

University of Missouri – Kansas City

Department of Physics

Kansas City, MO 64110-2499

UMKC-YIP-TR-2016 May 2016

Technical Report

Prompt Neutron Spectrometry for Identification of SNM in Unknown Shielding Configurations: FY16 ONR-YIP Final Report

by

C.B. Hoshor, E.R. Myers, S.M. Young, J.E. Currie, and A.N. Caruso*

University of Missouri – Kansas City, Kansas City, MO

*Please direct all technical correspondence to Anthony N. Caruso, carusoan@umkc.edu, 816-235-2505

Distribution Statement A: approved for public release.

Distribution

ADDRESSEE	DODAAD CODE	QUANTITY
Office of Naval Research ATTN: Ryan Hoffman ONR, 351 875 North Randolph St. Arlington, VA 22203-1995	N00014	1
Administrative Contracting Officer* *Transmittal letter only Fax: 312-353-6089	N62880	1
DEFENSE TECHNICAL INFORMATION CENTER 8725 JOHN J. KINGMAN RD SUITE 0944 FORT BELVOIR, VA 22060-6218 E-mail: TR@dtic.mil	HJ4701	1
DIRECTOR, NAVAL RESEARCH LAB ATTN: CODE 5596 4555 OVERLOOK AVENUE, SW WASHINGTON, D.C. 20375-5320 E-mail: reports@library.nrl.navy.mil	N00173	1

Distribution Statement A

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 31-05-2016		2. REPORT TYPE Final		3. DATES COVERED (From - To) MAR 2015 - MAR 2016	
4. TITLE AND SUBTITLE Prompt Neutron Spectrometry for Identification of SNM in Unknown Shielding Configurations: FY16 ONR-YIP Final Report				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER N00014-13-1-0402	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) C.B. Hoshor, E.R. Myers, S.M. Young, J.E. Currie and A.N. Caruso				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Missouri - Kansas City 5100 Rockhill Road Kansas City, MO 64110				8. PERFORMING ORGANIZATION REPORT NUMBER UMKC-YIP-TR-2016	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Office of Naval Research One Liberty Center, 875 N. Randolph Arlington, VA 22203				10. SPONSOR/MONITOR'S ACRONYM(S) ONR	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Distribution Statement A: approved for public release.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT This report represents the final documentation for the ONR FY12 Young Investigator Program (YIP) award, N00014-13-1-0402, titled "Prompt Neutron Spectrometry for Identification of SNM in Unknown Shielding Configurations". The technical portion of this report is not comprehensive, although the executive summary makes a high level attempt. The YIP effort reported here complements work completed under the mainline ONR Maritime-Weapons-of-Mass-Destruction-Detection (MWMDD) program (FY09 – FY14). A comprehensive technical understanding of all the related and prior work for this effort may be found from DTIC report "Solid State Neutron Spectrometer: FY14 ONR-MWMDD Final Report".					
15. SUBJECT TERMS neutron, gamma-ray, spectrometer, detector, WMD, VBSS, HRM, radiation, resolution, solid-state, active interrogation					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Anthony Caruso
U	U	U	UU	25	19b. TELEPHONE NUMBER (Include area code) 816-235-2505

Standard Form 298 (Rev. 8/98)
Prescribed by ANSI Std. Z39.18

Table of Contents

I. EXECUTIVE SUMMARY	7
II. GENERAL APPROACH TO DETECTOR AND MODERATOR GEOMETRIES	9
II.A. MODERATING SPECTROMETER OPTIMIZATION	9
II.B. RESPONSE UNIQUENESS AND SOLUTION STABILITY	9
II.C. PREVIOUS SPECTROMETER OPTIMIZATION PROCEDURES	11
II.D. VIRTUAL DETECTORS FOR REDUCED SIMULATION	12
II.E. VIRTUAL DETECTOR INPUT CARDS	12
II.F. RESULTS FROM VIRTUAL DETECTOR SIMULATIONS	13
II. G. GENETIC ALGORITHMS FOR RESPONSE FUNCTION SELECTION	16
II. H. DEFINING A FITNESS METRIC	16
II.I. GENETIC ALGORITHM EVOLUTION PARAMETERS	18
<i>The Pareto Frontier</i>	19
II.J. CLOSING REMARKS ON OPTIMIZATION	23
III. OUTREACH PROGRAM (SEAPERCH)	24
IV. REFERENCES	25

Acronyms

CBRN: chemical, biological, radiological, nuclear

DR SKO: dismounted reconnaissance set kit or outfit

VBSS: Visit, Boarding, Search and Seizure

DHS-DNDO: Department of Homeland Security – Domestic Nuclear Detection Office

DTRA-ONC: Defense Threat Reduction Agency – Contingency Operations Division

EOD: explosive ordinance disposal

DTIC: defense technical information center

ONR: Office of Naval research

UMKC: University of Missouri – Kansas City

MSND: Micro-structured Neutron Detector

HRM: Handheld Radiation Monitor

PHS: Pulse Height Spectrum

ANI: Active Neutron Interrogation

Administrative Information and Acknowledgements

Members of the University of Missouri – Kansas City Department of Physics and Naval Research Laboratory, Code 6770 were responsible for the concepts and/or work in this report. Code 6770 innovators include Stuart Jackson, Joseph Schumer, Robert Commisso, Bruce Weber, and Jacob Zier. This work builds on and is complementary to efforts funded by the Office of Naval Research under contract N00014-14-P-1115 “HRM Replacement”, grant N00014-11-01-0157 “High Efficiency Solid State Neutron Detection”, and (ONR-DURIP) N00014-13-1-0757 “Neutron Generator System for Warm and Fast Energy Neutron Interrogation”. The detector systems on which the neutron spectrometer concepts are built are derived from Kansas State University’s SMART Laboratory, headed by Douglas McGregor, Kansas State University’s Electronics Design Lab, headed by Tim Sobering, and Radiation Detection Technologies, Inc., headed by Steven Bellinger.

