Verification Technology: Unclassified Version

S. Drell (Chairman)
P. Banks
C. Callan
K. Case
J. Cornwall
F. Dyson
D. Eardley
N. Fortson
M. Freedman
R. Garwin
J. Katz
S. Koonin
R. LeLevier
C. Max
R. Muller
A. Peterson
W. Press
B. Richter
S. Ride
M. Ruderman
J. Sullivan
S. Treiman

October 1990

JSR-89-100A

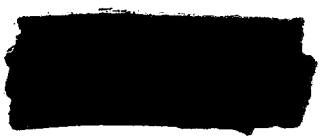

JASON
The MITRE Corporation
7525 Colshire Drive
McLean, Virginia 22102-3481
(703) 883-6997

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	October 15, 1990	

4. TITLE AND SUBTITLE

Verification Technology: Unclassified Version

5. FUNDING NUMBERS

PR - 8503Z

6. AUTHOR(S) S. Drell, P. Banks, C. Callan, K. Case, J. Cornwall, F. Dyson, D. Eardley, N. Fortson, M. Freedman, R. Garwin, J. Katz, S. Koonin, R. LeLevier, C. Max, R. Muller, A. Peterson, W. Press, B. Richter, S. Ride, M. Ruderman, J. Sullivan, S. Treiman

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

The MITRE Corporation
JASON Program Office A10
7525 Colshire Drive
McLean, VA 22102

8. PERFORMING ORGANIZATION REPORT NUMBER

JSR-89-100A

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Defense Advanced Research Projects Agency
1400 Wilson Boulevard
Arlington, VA 22209-2308

10. SPONSORING / MONITORING AGENCY REPORT NUMBER

JSR-89-100A

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION / AVAILABILITY STATEMENT

Government agencies and their contractors only;
all other requests to DoD controlling office.

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*

This report examines several technology issues relating to verification of nuclear weapons treaties. These include: non convertible design of cruise missiles, tags and seals, radiation detection and surveillance.

14. SUBJECT TERMS

Sea-launched cruise missile verification; PQC, tentpole design; inner shell design; tags and seals; detection of plutonium

15. NUMBER OF PAGES

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	SAR

Contents

1 INTRODUCTION AND SUMMARY

1.1 Introduction

This study analyzes several of the principal new challenges to effective verification of compliance to agreed limits to weapons now under discussion at the START negotiations. The new requirements are analyzed, new technologies are described, and specific proposals are presented for enhancing our capabilities to verify treaty compliance.

Prior to the ratification of the INF Treaty in 1987, the information on which the United States relied for verifying arms control treaties was derived primarily from three mutually recognized sources:

1. Formal data exchanges
2. Observation and analysis of systems tests for which alerting information was exchanged
3. National Technical Means (NTM) of observation.

(These were of course supplemented by human intelligence sources as available.)

At the INF negotiating table and in subsequent negotiations, four additional means and procedures have been added for verifying compliance with future treaties:

1. On-site inspection (OSI) of regulated activities and deployments
2. Perimeter portal monitoring (PPM) of declared facilities
3. Suspect site inspection (SSI) for illicit activities
4. Tags and seals for identifying legal systems.

These were judged to be required because of the added difficulties in verifying limits or bans on the broad range of weapons now being discussed: mobile ICBMs, cruise missiles that may be deployed on a wide variety of launchers at sea or in the air, and that are indistinguishable at a distance as to whether they are armed with conventional or nuclear warheads.

At the same time as we face increased complexity and difficulty in meeting the verification requirements of these new systems, it is important to keep in mind that costs–both operational and financial–as well as benefits come with elaborate and invasive means of verification. It is important that both the costs and benefits be carefully accounted for in deciding how much and what kind of verification is desirable.

Principal among the costs the US should recognize and try to minimize are:

1. Financial costs of staffing, equipment, operation and R&D, for US verification activities.

2. Security and inconvenience costs of intrusive Soviet verification for US military activities. An excessively intrusive inspection regime also raises serious concerns for the civilian industrial sector and its need to protect privileged processes and technologies. It can also conflict with the privacy protections of the 4th amendment to the US Constitution.

3. Political costs of overloading the US intelligence assessment apparatus with rigid verification requirements which results in frequent false alarms and acrimonious disputes over definitional and minor technical issues. An example of a definitional issue is the dispute between the two governments as to where to define the limits between allowed and forbidden activities in the ABM Treaty. A minor technical issue is exemplified by the prolonged arguments over the verification of the unratified Threshold Test-Ban Treaty (TTBT). The TTBT is itself of minor military importance to US security. Nevertheless, disputes over its verification have distracted the attention of intelligence analysts and political authorities from more important matters.

4. Instability costs of tying an international arms-control regime to a verification system which may be too burdensome and complex to withstand the shocks of international crises or major shifts in domestic

political climates. In a moment of crisis or loss of confidence a verification system that is seen to be excessively intrusive, complex, and not broadly accepted as fair and balanced may be repudiated and may bring down with it the whole fabric of arms control. Such an example would be the expulsion of an in-country team for on-site inspection if it is excessively invasive.

In developing the verification requirements for arms control treaties the primary purpose of arms control should be emphasized: it is to help stabilize the world in foul weather as well as fair. Also to make progress in arms control consistent with our natural interest we should not require more of verification than it can realistically do. Requirements for verification should be set, consistent with minimizing the four costs discussed above, to meet two necessary and sufficient conditions, neither more nor less:

- We must be able to detect violations of a scale that could upset the military balance and threaten our security
- We must be able to detect such violations soon enough to enable us to respond in a timely fashion.

These define the requirements of "effective verification." This study proceeds with this caution in mind.

1.2 Summary

Sections 2, 3, and 4 focus primarily on new challenges for verification that are presented by the current negotiations at START. Sections 5, 6, and 7 focus on current and developing technologies for increasing our verification capabilities.

In Section 2 we consider how one might design a new long-range air-launched cruise missile that is both unambiguously armed with conventional warheads and also not readily convertible to a nuclear delivery system. Such weapons could effectively attack a range of military targets from large stand-off distances if they achieve high accuracy (10's of feet). They raise a difficulty for arms control unless they are both unambiguously non-nuclear and also

not readily convertible. Otherwise the aircraft carrying them would be included among the accountable heavy bombers (AHB) according to counting rules under discussion at START. The technical difficulty created by this requirement results from the small size and weight of nuclear warheads with yields of the order of kilotons (a few hundred pounds) relative to conventional munitions (hundreds to a thousand pounds). This creates the need to design the structure of the new missile so that it contains no unobstructed chamber of diameter and length large enough to accomodate a small fission bomb. We will describe what can be done with the use of active transmission radiography.

In Section 3 we describe the problems and possibilities for verifying limits (or bans) on sea-launched cruise missile (SLCM) deployments, starting with an analysis of how much and what kind of verification is required in view of the nature and limits of the threat they pose. It is shown that effective verification would rely on a set of measures collectively designed to ensure compliance and to assure that militarily significant violations can be detected in time to be countered. No one measure alone can meet these requirements for SLCM verification. The verification regime for two hypothetical treaty limits on SLCMs are considered in detail: a ban on nuclear SLCMs of all ranges and separate sublimits on nuclear and conventional long-range SLCMs.

In Section 4 we address the problem of verifying limits on deployment of mobile ICBMs. A verification protocol relying on partitioning the deployment reservations and allowing for regular queries plus occasional challenges is described. It is designed to complement and make careful use of existing independent means of intelligence, to be minimally intrusive, and to ensure that any substantial violations would be at high risk of exposure. The protocol can be applied exclusively to allowed deployment reservations or nationwide to ensure zero deployments in proscribed regions.

In Section 5 we describe various technologies for tags and seals–including physical tags that must be made tamperproof and attached with reliable seals to individual treaty limited items (TLIs); "proximity tags" or means of registration that need not be sealed but serve only as a license to verify its proximity to a TLI when challenged; and Secure Registration Systems (SRS or "virtual tags") which are basically a set of procedures, such as an encrypted text, to verify TLIs upon remote electronic challenge. Various applications are considered for helping to verify compliance on deployed numbers (i.e.,

nuclear SLCMs) or with designated deployment regions (i.e., mobile ICBMs) without compromising the uncertainties in their actual locations on which they depend for survivability. A specific encryption scheme for an electronic SRS is proposed and discussed in detail.

Section 6 addresses the means of radiation detection that are applicable for determining the presence or absence of nuclear warheads on missiles and for X-raying cruise missile structures to determine whether or not they contain unobstructed chambers that could accommodate small nuclear warheads. Active methods of transmission radiography as well as the passive detection of nuclear radiations are reviewed and their applicability analyzed.

In Section 7 we review a broad range of recently developing and potentially new technologies to see what opportunities there are for making substantive improvements in our means of verification. The ideas considered include the potential of relatively small and inexpensive distributed sensor systems for frequent access and more highly survivable reconnaissance from space; the possibility of large focal length systems assembled in space to give medium resolution imagery from high altitudes; the use of active low-power laser illumination for night and daytime imaging; and the potential for "open skies" as recently proposed by President Bush as a supplement to NTMs for carrying out challenge inspections for mobile missile systems.

2 ADVANCED CONVENTIONAL AIR-LAUNCHED CRUISE MISSILES

2.1 General Remarks

An advanced air-launched cruise missile could provide improved long-range conventional standoff weapons for US strategic aircraft. If developed and deployed it would also raise a serious counting problem for START if its design permitted it to be readily converted from conventional to nuclear warheads.

The desiderata for such advanced conventional cruise missiles (CCM) include low observables; advanced propulsion to ranges in excess of 2000 nmi; improved guidance permitting CEPs of about 3 meters, which would ensure effectiveness of conventional munitions that weigh in the range of 1000 lbs and could be unitary, sub-munitions, or unitary penetrators; and autonomous operation.

There are important reasons to be able to confirm that such advanced air-launched cruise missiles (ALCMs) of the future be unambiguously armed with conventional warheads, and not be rapidly convertible from conventional to nuclear arming. The US should require these characteristics for its force so that neither the missile nor its carrier would come under the limits on nuclear weapons in a future arms control regime. We should also require the same characteristics for any future Soviet advanced conventional cruise missile that would not be accountable under START limits so that we can adequately assess their strategic nuclear strength. The general idea that future long-range conventional- and nuclear-armed cruise missiles should be distinguishable has already been accepted by the US and the Soviet Union.

Arms control issues raised by air-launched CMs differ from those raised by SLCMs; the latter will be discussed separately in the next section. All current ALCMs—the AGM-86B and the advanced cruise missile (ACM), i.e., the AGM-129 with a FY90 IOC—are nuclear and are counted as such in the draft START Treaty. In contrast, the Tomahawk SLCM comes in both conventional and nuclear versions which can be distinguished from one another

only by careful and intrusive OSI. The problem that we address in this section is that of specifying structural features of next generation air-launched CCMs to make them distinguishable from nuclear ones and to inhibit their being converted from conventional to nuclear. Possible airframe designs are discussed which contain no volume large enough to accommodate nuclear warheads of existing design. This is a tricky problem because of the small size of many nuclear warheads relative to 1000 lb conventional munitions. Figure (2-1) illustrates the situation in the current Tomahawk SLCM which can carry a 1000 lb. Bullpup conventional warhead or, to twice the range, the much smaller W-80 nuclear warhead. Indeed the US stockpile contains small nuclear warheads that fit in 155 millimeter and 8 inch diameter artillery shells. There is no principle that prevents the development of warheads (par. ex. of gun-type highly enriched U^{235} design) that can be fit into regions of still smaller diameter. They would be fission weapons of low efficiency and yields (perhaps $\approx$ 100 tons) but still highly destructive if delivered by accurate missiles as envisaged with CEPs of only a few meters. Issues raised by this possibility will be discussed at the end of this section.

2.2 Distinguishability of an Advanced CCM

There are two reasons why any new US air-launched CCM should be designed deliberately to be incompatible with loading on the external pylons with which the B-52H is currently equipped, and also incompatible with the internal common strategic rotary launchers (CSRL) with which they are also being outfitted. First, if or when the B-1B is loaded with the new CCM, its launchers would be distinct from the current ones on B-52Hs so that the B-1B would not be counted as an accountable heavy bomber, or AHB (i.e. one that carries long-range nuclear cruise missiles such as the B52H) under START counting rules. Second, neither would the B-2, a fact that could be ascertained solely by confirming the incompatibility of its CM-launcher with the current ALCM and ACM dimensions, but without requiring unacceptably intrusive inspection of its privileged technology (i.e. stealth).

There is also a complementary aspect to the conventional cruise missile problem: at START counting rules for *nuclear* CMs and their carriers are being elaborated. It is in the interest of the US that these rules permit an

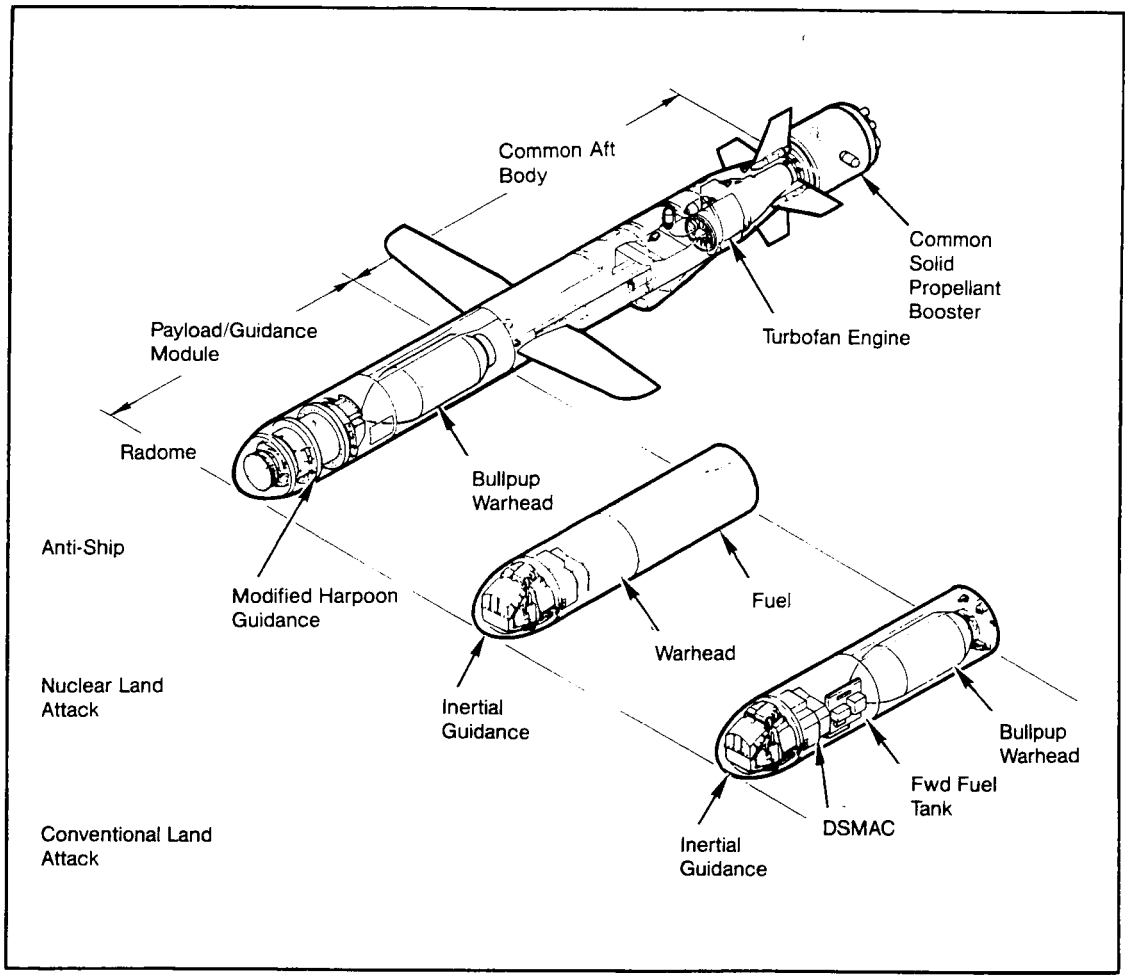

Figure 2-1. Versions of the Tomahawk showing different arming and fuel loading options.

accurate count without requiring unacceptably intrusive OSI for verification. A promising approach currently being discussed by the US is as follows.

The number of AHB's would be observed and counted at ALCM-capable and designated operational bomber bases together with their external pylons and the common strategic rotary launchers capable of loading missiles of 265 inch length. At non-ALCM bases challenge inspections could confirm that CSRLs are collared effectively and irreversibly (or at least not easily or quickly reversible) so as to preclude loading ALCMs and ACMs, and that there are no crew training activities or support facilities for ALCMs. This would protect B-1Bs from inclusion as AHBs in the START count. The same provisions would provide necessary assurances about Soviet activities.

Since all current ALCMs and ACMs (AGM-86B and AGM-129) are counted as nuclear[1] there is no need for intrusive OSI of them and of the radar suppression technology that is incorporated in their design. Protecting stealthy B-2 bombers in the future from intrusive inspections can be accomplished by equipping them with internal rotary launchers that are geometrically incompatible as ALCM/ACM carriers. This specific fact can be revealed by limited OSI of the launchers only without compromising privileged technology (i.e. stealth) of either the airframe or the cruise missiles.

Although it will be of distinguishable shape, new advanced CCMs will presumably have widths[2] no larger than the current cruise missiles—i.e. approximately 20 inches and 29 inches, respectively, for the AGM-86B and the AGM-129. The importance of such a size constraint will be apparent in the following discussion of how to design it so that its structure is incompatible with loading a nuclear warhead.

The continuing availability of telemetry from cruise missile test flights will be important for effectively monitoring future arms control provisions limiting new types of CCMs. An agreement to assure its availability should be included as one of the provisions at START. This would require the extension of the general agreement, already established in the START negotiations, to make such information available from test flights of ballistic missiles.

[1]This is also true of long-range Soviet ALCMs—i.e. the AS-15, their only ALCM of range $> 600\ km$.

[2]They are actually trapezoidal in cross section to allow efficient packing in the CSRL and the external wing pylons.

2.3 Non-convertibility of an Advanced CCM

Assuming that new advanced CCMs can be thus identified we turn to the issue of convertibility.

Obviously convertibility is a matter of degree. A "conversion" in the form of a complete remanufacture of the missile forward of its turbofan engine is not likely to be inhibitable by design specifications, since such a conversion amounts to throwing out the agreed upon treaty-compliant design and substituting a treaty-noncompliant one. Such a situation is essentially the same as covertly manufacturing a new, treaty-forbidden missile; verification safeguards deemed adequate for the latter case should apply to conversion-by-remanufacture also. No strategic planner would be willing to rely on a missile which had not been extensively tested in precisely its war configuration. Consequently a treaty regime which permits conventional cruise missiles would have to contain appropriate provisions for verifying cruise missile flight testing. This would effectively rule out the clandestine conversion of conventional to nuclear capability by means of a complete remanufacture of the missile.

Short of remanufacture, the kind of conversion that we want to inhibit is that which can be effected rapidly, in time of crisis, on deployed supposedly-conventional cruise missiles. We surely want it to be impossible to simply (screwdriver level) swap out a conventional warhead and insert a nuclear one. However, we also want to inhibit conversion in the shop facilities at a forward air base.

In their present form, cruise missiles, although not designed to be, are in fact eminently convertible. Their basic mechanical structure is a rigid, hollow tube of sand-cast aluminum with wings. Aft, the tube surrounds the turbofan engine. Engine thrust (in level flight equaling aerodynamic drag) is applied to the tube's circumference through the engine mount, around the circumference of the engine. The interior of the tube is essentially free of load-bearing components, allowing it to be filled with fuel cells, conventional warheads, or nuclear warheads, constrained only by weight and center-of-gravity (CG) considerations.

The wings of the ALCMs are attached at the bottom of the trapezoidal surface and are retracted prior to launch. In contrast the Tomahawk SLCMs

have a cylindrical design with the wings attached at the midpoint; they retract into the tube at its center. (See Figure 2-2.)

The diameter of the tube in current US designs is not large. The ALCM is constrained by the dimensions of the B52 rotary launcher to roughly 20 inches. A tube of this dimension accommodates a thermonuclear warhead of conventional design. However one could design a new CCM in such a way that its maximum unobstructed dimension is but a small fraction of its external diameter. For example the outer envelope of a new CCM might have the same size as the current designs, but the airframe structure could be changed from a tube to one in which load-bearing structures are distributed throughout the volume in such a way as to limit the maximum unobstructed dimension to a small fraction of the overall diameter. Since fuel and conventional munitions can be broken down into subunits, at essentially no detriment to their utility, such an unconventional airframe design would function perfectly well as a conventional cruise missile but not accommodate even a small, low-yield existing fission bomb.

There are several possible overall strategies for designing such an airframe, ranging from a single strong central structural member (tentpole) to a collection of weaker members distributed through the volume and deriving strength from their interconnections. While such concepts are not conventional or natural in airframe design, they should not significantly impact the aerodynamic properties of the missile. We will say a few words about the tentpole design first just to show how the basic requirements of structural integrity are met in an airframe.

2.3.1 Tentpole Design

In the standard cruise missile design, compression is applied to the tube (and its associated stiffening and load-bearing structure) by the engine thrust. There is also a bending moment due to the lift which is localized at the wings, near to the CG: the weight of the forward (warhead) and aft (engine) components is supported by the upper surface of the tube being under extension and the lower surface being under compression. Both because of the bending moment, and because of the thrust vector, it is a mathematical truism that a portion of the structure must be under longitudinal compression. The tentpole design of a new CCM would take two different forms depending

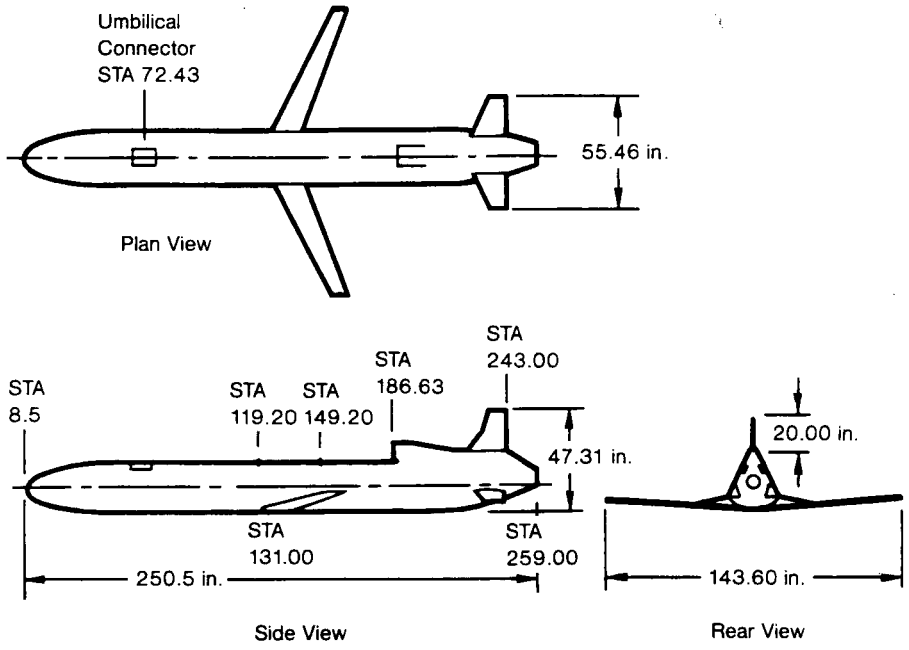

AGM-86B General Arrangement
With Folding Wing at the Bottom

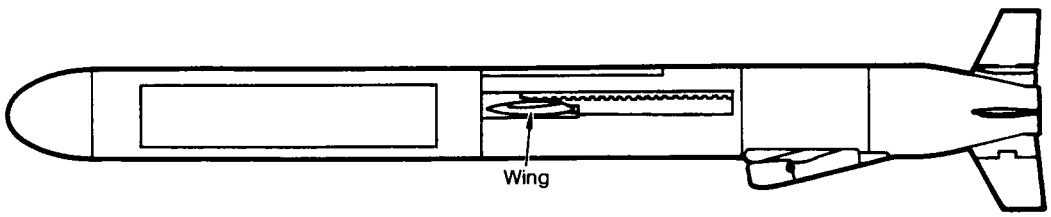

Tomahawk General Arrangement
With Folding Wing at Midpoint

Figure 2-2. Tomahawk general arrangements with folding wing at midpoint.

on the missile's wing configuration. If, as in the current ALCM design, the wings were to remain hinged at the bottom surface the tentpole CCM would have a longitudinal compression-bearing structure, forward of the engine, substantially only along the center line. The rear engine would apply thrust to a central rigid "tentpole" running the length of the missile forward. The tentpole would be supported against bending moments by (at the designer's option) a mixture of transverse disks and diagonal strut wires (under tension, not compression). Figure 2-3 sketches the idea. The role of transverse disks would also be to supply load-bearing connections for moving, handling, and loading the CMs on the pylons or rotary launchers. The tentpole would also provide the load-bearing joint for the wings in flight.

Alternatively, if the wings are retracted into the center of the missile tube as in the current Tomahawk SLCM configuration, the external missile tube—i.e. the surface—would remain the compression-bearing structure for the aft section of the missile into which the wings could retract during stowage. Forward of the wing compartment the tube would join a load-bearing transverse plate through which the load would be transferred to a central rigid "tentpole" running forward the remaining length of the missile. The outer skin of the missile in the forward section would not be load-bearing.

The net result is that the only nuclear weapons which could readily be carried by the cruise missile (i.e. carried without major structural modification) would have a diameter less than half the cruise missile diameter, and consequently a dramatically reduced yield. One could reduce the free dimension even more by making the load-bearing structure out of an interconnected latticework of individually rather weak rods (on the model of the cross-braced lattice of a tall antenna mast).

2.3.2 Inner Shell Design

A further reduction in size could be achieved if an inner compression-bearing cylindrical shell were positioned around a center tentpole rather than relying on a strong center tentpole alone. As illustrated in Figure 2-4 this structure restricts the cavity for loading a bomb to a still smaller transverse dimension. It also provides a strong structure to handle the transverse bending moments, in addition to transverse disks reaching to the surface to provide moving, handling, and loading joints. The volume between the inner

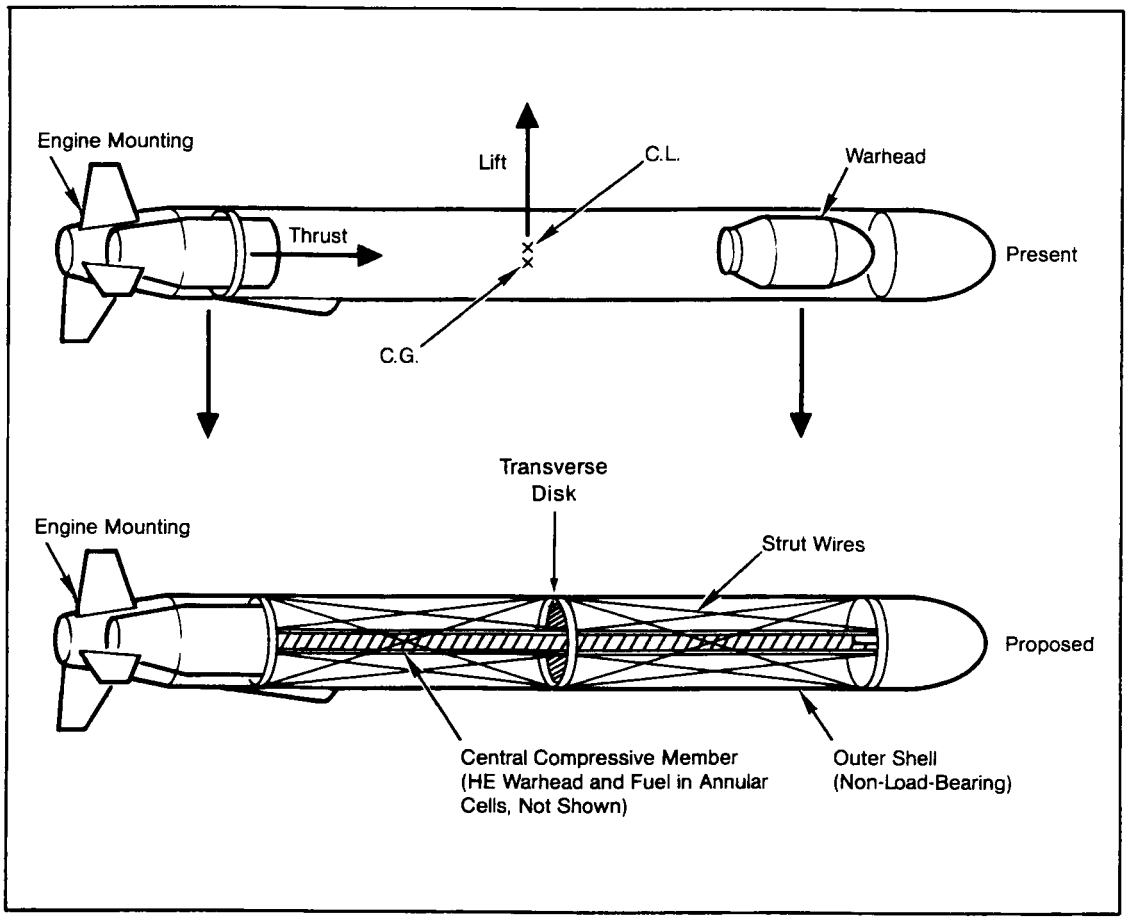

Figure 2-3. Tentpole design for a cruise missile.

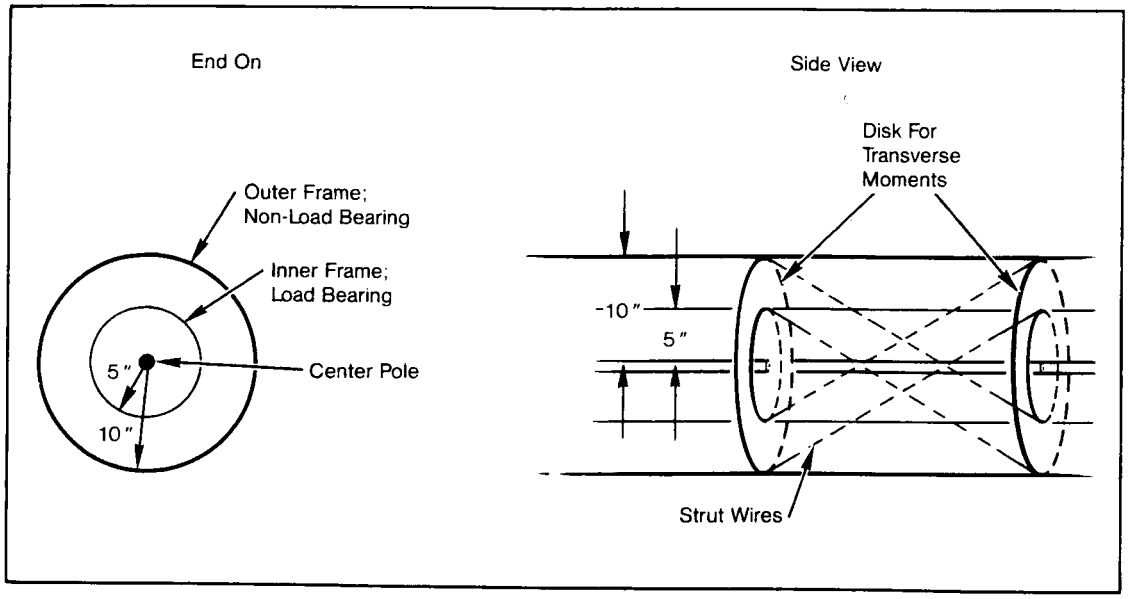

Figure 2-4. Inner shell design for cruise missile.

structural cylinder and the outer surface of the missile would be available for fuel, cables, and high explosives. Since the advanced CCM should be designed to ranges $\geq$ 2000 miles the overall volume of fuel, even allowing for significant improvements in propulsion over current ALCM design, will necessarily be at least as large as at present. Therefore such a missile will necessarily be larger than the current ALCM or be limited to a volume for high explosives no larger than currently available for the nuclear warhead (i.e. about 200 lbs.). Such an explosion will of course have a devastating effect on most structures except for highly hardened military command centers or underground ICBM silos if impact is achieved within the projected CEP of 3 meters. (See Section 2.4.)

The fuel could easily flow around the inner structure in either missile design. Also the explosive could be shaped[3] as desired and held in place, like the fuel, by the outer skin of the missile which, while not load-bearing, would be formed of composite material to reduce the radar cross-section.

