Toward a Secure Inspection System for Nuclear Warhead Verification Without Information Barrier

Alexander Glaser,* Boaz Barak,† and Robert J. Goldston‡

*Princeton University

†Microsoft Research New England

‡Princeton Plasma Physics Laboratory and Princeton University

ABSTRACT. We previously proposed an approach to nuclear warhead verification envisioning an inspection system that a priori avoids detection of sensitive information, using a so-called zero-knowledge protocol. Under such a protocol, the host can prove to an inspector that a warhead is authentic without revealing anything about its materials or design. The challenge remains, however, to demonstrate a practical implementation of such a system that can detect relevant violations and avoids even the possibility of snooping on electronic measurements as they are made. In this article, we examine the use of superheated drop (or "bubble") detectors to detect neutrons from active interrogation of an unclassified test object with 14.1-MeV neutrons. Zero-knowledge is achieved by the host pre-loading individual detectors so that they are "topped up" by the measurement itself to a previously agreed-upon, unclassified reference value. The required preloads are determined by the host prior to the inspection, but remain unknown to the inspector. The viability of the method is examined with MCNP5 Monte Carlo neutron transport calculations modeling the experimental setup, an investigation of a diversion scenario, and a mathematical analysis of the detected data.

Background

Existing nuclear arms-control agreements between the United States and Russia place limits on the number of *deployed* strategic nuclear weapons. Verification of these agreements can take advantage of the fact that deployed weapons are associated with unique and easily accountable delivery platforms, i.e., missile silos, submarines, and strategic bombers. The next round of nuclear arms-control negotiations, however, may begin also to include tactical weapons and non-deployed weapons. Both would require fundamentally new verification approaches, including authentication of nuclear warheads in storage and authentication of warheads entering the dismantlement queue. Dedicated inspection systems using radiation measurement techniques are likely to play a critical role in verifying such agreements, and different approaches have been proposed since the 1990s to accomplish this task.[1] In this context, the so-called template method is generally considered the most robust verification approach. It envisions the comparison of a complex fingerprint of an inspected item against the fingerprint of a reference item, or template, to confirm that both items are substantially identical. Radiation measurements on classified items would themselves be highly classified, however, and all proposed inspection systems thus far have had to rely on engineered information barriers to protect this data. In contrast, we propose a template verification approach that follows a zero-knowledge protocol under which a statement can be proven to be true without revealing why the statement is true.[2] This paper is an extension of our previous conceptual work in this area.[3]

Experimental Setup

We propose to use 14.1-MeV neutrons from a DT neutron generator to interrogate a test item in a staging area, allowing detailed radiographic imaging and also measurements of neutron intensities at large angles. Neutrons are collimated by 60 cm of polyethylene and illuminate a spherical container with the test item. An array of neutron detectors is placed in an appropriate position to make a radiographic image (Figure 1). Additional detectors (not shown) would be positioned at large angles to the beam to measure scattered neutrons and neutrons from fission events in test items containing nuclear material.