I. Executive Summary

This report represents the final documentation for the ONR FY12 Young Investigator Program (YIP) award, N00014-13-1-0402, titled “Prompt Neutron Spectrometry for Identification of SNM in Unknown Shielding Configurations”. The technical portion of this report is not comprehensive, although the executive summary makes a high level attempt. The YIP effort reported here complements work completed under the mainline ONR Maritime-Weapons-of-Mass-Destruction-Detection (MWMD) program (FY09 – FY14). A comprehensive technical understanding of all the related and prior work for this effort may be found from the Defense Technical Information Center report “Solid State Neutron Spectrometer: FY14 ONR-MWMD Final Report” and references therein [1].

Mitigating illicit transport of special nuclear material (SNM) in the maritime domain through advanced measurement methods was the original intent of the YIP work and it presented a unique challenge because of the variability in neutron scattering and shielding materials, as well as background neutron sources present in the maritime environment [2]. Gamma-ray spectroscopy has revolutionized radiation detection and has come to be regarded as a necessary passive Fleet tool for standoff and onboard SNM searches. To complement gamma spectroscopy, as well as to provide new functionality through the spectroscopy of all significant sources of free neutrons in the naval environment, high-resolution and high-intrinsic-efficiency neutron spectrometers for passive and active-interrogation-based neutron spectrometry applications were designed; the most promising designs were built and tested. The spectrometer was intended to yield sensitivity over the thermal to 14 MeV energy range, with special emphasis on resolving the properties of actively induced neutrons. Active interrogation brings its own set of challenges, including delay or recovery time after an interrogating bremsstrahlung pulse, as well as the short collection time (hundreds of nanoseconds to tens of milliseconds).

The main technical contributions which arose from this work did not end up taking active interrogation to a grand level, but rather, broke out horizontally, to show how instrument capabilities could be optimized for both passive and active applications. This concept was first explored, as detailed in prior reports, through heavily-constrained brute-force procedures aimed at simultaneously optimizing two instrument properties of great importance to both passive and active operational scenarios, intrinsic neutron detection efficiency and quantities proportional to energy resolution. A primary technical goal of this work was to build upon this early (and largely naïve) approach to moderating-type neutron spectrometer optimization; this goal was to be achieved by first taking one or more conceptual steps backward, in order to gain the perspective necessary to make substantial forward strides. This process began with a fresh in-depth survey of general instrumentation and methods for passive and active neutron spectrometry—including both well-established, time-tested methodologies and novel, boundary-stretching concepts—to provide the research team with a more comprehensive awareness of the current state of the art and a fuller understanding of the primary underlying difficulties and unresolved problems associated with both active and passive fields of neutron spectrometry. With this theoretical foundation refreshed and firmly established, it was immediately apparent that the heavy design constraints imposed in previous optimization procedures must be relaxed considerably to thoroughly explore the parameter space for potential instrument configurations. However, it was also clear that, since MCNP computation time was determined to be the greatest limiting factor in previous optimization attempts, design constraint relaxation would necessarily render brute-force simulation-heavy optimization approaches infeasible, necessitating the development of intelligent

alternative procedures to minimize the number of simulations required to determine an optimal instrument design. To this end, the virtual detector optimization method was conceived, which provides a much-needed simplification to parameter study (mcnp_pstudy) simulation set-up time and decreases the number of simulations necessary to explore the detector-position parameter space by many orders of magnitude, representing the first major technical advancement of this work. Further methodological improvements to moderating-type neutron optimization procedures borrowed, repurposed, and combined concepts from several other mathematical and scientific disciplines, including linear algebra, computer science, computational statistics, Bayesian inference, information theory, and machine learning. The condition number (borrowed from linear algebraic concepts) of a proposed instrument's response matrix, for example, was recognized and repurposed as an essential operational feedback tool analogous to neutron energy resolution when considering spectrum unfolding procedures based on Bayesian inference, and serves as one design optimization metric for this work, representing the second technical advancement of this work. The third significant advancement achieved in this work came about through the comparison of instrument configurations that optimized the aforementioned condition number metric to those that optimized the intrinsic neutron detection efficiency metric (i.e., the second operational metric considered here, consistent with previous optimization efforts); plotting the resultant values of these two metrics against one another for the "most optimal" instrument configurations under consideration generates what is known in computational statistics as a Pareto frontier, which allows for a more detailed analysis of the potential tradeoffs between the two desired optimization metrics. Perhaps the most significant technical advancement achieved in this work was in adapting the machine learning concept of genetic algorithms as a generalized method for moderating-type neutron spectrometer optimization, utilizing concepts from each of the aforementioned mathematical and scientific disciplines to establish a unified approach to the problem.

