6. Tags

Steve Fetter and Thomas Garwin

I. Introduction

An agreement on Conventional Armed Forces in Europe (CFE) may place numerical and geographical limits on more than 140,000 treaty-limited items (TLIs)[1] in 21 countries. Monitoring limits on such huge numbers of TLIs would be extremely difficult, as well as expensive and intrusive, with human inspectors alone. This chapter examines a promising way to effectively monitor limits while reducing cost and intrusiveness: the tagging of TLIs. The use of tags transforms a numerical limit into a ban on untagged items. The result is that many of the verification advantages of a complete ban can be retained for a numerical limit.

Tagging works by certifying that every TLI observed is one of those permitted under a numerical limit. A tagging system would involve the manufacture of a number of tags equal to the number of TLI, which would then be affixed to an essential part of each allowed TLI. If even one untagged TLI were ever seen—during on-site inspections (OSI), by national technical means (NTM), or even by nationals of the inspected party loyal to the treaty regime—then there would be *prima facie* evidence of a treaty violation. If properly designed, tags could also identify a TLI as belonging to a particular nation or as normally deployed in a particular region, which would make it easier to verify CFE sub-limits on national and regional deployments.

Other methods of counting a deployed force can only suggest that the allowed total is being exceeded, an indication that is unlikely to be conclusive and which might tend instead to cast doubt on the information going into the count. Tagging produces a much stronger impetus for political action in the event of a violation, because observation of an untagged system would provide unambiguous evidence of an overall violation. Since tagging allows the inventory of TLIs to be sampled, high confidence can be obtained with a smaller number of inspections.

All verification procedures have a common goal: to raise the political risk, the technical difficulty and the economic cost of cheating. No system can eliminate all possibility of cheating, but cheating can at least be made risky, difficult and expensive. For example, tags could not discover hidden stockpiles of undeclared weapons, but they could make it impossible to mix those weapons with tagged weapons. Depending on the facilities that would be open to inspection, this system of production, assembly, storage, testing, training, repair and deployment for its secret stockpile. Not only would the economic cost of such covert stockpiles be much higher than that of allowed TLI, but the risk of being caught—simply by an accident that exposed an undeclared TLI to the light of day—could well outweigh any military advantage that might otherwise have been gained from the undeclared inventory.

II. General characteristics of tags

Tagging systems have three key ingredients: a number of tags equal to the allowed number of TLI, a mechanism for associating a tag with a particular TLI and a protocol for verifying the authenticity of the tags. In most applications, checking tags would be an aspect of on-site inspections, but systems are conceivable in which the authenticity of tags would be checked remotely. Tagging systems should have the following general characteristics:

Counterfeiting

It must be impossible to counterfeit the tag without detection (or counterfeiting must be prohibitively expensive). Otherwise, the monitored party could simply produce counterfeit tags to cover TLI deployed in excess of the limit.

To make it more difficult for the monitored party to learn how to copy tags, the tags could be replaced at intervals with ones using different anticounterfeiting techniques. There is great scope for using very subtle features of tags to prevent their duplication if some fraction of the provided tags could be recovered and tested in a laboratory. Such subtle features could include altered isotopic composition of a particular element in a particular part of the tag, the deposit of a monoclonal antigen within a fiber, or seemingly random imperfections in a printing or manufacturing process. In some cases, the anticounterfeiting schemes that have been developed to protect national currencies could be used. A nation attempting to counterfeit such tags could never be sure that the copy duplicated all the identifying characteristics of the tag. As a test before a tagging system was agreed upon, a prize could be offered to anyone who succeeded in defeating the anticounterfeiting scheme.

Spoofing

It must be impossible to spoof the tagging system, or to fool it into thinking that a valid tag exists where there actually is none. For example, it must be impossible to reroute signals between the tag reader and a counterfeit tag so that the tag reader would actually receive a return signal from a valid tag at another location. Although preventing spoofing would be straightforward if an inspector had direct access to the tag, special precautions would have to be taken if tag reading was accomplished remotely.

