UCRL-JC-104584
PREPRINT

Evaluating Arms Control Treaty Verification Regimes: A Risk Analysis Approach

R. Scott Strait
Thomas A. Edmunds

This paper was prepared for submittal to the
Probablistic Safety Assessment and Management (PSAM), Conference
February 4-7, 1991, Beverly Hills, CA

February 4-7, 1991

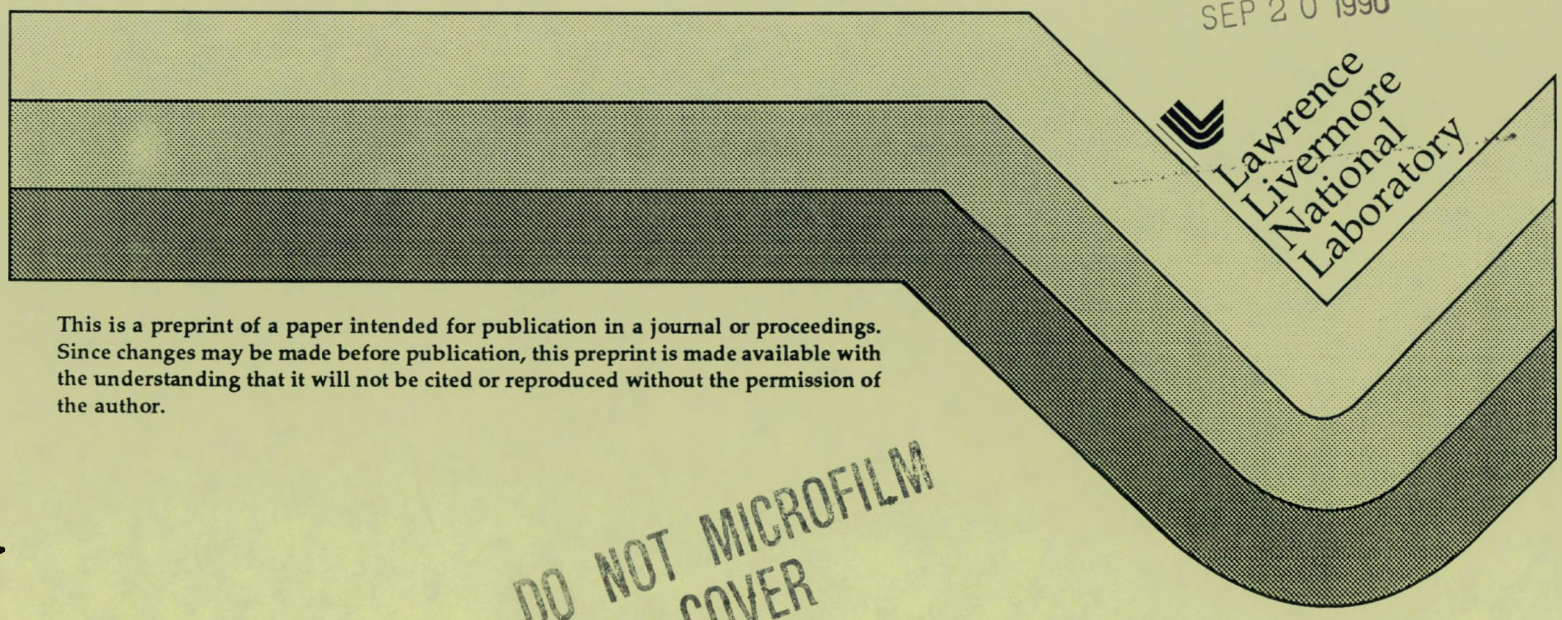

DISCLAIMER

EVALUATING ARMS CONTROL TREATY VERIFICATION REGIMES: A RISK ANALYSIS APPROACH

UCRL-JC--104584

DE90 017523

THOMAS A. EDMUNDS AND R. SCOTT STRAIT
Lawrence Livermore National Laboratory, P. O. Box 808, Livermore, CA 94550

ABSTRACT

We develop a quantitative risk analysis methodology to evaluate verification measures for a bilateral arms control treaty. The methodology is designed to accomplish an integrated evaluation of the total adversary evasion potential and the complete verification regime while considering the interaction among different verification measures. The method uses a network or a fault tree to: identify potential weaknesses in the overall treaty verification system; highlight the evasion and breakout strategies least likely to be detected or deterred; and determine the individual verification measures that offer the greatest benefit.

