

Monte Carlo simulations of a physical cryptographic
warhead verification protocol using nuclear
resonance fluorescence

by

Jayson Robert Vavrek

B.Sc. Honours Physics, University of Alberta (2014)

Submitted to the Department of Nuclear Science and Engineering
in partial fulfillment of the requirements for the degree of
Master of Science in Nuclear Science and Engineering
at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2016

© Massachusetts Institute of Technology 2016. All rights reserved.

Author

Jayson Robert Vavrek
Department of Nuclear Science and Engineering
May 20, 2016

Certified by.....

Areg Danagoulian
Assistant Professor of Nuclear Science and Engineering
Thesis Supervisor

Certified by.....

R. Scott Kemp
Norman C. Rasmussen Assistant Professor of Nuclear Science and
Engineering
Thesis Reader

Accepted by

Ju Li
Battelle Energy Alliance Professor of Nuclear Science and Engineering
Professor of Materials Science and Engineering
Chair, Department Committee on Graduate Students

Monte Carlo simulations of a physical cryptographic warhead verification protocol using nuclear resonance fluorescence

by

Jayson Robert Vavrek

Submitted to the Department of Nuclear Science and Engineering
on May 20, 2016, in partial fulfillment of the
requirements for the degree of
Master of Science in Nuclear Science and Engineering

Abstract

Future multilateral nuclear arms reduction efforts will require technologies for the verification of treaty compliance. In particular, warheads slated for dismantlement will need to be verified for authenticity without revealing any sensitive weapons design information to international inspectors. Recent efforts have investigated physical cryptographic verification protocols that attempt to solve this treaty verification problem by using physics processes rather than electronics to encrypt sensitive information. The physical cryptographic protocol simulated in this thesis exploits the isotope-specific nature of nuclear resonance fluorescence (NRF) measurements to provide a strong indicator of the authenticity of a warhead. To protect against sensitive information leakage, the NRF signal from the warhead is convoluted with that of an encrypting foil containing the same isotopes as the warhead but in unknown amounts. The convoluted spectrum from a candidate warhead is then statistically compared against that from an authenticated template warhead to determine whether the candidate itself is authentic.

This work presents the initial Geant4 Monte Carlo simulations of the physical cryptographic warhead verification protocol. Using a 2.7 MeV endpoint bremsstrahlung beam, a template warhead is interrogated. Several hoax geometries are also compared against the template to show the protocol's robustness against cheating. Isotopic hoaxes in which weapons-grade plutonium is replaced with reactor-grade plutonium or depleted uranium are shown to be detectable in realistic measurement times. An optimized geometric hoax that mimics the areal densities and attenuations of the authentic template warhead along one axis can also be detected with a second measurement under a different projection. Results of the simulations as well as future research objectives will be presented and discussed.

Thesis Supervisor: Areg Danagoulian

Title: Assistant Professor of Nuclear Science and Engineering

Thesis Reader: R. Scott Kemp

Title: Norman C. Rasmussen Assistant Professor of Nuclear Science and Engineering

“...instead of trying to make the crime impossible, they tried to provide in advance for the consequences...”

—Walter M Miller, Jr

Acknowledgments

First of all I must thank my parents Jody and Joanne, who inspired in me at an early age the love of learning that has helped me get to where I am today. They supported both my crazy idea of doing a physics undergrad and then my move across an entire continent to switch to nuclear science, and I will forever be appreciative of their encouragement to do what I love. I am also indebted to my grandfather Bob Foster, who was full of sharp advice and taught me my first algebra.

I don't think I would have been inspired to do physics in the first place were it not for Brent McDonough going out of his way to show a young student how magnificent the universe can be, as well as kind encouragement from Len Bonifacio and Patricia Milan.

James Pinfold helped me get my foot in the door by taking me on as an undergrad researcher at the University of Alberta, and the four years I spent there would not have been the same without many good friends. Spencer Axani and Tylan Murphy inspired in me a healthy sense of competition in the physics classroom and on the climbing wall, while Lee Pavelich has sparked great optimism in the future of humanity despite my field of research. I will always remember jam sessions with Levi Pardue and Jordon Rollison, as well as late-night discussions with them, Liam Innis, and Spencer Benterud. Finally, summer traditions with Adam and Noah Schwabe have been some of my fondest memories of back home.

In the Fishbowl, I must thank my fellow Canadian TD MacDonald (whose to-do list skills got me through quals) and hopefully-soon-to-be MAJ Jill Rahon. On other parts of Albany Street, Cody Dennett and Abdulla Alhajri have helped remind me that there is life outside the lab, especially post-quals, and Adam Kuang has taught me much about circuits, life, and reactor-relevant divertor diagnostics.

Zach Hartwig has provided exceptional assistance with his sleek Geant4 infrastructure and by needling me on my practice quals exam. His thesis has also been an invaluable model to my own, and I can only hope to someday write an Acknowledgments as inspiring as his. Fear is the mind-killer.

On other technical points I am also indebted to Bari Osmanov for bremsstrahlung simulations, Ruaridh Macdonald for MCNP simulations, Pablo Ducru for mathematical discussions, Derek Gaston for debugging assistance, Glen Warren for the G4NRF codebase, and he and Ken Jarman for red-teaming the project.

At MIT I have been fortunate to work with top-notch supervisors. Areg Danagoulian has been an excellent advisor, and reminds me on a daily basis the value of Fermi estimates. Scott Kemp's astounding breadth and depth of knowledge have encouraged me to think creatively about technical problems, and his willingness to revise my early drafts has brought some measure of coherency to the text. I look forward to working with you both in the years to come.

Long discussions with Dick Lanza have been helpful in setting out the future of this project (yes, Dick, I'm finally done writing) and Rachel Batista has been immensely helpful for her patience in dealing with my \$22 000 supercomputer bills and visitors from the White House.

Lastly, I'd like to thank Mark Watney, for teaching me what to do in the face of overwhelming odds.

Jayson Vavrek
Cambridge, MA, USA
May 2016

Contents

1	Introduction	17
1.1	Nuclear arms control	17
1.2	Nuclear weapon verification	18
1.2.1	Project Cloud Gap	19
1.2.2	The Black Sea experiment	20
1.2.3	Electronic information barriers	20
1.3	Physical cryptographic verification of nuclear weapons	22
1.3.1	Zero knowledge proof systems	22
1.3.2	The Princeton verification protocol	23
1.3.3	The MIT verification protocol	24
2	Nuclear resonance fluorescence	30
2.1	Pure Breit-Wigner NRF cross sections	31
2.2	Doppler-broadened NRF cross sections	32
2.3	Transmission vs backscatter NRF	36
2.4	Further corrections	37
2.4.1	Nuclear recoil	37
2.4.2	Lattice effects on Doppler broadening	39
2.5	Worked example: uranium-238	41
2.6	Angular dependence	42
2.7	NRF interaction rate in a measurement	43
2.8	Notch refill	46
2.9	Secondary bremsstrahlung background	50
3	Monte Carlo simulations	55
3.1	The G4NRF package	55
3.1.1	Nuclear data organization	56
3.1.2	Cross section evaluation	59
3.1.3	Final state sampling	60
3.1.4	Updates to the G4NRF code	61
3.2	Simulation details	64
3.2.1	Bremsstrahlung interrogation beam	64
3.2.2	Simulation geometries	66
3.3	Simulation results	69
3.3.1	Extrapolation to fixed measurement times	71

3.3.2	Statistical tests of spectra	76
3.3.3	Dose to the warhead	77
3.4	Conjectured nuclear level data	78
4	Discussion	83
4.1	Limitations	83
4.2	Future work	85
4.3	Conclusion	86
A	Standalone NRF line database	89

List of Figures

1-1	Schematic of the NRF measurement. Image: [23].	26
2-1	Illustration of an NRF transition involving the 0.680 MeV level of U-238. An incident photon of energy $E = 0.680$ MeV excites the U-238 nucleus from the ground state to the $E_r = 0.680$ MeV resonant energy level. This excited state subsequently decays with a characteristic lifetime of 0.035 ps, either directly to the ground state, emitting a photon of energy 0.680 MeV, or to the intermediate 0.0449 MeV state, emitting one photon of energy 0.6351 MeV followed by another of energy 0.0449 MeV. Image: [23].	31
2-2	Pure Breit-Wigner and Doppler-broadened cross sections for absorption by the $E_r = 2.176$ MeV U-238 resonance followed by decay to the ground state. The Doppler-broadened cross section peaks at approximately 32 b and has a width Δ_{eff} of around 1 eV.	41
2-3	Angular correlations $W(J_0, J_r, J_j, \theta)$ for common spin sequences. The pure dipole-dipole transition is denoted “dd” and the pure dipole-quadrupole transition by “dq”.	43
2-4	Simplified schematic of Fig. 1-1 used to derive Eq. 2.52, the photon energy spectrum $d^2n/d\Omega dE$ observed by the detector.	44
2-5	Plot of $R(E)$, a normalized ratio of average differential to total Compton cross sections, vs photon energy E . Over a suitable energy range, the value of $R(E)$ is very roughly 2/3. See Footnote 22.	48
2-6	Breakdown of full photon spectrum (black) into its photoelectric secondary bremsstrahlung (red), Compton scattering secondary bremsstrahlung (blue), pair lepton secondary bremsstrahlung (yellow), high-sampled NRF (green), and low-sampled NRF (violet) components for endpoint $E_{\text{end}} = 2.7$ MeV. Secondary electrons undergoing multiple interactions are shown in grey. The low-sampled NRF contribution (violet) to the NRF lines with small vertical error bars indicates a small amount of notch refill occurring in the foil itself, while the violet contribution to the NRF lines with larger vertical error bars indicates an NRF line that was not deliberately oversampled. Note that the bin at 2.42 MeV is an overlap of the highly-sampled Pu-239 decay at 2.423 MeV and the low-sampled U-238 decay line at 2.423 MeV arising from the 2.468 MeV resonance.	52

3-1	Hierarchy of G4NRF nuclear data organization. Each isotope is represented by a G4NRFNuclearLevelManager object, which contains isotope-relevant data and the vector of G4NRFNuclearLevel objects that represents the isotope's nuclear level scheme. Each G4NRFNuclearLevel contains further level information including the available decays and their respective branching ratios. The G4NRFNuclearLevelManager objects are stored in a singleton G4NRFNuclearLevelManagerStore object.	57
3-2	Smoothed Monte Carlo bremsstrahlung spectrum $p(E)$ (black) and sampling spectrum $s(E)$ (red) used in the simulations. The PDFs are normalized such that the sum rather than the true integral of their bin contents is each unity, but the distinction is immaterial when taking the ratio $w(E) = p(E)/s(E)$ as long as the distributions have identical binwidths. The spikes in $s(E)$ are narrow regions bracketing the NRF resonances of interest.	66
3-3	Graphical rendering of the Geant4 models of the Black Sea warhead (shown in cutaway, left) and the dual-slab geometric hoax (right). The plutonium layers are shown in orange, high explosive in green, uranium in blue, and reference foil in gray. The simulation only registers the events corresponding to photons emitted from the foil at $\theta \geq 135^\circ$ angles, relative to the beam axis.	69
3-4	Monte Carlo NRF spectra of the Black Sea control geometry (black curve) vs the identical Black Sea geometry with different Monte Carlo random seeds (red curve). The two spectra agree well within statistics. Lines corresponding to the four main NRF isotopes are marked by colored arrows. The errors are described in the text.	71
3-5	Monte Carlo NRF spectra of the Black Sea control (black) vs the Black Sea geometry with WGPu replaced by U-238 (red). The U-238 lines (light blue arrows) are strongly diminished in the hoax, while the Pu-239 lines (green arrows) and Pu-240 lines (orange arrows) are enhanced.	72
3-6	Monte Carlo NRF spectra of the Black Sea control (black) vs the Black Sea geometry with WGPu replaced by FGpu (red). The Pu-240 lines (orange arrows) are diminished in the hoax, but the change in Pu-239 lines (green arrows) is not easily detectable with the statistics available.	72
3-7	Monte Carlo NRF spectra of the Black Sea control (black) vs the unrotated geometric dual-slab hoax (red). The two spectra agree well within statistics.	73
3-8	Monte Carlo NRF spectra of the Black Sea control (black) vs the geometric dual-slab hoax rotated by 10 degrees (red). The two spectra no longer agree well within statistics.	73
3-9	Monte Carlo NRF spectra of the Black Sea control (black) vs the geometric dual-slab hoax rotated by 30 degrees (red). The two spectra disagree strongly in both the NRF lines and the continuum.	74

- 3-10 Plot of the ratio of the time required to achieve a desired NRF signal strength in the $E_r = 2.176$ MeV uranium-238 line with an arbitrary total foil thickness X versus the time required with $X = 2$ cm. For illustrative purposes, ratios are calculated via a modified version of Eq. 2.52 at $\theta = 0$ using only uranium-238 data with the μ_{NRF} terms divided by 4 to more closely model the NRF response of the simulated composite foil. For the $E_r = 2.176$ MeV decay to the ground state, the NRF decay cross section is $\sigma_{\text{NRF}}(E_r) = 31.5$ b and the non-resonant cross section is $\sigma_{\text{nr}}(E_r) = 18.8$ b. For $X \ll 1$ cm, the thin foil limit of Eq. 2.52 is reached, in which the NRF counts are directly proportional to X and therefore the time required to obtain a fixed number of counts is inversely proportional to X . For $X \gg 1$ cm, the saturation limit of Eq. 2.52 is reached, and the measurement time does not improve substantially. 75
- 3-11 NRF spectrum of the Black Sea control geometry. The spectrum has been scaled by a factor of ~ 254 , and the error bars have been suppressed for clarity. 76

List of Tables

2.1	Definitions of photon tags. The “isNRF” criterion corresponds to whether the photon was produced by an NRF event. The criteria logically add to unity assuming photoelectric, Compton, and pair bremsstrahlung events are the only possible non-NRF interactions that can result in a backwards photon.	52
3.1	Object, hoax, and reference foil geometries simulated in Geant4, using weapons-grade plutonium (WGPu), fuel-grade plutonium (FGPu), high explosive (HMX formula), highly-enriched uranium (HEU), U-238, and a four-isotope composite. The pair $(r_{\text{in}}, r_{\text{out}})$ denotes the inner and outer radii of the spherical shell, whereas ℓ denotes the sheet thickness along the beam axis. When elements are given rather than isotopes, the isotopic composition of each element is assumed to be natural. Isotopics are simplified from those in [49].	68
3.2	Detection probabilities for the strongest lines of the four isotopes in various hoax scenarios, as measured in units of σ . The values are computed for the strongest simulated NRF line of each isotope, using a 21 second run as described in the text. The “ $1 - P_{\text{alarm}}$ ” column indicates the probability that none of the four comparisons cross the 4σ threshold. The “Figure” column indicates the figure containing the Monte Carlo results that are scaled and then compared via Eq. 3.6 to compute the values in the isotope columns. The specific lines used are 2176 keV (U-238); 1957 keV (U-235); 2431 keV (Pu-239); and 2433 keV (Pu-240). The non-unity value for the $1 - P_{\text{alarm}}$ in the template vs authentic scenario is only due to the limited sampling in the Monte Carlo simulation.	77
A.1	Table of the 20 strongest NRF lines in the standalone database by integrated cross section with $92 \leq Z \leq 94$ and $1.7 \leq E/\text{MeV} \leq 2.7$. The cross sections quoted are specific to the E and thus are decay rather than absorption cross sections. Note that the peak Doppler-broadened cross sections $\sigma(E = E_r)$ were calculated using $T_{\text{eff}} = 300$ K unadjusted for the Debye temperatures θ and therefore may differ slightly from those computed in the main text.	90

Chapter 1

Introduction

1.1 Nuclear arms control

At the height of the Cold War nuclear arms race¹ in 1985, global nuclear weapon inventories contained $\sim 60\,000$ warheads, the United States holding 21 000 and the Soviet Union 39 000 at the time. These numbers have fallen to ~ 7000 each² [1] through arms reduction treaties such as SORT³ and New START⁴, as well as nuclear fuel conversion programs such as Megatons to Megawatts [4]. Although the majority of these warheads are stockpiled or retired, both the US and Russia each still have ~ 1800 warheads deployed and 900 on high alert, many more than the few hundred required to kill millions outright [5] and present an existential threat to humanity in the form of nuclear winter [6].

The further reduction of nuclear arsenals has thus been a major concern of policymakers, academics, and scientists. Moreover, the major nuclear powers have committed—at least in principle—to eliminating their nuclear weapons under the NPT⁵. To this end, the 2010 US Nuclear Posture Review Report outlines the official US commitment to the goal of eventual total nuclear disarmament, and identifies the

¹As measured by world total nuclear weapon stockpiles.

²These numbers are subject, however, to considerable uncertainty.

³Strategic Offensive Reductions Treaty [2].

⁴New Strategic Arms Reduction Treaty [3].

⁵Treaty on the Nonproliferation of Nuclear Weapons [7].

conditions required for the US to securely give up its arsenal [8].

In recent years, however, disarmament progress has been slow. Nuclear weapon states have been reluctant to reduce their arsenals, perceiving a loss of deterrence at low numbers of weapons. Furthermore, as the number of weapons decreases, the perceived marginal utility of a single weapon increases, increasing a state’s incentive to disregard a disarmament treaty by withdrawal, abrogation, or establishment of a covert nuclear reserve. The status quo of bilateral *indirectly*-verified disarmament⁶ therefore is unlikely to be maintained as the number of weapons decreases, and parties to future disarmament treaties will not dismantle their weapons until there exists a method of verifying that their counterparts are doing the same.

1.2 Nuclear weapon verification

One of the main obstacles to future disarmament is the current lack of “verification methods and technologies capable of detecting violations of disarmament obligations” [8] including techniques to verify the authenticity of warheads slated for dismantlement. Historically, there has been a perceived link between the intrusiveness and the confidence in a warhead verification method: a more confident test is bound to be more intrusive, revealing more classified information such as component dimensions and isotopic compositions (though the reverse is not necessarily true). Nuclear states have thus been highly resistant to existing warhead verification methods, citing national security concerns, while paradoxically insisting that future arms reduction treaties must be effectively verifiable. Drawing on a phrase common in arms control—the “verification problem” [10, 11]—this paradox can be called the *disarmament* verification problem:

How can a warhead be reliably identified as authentic without revealing classified information?

To date, no warhead verification procedure has been shown to solve this disarmament verification problem. The physical cryptographic warhead verification protocol

⁶The SALT II treaty, for instance, verifies delivery vehicles and not the warheads themselves [9].

described in this work, however, challenges the historical link between confidence and intrusiveness and provides a promising solution to the disarmament verification problem.

For historical context, Sections 1.2.1 and 1.2.2 outline two previous attempts at warhead verification protocols, Project Cloud Gap and the Black Sea experiment, that were influential in establishing the disarmament verification problem. Section 1.2.3 discusses modern attempts to circumvent the problem altogether by concealing all classified measurements and thus any sensitive weapons information behind electronic information barriers. Finally, Section 1.3 discusses more recent verification work that draws on the ideas of zero-knowledge proofs and cryptography, concluding with an overview of this work’s protocol in Section 1.3.3.

1.2.1 Project Cloud Gap

In 1967, the US Department of Defense conducted a mock warhead verification and dismantlement exercise, Project Cloud Gap, in which military officers acted as weapons inspectors [12]. The inspectors were given various levels of access to a suite of real and fake warheads, and tasked with judging the authenticity of each of the weapons. At the lowest level of access, the inspectors had access to the external weapon configuration only, and could make only crude, unspecific measurements such as the mass and moments of inertia of the weapon. With these tests, the inspectors correctly identified 85% of real warheads but only 3% of fake warheads. These rates improved to 95% and 62%, respectively, at the highest level of access, which included Geiger and neutron measurements, high-purity germanium (HPGe) spectroscopy, and x-ray imaging. Even at the lowest level of access, however, significant amounts of classified information were released, and the leaks only worsened at higher levels of access. Project Cloud Gap therefore established the notion that even simple verification procedures are susceptible to classified information leakage, and that stronger confidence of authenticity necessitates more intrusive measurements.

1.2.2 The Black Sea experiment

In 1989, American and Soviet scientists collaborated on a joint nuclear warhead verification exercise on the Soviet cruiser *Slava*, then located in the Black Sea [13]. The aim of the experiment was to test whether a surface-ship tactical nuclear weapons ban could be effectively verified. Using helicopter-borne He-3 neutron detectors, the Russian team was able to detect the enhanced neutron radiation from the warhead's fissile material with 3σ confidence up to 80 m away. The American team measured the warhead's photon energy spectrum at a distance of less than a meter, and later published their findings in the open literature.

The authors were able to derive from the photon spectrum the approximate isotopic composition of the Soviet warhead, including the fissile plutonium-239 and uranium-235, tritium breeder lithium-7, and reactor contaminant uranium-232. Furthermore, weapons scientists at Lawrence Livermore National Laboratory were able to combine the observed spectrum with their own advanced design knowledge to calculate even more precise design specifications [14]. Although these more precise values were kept classified, a significant amount of sensitive design information may have leaked as a direct result of the Black Sea experiment.

