

Abstract

HEVENER, RYNE ANDREW. Investigation of Energy Windowing Algorithms to Enhance Nuclear Material Screening in Radiation Portal Monitors. (Under the direction of Dr. Man-Sung Yim).

The illicit trafficking of nuclear materials has become one of the most serious security concerns in modern times. Radiation portal monitors used to detect the trafficking have to balance cost versus screening ability. This thesis investigated the use of energy windowing algorithms to improve the screening ability of a low cost variety of portal monitors, and created an algorithm to perform that very task.

The new algorithm consisted of energy windows of various sizes calibrated with specific target nuclides (when possible), a uniquely arranged alarm logic structure, and an alarm library designed to assist portal monitor installation workers during secondary inspections. The new algorithm's high energy masking sensitivity was then compared to two commercial energy windowing algorithms.

A simulated real-time emulator was created to assess the new algorithm's ability to correctly identify target nuclear materials. Actual vehicle data provided by Oak Ridge National Laboratory was combined with so-called injection sources to simulate the illicit trafficking of nuclear materials.

The results of the sensitivity analysis and simulated real-time emulation are discussed and recommendations for future work are presented.

Investigation of Energy Windowing Algorithms to Enhance Nuclear Material Screening in
Radiation Portal Monitors

by
Ryne Andrew Hevener

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

Nuclear Engineering

Raleigh, North Carolina

2011

APPROVED BY:

Dr. Man-Sung Yim
Committee Chair

Dr. David McNelis

Dr. Kimberly Weems

Dedication

To my mother, father, brother, sister, and soon-to-be wife: may God continue to bless us in all that we do.

Biography

Ryne Hevener was born in Bluefield, West Virginia on September 15th, 1987. He is the second of three children of Tom and Jane Hevener, along with his older brother Jordan and younger sister Riley. He graduated from North Carolina State University in Raleigh, North Carolina with a Bachelor of Science degree in Nuclear Engineering in May of 2010.

He continued his education at North Carolina State in pursuit of a Master of Science degree in Nuclear Engineering. During his time at NC State, he has worked under the direction of Dr. Man-Sung Yim and in conjunction with Oak Ridge National Laboratory.

He plans to marry his fiancée Heather Thompson in the summer of 2011.

Acknowledgements

First and foremost I want to thank God, without whom none of this would be possible. I also want to thank my parents, siblings, and fiancée for always being there for me, no matter what.

I would like to thank Dr. Man-Sung Yim for his constant support, enthusiasm, positive attitude, and friendship. I also wish to extend my gratitude to Dr. David McNelis and Dr. Kimberly Weems for their participation on my committee.

Final thanks go to the folks at Oak Ridge National Laboratory, including Ken Baird, Chris Blessinger, and Rob York. They have been a vast source of information and I appreciate their assistance.

Table of Contents

List of Figures	vi
List of Tables	viii
Chapter 1 – Introduction and Background.....	1
Chapter 2 – Literature Review	8
Chapter 3 – Scope of Work, Methods, and Approaches.....	32
Chapter 4 – Algorithm Development and Analysis	47
Chapter 5 – Discussion	124
Chapter 6 – Summary and Recommendations for Further Studies.....	130
References.....	132
Appendix.....	133
Appendix A – Sensitivity Analysis Plots.....	134
Appendix B – Simulated Real-Time Emulator Results Tables	161
Appendix C – Advanced Alarm Algorithm Statistics Tables	170

List of Figures

Figure 1.1: Example of a deployed RPM system (Kouzes, et al. 2005)	2
Figure 2.1: Measured spectra from a NaI(Tl) detector (Siciliano, et al. 2005).....	10
Figure 2.2: Measured spectra from a PVT detector (Siciliano, et al. 2005)	13
Figure 2.3: Example of average and sum buffers	14
Figure 2.4: Simulations of vehicle profiles showing the effects of background suppression.	18
Figure 2.5: Profiles of PRS for top panels, bottom panels, average and sum of all panels (Lo Presti, et al. 2006).....	29
Figure 3.1: Interaction of a source with a background spectrum.....	40
Figure 3.2: Example sensitivity plot	44
Figure 3.3: Simulated real-time alarm emulator diagram.....	46
Figure 4.1: HEU injected cargo and background spectra	50
Figure 4.2: S/\sqrt{B} versus channel calibration for HEU.....	50
Figure 4.3: ^{133}Ba injected cargo and background spectra	51
Figure 4.4: S/\sqrt{B} versus channel calibration for ^{133}Ba	52
Figure 4.5: ^{137}Cs injected cargo and background spectra	53
Figure 4.6: ^{137}Cs S/\sqrt{B} versus channel calibration	54
Figure 4.7: ^{60}Co injected cargo and background spectra	55
Figure 4.8: S/\sqrt{B} versus channel calibration curve for ^{60}Co	56
Figure 4.9: ^{40}K injected cargo and background spectra.....	57
Figure 4.10: ^{40}K S/\sqrt{B} versus channel calibration curve	58
Figure 4.11: Occupied Spectrum with Energy Windows Shown	60
Figure 4.12: Plot of standard deviation of the compensated ratio versus high energy window upper edge	62
Figure 4.13: Plot of S/\sqrt{B} versus channel	63
Figure 4.14: Advanced alarm algorithm alarm logic structure	64
Figure 4.15: Visual representation of normalized combined area	69
Figure 4.16: ^{57}Co injection source normalized combined area plot.....	72
Figure 4.17: HEU injection source normalized combined area plot.....	73
Figure 4.18: ^{133}Ba injection source normalized combined area plot	74
Figure 4.19: ^{137}Cs injection source normalized combined area plot.....	76
Figure 4.20: ^{60}Co injection source normalized combined area plot.....	77
Figure 4.21: ^{40}K injection source normalized combined area plot	78
Figure 4.22: Histograms of advanced algorithm window ratios.....	82
Figure 4.23: Histograms of advanced algorithm window thresholds	83
Figure 4.24: Background Suppression Histograms	88
Figure 4.25: Schematic of a point source traversing an RPM (Runkle, et al. 2005)	92
Figure 4.26: Detector yield from (a) f_i and (b) the spatial response function S_i	94
Figure 4.27: Plot of Sensitivity versus Injection Source Detected Count Rate for Standard Cargos.....	109
Figure 4.28: Plot of Sensitivity versus Injection Source Detected Count Rate, ^{40}K Cargos	109

Figure 4.29: Plot of Sensitivity versus Injection Source Detected Count Rate, Granite Cargos	110
Figure 4.30: Plot of False Alarm Rate versus Injection Source Detected Count Rate, Standard Cargos.....	111
Figure 4.31: Plot of False Alarm Rate versus Injection Source Detected Count Rate, ⁴⁰ K Cargos.....	111
Figure 4.32: Plot of False Alarm Rate versus Injection Source Detected Count Rate, Granite Cargos.....	112
Figure 4.33: Plot of Precision versus Injection Source Detected Count Rate, Standard Cargos	113
Figure 4.34: Plot of Precision versus Injection Source Detected Count Rate, ⁴⁰ K Cargos...	114
Figure 4.35: Plot of Precision versus Injection Source Detected Count Rate, Granite Cargos	114
Figure 4.36: Plot of Matthews Correlation Coefficient versus Injection Source Detected Count Rate, Standard Cargos	116
Figure 4.37: Plot of Matthews Correlation Coefficient versus Injection Source Detected Count Rate, ⁴⁰ K Cargos.....	116
Figure 4.38: Plot of Matthews Correlation Coefficient versus Injection Source Detected Count Rate, Granite Cargos	117
Figure 4.39: Plot of F1 Score versus Injection Source Detected Count Rate, Standard Cargos	118
Figure 4.40: Plot of F1 Score versus Injection Source Detected Count Rate, ⁴⁰ K Cargos ...	118
Figure 4.41: Plot of F1 Score versus Injection Source Detected Count Rate, Granite Cargos	119

List of Tables

Table 4.1: Energy windowing algorithm summary	59
Table 4.2: Full alarm truth table for advanced alarm algorithm	65
Table 4.3: Summary statistics for advanced algorithm window ratios	84
Table 4.4: Summary statistics for advanced algorithm window thresholds	84
Table 4.5: Background suppression summary statistics	89
Table 4.6: Emulator Results- Non-alarming POE Profiles	102
Table 4.7: Emulator Results- Alarming POE Profiles	102
Table 4.8: Example Emulator Results Table	105
Table 4.9: Confusion Matrix used for Advanced Alarm Algorithm Statistical Analysis	105
Table 4.10: Statistical Measures Derived from Confusion Matrix	107
Table 4.11: Alarm Library	121
Table 5.1: Alarm Library Human Factors Experiment	128

Chapter 1 – Introduction and Background

Radiation portal monitors (RPMs) are distributed across the globe with the goal of “detering, detecting, and interdicting” the illicit trafficking of nuclear material that can be used in nuclear weapons or radiological dispersal devices (RDDs or “dirty bombs”). Mainly distributed by the National Nuclear Security Administration’s Second Line of Defense program, RPMs are primarily located in former Soviet Union states and major international shipping ports (NNSA Public Affairs 2010).

Radiation portal monitors are large, pillar-like detection systems through which potential carriers of nuclear material pass. If an RPM occupier triggers an alarm in the initial pass-through (called the primary scan or inspection), the occupier is diverted to a secondary scan where a more detailed analysis of the occupier takes place. These monitors are designed to accept a specific mode of transportation and in their various types are deployed in airports, seaports, and border crossings. The three major modes of transportation RPMs are intended to investigate are pedestrians (using pedestrian monitors), cars/tractor trailers (vehicle monitors), and trains (rail monitors). Illicit nuclear materials will primarily be transported via one of these transport modes, thus it is necessary to have a system dedicated to each of these needs. Oak Ridge National Laboratory (ORNL) has provided large sets of vehicle monitor data, thus this study focuses on this type of RPM. Since all monitors based on polyvinyltoluene (PVT) gamma detectors share virtually identical detector setups, the results are applicable to any PVT style RPM.

RPMs are composed of two monolithic metal pillars (typically aluminum) that contain gamma ray and neutron detection equipment. Current models utilize polyvinyltoluene (PVT) organic plastic scintillators as the gamma detectors and helium-3 neutron detectors. Additionally, the master pillar contains signal electronics, the data acquisition system, and the communication system which records and transmits the data. For vehicle monitors,



Figure 1.1: Example of a deployed RPM system (Kouzes, et al. 2005)

each pillar typically contains two PVT scintillators stacked one on top of the other. Vehicles pass between the pillars at a relatively slow speed while the monitor collects data. If a significant increase in the count rate is detected while the monitor is occupied versus while it was unoccupied, an alarm will trigger. The algorithms that cause the alarms are the major focus of this study and will be discussed in depth at a later point. First, it is important to quickly address some of the attributes of PVT.

The PVT scintillator detection system is the lowest cost and most common type of gamma detector utilized in RPMs. PVT is a synthetic, organic polymer and can be made quite large for a relatively low cost. It is also resistant to harsh environmental factors (heat, cold, humidity), which is necessary for deployment in severe environments. The large detector volume increases detection efficiency, thus only one large or two smaller scintillators are required in a single pillar to provide sufficient detection capability. The main cause of concern for PVT hardware is the photomultiplier tubes affixed to the ends of the

scintillators. The behavior of the phototubes in extreme temperatures could affect RPM effectiveness, but modification of the site, such as installing solar shields, could abate the issue. Overall, the large size, low cost, and robust nature of PVT detectors make them excellent candidates for mass-deployment RPM projects.

PVT RPMs currently operate using simple gross-counting algorithms where a vehicle will trigger an alarm if it registers a count rate greater than a threshold count rate above the background radiation level. Typically, this threshold is 4σ above the mean background count rate; this allows for a false alarm rate of less than 1/30,000.

Unfortunately, PVT is not without its issues. Being a low density plastic, it lacks the energy resolution that spectrum-sensitive detectors like high purity germanium (HPGe) and thallium-doped sodium iodide, NaI(Tl), possess. PVT's energy resolution is so poor that gamma ray photopeaks do not appear in a gamma spectrum; as a result only the Compton continuum and Compton edge are visible. The Compton edge, however, is not a sharp edge but an extended sloping shelf. In short, PVT scintillators lack the natural spectroscopic capabilities of NaI(Tl) and HPGe detector systems.

The major radiation portal monitor problem is devising an inexpensive, robust system capable of discriminating between naturally-occurring radioactive material (NORM) and special nuclear material (SNM) without significantly disrupting the flow of commerce. This begs for the usage of detectors with fine spectroscopic detail, but in practice this is impossible because these detectors tend to be small and expensive.

The small size of the detectors implies that either a vehicle has to sit in the RPM for a significant amount of time to allow for adequate data collection, or many detectors are

needed per pillar to collect sufficient data to construct a detailed spectrum. The high cost of spectrum-sensitive detectors prevents the usage of multiple detectors per pillar, thus long count times have to be employed. RPMs are placed at major international shipping routes and interrupting the flow of commerce costs a great deal of money and frustrates both shippers and RPM operators, thus the resulting disruption in commercial traffic flow will not allow such systems to be employed for primary scans, and the high cost cannot be justified for their exclusive use in secondary scans. Also, installations where slow delays plague the operation put stress on drivers and RPM operators, often culminating in improper usage of the monitors, defeating the entire purpose of the program.

Given the seemingly devastating lack of detail from PVT spectrum data, it would appear as if no single RPM design is sufficient to achieve all the desired ends: low cost, minimal slowdown on initial scan, and good spectroscopic detail.

This study investigates the development of a method to upgrade existing, unsophisticated PVT RPMs, giving them the ability to perform spectroscopic analyses using advanced detection algorithms implemented solely through software modification. The approach uses a technique called energy windowing (EW) to detect differences between occupied-signal spectrum shapes and background-signal spectrum shapes to not only detect the presence of radioactive cargo, but to do a rough identification of the material present prior to a secondary inspection.

By giving an inspector some prior knowledge of what material to expect during the secondary inspection, the process can be sped up, decreasing the disruption in commercial flow. Secondary inspections are almost always performed using heavy, expensive, hand-held

spectroscopic detectors as well. If PVT RPMs can perform a sufficient spectroscopic analysis, they could possibly be used as secondary scanners themselves by making the vehicle sit in the monitor for an extended counting period.

The goal of EW algorithms is the ability to discriminate between SNM laden cargos and cargos carrying NORM. This discrimination depends on the afore-mentioned difference between spectrum shapes. Uranium-235/highly enriched uranium (HEU) and plutonium-239/weapons-grade plutonium (WGPu) are the two SNM sources of concern, and both materials have low-energy gamma peaks: HEU has a distinct 186 keV gamma peak and Pu-239 has approximately 10 large gamma peaks between 330 keV and 420 keV. This differs from the major NORM source ^{40}K , which has a preponderance of high-energy gamma rays. The other common NORM sources, ^{232}Th and ^{226}Ra , have gamma rays that spread across the energy spectrum.

Energy windowing works by dividing a vehicle's spectrum into low- and high-energy windows and comparing the vehicle's ratio of low-energy to high-energy counts to the ratio of low-energy to high-energy counts for a background spectrum taken just prior to the RPM's occupancy. If there are significantly more low-energy counts in the vehicle's spectrum than there are in the background spectrum, then it is possible that the vehicle is carrying an SNM source and the secondary inspection needs to be performed with great care. Vehicles carrying cargos that contain NORM sources, such as marble, granite, or bananas, would show more high-energy counts than the background spectrum, thus telling operators to expect a NORM source upon secondary inspection.

Energy windowing does not solve all problems though. With cargos mixed with low-energy gamma sources and high-energy gamma sources, the high-energy counts trick the algorithm into suspecting a NORM cargo and ignoring the low-energy gamma counts. This effect is known as NORM masking, high energy masking, or simply masking and is an inherent fault in EW algorithms. There are also questions about the accuracy of EW algorithm alarms due to the poor spectroscopic capabilities of PVT; in other words, can energy windowing really improve the NORM discrimination and SNM detection abilities of PVT RPMs?

By focusing solely on the gamma ray detection system, this research work investigates the use of energy windowing algorithms to provide a limited NORM identification ability while minimizing energy windowing's sensitivity to masking. After discussing the current gross counting algorithms and the problems associated with them, this study examines if, even with masking, the energy windowing/gross counting combined system would outperform a gross counting only system. A real-time emulator of the advanced EW algorithm is developed to show the effectiveness of the system by creating a visual of the hardware/software needed to implement the algorithm in an actual RPM system. This study does not examine the issue of neutron detection because energy discrimination is unnecessary; the presence of increased neutron count rates signifies the presence of a source requiring investigation.

Chapter 2 of the thesis is a review of the literature written in recent years about radiation portal monitors, specifically involving the development of energy windowing algorithms. Further detail about the construction and operation of RPMs is discussed.

Features of PVT and PVT RPMs are discussed in this chapter, especially with regards to the spectroscopic capability of the systems. Comparisons between PVT RPMs and NaI(Tl) RPMs are examined as well. Gross count algorithms and energy windowing algorithms are introduced and examined in detail. Background suppression is the final topic in chapter 2.

The methods, tools, and approaches used to perform this study are reviewed in chapter 3. Several programs were employed to complete the various tasks associated with parsing large datasets, manipulating the resulting data into a usable product, and analyzing the results. Details of the Perl scripts, MATLAB/Simulink codes, and Mathcad worksheets are necessary to grasp the scope of the inner working of the study.

Chapter 4 presents the development of energy windowing algorithms, including the thought process behind their construction and how they are implemented. This section includes a discussion on the standards used to development EW algorithms. The bulk of the chapter focuses on the development of a new, advanced energy windowing algorithm that seeks to heal the ills of previous conceptions. The three algorithms are also compared, looking at NORM discrimination ability, correct alarm performance, and sensitivity to masking. The experimental validation of the advanced algorithm is performed by a simulated real-time emulator and is discussed in chapter 4. This validation shows the effectiveness of the advanced algorithm in discriminating between NORM and SNM, and its ability to identify a limited range of selected nuclides.

The significance of the findings and the problems that occurred during the study are reviewed in chapter 5. Concluding arguments and suggestions for future work are contained within chapter 6.

Chapter 2 – Literature Review

Following the terrorist attacks of September 11, 2001, the threat of nuclear attack from sub-national actors, i.e. terrorist organizations, has become a serious concern for the United States government under the purview of the Department of Homeland Security and the National Nuclear Security Administration. In the effort to prevent the illicit trafficking of nuclear materials, radiation portal monitors have been installed at international border crossings and major seaports across the globe under the NNSA's Second Line of Defense program. The threat resides not only in the flow of complete nuclear weapons, but also the transport of special nuclear material as well as industrial and medical sources that could be used in radiological dispersal devices.

Originally, RPMs were tasked with detecting the movement of nuclear material out of SNM storage/processing facilities and detecting the presence of radioactive sources in scrap entering steel reprocessing facilities (Siciliano, et al. 2005). The new application of monitors to scan large volumes of cargos at busy ports has required a re-examination of RPM technology which is still ongoing.

As mentioned in chapter 1, RPMs consist of two monolithic pillars through which vehicles pass. Typical transit times range from about 5 seconds to 21 seconds regardless of the vehicle type or cargo (Lo Presti, et al. 2006); this implies that, as expected, passenger cars have short scan times due to their small size whereas tractor trailers, being much larger, take longer to drive through the pillars. This initial drive through the RPM is called the primary screening and acts as a trip-wire. If an alarm is triggered during the primary scan the vehicle

is diverted to a secondary screening area where a more detailed exam of the vehicle is performed. The source of the primary alarm is discovered and identified using spectroscopic equipment during the secondary inspection. The primary-secondary, two-stage screening process allows for efficient examination of vehicles (Kouzes, et al. 2005).

The main technical requirement of RPMs is to rapidly detect small, concentrated sources with high detection probability and an operationally acceptable false-alarm rate. Additional performance criteria identified by researchers include: adequate spatial coverage, rapid and accurate screening, NORM identification, durability, and affordable cost (Kouzes, et al. 2005). Of the five performance areas, rapid and accurate screening and NORM identification criteria can be directly addressed through the use of advanced detection algorithms.

Since vehicles usually take less than 30 seconds to pass through the RPM, the monitor needs to be able to detect an offending source with minimal data collection time. Counts are typically collected in 0.1 second time bins and a running sum of ten consecutive measurements constitutes a per-second count rate value (Kouzes, et al. 2005). With an MCA-based electronics package, the count rate data is also sorted into energy bins based on the pulse height response of the detector. The fast response and large available size of PVT scintillators allows PVT-based systems to make rapid decisions. The accuracy of primary screenings varies wildly with the detection algorithm used, as some algorithms can be sensitive to effects such as background suppression and NORM masking. The issue of background suppression has been addressed by researches and will be discussed later in chapter 2. NORM masking is an issue with energy windowing algorithms and has been

studied less extensively; it will be examined thoroughly in chapter 4. It has been shown that advanced algorithms (EW algorithms) can compensate for the effects of BG suppression and improve RPM accuracy (Kouzes, et al. 2005) (Ely, et al. 2006). NORM identification is a challenge for current PVT systems due to the lack of inherent spectroscopic capability of PVT scintillators and the primitive nature of gross counting algorithms. With the seeming lack of information PVT scintillators provide, it begs the question, “Why use PVT-based RPMs?”

The answer to this question lies in the fact that three of the five performance areas are easily met with PVT detectors: adequate spatial coverage, durability, and affordable cost. PVT scintillators can be produced in very large sizes for relatively low cost, thus providing these systems with large area coverage ability for much lower costs compared to NaI(Tl) (Siciliano, et al.

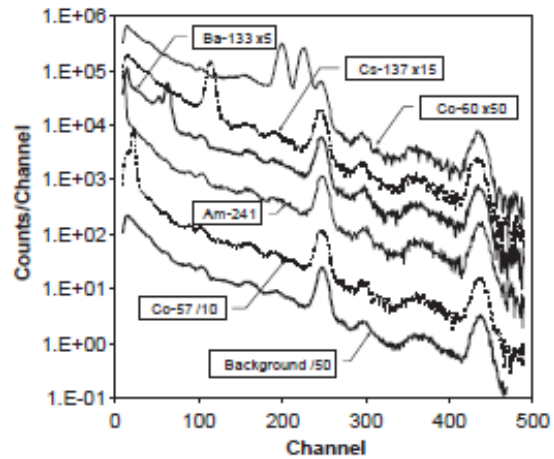


Figure 2.1: Measured spectra from a NaI(Tl) detector (Siciliano, et al. 2005)

2005). As mentioned earlier, NaI(Tl) does exhibit far better energy resolution than PVT and NaI(Tl) can actually show the gamma ray photopeaks not visible in PVT spectra. The smaller size and higher cost of NaI(Tl) scintillators force these so-called spectroscopic portal monitors (SPMs) to have longer counting times to collect adequate information to produce a fully resolved gamma spectrum (Siciliano, et al. 2005).

The benefit of better energy resolution has to be weighed against the detriment that longer counting times have towards the flow of commerce and the SPM's higher cost. To justify the use of SPMs over RPMs, three requirements have been identified that must be met by the SPMs: the SPM must have increased sensitivity over RPMs, the SPMs must reduce the operational burden on installation workers, and the SPM must have improved radionuclide identification capability over hand-held detectors used in secondary investigations (Siciliano, et al. 2005). Up to this point, PVT RPMs are still the most widely used models and the SPMs are only a minor player in the portal monitoring debate. With the development of advanced algorithms that increase the ability of PVT RPMs to reject NORM cargos and identify selected nuclides, the use of the cheaper PVT RPMs is likely to continue.

While PVT lacks inherent spectroscopic ability, it does have several desirable characteristics. PVT has high gamma ray efficiency (Kouzes, et al. 2005), effectively producing signals from gamma ray energies between 20 keV and 3000 keV. The light transmission ability of PVT is also superior to that of other plastic scintillators (Siciliano, et al. 2005). Detector response is fast so copious amounts of data can be collected in a short time. These properties coincide with the needs of portal monitors at border crossings and sea ports; NORM and SNM gamma ray energies fall mainly within the energy range PVT can effectively process and fast data collection is desired to minimize traffic slowdowns. While the spectroscopic ability of PVT has been reviewed previously, specific details about spectrum shape are necessary to understand how an energy windowing algorithm can increase the effectiveness of PVT RPMs.

Compton scattering is the most likely interaction of gamma rays with PVT for photon energies greater than 20 keV, and is the dominant interaction for energies greater than 80 keV. In contrast, photoelectric absorption is the dominant interaction until photons exceed 300 keV (Siciliano, et al. 2005). Photoelectric absorption is essentially a full-energy interaction in that the entire energy of the photon is absorbed in the NaI(Tl) material, thus a photopeak occurs in the energy spectrum at the specific energy of the incident photon (see Figure 2.2 above). Compton scattering is by no means a “full-energy” interaction; the energy of the Compton recoil electron is a function of both the incident photon energy and the scattering angle of the photon. The recoil electron energy is given by the equation:

$$E_{e^-} = hv - hv' = hv \left(\frac{(hv/m_0c^2)(1 - \cos \theta)}{1 + (hv/m_0c^2)(1 - \cos \theta)} \right) \quad \text{where:}$$

E_{e^-} = energy of recoil electron
 hv = incident photon energy (typically in keV) (2.1)
 hv' = scattered photon energy (typically in keV)
 m_0c^2 = mass-energy equivalence term (511 keV for an electron)

The minimum energy of a recoil electron occurs when a gamma ray “grazes” an electron, meaning that the scattering angle is approximately 0°. This results in a recoil electron with basically zero energy and a scattered gamma ray with nearly all the incident energy. The maximum energy transfer between a photon and an electron occurs when the photon backscatters 180° after a collision with an electron; this maximum recoil electron energy is called the Compton edge. The Compton edge is calculated from the equation:

$$E_{e^-}|_{\theta=\pi} = hv \left(\frac{2hv/m_0c^2}{1 + 2hv/m_0c^2} \right) \text{ where:} \quad (2.2)$$

$E_{e^-}|_{\theta=\pi}$ = maximum energy of recoil electron
 hv = incident photon energy (typically in keV)
 m_0c^2 = mass-energy equivalence term (511 keV for an electron)

In practice, gamma rays will scatter across all angles, so a detector dominated by Compton scattering interactions will output a signal between zero and the maximum recoil electron energy (called the Compton continuum). Thus, a 200 keV gamma ray source will create a signal in the range of 0 keV up to 88 keV in a PVT detector system. Unfortunately, the poor intrinsic resolution of the PVT scintillator results in the Compton edge “smearing” out from a sharp edge into a broad, sloped edge that can be seen in Figure 2.2 (Ely, et al. 2006).

Although the energy resolution is poor, crude information can be extracted from PVT spectra (Ely, et al. 2006). Since a given gamma ray source adds counts between 0 keV and the Compton edge, comparing the shape of this source-modified spectrum to the shape of a background spectrum can provide sufficient information to roughly identify certain nuclides. It is this




Figure 2.2: Measured spectra from a PVT detector (Siciliano, et al. 2005)

deviation of the occupied signal shape from the background signal shape acquired by PVT RPMs that is exploited in the advanced detector algorithm discussed in this paper, the energy windowing algorithm.

Current PVT RPMs utilize gross count algorithms to passively detect nuclear material. Gross count algorithms compare gross count rates while a vehicle is occupying a monitor to the average background count rate established prior to the vehicle entering the monitor. These algorithms can be implemented across the entire energy spectrum or in selected windows (a “window” being a group of consecutive channels binned together as a single, large “channel”) (Ely, et al. 2006). Gross count algorithm thresholds are calculated from the average background count rate, which is calculated by measuring background over a set time interval and averaging the count rate (Ely, et al. 2006).

In most implementations, RPMs operate in one of two modes: background mode and

occupied mode. Background mode occurs when an occupancy is not detected by the monitor. In this mode, the RPM records the count rate every five seconds and calculates the average background count rate (used in the threshold calculation) using an average buffer. The buffer is simply a moving average of the background count records, converted into a count rate (units of counts per second i.e. cps).

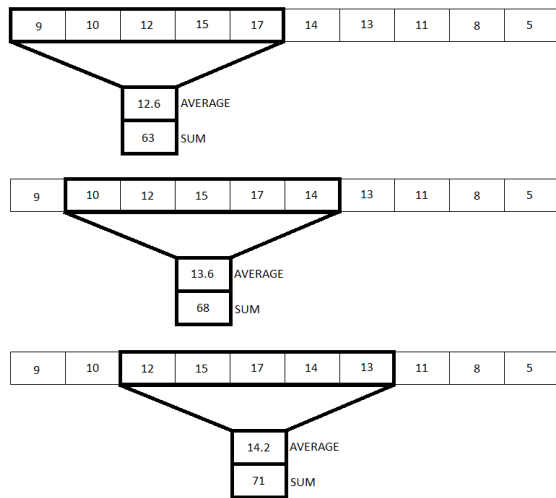


Figure 2.3: Example of average and sum buffers

Occupancy mode occurs when a vehicle triggers an occupancy sensor located in the RPM

pillars. In occupied mode, counts are recorded every 200 milliseconds. One-second count rates are calculated using a 1-second sum buffer, which is a moving sum of five 200-millisecond count records; this outputs a count rate with units of cps.

Given the average background count rate, the gross count threshold can be calculated by the equation:

$$T_{GC} = B + K\sqrt{B} \quad \text{where:} \quad (2.3)$$

T_{GC} = Gross count threshold (cps)
 B = Background count rate (cps)
 K = Sigma multiplier

The \sqrt{B} term is the standard deviation for the background count rate and the K term is the sigma multiplier that sets the threshold above the background level by a certain number of standard deviations. Probability distributions of gamma ray count rates indicate a positive-skewed Gaussian distribution. This observed Gaussian distribution is actually the large count rate limit of a Poisson distribution, thus the traditional Gaussian standard deviation is almost always replaced by the Poisson distribution formula for standard deviation (Ely, et al. 2006). The standard deviation used in gross count is thus the square root of the counts in a one-second interval (or square root of the count rate).

$$\sigma_{GC} = \sqrt{B} \quad (2.4)$$

The threshold is set to have a high detection efficiency while maintaining a low false alarm rate. A four sigma threshold is often used due to a false alarm rate of approximately 1 in 30,000.