INTRODUCTION

Bilateral arms control treaties place numerical and other limits on many different strategic weapon systems. Monitoring compliance with a treaty requires a complex system relying on many different verification measures. These measures vary significantly in terms of their effectiveness, intrusiveness, and economic costs. Similarly, the verification measures need to detect and deter a wide range of possible adversary treaty evasion strategies. Identification of acceptable, cost-effective verification measures is a challenging problem, as these measures must maximize the ability to detect and deter adversary violations and breakout attempts of greatest concern. An integrated evaluation of the complete arms control treaty verification system and the total adversary breakout potential is required. The evaluation should incorporate interactions among the different verification measures and the possible evasion strategies.

In this report, we describe a quantitative risk analysis methodology that addresses these needs. The methodology is based on a network or fault tree representation of a complete verification system and the possible adversary breakout strategies. It is designed to: identify potential weaknesses in an overall treaty verification system; highlight the evasion and breakout strategies least likely to be detected or deterred; and determine the individual verification measures that offer the greatest benefit.

In negotiating verification provisions for an arms control treaty, each party should consider the possibility of adversary evasion and what constitutes a militarily significant evasion. The level of evasion that is militarily significant can be expressed in terms of specific adversary objectives of covert production and deployment of various quantities of different weapon systems. Given a set of evasion objectives, we use the methodology to define an adversary evasion strategy as a sequence of steps that may be taken to achieve an objective. The steps include those necessary for design, production, testing, and deployment of the weapon system. There may be many feasible evasion strategies. We are most concerned with the evasion strategies that are least likely to be detected.

Possible treaty verification protocols may include a variety of different verification measures. Some measures may be more effective than others in detecting activities associated with the steps in a particular evasion strategy, and different verification technologies may complement one another in various ways. The uncertainty about whether a particular verification measure will be effective in detecting evasions is treated by assessing a probability of successful evasion at the steps where the measure is implemented. Our methodology models the interaction among verification measures, evasion strategies, and

Work performed under the auspices of the U.S. Department of Energy by the Lawrence Livermore National Laboratory under contract No. W-7405-ENG-48.

successful evasion probabilities. The methodology provides a useful tool for ranking possible verification measures in terms of their deterrence effect.

The methodology is composed of five phases: (1) Identify possible adversary evasion objectives; (2) Develop a model representing all evasion strategies that meet objectives; (3) For each step in each of these strategies, estimate probability of evasion associated with verification technologies in force at that step; (4) Use an algorithm to determine evasion strategies least likely to be detected; and (5) Analyze results and perform a sensitivity analysis, repeating phases 2, 3, and 4 for different sets of verification measures. The following sections discuss each of these steps individually.

PHASE 1: IDENTIFY ADVERSARY EVASION OBJECTIVES

In phase 1, the starting point for the analysis, we begin with the determination of what constitutes a militarily significant evasion of the treaty or potential for treaty breakout. Having determined the level of treaty evasion that constitutes a strategically significant advantage, we then identify specific adversary evasion objectives that might provide such an advantage.

One general class of objectives involves the production of complete weapon systems in excess of the limitations imposed by treaty. These complete weapon systems could be stored or deployed at covert or declared sites. A second general class of objectives involves production and stockpile of major components of weapon systems. These component stockpiles represent a threat because, if the treaty were abrogated, complete weapon systems could be rapidly assembled to gain a strategic advantage in a short time period. We refer to such a scenario as a breakout.

When more than one evasion objective is identified, we employ a weighting scheme in order to measure the relative desirability (for the adversary) of meeting these various objectives. Each objective is assigned a relative value on a numeric scale. This value can then be combined with the highest probability of successful evasion for that objective to obtain an overall measure of verification system effectiveness.

PHASE 2: DEVELOP MODEL OF EVASION STRATEGIES

In phase 2, we develop a model that represents all adversary strategies that may be used to evade verification measures in force and achieve the objectives. We represent an evasion strategy by a sequence of steps that must be performed in order to achieve a desired objective. Typically, these steps describe illegal production, testing, and deployment processes for a weapon. Thus, a means of completing all steps in this sequence without detection would correspond to an evasion strategy.

Determining how to divide the production, testing, and deployment of a weapon system into discrete steps depends on a number of factors. These include the production process and its choke points, the physical location of the different parts of the process, and the verification system to be evaluated. The monitoring points of the process for the different verification measures are very important in dividing the process into steps. Because different verification regimes may involve monitoring different steps, phase 2 may have to be repeated for each regime.

