Monitoring and Verification of Nuclear Weapons

by

Richard L. Garwin
IBM Fellow Emeritus, IBM Thomas J. Watson Research Center
P.O. Box 218, Yorktown Heights, NY 10598

Email: RLG2@us.ibm.com
URL: www.fas.org/RLG/

Presented in the *Short Course on*
Nuclear Weapons Issues in the 21st Century
Sponsors:
American Physical Society Forum on Physics/Society & GWU Elliott School

The George Washington University Elliott School of International Affairs

Washington, DC

November 3, 2013

I was assigned the topic of Monitoring Nuclear Weapons about which I will have something to say, but I will concentrate on Monitoring and Verification of Nuclear Warheads.

A *nuclear weapon* comprises not only the nuclear warhead (or a bomb) but also the delivery vehicle which might be an aircraft, cruise missile, or ballistic missile. The *weapon system* includes not only the weapon but also its basing (silo, submarine, or airfield) and the many other subsystems required to make it work—notably, command and control.

Monitoring involves the cooperative or uncooperative determination of location, identity, and perhaps readiness of the nuclear weapon. This is perhaps simplest in the case of fixed surface-based missiles, with increasing difficulty as one goes to submarine-launched missiles, cruise missiles for ships, submarines, or aircraft, silo-based and land-based cruise missiles. In addition, there are gravity bombs and artillery-launched nuclear projectiles, of which I remind you that the United States previously deployed thousands in Europe, atomic demolition munitions, nuclear-armed torpedoes and depth charges, and nuclear-armed anti-aircraft missiles that can be launched from the ground, from aircraft, or from ships. The nuclear weapons can be in deployed status, in transit, or in storage at central or dispersed locations.

Monitoring can be performed by imagery intelligence (IMINT), signals intelligence (SIGINT), and HUMINT. Uncooperative monitoring can have a substantial margin or error, and for large forces the value of monitoring is not very sensitive to small errors.

More interesting for this audience is probably the monitoring of *warheads* and especially monitoring and *verification* of arms limitation agreements, including partial or full disarmament. Under those circumstances, the monitoring would be done in cooperative fashion, with which the United States and Russia (and especially the old Soviet Union) have had a lot of experience. Typically, only specific systems or classes of nuclear weapons have been subject to agreement and hence to verification. Beginning with the 1972 Ballistic Missile Defence Treaty and the accompanying Limited Offensive Agreement, strategic nuclear weapons—those with a range >5500 km were subject to the Agreement. An integral part of the agreement was the permissibility of *National Technical Means—NTM* (photography from satellites—initially the CORONA film-return system, as well as electronic intelligence gathered by collection of *telemetry)*. Indeed, the agreements specified that a party should not interfere with these NTM or conceal the *treaty-limited items (TLI)* by visual shields or encryption of telemetry during missile tests.

Under the START treaty of July 31, 1991 between the Soviet Union and the United States, ingenious means were instituted for counting deployed warheads.
For instance, arrangements were made for on-demand removal of the shroud of a silo-based missile, and the demonstration that no more than the permitted number of multiple warheads was mounted on that missile. This could be done without any detection of the warheads themselves, but simply by the use of "soft covers" over

vacant portions of the “bus” that was to be used to deliver the Multiple Independently targeted Reentry Vehicles (MIRVs) to their individual targets. Instruments for verification that needed to operate in proximity to the verified items were not “NTM” and needed to be revealed in detail to the other side and perhaps jointly built. Even a gravity gradiometer was developed to determine the number of MIRVs in a payload, from just outside the missile shroud.

For instance, in support of the *INF Treaty* banning intermediate-range and shorter-range missiles possessed by the U.S. or the USSR, the United States had the right for 13 years[1] after entry into force of the Treaty to inspect the facility at Votkinsk which was used to manufacture the SS-25 ICBM to ensure that the missiles emerging from the *portal* of the plant at Votkinsk were SS-25s and not the shorter SS-20 IRBM previously manufactured there. Similarly, the USSR and then Russia maintained a permanent presence at the Hercules plant in Magna, Utah, which had earlier manufactured the Pershing II IRBM. This portal monitoring was part of a

[1] http://www.opbw.org/verex/docs/sess2/wps/BWC_CONF.III_VEREX_WP.49.pdf

perimeter and portal monitoring system, which is a traditional way of monitoring a finite number of points.