The next issue is to demonstrate that a missile of the above design can be constructed in a way that would make it quite difficult to convert it to one with a large enough unobstructed volume to accommodate an existing design nuclear warhead. One possible approach to such a conversion would be to saw a section out of the missile structure and replace it by a warhead inside a load-bearing tube with forward and aft load-bearing mounts. One would have to mate this tube both to the "tentpoles", and also to to the strut wires which previously traversed its volume. The missile designs described above would insure that this structural change would have to involve actual cutting and welding of the missile's principal structural members. Another approach would be to "glue" a nuclear weapon compartment to the front end of the missile. To maintain the CG over the wings, one would also have to attach a (much heavier) counterweight to the rear of the missile. The added weight and drag associated with such a modification would almost certainly greatly diminish the range and deployability of the missile. This possibility has to be looked at in more detail to see if it is at all practical. In any event, the modified missile amounts to a completely new class of vehicle in which no military planner would place his trust in the absence of an extensive program of flight testing.

[3]For solid explosives this can be done readily with little loss of explosive yield. Liquid explosives (such as, for example, methyl nitrate) are too volatile to consider in this context.

2.3.3 Honeycomb Design

Another possibility for the design of a future cruise missile that is unambiguously not convertible from conventional to nuclear would be to make a two-dimensional honeycomb of chemically strengthened glass, running the length of the forward portion of the cruise missile. The honeycomb might be a hexagonal tube about 4 inches across, pretty much filling the volume of the cruise missile. A specific implementation of this proposal is detailed in Appendix A to this section.

2.3.4 Verification Issues

To verify that the airframe is of the agreed-on kind, a deployed conventional force of these next-generation cruise missiles would be subject to a regime of inspection, either at their place of manufacture or on station. By transmission radiography, or x-ray, one would verify that (i) the outer aerodynamic shell of the forward section is not capable of bearing the load implied by a nuclear warhead, (ii) the central tentpole or inner load-bearing cylinder is a unitary structure, without any fittings allowing for a portion of it to be removed, and (iii) any strut wires likewise are not interrupted by fittings which would allow them to be easily shortened. A detailed discussion of radiography showing how to verify the internal structure by means of a Co^{60} source of less than one curie is given in Section 6.

Passive means of detecting nuclear radiation are not of much help in this problem. It tells nothing about internal structure or convertibility. In addition a fission bomb using highly enriched uranium as the fuel and tungsten, rather than depleted ^{238}U, as the tamper gives only low energy gammas that are easily absorbed.

An important point is that we would certainly insist on protecting some aspects of cruise missile technology (in particular, stealth-related items); this will complicate the inspection problem. However, radiographing the missile structure need not reveal any essentials of treatment of the missile surface for stealth, and indeed the missile could be shrouded in light material at all times during inspection. At the same time, we have to make sure that the only missiles which get extensive flight testing are of the agreed-on type, in

order to deny confidence in any possible illicit field modifications. This can best be handled by insuring availability of test flight telemetry.

2.4 How Small a Nuclear Warhead?

This brings us finally to the question of how small a nuclear warhead might be squeezed into such an advanced CCM design. Nuclear artillery shells exist with 8 inch and 155 mm = 6.1 inch diameters.

One can ask, in particular, what strategic value there would be if, for example, a 1 kT nuclear warhead could be squeezed into a new advanced CM that could deliver munitions with very high accuracy, such as the 3 meter CEP envisioned. One obvious question that arises is whether accurately delivered kiloton weapons would have strategic (i.e. silo-killing) value. The answer appears to be yes.

First consider mechanical kill by overpressure on the silo doors. The standard wisdom is that silos are hard to a few $\times 10^3$ psi. To be conservative we will assume that the silo is killed if the maximum overpressure exceeds 10^4 psi. The curves in "The Effects of Nuclear Weapons" by Glasstone and Dolan, Section 3, for the maximum overpressure P_{max} (on the ground) as a function of yield (Y) and distance (R) from a ground burst nicely fit, for small distances, the natural scaling law for the maximum overpressure at small miss distances

$$P_{max} \propto \frac{Y}{\frac{4\pi}{3}R^3} \,. \qquad (2-1)$$

By curve fitting, one finds (see Appendix B)

$$P_{max} = 10^4 \frac{Y(kT)}{\left(\frac{R(ft)}{70}\right)^3} \text{ psi.} \qquad (2-2)$$

Choosing the sure-kill criterion to be $P_{max} = 10^4$ psi, one has

$$R_{kill}(ft) = 70Y(kT)^{\frac{1}{3}} \qquad (2-3)$$

or $R_{kill} = 70$ ft for 1 kT or 40 ft for 0.2 kT. These aiming accuracies can be obtained currently, in the absence of defensive counter measures, by TERCOM augmented by DSMAC. A much higher accuracy of ~ 3 meters is the

goal for a future advanced CCM. In the Appendix B we discuss the scaling law (2–1) and the lethal effect of conventional warheads at these small CEPs. From conventional bomb data we show that 1000 pounds of conventional high explosives can be expected to create an on-target overpressure of approximately 500 psi if delivered with a CEP of roughly 3m, or 10 feet. Hardened strategic targets such as underground silos or command posts are rated at several thousand psi and would probably survive such an attack, although any other target would be severely damaged. On the other hand, for a small nuclear yield of 1 kT, Equations (2-2) and (2-3) show that even the most hardened targets would be destroyed; for R = 10 ft and Y = 1 kT, $P_{max} \approx 3.4 \times 10^6$ psi.

Finally, we consider, as an additional kill mechanism by nuclear warheads, melting of the nuclear warhead mounted on the missile in the silo by the neutrons released by the explosion of the attacking nuclear warhead. A 1 kT warhead releases 4×10^{12} joules. At 200 MeV per fission and 1.5 neutrons per fission, the explosion releases 2×10^{23} neutrons, which corresponds to a fluence of 1.6×10^{16} neutrons-cm^{-2} at a distance of 10 m. Making the approximately accurate assumption that the cross-section for Pu fission is 10^{-24} cm^2 (i.e., geometrical), the probability that any given nucleus in the target warhead undergoes fission is Prob $= (1.6 \times 10^{16})(10^{-24}) = 1.6 \times 10^{-8}$. Since each fission releases 200 MeV, the fission heating of the warhead material is roughly

$$\Delta E = (200\ MeV)\left(1.6 \times 10^{-8}\right) = 3.2\ eV \text{ per atom} . \qquad (2-4)$$

This gives a temperature rise of

$$\Delta T = \frac{3.2 \times 1.6 \times 10^{-12}}{3 \times 1.4 \times 10^{-16}} = 1.22 \times 10^4 K , \qquad (2-5)$$

where we have assumed a specific heat of 3 × (Boltzmann constant) for the atoms. A temperature rise of $10^3 K$ (actually more like $500K$) is enough to melt the metal of the warhead, so this seems to be a serious threat. Actually the silo gives a substantial shielding factor (easily a factor 10), so this melting mechanism will not enhance the kill radius substantially beyond that due to blast overpressure.

2.5 Summary

In order to meet the requirement that a new advanced cruise missile be unambiguously conventionally armed and not readily convertible to carrying a nuclear warhead, its structure must be designed so that it contains no unobstructed chamber of diameter comparable to the missile diameter. A nuclear warhead with a yield $\sim$ 1 kT, and delivered with a CEP of $\sim$ 10 ft, would create a maximum overpressure that can destroy the hardest targets. With active transmission radiography one can verify that the cruise missile airframe is constructed appropriately so that it cannot accommodate an existing fission weapon, and also cannot be readily converted to accept one. Finally we note that there is no reason, in principle, that one could not develop smaller or segmented fission bombs—particularly a simple gun type bomb using highly enriched uranium—if low yield underground nuclear testing continues.

APPENDIX A: HONEYCOMB DESIGN

First, the structure could be *cast* just as if it were metal, and in this form could have large passages molded through the glass, even more readily than if it were aluminum. It could also be machined with diamond tools very readily, in case one needed holes for a model change, and the like. After the structure is complete, it could be *chemically* strengthened, in a similar fashion to the strengthening and tempering of the side windows of automobiles. In the latter process, the glass is brought to the annealing range, and the surface suddenly cooled with jets of air. The surface becomes rigid, and it shrinks somewhat because of the thermal expansion of the glass with increasing temperature. Because the interior of the glass is still soft, when the surface shrinks, it does so without significant stress, and as the interior of the glass hardens, it does so without stressing the surface. But after the interior becomes rigid as well, further contraction of the interior leads to its putting itself into tension in the plane of the sheet, while the surface is in compression. Because the tension portion of the glass has no free surface, it has no great tendency to break, and the surface is now proof against the propagation of scratches, chips, and the like that normally weaken glass far below its theoretical strength. However, if the surface is cut or chipped to a depth greater than the thickness of the compression layer, so that the mar enters the tension layer, a crack instantly propagates and pulverizes the glass. This is what happens to side windows of cars when they are smashed on the streets.

The chemically strengthened glass is similar, but more convenient in a way because the surface is expanded by chemical treatment that replaces small ions with bigger ions. In this way, one could provide a strong structure that could not be modified in any way without totally dismantling the missile and treating it for a period of perhaps many months. It would be easier to build a new front end, and that is our purpose. Furthermore, unless the structure is strengthened, it would not be strong enough to serve as a structural material. Thus, one need not inspect the surface of the glass in detail, but could simply use tomography with the output limited to a reporting of this honeycomb skeleton, in order to confirm that there are no places within the missile large enough to contain nuclear warheads.

APPENDIX B: CALCULATION OF OVERPRESSURES

We have used the scaling law $P_{max} \propto YR^{-3}$ to estimate the maximum overpressure produced by low yield bombs at small miss distances. More generally, one can use simple dimensional arguments to show that at large (small) distances the overpressure scales as R^{-1} (R^{-3}) and to estimate the distance at which the transition between the two scaling laws occurs. The argument goes as follows.

Consider a blast wave produced in a fluid of ambient pressure p_0 by an explosion of total energy Y. At any fixed distance R from the blast the pressure signal will be a pulse that rises from zero to some maximum and then falls back to zero in some time interval. Typically we are only interested in the maximum overpressure and its dependence on distance rather than in the detailed shape of the pulse. On dimensional grounds alone, the dependence of P_{max} on distance must have the form

$$\frac{P_{max}}{p_0} = F\left(\frac{Y}{p_0 R^3}\right) \qquad (2-1)$$

where $F(x)$ is a dimensionless function of a dimensionless argument whose form has to be determined by physical considerations.

There are two limits in which the behavior of $F(x)$ is known. For small R the overpressure is large and the dependence on the ambient pressure should drop out of the expression for P_{max}. This requires that

$$F(x) \to x \qquad \text{as} \qquad x \to \infty \text{ , or } R \to 0.$$

For large R, the overpressure becomes small compared to the ambient pressure and the blast pulse should behave like an acoustic wave of small amplitude. For energy conservation reasons, the amplitude of a spherical acoustic wave decays with distance from its source as R^{-1}. The corresponding constraint on $F(x)$ is

$$F(x) \to x^{\frac{1}{3}} \qquad \text{as} \qquad x \to 0 \text{ , or } R \to \infty.$$

For intermediate distances $F(x)$ will interpolate between the two limiting power laws in some more or less complicated fashion that depends on the equation of state of the medium.[4]

The distance of crossover between the asymptotic regimes corresponds roughly to $x \sim 1$ or

$$R_{cross} \sim \left(\frac{Y}{p_0}\right)^{\frac{1}{3}} .$$

Needless to say, if we are interested in doing damage to a target, we want to be in the short distance regime where overpressure scales as $\frac{Y}{R^3}$! Some sense of how well the scaling laws work can be gotten from the two figures included here. Figure B–1 is a summary by Ted Postol (private communciation) of results from experiments on conventional bombs in the 500 to 2000 lb regime and shows an attempt to fit the results with a simple power law form for $F(x)$. The fit is moderately good and shows that, for these yields, distances between 10 and 100 feet lie in the transition regime between the two asymptotic scaling laws. The short distance, R^{-3} scaling regime is probably not accessible. Figure B–2 is a summary of the results tabulated in Glasstone and Dolan on the ground burst of a 1 kT weapon. One sees that accurate R^{-3} scaling is achieved only for $\frac{R}{R_{cross}}$ smaller than about .3 (i.e. for x larger than about $(1/0.3)^3 \sim 30$). Since R_{cross} is large (about 350 m for 1 kT), the inner scaling region is of interesting size.

The conventional bomb data in Figure B–1 show why a "zero CEP" conventional cruise missile is of some interest. Ten feet is probably a practical minimum CEP and 1000 lbs is probably a reasonable maximum HE payload for a long-range cruise missile. The on-target overpressure is then on the order of 500 psi. Hardened silos are rated at several thousand psi and would presumably survive such an attack, but any other target would be severely damaged.

[4] An approximate formula valid over a broad range of peak overpressures is given by Harold L. Brode in Annual Reviews of Nuclear Science (Vol 18, p. 153 (1968)): $P_{max} psi = 3300 \left(\frac{Y_{MT}}{R^3_{kft}}\right) + 192 \left(\frac{Y_{MT}}{R^3_{kft}}\right)^{1/2}$.

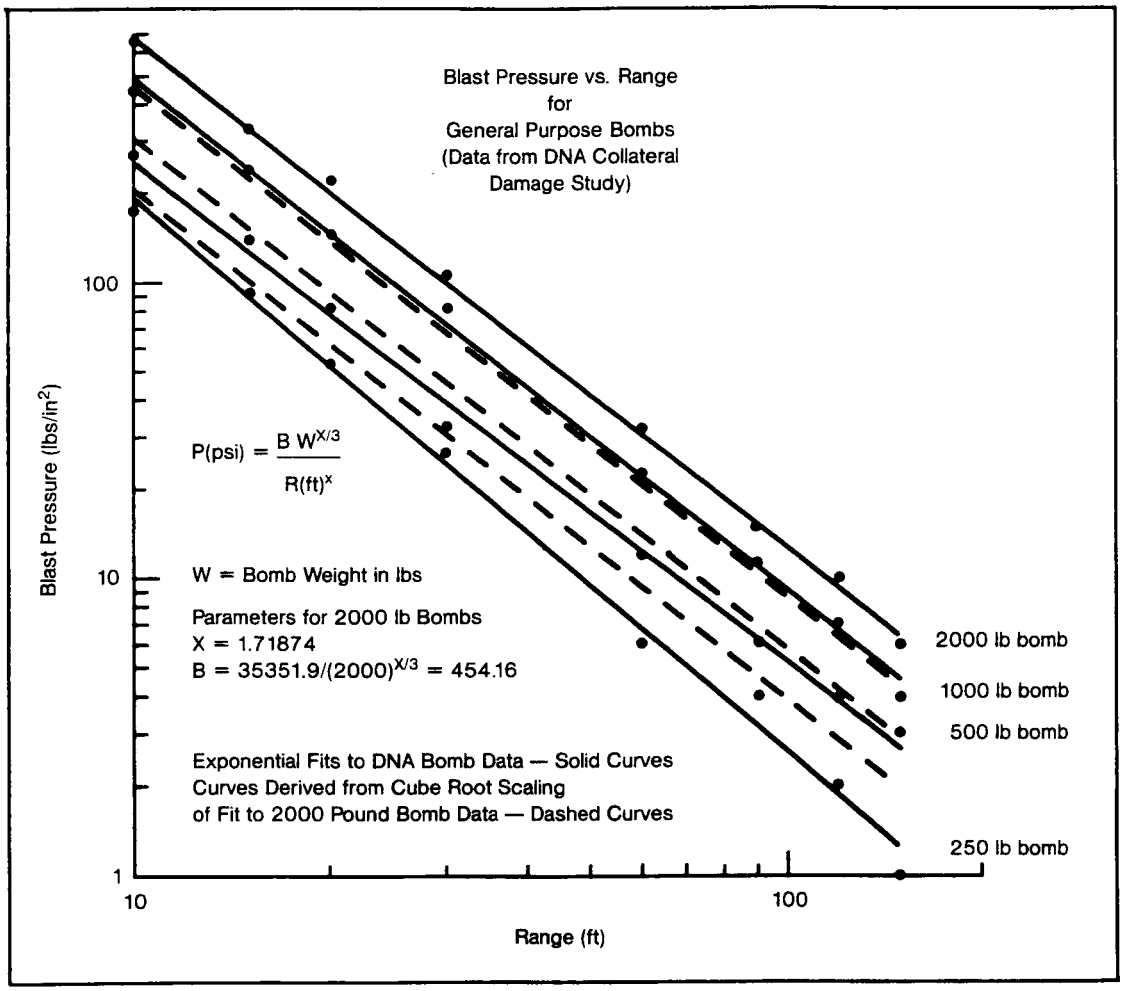

Figure B-1. From T. Postol. The weight of high explosive is taken to be one-half of that given for the bomb itself.

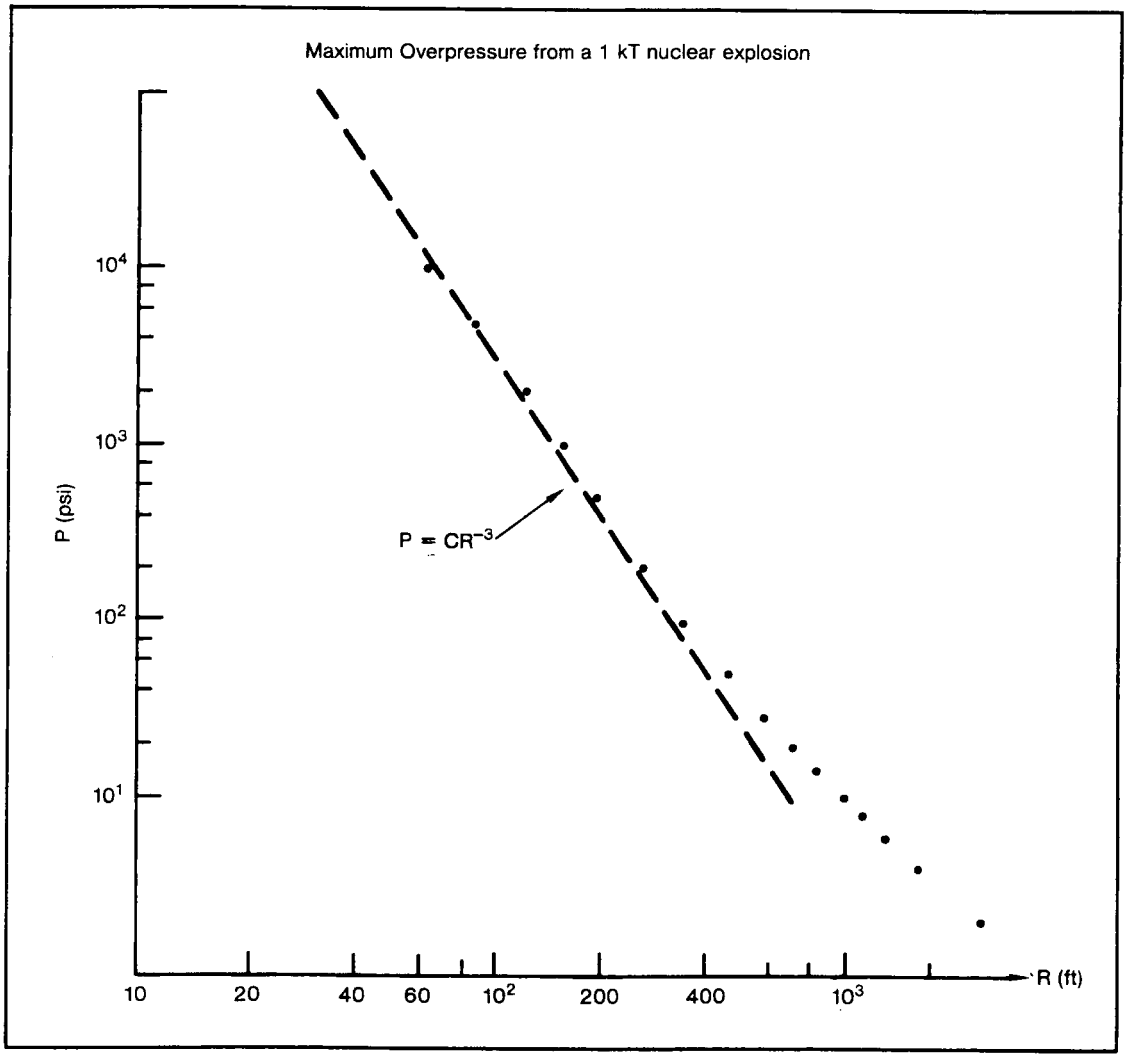

Figure B-2. The solid curve shows the $1/R_3$ dependence in Equation 2-1.

3 SEA-LAUNCHED CRUISE MISSILE VERIFICATION

3.1 General Remarks

Three characteristics of sea-launched cruise missiles (SLCMs), taken together, are responsible for the unique challenge they present to efforts to negotiate verifiable limits on their deployment:

1. Quantitative limits cannot be verified accurately because SLCMs are small, can be deployed widely, and are dual-purpose. Only by means of inspection at close range can nuclear-armed cruise missiles be distinguished from conventionally armed ones.

2. Once one accepts the possibility that at least several dozen long-range SLCMs may have been deployed with nuclear warheads, the potential of a leading edge attack against national command centers and/or bomber bases has to be addressed. This is based on the judgment that a small number of SLCMs might escape detection for all or a substantial portion of their flight. The US has at present no system to provide confident early warning of the launch of a cruise missile comparable to the launch warning of a ballistic missile attack that is provided by the DSP satellite from geosynchronous orbit. Technology does exist, however, to detect with confidence and to provide some tracking information on large numbers of such targets in flight, against ground clutter.[5]

3. Once the number of deployed long-range nuclear armed SLCMs exceeds, say, fifty or so, there is no crucial national security need to be able to count their numbers precisely. This is because cruise missiles are slow, and have very little, if any, potential for prompt large-scale, counterforce targeting.

[5]See for example the Report of the Ad Hoc Committee on Airships of the US Air Force Scientific Advisory Board, September 1987, (Dr. William H. Heiser, Chairman) and the JASON Report JSR-88-230 "Airships" S. Drell, J. Katz, and G. MacDonald and JSR-87-801, the 1987 JASON study on OTHB radars by G. MacDonald et al. These technologies should also work against low radar (stealthy) cruise missiles over water.

> The options for reliable counter-military targeting by cruise missiles are limited by technical characteristics of the missiles: their low velocity and their softness as targets for various kinds of interceptors. Massive, highly coordinated strikes against large, geographically diverse targets, such as ICBM fields, would be impractical because of the potentially relatively long warning times the missiles would afford were any of them detected. Moreover, the likelihood is high (from the attacker's conservative point-of-view) that more than a few of the large number of cruise missiles in such an attack would be detected relatively early. Once nuclear weapons started exploding in a vicinity, cruise missiles arriving later might not survive; because they are not hardened against various nuclear weapon effects nearly as well as are ICBM RVs, attack coordination presents an even more stressing problem than it does for ICBM attacks. Finally, the low velocity of cruise missiles precludes their use against targets that can reasonably be expected to move, for example, military units, mobile ICBMs, or command posts.

These characteristics determine what is needed for effective verification of SLCM deployments. In most general terms the challenge can be formulated as follows. There are three possible situations to consider:

- Case A is a treaty limiting, or banning, nuclear SLCMs that is faithfully obeyed by both the US and the Soviet Union.
- Case B is a treaty violated unilaterally by the Soviet Union.
- Case C is no treaty at all.

Other situations might also be considered, for example a declarative limit on numbers of nuclear SLCMs, agreed to by both countries, with no formal treaty and no formal verification process. There may be important benefits to be gained from declarative limits, but these are not relevant to the present discussion. So far as verification is concerned, the situation with declarative limits is equivalent to Case C.

Under the assumption of Case A a treaty banning all nuclear SLCMs would be of substantial value to the US. In particular it would remove the

possibility of a sneak attack against US coastal cities and military installations including the NCA in Washington DC; it would also remove the possibility of a sneak nuclear attack launched from Soviet submarines against surface ships of the US Navy. The US would still retain nonnuclear antiship SLCMs and the capability for long-range overland penetration by nuclear ALCMs. Therefore, under the assumption of Case A such a treaty would be highly preferable to Case C.

There is also a large geographical asymmetry between the US and the Soviet Union that favors Case A. In contrast to the US, the Soviet Union has no abundance of high value targets on land near to accessible sea coasts and within SLCM range unless the launchers are close offshore, nor does it have a surface navy of comparable importance to ours. A treaty banning nuclear SLCMs would negate this geographical advantage of the Soviets under the assumption of Case A. Furthermore such geographical advantages are permanent, whereas technological advantages in weapons such as now enjoyed by US Tomahawk SLCMs with their superior accuracy and platforms may only be temporary.

Next we consider Case B, in which unilateral violations are assumed to occur. It is pointless now to compare Case B with Case A. Obviously Case A is preferable, but the choice between A and B is made by the USSR and not by the US. The choice which is in the hands of the US is the choice between Case B, a treaty with violations, and Case C, the situation that exists if we have no treaty at all. We have to estimate whether a violated treaty is more dangerous than none.

Two types of consequences of unilateral violations must be considered.

- Type 1. An actual sneak attack on US territory or US ships using concealed nuclear SLCMs banned by treaty.
- Type 2. Political effects of a unilateral deployment of SLCMs in peacetime, and their incremental contribution to the possibility of blackmail based on the threat of nuclear attack.

There is no question that a sneak SLCM attack is a potential danger to US security and that in Case B such an attack is technically possible. However, the gravity of the danger is no greater in Case B than in Case C. Either in

Case B or in Case C the deterrence of SLCM attack rests on the totality of US strategic forces, not on the US SLCM force. The contribution of the US SLCM force, absent in Case B and present in Case C, will not materially affect the deterrent power of the remaining strategic forces of which it is only a small percentage (currently 370 warheads relative to more than 8000 under the currently envisioned START ceilings). The difference between Case B and Case C, so far as Type 1 consequences are concerned, is slight.

When we consider Type 2 consequences, based on political exploitation of a unilateral deployment of SLCMs in peacetime, we must assume that the violation of the treaty has become overt. A secret violation can only be exploited in a surprise attack, not in political pressure or blackmail. But, as soon as the violation becomes overt, the treaty is dead and Case B becomes almost identical with Case C. After an overt violation of a treaty, the US would be even more unlikely to submit to nuclear blackmail than in the absence of a treaty.

The conclusion of this analysis of Case B is that a unilaterally violated treaty would bring no military danger to the US that is not already present in the absence of a treaty. So far as military risks are concerned, Case B is hardly worse than Case C. The costs of a violated treaty would be largely intangible, following from a general breakdown of international stability rather than from any specific threat to US security.

To summarize the preceding argument, we have considered three possible alternatives. Case A is a treaty without violation, Case B a treaty with violations, and Case C, without any treaty. We reached two conclusions. Case A, an unviolated treaty banning nuclear SLCMs is substantially better for US security than Case C. In Case B, even with serious violations, the danger to US security is only slightly worse than in Case C. Putting all three cases together, we may say that X =(Advantage of A compared with C) is large compared with Y =(Disadvantage of B compared with C).

The purpose of verification is to reduce so far as possible the probability of violations. We may measure the effectiveness of verification by the probability P that a unilateral violation will occur. After a treaty is signed and ratified, the probability of Case B will be P and the probability of Case A will be $(1 - P)$. The net expectation of advantage derived from the treaty

can then be written symbolically as

$$E = (1 - P)X - PY \ . \qquad (3-1)$$

Since X is large compared with Y, the requirement of positive expectation from the treaty does not require that P be small. In fact, the treaty will be advantageous, better than no treaty at all, provided that

$$P < \frac{X}{X+Y} \ . \qquad (3-2)$$

The content of Equation (3–2) is illustrated in Figure 3–1. The shaded area corresponds to a positive expectation of advantage being derived from the treaty. The figure shows that even if the probability P of a unilateral violation is high, the treaty will be advantageous when Y/X is small. Another pictorial representation of this general result is shown in Figure 2–2.

This makes clear that a verification system can be useful, even if it provides only a slight inducement to honesty ($P < 1$), if it is part of a good treaty which offers greater benefits than risks ($X >> Y$).

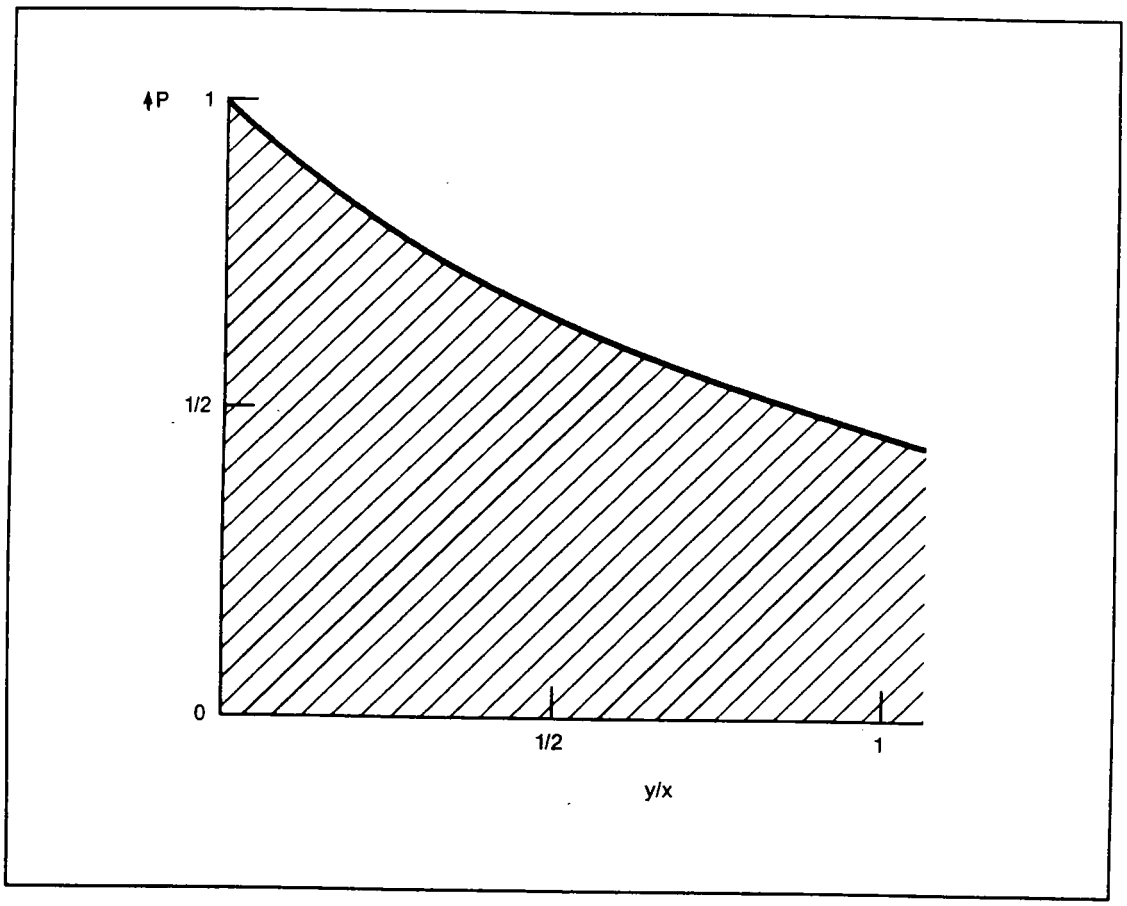

Figure 3-1. The shaded area corresponds to Eq. (3.2) for the positive expectation of advantage being derived from a treaty.