The issue with gross count algorithms is the lack of information able to be gleaned from an alarm. Recall that NORM sources produce high energy gamma rays and generate counts across all energies ranging from 0 keV up to the Compton edge. Therefore, it is highly likely that a NORM source (e.g. fertilizer, kitty litter, etc.) will generate a gross count alarm in all windows lower than the Compton edge, even a window designed specifically to detect SNM sources with energies well below 300 keV. This extreme lack of NORM discrimination ability renders gross count algorithms nearly useless in RPM installations that see frequent shipments of innocent radioactive cargo.

Gross count algorithms are also particularly sensitive to a phenomenon known as background suppression, which is also known as shadow shielding, baseline depression and/or baseline suppression. Gamma ray count rates collected from empty portal monitors are generally higher than the count rates measured during vehicle occupancy. This is a result of the shielding effect of the vehicle structure and cargo; the presence of a vehicle in the monitor shields the RPM's PVT panels from the environments gamma ray sources, namely the roadway, and concrete and brick structures (Lo Presti, et al. 2006). This effect varies greatly depending on the type of vehicle and the cargo it contains; trucks built with dense materials and vehicle carrying dense cargo tend to suppress background more than empty trucks (Lo Presti, et al. 2006). Also, background suppression tends to vary spatially as the truck moves through the portal. The engine block and axels consist of large amounts of high-Z materials so the suppression is increased, but the gap between the cab of a tractor trailer and the trailer itself shields very little. In fact, this gap can look like a radiation source if the time dependent spectrum is analyzed (Lo Presti, et al. 2006).

Background suppression poses a problem for gross count algorithms because it artificially raises thresholds. This effect prevents sources that would normally trigger an alarm if passed through the portal unshielded from actually alarming. The shielding effect can suppress the occupied count rates anywhere from 10% up to 30% (Lo Presti, et al. 2006), meaning that, given a threshold, a radiation source has to be potentially 30% more active to trigger the same alarm if background suppressed than if the suppression is not present. The quantity called the percent background suppression (PBS) is found from the formula (Lo Presti, et al. 2006):

$$PBS = \frac{(OccupiedCountRate - BackgroundCountRate)}{BackgroundCountRate} \cdot 100 \quad (2.5)$$

The PBS value indicates the extent to which background measurements are suppressed by a vehicle occupying a monitor. The formula is usually applied to individual time bins across the vehicle's time profile. The time profile, more often called the vehicle profile, is the plot of measured count rates versus time as the vehicle passes through the RPM. A negative value for the PBS implies that the occupied count rate is less than the background count rate; this will occur unless a source is present in the vehicle.

Figure 2.4 illustrates the background suppression problem effectively. The horizontal dashed line represents the gross count algorithm threshold from equation (2.3) and the noisy line represents the simulation of a vehicle profile.

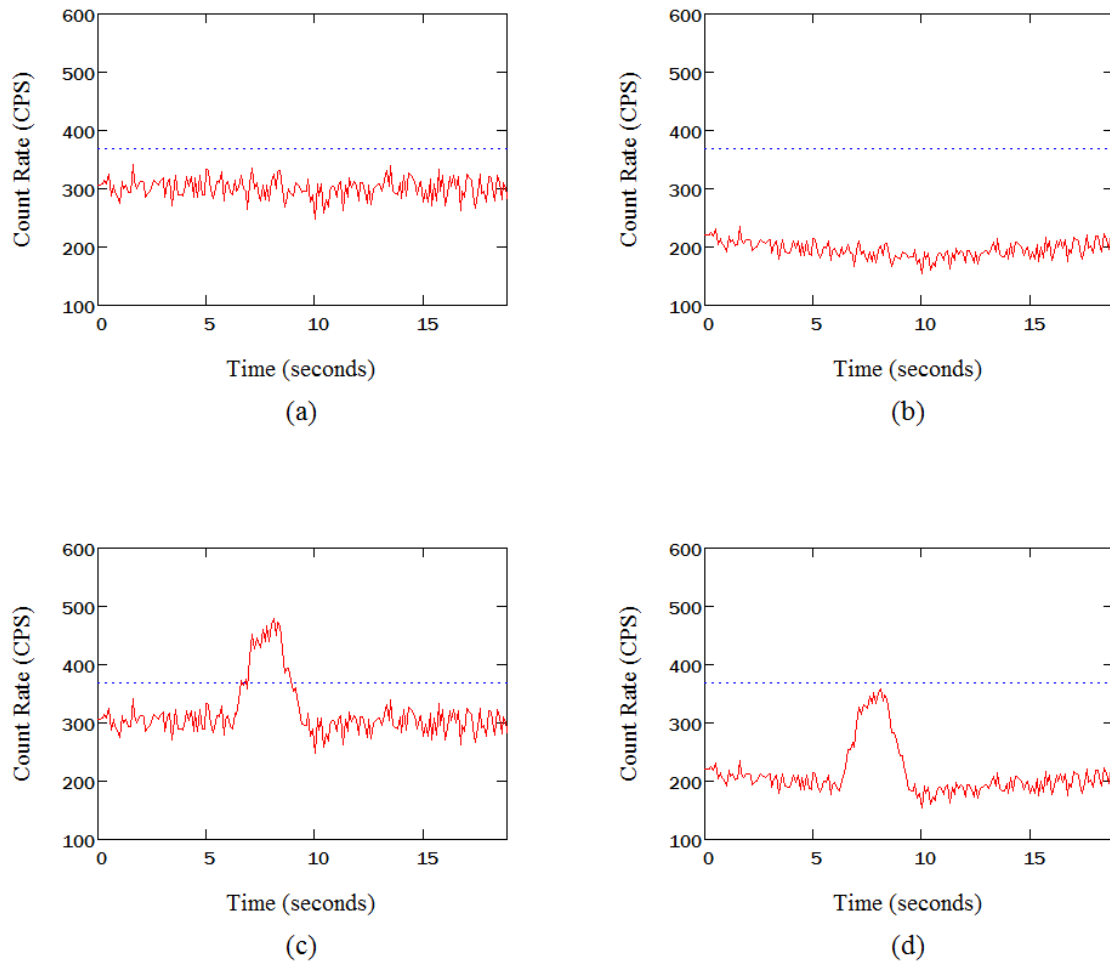


Figure 2.4: Simulations of vehicle profiles showing the effects of background suppression

- (a) Background time profile, no source, no background suppression**
- (b) Background time profile, no source, 30% average background suppression**
- (c) Vehicle profile, source present, no background suppression**
- (d) Vehicle profile, source present, 30% average background suppression**

Plots (c) and (d) reveal the nature of the background suppression problem: in plot (c), a gamma ray source will trigger an alarm by a significant margin if there is no background suppression (the vehicle is made of low density material e.g. fiberglass), whereas the same

source fails to trigger an alarm if the PBS is -30%. While this illustrates an extreme case, the detrimental effect of background suppression on gross count algorithms is clearly seen.

A key requirement for an advanced detection algorithm is its ability to mitigate the effects of background suppression. The energy windowing algorithms studied in this work are one such set of algorithms that diminish the impact of vehicle's shielding of PVT panels from background radiation (Ely, et al. 2006) (Lo Presti, et al. 2006). Prior to the explanation of how energy windowing accomplishes this, a detailed description of energy windowing itself is necessary.

Recall that gross count algorithms compare occupied count rates to background count rates in various windows and alarm if the occupied signal is above a calculated threshold. Energy windowing algorithms, rather than relying solely on gross count rates, compare the occupied signal shape to the shape of the background spectrum within certain windows. Energy windowing has the advantage of enhanced NORM discrimination and, as suggested in this paper, the potential to identify specific selected radionuclides.

Energy windowing's NORM discrimination advantage exploits the similarity between NORM source gamma ray spectra and background gamma spectra. By definition, NORM sources are naturally occurring radionuclides with the most commonly cited sources as ^{40}K , ^{232}Th , ^{226}Ra , and other uranium- and thorium-series nuclides (Ely, et al. 2006). Since they are natural, it is logical that a background radiation spectrum would mimic the shape of NORM source spectra. If a vehicle carrying a NORM source passes through an RPM, the energy windowing algorithm recognizes the similarity in shape between the occupied and

background signal and reports to the portal monitor operator that the radiation source is like NORM.

Man-made radiation sources, including radiopharmaceuticals the SNM sources ^{235}U (HEU) and ^{239}Pu (WGPu), have spectra dominated by low energy gamma radiation, which is in stark contrast to NORM sources which have a preponderance of high energy gamma rays. Energy windowing algorithms detect that an SNM enhanced spectra differs in shape from the background spectra, thus alarming and reporting to the portal monitor operator that a likely SNM source is inside the vehicle.

The specific details of energy windowing algorithms vary greatly depending on the specific equation used, but all algorithms utilize the same basic mechanics. A typical example, and the most common implementation, utilizes a two window system, where the entire energy spectrum is divided into two windows: a low energy window and a high energy window. The low energy window extends from the low level discriminator (set at the detector system's detection threshold) to a channel slightly above a selected nuclide's Compton edge (since the intrinsic resolution of PVT is poor, the Compton edge smears out and does not sit in a specific channel) (Ely, et al. 2006). The placement of the upper edge of the low energy window optimizes the statistical precision of the alarm algorithm (Ely, et al. 2006). The high energy window consists of the remaining portion of the energy spectrum.

As a vehicle passes through the RPM, counts are sorted into either the low energy or high energy windows. These values are then compared to background data which has also been sorted into the same windows. The occupied values are then compared to the background values to make alarm decisions (Ely, et al. 2006). One of the differences

between energy windowing algorithms is the method used to compare the occupied signal to the background signal. The three most commonly used equations are presented below.

The first equation is the simplest of the three approaches considered in the past, doing a direct comparison of the low energy window to the high energy window (Ely, et al. 2006):

$$R_{EW} = \frac{S_L}{S_H} \quad \text{and} \quad R_B = \frac{B_L}{B_H} \quad \text{where :}$$

$$R_{EW} = \text{Occupied ratio}$$

$$S_L = \text{Occupied count rate in low energy window}$$

$$S_H = \text{Occupied count rate in high energy window} \quad (2.6)$$

$$R_B = \text{Background ratio}$$

$$B_L = \text{Background count rate in low energy window}$$

$$B_{BH} = \text{Background count rate in high energy window}$$

This occupied ratio is compared to the background ratio, and if R_{EW} is greater than some threshold value above the background ratio, an alarm is sounded. The threshold comparison can also take many forms, but most often uses the same threshold style as gross count algorithms (Ely, et al. 2006):

$$\text{ALARM if: } R_{generic} > R_B + K\sigma_{R_B} \quad \text{where :}$$

$$R_{generic} = \text{Any generic occupied EW ratio} \quad (2.7)$$

$$K = \text{Sigma multiplier}$$

$$\sigma_{R_B} = \text{Standard deviation of } R_B$$

The equation for σ_{R_B} varies depending on the form of R_B and can be calculated using propagation of errors. The standard deviation, as expected, depends only on the background count rates.

An alternative form of the ratio in equation (2.6) uses the sum of the count rates in all windows as the denominator rather than only the high energy window (Ely, et al. 2006):

$$R'_{EW} = \frac{S_L}{\sum_{i=1}^n S_i} \quad \text{and} \quad R_B = \frac{B_L}{\sum_{i=1}^n B_i} \quad \text{where:}$$

R'_{EW} = Occupied adjusted ratio
 S_L = Occupied count rate in low energy window
 S_i = Occupied count rate in window i
 R_B = Background adjusted ratio
 B_L = Background count rate in low energy window
 B_i = Background count rate in window i
 n = Number of windows

(2.8)

The final equation presented varies in form from the previous two and is the most commonly used among RPM vendors. It takes the form of a “compensated ratio” (Ely, et al. 2006):

$$R_C = S_L - \frac{B_L}{B_H} \cdot S_H \quad \text{where:}$$

R_C = Compensated Ratio
 S_L = Occupied count rate in low energy window
 S_H = Occupied count rate in high energy window
 B_L = Background count rate in low energy window
 B_H = Background count rate in high energy window

(2.9)

In this formula, the occupied count rate in the high energy window is scaled by the ratio of low energy to high energy window background count rates. If the occupied signal mimics the shape of the background signal, the $\frac{B_L}{B_H}$ term scales the S_H term to a value close to S_L , thus making $R_C \approx 0$. Low energy counts result in a positive R_C value, whereas high energy counts cause R_C to be negative.

Since only one “ratio” is calculated when the compensated ratio equation is implemented, the threshold is calculated using the standard deviation of R_C . The derivation is predicated on two basic assumptions; the first is that, if no source is present during a vehicle scan, $S_H = B_H$ and $S_L = B_L$. The effect of background suppression is not taken into account because the effect varies wildly from vehicle to vehicle (Lo Presti, et al. 2006). The second assumption states that due to the long counting times for background, the fractional uncertainty of the background terms B_L and B_H is negligible (Blessinger 2010). Given the first assumption and recalling equation (2.4), the standard deviations for S_L and S_H are:

$$\sigma_{S_L} = \sqrt{B_L} \quad \text{and} \quad \sigma_{S_H} = \sqrt{B_H} \quad (2.10)$$

The sum of the squared variances of the R_C equation is:

$$\sigma_{R_C}^2 = \sigma_{S_L}^2 + \sigma_{\frac{B_L}{B_H} \cdot S_H}^2 \quad (2.11)$$

Recalling the second assumption that the fractional variances of the background terms are negligible, equation (2.11) reduces to:

$$\sigma_{R_C} = \sqrt{\sigma_{S_L}^2 + \left(\frac{B_L}{B_H}\right)^2 \sigma_{S_H}^2} \quad (2.12)$$

Implementing the results of equation (2.10) and reducing the equation results in:

$$\sigma_{R_C} = \sqrt{B_L + \frac{B_L^2}{B_H^2} B_H} \quad (2.13)$$

$$\sigma_{R_C} = \sqrt{B_L \left(1 + \frac{B_L}{B_H}\right)} \quad (2.14)$$

$$\sigma_{R_C} = \sqrt{B_L} \cdot \sqrt{1 + \frac{B_L}{B_H}} \quad (2.15)$$

Equation (2.15) gives the standard deviation used for compensated ratio threshold calculations (Blessinger 2010).

Studies indicate that there is little to no advantage of using one equation over another due to the overwhelming influence of systematic effects over the minor impact of individual equation statistics during operation (Ely, et al. 2006). A slightly modified version of the compensated ratio has been selected as the equation of choice in this paper due to the convenience of the ratio approaching zero when the background spectrum matches the occupied spectrum, the straightforward calculation of the standard deviation, and the prevalent use by RPM vendors. The development of this new equation is discussed in chapter 4.

Three pragmatic choices need to be made to optimize energy windowing algorithms:

- 1) What regions of energies define windows?
- 2) How many windows are needed?
- 3) What ratio, or combinations of ratios, best enables SNM/NORM discrimination?

Currently the ANSI standards of ^{57}Co and ^{133}Ba are used to set the low energy windows (Ely, et al. 2006). ^{57}Co produces a 122 keV gamma ray and is used as a proxy for HEU's 186 keV gamma. ^{133}Ba produces an 81 keV gamma ray and three gamma rays between 300 and 383 keV; these mimic ^{239}Pu 's series of gamma rays between 330 and 420 keV. New standards will be presented in this paper. This paper suggests that the selection of what gamma energy the low energy window is designed to detect is an immensely important factor in the optimization of energy windowing algorithms. The study finds that improperly selected low energy windows can destroy an algorithm's ability to detect certain SNM sources.

The number of windows depends on the number of nuclides one wishes the algorithm to detect. However, increasing the number of windows dilutes the ability of the algorithm to discriminate between radiation sources; between two and five windows appear to be practical numbers (Ely, et al. 2006).

With the ability to test multiple windows for high count rates within a single alarm algorithm, developers are able to create logic tables using combinations of windows to make accurate alarm decisions. The overall effectiveness of an algorithm largely depends on the individual developer's selections in regard to the three design choices mentioned above. The

logic table and major design decisions for the energy windowing algorithm developed in this paper are presented in chapter 4.

One of the major strengths of energy windowing algorithms is their ability to ignore the effects of background suppression. If one considers what background suppression is, it becomes obvious why energy windowing seems immune to it. Energy windowing compares the shape of an occupied spectrum to the shape of a background spectrum. If a vehicle is occupying the monitor, background radiation, primarily radiating from the ground, is attenuated in the vehicle and cargo, depressing the occupied signal strength below the signal levels of the background spectrum.

With a gross count algorithm, this artificially raised alarm thresholds and severely impacted alarm sensitivity. For energy windowing, this suppression of the occupied signal only lowers the magnitude of the count rates in the high and low energy window, but not the ratio between them. In other words, since a suppressed, occupied spectrum has the same shape as a background spectrum, the effect of background suppression on hurting alarm sensitivity is not seen. Reconsider equation (2.9), which is similar to the equation selected for this paper's algorithm:

$$R_C = S_L - \frac{B_L}{B_H} \cdot S_H$$

Take a background measurement and an occupied measurement with no background suppression: $B_L = 30$, $B_H = 15$, $S_L = 30$, and $S_H = 15$. The compensated ratio value is thus $R_C = 0$. Now assume that the occupied signal is background suppressed by 20% per window, giving occupied window values of $S_L = 24$ and $S_H = 12$. Notice that even though

the occupied count rates are depressed, the *ratio* of count rates and thus the compensated ratio does not change: $R_C = 0$. This example exhibits how background suppression is avoided by using energy windowing.

In order for this to be the case, it is assumed that high and low energy gamma rays attenuate virtually the same amount through a cargo carrying vehicle. This assumption appears to be valid in some studies (Ely, et al. 2006), but it is tested and verified in this paper as well.

For energy windowing ratios, the metric used to measure background suppression is different from the PBS values for gross counting. This value, called the percent energy window ratio suppression (PRS), is calculated by the formula (Lo Presti, et al. 2006):

$$PRS = \frac{(OccupiedRatio - BackgroundRatio)}{BackgroundRatio} \cdot 100 \quad (2.16)$$

If it is assumed that there is no source in a vehicle, a positive PRS indicates that the occupied ratio is higher than the background ratio and is artificially decreasing alarm thresholds; this can result in increased false alarms and is not desired. A negative PRS implies that the occupied ratio is being depressed and is artificially raising alarm thresholds; this decreases alarm sensitivity and is also not desirable. Since the compensated ratio does not have an occupied ratio and a background ratio, it is not possible to calculate a PRS value for compensated ratio energy windowing algorithms.

The assumption that gamma rays are attenuated equally across all energies does not imply that gamma rays are attenuated equally in all directions. In fact, pillars consisting of top and bottom PVT panels show different suppression ratios between panels. Top panels

tend to show slightly positive PRS values, indicating that a higher proportion of low energy counts are reaching the panels during an occupancy than during the background collection period. The converse is true of lower panels; the PRS values tend to be slightly negative, implying that a higher proportion of high energy counts are registered during an occupancy than during background collection. The important point is that the slightly positive and slightly negative PRS values sum to numbers very close to zero, proving that background suppression does not affect ratio calculations for energy windowing algorithms (Lo Presti, et al. 2006).

The likely cause of PRS differences between top and bottom PVT panels is differential scattering related to the paths of gamma rays through the cargo (Lo Presti, et al. 2006). Top panel gammas are differentially scattered to lower energy levels due to the longer path length through the cargo from the ground (where most background gammas originate) to the panel, resulting in proportionally more low energy gamma rays than high energy gammas being counted; this results in a higher compensated ratio being computed and positive PRS values. Lower panels have shorter path lengths through the cargo between the ground and the panel, thus fewer gamma rays are differentially scattered to low energies; therefore proportionally more high energy gamma rays are detected, lower compensated ratios are calculated, and negative PRS values result.

Figure 2.5 illustrates an example of energy windowing background suppression and the effects of differential scattering on top and bottom panel PRS values (Lo Presti, et al. 2006).

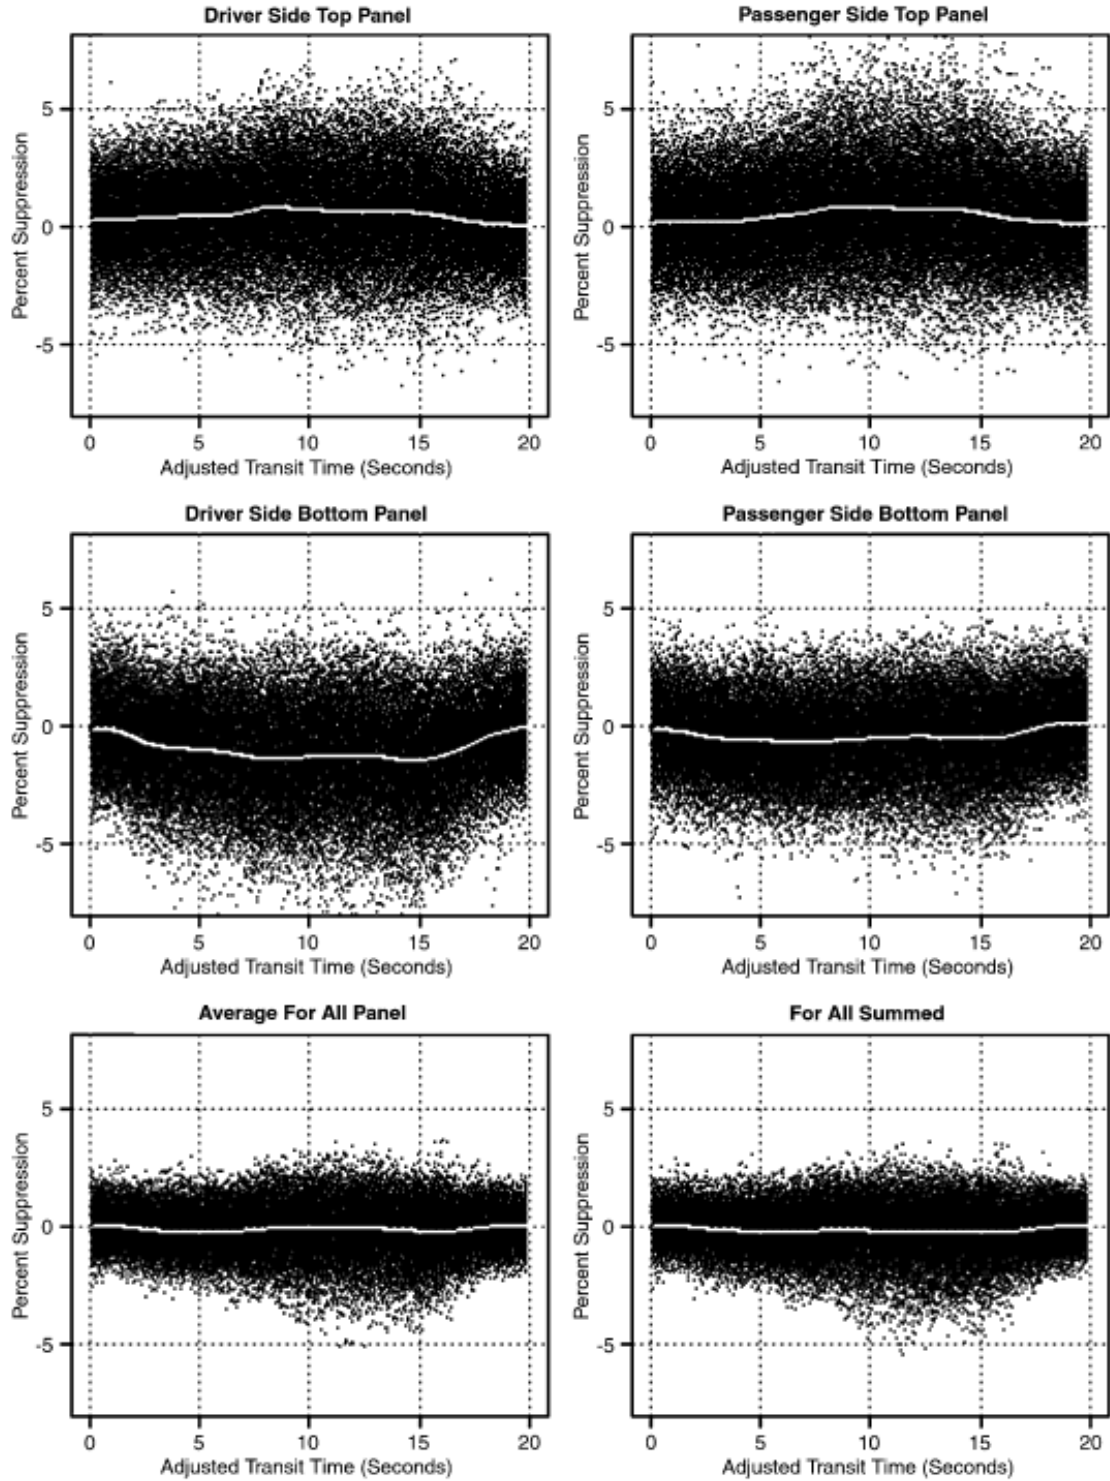


Figure 2.5: Profiles of PRS for top panels, bottom panels, average and sum of all panels (Lo Presti, et al. 2006)

While energy windowing algorithms seem to solve the background suppression problem, they are subjected to another phenomenon that could potentially severely hamper their effectiveness. This is the issue of high energy masking. When a high energy gamma source is present in a cargo, the source could skew the occupied ratio to a low value due to increased high energy counts. If a low energy source like SNM is present in the cargo, then the high energy source could “mask” the SNM and prevent an energy window alarm from triggering. This is easily explained using the compensated ratio equation (2.9):

$$R_C = S_L - \frac{B_L}{B_H} \cdot S_H$$

Assume the background count rates are $B_L = 30$ and $B_H = 15$. Now say an SNM source is present in the vehicle with non-radioactive cargo as well. This gives occupied window values of $S_L = 60$ and $S_H = 12$. Calculating the compensated ratio gives the value $R_C = 36$. Now say the background value is the same, but this time a high energy NORM source is in the cargo, giving occupied values of $S_L = 60$ and $S_H = 30$. Even though an SNM source is in the vehicle, the compensated ratio is $R_C = 0$. While this example is extremely simple, it accurately illustrates the masking problem.

This paper suggests that the mechanism to solve masking problems is to design a combined energy window/gross counting alarm algorithm. This algorithm utilizes energy windows to detect the most common NORM sources in addition to SNM and RDD sources, combined with gross counting windows to prevent any high activity source from passing unmolested through an RPM.

Masking is predicated on the assumption that neither the low energy nor high energy source, nor the sum of the sources, increases the count rate above a gross count alarm level. Otherwise, the vehicle would still be directed to secondary inspection where it is likely that the offending material would be detected.

With the literature review concluded, this paper will now detail the scope of this work and focus strongly on the methods by which the analysis was performed.

Chapter 3 – Scope of Work, Methods, and Approaches

The overarching goal of this research work is to propose an advanced radiation portal monitor alarm algorithm and fully explain the reasoning behind its development. Chapter 3 outlines the scope of the project while revealing the methods and approaches used to create the final product. Prior to the creation of an alarm algorithm, copious amounts of real-world data collected from a deployed RPM system needed to be sorted through; this data would provide the bedrock for the creation of the alarm algorithm.

Oak Ridge National Laboratory provided data from an RPM experiment at a United States port-of-entry. This data, called the POE data, consisted of information from 2129 vehicles passing through PVT RPMs along with dozens of calibration source data files. Twelve of the 2129 runs caused some sort of alarm; all the alarming vehicle carried innocent cargos. This installation tested an energy windowing algorithm, but the sensitive nature of this test precludes this study from knowing the specific detail of the algorithm. The system employed a 512 channel MCA with an energy range stated to be 0 keV to 3000 keV. The energy interval in each channel is assumed to be a constant 5.86 keV/channel. Thus, channel 0 (the first channel) collects gamma rays ranging from 0 keV to 5.86 keV, channel 1 collects gamma rays between 5.86 keV and 11.72 keV, etc. This constant energy interval can also be referred to as a linear energy response.

Each of the 2129 vehicle runs produced three data files: a .DAT file, a .PRO file, and a .HST file. The .DAT files are also known as data files; they contain window data, background accumulation time, background window counts, and vehicle scan data. The

vehicle scan data consists of window counts and time stamps. A Perl script was written to extract the background accumulation time and total vehicle scan time from the data files. Due to the unique formatting of the data files, this script is only of use for the POE test's .DAT files.

The .PRO files are vehicle profile files. These files contain raw vehicle time profile data. After five header lines, each subsequent line contains 513 space-delimited columns; column one contains the 0.1 second interval sum of all channel's counts, and columns two through 513 contain each MCA channel's individual counts. Several of these profiles were removed manually from the .PRO files, converted from 512 channel data into 256 channel data, and pasted into an MS Excel file. This Excel file is used as a data source for the simulated real-time alarm algorithm emulator developed in Simulink. This emulator will be discussed later in the chapter. The calibration sources were not driven through the RPM, therefore do not contain .PRO files. These sources were multiplied by a "spatial response function" to transform them from a static, count rate vector (1x256) into a matrix of time interval counts (???x256). The spatial response function will be explained in detail in chapter 4.

.HST files are histograms of the vehicle runs. For each of the 2129 runs, the histogram files contain 512 lines, one for each channel, containing: 1) the number of background counts in the given channel accumulated over the background counting period, and 2) the occupied counts in the given channel accumulated over the occupied time. These accumulated spectra were used to test three different energy windowing algorithms, two of which come from commercial RPM suppliers and the one developed in this paper. A Perl

script was written to extract the spectral information from the .HST files and background/occupied time from the .DAT files to calculate the average count rate per channel for each of the 2129 runs. This script also applied the three energy windowing algorithms, along with two additional gross count algorithms, to these count rate spectra. The Perl script output data into a text file specifically formatted for easy import into an Excel spreadsheet. The Perl script not only tested the algorithms on actual commercial traffic, but also allowed for a large statistical sample to be drawn for statistical studies. The large amount of data was easily extracted from the Excel file and read automatically into the computer algebra system Mathcad where further statistical analysis was performed. The summary statistics of the different ratio values from each energy windowing algorithm were calculated using built-in functions in Mathcad. The built-in functions calculate mean, median, variance, standard deviation, skew, and kurtosis instantly.