Of course, verification would have been simple if the U.S. had been allowed to do detailed radiography of the entire missile in its canister, to be compared with templates of the SS-25 and of the SS-20. This would have been totally unacceptable to the USSR, which didn't want to reveal any technical details beyond the overall external configuration, which, itself, was masked by the canister. The agreement, then, was for the U.S. to radiograph the canister in the region of an "inter-stage" of the SS-25, which would give little or no information on the SS-25, but would show a significant difference if there were an SS-20. The fact that the radiography would have revealed sensitive information about the SS-20 was of no consequence, because no SS-20 in its canister would be presented to the portal for verification. Thus, cooperative monitoring in support of verification is often intentionally limited in the detail it can provide, as will be seen in what follows.

Clearly, monitoring in support of verification is particular to the nature of the treaty or agreement itself.

As regards the nuclear warheads, there has been a great deal of effort and of writing on this subject, some of which I have perpetrated myself, as is evident from a simple Google search with

site:fas.org/RLG/   verification nuclear warheads

Frankly, to be tiresome about it, the *"site:"* qualifier is an essential tool for any netizen, as I have emphasized in many of my talks, but to little avail.   If you carry out this search, you will find a modest number of papers or presentations on my site, some of which are quite useful, and from which I largely draw the remainder of this talk.

There is little dissent from the proposition that the U.S. and Russia and the entire world would be a safer place if there were a lot fewer nuclear weapons. Practically, to reduce nuclear weapons elsewhere in the world necessitates great and verifiable reductions of the 16,000[2] warheads that constitute 90+% of the world's nuclear warheads and that are owned by these two major powers.

Many believe that the world would be safer if nuclear weapons couldn't exist, and many also that if all nuclear weapons were destroyed, international security would be enhanced. A lot of the other talks bear on this matter—but as befits the APS, there is little in the way of policy argument here, and much in the way of science and technology to help you make up your mind about the benefit and credibility of various policy approaches.

Traditionally two logical approaches have been proposed for verifying that a package asserted by side A to be one of a certain class of its warheads can be affirmed to the satisfaction of side B to be just that.

---

2 Global nuclear weapons inventories, 1945−2013, Hans M. Kristensen and Robert S. Norris, Bulletin of the Atomic Scientists 2013 69: 75, http://bos.sagepub.com/content/69/5/75

In fact, some agreement might be adopted to limit or to eliminate all nuclear warheads, and the verification problem under those circumstances would be simply to show that an item presented on its way to destruction is a nuclear warhead and not a particular type of nuclear warhead. The more difficult problem is to show that there are no nuclear warheads or weapon-usable material outside the verification system.

To verify warheads bound for destruction, it has been proposed to use the *attributes* of a nuclear warhead. For instance, that it contains U-235 or Pu-239 and more than a certain amount. The other approach that has been considered in detail is the *template*, for which a warhead is asserted to be definitely a W-76[3] and not some other warhead or something pretended to be a warhead.

THE VALIDATION OF DECLARATIONS

[3] The most numerous warhead carried on the U.S. Trident SLBM.

We have discussed cooperative monitoring and verification of individual warheads, which with the use of statistical sampling can be useful in validating declarations of numbers and locations of warheads. Now we consider the monitoring and verification not only of warheads but also of weapon-usable fissile materials, drawing on a 2005 study by the National Academy of Sciences's CISAC (Committee on International Security and Arms Control)[4].