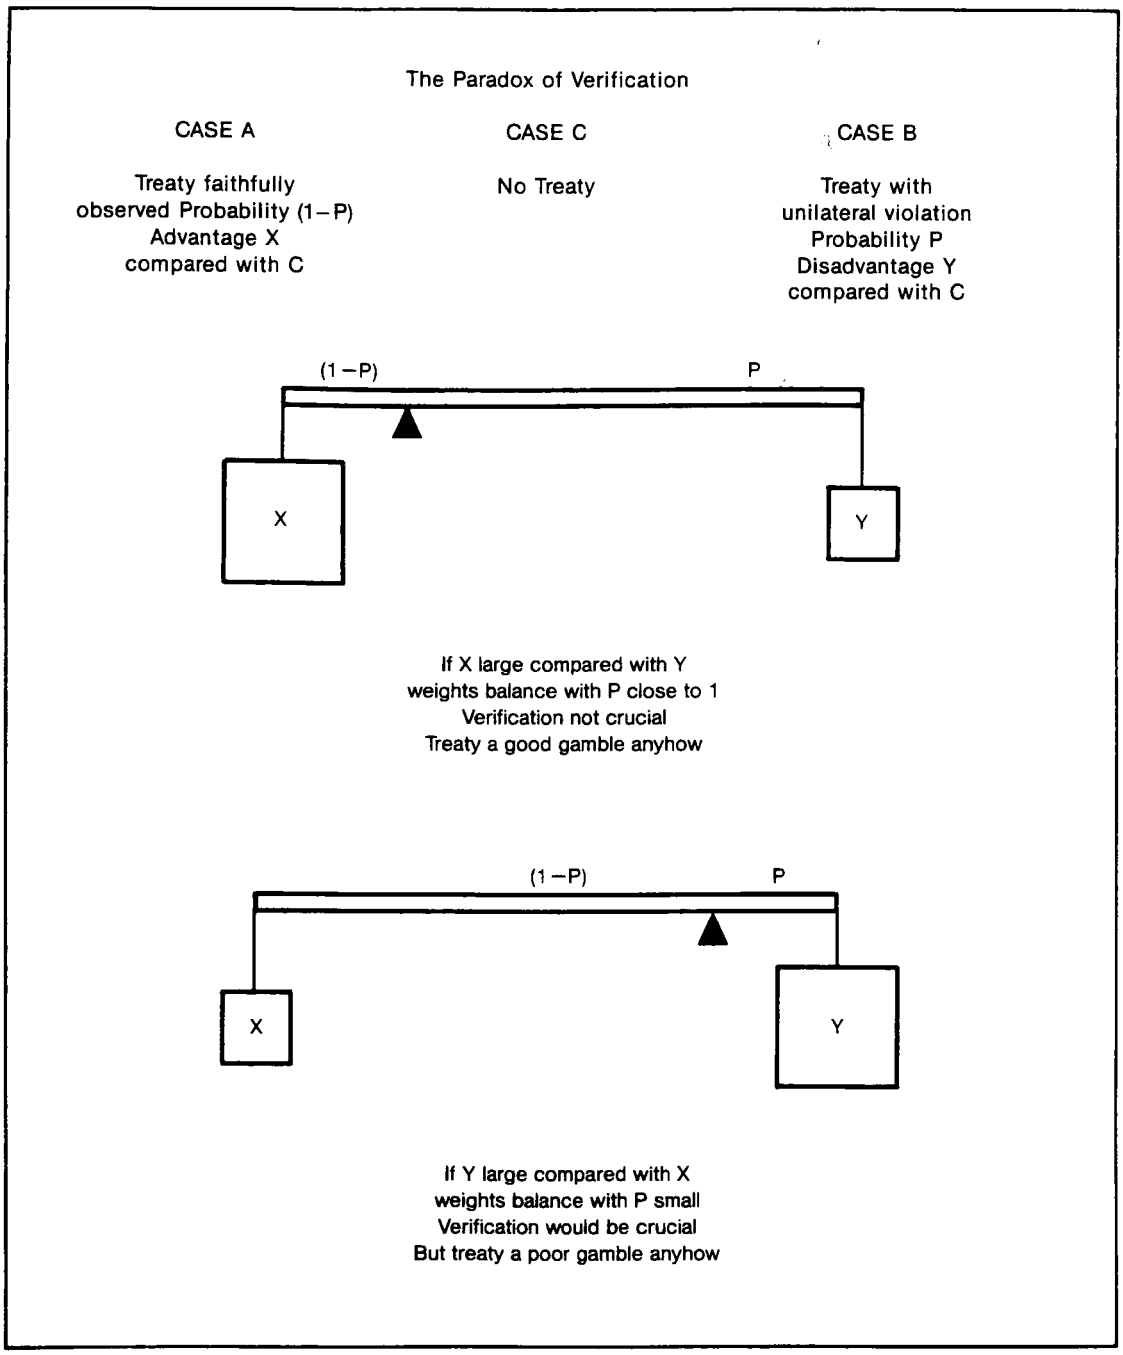

Figure 3-2. Moral: Strict verification is needed only for bad treaties. A good treaty can be useful even if verification is highly imperfect.

3.2 Characteristics of SLCMs

To understand the problems associated with verification of limits on SLCMs, we first discuss the characteristics of US and Soviet SLCMs and the numbers deployed by the two countries. There are several types of SLCMs, designed for ship attack or land attack, carrying nuclear or conventional warheads.

In 1983 the US began deployment of a new SLCM, the Tomahawk. The Tomahawk is a small, unpiloted jet aircraft, which flies subsonically and is capable of highly accurate delivery of nuclear or conventional warheads. The Tomahawk airframe is the basis for several SLCM variants. These variants have essentially identical airframes, but there are internal differences to accommodate the different warheads and different missions. The short-range, antiship variant, which carries only conventional warheads, has an operational range of approximately 450 km, and uses radar to home on its target. There are three long-range, land-attack variants: one carries a nuclear warhead, has a range of over 2500 km, and uses terrain contour matching to update its inertial guidance in order to achieve its accuracy; the other two carry conventional warheads (either submunitions or a unitary warhead), have about half the range (because the conventional warheads leave less room for fuel), and have additional guidance based on digital scene matching in the terminal phase to achieve the precise accuracy required for conventional munitions. All Tomahawk variants are deployed in canisters and can be launched in a variety of ways, and from a variety of platforms. They can be launched from the torpedo tubes or vertical launch systems of submarines, or the vertical launch systems or armored box launchers of appropriately outfitted surface ships. There are currently about 70 US surface ships and submarines capable of firing Tomahawks, and this number is planned to increase to nearly 200. The US has procured approximately 370 nuclear and 1650 conventional Tomahawks to date, and plans to purchase about 2000 more conventional Tomahawks between now and 1994. Procurement of additional nuclear Tomahawks is not planned until 1994. The only other SLCM deployed by the United States is the Harpoon, a conventionally armed, antiship weapon with a range of about 100 km.

The Soviet Union has deployed SLCMs since the early 1960s. Over the years the Soviets developed several short-range models, designed primarily

for ship attack. Their ranges vary from approximately 50 to 550 km, and most are capable of carrying either conventional or nuclear warheads. Some are rocket powered, some have air-breathing engines; some are supersonic, some are subsonic. All are larger than the Tomahawk, and are launched from surface ship or submarine launchers, not from torpedo tubes. It is estimated, based on a count of launchers, that the Soviet Union currently has approximately 1000 dual capable, short-range SLCMs deployed on a wide variety of surface ships and submarines. The number of short-range SLCMs which carry nuclear warheads is unknown, but it is estimated[6] that roughly 1/3 to 1/2 of these short-range Soviet SLCMs may be nuclear-armed.

In 1986 the Soviets began deploying their first long-range, land-attack SLCM, the SS-N-21. This missile appears to be quite similar to the Tomahawk. It has a small jet engine, flies subsonically and at low altitudes, and is small enough to be launched from torpedo tubes. At present there is only a nuclear-armed version of the SS-N-21, presumably because the guidance does not include either digital scene matching or a real-time satellite update which could give it the accuracy desired for long-distance delivery of conventional warheads. The Soviets have also begun testing a considerably larger, supersonic, long-range SLCM, the SS-NX-24. This missile is not yet operational.

3.3 Problems SLCMs Pose for Verification

The characteristics of SLCMs themselves raise problems for verification. We discuss these characteristics below, and attempt to put them in perspective for the verification discussion which follows.

1. SLCMs are relatively small and are manufactured and assembled at non-distinctive industrial facilities. Covert production facilities of this type would be difficult to detect with NTM. On the other hand, these facilities are part of a sophisticated industrial infrastructure and would be time-consuming and expensive to duplicate. The explosives and solid fuel booster make final assembly facilities more difficult to hide.

[6] "Soviet Nuclear Weapons" (Vol 4 of the Nuclear Weapons Databook) by T. Cochran, W. Arkin, R. Norris, and J. Sands for the National Resources Defense Council (Harper and Row 1989) p. 158.

Further, though SLCMs are small relative to other strategic delivery vehicles, they are not small relative to available space and other equipment the ships and submarines that they are deployed on. In fact, they are the largest weapons loaded onto or stored on US SLCM platforms. SLCMs on US ships are found only in their launchers or submarine torpedo rooms.

2. Nuclear and conventional SLCMs can employ the same airframe. This is true of the US Tomahawk, and the Soviet dual capable short-range SLCMs. Two implications of this are:

   (a) US nuclear and conventional SLCMs once in their launch canisters, cannot be distinguished by visual inspection alone. An inspector with nuclear detection or radiographic equipment would be required to identify a US SLCM as nuclear or conventional. The Soviets have only nuclear long-range SLCMs, so distinguishability is not yet an issue with these; it is not known whether their short-range nuclear ship-attack SLCMs are visually distinguishable from conventional counterparts.

   (b) It would, in principle, be possible to convert conventional SLCMs to nuclear SLCMs in the field. But it is important to point out that existing US SLCMs have not been designed for this, and cannot be readily converted. The Tomahawk is considered a complex system, not designed for field maintenance. Each missile is tested before it leaves the factory and then remains intact until it is fired or returned for maintenance or for recertification (which occurs every 3–4 years). During the time the missile is in the fleet, no maintenance is performed and electrical continuity is never broken. To change a variant from conventional to nuclear would require replacement of the entire front one-third of the missile, access to the avionics to change additional read-only memory, and re-establishment of electrical continuity and fuel-line integrity. Operational confidence would require a complete retest of the missile, which currently requires rather elaborate test equipment and well-trained technicians. Admiral Hostettler, head of the Cruise Missile Project Office from 1982-1986, made it clear that conversation at sea is not intended: "clearly this is beyond the scope of normal Navy maintenance concepts .... The capability to modify

variants in the fleet is not planned for the Tomahawk"[7].

It is not known whether Soviet short-range SLCMs are readily convertible from conventional to nuclear. It is, of course, important to determine whether they are. If Soviet SLCMs are also too difficult to convert in the field, then the problem of convertibility is not an immediate one, and could be addressed by placing treaty constraints on future SLCM designs.

3. SLCMs can, in principle, be deployed on any outfitted vessel. Other strategic systems (e.g. SLBMs and ALCMs) are deployed on dedicated strategic platforms that can be counted by NTM. Nuclear and conventional SLCMs are dispersed throughout the fleet on nonstrategic platforms. Verification efforts must therefore be directed at a variety of platforms. It should be noted, however, that in both the US and USSR the SLCMs are deployed only on particular classes of military vessels, and that NTM could determine whether other classes were being modified to carry SLCM launchers.

   There are several operational problems associated with covertly deploying nuclear SLCMs on nonmilitary vessels. It would be necessary to equip the ship with launchers, qualify the ship for use of the launchers, train the ships personnel to target and launch nuclear missiles (or dedicate trained military personnel to nonmilitary ships), and accept the reduced control over nuclear weapons. Though this is a cheating scenario which is often raised, the operational risk associated it with would be quite high.

3.4 Verification Measures

The verification of a SLCM arms control agreement would rely on a set of measures collectively designed to ensure compliance with the agreement, and to assure that militarily significant violations can be detected in time to be countered. A well-designed verification scheme should expose a potential cheater to multiple risks of detection. Verification measures should drive up the cost of cheating (for example, requiring the cheating party to

[7]Testimony before the Senate Armed Service Committee, March 15, 1985.

build a parallel covert infrastructure for production, maintenance and storage, or to drive the cheating party to consider risky forms of covert deployment). Verification measures should make circumvention costly, cumbersome and generally unattractive relative to other military options. The measures should also significantly increase the probability of detection of any attempted circumvention—they should make cheating difficult to get away with.

No single verification measure should be expected to carry the full burden of SLCM verification. As a corollary, no single verification measure need be foolproof. A successful verification scheme will put up a series of barriers, each of which must be circumvented by the cheating party, and which collectively present a significant deterrent to cheating.

1. NTM, so effective in monitoring ballistic missile silos, are not particularly useful for monitoring SLCM inventories, or detecting covert SLCM production facilities. NTM can, however, monitor SLCM launchers, and obtain data indicating testing of modernized SLCMs (which might have, for example, improved guidance or propulsion). An agreement prohibiting encoding telemetry from SLCM tests would ensure NTM a role, though covert SLCM tests need not use telemetry. NTM can also count the number of vessels which have been externally modified to carry SLCMs and the number of launchers on each of those platforms.

2. Data exchanges, and the validation of the data exchanged, would establish baseline conditions and support monitoring of the agreement. The actual data exchanged would depend on the provisions of the treaty, but could include location of SLCM facilities, SLCM design data, and information on the numbers of deployed SLCMs, their platforms and launchers. It should be noted that time is an ally of validation of SLCM numbers: over a period of three years, all US SLCMs cycle back to the factory for recertification; if deployed SLCMs are to be hidden, they must be hidden through the entire maintenance and redeployment cycle, and must evade all verification measures during that time.

3. Perimeter portal monitoring at facilities which are "choke points" in the path a weapon follows from production to deployment would permit a count of SLCMs. In particular, perimeter portal monitoring would be important at final production and final assembly facilities to identify

and count legally produced SLCMs. The establishment of portal monitoring at these locations would require a cheating party to establish a separate, covert infrastructure to manufacture SLCMs.

4. Radiation monitoring. [This topic will be discussed in detail in Section 6.] Some verification scenarios may require that, somewhere along the path from production to deployment, inspectors verify that a SLCM does not contain a nuclear warhead. If inspection occurs at locations where the inspector has close-in access to a canisterized SLCM, there are technically feasible approaches to the detection of nuclear material. For example, there is not enough free volume between the warhead and the canister to shield the warhead from an active neutron detector a few meters away.

   In contrast, remote detection of nuclear material is not considered feasible. Although it is perhaps possible to position a detector up to 50 meters outside a ship and detect nuclear material in the warhead of a cruise missile near the deck[8] successful detection depends critically on the nuclear materials used in the warhead, and on the cooperation of the other party: very modest shielding efforts could easily disguise the presence of nuclear material from either a passive or active detector.

5. Tags and seals. Tagging is the process of marking a missile, and/or its canister, so that it can be identified at some later time. A tag for Soviet missiles would be developed and produced by the US, and applied by US inspectors. The tags must be tamperproof, non-reproducible, and environmentally stable, and inspectors must be able to check the tags and describe results without transfer of sensitive technology. Technologies exist to make each tag unique (i.e. to "fingerprint" each SLCM), or identical (i.e. to label all the legal SLCMs with a common tag). Some of these technologies and their uses will be discussed in Section 5.

   Tagging by itself does little; but if used in conjunction with other verification provisions (PPM, short-notice inspections) it could be an effective means of monitoring legal SLCMs and isolating covert production lines. Covert SLCMs would have to bypass, or be hidden at, every

[8]The experiment performed cooperatively by Soviet scientists and a U.S. delegation organized by the National Resources Defense Council (July 1989) on the Soviet cruiser Slava showed a strong signal at 30 m and a marginal one at 56 m from a SS-12 warhead. (H. Lynch, private communication)

location that was subject to inspection, since without valid tags they could be immediately identified as illegal.

To aid in assuring that conventional SLCMs have not been converted to nuclear, a seal could be applied to conventional SLCMs. The seal would not prevent tampering or conversion, but on subsequent inspections would indicate whether it had occurred. The technology exists to manufacture tamperproof, unreproducible, durable materials which could be used to seal conventional SLCMs after they have been verified to be conventional.

6. Inspections. Inspections increase the probability of detection of violations and increase the effort required for a cheating party to hide its covert production and deployment. When inspections are used in conjunction with tagging, they can be effective in isolating clandestine production and deterring covert deployment. A covert SLCM, a SLCM without a tag, must be hidden at every stage where there is a threat of inspection.

   (a) Inspections to monitor deployments would play an important treaty role. A stumbling block in the development of this concept has been reluctance to allow shipboard inspections on the grounds that they are unacceptably intrusive and could compromise sensitive technology, techniques or operating practices. SLCMs deployed on declared SLCM carriers must be inspected, but we believe that the inspections could be accomplished without sending inspectors on board US ships.

   US SLCMs are deployed on battleships, cruisers, destroyers, and attack submarines. The surface ships carry SLCMs in their launchers — they do not carry reloads. There are sound reasons for this practice. These ships do not have storage space for objects the size of SLCMs (except in helicopter hangers, whose interiors are visible from outside); they do not have large elevators capable of transporting SLCMs between decks, or the moving and handling of equipment (such as rated cranes) for objects the size and weight of SLCMs. There is a need for stringent safety precautions onboard ships (e.g. fire suppression systems) when storing and handling explosives (like the solid rockets and fuel inside SLCMs); and there is a need for nuclear security precautions. Generally speaking, while it is not physically impossible to store a nuclear SLCM elsewhere

on one of these SLCM platforms, it is both difficult and hazardous to attempt to do so. The constraints are even more stringent for US attack submarines: SLCMs cannot be moved into or out of torpedo rooms of US submarines at sea.

Although US ships do not carry Tomahawk reloads, SLCMs could, in principle, be transferred onto these platforms from ammunition ships at sea. This is not current practice on US surface ships (in fact, US surface ships do not carry cranes certified for lifting objects as heavy as Tomahawks). We believe that reloading US attack submarines at sea would not be a viable practice, as these do not even reload torpedoes at sea.

Soviet Naval practices and procedures are not well known, but many of the operational and safety considerations discussed above must apply. In this regard, it is worth noting that Soviet short-range SLCMs are all larger (some much larger) and presumably heavier than the Tomahawk, and therefore would be even more difficult to store covertly and transport to launchers. Soviet practice is to keep nuclear warheads under even tighter control than in the US. A cheating scenario involving reloads or conversions of Soviet nuclear SLCMs at sea would run counter to long-established operational practice. Soviet aircraft carriers may require special treatment as they may be the one SLCM platform with storage space and equipment required for covert deployment.

In principal there is a possibility that covertly deployed nuclear SLCMs could be thrown overboard in order to avoid detection in a challenge inspection. However this would run counter to both US and Soviet strict accountability for nuclear warheads and weapons. One need only recall the distress caused by the unintentional loss of a few US nuclear warheads at sea to convince oneself that deliberately throwing a nuclear SLCM overboard would only be considered under singularly desperate circumstances.

The important point is that since under practical circumstances SLCMs are limited to their launchers or torpedo rooms, inspection of deployed SLCMs could be performed without sending an inspector on board the ship or submarine. An inspector on the dock could request that a specified launcher be unloaded, and its missile brought off the ship for inspection. That inspector could observe the unloading, then either check its tag and seal on the

dock, or follow it to a verification facility. Selection of missiles in a torpedo room could be done remotely using, for example, a video camera that was lowered through the weapons-loading hatch directly to the torpedo room.

The inspection procedure outlined above should have minimal impact on naval operations. For example, a violation consisting of the deployment of 50 illegal SLCMs randomly placed in 3000 available launch tubes would be detected at the 90% confidence level[9] by checking 138 of the missiles (276 for 99% confidence). This should be compared to the roughly 1000 Tomahawks that would be pulled from their launchers and returned to the factory each year for regular maintenance if the US deploys 4000 Tomahawks as planned.

Each party might also be permitted to conduct a certain number of shipboard inspections to verify that a ship declared not to be a SLCM carrier indeed had no launchers or SLCMs.

(b) On-site inspections at declared facilities are part of the verification protocol of the INF treaty, and could play an important role in SLCM verification as well. For example, a certain number of short-notice inspections could be permitted at facilities and installations declared to be associated with SLCM storage, maintenance and repair, testing, and elimination.

Flight test ranges could be monitored to insure that any SLCM being flight-tested was properly tagged, and therefore came from a legal production facility. This would make it more difficult for a party to qualify a covert production line by flight test, and require them to choose between the risk of being caught testing a covert SLCM and the risk of deploying an untested SLCM line.

(c) Suspect site inspections. Each side could have the right to inspect a limited number of sites which they suspect may be engaged in

[9]This is computed by noting that, for a random distribution, the relation between confidence, C, in detecting a violation; the total number, N, of missiles deployed; the number n, of illicit missiles; and the number of checks, S, is given by

$$S = \frac{N}{n} \ell n \left(\frac{1}{1-C} \right).$$

This formula applies for sampling with replacement–i.e., a particular item may be drawn more than once for inspection. The correction for sampling without replacement is negligible for large $N/n \gg 1$.

covert activity. These would not be "anytime, anywhere" inspections, but would allow the sides some latitude in uncovering covert activity.

None of these verification techniques, by itself, is sufficient to verify limitations on SLCMs. The methods may, however, be used in synergistic combination to increase overall confidence in the monitoring of a treaty.

The success of a verification scheme depends on its cumulative effect in deterring and detecting cheating. The potential cheater must do more than evade one verification measure, he must incur the costs and risks of evading all of them.

To illustrate the cumulative effect, consider a verification regime consisting of PPM, a verification facility where legally produced SLCMs are tagged and sealed, and a protocol of challenge inspections including inspections of deployed SLCMs and suspect sites. This is illustrated schematically in Figure 3-3. For each measure:

$P \equiv$ probability of successful evasion of a given measure
$C \equiv$ cost of evasion
$R =$ operational risk incurred as a result of evasion (that is, effect on weapon success)

Note that $P = P(C)$. The more you're willing to pay, the higher your probability of success, as drawn schematically in Figure 3-4.

1. At the manufacture and final assembly stage, a cheater has the choice of manufacturing "extra" SLCMs in the existing, legal facility and sneaking them out (probability of successfully evading $PPM \equiv P_{PPM}$), or incurring the cost of establishing covert manufacturing facilities ($\equiv C_M$). The probability that those covert facilities will not be detected (by NTM or suspect site inspections) is given by P_M.

2. The next verification measure is the checking, tagging and sealing of legal SLCMs. This has the effect of labeling legal missiles—any missile in the logistics chain without a tag is immediately identified as illegal. A cheater with illegally produced missiles must either sneak them into the verification facility as legal (with a probability of success of P_{SVF}), allow them (without tags) into the normal logistics flow (with

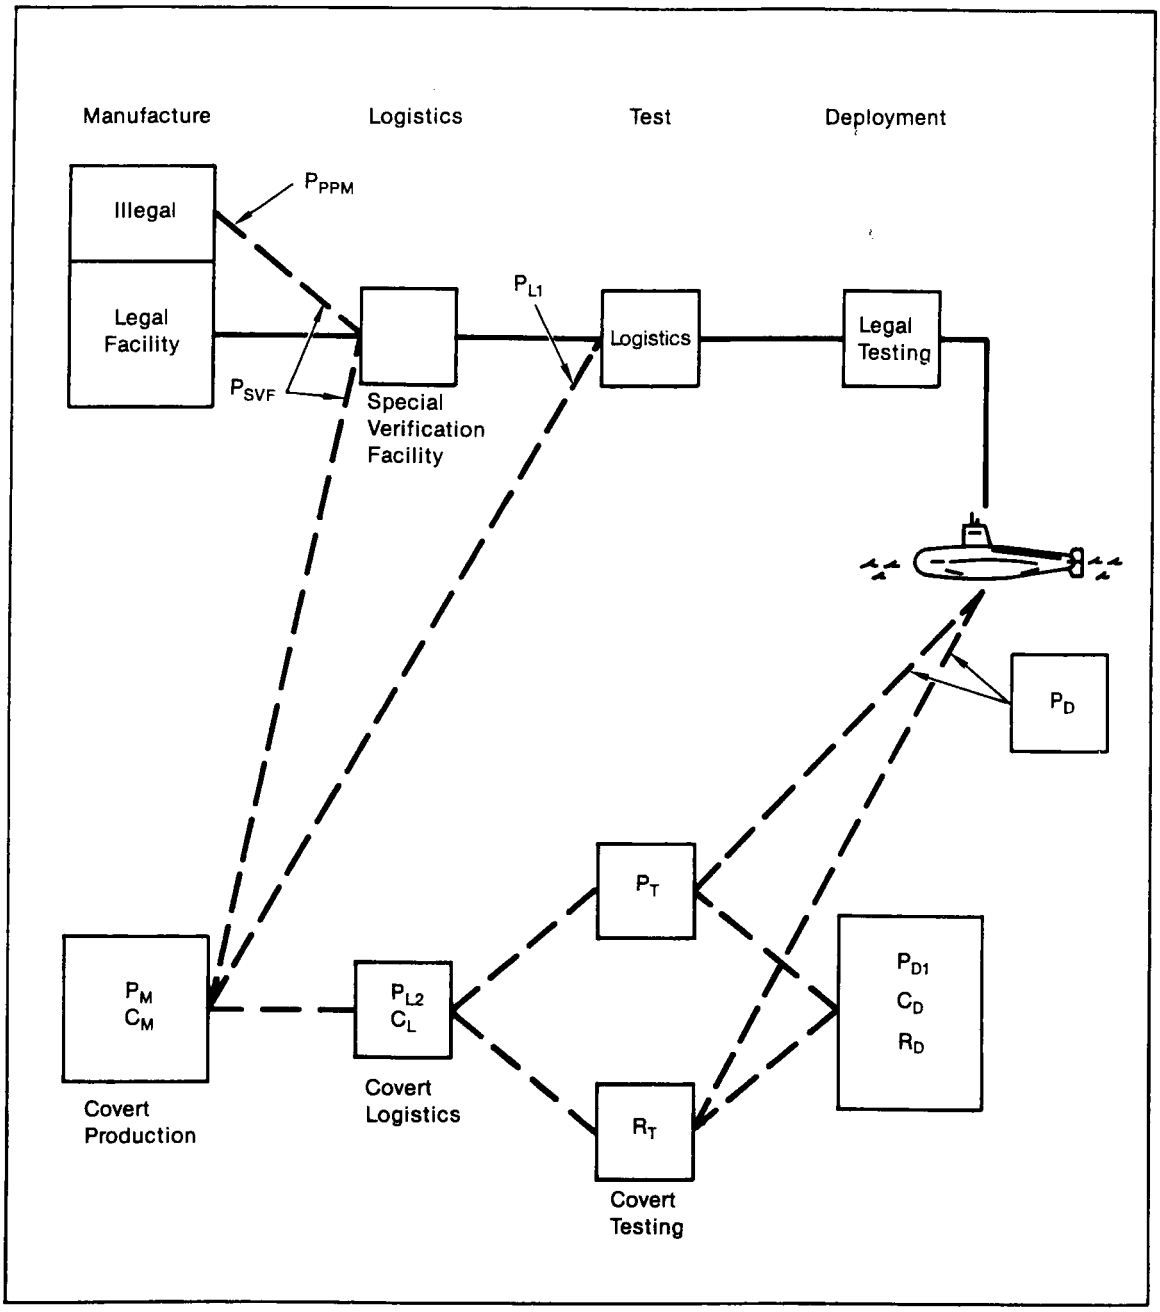

Figure 3-3. Possible cheating paths showing probability of evasion, costs, and operational risks incurred along the way.

probability of success $\equiv P_{L1}$), or establish a covert logistics chain (with probability of successfully evading NTM and suspect site inspections $\equiv P_{L2}$). There is a cost, C_L, in establishing and maintaining the covert logistics chain. Note that P_{SVF} can be made quite small by monitoring SLCMs' progress from the PPM station to the SVF.

3. An important, and separate, part of the path from production to deployment is the flight-testing of missiles. A cheater must decide whether to qualify a covert line through flight test (with an associated probability, P_T, of not being detected), or to bypass flight test and take the additional risk, R_T, that SLCMs produced via the covert line may have undetected problems.

4. The next phase is deployment. The cheater has the choice of deploying the illegal (untagged) SLCMs in launchers of declared SLCM platforms (with a probability of successful evasion $\equiv P_D$, which depends on the inspection sampling rate), or deploying them on non-SLCM platforms - an act which carries a cost $\equiv C_D$, an operational risk $\equiv R_D$, and a probability of successful evasion $\equiv P_{D1}$.

Figure 3-3 shows possible cheating paths, and indicates the probability of detection (or successful evasion) and both the cost and the possible loss of confidence in the weapon at each step along the way. The overall probability of successful evasion–both in not being detected and not introducing operational risk– is the product of probabilities of success at each stage, and the overall cost is the sum of costs incurred at each stage i.

$$\begin{array}{lll} \text{Probability of successful evasion} & P = \prod_i p_i & \\ \text{Operational risk as a result of evasion} & R = 1 - \prod_i (1 - r_i) & \\ \text{Cost of evasion} & C = \sum_i c_i & \text{individual steps i} \\ \text{Evasion} & & \end{array} \qquad (3-3)$$

This illustrates that no single verification measure need be 100% effective. In fact, a cheater must recognize that the probability of success is no

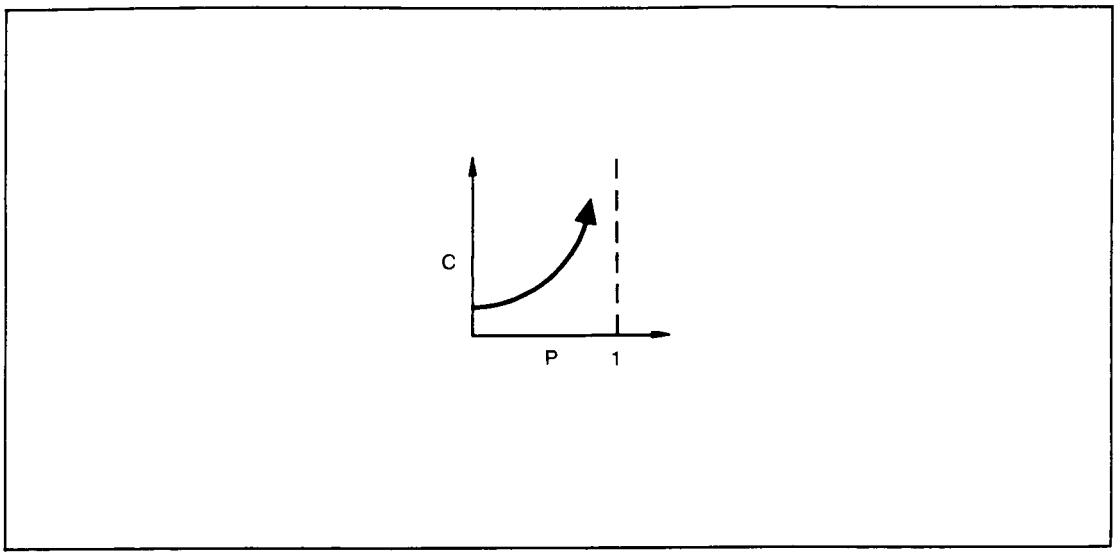

Figure 3-4. Sharply rising costs with increased probability of successful evasion.

greater than his probability of successfully evading the most effective verification measure. In deciding whether to cheat, a side must balance the overall probability of success (or, conversely, of detection) with both the additional cost and the reduced confidence in the weapon.

3.5 Verification Regimes for Two Hypothetical Treaty Limits on SLCMs

Since any verification regime for treaty limits on SLCMs is likely to employ several complementary verification measures, it is not easy to make general statements about how such a verification scheme would function. To give a feel for the quality of verification that might be achieved in different cases, it is necessary to consider concrete examples. In this section, we illustrate how the suite of verification measures described in Subsection 3.4 could work together for two specific hypothetical SLCM treaties: 1) A treaty which bans all nuclear SLCMs (both short and long range), but which contains no limits on conventional SLCMs; 2) A treaty which includes specific sub-limits on nuclear and conventional long-range SLCMs, but which does not limit short-range antiship SLCMs.

The discussion which follows is based in part on two recent Stanford papers discussing SLCM verification provisions[1,2]. We believe that the verification schemes outlined below can be implemented without compromising the policy of the US Navy to "neither confirm nor deny" the deployment of nuclear warheads on specific ships.

3.5.1 Verification of a Treaty Banning Nuclear SLCMs of All Ranges

A treaty which would simply ban nuclear SLCMs of all ranges has recently aroused interest because it has several practical advantages:

- Because the US has more long-range nuclear SLCMs and the Soviet Union has more short-range nuclear SLCMs, both sides would have to eliminate a comparable number of deployed weapons.

- Both sides would have a comparable number of Naval surface ships and submarines subject to verification. In particular, since the Soviet long-range SS-N-21 SLCM is currently believed to be deployed exclusively on submarines, by including short-range SLCMs in the ban one would bring Soviet surface ships under verification.

- This treaty language would not complicate the START regime by bringing conventional weapons (i.e., conventionally armed long-range SLCMs) into the treaty for the first time.

- The problem of determining how to verify a SLCM's operational range would not arise, since the treaty would cover all ranges.

- If all nuclear SLCMs were banned, it would become much easier to avoid compromising the US policy to "neither confirm nor deny" whether a given ship carries nuclear weapons. Indeed once US ASROCs and SUBROCs are retired in the early 1990's, the only nuclear weapons on board US Naval ships, other than aircraft carriers, would be the SLBMs aboard strategic submarines.

The verification regime would monitor the treaty-limited SLCMs in various ways from cradle to grave, starting with production and continuing through deployment and eventual destruction of SLCMs. The goal would be to force a side desiring to cheat to establish a whole new, parallel *covert* production and testing chain if it wanted to violate the treaty in a militarily significant way. The establishment of such a covert parallel infrastructure for SLCM production and testing would be both expensive and risky.

Since the details of the paths which Soviet SLCMs follow from production to deployment are not available to us, we shall discuss here procedures which would be applicable to monitoring US SLCM inventories. Analogous, but not identical, procedures would be developed for Soviet inventories once the data exchange described in the first paragraph below had occurred.

