ARTICLE

doi:10.1038/nature13457

A zero-knowledge protocol for nuclear warhead verification

Alexander Glaser[1], Boaz Barak[2] & Robert J. Goldston[3]

The verification of nuclear warheads for arms control involves a paradox: international inspectors will have to gain high confidence in the authenticity of submitted items while learning nothing about them. Proposed inspection systems featuring 'information barriers', designed to hide measurements stored in electronic systems, are at risk of tampering and snooping. Here we show the viability of a fundamentally new approach to nuclear warhead verification that incorporates a zero-knowledge protocol, which is designed in such a way that sensitive information is never measured and so does not need to be hidden. We interrogate submitted items with energetic neutrons, making, in effect, differential measurements of both neutron transmission and emission. Calculations for scenarios in which material is diverted from a test object show that a high degree of discrimination can be achieved while revealing zero information. Our ideas for a physical zero-knowledge system could have applications beyond the context of nuclear disarmament. The proposed technique suggests a way to perform comparisons or computations on personal or confidential data without measuring the data in the first place.

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 takes advantage of the fact that deployed weapons are associated with unique and easily accountable delivery platforms, that is, missile silos, submarines and strategic bombers, to which agreed numbers of warheads are attributed. The next round of nuclear arms-control agreements, however, may place limits on the total number of nuclear weapons and warheads in the arsenals. This would include tactical weapons as well as deployed and non-deployed weapons. Such agreements would require new verification approaches, including inspections of individual nuclear warheads in storage and warheads entering the dismantlement queue. This is a qualitatively new challenge because the design of nuclear weapons is highly classified information that cannot be exposed to international inspectors. A viable verification approach therefore has to resolve the tension between reliably verifying that the inspected warhead is authentic while avoiding disclosure of information about its design[1–4].

Practitioners and policy makers have been well aware of this conundrum, and prior work by national laboratories in the United States, Russia and the United Kingdom addressed it by using 'information barriers'[2,4]. These barriers consist of sophisticated automated systems that process highly classified information measured during an inspection, but only display results in a yes/no manner. Such systems are inherently complex, and require both parties to trust that they have no 'trapdoors' hidden from the inspector, which could be used to cause a system to declare invalid objects as authentic, nor side channels unknown to the host, which could leak classified information to the inspector or others. These concerns are serious obstacles to adopting such systems.

In this work we consider a fundamentally different approach to this problem. Rather than trying to acquire and analyse classified data behind an engineered information barrier, we use the cryptographic notion of zero-knowledge proofs to ensure that sensitive data are never measured in the first place.

Zero-knowledge proofs (with marbles)

These proofs, invented in the 1980s by Goldwasser, Micali and Rackoff[5], have become an important tool of modern cryptography. They achieve the paradoxical goal of allowing one to prove that a statement is true without revealing why it is true. Such proofs are extremely useful for many digital applications, including privacy-preserving data mining, electronic voting and online auctions[6]. To achieve zero knowledge, Goldwasser *et al.* extended the traditional notion of a proof from a static text to a protocol, which involves randomization and interaction between the prover and verifier. At the end of the protocol, the verifier has a high degree of confidence that the statement is correct, while the prover is guaranteed that the verifier did not learn anything about the data underlying the truth of the statement. For our application, the host submitting warheads for inspection takes the role of the prover and the inspector the role of the verifier.

Whereas classical zero-knowledge proofs are digital protocols, proving statements about mathematical objects, we illustrate the concept using a physical zero-knowledge protocol that is closely related to our proposed verification approach (Fig. 1):

Alice (the host) has two small cups both containing X marbles, where X is some number between 1 and 100. She wants to prove to Bob (the inspector) that both cups contain the same number of marbles, without revealing to him what this number X is. To do so, Alice prepares two buckets, which she claims each contain $(100 - X)$ marbles. Bob now randomly chooses into which bucket which cup is poured. Once this is done, Bob verifies that both buckets contain 100 marbles.

The protocol reveals no information on X because, regardless of the value of X, Bob always sees 100 marbles in both buckets. However, if the cups did not have the same number of marbles, then no matter how Alice prepares the buckets, with a probability of 50% after the pouring at least one of the buckets will not contain 100 marbles. If Alice and Bob repeat this game, say, five times, then if Alice consistently cheats she will be caught with probability $(1 - 2^{-5}) > 95\%$.