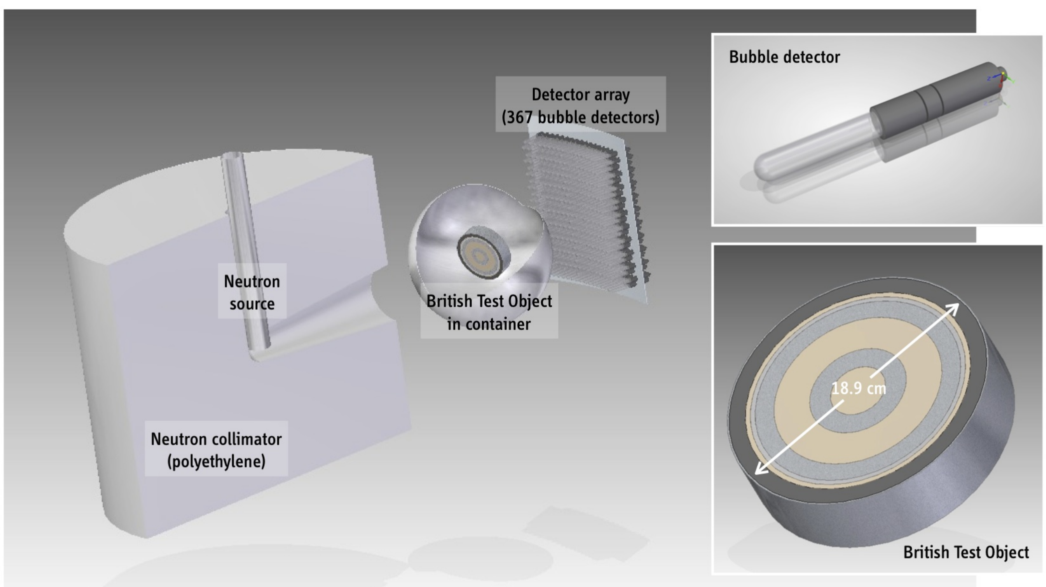

Figure 1. Experimental setup with neutron source, test item in container, and detector array (left). Insets show typical bubble detectors (top) and the British Test Object (bottom), which has an outer diameter of 18.9 cm and contains concentric rings of different materials including 7.75 kg of tungsten. Large-angle detectors are not shown. *3D models: Sébastien Philippe, Princeton University; Bubble Detectors: Bubble Technology Industries.*

Test item. The test item used for this project is a "British Test Object" (BTO), which consists of concentric shells of different materials, including polystyrene, tungsten, aluminum, graphite, and steel. The item does not contain special or other nuclear materials, but is used to develop and calibrate imaging systems for diagnostic analysis of nuclear weapons.[4] In the proposed setup, the BTO is suspended in a spherical aluminum container in order to avoid revealing to the inspector the initial orientation and, after arbitrary rotations during the inspection, therefore also any subsequent orientations of the test item inside the container.

Superheated Drop Detectors. Detectors that contain electronics are generally considered vulnerable to tampering and even snooping and are therefore difficult to certify and authenticate with high confidence. To overcome this problem, we propose using a detector technology that does not rely on electronic components at all; instead, we choose to measure neutron fluences using superheated drop or "bubble" detectors.[5] These neutron detectors are insensitive to gammas and cannot be used to measure neutron multiplicities, providing only information on neutron fluence. Their proper operation can be checked at any time by exposing them to an appropriate calibrated neutron source. Detector noise due to neutrons scattered from the surrounding environment ("room return") can be minimized by employing detectors that are sensitive only to neutrons above a specified energy threshold.

Two general types of measurements can be distinguished: direct transmission measurements, which produce a radiograph of the test item; and measurements at large angles, which detect scattered and fission neutrons and are particularly sensitive to material substitutions.[6] Only the radiographic data are discussed in further detail below. For the radiographic analysis, we work with a hexagonal array of 367 detectors. Information contained in these measurements is highly sensitive and must not leak. Below, we propose that such measurements, in effect, are only made on a differential basis, and so provide exactly zero information on the test items themselves if they are all identical. The proposed inspection protocol is summarized below before the results of a simulated inspection are discussed.

Inspection Protocol

Figure 2 illustrates the main elements of the inspection protocol. In our reference approach using bubble detectors, the key principle is the following:

> *Any measurement on the template or a valid test item will produce a previously agreed number of counts N_{max} with a Poisson distribution of statistical noise. Neither the signal nor the noise contain any information about the inspected item.*

Since N_{max} is known in advance to both sides, this means that the measurements do not reveal any new information. The reader might wonder why make these measurements if the result are predictable, but we will show that if there actually were a diversion then the measurements would differ from N_{max}. As a consequence, the host is only guaranteed confidentiality if he follows the protocol, which includes using substantially identical objects for the template and test items.

The main steps of the proposed protocol are as follows:[7]

1. The items offered for inspection are selected; typically, these will include warheads or warhead components already in storage. One or more templates will be randomly selected at a deployment site or sites in order to maximize confidence of their authenticity. All other items will be compared against the template(s). In the following, for simplicity, we assume one test item and one template, but the generalization to multiples of each is straightforward.

