this study. The current version of the APTDriver tool only simulates the creation of the concentration grid, not sampler collection results from user-defined routes. This functionality was incorporated into the tool for this work,

as well as the capability to dump the total dose, dose rate, and activity encountered by the sampler along each route to a comma-separated value (CSV) file. Thus, each project file could be run by the APTDriver tool and produce the data of interest without requiring additional post-processing. This data was further split to consider the total dose, average dose rate, maximum dose rate, average activity, maximum activity, and total activity encountered along each route.

Because the clouds from many events did not reach the top of the HYSPLIT domain (or, conversely, left little debris at low altitudes), the data was then filtered to exclude routes in which no activity was encountered. Unfortunately, two of the cases (Ivy King and Upshot-Knothole Dixie) encountered errors in their HYSPLIT dispersion model runs (likely due in part to the poor quality of their respective meteorological data sets), resulting in empty concentration grids. As a result, it was necessary to exclude this data from consideration. The conclusions drawn from this study, therefore, are based on results from the 27 events with meaningful data.

Note that only a specific set of HYSPLIT layer thickness cases were included in plots. These levels were chosen as representative of each data set, as the differences between adjacent cases were often negligible. Any cases not following this convention will be noted.

Though the data recorded from APTool was split into six types, the trends across cases for each of these data types were largely consistent for a given event. For example, the total dose in each case for Buster-Jangle Easy (shown in figure 2.15) indicates decreasing dose variability with increasing layer thickness, a pattern that is echoed in the average activity encountered in each case for Redwing Dakota (shown in figure 2.16), despite significant differences in each of the test event parameters used in this study. While higher variability in cases with smaller layer thicknesses is expected, the persistence of this phenomenon despite variations in data type and event parameters suggests that layer thickness can significantly impact concentration grid results.




Figure 2.15: Total dose received by a sampler across single-altitude routes for the Buster-Jangle Easy test event.




Figure 2.16: Average activity encountered by a sampler across single-altitude routes for the Redwing Dakota test event.




Figure 2.17: Maximum activity encountered by a sampler across each single-altitude routes for the Redwing Blackfoot test event.

This pattern of variability is generally consistent across the dataset, but when considering the relative effectiveness of similar layer thicknesses (particularly those with very fine spacing), results are more ambiguous. In some cases, like Redwing Dakota's shown in figure 2.16, fine grid spacing can result in inexplicable results. Note that the average activities between the 250 m and 500 m cases shows good agreement until 7 km in altitude, at which point the 250 m case drops to zero. The 500 m case also drops to zero activity at 14 km, even though clearly activity should have been encountered at these levels given the results from other cases. This phenomenon is inconsistent between test events, suggesting some incongruity between the HYSPLIT level definition, the route altitude, and the concentration grid that results in unreliable data. Note also that the 50 m thickness case provides non-zero results at much higher altitudes than the 500 m case, though these results are interspersed with several altitudes at which zero activity was recorded.

In other cases, such as that of the Redwing Blackfoot cases shown in figure 2.17, the differences between these finely-meshed concentration grids manifest themselves as variations between the results obtained at the same route altitude. At 4.5 km al-




Figure 2.18: Maximum dose rate encountered by a sampler across single-altitude routes for the Plumbbob Coulomb-B test event.

titude, the total activity encountered (i.e., the integrated activity over the entire route) using 250 m layer thicknesses is approximately 350 μCi , that encountered using 500 m layer thicknesses is approximately 310 μCi , and that encountered using 1,000 m layer thicknesses is approximately 280 μCi . Though these differences amount to little when considering the entire concentration grid domain, they also suggest that larger grid layer thicknesses can flatten out contours in the data that may be of interest. Again, however, the use of too fine a grid can have significant consequences: using 50 m layer thicknesses, non-zero activities are only recorded up to 0.5 km, highlighting the need to compromise between extremely fine and extremely coarse grid spacing for a given problem.

Determining such an optimal layer thickness is highly dependent on the event in question, but it is clear that the layer thickness range chosen for this study contains several thicknesses that could not be considered optimal. Throughout figures 2.15 to 2.18, layer thicknesses exceeding 2,500 m provide results well outside of expectations, as their layers encompass so many route altitudes that they can almost be considered to provide an altitude-averaged result across the domain. For most applications, such

a result is undesirable. In contrast, the 50 m layer thickness cases provide a level of fineness unsuitable to such a large domain, resulting in sparse datasets that do not resemble the concentration grid.

In the thickness range of 250 to 2,500 m, the optimal choice for a given event should be determined based primarily on the event yield. As is clear in figure 2.16, high-yield events offer poor results with meshing on the order of 250 to 500 m, while 1,000 and 2,500 m thicknesses provide consistent results throughout the domain. Meanwhile, figure 2.17 indicates that for lower-yield events, these large layer thicknesses can suppress some of the contours of the concentration grid by averaging the concentration over too much of the relevant altitude domain.

In general, the results obtained using 1,000 m layer thicknesses are the most consistent, with few anomalous “zeroed” altitudes and a reasonable level of accuracy throughout the domain. For cases in which the cloud is not expected to exceed 10 km (or in which the event yield $W < 10$ kt), 500 m layer thicknesses may be more appropriate in order to achieve a more accurate result and capture a larger portion of the concentration grid’s variability with altitude.

2.3 APTool

Though supplying only a few minor capabilities as a top layer for the DELFIC-HYSPLIT integration, APTool provides a seamless interface while granting the user the flexibility to adjust DELFIC and HYSPLIT input parameters without requiring the user to manually write input files.

Synergy between the models is facilitated through the DELFIC Adapter and Executor Model (DAEM) and HYSPLIT Adapter and Executor Model (HAEM), which interpret user input (whether through the use of the APTool GUI or via a JSON configuration file), run the models with these specifications, and then process the output in a manner specified by the user (an example of the typical route-planning workflow in APTool is



Figure 2.19: Typical workflow for a simulated problem in APTool. Radiological quantity computations are excluded for simplicity.

provided in figure 2.19). Though APTool does not require the user themselves to run either DELFIC or HYSPLIT, program executables are not provided by APTool; DAEM and HAEM are designed to serve as object-oriented interfaces between the purely procedural DELFIC and HYSPLIT and an arbitrary top-level interface (in this case, APTool).

One critical aspect of this top-level design is the standardization of meteorological data between the two models. As DELFIC's meteorological input requirements are much simpler than HYSPLIT's, DAEM is designed to extract the data it needs for a DELFIC run from the HYSPLIT meteorological data file, which serves the added purpose of ensuring that atmospheric conditions remain consistent throughout a given problem. DELFIC's CRM run requires only that the user define the event conditions and provide this meteorological file, with which APTool will create a DELFIC input file to the chosen specifications. The resulting stabilized cloud definition can then be formatted by HAEM and, along with any HYSPLIT-specific input information, passed to either HYSPLIT transport model (trajectory or dispersion) [16].

In a general sense, APTool's additional capabilities center around applying HYSPLIT's output concentration grid and DELFIC's OPM to simulate the distribution of radioactive material over the area of interest. Of these, the airborne route planning is the primary focus, but APTool also provides a module permitting the user to define ground-based air sample collection stations to determine the concentration of mass over the entire HYSPLIT trajectory and has proven effective in optimizing the location and number of samplers needed to deliver an air sample of high quality [3].

2.3.1 Radiological Concentration Grid using DELFIC's PAM

While APTool only requires users to explicitly run DELFIC's CRM once during the course of a problem, the OPM's capabilities are utilized at many points over the course of the workflow (though these calls are often kept out of sight of the user). The OPM's primary use in APTool (prior to this work) was in applying dose conversion factors (DCFs) to DELFIC's cloud in order to map out the distribution of radioactive material post-event. These DCFs encompass the dose rate coefficient data provided by the Environmental Protection Agency (EPA) as well as human-specific weighting factors provided by the International Commission on Radiological Protection (ICRP).

To ensure the accuracy of this radiological data, a supporting effort involved in this work was focused on updating the relevant DCFs used by DELFIC. In particular, the dose rate coefficients associated with total air immersion were updated to match the data provided in Federal Guidance Report (FGR)-15 [17], while organ weighting factors were adjusted to align with data in ICRP Publication 103 [18] (organ weighting factor improvements are outlined in table 2.3). Implementing the new DCFs in DELFIC was quite simple: the data is hard-coded into DELFIC's source code as a series of LUTs, which were overwritten to include the new data. A DELFIC build using the new DCFs was then set as the target executable for APTool.

The DCFs taken from FGR-15 are provided in units of $\text{Sv m}^3 \text{Bq}^{-1} \text{s}^{-1}$, or dose rate equivalent per unit volume source. Organ weighting factors are unitless values between zero and one. These combined coefficients, when applied to the cloud property data produced by DELFIC nuclide-by-nuclide, provide approximations of the exposure rate (rem h^{-1}) in the cloud at a given time. Because the rate of fallout activity decay is calculated within DELFIC for fission products over the time period of interest for this work (on the order of days and weeks post-event), this data can then be used to approximate dose, dose rate, or activity at any point within that time period [20].

Decay of fallout activity within 0.5 to 5,000 h post-event can be approximated using

Table 2.3: ICRP recommendations for tissue weighting factors in *Publication 26* [19] and *Publication 103* [18].

Tissue	Weighting Factor w_T	
	1977 <i>Pub. 26</i>	2007 <i>Pub. 103</i>
Bone surfaces	0.03	0.01
Bladder	—	0.04
Brain	—	0.01
Breast	0.15	0.12
Colon	—	0.12 ¹
Esophagus	—	0.04
Gonads	0.25	0.08 ²
Liver	—	0.04
Lungs	0.12	0.12
Red bone marrow	0.12	0.12
Salivary glands	—	0.01
Skin	—	0.01
Stomach	—	0.12
Thyroid	0.03	0.04
Remainder	0.30 ³	0.12 ⁴

¹ The dose to the colon is taken to be the mass-weighted mean of upper large intestine and lower large intestine doses.

² The w_T for gonads is applied to the mean of the doses to testes and ovaries.

³ The five most highly irradiated other organs and tissues are included in the remainder, each with $w_T = 0.06$.

⁴ The specified remainder tissues are: adrenals, extrathoracic tissue, gall bladder, heart, kidneys, lymphatic nodes, muscle, oral mucosa, pancreas, prostate (in men), uterus/cervix (in women), small intestine, spleen, and thymus.

Way-Wigner or $t^{-1.2}$ decay [21], expressed as

$$R_t \approx R_1 t^{-1.2} \quad (2.9)$$

where R_t is the dose rate at time t and R_1 is the *unit-time reference dose rate*, which is the dose rate at a time dependent on the units in which time is expressed (e.g., for time measured in hours, R_1 would be the dose rate 1 h post-event). As shown in figure 2.20, this approximation is quite accurate over the time period of interest. Accumulated dose and accumulated activity can be computed during the course of this time period by




Figure 2.20: Rate of decay of fallout dose rate over time post-event, up to approximately one month, as well as the Way-Wigner decay approximation. Dose rate data is expressed relative to the dose rate at 1 h post-event. Source: Adapted from [20, figure 9.16a].

integrating:

$$D \approx R_1 \int_{t_a}^{t_b} t^{-1.2} dt$$

$$D \approx 5R_1 \left(t_a^{-0.2} - t_b^{-0.2} \right) \quad (2.10)$$

where D is dose and R_1 is the unit-time reference dose rate (note that activity and unit-time reference activity can be used in a similar manner). Unit-time reference dose rates can be calculated using the cloud dose rate at the present time as in equation 2.9, then incorporated into the accumulation defined in equation 2.10. Thus, in combination with the grid traversal algorithm described in section 3.2.2.1, the time evolution of the cloud's activity can be followed throughout the duration of a sampler's route.



Figure 2.21: WC-135 aircraft on tarmac at Patrick AFB. Source: [23].

2.4 WC-135 Analysis Systems

The Boeing WC-135, shown in figure 2.21, is the purpose-fit form of the KC-135R aircraft used by the US Air Force to detect and sample nuclear clouds. There are currently two active aircraft using legacy systems, with three aircraft in the process of being re-fit with the systems described in this work [22]. Though the design of the aircraft itself plays little role in the model outside of its geometry (discussed in detail in section 3.2.2), its flight capabilities set some hard limits on the route planning simulation: stabilized clouds from events with yields > 100 kt typically exceed its flight ceiling of approximately 12,000 m (40,000 ft) [20], which could impact the selection of route waypoints, while its maximum flight speed (640 km h^{-1} or 403 mph) sets a lower limit on the amount of time required for the sampler to traverse a grid voxel.

The APTool route planner's personnel-centered model represents only one-third of the detection systems found on board the WC-135: in addition to the changes made to DELFIC's PAM to provide more accurate personnel dose rate data, this work's primary concern has been the accurate simulation of those as-yet unexplored External

Table 2.4: Relevant properties of detectors used by the ERD and DGS.


Detector	Geometry	Composition	Density (g cm ⁻³)	Dimensions (in.)	Energy Resolution (%)
CsI(Tl) Scintillator	Rect. Prism	>49.9 % Cs >49.9 % I <0.1 % TI	4.51×10^0	$3 \times 3 \times 14$	17.5
Geiger-Mueller Counter	Cylinder	98.7 % Ar >0.9 % H >0.3 % C	9.10×10^{-3}	2×9	—

Radiation Detector (ERD) and Directional Gamma Sensor (DGS) systems (shown on board the WC-135 geometric model in figure 2.22; note that although the aircraft cabin was modeled with a diameter of 12.8 m, this value is on the upper end of feasibility for a sampler cabin and is left as an adjustable parameter for the user). Despite inherent similarities between the two systems, their methods of use aboard the WC-135 are markedly different.


2.4.1 ERD System

The ERD system, as the name suggests, is the primary means through which the WC-135 performs its central mission task: cloud sample analysis. Paired with specialized filters designed to capture small particulates, the ERD provides a real-time indicator of sample composition during the course of a mission. The system consists of one thallium-activated cesium iodide (CsI(Tl)) scintillator and one Geiger-Mueller (G-M) counter sandwiched between ¼ in. plates of tungsten (on top) and stainless steel (on bottom). This design is depicted in figure 2.23 and the properties of each detector are summarized in table 2.4.

Because this system is meant to analyze the cloud directly, it is located outside the main body of the aircraft (the left and right sides of the aircraft operate independent ERD systems, but they will be referred to collectively as the “ERD system”). Efflu-



(a) Exterior.



(b) Schematic of assumed WC-135 analysis system layout used in this model with dimensions (not to scale).

Figure 2.22: WC-135 geometric model as used in verification suite. Layout consistent between model and verification suite.




Figure 2.23: ERD system design with dimensions. Source: Adapted from [24, figure 2].

ent air passes through the specialized filters during flight, catching any particulates of sufficient size. The ERD's response to the gamma-rays emitted by radioactive particulate isotopes on the filter is continuously measured and recorded, providing snapshots of the cloud's composition at a given time (the ERD's CsI(Tl) detector measures the response, with the G-M detector serving as a backup in case of CsI(Tl) detector failure [24]).

The mounting plates holding the detectors in place also provide a small measure of low energy (300 keV and below) gamma-ray attenuation, helping to minimize the contribution of background radiation to the detector response. In order to obtain as many samples as possible, filters are changed every 30 min while in the cloud; this has the additional benefit of ensuring that gamma-rays produced by particulates on the filter are not attenuated due to the interposition of other particulates between itself and the detector.

2.4.2 DGS System

The DGS system's primary objective is to maximize the quality and quantity of samples for the ERD system to analyze. Consisting of four CsI(Tl) scintillators (see table 2.4) mounted in a cruciform pattern and quartered by $\frac{1}{4}$ in. stainless steel plates (see figure 2.24), the DGS system's primary uses are to recognize the presence of a radioac-




Figure 2.24: DGS system design with dimensions. Additional aluminum framing not shown.

tive plume and, once inside, to direct the flight crew toward areas of high activity in order to maximize mission efficiency. The steel plates placed between each detector serve two purposes: to minimize crosstalk between detectors and to ensure that each detector is primarily sensitive to one quadrant in the yz -plane (as defined in figure 2.25).