1.2.3 Electronic information barriers

In order to prevent the leakage of such classified information, the US National Laboratories have researched the use of electronic information barriers in warhead verification measurements. This technique allows the direct measurement of classified warhead attributes behind a software, firmware, and/or hardware barrier, returning only binary (one-bit true/false), unclassified results as to the authenticity of the warhead. Classified warhead data generated in the course of the measurement is thus electronically encrypted, but the result is not.

Information barriers were first used to protect proprietary information in uranium enrichment measurements rather than classified information in warhead veri-

fication [15]. Using gamma-ray spectroscopy and x-ray fluorescence⁷ measurements, the system determines the enrichment levels of gaseous uranium hexafluoride (UF_6) in the header pipe of an operating centrifuge and returns only a binary true/false as to whether the measured enrichment level matches that declared ahead of time by the facility operator.

This ‘first-generation’ information barrier relied primarily on passive security, preventing the release of proprietary information by restricting access to the measurement data. A more active approach to security was introduced in the ‘second-generation’ information barriers designed specifically for warhead verification measurements. In these newer iterations, a ‘security watchdog’ not only provides a passive information barrier but actively monitors the computational and electronic systems for intrusions. Such an information barrier system was used in the Fissile Material Transparency Technology Demonstration (FMTTD), in which the Attribute Measurement System with Information Barrier (AMS/IB) was used to verify a series of nuclear weapon components to an uncleared Russian delegation at Los Alamos National Laboratory in 2000 [16].

As with any encryption system—physical or electronic—the system must be proven to be secure. Since an information barrier relies on electronic encryption to protect classified weapons information, the system must be certified against vulnerabilities such as software bugs, hacks, and backdoors that might have been introduced accidentally by programmer error or deliberately by a nefarious actor. Such certification is an arduous if not impossible task, and its efficacy lately called into question by discoveries that the US National Security Agency paid RSA Security to use and promote a cryptographic protocol to which the NSA had backdoor access [17].

A further problem with attribute systems is that the attributes must be highly specific to those of real warheads, and thus are liable to be classified themselves. Thus even establishing a list of what attributes are to be measured may reveal classified weapons information, such as the minimum mass of the plutonium pit. Relaxing the

⁷X-ray fluorescence (XRF) is the atomic analogue of the nuclear resonance fluorescence (NRF) process used as an interrogation method in this work.

specificity of the attributes broadens the space of objects that can pass an attribute verification test, opening the system up to hoaxes or diversions.

1.3 Physical cryptographic verification of nuclear weapons

1.3.1 Zero knowledge proof systems

To surmount the disarmament verification problem (and especially the problematic proposed solutions relying on electronic encryption) recent disarmament work has investigated zero knowledge (ZK) proof systems [18], a formal construct in computer science and cryptography that can prove a statement without revealing any additional information. A ZK proof is an iterative, interactive process between two parties termed the ‘prover’ and the ‘verifier’. The verifier queries the prover as to the truth of the statement in question, and the prover responds in such a way so as to convince the verifier without revealing *why* the statement is true. The process then iterates, allowing the verifier to increase his/her confidence in the prover’s statement.

This can be made more clear by considering a toy example in which the prover claims to know the number of needles on all the pine trees at any time in a forest, but may not disclose the actual number (or the method by which he/she knows the number) to the verifier. The verifier wishes to test the prover’s claim, so in an interactive zero-knowledge proof, the verifier will then remove a number of pine needles from the forest, keeping this number secret from the prover. Since the prover ostensibly knows the number of pine needles in the forest at any time, he/she can calculate the *difference* between the start and end of this step—equal to the number of pine needles removed by the verifier—and report this number to the verifier. If the number is correct, the verifier attains some measure of confidence in the prover’s claim that he/she knows the number of pine needles in the forest. By repeating this measurement, the verifier can attain an arbitrarily high degree of confidence that the prover is not just making fortuitous guesses. In this way, a ZK proof will not reveal

why the statement in question is true, converging *only* to a binary true/false value.

In its original cryptographic formulation, a ZK proof system must satisfy the following criteria [19]:

1. completeness: the system correctly identifies as true all true statements,
2. soundness: the system correctly identifies as false all false statements, and
3. zero-knowledge: for an honest prover, the system only reveals the truth or falsity of the statement.

In the context of warhead verification, a verification protocol based on a ZK proof system would verify the authenticity of a nuclear warhead without revealing the classified information as to *why* it is authentic, thus solving the disarmament verification problem. In this formulation, a ZK warhead verification protocol must satisfy the following criteria to the greatest extent possible:

1. completeness: the protocol correctly authenticates all real warheads,
2. soundness: the protocol correctly rejects all fake warheads,
3. zero-knowledge: for an honest host, the protocol only reveals the authenticity of the warhead.

In a warhead verification scenario, requirement 3 corresponds to the protection of classified weapons information.

1.3.2 The Princeton verification protocol

In 2014 the Princeton-based team of Glaser, Barak, and Goldston [20] proposed a ZK warhead verification protocol using transmission neutron measurements. A candidate warhead is irradiated with 14.1 MeV neutrons from a deuterium-tritium (DT) generator, and the transmission neutron radiograph of the warhead is captured by an array of neutron bubble detectors. The active volume of each detector is a superheated fluid that forms bubbles at a rate approximately proportional to the neutron dose

rate. Since the bubbles are stable, the final number of bubbles gives a measure of the integrated neutron flux on the detector. Such a radiograph would contain sensitive information about the design of the warhead, and thus cannot be made available to an inspector. However, if a radiograph of an authentic weapon is measured first, the neutron detectors in the candidate measurement can be preloaded with the *negative* of the authentic radiograph. The bubble detectors used for the candidate radiograph should therefore have a uniform number of bubbles (within statistical fluctuations) after irradiation if the candidate and authentic warheads are identical⁸.

Monte Carlo simulations in MCNP5 [21] suggest that the Princeton ZK protocol can detect both the diversion of tungsten (as a surrogate for weapon isotopes) and the replacement of tungsten with lead in a radiographed test object while passing authentic objects with high confidence. However, further simulations have shown that such neutron measurements are highly vulnerable to isotopic hoaxes using materials with very similar neutron cross sections [22], failing the soundness criterion. In particular, the Princeton ZK protocol cannot distinguish between weapons-usable plutonium-239 and depleted uranium, or between weapons-grade and fuel-grade plutonium⁹. Additionally, since the array of bubble detectors provides spatial resolution, the protocol (like all template verification systems) is fragile: small misalignments may result in bubble detectors that contain only the authentic radiograph, thus revealing classified information. Unless these weaknesses are addressed, the Princeton protocol cannot be used to verify the authenticity of nuclear warheads.

1.3.3 The MIT verification protocol

Rather than neutron interactions, which are not isotope-specific, the MIT verification protocol simulated in this work uses nuclear resonance fluorescence (NRF) photon measurements to explicitly measure the isotopic composition of the candidate and

⁸One cannot irradiate a neutron detector with negative counts, so the authentic radiograph is subtracted from a uniform, agreed-upon offset number of counts N_{\max} . The candidate radiograph and the negative preload should therefore sum to N_{\max} counts per detector if the objects are identical, and $\pm 1\sigma$ statistical fluctuations should fall within $N_{\max} \pm \sqrt{N_{\max}}$.

⁹The thicknesses of the hoax materials must be adjusted slightly to give the same total attenuation with different densities.

authentic warheads.

The verification protocol itself is a joint exercise drawing on the ideas of a zero-knowledge interactive proof. In this scenario, the prover is known as the *host* and is the country that owns the weapons being verified (and eventually dismantled) as well as the verification facility. Similarly, the verifier is known as the *inspector*. In bilateral arms reduction, the two countries may alternate measurements, trading the roles of host and inspector to incrementally reduce nuclear arsenals while maintaining approximate force parity. The inspector may also represent an independent third party such as the International Atomic Energy Agency (IAEA).

The measurement consists of five primary components: the interrogation beam, the template warhead, the candidate warhead, the reference foil, and the detectors; a schematic is shown in Fig. 1-1. The interrogation beam, which is owned by the host, is first generated by any of various means. In this work a bremsstrahlung beam from MeV-scale electrons is used, but future studies may use tunable quasimonochromatic photon sources as the technology matures. Once the beam is generated, it is used to interrogate the template warhead, which is also owned by the host. The template warhead must be shown to be authentic prior to the NRF measurement, which could be done for instance by removing it from an active intercontinental ballistic missile (ICBM). As will be discussed in Section 2.7, the interrogation beam is preferentially attenuated in the forward direction in narrow energy bands (‘notches’) due to NRF interactions in the warhead; this physically encodes a signature of the warhead into the flux that is transmitted through the weapon.

Although a measurement of the transmitted flux itself could in principle be used to determine the authenticity of the template warhead, there are two problems with doing this. Practically, the width of the spectral notches is $\mathcal{O}(\text{eV})$, which is far smaller than the $\mathcal{O}(\text{keV})$ resolution of the most precise high-purity germanium (HPGe) detectors. The notch would therefore reduce the measured flux in one energy bin by one part in a thousand, limiting the sensitivity of the test. Moreover, a direct measurement of the warhead, as in the Cloud Gap and Black Sea experiments, would reveal a significant amount of classified information.

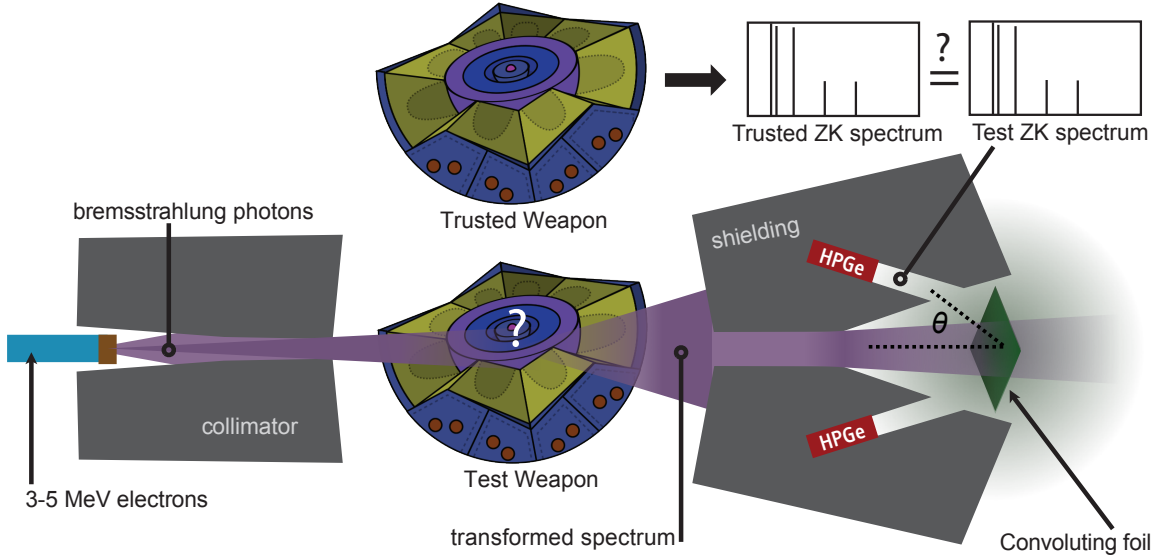


Figure 1-1: Schematic of the NRF measurement. Image: [23].

Rather than directly measure the NRF signature of the warhead, the transmitted spectrum next impinges on the reference foil (also host-owned), which contains all the isotopes of interest in the warhead but in unknown amounts—in this work, the isotopes of interest are defined to be uranium-235, uranium-238, plutonium-239, and plutonium-240. NRF interactions again occur with re-emission in 4π , and some NRF photons will interact to produce a signal in the HPGe detector(s), owned by the inspector. Since the energy spectrum as measured by the detectors is a function of both the warhead and the foil interactions, the individual contribution from and therefore the classified design information regarding each component is physically encrypted. This same NRF measurement is then made on the candidate warhead, and the spectra from the two measurements are compared for statistical identity, asking in broad terms whether

$$(\text{foil} \otimes \text{candidate}) \stackrel{?}{=} (\text{foil} \otimes \text{authentic}), \quad (1.1)$$

where \otimes denotes that the effects of the foil and the candidate/authentic warhead on the final spectra are in some sense ‘convoluted’¹⁰. Corresponding to the iterative con-

¹⁰Although not a convolution in the strict sense of the integral transform $(f \otimes g)(t) \equiv \int_{-\infty}^{\infty} f(\tau) \cdot g(t - \tau) d\tau$ with f representing the foil and g the warhead, the term is convenient for

vergence of a ZK proof to its final truth value, the pair of NRF measurements is then repeated at different warhead orientations, with the orientation the same between the candidate and authentic warheads but different for each set of measurements. If the pairs of convoluted spectra under each projection match within statistics, the candidate warhead is authenticated. Note that since the measured convoluted spectra are available to the inspector, and the inspector judges the warhead’s authenticity based on an analysis of the spectra, the protocol is not truly zero knowledge in the strict sense of a one-bit genuine/hoax result. To avoid confusion therefore, the MIT verification protocol, which nevertheless draws on the ideas of ZK proofs, is termed “physical cryptographic verification” for the encryption of sensitive weapon information by purely physical—rather than electronic—processes, namely the convolution of warhead and foil contributions to the spectra.

expressing the fact that the individual foil and warhead contributions are scrambled together to create the final signal.

Chapter 2

Nuclear resonance fluorescence

Nuclear resonance fluorescence (NRF) describes the resonant absorption and subsequent re-emission of a photon by a nucleus in its ground state. Upon absorption, an incident photon of energy sufficiently close to a higher nuclear energy level will induce a nuclear transition to this level, promoting the nucleus to an excited state [24]. This excited state will decay by photon emission through all possible channels, and the emitted photon energy will be approximately equal to the difference between the excited and final energy levels. In the case of a direct transition to the ground state, the emitted photon energy is approximately equal to the incident energy, while for decay via an intermediate level, the emitted photon energy is correspondingly reduced. In either case, the photon energies are specific to an individual set of nuclear levels, and thus to an individual isotope. As such, NRF measurements may be used to identify both the presence and concentration of specific isotopes in a sample. An illustration of NRF absorption and decay specific to U-238 can be seen in Fig. 2-1.

The formalism used to quantify NRF is described in the next sections. The following derivations follow closely those of Metzger [25], but introduce more explicit notation and make clearer various approximations and assumptions.




Figure 2-1: Illustration of an NRF transition involving the 0.680 MeV level of U-238. An incident photon of energy $E = 0.680$ MeV excites the U-238 nucleus from the ground state to the $E_r = 0.680$ MeV resonant energy level. This excited state subsequently decays with a characteristic lifetime of 0.035 ps, either directly to the ground state, emitting a photon of energy 0.680 MeV, or to the intermediate 0.0449 MeV state, emitting one photon of energy 0.6351 MeV followed by another of energy 0.0449 MeV. Image: [23].

2.1 Pure Breit-Wigner NRF cross sections

We begin with the total resonant absorption NRF cross section. This will allow the decoupling of the absorption and re-emission (decay) processes. For incoming photons of energy E , the pure quantum-mechanical Breit-Wigner cross section¹¹ for resonant absorption by a nucleus in its ground state is

$$\sigma_{0,r}(E) = \pi g_r \left(\frac{\hbar c}{E_r} \right)^2 \frac{\Gamma_{0,r} \Gamma_r}{(E - E_r)^2 + (\Gamma_r/2)^2}, \quad (2.1)$$

where the two-index $0, r$ subscript denotes a transition from the ground state (0) to the resonance level r that lies at energy E_r . The term $\Gamma_{0,r}$ is therefore the partial

¹¹This can be derived from the Schrödinger equation in a finite potential scattering problem [26, 27], but this is beyond the scope of this work.

width of the transition from ground to excited state, while Γ_r is the total width of the excited state, related to its lifetime τ_r by $\Gamma_r \simeq \hbar/\tau_r$. When ambiguity is possible, a *single* index will be used to describe a *level*, while a *double* index will be used to describe a *transition*. Lastly, g_r is a statistical factor, given by

$$g_r = \frac{2J_r + 1}{2(2J_0 + 1)}, \quad (2.2)$$

where J_r and J_0 are the total angular momenta of the nucleus at the resonance and ground states, respectively.

For a nucleus in a particular excited state r , the probability of its decay to a lower-energy state j by emission of a photon (of energy $E_r - E_j$, neglecting nuclear recoil, where $E_0 \equiv 0$) is given by its branching ratio

$$b_{r,j} \equiv \frac{\Gamma_{r,j}}{\Gamma_r}, \quad (2.3)$$

where summing over all possible intermediate states j (where implicitly $j < r$) gives

$$\sum_j \Gamma_{r,j} = \Gamma_r \quad \text{i.e.} \quad \sum_j b_{r,j} = 1. \quad (2.4)$$

The cross section for resonant absorption followed by photon emission through channel j is then simply the product of Eqs. 2.1 and 2.3:

$$\sigma_{0,r,j}(E) = \pi g_r \left(\frac{\hbar c}{E_r} \right)^2 \frac{\Gamma_{0,r} \Gamma_{r,j}}{(E - E_r)^2 + (\Gamma_r/2)^2}, \quad (2.5)$$

where the three-index $0, r, j$ subscript denotes the occurrence of two transitions.

2.2 Doppler-broadened NRF cross sections

Thermal motion of nuclei in the target medium induces a Doppler shift in the relative energies of the absorbing nuclei and the incoming photon, causing a broadening of spectral lines. The photon energy E' seen by a nucleus with velocity component

v along the photon direction (with $v > 0$ indicating motion toward the photon) is shifted from the stationary $v = 0$ case by, in the fully relativistic case,

$$E' = E \sqrt{\frac{1 + v/c}{1 - v/c}}. \quad (2.6)$$

The velocities of the target nuclei in the direction of the incoming photon are distributed according to the 1D Maxwell-Boltzmann distribution

$$w(v) dv = \sqrt{\frac{M}{2\pi kT}} \exp\left(-\frac{Mv^2}{2kT}\right) dv, \quad (2.7)$$

where M is the mass of the target nucleus and T is the temperature of the target medium. One can verify that $\int_{-\infty}^{+\infty} w(v) dv = 1$. Converting to a distribution in E' ,

$$w(v) dv = w(E') dE' = \sqrt{\frac{M}{2\pi kT}} \exp\left(-\frac{Mv^2}{2kT}\right) \frac{dv}{dE'} dE'. \quad (2.8)$$

To first order, Eq. 2.6 becomes

$$E' = E \left(1 + \frac{v}{c}\right) \implies dv = \frac{c}{E} dE', \quad v^2 = c^2 \left(\frac{E' - E}{E}\right)^2, \quad (2.9)$$

such that the distribution of energies E' is approximately

$$w(E') dE' = \sqrt{\frac{Mc^2}{2\pi kT}} \frac{1}{E} \exp\left[-\frac{Mc^2}{2kT} \frac{(E' - E)^2}{E^2}\right] dE'. \quad (2.10)$$

Writing the Doppler-broadened width¹² of this Gaussian distributions in E' as

$$\Delta \equiv E \sqrt{\frac{2kT}{Mc^2}}, \quad (2.11)$$

¹²One can show that $\Delta \sim 1$ eV is much smaller than the typical spacing between different resonances (~ 10 keV), and thus the use of single-level Breit-Wigner is justified over the multi-level Breit-Wigner or Reich-Moore formalisms.

the distribution becomes

$$w(E') dE' = \frac{1}{\sqrt{\pi}\Delta} \exp \left[-\frac{(E' - E)^2}{\Delta^2} \right] dE'. \quad (2.12)$$

Given the first-order expansion¹³ in Eq. 2.9, the bounds on the integral $\int w(E') dE'$ must be $E' \rightarrow \pm\infty$ rather than the more physical choices of 0 and $+\infty$. The distinction is insignificant, however, since the latter choice will result in a term of the form $\text{erf}\sqrt{Mc^2/2kT}$, which is very nearly 1. Physically, this means the distribution is quite sharply peaked and the difference in area under the $w(E')$ curve when using a lower bound of $-\infty$ is negligible. In this case,

$$\int_{-\infty}^{+\infty} w(E') dE' = 1 \quad (2.15)$$

is properly normalized.