Standing on the shoulders of the ONR subsidized SeaPerch program, outreach activities for elementary aged schoolchildren, in the form of underwater robotics, impacted more than 80 students over a three year period under this award. In all 40 SeaPerch kits were purchased through this award (many of the tools/parts were reused each year). The 4th-6th graders cut, drilled, waterproofed, assembled, tested and demonstrated their tethered robots. Demonstrations/competitions occurred at the high school pool where the students were challenged to underwater races and task manipulation. Lectures on voltage, buoyancy and other topics were provided, but the hands on work and engagement with parents is what ultimately made the difference; unfortunately, the statistics and convolution of child-rearing demographics/activities will not allow for a definite answer on whether more of these kids will go into a STEM related discipline. In the PI's experience, the SeaPerch program makes outreach manageable and effective; expansion and requirement of outreach on more 6.1 and 6.2 efforts would be a very good and small investment of funds.

At the university level, broader impact included the partial support, graduation, matriculation and/or training of 11 students, 2 postdocs, and 2 technicians. Two MS theses were produced, five non-provisional patents, two-licensed patents, two journal publications, follow on work with a DTRA SBIR (now in Phase II), and an R&D100 award. Manuscripts for journal publications are still being prepared and will be published over the next year.

In total, the YIP provided the base to learn/understand/traverse/catch-a-better-glimpse of the intellectual/political landscape of the Navy RDT&E enterprise, while also generating new intellectual capital and training of students of all ages.

II. General Approach to Detector and Moderator Geometries

II.A. Moderating Spectrometer Optimization

In order to design the most suitable moderating neutron spectrometer for a given application (e.g. SNM properties determination), the relationship between transport material configurations and detector positions requires optimization. However, optimization is not straightforward because we must rely on some property of the response matrix in order to determine how optimal a spectrometer design is. Furthermore, the properties of an ideal response matrix are not completely independent of the unfolding methods used to approximate the flux. This means that certain properties of the response matrix, such as determinedness, condition number, number of energy bins, etc. would ideally be optimized for the unfolding method. Despite this high specificity of spectrometer optimization, the following sections cover some general response matrix properties that are beneficial to unfolding.

II.B. Response Uniqueness and Solution Stability

Arguably the most important feature of multi-detector moderating neutron spectrometers is their ability to measure unique neutron responses from which the incident neutrons' energies may be approximated. Calculation of the neutron energy is very specific to the spectrometer transport material design and detector placement—that is, every moderating spectrometer has a unique set of response functions corresponding to its detector(s) and transport material design. These response functions must be determined before any neutron energy analysis is possible. Although the response functions would ideally be measured experimentally, tunable monoenergetic neutron sources do not exist over a large enough energy range to enable this. Instead, these response functions are calculated via MCNP by simulating mono-energetic neutron sources—covering the relevant energy range (generally 10^{-9} to 10^2 MeV)—incident on the spectrometer model. This set of mono-energetic detector responses makes a single neutron response matrix unique to the spectrometer model, having dimensions of energy and detector position (both discrete as the energy must be binned according to the desired resolution). This response matrix allows one to take a measurement of an unknown source and then approximate the incident flux, $\phi(E)$, by solving the following equation:

$$N(d) = R(d, E)\phi(E) \quad (1)$$

$N(d)$ is the measured counts on each detector and $R(d, E)$ is the response matrix for a set of detectors, d , in E energy bins. However, since the response functions are dependent on the messy process of neutron scattering and absorption within the transport material, these response functions are not, for lack of a better word, nice to work with. The response matrix, regardless of the material configurations, is virtually always non-invertible, and thus numerous solutions for the incident flux exist. Methods for solving such linear equations generally involve either maximizing the solution entropy and/or approximations of the response matrix inverse—a process referred to as *unfolding* the incident neutron spectrum. More detailed information on spectrum unfolding can be found in the work of Matzke [36].

A property of an ideal response matrix is its determinedness—the matrix is either underdetermined, determined, or overdetermined. It is determined when its number of energy bins is equal to the number of unique response functions (i.e. number of detectors), underdetermined when its number of detectors is less than the number of energy bins, and overdetermined when the number of detectors exceeds the number of energy bins. Although this property appears straightforward on the surface, it arises from a seemingly arbitrary decision made by the scientist, as both the number of detectors and the number of energy bins are largely up to the spectrometer designer. Either energy bins or number of detectors can be increased or decreased, somewhat arbitrarily. Of course, the arbitrary nature of these decisions has limits, because decreasing the number of energy bins below the desired energy resolution is counter-productive. For example, a response matrix with two energy bins can only be used to approximate the flux into those two

energy bins—one cannot expect to efficiently resolve multiple neutron energies with a single detector. Thus, the selection of an appropriate number of energy bins and detectors is highly application-specific, and there may be benefits to choosing fewer energy bins and/or detectors while still meeting the desired energy resolution. However, detectors and energy bins do have economic and/or computational costs, thus their numbers must be optimized, rather than arbitrarily maximized. On an intuitive level, determinedness can be understood by thinking about the response functions as sets of linear equations. When solving sets of equations, it's always beneficial to have an equal number of unknowns as you have equations, which—in effect—corresponds to response matrices having the same number of energy bins and detectors.