2. All items are placed in sealed storage containers prior to inspection and brought to a dedicated dismantlement facility using strong chain-of-custody protocols.

Figure 2. Proposed inspection protocol; see text for details.
Image source/credit: U.S. Department of Defense (top) and Paul Shambroom (bottom)

3. The inspector and the host agree on a reference bubble count N_{max}. This number can be established for the given inspection system and the allotted measurement time by the maximum bubble count that would be observed on any detector in the absence of any object in the measurement beam. N_{max} has to be large enough to be sensitive to meaningful diversion scenarios, but is limited by the size and design of the bubble detectors used in the inspection. If necessary, multiple detectors can be irradiated sequentially to increase the effective N_{max}.

4. The inspector announces the orientations of the container in front of the beam and detector positions that she intends to measure.

5. Since the host has complete knowledge of the warhead design, and all details of the measurement to be carried out, he can project the measurement values for all container orientations and detector positions. The host can generate the required database in

advance, for example, by making extensive measurements on real warheads and by running detailed computer simulations similar to those discussed below. The host can even make confirming measurements in the real test facility.

6. Knowing the expected measurement value N_{exp} for a particular orientation and direction, the host now offers preloaded sets of bubble detectors for use in the actual inspection so that the detectors will be "topped up" to N_{max} during the measurement, i.e., $N_{pre} + N_{exp} = N_{max}$. This preload must itself include Poisson statistical noise, so that the total measurement is Poisson distributed. The preload N_{pre} has to be considered as sensitive as the actual warhead design, and under no circumstance is the inspector allowed to have access to the preloaded bubble detectors prior to the measurement.[8] Detectors can be sealed inside of opaque covers, e.g. wrapped in black tape, or include other design features to prevent reading preloads when they are presented for an inspection.

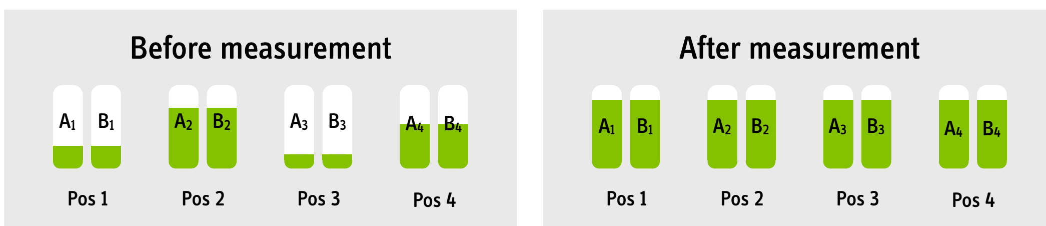

Figure 3. Pairs of bubble detectors before and after measurement. Depending on the orientation of the test item and the position in which the detectors will be used, the host preloads detectors pairwise so that every detector is "topped up" to N_{max} during the measurement. No information about the inspected item is revealed in the process.

7. Crucially, for every pair of detectors offered by the host, the inspector chooses which detector to use on the test item and which to use on the template. This strategy makes it impossible for the host to conceal a spoof by unequally initializing the detectors.

8. Measurements are carried out on both the template and the test item.[9]

9. If the test item and the template are substantially identical, then all detectors will on average have a bubble count of N_{max} at the end of the measurement, with an overall Poisson distribution arising from statistical noise. Host and inspecting party can agree on permissible deviations, which may be due to systematic and statistical measurement errors but also manufacturing tolerances, as part of the agreement on the specifics of the inspection protocol. In general, the optimum number of measurements will be a tradeoff between the speed and the sensitivity of the process; it will be further constrained by the maximum permissible neutron load on the warheads, although our estimate is that this will not be significant constraint.

Results of Monte Carlo Neutron Transport Simulations

To examine the prospects of successfully implementing the proposed verification approach, we have analyzed the experimental setup with a series of MCNP5 Monte Carlo simulations exploring a basic "diversion scenario" in which the tungsten rings of the BTO are replaced by lead rings of identical dimensions. In the example discussed below, the detectors are sensitive to neutron energies of 1 MeV and higher. Room return is not included in the calculation. The modeled maximum bubble count per detector is $N_{max} = 1000$, which provides acceptable counting statistics and should be technically achievable.[10] Figure 4 illustrates typical results for a valid item (template match) and an invalid item; for reference purposes, the radiograph of the test item is also shown, but this data is *never measured* in the inspection process because of the use of the preloaded detectors.