Proceeding via Eqs. 2.5 and 2.12, the Doppler-broadened NRF cross section is just the average of $\sigma_{0,r,j}$ over the distribution $w(E')$:

$$\begin{aligned} \sigma_{0,r,j}^D(E) &= \int_{-\infty}^{+\infty} \sigma_{0,r,j}(E') w(E') dE' \quad (2.16) \\ &= \sqrt{\pi} g_r \left(\frac{\hbar c}{E_r} \right)^2 \frac{1}{\Delta} \int_{-\infty}^{+\infty} \frac{\Gamma_{0,r} \Gamma_{r,j}}{(E' - E_r)^2 + (\Gamma_r/2)^2} \exp \left[-\frac{(E' - E)^2}{\Delta^2} \right] dE', \end{aligned} \quad (2.17)$$

where the choice of integral bounds corresponds to that in Eq. 2.15. The integral

¹³It may be of some interest to note that using the full Eq. 2.6 gives

$$\left(\frac{dE'}{dv} \right)_{\text{exact}} = \frac{E'}{c} \frac{1}{1 - v^2/c^2} \quad (2.13)$$

and thus

$$w_{\text{exact}}(E') dE' = \sqrt{\frac{M}{2\pi kT}} \exp \left(-\frac{Mc^2}{2kT} \left[\frac{(E'/E)^2 - 1}{(E'/E)^2 + 1} \right]^2 \right) \left| \frac{c}{E'} \left(1 - \left[\frac{(E'/E)^2 - 1}{(E'/E)^2 + 1} \right]^2 \right) \right| dE', \quad (2.14)$$

which is no longer Gaussian, has no analytical integral, and *does not numerically integrate to unity in general* on the interval $[0, +\infty)$. The first-order expansion in Eq. 2.9 is justified, however, in that v/c is small compared to unity (*e.g.* the most probable speed is $v/c = \sqrt{2kT/Mc^2} \sim 10^{-7}$ in uranium-238 at room temperature $kT \simeq 1/40$ eV) and higher-order expansions will involve only higher powers of (v/c) .

of Eq. 2.17, however, has no analytical solution. For numerical convenience, the substitutions

$$x \equiv 2(E - E_r)/\Gamma_r, \quad (2.18)$$

$$y \equiv 2(E' - E_r)/\Gamma_r, \quad (2.19)$$

$$t \equiv (\Delta/\Gamma_r)^2 \quad (2.20)$$

are made, transforming Eq. 2.17 to

$$\sigma_{0,r,j}^D(E) = 2\pi^{1/2} g_r \left(\frac{\hbar c}{E_r} \right)^2 \frac{\Gamma_{0,r} \Gamma_{r,j}}{\Gamma_r^2} \frac{1}{\sqrt{t}} \int_{-\infty}^{+\infty} \exp \left[-\frac{(x-y)^2}{4t} \right] \frac{dy}{1+y^2}. \quad (2.21)$$

Aside from further corrections discussed in Section 2.4, Eq. 2.21 is the final expression for the NRF $0 \rightarrow r \rightarrow j$ cross section and is used throughout this work. In the literature, however, questionable approximations are made to reduce the expression further. One tactic is to write $(x-y)^2 \simeq x^2 + y^2$, resulting in

$$\sigma_{0,r,j}^D(E) = 2\pi^{1/2} g_r \left(\frac{\hbar c}{E_r} \right)^2 \frac{\Gamma_{0,r} \Gamma_{r,j}}{\Gamma_r^2} \frac{1}{\sqrt{t}} \exp \left(-\frac{x^2}{4t} \right) \int_{-\infty}^{+\infty} \exp \left(-\frac{y^2}{4t} \right) \frac{dy}{1+y^2}. \quad (2.22)$$

The integral then evaluates to

$$\int_{-\infty}^{+\infty} \exp \left(-\frac{y^2}{4t} \right) \frac{dy}{1+y^2} = \exp \left(\frac{1}{4t} \right) \pi \operatorname{erfc} \left(\sqrt{\frac{1}{4t}} \right) \simeq \pi \quad (2.23)$$

for $t \equiv (\Delta/\Gamma_r)^2 \gg 1$. Thus, the Doppler-broadened cross section for the NRF process $0 \rightarrow r \rightarrow j$ is written as

$$\sigma_{0,r,j}^D(E) = 2\pi^{3/2} g_r \left(\frac{\hbar c}{E_r} \right)^2 \frac{\Gamma_{0,r} \Gamma_{r,j}}{\Gamma_r} \frac{1}{\Delta} \exp \left[-\frac{(E - E_r)^2}{\Delta^2} \right], \quad (2.24)$$

which is a Gaussian distribution in E centered at the resonance E_r with variance $\Delta^2/2$.

Proceeding without this questionable approximation, the total integrated cross section $\int \sigma_{0,r,j}^D(E) dE$ of this process, which quantifies the overall strength or likeli-

hood of the interaction, presents somewhat of a challenge. Recalling explicitly that $\Delta \equiv E\sqrt{2kT/Mc^2}$, the integral

$$\int_0^\infty \sigma_{0,r,j}^D(E) dE \quad (2.25)$$

exhibits an infrared divergence. Interestingly, though, the integrated non-Doppler-broadened cross section

$$\int_0^\infty \sigma_{0,r,j}(E) dE = \pi g_r \left(\frac{\hbar c}{E_r}\right)^2 \frac{\Gamma_{0,r}\Gamma_{r,j}}{\Gamma_r} \left[\pi + 2 \arctan\left(\frac{2E_r}{\Gamma_r}\right) \right] \quad (2.26)$$

$$= 2\pi^2 g_r \left(\frac{\hbar c}{E_r}\right)^2 \frac{\Gamma_{0,r}\Gamma_{r,j}}{\Gamma_r}, \quad (2.27)$$

where in the last line the approximation $E_r \gg \Gamma_r$ has been made, can be replicated by simply ignoring the E -dependence of Δ . In this way, treating Δ as constant with E , the total integrated Doppler-broadened cross section is

$$\int_0^\infty \sigma_{0,r,j}^D(E) dE = \pi^2 g_r \left(\frac{\hbar c}{E_r}\right)^2 \frac{\Gamma_{0,r}\Gamma_{r,j}}{\Gamma_r} \left[1 + \operatorname{erf}\left(\frac{E_r}{\Delta}\right) \right] \quad (2.28)$$

$$= 2\pi^2 g_r \left(\frac{\hbar c}{E_r}\right)^2 \frac{\Gamma_{0,r}\Gamma_{r,j}}{\Gamma_r}, \quad (2.29)$$

where the approximation $E_r \gg \Delta$ has been applied. In both cases, sending the arctan and erf functions to their limiting values is the same as extending the lower integral bound to $-\infty$. On physical grounds, it should be expected that the total integrated cross section is invariant under the Doppler effect and thus convergent. Furthermore, since the $\int w(E') dE'$ integral is just a change of variables from $\int w(v) dv$, which is known to converge, the underlying physics should not change.

2.3 Transmission vs backscatter NRF

Eq. 2.21 gives the Doppler-broadened cross section for a particular transition $0 \rightarrow r \rightarrow j$. Since these transitions are typically observed in backscatter experiments, their cross sections are referred to as backscatter cross sections. Conversely, the

transmission NRF cross section for a single resonance r describes the sum over all decay levels j ,

$$\sigma_{0,r}^D(E) = 2\pi^{1/2} g_r \left(\frac{\hbar c}{E_r} \right)^2 \frac{\Gamma_{0,r}}{\Delta} \int_{-\infty}^{+\infty} \exp \left[-\frac{(x-y)^2}{4t} \right] \frac{dy}{1+y^2}, \quad (2.30)$$

which differs from Eq. 2.21 only by a factor of the branching ratio, effectively decoupling the absorption step from the re-emission step. Since transmission NRF measurements are dependent on the amount of NRF absorption, these cross sections are referred to as both transmission and absorption cross sections. In the `G4NRF` simulation code used in this work, the transmission cross section is calculated by the function `NRF_xsec_calc()`, while the decay with branching ratio $b_{r,j} \equiv \Gamma_{r,j}/\Gamma_r$ is handled in the `PostStepDoIt()` function.

Finally, the sum over all resonances r gives the total NRF interaction cross section for a photon of energy E passing through a material:

$$\sigma_0^D(E) = \sum_r \sigma_{0,r}^D(E). \quad (2.31)$$

Given the $\Gamma_{0,r}$ and E_r factors, which do not reduce under the sum, this expression is best evaluated as a Monte Carlo sampling of resonances r —a process handled automatically by `Geant4`. Since the NRF level spacing is wide compared to the resonance width, however, there is very little overlap between resonances, and all but one term in the sum can typically be ignored in analytical calculations.

2.4 Further corrections

2.4.1 Nuclear recoil

To satisfy conservation of both energy and linear momentum simultaneously, the target nucleus must recoil during both the absorption and re-emission stages of an

NRF interaction. Using 1D kinematics¹⁴ to keep consistent with the collision model implicit in Eq. 2.9, the magnitude of the nuclear recoil upon absorption is

$$(\Delta E)_{\text{abs}} = \frac{E^2}{2Mc^2} \quad (2.35)$$

and upon re-emission with energy \tilde{E} ,

$$(\Delta E)_{\text{em}} = \frac{\tilde{E}^2}{2Mc^2}. \quad (2.36)$$

Since the nuclear recoils ΔE are small, $E \simeq \tilde{E}$, and thus the total recoil energy loss is approximately

$$\Delta E = (\Delta E)_{\text{abs}} + (\Delta E)_{\text{em}} \simeq \frac{E^2}{Mc^2}, \quad (2.37)$$

on the order of 10 eV for a uranium atom. In effect, this nuclear recoil shifts the emitted photon energy slightly lower than the resonant energy E_r by the amount ΔE : an incident photon exactly on-resonance, then, has an energy equal to the emitted energy *plus* the compensation for the nuclear recoil.

The magnitude of the nuclear recoil ΔE is small compared to the photon and level energies involved, and can be ignored in some analytical calculations. In particular,

¹⁴In the more general 2D case the recoil energy loss can be determined from the Compton-like formula

$$\tilde{E} - E = E \left[\frac{1}{1 + E(1 - \cos \chi)/Mc^2} - 1 \right] \quad (2.32)$$

where χ is the angle at which the photon is scattered relative to its incident direction. Making the approximation $E \ll Mc^2$, this reduces to

$$\tilde{E} - E \simeq \frac{E^2(1 - \cos \chi)}{Mc^2}. \quad (2.33)$$

Then at $\chi = \pi$ the magnitude of recoil energy loss reaches its maximum value of $2E^2/Mc^2$. Conversely, at $\chi = 0$ (forward scattering), $\tilde{E} - E = 0$, and there is no nuclear recoil. Averaging over χ , the average recoil loss in the case of isotropic scattering is

$$\langle \tilde{E} - E \rangle = \frac{\int_0^\pi (\tilde{E} - E) \sin \chi d\chi}{\int_0^\pi \sin \chi d\chi} \simeq \frac{E^2}{Mc^2}, \quad (2.34)$$

agreeing with the 1D result of Eq. 2.37.

the distinction between the true nuclear energy level E_r and the recoil-corrected emission energy $E = E_r - \Delta E$ can be disregarded if functions of these variables (such as Compton scattering cross sections) are slowly varying over the range $[E_r - \Delta E, E_r]$. As such, it is typically acceptable to speak of “the 2.176 MeV resonance” even though the 2.176 MeV measured in experiments and tabulated in databases is the emitted photon energy, while the true resonance lies ~ 10 eV higher¹⁵. Important exceptions to this are the NRF cross sections themselves—the downshift in emitted photon energy of $\Delta E \sim 10$ eV induced by the nuclear recoil is much larger than the Doppler-broadened width $\Delta \sim 1$ eV of the resonance¹⁶. An NRF photon of energy $E_r - \Delta E$ emitted by direct decay from the resonant level E_r to the ground state will therefore be unlikely to go on to induce a *second* NRF interaction at the same resonance (of a neighbouring nucleus), since it will be approximately 10 standard deviations below the center of the resonance (see for instance the later Fig. 2-2).

2.4.2 Lattice effects on Doppler broadening

Due to the presence of the atomic lattice, the speeds v of atoms in a solid are not truly distributed according to a Maxwell-Boltzmann distribution at temperature T . To good approximation, however, a Maxwell-Boltzmann distribution corrected using only an effective temperature T_{eff} can be used, where

$$T_{\text{eff}} = 3T \left(\frac{T}{\theta_D} \right)^3 \int_0^{\theta_D/T} x^3 \left(\frac{1}{e^x - 1} + \frac{1}{2} \right) dx \quad (2.38)$$

and θ_D is the Debye temperature of the material. For uranium ($\theta_D = 210$ K [28]) at $T = 300$ K, for instance, $T_{\text{eff}} = 307$ K, a small but potentially significant correction

¹⁵The precision of nuclear energy levels and emitted photon energies reported in the ENSDF databases is typically 100 eV or larger, which is greater than the recoil correction $\Delta E \sim 10$ eV. The level energies E_r and the recoil-corrected emitted photon energies (for direct decay to ground state) $E_r - \Delta E$ are thus equal within the reported precision. It is therefore ambiguous whether the given energy should be taken as correct for the level or for the emitted photon. If the former convention is chosen, the emitted photon energy must be decremented by ΔE ; if the latter is chosen, the level energy must be incremented by ΔE . In the G4NRF simulation code used in this work (see Section 3.1), the latter convention used by the code’s original authors is maintained for computational convenience.

¹⁶In theory, this energy downshift could be compensated for by placing the emitter on a rapidly-rotating centrifuge in order to blueshift the NRF photons back up to E_r [25].

given the presence of T in the exponential terms of the cross sections.

Thus, using Eqs. 2.21, 2.37, and 2.38, this formalism describes the following sequence:

1. Doppler-broadened NRF absorption of a photon with energy E ,
2. excitation to a resonant energy state E_r ,
3. emission of a photon of energy $E_r - E_j - \Delta E$,

with the cross section

$$\sigma_{0,r,j}^D(E) = 2\pi^{1/2} g_r \left(\frac{\hbar c}{E_r} \right)^2 \frac{\Gamma_{0,r} \Gamma_{r,j}}{\Gamma_r} \frac{1}{\Delta_{\text{eff}}} \int_{-\infty}^{+\infty} \exp \left[-\frac{(x-y)^2}{4t_{\text{eff}}} \right] \frac{dy}{1+y^2}, \quad (2.39)$$

or in an equivalent but perhaps more useful form,

$$\sigma_{0,r,j}^D(E) = 2\pi^{1/2} g_r \left(\frac{\hbar c}{E_r} \right)^2 \frac{b_{0,r} b_{r,j}}{\sqrt{t_{\text{eff}}}} \int_{-\infty}^{+\infty} \exp \left[-\frac{(x-y)^2}{4t_{\text{eff}}} \right] \frac{dy}{1+y^2}, \quad (2.40)$$

where Δ_{eff} and t_{eff} are the values of Δ and t at T_{eff} . In practice, this integral must be evaluated numerically. Since the integrand is still very roughly Gaussian, > 99.99% accuracy can be obtained by integrating between ± 4 standard deviations, i.e. $\pm 4\sqrt{2t_{\text{eff}}}$.

Useful simplifications can be made in the case of an incident photon that is exactly on-resonance. Then $x = 0$ and the integral evaluates to

$$\sigma_{0,r,j}^D(E_r) = 2\pi^{1/2} g_r \left(\frac{\hbar c}{E_r} \right)^2 \frac{\Gamma_{0,r} \Gamma_{r,j}}{\Gamma_r^2} \frac{1}{\sqrt{t_{\text{eff}}}} \pi \exp \left(\frac{1}{4t_{\text{eff}}} \right) \text{erfc} \left(\frac{1}{2\sqrt{t_{\text{eff}}}} \right) \quad (2.41)$$

$$= 2\pi^{3/2} g_r \left(\frac{\hbar c}{E_r} \right)^2 \frac{\Gamma_{0,r} \Gamma_{r,j}}{\Gamma_r} \frac{1}{\Delta_{\text{eff}}}, \quad (2.42)$$

similar to Eq. 2.23, where the exp and erfc terms are approximately 1 for $t_{\text{eff}} \equiv (\Delta_{\text{eff}}/\Gamma_r)^2 \gg 1$. This maximum cross section can also be related to the integrated cross section by Eq. 2.26:

$$\int_0^\infty \sigma_{0,r,j}^D(E) dE = \pi^{1/2} \Delta_{\text{eff}} \sigma_{0,r,j}^D(E_r). \quad (2.43)$$

2.5 Worked example: uranium-238

Uranium-238 has a prominent NRF line at $E_r \simeq 2.176$ MeV. The resonance width is $\Gamma_r = 0.058$ eV, and the decay can be directly to the ground state or to the 45 keV intermediate state with relative probability 52% [29]. Therefore the branching ratios for the two decays are $\{1.00, 0.52\}/(1.00 + 0.52)$ or $\{0.657, 0.342\}$. The angular momenta are $J_0 = 0$ and $J_r = 1$, giving $g_r = 3/2$. The mass of a uranium-238 atom is 2.217×10^5 MeV/ c^2 . With a Debye temperature of $\theta_D = 210$ K, $T_{\text{eff}} = 307$ K, and therefore $\Delta_{\text{eff}} = 1.06$ eV and $t_{\text{eff}} = 336.3$. Using this data, the pure NRF and Doppler-broadened decay cross sections (Eqs. 2.5 and 2.40) are shown in Fig. 2-2.




Figure 2-2: Pure Breit-Wigner and Doppler-broadened cross sections for absorption by the $E_r = 2.176$ MeV U-238 resonance followed by decay to the ground state. The Doppler-broadened cross section peaks at approximately 32 b and has a width Δ_{eff} of around 1 eV.

2.6 Angular dependence

The angular distribution of the outgoing photons, heretofore treated as isotropic in 4π , is described by the angular correlation function $W(J_0, J_r, J_j, \theta)$. For a given sequence of initial, intermediate, and final level spins $\{J_0, J_r, J_j\}$, this function gives the probability of emitting an NRF photon at polar angle $0 \leq \theta \leq \pi$ relative to the incident photon direction (and isotropic in azimuthal angle $0 \leq \phi \leq 2\pi$) as a sum of Legendre polynomials. By convention, W is normalized such that

$$\int_0^{2\pi} \int_0^\pi W(J_0, J_r, J_j, \theta) \sin \theta \, d\theta d\phi = 4\pi. \quad (2.44)$$

For common spin sequences $\{J_0, J_r, J_j\}$, W takes the following simple forms:

$$0 \rightarrow 1 \rightarrow 0 : \quad W(\theta) = \frac{3}{4}(1 + \cos^2 \theta) \quad (2.45)$$

$$0 \rightarrow 2 \rightarrow 0 : \quad W(\theta) = \frac{5}{4}(1 - 3 \cos^2 \theta + 4 \cos^4 \theta) \quad (2.46)$$

but becomes more complicated for spin sequences¹⁷ that exhibit mixing of multipole excitation modes [30]. For example, the $1 \rightarrow 2$ transition in the sequence $0 \rightarrow 1 \rightarrow 2$ may occur by a mixture of dipole and quadrupole transitions, in which case the equations [31]

$$0 \rightarrow 1 \rightarrow 2, \text{ dipole} - \text{dipole} : \quad W(\theta) = \frac{9}{8} \left(1 - \frac{1}{3} \cos^2 \theta \right) \quad (2.47)$$

$$0 \rightarrow 1 \rightarrow 2, \text{ dipole} - \text{quadrupole} : \quad W(\theta) = \frac{7}{8} \left(1 + \frac{3}{7} \cos^2 \theta \right) \quad (2.48)$$

describe the angular correlations for the pure dipole-dipole and pure dipole-quadrupole transitions, respectively. The amount of mixing between these two modes is given by the multipole mixing ratio δ [32, 33]. The angular correlations $W(\theta)$ for the above four spin sequences are shown in Fig. 2-3. In effect, the angular correlation function

¹⁷In particular, sequences involving transitions between two non-zero spin states.

$W(\theta)$ modifies the NRF emission (Eq. 2.21) to include a separable θ -dependence,

$$\sigma_{0,r,j}^D(E, \theta) = W(\theta) \cdot \sigma_{0,r,j}^D(E), \quad (2.49)$$

which will often cancel in relative calculations and is not used in the rest of this analytical treatment, though is included (when possible) in the Monte Carlo simulations.




Figure 2-3: Angular correlations $W(J_0, J_r, J_j, \theta)$ for common spin sequences. The pure dipole-dipole transition is denoted “dd” and the pure dipole-quadrupole transition by “dq”.

2.7 NRF interaction rate in a measurement

Having derived the NRF cross sections, it is now possible to calculate the evolution of a photon spectrum through the ZK system. Here we make use of Fig. 2-4 (a simplified version of the geometry of Fig. 1-1), in which the weapon and the foil are both homogeneous plates of the same single isotope.

The interrogation spectrum $\phi_0(E)$ (in units of photons/MeV), here a bremsstrahlung source, impinges on the weapon of thickness¹⁸ D . The flux transmitted $\phi_t(E)$ through the weapon is simply a result of both resonant and non-resonant attenuation,

¹⁸ D may be defined as a thickness or an areal density, depending on convention. See Footnote 26.




