SAND2015-2727R

Progress to Date on Developing the NUSAM Methodology Document

Mark Snell

Sandia National Laboratories

April 1, 2015

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Proposed Outline of NUSAM methodology (with assignments)

Note: Page numbers refer to the Analysis Working Group Document. Items in red font have been added to the basic Analysis Working Group Methodology Document

Section 8: Transportation Security Assessment

Introduction

The same basic steps occur when evaluating the physical protection of nuclear or radiological material in transport against theft or sabotage as are needed for protecting nuclear facilities and nuclear material in use or storage. There are some notable distinctions, however. There are limited layers of protection around material when in transport so that the analysis focuses more on scenario analysis than for fixed sites which may require formal path analysis if the facility has some complexity.

In many respects, ground transportation security is more challenging than security at a fixed site. Operation in the public domain is frequently required and the same degree of access limitation is not possible as it is in a protected fixed site. Because of these differences, response force personnel in transit play a more dominant role in the security of a mobile system than they do for a fixed site. In all cases, however, the system time delay that is required to provide the response force the time to react must be provided primarily by transportation vehicle technology elements.

There are several important aspects of transport security that must be considered in the effectiveness evaluation:

1. Movement and Vehicle States: Several possible movement states for a shipment are shown below

   At the Origination Site:
   - Loading
   - Origination Site Egress

   Urban movement
   Rural movement
   Engaged in Planned Stops
   Parked in a Safe Haven
   Engaged in Unplanned Stops
   At the Destination Site:
   - Destination Site Ingress
   - Unloading

   Figure 8-1 shows depicts some of these states ("Rolling" refers to the shipment moving) that correspond to operating states for a nuclear facility.

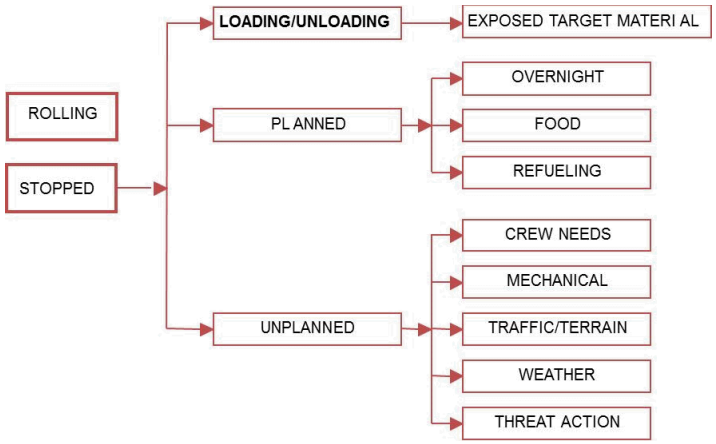

Figure 8-1. Vehicle States

Note that Intrasite shipments between Protected Areas at the same facility are treated as intersite shipments. (See Figure 8-2).

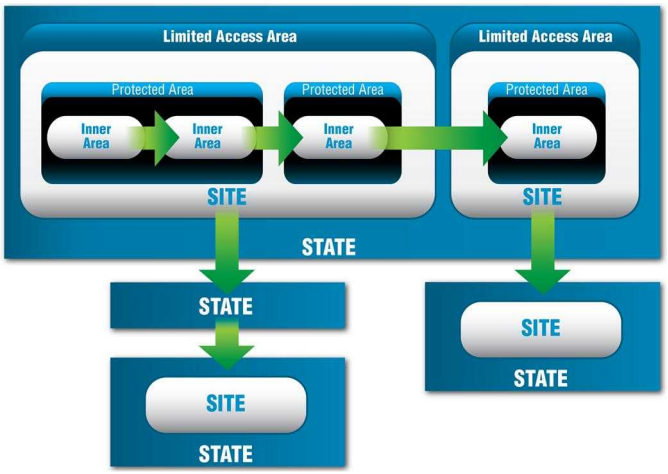

Figure 8-2. Combinations of Transportation Origination Points and Destinations

2. Location of an adversary attack along a route: An attack can occur anywhere along a route of up to several thousand miles, giving the adversary a wide choice of potential attack locations. The adversary could conceivably be able to attack in a location where it will be virtually impossible for any sizeable secondary response force to arrive within a useful period of time. Figure 8-3 depicts a hypothetical route where rural conditions exist over some distance (labeled Rural Movement).

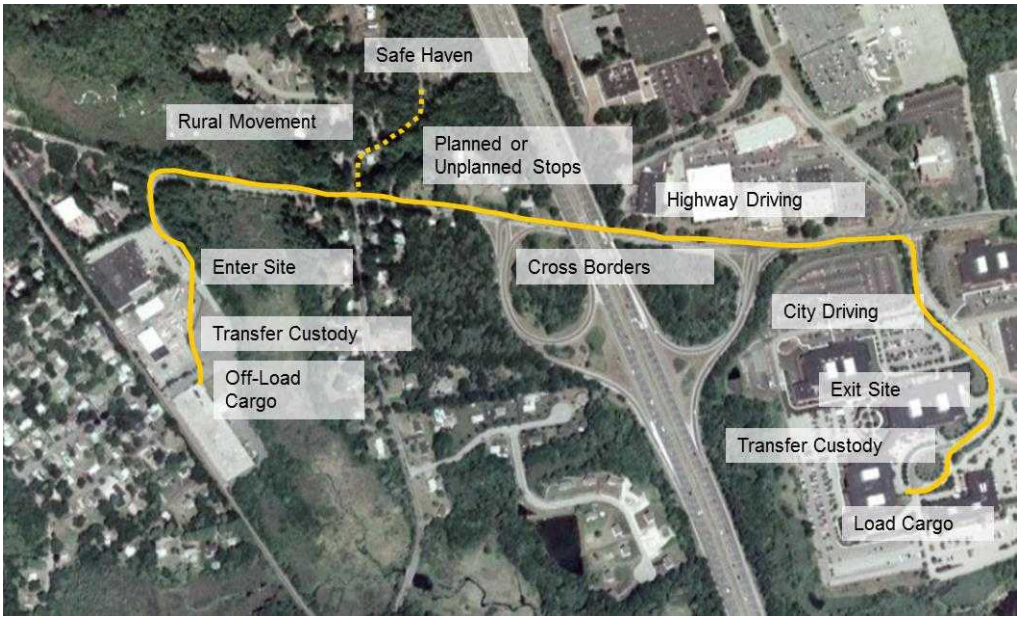

Figure 8-3. Possible Locations and Conditions Along a Route

3. Convoy configuration: The effectiveness of physical protection for a shipment will also depend upon how many vehicles will be involved in the movement, how protected they are against adversary weapons as defined in the Threat Assessment/DBT, and how vehicles in a convoy are configured (if more than one vehicle is involved).

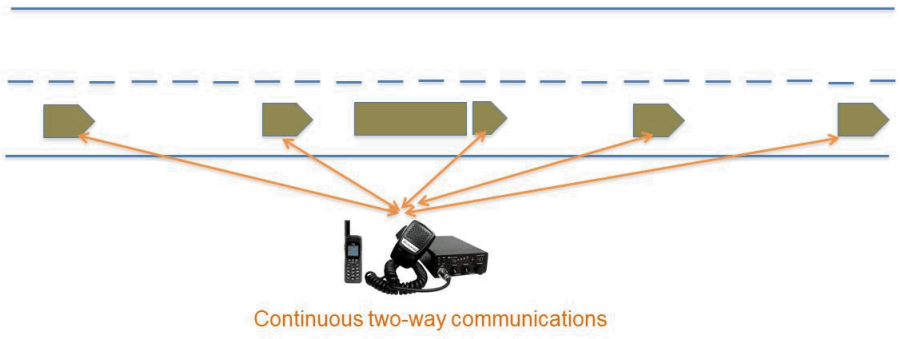

Figure 8-4. Possible Locations of Vehicles in a Convoy Along a Route

Some Considerations for Effectiveness Evaluation

Commonly in the international community, the only quantitative information available are estimates for delay times to defeat certain features of the transport itself or the target and estimates about response force times to respond to different points along the route. These estimates may be based on information collected from a fairly limited number of tests and exercise results.