DGS detector responses serve as inputs to the directional algorithm, which weighs the strength of the response in one detector relative to the other three and signals the result using a set of lights pointing in cruciform directions. The system is also time-windowed, ensuring that the indicator is providing the most up-to-date vectoring information possible. As shown in figure 2.22b, the DGS is placed near the center of the aircraft and is not exposed to cloud particulates directly (as is the case with the ERD); because the DGS's purpose is to detect the contours of the radioactive plume outside the WC-135, DGS detector response is designed to be mostly attributable to gamma-rays produced by particulates that surround the aircraft rather than those that have been captured by onboard filter systems.




Figure 2.25: DGS directionality quadrant boundaries.

2.5 Gamma-ray Transport and Detection

Providing accurate response simulations for these WC-135 systems requires rigorous and precise gamma-ray and scintillator models. These models were developed for this work from the ground up, based on deterministic radiation transport and a simplified detection model (to be described in greater detail in section 3.2.3) in order to provide an accurate but computationally efficient simulation capability.

2.5.1 Deterministic Radiation Transport

Though a stochastic (or random-walk) approach to transport simulation would provide a more rigorous and potentially more accurate result, the corresponding loss in computation speed makes the technique infeasible for this application (i.e., a long-range detection scenario with large spherical divergence between source and detector in a low-scatter medium) [25]. Instead, this work uses a ray tracing method to simulate the

attenuation of gamma-rays between source and detector as

$$I(E) = \frac{I_0(E)}{4\pi d^2} e^{-\mu(\rho, E)d} \quad (2.11)$$

where $I(E)$ is the intensity of gamma-rays of energy E at the detector, I_0 is their initial intensity at the source, d is the source-detector path length, $\mu(\rho, E)$ is the linear attenuation coefficient for gamma-rays of energy E through atmospheric air of density ρ , with the exponential $e^{-\mu(\rho, E)d}$ often referred to as a whole as the *attenuation factor* (note that this method assumes that particulates are emitting gamma-rays isotropically). While a stochastic model would model the emission and detection of individual gamma-rays in a manner reflective of real-world behavior, the concept of gamma-ray intensity is more abstract. In this context, intensity refers to the quantity of gamma-rays present with a given energy; thus, the source intensity $I_0(E)$ describes the number of gamma-rays emitted from a nuclide at energy E , while $I(E)$ is the number of gamma-rays that were able to traverse a distance d without being absorbed or scattered by the attenuating medium. As such, both “gamma-ray” and “gamma-rays” are used interchangeably throughout this work: “gamma-ray intensity” refers to the number of gamma-rays of the same energy emitted by a particulate per unit time, making the terms functionally equivalent in the context of deterministic gamma-ray transport. Furthermore, because this model is not time-invariant, all intensities can be considered to be provided in terms of the number of gamma-rays per unit time, unless otherwise stated.

In order to account for the possibility of detecting gamma-rays originating from different atmospheric layers, a summation term was added to the exponential term of equation 2.11 in order to capture the change in attenuation coefficient with altitude:

$$I(E) = \frac{I_0(E)}{4\pi d^2} e^{\left(-\frac{\mu(0, E)}{\rho} \sum_i \rho_i d_i\right)} \quad (2.12)$$

where $\frac{\mu(0, E)}{\rho}$ is the mass attenuation coefficient for air at sea level, ρ_i is the atmospheric

density of layer i , and d_i is the source-detector path length segment that passes through layer i .

2.5.1.1 *Buildup Factor*

This method achieves results within 1% of the uncollided point detector flux using stochastic methods [25], but does not consider the possibility of gamma-ray scatter into a detector region. Conventionally, deterministic approaches to gamma-ray transport will impose a *buildup factor* on equation 2.11, scaling the gamma-ray flux at the detector by a fixed value based on the energy, medium, path length, and solid angle of incidence between source and detector. This dependence on so many problem-specific variables makes determining the buildup factors for a given set of gamma-rays as computationally expensive as transporting the uncollided gamma-rays themselves; as a result, buildup factors are often pre-computed using stochastic simulations and accessed as LUTs in the deterministic transport algorithm. For model scenarios in which buildup factor parameters are not static, however, the number of pre-computations required can result in a memory-expensive LUT in which lookups are less computationally efficient than the calculation itself.

To balance the size of LUTs with computational efficiency, the model designed for this work computes buildup factors for gamma-rays emitted from collection filters (as the geometry between source and detector is static over the course of a route), but assumes that only uncollided gamma-rays will reach detectors from particulates located outside the aircraft. In worst-case scenarios, where all gamma-ray interactions between source and detector are scattering events in the direction of the detector, uncollided gamma-rays account for only 2% of the total gamma-ray flux at the detector [25]. More realistic scenarios, accounting for low atmospheric gamma-ray interaction probabilities and isotropic scattering, suggest instead that once-scattered gamma-rays contribute only 2% of the total flux. This work uses a hybrid approach, using buildup factors for

transport to the short-range, static ERD detectors and assuming they are negligible for the long-range transport between external particulates and the DGS.

2.5.2 Radiation Detection Model

As described in section 2.4, the WC-135 uses primarily CsI(Tl) scintillation detectors for analysis with G-M counter backups. Models were developed for each type with a level of detail corresponding to their complexity, with the spectroscopic CsI(Tl) requiring significantly greater fidelity than the basic G-M; in fact, the G-M counter's capabilities are modeled simply as a subset of that of the scintillator, as differences in interaction mechanisms between detector types are woven into each respective detector's intrinsic efficiency (described in section 3.2.3.2). As a result, only the more complex scintillation mechanism will be fully detailed here.

2.5.2.1 Scintillation Detection Model

In an infinitely large detector with 100% detection efficiency, each gamma-ray in the detector would deposit its full energy in a photoelectric interaction with a detector electron, resulting in a detector response that exactly mirrors the energy distribution of the incident gamma-rays produced by cloud particulates. The detectors used in this work, in contrast, are somewhat large, with detection efficiencies that vary with incident gamma-ray energy. Photoelectric interaction probabilities diminish with increasing energy, reducing the likelihood of a full-energy peak event, the electron energy registered by the detection system is subject to statistical noise effects and the detector medium's inherent energy resolution, and the finite size of the detector leads to additional losses due to photon escape. These imperfections in the detection system result in responses with little resemblance to the incident energy distribution, and so the simulation of these imperfections forms the core of the scintillation detector model used in this work.

The mode of interaction between a gamma-ray and a medium is probabilistic in




Figure 2.26: Relative probabilities for the three primary gamma-ray interaction mechanisms by energy and medium, highlighting the equivalent atomic number for a CsI(Tl) scintillator. Additional lines indicate where interaction probabilities are equivalent for the adjacent interaction types. Source: Adapted from [26, figure 25.1.1].

nature, with higher-energy gamma-rays exhibiting an increased likelihood of Compton scattering or pair production events. The relative probabilities for these three interaction types are shown in figure 2.26. For each incident gamma-ray, a interaction type was selected based on its probability at the given gamma-ray energy and the resulting process would be followed until the gamma-ray deposited all of its energy (or escaped the detector); the physics behind the interactions simulated for the purposes of this work are detailed in the following sections.

2.5.2.1.1 Photoelectric Absorption

Photoelectric absorption, in which the original gamma-ray transfers most of its kinetic energy to an electron in the detector volume, is the ideal gamma-ray interaction process from a detector response perspective. As long as the photoelectron does not




Figure 2.27: Example detector response in a 1×1 in. CsI(Tl) scintillator using a monoenergetic ^{137}Cs source. Backscatter peak is attributable to Compton scattering from materials surrounding detector.

escape, the full energy of the gamma-ray will be deposited in the detector using a single photoelectron (referred to as the *photopeak* in the response; see figure 2.27). Photoelectric interactions are typically prevalent for low-energy gamma-rays, but this probability decreases with growing energy.

2.5.2.1.2 Compton Scattering

Compton scattering events are the nearest gamma-ray analogue to elastic scattering interactions (in the laboratory frame of reference), with the energy transferred to detector electrons determined by the incident gamma-ray energy and the scattering angle θ [26]:

$$E_{e^-} = h\nu \left(\frac{\alpha(1 - \cos \theta)}{1 + \alpha(1 - \cos \theta)} \right) \quad (2.13)$$

where E_{e^-} is the energy of the electron, $h\nu$ is the energy of the incident gamma-ray,

and $\alpha = h\nu/m_0c^2$, the ratio of the incident gamma-ray energy to the electron's rest mass energy m_0c^2 (511 keV). In all Compton events, the scattered gamma-ray will maintain a kinetic energy $h\nu' = h\nu - E_{e^-}$ that must be accounted for; as a result, Compton scattering produces electrons with a spectrum of energies, which appear in a monoenergetic source's response as the Compton continuum. The edge of the Compton continuum is the energy for which the scattering angle $\theta \approx \pi$, in which the gamma-ray imparts as much of its kinetic energy as possible to the electron; the Compton continuum and the Compton edge are clear in the example shown in figure 2.27.

2.5.2.1.3 Pair Production

In pair production interactions, the incident gamma-ray creates an electron-positron pair in the detector volume, which then travel away to deposit their energy. The creation of the electron-positron pair requires an incident gamma-ray with kinetic energy exceeding the rest mass energy of the two particles combined; thus, pair production is only energetically possible if $h\nu > 2m_0c^2$, and any excess kinetic energy is split between the electron and positron:

$$E_{e^-} + E_{e^+} = h\nu - 2m_0c^2 \quad (2.14)$$

where E_{e^+} is the energy of the positron.

While the electron deposits its energy in the detector volume as with other interaction types, the positron deposits its energy more indirectly. After slowing down, if the positron remains in the detector medium, it will annihilate with a detector electron, creating two photons of energy m_0c^2 that are emitted in opposite directions, which can then result in more photoelectric interactions or Compton scatters in the detector medium. In small detectors, annihilation photons can often escape the detector volume prior to absorption, giving rise to the single and double escape peak features in

the detector response (located at $h\nu - m_0c^2$ and $h\nu - 2m_0c^2$ respectively) that correspond to deposition of only the electron and positron kinetic energies. These features are not present in the response shown in figure 2.27.

2.6 Model Nuclear Datasets

Though APTool does not fall under the Standardized Computer Analyses for Licensing Evaluation (SCALE) family of software, the ubiquity of SCALE at ORNL has led to the adoption of its data resources in many other tools (this data, in turn, is derived from Evaluated Nuclear Data File, Part B (ENDF/B)-VII.1). In particular, the resources provided by SCALE's submodule ORIGEN are ideal for this application, as the notation used in ORIGEN binary concentration (F71) files is consistent with that of ORIGEN's data; the three data files of particular interest were the decay, gamma-ray emission, and fission yield resources. As the fission yield resource is used as a part of DELFIC and did not require additional use in this work, the implementation of the decay and radiation emission data files in APTool will be the focus here (a summary of the nuclide decay and emission data is outlined in table 2.5). Notably, gamma-ray emission data is provided in units of γ disintegration⁻¹, remaining constant with respect to activity throughout a given problem. This gamma emission data was of particular interest, as it was first introduced to APTool due to this study; therefore, as part of the effort supporting this work, this gamma emission data was chosen for a deeper investigation.

2.6.1 Gamma-ray Emission Dataset Quality Analysis

The driving motivation behind investigating the gamma emission data used in this work centers around historical context: contemporary results are typically computed utilizing the most modern nuclear data available, data that any historical results used for comparison would not have had access to. Thus, it is often instructive to directly compare the contemporary datasets to touchstone historical works in order to determine whether

Table 2.5: Summary of nuclide data present in ORIGEN resources [7]. Sub-library data for activation products and actinides include only those nuclides not present in fission product dataset. Emission resource contains only gamma-ray emission data.

Decay Resource	Emission Resource	
2,237	1,277	nuclides (total)
1,675	1,277	nuclides (unique)
350	207	activation products
174	150	actinides
1,151	920	fission products

divergence in results is a product of data differences or errors in methodology.

The historical study chosen for comparison with this work focuses on the decay of gamma-ray emission spectra over a timescale comparable to that of this work, published in 1978 by Foster *et al.* [27]. Though this work was limited by the technological challenges of the time as well as the quality of the data available (particularly early-time decay data), it has maintained its relevance as a point of reference for delayed gamma-ray flux computations using common fissionable isotopes [28].

The methodology used by Foster *et al.* looks at the gamma-rays emitted from fission products produced by three fissionable nuclides (^{239}Pu , ^{235}U , and ^{238}U) using a constant flux of $10^{13} \text{ n cm}^{-2} \text{ s}^{-1}$ applied for 10^{-4} s using two different incident neutron energy spectra (fission neutrons and 14 MeV high-energy neutrons) over the time period from 0.1 to 1,000 s post-fission. The study primarily uses data from ENDF/B-IV (though the 14 MeV ^{239}Pu case uses data from ENDF/B-V) and produces time-dependent fission product decay data using CINDER-10 [29], a fission product transmutation and decay code similar in function to SCALE's ORIGEN. Technological and contractual limitations resulted in the use of only nine time intervals within 0.1 to 1,000 s, except for three cases in which 29 time intervals were used; table 2.6 details the energy-isotope cases, the number of time intervals available, and a comparison to data used in this work's investigation.


Much of this work's investigation was focused on attempting to replicate the results

Table 2.6: Energy-isotope abbreviations and number of time intervals used in Foster *et al.* [27, table 1] compared to this work.


Fissioning Nuclide	Neutron Energy	Abbreviation	Number of Decay Times, 0.1 to 1,000 s	
			Foster et al.	This Work
^{235}U	fission spectrum	25F	9	37
^{238}U	fission spectrum	28F	29	37
^{239}Pu	fission spectrum	49F	29	37
^{235}U	14 MeV	25H	9	37
^{238}U	14 MeV	28H	9	37
^{239}Pu	14 MeV	49H	29	37

presented in Foster *et al.*, figure 4 using modern data and tools. ORIGEN was used to produce the time-dependent fission product decay data, and with many fewer limitations on the allowed number of time intervals to request than with CINDER, nine time points per decade within 0.1 to 1,000 s were used for each case along with a final time point at 10,000 s (37 time points total). In addition to the increased number of time points used, the nuclear dataset available to ORIGEN is far more detailed than that available to Foster *et al.*; these data improvements are summarized in table 2.7.

Although mirroring the general methodology outlined by Foster *et al.* was rather simple, generating the composite gamma-ray spectra and normalizing the data proved more complex. As shown in table 2.7, Foster *et al.* were producing emission spectra for a total of 825 fission products using only the 180 nuclides with this data explicitly defined; to account for the missing spectra, it was assumed that the average spectrum of the remaining 645 nuclides was the same as that of the known 180. While ORIGEN uses a similar averaging scheme in its spectrum calculations [30], only 112 nuclides must use averaged spectra, significantly reducing the impact of this computation on results. In addition, Foster *et al.*, p. 12 claims to normalize calculated spectra to “the total beta or gamma energy-emission rate”, and yet figure 4 (data from which is included



(a) Spectra using fission-spectrum neutrons.



(b) Spectra using high-energy neutrons.

Figure 2.28: Relative gamma-ray emission rate per fission using fission-spectrum and high-energy neutrons for three fissionable isotopes. Plots on left adapted from [27, figure 4].

Table 2.7: Comparison of ENDF/B-IV data used by Foster *et al.* [27, table 2] and ENDF/B-VIII data used in this work.

Foster et al.	This Work	
825	1,151	nuclides (total, counting isomers separately)
112	927	isomers with half lives >0.1 s
42	51	different elements
96	109	different mass numbers
181	557	have neutron-interaction cross sections
180	939	have explicit data on beta and/or gamma spectra
712	1,003	are unstable and have average α , β , and γ energies and branching fractions
825	1,151	have fission yields for each of six fissionable nuclides for one or more ranges of neutron energy

on the left side of figure 2.28) indicates that ^{238}U fission produced more than 100% of the total gamma energy-emission rate. From this figure, it is clear that gamma spectra were in fact normalized to the energy-emission rate of the 25H case at 0.1 s post-fission, and it is this normalization scheme that was employed for the purposes of this work (the results of which occupy the right side of figure 2.28).

