Experimental Demonstration of a Physical Zero-Knowledge Protocol for Nuclear Warhead Verification

Sébastien Philippe,* Robert J. Goldston,*,† George Ascione,† Andrew Carpe,†
Francesco d'Errico,‡ Charles Gentile,† and Alexander Glaser.*

*Princeton University, Princeton, NJ

†Princeton Plasma Physics Laboratory, Princeton, NJ

‡Yale University, New Haven, CT

ABSTRACT. Zero-knowledge protocols are a class of interactive proof systems that yield nothing beyond the validity of the assertion being proved. Developed for computational cryptographic applications, their first non-trivial physical application was proposed by Glaser, Barak, and Goldston for the authentication of nuclear warheads. Here we report on the advancement of the experimental proof of concept of our physical zero-knowledge protocol for nuclear warhead verification. We present the first experimental demonstration of a zero-knowledge differential neutron radiographic protocol. The reference item (template warhead) and inspected items (treaty accountable items) are represented by objects of different fast neutron opacities made of steel and aluminum cubes. We illuminate items presented for inspection using a D-T14 MeV neutron generator, and record their radiographs on identical sets of super-heated emulsion (bubble) detectors, specially developed for our application and previously preloaded with the complement of the transmission image of the reference item. We find that, even with a modest number of bubbles, it is possible to discriminate between materials of different opacities and that zero-knowledge is leaked for inspected items identical to the reference item.

Background

A zero-knowledge proof is a proof that is both convincing to a verifier and, at the same time, does not yield any knowledge beyond its validity. Glaser, Barak and Goldston have proposed to apply this cryptographic concept to the physical problem of confirming the authenticity of a nuclear warhead.[1] Using a template-matching zero-knowledge protocol with non-electronic detectors, we are in principle able to address at the same time both the physical measurement and the information security challenges presented by a nuclear warhead inspection.[2] Our approach envisions the comparison of the geometry and material composition (including isotopic) of an inspected item against those of a reference item, or template, to confirm that items are substantially identical. Current efforts focus on the experimental proof of this concept using D-T neutron radiography, as well as the development of a p-Li active neutron interrogation technique, with much lower neutron energy, for differential isotopic measurement.[3]

In this paper, we present the experimental results of the first zero-knowledge differential neutron radiographic measurement. To our knowledge, this is the first demonstration of a physical zero-knowledge proof of physical properties,[4] and represents the first step towards the demonstration of an efficient zero-knowledge protocol for nuclear warhead authentication, where sensitive information is never measured.

Experiment and Apparatus

We wish to verify that items presented for inspection have an identical 14 MeV neutron radiograph to a reference item ("golden warhead" or template) and if so not learn anything but the fact that they are identical. The test items used in the experiment are combinations (or patterns) of 2 inch cubes of aluminum and steel. The materials have different opacities to 14 MeV neutron (44 % for 2" of steel and 27 % for 2" of aluminum) and the cubes can be arranged in different patterns on a 3 by 3 grid. Figure 1 shows the cube pattern used in the experiments to represent the reference item.

We wish to identify cases in which the cube pattern has been altered — without gaining any information about the configuration in cases where it has not. This is analogous to determining if a conjurer in a game of "cups and balls" has moved the ball from under one cup to under another, without ever determining the location of the ball if he has not moved it.[5]