An additional Perl script was written to calculate the background suppression per channel for each of the non-alarming runs; the alarm runs are outliers and would skew the data. The average percent background suppression (PBS), and slope and y-intercept of the background suppression versus channel curve is calculated by the script. Due to the limited number of counts occurring in higher channels the script only computes the above values for channels between 0 and 275 inclusive (recall the data is originally in 512 channel format). As examined in the spectra, occupied signals start to become extremely noisy between channels 250 and 300 because the low occurrence of high energy gamma rays and the small probability of interaction within the plastic scintillator. Extending the average, slope, and intercept calculations above this region drastically reduces the usefulness of the data.

The slope of the spectrum shows how the background suppression varies as gamma ray energy increases; this tests the previous assumption that background suppression is relatively uniform across all channels. The y-intercept of the spectrum gives the background suppression of channel 0, which is likely to be the most suppressed channel due to the fact that, while the slope is suspected to be small, it is also expected to be positive. The average PBS of each run is simply the average background suppression for each run; given that the PBS does not vary a great deal across channels, this value can be used as the token PBS value for each run. The values for the average PBS and the y-intercept should be very close to each other if the assumption is true that the background suppression does not decrease as gamma ray energy increases.

Perl was selected to do these file searches and computations because of its ease of use and versatility. Perl is extremely effective at processing large amounts of data and at searching through oddly formatted text documents for specific pieces of information. Fortran was also considered for use due to its familiarity and speed, but ultimately Perl's ability to search through hundreds of differently named files and extract specific lines contained within the files resulted in its selection as the data extraction tool. Since the calculations that are computed are simple algebraic problems, Perl is sufficient to simultaneously extract the data from the various file formats and generate the calculated data. While it was possible to integrate both the background suppression script and the energy windowing algorithm script into a single program, it was decided that two separate programs sharing much of the same code was a more efficient programming approach. Each task could then be handled individually and code manipulated separately without interfering with the other program.

While the POE data used a 512 channel data acquisition system, the two commercial energy windowing algorithms tested in this paper used 256 channel systems. Excel spreadsheets were used to quickly convert the 512 channel spectra extracted from the raw data files into 256 channel spectra. Calibration source data was extracted manually from the .HST files and inserted into spreadsheets as well. The calibration sources include ^{133}Ba (the ^{239}Pu proxy), ^{57}Co (the HEU proxy), ^{60}Co (a depleted uranium proxy), and ^{137}Cs (a possible RDD source). The calibration sources are used as “injection sources”, which means they are added into cargo spectra to simulate the effect of a source being carried by a vehicle. Sample cargo spectra from innocent POE runs can have the injection sources added to their occupied signal spectrum and tested with various alarm algorithms. The twelve alarming spectra from the POE dataset also serve as templates for algorithm testing since the alarming material was identified in secondary inspection. Two of the alarming vehicle carried ^{40}K -laden innocent cargos, and two carried granite cargos that contain ^{232}Th and ^{226}Ra . The non-alarming spectra, innocent alarming spectra, and injection sources serve as the dataset used during the entire alarm algorithm development process.

Excel spreadsheets also served as convenient platforms to create plots and occupied and background spectra. Excel is very easily integrated into both MATLAB/Simulink and Mathcad, and it provides a very simple method to format data and create plots; there was little debate as to what program would serve as the database for the higher level programs later implemented.

Recall that the .PRO files contain 2129 vehicle time profiles. The data from these files was removed manually and written into Excel files to be used in the simulated real-time

alarm algorithm emulator. A problem existed however, that the injection sources, being derived from the calibration sources, did not have corresponding .PRO files and thus could not be utilized in the emulator. As mentioned above, a “spatial response function” was used to create these source time profiles. The accumulated source spectra contained within the .HST files were converted into count rate vectors, then transformed using the spatial response function into time profiles which were injected into vehicle profiles found in the .PRO files of the POE dataset.

Five non-alarming and four alarming POE accumulated vehicle spectra were utilized for the source-injected cargo profiles. The six injection sources were: 1) an HEU source measured at Oak Ridge National Laboratory, 2) a ^{57}Co calibration/injection source, 3) a ^{133}Ba calibration/injection source, 4) a ^{60}Co calibration/injection source, 5) a ^{137}Cs calibration/injection source, and 6) a ^{40}K source created by subtracting the background counts from a ^{40}K -laden vehicle occupied spectrum. The injection sources, being count rate spectra, were then multiplied by the spatial response function and transformed from 1×256 vectors into $N \times 256$ matrices (where N is the number of rows depends on the corresponding vehicle profile occupied time). Each of the six injection source profile matrices were then added to each of the five non-alarming and four alarming vehicle profile matrices, resulting in 54 source-injected vehicle spectra. These profiles were of the same format as the profiles extracted from the .PRO files and were also written into the same Excel file used as the data source for the real-time alarm emulator.

The first step in the creation of the new alarm algorithm is to select which nuclides the algorithm seeks to detect. Recall that if too many sources are selected, the algorithm

loses sensitivity because of Compton edge overlap (Ely, et al. 2006). Using expert opinion from ORNL staff, four sources were chosen to have low energy windows calibrated specifically to detect them. These sources were the two SNM concerns ^{235}U and ^{239}Pu (HEU and WGPu, respectively), the RDD source ^{137}Cs , ^{60}Co which is both an RDD concern and a proxy for depleted uranium (used as a reflector in certain nuclear weapon designs), and the NORM source ^{40}K . The NORM sources ^{232}Th and ^{226}Ra increase count rates across the entire spectrum consistently, therefore a single window cannot be assigned to detect them. However, since their spectra shape largely mimics that of the background spectrum, no energy window alarms should trigger; rather a gross count alarm should occur.

Starting with the lowest energy nuclides, the upper edge of the low energy windows was set using one of two window calibration approaches. The first approach was a subjective visual test; a non-alarming cargo was injected with the chosen source and the count rate was plotted on the same graph with the corresponding background spectrum. The upper edge of the window was then set at the point on the graph where the Compton edge of the source drops down to the level of the background count rate. This process was continued for each of the nuclides moving from low energy to high energy.

The second approach is an objective formulaic approach comparing the signal to the background created by scientists at Oak Ridge National Laboratory. Figure 3.1 reveals the rationale behind the formula. It is desired that a given source will not fall within the variance of the background distribution, so the function below is maximized:

$$M(E) = \frac{\int_0^E S(E) dE}{\sqrt{\int_0^E B(E) dE}} \quad \text{where:}$$

$S(E)$ = Source count rate as a function of energy (channel) (3.1)
 $B(E)$ = Background count rate as a function of energy (channel)

Since RPM data is in discrete channel form, this equation is converted into a summation:

$$M_i = \frac{\sum_{k=2}^i S_k}{\sqrt{\sum_{k=2}^i B_k}} \quad \text{where:}$$

S_k = Source count rate in channel k (3.2)
 B_k = Background count rate in channel k

If the sum source spectrum count rate from channel 0 up to a given channel is close to the value of the sum background variance over the same range, the likelihood of the source being detected in that channel is small; essentially the source signal in this range is overwhelmed by the variance of the background. On the other hand, if the source count rate is significantly higher than the background variance in a given range, this window effectively “sees” more source counts than the background’s variance can “cover up”.




Figure 3.1: Interaction of a source with a background spectrum

In Figure 3.1, it can be seen that as S decreases or the variance \sqrt{B} increases, the probability that a source will fall within, say 4σ , of the variance of the background spectrum increases. Thus it is clear why a larger S/\sqrt{B} value is desirable; it optimizes the ability of a window to detect a given source. Looking back to equation (3.2), the low energy window's upper edge is set to the channel i that maximizes the value of M_i . While the sum always starts at $k = 0$, the lower bound of a higher energy window is always set one channel above the upper bound of the preceding window; this prevents window overlap.

To perform the two window calibration methods, a Mathcad worksheet was created. Mathcad was selected over MATLAB for this application because Mathcad allowed the plots to be generated and manipulated in real-time on the screen. Mathcad is a WYSIWYG (what-you-see-is-what-you-get) computer algebra system, so the convenience and speed of replotting, zooming, and tracing the graphs in Mathcad trumped MATLAB's plotting abilities in this case.

To finish the creation of and preliminarily test the newly suggested algorithm, an alarm simulator using post-processed data was created (please note the “alarm simulator” is entirely different from the “alarm emulator” mentioned earlier). Real-time data was not necessary for this application; accumulated occupied and background signals allow the ratio calculations and alarm logic to be tested in a simpler and faster manner. The accumulated data was in a single albeit large Excel file. The simulator was again implemented in Mathcad; it is preferred if the alarm simulator can change vehicle cargo and injection source combinations on the fly while quickly displaying the alarm results on screen, and Mathcad being a WYSIWYG program fits this application well. A MATLAB program could have also accomplished this task as well but was ultimately not selected as most other work was performed in Mathcad. A Perl script was also considered, but it was decided that a Mathcad worksheet would be more flexible to operate. The primary purpose of this alarm simulator was to create the alarm logic that would report to RPM operators the specific nuclide detected.

Using the window calibration Mathcad worksheet and the alarm simulator, the advanced alarm algorithm was fully developed but still required testing to determine whether it was an improvement over the commercial algorithms. Two programs remained to be created: a program to test the algorithm’s sensitivity to masking/overall effectiveness, and the simulated real-time alarm emulator. Masking was mentioned in chapter 2 as a major concern for energy widening algorithms, so it was vital to create a program to effectively test this particular algorithm’s sensitivity and compare it to commercial algorithm sensitivity.

The NORM masking sensitivity analysis program was written in Mathcad again due to the WYSIWYG format. This allowed source changes to be made on the fly, had the results of integrals listed on the screen in addition to generating and displaying the appropriate plots. MATLAB generated graphs in a different window and was more difficult to switch injection sources and cargos.

The sensitivity analysis program compares four algorithms simultaneously: 1) a commercial algorithm designed to detect only HEU using two windows, calibrated using ^{57}Co , 2) a commercial algorithm designed to detect ^{239}Pu and HEU using two windows, calibrated with ^{133}Ba , 3) the advanced alarm algorithm developed in this paper, using multiple windows and calibrated using multiple sources, and 4) a full spectrum gross count algorithm. The advanced algorithm, since it utilizes five windows, can actually produce up to five different sensitivity curves on a single plot; in other words, there is one curve per window.

The program plots the source count rate (as perceived by the detector) required to trigger each of the alarms given a certain cargo count rate. The cargos and sources can be changed so that the sensitivity of the algorithms to some source can be evaluated when the source is present in a NORM-laden cargo. Any point above the gross count sensitivity curve triggers a gross count alarm and any point above the energy windowing sensitivity curve triggers an energy window alarm. Any point below both lines fails to trigger an alarm.

The general shapes of the curves can be discovered through a thought experiment. First consider an energy windowing curve for an HEU source hidden in a ^{40}K -laden cargo like fertilizer. At zero cps cargo count rate ($x = 0$), an energy window alarm will trigger at

some positive source count rate, implying that $y > 0$. As the count rate of the cargo increases ($x > 0$), the energy window alarm will require more counts of HEU to alarm due to the high energy masking effect discussed in chapter 2. Thus, an energy windowing algorithm sensitivity curve is positive slope with a positive y-intercept if there is a high energy NORM cargo. If the cargo is not NORM laden, the energy windowing curve could actually have a very small negative slope, implying that the algorithm gets slightly more sensitive as the cargo count rate increases.

Now consider the gross count algorithm curve. At zero cps cargo ($x = 0$), a high HEU count rate is required for a gross count alarm, implying the curve has a large, positive y-intercept equal to the gross count threshold count rate. At zero cps HEU ($y = 0$), a large amount of ^{40}K is required to trigger the alarm, giving a large, positive x-intercept equal to the gross count threshold count rate. Therefore, a gross count sensitivity curve is negative sloped with positive x- and y-intercepts equal to the gross count threshold count rate. Since a full spectrum gross count algorithm is energy independent, changing the source and cargo will only change the magnitudes of the intercepts because the different background spectrums.

It is assumed that all three energy windowing algorithms are used in conjunction with the gross count algorithm (not necessarily the case). This means the area under the curves can be considered a metric for comparing the alarm sensitivities and thus overall algorithm effectiveness. The figure below shows an example sensitivity plot.

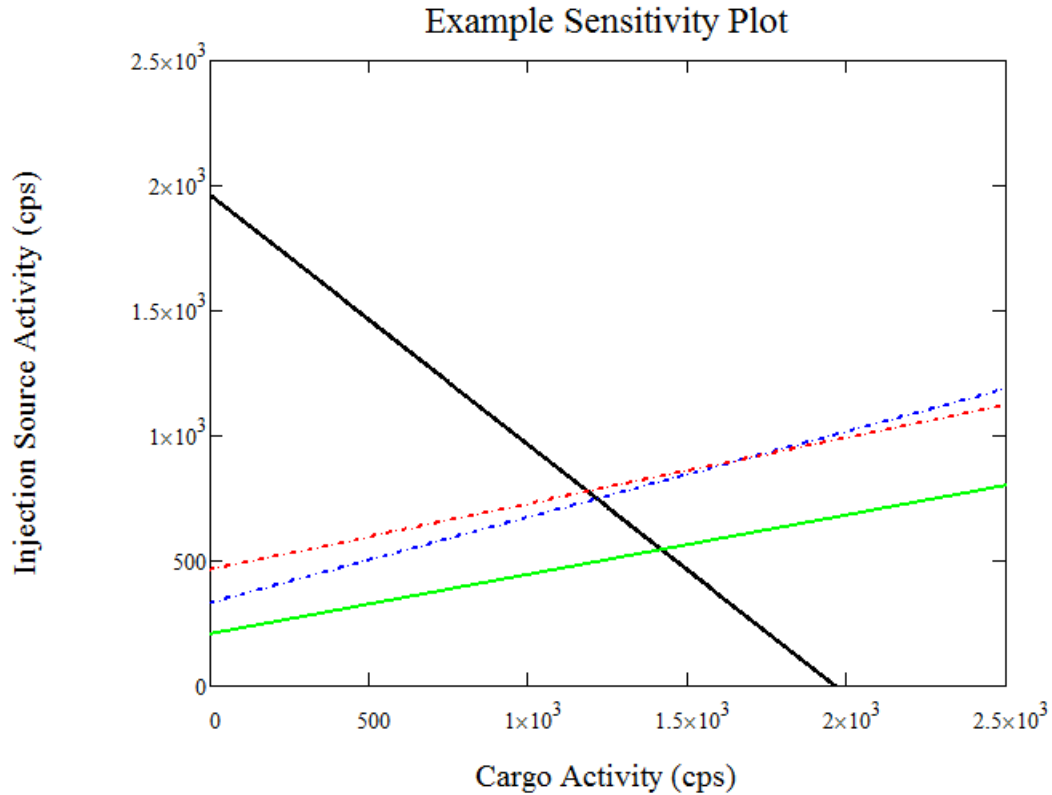


Figure 3.2: Example sensitivity plot

Black solid line: Gross count sensitivity curve

Green dashed line: Most sensitive EW algorithm

Blue dotted line: Second most sensitive EW algorithm

Red dot-dashed line: Least sensitive EW algorithm

The final assignment was to develop the simulated real-time alarm emulator. This task was performed in Simulink, a modeling/simulation environment within the MATLAB program. It is a graphical block programming tool that allows users to create system simulations using both customized function blocks and built-in blocks. This format is convenient because it fosters the creation of accurate system models that allow the user to see the logical flow of the information rather than running code through a black box. While it could be argued that individual Simulink blocks are black boxes, the blocks actually serve

very specific purposes and are transparent in their implementation. The Simulink model, with the proper software additions, could be converted into actual code to run true real-time simulations. The model, being visual in nature, also allows for the efficient transition from model to electronic implementation.

The Simulink model is referred to as “simulated real-time” because the data is not truly read into the emulator in real-time, but in a simulated real-time environment. This does not affect the results whatsoever since this simulated real-time is effectively accelerated real-time, meaning the software simply runs the code faster than a real RPM system would make alarm decisions. It is referred to as an “emulator” because the nature of an emulator is software being run by other software; the advanced alarm algorithm software model is being run in the Simulink environment. While this model may not be a true emulator, the reference to the model as an emulator serves to highlight the closeness between the Simulink model and an actual software implementation. It also differentiates it from the alarm algorithm simulator created in Mathcad.

As discussed prior, the emulator reads in the vehicle profile data from an Excel file. The background signal and occupied signals are read in in simulated real-time and processed just as the actual software would perform the assignment. Figure 3.3 shows the actual emulator in the Simulink environment. Examining the figure reveals the data buffer, the windowing functions, the ratio calculations and threshold tests, the alarm logic, and the alarm display.

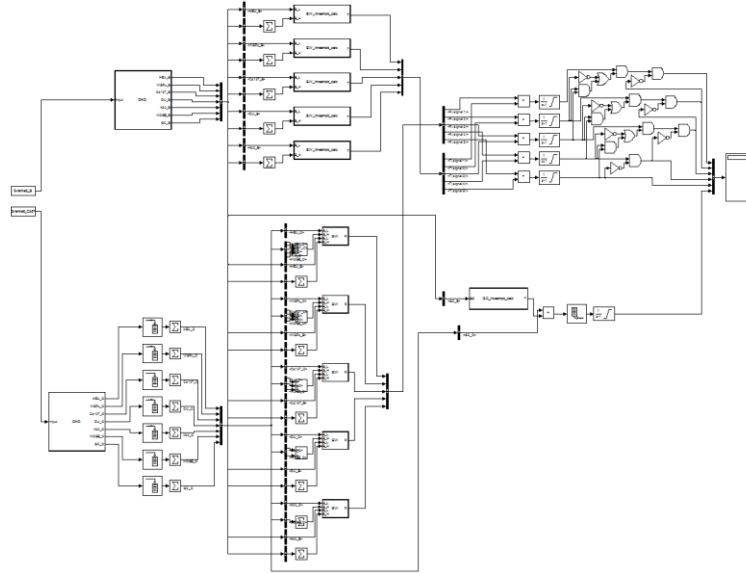


Figure 3.3: Simulated real-time alarm emulator diagram

The emulator dataset has 71 vehicles it can examine: 54 of them are source-injected alarming vehicles created using the spatial response function, 5 are non-alarming vehicle profiles from the POE data, and 12 are the alarming vehicle profiles from the POE data. After typing the vehicle name in the background and occupied data input blocks, the simulation can be started. It takes very little time (less than one second on average) to make the alarm decisions, and they are displayed in a display block directly on the Simulink screen. This allows for fast analysis of the vehicles without switching between windows.

Thus concludes chapter 3; the tasks necessary to create the advanced alarm algorithm have been discussed. Chapter 4 discusses the actual details of the algorithm and addresses the results of the window optimization and reveals the how the algorithm was optimized. The masking sensitivity will be compared between the three energy windowing algorithms as well. The experimental validation of the advanced alarm algorithm by the simulated real-time emulator will conclude chapter 4.

Chapter 4 – Algorithm Development and Analysis

The algorithm development process began with an examination of the two commercial algorithms that information was provided for. Both algorithms utilize the compensated ratio formula first seen in equation (2.9) and a basic, two-window design:

$$R_C = S_L - \frac{B_L}{B_H} \cdot S_H$$

All algorithm designs are based on a 256 channel data acquisition system with an energy range of 0 keV in channel 0 up to 3000 keV in channel 255; this results in each channel width of approximately 11.72 keV. To define the windows, the endpoints of the windows will be stated with the endpoints being inclusive.

With the electronic setup used in current RPMs, the gain is adjusted to shift the spectrum into the appropriate range so that channel 255 actually corresponds to 3000 keV. As gain increases, electronic noise is introduced into the spectrum. For the 0 to 3000 keV energy range, channels 0 and 1 are very noisy and have artificially increased count rates due to the gain. For this reason, neither the commercial algorithms nor the advanced algorithm utilizes channel 0 or channel 1 in either gross count or energy windowing calculations.

The first algorithm was calibrated using a ^{57}Co source, thus the low energy window is optimized to detect ^{57}Co , the HEU proxy. Henceforth, this algorithm will be known as the ^{57}Co algorithm. While it uses a two window design, the windows actually overlap; independent examination has not determined why the algorithm uses an overlapping design. The low energy channel contains windows 2 through 12; the high energy window covers channels 6 through 24. The alarm threshold for the ^{57}Co algorithm is unknown, so a standard

4σ threshold is used. Since the windows overlap, the standard deviation of R_C is no longer equation (2.15):

$$\sigma_{R_C} = \sqrt{B_L} \cdot \sqrt{1 + \frac{B_L}{B_H}}$$

The standard deviation actually takes the form of a three window system:

$$\sigma_{57Co} = \sqrt{B_1} \cdot \sqrt{1 + \frac{B_1^2 + B_2 \cdot B_3}{B_1 \cdot (B_2 + B_3)}} \quad \text{where:} \quad (4.1)$$

B_1 = Background count rate in channels 2 through 6
 B_2 = Background count rate in channels 7 through 12
 B_3 = Background count rate in channels 13 through 24

The second commercial algorithm is another two window design with the low energy window calibrated and thus optimized to detect ^{133}Ba , the ^{239}Pu proxy. This algorithm will now be known as the ^{133}Ba algorithm. The low energy window sums the count rates between channels 2 and 38, and the high energy window is the sum of channels 39 through 255. Due to the low energy window covering the range of ^{57}Co and ^{133}Ba , this single low energy window is designed to detect both SNM sources. Since these windows are non-overlapping, the standard deviation of the compensated ratio R_C is defined by equation (2.15):

$$\sigma_{R_C} = \sqrt{B_L} \cdot \sqrt{1 + \frac{B_L}{B_H}}$$

Notice the ^{133}Ba algorithm and ^{57}Co algorithm focus on detecting SNM material but do not make efforts to detect and identify NORM sources. This is one area where the advanced alarm algorithm differs.

Rather than creating windows only to detect SNM, this algorithm sets windows to find additional nuclides. These specific sources were mentioned in chapter 3, but the rationale for picking them was not elucidated. The five sources the advanced algorithm is designed to detect, each with their own two window modified compensated ratio equation, are: HEU, WGPu, ^{137}Cs , DU, and ^{40}K .

The selection of HEU as an individual window does not need in depth justification; HEU is a severe nuclear trafficking threat and must be detected if at all possible. Resources at Oak Ridge National Laboratory allowed an actual HEU source to be measured with a PVT RPM. This spectrum is used as an injection source and also used to calibrate the advanced algorithm specifically to detect it. While this opportunity may not be available to commercial RPM producers, the spectrum could easily be distributed and used to calibrate HEU windows. Injecting the HEU source into a non-alarming cargo and performing the S/\sqrt{B} channel optimization reveals the optimal location of the upper edge of the HEU low energy window. Figure 4.1 shows the spectra plots and Figure 4.2 shows the S/\sqrt{B} plot.




Figure 4.1: HEU injected cargo and background spectra

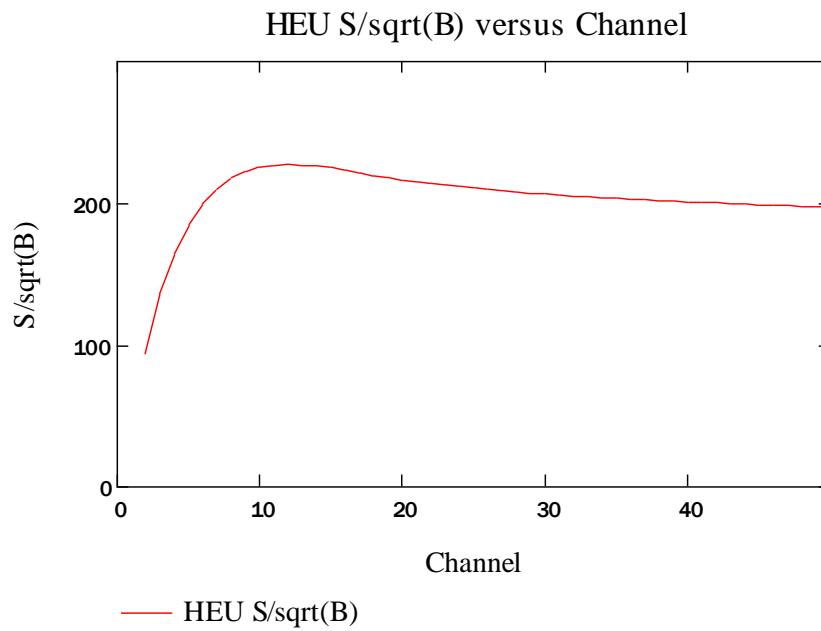


Figure 4.2: S/\sqrt{B} versus channel calibration for HEU

The optimized upper edge of the low energy window was found to be channel 12. The advanced algorithm's HEU low energy window ranges from channel 2 to channel 12. The high energy window contains the remaining channels, 13 through 255.

Weapons-grade plutonium (WGPu) is the other SNM source of concern. A ^{239}Pu source could not be obtained to take a measurement, so a ^{133}Ba source calibration was used to set the WGPu window. The injected cargo spectra and S/\sqrt{B} plot are seen below.

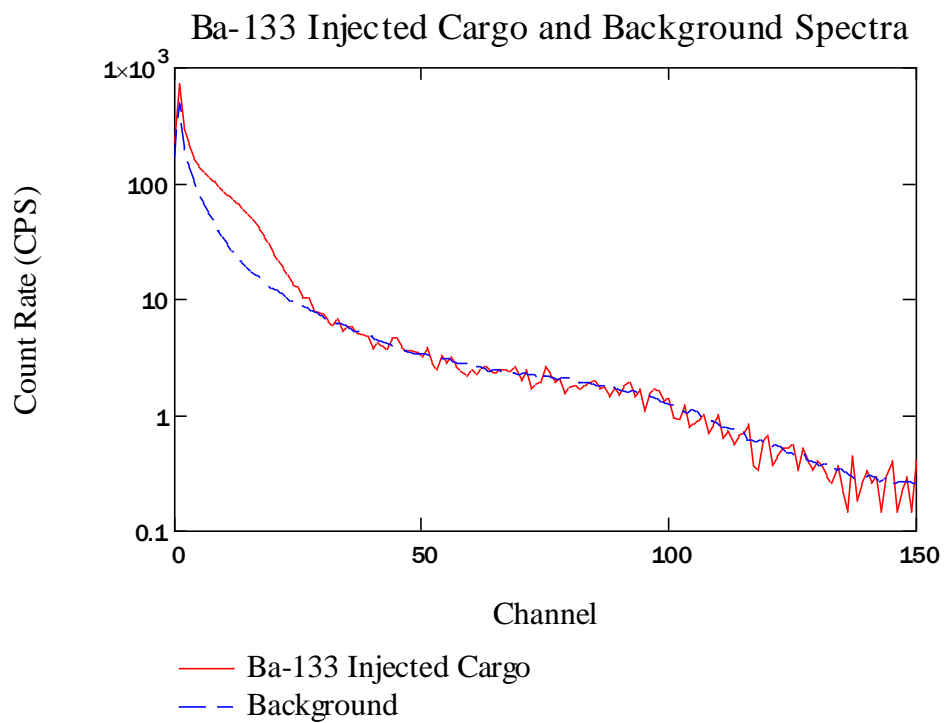


Figure 4.3: ^{133}Ba injected cargo and background spectra




Figure 4.4: S/\sqrt{B} versus channel calibration for ^{133}Ba

The S/\sqrt{B} calibration for ^{133}Ba optimizes WGPu’s low energy window upper edge in channel 25. Due to expert opinion of scientists at Oak Ridge National Laboratory, it was suggested that since the ^{133}Ba proxy produces slightly lower energy gamma rays than ^{239}Pu , the optimum channel should be increased by around 60 keV. This results in the WGPu low energy window upper edge increasing to channel 30. The final low energy window range for WGPu detection is channel 13 to channel 30. The high energy window covers channels 31 through 255.

The HEU low energy window (channels 2 through 12) is not included in the WGPu low energy window (nor does any other window cover the window below it) for one major reason. It is unnecessary to have a window cover multiple nuclides; it “muddies the water” by making a lower energy source, e.g. HEU, set off an alarm in a window designed to detect

a distinct, higher energy gamma producing nuclide, e.g. ^{239}Pu . This would conflate the logic table and render nuclide identification very difficult.

^{137}Cs is a potential RDD threat and produces a 662 keV gamma ray. This gamma ray sits at a sufficiently higher energy than WGPu's 300 keV gammas such that the Compton continuum can be resolved with a designated window. The injected cargo spectra and S/\sqrt{B} curve are seen below.




Figure 4.5: ^{137}Cs injected cargo and background spectra

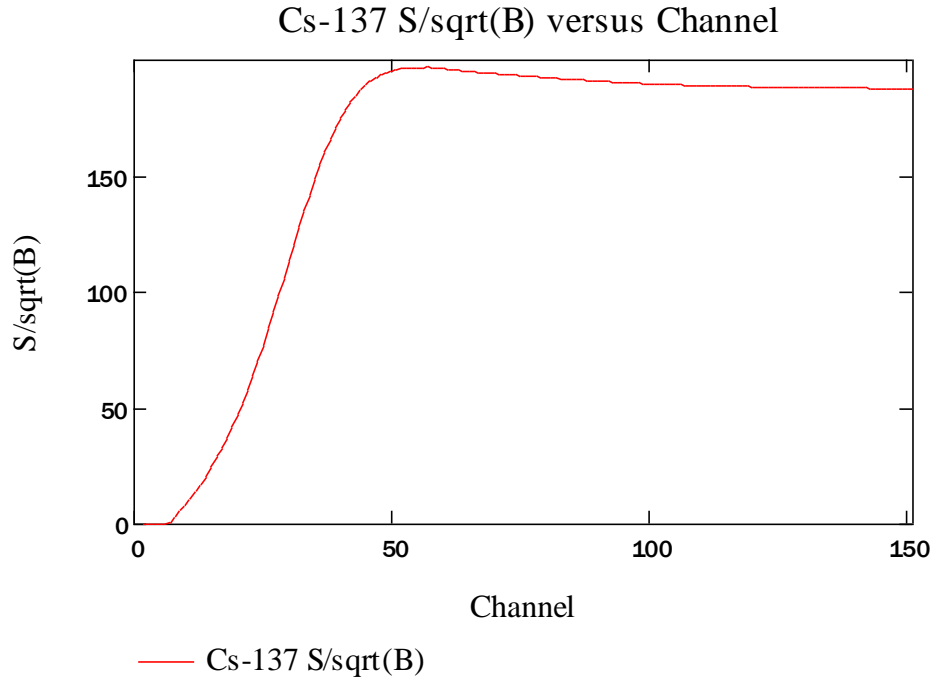


Figure 4.6: ^{137}Cs S/sqrt(B) versus channel calibration

The S/\sqrt{B} calibration reveals the optimum low energy window upper edge to be channel 57. Thus, the ^{137}Cs low energy window is the sum of count rates in channels 31 through 57. The high energy window covers channels 58 through 255.