Figure 2-4: Simplified schematic of Fig. 1-1 used to derive Eq. 2.52, the photon energy spectrum $d^2n/d\Omega dE$ observed by the detector.

since the NRF interactions in the weapon cause re-emission into 4π and effectively attenuate the forward-going flux:

$$\phi_t(E) = \phi_0(E) \exp[-D(\mu'_{\text{NRF}}(E) + \mu_{\text{nr}}(E))], \quad (2.50)$$

where the prime denotes the absorption attenuation coefficient¹⁹ rather than the decay/backscatter attenuation coefficient. Because $\mu'_{\text{NRF}}(E)$ is sharply energy-dependent with Doppler-broadened width $\Delta \sim 1$ eV, the NRF interactions in the weapon attenuate only narrow energy bands corresponding to NRF absorption lines, carving out a series of ‘notches’ in the spectrum. Note that if no weapon is present, $D = 0$ and $\phi_t(E) = \phi_0(E)$. The transmitted flux $\phi_t(E)$ then impinges on the reference foil of thickness X . Signal photons must undergo NRF in the foil to be observed; these will traverse a distance x before undergoing NRF with some emission angle θ and

¹⁹In general, the attenuation coefficient μ for a process is defined as $\mu \equiv N\sigma$ where $N = \rho N_A/A$ is the number density of target particles that may induce the process and σ is the microscopic cross section for the process. In this section, the primed $\mu'_{\text{NRF}}(E)$ is used to denote $N\sigma_{0,r}^D(E)$, while the unprimed $\mu_{\text{NRF}}(E)$ denotes $N\sigma_{0,r,j}^D(E)$. See also Footnote 26.

decay mass attenuation coefficient $\mu_{\text{NRF}}(E)$. The differential equation governing the double differential NRF spectrum $d^2n/d\Omega dE$ resulting from the interactions in an infinitesimal dx is therefore

$$\frac{d^3n}{d\Omega dE} = \phi_t(E) \mu_{\text{NRF}} \exp \left[-x \left(\mu_{\text{NRF}} + \mu_{\text{nr}} \frac{1 + \cos \theta}{\cos \theta} \right) \right] dx \quad (2.51)$$

where $\theta < \pi/2$ is the angle of NRF emission relative to the (backwards) beam direction, and the E -dependence of the μ terms has been suppressed for clarity. The first term of the exponential describes the resonant interaction after traversing a distance x , while the second term describes the non-resonant attenuation from traversing x plus the $x/\cos \theta$ the NRF photon must pass through to exit the foil. Integrating across the full extent of the foil from $x = 0$ to $x = X$, this has the solution²⁰

$$\frac{d^2n}{d\Omega dE} = \phi_t(E) \frac{\mu_{\text{NRF}}}{\mu_{\text{NRF}} + \mu_{\text{nr}} \frac{1 + \cos \theta}{\cos \theta}} \left\{ 1 - \exp \left[-X \left(\mu_{\text{NRF}} + \mu_{\text{nr}} \frac{1 + \cos \theta}{\cos \theta} \right) \right] \right\}, \quad \theta < \pi/2. \quad (2.52)$$

It is useful to consider the limiting case of Eq. 2.52 for thin reference foils and strong NRF cross sections ($\mu_{\text{NRF}} \gg \mu_{\text{nr}}$),

$$\frac{d^2n}{d\Omega dE} \simeq \phi_t(E) \mu_{\text{NRF}}(E) X, \quad (2.53)$$

which is linear in both the cross section (via μ_{NRF}) and the foil thickness X .

The NRF spectrum dn/dE in an angular region $\Omega = \{\theta_1 \leq \theta \leq \theta_2, \phi_1 \leq \phi \leq \phi_2\}$ is then²¹

$$\frac{dn}{dE} = \frac{1}{4\pi} \int_{\Omega} \frac{d^2n}{d\Omega dE} d\Omega = \frac{\Delta\phi}{4\pi} \int_{\theta_1}^{\theta_2} \frac{d^2n}{d\Omega dE} d \cos \theta, \quad (2.54)$$

²⁰Equation 2.52 describes the flux *incident* on a detector at angle θ . The flux *measured* by the detector is a function of both the incident flux and the detector response (especially the intrinsic efficiency ϵ_{int} of the detector).

²¹Though the angular correlation function $W(\theta)$ of Section 2.6 is left out here, the $1/4\pi$ factor by which it modifies the integration of the angular differential cross section in Eq. 2.54 must be included for absolute normalization. For full generality, this extra factor of $W(\theta)/4\pi$ should instead be inserted into Eq. 2.51.

where $\Delta\phi \equiv \phi_2 - \phi_1$, making use of the fact that $d\Omega \equiv d\cos\theta d\phi$ and that Eq. 2.52 does not depend on ϕ . There is no analytical solution to this integral, but for small intervals $\Delta\cos\theta \equiv \cos\theta_2 - \cos\theta_1$ over which the double differential spectrum does not vary significantly, the formula simplifies further to

$$\frac{dn}{dE} \simeq \frac{\Delta\phi}{4\pi} |\Delta\cos\theta| \left(\frac{d^2n}{d\Omega dE} \right)_{\theta=\theta_C}, \quad (2.55)$$

where $\theta_C \equiv (\theta_1 + \theta_2)/2$ is the central θ of the region.

In the verification protocol, the host wishes to keep the weapon thickness D unknown to the inspector. To prevent estimation of D by Eq. 2.52, the host must also protect the foil thickness X . The interrogation spectrum $\phi_0(E)$ may also be kept classified for additional secrecy.

2.8 Notch refill

The derivation leading up to Eq. 2.52 has considered the weapon to be a pure attenuator, and has therefore neglected the possibility of any source terms in the weapon. However, bremsstrahlung photons above the resonance energy E_r may undergo small-angle Compton scattering into the resonance, increasing the NRF interaction rate in the target. When this occurs in the weapon of thickness D , it in effect reduces the resonant attenuation by the weapon, refilling the notch in the transmitted spectrum.

The effects of notch refill can be quantified by considering the ratio of on-resonance flux (denoted $\phi(E_r)$ in this section, rather than the dn/dE of Section 2.7) accounting for notch refill to the same quantity ignoring notch refill. With this definition, notch refill introduces a single multiplicative correction factor

$$C_{\text{notch}} \equiv \frac{\phi_{\text{refill}}(E_r, \ell)}{\phi_{\text{no refill}}(E_r, \ell)} \geq 1. \quad (2.56)$$

The magnitude of this term can be estimated by examining the evolution of the on-resonance component $\phi(E_r, \ell)$ of the bremsstrahlung energy spectrum $\phi(E, \ell)$ that traverses a linear distance ℓ through a single homogeneous target (such as the warhead

to be interrogated). Accounting for notch refill, the differential equation governing $\phi(E_r, \ell)$ is

$$\frac{d\phi(E_r, \ell)}{d\ell} = -\phi(E_r, \ell)[\mu_C + \mu'_{\text{NRF}}] + \frac{2}{3}\mu_C \int_{E_r}^{E_{\text{max}}} \phi(E, \ell) \frac{dE}{E} \quad (2.57)$$

where the E -dependence of the μ terms has been suppressed for clarity, and E_{max} is the maximum photon energy that can scatter into the resonant region (derived below). The first term describes scattering out of the resonance by Compton and NRF processes with linear attenuation coefficients μ_C and μ_{NRF} , respectively. The second term describes the scattering into the resonance (notch refill), approximating the differential cross section for Compton scattering from $E \rightarrow E'$ as [34]²²

²² The authors in [34] do not cite or justify this approximation, but it can be derived from the kinematic [35] and Klein-Nishina [36] formulae for Compton scattering through angle θ from photon energy E to E' :

$$f(E, \theta) \equiv \frac{E'}{E} = \frac{1}{1 + E/m_e c^2 (1 - \cos \theta)}, \quad (2.58)$$

$$\frac{d\sigma}{d\Omega} = \frac{1}{2}\alpha^2 r_C^2 f^2(E, \theta) [f(E, \theta) + f^{-1}(E, \theta) - 1 + \cos^2 \theta], \quad (2.59)$$

where $\alpha \simeq 1/137$ is the fine structure constant and $r_C = \hbar/m_e c$ is the reduced Compton electron wavelength. Then

$$\frac{d\sigma}{dE'} = \frac{d\sigma}{d\theta} \left| \frac{d\theta}{dE'} \right| = \frac{d\sigma}{d\theta} \left| \frac{dE'}{d\theta} \right|^{-1} \quad \text{with} \quad \frac{dE'}{d\theta} = -\frac{E^2 \sin \theta / m_e c^2}{[1 + (1 - \cos \theta)E/m_e c^2]^2}, \quad (2.60)$$

and thus

$$\frac{d\sigma}{dE'} = 2\pi \frac{d\sigma}{d\Omega} \frac{[1 + (1 - \cos \theta)E/m_e c^2]^2}{E^2 \sin \theta / m_e c^2} \quad (2.61)$$

due to isotropy in azimuthal angle ϕ . Averaging over solid angle,

$$\left\langle \frac{d\sigma}{dE'} \right\rangle = 2\pi \int_0^\pi \frac{d\sigma}{d\Omega} \frac{[1 + (1 - \cos \theta)E/m_e c^2]^2}{E^2 / m_e c^2} d\theta \bigg/ \int_0^{2\pi} \int_0^\pi \sin \theta d\theta d\phi, \quad (2.62)$$

where the denominator is just 4π . Similarly, the total Compton cross section is

$$\sigma_C(E) = \int_0^{2\pi} \int_0^\pi \frac{d\sigma}{d\Omega} \sin \theta d\theta d\phi. \quad (2.63)$$

Defining the quantity

$$R(E) \equiv \frac{E}{\sigma_C(E)} \left\langle \frac{d\sigma}{dE'} \right\rangle, \quad (2.64)$$

it is apparent from Fig. 2-5 that to a very rough approximation, $R(E) \simeq 2/3$ over an appropriate




Figure 2-5: Plot of $R(E)$, a normalized ratio of average differential to total Compton cross sections, vs photon energy E . Over a suitable energy range, the value of $R(E)$ is very roughly $2/3$. See Footnote 22.

$$\frac{d\sigma_{\text{C}}(E \rightarrow E')}{dE'} \simeq \frac{2}{3} \frac{\sigma_{\text{C}}(E)}{E}. \quad (2.65)$$

The μ_{C} can be taken outside of the integral in Eq. 2.57, since $\mu_{\text{C}}(E)$ is nearly constant over the energy range $[E_r, E_{\text{max}}]$. Then the integral can be approximated as follows, writing $E_{\text{max}} \equiv E_r + \Delta E$ (the ΔE here differing from the recoil ΔE of Eq. 2.37):

$$\int_{E_r}^{E_r + \Delta E} \frac{\phi(E, \ell)}{E} dE \simeq \frac{\phi(E_r, \ell)}{E_r} \Delta E, \quad (2.66)$$

since $\phi(E, \ell)$ is also approximately constant over the (small) integrated energy range for a bremsstrahlung source. Note that E_{max} is limited by geometry and not the bremsstrahlung endpoint: since the Compton kinematic formula is one-to-one in angle and energy, only photon trajectories within a certain range of angles and therefore

energy range as claimed in Eq. 12 of the reference.

within a certain range of initial energies can Compton scatter into the notch. Then

$$\frac{d\phi(E_r, \ell)}{d\ell} = \left[-(\mu_C + \mu'_{\text{NRF}}) + \frac{2}{3}\mu_C \frac{\Delta E}{E_r} \right] \phi(E_r, \ell), \quad (2.67)$$

which is a simple first order ODE with solution

$$\phi(E_r, \ell) = \phi(E_r, 0) \exp \left\{ \left[-(\mu_C + \mu'_{\text{NRF}}) + \frac{2}{3}\mu_C \frac{\Delta E}{E_r} \right] \ell \right\}. \quad (2.68)$$

If notch refill had been neglected in Eq. 2.57, only the first term of this exponential would have resulted. Thus the notch refill factor from Eq. 2.56 can be written as

$$C_{\text{notch}} \equiv \frac{\phi_{\text{refill}}(E_r, \ell)}{\phi_{\text{no refill}}(E_r, \ell)} \simeq \exp \left(\frac{2}{3}\mu_C \frac{\Delta E}{E_r} \ell \right). \quad (2.69)$$

It then remains to estimate ΔE based on geometric constraints. Using the Compton scattering formula and approximating $1 - \cos \theta \simeq \theta^2/2$ for small θ , photons of maximum energy E will downscatter to E_r if

$$E_r \simeq \frac{E}{1 + E\theta_{\text{max}}^2/2m_e c^2}. \quad (2.70)$$

Then

$$\Delta E \equiv E - E_r \simeq \frac{E_r^2}{m_e c^2} \frac{\theta_{\text{max}}^2}{2}, \quad (2.71)$$

and applying $\mu = N\sigma$, where N is the target number density, Eq. 2.69 becomes

$$C_{\text{notch}}(\ell) \simeq \exp \left(\frac{2}{3}N\sigma_C(E_r) \frac{E_r}{m_e c^2} \frac{\theta_{\text{max}}^2}{2} \ell \right), \quad (2.72)$$

giving the final equation for the notch refill factor $C_{\text{notch}}(\ell)$. For a U-238 target of areal density 90 g/cm² (i.e. thickness $\ell = 4.7$ cm) and angular constraint $\theta_{\text{max}} = 0.1$, the $E_r = 2.176$ MeV resonance (with $\sigma_C(E_r) = 12.8$ b/atom [37]) will exhibit a notch refill factor of $C_{\text{notch}}(\ell) \simeq 1.04$. For targets of this approximate areal density, then, notch refill effects can be ignored to reasonable accuracy.

2.9 Secondary bremsstrahlung background

Eq. 2.52 predicts the NRF peak height, but neglects any secondary processes that may contribute to the background on top of which the peaks will sit. In particular, both on- and off-resonance photons striking the foil will undergo photoelectric absorption, Compton scattering, and pair production. The charged secondaries liberated or created by these processes may then undergo bremsstrahlung in the foil, emitting photons in a continuum of energies that will be seen by the detector. This is known as the secondary bremsstrahlung or EZ3DTM [38] continuum.

In general, the number of continuum photons increases with the average Z of the foil and decreases with increasing photon energy E to almost zero continuum near the endpoint E_{end} (in this work, $E_{\text{end}} = 2.7$ MeV) of the bremsstrahlung energy spectrum. This energy-dependence can be seen in part by examining the maximum energy thresholds E_{th} with which the secondary bremsstrahlung can be produced. Assuming a charged secondary can lose all its kinetic energy to bremsstrahlung, the maximum bremsstrahlung photon energies are exactly determined by the maximum possible energy transfer to the secondaries in each of the three above processes.

In the case of photoelectric absorption, nearly all the incident photon energy can be transferred to the photoelectron (minus the small atomic binding energy E_b of the photoelectron):

$$E_{\text{th,PA}} = E_{\text{end}} - E_b, \quad (2.73)$$

which is very nearly 2.7 MeV for $E_{\text{end}} = 2.7$ MeV. For Compton scattering, the maximum electron recoil energy occurs for photon scattering at $\theta = \pi$, in which case

$$E_{\text{th,CS}} = E_{\text{end}} \frac{2E_{\text{end}}/m_e c^2}{1 + 2E_{\text{end}}/m_e c^2}, \quad (2.74)$$

(now neglecting the electron binding energy), which evaluates to 2.47 MeV for $E_{\text{end}} = 2.7$ MeV. Note that the energy of a *direct* Compton scatter into a detector at $\theta = 3\pi/4$

is 0.269 MeV, which is well below the 1.7–2.7 MeV signal region used in this work. Finally, for pair production, the maximum kinetic energy transferred to the e^+e^- pair is $E_{\text{end}} - 2m_e c^2$, giving a threshold of

$$E_{\text{th,PP}} \simeq E_{\text{end}} - 2m_e c^2, \quad (2.75)$$

if one of the particles is emitted along the incident photon direction and thus receives the entire kinetic energy in the lab frame²³. This evaluates to 1.678 MeV for $E_{\text{end}} = 2.7$ MeV, but is often exceeded if this energy is transferred to an e^+ that annihilates inflight with an e^- in the foil.

Although the continuum extends very close up to the endpoint due to the contribution from photoelectric absorption, the combined contribution from both photoelectric absorption and Compton scattering is small compared to the NRF lines for energies ~ 500 keV below the endpoint E_{end} (see [30, Fig. 2.24] or Fig. 2-6, which shows the breakdown of the full spectrum from a foil-only geometry into its constituent parts based on the tag given to each photon as defined in Table 2.1). This gives rise to an energy window between $\sim E_{\text{end}} - 500$ keV and E_{end} in which there is very little continuum²⁴ and the NRF peaks are clearly visible. Below this energy window, the continuum increases, eventually drowning out the NRF peaks. The NRF peaks in a measurement must therefore lie within this energy window in order to obtain a statistically-significant signal. Future work will investigate if the continuum can be treated not as a background obscuring the NRF peaks, but as a signal itself—and if so, whether the continuum is vulnerable to sensitive information leakage.

²³This neglects the small nuclear recoil required for simultaneous conservation of energy and momentum. Additionally, the positron typically has a slightly higher kinetic energy than its electron partner due to its Coulomb repulsion from the positively-charged nucleus.

²⁴The standard Geant4 electromagnetic physics lists used here do not include Delbrück, nuclear Thomson, and Rayleigh scattering processes that will also contribute to the continuum in this region. The first two contribute only slightly to the continuum, while Rayleigh scattering primarily contributes very close to the endpoint E_{end} [30].

photon tag	criterion
photoelectric brems:	$(\neg \text{isNRF}) \cap (e^- \text{ from single photoelectric event})$
Compton brems:	$(\neg \text{isNRF}) \cap (e^- \text{ from single Compton event})$
pair brems/annihil:	$(\neg \text{isNRF}) \cap (e^\pm \text{ from single pair event})$
multiple events:	$(\neg \text{isNRF}) \cap (e^\pm \text{ from multiple events})$
high-sampled NRF:	$(\text{isNRF}) \cap (\text{primary } \gamma \text{ was highly-sampled})$
low-sampled NRF:	$(\text{isNRF}) \cap (\neg \text{primary } \gamma \text{ was highly-sampled})$

Table 2.1: Definitions of photon tags. The “isNRF” criterion corresponds to whether the photon was produced by an NRF event. The criteria logically add to unity assuming photoelectric, Compton, and pair bremsstrahlung events are the only possible non-NRF interactions that can result in a backwards photon.

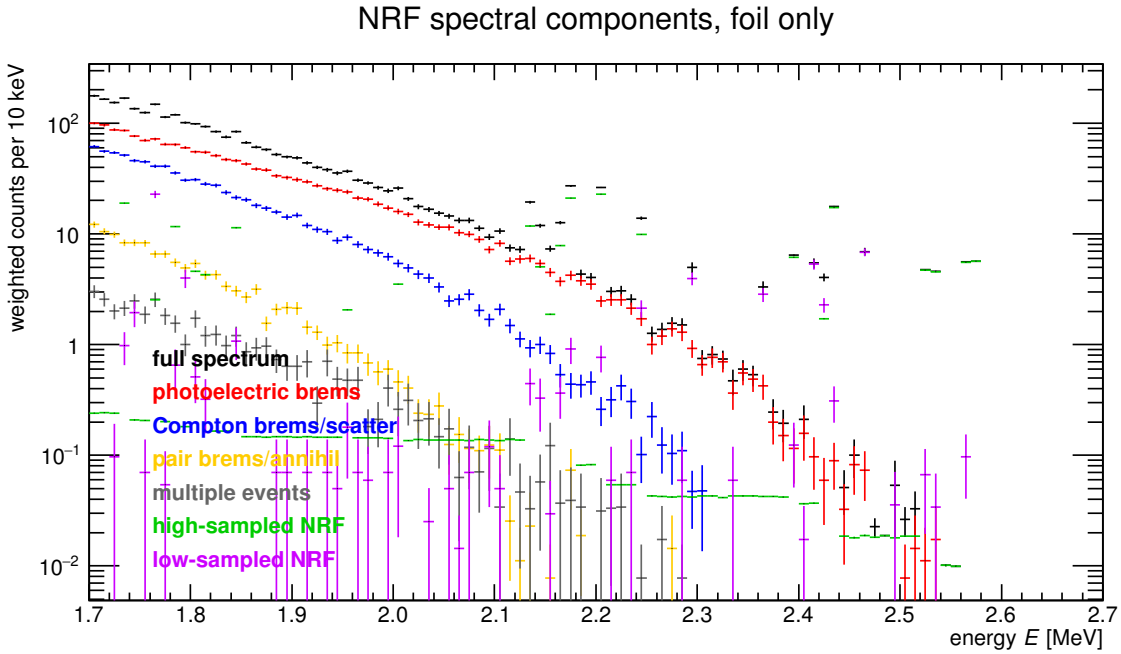


Figure 2-6: Breakdown of full photon spectrum (black) into its photoelectric secondary bremsstrahlung (red), Compton scattering secondary bremsstrahlung (blue), pair lepton secondary bremsstrahlung (yellow), high-sampled NRF (green), and low-sampled NRF (violet) components for endpoint $E_{\text{end}} = 2.7$ MeV. Secondary electrons undergoing multiple interactions are shown in grey. The low-sampled NRF contribution (violet) to the NRF lines with small vertical error bars indicates a small amount of notch refill occurring in the foil itself, while the violet contribution to the NRF lines with larger vertical error bars indicates an NRF line that was not deliberately oversampled. Note that the bin at 2.42 MeV is an overlap of the highly-sampled Pu-239 decay at 2.423 MeV and the low-sampled U-238 decay line at 2.423 MeV arising from the 2.468 MeV resonance.

Chapter 3

Monte Carlo simulations

As a proof of concept, several Monte Carlo simulations of nuclear resonance fluorescence interactions in nominal warhead and hoax object geometries are carried out using the Geant4 toolkit [39]. NRF interactions are modeled using an updated version of the `G4NRF` package [40] alongside the standard electromagnetic physics models. Each simulation was run for $\sim 20\,000$ core-hours on the Amazon EC2 (Elastic Compute Cloud) infrastructure, comprising nearly 10^{12} individual photon histories.