In reality, this analogy between response functions and sets of linear equations is over-simplified, and determinedness is by no means the most important feature of the response matrices. The nature of neutron scattering, which is inherently messy, results in sets of equations that virtually always contain multiple solutions. Thus, approximated solutions for the incident neutron flux using moderating spectrometer response matrices are rarely stable. That is to say: while different incident neutron fluences may provide the same or similar detector responses, a small change in the measured spectrometer response may result in a large change in calculated neutron spectra after unfolding. The reason for this instability lies within the broad and overlapping features of the response functions. These features can be seen in Figure 26, which shows a selection of 30 simulated response functions from detectors evenly spaced along the radius of a 20-cm sphere of HDPE.




Figure 26: Response functions selected from 30 evenly spaced detectors within a 20-cm spherical moderating neutron spectrometer.

Numerous response functions in Figure 26 are similar and/or their peak energy response exists over a broad range in energy. For example, detectors near the outer surface of the sphere, in the 15–19 cm range, all have a broad peak response in the 10^{-7} to 1 MeV range. These broad overlaps in detector response, in effect, allow for many solutions for the incident neutron fluence to result in the same, or very similar, measured response. For the purpose of this thesis, the broad and overlapping features of response functions will be referred to as a lack of uniqueness. Unique response functions would be represented by narrow peak responses at varying energies spanning the energy range of importance.

One property which, to an extent characterizes this lack of response uniqueness is the response matrix condition number. The condition of the matrix provides a numerical value indicative of the number of possible solutions. An ideal condition number of one means that there is only one solution to Equation 1, regardless of the count distribution. Larger condition numbers correspond to less uniqueness between the responses and more possible solutions. Generally speaking, moderating neutron spectrometers are far from ideal, and it is not uncommon for their response matrices to have condition numbers on the order of 10^3 – 10^7 .

To overcome the limitations of large condition numbers, there are two options: 1) mathematical methods used to unfold the incident neutron fluence from the poorly conditioned responses, and/or 2) modify the neutron transport material and detector configurations to improve the condition number. For 1), as previously mentioned, numerous methods exist including ranging from maximizing solution entropy to approximating the response matrix inverse. For 2), this requires a method to simulate numerous transport material and detector configurations and to compare their response matrices. The rest of this chapter will focus on a new approach using virtual detector simulations and genetic algorithms to optimize the transport material configuration, and the efficiency and condition number of corresponding response matrices.

II.C. Previous Spectrometer Optimization Procedures

Previously, multi-detector moderating spectrometer optimization has been performed mostly by trial and error. A set of transport material configurations were selected, along with a set of detector positions within the material, and their response functions were calculated and compared.

For example, a volume of transport material with detectors placed throughout would be modelled in MCNP—such as a cylinder of HDPE with detectors evenly distributed along its length. The transport material (HDPE) would then be swapped out with other neutron moderating or absorbing materials, and the spacing of the detectors changed. Then two or more sources would be simulated incident on the spectrometer designs. Their ability to resolve incident neutron spectra could then be approximated by computing the Pierson cross correlation score between the two difference sources for each of the simulated devices. However, the metrics used in this approach only allowed for a comparison of how well two or more spectrometers could differentiate two *specific* sources, rather than resolving general spectra. Thus, the source spectra used had a large impact on the results of these simulations. Additionally, the trial and error approach required at least four simulations to compare two different spectrometer designs.

Even if it only required a single simulation for each spectrometer configuration, the number of simulations to explore just the detector position space would far surpass what is feasible. For a single transport material configuration, finding the optimal detector positions would require N simulations in Equation 2, where p is the number of possible detector positions within the transport material configuration, and n is the number of detectors placed within this configuration.

$$N = \frac{p!}{(p-n)!(n)!} \rightarrow \frac{20!}{(20-5)!(5)!} \cong 15500 \quad (2)$$

Even with 20 possible detector positions and 5 detectors to be placed, simulating all combinations would require over 15500 simulations. Although the ideal number of detectors is still unknown, a reasonable number of detectors and positions, such as 30 detectors and 80 positions, would require over 10^{21} simulations. Granted, many of these simulations could likely be excluded purely on intuition by restricting regions of the spectrometer where we expect either poor efficiency or uniqueness. Regardless, the resources required to perform simulations and process results from even a fraction of the possible configurations is impractical.

II.D. Virtual Detectors for Reduced Simulation

The number of simulations required to brute-force optimize a moderating neutron spectrometer is unreasonable. To explore the full parameter space, the number of simulations must be reduced by approximation, and even then processing the results requires optimization algorithms. This section outlines a new approach to approximate the response functions of a given transport material configuration, dramatically reducing the number of simulations required.

The method described herein reduces the number of MCNP simulations required to approximate the entire parameter space of the moderating neutron to a single simulation for each transport material configuration. Rather than placing all of the transport materials *and* detectors in the simulation, this approach only simulates the transport materials and then approximates the detectors response as if they had actually been placed in the simulation. To do so, the transport material configuration is divided into voxels that are geometrically equivalent to the desired detectors (microstructured semiconducting neutron detectors, MSNDs). A flat distribution of neutrons are simulated incident on the transport material and the average neutron fluence through each of the voxels, or virtual detectors. The product of the average neutron fluence and the efficiency-corrected cross section for the detector material, integrated over all energies simulated, results in an approximation of the counts that would be measured if a detector had actually been placed in the simulation. This allows for any number of detector configurations to be explored with a single simulation, therefore reducing the number of required simulations to just the number of transport material configurations.

