Experimental demonstration of an isotope-sensitive warhead verification technique using nuclear resonance fluorescence

Citation: Vavrek, Jayson R., Brian S. Henderson, and Areg Danagoulian. "Experimental Demonstration of an Isotope-Sensitive Warhead Verification Technique Using Nuclear Resonance Fluorescence." Proceedings of the National Academy of Sciences 115, no. 17 (April 10, 2018): 4363–4368.

As Published: http://dx.doi.org/10.1073/PNAS.1721278115

Publisher: Proceedings of the National Academy of Sciences

Persistent URL: http://hdl.handle.net/1721.1/119658

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Experimental demonstration of an isotope-sensitive warhead verification technique using nuclear resonance fluorescence

Jayson R. Vavrek[a,1,2], Brian S. Henderson[a,1], and Areg Danagoulian[a]

[a]Laboratory for Nuclear Security and Policy, Massachusetts Institute of Technology, Cambridge, MA 02139

Edited by Roy Schwitters, The University of Texas at Austin, Austin, Texas, and accepted by Editorial Board Member Anthony Leggett March 13, 2018 (received for review December 7, 2017)

Future nuclear arms reduction efforts will require technologies to verify that warheads slated for dismantlement are authentic without revealing any sensitive weapons design information to international inspectors. Despite several decades of research, no technology has met these requirements simultaneously. Recent work by Kemp et al. [Kemp RS, Danagoulian A, Macdonald RR, Vavrek JR (2016) *Proc Natl Acad Sci USA* 113:8618–8623] has produced a novel physical cryptographic verification protocol that approaches this treaty verification problem by exploiting the isotope-specific nature of nuclear resonance fluorescence (NRF) measurements to verify the authenticity of a warhead. To protect sensitive information, the NRF signal from the warhead is convolved with that of an encryption foil that contains key warhead isotopes in amounts unknown to the inspector. The convolved spectrum from a candidate warhead is statistically compared against that from an authenticated template warhead to determine whether the candidate itself is authentic. Here we report on recent proof-of-concept warhead verification experiments conducted at the Massachusetts Institute of Technology. Using high-purity germanium (HPGe) detectors, we measured NRF spectra from the interrogation of proxy "genuine" and "hoax" objects by a 2.52 MeV endpoint bremsstrahlung beam. The observed differences in NRF intensities near 2.2 MeV indicate that the physical cryptographic protocol can distinguish between proxy genuine and hoax objects with high confidence in realistic measurement times.

physical cryptography | nuclear weapons | disarmament | verification

Nuclear arms reduction treaties have traditionally suffered from the disarmament verification problem: How can one confidently identify a warhead as authentic without having access to any sensitive design information that proves it is authentic? Rather than confront this apparent paradox, treaties such as the New Strategic Arms Reduction Treaty (New START) have relied on verification of warhead delivery vehicles—for example, missiles and bomber aircraft—rather than direct verification of warheads themselves. Future arms control agreements, however, may require some mechanism for the verification of individual warheads (1–4) to ensure that a country does not dispose of fraudulent or "hoax" warheads in a gambit to obtain a strategic nuclear advantage.

Warhead Verification

In a warhead verification protocol, a warhead owner ("host") attempts to prove ("inspector") that an object submitted for inspection and subsequent dismantlement and disposition is indeed a genuine nuclear warhead. An object successfully verified may then be dismantled by the host under a secure chain of custody (5) and counted toward the host's obligations under an arms reduction treaty. At the same time, the host seeks to prevent the inspector from learning any sensitive information about the design of the warhead, whether to prevent proliferation of nuclear weapons technology or disclosure of warhead architecture and vulnerabilities. Thus, the verification measurement must be designed and performed in such a way as to provide a strong test of authenticity while minimizing intrusiveness and maximizing information security. Nonauthentic warheads (hoaxes) fall into two broad categories: isotopic hoaxes, in which a valuable weapon component (e.g., the weapons-grade Pu fissile fuel) is replaced by a less valuable surrogate of similar geometry (e.g., reactor grade Pu), and geometric hoaxes, in which isotopes are present in their correct amounts but in a non–weapons-usable configuration (e.g., rough slabs of Pu rather than highly engineered spherical shells).

Past approaches to warhead verification have generally focused on the "attribute" approach, in which the protocol measures a set of key characteristics thought to define a warhead, such as the total mass of plutonium and the isotopic ratio of Pu-239 to Pu-240 in the object (6). Such measurements are highly intrusive and so are conducted behind an "information barrier" (IB), an electronic or software layer that shields the classified raw measurement data and presents the inspector with only a binary pass/fail answer for each of the attribute measurements (7). However, certifying that an electronic or (especially) software IB does not contain any hidden backdoors or functionalities—which a nefarious inspector could exploit to obtain sensitive information or a nefarious host could use to fraudulently simulate a "genuine" result—is exceedingly difficult and may never be satisfactorily proven. Moreover, attributes must be chosen specifically to describe real nuclear warheads and thus may constitute sensitive information themselves. Even then, the set of attributes may not be complete, opening the door to hoax objects that pass all of the attribute tests but nevertheless are not real warheads.