The fourth designated source is depleted uranium (DU), which uses a ^{60}Co proxy. DU generates a 1001 keV and a 766 keV gamma ray via decay of ^{234}Pa , a decay product of ^{238}U . ^{60}Co generates 1173 keV and 1332 keV gamma rays via decay of ^{60}Ni , the decay product of ^{60}Co . While the energies of the ^{60}Co 's gamma rays are not exactly the same as DU's, their proximity turns out to be sufficient with minimal adjustment. The Compton edge of ^{60}Co 's 1173 keV gamma is 963.2 keV and the Compton edge of DU's 1001 keV gamma is 797.5 keV. Recall that ^{137}Cs 's low energy window ends in channel 57 which corresponds to 668 keV. By making DU's low energy window stretch from the upper edge of ^{137}Cs 's low

energy window up to near the Compton edge of ^{60}Co 's 1173 keV gamma ray, almost the entirety of DU's Compton continuum can be collected. Below are the spectra of the ^{60}Co injection source and the S/\sqrt{B} calibration curve.

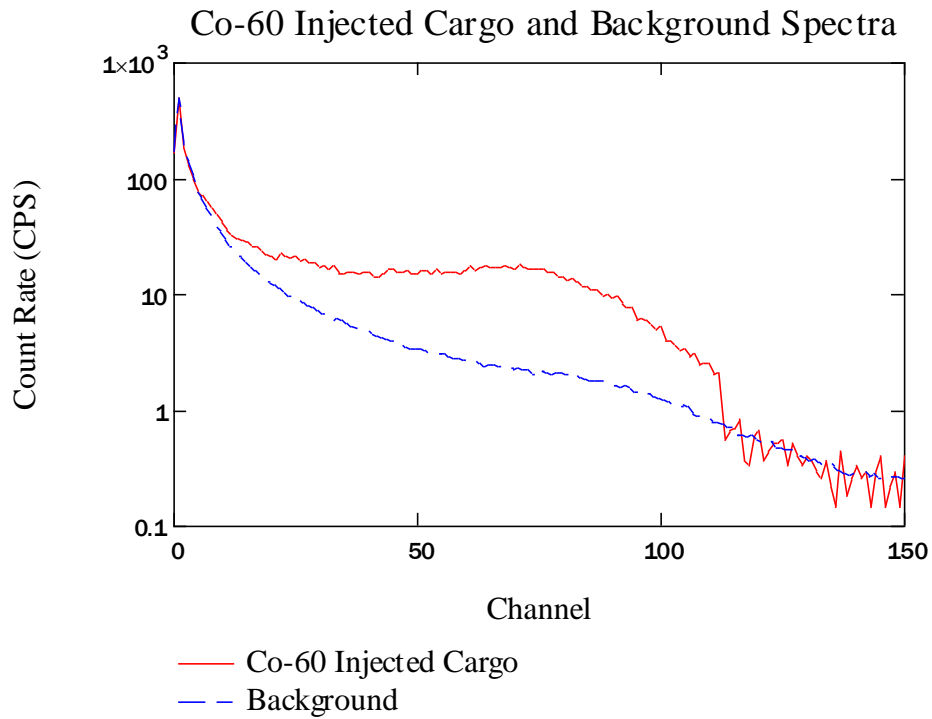


Figure 4.7: ^{60}Co injected cargo and background spectra




Figure 4.8: S/\sqrt{B} versus channel calibration curve for ^{60}Co

The S/\sqrt{B} calibration places the optimum upper edge of a ^{60}Co window at channel 112 (due to the 1332 keV gamma ray). By shifting the upper edge down around 319 keV, the difference between ^{60}Co 's 1332 keV Compton edge and DU's 1001 keV Compton edge, the window better captures DU rather than ^{60}Co . 319 keV corresponds to approximately 27 channels, giving the optimum channel for the DU upper edge to be channel 85. The DU low energy window range is thus channels 58 through 85. The high energy window is subsequently channels 86 through 255.

While all RPM installations will have different NORM cargos passing through, the most common NORM sources are ^{40}K , ^{232}Th , and ^{226}Ra . ^{40}K has a single, higher energy gamma ray (1460 keV) and thus makes a good candidate for setting a window. ^{232}Th and ^{226}Ra decay chains produce several high energy gamma rays and are more difficult to identify

with fine-tuned windows. A single energy window designed to detect ^{40}K should be sufficient since ^{40}K cargos are some of the more common innocent alarming cargos. A ^{40}K injected spectrum and S/\sqrt{B} curve are seen below.

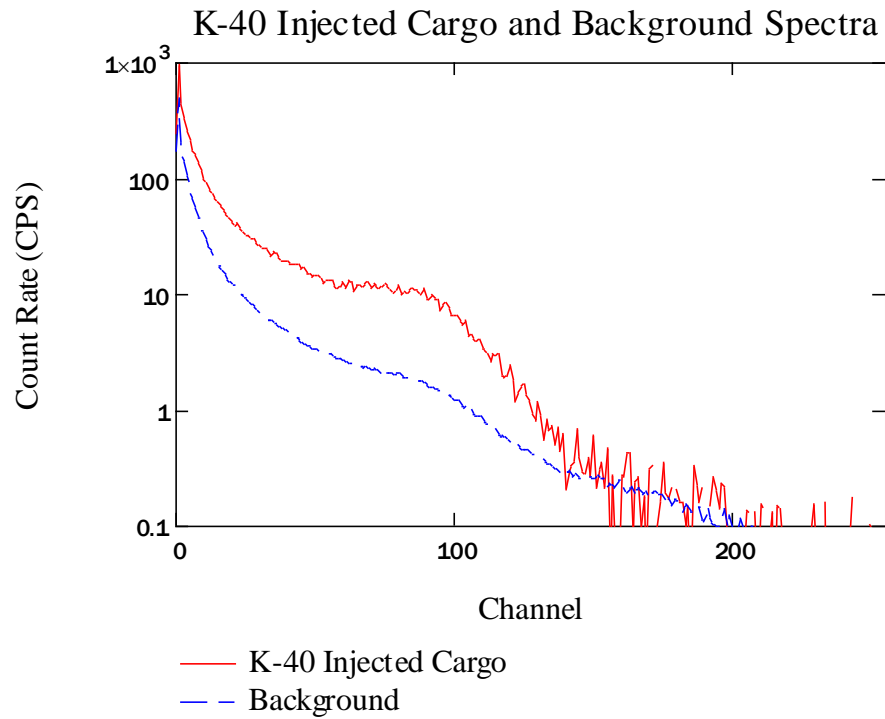


Figure 4.9: ^{40}K injected cargo and background spectra

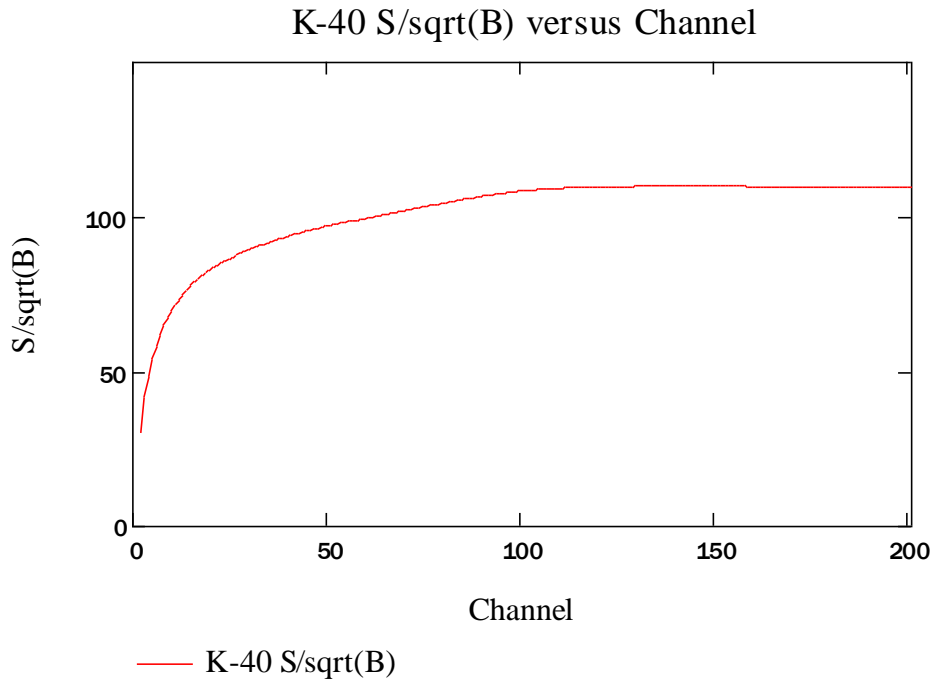


Figure 4.10: ⁴⁰K S/sqrt(B) versus channel calibration curve

The S/\sqrt{B} optimized upper edge of the ⁴⁰K low energy window was found to be channel 140. The margin between the optimized upper edge for ⁴⁰K (channel 140) and DU (channel 85) is small given the high gamma ray energies and broad Compton continuums, but the alarm logic accounts for this in an intelligent manner. It was mentioned above that ⁴⁰K and other NORM sources causes increased count rates in the DU window, and this is plainly visible in Figure 4.9. The high energy gamma rays necessary to add counts to the ⁴⁰K window will have a large Compton continuum that may increase the count rate in the DU window as well. A vehicle carrying NORM cargo could potentially cause an alarm in the DU and ⁴⁰K windows, which the alarm logic will state to be a NORM alarm (while the window is the “⁴⁰K” window, the alarm is a “NORM” alarm if the ⁴⁰K windows exceed its

thresholds). The ^{40}K low energy window consists of channels 86 through 140. The high energy window is channels 141 through 255.

A single window extends from channel 141 to the last channel, channel 255. This window is called the NOISE window because the occupied signal counts become sparse in this region. Since the gamma rays in this region are high energy, their probability of interaction in the PVT is small and thus the count rates fluctuate wildly unless the count time is very long (this is why background spectra are not as noisy in the higher energy channels). No individual alarm decisions are made based upon data solely in this window. The table below is a summary of all three energy windowing algorithms and lists the channels contained within the respective windows.

Table 4.1: Energy windowing algorithm summary

	Co-57 Algorithm	Ba-133 Algorithm	Advanced Algorithm				
	Co-57 Window	Ba-133 Window	HEU Window	WGPu Window	Cs-137 Window	DU Window	K-40 Window
Low Energy Window	2 - 12	2 - 38	2 - 12	13 - 30	31 - 57	58 - 85	86 - 140
High Energy Window	6 - 24	39 - 255	13 - 255	31 - 255	58 - 255	86 - 255	141 - 255

The plot below shows an occupied spectrum with the different energy windows highlighted. This reveals the relative size of the different windows in graphical form, allowing for easier interpretation of how the windows are summed to calculate the compensated ratio.

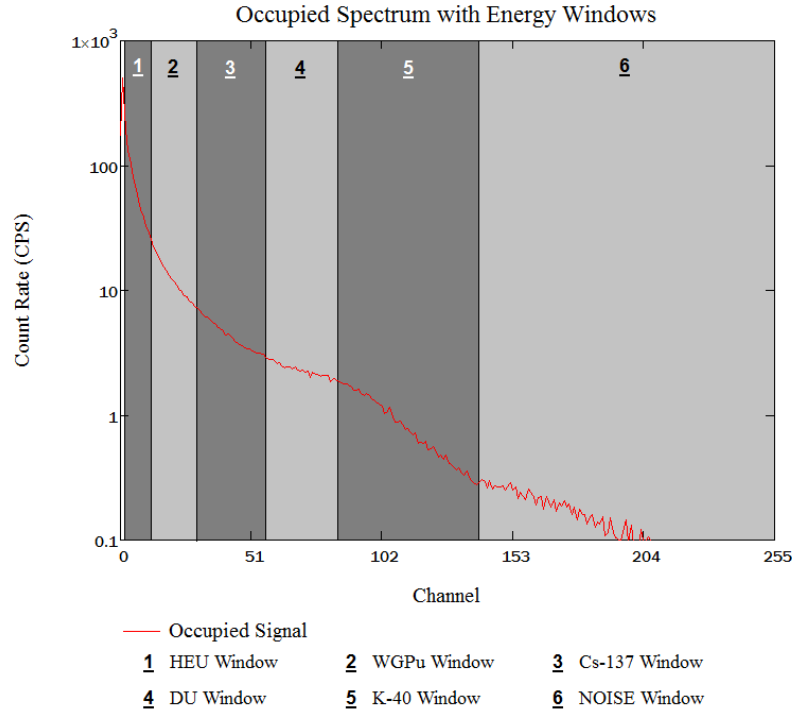


Figure 4.11: Occupied Spectrum with Energy Windows Shown

Equation (2.9) is the compensated ratio for a two window design:

$$R_C = S_L - \frac{B_L}{B_H} \cdot S_H$$

It has been mentioned that the newly proposed advanced algorithm employs multiple windows, so how is the compensated ratio adjusted to make the computation more efficient? There are two important properties of the advanced algorithm's window designs that allow efficient calculations of the ratios. The first property is that each ratio's high energy window extends from the upper edge of the low energy window to the last channel, channel 255. The second property is that every low energy window is contiguous with the low energy window below it, i.e. HEU's low energy window extends from channel 2 to channel 12, WGPu's low energy window extends from channel 13 to channel 30, ¹³⁷Cs from channels 31 to 57, DU

from channels 58 to 85, ^{40}K from channels 86 to 140, and the so-called NOISE window from channels 141 to 255. Using these two facts, the compensated ratio can be rewritten in the following form:

$$R_i = S_i - \frac{B_i}{\sum_{k=i+1}^6 B_k} \cdot \sum_{k=i+1}^6 S_k \quad \text{where:}$$

$i = 1, 2, \dots, 5$

Subscript 1 = HEU, Subscript 2 = WGPu, Subscript 3 = ^{137}Cs
 Subscript 4 = DU, Subscript 5 = ^{40}K , Subscript 6 = NOISE

R_i = Ratio calculation for window i (4.2)
 S_i = Occupied count rate in window i
 B_i = Background count rate in window i

$\sum_{k=i+1}^6 B_k$ = Sum of background count rates in windows $i + 1$ through 6
 $\sum_{k=i+1}^6 S_k$ = Sum of occupied count rates in windows $i + 1$ through 6

Note: While S_6 and B_6 exist, R_6 does not exist

The threshold for the advanced algorithm is a 4σ threshold, where the standard deviation takes the form:

$$\sigma_{R_i} = \sqrt{B_i} \cdot \sqrt{1 + \frac{B_i}{\sum_{k=i+1}^6 B_k}} \quad (4.3)$$

The width of the high energy windows for each ratio was not arbitrarily chosen, but was determined by minimizing σ_{R_i} . Assuming a fixed low energy window determined by the window optimization, the only possible way to minimize σ_{R_i} is by maximizing the summation term in the fraction's denominator. The only way to accomplish this is to contain

every window above the low energy window's upper edge in the summation. Figure 4.12 illustrates this minimization where the low energy window is channels 2 through 12 (HEU window) and the lower edge of the high energy window is channel 13. It is clear that increasing the upper edge of the high energy window, thereby increasing the width of the high energy window, minimizes the standard deviation and allowing for more sensitive alarms without increasing the false alarm rate.




Figure 4.12: Plot of standard deviation of the compensated ratio versus high energy window upper edge

Once the windows were optimized, the alarm logic could be developed with the assistance of the alarm simulator Mathcad worksheet. The truth table for the alarm logic was based on the physics of Compton scattering. Recalling that a Compton scattered gamma ray cannot exceed the energy of the Compton edge described by equation (2.2), it is physically impossible for a radiation source to generate an alarm in a window whose low energy lower

bound is above the source's Compton edge. The S/\sqrt{B} calibration ensured the smeared Compton edge resulting from PVT's poor energy resolution was accounted for. This was the first major assumption made when developing the alarm logic: a specific window's source cannot generate an alarm in the window above it. The second assumption again used the S/\sqrt{B} analysis to state that a window's source can generate an alarm in the window directly below it, but not in the next window down. The figure below is the compilation of all the S/\sqrt{B} calibration curves where the vertical dashed red lines represent the upper edge of each window.

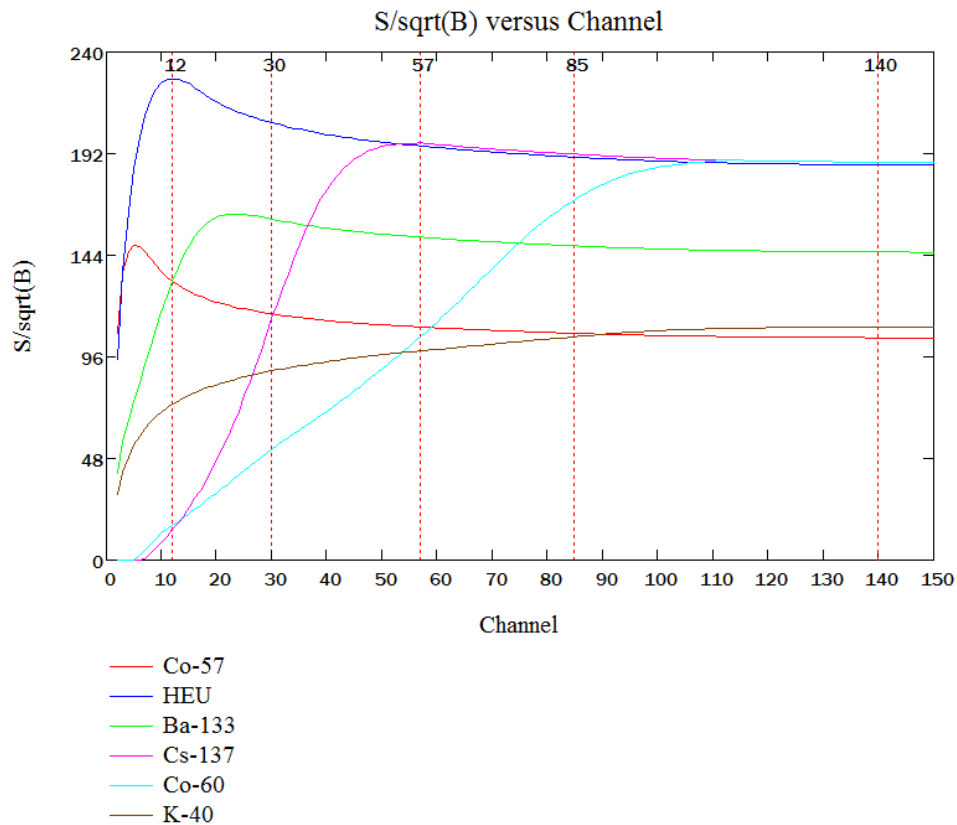


Figure 4.13: Plot of S/\sqrt{B} versus channel

The first assumption and second assumption can be verified by examining this plot. The S/\sqrt{B} curve for each window specific nuclide had high values in the window directly below it, but the effect was diminished greatly two windows down. This was seen very clearly for all windows except the ^{40}K source window. While the ^{40}K source S/\sqrt{B} signal still appeared to be strong two windows below the ^{40}K window (in the ^{137}Cs window), in practice ^{40}K sources do not generate sufficient counts in the ^{137}Cs window to generate an alarm.

With these two assumptions verified, the logic structure below was developed:

$A_{NORM} = 1 \text{ IF } (R_{40K} > T_{40K})$ $A_{NORM} = 0 \text{ OTHERWISE}$
$A_{DU} = 1 \text{ IF } (R_{DU} > T_{DU}) \wedge \neg (R_{40K} > T_{40K})$ $A_{DU} = 0 \text{ OTHERWISE}$
$A_{137Cs} = 1 \text{ IF } (R_{137Cs} > T_{137Cs}) \wedge \left[\neg (R_{DU} > T_{DU}) \vee \left\{ (R_{DU} > T_{DU}) \wedge (R_{40K} > T_{40K}) \right\} \right] \wedge \neg (A_{DU} = 1)$ $A_{137Cs} = 0 \text{ OTHERWISE}$
$A_{WGPu} = 1 \text{ IF } (R_{WGPu} > T_{WGPu}) \wedge \left[\neg (R_{137Cs} > T_{137Cs}) \vee \left\{ (R_{137Cs} > T_{137Cs}) \wedge (R_{DU} > T_{DU}) \right\} \right] \wedge \neg (A_{137Cs} = 1)$ $A_{WGPu} = 0 \text{ OTHERWISE}$
$A_{HEU} = 1 \text{ IF } (R_{HEU} > T_{HEU}) \wedge \left[\neg (R_{WGPu} > T_{WGPu}) \vee \left\{ (R_{WGPu} > T_{WGPu}) \wedge (R_{137Cs} > T_{137Cs}) \right\} \right] \wedge \neg (A_{WGPu} = 1)$ $A_{HEU} = 0 \text{ OTHERWISE}$

where:

R_i = Modified compensated ratio in window i

T_i = Threshold in window i

A_i = Alarm in window i if $A_i = 1$

Figure 4.14: Advanced alarm algorithm alarm logic structure

The full table form of the logic structure can be seen in the table below. This table lists the 32 possible combinations of threshold alarm inputs and the subsequent alarm outputs.

Table 4.2: Full alarm truth table for advanced alarm algorithm

INPUT					OUTPUT				
$R_{HEU} > T_{HEU}$	$R_{WGPu} > T_{WGPu}$	$R_{137Cs} > T_{137Cs}$	$R_{DU} > T_{DU}$	$R_{40K} > T_{40K}$	A_{HEU}	A_{WGPu}	A_{137Cs}	A_{DU}	A_{40K}
0	0	0	0	0	0	0	0	0	0
1	0	0	0	0	1	0	0	0	0
0	1	0	0	0	0	1	0	0	0
0	0	1	0	0	0	0	1	0	0
0	0	0	1	0	0	0	0	1	0
0	0	0	0	1	0	0	0	0	1
1	1	0	0	0	0	1	0	0	0
1	0	1	0	0	1	0	1	0	0
1	0	0	1	0	1	0	0	1	0
1	0	0	0	1	1	0	0	0	1
0	1	1	0	0	0	0	1	0	0
0	1	0	1	0	0	1	0	1	0
0	1	0	0	1	0	1	0	0	1
0	0	1	1	0	0	0	0	1	0
0	0	1	0	1	0	0	1	0	1
0	0	0	1	1	0	0	0	1	1
1	1	1	0	0	1	0	1	0	0
1	1	0	1	0	0	1	0	1	0
1	1	0	0	1	0	1	0	0	1
1	0	1	1	0	1	0	0	1	0
1	0	1	0	1	1	0	1	0	1
1	0	0	1	1	1	0	0	0	1
0	1	1	1	0	0	1	0	1	0
0	1	1	0	1	0	0	1	0	1
0	1	0	1	1	0	1	0	0	1
0	0	1	1	1	0	0	1	0	1
1	1	1	1	0	0	1	0	1	0
1	1	1	0	1	1	0	1	0	1
1	1	0	1	1	0	1	0	0	1
1	0	1	1	1	1	0	1	0	1
0	1	1	1	1	0	0	1	0	1
1	1	1	1	1	1	0	1	0	1

In making the logic structure, certain assumptions were made that would impact the decision making significantly. For example, if two consecutive windows alarmed, only the higher energy window's material was implicated as the material present. This was due to the fact that the offending material often caused the alarm in the window below its designated window. Unfortunately, this meant that if the lower window's material was actually present, it would not be identified explicitly. This issue extends to three, four, and five consecutive alarms. For three consecutive alarming windows, it was assumed that the highest energy window's material triggered the alarm in the window directly below it, and that the lowest alarming window's material was actually present. For four consecutive windows, it was again assumed that the highest energy window's material was present and triggered the window alarm directly below it, and the third lowest window's material was present and triggered the window alarm below it. This pattern of every other alarm triggering starting with the highest energy alarm extended to the five consecutive alarm case.

The implications of this logic became apparent immediately: a vehicle carrying sources that alarm in two consecutive windows will have only the higher energy source implicated, e.g. a vehicle carrying an HEU and WGPu source will only trigger a WGPu alarm. This situation is anticipated to be exceedingly rare, and if it arises, it is hoped the RPM workers, in secondary inspection, would identify all the illicit materials present. This human factors problem is outside the scope of this study and is not further addressed.

With the alarm logic completed, the advanced algorithm was fully developed; after the development, the masking sensitivity analysis was performed. The sensitivity analysis tested how sensitive energy windowing algorithms were to the effects of high energy

masking and also measured the overall effectiveness of the algorithms. Vehicle cargos included five non-alarming cargos, two ^{40}K -bearing cargos, and two granite-bearing cargos. The injection sources were: ^{57}Co , HEU, ^{133}Ba , ^{137}Cs , ^{60}Co , and ^{40}K . Each of the three algorithms' sensitivity to NORM was tested for each of the injection sources; the area under the energy windowing/gross count curve combinations was used as the metric. For example, the HEU source was injected into each of the nine available cargos and the area under the curves for each algorithm was recorded. The Mathcad worksheet could test all three energy windowing algorithms at once.

To create the sensitivity curves, an equation is needed to show how alarm sensitivity depends on NORM source count rate. Fortunately, all three algorithms utilized forms of the compensated ratio, so a single equation was derived to create the plots. Start with the compensated ratio, equation (2.9):

$$R_C = S_L - \frac{B_L}{B_H} \cdot S_H$$

The occupied signal count rates are made up of the sum of the injected source and the cargo. Therefore the occupied signal terms can be rewritten in terms of percent of source counts and percent of cargo counts:

$$R_C = \left(\% \text{Source} \cdot S_L^{\text{source}} + \% \text{Cargo} \cdot S_L^{\text{cargo}} \right) - \frac{B_L}{B_H} \cdot \left(\% \text{Source} \cdot S_H^{\text{source}} + \% \text{Cargo} \cdot S_H^{\text{cargo}} \right) \quad (4.4)$$

Since it is desired to plot the source count rate versus the cargo count rate, the equation above is solved for the percent source:

$$\% Source = \frac{R_C - \left(\% Cargo \cdot S_L^{cargo} \right) + \left(\frac{B_L}{B_H} \cdot \% Cargo \cdot S_H^{cargo} \right)}{S_L^{source} - \frac{B_L}{B_H} \cdot S_H^{source}} \quad (4.5)$$

The sensitivity plots need to show alarm sensitivities, therefore the R_C term is replaced with the actual threshold value. Since the thresholds vary slightly for each algorithm, the ratio threshold will simply be rewritten as R_{th} . Also, the “percent source” and “percent cargo” terms can be rewritten in terms of their respective gross counts (GC_C for cargo gross counts and GC_S for source gross counts).

$$\frac{Source}{GC_S} = \frac{R_{th} - \left(\frac{Cargo}{GC_C} \cdot S_L^{cargo} \right) + \left(\frac{B_L}{B_H} \cdot \frac{Cargo}{GC_C} \cdot S_H^{cargo} \right)}{S_L^{source} - \frac{B_L}{B_H} \cdot S_H^{source}} \quad (4.6)$$

The final equation simply converts (3.11) into functional form, replacing the “Cargo” variable with x , the cargo count rate, and replacing “Source” with $Alg(x)$, the injection source count rate as a function of cargo count rate. The count rates utilizes units of counts per second (cps).

$$Alg(x) = \frac{R_{th} - \left(\frac{x}{GC_C} \cdot S_L^{cargo} \right) + \left(\frac{B_L}{B_H} \cdot \frac{x}{GC_C} \cdot S_H^{cargo} \right)}{S_L^{source} - \frac{B_L}{B_H} \cdot S_H^{source}} \cdot GC_S \quad (4.7)$$

The gross count curve takes virtually the same form. It takes into account that the total gross count rate is made up of the source and cargo count rates. Again, the curve represents value that alarms, so the gross count threshold GC_{th} is employed:

$$GC_{th} = \%Source \cdot GC^{source} + \%Cargo \cdot GC^{cargo} \quad (4.8)$$

$$GC_{th} = \frac{Source}{GC^{source}} \cdot GC^{source} + \frac{Cargo}{GC^{cargo}} \cdot GC^{cargo} \quad (4.9)$$

$$GC_{curve}(x) = GC_{th} - x \quad (4.10)$$

Again, x replaces the cargo count rate and $GC_{curve}(x)$ replaces the source count rate.

The results of the NORM masking sensitivity analysis are seen in the series of figures below. Cargos labeled “Std #” are non-alarmed cargoes. “ ^{40}K -high” is a high activity ^{40}K -bearing cargo that triggered a gross count alarm at the POE site; it is suspected to serve as the most powerful masking source available. “ ^{40}K -low” is a low

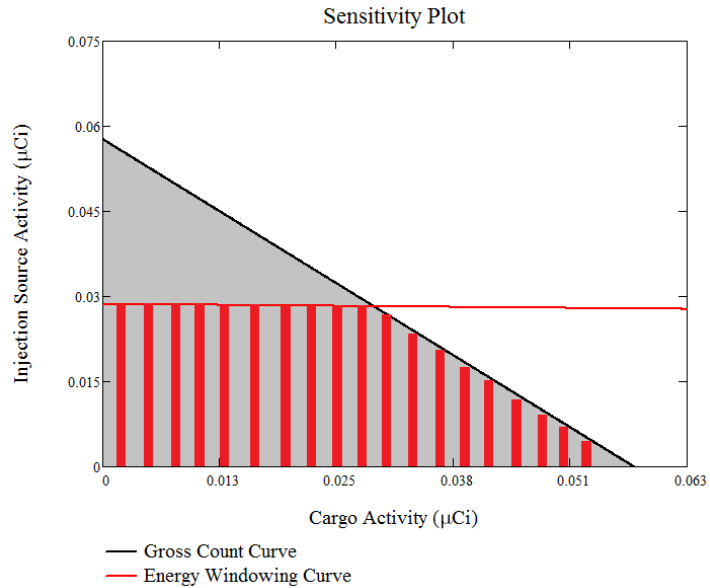


Figure 4.15: Visual representation of normalized combined area

The gray area represents the area only under the gross count algorithm curve.