3.1 The `G4NRF` package

By default, the photon interactions modelled in Geant4 include photoelectric absorption, Compton scattering, e^+e^- pair production, and other processes at much higher energies such as $\mu^+\mu^-$ pair production [41]. The first three processes are major sources of background in NRF measurements via e^\pm bremsstrahlung, as discussed in Section 2.9. NRF and other collective nuclear excitations, however, are not included in the standard distribution of Geant4²⁵. As a toolkit, rather than a full set of software, though, Geant4 allows the user to write custom code for various parts of the simulation including experiment geometry, data I/O, and physics processes. The `G4NRF` package was thus developed by David Jordan and Glen Warren to sim-

²⁵The decay physics of excited nuclei is available in the standard distribution, but the resonant absorption step required for NRF is not. Atomic or x-ray resonance fluorescence (XRF) is included in the standard Geant4 distribution.

ulate NRF physics in Geant4 [40]. Since Geant4 is object-oriented and thus highly modular, the inclusion of `G4NRF` is relatively simple—only ten or so lines of code are required to interface the package with standard Geant4 code.

All physics processes in Geant4, whether standard or custom, and regardless of internal complexity, must ultimately compute 1) a mean free path of interaction μ (typically computed as $\mu = N\sigma$ from a material number density N and a microscopic cross section σ) and 2) the final state kinematics for all daughter particles resulting from the interaction. For an NRF interaction, this amounts to evaluating the Doppler-broadened NRF absorption cross section $\sigma_{0,r}^D(E)$ of Eq. 2.30 (using nuclear level data from an ENSDF-derived database), and determining the energy and angle of the emitted photon by sampling the possible decay channels of the excited state. These calculations are described in more depth in the following sections.

3.1.1 Nuclear data organization

Nuclear data in `G4NRF` is organized primarily by isotope, and secondarily by individual level energy. Data common to an isotope, such as Z , A , ground state spin and parity J_0^π , Debye temperature θ_D , etc. is stored in a `G4NRFNuclearLevelManager` object corresponding to that particular isotope. These objects further include the level scheme for each isotope in the form of a vector of `G4NRFNuclearLevel` objects, which contain data specific to the level as identified by its energy E_r , including spin and parity J_r^π , width Γ_r , and branching ratios b_{rj} . As such, the collection of `G4NRFNuclearLevel` objects for each isotope is managed by that isotope’s `G4NRFNuclearLevelManager`; the collection of `G4NRFNuclearLevelManager` objects is in turn managed by a `G4NRFNuclearLevelStore` object, which is a global singleton that can be accessed at any point in the simulation. This hierarchy is shown in Fig. 3-1.

In the initialization phase, the `G4NRFNuclearLevelManager` and `G4NRFNuclearLevel` objects are constructed and grouped appropriately by `G4NRFNuclearLevelStore`. Their data members are then populated from five separate databases:

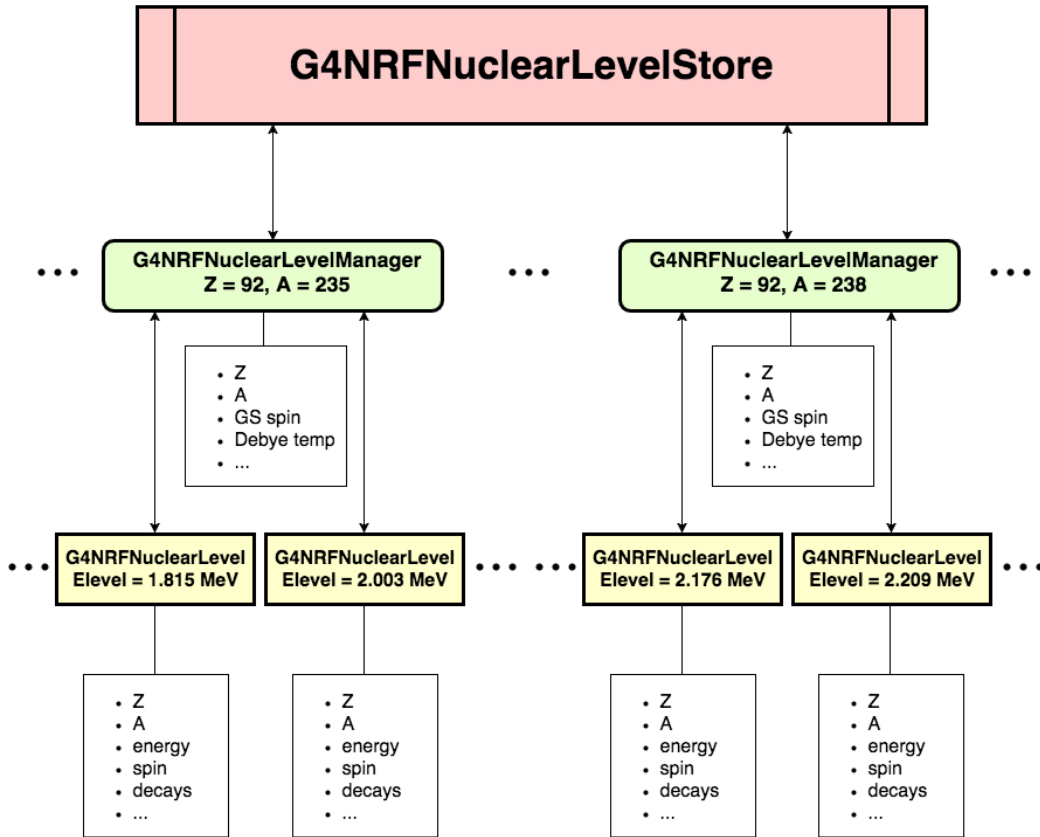


Figure 3-1: Hierarchy of G4NRF nuclear data organization. Each isotope is represented by a `G4NRFNuclearLevelManager` object, which contains isotope-relevant data and the vector of `G4NRFNuclearLevel` objects that represents the isotope’s nuclear level scheme. Each `G4NRFNuclearLevel` contains further level information including the available decays and their respective branching ratios. The `G4NRFNuclearLevelManager` objects are stored in a singleton `G4NRFNuclearLevelManagerStore` object.

1. `zZZ.aAAA` (`ZZ` and `AAA` denoting the set of nucleus (Z, A) values), which contains information on each level such as spin, parity, and width;
2. `level_table_AAA.dat`, which again contains information on each level such as spin, parity, and width;
3. `gamma_table_AAA.dat`, which contains multipolarity information for each photon;
4. `ground_state_properties.dat`, which contains the ground state spin and parity for each isotope;

5. `TDebye_data.dat`, which contains the Debye temperatures listed in both Kittel [42] and Ashcroft and Mermin [28] for each element.

While the nuclear level energy typically serves as a level’s unique identifier (given a particular Z and A), the uncertain nature of experimental data means that there are slight mismatches in level (and photon) energies across the first three (ENSDF-derived) databases. Matching data to its proper level during the population of the data members is therefore a delicate task. `G4NRF` admits a tolerance for energy mismatch, but it is impossible to choose a single energy tolerance that is strict enough to prevent inadvertent overlap but not so strict so as to fail to match two database entries that truly correspond to the same isotope, especially since the level spacing and the uncertainties in energy vary widely across the approximately 140 000 levels in the database. In practice, a compromise can be made by finding a fixed tolerance that retains most of the levels relevant to a simulation without introducing too many incorrect merges of nuclear data in other levels. In some cases it may be necessary to selectively adjust the databases to ensure crucial levels are not thus discarded. A level matching tolerance of 10 eV is used in this work, and any level that cannot be matched to this precision is deleted.

Updates to the original `G4NRF` code introduce secondary checks to ensure consistency after enforcing this 10 eV tolerance. Another useful identifier of a level is its spin; if two databases propose different spins for the same level (whether by truly different reported data or an incorrect merge due to the 10 eV tolerance), the level is deleted. Furthermore, if the level spin or width is empirically-unknown, the level is deleted. This last check, however, deletes most of the levels of U-235, Pu-239, and Pu-240 crucial to this work, and some approximate values must be entered by hand into the databases to ensure their retention—see Section 3.4.

The available Debye temperature data is also somewhat imprecise. For many elements, the Debye temperature θ_D —used in the numerical integration for the effective temperature T_{eff} (Eq. 2.38)—may only be known at low external temperatures T , may differ significantly between the two references consulted, or may not be known at all. In cases where at least some data is available, preference is given to the non-

low-temperature-determinations; if neither θ_D is singled out by this step, a simple average of the two values is used to calculate T_{eff} at initialization. If no data is available, T_{eff} is simply set to the real temperature of $T = 300$ K. The cross sections of the actinide isotopes considered in this work are not very sensitive to inaccuracies in θ_D , since these isotopes tend to have $\theta < 300$ K or equivalently by Eq. 2.38, $1 < T_{\text{eff}}/T < 1.05$, providing only a small correction factor to the real temperature of $T = 300$ K. In addition, the inclusion of the Debye temperature machinery allows for the later implementation of more precise values of θ_D should they become known.

3.1.2 Cross section evaluation

Inherent to Monte Carlo simulations is the evaluation of interaction cross sections and conversion from cross section to interaction probability at each step of the tracked particle. For a step size of Δx , the probability of undergoing any interaction at all is

$$P = 1 - \exp(-\mu\Delta x), \quad (3.1)$$

where $\mu \equiv \mu(E)$ is the total linear attenuation coefficient²⁶ i.e. the sum of the individual attenuation coefficients μ_k for each process k :

$$\mu = \sum_k \mu_k = \mu_{\text{NRF}} + \mu_{\text{Compton}} + \mu_{\text{pair prod.}} + \dots \quad (3.2)$$

Given that an interaction occurs in Δx , the probability that it is an interaction of type k is μ_k/μ .

In the updated version of G4NRF, the NRF absorption cross sections $\sigma_{0,r}^D$ are calculated at each step for each possible resonance r using Eq. 2.39 (implemented in `G4NRF::NRF_xsec_calc()`), then converted to a mean free path (i.e. the inverse of the linear attenuation coefficient $\mu_{0,r}^D$) and then to an interaction probability. Of par-

²⁶Strictly speaking, μ is also proportional to the material density ρ . It is often useful, therefore, to divide out the ρ -dependence and work instead with the *mass* attenuation coefficient: $\mu \rightarrow \mu/\rho$, which has units of cm^2/g . In this case, $\Delta x \rightarrow \rho\Delta x$ is called the *areal density* and has units g/cm^2 . In practice, the ρ is often left implicit, and it is up to the reader to determine whether a given μ is a linear or mass attenuation coefficient.

ticular note is that the full numerical integration is performed (using Simpson’s rule) at each step (and for each resonance of each isotope in the current volume). To some extent, this is the rate-limiting process in the simulation: the numerical integration in `NRF_xsec_calc()` takes up $\sim 20\%$ of the total simulation time, but is also only $\sim 20\%$ slower than the original Gaussian evaluation using Eq. 2.24. In future work, several different approximation schemes could be implemented to reduce the computational intensity of the cross section evaluation. For instance, since the final NRF spectrum is determined primarily by the integrated cross section $\sigma_{0,r,j}^{\text{int}} \equiv b_{r,j} \int \sigma_{0,r}(E) dE$ for each isotope (which in fact are the experimentally-known values), rather than the full energy-dependent cross section, the cross section could be modelled as a rectangular window centered at E_r with width 1 eV ($\simeq \Delta$ from Eq. 2.11) and height $\sigma_{0,r,j}^{\text{int}}/(1 \text{ eV})$ so that the integrated cross section in both cases is the same but no numerical integration is required. Another approach might be to use the questionable approximation of Eq. 2.24, which also preserves the integrated cross section and provides a Gaussian closed form. Finally, a lookup table of cross section evaluations could be generated ahead of time, and an interpolation routine used at runtime to approximate the value of the integral for incident photon energies that do not correspond to the exact mesh spacing used in the lookup table. Such a lookup table has been implemented as another update to `G4NRF`, but the results have not been validated against the final spectra generated by numerical integration on large-scale, high-statistics runs. A comparison of the accuracy of these methods is made in [30, Fig. 3.25].

3.1.3 Final state sampling

Since the cross section computed by `G4NRF` in `NRF_xsec_calc()` is the NRF absorption cross section, it remains to sample the photon decay of the excited nucleus and determine the energy (and angle) of the outgoing photon, essentially restoring the branching ratio factor $b_{r,j}$ in order to compute the full $\sigma_{0,r,j}^D(E)$. The `PostStepDoIt()` function loops through the list of possible final levels (comprising all possible intermediate levels plus the ground state) and selects one according to the branching

ratios of the corresponding `G4NRFNuclearLevel` object. This cascade may also select with non-zero probability a decay by internal conversion, emitting an atomic electron in place of an NRF photon. For the actinides considered in this work, the cascade is always a single choice between direct decay to ground state or decay through a single intermediate state [29, 43, 44]. More complicated cascades can be found in isotopes such as Fe-56 [45].

The angular distribution of emitted NRF photons is sampled in `G4NRF` using the angular correlation function $W(J_0, J_r, J_j, \theta)$. The polar angle $0 \leq \theta \leq \pi$ of the outgoing photon is thus distributed according to W , while the azimuthal angle $0 \leq \phi \leq 2\pi$ is uniformly distributed. If the level spins are unknown, `G4NRF` defaults to isotropic emission. Furthermore, it is possible to disable the angular correlation calculation altogether in `G4NRF` if simpler isotropic emission is desired.

In this work, the angular region of acceptance used is $3\pi/4 \leq \theta \leq \pi$, and the relevant spin sequences for U-238 are $0 \rightarrow 1 \rightarrow 0$ or $0 \rightarrow 1 \rightarrow 2$; the spins J_r for the U-235, Pu-239, and Pu-240 resonant levels are unknown but can be approximated (see Section 3.4) as $9/2$, $3/2$, and 1 , respectively. Plots of the transitions involving the postulated $J_r = 9/2$ can be found in [30, Fig. 2.44], and for $J_r = 3/2$, [30, Fig. 2.41]. In the specified angular region, the $W(\theta)$ correction for these transitions is typically on the order of 20% (and somewhat higher for the U-238 $0 \rightarrow 1 \rightarrow 0$ transition in Fig. 2-3). These terms can therefore be ignored to first order. The sequence $0 \rightarrow 2 \rightarrow 0$, which has rather large $W(\theta)$ corrections, does not appear in the four isotopes of interest.

3.1.4 Updates to the `G4NRF` code

In the course of this work, several updates were made to the `G4NRF` code, ranging from simple code modernizations to minor bugfixes to major updates:

1. Modernized C++ standards: added missing `#include`'s and `using`'s.
2. Modernized Geant4 standards: included key header files `G4SystemOfUnits.hh` and `G4PhysicalConstants.hh` in most of the source files.

3. Addition of a standalone NRF cross section library functionality. If the user enables this option, a database of the NRF photons known to each instance of `G4NRFNuclearLevelManager` is written to a text file (with a run time of 10–15 minutes for the entire database). A Python script `standalone.py` allows the user to search this library, filtering and sorting the output by ranges of Z , A , photon energy E , level energy E_r , and integrated cross section $\int \sigma_{0,r,j}(E) dE$. This standalone code is described further in Appendix A.
4. Enforcement of strict level deletion criteria as described in Section 3.1.1. The new `GetInvalidLevel()` and `SetInvalidLevel()` public methods of the class `G4NRFNuclearLevel` are implemented in order to access and set a private boolean flag `invalidLevel` that indicates whether the nuclear data for the level is suspect. The `G4NRFNuclearLevel` objects for which `invalidLevel = true` are deleted by the function `delete_bad_levels()` after `G4NRFNuclearLevelManager` reads the databases.
5. Workaround for levels tagged with a “+X” note (which indicates a potential level repetition [46]) in NNDC databases. In the plaintext databases, these level energies are incremented by 2.4 GeV then flipped in sign. If such an energy level of less than $E = -2.4$ GeV is found in the database search, these two steps are undone.
6. Bugfix in the search algorithm for the nearest level energy to a given photon energy. The original location of the `Ediff_min` variable declaration in the `RefreshGammas()` function meant that it was reset on each loop iteration, rather than only when a smaller energy difference was found.
7. Replacement of the cross section calculation function, `NRF_xsec_calc()`. By default, the function now performs the full numerical integration at each step and uses the material-dependent effective temperature T_{eff} (when possible) for Doppler broadening rather than the original 300 K. See Section 3.1.2.
8. Added basic debug functionality in the `G4NRF` class to abort the run upon cal-

ulation of a particular NRF cross section and print diagnostic level and photon information.

9. Replaced the database keypair generation function that would compute the string name of the database(s) required for a given cross section calculation. The original function recalculated the database name, using several slow string concatenations, at every step of the simulation. This was replaced with a vector of database names that is built once at initialization and can be accessed during run time. This update resulted in a speed increase of $\sim 20\%$.
10. Encapsulation of the `G4NRF` class in the `G4NRFPhysics` class, which allows inheritance from `G4VPhysicsConstructor` as per Geant4 standards for adding custom physics processes.
11. Use of `SetProcessType()` and `SetProcessSubType()` (both set to `fElectromagnetic`) in the `G4NRF` class to conform to further Geant4 custom physics process standards.

In addition to the above updates, several additions to the code were made following the completion of the main simulations:

1. Use of a cut to kill photons with $E < 1.7$ MeV (the lower limit of NRF lines from the four isotopes of interest) in the (user class) `steppingAction` rather than the `stackingAction`. This kills photons that scatter below 1.7 MeV, rather than killing only *new* tracks created below this energy, giving a $\sim 25\%$ speed improvement.
2. Use of a sharp energy threshold to rapidly return 0 cross section in the `G4NRF` class for photons far off-resonance to avoid numerical integration. In the notation of Section 2.2, the ratio

$$z \equiv \frac{x}{\sqrt{t}} = \frac{2(E - E_r)}{\Delta} \quad (3.3)$$

describes the energy displacement of a photon from the resonance condition at $E = E_r$. Setting a large tolerance of $|z| \leq 1.0 \times 10^3$ —which for $\Delta \sim 1$ eV

corresponds to photons with $|E - E_r| > 500$ eV—gives a 6% speed increase. Note that in the actual code, cuts are placed on the photon's $z^2 = x^2/t$ to avoid calling the costly `sqrt` function.

3. A cross section lookup table may be built at initialization to avoid on-the-fly numerical integration, but currently there is only a very basic interpolation between meshpoints and the results have not been validated against the full numerical integration.
4. Relaxation of the search boundaries in the `MakeWidth0()` method of the class `G4NRFNuclearLevel`. Initially, the upper acceptable bound `TAU_CUT_HI` on the level width Γ_r was set to to 1 eV. This was sufficient for the actinide isotopes of interest, which have $\Gamma_r \sim 50$ meV, but defaulted to a width of zero for levels wider than 1 eV in *e.g.* lead-208, causing the cross section evaluation to always return zero.

3.2 Simulation details

Author's note: the following sections discuss the actual implementation of the simulations of the NRF measurements, and were initially written for the Supplemental Information for an article [47] based on the same simulations.

3.2.1 Bremsstrahlung interrogation beam

A bremsstrahlung beam is used to interrogate the object and reference foil. This spectrum is first generated in a standalone Geant4 simulation in which a pencil beam of monoenergetic 2.7 MeV electrons impinges on a gold target of thickness 102 μm with a copper backing of thickness 20.1 mm. The subset of bremsstrahlung photons within 10° of the incident beam axis is then used to create the interrogation beam spectrum. This spectrum is smoothed using an interpolation routine to reduce the effects of discrete bin edges, then fed into the main Geant4 simulation as a pencil beam of photons directed towards the center of the warhead.

Because of computational limitations, the bremsstrahlung spectrum is not sampled directly—that is, the energy spectrum of particles injected into this simulation does not follow the distribution of the true bremsstrahlung spectrum. Although the integrated NRF cross sections (Eq. 2.26) are relatively strong ($\int \sigma(E) dE \sim \mathcal{O}(10 \text{ eV}\cdot\text{b})$), the energy-dependent NRF cross sections are exceptionally narrow with Doppler-broadened widths of $\Delta \sim \mathcal{O}(1 \text{ eV})$. If sampling from a broad-spectrum bremsstrahlung spectrum, an NRF interaction would be a rare event because the probability that the random number generator (or the nature of the bremsstrahlung process) will choose a photon energy sufficiently close to the resonant energy to excite an NRF transition is small. Thus, the computational effort required for obtaining a statistically significant NRF signal using a naïve sampling of the full bremsstrahlung beam is enormous. This is a classic variance reduction problem that can be mitigated by importance sampling.

Instead of directly and inefficiently simulating the incident bremsstrahlung beam $\phi_0(E)$, an artificial sampling spectrum $s(E)$ that emphasizes NRF transitions of interest²⁷ and de-emphasizes off-resonant photons is used [48, p. 165]. To adjust for this relative over- and under-sampling, respectively, each incident photon is tagged with a statistical weight factor $w(E) = p(E)/s(E)$, where $p(E)$ is the probability density function (PDF) of the true spectrum $\phi_0(E)$ and $s(E)$ is also normalized to unity. Plots of the $p(E)$ and $s(E)$ used in the simulations can be seen in Figure 3-2. This weight is then propagated throughout the simulation to the secondary particles produced by each primary photon. To recover the output NRF spectrum that would correspond to sampling from $p(E)$ rather than $s(E)$, every event’s contribution to the final recorded spectrum needs to be weighted by $w(E)$, where E is the initial energy of that event. This amounts to filling the final histogram not with one entry per signal photon, but with $w(E)$ entries per signal photon, where $w(E)$ is again the weight

²⁷Though each of the four isotopes of interest (^{238}U , ^{235}U , ^{240}Pu , and ^{239}Pu) has ten or so experimentally-observed lines, the ENSDF-derived database used by G4NRF lacks data for level spins J_r and widths Γ_r for many of the lines of the latter three isotopes. In some cases, reasonable estimates of these values can be made such that the observed integrated cross sections are replicated in the simulation. In other cases, such as the pair of ^{235}U transitions at 1769 and 1815 keV, this is not so, and the NRF signal of these incompletely-understood states is excluded from the model.

calculated from the *initial* photon energy. Since $s(E) \gg p(E)$ and thus $w(E) \ll 1$ for NRF photons of interest, this corrects for the oversampling of the NRF lines with respect to the true bremsstrahlung distribution.