[1]Department of Mechanical and Aerospace Engineering, Princeton University, E-Quad, Olden Street, Princeton, New Jersey 08544, USA. [2]Microsoft Research, 1 Memorial Drive, Cambridge, Massachusetts 02142, USA. [3]Princeton Plasma Physics Laboratory, Princeton University, PO Box 451, MS 41, Princeton, New Jersey 08540, USA.

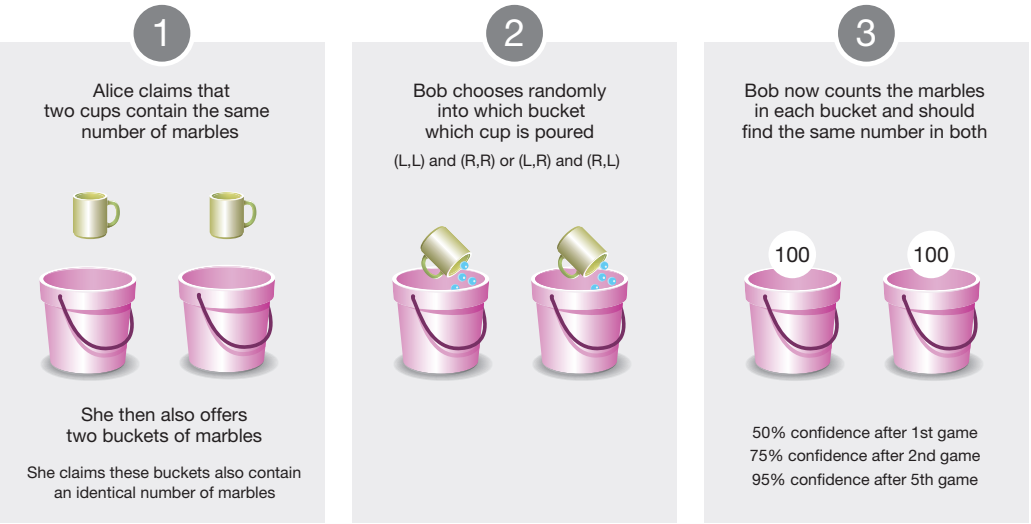

Figure 1 | A zero-knowledge protocol to prove that two cups contain the same number of marbles. L,L indicates left cup into left bucket, and so on. The confidence level increases with the number of games played. See main text for details.

From marbles to neutrons

The relevance of the above protocol to our setting is that we want to show that two or more putative warheads have identical neutron transmission and emission counts under irradiation by high-energy neutrons. We follow the template approach for warhead verification[2], in which a radiation measurement generates a complex and unique fingerprint of an inspected item. This fingerprint is then compared against the fingerprint of one or more templates to confirm that all items are materially identical. Template selection is a critical and challenging step. In practice, templates could be directly selected from deployed weapons so that the inspecting party has high confidence in their authenticity. Strong chain-of-custody measures would have to be in place to assure the inspectors that the templates have not been swapped out between visits. In the case of weapon systems that are not currently deployed, and in other cases where a trusted reference item may not be available, differential measurements could still be valuable. In these cases, measurements of a large number of warheads in a batch could confirm that they are all materially identical. Combined with some other supporting evidence (for example, records confirming the 'pathway' or 'provenance' of at least some of these items), this could provide confidence in the authenticity of all items in the inspected batch[7].

We compare the submitted items by recording the transmission pattern of 14-MeV neutrons, as well as recording the intensity of neutrons emitted at large angles from the items. Active interrogation of nuclear warheads or warhead components with high-energy neutrons and other types of radiation has been successfully demonstrated, but does require detailed safety analyses. Note that, even when exposed to strong neutron sources, the induced fission events in a nuclear warhead produce less than a milliwatt of heat, which is far less than the heat already generated by α-decay and spontaneous fission in items containing, for example, kilogram quantities of plutonium.

Neutron radiographic images of warheads contain highly classified information, but in our case they are actually never measured. Rather, by analogy to the marbles example, they are recorded using detectors that are preloaded with the negative of the radiograph. Preloaded values are not revealed to the inspector. As in the marbles example, after the measurement and if the host is telling the truth, the inspector always sees the same number of counts in every detector. Furthermore, as in the marbles example, preloads supplied with the submitted items are shuffled at random, so if the items actually differ, then no matter how the preloads are chosen, with significant probability the image will not be uniform, and a mismatch will be present on both items.

Unlike the marbles example, neutron measurements are inherently statistically noisy. To avoid conveying information through the noise distribution we use preloaded values that are noisy as well. In particular, since the signal added during interrogation will have a Poisson distribution, we also use a Poisson distribution for the preloads. Using the fact that the sum of two Poisson distributions is also Poisson in character, our protocol achieves the following: the neutron count obtained by any measurement on the template or on any valid submitted item is distributed according to the Poisson distribution with mean equal to a previously agreed-on value, N_{max}, and a standard deviation equal to $(N_{max})^{0.5}$.