Annex B: Techniques for Characterizing Performance Metrics

Introduction

Effectiveness evaluations depend on performance metrics such as delay times, detection probabilities, and PPS Response Times (PRT's). There are basically three methods for characterizing these performance metrics for systems and components: statistical data, use of expert judgement, and use of models and simulations. This section will focus on the first two topics.

Statistical Data

This method records one or more observations or values, typically through statistical sampling. Examples include: delay times for defeating a barrier, times for responders to perform some task, the number of times a Central Alarm Station operator properly identifies someone between perimeter fences, and whether an alarm is triggered when someone enters the area that the alarm is supposed to be monitoring. Where such data exists, classical statistical techniques can be used such as maximum likelihood estimators, confidence intervals, and hypothesis tests.

It is quite common to have only one observation for a barrier defeat time. In such cases, that observation can be treated as a point estimate. Typically, when a result is binary (e.g., the alarm successfully triggers or it fails to), the total number of successes, X, out of the total number of tests, N, is assumed to follow a binomial distribution with probability of success p. This information can be used to estimate p as the ratio X/N or to create a confidence interval for p. It is sometimes possible to estimate the components of PRT based on several exercises (for example, by having guards/response forces on different operational shifts perform the exercises several exercises per year). In such cases, a conservative approach to setting PRT is to select some percentile (e.g., 75th percentile value) rather than use an average value.

Simulations may need information on the status of the physical protection system at a particular time. A common way to represent this information is using what is called a "picture in time." A "picture-in-time" records information such as where guards and response forces were located, whether they were in vehicles, and the status of alarms, at a specific time that the information was collected.

Use of Expert Judgment

In some cases, either insufficient or no data exists or there is no safe way to correctly collect the data (as might happen when the event is some violent violation of a two-person rule). In such cases, the evaluation would depend on values elicited from experts.

Some models and simulations may only use point values; in such cases the experts would determine what the values were. In some cases the model/simulation requires distributions of values. There are several ways to develop such distributions using experts:

- For response or delay times, one approach is to have one or more expert estimate a minimum time and maximum time (or a 5^{th} and 95^{th} percentile time) subjectively. Then the expert(s) would subjectively determine other necessary properties about the distribution, such as the type (triangular, uniform, lognormal, and beta, etc.) and any variables that are not specified (such as the most likely value (mode) of the triangular distribution).

- For probabilities, it is common to assume that the expert's subjective distribution for the probability can be described by a beta distribution with parameters α and β. There are two approaches to fitting a beta distribution: 1) determine a 5^{th} and 95^{th} percentile for the probability – so that the expert's feels 90% sure that the true probability falls between these values – and to determine a and b from those limits; 2) have the expert determine what they think the probability should be – call it the value P* - and how many observations they believe their estimates is worth – call this value M – and then to set α = P*M and β = (1-P*)M.

In the 1970's and early 1980's U.S. security simulations assumed a triangular distribution for delay times with the mode set at the observed delay time taken from an actual test with the minimum and the maximum times set by experts. In such cases statistical data was combined with expert judgement to develop the three parameters needed for the triangular distribution. Other models assumed PRT and delay times follow normal distributions where the standard-deviation is some multiple of the mean (e.g., the standard deviation is .3 times the mean).

Sampled statistical data can also be combined using Bayesian statistics. As one example binomial counts for numbers of detections, X, out of N tests can be combined with values of α and β for a beta prior distribution determined by an expert to produce a posterior distribution for the probability that would follow a beta distribution with α'= X+α and β' = (N-X)+β. The analyst could then use the posterior distribution to develop an estimate of the probability as well as create tolerance intervals, analogous to confidence intervals developed solely for statistical data.

In some cases data may exist but may only cover a subset of the required values. As an example, transit times for adversaries may be collected during exercises but the values would be influenced by what the adversary is carrying, weather conditions, whether they are being fired upon, and how far they were going. In such cases expert judgement may be used to extrapolate transit times (or transit time distributions) from the data sets that do exist.

Annex E: Path Analysis

Introduction

Path Analysis proceeds, in a general way, to determine measures of effectiveness of a physical protection system based on comparison of an adversary timeline and one or more response timelines.

Path analysis focuses on the measure Probability of Interruption (P_I) as a key measure of PPS effectiveness against an adversary attack (other such measures will be discussed in a later section).

P_I is defined as the probability that the response forces will arrive and deploy in time before the adversary has completed their attack. P_I is calculated using an adversary timeline and a response timeline. The figure below depicts the adversary timeline at the top, indicating the Task Time it takes the adversary to complete all of his tasks, and also the sensing opportunities along the timeline which may cause the adversary to be detected. Below the adversary timeline there is a comparison between the PPS Response Time (PRT) and the Adversary Task Time Remaining on the path after first sensing at each possible sensing opportunity.

If PRT < Adversary Task Time Remaining After First Sensing then the corresponding sensing opportunity is considered timely; if this is not the case, then the opportunity is not timely.[1] P_I is equivalent to the probability that the adversary is detected at least one of the timely sensing opportunities. For the example in Figure E-1, the first two sensing opportunities are timely, so P_I = P(Detection at Sensing Opportunity 1 OR Sensing Opportunity 2). The Critical Detection Point or CDP is the last sensing opportunity on the adversary timeline that is timely, in this case Sensing Opportunity 2.

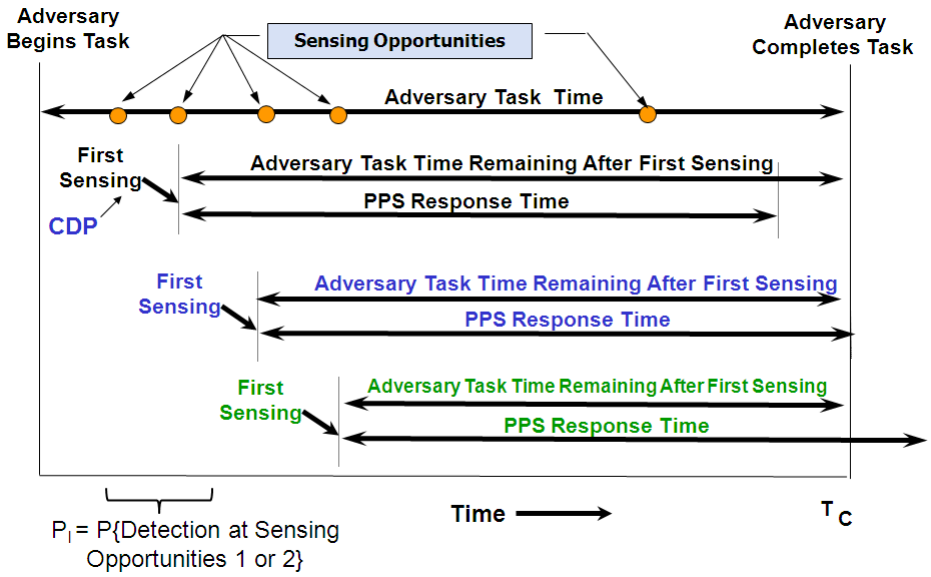

Figure E-1. Relationship between the Adversary Timeline and the Response Timeline

[1] This model is called "timely detection" and not "timely sensing" because the timing for the beginning of the detection process is the sensing event; hence from a timeline perspective timely detection equates to timely sensing.