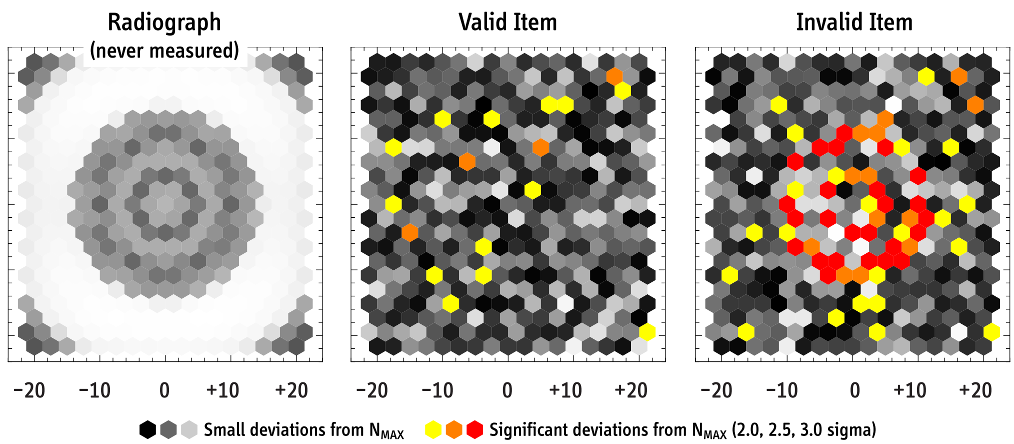

Figure 4. Results of MCNP5 simulations for a basic diversion scenario. The radiograph of the test item shown on the left is never measured, i.e., corresponds to a measurement without preloading the detectors. The other panels show total detector counts after measurements on a valid and an invalid item. Shades of gray and colors indicate absolute differences from $N_{max} = 1000$. The invalid item produces a larger number of suspicious data points, which are in this case also spatially correlated. *Concept: Charles Guo, Princeton University.*

As expected, in the case of the valid item, the bubble counts are distributed consistent with a Poisson distribution with expectation value N_{max}, and therefore root-mean-square deviation $\sqrt{N_{max}}$ (Figure 5). The measurements on the invalid item reveal significant differences from the expected distribution, which are clearly distinguishable in this simple diversion scenario. More sophisticated tests can be performed on the measured data to examine the consistency of the data with the expected distribution.

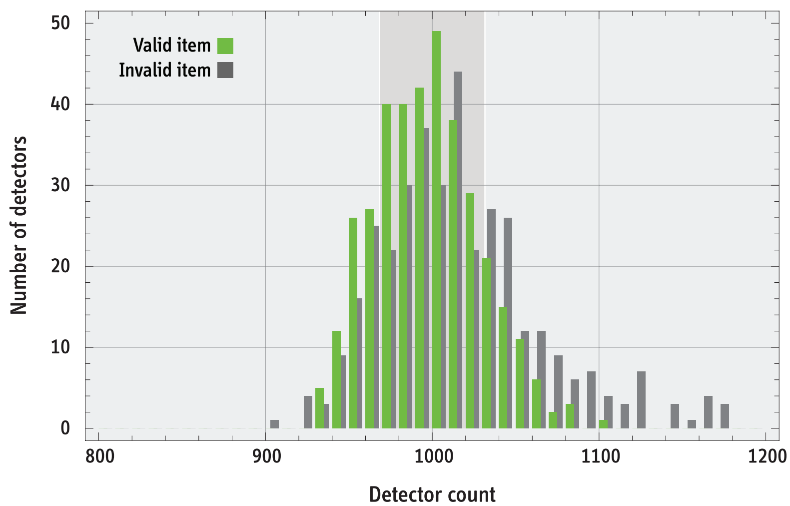

Figure 5. Histogram of detector counts. Detector counts for the valid item are consistent with a Poisson distribution with a mean of N_{max} and a root-mean square deviation of $\sqrt{N_{max}}$; in other words, there is no information in the signal or its noise. The distribution of counts for the invalid item is shifted towards higher values and has a skewed positive tail. The data collected for the invalid item does not pass goodness-of-fit tests.