Since N_{max} is known in advance by both sides, neither the measurement nor its noise reveals any new information. N_{max} for transmission could reasonably correspond to the maximum number of counts that is expected in the absence of a test item. If a submitted item varies from the true warhead (or the submitted preloads are not identical) an image may be seen that could contain sensitive information. This will be an additional strong incentive for the host not to cheat.

For simplicity of operation, we envision that the host places the detectors for each measurement in a removable board that forms part of the measurement system. Crucially, the inspector chooses which board to use with which test item. As in the marbles example, this means that if the host uses non-identical boards to try to mask invalid items, then with 50% probability the invalidity will be made more evident by the measurement with the mismatched boards. Since we expect that this 'game' will be repeated many times, even a risk-tolerant host would not accept the resulting low chance of success. We note that testing multiple warheads in parallel is an attractive option, because it makes the probability of detecting the use of non-identical preloads significantly higher.

Once the measurements have been completed and the detectors read out, the inspecting party can verify the functionality of the detectors by exposing them to additional neutrons. This is an important advantage of the proposed method. For inspection systems proposed so far, specialists from the inspecting country would not be allowed to examine any equipment once it has seen classified information.

Although we examine here neutron measurements using preloaded non-electronic detectors, there may be other non-electronic zero-knowledge protocols for warhead verification that can avoid the use of engineered information barriers. Indeed, such systems could be complementary to the neutron measurements discussed here.

Monte Carlo analysis

We now show how our approach can be implemented in practice, and that small differences between two objects can be reliably detected. We have analysed the approach with a series of simulations using the general Monte Carlo N-Particle (MCNP) transport code[8]. Construction of a physical experimental set-up is under way.

We propose to use 14-MeV neutrons from a deuterium-tritium neutron generator[9] to interrogate test items, allowing detailed transmission profile measurements and also measurements of neutron intensities at large angles due to elastic and inelastic scattering, fission and $(n,2n)$ reactions. The neutrons from the generator are collimated by 60 cm of polyethylene and illuminate the inspected item (Fig. 2). An array of neutron detectors placed at a distance of 50 cm behind the centre of the item provides the transmission measurements. Additional detectors (not shown) can be positioned with additional shielding at large angles to the beam, that is, in the shadow of the collimator, to measure neutrons emitted from the test item.

Test item

The test item used for this analysis is the unclassified 'British Test Object' (BTO), which consists of concentric rings of polystyrene, tungsten (two rings with a combined mass of 7.74 kg), aluminium, graphite, and steel. The BTO has an outer diameter of 18.9 cm and a height of 5 cm. This test object does not contain special or other nuclear materials, but is used to develop and calibrate imaging systems for diagnostic analysis of nuclear weapons[10]. $(n,2n)$ reactions in the tungsten used in this test object provide a reasonable approximation of induced fission events expected for a nuclear warhead or warhead component. Neutron multiplication in a real item would increase net neutron production rate by some finite amount, but the effect is extremely small for transmission measurements. Furthermore, the energy threshold (~10 MeV) used for the transmission detectors renders them insensitive to fission and $(n,2n)$ neutrons. The BTO is placed in a container in order to avoid revealing to the inspector the appearance or orientation of the inspected item inside the container.

Detector array

To assess the viability of our proposed protocol, we work with a board holding a hexagonal array of 367 detectors consisting of 21 rows of 17 or 18 detectors within an area of about 42 cm × 42 cm. The assumed area of each detector (pixel) is 2 cm^2. By rotating the BTO, the board can image it in any orientation.

In the analysis below, detectors are assumed to be sensitive to neutron energies >10 MeV. Neutrons scattered from the walls of the room are not included in the calculations, but preliminary studies indicate that room return has at most a small impact when using 10-MeV-threshold detectors, particularly if the room is specially prepared for the inspection—for example, with borated polyethylene in front of borated concrete walls.