It is useful to develop a geometric representation of the interrelationships among production steps, strategies, and objectives. To this end, we employ either a network or fault tree model. In order to illustrate this concept, we consider a network representation of the illegal weapon production and deployment process, depicted in Fig. 1a. The analogous fault tree representation is shown in Fig. 1b. In this paper we focus our discussion on network representations, but we refer to the fault tree approach where such additional explanation may be helpful.

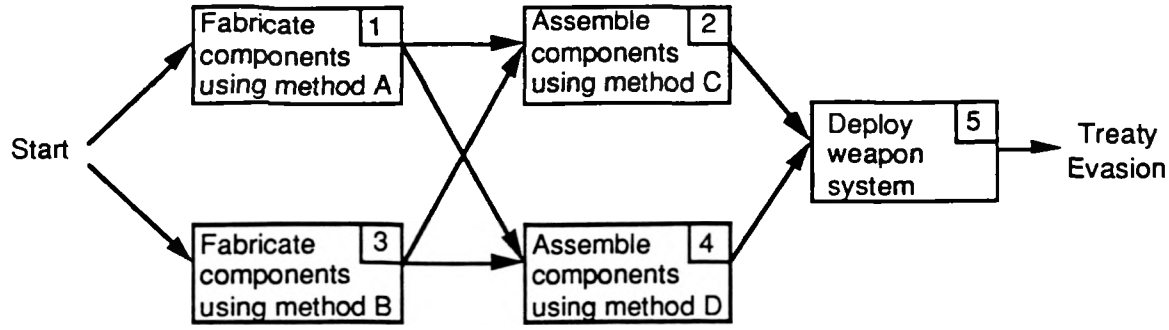

Figure 1a. Network representation of evasion strategies.

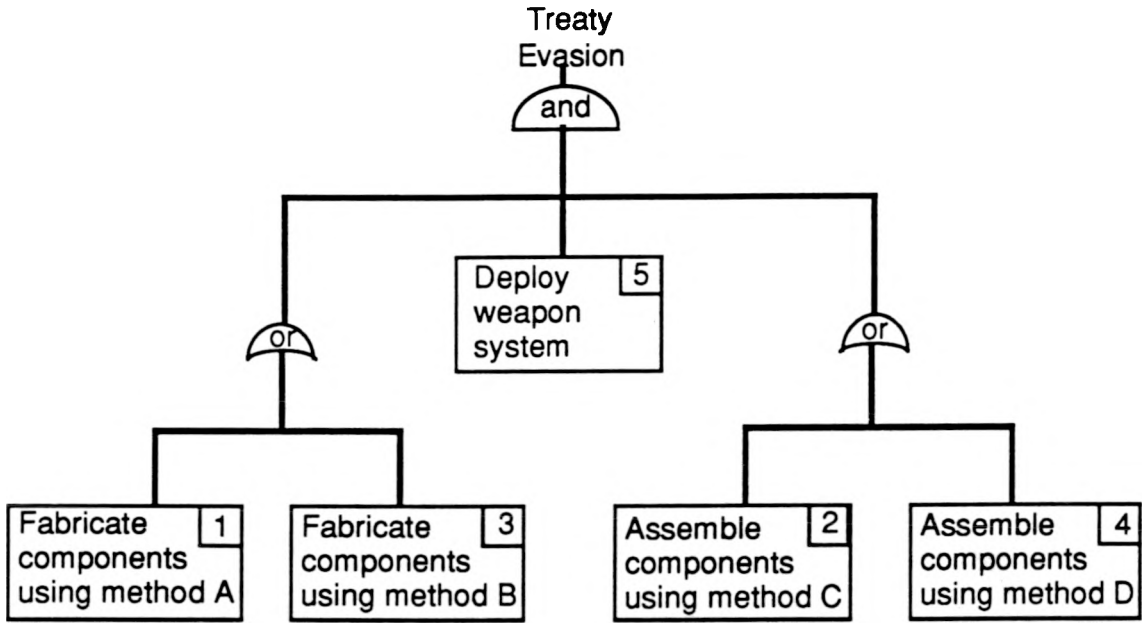

Figure 1b. Fault tree representation of evasion strategies.