Swapping

It must not be possible to move the tag from one TLI to another without the knowledge of the monitoring party. Tags need not be unremovable—they must only indicate in some obvious way that they had been removed. If tag swapping was possible, then valid tags could simply be moved to TLI being inspected at a particular time and place, or at least to those systems more susceptible to inspection by the other side. The tag should be placed on an essential part of the TLI, such as the turret of a tank, the gun barrel of an artillery piece, or the engine block of an armoured personnel carrier, so that a limit on that component is nearly equivalent to a limit on the entire TLI. If tags were glued onto TLIs, it could be arranged for part of the tag to change color or melt if exposed to the solvent required for the glue employed. An analog is in use in the U.S. domestic economy: to discourage the illegal parts business, cars now are made with serial number tags glued to their major sheet metal parts, and owners are warned not to attempt to remove these labels.

Espionage

The tag should reveal only that information required for verification. In other words, tags should not be agents of espionage that collect sensitive data about the TLI or its deployment patterns. Parties might be unwilling, for example, to emplace tags that could reveal low rates of readiness previously unknown to the other side.

In particular, the tagging system must not aid the monitoring party in locating TLI in real-time, since this could render tagged TLI vulnerable to pre-emptive attack. Such position information might even allow terminally guided munitions to home on the tags during an attack. For example, a radio beacon attached to tanks would certainly allow them to be counted by satellite receivers, but it could also allow attacking warheads to home on those targets.

Concerns about espionage could be alleviated if the physical details of the tags were exhaustively disclosed to the monitored party, but this would restrict the use of sensitive technologies and may make the tags easier to copy or spoof. On the other hand, the use of open tag technology would make it easier to publicize evidence of treaty violations, since no sensitive sources or methods would be compromised.

The monitored party could also be reassured by providing twice the number of tags, half of which could be selected at random, disassembled, and returned. It would be impossible to verify that there were no secret aspects of the tag (as noted above, some subtle secret aspects would be useful to prevent counterfeiting), but it should be easy to verify the absence of transmitters, explosives, bugs, cameras, etc.

In some contexts it may be desirable that tags not uniquely identify particular TLI. The monitored party may be concerned, for example, that valuable information could be gained if the monitoring nation were able to trace the deployment history of individual tanks. In this case the tag should simply indicate that a particular tank is an allowed tank, not that it is ‘allowed tank number 13647’. While this would rule out the use of ‘fingerprint’ tags (see below), it is possible to imagine tags that are identical under normal inspection but that incorporate anticounterfeiting measures that would be verified in more detailed, but much more limited, inspections.

Reliability

The tagging systems must be extremely reliable in the full range of environments that the TLI might experience during storage, testing, training, repair and deployment.

This may include extremes in temperature, vibration, humidity, radiation, etc., and some degree of deliberate abuse or tampering.

The tag must also have a very low false-alarm rate. False alarms not only undermine the mutual trust of parties which a treaty otherwise might engender, but, in sufficient number, they could create a background against which cheating would become easier. Designers of tagging systems should give some attention to reducing the possibility that the monitored party could deliberately act to increase the false-alarm rate as a prelude to an episode in which illegal TLIs would appear in transit or in repair and then concealed.

The physical size and power requirements of the tag should be such that the normal functioning of the tagged TLI would not be impaired in any way. Once again, the use of open tag technologies combined with the random inspection of tags should reassure the monitored party that the tag could not somehow interfere with or harm the TLI.

Cost

The tagging system must not be excessively costly. But even if a tag cost as much as a new car (about US $15,000), the total cost of tagging the 142,200 tanks, artillery pieces, armored combat vehicles, aircraft and helicopters that NATO would like to permit under CFE would only be 2 billion US$—a small fraction of the hundreds of billions of dollars presently spent each year in maintaining forces. It seems very likely that an effective tagging systems can be designed within this constraint.

III. Tag technologies

The simplest type of tag is a number painted on the side of a TLI. For example, allowed tanks could be given numbers 00001 through 20000. (Engraved serial numbers or bar codes printed on bumper stickers would also work.) The number could then be read by an on-site inspector with field glasses, by low-flying aircraft, or, if large enough, by a photoreconnaissance satellite. The problem with such simple schemes is obvious: how would the monitoring party know that each number or bar code was unique? If inspection occurs at a distance, counterfeiting would be easy.