Here the question is how to feasibly and affordably validate a disclosure or partial disclosure that will enable one side to prove to the other that it has declared numbers, types, and perhaps locations of its weapons. This is discussed extensively with the proposal that each state in possession of nuclear weapons maintain a detailed roster of weapons, with a "line" or paragraph in the roster detailing the type, location, and status of each weapon; the line would be updated as the weapon was transported and relocated, or even destroyed for refabrication. Of course this roster would be highly sensitive, with access strictly limited. But looking forward to the control, to

---

[4] *Monitoring Nuclear Weapons and Nuclear-Explosive Materials: An Assessment of Methods and Capabilities*. Washington, DC: The National Academies Press, 2005. http://www.nap.edu/catalog.php?record_id=11265

massive reductions and even elimination of NW, the roster would be very useful if it could be *sampled* for accuracy of a declaration, without revealing more than one weapon at a time. This could be achieved by the open publication of a a roster in which the lines are replaced by *message digests*, which are obtained from the lines by the application of a publically available *hash algorithm* such as SHA-256. The key points are (1) that the later confidential revelation of the unhashed line can be verified with certainty to be identical in content with the information published even years before, and (2) no method exists, even in principle, to determine the line from its message digest.

CAN THE DESTRUCTION OF NUCLEAR WARHEADS BE VERIFIED?

Nuclear disarmament was furthered by the treaty between the United States and the Soviet Union eliminating nuclear forces of intermediate range ("INF"), specifically all their land-based cruise missiles or ballistic missiles of 500-5500 km range worldwide, whether armed with nuclear or non-nuclear warheads. The verification measures introduced to ensure that these missiles were destroyed and that replacements were not manufactured were extensive and relatively costly. They were carried out successfully.

The nuclear warheads from these missiles were, however, not required to be destroyed, and testimony to the United States Senate implied that such warhead destruction could not be monitored with confidence. I disagreed with that view, and judged that warhead disassembly and destruction can be verified adequately, with a proper agreement and facility, and that this should be done in the future, although such verification of destruction should not preclude agreements to verify numbers of nuclear weapon and nuclear weapon carriers in the forces. My view is that there was no will to destroy the cruise-missile warheads, which might be repurposed for other delivery means.

In brief, in support of potential agreements to limit nuclear warheads, nuclear weapons in the actual forces should be identified and fitted with an identifying tag and a "seal" that would follow them and ease their verification and counting, even though there was no treaty in force to limit their numbers. Of course, without an agreement there is no requirement to report these numbers to the potential partner(s) in a future agreement, but such tags and seals can greatly assist a nation to manage and control its own nuclear weapons.

In support of an agreement to destroy some nuclear weapons, such a sealed nuclear weapon would be delivered to the dismantlement site of the country owning that nuclear weapon. It would be checked into the site and into a specific room, which would otherwise be free of fissile materials. After the dismantlement, three streams of materials would emerge—high explosive and perhaps other hazardous waste, non-nuclear elements such as metals and plastics, and fissile materials (plutonium and uranium). It is these last that would be monitored, and in any case, it would be verified that no fissile material was left in the disassembly room, and that it could not have been spirited away. It is also necessary to establish that the claimed nuclear warhead is indeed a warhead or bomb of the claimed class, and the best way to determine that is to select the warhead at random from the deployed forces, maintaining joint custody of the warhead (aided by its tag and seal) from its dismounting from the missile carrier until it enters the portal of the dismantling facility.

In this regard, there has been a lot of analysis and experimentation on the mechanism of monitored/verified dismantlement, including a SIPRI book of 2003, in which I have a relevant chapter)[5] and, especially, a UK-Norway exercise organized and reported by VERTIC[4]. From this latter report of September, 2010, we take the concluding sentences:

"It is therefore, possible to state that, despite a number of unsatisfactorily resolved hurdles, there is nothing to suggest that the verification of warhead dismantlement is not technically feasible. And nothing, moreover, to suggest that dismantlement

---

[5] "Technologies and procedures for verifying warhead status and dismantlement," by R.L. Garwin, chapter in "Transparency in Nuclear Warheads and Materials: The Political and Technical Dimensions," Edited by N. Zarimpas, SIPRI, Oxford University Press, 2003, pp. 151-164. Presented at SIPRI Workshop, Paris, 02/08-09/2001. This chapter available at www.fas.org/rlg/Science&ND%20v2d.pdf

verification cannot be kept within acceptable levels of tolerance—both in terms of intrusiveness and reliability."