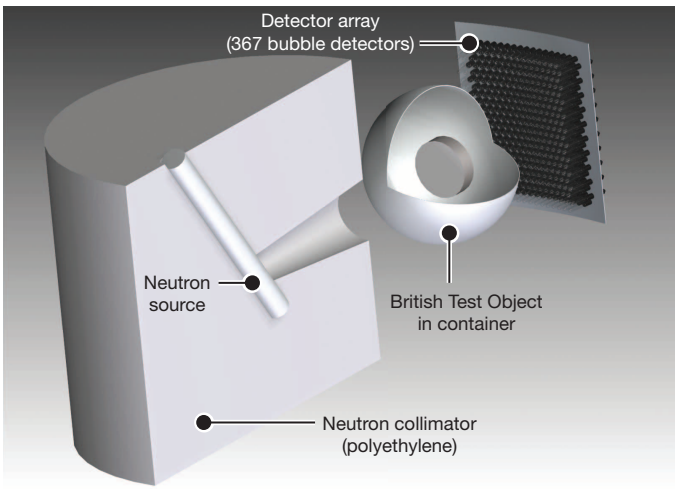

Figure 2 | Experimental set-up with neutron source, neutron collimator, British Test Object in container, and detector array. Large-angle detectors are not shown. See main text for details.

Figure 3 illustrates typical results from a template and a valid item in two different orientations; for reference purposes, the respective neutron radiographs of the test items are shown, but these data are never measured in the inspection, since only preloaded detectors are used. As expected, in the case of inspecting a valid item, detector counts are distributed consistent with a Poisson distribution with mean and variance N_{max}. In the following, only the more challenging side-view orientation of the BTO is used for analyses of four representative diversion scenarios.

Diversion scenarios

To examine diversion scenarios, in which material is removed or replaced, we need to define a decision rule to distinguish passed from failed tests. For our present purposes, we use a very simple rule looking for statistical outliers on predefined groups of pixels. If we denote individual detector counts by the numbers $X_1,\ldots, X_n$, then we can define new numbers $Y_1,\ldots,Y_k$, where every Y_j is the sum of a small number of the X_is, divided by the expected standard deviation of Y for a match case (that is, inspected item identical to template). We define the test to be positive (that is, diversion detected) if there is at least one j with $|Y_j| > T$, where T is a threshold chosen such that in the match case for every j the probability that $|Y_j| > T$ is at most p_{fp}/k where p_{fp} is our allowed false positive rate. Concretely, in our setting, we examine $k = 295$ non-disjoint seven-pixel windows defined by a central detector and its six nearest neighbours. In this case, to achieve a false positive rate $p_{fp} \leq 0.05$, the threshold can be computed numerically to be $T = 3.76$ standard deviations.

Sensitivity of the measurements to diversion scenarios increases with N_{max} and the associated improvements of counting statistics. We therefore examine in the following a series of different diversion scenarios and a range of values for N_{max} to determine system requirements (Table 1). In the full-removal scenario, both tungsten rings are removed from the BTO, which is easily detected even for very low detector counts. Similarly, if lead is used to substitute both tungsten rings, the diversion is clearly distinguishable even by simple visual inspection of the detector board (Fig. 4, top). Our proposed statistical test identifies the diversion in the full-substitution scenario with a probability of true positives, p_{tp}, of > 0.99, even for N_{max} as low as 1,000 detector counts.

The local-removal and local-substitution scenarios are more challenging. In these cases, a 36° sector of the outer tungsten ring is removed or replaced, which corresponds to a diversion of 543 g of tungsten contained in the BTO. To achieve a detection probability of 95%, an N_{max} of 5,000 is required in the case of the localized tungsten removal. When lead is used to substitute for tungsten in the 36° sector, N_{max} increases to 32,000 for the same detection probability.

Note that in these studies no use has been made of the emission detectors at large angles. The more realistic case of substitution of ^{238}U for ^{235}U in a nuclear weapon component results in a reduction by a factor of about two in the induced fission rate due to 14-MeV neutrons. Substitution of reactor-grade for weapon-grade plutonium has a small effect on the directly induced fission rate, but a large effect on the spontaneous fission rate, which could be detected passively by operating the side detectors in the absence of the neutron source. Thus the calculations presented here are conservative.

We note that 5% of the items will be flagged as invalid by our proposed test procedure due to the set 5% false positive rate, even in the case where all items are valid. Retesting flagged items will rapidly determine their validity. If a detection probability for invalid items of 95% is deemed too low, either routine retesting or a greater N_{max} can be implemented to increase this value. The optimization of any retesting scenario, and study of a wider range of host strategies for cheating, as well as inspector strategies for analysing signal patterns to find such cheating, will be the subject of future research.