More recent work has therefore focused on the "template" approach to verification, in which comparison with a known genuine object (the template) is used to certify subsequent objects presented for inspection (4, 8, 9). In such a protocol, the measurements of both the template and subsequent objects are

Significance

We present an experimental demonstration of an isotope-sensitive warhead verification protocol. The measurement is capable of detecting tampering with a warhead's material or geometry with high statistical confidence in realistically attainable measurement times, while simultaneously protecting sensitive warhead design information. Such a protocol could enable the verifiable elimination of nuclear warheads under a future arms reduction treaty.

Author contributions: J.R.V., B.S.H., and A.D. designed research, performed research, analyzed data, and wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission. R.S. is a guest editor invited by the Editorial Board.

Data deposition: The data and analysis code for this paper are available at https://github.com/jvavrek/PNAS2018.

[1]J.R.V. and B.S.H. contributed equally to this work.

[2]To whom correspondence should be addressed. Email: jvavrek@mit.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1721278115/-/DCSupplemental.

Published online April 10, 2018.

encrypted using the same method, so that only the encrypted signals (or "hashes") must be compared to authenticate. The hash should be unique to a particular combination of geometry and isotopic makeup (i.e., a particular warhead design), while containing no sensitive information about the object. As such, the hash is useless on its own and only has any use in comparison against the hash of another object—a warhead that is already known to be genuine. This authenticated template warhead could be established for instance via an unannounced visit by the inspector to a random launch facility in the host country and then by selecting a random warhead from an active-duty intercontinental ballistic missile.[†] A measurement of the authenticated template would then be used as the standard against which to compare the measurements from the same model of warhead covered by the arms control treaty.

Recent papers have put forth template verification protocols that aim to make a verification measurement of a warhead while protecting sensitive design information. A team of researchers at Princeton proposed and later experimentally demonstrated a verification protocol using superheated bubble detectors and fast neutron radiography (10, 11). In parallel, a team at the Massachusetts Institute of Technology (MIT) developed an alternative approach using isotopic tomography via transmission nuclear resonance fluorescence (tNRF) (12); the present work is an experimental demonstration of the MIT tNRF protocol. Further techniques using coded aperture-based passive neutron counters (13) and epithermal neutron resonance radiography (14), from Sandia National Laboratories and MIT, respectively, have been proposed in the past year.

The strengths and weaknesses of the aforementioned proposals can be compared by examining the three requirements of an ideal warhead verification protocol:

1. completeness: the ideal protocol must clear all real warheads;
2. soundness: the ideal protocol must raise an alarm on all hoax warheads; and
3. information security: the ideal protocol must be zero knowledge (15, 16)—for an honest host, it must not reveal anything beyond a binary genuine/hoax determination.

The Princeton protocol is essentially zero knowledge, returning a flat image (up to statistical variation) if the host has submitted a real warhead. In its original form (10), the measurement faces a challenge in the soundness requirement: Fast neutron radiography is insensitive to the isotopic or (in some cases) elemental composition of the object and cannot on its own distinguish between weapon materials and well-chosen hoax materials. Additional measurement modes using multiple incident neutron energies (17) have been proposed to increase the protocol's discrimination between fissionable and fissile isotopes. Similarly, work on the Sandia coded aperture protocol has focused on satisfying the completeness and information security aspects of the problem but has not demonstrated resistance to hoaxing by a neutron source of similar geometry and activity.

Unlike the Princeton and Sandia protocols, the two MIT protocols are highly sensitive to isotopics through their use of isotope-specific resonant phenomena, making them highly robust against a large class of hoaxes. While the MIT tNRF protocol is not zero knowledge (since the inspector has access to the hashed measurements rather than solely a binary genuine/hoax determination), and thus there are uncertainties about the extent of the information security of the MIT tNRF protocol, there are methods to make it sufficiently secure (ref. 12). This work demonstrates the core measurement of the MIT tNRF protocol and is an experimental implementation of an isotopically sensitive warhead verification measurement.

NRF Measurements

NRF describes the $\mathrm{X}(\gamma,\gamma')\mathrm{X}$ reaction in which a photon γ is resonantly absorbed by the nucleus X and then re-emitted as the excited nucleus subsequently transitions to its ground state (18, 19). The cross-section for an NRF interaction with absorption via the resonant energy level E_r is given by the Breit–Wigner distribution

$$\sigma_r^{\mathrm{NRF}}(E) = \pi g_r \left(\frac{\hbar c}{E_r}\right)^2 \frac{\Gamma_r \Gamma_{r,0}}{(E-E_r)^2 + (\Gamma_r/2)^2}, \qquad [1]$$

where Γ_r is the width of the level at E_r, $\Gamma_{r,0}$ is the partial width for transitions between E_r and the ground state, and g_r is a statistical factor as described in *SI Appendix*, section S1. For high-Z isotopes of interest, these fundamental widths are typically ~ 10 meV, but the effective width of the cross-section is increased to ~ 1 eV through Doppler broadening by thermal motion of the target nuclei. Imperfect detector resolution further broadens the measurable NRF resolution to widths of ~ 1 keV. Since the NRF lines of an isotope are still typically >10 keV apart, the set of resonance energies E_r provides a resolvable, one-to-one map between measurement space and isotopic space.