As indicated in Figure 1a, component fabrication, component assembly, and weapon deployment are required to achieve the objective. The component fabrication task may be accomplished by either evasion method A at step 1 or evasion method B at step 3, while the component assembly task may be accomplished by either evasion method C at step 2 or evasion method D at step 4. The assembled weapon is then covertly deployed at step 5. A strategy corresponds to a path through the network in Fig. 1a from the point labeled "start" to the point labeled "Treaty Evasion." The model is therefore able to represent all possible evasion strategies as a collection of paths through the network. In the corresponding fault tree model, each minimal cut set corresponds to a possible evasion strategy.

The methodology will display some paths that may be obviously inferior or impractical. The remaining phases of the methodology will reveal their inferior or impractical nature in the probabilities of successful evasion, and their inferior nature will be reflected in the results. They are included in the model because it is advisable to include all possible, and even inferior, strategies rather than risk eliminating important strategies.

PHASE 3: ESTIMATE EVASION PROBABILITIES

In phase 3, for each step in the network or basic event in the fault tree we estimate the probability that treaty evasions associated with that step will be undetected by verification measures in force at that step. These probabilities are by nature subjective judgments and may vary depending on one's perspective. They should, of course, be assessed by experts

familiar with verification technologies and weapon production processes and, to the extent possible, derived from theory or experiment.

Although there are refined techniques for assessing these probabilities through structured interviews, satisfactory results are likely to be obtained if the probabilities are based on agreement between multiple experts. Where experts disagree on these probabilities of successful evasion, the different opinions should be considered and separate evaluations performed. If the results of the evaluations are significantly different, then further investigation of the differences of opinion is warranted.

The probabilities of undetected evasion at a given step should be assessed separately for each verification measure in effect at that step. These probabilities are then combined into a single probability of undetected evasion for the step. If the evasions of the different verification measures are probabilistically independent, then the individual probabilities can simply be multiplied. However, if they are not probabilistically independent, then care should be taken to reflect the dependencies in the combination.

In our example, the evasion probability assigned to step 1 would reflect the likelihood that the adversary can fabricate illegal components without detection using method A. Note that if no verification measures are in force at a particular step, the corresponding evasion probability is 1.0. In Fig. 2, we have assigned probabilities of undetected evasion to each of the steps of the network in Fig. 1. For example, the probability that the adversary would be able to successfully evade the verification measures in place at step 1 has been assigned a value of 0.8.

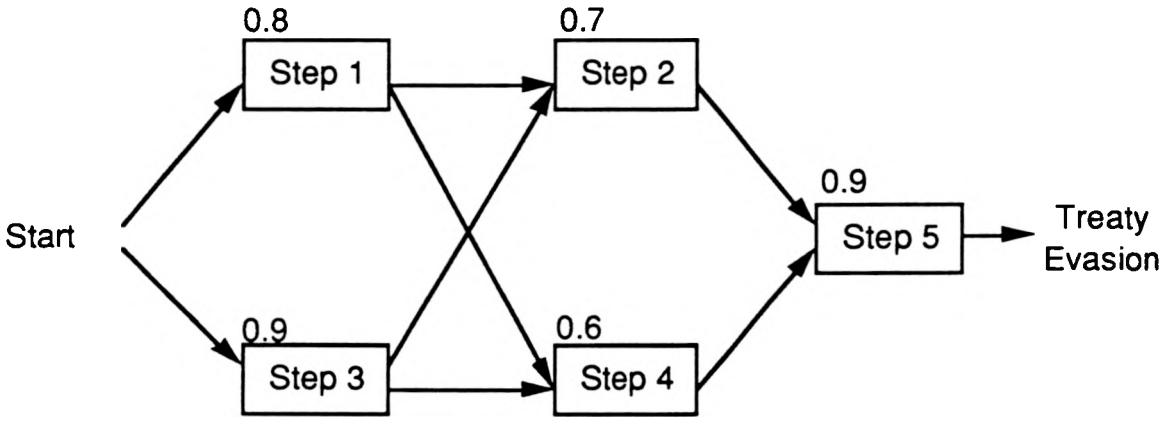

Figure 2. Example of network showing probabilities of undetected evasion for each step.

PHASE 4: DETERMINE EVASION STRATEGIES LEAST LIKELY TO BE DETECTED

In phase 4, we use the probabilities of undetected evasion for each step estimated in phase 3 to compute undetected evasion probabilities for all possible evasion strategies. The undetected evasion probability for a particular strategy equals the product of the undetected evasion probabilities for each of the steps in the strategy. The simplest algorithm for identifying strategies having the highest probability of successful evasion is path enumeration. This technique explicitly evaluates all possible evasion strategies, or paths through the network. We illustrate the technique with the example in Fig. 2, where we find the maximum probability path from "start" to "treaty evasion."