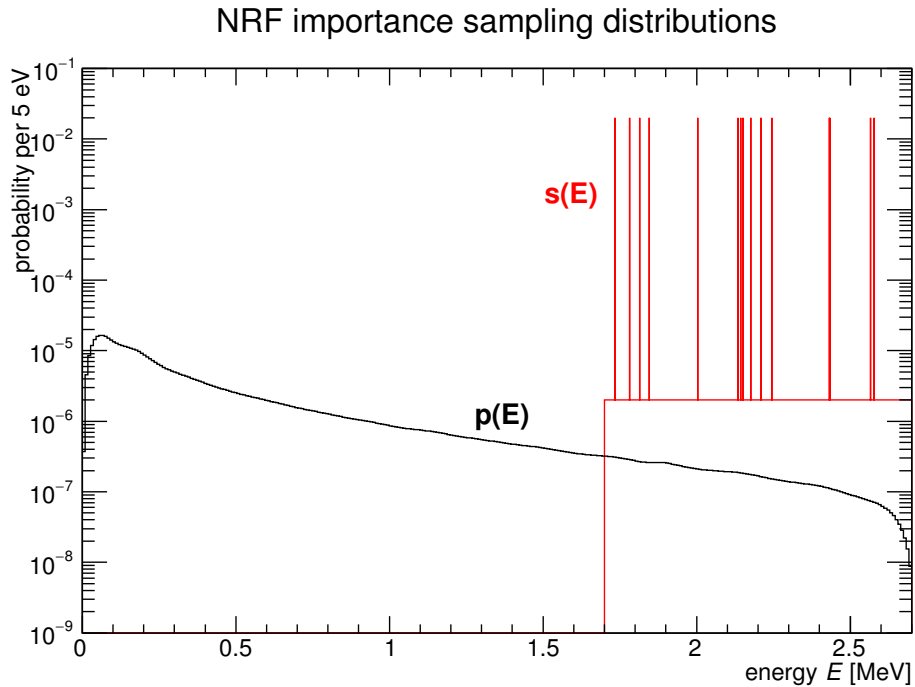


Figure 3-2: Smoothed Monte Carlo bremsstrahlung spectrum $p(E)$ (black) and sampling spectrum $s(E)$ (red) used in the simulations. The PDFs are normalized such that the sum rather than the true integral of their bin contents is each unity, but the distinction is immaterial when taking the ratio $w(E) = p(E)/s(E)$ as long as the distributions have identical binwidths. The spikes in $s(E)$ are narrow regions bracketing the NRF resonances of interest.

3.2.2 Simulation geometries

For the purposes of simulation, a nominal nuclear warhead geometry is constructed from an approximate characterization of a Soviet thermonuclear tactical warhead (hereafter the “Black Sea” warhead) as estimated from information in the public domain [13]. While the precise dimensions of the warhead are classified and thus unknown, the purpose of this work is to achieve a proof of concept, for which the

approximate dimensions are sufficient. Using simplified isotopic models²⁸ of the special nuclear material and the means of the ranges of dimensions quoted in the paper, the Black Sea geometry simulated in Geant4 consists of three concentric shells of weapons-grade plutonium (WGPu), high explosive (HE), and highly-enriched uranium (HEU).

Several hoax configurations are also considered. Two isotopic hoaxes are constructed by exchanging the WGPu for either fuel-grade plutonium²⁹ (FGPu) or pure U-238, the latter being a close analogue of depleted uranium. These two hoaxes test the ability of the system to detect material nonconformity in a candidate object. Additionally, a geometric hoax is constructed by arranging two sets of rectangular slabs to mimic the sequence of attenuations in the Black Sea warhead. The geometric hoax tests the ability of the system to detect geometric nonconformity in the candidate object.

The reference foil is a homogenized sheet of U-235, U-238, Pu-239, and Pu-240 (equal parts by mass), 2 cm thick with an overall density of $\rho = 19 \text{ g/cm}^3$. Perpendicular to the beam, the foil height and width are 50 cm, and the foil center is located 1 m downbeam from the origin at the center of the warhead.

Due to computational constraints we do not simulate energy deposition in a realistic detector geometry, which would result in most of the computed photons being discarded. Instead, we tally all photons with energy $1.7 \text{ MeV} \leq E \leq 2.7 \text{ MeV}$ exiting the foil with a backwards angle of $\theta > 3\pi/4$ (i.e. the backward-opening cone with solid angle 1.8 steradians) relative to the forward beam direction. To account for detector effects, the resulting photon counts are then multiplied by a realistic high purity germanium (HPGe) detector's geometric and intrinsic peak efficiencies.

To further improve computational efficiency, several cuts are made in particle

²⁸The weapon isotopics simulated include only the four isotopes of interest (uranium-235, uranium-238, plutonium-239, and plutonium-240)—see Table 3.1. A more detailed simulation might include the presence of minor isotopes such as uranium-234, decay products such as americium-241, and allotropic stabilizers such as gallium [49].

²⁹Since 1976, the U.S. Department of Energy has defined fuel grade plutonium as an intermediary between weapons-grade plutonium and reactor-grade plutonium. Fuel grade plutonium is defined as plutonium with between 7% and 19% Pu-240. Nuclear weapons made from fuel-grade plutonium have been successfully detonated. See [50].

tracking. All secondaries created in the weapon are killed, eliminating the tracking of electrons created by photoelectric absorption, Compton scattering, and pair production. This cut is not applied in the foil, however, as the secondary bremsstrahlung emitted by these charged particles is the primary contributor to the continuum underlying the NRF peaks—see Section 2.9. Any NRF photon emitted in the weapon at an angle greater than $\pi/4$ with respect to the beam direction is also killed, as it is unlikely to reach the foil (note: although NRF photons created in the weapon and emitted towards the foil are unlikely to induce a further NRF interaction in the foil due to the recoil energy deficit, they still may produce continuum photons and thus are not killed.)

Dimensions and material compositions of each geometry are summarized in Table 3.1, and diagrams of the Black Sea warhead, geometric hoax, and reference foils are shown in Figure 3-3. The full set of Monte Carlo results is shown in Figs. 3-4–3-9 and described in detail in Section 3.3.

Object	Material layers	Composition (by mass %)	Dimensions [cm]
Black Sea template	WGPu	94% Pu-239, 6% Pu-240	$(r_{\text{in}}, r_{\text{out}}) = (6.27, 6.7)$
	HMX	3% H, 16% C, 38% N, 43% O	$(r_{\text{in}}, r_{\text{out}}) = (6.7, 13.2)$
	HEU	95% U-235, 5% U-238	$(r_{\text{in}}, r_{\text{out}}) = (13.2, 13.45)$
Black Sea, with WGPu \rightarrow U-238 substitution hoax	U-238	100% U-238	$(r_{\text{in}}, r_{\text{out}}) = (6.27, 6.7)$
	HMX	3% H, 16% C, 38% N, 43% O	$(r_{\text{in}}, r_{\text{out}}) = (6.7, 13.2)$
	HEU	95% U-235, 5% U-238	$(r_{\text{in}}, r_{\text{out}}) = (13.2, 13.45)$
Black Sea, with WGPu \rightarrow FGPu substitution hoax	FGPu	86% Pu-239, 14% Pu-240	$(r_{\text{in}}, r_{\text{out}}) = (6.27, 6.7)$
	HMX	3% H, 16% C, 38% N, 43% O	$(r_{\text{in}}, r_{\text{out}}) = (6.7, 13.2)$
	HEU	95% U-235, 5% U-238	$(r_{\text{in}}, r_{\text{out}}) = (13.2, 13.45)$
Geometric hoax (see Figure 3-3)	WGPu	94% Pu-239, 6% Pu-240	$\ell = 2 \times 0.43$
	HMX	3% H, 16% C, 38% N, 43% O	$\ell = 2 \times 6.50$
	HEU	95% U-235, 5% U-238	$\ell = 2 \times 0.25$
Reference foil	Homogenous	25% each U-235, U-238, Pu-239, Pu-240	$\ell = 2$

Table 3.1: Object, hoax, and reference foil geometries simulated in Geant4, using weapons-grade plutonium (WGPu), fuel-grade plutonium (FGPu), high explosive (HMX formula), highly-enriched uranium (HEU), U-238, and a four-isotope composite. The pair $(r_{\text{in}}, r_{\text{out}})$ denotes the inner and outer radii of the spherical shell, whereas ℓ denotes the sheet thickness along the beam axis. When elements are given rather than isotopes, the isotopic composition of each element is assumed to be natural. Isotopics are simplified from those in [49].




Figure 3-3: Graphical rendering of the Geant4 models of the Black Sea warhead (shown in cutaway, left) and the dual-slab geometric hoax (right). The plutonium layers are shown in orange, high explosive in green, uranium in blue, and reference foil in gray. The simulation only registers the events corresponding to photons emitted from the foil at $\theta \geq 135^\circ$ angles, relative to the beam axis.

3.3 Simulation results

Figures 3-4 to 3-9 show the Monte Carlo simulated photon spectra in the angular acceptance region $\theta > 3\pi/4$ for various scenarios. In each plot, two weighted spectra resulting from $N_\gamma^{\text{sim}} = 8 \times 10^{11}$ incident photons are compared; the black curve is fixed to be the spectrum from the authentic Black Sea template, while the red curve is the spectrum of the particular candidate object. The error bars on each energy bin are given by the square root of the sum of the squares of weights $\text{err}_E = \sqrt{\sum_i w_i^2(E)}$ in that bin, where $w_i(E)$ is the weight of the i -th event contributing to the bin of energy E , as determined in Section 3.2.1. For the oversampled NRF lines this results in much smaller errors than simply the square root of the counts, as would be the case for simple Poisson errors. These errors represent the uncertainty on the expected signal strength as determined by the simulation, rather than the uncertainty due to counting statistics that would be seen in a real experiment. When calculating the discrepancies between authentic and candidate objects the value of err_E is replaced with the Poisson error, i.e. \sqrt{N} where $N = \sum_i w_i(E)$ is just the bin content, as outlined in Section 3.3.1.

In Figure 3-4, two authentic template objects are compared. Computationally, this amounts to simulating the Black Sea warhead twice, with different random number generator seeds used to sample the interrogation spectrum. The two spectra agree well within error bars, as expected.

In Figure 3-5, the first hoax scenario is considered: the hosts use the Black Sea warhead geometry but replace the WGPu component with pure U-238. This type of hoax will successfully avoid detection by neutron radiography systems [20] but will still be detected in NRF measurements (see Table 3.2). By Eq. 2.52, the presence of additional U-238 in the candidate object strongly *decreases* the U-238 NRF signal seen by the detector, since the object becomes a stronger filter of the lines corresponding to U-238. The corresponding loss of Pu-239 and Pu-240 in the candidate object *increases* the signal from these two isotopes when compared to the authentic template. Note that the continuum underlying the NRF lines also increases as the $Z = 94$ of the WGPu is replaced by the $Z = 92$ of the U-238; the continuum is a result of the photoelectron bremsstrahlung process which scales strongly with Z and thus is attenuated less in the hoax object.

In Figure 3-6, the WGPu is instead replaced with FGpu. In this substitution, the mass fraction of Pu-239 in the plutonium shell decreases from 94% to 86% (a change by a factor of 0.91) while the mass fraction of Pu-240 changes from 6% to 14% (a change by a factor of 2.3). This substantial increase in Pu-240 content manifests itself as a noticeable reduction in the Pu-240 lines of the candidate object.

In Figure 3-7, the geometric hoax (see Figure 3-3) that mimics the sequence of attenuations in the authentic Black Sea warhead is interrogated. By design, it is undetectable when perfectly aligned with the beam (rotation of 0°). The two spectra agree well within statistics, indicating a successful hoax at that particular angle.

In Figure 3-8, the geometric hoax of Figure 3-7 is rotated by 10° , such that more of the hoax material lies in the beam path. Most NRF lines considered show a weaker signal in the rotated hoax as compared to the rotated authentic object.

For completeness, in Figure 3-9, the geometric hoax is rotated to 30° . The further increase in the amount of material in the beam path causes the final signal to be

even more strongly attenuated than in Figure 3-8. Now, both the entire set of NRF lines and the continuum are substantially reduced in the candidate spectrum. In Section 3.3.1, it will be shown however that a rotation by 10° is sufficient to detect this particular hoax with 4σ confidence.

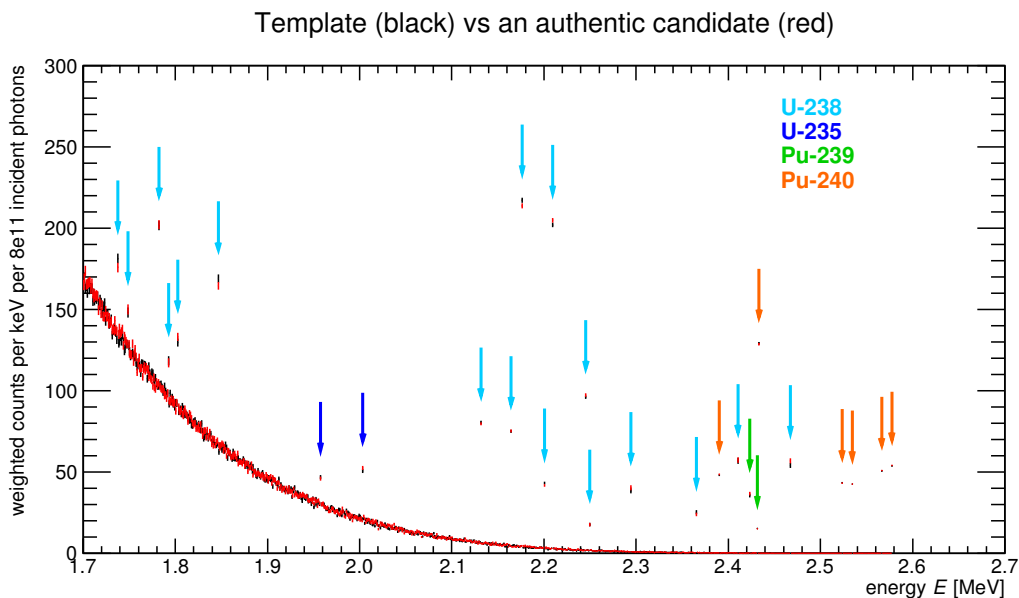


Figure 3-4: Monte Carlo NRF spectra of the Black Sea control geometry (black curve) vs the identical Black Sea geometry with different Monte Carlo random seeds (red curve). The two spectra agree well within statistics. Lines corresponding to the four main NRF isotopes are marked by colored arrows. The errors are described in the text.

3.3.1 Extrapolation to fixed measurement times

It remains now to extrapolate from the $N_\gamma^{\text{sim}} = 8 \times 10^{11}$ incident photons used for the weighted spectra in Section 3.3 to the number of photons N_γ^{exp} produced in a realistic experiment. The total number of electrons injected into an accelerator beam over a time t is $N_e = I_{\text{beam}}t/q$, where I_{beam} is the beam current (in Amperes) and $q = 1.602 \times 10^{-19}$ C is the (absolute) electron charge. The total number of photons produced across the entire bremsstrahlung energy spectrum is then just $N_\gamma^{\text{exp}} = f_{\text{conv}}N_e$, where f_{conv} is a dimensionless conversion factor, obtained from the Geant4 bremsstrahlung simulations, that gives the number of bremsstrahlung photons per electron




Figure 3-5: Monte Carlo NRF spectra of the Black Sea control (black) vs the Black Sea geometry with WGPU replaced by U-238 (red). The U-238 lines (light blue arrows) are strongly diminished in the hoax, while the Pu-239 lines (green arrows) and Pu-240 lines (orange arrows) are enhanced.

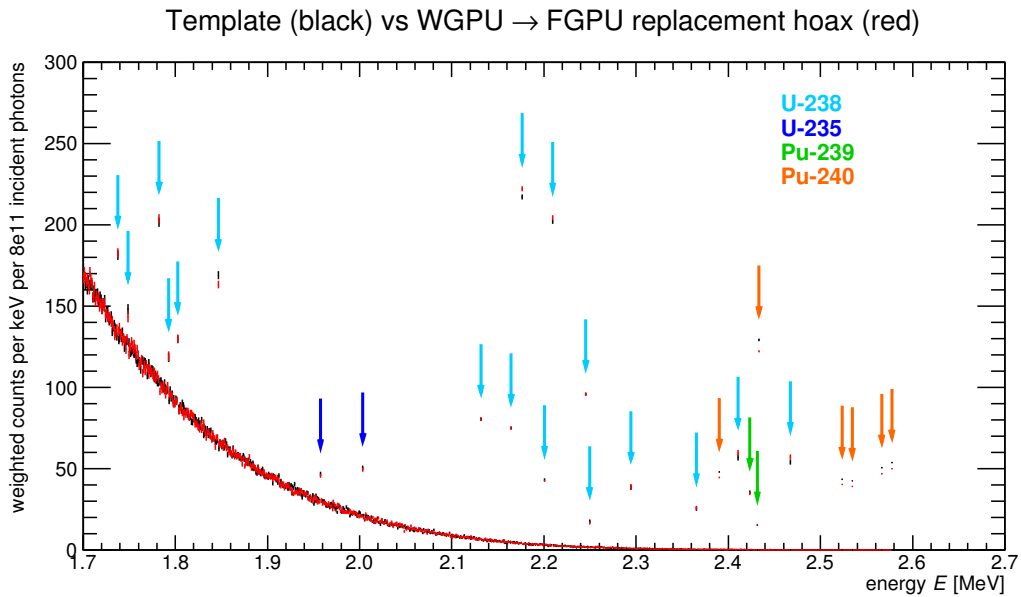


Figure 3-6: Monte Carlo NRF spectra of the Black Sea control (black) vs the Black Sea geometry with WGPU replaced by FGPU (red). The Pu-240 lines (orange arrows) are diminished in the hoax, but the change in Pu-239 lines (green arrows) is not easily detectable with the statistics available.

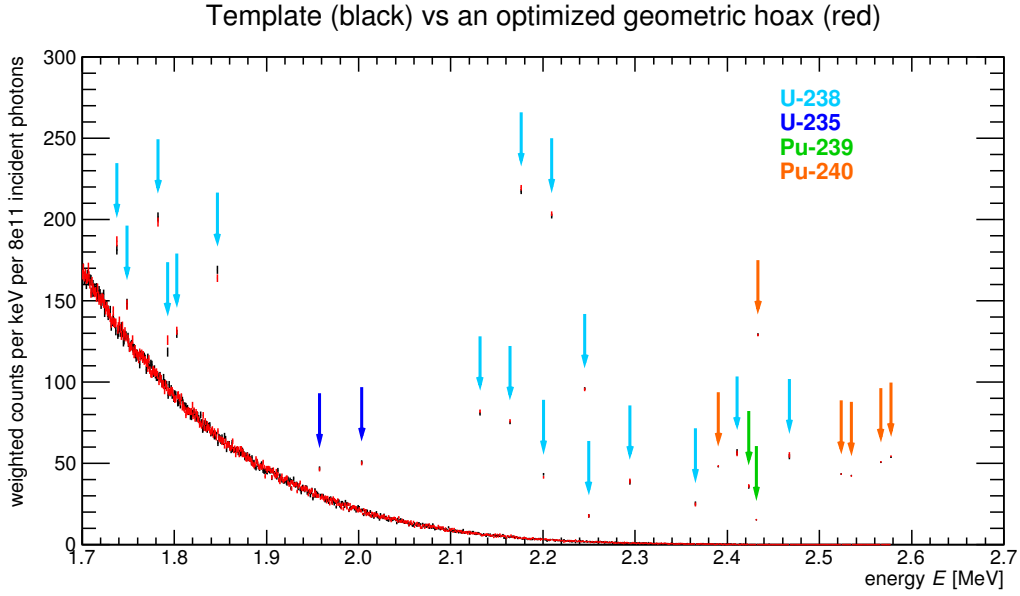


Figure 3-7: Monte Carlo NRF spectra of the Black Sea control (black) vs the unrotated geometric dual-slab hoax (red). The two spectra agree well within statistics.