The MIT verification protocol exploits the isotope-specific nature of NRF to make a template measurement of the mass and geometry of the isotopes of interest to the inspector. As discussed in the following section and illustrated in Fig. 1, the measurement uses a broad-spectrum bremsstrahlung photon source to irradiate the measurement object; NRF interactions in the object preferentially attenuate the photon flux at specific energies determined by the unique nuclear energy-level structure of each isotope according to how much of the isotope is present in the warhead. The remaining transmitted flux at these energies goes on to induce further NRF interactions in an encryption foil, leading to NRF emission into high-purity germanium (HPGe) photon detectors at an observed rate (*SI Appendix*, Eq. **S8**) that has been reduced by the presence of the NRF isotope in the warhead. The hashed measurements required for the template verification protocol are thus the recorded spectra, since it is impossible to precisely determine the warhead composition (i.e., the thickness D in *SI Appendix*, Eqs. **S7** and **S8**) from the height of the NRF peaks in the observed spectrum without knowledge of the detailed composition of the foil (i.e., the thickness X in *SI Appendix*, Eq. **S8**). The exact foil design is therefore decided by the host and kept secret from the inspector. The influence of the warhead composition on the height of the NRF peaks—and thus any sensitive warhead design information—is then said to be physically encrypted by the foil. This technique uses the laws of physics to mask sensitive information, rather than electronic or computer-based information barriers, making it substantially more robust against tampering and hoaxing than previously proposed techniques (7). Although the detailed construction of the foil is kept secret from the inspector to maintain the encryption, the mere presence of certain characteristic NRF lines in the detected spectrum corresponds to the presence of certain isotopes in the encryption foil, a fact the inspector may use to validate the utility of a given foil without breaking the encryption. The foil may also be placed under joint custody of the host and inspector to ensure it has not been altered between the template and candidate measurements. As an additional layer of information security, the host may add optional "encryption plates" of warhead materials to the measured object so that even if precise inference about the measured object is possible, it is impossible to infer anything about the warhead alone.

To protect against geometric hoaxes, the MIT protocol includes measurements of template and candidate warheads in random or multiple random orientations due to the difficulty for the host to engineer a hoax warhead that could mimic the

[†] In any template warhead verification protocol, the utility of every measurement hinges on the authenticity of the template. A complete solution to the question of first establishing such an authentic template will require classified knowledge of the chain of custody of a country's nuclear stockpile and therefore is an open question beyond the scope of this article.

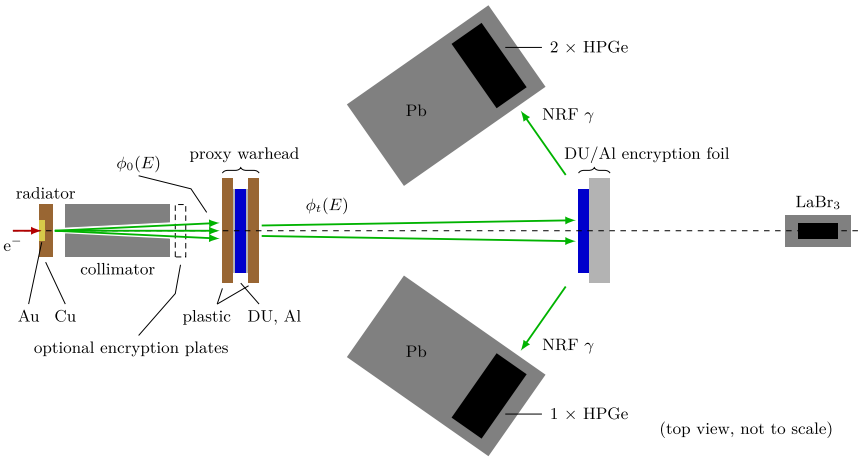

Fig. 1. Schematic of the physical cryptographic NRF measurement. As an information security measure, the large Pb shields prevent the HPGe detectors from directly observing the proxy warhead. Annotated photographs of the experiment geometry are shown in *SI Appendix*, Figs. S1 and S2.

template signal successfully along multiple projections. To increase the information security of this protocol, each orientation may be paired with a unique cryptographic foil to dilute the information content of the multiple measurements. Ref. 12 discusses the required complexity of such geometric hoaxes, which increases rapidly with the number of projections measured.

Experimental Design

Following the design depicted in Fig. 1, a bremsstrahlung beam was used to illuminate a circular section of the object undergoing interrogation. Since no real nuclear warheads were available in an academic setting, several proxy warheads were constructed. The proxy warheads were objects with a set of isotopes—U-238 and Al-27—that form the basis for proof-of-concept NRF experiments and subsequent extrapolations to more realistic settings involving weapon isotopes such as U-235, Pu-239, and Pu-240. The first proxy genuine target ("template I") was constructed from DU plates of total thickness 3.18 mm (wrapped in thin layers of Al foil, amounting to a total thickness of ~ 0.25 mm) encased between two 19 mm-thick layers of high-density polyethylene (HDPE) as proxy high explosives. In the first hoax target ("hoax Ia"), the DU was replaced by 5.29 mm of Pb sheets to match the nominal areal densities of high-Z material to better than 1%. A second experiments to emulate the verification of a genuine candidate warhead ("candidate Ig"). The Pb hoax was similarly remeasured ("hoax Ib"). To emulate measurements on different warhead designs, a second genuine target ("template II") with double the thickness of DU was also tested against a hoax with double the thickness of Pb ("hoax IIc") and against a partial hoax ("hoax IId") in which only half the DU was replaced. In total, seven measurements were conducted on five different targets (see Table 1 and *SI Appendix*, Figs. S10–S14).