The virtual detector approach makes one major assumption, that the effect of placing a detector within the moderator has a negligible effect on the neutron flux in other regions of the spectrometer. In extreme cases, this does not hold true when: 1) in any region of the spectrometer, the volume of detectors approaches that of the moderating volume, and 2) the detector area in one region completely shadow detector(s) in another region of the spectrometer. However, these exceptions to this approximation can be circumnavigated by bounding the possible detector positions to eliminate combinations of detector positions that would be non-physical. Additionally, for MSNDs, their high thermal efficiency and poor epithermal-fast efficiency helps to reduce the range of detector shadowing effects.

II.E. Virtual Detector Input Cards

A total of 16 transport material configurations were designed in MCNP. Each consisted of a spherical volume of HDPE divided into 400 0.5 mm thick spherical shells to approximate the thickness of an MSND (0.525 mm). Each shell represents a virtual detector. Optionally, the shells at radii of 5, 10, 15, and 20 cm were reserved for the optional placement of cadmium (as per each transport material configuration). **Table 2** shows the presence and location of these cadmium shells in each of the 16 simulated transport material configurations. These reserved positions for cadmium were designed to cover a wide range in possible locations in order to narrow the ideal locations for cadmium within the device.

Transport Configuration		01	02	03	04	05	06	07	08	09	10	11	12	13	14	15	16
Cadmium Positions (cm)	5.0		✓		✓		✓		✓		✓		✓		✓		✓
	10.0			✓	✓			✓	✓			✓	✓			✓	✓
	15.0					✓	✓	✓	✓					✓	✓	✓	✓
	20.0									✓	✓	✓	✓	✓	✓	✓	✓

Table 2. Locations of cadmium shells within the 20-cm sphere of HDPE totalling 16 transport material configurations models.

A point source of 10^{10} neutrons was simulated at a distance of 1 meter from the center of the transport material configuration. The neutrons were evenly distributed in energy ranging from 10^{-9} to 100 MeV divided into 100 logarithmically spaced energy bins, resulting in 5.625×10^5 neutrons incident on the transport material per energy bin. The average neutron fluence through each of the HDPE shells was calculated using an F4 particle-flux tally. The neutron fluence product with the efficiency-corrected cross-section of a single MSND was calculated to generate a set of 400 virtual detector response functions. These response functions were then plotted as a function of incident neutron energy, virtual response function (i.e., virtual detector position), and virtual response magnitude (i.e., virtual efficiency). Further analysis of the response matrices using genetic algorithms is discussed below. It is important to note that these simulations were designed as a proof of method, rather than a complete optimization study.

II.F. Results from Virtual Detector Simulations

Even without processing the virtual detector response matrices, intuitive results can be extracted. Figure 28 shows the results from transport material configurations 1, 3, 8, and 16 from Table 1. The response matrices were important for visualizing the detector responses of a given transport material configuration, providing a new insight into the messy nature of neutron scattering. For reference, the geometry of the transport material configurations 1, 3, 8, and 16 are shown in Figure 27.




Figure 17: Transport configurations 1, 3, 8, and 16. The outer surface of the HDPE sphere is represented by dashed red lines, and the Cd shells are represented by solid blue lines.




Figure 28: Efficiency as a function of radial position and incident neutron energy within transport material configurations 1 (TOP LEFT), 3 (TOP RIGHT), 8 (BOTTOM LEFT) and 16 (BOTTOM RIGHT).

A smooth transition in efficiency over thermal to fast neutron energies can be seen in the transport material configuration without Cd (configuration 1) where the response is unaltered by the thermal-neutron absorbing Cd. When a shell of the HDPE is replaced with cadmium (configuration 3 at 10 cm), the lower energy neutrons ($\sim 10^{-9}$ to 1 MeV) are absorbed by the cadmium and the virtual efficiency to these energies beyond that shell is reduced by many orders of magnitude. For neutrons exceeding 1 MeV, the effect of the cadmium on the efficiency occurs both inside and outside of the cadmium shell.

Intuitively, one would expect that the flow of neutrons would be largely unidirectional in the direction of incidence, and thus expect the neutron flux to be attenuated only behind a neutron absorbing material regardless of neutron energy. While this is true for the lower energy neutrons, for faster neutrons the efficiency is equally lost on both sides of the absorbing material. This is likely due to the nature of neutron scattering within the HDPE. Unlike thermal neutrons, which are already low enough in energy for detection by the MSNDs, fast neutrons must be down-scattered. During this process the fast neutrons are buzzing around in all directions within the moderating neutron spectrometers. This is not specifically surprising, as a neutron only requires two scattering events with ^1H to completely reverse its trajectory. This is analogous to the drift velocity of electrons through a semiconducting material in that the average flow of the neutrons through the HDPE in the direction of incidence is small relative to the erratic motion of the individual neutrons.