The gray/red striped area represents the combined area under both the energy windowing and gross count algorithm curves.

Dividing the gray/red striped area by the gray area results in the normalized combined area.

activity cargo that still triggered a gross count alarm. “Granite 1” and “Granite 2” are granite-bearing cargos that both triggered gross count alarms. The metric used to compare energy windowing algorithms is the area under the energy window/gross count combined algorithm curve divided by the area under only the gross count algorithm curve; this metric is called the normalized combined area. Figure 4.13 shows the visual representation of the normalized combined area. The normalization anchors the energy windowing algorithm areas to a static value, allowing instantaneous comparisons between them and the gross count algorithm that is the current practice. If the normalized combined area is close to 1.00 (approaching from the low side), then the energy windowing/gross count combined algorithm is not much more effective at detecting the injection source within the selected cargo than the gross count algorithm alone. If the normalized combined area is very small, the energy windowing/gross count combined algorithm is much more effective at detecting the injection source within the selected cargo. In comparing one energy windowing algorithm to another, the lower normalized combined area shows the more effective algorithm. This is how the sensitivity analysis program also measures overall algorithm effectiveness.

The results of the sensitivity analysis were quite telling. The first important finding was that the ^{57}Co and ^{133}Ba algorithm cannot detect ^{137}Cs or any other source producing only gamma rays of energy higher than ^{137}Cs 's. Since their low energy windows are far below where ^{137}Cs 's Compton continuum has a significant effect on the spectrum, the compensated ratio calculation results in a negative number and does not alarm. The sensitivity analysis program gives a negative result for the energy window/gross count algorithm integrals when

the energy windowing algorithm fails to detect the injection source; this means the value is nonsensical and not recorded.

The plots of normalized combined area versus cargo for each injection source (Figure 4.16 through 4.21) show the comparisons between algorithms and the effects of masking on each algorithm. For the advanced algorithm, these plots show how high energy masking affects each individual window. When the normalized combined area value increases for a cargo, it means the particular cargo decreases the sensitivity of the specific algorithm to the specific injection source. In some cases, e.g. the ^{40}K window of the advanced algorithm, a NORM cargo will increase the effectiveness of the algorithm because the window is specifically designed to detect NORM sources. Overall algorithm effectiveness can be examined by comparing the magnitudes of each algorithm's normalized combined area for the non-alarming cargos.

There is one plot per injection source; the normalized combined area on the y-axis and specific cargos on the x-axis. Directly below the picture is a table listing the actual normalized combined area values for each algorithm/cargo combination. The explanation of what each cargo is can be found earlier in this chapter. Notice only the ^{133}Ba algorithm and the advanced algorithm are seen on the ^{133}Ba injection source plot, and only the advanced algorithm on the ^{137}Cs , DU and ^{40}K plots. Each plot can be seen below with an explanation of them immediately underneath the figure. The final four points on each plot are the NORM cargos; for the ^{57}Co , HEU, and ^{133}Ba injection source plots, it is expected that the normalized combined areas for the four NORM cargos will be higher than for the five non-alarming cargos due to NORM masking effects.




Figure 4.16: ⁵⁷Co injection source normalized combined area plot

The ⁵⁷Co algorithm was very effective at detecting ⁵⁷Co, as expected. With ⁵⁷Co being an HEU proxy, the advanced algorithm was also effective at detecting ⁵⁷Co. The increased normalized combined area values for the high activity ⁴⁰K cargo and the other three NORM cargos for the ¹³³Ba and advanced algorithm were expected due to the increased counts in the high energy window (NORM masking effect). Given that the ⁵⁷Co algorithm high energy window only extends up to channel 24, it was surprising to see how much the NORM sources affected that algorithm. The ¹³³Ba algorithm did not perform nearly as well as the other two due to the extension of the low energy window above the range of ⁵⁷Co's Compton edge.




Figure 4.17: HEU injection source normalized combined area plot

The results of the HEU injection test were shocking. Figure 4.15 shows that the ^{57}Co algorithm was far less effective than both the advanced algorithm and the ^{133}Ba algorithm. Even though ^{57}Co is a proxy for HEU, the construction of the ^{57}Co algorithm caused it to lose sensitivity to HEU. Recall the ^{57}Co algorithm's low energy window covers windows 2 through 12, and the high energy window contains windows 6 through 24. The advanced algorithm, which calibrates the HEU window using an HEU source, has the low energy window containing channel 2 through 12, and the high energy window covering channels 13 through 255. Since the ^{57}Co algorithm uses overlapping windows, channels 6 through 12 appear in both the low and high energy windows. This was fine for detecting ^{57}Co , but recall HEU's gamma ray is 186 keV whereas ^{57}Co 's is 122 keV. HEU produces a significant signal in windows 6 through 12, so by counting this region in both the low and high energy

windows, the algorithm's sensitivity to HEU is severely diminished. Since the advanced algorithm has a window specifically tuned to detect HEU, it clearly outperformed the ^{133}Ba algorithm, which contains only a single window to detect HEU and WGPu. All three algorithms exhibited NORM masking, but the advanced algorithm experienced the effects to a slightly smaller degree than the ^{133}Ba algorithm; this was seen by examining the magnitude of the difference between the normalized combined areas.




Figure 4.18: ^{133}Ba injection source normalized combined area plot

The ^{57}Co algorithm was not able to detect the presence of the ^{133}Ba injection source; the count rate increase in its high energy window far outnumbered the count rate increase in the low energy window, resulting in a negative ratio and no alarm. The ^{133}Ba was more effective at detecting the ^{133}Ba injection source than the advanced algorithm, but the

effectiveness can be misleading. The ^{133}Ba algorithm utilizes a low energy window covering channels 2 through 38 whereas the advanced algorithm low energy windows only contains channels 13 through 30. The high probability 80 keV gamma ray from ^{133}Ba thus increases the count rate in the ^{133}Ba algorithm's low energy window and not in the advanced algorithm's low energy window. A gamma ray of this energy does not get produced with high probability from WGPu and thus artificially improves the effectiveness of the ^{133}Ba algorithm. The advanced algorithm was designed to detect the high energy gammas rays produced from WGPu using an adjusted ^{133}Ba calibration. Also, the ^{133}Ba algorithm was found to be far more sensitive to masking than the advanced algorithm. This was due to the ^{133}Ba algorithm's low energy window extending beyond the optimized upper edge for WGPu (channel 30), thus NORM sources increase the count rate in the low energy window of the ^{133}Ba algorithm more significantly than in the advanced algorithm. While failing to do so for a ^{133}Ba injection source, it is believed the advanced algorithm would more effectively detect WGPu than the ^{133}Ba algorithm. It has been shown above that the advanced algorithm was undoubtedly less sensitive to NORM masking than the ^{133}Ba algorithm.

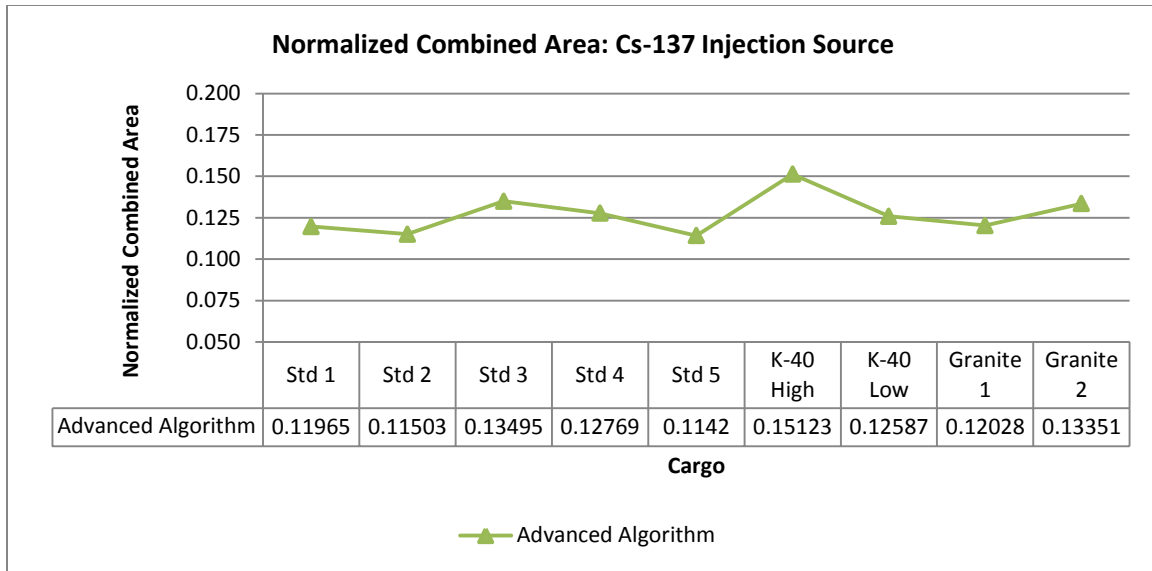


Figure 4.19: ^{137}Cs injection source normalized combined area plot

Notice neither the ^{57}Co nor ^{133}Ba algorithms were capable of detecting the ^{137}Cs injection source because the count rates in their low energy windows did not increase nearly as much as they did their high energy windows. The advanced algorithm, with its designated ^{137}Cs window, very effectively detects the injection source. Of important note was the magnitude of the normalized combined area values compared to the values for the ^{57}Co , HEU, and ^{133}Ba injection sources. The normalized combined areas were very small when the injection sources matched the sources used for calibration (HEU and ^{137}Cs) and were slightly higher when detecting proxy sources (^{57}Co and ^{133}Ba). The effect of masking was much less significant for the advanced algorithm's ^{137}Cs window than for the previous injection sources and in fact seems to improve algorithm effectiveness. This was due to the fact that the ^{137}Cs low energy window range extends into the portion of the spectrum where NORM sources cause increased count rates. Thus NORM cargos increase the count rates in both the low and high energy windows. Depending on the NORM source spectrum, the cargo may increase

the count rate more in one window than the other; this explains why some cargos seem to increase cause increased algorithm effectiveness while other decrease it. In conclusion, it appeared as if NORM masking effects were not severe for the advanced algorithm's ^{137}Cs window.




Figure 4.20: ^{60}Co injection source normalized combined area plot

Again, the only algorithm capable of detecting the ^{60}Co source was the advanced algorithm. The magnitudes of the normalized combined areas again showed that the advanced algorithm was very effective at detecting ^{60}Co even though the window was designed to detect DU. Although a DU injection source was not able to be generated, it is suspected that the advanced algorithm would detect DU very effectively as well. The effect of NORM masking on the DU window differed greatly from the effects on the lower energy windows (HEU, WGPu). Like with the ^{137}Cs window, NORM sources increased count rates in the low and high energy windows for the DU window, resulting in mixed masking effects.

The low normalized combined area seen for the high activity ^{40}K -laden cargo differed wildly from the high normalized combined area for the “Granite 2” cargo. While the NORM masking effect seemed to drastically increase or decrease algorithm effectiveness depending on the source, the effect was actually very small when examined in absolute terms. The lowest and highest normalized combined area values differed from the mean by only 0.054 and 0.040 respectively; the perceived radical masking effects were largely due to the scale of the plot required to highlight algorithm detail.




Figure 4.21: ^{40}K injection source normalized combined area plot

The advanced algorithm’s ^{40}K window exhibited behavior differed from the previous windows’ behavior. First, the magnitudes of the normalized combined areas were very large for the non-alarming cargos. In addition, the ^{40}K -laden cargos greatly decreased the normalized combined area. This effect was expected because this window was designed

specifically to detect such cargos. The granite cargos did not exhibit the effect to a significant extent; however, NORM cargos tend to be high activity so the likelihood of a NORM cargo to trigger a ^{40}K window alarm is high, even if the normalized combined area is large. This will be shown by applying the algorithm to the POE dataset. In summary, the advanced algorithm's ^{40}K window did not experience NORM masking effects because the window was designed to alarm for NORM cargos.

Appendix 1 shows each of the sensitivity plots. There is one plot per injection source/cargo combination, resulting in a total of 54 plots. Some interesting observations were made by studying these plots. For the lower energy windows (^{57}Co algorithm, ^{133}Ba algorithm, and advanced algorithm's HEU, WGPu, and ^{137}Cs windows), the slope of the energy windowing algorithm curves increased when sources were injected into NORM cargos versus injection in non-alarming cargos. The y-intercept, however, remained largely unchanged. The slope increased because the x-axis represented cargo count rate, and as the NORM cargo count rate increased, the count rates in the high energy windows increased, thus requiring larger source activities (higher y values) to increase the low energy window enough to trigger the alarms. The increase in normalized combined area was due almost entirely to the increase in slope. For the higher energy windows (advanced algorithm's DU and ^{40}K windows), NORM cargos caused the slope of the energy windowing curves to decrease for similar reasons. NORM cargos increased count rates in low energy windows (due to window design) more than the high energy windows, thus requiring smaller NORM injection sources to trigger alarms. From this examination, it can be said that NORM masking affects the slope of the energy windowing algorithm curves.

The y-intercept increased or decreased only by very small margins due to the background changing as the cargos were changed (each cargo spectrum has its own background spectrum). As the background changed, the energy window alarm threshold changed as well. Since the cargo spectra were all extracted from the POE dataset, the background spectra do not vary a significant amount, thus the y-intercept only slightly changes.

The value represented by the energy windowing algorithm y-intercept can be explained through a thought experiment. Assume the injection source spectrum shape triggers an energy window alarm. At 0 cps cargo count rate, the energy windowing algorithm alarms only if the source count rate (y-intercept) is high enough to exceed the 4σ threshold (too few counts would not develop a sufficient signal to trigger an alarm). The y-intercept was found to be exactly equal to the energy windowing alarm threshold. As the threshold is increased or decreased, the y-intercept does the same. In the sensitivity analysis program, when the sigma multiplier was changed from a 4σ threshold to a 5σ threshold, the y-intercept increased by exactly 1σ . The y-intercept decreased 1σ when the threshold was decreased to 3σ .

With the sensitivity analysis concluded, the advanced energy windowing/gross count algorithm was tested on a large scale using the POE dataset. A Perl script discussed in chapter 3 was run; it calculated the ratios and thresholds for the advanced alarm algorithm.

The advanced algorithm detected 12 alarms out of the 2129 vehicle runs. Of these 12 alarms, all 12 triggered the gross count alarm while only three runs triggered an energy window alarm. Nine of the 12 alarming runs contained innocent radioactive sources and

were thus “innocent alarms”. Two of the nine innocent alarms triggered the NORM energy windowing alarm; this was an encouraging result because the cargos were indeed identified to be NORM sources. One furniture cargo actually triggered the HEU window alarm while obviously containing no HEU (an illicit HEU cargo would only be reported in a classified document). Both the ^{57}Co and ^{133}Ba energy windowing algorithms also triggered an alarm for this cargo. It was not listed in the report what nuclide triggered the alarm, but after examining the spectrum it was clear that the HEU window did have an increased count rate. While only speculation, the high count rate, low-energy source present during that run tends to shift guesses towards a medical isotope injected in the driver or a passenger. This highlighted a known problem with energy windowing algorithms; medical isotopes cannot be identified and tend to trigger SNM alarms. These must be identified by secondary inspection. One paper suggests that medical isotopes may be the limiting operational factor for RPM border screening technologies (Kouzes, et al. 2005).

Three gross count alarming runs were not listed as alarming runs in the official report, meaning that they did not trigger an alarm during the test. This was likely due to the slight inaccuracy of using accumulated spectrum data to test the advanced alarm algorithm. False alarm rates are calculated on a per-decision basis using time profiles; since the Perl script used accumulated spectra rather than time profile spectra, the false alarm rate could not be calculated from these results. These three spectra, along with the other nine alarming runs, were tested in the simulated real-time advanced alarm algorithm emulator to test if they still trigger an alarm in the most accurate advanced alarm algorithm simulation available.

With the large number of runs available for study, a statistical analysis of the algorithm results was performed. The summary statistics provided details on the actual values of the ratio calculations. While this script was not effective in calculating false alarm rates or testing algorithm alarm accuracy, it did enable a statistical analysis to be performed on the ratio and threshold values. The 12 known alarm runs were removed for the statistical analysis to ensure the values were not skewed by the known outliers.

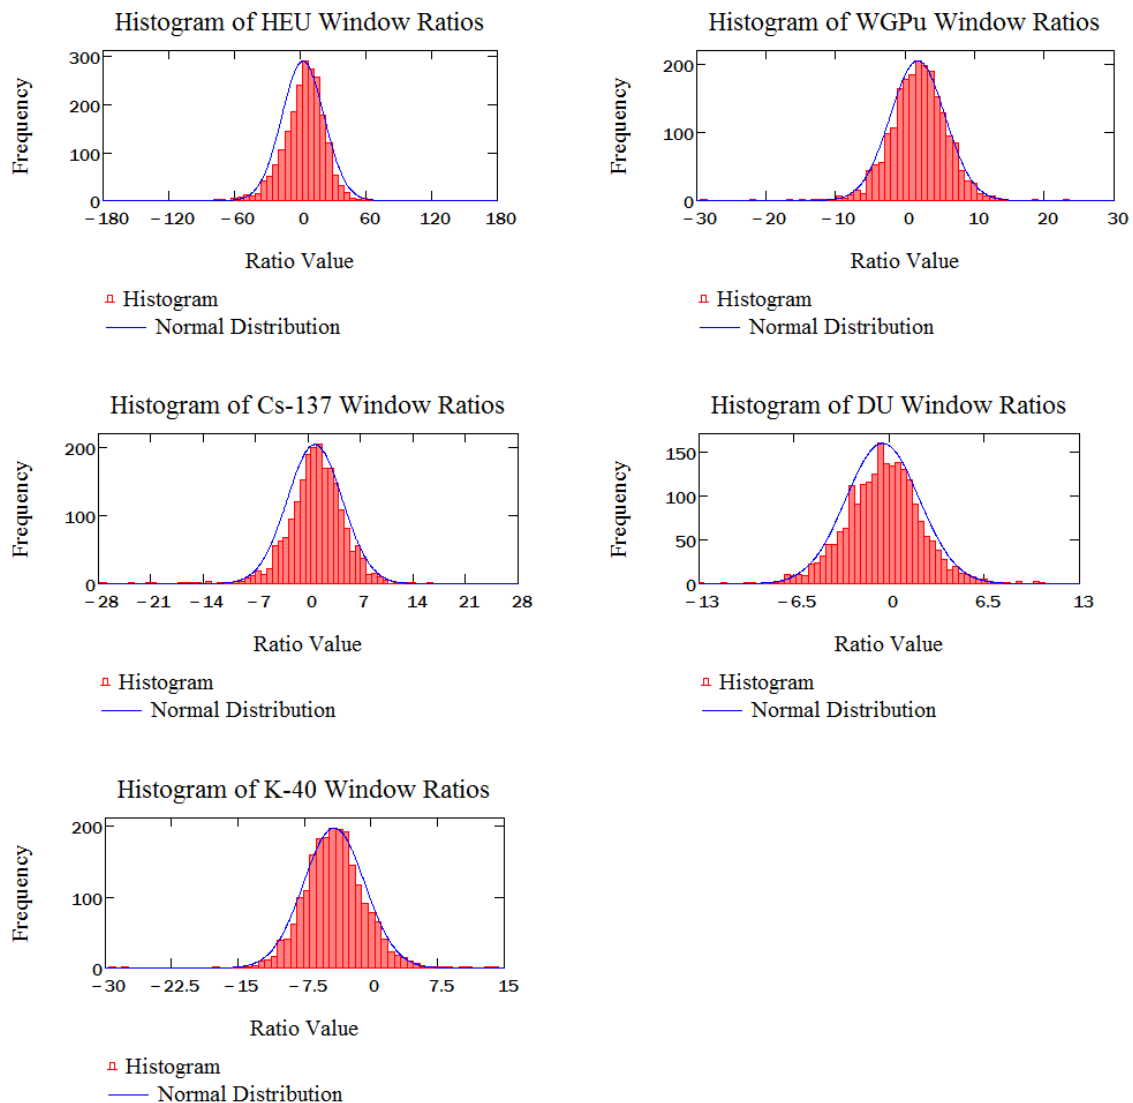


Figure 4.22: Histograms of advanced algorithm window ratios

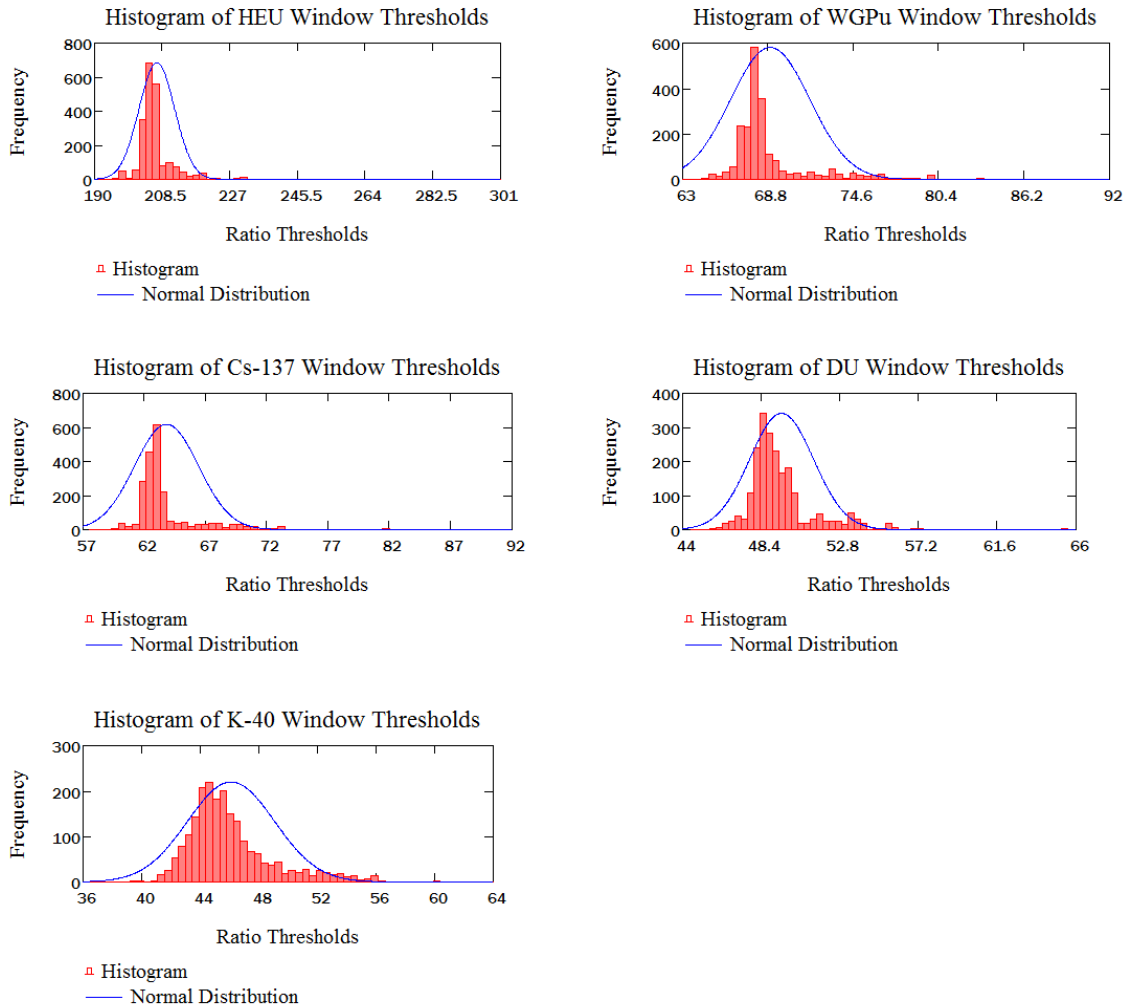


Figure 4.23: Histograms of advanced algorithm window thresholds

The distribution of ratio values and ratio thresholds for each advanced algorithm window can be seen in Figures 4.21 and 4.22, respectively. The 12 alarming runs were removed from the dataset to prevent known outliers from sullyng the results. Each individual plot shows a histogram of the ratio and a normal distribution curve with the same mean and standard deviation as the corresponding histogram. The details of the plots revealed that standard commercial vehicles produced a seemingly wide range of ratio values

while failing to trigger alarms. Tables 4.3 and 4.4 below show the summary statistics for each of the window's ratio values and thresholds.

Table 4.3: Summary statistics for advanced algorithm window ratios

Ratio Summary Statistics						
Window	Median	Mean	Variance	Standard Deviation	Skewness	Kurtosis
HEU	4.30214	2.58237	372.15523	19.29133	-0.929	7.64995
WGPu	1.79084	1.75979	15.52212	3.93981	-0.32827	3.33753
Cs-137	0.96284	0.87507	13.55493	3.6817	-0.8363	4.82591
DU	-3.82928	-3.69957	39.10663	6.25353	-0.0126	0.58875
K-40	-0.76032	-0.77065	1.12824	1.06219	0.03468	0.44529

Table 4.4: Summary statistics for advanced algorithm window thresholds

Threshold Summary Statistics						
Window	Median	Mean	Variance	Standard Deviation	Skewness	Kurtosis
HEU	206.30093	207.04304	23.83752	4.88237	4.78263	67.65764
WGPu	68.16597	68.9502	7.24236	2.69116	2.40707	7.40796
Cs-137	63.03049	63.8096	6.79163	2.60607	2.88861	13.66334
DU	49.08266	49.53504	3.19288	1.78686	2.0045	7.76445
K-40	45.42715	46.10144	8.84283	2.97369	1.42702	2.59292

The HEU window ratios had a very high variance compared to the other windows, but that variation did not result in any false HEU alarms. While the HEU window's standard deviation was approximately 19, the mean HEU threshold was approximately 207. An HEU ratio would need to be 10.60 standard deviations above its mean value to trigger the mean HEU threshold. This was found to be the case for every window, so much so that the standard deviation margins between the ratio means and threshold means were: 10.60 for HEU, 17.05 for WGPu, 17.10 for ¹³⁷Cs, 19.97 for DU, and 14.58 for ⁴⁰K.

The means and medians were all very close to zero for the ratio calculations, as expected from construction of the modified compensated ratio equation. In every window, zero was within one standard deviation of the mean. Skewness is a measure of the asymmetry of the distribution, and it was found that ratio values were only slightly skewed. The HEU, WGPu, ^{137}Cs , and DU windows were all slightly-to-moderately left skewed (negative skewness). The ^{40}K window was slightly right skewed (positive skewness). Skewness was difficult to see through visual inspection of the ratio histograms because the values were so small, however the normal distribution curves did show some degree of skewness.

Kurtosis is a measure of “peakedness” of a distribution where a value of 0 corresponds to the excess kurtosis of a normal distribution. A distribution with positive excess kurtosis is called leptokurtic and one with negative excess kurtosis is called platykurtic. The excess kurtosis for the HEU ratio was 7.65. This means the distribution was leptokurtic; it was narrowly distributed about the mean and had a few, very large deviations. Most of the variance was therefore explained by these extreme values. On the other hand, a highly platykurtic distribution would have indicated that values were more closely distributed around the mean than they are for a normal distribution, implying that the variance was mostly due to a high number of small deviations. The HEU, WGPu, and ^{137}Cs windows were all highly leptokurtic. The DU and ^{40}K windows were only slightly leptokurtic, meaning the “peakedness” of their distributions were similar to that of a normal distribution. Figure 4.21 revealed the histograms were more narrowly distributed than the normal distribution and that the leptokurtic nature of the ratio distributions was obvious through visual inspection.

The threshold data took a different shape from the ratio data. The variances of the thresholds were very small, much smaller (in relative terms) than the ratio variance. It can be inferred that this resulted from the relatively constant nature of background radiation. While rain, solar storms, and other environmental effects cause massive changes in background radiation, most days do not experience these phenomena and thus background tends to be steady. Since background radiation was the only variable used in calculating thresholds, it was most assuredly the driving factor for their histogram shape.

Examining the skewness of the threshold data reinforces the conclusions drawn from the variance data. Through both visual inspection and calculation, it was clear that threshold data are either strongly positive-skewed or moderately positive-skewed. This could be explained due to environmental factors. Rain is the most frequent disruptor of steady background radiation; it washes radon decay products onto the ground and significantly increases background radiation until the products decay away. Since rain increases background radiation, it would be expected that thresholds would be skewed to higher values when rain occurs. The positive skewness values seen in the threshold data appeared to validate the hypothesis. It could be suggested that the lower energy window thresholds would be affected more by the rain due to the low energy gamma rays produced by radon and its decay products. Since it was observed that the HEU, WGPu, and ^{137}Cs had higher skewness values than the higher energy windows, this also appeared to be validated.

The threshold distributions were found to be highly leptokurtic, implying that most of the variance was due to a few, very large deviations. The rain hypothesis put forward to explain the strong tendency towards positive-skewed distributions was also put forward to

explain the leptokurtic threshold distributions. Simple deductive reasoning was used to argue for the rain/kurtosis link:

- 1) Rain is an infrequent event
- 2) Rain causes large deviations in threshold values
- 3) The observed distributions have large, infrequent deviations
- 4) The high kurtosis values for the threshold distributions are due to rain

The argument was valid, the premises were sound, and direct knowledge of rain-detector response allowed the conclusion that rain was the major factor that contributed to the positive-skew and leptokurtic nature of the threshold distributions.