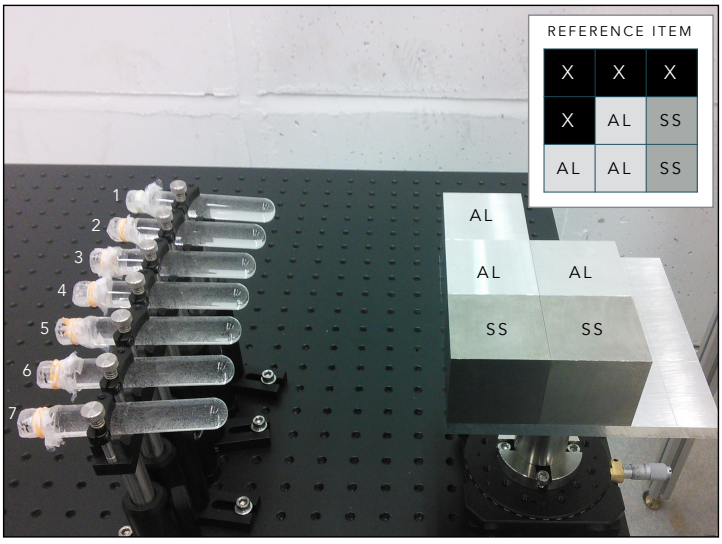

Figure 1. Cube pattern representing the reference item used in this experiment. SS: stainless steel cube, AL: aluminum cube and X: no cube. Upper right image: standard pattern representation used bellow. Note that the left side of the pattern is in front of the lower numbered detectors.

Apparatus. We use 14 MeV neutrons from a D-T neutron generator (Thermofisher B320) with a yield of $\sim 10^8$ neutrons per second to take radiographs of a test item. The item is mounted on a stand allowing precise, $\sim 10\mu m$, *(x,y)* translation and rotation about the z (vertical) axis. Neutrons emitted roughly isotropically from the generator are shielded by a borated polyethylene cylinder (1 m high $\times$ 1.2 m in diameter), except for a 5 cm high and 17 degree wide fan-shaped slot. A single row of 7 specially made superheated droplet neutron detectors (or "bubble" detectors) is placed 110 cm away from the source behind the test item (Figure 2). Detectors point towards the source and are spaced 1.6 degrees from each other. The odd-numbered detectors face the centers of the rear vertical edges of the blocks, while the even-numbered detectors face the centers of the rear faces of blocks. Both the detectors and the test item are mounted on a honeycomb aluminum optical table. A fast neutron counter (Eljen EJ-410 zinc sulfide scintillator) monitors the source fluence behind the bubble detectors during irradiation. The complete apparatus is located in a room shielded by borated concrete walls at the Princeton Plasma Physics Laboratory. The purpose of this simple apparatus is to produce one dimensional radiographs of various cube patterns and study the sensitivity of the system to differences in geometries and neutron opacities.

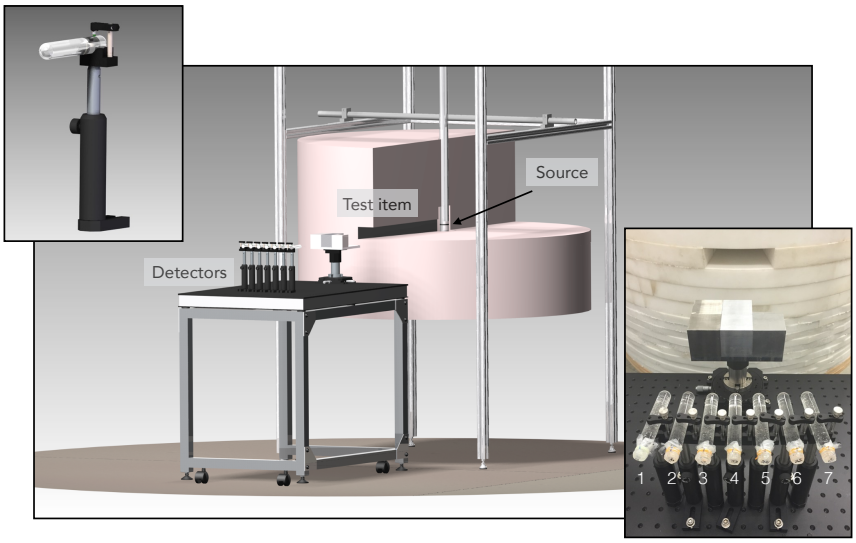

Figure 2. Main image: artist's representation of the experimental setup with neutron generator in collimator (right), test item (center) and single row of superheated detectors facing the source (left). The test item is a combination of 2" cubes of aluminum and steel. Upper left image: representation of a detector holder and vial filled with an emulsion of aqueous gel and superheated droplets. Lower right image: photograph showing the collimator opening, test items and detectors; numbers refers to detector positions.