Of particular note in figure 2.28 is the significantly lower gamma energy-emission rate for the 28F case between ENDF/B-IV and ENDF/B-VIII. Other authors have suggested that historical works did not properly account for photon emission anisotropy in ^{238}U composite spectra [31], which likely accounts for the discrepancy indicated here. Though this is a plausible explanation for the phenomenon and both ^{235}U and ^{239}Pu data align well with data reported by Foster *et al.*, the disconnect between the normalization scheme described on the page and that which is used in figures make drawing definitive conclusions difficult.

Though not strictly related to this investigation, the data used was also conveniently structured such that the limitations of the Way-Wigner decay approximation discussed in section 2.3.1 could be easily explored. Using 1,000 s as the reference time, figure 2.29 indicates that the Way-Wigner approximation can provide accurate results at




Figure 2.29: Relative gamma-ray emission rate per fission for three isotopes compared to Way-Wigner decay estimate. All data from ENDF/B-VIII. Unit reference time for $t^{-1.2}$ decay curves is 1,000 s.

post-event times as early as 1 min. Although this work is not primarily focused on such short timescales, this investigation indicates that the Way-Wigner approximation is generally robust at a wide range of time scales for an appropriate unit reference time.

CHAPTER 3

METHODOLOGY

Though the route planning model developed for this work is novel, the foundation on which it is constructed is a clean build of APTool version 1.4.0. Significant effort was expended to ensure that this model did not impact the core functionality of vanilla APTool; all capabilities up to and including the legacy APTool route planning model were left untouched by this work. All software was written in C++ (unless otherwise specified) using the Qt [32] framework for visualization.

Though this new model maintains compatibility with legacy APTool, this is not the case with the converse. Significant additions and edits to backend code were required to support the new capabilities this route planning model introduces; a description of this work will precede a discussion of the additions made to APTool itself. The description of the model will focus on additional concentration grid capabilities, cloud particulate simulation, implementation of the radiation transport and detection algorithms, along with miscellaneous other required additions. Details of the verification suite for the model will follow this description, including a discussion of the automated creation of Monte Carlo N-Particle (MCNP) transport software input files and the technique through which results from the route planning model are compared to those obtained using the stochastic MCNP. Finally, a description of the particular methodology used to obtain the results in chapter 4 will be provided, outlining the various parameters used in the APTool workflow in simulating the sample cases.

3.1 Nuclear Data Implementation

Though the legacy APTool route planning model can provide accurate results in radiological units through DELFIC's internalized DCFs (see section 2.3.1), this approach condenses all of the isotopic composition and gamma-ray emission information from a particulate into a single value, thus precluding the model from using this data to inform a radiation transport model. Thus, early efforts in this work centered around allowing the user to request particulate isotopics from DELFIC's OPM (as discussed in section 2.1.2) and internalizing isotopic decay and gamma-ray emission data into APTool.

3.1.1 Isotopic Composition of HYSPLIT Dispersed Cloud

As discussed in section 2.1.2, isotopic concentration files (known as F71 files) can be requested from DELFIC as part of a problem's output. Within APTool, this functionality has been extended to produce F71 files for particulates transported using HYSPLIT's trajectory model [16], but prior to this work had not been added to the dispersion model's workflow. Production of F71 files for dispersion model output required a circuitous workflow in order to avoid making major alterations to legacy APTool features: whereas the conventional workflow dictates only a single required DELFIC CRM run, production of F71 files is dependent on user-supplied time interval data provided to HYSPLIT, and so the necessary DELFIC OPM runs can only move forward after the associated HYSPLIT run.

After the user has executed the dispersion model, APTool's DAEM initiates the process of creating F71 files by reading the dispersion model request sent to HYSPLIT. From this request, the DAEM extracts sampling time increments and spatial information, combining this data with the pre-transport parcel definitions created after the initial DELFIC CRM run to inform the parameters of this specialized DELFIC run. Because the requirements for F71 file requests in the OPM are so stringent, they are set au-

tomatically from the DAEM rather than being left to the user: a 1x1m map with only a single grid cell is requested, with the HYSPLIT time intervals used as the F71 file time intervals and each DELFIC particulate size bin treated separately. APTool then runs DELFIC in the background and retains the resulting F71 file data in memory (in addition to creating the files for the user's external use).

Isotopic data is held in a hierarchical map object, where for a given sampling time, particulate size class, and nuclide identifier or *ZAID*, defined as

$$ZAID = 10000Z + 10A + M \quad (3.1)$$

where Z is nuclide atomic number, A the atomic mass, and M the isomeric state of the nuclide (with one corresponding to metastability), the concentration of the nuclide is split into three libraries: concentration of the nuclide as an *activation product* (also referred to as the *light element* concentration), as a *fission product*, or as an *actinide*. These distinctions are useful in application-specific analyses in which only one particular production mode is of interest [7]. As concentration grid data is also stored according to sampling time and particulate size (see section 3.2.1.1 for details), the production of F71 files effectively supplies an isotopic concentration grid in packed format that is optimal for the many table lookups that will be required by the radiation transport algorithm.

3.1.2 External Dataset Integration

To internalize the ORIGEN data resources (described in section 2.6), Python scripts were written to read in the data and format it into C++-compatible LUTs. These tables are indexed by nuclide identifier (to promote synergy with the concentration data discussed above) which are mapped to arrays containing the necessary data. Though the option to model beta particulate detection was removed at a sponsor's request, the model is built with this future scope expansion in mind: the emission data table in-

cludes lines for both gamma and beta (derived from Evaluated Nuclear Structure Data File (ENSDF) data [33]) radiation, and the decay resource accounts for all relevant nuclide branching fractions.

In total, the implementation of these nuclear data LUTs imposes only 3.1 to 3.4 MB (depending on whether beta emission data is included) of data overhead, allowing for quicker data access and removing the need for users to supply the data themselves at the cost of an almost-negligible increase in memory requirements. These minimal data needs are partially a product of excluding Bremsstrahlung radiation from model consideration: due to the energy range of the detectors involved (see table 2.4) and the additional data and computational overhead its implementation would impose as a result of its dependence on medium, its inclusion was deemed unnecessary.

3.2 APTool Model Additions

As this work describes a wholly new capability within APTool, model development consisted of building in new features from the lowest levels while also integrating them into the extant APTool structure. In the most basic form, these new features could be summarized as a discrete nuclear cloud particulate model and a detailed WC-135 model, with all other new capabilities subsumed into one of these two categories. The following sections detail the functionality of each of these models, with a subsequent discussion of the concentration grid additions and radiation transport and detection algorithms that tie these models together and into APTool.

3.2.1 Nuclear Cloud Particulate Model

The practical purpose of modeling cloud particulates explicitly in APTool's route planner is to connect the output of HYSPLIT's dispersion model to the isotopic composition data produced by DELFIC. Particulate size classes specified by the user in DELFIC's event definition are followed separate from one another for the duration of the APTool

workflow; HYSPLIT treats particulates of different size classes as different pollutants in dispersing the stabilized cloud, and so DELFIC's OPM can provide isotopic composition data for each size class in isolation from all others. The persistence of particulate data over the course of the model workflow permits the route planning simulation to freely extract information from any step in the process without the concern that the information will go out of scope or lose compatibility with the rest of APTool. Thus, the model can exploit the HYSPLIT concentration grid to, for example, map particulate locations and directionality (a subject discussed at greater length in section 3.2.1.1) while at the same time using the DELFIC definitions of particulate dimensions and composition in order to generate the particulate objects that will populate that grid; the seamless integration of DELFIC and HYSPLIT was a major factor not only in making the simulation of individual cloud particulates possible, but in making the model itself feasible.

Particulates are composed of a set of isotopes, which link together a nuclide's concentration, its decay data, and the gamma-rays it emits. Aggregating this data over all of the isotopes present in a particulate, the particulate's activity, its gamma-ray energy-emission rate, and other radiological quantities can be made available to APTool for use with the concentration grid. Gamma emission rates from each isotope are also aggregated and scaled by their activity, to be used as inputs to the radiation transport algorithm described in section 3.2.3. These gamma emission rates are calculated by combining each of the data sources described above as

$$A_i = N_A \lambda_i n_i$$

$$\varphi_{\gamma,i} = \nu_{\gamma,i} A_i \quad (3.2)$$

where A_i is the activity of nuclide i , λ is its decay constant (calculated using equation 2.7), $N_A = 6.022 \times 10^{23} \text{ mol}^{-1}$ is Avogadro's number, n is the nuclide abundance in mol (obtained from F71 concentration data) summed over activation product, actinide,

and fission product sub-libraries, ϕ_γ is the gamma emission rate, and v_γ is the number of photons emitted per decay event as defined in the emission data resource.

While this isotopic data has clear relevance to the radiation transport application described here, the necessity of particulate location information for this purpose is less apparent: given the HYSPLIT concentration grid and the isotopic data available, an integral approach in which gamma emission contributions from grid voxels adjacent to the sampler's location are aggregated and weighted by path length between the voxel and the aircraft could be utilized if radiation transport is the only use for particulate information. Because the WC-135 must also collect particulates for analysis as it passes through the cloud, however, the model must also provide a mechanism for simulating sample collection. Using particulate size and location information, the model can determine what should be collected by the ERD system's filters as the sampler passes through a given voxel (a process described in section 3.2.2.2), and as this model's algorithm is based on particulate proximity to the sampler, these discrete locations are essential to the accuracy of the ERD response.

One consequence of concentrating a particulate's nuclear data within a single object is the computational cost involved: three separate lookups in different LUTs, each of which include more than a thousand nuclides, are required in constructing each particulate, and so creation of new particulate objects is avoided as much as possible within the model (the particulate construction algorithm will be detailed in section 3.2.1.1). Full integration of this route planning model into APTool should alleviate this issue, however: separation of LUTs containing gamma-ray emission data from those housing decay data is a product of the "hands-off" approach with respect to APTool legacy code, which will not be necessary when the model is fully integrated.

Particulates are also designed to have time-dependent compositions, as the F71 concentration data it uses can be updated to coincide with the time interval of the route in question, effectively serving as a fission product decay mechanism. Because

sampling time intervals used in F71 files are typically much longer than the total time spent on a route, real-time decay of particulate isotopes during flight is tied to activity, which follows the simple exponential decay scheme described in equation 2.6.

where t_i is the initial time of interest, t_{i+1} is the final time of interest, and λ is the isotope's decay constant as defined in equation 2.7. This approach has produced a model capable of fully describing particulate fractionation both prior to and over the course of a given mission.

3.2.1.1 Concentration Grid Particulate Representation

Among the most significant additions to the extant APTool model was the introduction of discrete particulates to the concentration grid. These additions also provided one of the most intractable issues encountered throughout this work: constructing a large number of particulates across multiple grid voxels on the fly in a computationally efficient manner without sacrificing model accuracy. As discussed in section 2.2.1, default grid spacing in the latitudinal and longitudinal directions is equivalent to more than 100 km in each direction; even if this spacing is reduced to 0.01° , voxels will have volumes upwards of 1 km^3 while the largest cloud particulates will have volumes on the order of μm^3 . In order to produce an accurate number of particulates for voxels with high concentration, approximately 10^{20} particulates would need to be created with distinct locations, which is infeasible for a model seeking to maximize computational efficiency.

To this end, the memory- and processing-intensive tasks of combining the data that define each particulate are performed before grid traversal is initiated: for each size class, a single particulate object is constructed to define the initial route conditions for each particulate of that size in the cloud. These “base particulates” are the only particulate objects that are permanently held in memory during the traversal process; all other particulates are generated as copies of these base particulates and then placed as needed throughout the grid. Base particulates are also the only particulate objects that

are subject to radioactive decay, greatly reducing the computational burden involved by minimizing both the number of particulates requiring F71 composition updates to more than 1,000 nuclides each at the end of every time interval and the number of exponential operations necessary to simulate real-time decay via equation 2.6.

In the context of the model, cloud particulates generally fit one of two categories: particulates caught on collection filters and those that remain external to the sampler. A model incapable of reliably simulating the amount of material collected on filters will produce inaccurate detector responses from the ERD, while insufficiently populating the exterior of the sampler with cloud particulates will lead to inaccuracies in the DGS directional spectrum. Therefore, the model must ensure not only that particulate gamma-ray intensity is commensurate with the cloud concentration in a given voxel, but also that these particulates be placed in such a way that the conditions in and around the WC-135 are consistent with a real-world scenario.

In order to develop an algorithm capable of producing particulates appropriate to a voxel's concentration, a correlation between mass and particulate concentration must first be established. Given the voxel concentration and volume, the total mass of cloud material is a straightforward calculation. Determining the mass of individual particulates within the cloud is likewise quite simple, as DELFIC models each particulate within the stabilized cloud as a sphere with a default mass density of 2.6 g cm^{-3} [1, vol. 1] (roughly equivalent to the density of dry sand [20]); while the user may provide a different particulate density in the DELFIC input, it is applied uniformly over all particulates regardless of size. Because concentration grids are specific to each particulate size class, it is reasonable to assume that the particulates within a given concentration grid will be identical; in addition, the distribution of activity between particulate size classes [20] is taken into account by DELFIC in producing the stabilized cloud [1, vol. 1], eliminating the need to include it as part of this analysis. Thus, the total number of particulates within a given voxel may be calculated as

$$N_{i,p_j} = \frac{C_{i,j}V_i}{\rho_p V_{p_j}} \quad (3.3)$$

where N_{i,p_j} is the number of particulates of size class j in voxel i , $C_{i,j}$ is the mass concentration, V_i is the volume of the voxel, V_{p_j} is the volume of a particulate in this size class, and ρ_p is the particulate density. As noted previously in this section, however, voxel sizes are orders of magnitude greater than particulate sizes in the model, and so implementing this method directly is infeasible. A more efficient approach would be to use this result to scale the intensity of gamma-rays produced from a much smaller number of particulates while ensuring that these particulates retain the same spatial distribution as the much larger sample. Within the model, this *particulate weighting factor* is a simple scalar multiplier applied to each gamma-ray intensity prior to attenuation. The user is able to control this weighting factor by setting the desired number of weighted particulates to generate within each voxel, with the weighting factor decreasing as particulates are added to the system. While the model's default is to generate 100 weighted particulates per voxel, there is no upper limit to the number of particulates that can be generated if this default value is insufficient; however, due to the requirements of the directionality algorithm described in section 3.2.1.1.2, no fewer than four particulates may be generated.

3.2.1.1.1 Collection Filter Particulates

The technique used to generate filtered particulates for the purposes of this work is significantly more efficient than the method described above, although it is also much less grounded in physical reality. The scheme used to model particulate collection is summarized in figure 3.1; the cylinder in this figure represents the path of the ERD particulate collection filter through the voxel, pictured much larger than its true scale. Note that each side of the aircraft has a collection filter in the model, doubling the




Figure 3.1: Schematic of filter collection scheme used in route planning model, where R_f is the radius of the filter (6 in.), t is the track length of the sampler through the voxel, and Δx , Δy , and Δz are the voxel dimensions.

aircraft's collection capability.

As mentioned in section 2.2.1, cloud concentration is uniform within a given voxel; the path taken through a voxel has no effect on the number of cloud particulates encountered by the sampler, as this will be determined entirely by the sampler's track length. Thus, the amount of material subject to collection by the onboard filters is a function of the sampler's voxel track length, and the total volume of the cloud the filter passes through within the voxel is directly proportional to the total number of particulates the filter will encounter; this relationship can be summarized as

$$N_{f,i} = \frac{\pi R_f^2 t}{\Delta x \Delta y \Delta z} \sum_j^n N_{i,p_j} \quad (3.4)$$

where $N_{f,i}$ is the number of particulates collected by the filter across all particulate size classes; other quantities are defined in figure 3.1 and equation 3.3. The algorithm exploits this relationship by applying the fraction of filter coverage volume within a voxel to the total number of particulates in the voxel to determine the number of particulates that will be subject to the filter collection algorithm, with all remaining particulates designated as external to the sampler. Because the onboard filters are at a fixed location on the aircraft, positions of particulates to be collected can be tied directly to the filter's position and do not have to be placed in the voxel according to the sampling method described in section 3.2.1.1.2; not only does this reduce the computational burden of determining locations for these particulates, the particulate objects themselves can also be ignored, with the relevant particulate information instead being added directly to the filter's source term (described in section 3.2.2.2) for detection purposes. By merely comparing the voxel volume covered by the filter to the total voxel volume, a significant portion of the costly particulate construction process can be removed from the problem entirely.