Experiments were performed at MIT's High Voltage Research Laboratory (HVRL), which houses a continuous-wave Van de Graaff electron accelerator capable of producing electron kinetic energies of 2.0–3.0 MeV at beam currents of up to 30 μA. For the physical cryptography measurements, a 2.52 MeV electron beam at the maximum stable current (between 25 and 30 μA) was directed toward a water-cooled bremsstrahlung radiator consisting of a 126 μm-thick Au foil and ∼1 cm of Cu backing. The resulting 2.52 MeV endpoint bremsstrahlung photon beam was then collimated with a 20 cm-long conical collimator of entry diameter 9.86 mm and exit diameter 26.72 mm, producing an opening half-angle of about 5°. The beam configuration and stability are discussed in *SI Appendix*, section S2.

Optional encryption plates directly after the collimator may be included as an additional layer of information security. The encryption plates are composed of warhead materials in amounts unknown to the inspector, so that any inference about the warhead composition will in fact be an inference on the warhead plus encryption plates, thus protecting the warhead information. As with the encryption foil, the encryption plates must remain constant between the template and candidate measurements. In these experiments, no such encryption plates were included to maximize the available flux and thus the statistical precision and sensitivity of the measurements.

After passing through the proxy warhead or hoax, the transmitted flux then impinged on the encryption foil, which was constructed from 3.18 mm of DU plates followed by 63.5 mm of aluminum plates. The uranium and aluminum components demonstrate the verification measurement for high- and low-Z materials, respectively. Specifically, the measurements in this work are designed to show the detection of high-Z material diversions and the verification of low-Z material consistency.

The combined NRF signature of the measurement target plus encryption foil—at this point physically encrypted—was measured using three mechanically cooled Ortec 100% relative efficiency GEM P-type coaxial HPGe photon detectors. The detectors were placed ∼45 cm from the foil at an angle of 55° to the beam axis and surrounded by lead to shield against NRF photons directly from the warhead, as well as active backgrounds from the experimental setting, which would otherwise limit the performance of the detectors. The shielding moreover prevents the detectors from observing any passive photon spectra generated by radioactive material in the test objects. The lead shielding thickness ranged from 51 mm below the detectors to 254 to 305 mm along the line of sight from the collimator and warhead to the detectors. Only a 25.4-mm lead filter was placed between the detectors and encryption foil. This reduced by multiple orders of magnitude the low energy photon flux, which can cause pileup and dead time in the detectors, with only a moderate reduction in the NRF signal. Finally, Canberra Lynx Digital Signal Analyzers were used to record the photon spectra

Table 1. Proxy warhead verification measurements

Number	Scenario	Al-27 discrepancy ν (vs. template)	U-238 discrepancy ν (vs. template)
0	Template I	—	—
1	Hoax Ia (100% Pb)	$-0.051\ \sigma$	$7.9\ \sigma$
2	Genuine candidate Ig	$0.76\ \sigma$	$1.7\ \sigma$
3	Hoax Ib (100% Pb)	$1.7\ \sigma$	$9.8\ \sigma$
4	Template II	—	—
5	Hoax IIc (100% Pb)	$0.25\ \sigma$	$10.7\ \sigma$
6	Hoax IId (50% Pb)	$1.9\ \sigma$	$4.6\ \sigma$

ENGINEERING

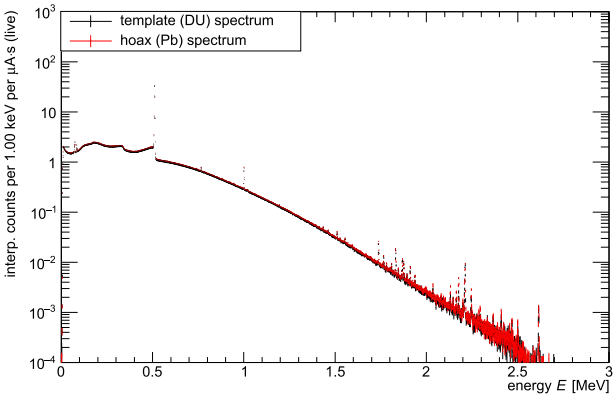

Fig. 2. Measured spectra for DU template II (black points) and Pb hoax IIc (red points). In this and subsequent figures, error bars are ±1 SD

in acquisition periods of 5 min (real time) to save the spectra for offline analysis and to estimate the detector dead time.

A 38.1-mm right square cylinder lanthanum bromide ($LaBr_3$) crystal was placed downstream from the foil as an independent diagnostic for the bremsstrahlung beam flux. It should be emphasized that such additional measurements are not part of the verification protocol. They are, however, useful in an experimental setting for determining the bremsstrahlung endpoint energy of 2.52 MeV (despite the 2.6 MV reading of the accelerator terminal voltage—see *SI Appendix*, section S2.2) as small shifts in electron energy can have a large effect on absolute photon flux (and thus measurement time) near the endpoint. The $LaBr_3$ scintillator was chosen for its extremely fast decay time (16 ns) and encased in a lead hut to avoid high pileup rates that could complicate the endpoint measurement—the detector was directly downbeam from the radiator, otherwise shielded only by the warhead and encryption foil. The detector was controlled using the ROOT-based (20) ADAQAcquisition software (21) and a CAEN DT-5790M digitizer.