The discussion below starts with a definition of adversary and response timelines based on a generalization of a path called an adversary action sequence or AAS. This more general abstraction will be used as it can be used to model both insider and outsider attacks and provides a linkage to simulation of an adversary attack plan. The discussion will then present formulas for determining P_I based on the two timelines and will then discuss how path analysis is performed.

Adversary and Response Timelines

The adversary timeline is composed of a sequence of times, each associated with a task that an adversary must complete to accomplish their objective of theft or sabotage. Each time represents how long it would take an adversary to complete that task, given characteristics about the adversary that might be specified within a Design Basis Threat, for example. Thus, the sum of the times represents how much time is required for carrying out all the tasks included in the adversary attack, from the start of the attack in a place where they are not likely to be detected (traditionally termed "offsite" in evaluation tools) until their objective is completed.

The list of tasks making up an adversary attack is defined as the "adversary action sequence." The adversary actions sequence is most generally defined as a time-ordered sequence of n tasks that the adversary has to complete. Typically, an action sequence can be thought of as a detailed plan of what a single adversary team or individual would need to accomplish to effect theft of nuclear or other radiological material, sabotage, or dispersal.

In carrying out the action sequence there are places on the timeline where sensing may occur. Sensing is defined as generation of some anomaly that could be evidence that an unauthorized adversary action is underway. The places on the timeline where sensing may occur are called "sensing opportunities." Each sensing opportunity has an associated probability of sensing, P_S, and an associated probability of assessment, P_A, which is the probability of a correct assessment conditioned on sensing occurring.

Traditionally, it has been assumed that each task has an associated sensing opportunity but this not a necessary assumption about AAS's. For the discussion below assume that there are N tasks, with times $\tau_1, \tau_2, \ldots\tau_N$ and that there are J sensing opportunities, with probabilities of sensing $P_{S1},\ldots, P_{SJ}$; J probabilities of assessment $P_{A1},\ldots,P_{AJ}$; and J probabilities of detection computed as $P_{Dj} = P_{Sj}*P_{Aj}$. To keep the discussion general, we will assume that there is a time T_{Rj} which is the time remaining on the adversary timeline after sensing opportunity j; the time remaining will depend on the task times, τ_n, in a way that will be discussed later.

Once sensing occurs (that is, an alarm is generated or an anomaly is noticed) , there are a set of actions that the response force will perform to counter the adversary; these actions are depicted on a response timeline. These actions will include 1) assessing the alarm/anomaly to determine if it is indeed due to an unauthorized act, 2) communicating with relevant response forces, and 3) deploying those forces to interrupt the adversary before they complete all tasks (see Figure 1). The total time from the alarm being generated (at T=0) until sufficient forces arrive to be able to interrupt (in this case, at T_3) is called the PPS Response Time (or PRT).

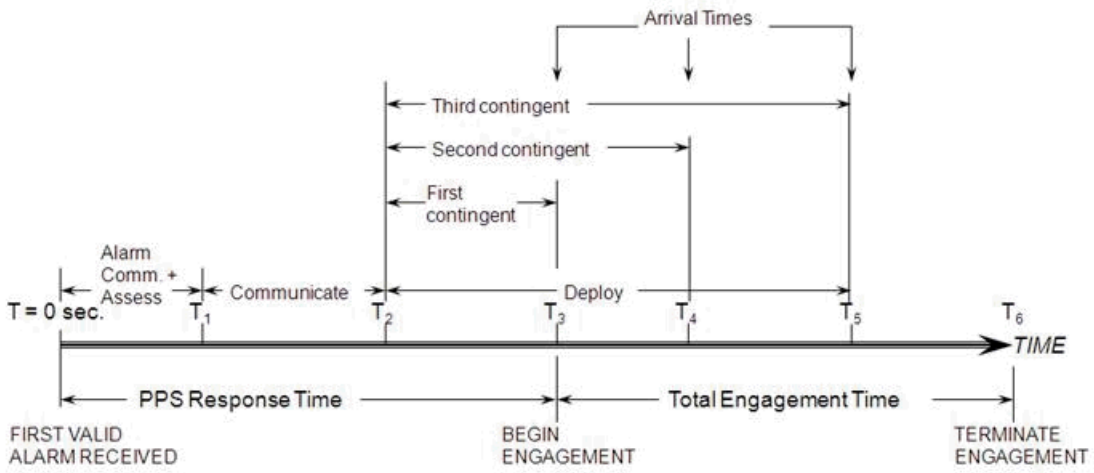

Figure E-2. Arrival times for response forces

In principle, every sensing opportunity may have its own unique Response Timeline (and associated PRT).

It should be noted that in some cases different response forces may arrive at different times; in Figure E-2 forces show up at different times (T_1, T_2, T_3); the forces that show up at each time are called contingents in this figure even if their arrivals at the same time are not coordinated. Three contingents are shown in Figure E-2, resulting in values PRT_1 (which is the PRT shown), PRT_2, and PRT_3.

From an analysis perspective, any of the contingent arrival times could be selected within the facility contingency plan as *the* PRT. Thus, if there are K responding contingents, each sensing opportunity could have K possible different PRTs. The notation would be PRT_{jk} = the kth PPS Response Time associated with detection occurring due to sensing at sensing location j[2].

Other Assumptions and Mathematical Definitions

Tasks can be viewed either generally as activities that are to be completed or more specifically as actions against physical protection measures (such as penetrating a wall or defeating a sensor) or as movement from one point or another. There is no requirement, however, that a task be performed in a particular place; for example a task might be to "learn the combination to the lock" which might occur in any one of a number of places.

The action sequence is assumed to be taking the adversary towards successfully completing the attack so that it is presumed that the state (however the adversary's "state" is defined) at the *end* of the task is "closer" in some sense to the objective than the state at the *beginning* of the task. As an example, the adversary could be physically closer to the target at the end of a transit task than at the beginning.

Delays along the AAS may be caused by armed engagements with guards and response forces.

[2] Am more general model would define PRT_{jkn}, where n is a task on the adversary timeline. This case will not be covered here for a number of reasons but a remark will be made about this topic at the end of this section.

Probabilities, delay times, and PPS response times can be assumed to be point values or to follow distributions.

Probability of interruption: the probability that the response arrives in time to defeat the adversary before the latter complete their AAS (we will show the equation for one contingent, also):

$$P_I = \sum_{j=1}^{J} P_{FDj}\, P\left(T_{Rj} - PRT_j > 0 \right)$$

where P_{FDj} = Probability of First Detection at sensing location j, defined as:

$$P_{FDj} = P_{Dj} \times \prod_{i=1}^{j} (1 - P_{Di})$$

Note: The product on the right is assumed to be equal to 1 when j =1, so $P_{FD1} = P_{D1}$.

Timely Detection: When the Time Remaining, T_{RJ}, and PRT_j are point values, those sensing opportunities, j, for which the Time Remaining, T_{Rj} exceeds PRT_j are said to be timely meaning that if detection occurs at one of those opportunities interruption will successfully occur before the adversary finishes all of their tasks.

Critical Detection Point: When the Time Remaining, T_{RJ}, and PRT_j are point values, the last sensing opportunity in the AAS that is timely is called the Critical Detection Point (CDP) – see Figure E-1. This point is considered critical in the sense that if detection does not occur before or at this opportunity then the adversary cannot be interrupted. An AAS does not necessarily have any timely sensing opportunities so that there may not be a CDP.