Even so, this simple tagging scheme can provide significant benefits by allowing the tank population to be sampled. For example, consider a population of 20,000 allowed tanks distributed at 50 declared sites (400 tanks/site), which is illegally increased by 10 per cent by painting duplicate numbers on the excess tanks. (No one site has tanks with duplicate numbers.) Suppose that each week two sites are randomly chosen for inspection, and during the inspections 10 per cent of the tanks are randomly chosen for inspection. The chance that two of the tanks chosen would have identical serial numbers (and thus indicate a violation) is about 1.4 per cent; after 51 inspections (i.e., one year) the probably of detecting a violation would be 50 per cent.[2] Of course, larger violations would be detected more quickly.

Fingerprint tags

Tags that are manifestly unique, on the other hand, reveal the real power of tags because each observation of a tag would give unambiguous evidence of compliance or violation. If unique tags were used in the above example, there would be a 98 per cent chance of detecting a 10 per cent violation after a single inspection at only one site.[3]

Unique tags could be based on the observation of an unreproducible pattern or ‘fingerprint,’ such as a three-dimensional image of fibers in a piece of fiberglass, the detailed roughness of a metal surface, or the reflection of light from flakes of glitter suspended in epoxy. The pattern itself could be public knowledge, since the principle of the tag is the unreproducibility of complex three-dimensional patterns. If the tag were an intrinsic characteristic of the TLI, duplication, spoofing, or swapping would be extremely difficult. The tag reading would be done with the same type of instrument used for the initial imaging, which would almost certainly require on-site inspection in close contact with the TLI. The drawback with all such schemes is that reading the tag would be time-consuming. The unique identification of TLI implied by this method could conceivably be objectionable to some parties.

A central requirement of fingerprint tags is that the pattern not change significantly with the passage of time. For example, a hole could be drilled in the tank turret and the roughness of the surface at the bottom of the hole recorded; the surface could be preserved by screwing a cap into the hole. Tampering could be prevented by recording the ultrasonic signature of the resulting cavity.

In practice, tags would combine a simple serial number with a unique pattern. Assume, for example, that an inspector visits a site with 100 artillery pieces. Inspectors could verify the presence of a bumper sticker or hand-painted number on all the artillery pieces, but only a small fraction of these would be selected for a more detailed inspection of the tag to verify its authenticity. For example, a detailed photograph of the original hand-painted number could be compared with the number under observation. Alternatively, rubber castings could be taken of engraved serial numbers or fiber patterns in a bumper sticker could be recorded. Sandia National Laboratory in the United States has developed the ‘glint’ tag, in which glitter paint is applied over a serial number; illuminating the glitter from several different combinations of angles then produces a unique pattern that can be analyzed with a camera and a computer.

Electronic tags

Another way to produce a unique tag is through the use of coded electronic signals. Electronic tags have many advantages: no pattern recognition is necessary—the tag could produce a simple ‘yes’ or ‘no’ answer; this answer could be readily determined without direct access to the TLI on which it is emplaced; and the identity of a particular item need not be divulged. On the other hand, it will be much more difficult to negotiate a treaty that makes use of electronic devices rather than a simple tag or fingerprint.

Electronic tags could be authenticated by a direct connection to a local or remote console, or the tag could be equipped with a low-power infrared (IR) transponder (much like a television's remote control), thereby allowing the tags to be queried from a few tens of meters away. Such tags are already in use in commercial assembly lines for inventory control. This would reduce the intrusiveness of OSI yet not provide a homing capability that would make the TLI vulnerable to attack. (If homing is a worry, the tags could be disabled during a crisis.) An extension of this idea would be to fly a pilotless airplane over the site to query tags.

If parties to an agreement do not object to tags that would identify individual TLIs, then the simplest electronic tag would work as the equivalent of a 'one-time pad.' Electronic access to the tag's memory would only be allowed following the input of a special code unique to the particular tag and to the number of times it had been read previously. Each tag would report its serial number, the number of times it had been inspected and a unique secret number for that serial number and index. The secret numbers would be compared with a master list to authenticate the reading. Each secret number would be erased after it was read, and the series of secret numbers would be different for each tag. The monitored party could know all the information that was transmitted to and from the tag during this process and yet could not use this information to counterfeit tags.