Results and Analysis

For each measured object, photon spectra[‡] from multiple acquisition periods and three separate detectors are combined into a single live charge-normalized[§] spectrum to improve the signal-to-noise ratio (see *SI Appendix*, section S3). Each spectrum is then fit with a series of Gaussian functions for the eight observed NRF peaks in the signal region near 2.1 to 2.3 MeV, on top of an exponentially decaying continuum background. U-238 contributes the 2.176, 2.209, and 2.245 MeV peaks; the branched decays 45 keV below each of these three; and a small peak with no branch at 2.146 MeV. Al-27 contributes the intense 2.212 MeV peak. The Pb isotopes have no NRF lines below 2.3 MeV. Altogether, the spectral fitting function is written as

$$f(E) = \exp(c_1 + c_2 E) + \sum_{k=1}^{8} \frac{a_k}{\sqrt{2\pi}\sigma_k} \exp\left[-\frac{(E-E_k)^2}{2\sigma_k^2}\right], \quad [2]$$

where c_1 and c_2 describe the shape of the continuum, and a_k, E_k, and σ_k are the area, mean, and SD fit parameters of the k^{th} peak. With eight sets of three peak parameters and two parameters for the continuum, this results in a total of 26 parameters per spectrum.

Once the 26-parameter fit (and set of associated fit parameter uncertainties) for each spectrum is computed using Eq. **2**, the detected NRF rate in each peak in counts per live μA·s, as predicted by integration of *SI Appendix*, Eq. **S8**, can be extracted as simply $A_k = a_k/\Delta E$, where a_k is the value of the area fit parameter for the k^{th} peak, and the division by the spectrum bin width ΔE enforces proper dimensions and normalization (ref. 22, p. 171). Similarly, the uncertainty in the NRF rate is $\delta A_k = \delta a_k/\Delta E$ (where δx is used to express the 1 SD uncertainty in a value x so as to distinguish it from other uses of the symbol σ) where δa_k is the uncertainty in the a_k fit parameter as reported by ROOT's TH1::Fit() subroutines (20).

One possible test statistic T for comparing the NRF peaks of a single isotope is the sum of net rates A_k (above the fit background) of the six U-238 peaks well-separated from the doublet: 2.131, 2.146, 2.164, 2.176, 2.200, and 2.245 MeV. The 2.209 MeV component of the 2.209 and 2.212 MeV doublet tends to have a larger uncertainty such that it does not contribute reliably to T and thus is excluded. Moreover, since the amount of Al-27 (and the total high-Z areal density) does not change between the warhead and hoax objects, the 2.212 MeV peak rate is consistent throughout the measurements (up to day-to-day beam variations—see *SI Appendix*, section S2.2). To compare the NRF spectrum of a candidate object to that of the genuine template, the discrepancy ν is defined as the difference in T divided by the uncertainty in the difference:

$$\nu \equiv \frac{T_{\text{cand}} - T_{\text{temp}}}{\sqrt{(\delta T_{\text{cand}})^2 + (\delta T_{\text{temp}})^2}}. \quad [3]$$

As the presence of an NRF isotope in the object reduces the corresponding observed NRF rate (and thus T), $\nu > 0$ indicates a possible diversion of the isotope in the candidate compared with the template, while $\nu < 0$ indicates a possible addition. Under the null hypothesis that the candidate object is a real nuclear warhead, $T_{\text{cand}} = T_{\text{temp}}$, so that (due to statistics alone) ν is normally distributed with mean 0 and SD 1: $\nu \sim \mathcal{N}(0,1)$. As such, ν measures the discrepancy from the null hypothesis in the number of standard deviations ("sigmas") where, for example, the probability of observing a $5\,\sigma$ discrepancy (regardless of sign) by chance alone—that is, $|\nu| > 5$—is 6×10^{-7}. Setting an alarm threshold $|\nu| > \nu^*$ by necessity trades-off the probability that the measurement declares a genuine warhead to be a hoax (type I error) and the probability that it declares a hoax warhead to be genuine (type II error). If low type I error is prioritized, a suitable alarm threshold may be $\nu^* = 5$, while $\nu^* = 3$ may be more suitable if low type II error is desired.

Figs. 2–4 show the culmination of the above analysis procedure for the fourth verification scenario listed in Table 1 (template II

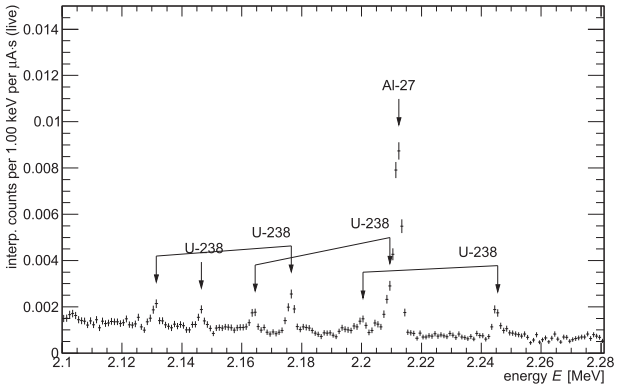

Fig. 3. Measured spectra for DU template II, zoomed to show the NRF signal region. For clarity, the spectrum of hoax IIc is not shown. Arrows indicate the branching relationships from the three main U-238 lines to the peaks 45 keV lower as well as the nonbranching 2.146 MeV U-238 and 2.212 MeV Al-27 peaks.

[‡] Data and analysis code are available at https://github.com/jvavrek/PNAS2018.