3.2.1.1.2 External Cloud Particulates

While the process of populating the collection filters could be streamlined, there are far fewer opportunities for improving the efficiency of constructing particulates outside the sampler; as the primary purpose of generating these particulates is to inform the directional detectors that make up the DGS system, appropriate distribution of particulate locations is prioritized over efficiency. The major challenges involved in generating adequately-distributed particulate locations consist of two inextricably linked issues: grid spatial resolution (discussed in section 2.2.1) and accounting for potential detector response contributions originating from outside the sampler's current voxel.

The generation of external particulate positions within the current voxel is primarily influenced by the dimensions of the voxel itself and the number of particulates the user

chooses to generate. The fundamental process is quite simple if voxels are of uniform dimensions: when the sampler enters the voxel of interest, pseudorandom numbers are sampled from uniform distributions in x , y , and z , and, assuming each random location is valid (i.e., none of the particulates' locations lie within the sampler's geometry and or on a voxel boundary; in these cases, new locations are sampled), particulate objects are placed at these locations. These particulates will then emit gamma-rays in the direction of the sampler's DGS detector as it traverses the voxel, updating the direction of emission as the sampler moves; which of the DGS detectors the emitted gamma-rays strike is dependent on the position of the particulate with respect to the sampler (see figure 2.25). If the sampler were to traverse the full width of the voxel through its center, the uniform conditions within the single voxel would, in theory, result in each of the four DGS detectors producing identical energy spectra and thus provide no information about the spatial distribution of activity within the cloud.

While the uniformity of the cloud over a given voxel facilitated the efficient simulation of cloud sample collection, it is detrimental to the DGS model: directional detection exploits variations in cloud concentration in order to identify and guide the sampler toward regions of high activity capable of yielding samples of the highest quality. Because this uniformity is inherent to the HYSPLIT concentration grid, the readily apparent solution to this dilemma would be to simply broaden the scope of the directional detection model beyond the single-voxel treatment to include all of the voxels immediately adjacent to the sampler's voxel. But although this approach was in fact adopted for the purposes of this work, the corresponding computational cost incurred by increasing the size of the working system to such an extent (depicted in figure 3.2) would have likely outweighed the utility it provides to the directional detection algorithm in the absence of additional applications within the model. As described in section 3.2.2.1, the inclusion of adjacent voxels in the model's working system ensures that contributions to onboard detectors' responses and dose to the aircrew take into account the location of the sampler; thus,




Figure 3.2: Increase in size of the working system within the concentration grid as a consequence of including adjacent voxels in the directional detection algorithm.

the significant expansion of the working system not only makes the directional detection algorithm more effective, it improves the quality of the simulation of all of the radiological quantities calculated for this work.

Though adding adjacent voxels to the system alleviates the issue of simulating directional detection in a voxel of uniform cloud concentration, the challenge of uniformly distributing a user-defined number of discrete particulate locations in voxels with user-defined dimensions was resolved less satisfactorily. The most significant obstacle to implementing an adequate sampling algorithm is ensuring that particulates are placed within a voxel such that its concentration is evenly distributed without the uneven dimensions producing source-detector path lengths that skew the directional detection method. One technique pursued early on in this work was to simply refrain from placing particulate objects in the grid and instead use knowledge of the cloud concentrations of voxels surrounding the sampler to “simulate” directional detection; though this method is simple, it exploits the omniscience of the model in a non-physical manner and merely

back-predicts the response in each DGS detector.

To avoid these issues, a modified uniform distribution sampling algorithm was implemented, utilizing the relative sizes of voxel dimensions provided by the user to bias the distribution: dimensions with larger grid spacings would tend to cluster particulate locations closer to the center, while those with smaller grid spacings would position them more evenly throughout the voxel. Though this biasing technique performed reasonably well if voxel dimensions were defined within a factor of two of one another, its effectiveness diminished as dimensions became more incongruous, as the sampling algorithm could not completely overcome extreme differences in source-detector path lengths. Though the model does not currently impose a restriction on the relative sizes of each voxel dimension, this would likely be the most effective technique to ensure reliable directional detection results.

It is due to these concerns over voxel sizing that the route optimization technique designed for this work has been deactivated. This method used the relative detector response among the four DGS CsI(Tl) scintillators and their respective angular “sectors” (see figure 2.25) to determine a suggested change in sampler direction. Changes in direction were implemented by updating the targeted sampler waypoint to a position that fit the suggested direction at the same sampler-waypoint distance. The optimization algorithm ignored any contributions from particulates located behind the sampler (i.e., opposite the current direction of travel) in order to avoid direction suggestions that would oscillate about a single region of high concentration, and the user is given significant flexibility over the extent and how often these suggestions influence the route they have set. Though this capability was deprecated from the final version of the model developed for this work, this utility could prove useful in the future development of sampler mission planning.

3.2.2 WC-135 Model

Encompassing the onboard detection systems and the geometric model of the aircraft itself, the WC-135 model is the core of the new route planning system. Implementation of sampling aircraft geometry required additional grid traversal considerations, as the current algorithm models the sampler as a point throughout the process. Similarly, the development of explicit models of onboard analysis systems paved the way for extensive changes to the method by which radiological grid quantities could be handled.

Dimensions used in modeling the WC-135 were obtained from publicly available sources [34], though only general information was made available; this data is summarized in figure 2.22b. Lacking a schematic or other detailed information, several assumptions were made regarding the dimensions of aircraft components: the nose and tail of the aircraft were each assumed to occupy 15% of the total aircraft length, the aircraft fuselage diameter was assumed to be 60% of the aircraft total height, the tailfin span was assumed to be $\frac{1}{8}$ as large as the wingspan, and the wing and tailfin locations were assumed to be at $\frac{1}{4}$ and $\frac{3}{4}$ the total length of the aircraft measured from the tip of the nose. As these dimensions play only a minor role in the model itself (as the ERD systems are located on the aircraft wings), these assumptions are largely arbitrary and serve primarily to ensure the model bears a passing resemblance to the aircraft itself.

3.2.2.1 WC-135 Grid Traversal

Though some additions were required, the underlying grid traversal algorithm in the legacy route planning model has been retained. The traversal method, as described by Amanatides and Woo [35], is highly efficient, requiring only two floating point comparisons and one floating point addition per grid voxel and returning the total path length necessary to cross it (see figure 3.3, where ι denotes the voxel track length). One con-




Figure 3.3: Legacy APTool simple grid traversal using ray tracing.

sequence of this simple traversal algorithm is its sole focus on voxel boundaries: no operations or comparisons are performed as the sampler crosses a voxel, and all relevant quantities are simply integrated over the sampler's path at each boundary. This method is quite effective when quantities of interest within the grid are uniform throughout a given voxel, but the non-homogeneities introduced by modeling cloud particulates explicitly (as discussed in section 3.2.3) require additional considerations.

Though a path integral approach to grid traversal would yield the most accurate result, applying such a method in the new route planning model was infeasible: the sampler moving through a given voxel will have continuously-changing path lengths from source particulates to the onboard detectors that would be computationally intensive to account for. Instead, voxel traversal has been split into four sections, small enough to ensure that the difference in source-detector path length from the beginning to the end of the section would not substantially affect the attenuation factor, but also




Figure 3.4: Grid traversal using sampler with physical model.

large enough to ensure a reasonably efficient calculation (see figure 3.4). Sample air-speed is used to determine the number of gamma-rays emitted over the course of one track length section by calculating the associated traversal time and applying it to the gamma-ray emission rate.

The other grid traversal component differentiating the new route planning model from the legacy version is consideration of the physical aircraft. At each calculation point (i.e., the end point of a track length section), the radiation transport algorithm is applied from each particulate in or around the aircraft. For those external to the aircraft, attenuation through the skin of the aircraft must be considered: though a gamma's path length through air is often much longer than within metallic fuselage material, the mean free path within the fuselage is also significantly less than within air, and so the attenuation factor for the fuselage is often non-negligible. Additionally, for gamma-rays detected by the DGS system, attenuation through the cabin of the aircraft is necessary;

as the cabin is pressurized, the density of the air inside the cabin is often substantially greater than the surrounding atmospheric air, yielding yet another potentially significant attenuation factor to consider.

In addition to consideration of the physical aircraft, the treatment of the physical cloud with respect to the concentration grid was also improved. As currently constituted, the APTool route planning model uses the current voxel's mass concentration, the time spent traversing the voxel, and the DCFs to approximate activity and exposure rates for the sampler within each voxel, as described in section 2.3.1. One of this model's core assumptions is that voxel sizes will be large enough such that at any given time, all contributions to onboard detectors or dose to personnel would originate from the current voxel only; whether this assumption is valid for a given problem, however, is contingent on the user's specifications about grid spacing and placement of sampler route waypoints. In cases where the sampler passes along the edge of a voxel, the model does not adjust to take into account contributions originating from the opposite side of the voxel boundary.

As in the exaggerated example provided in figure 3.5, this approach simultaneously overpredicts contributions from cloud particulates that are far from the sampler but within the same voxel and underpredicts contributions from nearby particulates that lie across a voxel boundary. While this would suggest that it may be possible to add only the adjacent voxel nearest to the sampler's track to the working system, this ignores cases in which the sampler passes through a voxel corner in which two adjacent voxels are nearby, as well as the possibility that a user will choose to change the sampler's direction within a voxel by setting a waypoint. Thus, the model's working system has been expanded such that the sampler's current voxel and any adjacent voxels with non-zero cloud mass concentration are included.




Figure 3.5: Loss of incident detector gamma-ray flux using a single voxel system. Activity and dose rate losses of similar magnitude.

3.2.2.2 Cloud Sample Collection

As was briefly discussed in section 3.2.1, modeling of cloud particulate collection by the WC-135 is performed explicitly by including filter objects within the aircraft object definition (sampler filter locations are shown in figure 2.22b). The purpose of filters within the context of the model is to accumulate cloud particulates on its surface and aggregate their data, which can then be used to define a planar radiation source for the transport and detection algorithms (detailed in section 3.2.3.1).

Particulate accumulation on filters is modeled as a deterministic process. Filter

capture efficiency is dependent on particulate size, as filters are ineffective at catching sub-micron particulates but nearly 100% efficient when encountering particulate sizes on the order of 10 μm . Rather than using random numbers to determine if a given particulate is caught, however, the filter efficiency is instead used to scale the intensity of gamma-ray emissions from the particulate in question. This approach allows the filter simulation to use nearly every particulate that passes through it (as long as its capture efficiency is greater than zero), improving the computational efficiency of the model in the process. The filter also has a defined maximum capacity (defined in terms of the number of particulates present on the filter and their respective sizes) that, when reached, forces the filter to capture every particulate that contacts it at 100% efficiency, as the filter has become entirely opaque to particulates of all sizes.

Particulates that are successfully caught by the filter merge their gamma-ray data in order to streamline the radiation transport calculation. While data from particulates that remain external to aircraft can be thrown out when the aircraft moves away from the voxel in question, data from particulates caught on filters must be retained in memory for as long as the filters remain in place in order to contribute to the ERD response. To maximize computational efficiency, the model treats the filtered particulates as a spatially-homogeneous disc source emitting the gamma lines of each of its constituent particulates. This allows the particulate objects themselves to be purged from memory once collected, with only their component isotopes and particulate size classes kept in memory to facilitate radioactive decay of filtered material (following the same procedure as described in section 3.2.1). This approach significantly reduces the computational burden of detecting collected samples by replacing distributed radiation sources of varying source-detector solid angle with a single source of unchanging geometry.

Though filter models continue to accumulate particulates even after reaching maximum capacity, this behavior is undesirable in a real sampling scenario. To avoid this, the ERD systems each begin with 12 clean filters per mission that can be swapped

in and out on the fly. Because the details of the filter swapping protocol remain unknown to the public, the model currently switches out filters every 30 min of flight time regardless of the amount of cloud material collected. The method by which dirty filters are stored on the aircraft is likewise unclear: as shown in figure 2.22b, the model is designed to simply store all filters not currently in use by the ERD system in a single container. While this approach is highly unlikely to be used on real-world samplers due to the possibility of cross-contamination between clean and dirty filters, the modeled system is concerned only with the dirty filters in storage and thus any other clean filter storage can be safely ignored (except to track how many filters the ERD system has used in a given mission).

The practice of keeping dirty filters on board real-world aircraft has its roots in the methodology of the earliest cloud sampling missions during US nuclear tests in the 1950s and 1960s [36], in which these filters would be taken away for radiochemical analysis after mission completion (and which still has value even in an age when real-time analysis is possible); though this post-mission analysis lies outside the scope of this work, modeled filters remain in memory even after being swapped out because the samples they have collected still contribute to aircraft analysis system responses. In the context of the model, stored filter data is merged together in a manner similar to the merging of filtered particulate data: stored filters are represented as a disc radiation source comprising the accumulation of all gamma lines from all particulates on all stored filters, though in this case the disc source is surrounded by a container that provides some measure of gamma-ray shielding (the exact properties of the filter storage container are unknown, but assumed to be equivalent to $\frac{1}{4}$ in. of aluminum for the purposes of this work).

3.2.3 Radiation Transport and Detection

Though the particulate and WC-135 models are the most significant objects added to the APTool route planner, the core of the new design described in this work is the radiation physics treatment. The basic functional characteristics of the transport and detection algorithms are outlined in sections 2.5.1 and 2.5.2, but additional considerations were required in order to account for differences in source-detector geometry, nuances within the concentration grid, and to improve computational efficiency. In order to fully capture these distinctions, a general description of the transport and detection algorithms used to simulate the responses of the CsI(Tl) scintillators and G-M counters will be followed by discussion of the specific methods used to model the behavior of the ERD and DGS systems and a brief note regarding adjustments made to the calculation of dose and activity encountered by the sampler during flight.

3.2.3.1 Source-Detector Radiation Transport

Radiation transport from source particulates to the ERD and DGS systems through the concentration grid is controlled via the application of equation 2.12, but special considerations must be made in each of these cases. The simplest of these is the transport of gamma-rays from cloud samples collected on wing-mounted particulate filters to the ERD system detectors (the source-detector geometry is shown in figure 2.23).

Consisting of a disc radiation source and rectangular or cylindrical detector (depending on whether the CsI(Tl) or G-M detector is being considered) separated by a fixed distance that does not cross a voxel boundary in the concentration grid, this source-detector problem can be solved using equation 2.11 with only minor alterations. In particular, while equation 2.11 assumes an isotropic point radiation source and a point detector (with *geometry factor* $\frac{1}{4\pi d^2}$, where d is source-detector path length), the geometry of this problem is slightly more complex.

Because the distance between the collection filter and the ERD system detectors does not change throughout the problem, determining an analytical solution for the solid angle of this geometry was deemed unnecessary. Instead, the solid angle was computed via a simple stochastic simulation using MCNP: using a disc source emitting gamma-rays isotropically and a rectangular detector, a simple surface flux computation accounting for the surface area of the detector face was carried out (problem geometry was defined as in figure 2.23). While the geometry factor for a point source-point detector problem with the same separation distance as between the collection filter and the ERD detectors would suggest roughly one in 30,000 source gamma-rays would strike the detector, the disc source-rectangular (or cylindrical) detector geometry suggests an interaction approximately every 1,000 gamma-rays instead. Once the problem's geometry factor has been accounted for, the modified form of equation 2.11 may be applied, taking into consideration the density of atmospheric air at the sampler's altitude as well as the energy of each gamma-ray in order to determine the attenuation coefficient μ .

Computing radiation transport from particulates external to the sampler to the DGS system is slightly more complex (the scheme is depicted in simplified form in figure 3.6). The first step involved determining which DGS detector the particulate's gamma-rays would interact with, which involved simply calculating the path length between the source and the center of each detector, with the closest detector selected. To avoid the computationally-expensive process of computing geometry factors for each particulate emitting gamma-rays toward DGS detectors, the geometry factor for a particulate more than 10 m (approximately 10 times the length of the DGS detectors plus the distance from the DGS system to the exterior of the aircraft) from the closest DGS detector was assumed to be the same as that of a point source and a point detector at the same distance. Within that range, however, geometry factors were again determined using a stochastic approach: MCNP input files were created to sample source particulate locations radially about a DGS detector at specified angles and at varying distances, with

surface flux calculations used to determine the geometry factor for a particular case. These values were then placed into a LUT, where the DGS detector in question would inform the chosen angle (according to the formulation as defined in figure 2.25) and the particulate distance would be used as the table's interpolation factor.