Remark: It is typically assumed that all of the sensing opportunities before the CDP are also timely. While the Time Remaining, T_{Rj}, stays the same or decreases further along the AAS, the PRT_j's do not necessarily vary in such a way that all opportunities are timely before the CDP. The only simple sufficient condition for achieving this assumption is that $PRT_j \leq PRT_{CDP}$ for sensing opportunities before the CDP.[3]

Progression from Analysis of Timelines to Path Analysis

Path analysis looks at effectiveness of the physical protection system against paths as opposed to AAS's. A path is a time-ordered sequence of locations that an adversary proceeds to during an attack. The paths may be defined by a sequence of elements and areas from an adversary sequence diagram or by a sequence of actions performed by an insider from an adversary action sequence diagram. The same type of metrics, such as P_I, can be calculated for paths as are calculated for AAS's.

This section will discuss the relationship between paths and AAS's starting with adversary action sequences.

[3] $TR_{CDP} > PRT_{CDP}$, which means that $T_{Rj} > PRT_{CDP}$ for all sensing opportunities before the CDP.

In principle, it should be possible to find the most-vulnerable AAS, defined as an AAS that minimizing the one or more metrics over all possible AAS's from some starting point outside the facility to the target(s). This is impractical for a number of reasons:

- One AAS can differ from another by including different numbers of tasks
- Two AAS's can be identical except that the adversary performs a single task against a single physical protection measure (such as a fence) using different defeat methods (such as cutting through the fence versus climbing over it) or using different tactics, such as force, stealth, and deceit;
- The performance data for an AAS (P_{Dj}, t_n, T_{Rj}, and PRT_j) will change depending on specifically where the adversaries are located, where they are going (and how quickly);
- Performance data for an AAS may vary based on the time(s) of day, operational conditions(s), weather condition(s), etc., under which the AAS is being performed[4].

There are a number of ways of addressing these issues:

- Categorize AAS's by sets of locations that the adversary moves through and perform the search only over each set. To accomplish this, the set of AAS's that proceed through the same set of locations would be said to follow the same path. As an example, an adversary path through an adversary sequence diagram might consist of the adversary: penetrating a fence, crossing a protected area, penetrating a door, crossing a building interior, penetrating a certain wall, crossing a vital area, and sabotaging a pump. This path "includes" all AAS's that go to these locations however defined by the analyst (for example, penetrating a fence might refer to crossing a perimeter fence anywhere along a 3 kilometer boundary). Thus, a large number of AAS's are represented by a set of paths that can be searched to find the one with the lowest Probability of Interruption, etc.
- Determine conservatively (low) estimates of performance metrics by using minimum probabilities of detection and delay times across defeat methods and operating conditions. These minimum values may be chosen by the analyst but they may also be chosen based strategies that the adversary might use (such as minimize detection down to a certain task on the AAS and then minimize delay thereafter).
- Perform analyses for each of several facility states, where the "state" refers to operational condition(s), weather condition(s), etc., and facility targets.

Path analysis, then, includes searching over all paths looking for the one with the lowest P_I, etc. To find the best path, the other two issues need to be addressed. For example, some decision needs to be made about assigning detection and delay times based on all the different defeat methods that the adversary has at each step in the path. Finally, all facility states need to be addressed in some reasonable fashion. These issues will be discussed below.

[4] In some AAS's the tasks cumulatively may extend over long periods of time such as hours or days, resulting in multiple times, states, and weather conditions.

Path Searches Over Networks

Path searches are typically performed over some sort of network that the adversary must pass through. Several networks in use are described below. For example paths can be defined on one type of network called an Adversary Sequence Diagram (ASDs) (see Figure E-3). In this diagram, the long rectangles represent security areas where an adversary can travel while the squares represent security features that the adversary must defeat such as gates (GATs) and Doors (DORs).

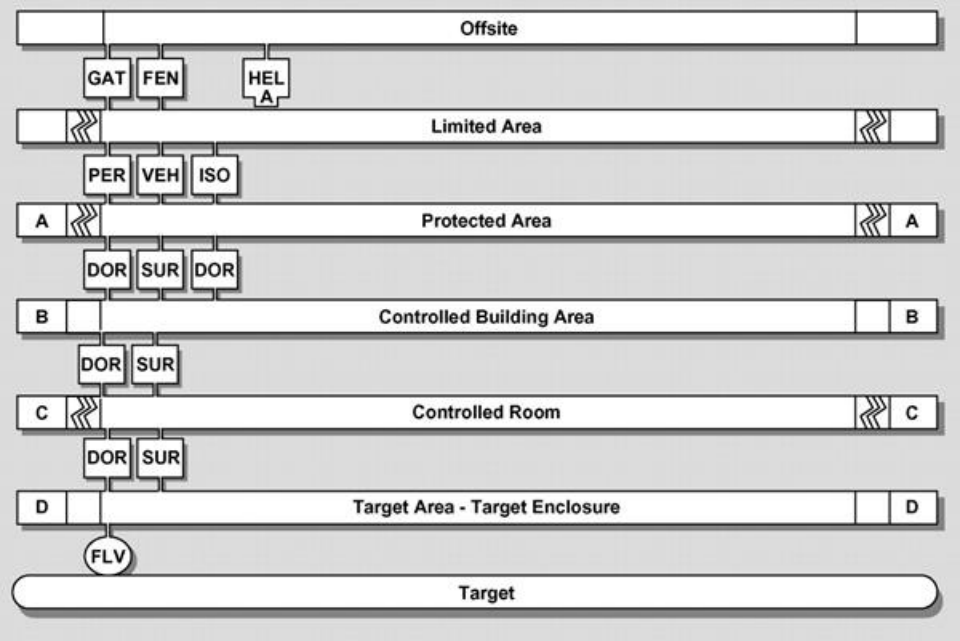

Figure E-3. Example of an Adversary Sequence Diagram

Note that this ASD could be simplified just to show the boundaries of formal security areas, such as limited, protected, inner, and vital areas.[5]

Alternatively, the paths through the facility can be represented on sort of mesh or grid as shown in Figure E-4. Note that the mesh may consist of different types of polygons (such as squares, hexagons, and triangles) and the polygons may be regular (that is with identical sides and angles) or irregular where these sides and angles vary between polygons. The mesh or grid may be two-dimensional or three-dimensional. Two paths could differ merely by going through different grid points even though the physical protection measures that are attacked, such as walls and sensors, are attacked are identical.

There are two important issues that arise with respect to performing path analysis on these networks:

- How does one ensure that the most-vulnerable path (MVP) through the network (from the defenders' concern about, for example, low P_I) is actually identified?

[5] In this example, everything within the controlled room boundary, between it and the controlled building area, might be in a vital area.

- How does the analyst deal with mobile elements of the physical protection system, such as guards and response forces, that might interact with the adversary on the path?