[§] The term "live" is used to denote measurement times calculated using live time—that is, the real time minus the detector's dead time. "Live charge" therefore corresponds to the product of beam current with live time.

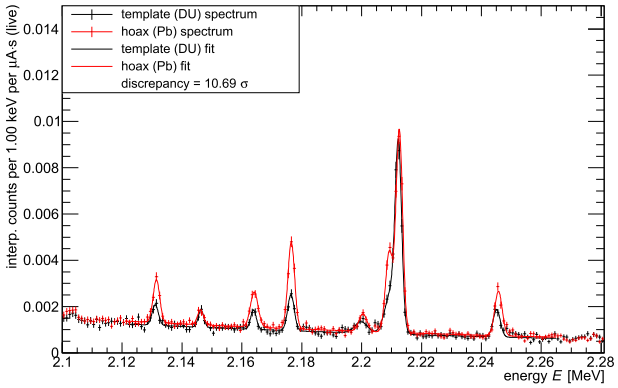

Fig. 4. 26-parameter Gaussian peak plus exponential background fits to the spectra of template II (black points and curve) and hoax IIc (red points and curve). A comparison of spectra for all verification measurements is shown in Table 1.

vs. hoax IIc). Fig. 2 contains the two combined spectra measured for the template II (DU) and hoax IIc (Pb) proxy warheads; in this unzoomed energy range, the genuine and hoax spectra at first appear to match quite closely, with no obvious distinguishing features. Focusing on the NRF signal region in Fig. 3 (where only the template II spectrum is shown for clarity), the NRF peaks from U-238 and Al-27 become visible; Fig. 4 subsequently shows the 26-parameter fits to the two spectra and the computed discrepancy of $\nu = 10.7$. The discrepancies for all verification scenarios are shown in Table 1 (see also *SI Appendix*, Table S1). In all four hoax scenarios, a discrepancy in T greater than an alarm threshold of $\nu^* = 3$ was attained in ~ 20 μA·h (live, on three detectors) per measured object, indicating diversions in the uranium component. In the genuine candidate scenario, the $1.7\,\sigma$ discrepancy in uranium (primarily a result of day-to-day beam variations) does not trigger the alarm at $\nu^* = 3$ and is clearly delineated from the much larger observed discrepancies in the hoax cases. Similarly, the Al-27 comparisons all exhibit $|\nu| < 2$, indicating consistency in the aluminum component across all measurement scenarios.

The continua underlying the peaks—generated from both pileup and secondary electron bremsstrahlung in the foil—also provide some insight. For the spectra in Fig. 2, the integrals from 1–2 MeV differ by 5%. The differences rise to 6 to 10% when comparing measurements performed on different days due to beam variations but are only 2% in the other same-day measurements. Though these small differences are significant given the high statistics at low energies, the close matching of continua between the template and hoax scenarios suggests that the continuum background may not encode any appreciable information about the isotopic content of the weapon. This lack of distinguishing information in the majority of the spectrum also may indicate that nonresonant photon transmission measurements such as radiographs would likely fail to detect hoaxes of the same areal density.

Discussion

Extrapolation to Future Systems. The proxy warheads used in this article are relatively thin—templates I and II have total on-axis areal densities of ~ 11 and 17 g/cm^2, respectively—and do not accurately represent typical areal densities of real warheads. More realistic warhead models in the open literature range from the compact (~ 50 g/cm^2) Black Sea-type warhead used in ref. 12 to the thicker (~ 200 g/cm^2) models of Fetter et al. (23). Moreover, verification measurements conducted on real warheads will use the NRF lines associated with the fuel isotopes U-235 or Pu-239, whose strongest lines are 2 to 5× less intense than the U-238 lines considered in this work (24). In the present experimental design, verification of realistic weapon designs would therefore require several orders of magnitude longer measurement times than the ~ 60 detector live μA·h used here (see Table 2). In a dedicated warhead verification facility, however, these unrealistically long measurement times could be ameliorated by increasing the electron beam current and the number of detectors (see *SI Appendix*, section S4). A modern commercially available electron accelerator may have a continuous wave beam current of at least 5 mA at ~3 MeV (25), a factor of 200× improvement over the 25 μA used in this experiment. While this increase in beam current would affect the event rate in the detectors (thus reducing the effective live time), the increase in the event rate is only ~30× due to the increased attenuation of realistic warheads (see *SI Appendix*, section S4). This increase may be mitigated by reducing the detector sizes and the operating with more detectors, by optimizing the balance of the detector event rate and the available measurement time, or by taking advantage of future developments of high-rate HPGe detectors capable of operating at MHz rates (26). Additionally, the shielding and detector filters used in this experiment could be significantly optimized to reduce the low-energy photon rate in the HPGe detectors to further alleviate this effect. Extending the array of detectors from 3 to 30 would provide another factor of 10× reduction in measurement time and would provide the additional benefit of reducing the dose to the warhead—here estimated at ~ 30 Gy per 1 h measurement at 25 μA for template I—required to achieve the same confidence. Doses for other warhead configurations are presented in ref. 12. Such a dedicated verification system would be capable of attaining $5\,\sigma$ confidence in a single NRF line in a Pb hoax scenario involving the Fetter et al. uranium–uranium model of Table 2 in ~ 15 to 20 min per projection per object, for a capital cost on the order of US$5 million. This required runtime increases to ~ 5 h for the worst-case plutonium–uranium model in Table 2. For thinner warheads, or for warheads that have been partially disassembled, even lower measurement times would be required, creating opportunities for measurements at multiple warhead orientations, for measurements of isotopes with weaker NRF lines, or for ruling out less discernible hoaxes. More information on the calculation of the required runtimes for realistic warhead configurations may be found in *SI Appendix*, section S4.