If the tags cannot indicate their identity then their inputs and outputs must be identical; this is possible with modern cryptographic schemes. In either case the tag would have to be protected against nonelectronic means of discovering the series of secret numbers. For example, the chip could be shielded and provided with a membrane that would trigger a self-destruct mechanism if violated. (Similar mechanisms are used by the USA to protect nuclear weapons against unauthorized use.) In the case of unique tags, each tag need only protect its codes against tampering that does not leave any indication of misuse. In the case of identical tags, one must prevent even destructive means of discovering the secret codes, since a few tags could be sacrificed in a counterfeiting effort.

IV. Tagging treaty-limited items

The following examples are intended to show the weaknesses as well as the strengths of the tagging concept.

Soldiers

Limits on the number of troops in Central Europe have been under discussion since the mid-1950s, as have schemes for monitoring such an agreement. Usually a continual presence of inspectors or remote monitoring equipment at checkpoints supplemented by occasional forays by human inspectors has been thought to be required. It is unclear, though, how the observation of an unusually large number of troops in a particular region would be anything more than an occasion for suspicions that could not be easily resolved. Conversely, the intrusiveness required to monitor agreed force dispositions might yield evidence of force weaknesses that in a crisis could make the military balance less, rather than more, stable. With a tagging system, however, the discovery of a single soldier without proper identification (i.e., without a tag) would be conclusive evidence of a violation, yet no information need be collected about either the overall number of troops in the region or their disposition.

A tagging system for troops might work in the following manner. Suppose that limitations were imposed on the total number of active military personnel in each of several zones. At random or fixed intervals (say every six months) the monitoring party would supply enough ID cards (tags) so that the monitored party could issue one to each soldier in the zone. The ID cards would have a section where a thumb print could be registered within two or three days of the issue of the card.(Chemicals in the card could ensure that after this active period either the thumb print would no longer register or the card itself would indicate that a longer delay had occurred.) Every soldier in a controlled zone would be required to carry the appropriate ID card with his or her own thumb-print. Transfers of soldiers could be accommodated by the exchange of used ID cards for new ones.

With such a system in place, if an inspector ever found a soldier without a valid ID card, there would be a clear violation that could be investigated directly. Moreover, soldiers without valid ID cards might be aware of this fact themselves, and might chose to reveal this to inspectors. In most other verification schemes the total number of soldiers in the zone would be inferred from the number and types of units observed to be deployed there, or from some other set of imprecise measures that would not identify any specific individual as constituting a breach of the limit, even if they gave some general indication of a violation.

In practice this tagging scheme would have to be elaborated in great detail. Most obviously, there is the technical design of an ID card that could be personalized by thumb print within the required time period and not simply provided to troops just before an inspection. Since the cards would be provided by the monitoring party, and since inspectors could randomly recover a small fraction of the ID cards and return them to the laboratory for detailed analysis, occasional changes in the details of card technology could be used to ensure over time that there were was no counterfeiting or misuse of the tags.

The example suggests several other aspects of any tagging system. First, as implied by the mention of inspectors, tagging only works if there is some chance of observing the controlled items and the presence or absence of associated tags. In the case of ground troops, it is assumed that inspectors would be given fairly free access to transit routes, if not to all military bases. The personalized quality of the ID card would ensure that no single tag could be used to provide safe transit for a succession of soldiers, who would then disappear into uninspectable bases or other safe havens.

Finally, the example suggests that while tags can help ensure that a precisely defined limit is not exceeded, there are many potential verification problems for which tags would provide no help at all. If an inspector, for example, came upon an individual in uniform with an automatic weapon but no tag, it might be explained that the person was a police officer or some other quasi-military officer (e.g., a customs agent) and not a soldier at all. Moreover, soldiers travelling out of uniform and separately from their weapons on public transport or in cars may not be identifiable as such. Tagging cannot remedy imprecise definitions of what is controlled by an agreement. Only if the parties can agree on a clear definition of who is a ‘soldier’ can numbers of ‘soldiers’ be controlled.

Tanks, artillery and armored combat vehicles

Compared to limits on people, limits on hardware are in some ways easier and in some ways harder to verify by tagging. Certain major classes of military hardware have no civilian use and so cannot merely blend into the civilian landscape. Such specialized hardware includes tanks, artillery, armored combat vehicles (ACVs), fighter-bombers, most bridging equipment and most munitions, but not jeeps, trucks, buses and transport aircraft. Observation of a tank leaves little question that it is a controlled item. For the same reason of singularity, though, close-up scrutiny of a soldier would reveal fewer military secrets than close-up scrutiny of an advanced weapon.