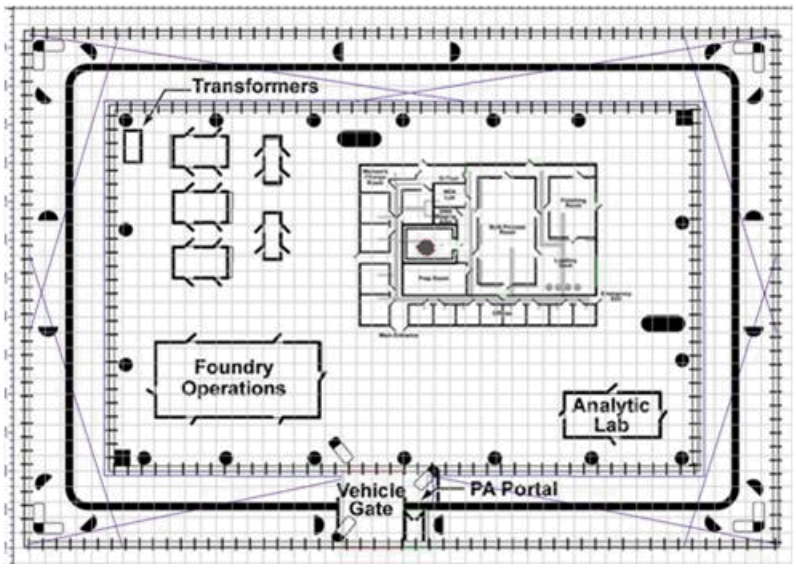

Figure E-4. Example of a Mesh Associated with a Facility

In some cases shortest path algorithms, such as Djikstra's or A* methods, can be used to find the MVP. These algorithms can only be used, however, under certain conditions that must be verified in the underlying model. For example, such algorithms typically require detection probabilities, delay times, and PRT's need to be point values (as opposed to following distributions)[6].

In other cases, though, such algorithms cannot be proven to work. In such cases, one of several approaches can be taken:

- Have the analyst determine the path;
- Keep the network small enough that an algorithm can review all of the paths (as is done with ASD's); and
- Perform some global search method that is likely to give the MVP, such as genetic algorithms.

Determining Worst-Case Probabilities and Times

Even for a single path, there is still the issue of how the adversary performs each task to defeat individual physical protection measures. Several types of decisions come up: For example, if the task is penetrating a fence, do they attempt to climb a fence or cut through it? If they decide to cut through it, what tool might they use, and what delay time against the DBT should be used)? If they use wire cutters to cut through it what is the associated probability of detection?

[6] It is possible to sample probabilities and times from distributions N times and to solve MVP's through N networks treating the values as if they are point values. This approach comes up with N MVP's calculated under slightly different assumptions, which can provide information about how the uncertainty in data affects results. This is different, however, than trying to find the MVP taking the distributions into account.

There are two main ways of making these decisions:

- Use expert judgement: In this case, one or more experts decide what the defeat methods are and what the associated delay times and probabilities are.
- Use information about where the Critical Detection Point (CDP) is on the path. In this approach the analyst selects defeat methods that minimize delay starting at the end of the path until a CDP is found; and then minimizes detection back to the start of the path.

As discussed earlier, the CDP can only be defined when detection probabilities and times are point values.

Addressing Multiple Facility States and Targets

An effectiveness evaluation may need to address effectiveness during each of several facility states, where the "state" refers to operational condition(s), weather condition(s), etc., and different targets. The combination of states and targets actually evaluated are determined by the analyst based on judgement about such factors as: how important the target is, how often the state occurs and whether its occurrence can be predicted, and whether physical protection for one facility state or target can be considered more effective than for another based on expert judgment.

Annex F: Methods for Determining Critical Systems in Evaluations

Introduction

This section discusses several generic methods for determining how to defeat components, subsystems, and entire physical protection systems. These can be classified generally as logic diagrams (which include fault trees, physical protection logic trees, and attack trees) as well as failure mode and effects.

Conceptually, logic diagrams and failure mode and effects are deductive and inductive methods of analysis, respectively. Deductive methods answer the question "how can some system failure state occur" while inductive methods answer the question "what happens if…?" Logic diagrams start with some top-level undesirable event - such as release of nuclear material (as in vital area identification) or failure of the communications systems with the offsite response – and then represent the combinations of component and subsystem events that can cause that top-level event. Inductive methods, such as failure mode and effects, start with some possible event, such as cutting an alarm cable or defeating certain equipment in a vital area and then try to deduce all of the effects that may occur as a result of that event. The two methods are complementary, as logic diagrams can be used to determine what types of adversary defeat methods to apply failure mode and effects analysis to, while failure mode and effects analysis can verify, to a certain extent, the logic diagram.

Logic Diagrams

The logic diagram is a useful tool for evaluating security for a nuclear facility. Several types of logic diagrams will be discussed in this section after a basic discussion of logic diagrams:

- Fault trees
- Physical protection logic trees
- Attack trees

As physical protection logic trees and attack trees are special cases of fault trees[7], fault trees will be covered first.

Overview of Logic Diagrams

The logic diagram is a graphical representation of combinations of events that can result in a specified state or event. For our example, the specified state is a hypothetical release of significant amounts of radioactive material from a notional nuclear plant that causes High Radiological Consequences as a result of sabotage of critical components.

Figure F-1 illustrates the symbols that are used in logic diagrams. The logic diagram shown represents relationships between events. Each event will have a written description in the large rectangle in the logic diagram. A smaller rectangle placed immediately under the description will show the event name or label. Event names are brief and are formed from combinations of letters and numbers.

[7] Note that in practice attack trees and fault trees are often used in different ways (e.g., adversary costs/resources might be associated with nodes in attack trees while probabilities might be associated with leaves in fault trees) but, in principle, probabilities could be applied to attack trees while costs/resources could be applied to fault trees.

The symbols of the logic diagram shown here in Figure F-1 will be discussed in detail. These include symbols for:

- logic gates
- events
- transfer operations

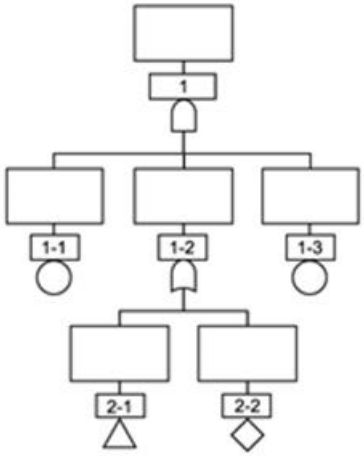

Figure F-1. Logic Diagram Symbols

Logic Gates

Two kinds of logic gates, the AND gate and the OR gate, are used in the logic diagrams. Gates have inputs and may or may not have an output. Inputs enter the bottom of the gate; outputs exit the top of the description rectangle above the gate.

AND Gate

The shape of the AND gate is a round arch with a flat bottom (see Figure F-2). For the undesired event described above the AND gate to occur, all of the events that have an input into the AND gate must occur. Thus, if any one of the input events can be prevented, the event described above the AND gate will be prevented. In this example the top event is generated by an AND gate whose inputs are Events 1-1, 1-2, and 1-3. Thus, the top event will occur if, and only if, Events 1-1, 1-2, and 1-3 all occur.

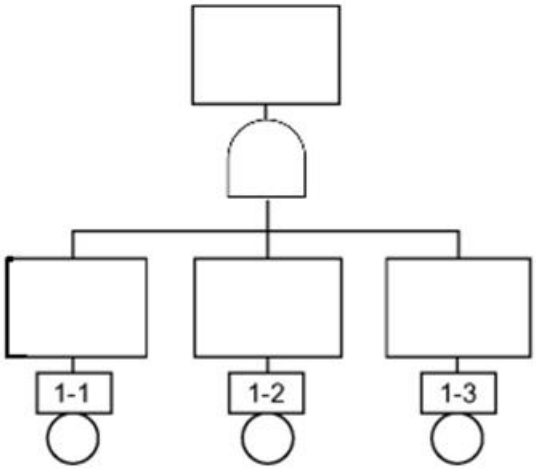

Figure F-2. Example of an AND Gate

Or Gate

The shape of the OR gate is a pointed arch with a curved bottom (see Figure F-3). For the undesired event described above the OR gate to occur, any one (or more) of the events that input to the OR gate must occur. All of the input events must be prevented in order to prevent the event described above the OR gate. For example, in Figure F-3 the top event has an OR gate whose inputs are Events 1-1, 1-2, and 1-3. Thus, the top event occurs if one or more of Events 1-1, 1-2, or 1-3 occur.