A proposed verification regime for a treaty banning all nuclear SLCMs would include the following elements:

1. An extensive data exchange, to establish baseline conditions for verification. Both sides would identify the types of SLCMs in their stockpiles and under development, declare locations of facilities for SLCM

airframe production, assembly, and testing, nuclear warhead mating facilities, SLCM storage and maintenance depots, flight-test ranges, ship loading depots, etc. Both sides would provide baseline data on the number of SLCMs at each declared site, and would give sufficient transit information to allow the cruise-missile logistics flow to be monitored. Both sides would provide specific design data (weight, size, range) on SLCMs and their canisters and would declare which classes of ships carry the treaty-limited SLCMs. Finally, one would agree on procedures to monitor the destruction of existing SLCMs as required by the treaty.

2. Perimeter-portal monitoring of all declared facilities that assemble conventional SLCMs, and of facilities where nuclear SLCMs used to be produced, if these are separate. (The US has 2 such production facilities for the Tomahawk.)

3. Monitoring of all SLCMs leaving the production facilities, until they reach a facility at which they would be verified not to carry nuclear warheads. As an alternative to monitoring in transit, one could tag and seal the SLCM shipping containers, and then verify that only tagged and sealed containers enter the verification facility, and that the same number enter as had left the production facility.

4. A facility at which *all* legally produced SLCMs would be checked to verify that they do not carry nuclear warheads. This verification facility could be at the portal of the production facility or some distance away. A variety of methods, both active and passive, would provide acceptable signals if nuclear materials were present in quantities used in nuclear warheads. This test could be done on a SLCM in its canister, if prior data exchanges had indicated that there was not room for significant shielding between the SLCM itself and the canister wall. Once a SLCM had been verified to be conventional, its canister would be tagged for identification and sealed to permit detection of tampering. Any SLCM discovered during subsequent challenge inspections or routine monitoring that did not have a legal tag and unbroken seal would be presumed to be in violation of the treaty.

5. Inspections of a sampling of deployed SLCMs, to verify that only treaty-approved SLCMs were being deployed. The regime for inspection of *ships* is a sensitive topic, as the US Navy does not want Soviet inspec-

tors on board its ships. It would be possible to avoid ship inspection altogether if PPM could be established around Naval ports.

However, this would likely be both impractical and undesirable. Thus a limited number of short-notice *challenge* inspections of ships in port, both Naval and merchant-marine, appears necessary. One could envision a different number of challenge inspections allowed for ships and submarines declared to be SLCM platforms and those that are not supposed to be carrying SLCMs.

It might be possible, however, to inspect SLCMs on declared platforms without sending inspectors on-board. As discussed in Subsection 3.4.6, if one examines US Naval operational procedures and SLCM platforms, a case can be made that challenge inspections of US Navy ships in port might be accomplished without sending inspectors on to ships or submarines. More concrete data regarding the Soviet Navy's operational practices are needed before we can make a similar assessment regarding challenge inspections of the Soviet Navy's ships. But based on our present knowledge we think that Soviet Naval ships[10] might be inspected remotely as well, since there are typically few large storage areas with enough volume and access to hide a SLCM. The inspection of a sampling of deployed SLCMs would verify that those sampled had legal tags and undisturbed seals.

Challenge inspections of merchant ships suspected to be carrying undeclared SLCMs would have to take place with inspectors onboard ship, but since these ships are not supposed to have military missions, presumably there would not be cause for the side being inspected to object. (An exception might be Soviet trawlers on offshore intelligence missions, but this case could be taken care of via a limited veto-power over challenge inspections of merchant ships.)

6. A limited number of short-notice challenge inspections would supplement the monitoring of the production chain. Declared SLCM infrastructure facilities subject to challenge inspections could include production, storage, and maintenance facilities, declared test ranges, etc.

[10]If a treaty included a limit, not a ban, on nuclear SLCMs it might be reasonable to eliminate inspections of subs because the incentive for cheating by adding additional SLCMs to these platforms is rather low: since there is limited space available, deployment of additional SLCMs could only be done at the expense of other important military systems.

In addition, undeclared suspect *land*-based sites would be subject to a limited number of short-notice challenge inspections. The purpose of these would be to deter or detect the development of a covert production chain for SLCMs.

7. Verification of elimination of existing nuclear SLCMs. The inventory of treaty-limited SLCMs that is already deployed must also be tagged and sealed. This could be done at a declared facility and on a time-scale agreed upon by both sides. In addition, US cruise missiles are returned to a service facility every few years for maintenance and testing. PPM inspectors at the declared service facility would verify the removal of the tag and seal from an incoming SLCM canister on the way in, and re-tag and seal it on the way out. Depending on the specific language of the treaty, SLCMs in the service facility might be removed from the accounting system once their tags had been removed.

Most of these elements have already been discussed in more detail in Subsection 3.4. For this treaty regime, convertibility of conventional SLCMs to nuclear ones is a pivotal issue. However because there are no permitted nuclear SLCMs in this regime, many of the steps of the verification process become considerably less complicated.

In the present study we did not undertake to decide whether a treaty regime with no nuclear SLCMs would be acceptable or desirable from a military point of view. However, we were struck by the considerable diversity of opinion we encountered on this subject among the military staff with whom we spoke. Several individuals emphasized that nuclear SLCMs give the US Navy added operational flexibility. But we heard no clear consensus on either a strategic or theater mission for long-range nuclear land-attack SLCMs, and since US nuclear SLCMs are not currently in the SIOP, their role in the overall nuclear deterrent posture is not clear at present. Perhaps we were not in an appropriate forum to hear a full and authoritative discussion of these issues. Nevertheless, we think it would be valuable at this time to undertake a careful reassessment of the military role of nuclear land-attack SLCMs, relative to the military (and political) advantages to be gained from eliminating the comparable Soviet threat.

3.5.2 Verification of a Treaty with Separate Sublimits on Nuclear and Conventional Long-Range SLCMs

The treaty would not apply to short-range SLCMs, so provision must be made for verification of the ranges of declared SLCMs. A crucial component would be the specific range limit below which the treaty would allow SLCMs to operate. This range limit should be low enough that one could not covertly convert a warhead from conventional to nuclear, and use the consequent enhanced range (due to the smaller volume occupied by the nuclear warhead) to exceed the agreed-upon limit. The range limit should also be low enough that one could not covertly enhance the propulsion system of a short-range SLCM, either nuclear or conventional, to boost it into the long-range category. One would need to agree upon a clear criterion, such as an "as-tested" rule, that would establish the range of a SLCM class for purposes of the treaty. Prohibition of the encryption of flight-test data would aid in range verification if an "as-tested" rule were adopted.

The monitoring regime would begin at the point of SLCM production. All completed long-range SLCM airframes emerging from declared production facilities would be given tags, which would be checked at a perimeter-portal monitoring station and counted under the treaty limits.

In current US practice all of the SLCMs, both conventional and nuclear, then follow roughly the same route, although they are mated with their warheads at different points along that route. The first stop after the production facility is a final assembly facility where the airframes are fueled and attached to their solid-rocket launch boosters. Conventional SLCMs also receive their conventional warheads at this facility. Since the facility handles high explosives, flammable fuel, and solid-rocket motors, it is of necessity distinctive in that it is specially equipped and located away from populated areas. Both conventional and nuclear SLCMs are loaded into their canisters before emerging from this facility, although the nuclear versions are still without warheads. The closed canisters would be tagged leaving the facility. PPM would count the canisters leaving the facility, and verify that all the tags on the airframes entering and on the canisters leaving the facility were intact.

The producing country would declare canisters leaving the facility to be either nuclear or conventional. (Recall that nuclear SLCMs would only receive

their warheads at a point later in the chain, but for the purposes of counting they would be *declared* to be nuclear when leaving the fueling facility.) It is desirable from the point of view of the security of nuclear warhead design to avoid radiographic or other close probing of nuclear SLCMs. Therefore, limits on *nuclear* SLCMs could be verified by counting *total* tagged SLCMs leaving the facility, and by subtracting the number of verified *conventional* SLCMs from the total.

PPM inspections using nondestructive techniques, such as active or passive radiation detection, would verify that those SLCMs declared to be conventional do not contain special nuclear materials. This might be done at a facility adjacent to the portal of the fueling facility. A variety of radiation detection techniques would be technically feasible (see Section 5 of the present report), if inspectors were given access to the exterior of the SLCM canister. Once a canister containing a conventional SLCM was verified to be nonnuclear, it would be tagged, sealed with a tamperproof seal and counted. PPM would count the tagged and sealed canisters leaving the facility.

Those SLCMs destined to receive nuclear warheads would be declared nuclear, and their canisters would be tagged but not sealed. They would then be counted under the sublimits of the treaty. They would then be transported to the facility at which the nuclear warhead is mated with the SLCM airframe. At the entrance to this mating facility, PPM would verify that all entering SLCM canisters had authentic tags, so that covertly produced missiles could not be brought into the facility for nuclear warhead mating. When exiting from the mating facility, the canister of a nuclear SLCM would be given a tamperproof seal, whose intactness would be verified by PPM inspectors. After this stage all SLCMs, both conventional and nuclear, would be tagged and sealed.

In this scenario it would be necessary to make the tags of canisters containing nuclear and conventional SLCMs indistinguishable from one another, in order to avoid compromising the "neither confirm nor deny" policy of the US regarding use of nuclear weapons on specific ships. Since all tagged SLCMs would already have been counted against the treaty limits during the production chain, in a system using identical tags the goal of an inspection would be to verify that all SLCMs had legitimate tags and intact seals.

A potential weak point of the above verification scenario is the possibility of conversion of legitimate conventional SLCMs to illegitimate nuclear ones in time of escalating crisis. In times of normal relations between countries, it is likely that the application of seals, together with the threat posed by challenge inspections of ships, would inhibit conversion of conventional SLCMs to nuclear ones. But as long as conventional SLCM designs have enough space inside to allow installation of a nuclear warhead (as does the Tomahawk), the possibility of conversion to nuclear in time of escalating crisis is a potential problem. PPM and NTM might deter a covert production line for *airframes*, but there is nothing in this verification scenario to prevent the production and stockpiling of extra *nuclear warheads* for SLCMs. In a crisis, the seals on conventional SLCM canisters could then be deliberately broken, and in principle the extra stockpiled nuclear warheads could be used to convert conventional SLCMs to nuclear ones, creating a scenario for rapid breakout from the treaty regime. This is not practical with the current U.S. Tomahawk design, but it is important to establish whether it is impractical for the Soviets, and to take steps to ensure that it will be impractical in the future.

Most importantly, it will be very important to get more information about whether the convertibility of Soviet SLCMs is an issue. Although most of the older Soviet SLCM models exist in both nuclear and conventional implementations and are hence potentially convertible, these SLCMs are short-range (from 50 to 550 km). If the range limits of the treaty are such that they are not included in the treaty limits, then their convertibility is not an issue for the purposes of a treaty limiting *long*-range SLCMs. The new Soviet long-range land-attack SLCM, the SS-N-21, presently exists only in a nuclear version. If no conventional version is planned, then convertibility is also not an issue for the present *long*-range Soviet SLCM force.

Secondly, it will be important to decrease the inherent convertibility in the design of conventional SLCMs. As discussed in some detail in Subsection 3.2, the US does not at present have the capability to convert a conventional Tomahawk to a nuclear one in the field, and US Navy operational personnel we have spoken with say that they find it quite implausible to imagine doing so in the future, for operational reasons. Nevertheless, it would be an important step to include in the treaty language a provision that *future* designs of conventional SLCMs should be inherently nonconvertible to nuclear, as discussed for ALCMs in Section 1 of this report.

It should be noted, however, that the question of the convertibility of the present short-range Soviet SLCM force re-emphasizes the importance of careful consideration of the range limit for a treaty limiting long-range SLCMs only. It would be desirable to have a gap between the upper limit of the range of Soviet short-range SLCMs (around 550 km) and the lower limit for the definition of "long-range" treaty-limited SLCMs, so that the Soviets could not covertly convert a conventional short-range SLCM to a nuclear version, and use the consequent enhanced range to turn it into a long-range SLCM. Concrete schemes to verify SLCM range limits are an area that needs further thought and analysis.

It is apparent from the length of this discussion that the overall verification regime for a treaty with sublimits on nuclear and conventional long-range SLCMs is going to be complex and intrusive. This stems from the fact that the delivery vehicles are small in size, are manufactured in industrial facilities that are externally similar to many in the civilian economy, and have many platforms from which they could potentially be launched. None of the verification measures by itself is adequate for such a situation; hence the complexity.

3.6 Summary

We have described the problems and possiblilities for verifying limits (or bans) on SLCM deployments, starting with an analysis of how much and what kind of verification is required in view of the nature and limits of the threat they pose. It is shown that effective verification would rely on a set of measures collectively designed to ensure compliance and to assure that militarily significant violations can be detected in time to be countered. No one measure alone can meet these requirements for SLCM verification. The verification regime for two hypothetical treaty limits on SLCMs are considered in detail: a ban on nuclear SLCMs of all ranges and separate sublimits on nuclear and conventional long-range SLCMs.

REFERENCES FOR SECTION 3

1. "Potential Verification Provisions for Long-Range, Nuclear-Armed Sea-Launched Cruise Missiles", Workshop Report, Center for International Security and Arms Control, Stanford University, July 1988.

2. "A Proposal for a Ban on Nuclear SLCMs of All Ranges", G. N. Lewis, S. K. Ride, and J. S. Townsend, Center for International Security and Arms Control, Stanford University, June 1989.

4 A FORMAL DIALOGUE (PQC) TO ASSIST VERIFICATION OF MOBILE LAND BASED MISSILES

4.1 Application of PQC to Reservations on which Mobile ICBMs are Deployed

The U.S. requires reliable knowledge of Soviet compliance with the restrictions that START would place on mobile ICBMs. This section describes a verfication protocol (called PQC for partition, query, and challenge) to meet this requirement. It is designed to complement our existing intelligence[11] and be minimally intrusive. At the same time, it would place any Soviet violation at risk of exposure; for material violations that risk would be great.

PQC specifies a U.S.-USSR dialogue which exploits our necessarily limited knowledge of Soviet activities. Engaging in this colloquy magnifies the probability of exposure if subterfuge is present. The situation is familiar to all prosecutors: a well conducted investigation of persistent wrong doing should yield an indictment - at least one for perjury.

Important information is available from intelligence sources and National Technical Means (NTM). This information is independent of anything new that the treaty apparatus will provide. From the standpoint of deterring violations, it is particularly useful since the Soviets will not know exactly what it reveals. We should structure verification to exploit this information to the fullest and supplement it where necessary by a treaty sanctioned exchange of information.

The verification protocol will be symmetrical in outline and philosophy (although different in detail if the two sides adopt different mobile missile basing modes), so we treat only American verification of Soviet missiles. We presume START will specify reservations to which MICBMs will be restricted. PQC can be formulated to apply only within these reservations or alternatively it may be applied to the entire national territories. We consider these cases sequentially.

[11] Our national technical means (NTM) already provides an accurate count of ICBMs in fixed silos.

Suppose START permits N operational mobile ICBM's (MICBM's) with various sublimits dependent on MIRVing. Once an MICBM is mated to its launcher it is presumed operational and would be restricted to specified reservations comprising the MICBM deployment areas and the transportation routes connecting these to production, repair, and storage facilities. The Soviets would partition their MICBM reservations into a total number M of disjoint blocks. The number of blocks M would be specified by the treaty, but the Soviets would draw the partition. To fix the idea, we will take some "ball-park" numbers. Let $N = 200$ be the total permitted number of operational MICBMs, and $M = 40$. Up to, say, 12 times a year, but no more than twice in any 30-day period, we would ask the Soviets to label their partition. Within 6 hours they must supply the exact number N_i of MICBMs in each block B_i (of course $\sum_{i=1}^{M} N_i \leq N$) and guarantee that these numbers will not change for three hours. At this point, the Soviets have given us a block-by-block census but no missile coordinates. Within an hour the US could choose a box to query. Then the Soviets would have one hour to supply us with the exact locations (and any other agreed upon information, e.g., MIRVing counts and/or identification tag numbers) of the N_i MICBM's in that block.[12] If this information is consistent with our expectations based on NTM and other knowledge, the cycle may end there. If however, we have independent information of suspicious objects or activities elsewhere in the block we have selected, or occasionally at random, the U.S. would move to the next step, a challenge.

While queries are roughly monthly, challenges would be limited to a smaller number, perhaps 15 in a 5-year period. The U.S.S.R must fix any ICBM within the challenged block for two weeks and during that period allow our inspectors access to that block for the purpose of investigating any aspect of possible START treaty violations - particularly the presence of an undisclosed missile. Such access should be limited only by the USSR's legitimate needs to protect military secrets not connected with potential START violations. The rapid pacing of the steps labeling the partition, query, and challenge ensures that normal operational motions of ICBMs will not be materially disrupted except in the event of a challenge. And even then only the objects in the challenged block (in our example, perhaps one fortieth or 2-1/2 % of the force) are pinned. In a challenge the USSR would undertake to

[12] In our numerical example, on the order of only $5 = 200/40$ missile locations would be revealed during a query, i.e. $N_i \approx 5$.

expedite rapid access to any points within the block and a small grace-region beyond its perimeter. An advance party of the inspection team should arrive at the point of concern as soon as possible to freeze "ground truth".

The protocol PQC would necessitate few operational restrictions except within a challenged block. In particular, giving the USSR 6 hours to take the census of MICBMs in each block allows them to maintain a highly decentralized knowledge of specific missile locations. Such decentralization might be regarded as essential to security.

The Soviets would prefer to keep a rough parity among the number of missiles in each block since by challenging a block labeled with a large N_i the U.S. could temporarily pin that number of missiles. In our example, the Soviets might decide to keep each N_i between 3 and 7. Bunching of missiles can be avoided through logistics or by changing the partition. The Soviets would be free at any time to change the partition while retaining the total number of blocks M. They could respond to the next request for labels by presenting a new partition complete with numerical labels N_i. A particularly straightforward, but in no way mandatory, organization of the MICBM force would be to divide the command into 5 wings each comprised of 8 squadrons. Each squadron, though operating in close proximity (within several kilometers) to the 7 others in its wing would stay in its own block B_i, making each $N_i = 5$ again on the assumption $M = 40$ and $N = 200$. If missiles of different blocks are near each other, as in this example, it will be necessary to define the partition with some accuracy.

The definition of an accountable MICBM must be precise. Both sides will wish to retain many extra missiles at production facilities as spares and for test flights. These would not be included in the block totals N_i since they are not operational. It must, however, be ensured that these "extras" cannot rapidly become operational. For this launchers must be tightly controlled: very few spares can be permitted and those that do exist should not be permitted near missile storage areas. Also the storage of extras and the design of launchers should be restricted so that rapid reloading is not possible. In this way, the MICBM coupled to its launcher becomes a natural unit of threat and one susceptible to regulation.

Does PQC meet the rather complex constraints of secrecy? The locations of mobile missiles are supposed to be secret - that is why they are mobile. Yet,

the inspected side must be willing to convey information on its mobile missile deployments (and on other matters related to production and stockpiling) if they wish to establish that numerical ceilings have been respected. However, they will be reluctant, and properly so, to convey information which could be used to target their mobile missiles. Inspection demands a compromise - it is difficult to imagine inspecting or verifying a missile without knowing its location. The inspected side must be willing to disclose the location of a small and rotating fraction (perhaps 2 % to 5 %) of its mobile missile forces. It is worth noting that both navies live with comparable but greater vulnerabilities as each ballistic missile submarine spends a substantial fraction of its life cycle in port. We will argue through a simplified model that revealing the partition and its labels does not increase the vulnerabilities of the USSR's (or the U.S.'s) missile force.

There is another aspect of secrecy. For the inspecting side it is costly to reveal a suspicion derived from intelligence when this might be traced to its sources. Also we do not wish to instruct the Soviets in the strengths and weaknesses of our NTM by establishing a routine in which we frequently pinpoint the site of greatest concern. The preceding protocol would enable our intelligence community to pursue a suspected violation without revealing its exact location. For example, suppose a possible illicit missile came to our attention but upon making the appropriate block query its coordinates were announced in good order; then no further action would be required. Only if the results of the query appeared to confirm a suspicion, e.g., if the "hidden missile" was not associated to any of the supplied coordinates, would we have to visit the site. Thus we avoid disclosing what is bothering us - and thereby revealing aspects of NTM and/or other intelligence methods - in the first round. This substantially reduces the exposure to our intelligence methods. If the suspicious object lay outside the Soviet MICBM reservation PQC, as described above, would not play a role. Such concerns would be brought to the appropriate consultatory committee. Alternatively START could anticipate such concerns by incorporating a more extensive application of PQC in which the partition blocks, taken together, covered the entire Soviet Union. The costs and benefits of an extended nationwide PQC are considered in Subsection 4.2.

We consider two questions quantitatively:

1. How much is the vulnerability of the deployed force of MICBMs increased by information revealed by the PQC protocol?, and

2. Would violations be discovered with high confidence in a timely fashion?

To answer the first question we must specify a hypothetical attack. This is done in the context of a simple two-person game. In this game, mobile missiles become tacks on a game board which we take to be 40 squares wide and 100 squares long. Each of the 40 columns corresponds to one block B_i and each of the 100 squares per block is assumed to represent a possible aim point for an attacker. The mobile force of $N = 200$ missiles is scattered over the 40 x 100 = 4,000 squares. Taking two scenarios, we consider attacks of 1,000 and 2,000 warheads and assume unit (100 %) probability of killing any MICBM on a warhead-targeted square. Player 1 is faced with the problem of maintaining the security of his MICBM force, i.e. tacks, while Player 2 may at any time release a first strike against these. Player 1 hides 200 tacks on squares of his choosing (one tack per square) and at any moment Player 2 may "attack" by declaring, all at once, 1,000 coordinates, (or for the larger first strike, 2,000 coordinates) among the 4,000 possible. Tacks on these squares are removed. After the attack the expected number of tacks remaining is 150 and the standard deviation of this distribution (so-called "hypergeometric") is computed in the Appendix (Point 1), to be 6.0 (in the second case one expects 100 tacks remain and the standard deviation is 6.9). If Player 1 is not restricted to placing only one tack per square, the distribution becomes Gaussian with mean still 150 tacks remaining, and standard deviation 6.9, (in the second case, 100 and 9.7 respectively). These are the results in the absence of the PQC protocol.

Now applying the protocol suppose Player 1 agrees, upon request, to label his rows, that is to tell Player 2 exactly how many tacks are on each row of 100 squares, but not on which squares. Suppose, for simplicity, Player 1 elects to put 5 tacks on each row and divulges this information to Player 2. There are two observations:

(1) In informing Player 2 that the deployment of tacks consists of exactly five on each row Player 1 has revealed information. This information and an honest response to a query facilitates verification but it cannot be used by Player 2 to increase the expected number of tacks he can remove by more

than 4.75 (see Appendix, Point 2). It is true that the revealed information allows Player 2 to adopt "low variance" strategies, for example, by sweeping 10 (or in the second case 20) complete rows, but he cannot materially increase his expected return. In any real world application, this possibility of variance reduction is certainly masked by larger uncertainties not incorporated into the model. Thus the agreement to label rows is verifiable without material compromise in security.

If the row labels are unequal, i.e. for some $N_i > 5$, the expected harvest of tacks is higher. This is for two reasons. First, a relatively larger number of locations can be discovered by query. And second, the attacker could saturate those rows with large labels, $N_i > 5$. The second effect, while quite strong in the model, would be less sharp in practice; different partition blocks would contain basing areas of different sizes and shapes; some would require more warheads to saturate than others. Nevertheless, there will be an incentive to adjust either the announced partition or the operations of MICBMs to produce a rough equality of missiles per block.

(2) Turning to the second question, Player 1's compliance with this agreement can be verified with rapidly (exponentially) increasing certainty if at intervals Player 1 is required to reveal a row of Player 2's choosing so that both can see if indeed only 5 tacks are on the row. We assume that Player 1 may redistribute his tacks after each inspection. If even one extra tack is present, there is a 50 % chance of finding it on the row revealed by Player 2's query after 27 random inspections (see Appendix, Point 3). If even as few as 10 extra tacks are hidden randomly in this model, only 8 inspections will reveal a violation with 90 % confidence.

Of course, finding illicit missiles will not be so easy as in this elementary model. Illicit missiles would be handled very carefully. They would not be randomly distributed, and would not be easy to locate within a challenged block. In fact, neither the hiding nor the seeking would be random since our queries would be guided by intelligence, which was ignored in the preceding model. Nevertheless, Soviet query responses will rapidly become a body of evidence from which inferences may be drawn. If our NTM spots even a single mobile missile per month and upon our query its coordinates are confirmed, then after a year we would begin to believe that either the Soviets are within about 10% of their ceiling (see Appendix, Point 4), are very good at hiding only an (illicit) subset of their missiles, or are very good at knowing just what

missiles we have and have not seen. The third possibility credits the Soviets with an unrealistic ability to penetrate or manipulate our NTM. The second possibility should be approached in a quantitative manner. If, for example, we thought the illicit missiles were 5 times harder to spot, after three years we would have 95% confidence that the percentage of illicit missiles was smaller than 30% (see Appendix, Point 5). After one year we would have 63% confidence. This considers the query only.

A challenge inspection always has the potential to turn up something unexpected. The area of a block might be extensive but the relevant linear measures, miles of road and rail, would not be great. Suppose that there is at least a 50% chance that an illicit missile will be discovered[13] during a challenge inspection if one or more illicit missiles is hidden within the block. If only 25% of the blocks contain hidden missiles the chances that at least one would be found during 12 random inspections is 80% (see Appendix, Point 6). Thus, query and challenge reinforce each other.

While elementary statistical models give only clues to the strengths and weaknesses of a verification scheme, they suggest that the partition-query-challenge format (PQC) would operate effectively within a complete verification regime for MICBM's.

Since there are many ideas of what verification means and how it might be accomplished, it is instructive to augment the issues of secrecy and certainty already discussed with a thumbnail list of other virtues and vices.

It is undersirable to establish verification procedures which:

1. are likely to generate false alarms,
2. are so invasive as to invite accusations of collateral spying,
3. are so demanding or complicated that accidental abrogation is likely, or
4. would be suspended in a crisis - increasing suspicions at precisely the wrong time.

[13]The Soviets would almost certainly prevent the literal "discovery" of an illegal missile by a ground team. But the required obstructive behavior would tell its own story. Thus the 50% should be treated as probability that the Soviets would fail to fool our inspectors.

The list of virtues is shorter: (1) transparency - it should not be difficult to understand how (or that) the protocols work, and (2) economy - the less cost the better.

PQC must be judged against these standards when applied to real MICBM deployments which might be selected by the U.S. and USSR, viz. carry-hard, rail-mobile, road mobile within reservations of various sizes. Whichever deployment scheme(s) the USSR chooses for basing its MICBMs, we must be convinced that illicit missiles cannot be surreptitiously shuffled out of sight, perhaps into an adjacent block, when we make a challenge. In our favor is that the challenger controls the time and place (and therefore the weather conditions) of the challenge. If they illegally overdeployed, the Soviets could have no confidence of winning such a shell game. Furthermore they would have to win every time to avoid detection.

Each type of deployment presents different practical problems for PQC since the partitions will have different geometries, and the number of, and the separations between, individual missile aim-points will vary. As illustrated by the above numerical examples, the conclusions are not very sensitive to such details. Furthermore, we have described PQC as a stand-alone protocol for counting MICBMs. In order to provide added confidence that no illicit MICBMs are being deployed, the PQC protocol could be supplemented by perimeter-portal monitoring of declared production sites, with all deployed missiles being tagged by one of the means discussed in Section 5. Also, special notification of transit could be required to move missiles from their production facilities to their deployment areas. We would practice watching these movements and as our skill increased so would our ability to detect any similar movements between any illegal production facilities and illegal deployment areas.

So far we have considered only verifying limits on the number of operational missiles, not the number of warheads carried on each one. Warhead numbers in the past have been determined primarily by counting the maximum number deployed during test flights. The current proposals to count the actual number deployed will require more intrusive and sophisticated measures, such as nuclear detection, or radiographic tomography as described in Section 6. This raises difficult issues of protecting bomb design secrets and of ensuring that the MIRV count cannot be substantially increased either

covertly or in a rapid "break out". Because of their suceptibility to detection by NTM, missiles on launchers are among the easiest objects to regulate with a PQC protocol. To regulate warheads with PQC would require much adaptation. The difficult problem of warhead limits was not futher pursued in this study.

4.2 Extension of PQC to Nationwide Partitions

We come to the second more extensive formulation of PQC. The Soviets (and symmetrically the U.S.) would be required to partition their entire national territory into M blocks (again M might be 40). Each block would contain a small piece of a missile reservation and a large chunk of other territory. A typical block B_i need not be connected. Large country-covering blocks would have ramifications beyond MICBM counting. A block query with the possibility of a follow-on challenge could probe any aspect of a START violation with more or less subtlety. Suppose, for example, a factory with a possibly illicit output came to our attention. We might query its block and watch for visible changes in its operations. If illicit activities are in progress, the query - and more so a challenge - would be very stressful to the Soviets since they would not know until the last moment if their factory had been identified or if the block had been chosen at random. As before, the indirectness of PQC protects intelligence sources.

Whatever the strengths of a more comprehensive PQC, these must be balanced against its liabilities. Inspectors roaming large territorial blocks represent a double edged problem. The inspectors must believe that nothing important to their work is being concealed while at the same time the host country must protect legitimate military and proprietory secrets. The potential for generating false issues and suspicions is great and the danger that these could lead to a breakdown of the treaty cannot be discounted. Verification is the foundation of START and as such should be possessed of the maximum stability. In particular, it should be capable of withstanding the collapse of Glasnost and Perestroika. If it is judged that an application of PQC limited to missile reservations is sufficient to achieve militarily effective verification, then this course is safer in the long run than more extensive and intrusive schemes. As a means of implementing the PQC protocol, overflights by aircraft, as recently proposed by President Bush (see Section 7) would be

useful—and particularly if nationwide application of PQC were judged to be necessary, such overflights would be an important supplement to NTM.

4.3 Summary

PQC engages the participants in an exchange of information for verifying limits on deployment of MICBMs. The PQC protocol relies on partitioning the deployment reservations and allowing for regular queries plus occasional challenges. It is designed to complement and make careful use of existing independent means of intelligence, to be minimally intrusive, and to ensure that any substantial violations would be at high risk of exposure without revealing targeting information that would compromise security of the MICBM force. The protocol can be applied exclusively to allowed deployment reservations or nationwide.

Appendix

Point 1

We use E to denote the average or expected value of a random variable. The variance is the expected value of the square of the variable minus its mean. Variance describes the "spread" of a distribution. Similarily, covariance describes the relatedness of two random variables. The standard deviation of a variable is the square root of its variance.

Suppose we have a total of N squares and m squares are marked with a tack. Suppose we make n "blind draws without replacement". Let T be the total number of tacks drawn. In our model $N = 4,000, m = 200$, and $n = 1,000$. T is the sum of n random variables, $T = \sum_1^n I_k$, where $I_k = 1$ if a tack is drawn on the k^{th} try and $I_k = 0$ otherwise.

$$\begin{aligned} E(T) &= \sum_1^n E(I_k) = \frac{nm}{N} \\ \text{var}(t) &= \sum_1^n \text{var}(I_k) + \sum_{j \neq k} \text{cov}(I_j, I_k) \end{aligned}$$

but

$$\text{var}(I_k) = \frac{m}{N}(1 - \frac{m}{N})$$

and

$$\begin{aligned} \text{cov}(Ij, Ik) &= E(I_j - E(I_j))(I_k - E(I_k)) = E(I_j I_k - I_j E(I_k) - I_k E(I_j) + E(I_j)E(I_k)) \\ &= E(I_j I_k) - E(I_j)E(I_k) = \text{Prob}(I_j = 1)\text{Prob}(I_k = 1 | I_j = 1) - \frac{m^2}{N^2} \\ &= \left(\frac{m}{N}\frac{m-1}{N-1}\right) - \frac{m^2}{N^2} \end{aligned}$$

So

$$\begin{aligned} \text{var}(T) &= n\frac{m}{N}\left[\left(1 - \frac{m}{N} + (n-1)\right)\left(\frac{m-1}{N-1} - \frac{m}{N}\right)\right] \\ &\approx \frac{nm}{N}\left(1 - \frac{m}{N}\right)\left(1 - \frac{n-1}{N-1}\right) \\ &\quad \text{"Gaussian part"} \quad \text{"nonreplacement part"} \end{aligned}$$

In our model the variance is $\text{var}(T) = 1,000(0.05)(0.95)(3,000/3,999)$. The standard deviation is $s = (\text{var}(T))^{1/2} \approx 6.0$.

Point 2

If the location of 5 tacks is known on a board with 4,000 squares, then the expected number of tacks removed by an attack covering 1,000 (or 2,000) squares is:

$$5 + 995\frac{195}{3,900} = 54.75 \qquad \left(5 + 1,995\frac{195}{3,900} = 104.75\right).$$

Knowing the location of 5 tacks increases expected number of tacks removed by 4.75. In this arithmetic, the denominators are 3,900 since the squares on the revealed row - except for the 5 with tacks - need not be attacked.