Several other differences make hardware harder to control through tags. First, there is less reason for any particular piece of hardware to emerge from hiding or to be involved in exercises and training. An opponent determined to violate an agreement could maintain a stock of tagged equipment to be used in peacetime operations and an untagged stockpile that would be kept out of view until shortly before the outbreak of hostilities. This problem is conceptually similar to the possibility of unknown stockpiles in an absolute ban on a class of weapons. A decisive difference, though, is that, in the case of a numerical limitation, troops would have the opportunity to train with the legal weapons of the same type.

Second, each hardware item has less of an essential identity than a person. With a thumb-print tag, one can be sure that the monitored party is not using a single tag and ‘transplantable thumbs’ to cover the transit of multiple people across inspected areas. With hardware, some care would have to be taken to design ways to take the equivalent of a fingerprint for each TLI, or to attach absolutely nonremovable tags to crucial pieces of the item. The complexity of the problem is indicated by the fact that a nonremovable tag on the fender of a tank would be of little help because the same fender could be unbolted and used *seriatim* to transfer large numbers of tanks to unknown storage warehouses. If the turret of a tank represented a large part of the value of the tank and the turret could not be easily removed or concealed, however, then a nontransferrable tag on the turret would suffice, because a limit on tank turrets would be equivalent to a limit on tanks.

As a technical matter, a nontransferrable, noncopyable tag for a tank turret is not hard to devise. As mentioned above, the surface roughness of a region of the turret (e.g., around an engraved serial number) would certainly be non-transferrable and difficult to copy, although it would also be difficult to authenticate and special precautions would have to be taken to prevent the fingerprint from changing with time. Electronic tags could use a capacitance, contact, or ultra-sound sensor, a fiber-optic cable seal, or two electrically communicating devices on opposite sides of the turret, to ensure that once emplaced it could not be removed without providing a tell-tale record. More simply, a limited amount of special epoxy glue made with unstable components and identifying trace elements or isotopes might be provided, or tags might be emplaced with ordinary glue and an ultrasound fingerprint of the resulting assembly recorded.

If several classes of tags were provided, or if the tags had serial numbers, then tagging could be used to control the number of tanks, artillery, or ACVs in each of several zones of interest, as well as the total in the overall region. The tags for each zone might be different colors and shapes, so that close inspection would not be required to ascertain that a tank was in an allowed area; only a small number of random close-in inspections would be required to verify that the tags were authentic. Such a scheme would be complicated, though, if tanks were routinely moved between zones.

If details about tank dispositions were not considered sensitive, the monitored party might simply be responsible for turning over to the monitoring party a roster of which tag serial numbers were in each zone prior to the beginning of each inspection period. Even if the dispositions were sensitive, the roster idea could be adapted using cryptographic techniques so that the monitored party could keep the overall roster secret while still providing assurance that any particular observed system was within a sublimit for a particular zone. Cryptographic or electronic means could be used to produce the equivalent of a system where the roster is deposited with a neutral and confidential judge who responds ‘yes’ or ‘no’ to queries of the form, ‘Is tank number 1197 allowed in zone 2?’.

Combat aircraft and helicopters

Tags could find similar uses in monitoring aircraft and helicopters, although there are some important differences. Although tagging jet engines would probably be sufficient in the case of combat aircraft, it is not so obvious what component of a helicopter could be tagged, since there are no major external components. Occasional detailed inspections of fingerprint tags on the motor may be required to ensure the authenticity of simple tags on the housing.

The most obvious difference between aircraft and heavy ground equipment is that aircraft are far more mobile; they can transport themselves into or out of a zone in a matter of minutes or hours. It is difficult to see how human inspectors, with or without tags on the aircraft, could deal with this fundamental problem through on-site inspection. Remotely-monitored tags or tags read by automatic sensors at the end of runways may be the only solutions (see below).

V. The role of tags in a verification system

The role tags play in a verification system is mostly determined by how and when the authenticity of tags would be checked: during on-site inspections, through remote telemetry, or at natural choke points or artificial portals in the monitored country.