Superheated droplet detectors. Our inspection protocol requires the use of pre-loadable non-electronic detectors. This way we avoid the certification and authentication problems associated with electronic measurement systems which, may be subject to snooping or tampering. As a consequence, the host is able to provide the pre-loaded detectors for the measurement while the inspecting party can have high confidence in the results (soundness of the proof in the cryptographic literature). For example, even after the measurements have been made, the inspector can demonstrate to his satisfaction that the detectors operate as expected. In the experiments reported here, the detectors comprise standard glass vials (test tubes) filled with an emulsion of C-318 fluorocarbon (C_4F_8, Octafluorocyclobutane) superheated droplets homogeneously dispersed in an aqueous gel matrix.[6] There are close to 4000 droplets of 100 μm diameter per cm^3. The detectors were developed with characteristics tailored to our application.[7] In particular, the detectors are insensitive to incident neutrons with energies below 1 MeV as well as to gamma radiation. When a metastable droplet vaporizes and explodes due to an energetic enough neutron interaction, it expands into a stable bubble about six times larger. The detectors can be irradiated several times and record the total fluence to which they are exposed. If an isostatic pressure of 500 psi is applied to the aqueous gel for about 10 minutes, the bubbles are recondensed into droplets.

Counting Approach. The absolute efficiency of the detectors is of the order of 4×10^{-4} bubbles per crossing neutron, in an active (or counting) volume of 6.5 cm^3. After irradiation, pictures of the detectors are taken using a commercial instrument (BDR-III reader, Bubble Technology Industries) at multiple angles. The bubbles are then counted with the BDR-III commercial software with modified settings to accommodate our vials. The results are compared with visual counting using large printed copies of the images to assess the counting performance of the image processing. The main challenge stems from the correct segmentation of bubbles that overlap in the 2D image. The optical 2D counting method is subject to a reproducible saturation mechanism of bubbles occluding one another, but remains a viable method for this proof-of-principle experiment when the saturation is taken into account. To obtain the statistics required to detect more subtle changes in opacity, thousands of bubbles will be needed.[1] This will require the development of alternative counting methods possibly involving 3D imaging and/or larger detectors set further away from the source.

Results

We first characterized the bubble detector fluence response, $B = f(C)$, for our apparatus, where B is the number of bubbles given by the reading system and C the counts from the fast neutron detector. All bubble detectors were from a same batch and are assumed to have equal efficiency. From the results, presented in Figure 3, we could not infer different efficiencies of individual detectors, nor different fluences at the 7 different detector locations. Data were acquired in two ways to study the effect of recompression on the detector response. One detector batch corresponding to 5000, 10,000, etc., counts was irradiated, measured, recompressed and then irradiated to the next higher level. A second batch was irradiated, counted, and then further irradiated to provide data at 2500, 7500, etc., counts. Figure 3 shows that recompression had no effect on the measurements.

Importantly, the response curve presents a reproducible saturation effect which originates from the bubble counting software's limited ability to segment and identify bubbles touching or occluding each other on a 2D image.[8] Other potential sources of additional nonlinearity are also being investigated. These are estimated to be in the few-percent range and include reduced bubble formation due to reduced emulsion fraction and bubbles being forced out of the control volume. Effects due to simple droplet depletion should be $\sim 1\%$ as the control volume contains about 24,000 droplets before irradiation.

For the present experiments, we set $Nmax = 154$, which corresponds to the mean number of bubbles expected in the detectors if they were exposed directly (without the presence of a test item) to the source for about 6 minutes. After the protocol ends and if the host presented a valid item and a matching pre-load, the total count in each detector is expected to be the same as if no object were present. This $Nmax$ value was chosen so that the bubbles could be counted optically with only moderate saturation. It is nevertheless large enough to yield statistically significant results in this proof of concept.