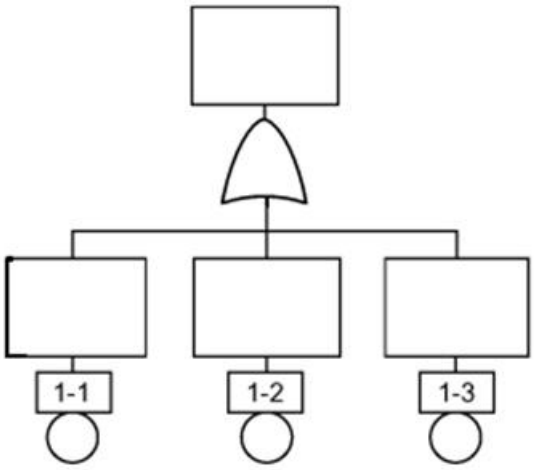

Figure 3. Example of an OR Gate

Events

There are several types of events in logic diagrams. They include:

- end events
- intermediate events
- primary events

End Events

If an event is not used as input to another gate, it is called an end event. Logic diagrams have only one end event, the topmost event of the tree. In Figure 1, Event 1 is the end event. Sometimes this event is also called the treetop.

Intermediate Events

Events that have both inputs and outputs are called intermediate events. In Figure 1, Event 1-2 is an intermediate event.

Primary Events

Events that do not have an input are called primary events. They represent the start of actions that ultimately generate the end event. Two types of primary events are distinguished by the symbol that appears immediately below the name of the primary event: the basic event and the undeveloped event.

The basic event is symbolized by a circle below the rectangle. A basic event can be understood and evaluated qualitatively or quantitatively, depending on the purpose of the analysis, without further development of the event into causes or specific cases. In Figures F-2 or F-3, Events 1-1 and 1-3 are basic events.

The undeveloped event is symbolized by a diamond below the rectangle. An undeveloped event is an event whose causes are insufficiently understood to be included in the logic diagram. For the purpose of evaluation, the undeveloped event is treated as a basic event. The conclusions drawn from the analysis of a tree that contains an undeveloped event are tentative and subject to revision. In Figure F-1, Event 2-2 is an undeveloped event.

Transfer Operation

The transfer operation is represented by an upright triangle. The transfer operation is used to make the graphic display of the logic tree more compact and readable. Since many logic diagrams, as they are developed, occupy a wide left-to-right space across a page, it might be necessary to disconnect the development of an event and place it at a more convenient position on the page or on another page

Fault Trees

The logic diagram is a useful tool for determining the potential theft and sabotage targets for a nuclear facility. One type of logic diagram, called a fault tree, graphically represents the combinations of component and subsystem events that can result in a specified undesired state.

One type of logic diagram, called a fault tree, graphically represents the combinations of component and subsystem events that can result in a specified undesired state. Among other things, the fault tree is a useful tool for determining sabotage targets for a nuclear facility.

For the example discussed here, the undesired state (or event) is a release of significant amounts of radioactive material from the plant as a result of sabotage of critical components. The physical protection system is intended to prevent sabotage of these components. Logic diagrams that are intended to identify the sets of components an adversary would have to sabotage to cause the radioactive release are called

sabotage fault trees. Sabotage fault trees are used for identifying vital area. They describe, in this case, the hypothetical actions an adversary must accomplish to cause sabotage, and they can be used to identify the areas (locations) to be protected in order to prevent sabotage.

Figure F-4 shows the top portion of a sabotage fault tree for a hypothetical pressurized water reactor.

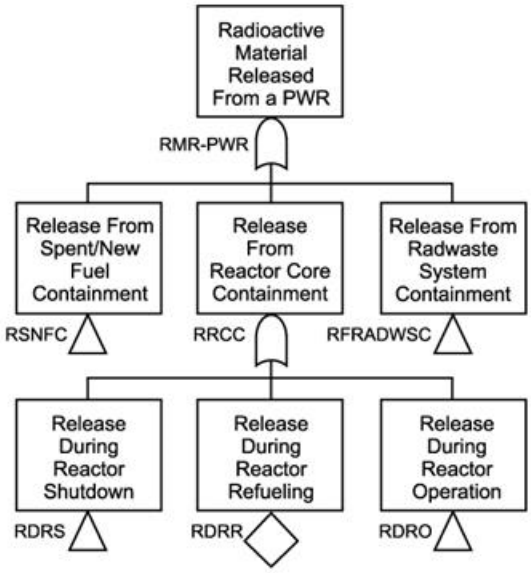

Figure F-4. Top Portion of a Sabotage Fault Tree for a Hypothetical Pressurized Water Reactor

Physical Protection Logic Trees

Physical protection logic trees were developed to depict all the ways for defeating areas and protection layers of physical protection systems. The figure below is a logic tree indicating several ways to defeat a perimeter gate using force, stealth, or deceit. Traditionally, this tree would be part of a "Boundary" logic tree covering all of the ways to defeat the entire facility perimeter. Note that the tree does not yet describe how the lowest-level events are caused: for example, we have not specified how the adversary will attempt to avoid detection going under the gate.

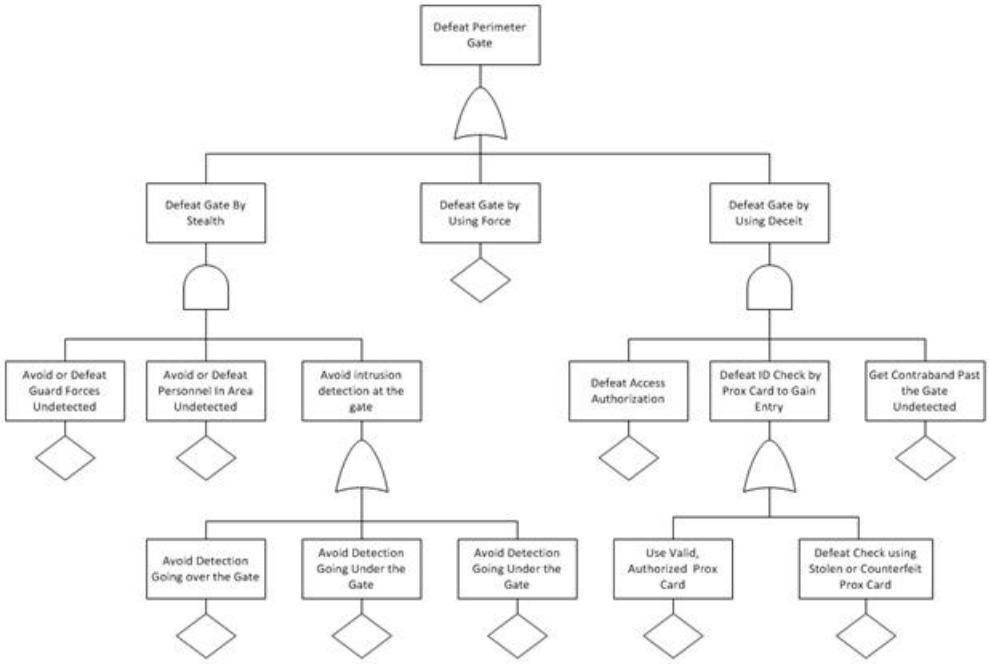

Figure F-5. Example of a Physical Protection Logic Tree

Attack Trees

Attack Trees are logic diagrams showing how some top level event can be caused (in this sense they are deductive like fault trees). The attack tree has a root representing some undesirable event, child nodes that if true make the direct parent true, ultimately resulting in leaf nodes, which are the lowest level nodes in the tree. In the following figure, the top node is "Defeat Prox Card Authentication" which organizes possible ways to attack a hypothetical prox card. To better understand the diagram, note that:

- Any of the child nodes being true causes a parent node in the figure except where an "and" is indicated in the diagram (such as "Obtain PIN" AND "Obtain Card") where all the specified child nodes must be true.