Point 3

If one extra tack is present the chance that a random inspection will not reveal it is $1 - 1/40 = 0.975$. Thus the chance of passing 27 consecutive inspections is $= (0.975)^{27} = 0.50$. For 10 extra tacks and 8 inspections the corresponding numbers are $1 - 10/40 = 0.75$ and $(0.75)^8 = 0.10$.

Point 4

In an undifferentiated population of $(1.1)N$ missiles , i.e. with an illicit population of 10%, the chances of a query response legitimizing an observed missile is $N/(1.1)N = 10/11 = 0.9091$. Since $(10/11)^{12} = 0.32$, the chance that twelve consecutive missiles would be legitimized is 32%. If the population were larger than $(1.1)N$, the corresponding probability would be even smaller.

Point 5

We assume that 30% of the MICBMs are part of an "illegal" population and that at least one missile a month is observed. The block containing that missile is queried. The chance that an observed missile is legal$=\frac{5\times 70}{5\times 70+30} \approx$ 0.921. The chance that over 36 months, the queried block will report only legal missiles is $(0.9210)^{36} \approx 0.05$; the chance is about 0.37 after one year.

Point 6

According to our assumptions, the chance of passing one challenge is $\leq (1 - (0.25)(0.5)) = 0.875$. The chance of passing 12 consecutive challenges is $= (0.875)^{12} \approx 0.20$.

Point 7

Finally technical comment is in order. In a purely mathematical sense, the

first player in our game will reveal approximately 8.4% of the total information needed to specify his deployment upon transmitting a labeling indicating that 5 tacks are present in each row. Information, in this sense, is measured as the minimal length of a number which could possibly code a given situation. As we have seen, this information by itself is of absolutely no use in targeting. However, it is not impossible that in the presence of some other, unspecified information, the labeling might have a slight (no more than and almost certainly much less than 8.4%) incremental effect on the expected return. For this reason it is reassuring to make the following calculation.

Using Sterling's formula, the $\log_{10}$ of the number of ways of placing 200 tacks in 4,000 squares is:

$$\begin{aligned} \log_{10}\frac{4,000!}{200!3,800!} &\approx \frac{4,000^{200}}{200!} \approx \log_{10}\frac{10^{720.4}}{10^{460}e^{-200}\sqrt{2\pi\cdot 200}} \\ &\approx 260.4 + 87.2 - 1.5 \approx 346 \end{aligned}$$

The $\log_{10}$ of the number of ways of placing 200 tacks, with 5 tacks on each row is:

$$\approx 40\left(\log_{10}\frac{100^5}{5!}\right) \approx 40 \times 7.92 \approx 317.$$

5 TAGS AND SEALS

5.1 Introduction

This section addresses possible applications and various technologies of tags and seals. When applied to nuclear SLCMs, the role of tags and seals is simply to verify limits on numbers, including a possible ban as discussed in Section 3. Other applications present more difficult problems that are not amenable to such a direct approach. One of these of current major interest at START is verifying that no mobile ICBMs are illegally deployed. This is a problem of determining numerical limits in designated deployment regions without compromising the uncertainties in the actual locations of mobile ICBMs, on which they depend for survivability, and at the same time verifying that none are illegally deployed in proscribed deployment areas. We describe a range of tagging concepts and technologies that might be employed in either or both of these applications.

A tag is a unique identifier, impossible to duplicate, associated with each treaty-limited item (TLI) of an arms control agreement. A seal is a mechanism for attaching a tag to the TLI in such a way that any attempt to remove the tag will alter it in a permanent way that is detectable when the tag is next read. It is not always necessary, or even desirable, to seal a tag to the TLI. A good name for an unsealed tag is "proximity tag". A proximity tag must be kept close to the TLI (just as your driver's license must be kept with you when you are driving) but it need not be permanently or even physically attached to the TLI. The proximity tag must be produced upon a legitimate "query" from the opposing side (as when the traffic cop asks for your license). The required proximity depends on the time allowed between the challenge and the verification, and this depends on the system being verified.

A tag can either be a "unique tag" or a "class tag". A unique tag, applied to a TLI, distinguishes that TLI from all others, even of the same type. It is a true "serial number". By contrast, a class tag simply identifies its TLI as one of a certain class of TLIs, for example SS-24's or Mobile MX's. There are arguments both for and against implementing each kind of tag. Unique tags furnish a greater degree of verification confidence, but they also reveal to the

inspector relatively more intrusive details of a TLI's logistic trail — where a particular TLI has been, how often it is moved, overhauled, etc. Below we will comment on which tag technologies lend themselves to unique versus class tags. In general, it is hard to have physical tags avoid being unique — even if unintentionally — since the inspecting side can make surreptitious note of small physical differences among tags. On the other hand, proximity tags (as defined above) are intrinsically nonunique (class) tags, since they validate whichever single TLI, of the designated class, they are in proximity to. Even if the proximity tags are uniquely identifiable, the tag holder may choose to have one less TLI than the number of tags. The floating of this loose, legitimate tag throughout the system will free the proximity tags for a vast shuffle, if the holder thinks such is desirable.

"Electronic tags" are physical tags based on the established technology of tamperproof microchips. By a combination of passivating and antietch coatings, the information stored in a chip can be rendered secure, even against sophisticated laboratory attack. The tag can be powered by batteries or else by induction fields only at the time of interrogation, and it is the size of a wristwatch. It can be designed for remote readout. All such verification tags for both the US and the Soviet Union could be physically identical and made to specifications openly shared and inspected. Electronic tags can use either a cryptographic algorithm, or else a one-time pad. An illustrative example of how they would operate with the latter is as follows: each tag would be loaded at the factory, by means of an interface similar to the IR remote control for a TV set, with two lists of 1000 random numbers, each 8 digits (bytes) long. In principle, the US would load twice as many chips as are required to attach to the Soviet TLIs and for each chip attached would keep an identical chip loaded in parallel with the same numbers and with the same serial number. After 1000 number pairs were loaded, the chip would switch automatically to another mode in which it would remain forever, in effect "burning the bridges" which made loading possible. In this verification mode, the chip would respond with one number in the second list when queried with the corresponding number in the first list. The tag would thus have a lifetime of 1000 queries.

A "virtual tag" is a set of procedures that can substitute for a physical tag. Recently several groups have begun to use the term Secure Registration System, or "SRS" in place of the term virtual tag. It can perform the same tasks required of the physical tag, without requiring that any object be placed

on or near the treaty-limited item (TLI). A virtual tag might consist of an encrypted text containing the identity and locations of a TLI—conceptually, a single line in a table of data. If one side (say the U.S.) requested that the Soviet Union demonstrate the compliance of a particular TLI at a particular site, the S.U. could comply by supplying the key that decrypts the portion of the text that contains the required information about that particular TLI, or about a TLI nearest to a specified location. Since each TLI would have its own key, decryption of the location of a particular TLI would not help allow the decryption of the locations of the other TLIs. Of course there must be strong safeguards that the encrypted virtual tag neither compromises a TLI's location (that is, before a challenge) nor can be made to validate more than one location (that is, after a challenge). We discuss these issues below.

A "seal" provides a means of ensuring that a tag remains attached to the TLI. In most cases, the seal is simply some kind of physical glue of a sort believed to be unremovable by surreptious means. There are established sophisticated technologies for seals, utilized by the diplomatic and intelligence communities. The purpose of the seal can be either simply to attach the tag, or also to insure that some component of the TLI itself has not been disturbed (e. g. that a weapons compartment has not been opened). In the latter role, a seal based on current technology might be a multilayer adhesive tape with a hologram (like those on credit cards) embedded in it, designed to tear apart if the tape is tampered, and perhaps also with a unique fluorescent signature (see below).

For electronic tags, fiber optic technologies might be utilized to make seals that are more highly tamper resistant. The underlying idea is to have a loop of fiber optics with both ends terminated on a tamperproof, powered, microchip. The chip sends coded interrogation pulses, one every few microseconds, through the loop of fiber. If the fiber is ever broken (even for a fraction of a millisecond), the chip permanently erases itself and powers off.

The fiber optic loop can, at installation, be threaded through any desired path; in effect, a hatch door can be "sewn" shut by the seal. Alternatively, a fiber-optic "stringbag" knotted with a gross mesh of several inches and with the fiber clad with a layer of plastic as protection against dirt, water, etc., could be a generally useful type of physical seal. The "stringbag" or purse could enclose a missile, or missile canister, with its mesh (several inches)

large enough to accommodate mounting or handling bolts for servicing, but sufficiently small to prevent changes of TLIs, such as the substitution of nuclear for conventional warheads in cruise missiles.

5.2 Attached Physical Tags

Many types of physical tags are under development, and have been briefed to us. We include a partial list below, with a short comment on each one:

- reflective particle tag (RPT), also known as "glitter paint", small flakes embedded in a plastic matrix. The most widely discussed of the tags, largely because it is inexpensive, and has received the most effort.
- scanning electron microscope (SEM) images of tags or TLI surface. Present technology cannot duplicate the sub-micron structure visible in these images. Portable SEMs are under development.
- holographic correlation. A holograph of a TLI surface is compared with the original; differences smaller than a wavelength of light can be discovered by the distortion of interference fringes. Its vulnerability is discussed below.
- subsurface ultrasonics, shows the structure of a seal in three dimensions. This method can also help assure a tag has not been removed.
- eddy current scanning, shows voltonic structure in 3-D. Three dimensions are more difficult to duplicate than two.
- geologic crystal acoustic microscopy; flaws in crystals cannot be duplicated with any known method. Sealing presents a problem because of the large size of the crystal; most of the crystal can be removed.
- fluorescent fingerprint. Ratios of spectral lines when illuminated under different wavelengths depend on the physical history of the tag as well as the chemical makeup. The *amount* as well as the spectrum must be measured to make sure material from one tag has not been shared.
- DNA signature.

- electronic tagging, possibly using cryptographic methods, discussed at greater length in what follows.

The goal of a tagging system is a tag that is too expensive to be duplicated or otherwise defeat. Rather than asserting that a "tag is impossible to duplicate" it is important to try to understand and assess the effort necessary to negate it. Once a tag is developed, it is important that a "red team" attempt to remove the tag from the surface, and/or see if they can duplicate the tag. A red team differs from the design team in that they receive credit and recognition by breaking the tag. They are thus fully motivated to apply their creative abilities to showing that a tag is insecure. This is essential because we must anticipate that another Party to the treaty will eventually have a red team attacking the tag system, and that it could be of substantial military value to them to defeat the tagging.

Red-teaming must be used to attack not only the physical tag itself, but also the procedures that are to be used to verify the tags. For example, if an instrument is brought into the field in order to provide a semi-automated check of a tag, it might be easier for the tag holder to alter the instrument to give a false positive reading (claiming a tag is acceptable when it is a phony). The holographic tag, for example, depends on the presence of "fringes" in the superimposed images; it might be possible to attack the imaging system to produce false fringes. We do not believe that simple protocols (such as keeping the instruments under the control of the inspecting team) can satisfactorily address this issue. As anyone who has studied the art of stage magic knows, protocols can give a deceptive sense of security, and an experienced magician knows how to "misdirect" so that the subject is entirely unaware of the fact that the protocols have been violated. The red-team must include a person skilled in misdirection to assess the security of semi-automated readout schemes — with the one exception of electronic tags, which do not involve instruments on-site.

So far the only tag of those listed above that has received substantial red-teaming is the reflective-particle tag (RPT). Red-teaming can continue even after a treaty is in place; the treaty should have a provision for replacing tagging systems which are demonstrably insecure.

Although electronic tags have not been extensively investigated, the methods of cryptography have received many billions of dollars of effort. A tremen-

dous amount is known about what it takes to break an encrypted message. Our national security already depends on the assessment that our highest level encryption methods are effectively unbreakable. As we discuss further in the section on "virtual tags", encryption for these tags is even more secure. Thus we think it is fair to say that encryption methods have received an enormous amount of "de-facto" red-teaming, much more than any of the other tagging methods listed. Because this enormous experience can be immediately brought to bear on the electronic tagging problem, because direct access to the tag need not involve physical contact, and because of the simplicity of the electronic tags, they are a very promising approach.

Since the goal of tagging is to make it prohibitively expensive for the tag holder to duplicate, it is best not to depend on any one technology but instead to use a combination. For example, one might use a "triad" of RPT, fluorescent fingerprint, and electronic tagging. A portable scanning electron microscope could be used to image the sub-micron surface, but this would be a last resort, if our other tags indicate something is wrong, but the tag holder insists they are in compliance. A scientific or technological breakthrough might compromise any one tagging method, but it is unlikely that several different methods, based on different physical principles, will all be broken.

Various "surface feature" tags have been proposed, but, over time, techniques for reproducing surface features on the sub-micron scale will be developed, facilitating the forgery of tags based on surface features.

Sub-surface acoustic techniques should be pursued vigorously as a countermeasure to undercutting and removal of physical tags. We recommend further Red Team experiments on undercutting.

5.3 "Proximity" Tags

Consider next the application of a "proximity system" in which the tag is not physically attached to the TLI. In this case what is required is to verify the proximity of the tag to one and only one TLI–hence the alternative name "inertial seal."

The key advantage of a proximity tag (apart from the fact that it is intrinsically a class tag–if this is deemed advantageous) is that it can be larger than an attached tag, say briefcase sized. As a result it can contain secure communicating and/or tamperproof motion-sensing electronics to support the challenge protocols.

Having a tag which can communicate to the outside world (only when the host country specifically enables it to do so, of course), has very significant verification advantages; it allows challenge verifications to be cheap, timely, and frequent. The in-country inspection team need not be transported to the site of a challenge; or (at its option) it can be transported to a small fraction of challenged sites.

For example a proximity tag could consist of a small module containing an electronic microchip and a UHF transponder. The treaty would require that the proximity tag, with its unique microchip, be kept within some distance (say tens or a hundred meters) of the TLI. The transponder could respond to queries simply via a whip antenna (no microwave dish is required). Because of the simplicity of the transponder, it is reasonable to require that, when a challenge is issued, the host country be required to power it up and ready it for remote interrogation within a few minutes. This quick response, of course, is not only possible but it is necessary, so that there is insufficient time to transport a proximity tag from a remote location. Quick compliance guarantees that the proximity tag is close to the missile, and makes sealing unnecessary.

The location of the transponder with its proximity tag (or proximity tags) could be determined by an orbiting system similar to Geostar, consisting of enough satellites that one or two of them is within receiving range. From the known position of the satellites, and the round-trip travel time of the signals, a computer on the ground could verify that the transceiver is at the specified location to the required accuracy of a few tens of meters. (Note that it is not necessary to locate the license, only to verify that it is roughly the expected distance from the satellite.) All the processing to determine position is done with a computer on the ground. Not many bits are required, so a very low power ground transceiver (a few watts) is sufficient. By doing pre-computation, the tag should be able to respond in nanoseconds.

The system could be implemented with just one or two satellites, or with piggyback systems on other satellites (as is done with Geostar). Full coverage is not necessary, because the U.S. gets to pick the time and place that it makes the queries, so it will make the queries when it has adequate coverage with the existing satellites. It would be impractical for the other party to try to move tags to match the satellite coverage.

An alternative to allowing the proximity tag to communicate is to give it a tamperproof motion sensor and clock. In that case, when the inspection team finally does arrive at the challenged site, they verify that the proximity tag was not moved to its present location after the time of the challenge.

5.4 Secure Registration System/Virtual Tags

Finally we consider Secure Registration Systems (SRS), i.e. a procedure that does not require that any object be placed on or near the TLI and yet performs the same tasks as a physical tag. The idea was introduced by Thomas Garwin (OTA report AAC-TR-10401/80; February 1980), who called it a "virtual tag." A virtual tag might consist of an encrypted text containing the identity and location of a TLI. If one side (say the U.S.) requested that the Soviet Union demonstrate the compliance of a particular TLI at a particular site, the S.U. would supply the key that decrypts the portion of the text that contains the required information about that particular TLI. Alternatively, the U.S. could request that the Soviets supply the key applicable to the TLI nearest to a specified location. Each TLI would have its own key, so decryption of the location of a particular TLI would not help allow the decryption of the locations of the other TLIs.

SRSs have the advantage that they are inexpensive, easy to implement, and based on a technology (encryption) that has been extensively studied and red-teamed. Because they do not require on-site inspection (except as an adjunct) they have many of the advantages of the proximity-system electronic tags.

There are at least three possible objections to SRSs that must be answered:

1. We would be handing over to the other side encrypted information on the location of all of our TLIs in a certain class. If they could break the encryption then they could use this information to target our TLIs.

2. There may be a way to cheat by letting the same encrypted message stand for two or more TLIs rather than just one.

3. To implement SRSs we must give the other side some knowledge of what encryption schemes we consider unbreakable, and such information could be useful to them.

Of course the national security of the U.S. already depends on the security of encryption, since many Top Secrets are protected only that way. Furthermore, several important features of the SRS scheme make it more secure than most other problems that use encryption, and which answer the objections listed above. They are:

1. Each message has its own key. The breaking of a single key would give the location of only a single missile, and would not help in the decryption of the locations of the other missiles.

2. There is no need to distribute the keys. In ordinary cryptography, the keys must be known by both the sender and the receiver, and thus there must be at least one copy for everybody that is communicating. Key distribution is a primary vulnerability for one-time pads. For the SRS problem, a key would never be distributed until a valid query was made by the other side, and then it could be sent openly. Until then, each key could be kept at the location at which it was generated. (One could arrange that not more than 10% of the keys were generated at any one location.) Since the keys would be kept only by those who already knew the locations of the TLIs, the existence of these keys does not increase the security risk.

3. The messages and the keys can be changed at frequent intervals, perhaps once every hour, or once every day. This is particularly easy to do since the keys do not have to be distributed, and the encrypted messages can be sent over completely open channels. The changes record updates in the position of the TLI, although they should also contain minor changes in the text (such as the time at which the message is

encrypted) that would change the ciphertext even if the location were unchanged. The information contained in the message only has value for a limited time, i.e. until the TLI is moved. Thus any scheme to break the encryption, if it is to have value to the enemy, would have to be accomplished in a time short compared to the time it takes to move a TLI.

4. The message could be encrypted with a non-invertible (one-way) procedure. This is similar to encryption, but it has the feature that there is no known way (other than message exhaustion) to invert it. In other words, given the key and the encrypted message, the message itself still cannot be recovered; however application of the key to the true message can be used to authenticate the validity of the encrypted message. We might require that Soviet messages be passed through an American encryption method (such as the DES) prior to the application of the Soviet non-invertible method; this would counter the fear that, in principle, the other side might deliberately devise a method whereby two potential missile locations could be encrypted (using different keys) to produce an identical output, thereby using each line of text to validate two missiles.

Actually, there is a continuous spectrum of possibilities between physical proximity tags (at one end of the spectrum) and SRSs (at the other). Start, for example, with the motion-sensing proximity tag described above. Now *instead* of incorporating a motion sensor, it is exactly equivalent to require (by treaty memo of understanding, or MOU) the following procedure: Whenever the host country moves the proximity tag to a new location, it must "tell" the tag its exact new location (for example in the geographical coordinates of the Global Positioning System or its Soviet counterpart). The tag, which has a tamperproof clock, records this information in secure fashion. How does the inspecting party know that the proximity tag was not moved? Simply by verifying, at the time of physical inspection, that the tag is in fact where it thinks it is–and that the location is timestamped prior to the challenge.

One quickly realizes that the location entered into the proximity tag can in fact be the location of the TLI validated by the proximity tag, not its own location. The inspection team then verifies that a challenged TLI's actual location was timestamped by that TLI's proximity tag. But now it makes no difference whether the tag is close to the TLI or not! All the proximity

tags could be at a single location, in Moscow (say), or (for the US) in SAC Headquarters in Omaha.

Imagine this room full of briefcase-sized proximity tags, one per TLI. What services are they in fact providing? Only these three: (1) they are "tokens" in one-to-one correspondence with TLIs; (2) they allow the respondent (host) country to change *and record* the location of his TLIs at will, without telling the challenging country that he has done so; and (3) they allow the challenging country to verify—but only after a challenge— that the location was in fact correctly, previously, recorded.

The final step to a SRS is the realization that these three functions can be performed without using any hardware at all, purely as a protocol that exchanges encrypted "digital signatures." The "tokens" corresponding to TLIs are simply a fixed number of entries into a data base maintained by the host country at a location of its own choosing. The host country enters the location of his TLIs (or an encrypted form of that location) into the database whenever he relocates a TLI. The challenging country receives regular "digital signatures" for each entry in the database. These do not allow him to know what is in the database, but they do allow him—subsequent to a challenge—to verify that the missile location was previously recorded.

Since the concept of SRS can take some getting used to, we will elaborate with some level of detail:

Suppose that there are N TLIs in some particular class (e.g. rail-mobile missile), numbered by $i = 1, 2, \ldots, N$. The tag for each TLI is a (unique, random) bit sequence T_i. The sequences T_i are public, and known to both parties.

The location of each TLI is expressed as a bit sequence L_i. Each party, in his role as respondent, maintains current knowledge of the locations of his TLIs; that is, he keeps up to date the N sequences (T_i, L_i). (For security reasons, he may not wish to keep all the sequences at any one location or accessible to any single individual; that is allowed.)

Whenever the respondent changes the location L_i of a TLI (or more often if he wishes to prevent "traffic analysis") he is required to telecommunicate to the challenger a "digital signature" of the new sequence (T_i, L_i). This digital signature does not enable the challenger, under any conceivable circumstance,

to determine L_i for himself; in fact, the set of L_is for mobile TLIs are the respondent's protected secrets of the highest order. The digital signature does allow the challenger to verify, later, after he has issued a challenge and the respondent has produced the "tag" in the form of the sequence (T_i, L_i), that the produced tag has not been altered from what it was when its digital signature was last sent (before the challenge).

The problem of generating digital signatures so secure that they can be trusted not to convey the information that they authenticate (i.e. not to reveal the location of our mobile missiles) is—perhaps surprisingly—a straightforward cryptological exercise. A possible general technique, is first, to encrypt the sequence (T_i, L_i) by an extremely secure encryption technique (e.g. one now certified for secrets of the highest sensitivity) and, second, to send as the signature only a small fraction of the encrypted bits. This ensures that even in the extremely unlikely event that an adversary were able to find a flaw in the encryption algorithm, he would still be lacking a valid ciphertext to work backward from. It is important to note that each TLI's location is encrypted with a *different* key. Multiple encryption is allowed. If desired, the list of TLI's can be divided into ten (say) sublists, with each list's digital signature generated by a different mathematical algorithm or combination of algorithms. Below, we give some additional technical discussion on the generation of digital signatures.

For definiteness, here are some specific scenarios under the protocol proposed here.

1. So as to avoid vulnerability to espionage, it is prudent that exact current locations of US mobile missiles will not be known centrally. Suppose that each Wing Commander knows the location of missiles under his command. Whenever a missile is repositioned, he forwards to NCA digital signatures of the new location, along with some additional validation bits verifying to NCA that the signatures have been correctly computed.

2. NCA removes the additional validation bits and adds an additional layer of encryption of its own, whose key is changed hourly. Once per hour, on the hour, NCA transmits the list of digital signatures, through treaty-agreed channels, to Soviet verification authorities. Since NCA's encryption key changes between each transmission, all signatures are

different on each transmission. This denies the Soviet side any knowledge of whether US missiles have or have not been relocated. Crypto keys are of course *not* transmitted. (A Wing Commander's keys are not even known to NCA.)

3. Every hour on the half hour (say), NCA or appropriate US authorities receive the Soviet list of digital signatures.

4. Suppose the US challenges a particular missile ("where is it?"), or geographic location ("what is the valid tag number of the missile there?"). The USSR must then provide an exact location, tag number, *and crypto key* corresponding to that missile's most recently transmitted signature. We use these quantities to compute a signature and to verify that it matches the transmitted one. On-site inspection may follow at our option up to a maximum number agreed to at the negotiating table. For this procedure to work it is crucial that, in practice, a single tag could not validate more than one missile and location (see section below). Since digital challenges are much cheaper than on site inspections, we can use many more of them, and achieve a much higher confidence in Soviet compliance.

5. If the USSR issues a valid challenge under the treaty, the US Wing command is queried by NCA to provide a specific location and key. These, along with the additional key added by NCA, are forwarded to Soviet authorities, who verify that the data match our previously transmitted signature and that we are treaty compliant.

5.4.1 Fuzzy Locations: An Additional Security Against Targeting

The encrypted message need not contain sufficient information for the Soviets to target our missiles, but only enough information for them to be able to verify treaty compliance. There are several ways to do this. Conceptually the easiest is to give only partial information about the location of the TLI in the encrypted message. For example we could give them just the latitude but not the longitude of the TLI. (We don't advise this one; it is just the simplest.) Another possibility is as follows: instead of giving the location of the TLI, the message would only state that the TLI is within a specified rectangle, with dimensions of 10 meters by 10 kilometers; this is useful for

verification, but too broad a region to target except with many warheads. Alternatively, the message could give ten well-separated small locations, but not say which one contains the TLI. An incompletely-specified location is called a "fuzzy location." If we ask the Soviets to validate the compliance of a TLI that we have spotted at a particular location, they must immediately (within a few minutes) deliver to us the key to decrypt the message that contains that location. We might then want to check the entire fuzzy region (using intrusive measures, such as on-site inspection, if necessary) to make sure that there were no other TLIs present. (If there were, their number should have been disclosed to us in the decrypted message.) Most of the time on-site inspection would not be necessary, since we get to pick the areas for verification. We could choose at our discretion an area that is entirely visible to NTM.

A second way to give the Soviets sufficient information to verify but not to target, is to deliver to them only a subset of the required bits. Instead of delivering just the key to the Soviets when they query a particular location or TLI, we would then deliver to them the missing bits and the key. (It is still necessary that the sum of the number of missing bits and bits in the key be less than the number of bits in the message that was sent, in order to assure the validity of the original message.) A location might be presented as 32 bits (16 each) for latitude and longitude; the eventual digital signature might be several times this length.

5.5 "Tell-Me-Your-Closest" Protocol

"Tell-Me-Your-Closest" is an example of a protocol that relies on a combination of physical tags and challenge inspections to verify limits on MICBMs. It differs from the SRS scheme of the preceding section which relies primarily on electronic data exchange via cryptographic keys rather than verifying authenticity of physical tags. In this sense "Tell-Me-Your-Closest" is another concept of MICBM verification that has both challenge and tags as essential physical elements. As in the PQC protocol it is designed to make effective use of independent intelligence information but in a manner that protects both sources and methods.

Under this scheme, the parties to the treaty agree as follows: (i) They agree on a common geodetic model, so that locations on the Earth, and distances between locations, can be specified to an accuracy of less than 10 m. (ii) Each agrees to maintain real-time knowledge of the exact location of *his own* TLI's. (iii) They agree on a number N (or fraction F) of TLI's whose locations will be revealed by the protocol to be described; and on a time T sufficient to relocate, and hide, a TLI of known location. Typical values might be $F = 10\%$, $T = 30$ days.

The challenge-and-response is as follows: The challenger, at a time of his choosing, specifies an exact location on Earth. This location can, *but need not*, be close to the location of a suspected TLI. The respondent must (i) immediately (e. g. within seconds or minutes) provide the location and tag number of the geodesically closest TLI to the challenged location — whether this location is close to, or is distant from (e. g. hundreds or thousands of kilometers) the specified point; (ii) Within a specified, prompt, time allow on-site verification (or remote interrogation of a proximity tag) to determine that the revealed TLI is a licensed one.

The declared TLI now becomes a part of the location-revealed fraction F. It remains so — and may not be relocated — until released by the challenger. Once the challenger has filled his full quota F, he can make new challenges only be releasing a TLI and waiting a time T (30 days) for its location to be deemed uncertain. The number of allowed challenges per unit time is readily calculated: for a force of 1000 tagged TLIs, with $F = 10\%$ and $T = 30$ days, about 3 challenges per day could be issued. Challenges will be viewed as routine occurrences, not as accusations of violation.

To see how this protocol works, we consider several different scenarios and their possible challenge strategies:

- Suppose that a general deployment area is known to the challenger, but he suspects that untagged, illegal missiles are mixed with the legal in that area. Over a period of days (say) the challenger issues repeated challenges of a single point, approximately centered in the missile field. As successive closest (tagged) missiles are revealed, the radius of the region susceptible to on-site inspection grows. When the challenger believes it to be large enough to contain an illegal missile, he invokes an on-site inspection. Since the challenger uses up more and more of

his pool of location-revealed missiles with successive challenges, he is deterred from extending the radius more than he needs to.

Arguably a side intending to cheat would remove any illegal missiles before the inspection team could arrive on-site. The protocol, however, forces the removal of all illegals at a time and geographic location of the challenger's choosing, without warning, repeatedly over time. Such surreptitious removals would be very expensive compared with the cost of a challenge (which might or might not be followed up by on-site inspection), and they could not reliably escape eventual detection by overhead surveillance and other means (including on-site evidence of recent missile siting). Cheating should thus be deterred.

- Suppose that by overhead surveillance, or HUMINT, a challenger knows the location of a particular illegal missile, but does not want, at that time, to reveal his knowledge of that location with any precision. He issues a sequence of challenges to "random" points, including one close to — but not closer than should occur by chance — the illegal missile's location. When the respondent fails to declare the illegal missile (by instead declaring a legal missile more distant from the challenge point), the challenger has begun to build a case that a violation exists.

  In fact, a graduated series of implicit messages can be sent to the respondent by returning and challenging nearby points — closer than could occur by chance 10% of the time, followed by closer than could occur by chance 1% of the time, etc. Deniability is maintained during such a series and specific capabilities are not revealed — such series could be conducted from time to time around points of no particular interest — but the respondent will be brought to appreciate that he is caught in a violation. On-site inspection could reasonably follow.

- Challenge points can deliberately be chosen to be as *distant* from any known missile deployment areas as possible. The missile declared in response will thus be quite far away, defining a large circle that is represented as having no missiles in it (tagged or untagged). The observation by any means of *any* missile within this area then becomes a treaty violation.

  In summary, a "Tell-Me-Your-Closest" protocol, provides a means for leveraging surveillance capabilities and (relatively expensive) on-site inspections: First, because even a modest degree of surveillance capability can be an effective deterrent against cheating — this is because

the respondent does not know whether any particular challenge location is backed up by a surveillance observation, or indeed whether the challenged point has any particular geometric relation to the point at which an illegal missile is suspected. Second, because some sequences of challenges will putatively precede an on-site inspection — thus forcing a cheater to respond to a (cheap) sequence of challenges with an (elaborate and potentially observable) relocation of illegal missiles, even when no on-site inspection subsequently takes place, and repeatedly.

5.6 Summary

The actual implementation of tags and seals would depend on what kind of TLI is being verified. They could be a useful supplement to PPM to confirm that nonnuclear SLCMs have not been converted to treaty limited nuclear SLCMs or that unauthorized numbers of nuclear SLCMs have not been deployed, as a result of unknown production sites or of in-service conversion of warhead types. Possible implementations of such a tag and seal procedure for verifying limits on nuclear SLCMs were discussed in Section 3. They can be used for verifying limits on MICBMs without providing essential targeting information that would compromise the survivability of the missile force. In this application virtual seals or physical seals joined with challenge inspections are applicable.

In designing a tagging and sealing scheme it is important to avoid operational and technical complexity and high costs while at the same time maintaining confident, tamper-proof, and non-overly intrusive procedures. This is not easy, but it is important to begin understanding the full potential of tags and seals which are likely to be integral components of future verification schemes for START.

6 RADIATION DETECTION

6.1 Introduction

Radiation monitoring will likely be an important element of any treaty limiting nuclear SLCMs. Penetrating radiation (neutrons or γ-rays with energies greater than 100 keV) can be used to determine that a SLCM warhead is nuclear or nonnuclear without disassembly of the weapon or damage to its electronics. Radiation monitoring can also be used to determine the number of warheads on a MIRVed ballistic missile. We do not consider that application here.

Radiation monitoring can be passive or active. In the former case, intrinsic radiation emitted by uranium or plutonium in the weapon is detected. In active detection, the warhead compartment is exposed to interrogating radiation generated by an external accelerator or radioactive source, enabling one to ascertain the presence or absence of a nuclear warhead by the peculiar response of special nuclear materials to the incident radiation. Radiation monitoring may or may not produce an image of the warhead and its surroundings, depending on the techniques used. As with all verification techniques, there exists some tension between the desire to positively identify a treaty limited item, and the reluctance to reveal technical information about such items. We discuss these issues further below after describing the various options for radiation detection.

In this discussion we address primarily the technology for determining whether or not a given SLCM is carrying a nuclear warhead. Frequent false negatives will undermine confidence in the verification procedures, while frequent false positives will cause more challenges than might be desirable. Several factors determine the utility of any particular radiation scheme for SLCM verification:

- Counting time required (a function of the sensitivity of detectors, signal strengths, and background levels)
- Access required (proximity to the warhead required for emplacing detectors, etc.)