Preloadable non-electronic detectors

Perhaps the most critical aspect of a viable implementation of the proposed verification approach is the choice of the detector technology. The

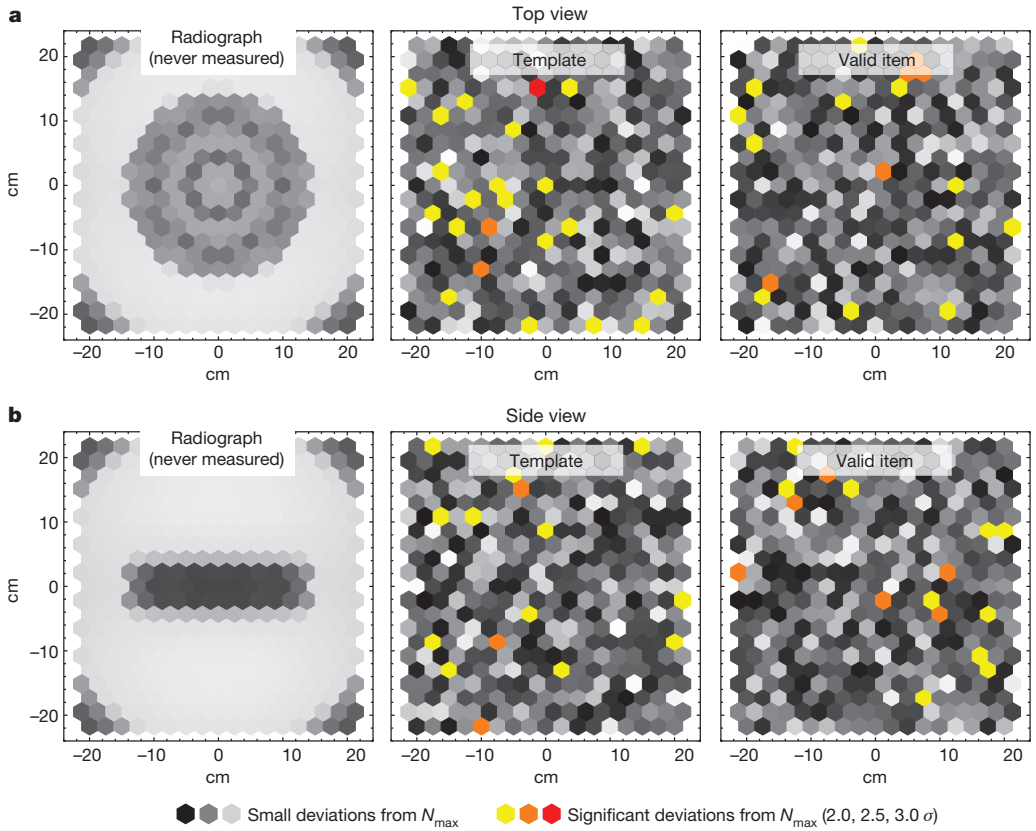

Figure 3 | Results of MCNP5 simulations for interrogations of the British Test Object in two different orientations. Inspection of valid items reveals no information about them. **a**, View from the top; **b**, view from the side. Leftmost panels, radiographs of the items are never measured; they correspond to measurements without preloading the detectors. Other panels visualize total detector counts, including the host's preloads, after the measurements on the template (middle panels) and on an identical (valid) item (rightmost panels), and hence are simply independent Poisson random variables with mean N_{max}. Shades of grey and colours indicate differences from N_{max} scaled to $(N_{max})^{0.5}$. For this calculation, $N_{max} = 5{,}000$.

detectors must have the capability to be preloaded with a desired neutron count before the inspection. At a minimum, this preload has to persist for hours or days and its decay rate, if present, be well characterized. Preloaded counts must be indistinguishable from counts accumulated during irradiation of the test items. Detectors should be energy selective so that the effect of low energy neutrons returning from room walls can be minimized. They should be insensitive to γ-rays, have high efficiency, and permit total counts in the range discussed above. Finally, relying on a non-electronic detection mechanism is highly advantageous given that complex electronic components and circuits are potentially vulnerable to tampering and snooping. We find that at least two detector technologies can meet these criteria: superheated emulsions ('bubble detectors') and neutron activation analysis detectors.

In superheated emulsions, neutron recoil particles trigger the formation of macroscopically observable bubbles from microscopic droplets that are dispersed in an inert matrix[11]. These detectors can be configured to have essentially any desired energy threshold from 10 keV to 10 MeV. Commercially available, polymer-based bubble detectors are limited to a maximum bubble count of the order of a few hundred bubbles, beyond which camera-based imaging techniques cannot resolve bubbles individually. Superheated drop detectors produced with an aqueous gel can be used up to much higher bubble counts. Either optical tomography or magnetic resonance imaging allow the counting of bubbles hidden in the depth of the fluid[12,13]. If the highest N_{max} is desired, multiple detectors can be exposed in series. By the proper choice of a compliant matrix, detectable ageing (growth) of bubbles can be eliminated. Net detection efficiency of the order of 1% can be easily achieved. The emulsions can be contained in opaque containers so that a preload is not visible to the inspector.