Tags as an aid to on-site inspection

The most straightforward way to use tags would be in conjunction with OSI. The use of tags would provide a clear way in which information gained during individual OSIs could contribute to an overall judgment concerning compliance with a treaty. Without tags, OSI cannot produce much direct information about the total number of TLIs deployed unless all sites are inspected simultaneously—an extremely costly and intrusive option.

With conventional OSI inspection schemes (i.e., without tagging), one might learn that 30 aircraft were at site A in January, 40 at site B in June and 50 at site C in December, but there would no way to conclude whether or not the total number of aircraft at all sites at any one time exceeded the permissible limit. A periodically declared roster of how many aircraft were at each site would reduce this problem, but such a roster would not give confidence about the completeness of the count at a given site. If all allowed aircraft were tagged, however, one could tell if every aircraft found at whatever facility was part of the allowed inventory.

Tags could be affixed to an essential part of each TLI during an initial round of baseline OSIs. If the anti-swapping and anti-counterfeiting measures were sufficiently foolproof, one could simplify the process greatly by passing out the allowed number of tags to the monitored party, whose own personnel would affix them. There should be strong incentives for the monitored party to affix the tags promptly and properly, since not to do so would increase the chance that an untagged system would be discovered.

During an OSI, inspectors would locate TLI and attempt to verify the authenticity of their tags and to verify that the tags had never been removed. If tag reading was difficult, as would be likely for fingerprint tags, a random sample of the tags could be checked. Procedures would have to be worked out for the return of a tag when a TLI was destroyed or otherwise removed from the inventory.

If the verification regime permitted inspections on short-notice at the option of the monitoring party, they could be timed to take maximum advantage of national intelligence capabilities. The movement of TLI into or out of the facility to be inspected could be monitored closely by NTM or by special cooperative measures just prior to the event. Even if an untagged TLI were never actually found during an OSI, tagging could force a cheater into more obviously suspicious behavior. Moreover, because untagged TLI would have to receive special handling at all times, tags might greatly increase the number of people who knew of a treaty violation on the part of their country, increasing the likelihood that the violation would become widely known.

The use of on-site checking of tags in providing evidence of cheating—indeed, the use of any type of OSI for this purpose—should not be oversold, because access to the evidence would always be in the control of the monitored party. Although it is true that the detection of a single untagged TLI would be evidence of a violation, the monitored party would be unlikely to allow an OSI when such a possibility existed. It would always be advantageous from the cheater's perspective to make up excuses for delaying or denying an OSI rather than risk discovery of a 'smoking gun.' This may lead to a paradox of sorts, because if a tagging system were implemented the lack of a tag could become, in the eyes of the world community, the *only* acceptable evidence of a violation.

Thus, even though tags could provide unambiguous evidence of a violation with just a single observation, it is unlikely that this would ever happen during an OSI. The monitoring party probably would have to act on more ambiguous evidence, such as a refusal or delay of OSIs, surreptitious movement of TLIs out of declared facilities, tag tampering, or other suspicious behavior. However, tagging would have played a role in eliciting this suspicious behavior. Moreover, because of tagging's relative efficiency in detecting violations, tags should reduce the likelihood that a country would decide to cheat in the first place (which is presumably the main purpose of verification).

Tags read on-site could be an excellent way to help build confidence between parties who are in compliance with an agreement. Because tags make inspections more efficient, they would have the virtue of minimizing the number of inspections required for a given level of confidence. Tags also could reduce the chance that false claims of treaty violation would be used for political reasons.

Remotely monitored tags

Supplementing OSI may be the most obvious role for tags, but it is by no means the only conceivable role that tags could play. In fact, certain verification regimes may be easier to negotiate if requirements for OSIs, especially when involving trained foreign personnel at sensitive military locations, are minimized. If tags could be read remotely, routine OSIs may not be needed. Three basic schemes using remote reading come to mind: the tag could transmit a continuous or intermittent signal, the tag could be provided with a two-way communication link, or the tag could record position information for later interrogation.

The most obvious remote sensing method is for every tag to transmit a coded set of high-frequency radio pulses. The location of the tag could then be determined by satellite receivers using time-of-flight measurements. The obvious drawback of this scheme is that one might be able to home on the beacons during an attack. The monitored party might be given the ability to switch off the transmitters in time of crisis to ease this problem, but such a system might have the paradoxical effect of aggravating a crisis since switching off the beacons could be taken to indicate that the monitored party was preparing for war. Even worse, there could be pressures to launch an attack while the beacons were still on, or shortly after they were switched off, when the approximate location of the TLI would still be known. The very necessity of making such decisions would distract leaders from dealing with more substantial issues.