Figure 3.6: Schematic of radiation transport from external particulate to DGS system detector (not to scale). Dashed box indicates boundary between concentration grid altitude layers.

As is shown in figure 3.6, additional attenuation terms were also required in order to apply the grid-dependent equation 2.12. The first step in this process was to determine whether a particulate lay in a different grid altitude layer than the sampler; if so, the path length between the particulate and the voxel boundary and the atmospheric density at that altitude were computed first (particulates in different voxels at the same atmospheric layer as the sampler could skip this computation). Then, attenuation between either the particulate itself (if no other grid layers were involved) or the voxel

boundary was computed using the density of atmospheric air at the sampler's layer. Attenuation through the aircraft skin followed; because the composition of the WC-135 skin was not provided, the model currently assumes that the material used has properties similar to aluminum and is approximately $\frac{1}{2}$ in. thick. While this thickness likely exceeds the true aircraft skin thickness, the assumption of an aluminum fuselage permits the computation to combine the attenuation calculation for the fuselage with that of the housing surrounding the DGS system (not shown in figure 3.6); these computations will be split if the composition of the WC-135 is clarified in the future. Finally, attenuation between the wall of the aircraft and the DGS detector itself is computed, using the path length between wall and detector as well as the density of air in the pressurized cabin (assumed to be approximately that of sea level air) to calculate this final attenuation factor.

The exponential terms for these attenuation factors (i.e., cross-section multiplied by path length) are computed prior to applying the exponentials themselves, as the multiplicative nature of exponential attenuation allows for these terms to be summed before exponentiation (as in equation 2.12). As with the ERD computations, this process is repeated for each gamma-ray produced by each particulate before the transport algorithm is completed.

3.2.3.2 Radiation Detection Algorithm

As mentioned in section 2.5.2, the necessity of a detailed scintillation detection model inherently subsumed a significant quantity of the methods required for accurate G-M counter simulation; thus, a description of the general detection algorithm common to both detector types will be followed by a brief summary of the complexities specific to CsI(Tl) and their implementation. Each detector object is defined in terms of its location, dimensions, effective energy range, energy resolution (only of use for CsI(Tl)), and detection efficiency; this information is summarized in table 2.4 for the G-M and CsI(Tl)

detectors.

Detector locations are set by the aircraft model itself according to the layout shown in figure 2.22b; fully describing the ERD and DGS systems requires a total of eight detectors, four of which are the DGS CsI(Tl) directional detectors and four of which are the ERD CsI(Tl) detectors and G-M counters (one pair each per ERD system). Specifying detector dimensions is more speculative: while DGS detector sizes have been provided, the dimensions given in figure 2.23 were inferred from sponsor-provided images in the case of G-M counters or simply assumed to be equivalent in size to those in the DGS system in the case of ERD scintillators.

Similarly, effective energy ranges for detector models were not provided for all detector types, but in this case the ambiguity had little effect on the model. Both the ERD and DGS systems have been designed with default energy windows extending from 50 keV to 3 MeV [37], comfortably within the effective energy ranges of CsI(Tl) [38] and G-M [39] detectors. However, a detector's ability to provide a response does not guarantee that each incident gamma-ray will produce a response, as the *intrinsic efficiency* of each detector depends on the detector type and the energy of the incident gamma-ray. In general terms, intrinsic efficiency is defined as

$$\epsilon_{\text{int}} = \frac{C}{\phi_{\gamma}} \quad (3.5)$$

where C is the total number of counts recorded and ϕ_{γ} is the number of gamma-rays incident on a detector. Intrinsic efficiency is determined primarily by the probability that a gamma-ray of a given energy will interact within the detector volume and, given the volume is large enough, that these interactions will result in the gamma-ray being absorbed; the absorption efficiencies for the detectors used in this work are provided in figure 3.7 and are included in the model as LUTs.

In the context of the model, a detector's intrinsic efficiency is simulated in a manner similar to that of the filter capture efficiency described in section 3.2.2.2: as a




Figure 3.7: Gamma-ray absorption efficiencies for Csl(Tl) and G-M detectors over the default energy window for the ERD and DGS systems. Csl(Tl) efficiencies adapted from [40, figure 9]; G-M counter efficiencies adapted from [39, chapter 5, figure 11].

scalar multiplier of gamma-ray intensity based on incident energy. Thus each impinging gamma-ray can potentially contribute to the detector response, resulting in a more computationally efficient simulation.

After a gamma-ray has struck the detector, the interaction process is quite simple. Gamma-rays entering a Csl(Tl) will be subject to the interaction procedures described in section 2.5.2.1 (the algorithm for which will be detailed in section 3.2.3.2.1), while those interacting with a G-M will not (gamma-rays in scintillators will also be subject to Gaussian energy sampling, also described in section 3.2.3.2.1). After this process, gamma-rays (or secondary particles created by gamma-rays) of energy outside the allowable window (i.e., the ERD and DGS system default energy window) are thrown out immediately.

All gamma-rays of appropriate energy are then subject to the system's detection threshold, used to filter out all gamma-rays of intensity insufficient for the production of

pulses in the detection system. This threshold is necessary in a deterministic system based around exponential gamma-ray attenuation (see equation 2.11) as the exponential term ensures that some fraction of the emitted gamma-rays will reach the detector even in the most extreme scenarios; thus, the model must define a lower limit of detectability so as to only count the gamma-rays that would realistically reach the detector. Ideally, this threshold would be associated with the sensitivity of the detection system as a whole in order to align model behavior with real-world operation; as the sponsor has not released this information, the detection threshold is left as an adjustable parameter within the model. By default, the model will reject gamma-ray intensities less than $10^{-10} \gamma s^{-1}$; as this corresponds to approximately one gamma-ray event every 300 yr, this threshold is quite conservative, and will likely be adjusted when the model has concluded its proof-of-concept stage (and when the necessary information is received) to better fit the onboard detection systems it is simulating.

Gamma-rays with intensities above the detection threshold will register a number of counts in the detector corresponding to their intensities. Due to the design of the grid traversal scheme in the context of this work (described in section 3.2.2.1), gamma-ray intensities must be converted from γs^{-1} to γ by multiplying the intensity by the traversal time since the last point where detector counts were registered (i.e., one of the four points along the aircraft's track within a voxel). After counts have been registered, a series of quality checks are initiated to ensure that detector data integrity is maintained; the most notable of these checks looks for evidence of underflow errors, in which gamma-rays have been attenuated to an extent such that their intensity is less than the minimum 64-bit double-precision value in the C standard programming library (approximately 2.23×10^{-308}). Behavior when surpassing this value varies by system, but the most common outcome is the system rolling it over to the maximum 64-bit number (approximately 1.80×10^{308}), an implausible number of counts for any realistic detection scenario; thus, a check is in place to ensure that these underflows

are caught and thrown out before being recorded.

3.2.3.2.1 Csl(Tl) Detection Algorithm

While the algorithm governing G-M counter simulation is largely similar to the general discussion in section 3.2.3.2, the Csl(Tl) system is significantly more complex, primarily due to the modes of operation for the two detector types: while the G-M counter is used only to detect the presence of gamma-rays, scintillators produce light pulses of amplitude proportional to the energy of the originating gamma-ray, and thus can provide identifying information about gamma-emitting nuclides in their vicinity. The primary modes of interaction between gamma-rays and scintillators are described in section 2.5.2.1 and their energy dependence is outlined in figure 2.26. While this plot can provide some general information on the energy ranges over which each interaction type is dominant, however, the interaction probabilities for each at a given incident gamma-ray energy are functions of the scintillator's composition and location of interaction inside the crystal. Therefore, the creation of an entirely precise LUT from which the model could determine the most likely interaction type (as has been possible with the particulate collection and detector intrinsic efficiency algorithms described in this section) is infeasible; and while a deterministic approach to mutually-exclusive outcomes has governed much of the model to this point, using this technique in the context of interactions with highly-divergent outcomes risks pulling the simulation into non-physical behavior (e.g., extreme event over- or under-counting).

To avoid this, a new technique was introduced in order to provide a reasonably accurate simulation of scintillator energy deposition. Taking the *linear attenuation coefficient* μ_ℓ defined in terms of the three principal interaction modes' probabilities [26]

$$\mu_\ell(E) = \sigma(E) + \tau(E) + \kappa(E) \quad (3.6)$$




Figure 3.8: CsI(Tl) interaction probabilities as a function of gamma-ray energy. Note iodine K-edge feature located at approximately 33 keV; the default energy window is designed to filter out any contributions due to this phenomenon.

where $\sigma(E)$ is the photoelectric absorption probability at gamma-ray energy E , $\tau(E)$ is the Compton scattering probability, and $\kappa(E)$ is the pair production probability; the variation of these quantities with gamma-ray energy is shown in figure 3.8 (note that the attenuation coefficient provided in equation 2.11 is that of atmospheric air, not CsI(Tl)). Using the total attenuation coefficient, the relative probabilities of each interaction mode can be determined, and these serve as the basis of the energy deposition technique.

Gamma-rays reaching the detector with intensities greater than the detection threshold are passed through the energy deposition algorithm, which will determine the gamma-ray energy that the CsI(Tl) detector “sees”. To ensure that each gamma-ray is only subject to one interaction at a time (in contrast to the superposition of interactions in the filter collection and intrinsic efficiency algorithms), a pseudorandom number is chosen between zero and one and compared to the probabilities of each interaction mode

at the gamma-ray energy; the more probable the interaction, the greater the likelihood that the pseudorandom number will land in its range. Once an interaction has been chosen, the interaction process proceeds as described in section 2.5.2, though in each case any secondary or scattered gamma-rays are followed to ensure that the original gamma-ray's energy is accounted for from the time it enters the detector until it either deposits its full energy or escapes the detector. Detector escape is assumed to occur when a gamma-ray has traveled a sufficient number of mean free paths such that its total path length exceeds the smallest detector dimension. Mean free path is defined as [26]

$$MFP = \frac{\int_0^\infty x e^{-\mu_\ell x} dx}{\int_0^\infty e^{-\mu_\ell x} dx} = \frac{1}{\mu_\ell} \quad (3.7)$$

where x is the distance traveled through a medium with linear attenuation coefficient μ_ℓ . The interaction process for photoelectric absorption is simple: the gamma-ray deposits its full energy in the detector and the detection algorithm resumes. In Compton scattering and pair production, however, production of secondaries must be considered. For Compton scattering, the cosine of the scattering angle θ is selected via another pseudorandom number ranging from -1 to 1 ; the application of the Klein-Nishina formula [41] to determine the scattering angle based on gamma-ray energy was deemed unnecessary, particularly due to the computational expense it imposes. The outgoing electron energy is then determined using equation 2.13, with the remainder left to the scattered gamma-ray; this gamma-ray is then able to interact with the detector at its new energy, provided it does not escape the detector. In scenarios in which the gamma-ray energy exceeds 1.022 MeV , pair production must be considered, presenting an even more complex array of interaction options. The original gamma-ray disappears and is replaced by an electron-positron pair, with the electron depositing its energy and the positron moving a short distance before annihilating into two 511 keV gamma-rays, which can then interact with the detector via either photoelectric absorption or Comp-

ton scattering, or some combination of detector escape and further interaction. Thus, depending on the energies in question, it is possible for a single gamma-ray to produce several electrons within the detector volume, each of which must be accounted for by the detector when the energy deposition algorithm is completed.

Though the energy deposition algorithm will provide the necessary electron energies that produce detector pulses, the detection algorithm must also account for the energy resolution of the CsI(Tl) system. Because pulses are recorded in a detection system via the migration of a discrete number of charge carriers (in this case, electrons), a system in which the number of charge carriers created by gamma-rays of a given incident energy is inconsistent and will register the same deposited gamma-ray energy as slightly higher or lower energies in some cases. Thus, the concept of energy resolution has been established to describe the range of attributed energies about the true deposited energy that a particular scintillator will produce. Mathematically, energy resolution R is defined as

$$R = \frac{FWHM}{E_0} \quad (3.8)$$

where full width at half-maximum (FWHM) is the width of the distribution of energies about the true deposited energy E_0 at half its peak height. These parameters are illustrated in figure 3.9. As the statistical fluctuation in the number of charge carriers is the primary contributing factor to the detector's resolution, its response to any discrete gamma-ray energy will have a Gaussian distribution about the true value E_0 with a standard deviation of σ . The standard deviation of the Gaussian is related to the FWHM via the relation

$$FWHM = 2.35\sigma \quad (3.9)$$




Figure 3.9: Factors contributing to detector energy resolution, with σ defined as the standard deviation of the energy distribution. Adapted from [42, figure 4.5].

and the Gaussian distribution is defined as

$$f(E) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{E-E_0}{\sigma}\right)^2} \quad (3.10)$$

Within the model, the Gaussian standard deviation varies with E_0 , but can be determined at the energy of interest via applying fitting functions to experimental spectra; using this procedure, σ was calculated for energies within the CsI(Tl) detection window. For each electron produced by incoming gamma-rays, a combination of the relevant σ value and equations 3.8 and 3.9 can be used to obtain the appropriate energy resolution value. The Gaussian distribution is then constructed using equation 3.10 with E_0 as the electron energy and a sample E is drawn from it; this sampled energy is then


used to populate its corresponding energy bin. Each CsI(Tl) detector's energy window is split into 1,024 channels with uniform widths of approximately 2.88 keV each. The number of counts associated with the sampled energy are added to its corresponding bin, after which the quality checks described in section 3.2.3.2 are performed and the algorithm completes.

Finally, it is important to note that although figure 2.27 is useful in identifying some common features of gamma-ray detector responses, such monoenergetic energy spectra are rare when analyzing nuclear cloud samples; unstable nuclides in a sample produce gamma-rays that can vary in energy from the sub-keV range to the tens of MeV, all of which can be emitted simultaneously and are subject to all of the interaction mechanisms described above. As a result, the single clean energy spectrum produced using a monoenergetic source is replaced by many overlapping and inseparable energy spectra that must be analyzed closely in order to identify the discrete gamma-ray energies that produced it.


However, the intractable and tedious process of identifying and analyzing each peak in a detector response is eased somewhat in a fission product context. The probability distribution of fission product nuclides is well known (see figure 2.12) with several prominent gamma-ray emitters among them; a selection of these is provided in figure 3.10, along with peaks produced due to secondary interactions within the detection system. In a detector with sufficient energy resolution, these peaks may be uniquely identified by isotope, leaving only a few gamma-rays to be identified by other means.

3.2.3.2.2 Sampler Dose and Activity Calculation

Though the implementation of a dedicated radiation transport algorithm opened the possibility for a finely-detailed dose and activity calculation method to replace that of the legacy APTool route planning model, it has instead been left largely intact in this work. Early efforts to utilize particulates as the basis for these computations yielded results



(a) CsI(Tl) detector response at 36 h post-event within WC-135 default energy window; some notable fission-product nuclides with $E_\gamma > 1$ MeV marked.



(b) Detector response focusing between 0 MeV to 1 MeV with some notable fission-product nuclides and corresponding E_γ marked.

Figure 3.10: CsI(Tl) detector response to fission product gamma-rays with some prominent E_γ included.

that differed only slightly from those obtained using the legacy method (attributable to the uniformity of the concentration grid as well as both DCF and particulate weight dependence on concentration), but at the expense of significantly longer computation times.

Thus, modifications to the dose and activity computations were minor. For the activity computation, the most significant change was to include the activity of sampled material on collection filters, whether during sample collection or after they are stored in the aircraft for additional analysis. Including this activity was a straightforward process of merely summing each filter's aggregate particle activity and combining this value with the activity of external material encountered by the aircraft at each time step; however, this filter activity is not subject to the $t^{-1.2}$ decay correction applied to external activity (see section 2.3.1 for details), as the time decay of filtered material is performed explicitly via equation 2.6.

Changes to dose calculations were likewise minor, though slightly more extensive than those made to activity computations. As was briefly discussed in section 3.2.2.1, dose due to cloud material in adjacent voxels was introduced to the model, with dose contributions from these voxels attenuated according to the sampler's distance from its shared boundary. As these path lengths were pre-computed in order to inform the radiation transport algorithm (see section 3.2.3.1), this addition had little impact on the computational efficiency of the model and brought the dose calculation into line with the modified grid traversal technique introduced by this work.