- Ease of setting up necessary equipment
- Vulnerability to spoofing (concealment of a nuclear warhead)
- The degree to which design details of a nuclear warhead might be revealed.

6.2 Passive Detection

6.2.1 Detection of Plutonium

Passive detection of plutonium in a nuclear warhead is best done by searching for neutrons emitted by the isotope ^{240}Pu. This isotope has a spontaneous fission half-life of 1.3×10^{11} years and emits on the average 2.15 neutrons with each fission[14]. Five kilograms of Pu in a warhead thus generate $4.5 \times 10^6 \eta$ neutrons per second, where η is the isotopic fraction of ^{240}Pu. Although the fissile isotope of Pu is ^{239}Pu, ^{240}Pu is invariably present as a contaminant. Typical levels are $\eta > 4 \times 10^{-2}$, implying the production of 10^5 neutrons per second or more in a typical nuclear warhead.

The fission neutrons are emitted with characteristic energy $\sim$ 1 MeV, but are thermalized and attenuated in the material surrounding the warhead. About 10 % of the neutrons will escape the warhead.[1,2] As the natural background flux of thermal neutrons is 10^{-2}–10^{-3}/ cm^{2}sec, a reliable detection of a warhead with Pu can be made in $\sim$ 1 sec at a distance of $\sim$ 1 m. This has been demonstrated in measurements on actual warheads.

Passive neutron detection has the advantage of not revealing weapon design details, as the diffusion-like process the neutrons must undergo to escape the warhead blurs all but the coarsest geometrical information. However, it is possible to shield these neutrons. For example, depending upon the degree of moderation already provided by the high-explosive and other materials around the plutonium, several centimeters of a boron-loaded hydrogenous

[14]This is the most important radiation for passive detection. Gamma-rays are also emitted by plutonium but suffer considerable self-attenuation if the bomb of plutonium contains a significant amount of depleted uranium because of their relatively soft energies. Table 6-1 gives the predominant emissions for uranium and plutonium isotopes. Linear absoprtion coefficients for the gamma-ray lines are listed in Table 6-2. The attenuation length for gamma-rays is the reciprocal of the linear absorption coefficient.

Table 6-1

PRINCIPAL RADIATIONS FROM URANIUM AND PLUTONIUM ISOTOPES

	^{235}U	^{238}U	^{239}Pu	^{240}Pu
Gamma-Ray energy in keV (intensity in γ/s per gram of isotope)	144 (7.8×10^3) 163 (3.7×10^3) 186 (4.3×10^4) 202 (8.0×10^2) 205 (4.10×10^3)	743 (7.1) 766 (2.6) 786 (4.3) 1001 (7.5)	129 (1.4×10^5) 375 (3.6×10^4) 414 (3.5×10^4)	160 (3.4×10^4) 642 (1.1×10^3)
Neutron fission spectrum (intensity) in n/s per gram of iostope	1.1×10^{-5}	1.4×10^{-2}	2.3×10^{-2}	9.9×10^2
Other Radiations	Bremsstrahlung	Bremsstrahlung	Bremsstrahlung	Bremsstrahlung

Table 6-2

GAMMA RAY LINEAR ATTENUATION COEFFICIENTS

Gamma Ray		Absorbing Material	
Energy in keV	(Source Isotope)	High Z Mat. (typically U)	High Explosive (C,N,H,O,)
141	(^{239}Pu)	$\mu = 60$ cm^{-1}	$\mu = 0.27$ cm^{-1}
186	(^{235}U)	32	0.24
375	(^{239}Pu)	5.9	0.19
414	(^{239}Pu)	4.9	0.18
642	(^{240}Pu)	2.4	0.15
766	(^{238}U)	1.9	0.14
1,001	(^{238}U)	1.4	0.12

material could be used to attenuate the neutrons significantly. It might also be possible to avoid detection of a warhead containing plutonium by using Pu depleted in ^{240}Pu, although great effort would be required to reduce η below a few percent. Finally, it is also possible to utilize only highly enriched uranium in the fission stage, thereby suppressing a neutron signature completely although at a cost in weapon efficiency as measured by the yield-to-weight ratio.

6.2.2 Detection of Uranium

For uranium detection the best passive technique is to measure gamma-ray emissions[15]. One wants to focus on the highest energy emissions since these are least affected by self-shielding in the primary and are best able to penetrate the material surrounding it. The fissile isotope is ^{235}U but it emits no high energy gamma-rays (Table 6-1). In contrast the isotope ^{238}U emits 7.5 hard gammas (1 MeV) per gram per second. Both enriched (^{238}U $\sim$ 7%) and depleted (^{238}U $\sim$ 100%) uranium are used in nuclear devices, generating 0.5 and 7 gammas per gram, respectively (Table 6-2). The absorption lengths of these gamma-rays in iron, lead, and uranium are 2.3 cm, 1.4 cm, and 0.84 cm, respectively. Simple models[1,2] of warheads predict emergent fluxes of roughly 10^6 gamma-rays per second, a level that can be detected easily using either scintillation or semiconductor detectors.

Shielding of 1 MeV gamma-rays is difficult in view of the ranges given above and the required attenuations of 10^{6-7} needed to force counting times beyond 1 minute. Another way to suppress the gamma-ray signal from ^{238}U is to use another material, viz tungsten or lead, instead of ^{238}U in the weapon where high density material is required.

A potential problem of exploiting the 1 MeV signal to detect nuclear warheads that contain ^{238}U is the false positives that will result from non-nuclear weapon related depleted uranium in environments where verification is likely to be carried out. For example, the bullets of ship air defense systems often employ depleted ^{238}U because of its high density.

The combination of neutron and gamma-ray detection is an attractive possibility for passive monitoring since most modern warheads contain both

[15]The terms gamma-rays and X-rays are used interchangably here.

plutonium and uranium. A dual system would be less vulnerable to spoofing than either method separately.

6.3 Active Techniques

Active techniques offer greater resistance to spoofing, although with the penalties of greater complexity, cost, and radiation concerns associated with the active source. If used to detect the presence of nuclear warheads they can raise concerns about revealing details of warhead design; there are no such concerns if active techniques are used only to confirm the absence of nuclear warheads. Gamma-ray transmission (radiography), gamma-neutron threshold analysis, neutron transmission, and photon or neutron interrogation are all possibilities. Research and development on all of these are being pursued at various DoE laboratories, as shown in Table 6-3. Active techniques can be employed to produce a low-resolution image of the primary of a nuclear warhead or to induce radioactive emissions from the warhead. A detailed presentation of each of these methods is not possible here. We will focus our discussion on two particular techniques that already appear promising: (1) transmission radiography and (2) the detection of delayed fission gammas following photofission.

6.3.1 Transmission Radiography

General Considerations

Transmission radiography of cruise missiles at gamma-ray energies provides a potentially sensitive and selective way of distinguishing between missiles with nuclear and nonnuclear warheads. High-Z materials typical of a nuclear warhead can be detected by the added attenuation they cause, relative to the lower-Z components of conventional explosives. Radiography does not reveal the presence of plutonium or uranium specifically but does demonstrate unambigiously the presence or *absence* of high Z -absorbing material.

To provide more refined information, radiographs could be performed at two gamma-ray energies; with the measurements being taken one energy at a time. Such data would allow a rough characterization of the absorbing

Table 6-3.

Laboratory	Radiation Technology
Argonne National Laboratory:	• Hodoscope – Gamma Transmission Hodoscope – Neutron-reaction Hodoscope – Associated-particle Hodoscope – Californium-correlation Hodoscope
Idaho National Engineering Laboratory:	• Fission Assay Tomography System • Gamma-neutron Threshold Technique
Los Alamos National Laboratory:	• 14.7 MeV Neutron Transmission Radiography • Associated Particle Technique for One-sided Imaging • Small Radiation Detecting Instruments – Hand-held Gamma-ray Verification Instrument – Hand-held Gamma-ray Instrument with extendable boom – Hand-held Neutron Verification Instrument – Portable Unattended Neutron Monitoring System based on IAEA Reactor Power Monitor – Portable Neutron Briefcase • Neutron Source Imaging Detector • SNM Identification Case • Automated, Unmanned Portal Radiation Monitor • Fourier Transform Camera for Gamma-ray Imaging • Unintrusive Verification of Specific Nuclear Weapon Systems
Lawrence Livermore National Laboratory:	• Induced X-ray Fluorescence • Gamma-ray Telescope • Neutron Interrogation
Oak Ridge National Laboratory:	• Nuclear Weapon Identification System
Battelle Pacific Northwest Laboratory:	• Optically Stimulated Luminescence • Ne-213 System • Portable Passive Neutron and Gamma Imaging System • PNL-Directed Neutron Sensor
Sandia National Laboratory:	• Passive Gamma-ray Sensor Assessment • INF Passive Fast Neutron Detector Upgrade • Simultaneous Detection of Neutrons in a Gamma-ray Spectrometer • Encrypted Verification Scheme • High Z XRF Unique Response Detector • Compact Neutron Spectrometer • Low Energy Neutron Interrogation for Warhead Discrimination

Also, technologies being developed for DoD and other existing technologies, as well as new proposals from DOE, are being addressed by the RDP-AVT Panel.

material on the cruise missile, e.g., one could tell whether an absorbing region contained a large thickness of a low Z material (Al or Fe for example) or a small thickness of a high Z material (Pb, U, etc.). This distinction is possible because the gamma-ray absorption coefficients of high Z and low Z materials have a different dependence on gamma-ray energies in the 0.1–1 MeV region.[16] To exploit this discriminant best, one would like to choose the lower of the two gamma energies to be well below 1 MeV. However, radiography at such lower energies would require a stronger source, a potentially serious drawback.[17] Of course if transmission radiography were being used only to confirm absence of special nuclear materials by the absence of high absorption there would be no need for two energies.

Spatial Resolution

If transmission radiography is permitted for confirming the presence of a nuclear warhead[18] in *nuclear* SLCMs, it will be necessary to constrain its spatial resolution so that sensitive weapons design information is not revealed. This appears to be feasible. If only *non-nuclear* SLCMs are to be radiographed, resolution is not a problem. In fact high resolution can be useful in verifying non convertability as we discuss later in this section.

Consider the conventional approach to radiography. Here we would illuminate a SLCM from the side using a gamma-ray source, and perform a scan and obtain a radiographic image using an imaging detector on the opposite side of the SLCM warhead. In this approach, there would be the potential of revealing sensitive design information about the warhead. To prevent this, one could degrade the image by deliberately blurring or defocusing the radiography, using a variety of standard technical means. However, if the *short-wavelength* information was recorded in blurred form, there would be the possibility that a high-resolution image might be reconstructed using

[16]This difference in energy dependence can be traced largely to the rapid falloff of absorption above the characteristic K-edges that are strong features for high Z materials in the region near 100 keV. (See Figures 6-1 and 6-2.)

[17]One might consider doing transmission radiography with gamma-ray energies chosen to bracket the K-edge of a particular element, say uranium, so that a more refined Z determination could be made. This does not look attractive, however, because the extremely high absorption coefficients in the K-edge regions of high Z materials would force one to extremely intense sources.

[18]This also applies to counting the number of deployed warheads on MIRVed ballistic missiles.

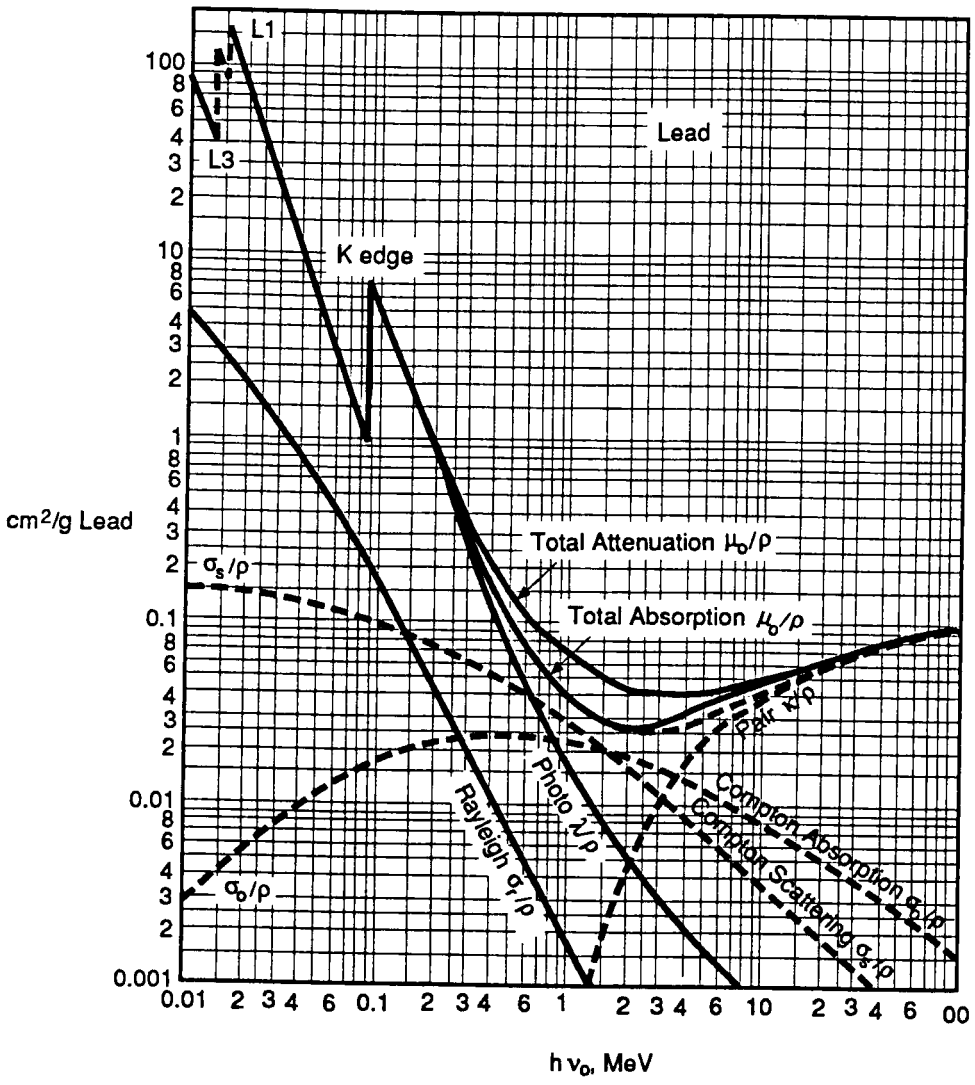
L1
L3
Lead
K edge
Total Attenuation μ_0/ρ
Total Absorption μ_a/ρ
σ_s/ρ
σ_a/ρ
Pair κ/ρ
Photo τ/ρ
Rayleigh σ_r/ρ
Compton Absorption σ_{ca}/ρ
Compton Scattering σ_{cs}/ρ
cm²/g Lead
100
10
1
0.1
0.01
0.001
0.01 2 3 4 6 0.1 2 3 4 6 1 2 3 4 6 10 2 3 4 6 00
$h\nu_0$, MeV

Figure 6-1.

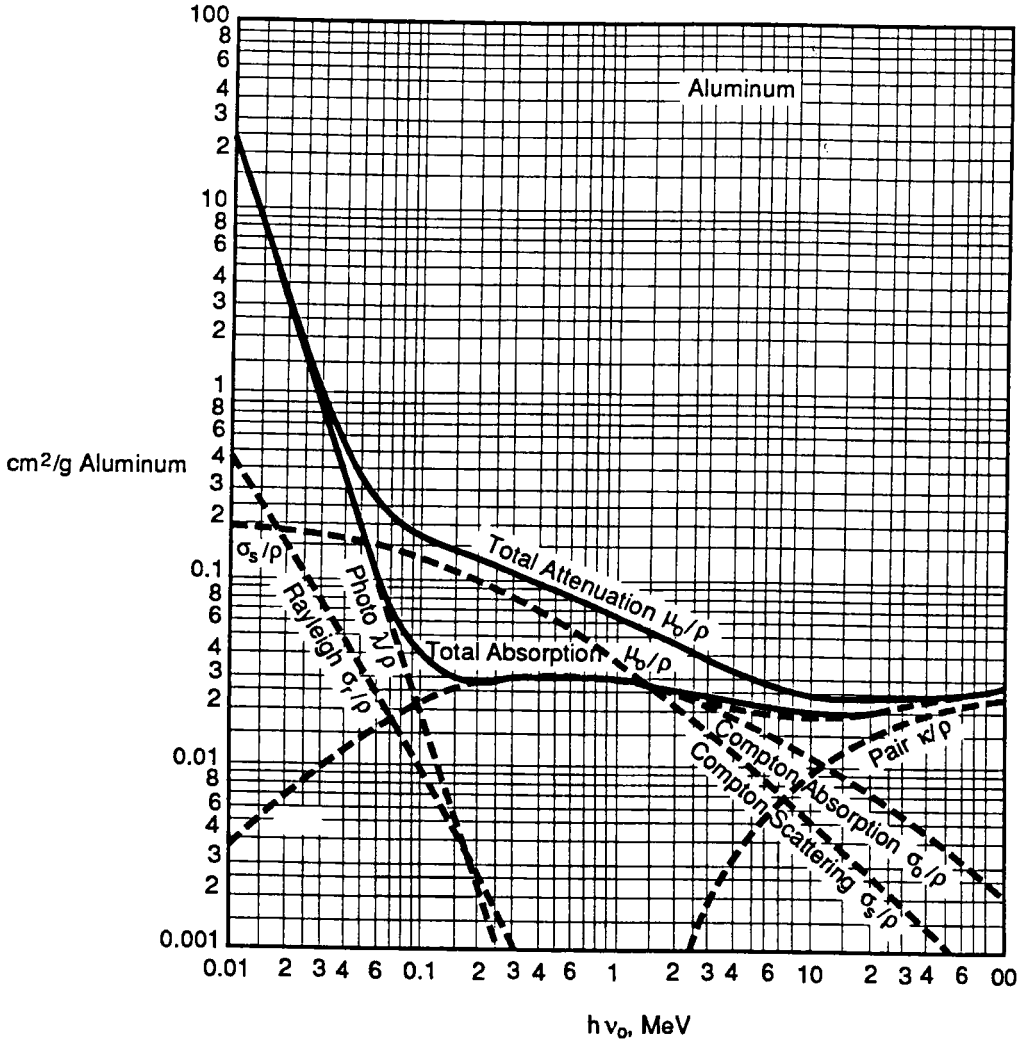
Aluminum
cm²/g Aluminum
Total Attenuation μ_a/ρ
Total Absorption μ_b/ρ
σ_s/ρ
Photo τ/ρ
Rayleigh σ_r/ρ
Pair κ/ρ
Compton Absorption σ_b/ρ
Compton Scattering σ_s/ρ
100
8
6
4
3
2
10
8
6
4
3
2
1
8
6
4
3
2
0.1
8
6
4
3
2
0.01
8
6
4
3
2
0.001
0.01 2 3 4 6 0.1 2 3 4 6 1 2 3 4 6 10 2 3 4 6 00
$h\nu_o$, MeV

Figure 6-2.

image-reconstruction algorithms if the original image had enough signal-to-noise ratio. If this risk is judged unacceptable an alternative approach can be used.

An alternative approach would be to not permit imaging of the entire warhead. One might use a small, well-collimated transmitting beam together with a receiver consisting of only a single, unsegmented detector (one pixel). The source-receiver pair for this system would be restricted to attenuation measurements on a coarse spatial grid; the system would *not* have imaging capabilities.

One would take attenuation measurements on a relatively coarse grid of discrete paths across the airframe of the cruise missile, and perhaps also along the airframe axis form front to back. The spacing of the grid points could be agreed upon ahead of time by the two parties, and monitored during the inspection and enforced by mechanical means. The optimum grid layout would guarantee that one or two lines of sight would pass through the warhead compartment but that any two adjacent lines of sight would be far enough apart that detailed structural information about the warhead would not be revealed. Additionally, it would be necessary that the registering of the grid not shift from one missile of the same type to the next. The concern is that if the grid were permitted to shift up or down slightly from one missile to the next, one might be able to reconstruct the overall internal configuration of a warhead by combining radiographs of many missiles of the same type.

6.3.2 Radiographic Scanning and Photofission

The detection of nuclear warheads by their own spontaneous gamma-ray emission was mentioned above. Here we discuss late time ($\gtrsim$ 1 min) detection via delayed gamma-ray emission due to the photofission of uranium, plutonium or thorium. Fission would be induced using bremsstrahlung gamma-rays from an electron linear accelerator machine of the type that is used to image cartons as agreed to in the INF Treaty.

The photofission cross-section for U, Pu, and Th has a threshold between 5.0 and 5.5 MeV and is largely due to the giant dipole resonance. The cross-section for ^{238}U has a giant resonance shape with a peak of 125 millibarns at

14 MeV and a full-width at half maximum of 8.8 MeV. The cross-section for the other two elements is similar.

The imaging system agreed to by the US and the Soviets for the INF treaty will be used as a model in our calculations. A transporter would move a container containing a SLCM through the imager in some 10 to 20 seconds and then take it out of the imaging structure in a minute or so. The detection times of interest occur after the background caused by the pulsing of the radiographic electron linear accelerator is over.

The basic function of the radiographic scanner would be to identify the contents of a shipping container. Should the container have nuclear material, uranium, plutonium, or thorium, the photons from the imaging system will induce photofission and the resultant delayed fission gamma-ray and neutron emissions will be detectable.

We will estimate here the rate of photofission-production by a Varian LINATRON 3000, a 9 MeV accelerator producing a radiation flow that is described by the manufacturer as "in excess of 3000 $R \cdot \mathrm{m}^2 \cdot \mathrm{min}^{-1}$ at its maximum output rate." (Here R stands for 1 roentgen = 1 rad = 10^{-2} J/kg = 100 ergs/gm.) The accelerator produces a gamma-ray beam that fills a cone with half cone angle of 14°; at a distance of 7 m it covers an area of 10 × 10 square feet = 9.3 m^2. Thus the "potency" of the linac is

$$\Gamma = \frac{3000R \cdot \mathrm{m}^2 \cdot \mathrm{min}^{-1}}{9.3\ \mathrm{m}^2 \times 60\ \mathrm{sec/min}} = 5.4\ \mathrm{rad/sec.} \qquad (6-1)$$

Since the energy spectrum of bremsstrahlung is approximately flat up to the maximum photon energy E_o, the gamma-ray number flux has the form

$$dN = C_o \frac{dE}{E},\ E < E_o \qquad (6-2)$$

where C_o is a constant. This gives for the energy deposition rate or potency, Γ,

$$\Gamma = \int_o^{E_o} EC_o\mu_m \frac{dE}{E} \qquad (6-3)$$

where μ_m is the mass attenuation coefficient, about 0.05 cm^2/gm in heavy metals such as lead or uranium. Thus,

$$\Gamma = E_o C_o \mu_m, \qquad (6\text{-}4)$$

$$= C_o \times 9\ \text{MeV} \times 1.6 \times 10^{-6} \frac{\text{ergs}}{\text{MeV}} \times \frac{0.05\ \text{cm}^2}{\text{gm}} \times \frac{\text{gm rad}}{10^2\ \text{ergs}}$$

$$\text{so } C_o = 3 \times 10^9 \text{photons/cm}^2\text{sec} .$$

To estimate the rate of photofission assume an area of fissile material of cross-section $A = 300$ cm^2, which gives a total photon fluence of 9×10^{11} photons/sec impinging on the warhead. We next calculate the total photofission cross secton for these incident photons. As discussed previously, the photofission cross-section has a peak of 125 mb at 14 MeV and a threshold of about 5 MeV. We approximate the cross-section simply by

$$\sigma_{pf} = 1.25 \times 10^{-25} \left(\frac{E-5}{9}\right) \text{cm}^2; \quad 5 \leq E \leq 14\text{MeV} \qquad (6-5)$$

where E is the photon energy in MeV. Integrating from 5 to 9 MeV we have then

$$\sum \equiv \int_5^9 \sigma_{pf} \frac{dE}{E} = 1.5 \times 10^{-26}\ \text{cm}^2 . \qquad (6-6)$$

Denoting the atomic weight of the target (uranium) by $M = 238$ gms, we have for the rate of photofission (N_o = avagadros number)

$$F = \left(\frac{AN_o}{\mu_m M}\right) C_o \sum$$

$$= \left(\frac{300\ \text{cm}^2 \times 6.02 \times 10^{23}}{0.05\ \frac{\text{cm}^2}{\text{gm}} \times 238\ \text{gm}}\right) \times 3 \times 10^9\ \frac{\text{photons}}{\text{cm}^2\text{sec}} \times 1.5 \times 10^{-26}\ \text{cm}^2$$

$$= 6.8 \times 10^8\ \frac{\text{fission}}{\text{sec}} . \qquad (6\text{-}7)$$

The transporter moves at nominal rate of 0.7 inches per second so a warhead is exposed for about 17 seconds in the imager giving a total $N_f = 1 \times 10^{10}$ fissions.

As a result of the induced photofission, the decaying fission products will emit delayed gamma-rays and neutrons. We now estimate the delayed gamma-ray counting rate. The delayed gamma rays represent about 3%, or 6 MeV, of the total fission energy. This corresponds roughly to $n_\gamma = 6$ delayed gamma-rays since the mean energy of the delayed gammas is about 1 MeV. Let $q(t)$ be the disintegration rate of the fission products per fission ($\int_o^\infty q(t)dt = 1$). $q(t)$ is approximately

$$q(t) = \frac{1}{6}; \quad 0 < t < 1\ \text{sec} \qquad (6\text{-}8)$$

$$= \frac{1}{6} t^{-1.2}; \quad t > 1\ \text{sec.}$$

Thus, the delayed gamma-ray emission rate would be

$$\dot{C} = n_\gamma N_f q(t) \tag{6-9}$$

$$= 10^{10} \times \frac{1}{t^{1.2}} \text{ sec}^{-1} \quad t > 1 \text{ sec.} \tag{6-10}$$

The delayed gamma-ray emission rate at one minute after imaging would be of the order 7×10^7 gammas/sec. This estimate neglects the self-shielding effects and the attenuation due to any intervening material. Fetter, *et al*[1], estimate that the fraction of the delayed gammas leaving the jacket of three variants of nuclear weapons designs that employ ^{238}U varies from 0.06 to 0.15. Thus, the delayed gamma-ray emission rate could vary from 4.2×10^6 to 1.1×10^7 gammas/sec. If there were additional shielding the number of delayed gammas reaching the outside would be further reduced. These levels of source strengths are detectable by standard gamma-ray counters.

There are also delayed neutrons which are emitted following the fission process. They are not as robust a signature as the delayed gammas. The delayed neutrons comprise about 1% of the total fission energy, and are in the 1 MeV range. They are partially shielded by the high explosive surrounding an actual primary and can be further reduced by adding a boron shield. Shielding reduces the energy of escaping neutrons to the order of kilovolts[2].

6.3.3 Alternative Sources

There is a question of radiography using a ^{60}C0 source which has strong gamma-ray lines at 1.17 and 1.33 MeV with a half-life of 5.21 years. The use of cobalt raises safety questions, particularly in plants which deal with high explosives. In addition, ^{60}C0 gammas are well below the photofission threshold and thus are of no interest for this technique.

We should finally remark that if we use a more energetic linear accelerator that covers the full giant resonance, say a 20 MeV machine, the photofission yield would be up by a factor of about 8.

6.4 Radiation Detection in a Treaty Context

The utility of radiation detection (passive or active) for SLCM verification depends on the proposed site of the measurement, and on the treaty context in which it is being carried out. The comments in this subsection are addenda to the discussion of Section 3.

6.4.1 Transmission Radiography at the Point of SLCM Final Assembly

One could envision using transmission radiography to distinguish between nuclear and conventional SLCMs at a perimeter-portal monitoring station outside the facility where the SLCM airframe and warhead are joined together. This would make sense in the context of a treaty which limits the number of nuclear SLCMs, without necessarily limiting the number of conventional SLCMs. Each side would declare that a certain number of its SLCMs were nuclear. Those SLCMs which had been "declared" to be nuclear would then not need to be radiographed; they would be presumed to be nuclear and counted as such. One would want to verify that the remainder of the SLCMs, which had been "declared" to be conventional, were in fact not nuclear. Access to the SLCM would not be a difficulty, since the SLCM canister would not yet be installed in a launch tube, torpedo tube, or other launch structure. One could easily ascertain via an attenuation measurement that the warhead did not contain large amounts of high-Z materials. The spatial resolution of the radiography measurements could be as high as the verifying country desired, because if the warhead were indeed conventional as claimed, there would be no nuclear design information revealed by the radiography. If it were felt that too much structural information about the conventional airframe or warhead were being revealed, the spatial resolution could also be deliberately degraded by using a small collimated beam on a predetermined grid of sight lines, as described above.

6.4.2 Transmission Radiography Measurements Made Onboard Ships

Practical aspects of conducting transmission measurements onboard a ship are difficult. One example is the geometry of existing launch and storage configurations. In some cases, below-deck access for the inspectors might be required. In other cases, the only practical access available might be in the narrow space between the canister or SLCM and the launch tube; this would seriously constrain the instrumentation packaging. While these restrictions are not necessarily prohibitive, they do suggest that access to sensitive parts of a ship may be needed in some cases in order to obtain adequate transmission radiography measurements. The navies of the United States or the Soviet Union might argue strongly against this type of relatively intrusive onboard inspection. One might deal with this issue by selecting a few SLCMs randomly for radiography, relying on remote sensors as described in Section 3 to verify that the selected missile has not been switched. Those selected would then be pulled from their launchers for inspection elsewhere.

6.4.3 More Radiography

We would like to stress that radiography of *nuclear* SLCMs is not necessary for effective verification of a limit on the number of nuclear SLCMs. In a scenario in which there are inspections of SLCMs at the factory where they originate, it would be sufficient to radiograph or inspect only those missiles which have been declared to be *nonnuclear*, and to presume that all SLCMs which have been *declared* to be nuclear are in fact nuclear. Once this has been done, all SLCMs could be given identical tags to certify that they had been counted by the inspectors. If the tags were tamperproof, then subsequent inspections at choke points and servicing installations as described in Section 3 would need only confirm that all SLCMs were tagged and sealed. As a consequence, transmission radiography of nuclear SLCMs would not be necessary, either onboard ship or at the factory of origination.

An important consideration for any attenuation measurement to detect nuclear warheads is the possible presence of depleted uranium in conventional

SLCM warheads. Since depleted uranium would also produce high attenuation, it will be important to ascertain via data exchanges whether conventional US or Soviet SLCMs contain depleted uranium in a geometry or quantity that would make attenuation-based measurement ambiguous.

6.5 Radiography of Struts

Finally we discuss in this subsection a simple application of radiography using a weak source to verify that the internal structure of a conventional SLCM precluded its conversion to a nuclear SLCM (see earlier discussion of this matter in Section 2.) Let us suppose we seek to verify the presence of a 1 cm-wide strut inside a cylindrical cruise missile 50 cm in diameter. We will use radiography with an array of 1 cm × 1 cm detectors 1 meter from a source of 1 MeV photons (e.g., ^{60}C0), which have an attenuation length of 2.3 cm in iron and 6.7 cm in aluminum. This is illustrated schematically in Figure 6-3.

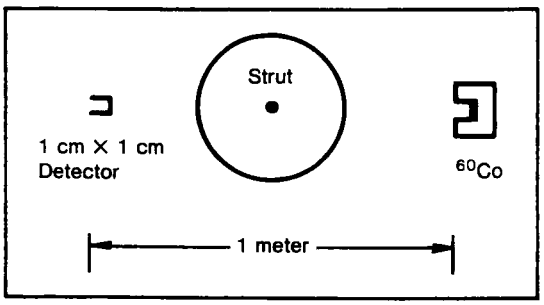

Figure 6-3.

We assume that the aeroshell, canister, etc. are equivalent to 1 cm of Al, so that the "baseline" attenuation is $e^{-2\text{ cm}/6.7\text{ cm}} \approx 0.74$. The presence of a strut 1 cm thick introduces an additional attenuation of $e^{-1\text{ cm}/6.7\text{ cm}} \approx 0.86$. To detect this difference with confidence requires some 1000 baseline counts (860 counts when the strut is present). Thus, if the measurement time is t, the baseline count rate is $1000/t$. Taking into account the baseline attenuation (0.74) the solid angle subtended by the detector ($1\text{ cm}^2/4\pi \cdot 1\text{ m}^2 \approx 8 \times 10^{-6}$), and the efficiency of the detector ($\sim$ 10%) we need a source strength of $(2 \times 10^9)/(\text{t (sec)}) \tilde{-} (0.05\text{ Ci})/(\text{t (sec)})$.