Figure 4 shows radiographs of the reference item and complementary radiographs that were preloaded in several sets of detectors. Results for the radiograph are consistent with MCNP predictions used to design the experiment. After inverting the detector response function, the complementary radiographs were obtained by directly irradiating the detectors in the absence of items and removing them when they had been exposed to a fluence of $f^{-1}(N_{pre})$ corresponding to a preload bubble count equal to $N_{pre} = N_{max} - N_{ref}$. This procedure does not account, however, for the nonlinearity of the detector response, which requires $N_{pre} = f(f^{-1}(Nmax) - f^{-1}(N_{ref}))$. We corrected the final results (figures 5 and 6) to account for this $\sim 1\,\sigma$ effect, using the measured calibration curve (figure 4). Future experiments will implement this approach directly.

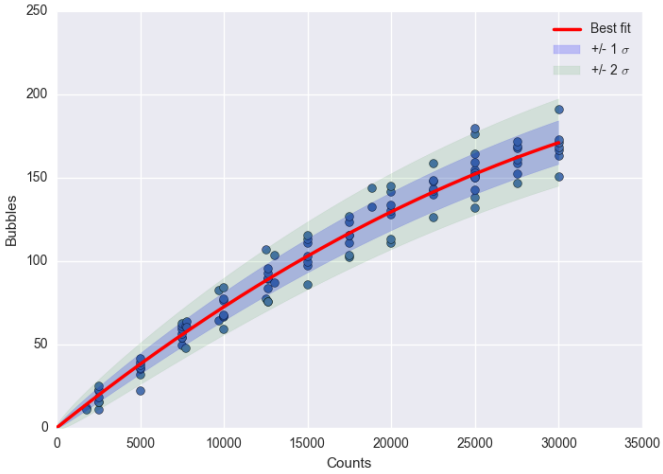

Figure 3. Detector calibration curve, $B = f(C)$, in the absence of test items. Vertical axis is bubble counts, horizontal axis is counts in the fast neutron detector. Blue dots are data points and include detectors at different positions. The red line is the best fit of a second-order polynomial forced through the origin. 25000 counts in the fast neutron detector corresponds to about 6 minutes of irradiation. The noise level is reasonably well estimated by $\sqrt{N}$ despite possible effects due to the saturation mechanism, detector differences and position.

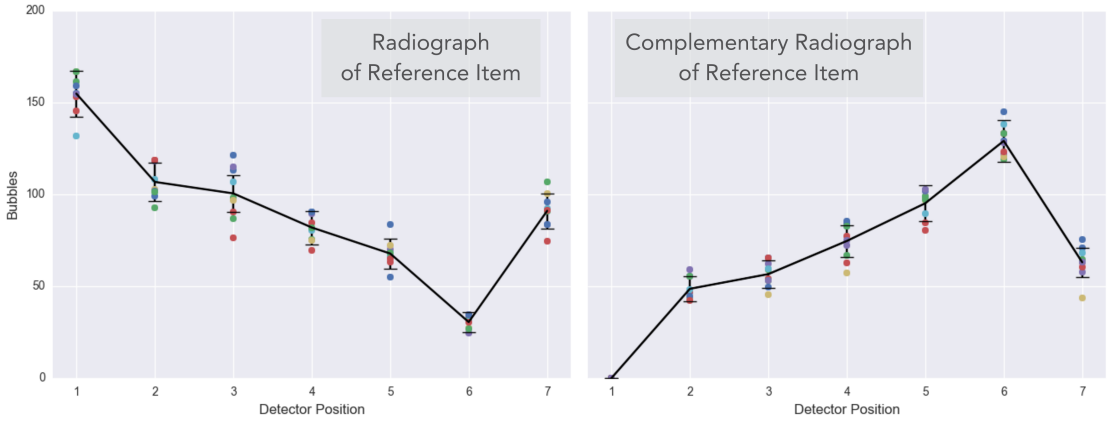

Figure 4. Radiograph (left) and complementary radiograph or preload (right) for the reference item, object A. Nine separate radiographs were measured and averaged to determine the complement. Ten such complements or preloads were prepared. Data points show the bubble count at each position. The black line connects the means of the measurements at each position. Vertical error bars are 1 $\sigma = \sqrt{N}$ deviation from the means.