For neutron activation analysis, imaging can be undertaken using, for example, an array of hexagonal prisms made of zirconium[14]. ^{90}Zr has a neutron activation threshold of 12 MeV through an $(n,2n)$ reaction. The resulting ^{89}Zr has a half-life of 3.3 days, which must be taken into account to determine the required level of preloading. Counting the γ-rays from ^{89}Zr decay in high-purity germanium well detectors should give a net detection efficiency of about 0.25%. For 3-cm-long prisms, with a cross-sectional area of 2 cm^2, this would provide an N_{max} in the range of 20,000 within one hour, for a commercially available DT neutron generator producing 3×10^8 neutrons s^{-1} (ref. 15). Indium has an appropriate activation response to fission neutrons, with reduced sensitivity in the range of 14 MeV, for use in the side detectors. It will be important to ensure that unshielded, preloaded activation samples are not in the presence of γ-ray detectors before their final exposure in order to avoid providing the inspector with any information about the preloads.

Detectors can be preloaded with counts with the appropriate statistical properties by exposing them to energetic neutrons for a pre-calculated period of time and/or through a pre-calculated depth of shielding. As discussed above, statistical noise in the measurement will not reveal any information. However any systematic measurement errors must be well understood, such that while one detector may be characterized by a different

Table 1 | Probability of an item being flagged as 'invalid' as a function of N_{max} with 10,000 realizations for each case

Item	N_{max}					
	500	1,000	2,000	5,000	10,000	32,000
Valid item*	≤5%	≤5%	≤5%	≤5%	≤5%	≤5%
Full removal of both tungsten rings	>99.9%	>99.9%	>99.9%	>99.9%	>99.9%	>99.9%
Full substitution of tungsten rings by lead	77.7%	99.5%	>99.9%	>99.9%	>99.9%	>99.9%
Local removal of tungsten	Undetectable	15.7%	41.7%	94.6%	>99.9%	>99.9%
Local substitution of tungsten by lead	Undetectable	Undetectable	6.0%	11.7%	30.2%	95.5%

* By design.