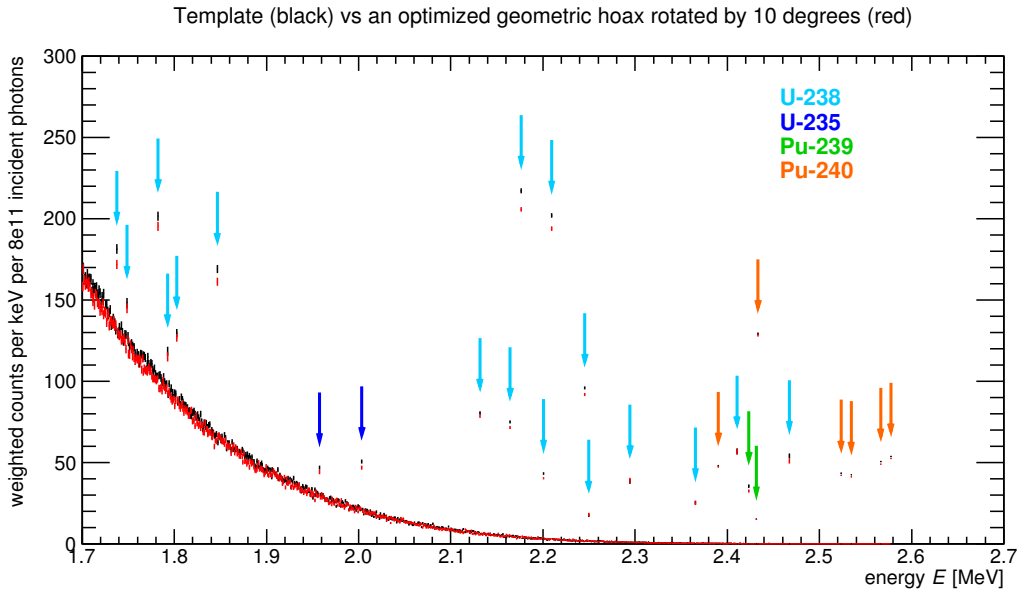


Figure 3-8: Monte Carlo NRF spectra of the Black Sea control (black) vs the geometric dual-slab hoax rotated by 10 degrees (red). The two spectra no longer agree well within statistics.

incident on the radiator. For the angular acceptance of 10° used here, f_{conv} has been determined to be $f_{\text{conv}} \simeq 0.0144$, independent of the artificial sampling distribution $s(E)$. Altogether, the total number of incident photons N_γ^{exp} produced in a single

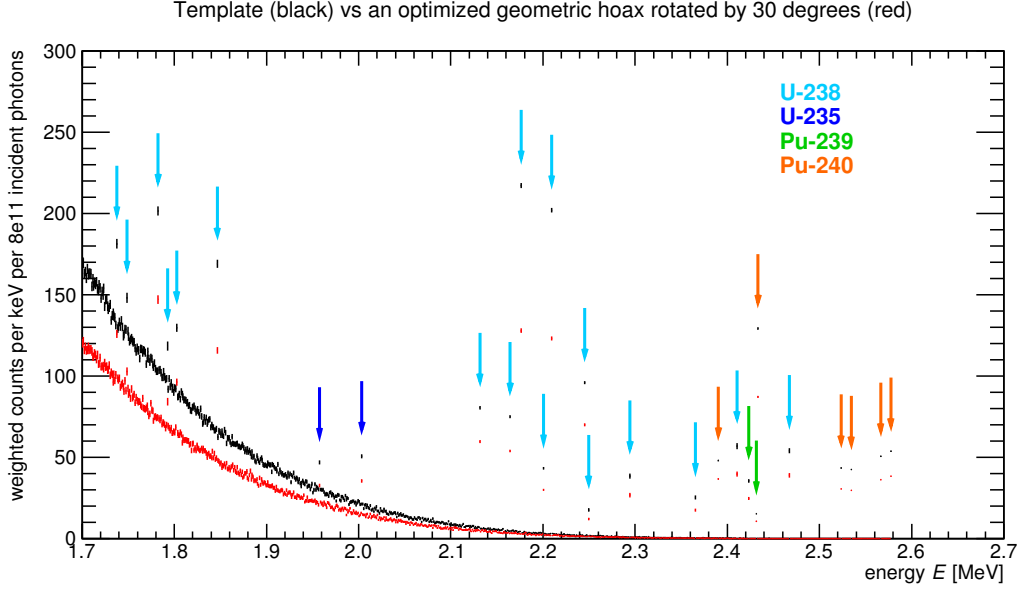


Figure 3-9: Monte Carlo NRF spectra of the Black Sea control (black) vs the geometric dual-slab hoax rotated by 30 degrees (red). The two spectra disagree strongly in both the NRF lines and the continuum.

accelerator run is

$$N_{\gamma}^{\text{exp}} = f_{\text{conv}} I_{\text{beam}} t / q, \quad (3.4)$$

which amounts to a rate of $N_{\gamma}^{\text{exp}}/1\text{ s} = 3.6 \times 10^{14}$ incident photons per second with a proposed beam current of 4 mA [51]. To extrapolate from the Monte Carlo results of Figures 3-4–3-9 to the counts and therefore the beam time needed for a realistic experimental setup, the spectra are scaled linearly by a factor

$$k \equiv \epsilon_{\text{int}} \epsilon_{\text{geom}} M_{\text{det}} N_{\gamma}^{\text{exp}} / N_{\gamma}^{\text{sim}}, \quad (3.5)$$

where the geometric and intrinsic efficiencies of a number M_{det} of realistic HPGe detectors have also been taken into account. For a cylindrical detector of radius 5 cm positioned 1 m from the foil, the geometric efficiency is $\epsilon_{\text{geom}} \simeq 0.008/1.8/\text{detector}$ (to adjust from the 1.8 steradians of the $\theta > 3\pi/4$ signal region to the realistic solid angle) and the intrinsic efficiency is $\epsilon_{\text{int}} \simeq 0.20$ [52]. The value of this scale factor is set to be $k \simeq 254$, assuming $M_{\text{det}} = 30$ such detectors, as for these values, the significances in

Table 3.2 will be reached with $t = 21$ s. Note that these results are obtained for a foil thickness of $X = 2$ cm; if a thinner foil is required (see Section 4.1), fewer counts per second will be generated in the foil and the measurement times will increase roughly as shown in Figure 3-10. Upon scaling, the Monte Carlo error bars are replaced by the Poisson error bars that would be seen in a real experiment. Explicitly, an NRF line with n counts and error bars of $\pm\delta n$ in the Monte Carlo simulation (where $\delta n \neq \sqrt{n}$ in general) is scaled to a value of kn with error bars of $\pm\sqrt{kn}$. The plot of the Black Sea NRF spectrum under such scaling is shown in Figure 3-11.




Figure 3-10: Plot of the ratio of the time required to achieve a desired NRF signal strength in the $E_r = 2.176$ MeV uranium-238 line with an arbitrary total foil thickness X versus the time required with $X = 2$ cm. For illustrative purposes, ratios are calculated via a modified version of Eq. 2.52 at $\theta = 0$ using only uranium-238 data with the μ_{NRF} terms divided by 4 to more closely model the NRF response of the simulated composite foil. For the $E_r = 2.176$ MeV decay to the ground state, the NRF decay cross section is $\sigma_{\text{NRF}}(E_r) = 31.5$ b and the non-resonant cross section is $\sigma_{\text{nr}}(E_r) = 18.8$ b. For $X \ll 1$ cm, the thin foil limit of Eq. 2.52 is reached, in which the NRF counts are directly proportional to X and therefore the time required to obtain a fixed number of counts is inversely proportional to X . For $X \gg 1$ cm, the saturation limit of Eq. 2.52 is reached, and the measurement time does not improve substantially.

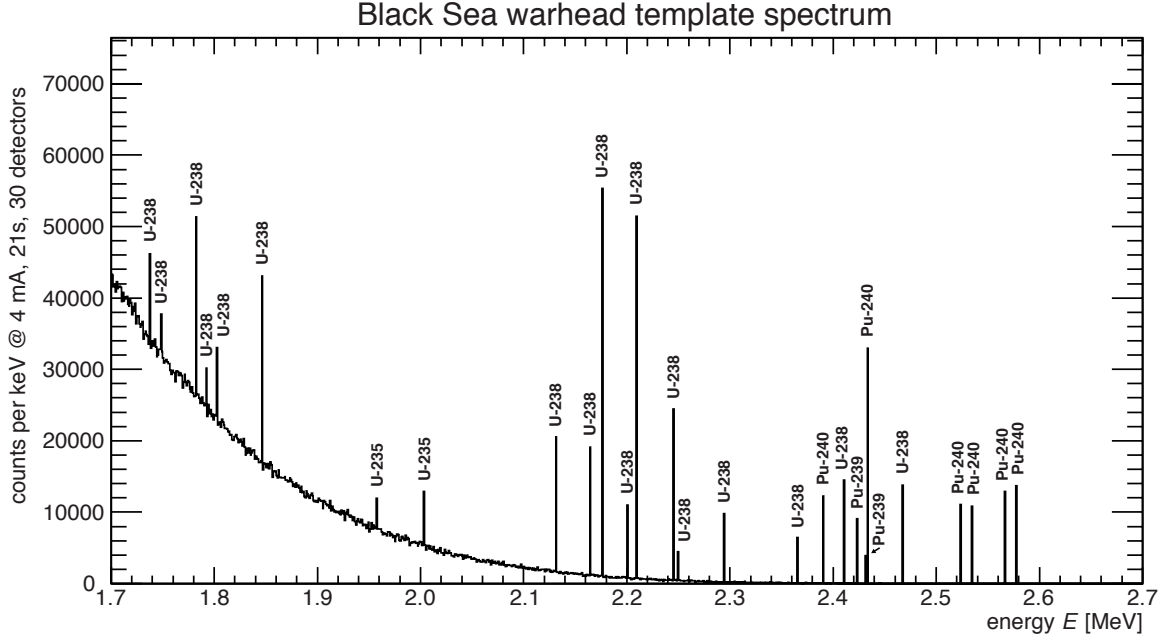


Figure 3-11: NRF spectrum of the Black Sea control geometry. The spectrum has been scaled by a factor of ~ 254 , and the error bars have been suppressed for clarity.

3.3.2 Statistical tests of spectra

The NRF lines in spectra from a candidate object are compared to the NRF lines in the spectra of the template object to determine if the two objects match. The deviation in the strength of the NRF lines can be expressed in units of σ by

$$\Delta = \frac{c_{\text{can}} - c_{\text{tem}}}{\sqrt{\sigma_{\text{can}}^2 + \sigma_{\text{tem}}^2}}, \quad (3.6)$$

where c_{can} and c_{tem} denote the counts (or scaled counts if using simulated data) in the NRF energy bin of interest for the candidate and authentic template object, respectively, and σ_{can} and σ_{tem} are the corresponding errors. For the lines with little or no continuum background, the errors are approximately the square root of counts. For the hoax scenarios considered, the differences between the hoax and template for at least one line—see Table 3.2—always exceeded 7σ in ~ 20 array-seconds.

Note that this comparison, shown numerically in Table 3.2, simply compares the single strongest observed line from each of the four isotopes. This data reduction technique thus ignores all the other observed NRF lines, many of which are still quite

Scenario	^{238}U	^{235}U	^{239}Pu	^{240}Pu	$1 - P_{\text{alarm}}$ (4σ threshold)	Figure (in text)
Template vs authentic	1.41	0.454	1.47	1.69	0.98	3-4
Template vs WGPu \rightarrow U-238	127	3.49	14.0	1.48	0.0	3-5
Template vs WGPu \rightarrow FGPu	3.75	1.93	1.49	7.16	4.6×10^{-4}	3-6
Template vs geometric hoax 0° rotation (perfect alignment)	1.81	0.96	0.747	0.320	0.98	3-7
Template vs geometric hoax, 10° rotation	8.83	4.40	0.0796	0.964	2.3×10^{-7}	3-8
Template vs geometric hoax, 30° rotation	76.7	25.3	14.2	45.6	0.00	3-9

Table 3.2: Detection probabilities for the strongest lines of the four isotopes in various hoax scenarios, as measured in units of σ . The values are computed for the strongest simulated NRF line of each isotope, using a 21 second run as described in the text. The “ $1 - P_{\text{alarm}}$ ” column indicates the probability that none of the four comparisons cross the 4σ threshold. The “Figure” column indicates the figure containing the Monte Carlo results that are scaled and then compared via Eq. 3.6 to compute the values in the isotope columns. The specific lines used are 2176 keV (U-238); 1957 keV (U-235); 2431 keV (Pu-239); and 2433 keV (Pu-240). The non-unity value for the $1 - P_{\text{alarm}}$ in the template vs authentic scenario is only due to the limited sampling in the Monte Carlo simulation.

strong, as well as the secondary bremsstrahlung continuum underlying the peaks. By combining multiple lines and potentially even the continuum, more elaborate comparison methods may allow future tests to achieve stronger confidence with the same number of photons, or conversely, the same confidence with fewer photons, significantly reducing the measurement or simulation time.

3.3.3 Dose to the warhead

The radiation dose imparted to the warhead during the course of the NRF measurement may be of interest, as some components of the warhead may be compromised at high doses. Although not estimated in the initial full-scale simulations of Figs. 3-4–3-9, the dose estimation does not require high statistics in the NRF bins, and thus can be extrapolated from smaller-scale simulations.

For the dose calculation, the tracking cuts described in Sections 3.2.2 and 3.1.4 are disabled and the sampling distribution $s(E)$ is extended to cover the entire 0–2.7 MeV energy range, as low-energy photons contribute significantly and nonlinearly

to the total dose. Using the `sensitiveDetector` functionality of Geant4, the total weighted dose to the entire outer uranium layer (of mass 10.4 kg) is computed as 3.2×10^{-9} Gy per 1×10^6 incident photons. To adjust for the 10° beam spread that would be present in a realistic beam, an upper bound on the *local* dose at the beam entry point on the warhead can be estimated by assuming that the entire energy deposition in the uranium shell occurs in a uranium cone of half-angle 10° and height 0.25 cm (thus with mass ~ 0.0097 g) as per Table 3.1. This amounts to scaling the 3.2×10^{-9} Gy by a factor of $\sim 10^6$. Furthermore, to extrapolate from the 1×10^6 incident photons to the full measurement time, the dose is scaled as in Section 3.3.1 by a factor of $8 \times 10^{11} / (1 \times 10^6) \times 254$ to give an approximate upper bound on the local dose of ~ 700 kGy for the 2.4 C integrated beam charge used in the measurement. Nearly all of this dose is due to broad-spectrum bremsstrahlung photons that do not contribute to the NRF signal, and thus the dose could be substantially reduced by using monochromatic photon sources as discussed in Section 4.1.

Although this value of 700 kGy has been computed specifically using energy deposition in the uranium shell, it indicates that the dose to other weapon components will also be quite large. The fissile plutonium is not likely to be sensitive to such large doses, as the ~ 10 W of heat produced by alpha decay in a 4.5 kg plutonium mass corresponds to a uniform dose rate of 2.2 Gy/s, or 700 kGy in only 3.7 days, much shorter than the operational lifetime of a US nuclear warhead. The large dose induced by the NRF measurement, however, will destroy even radiation-hardened electronics in the warhead [53, 54].

3.4 Conjectured nuclear level data

Much nuclear level information for vital NRF isotopes of interest including U-235, Pu-239, and Pu-240, is either unknown or not published in the open literature. In particular, the excited state spins J_r and widths Γ_r are unavailable for many known NRF transitions. This is due to the fact that most NRF experiments are only sensitive

to the integrated cross section (Eq. 2.26),

$$\int \sigma_{0,r,j}(E) dE = 2\pi^2 g_r \left(\frac{\hbar c}{E_r} \right)^2 \frac{\Gamma_{0,r} \Gamma_{r,j}}{\Gamma_r}, \quad (3.7)$$

where $g_r = (2J_r + 1)/[2(2J_0 + 1)]$ is the statistical spin factor, since the resolution of a typical HPGe detector (~ 1 keV) is much too large to determine the exact shape and thus the width of the resonance (~ 10 meV before Doppler broadening). In such experiments, both J_r and Γ_r in Eq. 3.7 are unknown, while the integrated decay cross section $\int \sigma_{0,r,j}(E) dE$ (or sometimes the absorption cross section $\int \sigma_{0,r}(E) dE$, having summed over all decay modes j) is measured and known [55, 56, 57]. The $\Gamma_{0,r}$ and $\Gamma_{r,j}$ terms can be rewritten as known decay branching ratios $b_{r,j} \equiv \Gamma_{r,j}/\Gamma_r$. The **G4NRF** code, however, requires its databases have values for both J_r and Γ_r for each level of interest.

Although the J_r and Γ_r are unknown, reasonable estimates for the J_r can be made based on nuclear spin selection rules. The Γ_r can then be back-calculated by rearranging Eq. 3.7 to

$$\Gamma_r = \frac{1}{b_{r,0} b_{r,j}} \frac{1}{2\pi^2 g_r} \left(\frac{E_r}{\hbar c} \right)^2 \int \sigma_{0,r,j}(E) dE. \quad (3.8)$$

Since it is typically the strongest NRF lines that are of interest, the transition between excited and ground states is assumed to be of type E1 or M1, i.e. that the difference $J_r - J_0 = 1$. This can be seen via the Weisskopf estimates [58] for gamma decay rates, which state that in general, the higher the spin difference, the slower the rate of transition. In this context, higher spin differences give rise to larger transition lifetimes, therefore smaller partial widths $\Gamma_{0,r}$ (since $\tau \propto 1/\Gamma$), therefore smaller integrated cross sections³⁰. For U-235, Pu-239, and Pu-240, this assumption gives spins J_r of 9/2, 3/2, and 1, respectively. With these assumed resonant level spins, the conjectured level widths Γ_r can be found. Note that while the conjectured spins and widths may differ substantially from their true (but unknown) values, this calculation

³⁰Although the g_r in Eq. 3.7 introduces a term proportional to J_r , any increase in the integrated cross section via g_r due to an increase in J_r is counteracted strongly by the decrease in $\Gamma_{0,r}$ from the Weisskopf rules.

preserves the integrated cross section.

The calculation becomes slightly more complicated due to a peculiarity of some odd- A nuclei in which many of the energy levels E_r have a “partner level” at $E_{r+1} \equiv E_r + E_1$, where E_1 denotes the energy of the first excited state (note: the intermediate level E_1 is not always the *first* excited state but can be treated as such for this discussion), possibly arising from a coupling of single-particle and rotational excitations [59]. Then (neglecting recoil) a photon of energy E_r can arise from the decay $r \rightarrow 0$ but also from $r + 1 \rightarrow 1$. For example, the 2.135 MeV photon in Pu-239 can arise from a decay from the 2.135 MeV level direct to the ground state or from the 2.143 MeV level decaying to the 8 keV first excited state. To find the width of the 2.135 MeV level, one can write the observed integrated cross section (4 eV·b from Bertozzi [56]) as the sum of the two decays:

$$\begin{aligned}
 4 \text{ eV}\cdot\text{b} &= \sum_r \int \sigma_{0,r,j}(E) dE & (3.9) \\
 &= 2\pi^2 g_r \left(\frac{\hbar c}{2.135 \text{ MeV}} \right)^2 \Gamma_{2.135 \text{ MeV}} & (3.10) \\
 &+ 2\pi^2 g_r \left(\frac{\hbar c}{2.143 \text{ MeV}} \right)^2 \frac{1.0}{1.0 + 0.31} \frac{0.31}{1.0 + 0.31} \Gamma_{2.143 \text{ MeV}}
 \end{aligned}$$

where $\Gamma_{2.143 \text{ MeV}}$ can be found using the simpler technique above, the assumed g_r factors are the same for both decays (since g_r does not involve the spin of the final level), and the branching ratios (1.0 for the 2.135 MeV decay from the 2.135 MeV state) are from NNDC [43]. This calculation gives a value of $\Gamma_{2.135 \text{ MeV}} = 0.00233921 \text{ eV}$. Note that rather than relying on these more complicated two-term calculations, the poorly-understood NRF lines for which this technique is required were excluded from the model, as inconsistencies and experimental uncertainties in tabulated branching ratios may result in negative values of Γ_r . NRF lines with widths following directly from Eq. 3.8 and an assumed J_r were not excluded.

Chapter 4

Discussion

4.1 Limitations

The results of Section 3.3 show, as a proof-of-concept, that the physical cryptographic NRF measurement could reliably verify the authenticity of a nuclear warhead. These simulations suggest that a future, optimally-engineered real system may, in the language of Section 1.3, satisfy the completeness, soundness, and zero-knowledge criteria required of a practical warhead verification technique. To move beyond the current proof-of-concept stage, however, several limitations must be addressed.

First, the bremsstrahlung pencil beam used has no spatial extent or beam divergence as would be observed in a realistic beam. Therefore, the pencil beam only illuminates a small portion of the warhead, and thus does not eliminate all possible geometric hoaxes. The simulations have also only been completed at a single bremsstrahlung endpoint energy of $E_{\text{end}} = 2.7$ MeV, which can only access the NRF lines of the four weapon isotopes of interest. The NRF lines of the carbon and nitrogen isotopes in the high explosives, for instance, will require simulations at $E_{\text{end}} \gtrsim 5$ MeV [60, 61, 62].

Second, the model warhead is highly simplified in both its isotopics and geometries. It has ignored, for instance, minor isotopes [49] and the gallium stabilizer [63] in the fissile material that may produce a small but measurable NRF signal. Moreover, the

model neglects the secondary stage, the tamper/radiation case, and other warhead components (any of which may significantly attenuate the beam), and assumes that the three layers are both concentric and perfectly spherical.