There are a few major conclusions to be drawn from the statistical analysis. Threshold distributions were found to be very narrowly distributed about the mean and have small variances. Rain and other background-disrupting effects do change threshold values, but these changes are infrequent short-lived. The ratio values had higher variances due to the large number of variables that can affect ratio calculations such as cargo, vehicle size, vehicle material, etc. Even with these variances, ratio values were found to be significantly separated from the threshold values which implied that statistical false alarms were more unlikely than the 1/30,000 decisions false alarm rate expected by using 4σ thresholds. This separation allowed a high degree of confidence to be placed into the advanced algorithm's ability to detect illicit radioactive source trafficking without significant false alarm interruption.

Looking back to chapter 3, there was an additional Perl script used to examine the background suppression of the port-of-entry data. This script calculated the average percent

background suppression (PBS) between channels 0 and 275 for each non-alarming vehicle (recall the data was in 512 channel format for use in this script). It performed a linear regression of the PBS versus channel data, calculated and outputted the slope and intercept of the line; this was in order to examine how much the PBS changed with increasing energy. The histograms below show the distributions of the average PBS values, the y-intercept values, and the slope values.

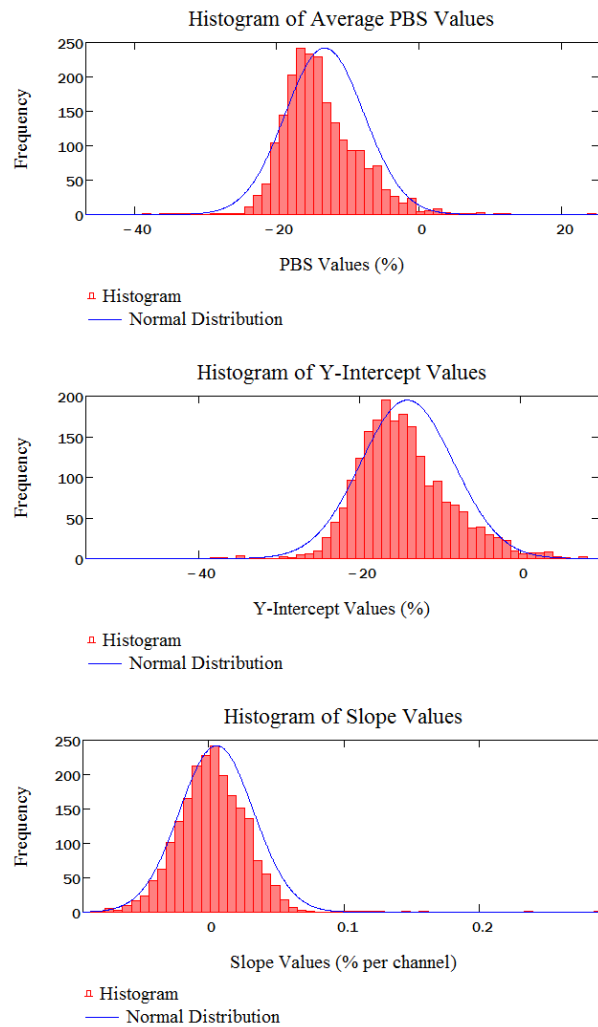


Figure 4.24: Background Suppression Histograms

Below, a table of the summary statistics of the background suppression data can be seen.

Table 4.5: Background suppression summary statistics

Background Suppression Summary Statistics						
	Median	Mean	Variance	Standard Deviation	Skewness	Kurtosis
Average PBS	-14.16744	-13.32727	30.50622	5.52324	0.8667	4.14955
Y-Intercept	-14.78163	-14.09606	33.31332	5.77177	0.38114	1.86717
Slope	0.00504	0.00559	0.00075	0.0274	1.63786	15.01253

The mean and median of average PBS value were very reasonable; typical PBS values range from 10% to 30%, so a mean of near 14% was encouraging. With a standard deviation of 5.52%, the mean average PBS value was found to be 2.41 standard deviations below zero. Using Chebyshev's inequality and assuming only the mean and standard deviation are known, it was calculated that at least 82.8% of the average PBS values were less than 0, meaning that almost 83% of cargos experienced at least some degree of background suppression. Looking at the individual values from the data set, 98.4% of cargos had average PBS values less than 0. The distribution was positive-skew and highly leptokurtic; this was due to the infrequent, yet extreme, positive average PBS values. This begged the question, how could there be a cargo with such a large, positive average PBS?

The cause of the infrequent positive average PBS values was due to short background counting times. Long background counting times are important because thresholds and ratio calculations depend on accurate background data. Short counting times for occupancies are not as critical because statistical fluctuations are taken care of by using a 4σ alarm threshold. If background counting times are short, thresholds can be artificially high or low, impairing both monitor effectiveness and altering background suppression calculations. If

two occupancies occurred within a short time frame, say 10 seconds between trucks, the background counting statistics were very poor. These short counting times often resulted in extreme and incorrect background count rate calculations. If the background count rate was artificially very low or high, the occupied count rate calculations could differ wildly since occupancies have inherently shorter counting times. A combination of low background count rate and high occupied count rate would necessarily result in a large, positive average PBS value; the converse is true if the background count rate is high and the occupied count rate is low. Short background counting times resulted in the extreme values that accounted for the bulk of the variance in the average PBS distribution.

To verify the assumption that background suppression does not vary greatly by channel, the y-intercept and slope of the PBS versus channel linear regression line was calculated for each run. The distributions of the slopes and y-intercepts are seen in the histogram plot above. First, it is important to point out that “increasing PBS” means the PBS value becoming less negative or more positive, implying background suppression is decreasing; a “decreasing PBS” means the PBS value is becoming less positive or more negative and that the background is being suppressed more.

The mean of the slopes was found to be 0.0056 PBS% per channel, meaning that the PBS from channel 0 to channel 275 increased, on average, approximately 1.54%. By extrapolating out to channel 511, the PBS increased 2.87%. Increasing PBS with increasing channel number/energy (a positive slope) implied that higher energy gamma rays were suppressed less than lower energy gamma rays, but not to a significant extent. In fact, a significant fraction of slopes, 41.4%, were actually negative, implying that high energy

gamma rays attenuated more than low energy gammas. However, a 99% confidence interval does not contain zero, therefore the slope was treated as positive at a 99% confidence level. This led to the conclusion that the background suppression decreased with increasing gamma ray energy and consequently the PBS value increased with increasing energy. However, the degree to which the PBS value increased was very slight, only a 3% increase in PBS across the entire spectrum. It was concluded that this increase did not significantly harden occupied signal spectrums by attenuating lower energy gamma rays preferentially. Thus, background suppression did not skew energy windowing data towards negative ratios by artificially decreasing the count rate in the low energy windows.

Checking the proximity of the y-intercept mean and average PBS mean was an easy way to supplement the slope conclusion that the background suppression did not change greatly as energy increased; this calculation served no other practical purpose. The y-intercept values were very similar to the average PBS values, with the mean and media differing by less than one percentage point each. If the slope (PBS per channel) was large and positive, the y-intercept would have necessarily been much smaller (more negative) than the mean average PBS. This affirmed the conclusion made in the previous paragraph.

In addition to verifying the assumption that background suppressions varies minimally by channel, it was also useful to know the typical values that average PBS took on for the POE dataset. Now to look back at what has been completed thus far: the advanced alarm algorithm has been fully created, its sensitivity was compared to other commercial algorithms, a statistical analysis of the thresholds and ratios it calculated was completed, and an examination of the POE data's background suppression statistics was performed.

Prior to the implementation of the simulated real-time alarm emulator, the injection source spectra needed to be transformed into vehicle time profiles. This required the use of the afore-mentioned spatial response function. Recall the total accumulated counts extracted from the .HST files were converted to count rates using the detection time; some count rates were given for other injection sources. The spatial response function's sole purpose was to convert source count rate data into time profile data where the profile was based off of point source geometry. Since the injection sources were count rate data only, they were the only sources that needed this transformation; the cargos into which the sources were injected were all based off of actual POE vehicle runs and thus have .PRO files with the corresponding time profiles already generated.

The spatial response function (SRF) was developed from the




Figure 4.25: Schematic of a point source traversing an RPM (Runkle, et al. 2005)

model of a point source passing through an RPM (Runkle, et al. 2005). The continuous function below describes this process by using the basic physical principal of $1/r^2$ to closely approximate the detector yield from a bare point source's transit through a portal monitor (Runkle, et al. 2005):

$$f(x) = \left(\frac{d}{r}\right)^2 \cdot \cos(\theta) = \left(\frac{d}{\sqrt{d^2 + x^2}}\right)^3 \quad \text{where:} \quad (4.11)$$

$f(x) =$ Detector yield fraction

See Figure 4.22 for schematic explanation of variables in the equation above. Since counts were recorded in an RPM system in discrete time intervals, converting this equation into a discrete time series yielded more applicable equations:

$$t_i = \frac{-L}{2 \cdot v} + i \cdot \text{intvl} \quad \text{where:}$$

t_i = Time relative to center axis of portal
 intvl = RPM occupied detection interval
 i = interval number, $i = 0, 1, \dots, \frac{L}{v \cdot \text{intvl}}$
 L = Length of tractor trailer
 v = Velocity of tractor trailer

(4.12)

$$f_i = \left(\frac{d}{\sqrt{d^2 + (v \cdot t_i)^2}} \right)^3 \quad \text{where:}$$

f_i = Detector yield fraction at time t_i
 t_i = Time relative to center axis of portal

(4.13)

When $t_i = 0$, the point source is at the center axis of the radiation portal monitor and $f_i = 1$. This essentially means 100% of the count rate is seen by the detector at the RPM center axis. Simply multiplying the source count rate by f_i would be incorrect for the detector response because the detection interval was not a full second, but only 0.1 seconds. For example, assume the count rate in channel 5 was 200 cps. When the source passes the RPM's center axis, $f_i = 1$ and thus the detector response would be 200 counts in channel 5 with an un-scaled spatial response function; this is incorrect. Rather, the number of counts collected when the source passes the center axis should be $f_i = 1$ times the detection interval, 0.1 seconds, times the source count rate, 200 cps in channel 5. Thus, the number of counts collected during this interval would be 20 counts. Therefore, f_i needed to be scaled by the detection interval to preserve detection interval/source count rate agreement. This was

accomplished by multiplying f_i by the detection interval. This produced the final source spatial response function:

$$S_i = \text{intvl} \cdot f_i \quad \text{where:} \quad (4.14)$$

$S_i =$ Spatial response function

See the figure below for a comparison of f_i and the scaled spatial response function S_i .

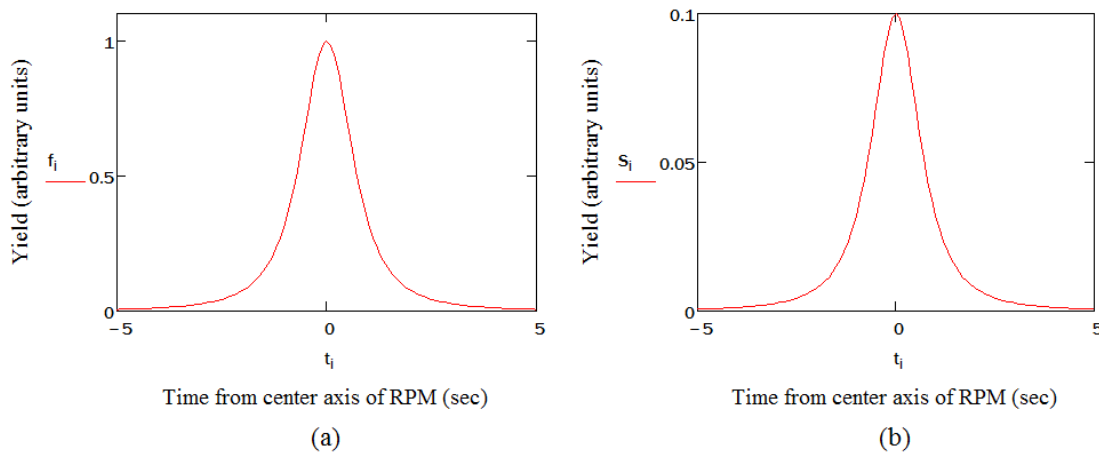


Figure 4.26: Detector yield from (a) f_i and (b) the spatial response function S_i

The calculation of the SRF used several objective real-world parameters in order to be calculated. Objective values included truck length, pillar spacing, occupied time and data collection time interval. The length of a tractor trailer varies, but one of the most common sizes is 70 feet. The typical width between RPM pillars is 2 meters. The detection interval was selected to be 0.1 seconds, the standard interval used in modern portal monitor installations. The first variable parameter chosen was occupied time. Since the speed varies by driver (they do not often observe the speed limit), this value is inherently changing. The occupied time for the vehicle runs was listed in the .DAT files and was extracted from them

manually for each of the nine vehicle profiles used. As an example, for a 70 foot truck with an occupied time of 10 seconds (resulting in a 7 feet per second vehicle velocity) and a data collection interval of 0.1 seconds, the SRF produced 101 discrete values used to scale count rate data vectors with dimensions 1x256 into vehicle profile matrices with dimensions 101x256. The number of rows in a SRF generated source profile depends on the occupied time. These injection source profiles were simply added to the already existing vehicle profiles extracted from the .PRO files.

To thoroughly test the simulated real-time advanced alarm algorithm emulator, the injection sources were scaled to eight different count rate levels: 277.5 cps, 555 cps, 1110 cps, 2220 cps, 3330 cps, 4440 cps, 5550 cps, and 6660 cps. These count rate levels represent the source count rate as seen by the detector (henceforth called “detected count rate”) and make no specific statement on the actual size of the source. This paper suggests this to be a more effective approach than stating a source activity. Declaring that an algorithm can detect an X μCi source is impossible in practice because a vehicle’s structural materials and cargo can radically affect the attenuation of source gamma rays and thus the RPM’s ability to detect a source. Inferring a source activity from the detector’s count rate is an extremely complex inverse problem requiring the vehicle’s geometric details, constituent materials and their geometry, and a sophisticated modeling program. Current research at Oak Ridge National Laboratory is investigating MCNP models of this problem. Many issues arise when attempting to model this situation; environments are complex with many variables (weather, brick buildings, etc.), vehicles are highly non-uniform and drastically affect source detection depending on their structure and cargo, source geometry is very difficult to extract from the

low resolution PVT spectra, and the light transmission/pulse response of PVT detectors is extremely difficult to model. This modeling work is outside the scope of this project and thus no attempt was made to quantify source activity from detector count rate.

Using the detected count rate rather than the source strength (activity), the results of the simulated real-time emulator give rise to conclusion statements like the following: “The advanced alarm algorithm will trigger a gross count alarm and an HEU alarm given a 3330 cps detected source count rate from an HEU source injected into the ‘Std 1’ cargo” or “The advanced alarm algorithm will trigger only a gross count alarm given a 1110 cps detected source count rate from HEU source injected into the ‘⁴⁰K High’ cargo”. Statements such as these are independent of actual source activity; whether the source is a heavily shielded 25 kg HEU source or an unshielded 2 kg HEU source, the statements only draw conclusions based upon the count rate seen by the detector. While these conclusions are more general than would be preferred, they are accurate as far as the simulated real-time emulator is accurate.

With the five non-alarming “Std #” cargos, two ⁴⁰K cargos, two granite cargos, six injection sources, and eight count rate levels, the total number of injection source/cargo combinations comes out to be 432. Having outlined the vehicle profiles to be tested and the scope of the conclusions made by the emulator, the development of the simulated real-time advanced alarm algorithm emulator can be discussed.

There were essentially four components to the emulator: the data acquisition system, the threshold calculator, the EW ratio calculator, and the alarm logic system. Each component was designed to mimic the behavior of an actual RPM system.

The data acquisition system (DAQ) contained the inputs, split the spectra into the individual windows, and buffered the occupied spectrum. There were two inputs into the emulator: the background spectra and the occupied spectra. The background spectra were a single 1x256 vector of background count rates extracted from the .HST files. In actuality, the .HST files contained the total accumulated background counts; these were easily converted into background count rates using the background count time and placed into an Excel file for input into the emulator. There was one background spectra per cargo. The name of the cargo was simply typed into an input block and when the emulator was run, the spectrum was automatically extracted. A DAQ function block was created that split the background spectrum into the energy windows; this block took the count rate vector as its input, summed the channels to create the window background count rates, and outputted seven signals, each for a specific energy window. Six of the signals were the HEU, WGPu, ^{137}Cs , DU, ^{40}K , and NOISE windows, while the seventh signal was the entire spectrum for gross count calculations. The DAQ block sampled this signal every 0.1 seconds; since the background signal was static, the sample rate was irrelevant. These seven signals were multiplexed to reduce the number of “wires” in the emulator. The multiplexed background signal was fed to two different emulator components: the threshold calculator and the EW ratio calculator.

The occupied spectra were the matrices generated by the spatial response function. The name of the cargo/injection source combination was typed into an input block and the emulator extracted the signal data from the source injected cargo Excel data file. The DAQ block again took the channel data as input and summed it, outputting seven signals that contained the window count rate. Recall that the sample rate of an RPM occupancy was 0.1

seconds, and consequently each source injected cargo matrix row was constructed to represent a 0.1 second accumulated count. Therefore, the DAQ components were set to sample every 0.1 seconds (this was why the background DAQ components sampled every 0.1 seconds, consistency with the occupied section). This was the method by which the emulator ran in simulated real-time; the sample rate set in the DAQ blocks and the time intervals in the first column of the source injected cargo matrices were treated by Simulink as real-time data. The program processed the data much faster than in real time, but treated the input information and the decision making as if the data was being processed in real-time. Thus, a 60 second simulation was run in fewer than three seconds, allowing for the various cargo/injection source combinations to be simulated in a short time.

After being split, these seven DAQ signals were passed into the buffer, where ten consecutive samples were summed to generate the 1-second count rates. These seven buffered signals were then multiplexed, again to reduce the emulator clutter. The multiplexed occupied signals were passed to two other emulator components: the energy window ratio calculator and the alarm logic system (only for the gross count alarm decisions).

The threshold calculator took the background signal and calculated the energy window and gross count thresholds. The multiplexed background signal was demultiplexed and fed into a function block that calculated each window's threshold. Each function block required the low window count rate and high window count rate. In addition, the high window count rate was the sum of the window count rates above the target window; therefore sum blocks were utilized to perform this action prior to being fed into the threshold

calculator function block. The energy windowing threshold function blocks utilized equation (4.3) and 4 as the sigma multiplier:

$$threshold = 4 \cdot \sigma_{R_i} = 4 \cdot \sqrt{B_i} \cdot \sqrt{1 + \frac{B_i}{\sum_{k=i+1}^6 B_k}} \quad (4.15)$$

Six thresholds were calculated, five for the energy windowing windows (HEU, WGPu, ¹³⁷Cs, DU, and ⁴⁰K), and one for the gross count threshold. These thresholds were multiplexed into a single wire and passed to the alarm logic system.

The energy windowing ratio calculator was very similar in design to the threshold calculator in that multiplexed signals are demultiplexed, passed to function blocks which output values that are then multiplexed. The energy windowing ratio function blocks require four inputs: low background window, high background window, low occupied window, and high occupied window. Recall that the low windows were simply the target windows, e.g. HEU window, while high windows require several high energy windows to be summed together, e.g. WGPu, ¹³⁷Cs, DU, ⁴⁰K, and NOISE windows. Sum blocks performed this action prior to the high windows being passed into the energy windowing ratio calculator function block. This block used the modified compensated ratio, equation (4.2):

$$R_i = S_i - \frac{B_i}{\sum_{k=i+1}^6 B_k} \cdot \sum_{k=i+1}^6 S_k$$

There were five ratio calculator blocks, one for each window tested. The outputs from these blocks were the ratio values to be tested against the threshold calculated by the threshold

calculator section of the emulator. The ratio values were multiplexed and passed on to the alarm logic system.

The alarm logic system tested both the gross count and energy windowing alarms. For the gross count alarm section, the gross count occupied signal and gross count threshold signal were fed into a “greater than” relational operator that outputted a 1 if the signal exceeded the threshold and a 0 otherwise. This signal was fed into an integrator that saturated at a value greater than or equal to 1; this assured that if the gross count occupied signal exceeded the gross count threshold at any time, the signal would remain at 1 and therefore continuously alarm.

For the five energy windowing alarms, each window’s ratio value and threshold signal also fed into a “greater than” relational operator block that again outputted a 1 if the ratio exceeded the threshold, and a 0 otherwise. The outputs from the relational operator block passed directly into integrators that saturated at 1. This was done for the reason stated above: to assure that any single occurrence of a ratio exceeding a threshold continuously outputs a 1. The outputs of the integrator blocks fed into the alarm logic blocks that implement the logic framework derived earlier in this chapter (see Table 4.2 and Figure 4.13 for the truth table and formal logic statements, respectively). The energy windowing logic blocks and the gross count integrator block sent the alarm decisions to the output display block; this block had six display lines, one for each available alarm. The display showed a 1 for each alarm that triggered and a 0 if the alarm does not trigger.

With the emulator fully developed, the results of the test can be discussed. There were essentially two groups of profiles tested: non-injection source POE runs and source-injected cargo profiles. The first group tested was the non-injection source POE profiles.

The non-injection source POE profiles consisted of a mix of alarming and non-alarming profiles. The five “Std #” cargos were tested to ensure that they did not alarm without an injection source so they could be considered “standard” non-alarming cargos. An addition five cargos were also tested; both these and the “standard” cargos were randomly selected from the list of non-alarming POE dataset (this non-alarming list refers to the 2117 POE cargos that did not alarm during the POE tests).

The 12 alarming POE cargos were also tested to see in which window(s) the alarms occurred. Of these 12, four were selected as the alarming cargos in which the injection sources were injected; two were the high- and low-activity ^{40}K cargos and two were the granite-bearing cargos. Testing the NORM-bearing cargos before testing the source-injected cargos allowed predictions to be made for the source-injected cargos. In other words, by having expectations for each source-injected cargo’s alarms, the advanced alarm algorithm’s accuracy was tested.

The tables below summarize the results of the advanced alarm algorithm simulated real-time emulator for the non-injection source POE profiles. The tables separate the non-alarming cargos from the alarming ones.

Table 4.6: Emulator Results- Non-alarming POE Profiles

	R00122	R00490	R001497	R03304	R05007
HEU Alarm	0	0	0	0	0
WGPu Alarm	0	0	0	0	0
Cs-137 Alarm	0	0	0	0	0
DU Alarm	0	0	0	0	0
NORM Alarm	0	0	0	0	0
GC Alarm	0	0	0	0	0
	Std 1	Std 2	Std 3	Std 4	Std 5
HEU Alarm	0	0	0	0	0
WGPu Alarm	0	0	0	0	0
Cs-137 Alarm	0	0	0	0	0
DU Alarm	0	0	0	0	0
NORM Alarm	0	0	0	0	0
GC Alarm	0	0	0	0	0

Table 4.7: Emulator Results- Alarming POE Profiles

	R04043/K-40 Low	R05015/K-40 High	R00436/Granite 1	R05027/Granite 2	R01287	R01507
HEU Alarm	0	0	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	1	1	0	0	0	0
GC Alarm	1	1	1	1	1	1
	R01590	R02043	R02119	R02168	R04424	R04425
HEU Alarm	1	0	0	0	0	0
WGPu Alarm	0	0	0	0	0	0
Cs-137 Alarm	0	0	0	0	1	1
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	1	1	0	0	1
GC Alarm	1	1	1	1	1	1

The tables above reveal that the advanced alarm algorithm did not alarm for any of the 10 non-alarming profiles. Since the five standard cargos were confirmed to be non-

alarming when tested without any injection sources, any alarms arising from a standard cargo/injection source combination were attributable to the injection source alone.

The advanced alarm algorithm did alarm for the alarming POE runs. There were several important observations made: 1) each cargo triggered the gross count alarm, confirming that the advanced algorithm's gross count alarm component performed correctly; 2) the ^{40}K cargos triggered NORM alarms, confirming that the advanced algorithm correctly identifies ^{40}K cargos as NORM-laden; and 3) the granite cargos caused unexpected energy windowing alarms. The advanced algorithm's ability to identify ^{40}K -laden cargos as NORM-bearing vehicles was a major design focus and this observation confirms the algorithm's ability to do so for basic ^{40}K -laden cargos (the injection source/cargo combination tests tested the ability to detect NORM while simultaneously detecting other radioactive materials). The two ^{40}K cargos selected for the injection source tests were confirmed to trigger gross count alarms and NORM alarms.

The third observation was the cause of most concern. Examination of accumulated granite spectra, including implementing energy window algorithms and visual inspections, did not reveal any energy windowing alarms. Therefore, it was assumed that granite cargos did not cause sufficient changes in spectrum-shape to cause energy windowing alarms. After running the advanced alarm algorithm emulator however, it was observed that granite cargos caused unanticipated energy windowing alarms in simulated real-time. Accumulating count rates then generating a spectrum, rather than checking spectrum shape in real-time, did not capture how the spectrum changes in time; this time-varying spectrum shifting resulted in different alarms than the alarms generated from an accumulated spectrum. It was concluded

that the real-time energy windowing calculations captured more detail and thus generated more accurate alarms than using accumulated count rate spectra.

The granite cargos most often triggered ^{137}Cs alarms but were also observed triggering WGPu alarms. The “Granite 1” and “Granite 2” cargos used for the injection source tests triggered WGPu and ^{137}Cs alarms, respectively. In addition, they both triggered gross count alarms. Given that most granite-bearing cargos would trigger ^{137}Cs alarms, the “Granite 2” cargo could be considered more representative of a typical granite-bearing cargo. While the “Granite 1” cargo profile, which triggers a WGPu alarm, may not occur as often as the “Granite 2” cargo profile, its inclusion in the injection source tests reveals the alarm implications of a granite cargo that causes a WGPu alarm. The granite alarms posed the most serious problem to the advanced alarm algorithm’s performance due to the WGPu and ^{137}Cs alarms the granite cargos can trigger. These alarms interfered with the HEU, WGPu, and ^{137}Cs alarm logic, resulting in potential false or missed alarms.

With the cargos tested, the injection source/cargo combinations could now be examined. For five standard cargos, two ^{40}K cargos, and two granite cargos, each testing six injection sources (^{57}Co , HEU, ^{133}Ba , ^{137}Cs , ^{60}Co , ^{40}K) at eight different count rates, 432 injection source/cargo combinations were tested. The results of the emulations were recorded in tables of the form seen below.

Table 4.8: Example Emulator Results Table

Example Emulator Results Table						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	1	0	0	0
WGPu Alarm	0	0	0	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	1	0
NORM Alarm	0	0	0	0	0	0
GC Alarm	0	1	0	1	1	1

The alarms can be seen on the left side while the injection sources are along the top of the table. An entry of “1” means an alarm was triggered and an entry of “0” means the alarm did not trigger. There was one table filled out for each cargo/injection source detected count rate combination, resulting in 72 tables; these tables can be found in the appendix.

To perform the statistical analysis of the emulator results, a confusion matrix was employed. A confusion matrix allows for the comparison between the advanced algorithm’s alarm results and the actual material contents of the cargos. The confusion matrix used in the analysis is seen below.

Table 4.9: Confusion Matrix used for Advanced Alarm Algorithm Statistical Analysis

		Advanced Algorithm Alarm Results	
		Alarm	No Alarm
Actual Vehicle Contents	Material Present	Correct Identification "True Positive"	Miss "False Negative"
	Material Absent	False Alarm "False Positive"	Correct Rejection "True Negative"

The confusion matrix was applied to each results table (e.g. Table 4.8) and produced the number of correct identifications, misses, false alarms, and correct rejections. The actual vehicle contents were determined by both the injection sources and the expected alarms for the cargo. The results of the non-injection source POE cargos served as the basis for comparing the injection source/cargo combination runs. Table 4.6 reveals that non-injection source standard cargos do not alarm. Therefore, the only expected alarms result from the injection sources. Table 4.7 shows the alarms expected for the granite and ^{40}K cargos when no injection source is present. For these cargos, it is expected that the alarm algorithm will identify both the cargo source and the injection source.

The pieces of data from the confusion matrix (correct identifications, false alarms, etc.) were compiled according to injection source detected count rate and cargo class (standard, ^{40}K , and granite) for each of the energy window alarms. The compiled numbers were used to calculate various measures of the advanced algorithm's performance. The table below lists the statistical measures calculated, their fundamental definitions, their interpretations within the context of this application, and the formulas used to calculate them.

Table 4.10: Statistical Measures Derived from Confusion Matrix

Statistic	Definition	Interpretation	Formula
Sensitivity	Proportion of positives that were correctly identified as positives	Probability that a material present in the vehicle will be identified correctly	$\frac{TP}{TP + FN}$
False Alarm Rate	Proportion of negatives that were incorrectly identified as positives	Probability that a material absent from the vehicle will be incorrectly identified as being present	$\frac{FP}{FP + TN}$
False Negative Rate	Proportion of positives that were incorrectly identified as negatives	Probability that a material present in the vehicle will not cause the correct alarm	$\frac{FN}{FN + TP}$
Specificity	Proportion of negatives that were correctly identified as negatives	Probability that a material absent from the vehicle will not be implicated with an alarm	$\frac{TN}{TN + FP}$
Precision	Proportion identified as positive that were indeed positive	Probability that an alarm correctly identifies that specific material	$\frac{TP}{TP + FP}$
Accuracy	Proportion of true results	Probability that an alarm or a non-alarm is correct	$\frac{TP + TN}{TP + TN + FP + FN}$
Matthews Correlation Coefficient	Correlation coefficient between the observed and predicted classifications	0 is average random prediction, 1 is perfect prediction, and -1 is perfect inverse prediction	$\frac{TP \cdot TN - FP \cdot FN}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}}$
F1 Score	Measure of the algorithm's performance	Weighted average of the precision and sensitivity	$\frac{2 \cdot TP}{(TP + FN) + (TP + FP)}$

Of these measures, sensitivity, false alarm rate, precision, Matthews correlation coefficient, and F1 score were considered to be the most important. The sensitivity and precision are both measures that assess the ability of the advanced algorithm to identify the contents of a vehicle. False alarm rate measures the tendency of the algorithm to incorrectly identify materials; this is important when considering the effect that false alarms have on RPM workers. The Matthews correlation coefficient and F1 score are summary measures

that attempt to evaluate overall algorithm performance with a single value. The tables of these statistics can be found in the appendix.