Here, we briefly explore the question of how many data points have to be generated so that false-positive and false-negative rates are below acceptable mutually-agreed levels.

In the case that there is no cheating, the measurements should correspond to a Poisson distribution with mean of N_{max}. There are a number of well-established goodness-of-fit tests for comparing measured distributions with ideal ones.[11] To compare these tests with each other, we choose a threshold for flagging items as invalid such that exactly 5% of the valid items fail a particular test as false positives. With this calibration, Figure 6 shows the corresponding false-negative rates (i.e., the probability that an invalid item is flagged as a valid one) for several standard tests as a function of the number of detectors used. All flagged items would be retested to eliminate the false positives. This process should converge rapidly to identifying invalid items with high confidence.

A simple square-difference test, in which the mean of the squared differences from N_{max} of the measured distribution is compared with the expected value N_{max}, is remarkably effective for this diversion case, giving less than 1% false negatives with only 60–65 detectors. With more detectors, there should be room to reduce both the false-negative and false-positive rates. Future research will include examination of the most effective distribution tests under various diversion scenarios, as well as studies of the most efficient means of maximizing the discrimination power of the retesting process.

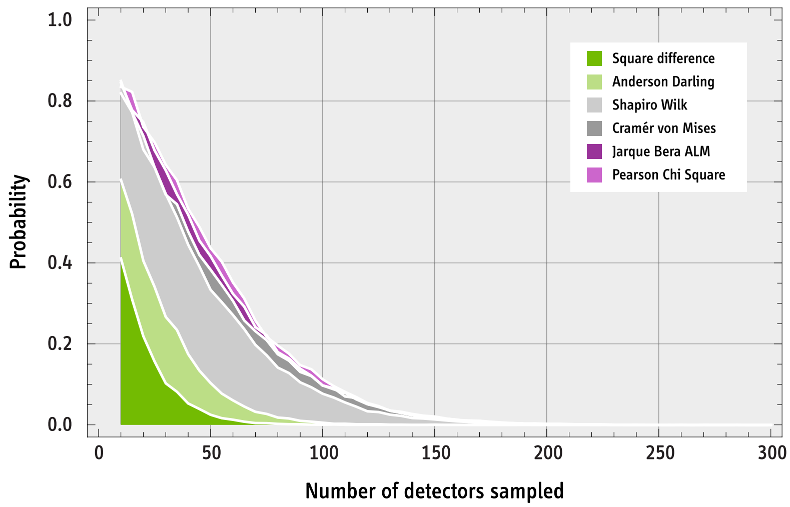

Figure 6. Probability of incorrectly tagging the invalid item as a "template match" as a function of the total number of detectors sampled. Depending on the type of test, 60–150 data points are sufficient to reduce the false-negative rate below 1%.

Implementation Challenges

For this protocol to be effective, there are some important requirements on the measurement process. Most centrally, the neutron source must be well controlled and measurable, so that there is no significant variation in the neutron field produced nor in the total number of neutrons emitted when irradiating different items. An accurate neutron flux monitor can be used to set the irradiation time, so perfect reproducibility is not required in the rate of neutron production. Since coincidence counting is not possible with passive detectors, it is also important that room-return neutrons be acceptably minimized and reproducible. Furthermore, it must not be possible, when the final counting is performed, to distinguish preloaded bubbles from bubbles produced in the measurement process. Over the period between when the pre-loaded bubbles are formed and the measurements are made, there must not be any detectable aging.

Statistical noise in the measurement will not reveal any information if the noise in both the preload and the actual measurements are characterized by Poisson statistics. However, any systematic measurement errors must be well understood, such that while one detector may be characterized by a different efficiency than another, which can be calibrated out at the end of the process, this efficiency must not vary significantly between the preload and the measurement processes. For example, it will be important to maintain control over the temperature of the detector arrays.