- The coloring in the diagram attempts to convey the same information without having to duplicate subordinate attack trees. For example, the black arrows proceeding into and out of the Target User node indicate that the adversary can obtain the card by targeting the user by threatening, bribing, or blackmailing them; the red lines indicate that the adversary can target the user to obtain the pin the same way but can also shoulder surf or carry out social engineering, which is impractical for obtaining the card.

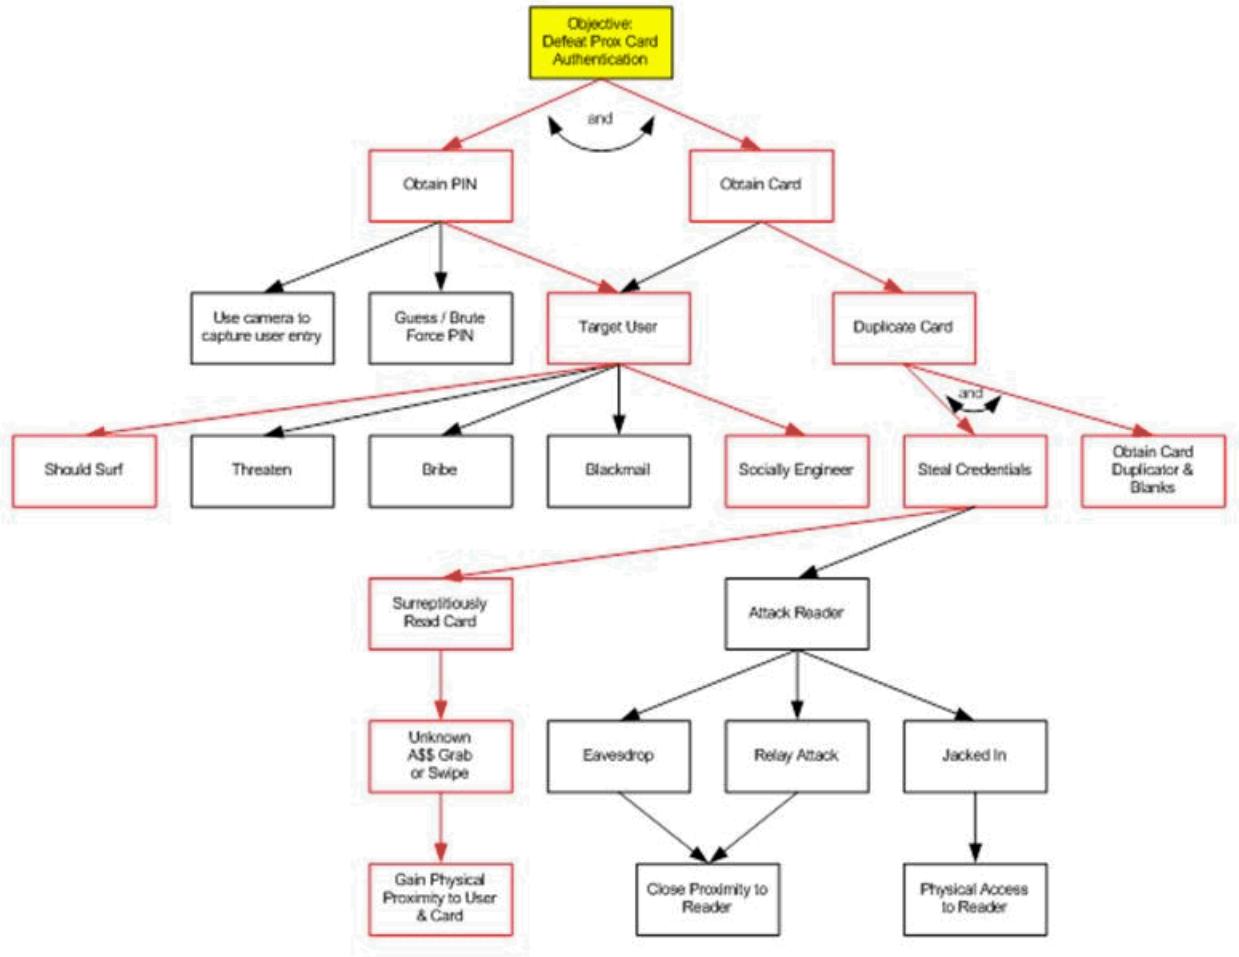

Figure F-6. Example of a Hypothetical Attack Tree to Defeat Prox Card Authentication

Notice that the highest level node in this hypothetical attack tree "Defeat Prox Card Authentication" is one way to cause the lowest-level event "Defeat Check Using Stolen or Counterfeit Prox Card" in the example hypothetical physical protection logic tree displayed above (another way to defeat the check with a stolen prox card is make up a convincing story why the adversary has a prox card but can't get in so that a gullible employee would let the adversary in).

Failure Mode and Effects

To be filled in.

Annex G: Develop Adversary Scenarios

To be developed. Here is some partial content:

Introduction

When developing attack scenarios, scenarios should be chosen to challenge the plant security and operations to the maximum extent practicable within the constraints established by a TA/DBT. The TA/DBT is used as a management and design tool that helps facilitate informed decision-making and establishes technical requirements for security designers.

Thus, the scenarios selected should be those to which the facility or transport is judged to be most vulnerable. In a evaluation, the target set and sabotage scenario should be selected to exercise as many of the deficient aspects of the security program as practicable. Absent such an assessment, target set and scenario selection considerations would include the following:

- Minimization of the number of physical barriers that an adversary must overcome (e.g., selection of target sets with the minimal number of areas).
- Minimization of likelihood that adversary actions would be detected and correctly assessed.
- Minimization of likelihood of timely and effective security response

Identify operational vulnerabilities

In order to identify site vulnerabilities across various operational conditions and states, consider different:

- Operational conditions (operational versus non-operational)
- Target material configurations (reactor load-out versus operations)
- Response force alert levels and personnel “crews”
- Different upgrade packages

Exploiting the Vulnerabilities

When promising vulnerabilities have been identified, it will be required to develop an action plan how the each vulnerability will be exploited. The action plan will need to have the attention to detail and organization in how the attack will be executed. The following steps can be followed (note that the scenario is hypothetical):

- First creating a list of essential tasks that have to be accomplished for the attack based on that vulnerability to succeed. Such a list might look like the following for a target:
    - Task 1: Enter building XYZ
    - Task 2: Collect 20 Kg of U235 in storage containers
    - Task 3: Leave site with material without pursuit by response forces
    - Task 4: Arrive undetected at safe house in city ABC
    - Task 5: Hold off responding units so that tasks 1 through 3 are accomplished

  These tasks should be kept as simple as possible.

- Next, creating sub-plans that describe how one or more teams of attackers can perform each task within resource constraints. These sub-plans should describe:
    - Who is involved?
    - What are they doing as a function of time?
    - How are they performing each step?
    - What equipment are they using?
    - How are they transporting the equipment?

- Finally, combine these sub-plans into a master attack plan/scenario description, adjusting sub-plans to meet overall constraints imposed by the DBT and perhaps the site as well as to achieve synchronization between teams.