An important note about this statistical analysis was that not every measure could be calculated for every energy windowing alarm/detected count rate/cargo class combination. This was a result of some alarms having empty figures from the confusion matrix, i.e. some tests resulted in only correct rejections with no correct identifications, false alarms, or false negatives. In these situations, many of the statistics were not able to be calculated due to the denominators being equal to zero. This was a somewhat common occurrence for the DU alarms because ^{60}Co was an imperfect DU proxy source, thus no injection source was capable of producing expected DU alarms. If there were no cargos with DU “actually present” in the vehicles, there could be no true statistical analysis of the results.

The five critical statistics for each energy window alarm were plotted against injection source detected count rate to assess how alarm effectiveness depended on the injection source count rate. By doing this, the advanced algorithm could be stated as having “X percent sensitivity to HEU when the HEU detected count rate is above Y counts per second”. Statements of this type are extremely useful in communicating the algorithm’s effectiveness. A single plot contains the data for all five energy windowing alarms at a single injection source detected count rate and a single cargo class, resulting in 15 “statistic” versus injection source detected count rate plots. These plots and their analyses can be seen on the next several pages.




Figure 4.27: Plot of Sensitivity versus Injection Source Detected Count Rate for Standard Cargos




Figure 4.28: Plot of Sensitivity versus Injection Source Detected Count Rate, ⁴⁰K Cargos




Figure 4.29: Plot of Sensitivity versus Injection Source Detected Count Rate, Granite Cargos

The three sensitivity plots revealed different algorithm response between the energy window alarms and the cargo classes. The most common cargo class, standard cargos, showed that the sensitivity increased rapidly with increasing injection source detected count rate. Both the HEU and WGPu alarms reached 100% sensitivity between 555 and 1110 cps detected count rate, while the NORM alarm required more and the ¹³⁷Cs alarm required less. For ⁴⁰K cargos, it was shown that masking effects increased the number of counts required for 100% sensitivity to between 2220 and 3330 for the HEU alarm, and 1110 and 2220 for the WGPu alarm. It also showed that masking did not significantly affect alarms above the ¹³⁷Cs window. Clearly, the granite cargos caused serious identification problems, resulting in less than 60% sensitivity for the HEU and WGPu alarms. This was caused by the WGPu and ¹³⁷Cs alarms the granite cargo itself caused; these alarms impede the alarm logic from identifying HEU and WGPu correctly.




Figure 4.30: Plot of False Alarm Rate versus Injection Source Detected Count Rate, Standard Cargos




Figure 4.31: Plot of False Alarm Rate versus Injection Source Detected Count Rate, ⁴⁰K Cargos




Figure 4.32: Plot of False Alarm Rate versus Injection Source Detected Count Rate, Granite Cargos

There was not a tremendous amount of detail to be gleaned from the false alarm rate plots, but the results were positive. For the standard cargos there was only one alarm with a non-zero false alarm rate; unfortunately, this was the HEU alarm which was arguably the most important alarm. The HEU false alarms were caused by a low count rate ^{133}Ba source, the WGPu proxy. When discussing the appropriateness of ^{133}Ba as a WGPu proxy, it was determined that ^{133}Ba has more low energy gamma rays than true WGPu, resulting in increased counts in the HEU window. It is not expected that this situation will arise during actual transportation, so these HEU false alarms do not cause much concern. There were no false alarms for the ^{40}K cargos due to the low count rate ^{133}Ba source being masked by the ^{40}K 's high energy gamma rays. The granite cargos caused a 12.5% false alarm rate for HEU, again due to the ^{137}Cs and WGPu alarms caused by the granite cargos themselves. The alarm

logic interference can result in both ^{133}Ba and ^{137}Cs causing false HEU alarms. The implication of this effect will be discussed further in the Chapter 5.




Figure 4.33: Plot of Precision versus Injection Source Detected Count Rate, Standard Cargos




Figure 4.34: Plot of Precision versus Injection Source Detected Count Rate, ⁴⁰K Cargos




Figure 4.35: Plot of Precision versus Injection Source Detected Count Rate, Granite Cargos

Precision, for this application, was the probability that an energy window alarm correctly triggered, e.g. for a 555 cps injection source detected count rate in a standard cargo, 90% of HEU alarms were a result of an HEU source being present. The precision statistics were very good for standard cargos; the HEU algorithm reached 100% precision between 555 and 1110 cps, and the other algorithms reached 100% precision for injection source detected count rates less than or equal to 277.5 cps. For ^{40}K cargos, the precision statistic was at 100% for the HEU, WGPu, and ^{137}Cs window alarms. The DU alarm was incalculable do to a lack of correct identifications, but was expected to have a precision of 0% across the board. This is due to the alarm logic; it is fundamentally impossible for both the DU and NORM alarms to trigger at the same time, thus a sufficient amount of ^{40}K can always prevent a DU alarm from triggering. This same alarm logic problem was seen in the 67% precision for HEU alarms for granite cargos. The WGPu alarm triggered by some granite cargos will always prevent an HEU source from triggering an alarm.




Figure 4.36: Plot of Matthews Correlation Coefficient versus Injection Source Detected Count Rate, Standard Cargos




Figure 4.37: Plot of Matthews Correlation Coefficient versus Injection Source Detected Count Rate, ⁴⁰K Cargos




Figure 4.38: Plot of Matthews Correlation Coefficient versus Injection Source Detected Count Rate, Granite Cargos

For standard cargoes, the HEU and WGPu alarms appeared to reach a Matthews correlation coefficient of 1 between 555 and 1110 cps detected count rate, while the NORM alarm reached 1 between 1110 and 2220 cps. For ^{40}K cargoes, it took higher count rates for the HEU, WGPu, and HEU alarms to reach maximum detection ability. It was clear that the algorithm faced more serious detection issues for granite cargoes due to the coefficient values being farther from 1 for the HEU, WGPu, and ^{137}Cs . The alarm logic issues discussed earlier clearly impacted this statistic for the granite cargoes. Considering the small proportion of ^{40}K and granite cargoes compared to standard cargoes, and the effectiveness of the algorithm for standard and ^{40}K cargoes, the analysis of the Matthews correlation coefficient showed the algorithm to be effective in identifying the ^{40}K cargoes/injection sources as NORM and the illicit materials as such.




Figure 4.39: Plot of F1 Score versus Injection Source Detected Count Rate, Standard Cargos




Figure 4.40: Plot of F1 Score versus Injection Source Detected Count Rate, ⁴⁰K Cargos




Figure 4.41: Plot of F1 Score versus Injection Source Detected Count Rate, Granite Cargos

The F1 score, being a weighted average of precision and sensitivity, is an effective measure in assessing the advanced algorithm’s identification ability. In addition, due to the way it is calculated, the F1 score revealed more complete results than the Matthews correlation coefficient. The F1 score results did appear to match the Matthews correlation coefficient results closely, revealing that identification ability for standard cargoes reached its maximum with fewer than 2220 cps for all alarms. For ^{40}K cargoes, the alarms reached maximum effectiveness with fewer than 3330 cps for all alarms. Granite cargoes again did not perform as well as the other two, but this was expected.

After performing the statistical analysis on the advanced alarm algorithm’s simulated real-time emulator, it was determined that the advanced algorithm was very effective at detecting and identifying both illicit nuclear materials and ^{40}K NORM sources. The

algorithm proved most effective for standard cargos, the most common variety. The illicit material identification performance somewhat decreased for ^{40}K cargos due to masking effects; the injection source detected count rates required for 100% HEU and WGPu alarm sensitivity increased to approximately 3330 cps and 2220 cps respectively. While the percent increases were drastic (200% and 100% respectively), the detected count rates were still relatively low when considered in perspective; a 2 kg equivalent unshielded HEU source produced a count rate of over 7800 cps. Granite cargos were shown to decrease identification ability significantly. This was not a result of masking, but of alarm logic interference. As stated above several times, the WGPu and ^{137}Cs alarms triggered by granite cargos can prevent HEU and WGPu sources from triggering the proper alarms. Combining an HEU injection source with a granite cargo that triggers a ^{137}Cs alarm will cause an HEU and ^{137}Cs alarm (both correct), but combining an HEU source with a granite cargo that generates a WGPu alarm will result in only the WGPu alarm. While the WGPu alarm is correct, the HEU alarm is missed. Similarly, a WGPu source combined with a granite cargo that triggers a ^{137}Cs will trigger the ^{137}Cs and HEU alarm, missing the WGPu alarm. The triggering of the HEU alarm by this combination was likely due to the fact that ^{133}Ba was an imperfect WGPu proxy. The standards/calibration source problem will be discussed in chapter 5.

Given the complications from the alarm logic and the various combinations of alarms that arise from different cargo combinations, an “alarm library” was developed to help solve this problem. The library consists of possible vehicle contents given certain alarm combinations. While there are nearly endless combinations of potential sources, the library

is concerned with the target nuclides and potential NORM sources. The table will highlight two categories of materials: likely materials- materials implicated by the advanced algorithm's alarms and likely present in the vehicle; and possible materials- materials that could possibly be in the cargo but were not explicitly identified by the algorithm. The possible materials are simply the target material in the window directly below the alarming window. For example, a WGPu alarm will have WGPu and granite as the "likely" materials, and HEU as a "possible" material. In this organization, there is higher confidence that the "likely" materials are present than the "possible" materials because the "likely" materials were directly implicated and the algorithm is assumed to be an effective identifier of vehicle contents. The "possible" material decisions are based on the shortcomings of the algorithm's alarm logic and not on concrete evidence; this in essence loses a degree of freedom for the "possible" decisions, decreasing the amount of confidence that can be placed in them.

Table 4.11: Alarm Library

	Alarms	Likely Material	Possible Material
Single Alarms	HEU	HEU	-
	WGPu	WGPu, granite	HEU
	Cs-137	Cs-137, granite	WGPu
	DU	DU	Cs-137, granite
	NORM	K-40	DU
Multiple Alarms	HEU/Cs-137	Cs-137, granite, HEU	WGPu
	HEU/DU	DU, HEU	Cs-137, granite
	HEU/NORM	K-40, HEU	DU
	WGPu/DU	DU, WGPu, granite	HEU, Cs-137
	WGPu/NORM	WGPu, granite, K-40	HEU, DU
	Cs-137/NORM	K-40, Cs-137, granite	WGPu, DU
	HEU/Cs-137/NORM	HEU, Cs-137, granite, K-40	WGPu, DU

It is important to highlight the alarm library is intended to be a resource to assist RPM installation workers during secondary inspections. It is not a definitive list of what is or is not in a vehicle. The alarm library is seen in the table above.

An effective method to confirm WGPU presence in a vehicle is the neutron signature detected by the RPM. Neutron background radiation levels are typically very low, and the neutron signature from WGPU makes it more difficult to transport undetected through RPMs. Thus, using neutron alarms to discriminate between WGPU-triggered and granite-triggered WGPU alarms could be an effective addition to future algorithms; this mechanism would be known as the “WGPU neutron discriminator”. A simple adjustment to the alarm logic would provide the discriminator (neutron data was not provided so the WGPU neutron discriminator could not be implemented in the current iteration of the advanced algorithm). The alarm logic statements below give an example of how a WGPU neutron discriminator would be seamlessly added to the advanced algorithm, resulting in a new “Granite” alarm and a “WGPU” alarm.

$$A_{NORM} = 1 \text{ IF } (R_{40K} > T_{40K})$$

$$A_{NORM} = 0 \text{ OTHERWISE}$$

$$A_{DU} = 1 \text{ IF } (R_{DU} > T_{DU}) \wedge \neg(R_{40K} > T_{40K})$$

$$A_{DU} = 0 \text{ OTHERWISE}$$

$$A_{137Cs} = 1 \text{ IF } (R_{137Cs} > T_{137Cs}) \wedge \left[\neg(R_{DU} > T_{DU}) \vee \left\{ (R_{DU} > T_{DU}) \wedge (R_{40K} > T_{40K}) \right\} \right] \wedge \neg(A_{DU} = 1)$$

$$A_{137Cs} = 0 \text{ OTHERWISE}$$

$$A_{WGPu} = 1 \text{ IF } (R_{WGPu} > T_{WGPu}) \wedge (A_{neutron} = 1) \wedge \left[\neg(R_{137Cs} > T_{137Cs}) \vee \left\{ (R_{137Cs} > T_{137Cs}) \wedge (R_{DU} > T_{DU}) \right\} \right]$$

$$A_{WGPu} = 0 \text{ OTHERWISE}$$

$$A_{Granite} = 1 \text{ IF } (R_{WGPu} > T_{WGPu}) \wedge (A_{neutron} = 0) \wedge \left[\neg(R_{137Cs} > T_{137Cs}) \vee \left\{ (R_{137Cs} > T_{137Cs}) \wedge (R_{DU} > T_{DU}) \right\} \right] \wedge \neg(A_{137Cs} = 1)$$

$$A_{Granite} = 0 \text{ OTHERWISE}$$

$$A_{HEU} = 1 \text{ IF } (R_{HEU} > T_{HEU}) \wedge \left[\neg(R_{WGPu} > T_{WGPu}) \vee \left\{ (R_{WGPu} > T_{WGPu}) \wedge (R_{137Cs} > T_{137Cs}) \right\} \right] \wedge \neg[(A_{WGPu} = 1) \vee (A_{Granite} = 1)]$$

$$A_{HEU} = 0 \text{ OTHERWISE}$$

where:

R_i = Modified compensated ratio in window i

T_i = Threshold in window i

A_i = Alarm in window i if $A_i = 1$

This concludes the development and analysis of the advanced alarm algorithm. The major concern that arose from the emulator study was the negative influence of the granite cargos' unexpected WGPu and ^{137}Cs alarms; this is addressed at length in the following chapter. In addition, further elaboration on the successes and failings of the advanced alarm algorithm, including comparisons between it and the ^{57}Co and ^{133}Ba algorithms and comments on calibration source standards, are discussed.

Chapter 5 – Discussion

With the conclusion of the sensitivity analysis, it was observed that the advanced algorithm performance was superior to the two commercial algorithms in several ways. The advanced algorithm's HEU detection prowess far surpassed that of either the ^{57}Co or ^{133}Ba algorithms. Detection of HEU is a major goal of RPM installations, so confirming the advanced algorithm's effectiveness in this aspect shows great promise for the future implementation of this algorithm. While the ^{133}Ba algorithm more effectively detected the ^{133}Ba injection source, it is believed the advanced algorithm would detect WGPu more effectively.

The primary reason for the advanced algorithm's superior HEU detection ability is suspected to be the precise calibration of the HEU window using an actual HEU source. The algorithm calibrated with a ^{57}Co source was considerably less effective at detecting HEU. From the emulator results, it was clear the ^{137}Cs and ^{40}K windows also performed very well; again attributable to the fact that the ^{137}Cs and ^{40}K windows were calibrated using actual ^{137}Cs and ^{40}K sources. These observations lead to the conclusion that, to optimally detect a specific nuclide, it is of the utmost importance to calibrate the window with the target nuclide.

This highlights a deficiency in the advanced algorithm's simulated real-time emulator tests. The WGPu and DU windows were not tested as adequately as desired because WGPu and DU injection source spectra were not available. If WGPu and DU spectra can be taken with a PVT detection system and the data is provided, recalibration of the WGPu and DU windows can be readily performed to be more precise. Injection sources made from these

spectra would then be used in the emulator, allowing more accurate statistics for these two windows to be obtained.

An obstacle preventing the use of HEU, WGPu, and DU as standard calibration sources is the sensitive nature of these nuclides; it is impossible for these sources to be readily available for use, even by RPM manufacturers. Consequently, an alternative calibration technique could be developed to avoid the need for these sensitive materials. This is a difficult task due to each RPM manufacturer's individual electronic components. For example, RPMs tested at ORNL often use a ^{137}Cs calibration source and adjust the gain such that a point 2/3 down the Compton edge (recall it is more of a slope) corresponds to channel 50 in a 256 channel data acquisition system (channel 100 in a 512 channel data acquisition system, etc.). This method is used to decrease the prevalence of electronic noise in the lowest channels and make the detector response uniform from one test to another, but could be used as a starting point for an alternative calibration standard. Using the assumption that PVT/phototube response is linear with regard to incident gamma ray energy for all PVT-based systems, all systems calibrated with the ^{137}Cs Compton edge in channel 50 for 256 channel MCAs could use the same window discriminators as the advanced algorithm (simple scaling results in the correct window discriminators for 512, 1024, etc. channel MCAs). This proposed standard could avoid the need to calibrate each individual PVT RPM for every target nuclide, but requires much further investigation before it can come under consideration as an alternative.

The most significant issue the advanced alarm algorithm faced was the granite problem: the fact that granite-bearing cargos triggered WGPu and ^{137}Cs alarms unexpectedly.

This is a difficult problem to solve, but there are some partial solutions available. A partial solution, mentioned briefly at the end of chapter 4, is the addition of a WGPu neutron discriminator. By linking the WGPu alarm logic to a neutron alarm, the algorithm could distinguish between WGPu alarm-triggering granite cargos and actual WGPu. Intuitively, this addition would allow better WGPu discrimination by taking advantage of the spontaneous neutrons generated by ^{239}Pu (21.8 neutrons per second-kilogram). Given the relative difficulty of shielding neutrons and naturally low neutron background count rates, this neutron discriminator could be extremely useful.

Dealing with granite cargos that trigger ^{137}Cs alarms poses a more difficult problem as ^{137}Cs sources are gamma emitters and not accompanied by neutron emission. With current PVT RPM technology, there does not appear to be an effective method to discriminate between ^{137}Cs alarm-triggering granite cargos and actual ^{137}Cs sources during the primary inspection (that is, a vehicle's first pass through the RPM). A secondary inspection utilizing higher resolution detectors should easily be able to detect whether the cargo is granite or ^{137}Cs . Given the *a priori* knowledge provided by the advanced algorithm, RPM installation workers would be searching in particular for the ^{137}Cs 662 keV photopeak or a series of gamma rays caused by natural decay series. The software used with the secondary inspection detectors identifies the nuclides automatically, but providing additional knowledge to the workers from primary inspections that agrees or disagrees with the secondary inspection results could increase the confidence of workers that their conclusions are either correct or require even further investigation.

If a WGPu neutron discriminator is not employed and granite cargos are conflated with WGPu and ^{137}Cs sources, a more thorough investigation of WGPu and ^{137}Cs alarms should be implemented in any case for the following reason: granite cargos are very effective at shielding SNM sources, posing a serious masking threat. This situation is illustrative of how the alarm library should be used by RPM workers. Given certain alarm outputs by the advanced algorithm, RPM workers should focus their analysis on the “likely” and “possible” sources identified in the library during secondary inspection. The alarm library attempts to ensure that target sources potentially masked by either high energy masking or alarm logic interference are not ignored by workers.

Some of the most critical future research to improve the field use of RPMs lies in the area of human factors engineering. Adding nuclide identification capabilities to RPM installations is only useful if the information collected by the advanced algorithm is effectively transmitted to the cargo inspectors. Alarm interference and, to a smaller extent, high energy masking further complicate the issue by suppressing alarms for materials that could be present. The application of the alarm library attempts to solve the problem by listing “likely” and “possible” sources, but the effects of the alarm library on inspector behavior has yet to be analyzed. An ideal human factors test is presented in the table below. To test the various capabilities of the alarm algorithm and alarm library to identify the cargo materials correctly, examining different combinations of sources and cargos would be necessary.

Table 5.1: Alarm Library Human Factors Experiment

Vehicle	Secondary Inspector Group (block)							
	Group 1		Group 2		Group ...		Group b	
	Alarm Library		Alarm Library		Alarm Library		Alarm Library	
	Used	Not Used	Used	Not Used	Used	Not Used	Used	Not Used
1	#	#	#	#	#	#	#	#
2	#	#	#	#	#	#	#	#
...	#	#	#	#	#	#	#	#
a	#	#	#	#	#	#	#	#

To test if the alarm libraries had a significant effect on the secondary inspectors' cargo identification capability and if there is an interaction between the cargo contents and the alarm library's effect, a blocked factorial design is suggested as shown in Table 6.1. The results of the experiment could then be used to determine which cargos' identification rates were positively affected by the alarm library. The experiment seeks to ignore variability due to different secondary inspectors by blocking; the hypothesis that the inspectors had no effect would be tested only marginally due to the experimental design. This experiment is resource intensive (time, people, sources, trucks, etc.) so the likelihood of adjustments to the design to fit a budget is high. However, it is believed this experiment would confirm whether the alarm library itself is effective.

A very similar design to this one could be used to test whether there is a significant difference between correct cargo identification for the full advanced alarm algorithm and a pure gross count algorithm. The issue is that these experiments are expensive, whereas upgrading RPMs with the advanced algorithm is very low cost. It is also highly doubtful that

the additional information provided by the advanced algorithm would detrimentally affect the ability of secondary inspectors to identify radioactive sources in vehicles.

Chapter 6 – Summary and Recommendations for Further Studies

This study investigated the development of an advanced alarm algorithm based on energy windowing to enhance the RPM's cargo screening capability. An advanced alarm algorithm was successfully developed and, in comparison to the status quo of pure gross counting RPMs, the new algorithm was shown to surpass them in effectiveness. The advanced alarm algorithm has been shown to be an effective mechanism to increase the identification capability of PVT radiation portal monitors. The inclusion of a gross count alarm in the advanced algorithm ensured that high activity cargos, whether benign or malicious, would trigger alarms and provoke suspicion. The addition of the energy windowing alarms added the ability to identify certain target nuclides, providing more information from the primary inspection than would ever be possible from gross count only RPM algorithms. It is believed the identification ability provided by the advanced algorithm would increase secondary inspection speed and efficiency by granting RPM workers *a priori* knowledge about the potential contents of the alarming vehicles. By utilizing the alarm library, RPM workers would know specific sources to be aware of; these could be sources directly implicated by the alarms or sources the RPM cannot "see" due to either alarm interference or high energy masking.

Possibly the most significant aspect of the advanced alarm algorithm is its ease of implementation. Theoretically, any RPM system can be upgraded with the advanced algorithm with simple software modifications. In essence, the core alarm program requires only basic modifications to the code to transform the system from a pure gross count algorithm to the advanced alarm algorithm. The transition from gross count algorithm to

advanced algorithm would be a very low cost upgrade, especially when compared to the alternative of adding spectroscopic portal monitors for secondary inspections to RPM installations.

The low cost improvement requires no hardware modifications and vastly increases the ability of PVT systems to discriminate between ^{40}K NORM bearing cargos and cargos carrying illicit nuclear materials. The additional information provided by the advanced algorithm could also increase the speed and efficiency of secondary inspections, reducing the disruption of commercial traffic. With no perfect solution, the advanced alarm algorithm is believed to be one of the most effective methods in improving PVT RPMs and avoiding the need for more expensive portal monitoring technologies.

While the simulated real-time emulator was an effective proving ground for the advanced alarm algorithm, it is necessary to test the algorithm in an actual RPM installation in preparation for practical implementation. It is expected that adding the advanced algorithm to an existing RPM system is not difficult, thus the prospect for future laboratory testing is high.

References

- Blessinger, C. "Appendix C: Derivation of the Statistical Uncertainty of the Energy Windowing Algorithm." 2010.
- Ely, James, Richard Kouzes, John Schweppe, Edward Siciliano, Denis Strachan, and Dennis Weier. "The use of energy windowing to discriminate SNM from NORM in radiation portal monitors." *Nuclear Instruments and Methods in Physics Research*, 2006: 373-387.
- Kouzes, R T, J H Ely, B D Milbrath, J E Schweppe, E R Siciliano, and D C Stromswold. "Spectroscopic and Non-Spectroscopic Radiation Portal Applications to Border Security." *IEEE Nuclear Science Symposium Conference*. Puerto Rico, 2005. 321-325.
- Lo Presti, Charles A, Dennis R Weier, Richard T Kouzes, and John E Schweppe. "Baseline suppression of vehicle portal monitor gamma count profiles: A characterization study." *Nuclear Instruments and Methods in Physics Research*, 2006: 281-297.
- NNSA Public Affairs. "Fact Sheet: NNSA's Second Line of Defense Program." *National Nuclear Security Administration Website*. February 2, 2010.
<http://nnsa.energy.gov/mediaroom/factsheets/nnsassecondlineofdefenseprogram> (accessed September 2010).
- Runkle, Robert C, Theresa M Mercier, Kevin K Anderson, and Deborah K Carlson. "Point Source Detection and Characterization for Vehicle Radiation Portal Monitors." *IEEE Transaction on Nuclear Science* 52, no. 6 (2005).
- Siciliano, E R, J H Ely, R T Kouzes, B D Milbrath, J E Schweppe, and D C Stromswold. "Comparison of PVT and NaI(Tl) scintillators for vehicle portal monitor applications." *Nuclear Instruments and Methods in Physics Research*, 2005: 647-674.

Appendix

Appendix A – Sensitivity Analysis Plots




Figure A.1: Sensitivity plot of ^{57}Co injection source and Std1 cargo combination




Figure A.2: Sensitivity plot of ^{57}Co injection source and Std2 cargo combination




Figure A.3: Sensitivity plot of ^{57}Co injection source and Std3 cargo combination




Figure A.4: Sensitivity plot of ^{57}Co injection source and Std4 cargo combination




Figure A.5: Sensitivity plot of ^{57}Co injection source and Std5 cargo combination




Figure A.6: Sensitivity plot of ^{57}Co injection source and ^{40}K High cargo combination




Figure A.7: Sensitivity plot of ⁵⁷Co injection source and ⁴⁰K Low cargo combination




Figure A.8: Sensitivity plot of ⁵⁷Co injection source and Granite 1 cargo combination




Figure A.9: Sensitivity plot of ⁵⁷Co injection source and Granite 2 cargo combination




Figure A.10: Sensitivity plot of HEU injection source and Std1 cargo combination




Figure A.11: Sensitivity plot of HEU injection source and Std2 cargo combination




Figure A.12: Sensitivity plot of HEU injection source and Std3 cargo combination




Figure A.13: Sensitivity plot of HEU injection source and Std4 cargo combination




Figure A.14: Sensitivity plot of HEU injection source and Std5 cargo combination




Figure A.15: Sensitivity plot of HEU injection source and ⁴⁰K High cargo combination

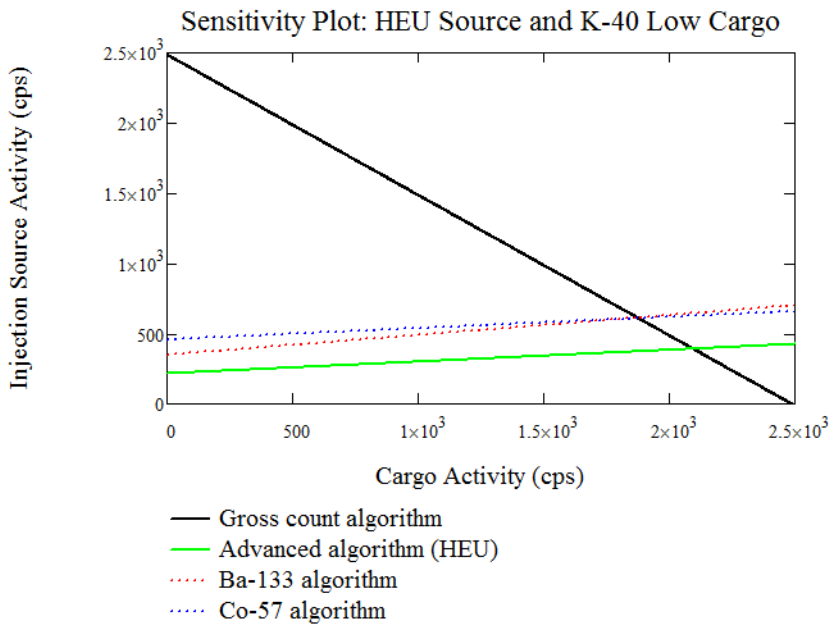


Figure A.16: Sensitivity plot of HEU injection source and ⁴⁰K Low cargo combination




Figure A.17: Sensitivity plot of HEU injection source and Granite 1 cargo combination




Figure A.18: Sensitivity plot of HEU injection source and Granite 2 cargo combination




Figure A.19: Sensitivity plot of ^{133}Ba injection source and Std1 cargo combination




Figure A.20: Sensitivity plot of ^{133}Ba injection source and Std2 cargo combination




Figure A.21: Sensitivity plot of ^{133}Ba injection source and Std3 cargo combination




Figure A.22: Sensitivity plot of ^{133}Ba injection source and Std4 cargo combination




Figure A.23: Sensitivity plot of ¹³³Ba injection source and Std5 cargo combination




Figure A.24: Sensitivity plot of ¹³³Ba injection source and ⁴⁰K High cargo combination




Figure A.25: Sensitivity plot of ¹³³Ba injection source and ⁴⁰K Low cargo combination




Figure A.26: Sensitivity plot of ¹³³Ba injection source and Granite 1 cargo combination




Figure A.27: Sensitivity plot of ¹³³Ba injection source and Granite 2 cargo combination




Figure A.28: Sensitivity plot of ¹³⁷Cs injection source and Std1 cargo combination




Figure A.29: Sensitivity plot of ^{137}Cs injection source and Std2 cargo combination




Figure A.30: Sensitivity plot of ^{137}Cs injection source and Std3 cargo combination




Figure A.31: Sensitivity plot of ^{137}Cs injection source and Std4 cargo combination




Figure A.32: Sensitivity plot of ^{137}Cs injection source and Std5 cargo combination




Figure A.33: Sensitivity plot of ¹³⁷Cs injection source and ⁴⁰K High cargo combination




Figure A.34: Sensitivity plot of ¹³⁷Cs injection source and ⁴⁰K Low cargo combination