Finally, we can set t by requiring that the entire SLCM of length ~ 6 m be scanned in 5 minutes. This implies that each cm of length is exposed to $t = 0.5$ sec of radiation. Thus, if a 50 element line array of detectors is used, the interior can be imaged with 0.10 Ci of ^{60}C0. We note that such a source is quite feeble relative to the kilo-Curie sources used routinely in hospitals and so should present no health hazard if routine radiation safety procedures are followed.

6.6 Summary

Active and passive means of radiation detection can be used to determine the presence or absence of nuclear bombs. In particular counting delayed gammas resulting from photofission induced in uranium or plutonium is a sensitive signature to confirm the presence (or absence) of nuclear warheads in cruise missiles. Simple radiography using a ^{60}C0 source of only ≈ 0.1 Ci can detect the presence of internal structural elements in cruise missiles that are designed to make unobstructed regions sufficiently small that it would be impossible to arm them with existing small fission bombs.

REFERENCES

1. S. Fetter, V. A. Frolus, M. Miller, R. Mogley, O. F. Prilutskii, S. N. Rodmov, and R. Z. Sagdeev, Detection Nuclear Warheads, preprint of the FAS/Soviet Scientists Working Group on Verification, July 18, 1988 (Unclassified).

2. R. Z. Sagdeev, O. F. Prilutskii, and V. A. Frolov, Problems of Monitoring Sea-Based Cruise Missiles Bearing Nuclear Warheads, Report of the Committee of Soviet Scientists for Protecting the World from the Nuclear Threat, Key West Conference, February 1988.

7 SURVEILLANCE TECHNOLOGIES

7.1 Introduction

This chapter is devoted to a discussion of several ideas for applying new and newly developing technologies to extend the U.S. capabilities for overhead surveillance and thereby to improve the means of verification. The emphasis here is on near-term options. The concepts described will require detailed engineering and system analysis in order to evaluate them more fully.

We first describe a constellation of relatively small and inexpensive photo reconnaissance satellites in low earth orbit ($\sim$ 300-400 km altitudes). The primary attractiveness of such a system is two-fold: it enhances our capability for activity monitoring by making frequent overflights of all sites of potential interest; and it is more survivable against an anti-satellite threat since it presents many targets rather than one or only a few very high value ones. In order to keep the optical systems and satellites themselves relatively simple and light we settle for moderate ground resolution ($\sim$ 1 meter) imagery, which is adequate for many intelligence purposes, particularly with the emphasis on activity monitoring for treaty verification and other general needs.

A second proposal that we describe is to achieve longer dwell times over target and better survivability against primitive ASAT threats by deploying a few large optical observing platforms at high altitudes ($\geq$ 5,000 km). A specific implementation of this idea with a large long-focal length refracting lens telescope is described.

Next we discuss critical issues raised by the possibility of equipping surveillance satellites with lasers or radars to illuminate the ground. Finally, we review questions related to the recent proposal of President Bush for the U.S. and the Soviet Union to allow aircraft overflights to enhance surveillance. This resurrects the 1956 "open skies" proposal of President Eisenhower.

7.2 Small Satellite Reconnaissance Fleet

Recent technological developments leading to miniaturization of sensors and communications links and to reductions in required power levels have the potential to reduce the cost and size of essential components of reconnaissance satellites. As examples of such progress achieved in a number of programs, including in particular SDI, we mention: fiber-optic gyros with < 0.1 deg/hr drift, star trackers accurate to 100 μrad with a 60° field of view, on-board computing power in the range of 10-20 MIPS, laser diode arrays producing 5 to 10 Watts per array at 30% overall power efficiency and CCD arrays with of the order of 10^6 pixels of individual dimension 10 μm $\times$ 10 μm. These advances lead us to consider a constellation of relatively small and simple reconnaissance satellites that achieve medium ground resolution ($\sim$ 1 meter) from LEO ($\sim$ 300-400 km).

Atmospheric drag limits how low an altitude such a system can operate at economically. In the illustrative examples that we give in the following discussion we choose an altitude H = 300 km, which is consistent with a two-year lifetime for a satellite weighing $\approx$ 1000 pounds during periods of maximum sun spot activity.

The advantages of such a system for intelligence and verification include:

- its ability to provide frequent coverage for monitoring activities of high intelligence value or as required for verifying compliance with arms control treaties. It is of course not necessary for a satellite to photograph everything in its field of view on each overflight; the very fact that it presents the possibility of such coverage can immensely complicate, if not discourage, the scheduling of large-scale activities which are illicit or would provide evidence of high intelligence value.
- its robustness against an ASAT threat since it presents many targets rather than a very few, each of high value.

We begin with estimates of the number of satellites and of the total data transmission rate as a function of the frequency of overpasses and the fraction of available imagery returned. Next we discuss the basic requirements to be met by the optical telescope and the detectors. Finally we turn to the design

and cost of the other major components of the small satellite, including solar panel power for housekeeping and stabilization thrusters to make up drag, batteries, communication links (possibly laser), attitude control, etc.

7.2.1 Frequency of Coverage and Numbers of Satellites

Key issues are the total number N_{total} of satellites in the whole constellation and the fraction of the total imagery that can be transmitted.

In order to illustrate the idea with a specific set of numbers we choose H=300 km for the altitude of the satellite fleet in circular orbits and specify a ground resolution of 1 meter at 45° slant range, $\sqrt{2}$ H, for light of 0.6 μm wavelength. For the telescope operating at the diffraction limit this requires an aperture of diameter

$$D = \frac{1.22\lambda \times \sqrt{2}\mathrm{H}}{1\ \mathrm{m}} \approx 30\mathrm{cm}. \qquad (7-1)$$

This assumes that resolution, not light grasp, drives the design as is true for daytime viewing by modern CCD detectors.

The satellite constellation can be organized into planes, with each plane at fairly steep inclination. For instance the planes could define sun synchronous orbits although this is not necessary. For good coverage in clear weather at mid latitudes, the planes must be spaced about 2H apart at latitude 45°, so that satellites in adjacent planes can among themselves cover all the intervening territory out to their limiting slant ranges. Therefore the number of planes is

$$N_{\text{planes}} \approx 24 \left(\frac{300\ \mathrm{km}}{H}\right). \qquad (7-2)$$

If each plane contains N_{sat} satellites, then the revisit time is

$$T_{\text{revisit}} \approx \frac{(90\ \mathrm{min})}{N_{\text{sat}}} \left[1 + \frac{1}{13}\left(\frac{H}{300\ \mathrm{km}}\right)\right] \qquad (7-3)$$

and the total number of satellites in the whole constellation is (with $N_{\text{sat}} \geq 1$)

$$\begin{aligned} N_{\text{total}} &\approx N_{\text{planes}} \cdot N_{\text{sat}} \\ &\approx \frac{(90 \text{ min})}{T_{\text{revisit}}} \left[24 \left(\frac{300 \text{ km}}{H} \right) + 1.8 \right]. \end{aligned} \tag{7-4}$$

For example, at $H = 300$ km, 24 satellites (arranged with one in each of 24 orbital planes) would revisit each spot about every 100 minutes.

If we use fewer planes than specified in Equation (7-2) (or equivalently, $N_{\text{sat}} < 1$), then the average revisit rate is still given in Equation (7-3), with fluctuations that can be made reasonably small by judicious phasing of satellites in each orbital plane. Thus, at $H = 300$ km, 12 satellites could be arranged in 12 planes to revisit each spot about every 190 minutes.

Even if we make N_{total} so large as to permit total coverage imagery in real time, the data rate for such coverage would be too large to be handled in practice. Thus if we take the Soviet area of interest as $\approx 10^7$ km^2, this is equivalent to 10^{13} pixels at a resolution of 1 m. If we take "real time" to mean one image every 10 s, and if each pixel requires 10 bits, the data rate required for total imagery would be 10^4 Gbit/sec, far too high to be practical.

The system therefore has to task its satellites to return images only of selected areas at selected times. Cutting the total system data rate down to about 1 Gbit/sec (a conservative figure) would necessitate selected targeting of a small fraction of the area of interest, or about 100 km^2/sec. If this were broken down into 1000 frames in 10 seconds of size 1 km^2 each, for instance, this system would produce 100 frames/sec, with each frame comprising 10^7 bits.

This estimate neglects any reduction in data transmission resulting from application of image compression techniques, e.g., the method of vector quantization. VQ has been exploited by GlobeSat, Inc. for its proposed tactical imaging satellite with a resulting 12:1 reduction in the number of bits transmitted per image. In general, the amount of compression will depend on system performance criteria such as the number of bits per sample, the degree of reconstruction precision required relative to the initial quantization error in encoding the original image, etc.

7.2.2 Telescope and Detector Requirements for Small Satellite Coverage

Each satellite at height H carries a telescope designed to survey the field of view (area $\cong \pi H^2$) beneath it at any time. The telescope should have the capability of imaging areas located anywhere within this field of view. We outline here some generic characteristics and possible limitations of such a telescope.

Some portion of the image plane–small enough to avoid aberration–would be covered with CCD detectors onto which steerable viewing optics would guide images of a succession of selected areas. In one approach, an array of CCD's would map out a swath along the ground track of the moving satellite, while cross-track steering mirrors would move the swath from side to side.

In order to keep the data rate at a manageable level, the constellation would return images of only a small selected fraction of the entire viewing area at any one time. Sometimes a large fraction of an individual satellite's viewing area might be desired, depending upon what targets were momentarily within range. Other times, simply searching along a single road or railway might suffice.

First we estimate the dwell time on target required to secure an image. Assume a given resolution element on the earth is imaged into a single detector pixel. The size of the telescope aperture is determined by the size of the resolution element and the altitude of the satellite. At $H = 300$ km the lens diameter is 30 cm and the focal length of the telescope is ~ 4 meters in order to focus a 1 meter resolution element at slant range $\sqrt{2}$ H onto a 10 μm CCD. Using Equation (7-1), we find that an upward scattered intensity at the earth of I Watt/Area, uniformly directed into 2π steradians, produces a photon counting rate in each pixel of approximately

$$dn/dt \cong 0.9\epsilon I\lambda^2/2\pi h\nu = 0.9\epsilon I\lambda^3/2\pi hc \qquad (7-5)$$

to be averaged over the wavelength λ of the light, where ϵ is the detector efficiency. Note that Equation (7-5) is independent of the altitude and surface resolution. Assuming I is about one-tenth solar illumination to account for

average solar elevation and surface reflectivity, i.e. $I \cong 0.2$ kW/m^2 we find, with $< \lambda > \cong 0.6$ μm and $\epsilon = 0.6$,

$$dn/dt = 2 \times 10^7 \text{photons/sec per pixel.} \qquad (7-6)$$

To measure a pixel to pixel contrast of 3% would require 10^3 photons/pixel under shot noise conditions (good CCD's are approaching detector noise <10 counts/pixel, i.e. consistent with this level), requiring less than 10^{-4} sec dwell-time at each pixel. This dwell-time has a nice match to the typical ground speed ($\cong$ 7 km/sec) of a LEO satellite, which covers the desired resolution distance of 1 meter in 1.4×10^{-4} sec.

If there were no practical limitation on the steering speed of the optics, then this short dwell-time per pixel would allow the entire satellite field of view to be covered by a very modest area of CCD's. As an example, at an elevation of H = 300 km, there are about $\pi(300 \text{ km})^2/ 1 \text{ m}^2 \cong 3 \times 10^{11}$ resolution elements in the field of view. If we wished to cover this entire field in 30 seconds, with 10^{-4} sec to record each element in a detector pixel, we would require 10^6 pixels. With the 10 μm $\times$ 10μm CCD pixel size, this means a CCD area of 1 cm^2, readily accommodated with acceptable aberration in the image plane of the 30 cm diameter focusing lens needed at this altitude. Since the instantaneous field of view for such an array is $\approx (10^3 m)^2$, it would require moving a steering mirror of 30 cm diameter at angular velocities > 40 rad/sec in order to cover the entire satellite field of view. This is unrealistic. By contrast, to follow a single railway headed 45° relative to the ground track would require a steering angular velocity of only $\cong 10^{-2}$ rad/sec, well within reasonable capabilities.

Given these limitations on steering, together with the limitations on the rate of reconstructing data from successive frames, as discussed in Subsection 7.2.1, each satellite would usually be tasked to photograph a few fixed sites or one to two roads or railways within a given field of view on each overpass. One might also consider the possibility of equipping the satellite with a low-resolution viewer in order to provide input for precision aiming and image reconstruct with 1 m resolution. This may be particularly useful in the presence of cloud cover.

Another important consideration is the detector array. Detector technology has developed in two directions: (a) to narrow CCD arrays and (b) to CCD square arrays. The narrow arrays are used in a "push-broom" mode to

obtain images by accumulating time sequential data as the satellite passes over the target region. Pixels in the linear array are sampled rapidly to give high resolution in the line of flight, while resolution in the perpendicular direction is determined by the optical characteristics of the telescope and the size of the CCD elements. Square CCD arrays are most often operated in a snapshot or "staring" mode. Signals are collected by each pixel element of the array during a sample period. This period must be short enough (because of the motion of the surface field of view) to give the desired individual pixel resolution. To form a useful image, the push-broom system must be operated continuously, yielding a strip image of the terrain under surveillance. The staring system, in contrast, needs to repeat its signal-gathering only when a new field of view is present on the detector.

Various factors influence the choice of push-broom or staring modes for satellite imaging. Foremost among these is the question of image quality. Owing to external satellite drag and internal distortions arising from thermal effects, the surface field of view is continually distorted and varying. With push-broom systems, yaw of the satellite will give complex side to side motions of the field of view, greatly complicating the post-flight rectification of the data into a usable image. Satellite pitch (up and down of the longitudinal axis) gives backwards and forwards motions which are almost impossible to overcome in image rectification. For 1 meter resolution, changes of pointing angle on the order of 1 to 10 μradians will displace the pixel surface locations sufficiently to cause difficulty in image processing. Of course, it is the rate at which such displacements occur that will determine the actual sideways displacement of the image. Thus, there is a direct relationship between the satellite attitude control system and the image quality. Rapid drift of satellite attitude in a dead-band between control limits will affect the quality of the image and the degree to which extensive post-image acquisition processing is necessary.

For large satellite remote sensing systems the push-broom method of image acquisition is preferred. This is the case because large satellites, such as SPOT or LandSat, can accommodate precision pointing and attitude control systems to compensate for orbital torques, system mechanical distortions associated with thermal balance, and pointing mirror motions associated with selection of ground targets.

For a small satellite with constrained resources, it may be very difficult to

avoid having substantial platform motions. The smaller overall satellite mass, the smaller ratio between the moments of inertia of the overall satellite and the mirror pointing system, and the lower level of thermal protection afforded by a smaller thermal mass will most likely make it difficult to accommodate the push-broom mode of image acquisition. We note that a staring system can also be operated in a push-broom mode depending on how the pixel outputs are sampled.

Thus a reasonable choice of detector seems to be a square array of about 1000×1000 pixels. A snapshot of a given site might involve actually transmitting only a small fraction of the available pixel signals, depending upon the area of the site and the accuracy to which the pointing angle is known.

7.2.3 System Considerations and Costs

Satellites of the verification surveillance system must be reliable, available at relatively low cost, be capable of being launched to LEO by low cost rocket boosters, and have an expected operational lifetime of several years. In the spirit of treaty verification, hardness and survivability against radiation and other aggressive influences is not essential. A willful act of destruction of a verification surveillance satellite would be tantamount to a violation of the underlying treaty agreement. To achieve low cost, moderate lifetime, and ground resolution of ~ 1 m requires a delicate balancing of the satellite sub-systems supporting the operation of the optical sensor. Figure (7-1) illustrates the important subsystems which influence the overall satellite operation. For example, the sensor system will require power, communication computing, attitude control and data storage. With a requirement for 1 meter resolution and perhaps 25 meter pointing accuracy, the attitude control system must provide very high quality information about the platform location, attitude and rates of drift of the telescope field of view. This requires state of the art star sensors , control moment gyros, a high quality thruster system, and supporting electronics. Likewise, to support image data transfer rates up to 100 mbps, a sophisticated communication system must be incorporated into the satellite design.

In looking towards the future, we can see that certain technological developments presently taking place could profoundly affect the design of relatively small surveillance satellites. These include the following:

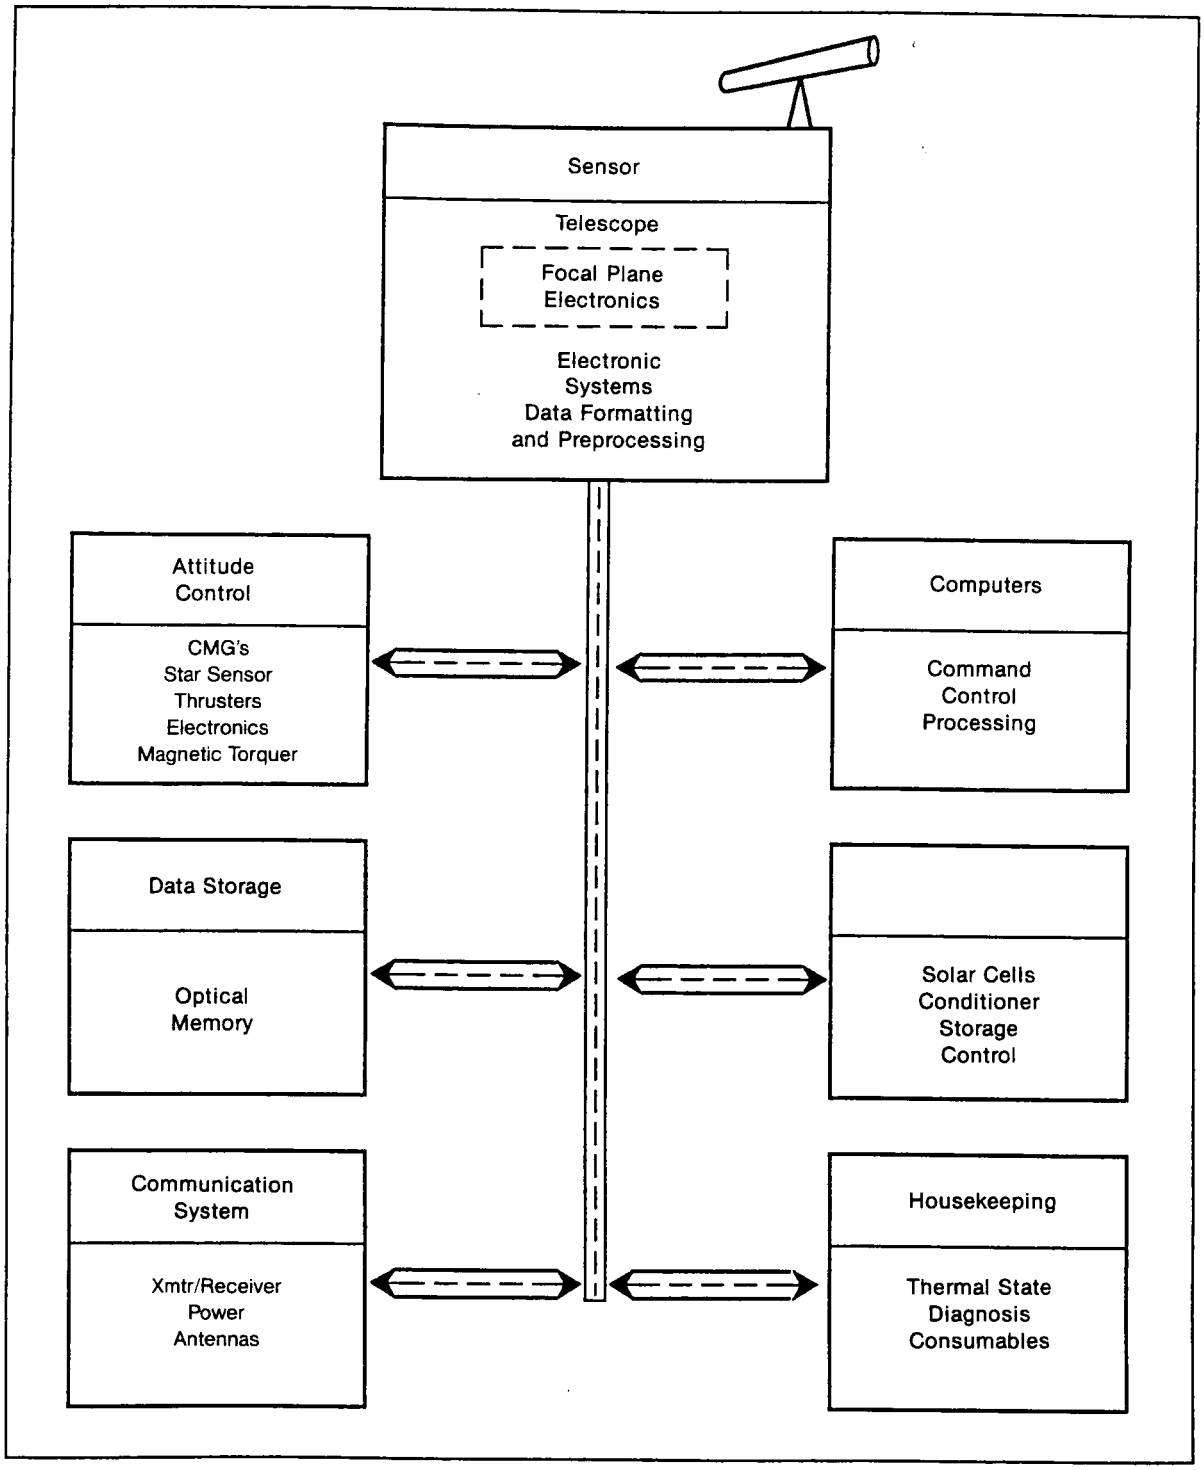

Figure 7-1. Verification Satellite Surveillance System.

1. Attitude Determination and Control System

   - Low cost, high precision star sensors
   - Fiber optic gyro systems
   - High speed computers

2. Data Storage

   - $10^9 - 10^{11}$ bit storage with 50-100 mbps read/write times

3. Communication Systems

   - Laser satellite - satellite data transfer at 100 mbps

4. Computers

   - RISC/SPARC high speed processing with low power, volume, and mass hardened against natural background radiation

5. Sensor System

   - large, low power CCD arrays with robotic, high yield manufacturing, test and calibration
   - large mirrors and lenses

Based on previous small satellite studies (e.g., a 1985 design for MAP-SAT by Itek), it is possible to provide an estimate of the mass budget for a conventional small remote sensing satellite and to extrapolate trends to a future small surveillance system. This information is presented in Table 7-1. It is seen that at present, such a satellite would have a mass in the range 800 to 1400 kg, depending upon whether refractive (light weight) or folded reflective optics (heavy system) are used. The large decrease in future mass for communication equipment reflects the weight savings from using a laser system, such as the laser diodes described in Subsection 7.5, rather than the usual rf transmitter and antenna. The much lighter attitude control system would result from the fiber-optic laser gyros and compact star trackers being developed within SDIO.

With respect to electrical power, estimates have also been made of standby and data-gathering phases of standard satellite operations (see Table 7-2) and gains which can be expected in the future.

Table 7-1 Total Mass Budget

	Present	Future
Power System	50 kg	40 kg
Sensor System		
Petzval (refractive)	250 kg	200 kg
Schmidt (reflective)	300 kg	
Communication system	160 kg	30 kg
Attitude system	160 kg	80 kg
Computer	10 kg	5 kg
Structure	190 kg	100 kg
	1120 ± 25% kg	455 kg

Table 7-2 Satellite Power Budget (Watts)

	Stand-by	Operations	Future
Sensor	26	180	50
Attitude System	55	35	40
Computers	77	77	20
Communications	25	250	40
Image Compression	3	28	10
Power System	50	50	25
	240 W	620 W	185 W

Power System		
Batteries (2)	25 kg	
Solar Cells	12.5 kg	(5 m^2)
Electronics	10.0 kg	
Stunts (2)	3.0 kg	
	50.5 kg	

Perhaps the foremost cost issue involves the expense of manufacturing a high quality, space-qualified optical system capable of providing 1 meter resolution. The telescope size used in Table 7-1 would be appropriate for 1-meter resolution at 300 km altitude. Higher altitude satellites would require larger (and more expensive) telescopes. In any event, these telescopes must be rugged in construction, yet precise in operation. Construction of similar systems in universities for NASA science missions shows that the principal costs of development are engineering, rather than capital equipment. The tasks of assembly, verification, testing, and calibration are human-intensive tasks which scale in proportion to the number of systems being developed. The economies of scale, such as might come from fabrication of multiple optical mounts or detectors, seems to be small in this case. In fact, no industrial experience has yet emerged that might help reduce the costs for these types of complex instruments. Other system components, such as the fine vernier thrusters, the star sensors, the communications systems, and so forth do admit to some economy of scale.

Using methods of estimating costs developed by Itek, Globesat Inc. and others for small satellite systems, the weight and power reductions indicated in Tables 7-1 and 7-2 should translate into a significant, perhaps 40%, savings in cost. One billion dollars for all the satellites in an entire reconnaissance fleet may be a possible cost target.[19]

Operations costs should also be taken into account when considering the feasibility of a small satellite surveillance system. In particular, the cost of operating the data communications system adequate to acquire and forward a large number of images per day must be estimated, especially when one takes into account that current DoD and NASA programs have almost saturated the current TDRSS capabilities. On the other hand, if a store and forward

[19]In the 1985 ITEK design for MAPSAT, the Petzval telescope had a 30 cm lens with a 1.5 m focal length and a 13 μ pixel size for the CCD element, designed to give 10 m ground resolution from an altitude of 1000 km. The telescope weight was approximately 60 kg and that of the satellite was 1000 kg. The estimated cost per satellite was $ 80 M in 1984 dollars. In the 1989 Globesat Inc. designs for a tactical imaging satellite, a Schmidt telescope system was proposed to give 5 m ground resolution from an altitude of 700 km. Special image compression hardware and a proprietary computer system was included to reduce UHF communication system bandwidth and computational system energy requirements. The satellite was physically configured to use the Pegasus launch system, resulting in a total satellite mass of 420 lbs and an overall length of 1.8 m. The cost of the satellite and launcher was estimated to be less than $ 15 M for the first prototype.

method of ground delivery is to be used, one needs to know the impact this will have on various foreign communications sites.

In the end, one of the important opportunities for the small satellite approach is for it to avoid the high cost syndrome which impacts all major U.S. satellite systems. It is well-known that unreasonable demands for substantial on-orbit lifetime in the face of hypothetical threats greatly multiplies the costs of building and flying military space systems. By adopting a short lifetime and permitting some reasonable level of operational vulnerability to intentional disablement it might be possible to achieve significant reductions in the overall system costs. In fact, since this is a surveillance system intended for peacetime operation, such arguments may prove successful.

Finally, it is worth mentioning the advantage to be gained by adopting a system philosophy which recognizes and supports the evolution of system capabilities consistent with current technological capabilities and costs. By avoiding or deferring high cost technology drivers in the early satellites, total system cost (design, acquisition, operations, and replacement) can be kept under control. The key point is to keep system designers and financial managers in close contact with the potential users of the system. This can prevent inadvertent escalation of costs due to hypothetical system operation requirements being taken literally by the design engineers.

7.3 High Altitude Surveillance

We consider possible designs of a satellite telescope for higher elevations ($\geq$ 5000 km), capable of keeping a large ground area of interest under surveillance during daylight at a resolution of 1 meter. Placing such a telescope on a few high altitude satellites could make essentially continuous coverage of a country possible, weather and daylight (or laser illumination) permitting. The required number of satellites can be found by an analysis similar to that of Subsection 7.1.1, but taking into account the curvature and rotation of the earth; this number ranges from fewer than 6 at 5000 km elevation to 1 at GEO. In addition to providing constant coverage, such platforms would be less vulnerable to ASAT threats.

The telescope could use either a refracting lens or a reflector to focus. In either case it would need a large aperture to attain 1 meter ground resolution,

viz. a diameter of 5 meters at 5000 km elevation. A reflector is more sensitive to distortions in its shape than a refractor of long focal length; however a refractor is heavier and must be corrected for chromatic aberration. We discuss these issues in subsections below.

Using an unfilled aperture (such as a ring) would save in weight and complexity, especially in the case of a refractor since only the thin circular edge would be needed. We therefore include a brief discussion of unfilled apertures and of issues relating to their optical performance.

NASA has already put considerable effort into learning how to build large space structures. We will not consider construction issues in any detail, except to emphasize mechanical simplicity and lightness of weight.

A potential alternative to reflectors or refractors would be to fill the aperture with lasers or detectors to obtain images by coherent processing, thereby eliminating the need for an extended third (focal) dimension in the telescope. For example, if an onboard laser is used to illuminate a ground patch, the returning light would produce phased signals in an array of detectors in the pupil-plane, which could be analyzed to reconstruct an image of the ground. This is an active, but not yet mature, technology and we do not consider it here any further.

7.3.1 Refracting Lens Telescope

First we demonstrate the degree to which the control of distortions poses a less serious problem for a refractive lens than for a reflector of the same aperture and f-number. Consider first a reflector of diameter D, focal length L, hence f-number $f = L/D$. Displace an element of the reflector by a distance δ parallel to the axis. At the focus, for $f > 1$, the change in phase of the ray from the displaced element is of order δ/λ. If $\delta_{\rm rms}$ is the mean displacement from a parabola, these distortions introduce an angular spread $\Delta\Theta \approx \delta_{\rm rms}/D$. This is to be compared with the diffraction limit $\Delta\Theta_{\mathcal{A}} \approx \lambda/D$. To achieve the diffraction limit, we must therefore keep the mean displacement $\delta_{\rm rms}$ small compared to the wavelength.

Now consider the situation for a refractor of the same diameter D and focal length L. For a lens of large f-number, the requirement to control

distortions is less severe than for a reflector. If we displace an element of the lens by a distance δ in a direction parallel to the axis, the phase change at the focus will be of order $\delta/f^2\lambda$. If the displacement is perpendicular to the axis, the phase change is larger, of order $\delta/f\lambda$. To achieve the diffraction limit, we must therefore keep the mean displacement $\delta_{\rm rms}$ small compared to $f\lambda$. Clearly, for large values of f this is a less severe demand than we encounter in the case of a reflector. For the applications under discussion we will want f very large indeed, perhaps as large as $f = 100$. Of course, the larger the f-value the thinner, hence lighter the lens. The overall lens thickness must be accurate to better than a wavelength λ but this requirement is not so difficult.

Large f-number also helps with the problem of chromatic aberration which is nevertheless severe as we now show. Let $dn/d\lambda$ be the change with wavelength of the refractive index of the lens material. Then a spread in wavelengths $\Delta\lambda$ will smear out the focal length of the lens by:

$$\Delta L \cong \frac{fD}{n-1}\frac{dn}{d\lambda}\Delta\lambda \qquad (7-7)$$

which must be less than the length of the focal region (i.e., $\Delta L < f^2\lambda$) in order that the image be in focus for all wavelengths within $\Delta\lambda$. Thus the fractional range of wavelengths focused to a single image by a lens that is uncompensated for chromatic aberration is:

$$\frac{\Delta\lambda}{\lambda} < \frac{f}{D}\frac{n-1}{dn/d\lambda}. \qquad (7-8)$$

To a crude approximation this is also the fraction of sunlight that will be focused without image distortion. Typically, $dn/d\lambda \cong 10^{-5}(n-1)$ per angstrom for glass, so if $D = 5$ meters and $f = 50$, we would be limited to about only 10^{-4} of the available sunlight.

Of course optical designers have learned how to compensate chromatic aberration to high precision. For example, one can add a defocusing lens of smaller absolute power made of a material with larger $dn/d\lambda$, or use a number of mutually compensating lenses to match out the second derivatives, etc. In order to keep the added weight to a minimum, the compensation is probably best done by using a single large primary lens, and then near its focus forming a small image of the primary lens where small compensating lenses can be

located, finally using a small reflector to form an achromatic image of the ground.

In what follows, we assume chromatic aberration to be compensated, but it is clearly the major obstacle to using lenses with such large apertures.

An illustrative system is a refractive lens of diameter $D = 5$ m and f-number 50, hence focal length 250 m, at an altitude of $H = 5000$ km. The detector package containing the compensating optics for chromatic aberration as well as the focal plane detector and electronics would be stably tethered to the main satellite unit by a cable about 250 m long.

Due to their different altitudes and Keplerian orbit periods, the tension in the tether produced by the detector package of mass M_{det} would be:

$$T \cong \frac{3fDR_{\text{earth}}^2}{(R_{\text{earth}} + H)^3} M_{\text{det}} g \qquad (7-9)$$

where $g = 9.8$ m/s^2. For our illustrative system, $T \cong 2 \times 10^{-5}$ $M_{\text{det}} g$, which can be supported by a tether of negligible weight. Unless controlled or damped, there could be slow oscillations (of order the orbital period) of the detector package as well as mechanical oscillations from waves on the tether.

Theoretical studies[20] show that tether damping properties are extremely important for the overall pointing accuracy. The damping properties of lateral modes, in particular, are important since these have very slow natural damping rates (up to several years). Thus, active control of the focal plane detector attitude with control moment gyros will be essential. This introduces a number of technical complications since the pointing accuracy achievable with a given system will depend on the bandwidth of the attitude control loop.