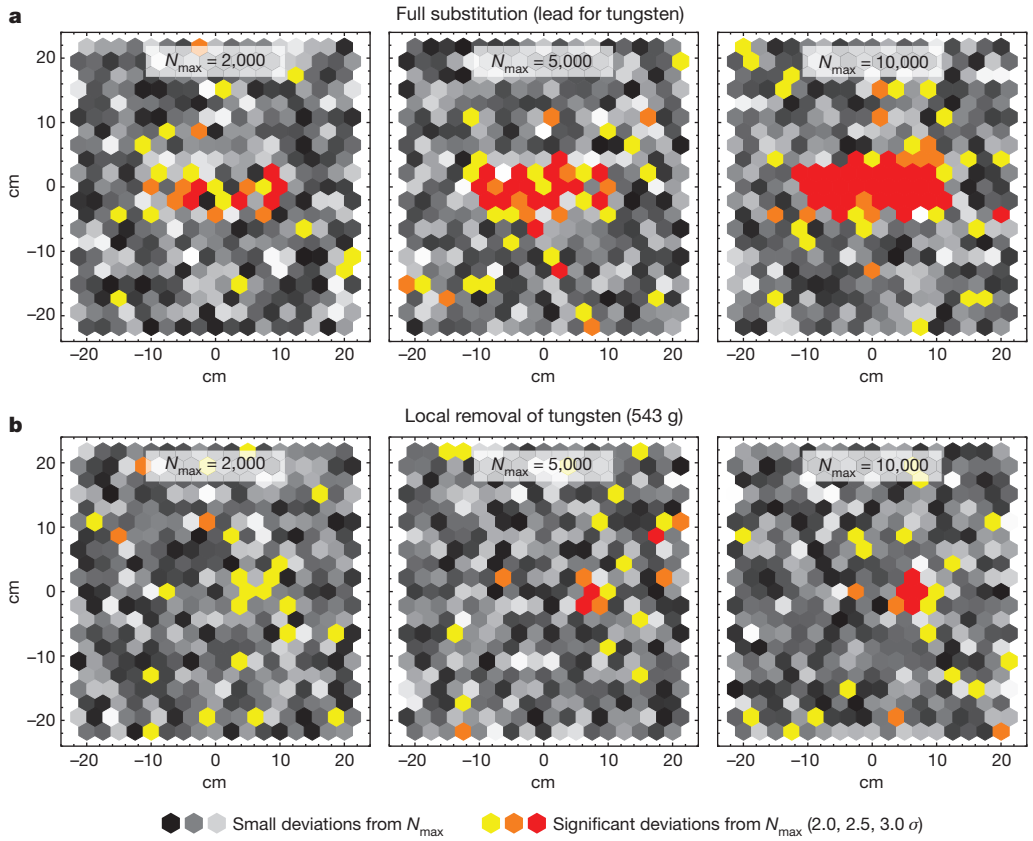

Figure 4 | Results of MCNP5 simulations for two notional diversion scenarios. **a**, Full substitution: tungsten rings in the British Test Object (BTO) are replaced by lead rings. **b**, Local removal: a 36° sector of the outer tungsten ring (mass 543 g) is removed. Inspection of invalid items results in anomalies of the detector counts that become more pronounced with increasing N_{max}. Shades of grey and colours indicate differences from the selected values of N_{max} = 2,000, 5,000 and 10,000 (leftmost, middle and rightmost panels respectively) scaled to $(N_{max})^{0.5}$.

efficiency from another, which can be calibrated out, this efficiency must not vary significantly between the preload and the measurement processes. For example, it is important to maintain control over the temperature of bubble detectors during irradiation. The DT neutron generator must also be well controlled and measurable, so that there is no significant variation in the shape of the neutron field produced nor in the total number of neutrons emitted when irradiating 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. We anticipate that these requirements can be met, but the techniques to achieve the necessary degree of control need to be demonstrated and validated.

The steps following a measurement should be relatively straightforward. Since the information contained in the detectors is in principle unclassified, protocols can be devised that permit using both host-provided and inspector-provided measurement tools. We anticipate that, depending on the strength of the neutron source, the measurements themselves could be completed within hours. Readout would be very quick in the case of bubble detectors, but could take days in the case of some activation detectors. This is not a major constraint, since readout could take place in parallel with other measurements and activities at the site. The authentication process could be accelerated dramatically if N warheads are processed simultaneously (including, for example, $M < N$ reference items). Typically, authentication of one warhead per day can be considered a reasonable target, especially since dismantlement of the warhead itself (including recovery of fissile material and removal of classified features) would take much longer. Authentication is therefore not a significant bottleneck in the process.

Developing a practical inspection system for nuclear warhead verification will be a major undertaking. Ideally, such a system should be jointly developed by partners from weapon and perhaps also non-weapon states. The successful UK–Norway Initiative has shown that such collaborations are possible[16]. Similar efforts could be undertaken for the system proposed here. They would help refine and demonstrate to the satisfaction of all parties the robustness of the method in practical situations, where systematic errors in measurements, small misalignments, or variations in environmental conditions may pose additional challenges that are difficult to anticipate with computer simulations.

More generally, we believe that our approach and techniques could have other applications beyond the area of nuclear disarmament. Once data are measured and converted to digital form, secure comparisons and computations can be performed using many cryptographic tools[17,18]. However, if the measurement device itself cannot be trusted, it is best to ensure the data are never measured in the first place. This could be the case not just for state secrets, but also for personal data, such as biometric data or results of medical tests. For example, following earlier versions of the present work, it was proposed[19] that similar ideas could be used to compare DNA found in a crime scene with a suspect's DNA without actually measuring the latter and creating a DNA profile. Exploring other such applications is an exciting future direction of research on zero-knowledge proofs.

Conclusion

Authenticating nuclear warheads without revealing classified information represents a qualitatively new challenge for international arms-control inspection. Here we have shown an example of a zero-knowledge protocol based on non-electronic differential measurements of transmitted and emitted neutrons that can detect small diversions of heavy metal from a representative test object. This technique will reveal no information about the composition or design of nuclear weapons when only true warheads are submitted for authentication, and so does not require an engineered information barrier. The zero-knowledge approach has the potential to remove a major technical obstacle for verifying deep cuts in the nuclear arsenals, which will probably require verification of individual warheads, rather than whole delivery systems. Timely demonstration of the viability of such an approach could be critical for future rounds of arms-control negotiations.

Received 30 January; accepted 28 April 2014.