Information Security. The equation for the predicted NRF rates (*SI Appendix*, Eq. **S8** or its integrated form) contains multiple quantities that are kept secret from the inspector and thus cannot be used alone to infer the warhead thickness D from a physically encrypted spectrum. However, it may be possible to construct a system of equations from *SI Appendix*, Eq. **S8**—one equation per NRF peak—and make a series of simplifying approximations, in which case there may be at least as many equations as unknowns and inference may be possible. As previously shown in Fig. 1 and described in ref. 12, a solution to this multiline inference problem is to include optional encryption plates of relevant materials of unknown thickness ΔD to the collimator output. As such, any inference on the isotope of interest will estimate only an upper bound $D + \Delta D$. In fact, if such encryption plates are used, the foil parameter X no longer needs to be kept secret from the inspector, eliminating the information security complexities of ensuring that the foil has not been nefariously designed.

Table 2. Warhead geometries and approximate detector live charges required for Pb replacement hoax detection at 5 σ

Comparison (model ref.)	NRF line	Foil	Detector live μA·h
WGU + W vs. Pb + W (23)	U-235 1.733 MeV	WGU	25×10^3
WGU + DU vs. Pb + DU (23)	U-235 1.733 MeV	WGU	40×10^3
WGPu + W vs. Pb + W (23)	Pu-239 2.431 MeV	WGPu	600×10^3
WGPu + DU vs. Pb + DU (23)	Pu-239 2.431 MeV	WGPu	800×10^3
WGU vs. Pb (12)	U-235 1.733 MeV	WGU	0.15×10^3
WGPu vs. Pb (12)	Pu-239 2.431 MeV	WGPu	3.5×10^3

ENGINEERING

Lastly, the continuum background may contain sensitive information, especially given the large number of photons it comprises over the entire range of the spectrum (see Fig. 2). The "logarithmic slope" parameter c_2 in Eq. **2**, for instance, depends moderately on the atomic number Z of the foil materials (27). As discussed above, however, the continuum appears to encode very little information about the Z of the warhead materials for a fixed areal density. A thorough analysis of the continuum information content is therefore a vital next step in the analysis of the physical cryptographic NRF protocol. For a more complete discussion of information security issues and possible solutions, the reader is referred to ref. 12.

Conclusions and Future Work

We have reported on the successful demonstration of the MIT tNRF physical cryptographic warhead verification protocol. The isotope-sensitive tNRF measurement is capable of distinguishing proxy nuclear warheads from hoax objects with high confidence in total measurement times of around 1 h per object. Extrapolations to more realistic warhead designs indicate that a dedicated warhead verification facility could conduct $5\,\sigma$ verification measurements in less than an hour while protecting sensitive warhead design information.

The NRF verification technique may be expanded to other isotopes that may be found in nuclear weapons (beyond U-238 and Al-27) such as U-235 or Pu-239 in the fissile fuel and nitrogen and carbon isotopes in the high explosives (28). Similarly, testing the measurement's sensitivity to geometric hoaxes would be a useful development. Finally, an additional layer of information security may be added through analog-to-digital converters (ADCs) with nonuniform binning, which are currently being developed. Such ADCs would act as very low-level, more easily verifiable information barriers. If installed in the acquisition systems of the HPGe detectors, such ADCs could be used to remove all spectral features except one NRF line per isotope from the observed spectrum, thus eliminating possible information security concerns such as the continuum and the multiline inference problem.

In a broader context, the implementation of any warhead verification protocol in a real arms control agreement faces two challenges. First, an assessment of the protocol's utility and security must be made by nuclear weapons laboratories. To this end, future work on any warhead verification protocol should involve collaboration with the US and possibly Russian national laboratories, and possibly combining multiple proposed verification techniques as part of an overarching protocol. Such a joint effort will enable research that otherwise could not be conducted in academic settings, such as the aforementioned measurements involving weapons isotopes and realistic, classified weapon geometries. Finally, the implementation of a warhead verification protocol is predicated on the existence of future arms control frameworks and thus requires a commitment to the goal of deep reductions in the world's nuclear arsenals.

ACKNOWLEDGMENTS. The authors thank Chathan Cooke for operating the HVRL accelerator, Charles Epstein and Richard Milner for their electron beam diagnostics, Rob Goldston for useful discussions, and R. Scott Kemp for offering comments on the manuscript. This work was funded in part by the Consortium for Verification Technology under Department of Energy National Nuclear Security Administration Award DE-NA0002534. B.S.H. gratefully acknowledges the support of the Stanton Foundation's Nuclear Security Fellowship program.