Requirements for the tethered detector system can be determined by assuming that the pointing error should give a lateral displacement at the detector no greater than about 1/3 of the detector width; i.e., about 300 m on the ground. At a distance of 5000 km, this corresponds to a pointing error less than 60 μ radians. Calculations for a different system from what has been considered here indicate that the pointing error bound caused by

[20]Xiaohua He, Attitude Control of Tethered Satellites, Ph.D. Dissertation, Department of Aeronautics and Astronautics, Stanford University, Stanford, CA, 94305, October 1989.

disturbance torques associated with various natural sources can be on the order of several 10's of arcseconds; i.e., on the order of 100's of μ radians. More detailed analyses are required in order to determine whether these natural fluctuations can be controlled by active means adequate to give the stability required for high resolution imaging.

As in the earlier discussion (Subsection 7.1.2), the focal plane detector could be composed of about 10^6 CCD pixels for each square km being imaged on the earths surface. The diameter of the beam waist at focus is $f\lambda$, so a one-meter resolution element on the earths surface is focused to a convenient CCD pixel size of about 30 microns. Without allowance for any losses due to chromatic aberration, the fluence per detector pixel arising from sunlight reflected at the earth's surface would be, as in Subsection 7.1.2, about 10^7 photons/sec.

If the 5 m refractive lens were to be made of glass, the weight for $f = 50$ would be about 350 kg. [If the lens is composed of gaseous hydrogen, the weight of the gas comes to 40 kg, but the weight of the bag, if it is to sustain the stresses at a temperature of 300° K, must amount to about 400 kg.] Although 350 kg is not necessarily too heavy, if desired the lens weight could be reduced by using a longer tether and hence greater f-number.

For higher altitudes, the lens weight increases as D^3 and therefore as H^3 for a constant f-number. One might then wish to consider using a ring (annular) lens to reduce the weight (but with some loss in imaging quality; see Subsection 7.3.3).

To illustrate, take $b \ll D$, where b is the width of the annulus. If t is the thickness of the lens at the inner edge, and n the index of refraction of the lens material, then

$$(n-1)t = b/2f. \qquad (7-10)$$

The mass of such a lens is

$$M = \pi\rho Dbt/2 = \pi\rho Db^2/4f(n-1). \qquad (7-11)$$

The area filling factor is given by $F = 4b/D$. Thus, for $F = 0.2$ we take $b = 0.05D$. Relative to the case of the filled aperture, this reduces the weight by a factor of 17, while reducing the focused light by a smaller factor, 5. At 5000 km altitude, such a ring lens would weigh only 20 kg.

7.3.2 Reflector Telescope

Reflectors have long been the method of choice in astronomy for large-aperture telescopes, and may well offer the best choice for high altitude surveillance because they avoid chromatic aberration.

The major problem, as discussed early in the previous subsection, is that of shape distortion degrading the sharpness of the image. We limit our discussion of reflectors to a few comments on this important problem, which is somewhat comparable to the issue of chromatic aberration in refractors. A major effort has been–and continues to be–invested in adaptive optics, which could be used to compensate shape distortions in a reflector ≥ 5 m diameter.

For such a large aperture, the reflector would be assembled from a number of smaller mirrors on some flexible mount, and the positions of these mirrors would be individually adjusted. However, it should be a difficult and costly technical challenge to control the mirror positions to optical tolerances.

The simplest and least expensive alternative might well be to control the primary reflector structure to at best millimeter tolerances, and carry out the optical correction in a relatively small flexible mirror system near the focus. Algorithms for adjusting the shape of the small mirror to maintain sharp, nearly diffraction–limited, images have been developed and tested for some time now [cf R.A. Muller and A. Buffington, J. of Opt. Soc. of Am. 64, 1200 (1974)]. This sort of adaptive optics seems to be quite a promising approach for operating large, light-weight reflectors in space.

7.3.3 Unfilled (Ring) Apertures

Properties of ring apertures were explored in an earlier JASON report (JSR-85-503, July 17, 1985), from which we borrow liberally here. The ring aperture fills the region between two concentric circles of radii r_1 and r_2, with r_2 the larger radius. The filling factor F is the fraction of the full aperture area that is used;

$$F = (r_2^2 - r_1^2)/r_2^2. \qquad (7-12)$$

In Figure 7-2 the diffraction pattern for an extreme ring geometry ($F =$

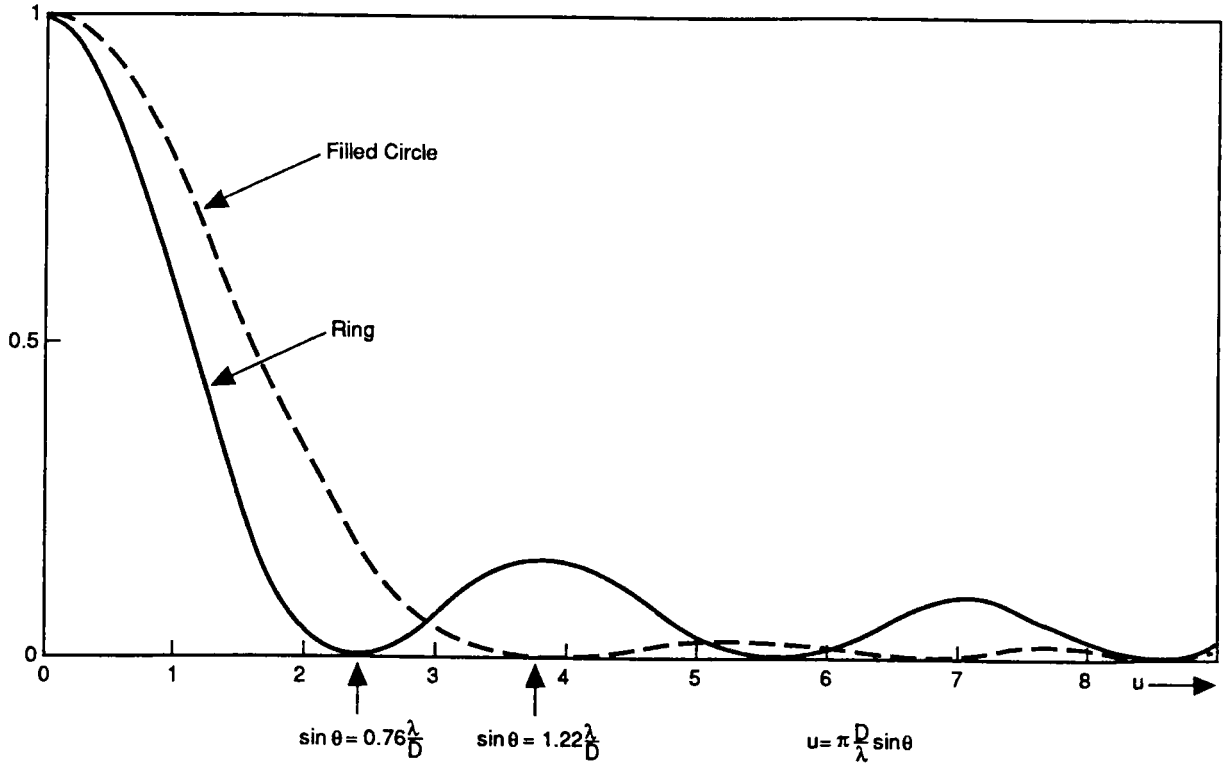

Figure 7-2. Diffraction patterns for filled circular aperture $\left(\frac{2J_1(u)^2}{u}\right)$, and for a ring aperature $(J_0(u))^2$ with area filling factor 0.04.

0.04) is shown together with the pattern for the more familiar filled circular aperture of the same diameter $D = 2r_2$, normalized to unity at the origin in each case. The intensity is plotted as a function of $u = \pi D(\sin\Theta)/\lambda$, where Θ is the viewing angle relative to the ring axis. The envelope of the side lobe intensity falls off as $1/u$, so much of the light appears in the side lobes. Nevertheless the central diffraction peak in principle allows high resolution images near the limits set by the size of the outer radius. The fractional power in this peak is roughly equal to F, while the total intensity is reduced relative to the filled aperture by another factor of F, so the contrast intensity in an image resolved by the central peak will be reduced by a factor of F^2 compared with the filled aperture.

A more precise way of showing the response of the ring aperture to spatial frequencies in the object field is to find the modulation transfer function (MTF). In Figure 7-3 we plot the MTF of a ring aperture for a variety of filling factors F. Note the drop in MTF compared with the full aperture over a broad range of the middle frequencies (those near 1/2 the maximum frequency).

Let's apply these considerations to the signal obtainable from a resolution element on the earth as discussed already in Subsection 7.1.2. The photon count rate given by Equation (7-5) must be reduced by roughly the factor F^2, implying a contrast intensity of about $2 \times 10^7 F^2$ detected photons/sec per pixel for average solar illumination of the earth. Thus, a filling factor of $F = 0.2$ (which as we saw in Subsection 7.3.1 would save a factor of 17 in the weight of a refractive lens) would still have a contrast intensity of 10^5 detected photon/sec per pixel, giving a 2% pixel to pixel contrast in 10 msec.

A major issue with unfilled apertures in viewing a complicated field is the degree to which there would be cross-talk due to the side lobes from different parts of the field. Image processing might be required to make pictures readily interpretable to the eye.

7.4 GEO Radar Transmitter

Here we give a brief description of the concept of a radar in space to observe objects on ground. A slightly more detailed analysis of a similar

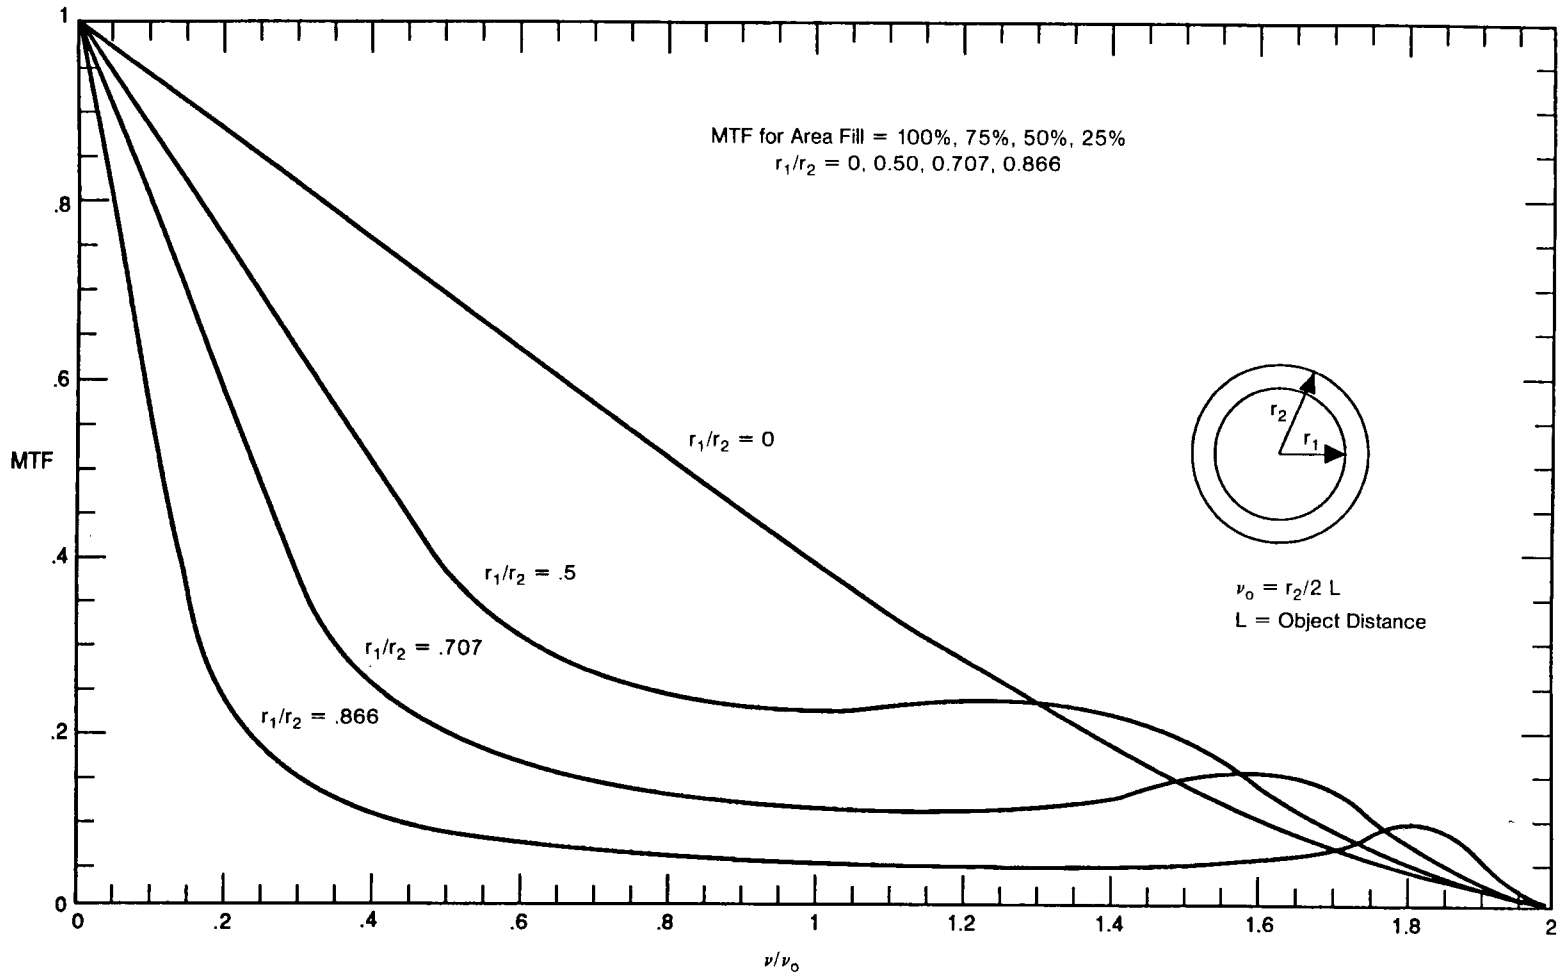

Figure 7-3. The Modulation Transfer Function (MTF) as a function of ν, the spatial frequency of the object field, for ring apertures of various area filling factors.

concept for a radar to observe objects in space is given in JASON Report JSR-89-900.

The idea (see Figure 7-4) is to locate a single radar transmitter at high altitude, possibly at GEO, to give it wide coverage of the earth, and to locate receivers in a fleet of satellites at LEO where the return signal from illuminated ground targets is still large enough to detect.

The transmitter could direct it's beam to any point of interest on the ground within observing range of one of the LEO receivers. Such a system obviates the need for carrying a bulky transmitter/power package on each LEO satellite, thus keeping these satellites small and economical.

The high altitude transmitter would be an X-band phased array powered at MW levels from the ground by a microwave beam broadcast from a roughly 10 km $\times$ 10 km antenna farm on earth. Such a beam, acting as a filled array, would have a main lobe 200 m wide at GEO. In practice, the ground antenna is likely to be a sparse array with significant power in the sidelobes, but generating a few MW of power on the ground is cheap and we will not be concerned with inefficiencies in getting this power to GEO.

The GEO dish obviously could simply reform the beam and steer it to a chosen area, 10 km $\times$ 10 km on earth, but such a small spot size with so much power is unnecessary. With 1 MW of downlink power, a spot 100 $\times$ 100 km could be illuminated at the same power per unit area as a typical SAR satellite like SEASAT can apply to a 10 km $\times$ 10 km area with a 10 kW transmitter.

The GEO dish would most probably not be a simple reflector of power, since the forming and steering, pulse shaping, and possibly even conversion of frequencies may have to be done on the dish. The incoming power could be converted directly on the antenna to electrical power by rectennas, or one could simply pick up the power at the focus, for further conversion. After conversion, the radiated power goes to a phased array on the disk and then to the desired ground spot. A small portion of the downlink beam can be directed to the LEO receiver or airplane in order to give a phase reference signal for coherent processing. Of course, one will have to broadcast a beam with the correct ambiguity function to avoid phase or range ambiguities in the coherently-processed signal, but this technology will be very similar to that of SEASAT, etc. Note that a beam directed from GEO to the Soviet

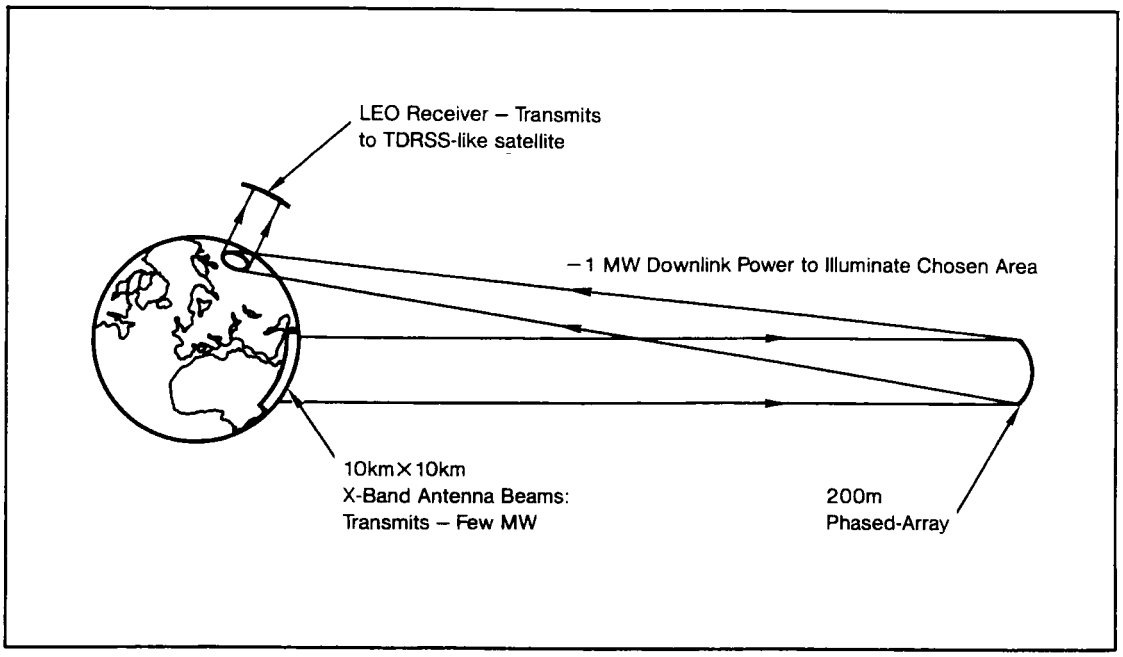
LEO Receiver – Transmits
to TDRSS-like satellite
– 1 MW Downlink Power to Illuminate Chosen Area
10km×10km
X-Band Antenna Beams:
Transmits – Few MW
200m
Phased-Array

Figure 7-4.

Union amounts to a side-looking radar beam, and it would be possible in principle at least, to count the axles of railcars to distinguish an SS-24 car from something else.

It is not clear that a GEO transmitter and a LEO receiver fleet is the optimum arrangement. A few transmitters at half-synchronous altitude, with a smaller fleet of receivers at several thousand km altitude, may be a more economical choice, or one might even find that colocated transmitters and receivers work best. There is nothing sacred about using GEO to station a satellite receiving microwave power from Earth; at lower altitudes the duty cycle for power transmission goes down because of the geometry, but in return the satellite can be made smaller.

7.4.1 Implications for Mobile ICBMs

So far, we have discussed this quasi-continuous surveillance system (radar or optical, or both) as if it were an agreed-upon tool for treaty verification. However, its worth in verifying the location of mobile ICBMs is coupled with equal utility in threatening their survival. Given frequent coverage, so that the position of a mobile ICBM is known at frequent intervals, it could be very risky to argue for its survivability based on dashing to a new location at the advent of an ICBM attack from the other side. The reason is that nearly all the technology needed is now at hand to create a new generation of smart RVs, essentially carrying their own PBVs (cf the PBALL concept of the sixties). These RVs would have small, reasonably good IMUs, on-board computers, communication links to (for example) other satellites, and 100 kg or so of rocket fuel and motor. They could chase down a mobile ICBM by maneuvering in space (not the atmosphere), with current locations given by the ostensible verification system. The counters of launch under attack or attempting to thwart verification by hiding have their own obvious instabilities and drawbacks.

One might hope to decouple verification and survivability to some degree by deliberately arranging for coverage gaps of, say, a few hours by controlling the size and disposition of the LEO satellite constellation, or by arranging the schedule of allowed airplane overflights (see Subsection 7.6).

7.5 Surveillance Using Laser Illumination

Equipping satellites with lasers for illuminating the ground would extend their optical viewing capability to the nighttime, a particularly important advantage at higher latitudes during winter. Major issues are the laser power needed to produce useful imagery, perceived intrusiveness over foreign territory, and eye-safety on the ground.

The same telescope could be used to image the ground in the daytime by sunlight and at night by laser light. [With laser illumination, an alternative to using a conventional telescope is to obtain images by processing unfocused (pupil-plane) signals. This is an active area of research at present.]

The prospects for using laser illumination have been enhanced considerably by recent improvements in laser diode technology, especially the development of coherent laser diode arrays. Compact arrays are now readily available with a power output of 5 watts, and overall efficiency of 30%. Gallium arsenide laser diodes produce light in the 0.7-0.9 μm region where Si CCD detectors have maximum efficiency.

To illustrate the possibilities, we imagine a satellite at 500 km altitude with 20 laser diode arrays for illuminating, say a 10 m $\times$ 10 m patch on the earth with a total power of 100 watts. The beam divergence angle of 20 μrad is attainable with present diode arrays and would require a transmitter aperture < 5 cm. The entire transmitter would weigh a few kilograms and consume a few hundred watts.

As in Subsection 7.2.2 we use Equation (7-5) to estimate the photon count rate in each detector pixel due to the ground return signal. An incident surface illumination of 1 W/m^2, with an assumed average surface reflectivity of 0.2, yields an upward scattered intensity of $I = 0.2$ W/m^2. Taking $\lambda = 0.9$ μm (invisible to the eye) and $\epsilon = 0.7$ at that wavelength, we find:

$$dn/dt \cong 10^5 \text{photons/sec per pixel.} \qquad (7-13)$$

Thus a 3% pixel contrast resolution would require an exposure of about 20 msec, short enough to obtain snapshots of vehicles moving up to speeds somewhat greater than 60 mph. Though not as good as with sunlight illumination, this is still a very interesting level of performance.

An important question is how intrusive lasers would be over foreign territory. One has the option of using visible or invisible light, whichever would prove the most acceptable. Eye safety would of course be a major concern; the cause of damage at the wavelengths of interest here would be overheating of the retina at the point of focus. We list in Table 7-3a the American National Standards Institute maximum permissible exposure (MPE) recommendations, as explained and reported in Safety With Lasers and Other Optical Sources by David Sliney and Myron Wolbarsht (Plenum, 1980) p. 261 ff (MPE is defined as 10% of the level believed to cause damage to the human eye).

For pulses shorter than about 2×10^{-5} sec, the local heating of the retina is independent of the pulse length because thermal conduction is too slow to spread the heat on this time scale. For longer pulses, thermal conduction begins to play a significant role, enabling the retina to absorb more total energy without the point of focus overheating. Thus, MPE increases with pulse length, or exposure time t, varying as $t^{3/4}$ for $t > 1.8 \times 10^{-5}$ sec. This behaviour is seen in Table 7-3a for the wavelength range 400-1049 nm. There is some variation of MPE with wavelength in this region due to focal spot size, absorbtivity etc, that is shown in Table 7-3b.

The level of illumination (1 W/m^2) we used for illustration amounts to 2×10^{-6} J/cm^2 in the 20 msec snapshot contemplated, which is well below the MPE calculated from Table 7-3 for this value of t for any wavelength in the 400-1049 nm region.

If several orders of magnitude higher laser illumination intensity were to be contemplated, eye safety could be maintained by using wavelengths in the 1.6–1.7 μm region, where the atmosphere has a transparent window, but where the water (aqueous humor) in the eyeball has an absorption length of 1 mm, allowing only about 10^{-4} transmission to the retina (consistent with the large MPE in the bottom rows of Table 7-3a). Laser diodes also operate at this wavelength, but the longer wavelength would entail some loss of angular resolution and a factor of 30 loss in photon counting efficiency with the best detectors presently available.

In summary, laser diode arrays in the 0.7–0.9 μm region now make laser illumination a promising satellite reconnaissance tool for optical viewing at night.

Spectral Region	Wave Length	Exposure Time, (t) Seconds	Exposure Limits
UVC	200 nm to 280 nm	10^{-9} to 3×10^{4}	$3mJ/cm^2$
UVB	280 nm to 302 nm	"	3 "
	303 nm	"	4 "
	304 nm	"	6 "
	305 nm	"	10 "
	306 nm	"	16 "
	307 nm	"	25 "
	308 nm	"	40 " not to
	309 nm	"	63 " exceed $0.56t^{1.4}$
	310 nm	"	100 " J/cm^2
	311 nm	"	160 "
	312 nm	"	250 "
	313 nm	"	400 "
	314 nm	"	630 "
	315 nm	"	$1.0J/cm^2$
	315 nm to 400 nm	10^{-4} to 10	$0.56t^{1.4}J/cm^2$
UVA	315 nm to 400 nm	10 to 10^{3}	$1.0J/cm^2$
	315 nm to 400 nm	10^{3} to 3×10^{4}	$1.0\ mW/cm^2$
Light	400 nm to 700 nm	10^{-9} to 1.8×10^{-5}	$5 \times 10^{-7}\ J/cm^2$
	400 nm to 700 nm	1.8×10^{-5} to 10	$1.8(t/^{4}\backslash t)\ mJ/cm^2$
	400 nm to 549 nm	10 to 10^{4}	$10mJ/cm^2$
	550 nm to 700 nm	10 to T_1	$1.8(t/^{4}\backslash t)\ mJ/cm^2$
	550 nm to 700 nm	T_1 to 10^{4}	$10C_B\ mJ/cm^2$
	400 nm to 700 nm	10^{4} to 3×10^{4}	$C_B\mu W/cm^2$
IR-A	700 nm to 1049 nm	10^{-9} to 1.8×10^{-5}	$5C_A \times 10^{-7}J/cm^2$
	700 nm to 1049 nm	1.8×10^{-5} to 10^{3}	$1.8C_A(t/^{4}\backslash t)\ mJ/cm^2$
	1050 nm to 1400 nm	10^{-9} to 10^{-4}	$5 \times 10^{-6}\ J/cm^2$
	1050 nm to 1400 nm	10^{-4} to 10^{3}	$9(t/^{4}\backslash t)\ mJ/cm^2$
	700 nm to 1400 nm	10^{3} to 3×10^{4}	$320C_A\mu W/cm^2$
IR-B & C	1.4μm to $10^{3}\mu$m	$10^{-9}$to 10^{-7}	$10^{-2}\ J/cm^2$
	1.4μm to $10^{3}\mu$m	10^{-7} to 10	$0.56\ (^{4}\backslash t)\ J/cm^2$
	1.4μm to $10^{3}\mu$m	10 to 3×10^{4}	$0.1\ W/cm^2$

C_A– See Fig. 7-5 (b), Laser EL listing.

$C_B = 1$ for λ = 400 to 550 nm; $C_B = 10^{[0.015(\lambda-550)]}$ for λ = 550 to 700 nm.

T_1 = 10s for λ = 400 to 550 nm; $T_1 = 10 \times 10^{[0.02(\lambda-550)]}$ for λ = 550 to 700 nm.

For λ = 1.5 to 1.6 μm increase EL by 100.

Table 7-3a. Maximum Permissible Exposure (MPE) for Direct Ocular Exposures (Intrabeam Viewing) from a Laser Beam.

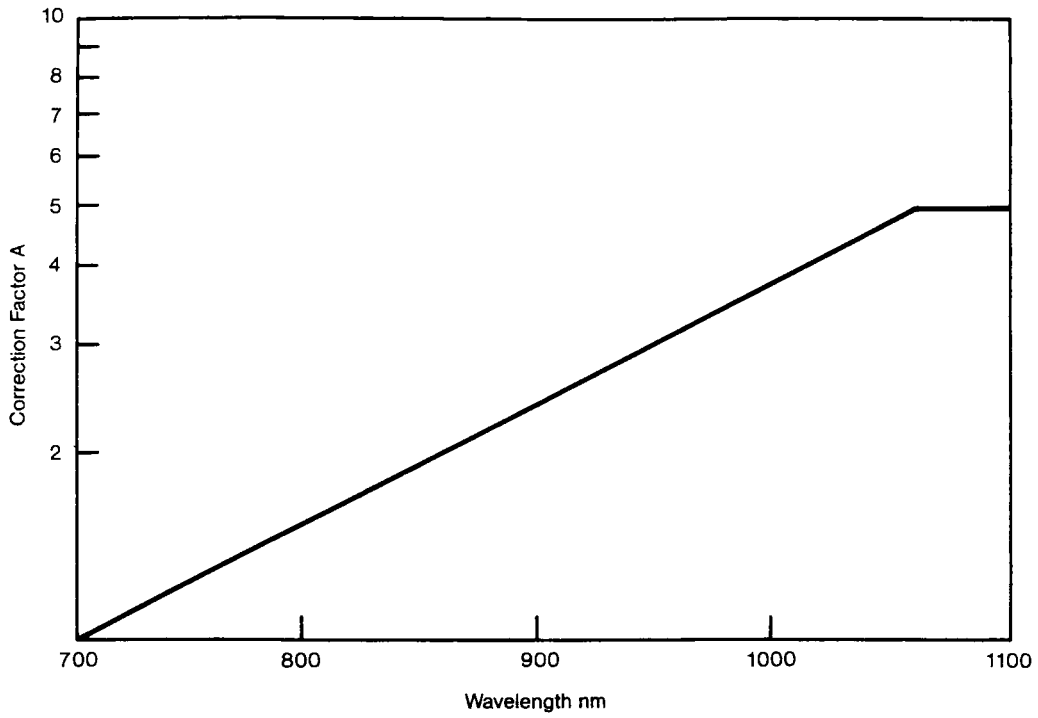

Table 7-3b. Correction Factor A, C_A. The formula for C_A is: $C_A = 1$ for wavelengths (λ) of 400 nm to 700 nm; $C_A = 10^{[0.002(\lambda - 700\ nm)]}$ for 700 nm $< \lambda <$ 1050 nm; and $C_A = 5$ for 1050 $< \lambda <$ 1400 nm.

7.6 "Open Skies:" Aircraft Overflights

President Bush recently proposed that "open skies," meaning permitted surveillance overflights, be revisited as a possible verification and confidence-building measure. The proposal echoes similar Eisenhower administration proposals dating back to 1956. What is different now is not only new Soviet openness, but also the fact that both the US and the Soviet Union accept NTM satellite overflight as routine–in fact protected by treaty. It therefore may be possible to find regimes of allowed overflight which are only incrementally more intrusive than existing NTMs, but which more than repay the increased intrusiveness with mutually useful verification and confidence.[21]

To illustrate, let us give one example. The existence and utility of reconnaissance satellites is accepted by both sides. Satellite orbits are highly predictable. It is taken as a given by each side that the other will refrain from some activities, which would otherwise be observable, during a satellite pass-once or a few times per day, say, for a total of 20 minutes. The long advance predictability of reconnaissance coverage makes it possible to hide, by careful advance scheduling, even very large and elaborate activities. Each side might worry, in the extreme case, that preparations for war or treaty breakout could be thus hidden.

Suppose that aircraft overflights were allowed with the following conditions: (1) They would be limited in number. (2) They would be single straight-line flights over the "host" country a limited time in advance, say two hours. (4) Camera apertures would be limited in number (e.g., to one) and diameter (e.g., to 10 cm), so that resolution would be not substantially greater than existing NTM.

What would such flights accomplish? They would still allow the host country to hide–on random two hour notice–a small number of its most sensitive activities. However, because of the inherent "friction" of rapidly disseminating information and communicating orders, the host country would be unlikely to be able to hide large, complicated, pre-planned activities taking place at many locations. But these are just the kinds of activities that are the most worrisome.

[21] We note that the 1987 Stockholm accords on confidence and security building measures sanctions overflights for observing any large scale military exercises and maneuvers.

Short notice overflights would thus seem to add a different and useful dimension to confidence-building.

7.7 Summary

We have reviewed a broad range of recently developing and potentially new technologies that can add to our means of verification. Of particular interest is the possibility of a fleet of relatively small and inexpensive satellites at low altitude for providing a frequent revisit capability with medium (1 meter) ground resolution.

REFERENCES

- Xiaohua He, Attitude Control of Tethered Satellites, Ph.D. Dissertation, Department of Aeronautics and Astronautics, Stanford University, Stanford, CA 94305, October, 1989.