The evasion probability associated with each step is shown above the corresponding box. This example network contains four paths and associated evasion strategy probabilities. For example, the probability of undetected evasion associated with the path using steps 1, 2, and 5 is $0.8 \times 0.7 \times 0.9 = 0.504$. Using the path enumeration technique, one can determine

that the evasion strategy with the highest probability of undetected evasion. In this example, the lowest probability of detection is 3 → 2 → 5, with an associated probability of undetected evasion equal to 0.567. This path is shown in Fig. 3.

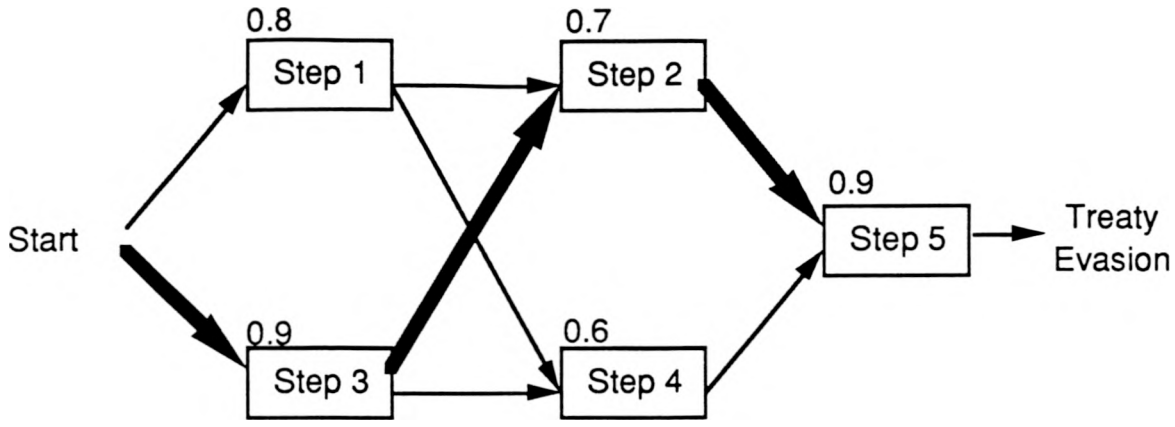

Figure 3. Example of network showing path with highest probability of undetected evasion.

The advantage of the path enumeration approach is that it explicitly represents all possible evasion strategies. Such an approach may be warranted if the model contains relatively few manufacturing steps. The disadvantage of this approach is that the number of paths through the network grows rapidly with network size. More efficient algorithms are available that selectively evaluate those paths with the highest probabilities of undetected evasion. The fault tree model can also be used for the analysis. In general, the fault tree approach is preferred if there are many, complex logical interrelationships among process steps that make a network representation impractical.

PHASE 5: ANALYZE RESULTS AND PERFORM SENSITIVITY ANALYSIS

Finally, in phase 5 we analyze and interpret the results to identify the evasion strategies least likely to be detected and to identify the single verification measures and verification regimes that are most effective. As part of this effort, we perform a sensitivity analysis by repeating phases 3 and 4 for different verification regimes. Additionally, we may need to review phase 2 to ensure that all possible evasion strategies that may be employed to defeat the verification technologies are represented in the model.

The results of the algorithm represent a figure of merit, which is used to compare the deterrent effect of different verification measures. It is computed for each evasion objective by multiplying the objective's undetected evasion probabilities obtained in phase 4 by the objective's relative weights determined in phase 1. We compare the figure of merit when a particular verification technology is in force with the figure of merit obtained when that technology is not in force. This provides a quantitative measure of the value of that particular verification technique, taking into consideration changes in evasion strategies in response to the introduced verification technology. In this manner, we identify the verification techniques that are most effective in detecting adversary treaty violations.

At this phase of the methodology, we also test the sensitivity of the results to any inputs whose values may be controversial or subject to disagreement. In many cases the results will not change significantly when the inputs are changed. In such cases no further analysis is required. In cases where the results are highly sensitive to the input values, we need to further analyze the input values to resolve the important uncertainties. If this is not practical, then the results should be presented so as to reflect important differences of opinion.