No one said that it would be easy, and much more work remains to be done. Fundamentally, what is required is assurance that the object that enters the dismantlement facility is indeed a nuclear weapon of the proclaimed type, and that the fissile materials are verified to leave the facility in a stream that is henceforth irrevocably committed to civil purposes.

ENTRY INTO THE VERIFICATION REGIME

A warhead or an amount of SNM in a container could be entered into the verification regime by affixing a simple tag and noting the type and sub-type of warhead, etc. It might also be necessary to note the type of container and the orientation of the warhead within the container. Deployed warheads, or warheads that are taken from containers and that become deployed, might have a tag affixed in an approved fashion and could also have a seal (such as the fiber optic purse of the referenced documents) to provide assurance that the tag still refers to the same warhead. At some later time, the system might acquire detailed information for validating the identity.

To be specific, assume that a number of containers assert that their contents is a W-88 warhead in a Mk-5 reentry vehicle. Two approaches are discussed for verifying that a TLI contains the warhead as stated. One of those is the measurement of "attributes", which might include at least a minimum mass of SNM, etc. I believe that the attributes are, perforce, so inaccurate that they do not provide reasonable verification or even transparency. For instance, if some warheads have as little as 4 kg of Pu, whereas others have 6 kg, the attribute for Pu would have to be set well below 4 kg, in order not to have too many rejections of those low-Pu warheads. A 6-kg primary could be converted into a 4-kg warhead (if that were possible) and 2-kg of Pu sold. To the extent that an attribute system depends upon tags and seals, further measurement of attributes can be dispensed with, because they add little, in the case of nuclear warheads.

Considerably more confidence can be placed in a *Passport* (Russian term) or *Template* (American term) approach which makes use of detailed and precise measurements of radiation characteristics of the TLI. The data obtained are sufficiently detailed so that they would provide useful weapon design information. While the U.S. and Russia might ultimately be willing to exchange such information, to release it into an international system would foment nuclear weapon proliferation rather than inhibit it. So the idea is to have precision measurement combined with an information barrier. For instance, in 1989 Brookhaven National Laboratory demonstrated the Controlled Intrusiveness Verification Technology (CIVET) approach, using computers without persistent memories to make decisions without the release of sensitive data. This was demonstrated to Russian personnel in 1997 at Oak Ridge in preparation for the Mayak storage facility. Similar systems have been demonstrated at Sandia National Laboratory and at the Pantex facility. More recent measurements taken with Sandia's Trusted Radiation Inspection System (TRIS) show the power of the (template) system in discriminating among 15 objects (8 pits, 5 fully functional bombs or reentry vehicles, and 2 secondaries).

Fig. 1 indicates the gamma ray spectrum of plutonium. The various lines can be used to help distinguish (in attribute or template fashion) so-called weapon plutonium from civil plutonium, although that distinction is not important in the weapon usability of plutonium. Nevertheless, it could prevent the substitution of civil plutonium for something that claims to have come from a nuclear weapon.

The energy region 635 to 665 keV is expanded in the lower right of the curve.

Pu-600—Weapons-grade Pu / Presence of Pu

- **Weapons-grade plutonium:**
  - For weapons-grade plutonium: $^{239}Pu + ^{240}Pu \approx$ Total Pu
  - Therefore, if the ratio $^{240}Pu/^{239}Pu$ is low, the material is weapons-grade
  - This value is also used in conjunction with the neutron multiplicity data to determine the plutonium mass
- **Presence of plutonium:**
  - 646 and 659 keV peaks.
  - Determination of presence requires *both* the 345 keV peak from Pu-300 and the 646 and 659 keV peaks from Pu-600