The dose calculation was also modified to implement a shielding factor for the aircraft, which is meant to account for the reduction in dose to personnel as a result of gamma-ray attenuation through the WC-135's fuselage. In effect, this shielding factor is the attenuation factor of the aircraft's fuselage material averaged over all incident gamma-ray energies and is used as a blanket reduction factor for the dose encountered on a mission. Though these changes to activity and dose calculations are in large

part simply necessary byproducts of the more extensive changes to the APTool route planning model, the cumulative effect of improvements to DELFIC DCFs, increased computation fidelity due to the use of shorter grid traversal steps, and these calculation adjustments have augmented the simulation of these quantities while providing little additional computational overhead.

3.3 MCNP Verification Technique

As with any model centered around simulating nuclear event effects, supporting field experiments could not be performed in order to verify simulation results, and although the US military has published reports containing airborne sampling data of tests events dating back to Operation Crossroads in 1946 [36], many of the sampling parameters central to this work (e.g., the particular route taken through the cloud, airspeeds, detector responses) were not available within these datasets. Therefore, the task of verifying model results was undertaken using the established nuclear transport software MCNP.

MCNP employs a stochastic transport algorithm, following individual radiation particles from their creation at a source through whatever interactions they may have with the media surrounding them until they are absorbed or otherwise destroyed. Particle interactions (including the conditions of the particle's creation at the source) are weighted by relative probabilities with outcomes determined via random number; in essence, each particle history represents a single pass through one path of one particle's decision tree, and so this process must be repeated many times such that the phase space consisting of each particle's decision tree is fully sampled. This differs significantly from deterministic transport, which avoids simulating interactions on an individual basis, considers the superposition of probabilistic events rather than each individually, and is typically employed in simulating the transport of many individual particles at once (for example, gamma-rays of the same energy emitted from a particulate, as in this work). In short, MCNP provides simulations that are more faithful to the real-world physics

involved in a given problem, but are often much more computationally-intensive than deterministic simulations due to the number of particle histories required to fill out the phase space. For the purposes of this work, comparing against MCNP results will provide crucial information regarding the validity of the assumptions made in developing the radiation transport and detection algorithms that form the core of this model.

As MCNP is designed to model radiation transport in a static geometry, verifying the results of a sampler's entire route in a single simulation would be infeasible. Instead, several individual time slices of a sampler's route were replicated within MCNP to assess the performance of the radiation transport and detection algorithms in isolation from the rest of the model. In order to do this, several additional C++ modules were created to run alongside the route planning model, extracting the working system of the sampler within the concentration grid and recreating it within MCNP, where the responses from each detector in the ERD and DGS could be simulated and compared to the route planning model detector response simulation for that same point in time.

Though translating the components of the route planning model to MCNP on an individual basis is quite simple, automating this process was significantly more complex. Each MCNP simulation required a few basic geometric components: the sampler itself, its detectors, the collection filters, the particulates external to the sampler, and the atmospheric air surrounding them (MCNP runs included only what was present in the sampler's current voxel and those adjacent to it). Positions for each of these objects were set relative to the sampler itself, which was set at the center of the problem geometry. These objects were assigned materials reflecting their true isotopic composition and mass density, with the composition of each particulate set according to its F71 data and atmospheric composition and density determined as functions of the sampler's altitude.

Definition of the radiation source term involved specifying both the physical properties (e.g., size, location, composition) of each cloud particulate relevant to the sampler

at a given time as well as its list of gamma-ray energies and emission probabilities; collectively, particulates are defined as a distributed volumetric source. Collection filters were also included as disc sources that were likewise added to the distributed source definition. In order to avoid tracking particles emitted away from any onboard detection systems, source angles of emission were biased toward these detectors. Source position for a given event history was determined through random sampling of the particulate and filter objects; with many different source terms and different energy spectra associated with them, fully sampling the phase space of this problem required many event histories to be executed.

While the user can manually input the number of particle histories to follow, the user may also choose to define a stop condition based on the relative uncertainty of an output quantity bin (e.g., when the uncertainty of the total energy bin for a particular tally reaches 5%), or define multiple stop conditions that end the run when one of them is met. For the purposes of this work, the option to use multiple conditions was chosen: if the relative uncertainty of the total energy bin for any DGS detector pulse-height tally fell below 0.5 % or the number of particle histories exceeded 1×10^8 , the run was concluded. The use of DGS detector responses for this condition (rather than an ERD detector) was due to the significantly longer source-detector path length between external particulates and the DGS system than for the collection filters and the ERD system, and a cumulative uncertainty value of 0.5 % was sufficient to ensure that every constituent energy bin for the pulse-height tally had a relative uncertainty below 10 % (the uncertainty threshold defined by MCNP developers as adequate for this tally type [43]). The second condition defining a maximum of 1×10^8 particle histories was determined via tests of the verification model, where even in the most extreme source-detector path length cases, the uncertainty condition would be met before exceeding this value; thus, this condition serves to catch cases in which the MCNP geometry or source conditions were incorrectly defined, as without this condition MCNP would run

indefinitely.

Finally, MCNP requires users to specify what results to produce at the end of a run as part of their input file. For the purposes of this work, quantities of interest included the pulse-height response in the CsI(Tl) detectors (to be compared to the route planning model's simulated detector response) and the gamma-ray flux both at the detectors and at the interface between the sampler body and the atmosphere. Simulating the flux at the outside of the sampler provided a simple measure of the total amount of radiation entering the aircraft's cabin, to be used to ensure the scenario being simulated within MCNP matches up to the corresponding route planning model run.

3.4 Sample Case Parameters

In order to assess the capabilities of the new route planning model, a set of four historical US test events (outlined in table 3.1) were selected to represent a wide range of event yields while producing clouds in which a large proportion of the radioactive material concentration would lie below the WC-135's ceiling during the time period of interest; as a result, this analysis will focus exclusively on events classified as surface bursts by DELFIC. All cases were first run using the legacy APTool route planning model (which also utilized the DELFIC executable containing legacy DCFs; see section 2.3.1) in order to compare the resulting dose and activity data to that produced using the updated route planning model.

Aside from the event definition parameters specific to each case, APTool input was standardized across these cases to ensure consistent results for each. 50 particulate size classes were used in each case's DELFIC CRM run, and the meteorological data used by both DELFIC and HYSPLIT was taken from the National Centers for Environmental Prediction (NCEP)/National Center for Atmospheric Research (NCAR) reanalysis archive [44].

HYSPLIT dispersion model parameters were tailored specifically for route planning

Table 3.1: US surface and atmospheric test events selected to assess the capabilities of the improved route planning model. $N=4$ events.

Operation	Event	Event Yield (kt) [*]	Location	Event Classification
Tumbler-Snapper	George	1.50×10^1	NNSS	Surface
Greenhouse	George	2.25×10^2	Enewetak Atoll ^{**}	Surface
Hardtack I	Elder	8.80×10^2	Enewetak Atoll ^{**}	Surface
Redwing	Tewa	5×10^3	Bikini Atoll ^{**}	Surface

^{*} Data extracted from DOE/NV-209 [15]

^{**} Event occurred over water

purposes: concentration grid voxels were spaced 1 km apart in each direction (approximately 0.01° in latitude and longitude), though an additional altitude layer was set at 100 m so as to separate material that had been deposited on the ground from that which remained airborne, and the altitude grid was cut off above the WC-135's ceiling of 12 km; sampling time intervals were set at 12 h and the total dispersion runtime was set to 96 h or 4 d. This relatively short runtime is a product of compromise between the size of the concentration grid and the desired time scale for the model: to achieve reasonable HYSPLIT dispersion model run times, the concentration grid was defined with 400 voxels in latitude and longitude, amounting to only a $4^\circ \times 4^\circ$ grid centered on the event GZ that would contain no airborne debris after a few days of dispersion, and thus a shorter runtime was selected. All other dispersion model inputs were set to default values.

Inputs for the routes the sampler would fly for each event necessarily varied between events in most cases, but the criteria used to select route waypoints were kept consistent to the extent possible. Three separate routes were flown for each event: the first after 12 h of dispersion, the other two varying by event. In each case, the third route would be flown at the latest time interval for which any cloud concentration could be found above ground level, with the second route flown at a time interval approximately halfway between the first and third routes. In the lowest-yield case, all cloud material

had deposited by the third day post-event, and thus the routes were separated by only a day; in the highest-yield case, cloud material remained airborne for the entirety of the dispersion runtime, and thus the routes were set for 1, 2, and 4 d post-event.

Positional parameters for the waypoint locations within these routes were selected using a similar philosophy. Each route consisted of three waypoints: one in the highest altitude layer for which cloud material could be found, one in a median altitude layer, and one in the lowest altitude layer for which cloud material could be found; in some cases, this methodology permitted the entirety of the cloud's height to be sampled during the mission, while in other cases the sampler only traversed three or four layers (the latter cases primarily occurring at the very end of the dispersion runtime). In each case, latitude and longitude were selected based on the region of highest concentration in the layer of interest. In addition, the speed of the sampler was varied by waypoint: at the highest altitude, the sampler airspeed was set to half its maximum value (202 mph or 325 km h^{-1}), increased to 300 mph (482 km h^{-1}), and then set to its maximum airspeed of 403 mph (648 km h^{-1}) for the final waypoint. This variation in flight speed allowed the sampler a small amount of extra time per voxel in the uppermost portions of the cloud where responses were expected to be less intense than the more densely-concentrated cloud regions near to the ground.

CHAPTER 4

RESULTS

As described in chapter 1, developing a flexible model design capable of fitting into future sponsor applications resulted in a primary focus on the functionality of the model at the expense of producing results specific to a single application. This study reflects this philosophy, centering on the description of the model's capabilities alongside sample results demonstrating its functionality (the absence of real-world experimental data from WC-135 sampling missions was an additional factor in this approach). Thus, the following sections will focus on model performance in comparison to its legacy counterpart and verification of detector response results obtained using MCNP.

In order to assess the level of improvement between the legacy APTool route planning model and this work, verifying that the radiological quantities simulated by the legacy model could be brought into better alignment with real-world expectations using the refinements described in chapter 3 was of utmost importance. Quantifying the performance of the additions to the legacy model, however, will necessarily remove the legacy model from the analysis and focus instead on the detector response results from the detailed model developed for this work in comparison to those simulated using MCNP. These analyses will also consider how the performance of the model varies as a result of changing event parameters and suggest future model improvements.

4.1 Legacy APTool Capability Improvements

Though the scope of this work did not emphasize improving the simulation capability of the route planner's radiological quantities (as detailed in section 3.2.3.2.2), the fact that this work constituted an entire overhaul of the model was expected to produce




Figure 4.1: Sampler route and mission accumulated dose comparison between legacy and updated models for Redwing Tewa.

non-negligible differences in results, while also providing a direct method by which the capabilities of each model can be quantifiably compared. Therefore, each of the sample test events outlined in table 3.1 were simulated using identical parameters in both legacy APTool and an APTool build utilizing the new route planning model, producing in each case plots of accumulated dose, instantaneous dose rate, activity, and accumulated activity over route distance.

Overall, dose accumulation results are not significantly affected by the changes made to the dose calculation method; while including dose contributions from adjacent voxels produces a noticeable increase in dose received over the legacy model (as shown in figure 4.1), the additional features of the dose calculation implemented in the updated model do not introduce new mechanisms of conveying dose, and thus the shape of the dose profile remains mostly intact.

Though dose results were largely unaffected by new model capabilities, the rate of change of dose with time is subject to the modified grid traversal algorithm that emphasizes shorter stride lengths through grid voxels. Because dose rate is calculated




Figure 4.2: Dose rate time dependence comparison between legacy and updated models for Hardtack I Elder. $t^{-1.2}$ decay curve provided as reference.

at the end of each step via equation 2.9, this value is updated more frequently than in the legacy model. While this modification has only a subtle impact on simulated dose rate with respect to the legacy model (as shown in figure 4.2), the change in dose rate over time follows a trajectory independent of the legacy model, while also providing a somewhat closer fit to $t^{-1.2}$ decay, particularly within a day of the event time.

Despite the expectation that the updated model would have a more pronounced effect on activity results than with dose, activity instead proved to be similarly resilient, as shown in figure 4.3. This is particularly notable given the inclusion of the debris present on stored collection filters in the analysis, which produced a non-negligible but largely insignificant increase in total sampler activity over the legacy model; intuitively, the on-board storage of concentrated radioactive cloud material would be expected to present a significant hazard. That this result is absent is a product of real-world sampler protocols dictating model performance: as discussed in section 2.4.1, filters are currently changed every 30 min while in the cloud, and thus very brief routes will see no change in activity. However, even for longer routes in which filters do contribute to onboard activity, their storage container provides substantial shielding in the interest of aircrew




Figure 4.3: Sampler route and mission accumulated activity comparison between legacy and updated models for Greenhouse George.

safety, which has the ultimate effect of suppressing (but not eliminating) contributions from stored filters.

Though updated model radiological results were not expected to change dramatically from those produced by the legacy APTool route planning model, the modifications implemented to improve the quality of these simulations did little to affect the end results. The introduction of previously-excluded sources of dose and activity to the model proved to have the most significant impact on results, though their effect was purely additive; without structural changes regarding the dose and activity calculations themselves, large discrepancies in radiological results between models are unlikely, even in extreme cases. While this suggests that the updated model succeeded in ensuring that simulations are reflective of real-world results, much of the work directed at this objective played little part in achieving it.


4.2 WC-135 Onboard Detector Responses

As the fundamental mission of the WC-135 and other sampling aircraft involves both the collection and detection of radioactive cloud material, much of the utility of the model developed for this work lies in its detector response simulation capability; indeed, nearly all of the functional parts added to the model for the purposes of this work (and in the case of the automated MCNP verification procedures, the development of an entirely self-contained software package) were implemented in support of this objective.


For each voxel point in which the WC-135 encountered detectable radiation, each of the eight detectors (two ERD CsI(Tl) scintillators, two ERD G-M counters, four DGS CsI(Tl) scintillators) registering counts would dump its response along with a time tag (in order to sync these responses with the corresponding MCNP verification run), resulting in full mission histories for each detector that can be used to improve mission quality in a variety of ways; for example, in figure 4.4a, both the second and third routes show noticeable gaps in the response, indicating suboptimal waypoint placement by the planner. While the legacy model would indicate that no dose or activity was registered during this portion of the mission, the time-dependent detector response results provide

These detector response histories also provide a way for mission planners to investigate the qualities of an individual route by tracking fluctuations in response and notifying them of any points of “dead air” over the course of the route, such as in figure 4.5. By continuously dumping this data to the user over the course of a mission, the planner is given total control over the level of scrutiny they apply on a given mission plan.

As was also briefly described in section 3.2.3.2.1, the ability to determine a gamma-ray’s originating isotope is a fundamental concern for any software used in a nuclear forensics context. Often, this functionality takes the form of a post-processing peak-finding technique or signal processing algorithm that can offer an educated guess about



(a) Time profile. Width of each band indicates route flight time.



(b) Three-dimensional view.

Figure 4.4: WC-135 left wing ERD responses over time for the three route mission flown for the Tumbler-Snapper George event.




Figure 4.5: WC-135 upper quadrant DGS detector responses over time for the third route flown for the Tumbler-Snapper George event, corresponding to the top band in figure 4.4a.

the identity of the isotope that produced a gamma-ray of interest to the user. However, because this model was built from the ground up with concerns such as these in mind, debug options were implemented to allow users to tag gamma-rays produced from isotopes of interest and output energy spectra produced by these gamma-rays separately; an example collection of these isotope-specific emission lines is included in figure 4.6.

In these spectra, annihilation photons were split out from this categorization and counted separately, as the interaction is isotope-independent (assuming the 1.022 keV energy threshold is exceeded) and because these photons consistently deposited their full energy in the detection volume. Given the ubiquity of fission product gamma-rays of sufficient energy, that the 511 keV energy bin had the greatest number of counts for almost all of the detector responses recorded from these sample cases was the expected outcome.