The effective range of the thermal absorber, cadmium, for attenuating fast neutrons is approximately 2–5 cm in HDPE. Thus, placing detectors within this range of the cadmium is counterproductive to maximizing the fast neutron efficiency. This is particularly important for designing a spectrometer capable of detecting SNM, as much of the usable fission neutrons will fall within the ≥ 1 MeV energy range. However, the objective is not always so straightforward, and the desired response uniqueness may outweigh the benefits of increased efficiency to the fast neutron range. Therefore, we need an effective way to analyze the tradeoff between response matrix condition number and detection efficiency. The following section covers one such method for analyzing the transport material configurations by using a genetic algorithm to select response functions based on their condition number and efficiency.

II. G. Genetic Algorithms for Response Function Selection

The virtual detector method generates a virtual response matrix containing all of the possible response functions for a single transport material configuration. However, selecting response functions from a single transport material configuration and comparing the resulting response matrices across multiple transport material configurations is non-trivial. This section covers the use of a genetic algorithm to extract the most optimal combinations of response functions from the virtual detector simulations, and a method to compare the efficiency and condition number of the optimal response matrices across multiple transport material configurations.

The virtual detector approach generates an approximation of all of the possible response functions for a given transport material configuration; the task here is to process these response functions to select the optimal combinations thereof. The optimization parameter space consists of a number of possible response functions (i.e. detector positions), and a desired number of responses (detectors), which is easily translated into binary by associating *on* (1) with the selection of a response matrix, and *off* (0) with the exclusion of that response from the selected response matrix. The equation below demonstrates the binary translation of selection of 2 response functions, r_2 and r_1 , to form a selected response matrix.

$$\begin{array}{c}
 0 \\
 1 \\
 1 \\
 0 \\
 0
 \end{array}
 \begin{array}{c}
 r_1 \\
 r_2 \\
 r_3 \\
 r_4 \\
 r_5
 \end{array}
 \begin{bmatrix}
 a & b & c & f \\
 g & h & i & j \\
 k & l & m & n \\
 o & p & q & r \\
 s & t & u & v
 \end{bmatrix}
 \rightarrow
 \begin{array}{c}
 r_2 \\
 r_3
 \end{array}
 \begin{bmatrix}
 g & h & i & j \\
 k & l & m & n
 \end{bmatrix}
 \quad (3)$$

on/off *virtual response matrix* *selected response matrix*

Evolutionary algorithms, such as genetic algorithms (GAs), are particularly well suited to optimize this binary space [37]. Given a set of possible response functions, the GA's job is to: 1) turn a set of detectors *on* (i.e. add a set of response functions to the response matrix), 2) calculate a desired quantity—the fitness metric—from this selected response matrix, 3) reject the least fit response matrices, and 4) and breed and mutate (i.e. swap response functions) response matrices of the most fit. In GA jargon, the genes of the individuals consist of a selection of virtual response functions from the virtual response matrix (i.e., the gene pool). In a sense, each individual is a virtual spectrometer design, with an associated transport material configuration and a set of detector positions corresponding to the selected response functions.

II. H. Defining a Fitness Metric

The difficulty with GA optimization is selecting and balancing the fitness metric to drive the evolution of the population. This requires describing the sets of response functions with a metric that provides the most desirable properties. Although fitness metrics are straight-forward when only a single property is desired, such as the efficiency, there are numerous desired properties of an ideal moderating spectrometer. These properties include: 1) the neutron efficiency to a particular energy or energy spectrum, 2) the energy range to be resolved, 3) the resolution for specific energies and/or over some range (note that neutron energy resolution has yet to be formally defined), and 4) the practical and monetary cost of the system (i.e. the

weight/geometry of the device and the total detector area/volume). The efficiency to a neutron spectrum is easily calculated from the response matrix by folding (i.e. taking the product) the desired spectrum into the response matrix and taking its sum. The other three properties are not so easily defined. Specifically, the energy range to be resolved, (2), and the energy-specific resolution (3) are the most difficult to quantify. Both of these properties are—to some extent—wrapped into the condition number and the energy binning structure. Manipulating the energy bins—for example, using narrow energy bins in regions where the most resolution is required, and broad bins in regions of little importance—may emphasize certain energy ranges in the condition number. Further discussion of this topic can be found in Chapter 5. For simplicity, the practical and monetary costs (4) were not considered outside of maintaining a reasonable number of detectors.

The number of response functions to select was initially left as a variable for the GA to determine. However, it quickly became apparent that allowing the GA to determine the number of responses was problematic. The condition number drove the evolution toward minimizing the number of detectors. This is believed to be in part a result of calculating a condition number of an undetermined matrix, which requires use of the pseudoinverse, minimizing the condition number as the response matrix tends toward a single response function. On the other hand, efficiency drives toward a maximum number of response functions to improve the total response (more detectors → greater device efficiency). To prevent the GA from converging on either boundary (all possible response functions or a single response function) the number of detectors was fixed at 30. This number of detectors was chosen to provide a reasonable amount of variability within the device, and still allow for a physically relevant average spacing between detectors of ~0.67 cm. Ideally, the GA would involve a weighting function to balance the number of detectors within some range, however, this was not achieved.