1. Drell S, Callan C, Cornwall J, Dyson F, Eardley D (1993) *Verification of Dismantlement of Nuclear Warheads and Controls on Nuclear Materials* (MITRE Corp JASON Program Office, McLean, VA), Technical Report JSR-92–331.
2. Comley C, Comley M, Eggins P (2000) *Confidence, Security and Verification: The Challenge of Global Nuclear Weapons Arms Control* (United Kingdom Ministry of Defence, Whitehall, London), Technical Report AWE/TR/2000/001.
3. Spears D (2001) *Technology R&D for Arms Control: Arms Control and Nonproliferation Technologies* (US Department of Energy, National Nuclear Security Administration, Defense Nuclear Nonproliferation Programs, Washington, DC).
4. Fuller J (2010) Verification on the road to zero: Issues for nuclear warhead dismantlement. *Arms Control Today* 40:19–27.
5. Bunch KJ, Jones M, Ramuhalli P, Benz J, Denlinger LS (2014) Supporting technology for chain of custody of nuclear weapons and materials throughout the dismantlement and disposition processes. *Sci Global Secur* 22:111–134.
6. Defense Threat Reduction Agency (2001) *Technical Overview of Fissile Material Transparency Technology Demonstration* (Los Alamos National Laboratory, Los Alamos, NM), Technical Report. Available at www.lanl.gov/orgs/n/n1/FMTTD/presentations/pdf_docs/exec_sum.pdf. Accessed December 4, 2017.
7. Close D, MacArthur D, Nicholas N (2001) Information Barriers—A Historical Perspective (Los Alamos National Laboratory, Los Alamos, NM), Technical Report LA-UR-01–2180. Available at lib-_www.lanl.gov/la-_pubs/00796106.pdf. Accessed December 4, 2017.
8. Yan J, Glaser A (2015) Nuclear warhead verification: A review of attribute and template systems. *Sci Global Secur* 23:157–170.
9. Marleau P, et al. (2015) Report on a zero-knowledge protocol tabletop exercise (Sandia National Laboratories, Livermore, CA, Los Alamos National Laboratories, Los Alamos, NM), Technical Report SAND2015–5075.
10. Glaser A, Barak B, Goldston R (2014) A zero-knowledge protocol for nuclear warhead verification. *Nature* 510:497–502.
11. Philippe S, Goldston RJ, Glaser A, d'Errico F (2016) A physical zero-knowledge object-comparison system for nuclear warhead verification. *Nat Commun* 7:12890–12896.
12. Kemp RS, Danagoulian A, Macdonald RR, Vavrek JR (2016) Physical cryptographic verification of nuclear warheads. *Proc Natl Acad Sci USA* 113:8618–8623.
13. Marleau P, Krentz-Wee R (2017) Investigation into Practical Implementations of a Zero Knowledge Protocol (Sandia National Laboratories, Livermore, CA), Technical Report SAND2017–1649.
14. Hecla JJ, Danagoulian A (2017) Nuclear disarmament verification via resonant phenomena. arXiv:1709.09736.
15. Goldwasser S, Micali S, Rackoff C (1989) The knowledge complexity of interactive proof systems. *SIAM J Comput* 18:186–208.
16. Blum M, Feldman P, Micali S (1988) Non-interactive zero-knowledge and its applications. *Proceedings of the 20th Annual Association for Computing Machinery Symposium on Theory of Computing* (ACM, Chicago), pp 103–112.
17. Yan J, Glaser A (2015) Two-color neutron detection for zero-knowledge nuclear warhead verification. *Proceedings of 56th Annual INMM Meeting* Available at www.princeton.edu/∼aglaser/PU105-_Yan-_Jie-_Glaser-_2015.pdf. Accessed October 23, 2017.
18. Metzger F (1959) Resonance fluorescence in nuclei. *Prog Nucl Phys* 7:54–88.
19. Kneissl U, Pitz H, Zilges A (1996) Investigation of nuclear structure by resonance fluorescence scattering. *Prog Part Nucl Phys* 37:349– 433.
20. Brun R, Rademakers F (1997) ROOT—An object oriented data analysis framework, *Nucl Instrum Methods Phys Res Sect A* 389:81–86.
21. Hartwig ZS (2016) The ADAQ framework: An integrated toolkit for data acquisition and analysis with real and simulated radiation detectors. *Nucl Instrum Methods Phys Res Sect A* 815:42–49.
22. Bevington PR, Robinson DK (2003) *Data Reduction and Error Analysis for the Physical Sciences* (McGraw-Hill, New York).
23. Fetter S, et al. (1990) Detecting nuclear warheads. *Sci Global Secur* 1:225–253.
24. Bertozzi W, et al. (2008) Nuclear resonance fluorescence excitations near 2 MeV in ^{235}U and ^{239}Pu. *Phys Rev C* 78:041601.
25. IBA Industrial (2017) IBA Dynamitron. Available at www.iba-_industrial.com/accelerators#dynamitron-_e-_beam-_accelerator. Accessed October 20, 2017.
26. Fast J, et al. (2013) Performance of the Ultra-High Rate Germanium (UHRGe) System (Pacific Northwest National Laboratory, Richland, WA), Technical Report PNNL-23084.
27. Bertozzi W, Korbly SE, Ledoux RJ, Park W (2007) Nuclear resonance fluorescence and effective *Z* determination applied to detection and imaging of special nuclear material, explosives, toxic substances and contraband. *Nucl Instrum Methods Phys Res Sect B* 261:331–336.
28. Caggiano JA, et al. (2007) Nuclear resonance fluorescence measurements of high explosives. *Nuclear Science Symposium Conference Record, 2007. NSS'07* (IEEE, Honolulu, HI), Vol 3, pp 2045–2046.