1. Comley, C, *et al. Confidence, Security & Verification: The Challenge of Global Nuclear Weapons Arms Control* http://www.fissilematerials.org/library/awe00.pdf (Atomic Weapons Establishment, Aldermaston, UK, 2000).
2. Spears, D. (ed.) *Technology R&D for Arms Control* http://www.fissilematerials.org/library/doe01b.pdf (US Department of Energy, Office of Nonproliferation Research and Engineering, Washington DC, 2001).
3. Fuller, J. in *Cultivating Confidence: Verification, Monitoring, and Enforcement for a World Free of Nuclear Weapons* (ed. Hinderstein, C.) Ch. 4 (Hoover Institution Press, 2010).
4. Anderson, B. *et al. Verification of Nuclear Weapon Dismantlement: Peer Review of the UK MoD Programme* (British Pugwash Group, London, 2012).
5. Goldwasser, S., Micali, S. & Rackoff, C. The knowledge complexity of interactive proof-systems. *SIAM J. Comput.* **18,** 186–208 (1989).
6. Chazelle, B. The security of knowing nothing. *Nature* **446,** 992–993 (2007).
7. Fuller, J. Verification on the road to zero: issues for nuclear warhead dismantlement. *Arms Control Today* **40** (10), 19–27 (December 2010).
8. A general Monte Carlo N-particle (MCNP) transport code. Release MCNP5-1.40 http://mcnp.lanl.gov (Los Alamos National Laboratory, 2005).
9. Roquemore, A. L., Jassby, D. L., Johnson, L. C., Strachan, J. D. & Barnes, C. W. Performance of a 14-MeV neutron generator as an in situ calibration source for TFTR. In *15th IEEE/NPSS Symposium on Fusion Engineering* 114–118 (IEEE, 1993).
10. Hall, J. Uncovering hidden defects with neutrons. *Sci. Technol. Rev.* www.llnl.gov/str/May01/Hall.html, 4–11 (May 2001).
11. d'Errico, F. Radiation dosimetry and spectrometry with superheated emulsions. *Nucl. Instrum. Methods Phys. Res. B* **184,** 229–254 (2001).
12. d'Errico, F., Nath, R., Lamba, M. & Holland, S. K. A position sensitive superheated emulsion chamber for three-dimensional photon dosimetry. *Phys. Med. Biol.* **43,** 1147–1158 (1998).
13. d'Errico, F., Di Fulvio, A., Maryañski, M., Selici, S. & Torrigiani, M. Optical readout of superheated emulsions. *Radiat. Meas.* **43,** 432–436 (2008).
14. Bleuel, D. L. *et al.* Neutron activation analysis at the National Ignition Facility. *Rev. Sci. Instrum.* **83,** 10D313 (2012).
15. Thermo Scientific P 385 Neutron Generator. www.thermoscientific.com/en/product/p-385-neutron-generator.html.
16. Allen, K. *et al.* UK-Norway Initiative (UKNI) approach for the development of a gamma ray attribute measurement system with an integrated information barrier. In *Proc. ESARDA 35th Annual Meeting (Ispra, Italy, June 2013)* (ed. Sevini, F.) (European Commission, 2013).
17. Goldreich, O., Micali, S. & Wigderson, A. in *Proc. 19th Annual ACM Conference on Theory of Computing* (ed. Aho, A. V.) 218–229 (ACM Press, 1987).
18. Lindell, Y. & Pinkas, B. Secure multiparty computation for privacy-preserving data mining. *J. Privacy Confident.* **1,** 59–98 (2009).
19. Fisch, B., Freund, D. & Naor, M. Physical zero-knowledge proofs of physical properties. In *Proceedings of CRYPTO 2014* (eds Garay, J. & Gennaro, R.) (in the press); http://www.iacr.org/conferences/crypto2014/acceptedpapers.html.

Acknowledgements This project was supported by Global Zero and the US Department of State, the Princeton Plasma Physics Laboratory (DOE contract DE-AC02-09CH11466) and in-kind contributions from Microsoft Research New England. We thank D. Dobkin, F. d'Errico, J. Fuller, D. MacArthur, J. Mihalczo, S. Philippe and M. Walker for discussions and feedback. We thank S. Philippe (Princeton University) for graphics in Fig. 2. All simulations were run on Princeton University's High Performance Cluster.

Author Contributions A.G. was project leader and performed all of the MCNP5 calculations. B.B. and R.J.G. contributed, along with A.G., to developing the overall approach and the application of zero-knowledge proofs to warhead verification.

Author Information Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing financial interests. Readers are welcome to comment on the online version of the paper. Correspondence and requests for materials should be addressed to A.G. (aglaser@princeton.edu).