While holding the total number of response functions constant, the fitness metric used was a linear combination of the condition number and the efficiency, as shown in Equations 4–6. Since the condition number is often many orders of magnitude greater than the efficiency, the $\log_{10}()$ of the condition number was used.

$$E_w = \sum_e \sum_i^{30} [R_{ei} \cdot W(e)] \quad (4)$$

$$C = \|R(e, i)^{-1}\| \cdot \|R(e, i)\| \quad (5)$$

$$F = \frac{7 - \log_{10} C}{7} \cdot B_c + E_w \cdot B_E \quad (6)$$

E_w is the relative efficiency

$R(e, i)$ is the response matrix

$W(e)$ is an energy-dependent weighting function

$A(e)$ is a response-dependent weighting function

C is the condition number

F is the fitness metric

B_C and B_E are balancing factors for C and E respectively

The weighting factors provide another degree of freedom for the GA architecture. The energy dependent weighting function, $W(e)$ is used to specify the neutron spectrum to calculate the efficiency to. The function used for $W(e)$ was the Watt spectrum for U^{235} fission neutrons. B_C and B_E are factors used to balance the importance of the condition number and efficiency to the fitness metric. If B_C is zero, the GA will drive only toward improving the total efficiency. On the other hand, if B_E is zero, the GA will drive toward improving only the condition number. Variations between zero and infinity allow for a custom balance

between the two properties of the response matrix. This is a very important tool for analyzing the relationship between condition number and efficiency.

II.I. Genetic Algorithm Evolution Parameters

The genetic algorithm was written in Python 2.7 and the results described herein used a population size of 100, evolved over 500 generations. These values were selected to ensure that the initial random population covered a sufficient portion of the parameter space, and sufficient generations and mutations occurred for each GA to converge on global maximum fitness. Rather than using all 400 possible response functions produced by each virtual detector simulation, the GA was restricted to only select response functions from every 5th possible response function to ensure a radial spacing minimum of 2 mm between any two detectors, reducing the number of possible response functions from 400 to 80. This is illustrated in Figure 29, showing all 80 of the possible detector radii from which the GA selects. The dashed red-lines correspond to radii which are reserved for Cd which were off limits to the GA, bringing the actual number of available responses to a value between 76 and 80, depending on the presence of Cd.




Figure 29: Physical representation of the 80 radial detector positions (BLACK SOLID) from which the genetic algorithm selects response functions, and the four positions reserved for optional Cd placements (DASHED RED).

The first generation of 100 individuals was created by selecting at random response functions from the 80 available. The fitness of each of these individuals was then calculated via Equations 4–6, and the response matrices whose fitnesses fell within the bottom 50% were discarded. The process of gene selection is demonstrated by Figure 30.




Figure 30: Genetic algorithm process beginning with the initial population of random response functions.

One of the response matrices with a fitness in the top 20% was bred with another from the top 50% with a probability proportional to their fitnesses. The breeding process was performed gene-by-gene—that is, response function by response function—based upon the two mates’ fitnesses. For each of the individuals’ genes, a dice weighted by the mates’ fitnesses was rolled, and the winner passed on the corresponding gene (response function) to the child response matrix. This breeding was repeated 50 times to replace the 50 discarded individuals. Each new child had a 10% chance to be selected for random mutation. If selected, a die was rolled for each gene with a 1% probability for swapping out the corresponding response function for another at random to maintain a total of 30 response functions. In GA jargon, the algorithm used a 10% individual mutation rate, and a 1% gene mutation rate, and strong elitism (50%). The results from the GA ran on the virtual response matrices from the 16 transport material configurations in **Table 2** are discussed in the following section.

The Pareto Frontier

Although some improvement in both the condition number and the efficiency of the designs were seen over the course of evolution within the GA, it remained unclear as to whether or not the changes seen were indeed improvements over real devices. Each time the genetic algorithm was run on a virtual response matrix, the results varied significantly depending on the weighting factors of the condition number and efficiency. Figure 31 demonstrates the varying results from each run of the GA on the same virtual response matrix, optimizing for either condition number (circles) or efficiency (triangles). Although the transport material configurations—and thus the virtual response matrix gene pool—were identical, the GA would arrive at different solutions solely based on the weighting factors used.




Figure 31: Condition number and efficiency of individual selected response matrices from two genetic algorithm runs: One using efficiency as the fitness metric (TRIANGLES) and the other using condition number as the fitness metric (CIRCLES).

The solutions for efficiency vs. condition number driven evolution had two distinct features: 1) those from efficiency driven evolution favored an increase the detected density near the surface of the sphere and low detector density near the center, and 2) those from condition driven evolution favored a balanced distribution of detectors throughout the volume of the sphere, as shown in Figure 32 for the solid HDPE configuration (configuration 1). This is as expected, because the thermal neutron flux from a Watt spectrum peaks in the outer 5 cm of the HDPE sphere, which also correspond to very similar response functions. Thus we expect the condition number for a response matrix containing numerous response functions in the outer 5 cm of HDPE to be greater, corresponding to less resolution. However, the high thermal flux in this region corresponds to improved device efficiency. Hence, when the fitness metric is solely dependent upon the efficiency or the condition number, the GA converges on two vastly different solutions. Interestingly, when the efficiency and condition number were balanced in the fitness metric, the GA arrived at solutions that blended favoring detector density in the outermost radii with distributing detectors throughout the volume. This demonstrated on an intuitive level that the GA was indeed arriving at solutions for the detector positions which were optimizing the efficiency, condition number, or a combination thereof.