A better plan would be to have the beacons emit signals randomly and infrequently in time, so one would never know the location of a large fraction of the tagged TLI at any one time. For example, an inventory of aircraft could be equipped with beacons that emitted a signal once every ten days. If the aircraft were moved once per day, then the monitoring party would only know the location of ten percent of the inventory at any one time.

In another remote-monitoring scheme, each tag would contain a receiver that recorded position information given by a navigation system. This system has the advantage that the quality of the location information could be controlled. If, for example, the resolution of the navigation system is too great, then the system's output could be filtered to report only the number of a map square in which the tag could be found. After a period of time, the degraded information stored in the tags could be transmitted to the monitoring party. This transmission could be encrypted and security codes added to ensure the authenticity of the data. If the time delay were short (a few days), this idea would be operationally similar to the beacon scheme. Alternatively, the tags could be collected and sent back to the monitoring party and new tags issued. The tags themselves would then constitute a time-lagged data base of the position of every allowed aircraft. Tags of this type could be used to enforce regional limitations on TLI, such as the number of aircraft or tanks near the central front in Europe.

Of course, neither of these tagging systems could detect undeclared TLIs. The presence or absence of undeclared TLI would be verified by comparing the location information supplied by the tags to data collected by photoreconnaissance aircraft or satellites. For example, a photograph that showed a TLI at a location that was not recorded by any of the tags at that time would be evidence of a treaty violation.

These systems do not resolve all problems, however, and they create some of their own. First, they rely on NTM to detect violations. Since a cheater would be very careful not to expose untagged TLI to photoreconnaissance, the probability of observing a violation would be very small. One would probably have to depend on accidents to expose (or deter) cheating. Second, the system would place high demands on technology. It may not be possible to produce the type of tag described here—the receiver or beacon may simply be too large or require too much power. It also may not be possible to develop tags that are sufficiently reliable. If a photograph shows the location of a TLI but the tag records the position information inaccurately, then a false indication of a treaty violation would occur.

Monitoring tags at choke points or portals

Tags also could be used effectively at natural choke points or artificial portals—places through which TLI must pass at least occasionally. As an example of a natural choke point, consider a rail spur that leads into an army depot or a warehouse where equipment is stored. Imagine, for example, that a train loaded with tanks on flat-bed cars was approaching a choke point equipped with sensors. If the tanks had a valid tag, the sensors could authenticate them (perhaps using a bar-code reader or the IR transponder mentioned above). If an untagged tank tried to pass through the choke point, other sensors, such as scales, video cameras with pattern-recognition software, or x-ray machines, would determine that the object could be a tank. The monitored party would then be required under the verification regime to allow more intrusive inspection to prove that the object was not a TLI. As another example, tag readers could be placed at the ends of runways; whenever sensors detected the landing or takeoff of a treaty-limited aircraft (acoustic or IR signatures might distinguish combat aircraft from other non-TLI aircraft), they could expect to read a valid tag on the underside of the plane.

If a natural choke point could not be found, one could be created by surrounding declared facilities with monitored fences which force the movement of TLIs through a gate or portal where they could be observed and counted. The declared facilities could be any combination of production, assembly, storage, testing, training, repair and deployment areas. The fence, or perimeter, would be a two-dimensional barrier around the monitored party's facilities that could not be violated without detection. A wide variety of perimeter sensors could be used, including seismic detectors, microwave intrusion detectors, acoustic sensors, video and IR cameras, metal detectors, short-range radars, or pressure sensors. Possible monitoring devices at the portal might be video or IR cameras, weighing scales, x-ray, gamma-ray, neutron, or ultra-sound imaging devices, metal detectors and human inspectors. The perimeter/portal data could be transmitted to the monitoring party in a secure mode or interpreted by human inspectors stationed at the site.

Consider the case of a perimeter/portal system at an production or assembly plant. When a finished TLI was ready to leave the assembly plant, the monitored party could simply declare the TLI and the count of deployed systems would be increased. If the monitored party did not declare the TLI, monitoring devices at the portal would determine that the object *could* be a TLI. Unless further inspection was permitted to determine that the object was *not* a TLI, the monitored party would be in violation of the treaty when the TLI left the facility.