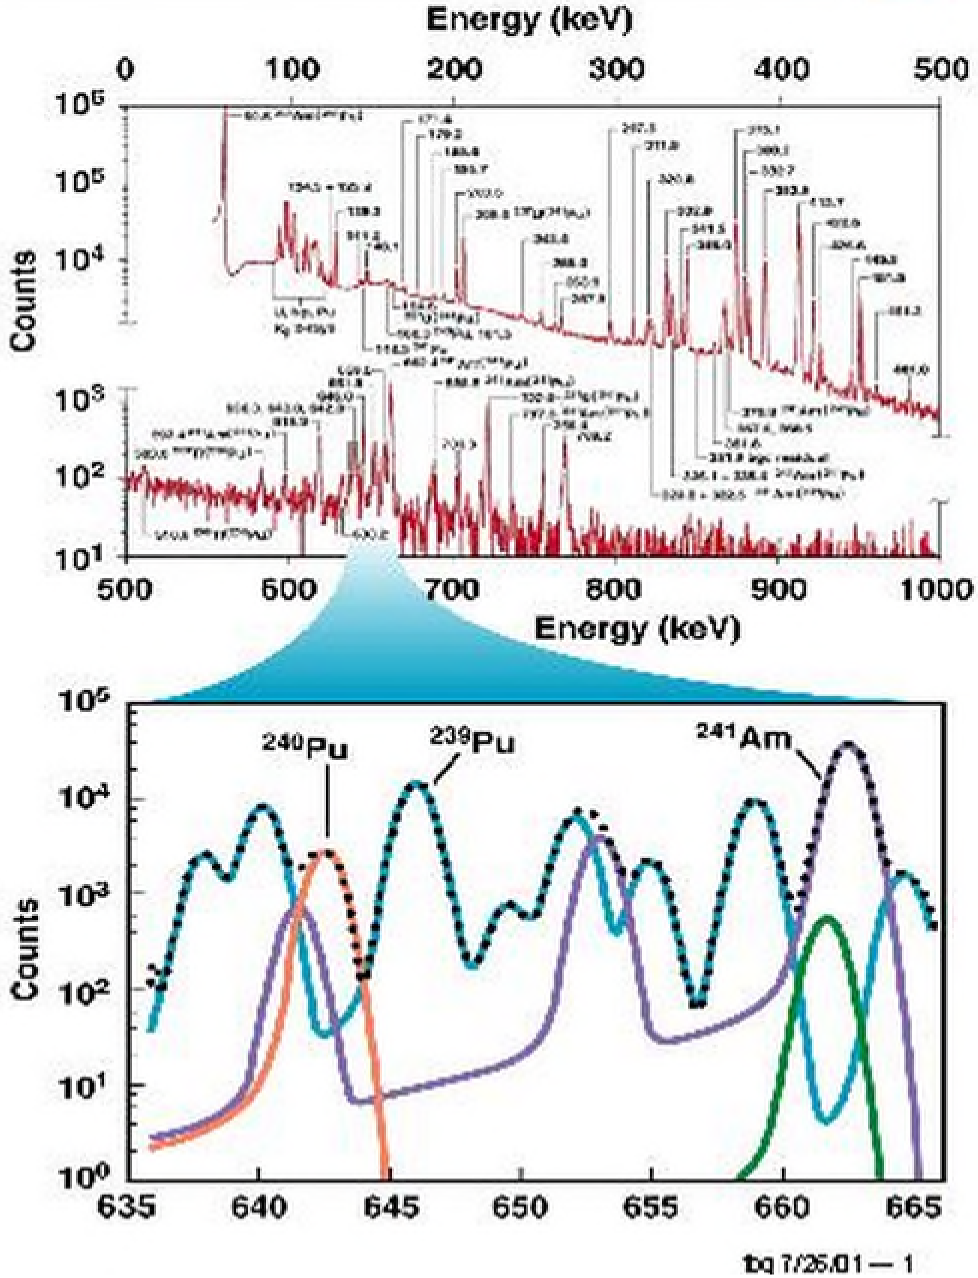

Fig. 2 [labeled "10"] is a simple illustration of the external observation via neutron counting of a sphere and a flat disk of plutonium. The two can readily be distinguished.