Third, the encrypting foil thickness is $X = 2$ cm, which produces a strong NRF signal and is thus computationally convenient, but such a thickness may be unworkable in practice. The $50 \times 50 \times 2$ cm foil of density $\rho \simeq 19$ g/cm³ used in the Geant4 simulations, for instance, will have an impractically-large mass of 95 kg including several significant quantities of plutonium-239, and have a neutron multiplication factor $k_{\text{eff}} = 0.45$ [22]. Preliminary investigations of the information leakage in the system also indicate that the error in estimating the foil thickness under a certain class of approximations is minimized at foil thicknesses of ~ 1 cm. In a real system, foil thicknesses may therefore be on the order of $X \sim 1$ mm or lower, reducing count rates and increasing measurement time as per Fig. 3-10.

Fourth, no detector response was simulated in order to save on computational resources. Although intrinsic peak detector efficiencies were applied via Eq. 3.5, the broadening of the peaks into approximately Gaussian line shapes due to the fundamental detector resolution was not accounted for in the final scaled spectra and extrapolation to experimental measurement times. An analysis of experimental data must therefore integrate the areas under the NRF peaks as a signal rather than solely take the counts in single peak energy bins to achieve the same statistics. Additionally, spectral features seen in real detector systems such as the Compton edge and continuum, the backscatter, escape, and 511 keV peaks, natural background, and pileup have not been simulated. Note that since these spectral features are produced after encryption by the foil, they do not change the information security of the measurement.

Fifth, as outlined in Section 3.3.3 the dose to the warhead is also quite high, on the order of 1 MGy for the measurement used in this work. The dose will also be increased if the foil is thinner, as longer measurement times will be necessary to achieve the same statistics with less NRF production in the foil.

4.2 Future work

The above limitations should be addressed in future work. Further Monte Carlo simulations will be used to determine first the effect of a more realistic cone beam rather than a pencil beam. The protocol's sensitivity to small misalignments or misrotations must be quantified. More elaborate hoaxing scenarios, as well as the diversion of small amounts of special nuclear material, will also be studied. A more realistic model that accounts for the rarer background processes noted in Footnote 24, simulates more complicated isotopic vectors, and includes additional attenuating material in the test object will also be of use. In the last case especially, the computational effort required presents a significant challenge, so further variance reduction techniques, particle tracking cuts, and cross section approximation methods should be investigated.

A further combination of Monte Carlo simulations and analytical methods will be useful in a vulnerability analysis of the protocol. Though the protocol is in principle information-secure due to the underdetermined set of Eqs. 2.52, it may be possible to make approximations that reduce the number of free parameters and divulge approximate values of the classified weapon areal density D or foil areal density X .

Perhaps most importantly, the physical cryptographic protocol must be demonstrated in the laboratory. To this end, an experimental campaign is underway at the High Voltage Research Laboratory (HVRL) at MIT. The HVRL bremsstrahlung beam at $E_{\text{end}} = 2.7$ MeV has been measured using a LYSO scintillator and its energy spectrum found to agree well with the $E_{\text{end}} = 2.7$ MeV simulations used in this work. With the bremsstrahlung source quantified, it will be possible to conduct NRF measurements of uranium-238. Validating the protocol for uranium-235, plutonium-239, and plutonium-240 will require cooperation from the National Laboratories. Other non-fissionable weapon isotopes, such as beryllium in the neutron reflector, copper in the electronics, and iron in the steel structural elements, will require higher- E_{end} accelerators, as the energy of the main NRF transitions increases with decreasing mass number A [64]. A potential facility for these experiments is the 6 MeV electron

accelerator at UC Santa Barbara.

The experimental infrastructure may also allow for the measurement of the poorly-known nuclear level data (the level spins J_r and widths Γ_r) in uranium-235, plutonium-239, and plutonium-240 that determines the calculation of the cross section in Eq. 2.21. By placing detectors at a series of polar angles θ , it should be possible to determine the J_r via a Bayesian estimate using *e.g.* Eqs. 2.45 or 2.46. Given the J_r and the measured integrated cross section, the width Γ_r can be determined as discussed in Section 3.4.

Future systems may also be designed using quasimonochromatic sources, which would selectively generate only photons close to the NRF resonance energies, reducing the total dose to the warhead. Contemporary monochromatic sources typically either exploit the kinematics of interactions to generate (nearly) single-energy photons (*e.g.* [65]) or various filtering mechanisms (*e.g.* [66]) to eliminate photons of all but the desired energy from an initial broad spectrum. In both cases, the sources tend to suffer from low intensities per unit energy relative to bremsstrahlung, and significant work will be required in order to develop a monochromatic system that could be practically used for the physical cryptographic protocol.

4.3 Conclusion

The use of nuclear resonance fluorescence in a physical cryptographic nuclear warhead measurement is a promising development for future treaty verification. The protocol as simulated accepts genuine warheads and rejects hoax warheads with high confidence in realistic measurement times while protecting sensitive weapons design information. Although much work remains to be done in both the technical and political domains before more general conclusions can be drawn and the protocol can be implemented, the results of this thesis provide an encouraging path forward towards eliminating warhead verification as a major obstacle to future nuclear disarmament.

Appendix A

Standalone NRF line database

A standalone database of NRF lines was developed in the course of the Monte Carlo simulations for access to valuable NRF data outside of the Geant4 environment. The nuclear data organization scheme of the G4NRF code (see Section 3.1.1) is first used to print NRF transition data to a separate text file during the initialization phase of the Geant4 simulations. For the full database of $\sim 140\,000$ levels across all isotopes, this takes ~ 15 minutes on a single core of a Late 2011 MacBook Pro with a 2.4 GHz i5 processor. Once the database is generated, it is stored persistently and the database write functionality can be disabled for later Geant4 simulations.

A Python script `standalone.py` then applies filtering and sorting rules to the transition data, reducing the data to a more intelligible format. Available filters include:

1. atomic number Z or minimum and maximum Z_{\min} and Z_{\max} ,
2. mass number A or minimum and maximum A_{\min} and A_{\max} ,
3. minimum and maximum photon energies $E_{\gamma,\min}$ and $E_{\gamma,\max}$,
4. minimum and maximum level energies $E_{r,\min}$ and $E_{r,\max}$,
5. minimum and maximum integrated cross sections $\int \sigma(E) dE|_{\min}$ and $\int \sigma(E) dE|_{\max}$,

while sorting options include:

Z	A	E_r [MeV]	E [MeV]	$\int \sigma(E) dE$ [eV·b]	$\sigma(E = E_r)$ [b]
94	240	2.4332	2.4332	68.98856	33.26208
92	238	2.1760	2.1760	61.12479	32.81645
92	238	2.2090	2.2090	59.31216	31.36758
94	240	2.5775	2.5349	39.30047	17.88748
94	240	2.5775	2.5775	39.30047	17.88748
94	240	2.5664	2.5233	39.09976	17.87310
94	240	2.5664	2.5664	36.09209	16.49825
94	240	2.4332	2.3904	35.53956	17.13501
92	238	2.1760	2.1310	31.78486	17.06453
92	238	2.2090	2.1640	30.84229	16.31112
92	238	2.2450	2.2450	30.69278	15.97176
92	235	1.7693	1.7693	30.69091	20.13663
92	235	1.7335	1.7335	29.80212	19.95684
92	238	2.4680	2.4680	26.91979	12.74264
92	238	2.4100	2.4100	23.43565	11.36038
92	238	1.8460	1.8460	21.83718	13.81967
92	238	1.7820	1.7820	20.89200	13.69636
92	238	2.2450	2.2000	15.96023	8.30531
92	238	2.4680	2.4230	13.45987	6.37131
92	238	2.2950	2.2950	13.45334	6.84826

Table A.1: Table of the 20 strongest NRF lines in the standalone database by integrated cross section with $92 \leq Z \leq 94$ and $1.7 \leq E/\text{MeV} \leq 2.7$. The cross sections quoted are specific to the E and thus are decay rather than absorption cross sections. Note that the peak Doppler-broadened cross sections $\sigma(E = E_r)$ were calculated using $T_{\text{eff}} = 300$ K unadjusted for the Debye temperatures θ and therefore may differ slightly from those computed in the main text.

1. integrated cross section $\int \sigma(E) dE$,
2. photon energy E ,
3. descending or ascending order,
4. the number of transitions to display,

all of which may be specified at the command line via `./standalone.py <options>`. Further documentation is available using the help flag: `./standalone.py -h`. Table A.1 shows example output for the 20 strongest lines of the four isotopes of interest in the energy region used in this work via `./standalone.py --Zmin 92 --Zmax 94 -n 20 --Egmin 1.7 --Egmax 2.7`.

Bibliography

- [1] Federation of American Scientists. Status of world nuclear forces, 2016. Retrieved from <http://fas.org/issues/nuclear-weapons/status-world-nuclear-forces/> on April 1, 2016.
- [2] Treaty between the United States of America and the Russian Federation on strategic offensive reductions (the Moscow treaty), 2002. Retrieved from <http://www.state.gov/t/isn/10527.htm> on April 30, 2016.
- [3] Treaty between the United States of America and the Russian Federation on measures for the further reduction and limitation of strategic offensive arms, 2010. Retrieved from <http://www.state.gov/documents/organization/140035.pdf> on April 30, 2016.
- [4] Russian-U.S. agreement concerning the disposition of highly enriched uranium extracted from nuclear weapons, 1993. Retrieved from <http://www.armscontrol.ru/start/docs/heu93t.htm> on April 12, 2016.
- [5] Theodore A Postol. Possible fatalities from superfires following nuclear attacks in or near urban areas. In Robert Q Marston and Fredric Solomon, editors, *The Medical Implications of Nuclear War*. National Academy Press, 1986.
- [6] Owen B Toon, Alan Robock, and Richard P Turco. Environmental consequences of nuclear war. *Physics Today*, 61:37–42, 2008.
- [7] International Atomic Energy Agency. Information circular: Treaty on the Non-Proliferation of Nuclear Weapons, 1970. Retrieved from <https://www.iaea.org/sites/default/files/publications/documents/infcircs/1970/infcirc140.pdf> on March 31, 2015.
- [8] US Department of Defense. Nuclear posture review report, 2010. Retrieved from http://www.defense.gov/Portals/1/features/defenseReviews/NPR/2010_Nuclear_Posture_Review_Report.pdf on March 31, 2016.
- [9] Treaty between the United States of America and the Union of Soviet Socialist Republics on the limitation of strategic offensive arms, 1979. Retrieved from <http://www.state.gov/www/global/arms/treaties/salt2-1.html> on April 30, 2016.

- [10] Kosta Tsipis. *The Arms Race in an Era of Negotiations*, chapter Nuclear Cruise Missiles: The Problem of Verification, pages 100–114. Palgrave Macmillan UK, London, 1991.
- [11] R. Wolfson. *Nuclear Choices: A Citizen's Guide to Nuclear Technology*. New liberal arts series. MIT Press, 1993.
- [12] US Arms Control and Disarmament Agency. Final Report - Volume I. Field Test FT-34. Demonstrated Destruction of Nuclear Weapons (U), 1969. Declassified March 30, 1999. Retrieved from <http://fas.org/nuke/guide/usa/cloudgap/ft-34.pdf> on April 1, 2016.
- [13] Steve Fetter, Thomas B Cochran, Lee Grodzins, Harvey L Lynch, and Martin S Zucker. Gamma-ray measurements of a Soviet cruise-missile warhead. *Science*, 248(4957):828–834, 1990.
- [14] Thomas B Cochran. The Black Sea experiment. In *From Reykjavik to New START: Science Diplomacy for Nuclear Security in the 21st Century*, 2011. Retrieved from http://docs.nrdc.org/nuclear/files/nuc_11020401a.pdf on May 9, 2015.
- [15] DA Close, DW MacArthur, and NJ Nicholas. Information barriers - a historical perspective. Technical report, Los Alamos National Laboratory, 2001. Retrieved from <http://lib-www.lanl.gov/la-pubs/00796106.pdf> on April 30, 2016.
- [16] Defense Threat Reduction Agency. Technical overview of Fissile Material Transparency Technology Demonstration. Technical report, Los Alamos National Laboratory, 2001. Retrieved from http://www.lanl.gov/orgs/n/n1/FMTTD/presentations/pdf_docs/exec_sum.pdf on April 30, 2016.
- [17] Joseph Menn. Exclusive: Secret contract tied NSA and security industry pioneer. Reuters, 2013. Retrieved from <http://www.reuters.com/article/us-usa-security-rsa-idUSBRE9BJ1C220131220> on May 5, 2016.
- [18] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 18(1):186–208, 1989.
- [19] Manuel Blum, Paul Feldman, and Silvio Micali. Non-interactive zero-knowledge and its applications. In *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing*, STOC '88, pages 103–112, New York, NY, USA, 1988. ACM.
- [20] A. Glaser, B. Barak, and R.J. Goldston. A zero-knowledge protocol for nuclear warhead verification. *Nature*, 510(7506):497–502, 06 2014.

- [21] X-5 Monte Carlo Team. MCNP - Version 5, Vol. I: Overview and Theory. Technical report, 2003. LA-UR-03-1987.
- [22] Ruairidh Macdonald, 2015–2016. Private communications.
- [23] R. Scott Kemp. Private communication.
- [24] U. Kneissl, H.H. Pitz, and A. Zilges. Investigation of nuclear structure by resonance fluorescence scattering. *Progress in Particle and Nuclear Physics*, 37:349 – 433, 1996.
- [25] Franz R Metzger. Resonance fluorescence in nuclei. *Prog. in Nuc. Phys*, 7:54, 1959.
- [26] Richard Fitzpatrick. *Quantum Mechanics*, chapter Resonances. 2010. Retrieved from <http://farside.ph.utexas.edu/teaching/qmech/Quantum/node138.html> on March 4, 2016.
- [27] K.S. Krane. *Introductory Nuclear Physics*, chapter The force between nucleons. Wiley, 1987.
- [28] Neil W Ashcroft and N David Mermin. *Solid state physics*. Holt, Rinehart and Winston, 2005.
- [29] E Browne and JK Tuli. *Nuclear Data Sheets*, 127, 2015.
- [30] Brian Quiter. *Nuclear resonance fluorescence for nuclear materials assay*. PhD thesis, University of California, Berkeley, 2010.
- [31] E Brady and M Deutsch. Angular correlation of successive gamma-rays. *Physical Review*, 78(5):558, 1950.
- [32] Kenneth S. Krane. E2, M1 multipole mixing ratios in even-even nuclei, $A \geq 152$. *Atomic Data and Nuclear Data Tables*, 16(4):383 – 408, 1975.
- [33] Kenneth S. Krane. E2, M1 multipole mixing ratios in odd-mass nuclei, $A > 150$. *Atomic Data and Nuclear Data Tables*, 18(2):137 – 203, 1976.
- [34] J Pruet, DP McNabb, CA Hagmann, FV Hartemann, and CPJ Barty. Detecting clandestine material with nuclear resonance fluorescence. *Journal of Applied Physics*, 99(12):123102, 2006.
- [35] Glenn F Knoll. *Radiation detection and measurement*, chapter Radiation interactions. John Wiley & Sons, Inc., 4 edition, 2010.
- [36] S. Weinberg. *The Quantum Theory of Fields*. Number v. 1 in The Quantum Theory of Fields 3 Volume Hardback Set. Cambridge University Press, 1995.

- [37] MJ Berger, JH Hubbell, SM Seltzer, J Chang, JS Coursey, R Sukumar, DS Zucker, and K Olsen. XCOM: Photon cross sections database. *NIST Standard Reference Database*, 8, 2013.
- [38] William Bertozzi, Richard Hasty, Alexei Klimenko, Stephen E Korbly, Robert J Ledoux, and William Park. Imaging and radiography with nuclear resonance fluorescence and Effective-Z (EZ-3DTM) determination; SNM detection using prompt neutrons from photon induced fission. In *American Institute of Physics Conference Proceedings*, volume 1099. AIP, 2009.
- [39] S Agostinelli, John Allison, K Amako, J Apostolakis, H Araujo, P Arce, M Asai, D Axen, S Banerjee, G Barrand, et al. Geant4—a simulation toolkit. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 506(3):250–303, 2003.
- [40] David V Jordan and Glen A Warren. Simulation of nuclear resonance fluorescence in Geant4. In *Nuclear Science Symposium Conference Record, 2007. NSS'07. IEEE*, volume 2, pages 1185–1190. IEEE, 2007.
- [41] Geant4 Collaboration. Physics reference manual, 2007.
- [42] Charles Kittel. *Introduction to solid state physics*. Wiley, 2005.
- [43] E Browne and JK Tuli. *Nuclear Data Sheets*, 122, 2014.
- [44] E Browne and B Singh. *Nuclear Data Sheets*, 109, 2008.
- [45] Huo Junde, Huo Su, and Yang Dong. *Nuclear Data Sheets*, 112, 2011.
- [46] J K Tuli. Evaluated nuclear structure data file: A manual for preparation of data sets, 1983.
- [47] R. Scott Kemp, Areg Danagoulian, Ruaridh Macdonald, and Jayson Vavrek. Physical cryptographic verification of nuclear warheads. *Proceedings of the National Academy of Sciences*. Under review.
- [48] Ian H Hutchinson. *A student's guide to numerical methods*. Cambridge University Press, 2015.
- [49] RJ McConn Jr, CJ Gesh, RT Pagh, RA Rucker, and RG Williams III. Compendium of material composition data for radiation transport modeling, revision 1. Technical report, Pacific Northwest National Laboratory, 2011.
- [50] Department of Energy. Additional information concerning underground nuclear weapon test of reactor-grade plutonium, 1994. Retrieved from <http://permanent.access.gpo.gov/websites/osti.gov/www.osti.gov/html/osti/opennet/document/press/pc29.html#ZZ0> on May 10, 2016.
- [51] IBA Industrial. Rhodotron[®] E-beam accelerator: Capacity on demand. Technical report, 2014.

- [52] Glenn F Knoll. *Radiation detection and measurement*, chapter Germanium gamma-ray detectors. John Wiley & Sons, Inc., 4 edition, 2010.
- [53] Kari Leppälä, Kirgiz Ssr, and Raimo Verkasalo. Protection of instrument control computers against soft and hard errors and cosmic ray effects. In *International Seminar on Space Scientific Engineering*, 1989.
- [54] David R Meyers. A brief discussion of radiation hardening of CMOS microelectronics, 1998. Retrieved from <http://www.osti.gov/scitech/servlets/purl/2453> on May 8, 2016.
- [55] Brian J Quiter, Bernhard A Ludewigt, Vladimir V Mozin, Cody Wilson, and Steve Korbly. Transmission nuclear resonance fluorescence measurements of ^{238}U in thick targets. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 269(10):1130–1139, 2011.
- [56] William Bertozzi, Joseph A Caggiano, Walter K Hensley, Micah S Johnson, SE Korbly, RJ Ledoux, Dennis P McNabb, EB Norman, William H Park, and Glen A Warren. Nuclear resonance fluorescence excitations near 2 MeV in U-235 and Pu-239. *Physical Review C*, 78(4):041601, 2008.
- [57] B. J. Quiter, T. Laplace, B. A. Ludewigt, S. D. Ambers, B. L. Goldblum, S. Korbly, C. Hicks, and C. Wilson. Nuclear resonance fluorescence in ^{240}Pu . *Phys. Rev. C*, 86:034307, Sep 2012.
- [58] Leon van Dommelen. *Quantum mechanics for engineers*, chapter 14.20.4 Weisskopf estimates. 2013. Retrieved from https://www.eng.fsu.edu/~dommelen/quantum/style_a/ntgd.html on April 28, 2016.
- [59] K.S. Krane. *Introductory Nuclear Physics*, chapter Nuclear models. Wiley, 1987.
- [60] Takehito Hayakawa, Hideaki Ohgaki, Toshiyuki Shizuma, Ryoichi Hajima, Nobuhiro Kikuzawa, Eisuke Minehara, Toshiteru Kii, and Hiroyuki Toyokawa. Nondestructive detection of hidden chemical compounds with laser Compton-scattering gamma rays. *Review of Scientific Instruments*, 80(4), 2009.
- [61] F. Ajzenberg-Selove and J.H. Kelley. *Nuclear Physics A*, 506, 1990.
- [62] F. Ajzenberg-Selove. *Nuclear Physics A*, 523, 1991.
- [63] T. S. Rudisill and M. L. Crowder. Characterization of δ phase plutonium metal. Technical report, Westinghouse Savannah River Company, 2000. Retrieved from <http://sti.srs.gov/fulltext/tr9900448/tr9900448.html> on May 4, 2016.

- [64] Leon van Dommelen. *Quantum mechanics for engineers*, chapter 14.13 Collective structure. 2013. Retrieved from https://www.eng.fsu.edu/~dommelen/quantum/style_a/ntgd.html on April 28, 2016.
- [65] T. Scott Carman, Vladimir Litveninko, John Madey, Charles Neuman, Blaine Norum, Patrick G. O’Shea, N. Russell Roberson, Carol Y. Scarlett, Eric Schreiber, and Henry R. Weller. The TUNL-FELL inverse Compton γ -ray source as a nuclear physics facility. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 378(1–2):1 – 20, 1996.
- [66] J. Jolie and M. Bertschy. A tunable monochromatic gamma-ray source: Part 1. Concept of the source. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 95(3):431 – 436, 1995.