This system would have the advantage that declared TLI would not be inspected by intrusive devices at the portal. But in order to retain this advantage, TLI would have to be tagged before leaving the facility so that they could be returned for maintenance. Without tags, the monitoring party would have to inspect any returned TLI to ensure that the monitored party was not returning bogus TLI and replacing them with real TLI. Tags also would prevent covertly produced TLI from having access to declared production and assembly plants.

There are several disadvantages to perimeter/portal systems, especially when they are applied to a wide variety of declared facilities. First, both sides may be reluctant to allow the other to construct a perimeter composed of a wide variety of sensors around sensitive military areas and to allow intrusive inspections of any entering or exiting objects that the monitoring party claimed *could* be a TLI. The potential for gathering intelligence information that was not required for verification purposes would be obvious. Second, the perimeter/portal systems would be expensive—even more so if supplemented by a human presence. Third, such a system would necessarily be very complex, requiring perhaps hundreds of agreed rules governing the interpretation of data. Finally, perimeter/ portal systems probably would disturb the normal functioning of declared facilities. However, tagging provides a natural complement to perimeter/portal systems, allowing reduced intrusiveness within the controlled facilities.

VI. Conclusions

Tags are a technical fix—a gimmick that will only aid the negotiating process to the degree that those technical difficulties with verification that tags could ameliorate are delaying the completion of treaties. Even if tags could make limits on certain TLIs easier to verify, there may be other barriers to agreement. In so far as this is the case, tags could become part of the problem instead of part of the solution—a source of endless detailed technical discussion that could be used to obfuscate more fundamental differences. An agreement incorporating tags would undoubtedly be far more detailed and more difficult to negotiate than one without tags, especially if electronic tags and/or remote-monitoring schemes are used. Although the United States and the Soviet Union have shown an ability to negotiate technically complex treaties, such complications should only be introduced when an agreement would be impossible without them.

The authors are optimistic that tags could be designed that meet all of the generic requirements outlined above: resistance to counterfeiting, spoofing, swapping, espionage, homing, etc. More work is needed to explore the feasibility of tagging concepts and to define the overall verification system of which tags could be a part, because the tag technology needed will depend much more on the verification regime as a whole than on any general requirements that tags must meet. Once a promising verification system is defined that requires a certain type of tag, then the development of specific tagging hardware could go forward productively.

In summary, while tags are not a panacea for the problems of monitoring numerical limits on concealable weapons, they could have much to offer if part of a carefully designed system. To be truly available for CFE verification, tagging systems will have to be the subject of detailed discussion among the parties, including parallel technical research and development on both sides.

Notes and references

[1] In the CFE Negotiation, TLIs include tanks, artillery, armoured combat vehicles (ACVs), combat aircraft and combat helicopters. The NATO position is that the number of these. TLIs in the Atlantic-to-the-Urals zone should be reduced from about 270 000 to 147 200 (40 000 tanks, 33 000 artillery, 60 000 ACVs, 10 400 aircraft and 3800 helicopters). According to the WTO, the current number of these items is about 350 000; in addition, they propose limits of 40 000 for artillery, 56 000 for ACVs, 9400 for aircraft and 1.4-1.5 million combat personnel (Arms Control Today, vol. 20, no. 3 (Apr. 1990), p.5.

[2] On average, about 36 legal and 4 illegal tanks would be selected during each inspection at each site. The probability that none of the four illegal tank numbers at one site would match the number on the first legal tank selected at the other site is (20 000 – 4)/20 000, the probability that none would match the second legal tank is (19 999 – 4)/19 999, and so on. The probability that the excess tanks at both sites would have numbers that would match at least one of the numbers on allowed tanks at the other site is approximately $1 - [(19982 - 4)/19982]^{72} = 0.014$ per inspection, and the probability that n inspections would reveal a violation is $1 - (1 - 0.014)^{n}$.

[3] If there are 400 legal and 40 illegal tanks at a given site, then the probability of choosing a legal tank at random is (10/11). The probability of choosing 40 legal tanks at random (10 per cent of the legal inventory) is $(10/11)^{40} = 0.022$.