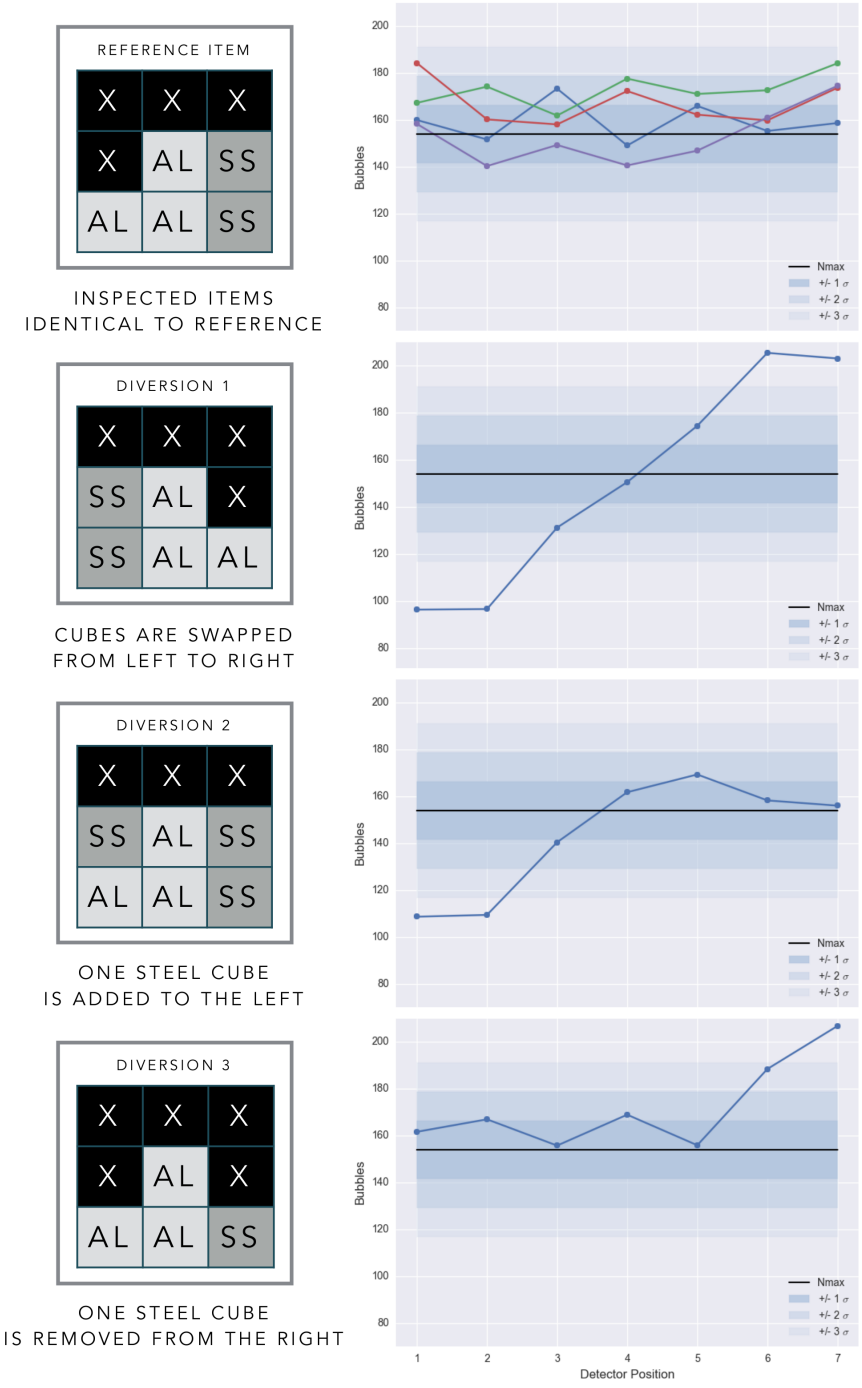

Figure 5. Results of inspections for valid items and diversion scenarios. Error bands are given by $\sigma = \sqrt{N_{max}}$. In the honest host scenario (top), four inspected items cannot be distinguished from the reference item and the results are zero-knowledge. In the three diversion scenarios (invalid item radiographed using a complementary radiograph of the reference item as preload), the cheating host is caught and the final images are no longer zero-knowledge.

Once preloaded detectors have been prepared, the next step is to irradiate them again in the presence of the test items. Figure 5 shows the results for valid items identical to the reference item as well as the results of three diversion scenarios where the cube pattern of the reference item was altered.

In the "honest host" scenario (valid items), the radiograph of the reference item is added to a set of detectors previously preloaded with its complementary radiograph. In this case, the data are randomly distributed and the result of the inspection is zero-knowledge. Specifically, one can gain no knowledge about the opacity profile of the reference item from this measurement. We cannot tell where the conjurer has placed the ball, but we cannot claim that he has moved it. Error bands are calculated from the expected value of *Nmax*. These data also suggest some day-to-day variability that we plan to investigate.