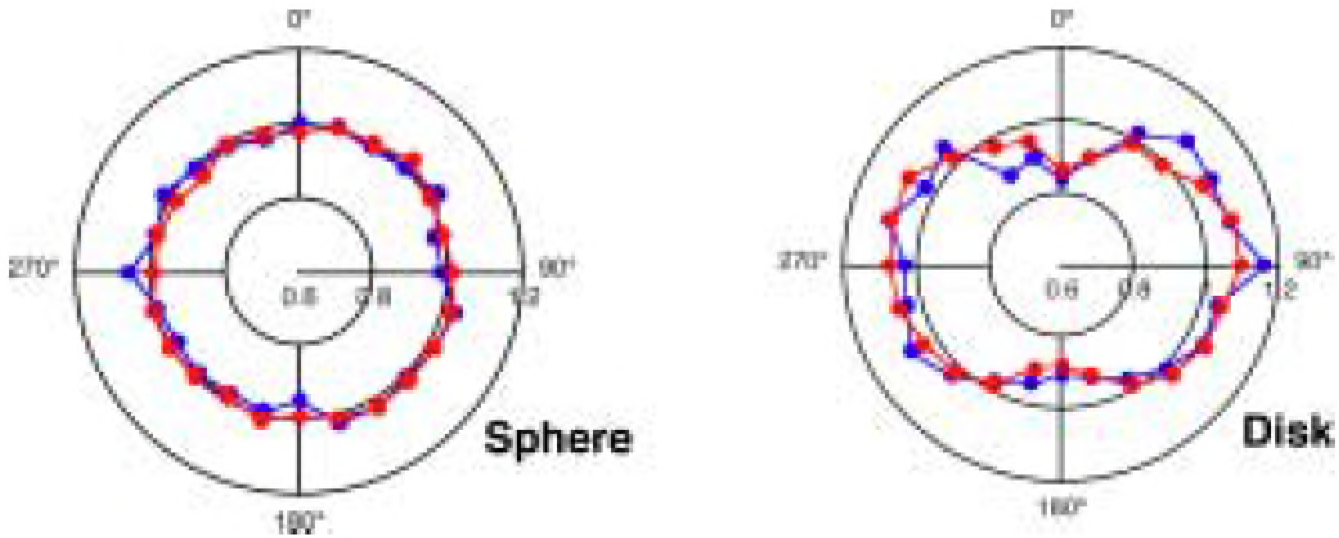

Figure 10. (a) Polar plot of the neutron counts recorded from a plutonium sphere in an AT400R container during the 1996 joint experiments. The circular pattern indicates the presence of a cylindrically symmetric object in the container. (b) This polar plot produced from counts recorded from a thick plutonium disk shows that the neutron field is anisotropic, indicating asymmetry.

Fig. 3 shows the regions used in a template approach with a high-resolution germanium detector in order to verify that a given item claiming to be one or another nuclear device is what it claims to be.

Fig. 3 - Energy group structure used to analyze low-resolution spectral data

Energy Range (keV)	Principal Significance of the Energy Group	Template Uncertainty (%)
80 - 120	U and Pu x-rays	10
120 - 160	continuum	1
160 - 172	sensitivity to energy-calibration error	exclude
172 - 198	^{235}U at 186 keV	1
198 - 230	^{237}U at 208 keV, variable in plutonium	exclude
230 - 290	continuum	1
290 - 350 plus 390 - 500	^{239}Pu full-energy peak region (change in sum of counts is insensitive to energy calibration error)	1
350 - 390	^{239}Pu full-energy peak region	1
500 - 600	continuum	10
600 - 711	^{241}Am at 662 keV, variable in plutonium	20
711 - 821	^{238}U at 766 keV	2
821 - 936	continuum	20
936 - 1090	^{238}U at 1001 keV	1
1090 - 1200	continuum	5

1200 - 2480	continuum from ^{238}U and ^{232}U	20
2480 - 2750	^{232}U at 2614 keV, variable in HEU	30

In Fig. 4, the templates are arrayed in the first row, with "Px" one of eight pits; the five "Fx" fully functional bombs of different types; and the two "Sx" being two secondaries in their canisters.