Distribution Statement A




Figure 32: Radial detector positions of response functions for transport material configuration 1 (solid HDPE) corresponding to a efficiency driven evolution (TOP) vs. condition number driven evolution (BOTTOM). The optimal solution for condition number-driven evolution is characterized by detectors distributed throughout the volume, where efficiency-driven optimization arrives at solutions containing dense placement of detectors in the outermost radii.

Attempts to balance the efficiency and condition number were largely unsuccessful: each run of the GA arrived at different fitnesses even when small changes were made to the importances of its constituents. Naively, we believed in the ability of the GA to arrive at a configuration which maximized both properties simultaneously. This was true for the initial generations of the GA: configurations were found that had improved condition number and efficiency over the randomly generated initial configuration. However, at

some point, the condition number and efficiency became negatively correlated—increasing one property required a sacrifice in the other. Because of this, using both properties in the fitness metric meant that no single optimal configuration existed, but rather a maximum boundary existed at which there is a tradeoff between the condition number and efficiency.

We had unknowingly stumbled upon a fairly common phenomena when attempting to optimize multiple parameters. If the two properties rely on a common limited resource—such as a fixed number of response functions—and these two properties have a negative relationship, a Pareto frontier exists. The Pareto frontier represents the maximum output of a system; in our case, this maximum is a combination of the optimum condition number and efficiency. By changing the importance weights of the condition number and efficiency, we were unknowingly exploring the Pareto frontier for these transport material configuration.

We explored the Pareto frontier for each of the transport material configuration simulated by varying the importance weights between efficiency and condition number. A total of 272 GA runs were performed using 17 varying importances on the 16 virtual response matrices. For each virtual response matrix, 2 of the 17 GAs explored the boundaries of the Pareto frontier—placing zero importance on either condition number or efficiency. The other 15 GAs used importance weights that covered the range between condition number and efficiency. The results for these GAs are shown in Figure 33. Each curve represents the best fit to the optimal response matrix.




Figure 33: Pareto fronts for the 16 transport material configurations generated from 272 runs of the genetic algorithm while varying the importance of efficiency and condition number.

Interestingly, the simplest transport material configuration—the configuration without any Cd—demonstrated the largest Pareto-frontier across the board. Regardless of whether this design was optimized for efficiency, condition number, or a combination thereof, the solid HDPE configuration achieved greater fitness than any other configuration. Although this is somewhat of a null result in the sense that the simplest design was optimal, it is important to note that the 16 transport material configurations explored in this text were by no means comprehensive. These do, however, represent an initial step toward exploring the parameter space of multi-detector moderating neutron spectrometers.




II.J. Closing Remarks on Optimization

Virtual detector simulations via MCNP offer a significant reduction in the time required to simulate a moderating spectrometer transport material configuration. Although the transport material configurations simulated via the virtual detector method were limited in their scope, they represent an initial study of a small portion of the transport material parameter space, laying the foundation for further exploration of spectrometer materials and geometries. The stand alone results from these simulations, although preliminary, provide a new way to both visualize and intuitively interpret the neutron scattering and absorption profile of transport materials.

In combination with genetic algorithms—or other optimization algorithms—the virtual response simulations allow us to explore the optimal properties of any given transport material configuration based upon a measure of fitness. In our case, condition number and efficiency were selected to drive the evolution, however, these are by no means the only telling fitness metrics. Although the condition number correlates with energy resolution, the fitness metric should ideally involve a measure of accuracy of a specific energy-unfolding method. However, to avoid running a simulation for every individual in the GA, the unfolding method would need to be efficient and only require a response matrix and a selected neutron spectrum. Despite the limited scope of the work discussed herein, these two tools provide future works with a method for application-driven optimization.

III. Outreach Program (SeaPerch)

There were 22 new participants in the outreach program (SeaPerch) this year. Equipment for this program was funded through the YIP award and subsidized through the AUVSI foundation. The students range from fourth to sixth grade and for many of them, this is the first time they've picked up a drill, soldered, or thought about conservation and transfer of energy.



IV. References

- [1] C.B. Hoshor, E.R. Myers, B.J. Rogers, S.M. Young, J.E. Currie, J.A. Crow, S.L. Bellinger, T.J. Sobering and D.S. McGregor, and A.N. Caruso, “Solid State Neutron Spectrometer: FY14 ONR-MWMDD Final Report”, UMKC-0001AM-TR-2014, JUN 2015 (U // FOUO / Export Controlled / Distribution Statement C). DTIC Accession No. ADB408414.
- [2] J.S. Schumer, A.N. Caruso, Z.W. Bell, B.F. Philips, R.A. Kroeger and R.J. Commisso, “Proceedings of the ONR Gamma and Neutron Background Workshop (1-2 Feb 2012)”, NRL Memorandum Report, 2013 (Distribution Statement C).
- [3] M. Matzke, “Propagation of uncertainties in unfolding procedures,” *Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip.*, vol. 476, no. 1–2, pp. 230–241, Jan. 2002.
- [4] M. Melanie, “An introduction to genetic algorithms,” *Cambridge, Massachusetts London, England*, p. 162, 1996.