In the first diversion scenario, the cubes of the inspected item are swapped from left to right. The effect of irradiating this item with the preload of the reference item is rather dramatic and can be easily identified. Two additional scenarios consisting of altering the reference item pattern by either adding (increased opacity) or removing (decreased opacity) a single steel cube. Again, both cases would be caught with *Nmax* as low as 154. The results for a smaller diversion, where a steel cube was replaced with aluminum, are shown in Figure 6. Here, individual inspections do not always produce clear outcomes. The experiment was repeated five times and the data averaged to catch the reduction in opacity.

Additional diversion scenarios are under investigation and consist of more subtle changes in opacity or geometry. However, the statistics needed to catch small variations from the reference item are likely to require higher values of *Nmax*, or multiple exposures following an optimized testing strategy.[9]

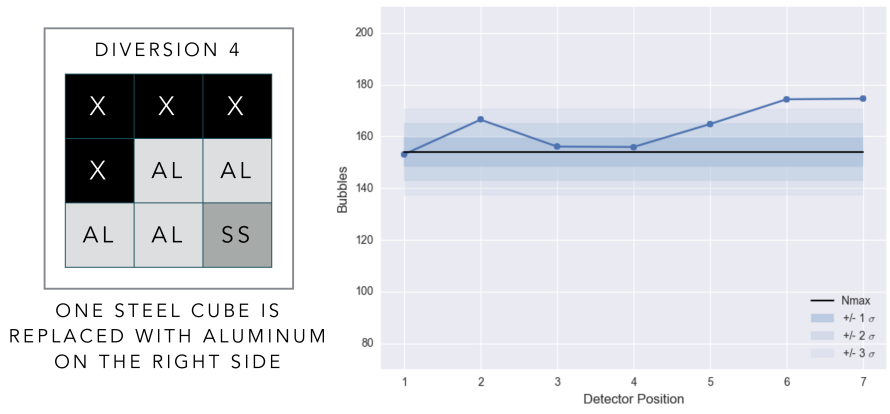

Figure 6. Results of inspections for the fourth diversion scenario (small reduction in opacity on the right hand side). Five inspections were averaged. Error bands are given by $\sigma = \sqrt{N_{max}/5}$.

Future Work

The current value of $Nmax \sim 150$ limits our ability to detect subtle changes in geometry and opacity. Future iterations of our experimental apparatus will therefore need alternative methods for counting larger number of bubbles, without introducing unacceptable nonlinearity in the detector response curve. Appropriate technologies will not only help to scale up $Nmax$ but also the number of "pixels" or detectors for our radiographs.