Fig. 4 - Average χ^2 for Comparisons of Measurements with Empirical Templates

		Template															
Source	**#**	**Back**	**PA**	**PA***	**PB**	**PC**	**PD**	**PE**	**PF**	**PG**	**FB**	**FC**	**FD**	**FE**	**FF**	**SB**	**SF**
PA	1	1285	**.6**	30	86	102	98	89	140	98	1039	75	414	680	782	5223	789
PA	2	1247	**1.7**	26	78	101	96	86	136	97	940	69	378	635	766	4565	765
PA	3	1298	**1.3**	31	85	97	92	86	141	92	1057	75	436	702	814	5056	817
PA	4	1320	**.9**	34	93	105	100	92	143	100	1123	84	457	728	817	5544	827
PA*	1	1034	47	**.7**	61	120	118	83	130	111	572	25	180	394	560	3391	547
PA*	2	1021	42	**1.1**	60	122	119	83	124	112	550	29	169	375	536	3417	524
PA*	3	1030	43	**.9**	65	120	118	87	135	111	569	25	171	379	533	3594	524
PB	1	1009	121	107	**1.2**	93	91	15	27	81	558	91	319	547	794	1380	760
PB	2	1008	117	101	**1.0**	95	93	14	25	84	548	86	304	528	771	1427	740
PB	3	1000	119	103	**1.0**	95	93	14	25	83	542	88	305	526	770	1364	739
PB	4	996	117	101	**.9**	95	93	15	25	83	541	89	305	528	772	1361	740
PC	1	2023	497	740	698	**.7**	**1.0**	263	398	7.8	2126	808	1777	1762	1860	4010	1846
PD	1	2016	496	733	681	**1.2**	**.8**	253	385	5.8	2112	794	1763	1755	1857	3985	1842
PE	1	1328	179	195	32	84	82	**.8**	25	70	923	177	605	863	1108	1994	1081
PF	1	1284	172	190	36	128	124	20	**.7**	112	858	203	566	821	1071	1836	1045
PG	1	2003	492	710	623	9.0	6.8	221	342	**.5**	1982	761	1665	1719	1862	3527	1843
FB	1	317	129	84	116	156	156	131	139	152	**.8**	92	32	7.7	42	400	27

FB	2	312	129	82	113	154	153	128	137	150	**.9**	90	31	8.2	45	381	29
FB	3	312	123	83	113	152	151	128	135	149	**.8**	91	32	8.5	45	386	29
FC	1	973	110	29	84	123	124	99	181	115	496	**.8**	140	336	491	3102	480
FD	1	540	113	51	90	139	139	111	142	134	63	43	**.9**	34	128	1217	113
FE	1	557	157	99	145	189	188	161	183	184	11	102	26	**.6**	46	790	33
FF	1	511	211	154	205	260	257	223	227	254	55	174	86	31	**1.0**	1477	6
SB	1	123	134	112	125	142	141	130	130	140	52	118	88	64	66	**.8**	57
SB	2	118	136	113	125	142	142	130	130	140	54	118	89	66	66	**1.3**	58
SB	3	130	135	114	128	144	143	133	134	142	53	120	90	65	65	**.8**	57
SB	4	121	135	113	126	143	142	131	131	141	54	119	90	66	66	**1.5**	58
SF	1	414	189	139	181	229	227	197	199	225	27	156	77	22	6.4	861	**.5**

* Pit A inside a shipping container (all other pits are in AL-R8 containers).

The first column indicates the actual source presented to the counting system. All templates and sources were in a standard "AL-R8" container, while PA* was inside a shipping container instead.

The "goodness of fit" of the measurement vs. the templates is shown in the cells of the illustration. A goodness of fit below 2 is never seen except when a particular source is compared with the proper template, with the exception of PC and PD. These two pits are almost identical and hence are indistinguishable by this technique.