Figure A.35: Sensitivity plot of ¹³⁷Cs injection source and Granite 1 cargo combination




Figure 1.36: Sensitivity plot of ¹³⁷Cs injection source and Granite 2 cargo combination




Figure A.37: Sensitivity plot of ^{60}Co injection source and Std1 cargo combination




Figure A.38: Sensitivity plot of ^{60}Co injection source and Std2 cargo combination




Figure A.39: Sensitivity plot of ⁶⁰Co injection source and Std3 cargo combination




Figure A.40: Sensitivity plot of ⁶⁰Co injection source and Std4 cargo combination




Figure A.41: Sensitivity plot of ^{60}Co injection source and Std5 cargo combination




Figure A.42: Sensitivity plot of ^{60}Co injection source and ^{40}K High cargo combination




Figure A.43: Sensitivity plot of ⁶⁰Co injection source and ⁴⁰K Low cargo combination




Figure A.44: Sensitivity plot of ⁶⁰Co injection source and Granite 1 cargo combination




Figure A.45: Sensitivity plot of ^{60}Co injection source and Granite 2 cargo combination




Figure A.46: Sensitivity plot of ^{40}K injection source and Std1 cargo combination

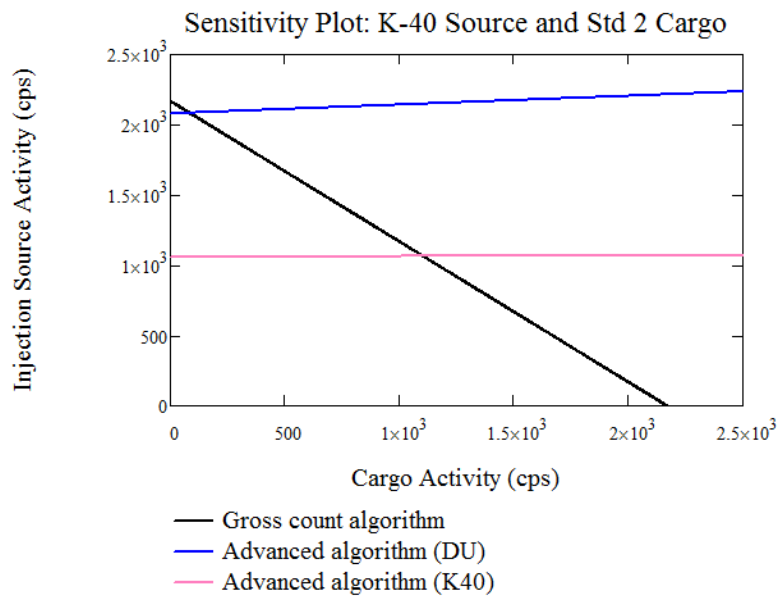


Figure A.47: Sensitivity plot of ⁴⁰K injection source and Std2 cargo combination




Figure A.48: Sensitivity plot of ⁴⁰K injection source and Std3 cargo combination




Figure A.49: Sensitivity plot of ⁴⁰K injection source and Std4 cargo combination




Figure A.50: Sensitivity plot of ⁴⁰K injection source and Std5 cargo combination




Figure A.51: Sensitivity plot of ⁴⁰K injection source and ⁴⁰K High cargo combination

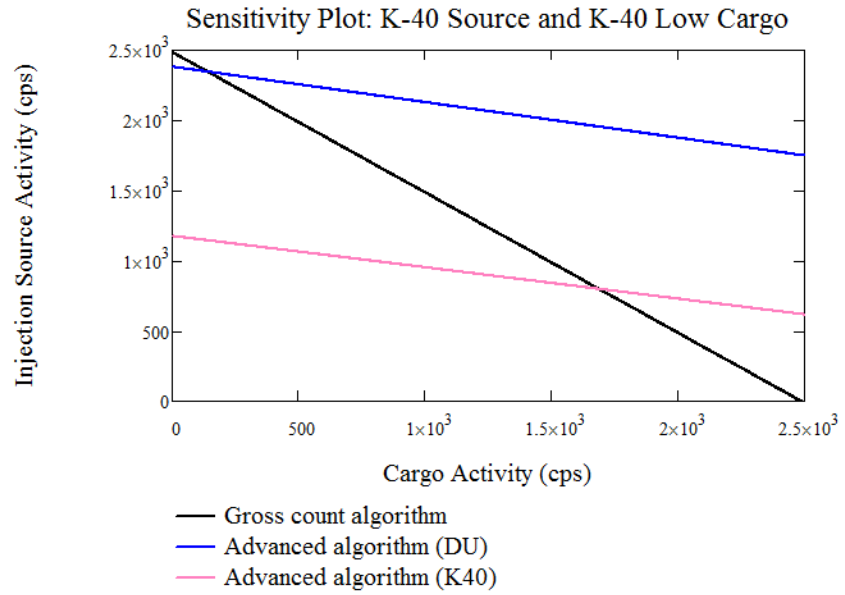


Figure A.52: Sensitivity plot of ⁴⁰K injection source and ⁴⁰K Low cargo combination




Figure A.53: Sensitivity plot of ⁴⁰K injection source and Granite 1 cargo combination




Figure A.54: Sensitivity plot of ⁴⁰K injection source and Granite 2 cargo combination

Appendix B – Simulated Real-Time Emulator Results Tables

Table B.1: Standard 1 Cargo Results, All Injection Source Count Rates

Std 1 Cargo, 277.5 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	1	0	0	0
WGPu Alarm	0	0	0	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	1	0
NORM Alarm	0	0	0	0	0	0
GC Alarm	0	1	0	1	1	1

Std 1 Cargo, 555 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	0
GC Alarm	1	1	1	1	1	1

Std 1 Cargo, 1110 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Std 1 Cargo, 2220 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Std 1 Cargo, 3330 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Std 1 Cargo, 4440 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Std 1 Cargo, 5550 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Std 1 Cargo, 6660 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Table B.2: Standard 2 Cargo Results, All Injection Source Count Rates

Std 2 Cargo, 277.5 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	0	0	0	0	0	0
WGPu Alarm	0	0	0	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	1	0
NORM Alarm	0	0	0	0	0	0
GC Alarm	0	0	0	0	0	0

Std 2 Cargo, 555 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	0
GC Alarm	0	0	0	0	0	0

Std 2 Cargo, 1110 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Std 2 Cargo, 2220 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Std 2 Cargo, 3330 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Std 2 Cargo, 4440 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Std 2 Cargo, 5550 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Std 2 Cargo, 6660 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Table B.3: Standard 3 Cargo Results, All Injection Source Count Rates

Std 3 Cargo, 277.5 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	0	0	0	0	0	0
WGPu Alarm	0	0	0	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	1	0
NORM Alarm	0	0	0	0	0	0
GC Alarm	0	0	0	0	0	0

Std 3 Cargo, 555 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	0	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	1	0
NORM Alarm	0	0	0	0	0	0
GC Alarm	0	0	0	0	0	0

Std 3 Cargo, 1110 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	0
GC Alarm	0	1	1	1	1	1

Std 3 Cargo, 2220 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Std 3 Cargo, 3330 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Std 3 Cargo, 4440 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Std 3 Cargo, 5550 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Std 3 Cargo, 6660 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Table B.4: Standard 4 Cargo Results, All Injection Source Count Rates

Std 4 Cargo, 277.5 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	0	1	0	0	0	0
WGPu Alarm	0	0	0	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	1	0
NORM Alarm	0	0	0	0	0	0
GC Alarm	0	1	0	1	1	1

Std 4 Cargo, 555 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	1	0	0	0
WGPu Alarm	0	0	0	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	0
GC Alarm	1	1	1	1	1	1

Std 4 Cargo, 1110 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Std 4 Cargo, 2220 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Std 4 Cargo, 3330 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Std 4 Cargo, 4440 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Std 4 Cargo, 5550 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Std 4 Cargo, 6660 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Table B.5: Standard 5 Cargo Results, All Injection Source Count Rates

Std 5 Cargo, 277.5 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	1	0
NORM Alarm	0	0	0	0	0	0
GC Alarm	0	1	1	1	1	1

Std 5 Cargo, 555 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	0
GC Alarm	1	1	1	1	1	1

Std 5 Cargo, 1110 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Std 5 Cargo, 2220 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Std 5 Cargo, 3330 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Std 5 Cargo, 4440 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Std 5 Cargo, 5550 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Std 5 Cargo, 6660 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Table B.6: ⁴⁰K High Cargo Results, All Injection Source Count Rates

K-40 High Cargo, 277.5 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	0	0	0	0	0	0
WGPu Alarm	0	0	0	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	1	1	1	1	1	1
GC Alarm	1	1	1	1	1	1

K-40 High Cargo, 555 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	0	0	0	0	0	0
WGPu Alarm	0	0	0	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	1	1	1	1	1	1
GC Alarm	1	1	1	1	1	1

K-40 High Cargo, 1110 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	0	0	0	0	0	0
WGPu Alarm	0	0	0	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	1	1	1	1	1	1
GC Alarm	1	1	1	1	1	1

K-40 High Cargo, 2220 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	0	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	1	1	1	1	1	1
GC Alarm	1	1	1	1	1	1

K-40 High Cargo, 3330 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	1	1	1	1	1	1
GC Alarm	1	1	1	1	1	1

K-40 High Cargo, 4440 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	1	1	1	1	1	1
GC Alarm	1	1	1	1	1	1

K-40 High Cargo, 5550 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	1	1	1	1	1	1
GC Alarm	1	1	1	1	1	1

K-40 High Cargo, 6660 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	1	1	1	1	1	1
GC Alarm	1	1	1	1	1	1

Table B.7: ⁴⁰K Low Cargo Results, All Injection Source Count Rates

K-40 Low Cargo, 277.5 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	0	0	0	0	0	0
WGPu Alarm	0	0	0	0	0	0
Cs-137 Alarm	0	0	0	0	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	1	1	1	1	1	1
GC Alarm	1	1	1	1	1	1

K-40 Low Cargo, 555 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	0	0	0	0	0	0
WGPu Alarm	0	0	0	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	1	1	1	1	1	1
GC Alarm	1	1	1	1	1	1

K-40 Low Cargo, 1110 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	0	1	0	0	0	0
WGPu Alarm	0	0	0	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	1	1	1	1	1	1
GC Alarm	1	1	1	1	1	1

K-40 Low Cargo, 2220 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	1	1	1	1	1	1
GC Alarm	1	1	1	1	1	1

K-40 Low Cargo, 3330 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	1	1	1	1	1	1
GC Alarm	1	1	1	1	1	1

K-40 Low Cargo, 4440 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	1	1	1	1	1	1
GC Alarm	1	1	1	1	1	1

K-40 Low Cargo, 5550 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	1	1	1	1	1	1
GC Alarm	1	1	1	1	1	1

K-40 Low Cargo, 6660 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	0	0	0	0
WGPu Alarm	0	0	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	1	1	1	1	1	1
GC Alarm	1	1	1	1	1	1

Table B.8: Granite 1 Cargo Results, All Injection Source Count Rates

Granite 1 Cargo, 277.5 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	0	0	0	0	0	0
WGPu Alarm	1	1	1	0	0	1
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	1	0
NORM Alarm	0	0	0	0	0	0
GC Alarm	1	1	1	1	1	1

Granite 1 Cargo, 555 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	0	0	0	0	0	0
WGPu Alarm	1	1	1	0	0	1
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	0
GC Alarm	1	1	1	1	1	1

Granite 1 Cargo, 1110 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	0	0	0	0	0	0
WGPu Alarm	1	1	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Granite 1 Cargo, 2220 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	0	0	0	0	0	0
WGPu Alarm	1	1	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Granite 1 Cargo, 3330 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	0	0	0	0	0	0
WGPu Alarm	1	1	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Granite 1 Cargo, 4440 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	0	0	0	0	0	0
WGPu Alarm	1	1	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Granite 1 Cargo, 5550 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	0	0	0	0	0	0
WGPu Alarm	1	1	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Granite 1 Cargo, 6660 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	0	0	0	0	0	0
WGPu Alarm	1	1	1	0	0	0
Cs-137 Alarm	0	0	0	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Table B.9: Granite 2 Cargo Results, All Injection Source Count Rates

Granite 2 Cargo, 277.5 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	1	0	0	0
WGPu Alarm	0	0	0	0	0	0
Cs-137 Alarm	1	1	1	1	0	1
DU Alarm	0	0	0	0	1	0
NORM Alarm	0	0	0	0	0	0
GC Alarm	1	1	1	1	1	1

Granite 2 Cargo, 555 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	1	0	0	0
WGPu Alarm	0	0	0	0	0	0
Cs-137 Alarm	1	1	1	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	0
GC Alarm	1	1	1	1	1	1

Granite 2 Cargo, 1110 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	1	0	0	0
WGPu Alarm	0	0	0	0	0	0
Cs-137 Alarm	1	1	1	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Granite 2 Cargo, 2220 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	1	0	0	0
WGPu Alarm	0	0	0	0	0	0
Cs-137 Alarm	1	1	1	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Granite 2 Cargo, 3330 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	1	0	0	0
WGPu Alarm	0	0	0	0	0	0
Cs-137 Alarm	1	1	1	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Granite 2 Cargo, 4440 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	1	0	0	0
WGPu Alarm	0	0	0	0	0	0
Cs-137 Alarm	1	1	1	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Granite 2 Cargo, 5550 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	1	0	0	0
WGPu Alarm	0	0	0	0	0	0
Cs-137 Alarm	1	1	1	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Granite 2 Cargo, 6660 cps Injection Source						
	Co-57	HEU	Ba-133	Cs-137	Co-60	K-40
HEU Alarm	1	1	1	0	0	0
WGPu Alarm	0	0	0	0	0	0
Cs-137 Alarm	1	1	1	1	0	0
DU Alarm	0	0	0	0	0	0
NORM Alarm	0	0	0	0	1	1
GC Alarm	1	1	1	1	1	1

Appendix C – Advanced Alarm Algorithm Statistics Tables

Table C.10: Alarm Statistics for Standard Cargos, 277.5 cps Injection Sources

Statistics, Standard Cargos, 277.5 cps Injection Source					
	HEU	WGPu	Cs-137	DU	NORM
Sensitivity	0.5	0.2	1	1	0
False Alarm Rate	0.05	0	0	0	0
False Negative Rate	0.5	0.8	0	0	1
Specificity	0.95	1	1	1	1
Precision	0.8333333333	1	1	1	#DIV/0!
Accuracy	0.8	0.8666666667	1	1	0.8333333333
MCC	0.530330086	0.415227399	1	1	#DIV/0!
F1 Score	0.625	0.3333333333	1	1	0

Table C.11: Alarm Statistics for Standard Cargos, 555 cps Injection Sources

Statistics, Standard Cargos, 555 cps Injection Source					
	HEU	WGPu	Cs-137	DU	NORM
Sensitivity	0.9	0.8	1	1	0.4444444444
False Alarm Rate	0.05	0	0	0	0
False Negative Rate	0.1	0.2	0	0	0.5555555556
Specificity	0.95	1	1	1	1
Precision	0.9	1	1	1	1
Accuracy	0.9333333333	0.9666666667	1	1	0.8333333333
MCC	0.85	0.877058019	1	1	0.59914469
F1 Score	0.9	0.888888889	1	1	0.615384615

Table C.12: Alarm Statistics for Standard Cargos, 1110 cps Injection Sources

Statistics, Standard Cargos, 1110 cps Injection Source					
	HEU	WGPu	Cs-137	DU	NORM
Sensitivity	1	1	1	#DIV/0!	0.9
False Alarm Rate	0	0	0	0	0
False Negative Rate	0	0	0	#DIV/0!	0.1
Specificity	1	1	1	1	1
Precision	1	1	1	#DIV/0!	1
Accuracy	1	1	1	1	0.9666666667
MCC	1	1	1	#DIV/0!	0.9258201
F1 Score	1	1	1	#DIV/0!	0.947368421

Table C.13: Alarm Statistics for Standard Cargos, 2220 cps Injection Sources

Statistics, Standard Cargos, 2220 cps Injection Source					
	HEU	WGPu	Cs-137	DU	NORM
Sensitivity	1	1	1	#DIV/0!	1
False Alarm Rate	0	0	0	0	0
False Negative Rate	0	0	0	#DIV/0!	0
Specificity	1	1	1	1	1
Precision	1	1	1	#DIV/0!	1
Accuracy	1	1	1	1	1
MCC	1	1	1	#DIV/0!	1
F1 Score	1	1	1	#DIV/0!	1

Table C. 14: Alarm Statistics for Standard Cargos, 3330 cps Injection Sources

Statistics, Standard Cargos, 3330 cps Injection Source					
	HEU	WGPu	Cs-137	DU	NORM
Sensitivity	1	1	1	#DIV/0!	1
False Alarm Rate	0	0	0	0	0
False Negative Rate	0	0	0	#DIV/0!	0
Specificity	1	1	1	1	1
Precision	1	1	1	#DIV/0!	1
Accuracy	1	1	1	1	1
MCC	1	1	1	#DIV/0!	1
F1 Score	1	1	1	#DIV/0!	1

Table C. 15: Alarm Statistics for Standard Cargos, 4440 cps Injection Sources

Statistics, Standard Cargos, 4440 cps Injection Source					
	HEU	WGPu	Cs-137	DU	NORM
Sensitivity	1	1	1	#DIV/0!	1
False Alarm Rate	0	0	0	0	0
False Negative Rate	0	0	0	#DIV/0!	0
Specificity	1	1	1	1	1
Precision	1	1	1	#DIV/0!	1
Accuracy	1	1	1	1	1
MCC	1	1	1	#DIV/0!	1
F1 Score	1	1	1	#DIV/0!	1

Table C. 16: Alarm Statistics for Standard Cargos, 5550 cps Injection Sources

Statistics, Standard Cargos, 5550 cps Injection Source					
	HEU	WGPu	Cs-137	DU	NORM
Sensitivity	1	1	1	#DIV/0!	1
False Alarm Rate	0	0	0	0	0
False Negative Rate	0	0	0	#DIV/0!	0
Specificity	1	1	1	1	1
Precision	1	1	1	#DIV/0!	1
Accuracy	1	1	1	1	1
MCC	1	1	1	#DIV/0!	1
F1 Score	1	1	1	#DIV/0!	1

Table C. 17: Alarm Statistics for Standard Cargos, 6660 cps Injection Sources

Statistics, Standard Cargos, 6660 cps Injection Source					
	HEU	WGPu	Cs-137	DU	NORM
Sensitivity	1	1	1	#DIV/0!	1
False Alarm Rate	0	0	0	0	0
False Negative Rate	0	0	0	#DIV/0!	0
Specificity	1	1	1	1	1
Precision	1	1	1	#DIV/0!	1
Accuracy	1	1	1	1	1
MCC	1	1	1	#DIV/0!	1
F1 Score	1	1	1	#DIV/0!	1

Table C. 18: Alarm Statistics for ⁴⁰K Cargos, 277.5 cps Injection Sources

Statistics, K-40 Cargos, 277.5 cps Injection Source					
	HEU	WGPu	Cs-137	DU	NORM
Sensitivity	0	0	0.5	#DIV/0!	1
False Alarm Rate	0	0	0	0	#DIV/0!
False Negative Rate	1	1	0.5	#DIV/0!	0
Specificity	1	1	1	1	#DIV/0!
Precision	#DIV/0!	#DIV/0!	1	#DIV/0!	1
Accuracy	0.666666667	0.833333333	0.916666667	1	1
MCC	#DIV/0!	#DIV/0!	0.674199862	#DIV/0!	#DIV/0!
F1 Score	0	0	0.666666667	#DIV/0!	1

Table C. 19: Alarm Statistics for ⁴⁰K Cargos, 555 cps Injection Sources

Statistics, K-40 Cargos, 555 cps Injection Source					
	HEU	WGPu	Cs-137	DU	NORM
Sensitivity	0	0	1	#DIV/0!	1
False Alarm Rate	0	0	0	0	#DIV/0!
False Negative Rate	1	1	0	#DIV/0!	0
Specificity	1	1	1	1	#DIV/0!
Precision	#DIV/0!	#DIV/0!	1	#DIV/0!	1
Accuracy	0.666666667	0.833333333	1	1	1
MCC	#DIV/0!	#DIV/0!	1	#DIV/0!	#DIV/0!
F1 Score	0	0	1	#DIV/0!	1

Table C. 20: Alarm Statistics for ⁴⁰K Cargos, 1110 cps Injection Sources

Statistics, K-40 Cargos, 1110 cps Injection Source					
	HEU	WGPu	Cs-137	DU	NORM
Sensitivity	0.25	0	1	#DIV/0!	1
False Alarm Rate	0	0	0	0	#DIV/0!
False Negative Rate	0.75	1	0	#DIV/0!	0
Specificity	1	1	1	1	#DIV/0!
Precision	1	#DIV/0!	1	#DIV/0!	1
Accuracy	0.75	0.833333333	1	1	1
MCC	0.426401433	#DIV/0!	1	#DIV/0!	#DIV/0!
F1 Score	0.4	0	1	#DIV/0!	1

Table C. 21: Alarm Statistics for ⁴⁰K Cargos, 2220 cps Injection Sources

Statistics, K-40 Cargos, 2220 cps Injection Source					
	HEU	WGPu	Cs-137	DU	NORM
Sensitivity	0.75	1	1	#DIV/0!	1
False Alarm Rate	0	0	0	0	#DIV/0!
False Negative Rate	0.25	0	0	#DIV/0!	0
Specificity	1	1	1	1	#DIV/0!
Precision	1	1	1	#DIV/0!	1
Accuracy	0.916666667	1	1	1	1
MCC	0.816496581	1	1	#DIV/0!	#DIV/0!
F1 Score	0.857142857	1	1	#DIV/0!	1

Table C. 22: Alarm Statistics for ⁴⁰K Cargos, 3330 cps Injection Sources

Statistics, K-40 Cargos, 3330 cps Injection Source					
	HEU	WGPu	Cs-137	DU	NORM
Sensitivity	1	1	1	#DIV/0!	1
False Alarm Rate	0	0	0	0	#DIV/0!
False Negative Rate	0	0	0	#DIV/0!	0
Specificity	1	1	1	1	#DIV/0!
Precision	1	1	1	#DIV/0!	1
Accuracy	1	1	1	1	1
MCC	1	1	1	#DIV/0!	#DIV/0!
F1 Score	1	1	1	#DIV/0!	1

Table C. 23: Alarm Statistics for ⁴⁰K Cargos, 4440 cps Injection Sources

Statistics, K-40 Cargos, 4440 cps Injection Source					
	HEU	WGPu	Cs-137	DU	NORM
Sensitivity	1	1	1	#DIV/0!	1
False Alarm Rate	0	0	0	0	#DIV/0!
False Negative Rate	0	0	0	#DIV/0!	0
Specificity	1	1	1	1	#DIV/0!
Precision	1	1	1	#DIV/0!	1
Accuracy	1	1	1	1	1
MCC	1	1	1	#DIV/0!	#DIV/0!
F1 Score	1	1	1	#DIV/0!	1

Table C. 24: Alarm Statistics for ⁴⁰K Cargos, 5550 cps Injection Sources

Statistics, K-40 Cargos, 5550 cps Injection Source					
	HEU	WGPu	Cs-137	DU	NORM
Sensitivity	1	1	1	#DIV/0!	1
False Alarm Rate	0	0	0	0	#DIV/0!
False Negative Rate	0	0	0	#DIV/0!	0
Specificity	1	1	1	1	#DIV/0!
Precision	1	1	1	#DIV/0!	1
Accuracy	1	1	1	1	1
MCC	1	1	1	#DIV/0!	#DIV/0!
F1 Score	1	1	1	#DIV/0!	1

Table C. 25: Alarm Statistics for ⁴⁰K Cargos, 6660 cps Injection Sources

Statistics, K-40 Cargos, 6660 cps Injection Source					
	HEU	WGPu	Cs-137	DU	NORM
Sensitivity	1	1	1	#DIV/0!	1
False Alarm Rate	0	0	0	0	#DIV/0!
False Negative Rate	0	0	0	#DIV/0!	0
Specificity	1	1	1	1	#DIV/0!
Precision	1	1	1	#DIV/0!	1
Accuracy	1	1	1	1	1
MCC	1	1	1	#DIV/0!	#DIV/0!
F1 Score	1	1	1	#DIV/0!	1

Table C. 26: Alarm Statistics for Granite Cargos, 277.5 cps Injection Sources

Statistics, Granite Cargos, 277.5 cps Injection Source					
	HEU	WGPu	Cs-137	DU	NORM
Sensitivity	0.5	0.571428571	0.857142857	1	0
False Alarm Rate	0.125	0	0	0	0
False Negative Rate	0.5	0.428571429	0.142857143	0	1
Specificity	0.875	1	1	1	1
Precision	0.666666667	1	1	1	#DIV/0!
Accuracy	0.75	0.75	0.916666667	1	0.833333333
MCC	0.40824829	0.597614305	0.845154255	1	#DIV/0!
F1 Score	0.571428571	0.727272727	0.923076923	1	0

Table C. 27: Alarm Statistics for Granite Cargos, 555 cps Injection Sources

Statistics, Granite Cargos, 555 cps Injection Source					
	HEU	WGPu	Cs-137	DU	NORM
Sensitivity	0.5	0.571428571	0.714285714	#DIV/0!	0.5
False Alarm Rate	0.125	0	0	0	0
False Negative Rate	0.5	0.428571429	0.285714286	#DIV/0!	0.5
Specificity	0.875	1	1	1	1
Precision	0.666666667	1	1	#DIV/0!	1
Accuracy	0.75	0.75	0.833333333	1	0.833333333
MCC	0.40824829	0.597614305	0.714285714	#DIV/0!	0.632455532
F1 Score	0.571428571	0.727272727	0.833333333	#DIV/0!	0.666666667

Table C. 28: Alarm Statistics for Granite Cargos, 1110 cps Injection Sources

Statistics, Granite Cargos, 1110 cps Injection Source					
	HEU	WGPu	Cs-137	DU	NORM
Sensitivity	0.5	0.428571429	0.714285714	#DIV/0!	1
False Alarm Rate	0.125	0	0	0	0
False Negative Rate	0.5	0.571428571	0.285714286	#DIV/0!	0
Specificity	0.875	1	1	1	1
Precision	0.666666667	1	1	#DIV/0!	1
Accuracy	0.75	0.666666667	0.833333333	1	1
MCC	0.40824829	0.487950036	0.714285714	#DIV/0!	1
F1 Score	0.571428571	0.6	0.833333333	#DIV/0!	1

Table C. 29: Alarm Statistics for Granite Cargos, 2220 cps Injection Sources

Statistics, Granite Cargos, 2220 cps Injection Source					
	HEU	WGPu	Cs-137	DU	NORM
Sensitivity	0.5	0.428571429	0.714285714	#DIV/0!	1
False Alarm Rate	0.125	0	0	0	0
False Negative Rate	0.5	0.571428571	0.285714286	#DIV/0!	0
Specificity	0.875	1	1	1	1
Precision	0.666666667	1	1	#DIV/0!	1
Accuracy	0.75	0.666666667	0.833333333	1	1
MCC	0.40824829	0.487950036	0.714285714	#DIV/0!	1
F1 Score	0.571428571	0.6	0.833333333	#DIV/0!	1

Table C. 30: Alarm Statistics for Granite Cargos, 3330 cps Injection Sources

Statistics, Granite Cargos, 3330 cps Injection Source					
	HEU	WGPu	Cs-137	DU	NORM
Sensitivity	0.5	0.4286	0.7142857	#DIV/0!	1
False Alarm Rate	0.125	0	0	0	0
False Negative Rate	0.5	0.5714	0.2857143	#DIV/0!	0
Specificity	0.875	1	1	1	1
Precision	0.66667	1	1	#DIV/0!	1
Accuracy	0.75	0.6667	0.8333333	1	1
MCC	0.40825	0.488	0.7142857	#DIV/0!	1
F1 Score	0.57143	0.6	0.8333333	#DIV/0!	1

Table C. 31: Alarm Statistics for Granite Cargos, 4440 cps Injection Sources

Statistics, Granite Cargos, 4440 cps Injection Source					
	HEU	WGpu	Cs-137	DU	NORM
Sensitivity	0.5	0.4286	0.7142857	#DIV/0!	1
False Alarm Rate	0.125	0	0	0	0
False Negative Rate	0.5	0.5714	0.2857143	#DIV/0!	0
Specificity	0.875	1	1	1	1
Precision	0.66667	1	1	#DIV/0!	1
Accuracy	0.75	0.6667	0.8333333	1	1
MCC	0.40825	0.488	0.7142857	#DIV/0!	1
F1 Score	0.57143	0.6	0.8333333	#DIV/0!	1

Table C. 32: Alarm Statistics for Granite Cargos, 5550 cps Injection Sources

Statistics, Granite Cargos, 5550 cps Injection Source					
	HEU	WGpu	Cs-137	DU	NORM
Sensitivity	0.5	0.4286	0.7142857	#DIV/0!	1
False Alarm Rate	0.125	0	0	0	0
False Negative Rate	0.5	0.5714	0.2857143	#DIV/0!	0
Specificity	0.875	1	1	1	1
Precision	0.66667	1	1	#DIV/0!	1
Accuracy	0.75	0.6667	0.8333333	1	1
MCC	0.40825	0.488	0.7142857	#DIV/0!	1
F1 Score	0.57143	0.6	0.8333333	#DIV/0!	1

Table C. 33: Alarm Statistics for Granite Cargos, 6660 cps Injection Sources

Statistics, Granite Cargos, 6660 cps Injection Source					
	HEU	WGpu	Cs-137	DU	NORM
Sensitivity	0.5	0.4286	0.7142857	#DIV/0!	1
False Alarm Rate	0.125	0	0	0	0
False Negative Rate	0.5	0.5714	0.2857143	#DIV/0!	0
Specificity	0.875	1	1	1	1
Precision	0.66667	1	1	#DIV/0!	1
Accuracy	0.75	0.6667	0.8333333	1	1
MCC	0.40825	0.488	0.7142857	#DIV/0!	1
F1 Score	0.57143	0.6	0.8333333	#DIV/0!	1