Though these spectra are not particularly revealing for any of the sample cases discussed in this work, they could provide valuable information in nuclear forensics contexts where isolating the presence of specific nuclides is a primary focus. For end users without forensics applications in mind, however, the use of this feature is discouraged; depending on the number of isotopes selected for this analysis, output data file sizes can be on the order of gigabytes.

4.2.1 Sample Case Verification

In assessing a typical radiation detection algorithm, experimental detector responses can be used to validate results produced using a detection algorithm. In the absence of experimental data, researchers often turn to other detection algorithms that have already been validated in order to evaluate their own method. As described in section 3.3, MCNP was the proxy software chosen for the purposes of this work. MCNP's pulse-height tally has been shown to be a valid surrogate for detector response when physical detector properties are correctly accounted for [45, 46], and therefore these






Figure 4.6: Sample raw output of WC-135 right wing ERD detector responses by isotope at 24 h following the Redwing Tewa event.




Figure 4.7: Route planning model left wing ERD detector response compared to MCNP pulse-height tally for Tumbler-Snapper George at 36 h post-event.

tally results will serve as the benchmark for the APTool route planning model results presented here.

However, for verification purposes, such a comprehensive view is unnecessary: as verification runs are exact replicas of individual sampler grid traversal points, each individual verification holds as much weight in assessing the model as the cumulation of all verification runs would (while also being significantly less memory-intensive). Each of the four cases summarized in table 3.1 has a sample detector response included in figures 4.7 to 4.10, with each plotted against its corresponding MCNP pulse-height tally.

As with the supporting work described in chapter 2, agreement statistics are also included as defined in equations 2.4a and 2.4b. As MCNP is used as the verification model for the purposes of this work, pulse-height tally data serve as the “observed” results. Because the two models pass output to the user on different bases (MCNP per particle history, route planning model as raw counts), a normalization procedure was required in order to directly compare their data: plots were normalized separately to their own maximum bin value. The maximum was selected rather than the sum for this




Figure 4.8: Route planning model left wing ERD detector response compared to MCNP pulse-height tally for Greenhouse George at 48 h post-event.

purpose to ensure consistent axis presentation across these figures.

Of these cases, Tumbler-Snapper George's modeled response provided the best fit to its MCNP counterpart, such that including the plots on the same axes made them virtually indistinguishable. Agreement between the updated model results and those from MCNP deteriorated roughly linearly with increasing event yield, a phenomenon at least partially attributable to the WC-135's flight ceiling preventing the sampler from accessing all but the bottom of the dispersing cloud (the normalization of the data also contributes to this effect, as the peak-to-total ratio also decreases as event yield rises).

As discussed in section 3.2.3.2.1, Gaussian energy broadening was applied to energy deposition interactions in CsI(Tl) within the model to approximate statistical fluctuations in the number of charge carriers per interaction in a physical detector. MCNP also provides the option to apply Gaussian broadening to pulse-height tally results based on user-defined coefficients that are used in a Gaussian standard deviation fitting function [47]. Despite defining these coefficients based on experimentally-determined Gaussian standard deviation values, however, MCNP results indicate that approximately half of the requested broadening was actually applied, producing spectra




Figure 4.9: Route planning model left wing ERD detector response compared to MCNP pulse-height tally for Hardtack I Elder at 12 h post-event.

with energy resolution values far below those of typical CsI(Tl) detectors.

Though these individual response comparisons reveal some of the more nuanced distinctions between the model developed here and that utilized by MCNP, the primary interest of this verification study was the relative performance of the model across multiple cases and routes. Figures 4.11 and 4.12 show the agreement statistics for the sample cases studied for the purposes of this work; data for each voxel point are included, but spread across the entire route time for visualization purposes.

Overall, results are quite consistent across these cases in that each shows significant variability in model and MCNP response over time. These fluctuations are far more prominent for response FRMS than with FMD, however; because differences in results in individual energy bins of the response have a greater impact on FRMS, slight discrepancies between model detection algorithms can produce very different results for gamma-rays of the same energy (e.g., a slight difference in the Compton scattering probability for a 0.8 MeV gamma-ray could result in one model sampling a photoelectric interaction while the other samples a Compton scatter, resulting in a full energy event for the first and many possible event energies for the second). While these in-








Figure 4.10: Route planning model left wing ERD detector response compared to MCNP pulse-height tally for Redwing Tewa at 24 h post-event.

dividual differences would smooth out over the course of many thousands of detection events if considered on a continuum, the use of individual and discrete energy bins instead magnifies these issues in the final analysis.





In contrast, the fractional mean deviation results for each case indicate that these differences have little cumulative effect on the predictive value of APTool route planning model detectors responses when compared to MCNP simulations. While there remains significant variation across time points within a given case, in general, responses maintain the same shape between the deterministic model and the equivalent stochastic simulation. However, results also indicate consistent response overprediction using the model developed for this work; one factor contributing to this phenomenon is the use of an arbitrary detection threshold for the deterministic model (as discussed in section 3.2.3.2), but likely the more significant issue is the lack of sufficient buildup factor correction for gamma-rays produced by particulates external to the aircraft (described in section 2.5.1.1). In both of these cases, future work could leverage this MCNP verification suite to improve the quality of the deterministic simulation: buildup factors can be computed for each time slice of each case alongside the simulated detector response



(a) Tumbler-Snapper George event.

(b) Greenhouse George event.

Figure 4.11: Agreement statistics between route planning model detector responses and MCNP pulse-height tallies for Tumbler-Snapper George and Greenhouse George events. Data spread over entire route travel time and averaged over ERD and DGS system detectors.



(a) Hardtack I Elder event.

(b) Redwing Tewa event.

Figure 4.12: Agreement statistics between route planning model detector responses and MCNP pulse-height tallies for Hardtack I Elder and Redwing Tewa events. Data spread over entire route travel time and averaged over ERD and DGS system detectors.

without reducing computational efficiency, while FMD results themselves can signal to the user whether the threshold is too high or too low by its sign.

While these results indicate relatively strong agreement between MCNP pulse-height tally and APTool route planning model detector response results, the discrepancies between them clearly suggest that the deterministic model developed for this work does not provide the level of fidelity that an equivalent MCNP simulation is capable of. Several potential areas of improvement for the deterministic model have been identified for future work, however, and although this model is not expected to compete with MCNP's rigorous detector response algorithm in terms of accuracy, the APTool route planning model provides a good approximation of the detector response from atmospheric and collected cloud particulates without introducing additional external software to APTool.

4.3 Model Computational Efficiency

Though significant portions of the APTool route planning model were developed and modified for the explicit purpose of improving computational efficiency, model run times still fall far short of expectations. Computation time comparisons were performed as part of the data collection procedure for the results presented in section 4.1 in order to quantify the efficiency loss incurred due to the implementation of the model functionality described in this work. In each case, the same routes took under 5 s for the legacy model to produce dose and activity data, while under the updated model, obtaining these same results took more than 1 h (data provided in table 4.1).

Because the legacy APTool model was not affected by the implementation of new capabilities, the current approach to the dose and activity calculations is likely preferable to the updated model despite the slight improvements it introduces; users that have no interest in investigating the response of onboard detectors are encouraged to continue to utilize the legacy model for purely dose and activity-related computa-

Table 4.1: Legacy APTool route planning model, updated model, and MCNP verification computation time for sample cases.

Operation	Event	Legacy Model Computation Time (s)*	Updated Model Computation Time (h)*	MCNP Verification Computation Time (h)**
Tumbler-Snapper	George	3.23	1.24	4.11
Greenhouse	George	4.97	1.34	4.54
Hardtack I	Elder	6.61	1.78	4.96
Redwing	Tewa	7.07	2.23	5.05

* Executed on a single AMD Ryzen 3 1200 CPU core (3.1 GHz)

** Executed on 32 AMD Opteron 6320 CPU cores (2.8 GHz) in parallel

tions. As the design of the updated model introduces many complex features, it is unlikely that computation times will ever approach those of the legacy model; however, a similar relationship exists between the detector response results achieved with the updated model and their corresponding MCNP equivalents. Though MCNP response results are unquestionably more accurate than those achieved using the deterministic updated model, computation times were significantly longer for the stochastic runs while also providing none of the debug and auxiliary forensics output options unique to the deterministic model developed for this work.

Improving the computation time for this model is among the most important components of this work that will require improvements in future iterations. Of these, the implementation of parallel execution functionality will provide drastic improvements in model wall-clock run times; as this feature is already present for several steps in the APTool workflow [16], parallel execution will likely be implemented prior to the model reaching end users. However, running the model in parallel does not alleviate computation time issues. Solving these issues will involve significant code modularization, implementation of improved memory management techniques, and conversion to asynchronous data output methods during model execution. In the interest of maintaining the integrity of the underlying APTool code, many of these inefficient code structuring

methods were implemented intentionally, as inserting more efficient methods would necessitate legacy model changes in response. Therefore, although the updated APTool route planning model is currently highly inefficient, many of the techniques necessary to improve efficiency are already developed in preparation for its full integration into the master APTool source code.

CHAPTER 5

CONCLUSIONS

Atmospheric sample collection aircraft provide a relatively safe and reliable method of ensuring compliance with international treaty obligations regarding the responsible use of nuclear material. The time sensitivity of such missions requires that adequate models are established such that prospective mission plans can be rigorously tested prior to deployment. Though current modeling techniques are capable of ensuring adequate aircraft personnel safety and producing estimates of total activity encountered during a mission, more advanced capabilities such as modeling of onboard detection systems and providing recommendations on improving mission plans have yet to be implemented. Using the ORNL-developed APTool in conjunction with DELFIC and NOAA's HYSPLIT, a simulation tool developed to provide these capabilities has been designed and implemented.

Combining ENDF/B-VIII data, an efficient deterministic gamma-ray transport algorithm, and detailed models of the radiation detection systems onboard the WC-135 airborne sampler, this software provides all of the functionality required to fully simulate a nuclear cloud sampling mission, from the physical traversal of the aircraft through the cloud to the directional detection of cloud material to influence flight path. The model was created with flexibility as its core principle, and is designed to fulfill the needs of both end users developing mission plans as well as researchers interested in forensic analysis of nuclear cloud material. Deep investigations into the capabilities of DELFIC and HYSPLIT provided further improvements to the model via parameter optimization and, in DELFIC's case, resulted in the implementation of improvements to this extant software in support of this work. Refinement of the HYSPLIT dispersion model concentration grid traversal algorithm, efficient implementation of nuclear data as LUTs,

and explicit modeling of nuclear cloud particulates provided a reasonably fast and accurate method of predicting detector response results with high fidelity. The ground-up development approach also permitted the inclusion of advanced debugging and forensics output options, giving users the capability to follow not just the detector response history over the course of a sampling mission, but also the response contributions from individual isotopes and interaction types, providing a wealth of data to mine in future research applications.

Though the model was not built to provide improvements to the simulation of radiological quantities (such as dose rate to onboard personnel and the activity of encountered material) that is the sole mission of the legacy APTool route planning model, the enhancements instituted by this work resulted in marginal improvements to these results as well. While this model is meant to provide a route planning solution of greater fidelity than the legacy system, this model was able to be developed without impacting the features present in APTool currently, offering users the flexibility to quickly test a sampling plan in the legacy model and execute a more detailed simulation within the same workflow.

In addition to the development of this model, a detector response validation framework was constructed in support using MCNP. This framework operates in tandem with the model, producing MCNP input files at each calculation point during the simulation that exactly reflect the position of the WC-135, particulates present on its filters and in the immediate vicinity of the aircraft, the atmospheric conditions, and the onboard detection systems, each of which is time-tagged to correspond with the detector response results produced for that same sampling scenario within the route planning model. Validation results for the four sample cases described in this study indicate general agreement between the stochastic MCNP and this work's deterministic model, and the success of this integration opens the door for potential future applications of MCNP results to inform and improve the quality of the route planning model's simulations.

Though this work succeeded in fulfilling its objectives, additional work remains in the future. Despite a concerted effort to maximize model computational efficiency, model changes resulted in significantly longer run times than are achievable with the simpler legacy model. Full integration into APTool will alleviate many of these issues, however, and efforts to streamline code structure and parallelize model execution are underway. Even with these limitations, however, the route planning model remains a more efficient transport simulation option than MCNP while also providing more accurate and higher-fidelity radiological simulations than the legacy model it was built from.

One of the most significant obstacles to new technological development today is the difficulty involved in simultaneously designing and implementing a simulation framework to support the new technology and train prospective users. But with the Air Force's rollout of its new and improved line of WC-135 aircraft [22] coinciding neatly with the complete overhaul of the airborne route planning model used to simulate its missions, the challenges of nuclear treaty verification in the modern age can be met head-on.

Appendices

Appendix A.

Nuclear Test Data

Table A.1: US surface and atmospheric test events with sufficient observed cloud rise data for the analysis described in section 2.1.1.1 ($N=32$). All heights given in reference to MSL.

Operation	Event	Event Yield (kt) [†]	Stabilized Cloud Top Height (ft) ^{††}
Buster-Jangle	Able	3.50×10^0	31,700
Buster-Jangle	Charlie	1.40×10^1	41,000
Buster-Jangle	Sugar	1.20×10^0	15,000
Castle	Nectar	1.69×10^3	71,000
Crossroads	Able	2.10×10^1	40,000
Greenhouse	Dog	8.10×10^1	56,000
Greenhouse	Easy	4.70×10^1	41,000
Greenhouse	George	2.25×10^2	56,000
Greenhouse	Item	4.55×10^1	40,000
Hardtack I	Koa	1.37×10^3	72,200
Ivy	King	5×10^2	67,000
Ivy	Mike	1.04×10^4	98,000
Redwing	Apache	1.85×10^3	66,700
Redwing	Cherokee	3.80×10^3	94,000
Redwing	Dakota	1.10×10^3	75,000
Redwing	Flathead	3.65×10^2	65,700

[†]Data extracted from DOE/NV-209 [15].

^{††}Data extracted from DASA-1251 [5].

Table A.1 (continued).

Operation	Event	Event Yield (kt) [†]	Stabilized Cloud Top Height (ft) ^{††}
Redwing	Lacrosse	4×10^1	38,000
Redwing	Mohawk	3.60×10^2	65,000
Redwing	Yuma	1.90×10^{-1}	8,000
Redwing	Zuni	3.50×10^3	79,000
Sandstone	Yoke	4.90×10^1	56,000
Teapot	Apple-2	2.90×10^1	51,000
Teapot	Hornet	4×10^0	37,000
Teapot	Tesla	7×10^0	30,000
Teapot	Turk	4.30×10^1	44,700
Teapot	Zucchini	2.80×10^1	40,000
Tumbler-Snapper	Fox	1.10×10^1	41,000
Upshot-Knothole	Annie	1.60×10^1	41,000
Upshot-Knothole	Badger	2.30×10^1	36,000
Upshot-Knothole	Harry	3.20×10^1	42,500
Upshot-Knothole	Nancy	2.40×10^1	41,500
Upshot-Knothole	Ray	2×10^{-1}	12,800

[†]Data extracted from DOE/NV-209 [15].

^{††}Data extracted from DASA-1251 [5].

Table A.2: US surface and atmospheric test events with sufficient observed stabilized cloud top data for the analysis described in section 2.1.1.2 ($N=93$). All heights given in reference to MSL.

Operation	Event	Stabilized Cloud Top Height (ft) [†]	Operation	Event	Stabilized Cloud Top Height (ft) [†]
Antler	Round 2*	24,000	Hardtack II	Socorro	26,000
Antler	Round 3*	22,000	Ivy	King	67,000
Buffalo	Round 1*	28,400	Ivy	Mike	98,000
Buffalo	Round 2*	13,000	Mosaic	Round 2*	47,000
Buffalo	Round 3*	14,500	Plumbbob	Boltzmann	33,000
Buster-Jangle	Baker	31,700	Plumbbob	Coulomb-B	18,000
Buster-Jangle	Charlie	41,000	Plumbbob	Diablo	32,000
Buster-Jangle	Dog	46,000	Plumbbob	Franklin Prime	32,000
Buster-Jangle	Sugar	15,000	Plumbbob	Hood	48,000
Castle	Koon	53,000	Plumbbob	Kepler	28,000
Castle	Nectar	71,000	Plumbbob	Morgan	40,000
Crossroads	Able	40,000	Plumbbob	Owens	35,000
Greenhouse	Dog	56,000	Plumbbob	Priscilla	43,000
Greenhouse	Easy	41,000	Plumbbob	Smoky	38,000
Greenhouse	George	56,000	Plumbbob	Stokes	37,000
Greenhouse	Item	40,000	Plumbbob	Wheeler	17,000
Hardtack I	Aspen	48,600	Ranger	Fox	43,000
Hardtack I	Butternut	35,000	Redwing	Apache	66,700
Hardtack I	Cactus	19,000	Redwing	Cherokee	94,000
Hardtack I	Cedar	50,000	Redwing	Dakota	75,000

[†] Data extracted from DASA-1251 [5].