Since a template measurement system with an information barrier (TMS/IB) might simply illuminate a red light if the claimed TLI was not one of the claimed class, and a green light if it was, no secret information could be legitimately obtained in the process. However, the inspected party would want to make sure that detailed clandestine measurements were not being taken, and for this reason the measuring instrument is always assumed to remain with the inspected party.

However, the inspecting party has much greater concern: First, that the green light does not automatically come on after a "counting interval" whatever the content of the container. Second, if the TMS/IB is truly making measurements of the claimed TLI, then how can one be sure that the template has not been changed to accord with what the TLI actually is?

The approach is one of joint preparation of templates in a "trusted system" which is fully understood by both sides and perhaps chosen among several available for the task. The result is a template prepared either from a "golden warhead" (in analogy to the "golden chip" for automated inspection of semiconductor products), or as an average of several warheads claimed to be identical and measured by the TMS/IB in a cooperative fashion to be similar in characteristics.

Again, the Secure Hash Algorithm SHA-1 comes into play in assuring the inspecting party that the template that remains in possession of the inspected party has not changed.

Much more can be said about such measurements and their supplementation by heat generation from a Pu pit, container weight, and the like. Similarly, a lot can be said about tags and seals, but this is no place for an exhaustive presentation.

PROBLEMS RAISED BY LARGE FORCES AND SMALL FORCES

Especially in support of the total elimination of nuclear warheads, there is and should be concern that some small fraction have been hidden and kept out of the overall agreement. This is a valid concern addressed but not fully solved in the 2005 CISAC study.

Another problem is that raised by small forces of nuclear weapons, such as those held by North Korea and with a considerably larger number, China. Here it has been pointed out to me[6] that if there are very few warheads of each type, the template approach is less useful or may even lack any utility. Thus if each warhead is handcrafted and quite different, there is no template that would fit more than one

[6] Personal communication from Zhu Jianyu 10/29/2013.

warhead and if warheads are presented for verification on the way to destruction, some other approach besides the template or the "golden warhead" must be used.

This observation is valid and leads one to ask why the template approach is accepted as viable for states with large numbers of warheads. This stems from our assumptions that the template would be created from one of a number of claimed-equivalent warheads that would be demounted from the delivery vehicle for the purpose of creating a template. Each of the warheads would be accompanied by a statement of when it was manufactured and when mounted on the delivery vehicle. In fact, a template might be created from the average of the detailed observations of, say, five individual warheads. Still on the "red" side of the monitoring equipment for the creation of the template, several such images or data sets could be obtained and not only averaged, but also their dispersion measured. This relates to the "*initialization process*" for verification using an *"information barrier."* A lot of this hand waving would need to be replaced by solid analysis and observation beyond those simple cases that I cite.

A more recent experiment carried out by teams including the UK Atomic Weapons Establishment—AWE-- and the Norwegian Defence Research Establishment supported by VERTIC has validated elements of cooperative verification of dismantlement of a simulated nuclear warhead, and has also shown problems that must be overcome with some of the elements of verification approaches[7]. The experiment was oriented toward verification that relied on attributes of a nuclear weapon and not precise templates of specific nuclear weapons. Nevertheless, it explored the "strands" of managed access and information-barrier technology.

Nevertheless, the small-number problem is real. It could be addressed by an *attribute* approach, tightened to include other aspects of the claimed weaponry, and also a detailed and perhaps verifiable history of the particular warhead, including when it was mounted on its carrier, and when it was moved most recently.

---

[7] *Verifying Warhead Dismantlement: Past, present, future*, by David Cliff, Hassan Elbahtimy and Andreas Persbo, VERTIC RESEARCH REPORTS NUMBER 9, SEPTEMBER 2010, http://www.vertic.org/media/assets/Publications/VM9.pdf

In any case, the verification of absence of weapon-usable fissile material must also be considered for those with small forces as well as those with many nuclear weapons. This was a major concern of the 2005 CISAC report, which provides an extensive discussion and many references for those who wish to delve deeper.

In general, this presentation should be regarded as a pointer to many documents now widely available on the web, which provide more details than I can present here or than you can take away.