We anticipate that these requirements can be met, but the techniques to achieve the necessary degree of control need to be demonstrated and validated.

Conclusion

Dedicated inspection systems are likely to play a critical role in verifying future arms control agreements, which may cover both tactical and non-deployed nuclear weapons and require verified warhead dismantlement. A major new verification challenge will be to authenticate nuclear warheads offered for inspection without divulging classified information. Using so-called information barriers is one possibility to accomplish this task, but such barriers result in complex inspection systems that are difficult to certify and authenticate. As an alternative, a zero-knowledge protocol for nuclear warhead verification has been proposed that avoids detector-side electronics and does not require a technological information barrier to protect classified information. The proposed verification approach meets the dual requirements that measurements can be made to any desired accuracy, and that the information that is produced by the measurement, in the case that the weapons are all real, contains no sensitive information whatsoever.

Ongoing experimental and theoretical work focuses on a preliminary demonstration of the proposed scheme, including optimization of the measurement protocol and distribution tests over a range of diversion scenarios, followed by assessments of approaches to resolving each of the implementation challenges.

Acknowledgement. *This project is supported by generous grants from Global Zero and the U.S. Department of State, funding from the Princeton Plasma Physics Laboratory, supported by DOE Contract DE-AC02-09CH11466, and in-kind contributions from Microsoft Research New England. All simulations were run on Princeton University's High Performance Cluster.*

Endnotes

[1] David Spears (ed.), *Technology R&D for Arms Control,* U.S. Department of Energy, Office of Nonproliferation Research and Engineering, Washington, DC, Spring 2001, www.ipfmlibrary.org/doe01b.pdf.

[2] Bernard Chazelle, "The Security of Knowing Nothing," *Nature,* 446, 26 April 2007, pp. 992–993.

[3] A. Glaser, B. Barak, R. J. Goldston, "A New Approach to Nuclear Warhead Verification Using a Zero-Knowledge Protocol," *53rd Annual INMM Meeting,* Institute of Nuclear Materials Management, July 15–19, 2012, Orlando, Florida.

[4]James Hall, "Uncovering Hidden Defects with Neutrons," *Science & Technology Review,* May 2001, www.llnl.gov/str/May01/Hall.html.

[5]Francesco d'Errico, "Radiation Dosimetry and Spectrometry with Superheated Emulsions," *Nuclear Instruments and Methods in Physics Research B*, 184 (2001), 229–254. The current experimental setup uses detectors from *Bubble Technology Industries,* www.bubbletech.ca.

[6]The fission cross section for 14.1 MeV neutrons differs little between fissile and fertile targets, e.g. between U-235 and U-238. However, $1/(1-k)$ neutron multiplication will happen within a warhead, and will differ between fissile and fertile materials. It may nonetheless prove to be desirable to use lower energy (300–600 keV) neutrons, perhaps only for the low-spatial-resolution side measurements. These can be produced using the (α, n) reaction on lithium-7 or by moderation of the 14.1-MeV neutron source.

[7]Variations of Steps 1 and 2 are expected for all verification approaches based on the template method. All subsequent steps are unique to the zero-knowledge verification approach proposed here.

[8]Of course, there are many other objects in the possession of the host to which the inspector does not have direct access, such as the weapons being inspected themselves.

[9]In fact, the inspector may choose to make only a few selected measurements on the template; she could even decide not to make any measurements on the template at all. The opportunity to use detectors on the template to validate preload may generally be sufficient to deter the host from cheating.

[10]Francesco d'Errico, Yale University, personal communication, April 2013.

[11]It is worth noting that these distribution tests only consider the histogram of counts, and ignore spatial information. Thus they can be performed after the host mixes up the detectors so that the inspector does not know which orientation each detector corresponds to. This can help minimize information leakage in case there are some small differences between the test and template objects (e.g. due to manufacturing variations or misaligned positions). The inspector may learn, with some uncertainty, from the bubble counts that these differences exist, but will not learn their location.

Submitted, Revision 6, June 2013