The present method could be improved by using a videorecording of the detector rotating through 360 degrees. Other potential candidates include simple methods such as measuring the macroscopic change in gel volume due to the creation of bubbles (pycnometer) or the change in light intensity scattered of bubbles. A more complex 3D approach could be based on magnetic resonance imaging (MRI) or sonography. Additionally, we plan to employ superheated emulsion detectors with higher efficiency (up to 1 %) to maintain short irradiation times and to limit the irradiation of inspected items.

Our warhead authentication approach requires the ability to distinguish fissile materials of different isotopic compositions. To do so, we are considering an alternative lower energy p-Li neutron source.[3] We are also investigating beam filters for use with a more powerful D-T neutron generator and may perform measurements on uranium objects of various levels of enrichment.

Conclusion

We report results from the first zero-knowledge differential neutron radiographic measurement. To our knowledge, this is the first demonstration of a physical zero-knowledge proof of physical properties. This proof of concept constitutes a small but important first step towards an efficient zero-knowledge protocol for nuclear warhead authentication where sensitive information is never measured.

In this paper, we have realized the various steps of a zero-knowledge inspection. We have developed a set of methodologies to perform 14 MeV D-T neutron radiography of test items represented by patterns of steel and aluminum cubes. Radiographs are obtained by counting macroscopic bubbles from irradiated non-electronic superheated emulsion detectors. We were able to pre-load the detectors with the complementary radiograph of a reference item and, by subsequent irradiation, verify whether items presented for inspection were identical to the reference item or not. We confirmed that the results yielded zero-knowledge when valid items were tested and successfully identified tampered items for four different diversion scenarios.

Acknowledgement. *The authors thank the Health Physics team of the Princeton Plasma Physics Laboratory for their support during the conduct of the experiments and Margarita Gattas-Sethi (Yale University) for manufacturing the detectors. Funding for this work was provided by DOE/NNSA through the Consortium for Verification Technology to Princeton University, Yale University and the Princeton Plasma Physics Laboratory. Financial support was also provided by the MacArthur foundation.*

Endnotes

[1]A. Glaser, B. Barak, and R. Goldston, "A Zero-knowledge Protocol for Nuclear Warhead Verification," *Nature*, 510, June 2014, pp. 497-502

[2]S. Philippe, B. Barak and A. Glaser, "Designing Protocols for Nuclear Weapons Verification," Proceedings of the 56th Annual INMM Meeting, Indian Wells, CA, July 2015.

[3]J. Yan and A. Glaser, "Two-Color Neutron Detection for Zero-Knowledge Nuclear Warhead Verification," Proceedings of the 56th Annual INMM Meeting, Indian Wells, CA, July 2015.

[4]B. Fisch, D. Freund and M. Naor, "Physical Zero-Knowledge Proofs of Physical Properties," CRYPTO 2014, volume 8617, pages 313-336, Springer, Aug. 17-21, 2014.

[5]Bosch, Hieronymus. The Conjurer, oil on wood, 1502 (Musée Municipal, St.-Germain-en-Laye). http://art-in-europe.info/2014/10/20/hieronymus-bosch-the-conjurer-1502-in-hd/

[6]F. d'Errico, "Radiation Dosimetry and Spectrometry with Superheated Emulsions," Nuclear Instruments and Methods in Physics Research B, 184 (2001), 229-254.

[7]R. J. Goldston, F. d'Errico, A. di Fulvio, A. Glaser, S. Philippe and M. Walker, "Zero-Knowledge Warhead Verification: System Requirements and Detector Technology," 55th Annual INMM Meeting, 20-24 July 2014, Atlanta, Georgia

[8]These data do not necessarily imply that BTI bubble detectors read with BTI equipment produce non-linearities. These results apply only to our own detectors after we modified significantly the standard factory settings to accommodate our specific conditions.

[9]S. Philippe, M. Kutt, B. Barak, A. Glaser and R. J. Goldston, "Testing Strategies for the Authentication of Nuclear Weapons with High Confidence," pre-print manuscript.