* Denotes joint US-UK test event with data extracted from individual test reports.

Table A.2 (continued).

Operation	Event	Stabilized Cloud Top Height (ft) [†]	Operation	Event	Stabilized Cloud Top Height (ft) [†]
Hardtack I	Dogwood	58,000	Redwing	Erie	32,000
Hardtack I	Elder	50,000	Redwing	Flathead	65,700
Hardtack I	Fir	90,000	Redwing	Huron	54,000
Hardtack I	Hickory	24,000	Redwing	Inca	42,000
Hardtack I	Holly	15,000	Redwing	Kickapoo	16,000
Hardtack I	Juniper	40,000	Redwing	Mohawk	65,000
Hardtack I	Koa	72,200	Redwing	Osage	21,000
Hardtack I	Linden	20,000	Redwing	Seminole	16,000
Hardtack I	Magnolia	44,000	Redwing	Tewa	99,000
Hardtack I	Maple	40,000	Redwing	Yuma	8,000
Hardtack I	Nutmeg	20,000	Redwing	Zuni	79,000
Hardtack I	Olive	50,000	Sandstone	Yoke	56,000
Hardtack I	Pisonia	55,000	Sunbeam	Johnnie Boy	17,000
Hardtack I	Redwood	51,000	Teapot	Apple-2	51,000
Hardtack I	Sequoia	17,000	Teapot	Hornet	37,000
Hardtack I	Sycamore	46,000	Teapot	Tesla	30,000
Hardtack I	Tobacco	18,000	Teapot	Turk	44,700
Hardtack I	Walnut	61,000	Teapot	Zucchini	40,000
Hardtack I	Yellowwood	50,000	Tumbler-Snapper	Charlie	42,000
Hardtack II	De Baca	17,500	Tumbler-Snapper	Fox	41,000
Hardtack II	Humboldt	7,500	Upshot-Knothole	Annie	41,000

[†] Data extracted from DASA-1251 [5].

* Denotes joint US-UK test event with data extracted from individual test reports.

Table A.2 (continued).

Operation	Event	Stabilized Cloud Top Height (ft) [†]	Operation	Event	Stabilized Cloud Top Height (ft) [†]
Hardtack II	Juno	5,500	Upshot-Knothole	Badger	36,000
Hardtack II	Lea	17,000	Upshot-Knothole	Climax	42,700
Hardtack II	Mora	18,500	Upshot-Knothole	Harry	42,500
Hardtack II	Rio Arriba	13,500	Upshot-Knothole	Nancy	41,500
Hardtack II	Sanford	26,000	Upshot-Knothole	Ray	12,800
Hardtack II	Santa Fe	18,000			

[†] Data extracted from DASA-1251 [5].

* Denotes joint US-UK test event with data extracted from individual test reports.

Table A.3: US surface and atmospheric test events selected for the analysis described in section 2.2.2. $N=29$ events.

Operation	Event	Event Yield (kt) [†]	Location	Event Classification
Buster-Jangle	Easy	3.10×10^1	NNSS	Air burst
Buster-Jangle	Sugar	1.20×10^0	NNSS	Surface burst
Castle	Bravo	1.50×10^4	Bikini Atoll [*]	Surface
Castle	Koon	1.10×10^2	Bikini	Surface
Greenhouse	Easy	4.70×10^1	Enewetak Atoll	Surface
Hardtack I	Elder	8.80×10^2	Enewetak	Surface
Hardtack I	Fig	2×10^{-1}	Enewetak	Surface
Hardtack I	Oak	8.90×10^3	Enewetak	Surface
Hardtack I	Redwood	4.10×10^2	Bikini	Surface
Hardtack I	Sequoia	5.20×10^0	Enewetak	Surface
Hardtack I	Sycamore	9.20×10^1	Bikini	Surface
Hardtack II	Wrangell	1.20×10^{-1}	NNSS	Air
Ivy	King ^{**}	5×10^2	Enewetak	Air
Plumbbob	Coulomb-B	3×10^{-1}	NNSS	Surface
Plumbbob	Hood	7.40×10^1	NNSS	Air
Ranger	Baker	8×10^0	NNSS	Air
Redwing	Blackfoot	8×10^0	Enewetak	Surface
Redwing	Cherokee	3.80×10^3	Bikini	Air
Redwing	Dakota	1.10×10^3	Bikini	Surface

[†]Data extracted from DOE/NV-209 [15]

^{*} Events at Bikini and Enewetak Atolls are considered bursts over water

^{**} Simulated event failed to produce HYSPLIT concentration grid

Table A.3 (continued).

Operation	Event	Event Yield (kt) [†]	Location	Event Classification
Redwing	Kickapoo	1.50×10^0	Enewetak	Air
Redwing	Osage	1.70×10^0	Enewetak	Air
Redwing	Yuma	1.90×10^{-1}	Enewetak	Air
Sandstone	Zebra	1.80×10^1	Enewetak	Surface
Teapot	Tesla	7×10^0	NNSS	Surface
Teapot	Turk	4.30×10^1	NNSS	Surface
Tumbler-Snapper	Able	1×10^0	NNSS	Air
Tumbler-Snapper	Fox	1.10×10^1	NNSS	Surface
Upshot-Knothole	Dixie ^{**}	1.10×10^1	NNSS	Air
Upshot-Knothole	Nancy	2.40×10^1	NNSS	Surface

[†]Data extracted from DOE/NV-209 [15]

* Events at Bikini and Enewetak Atolls are considered bursts over water

** Simulated event failed to produce HYSPLIT concentration grid

REFERENCES

- [1] H. G. Norment, Ed., *DELFI C: Department of Defense Fallout Prediction System*, 3 vols., DASA-1800, Bedford, MA: Atmospheric Science Associates, 1970.
- [2] J. P. Lefebvre, D. A. Hooper, M. W. Swinney, E. D. Kabel a, and V. J. Jodoin, “The Airborne Debris Collection Planning Tool,” Oak Ridge National Laboratory, Oak Ridge, TN, Tech. Rep. ORNL/TM-2017/741, 2017.
- [3] D. A. Hooper, E. D. Kabel a, J. P. Lefebvre, S. J. Cope, and V. J. Jodoin, “Ground-based collection of volatile samples: Optimal sampler networks,” Oak Ridge National Laboratory, Oak Ridge, TN, Tech. Rep. ORNL/TM-2018/1101, 2018.
- [4] J. R. Knowles and V. J. Jodoin, “Predicting nuclear yield from stabilized cloud top height using DELFI C and observed entrainment parameters,” Oak Ridge National Laboratory, Oak Ridge, TN, Tech. Rep. ORNL/SR-2017/177, 2017.
- [5] H. A. Hawthorne, Ed., *Compilation of Local Fallout Data from Test Detonations 1945-1962. Continental US and Oceanic Tests*, 2 vols., DASA-1251-EX, Santa Barbara, CA: General Electric Co., 1979.
- [6] G. Spriggs and J. Moye, “Weapon physicist declassifies rescued nuclear test films,” Lawrence Livermore National Laboratory, Livermore, CA, 2017.
- [7] B. T. Rearden and M. A. Jessee, “SCALE code system version 6.2.2,” Oak Ridge National Laboratory, Oak Ridge, TN, Tech. Rep. ORNL/TM-2005/39, 2017, Available from Radiation Safety Information Computational Center as CCC-834.
- [8] H. Bateman, “The solution of a system of differential equations occurring in the theory of radioactive transformations,” *Proc. Cambridge Phil. Soc.*, vol. 15, pp. 423–427, 1908.
- [9] C. J. Bridgman, *Introduction to the Physics of Nuclear Weapons Effects*. Ft. Belvoir, VA: Defense Threat Reduction Agency, 2001.
- [10] D. Brown, M. Chadwick, R. Capote, *et al.*, “ENDF/B-VIII.0: The 8th major release of the nuclear reaction data library with CIELO-project cross sections, new standards and thermal scattering data,” *Nucl. Data Sheets*, vol. 148, pp. 1–142, 2018, Special Issue on Nuclear Reaction Data.
- [11] E. C. Freiling, “Radionuclide fractionation in bomb debris,” *Science*, vol. 133, no. 3469, pp. 1991–1998, 1961.

- [12] R. R. Draxler and G. D. Hess, "An overview of the HYSPLIT_4 modelling system for trajectories, dispersion, and deposition," *Aus. Met. Mag.*, vol. 47, pp. 295–308, 1998.
- [13] A. Staniforth and J. Côté, "Semi-Lagrangian integration schemes for atmospheric models – A review," *Mon. Weather Rev.*, vol. 119, no. 9, pp. 2206–2223, 1991.
- [14] R. Draxler, B. Stunder, G. Rolph, A. Stein, and A. Taylor, "HYSPLIT user's guide," NOAA Air Resources Laboratory, Silver Springs, MD, Tech. Rep. NOAA ERL ARL-230-5, 2020.
- [15] US Department of Energy, "United States Nuclear Tests: July 1945 through September 1992," National Nuclear Security Administration Nevada Field Office, Las Vegas, NV, Tech. Rep. DOE/NV-209-REV-15, 2000.
- [16] J. W. Inman, J. P. Lefebvre, T. P. Norby, I. J. Huff, V. J. Jodoin, and M. R. Brown, "Airborne Debris Collection Planning Tool user manual," Oak Ridge National Laboratory, Oak Ridge, TN, Tech. Rep. ORNL/TM-2020/1838, 2021.
- [17] M. B. Bellamy, S. A. Dewji, R. W. Leggett, *et al.*, "Federal Guidance Report no. 15: External exposure to radionuclides in air, water, and soil," Office of Radiation and Indoor Air, US Environmental Protection Agency, Washington, DC, Tech. Rep. EPA 402-R-19-002, 2019.
- [18] ICRP, "The 2007 recommendations of the International Commission on Radiological Protection," *Ann. ICRP*, vol. 37, no. 103, 2007.
- [19] ICRP, "Recommendations of the International Commission on Radiological Protection," *Ann. ICRP*, vol. 1, no. 26, 1977.
- [20] S. Glasstone and P. J. Dolan, *The Effects of Nuclear Weapons*, 3rd ed. Washington, D.C.: US Department of Defense and Energy Research and Development Administration, 1977.
- [21] K. Way and E. P. Wigner, "The rate of decay of fission products," *Phys. Rev.*, vol. 73, no. 11, pp. 1318–1330, 1948.
- [22] V. Insinna, "Air Force to start transforming tankers into WC-135 'nuke sniffers' in FY19," *Defense News*, 2018, Available: [defensenews.com](https://www.defensenews.com).
- [23] S. A. Romano. (2012). Nuclear treaty monitoring aircraft visits Patrick AFB, Air Force Technical Applications Center. Available: [af.mil](https://www.af.mil).
- [24] J. Moore, M. Myjak, and D. Lee, "Theory of operation/Product specification, ERD - External Radiation Detector - Advanced Atmospheric Research Equipment,"

Kansas City National Security Campus, Kansas City, MO, Tech. Rep. CD1-ERD-0001, 2020, Draft document.

- [25] J. W. Inman, B. A. Wilson, and V. J. Jodoin, “Atmospheric gamma-ray transport from a radioactive cloud to a low earth orbit satellite using MCNP,” *Trans. Am. Nucl. Soc.*, vol. 121, no. 1, pp. 1296–1299, 2019.
- [26] R. D. Evans, *The Atomic Nucleus*. Bombay – New Delhi: Tata McGraw-Hill, 1955.
- [27] D. G. Foster, T. R. England, and N. L. Whittemore, “Experimental-series parameters for the decay of multigroup beta and gamma spectra from 0.1 to 1000 seconds after a fission burst,” Air Force Weapons Laboratory, Kirtland AFB, NM, Tech. Rep. AFWL-TR-78-4, 1978.
- [28] K. E. Mesick, “Delayed gamma signature accounting for field line transport of debris ions after a nuclear explosion,” Los Alamos National Laboratory, Los Alamos, NM, Tech. Rep. LA-UR-19-29657, 2019.
- [29] T. R. England, “CINDER - a one-point depletion and fission product program,” Bettis Atomic Power Laboratory, Pittsburgh, PA, Tech. Rep. WAPD-TM-334, 1962.
- [30] *SCALE Code System Version 4.4*, 3 vols., ORNL/NUREG/CSD-2/R6, Washington, DC: Nuclear Regulatory Commission, 2000.
- [31] T. K. Lane and E. J. Parma, “Delayed fission gamma-ray characteristics of ^{232}Th , ^{233}U , ^{235}U , ^{238}U , and ^{239}Pu ,” Sandia National Laboratory, Albuquerque, NM, Tech. Rep. SAND2015-7024, 2015.
- [32] The Qt Company Ltd., *Qt*, version 5.12.8, Available: qt.io.
- [33] J. K. Tuli, “Evaluated Nuclear Structure Data File,” Brookhaven National Laboratory, Upton, NY, Tech. Rep. BNL-60875, 1995.
- [34] *WC-135 Constant Phoenix fact sheet*, No. 104494, US Air Force, 2005, Available: af.mil.
- [35] J. Amanatides and A. Woo, “A fast voxel traversal algorithm for ray tracing,” *Eurographics*, vol. 87, no. 3, pp. 3–10, 1987.
- [36] L. B. Taylor, Ed., *History of Air Force Atomic Cloud Sampling*, 1 vols., Air Force Systems Command Historical Series 61-142-1, Kirtland AFB, Albuquerque, NM, 1963.

- [37] M. Myjak, "Analysis algorithms for External Radiation Detector and Directional Gamma Sensor," Pacific Northwest National Laboratory, Richland, WA, Tech. Rep. PNNL-29851, 2020.
- [38] C. Schmidt, "Cesium iodide as a gamma-ray spectrometer," *IRE Trans. Nucl. Sci.*, vol. 7, no. 2/3, 1960.
- [39] W. K. Sinclair, *Radiation Dosimetry*, G. L. Brownell and G. J. Hine, Eds. New York: Academic Press, 1956.
- [40] Saint-Gobain Crystals Division, *Efficiency calculations for selected scintillators*, Brochure, 2016, Available: saint-gobain.com.
- [41] O. Klein and Y. Nishina, "On the scattering of radiation by free electrons according to Dirac's new relativistic quantum dynamics," *Nature*, vol. 122, pp. 398–399, 1928.
- [42] G. F. Knoll, *Radiation Detection and Measurement*, 4th ed. Hoboken, NJ: John Wiley & Sons, Inc., 2010.
- [43] C. J. Werner *et al.*, "MCNP users manual - code version 6.2," Los Alamos National Laboratory, Los Alamos, NM, Tech. Rep. LA-UR-17-29981, 2017.
- [44] E. Kalnay, M. Kanamitsu, R. Kistler, *et al.*, "The NCEP/NCAR 40-year reanalysis project," *Bull. Amer. Meteor. Soc.*, vol. 77, no. 3, pp. 437–471, 1996, Available: <ftp.arl.noaa.gov>.
- [45] A. Sood, R. A. Forster, B. J. Adams, and M. C. White, "Verification of the pulse height tally in MCNP 5," *Nucl. Instrum. Methods in Phys. Res. B*, vol. 213, pp. 167–171, 2004.
- [46] A. Sood, M. S. Reed, and R. A. Forster, "New Results for Pulse Height Tally Verification in MCNP5," Los Alamos National Laboratory, Los Alamos, NM, Tech. Rep. LA-UR-03-4271, 2003.
- [47] W. A. Metwally, R. P. Gardner, and A. Sood, "Gaussian broadening of MCNP pulse height spectra," *Trans. Am. Nucl. Soc.*, vol. 91, pp. 789–790, 2004.