Adding Supporting Team Sub-Plans to Scenarios

Supporting teams can be assigned to complete other essential tasks or to aid the main team directly. Often, the remaining tasks look like: "Hold off responding units so ..." or "Neutralize offsite response..." Thus, one good use of supporting teams is to delay or incapacitate the response through setting ambushes, creating diversions, and attempting to confuse the response.

Using Path Analysis for Scenario Development

Path analysis can suggest sub-plans that serve as the main or "direct" part of the attack (direct in the sense of going to the target). Such plans might be based on the minimum delay, minimum P_I, or minimum $P_I * P_N$ paths

Details can be added to these path descriptions to fill out the scenario. For example, instead of the step "Penetrate Fence" found in the path analysis, the scenario description might consist of: "Four adversaries bridge fence using ladder carried in from vehicle parked outside at night during a storm. Last adversary monitors radio traffic."

Of course, multiple scenarios can be developed for a single path by slightly varying the method by which the adversary attacks different protection elements along the path.

Be aware, though, the most-vulnerable path (MVP) from path analysis may be a poor basis for creating a scenario. This may occur because typically low P_I paths should be corrected with upgrades during the path analysis phase. After such upgrades, the MVP should now have a high P_I rendering that path less desirable At this stage scenario analysis might more profitably consider factors not found in path analysis: preventing neutralization and employing other teams to prevent interruption.

Appendix O. Additional Technical Information on Performance Metrics for Timeline Analysis

Metrics are methods for measuring things, in this case some facet of how well the security system performs against an adversary action sequence (AAS). Metrics can be categorized by whether they are determined quantitatively or qualitatively.

Quantitative Measures

There are several metrics that might be calculated for an AAS:

- Cumulative Probability of Detection along AAS up to some sensing location j:

$$P_{DCumulative}(j) = \sum_{i=1}^{j} P_{FDi}$$

- The remaining task time on an AAS from task n to the end (inclusive of task n):

$$T_{Cumulative}(n) = \sum_{i=n}^{N} \tau_i$$

  Note: The delays after each sensing opportunity j were denoted by T_{Rj} earlier.

- Expected time to complete the AAS (with E() as the expectation operator):

$$E(\text{AAS Time}) = \sum_{j=1}^{J} P_{FDj}\, T_{Rj}$$

- Also, expected margin and expected delay deficiency. Note that these expected times can be computed conditionally or unconditionally. For simplicity, we will assume only one response contingent (K=1).

$$E(\text{Margin}) = \sum_{j=1}^{J} P_{FDj}\, E\left(\text{maximum of 0 and } \{T_{Rj} - PRT_j\}\right)$$

$$E(\text{Delay Deficiency}) = \sum_{j=1}^{J} P_{FDj}\, E\left(\text{maximum of 0 and } \{PRT_j - T_{Rj}\}\right)$$

Note that the Expected Margin measures, roughly, the average amount of time between the time remaining on an AAS and the PRT (assuming K=1), as long as the time remaining exceeds the PRT. The Expected Delay Deficiency looks at the average amount of delay between the PRT and the time remaining as long as the former and greater than the latter. These delays can be calculated here in an unconditional way or can be calculated conditionally as follows:

$$E(\text{Margin}|\text{Margin Exists}) = \frac{\sum_{j=1}^{J} P_{FDj}\, E\left(\text{maximum of 0 and } \{T_{Rj} - PRT_j\}\right)}{\sum_{j=1}^{J} P_{FDj}\, P\left(T_{Rj} - PRT_j > 0\right)}$$

$$E(\text{Delay Deficiency}|\text{Delay Deficiency Exists}) = \frac{\sum_{j=1}^{J} P_{FDj}\, E\left(\text{maximum of 0 and } \{PRT_j - T_{Rj}\}\right)}{\sum_{j=1}^{J} P_{FDj}\, P\left(PRT_j - T_{Rj} > 0\right)}$$

These are conditional expectations: for example, E(Margin|Margin Exists) means the expected value of the Margin $T_{Rj} - PRT_j$ given that the Margin "exists", that is, is positive. Conditional expectations will be defined as 0 if the probability in the denominator is 0.

- Probability of interruption: This is the probability that the response arrives in time to defeat the adversary before the latter complete their AAS (we will show the equation for one contingent, also):

$$P_I = \sum_{j=1}^{J} P_{FDj}\, P\left(T_{Rj} - PRT_j > 0\right)$$

Note that this is the denominator for the conditional expectation: Expected Margin given that the Margin Exists.

- Probability of System Effectiveness, P_E:

  For this, define P_{Nj} = Probability of Neutralization, given detection occurs at sensing location j.

$$P_E = \sum_{j=1}^{J} P_{FDj} P_{Nj}$$

  Neutralization is defined as occurring if the response force prevents the adversary from completing the attack, either by driving them away, arresting them, or killing them.

  Conceptually, P_{NJ} is the probability that the response would neutralize the adversary in an exercise assuming first detection of the adversary occurred due to an alarm at sensing opportunity j.

Qualitative Models

Some of the metrics that can be expressed quantitatively also have an equivalent qualitative version. A common approach is to probabilities are treated qualitatively while delays and PRT's are expressed quantitatively). Qualitative probabilities are typically assigned as "Very High" down to "Very Low."

For example:

- One can speak about assigning a qualitative probability of sensing and probability of assessment for sensing location j;

- A qualitative probability of detection at a sensing opportunity j would then be some function of the qualitative probability of sensing and a qualitative probability of assessment.
- Instead of trying to calculate $P(T_{Rj} - PRT_j > 0)$, a qualitative approach might be to assign a qualitative interruption score, indicating whether the response should be able to interrupt the adversary somewhere along the AAS, given detection at sensing opportunity j. This qualitative interruption score may be based on knowledge about how the time remaining, T_{Rj}, compares to PRT_j, but also may factor in whether it is easy for the response to find the adversary, for example, an issue not explicitly considered when determining quantitative P_I's.
- For neutralization, a qualitative score for neutralization can be assigned assuming interruption has occurred. That is, given that sensing location j generated an alarm that caused a positive assessment AND the response arrived in time to interrupt, how likely is it that the adversary was neutralized?
- The contribution of sensing location j would then be a function of the qualitative scores for sensing, assessment, interruption, and neutralization (given interruption) determined for sensing location j. This would contrast with the qualitative model where the contribution of sensing location j would be $P_{FDj}P_{Nj}$.

Remark: The quantitative P_E formula *could* include terms of the form:

P{Neutralization given T_{Rj}-PRT_j >0 AND detection at sensing opportunity j}

but this is typically impractical. Instead, a different approach is typically taken, which will be discussed shortly.

Qualitative $P_{Dcumulative}$, $P_{Interruption}$, and P_E can be determined using the metrics for each sensing opportunity as long as there are rules for assigning qualitative probabilities P(A or B) and P(A and B) based on qualitative P(A) and P(B) for events A and B, respectively. At the same time, it is impractical to combine qualitative probabilities with numerical times so there are no qualitative versions of Expected Time Remaining, Expected Margin, or Expected Deficiency.

Timely Detection

When delay times and PPS Response Times are point values then some sensing opportunities are timely in the sense that if detection occurs at one of those opportunities then interruption will successfully occur before the adversary finishes all of their tasks. One can then define P_E as

$$P_E = \sum_{\substack{j=1,\ldots,J \\ \text{that are timely}}} P_{FDj}P_{Nj}$$

This formula does match more closely with qualitative methods for setting P_E as P_{Nj} is now only counted if interruption occurs.

Appendix P: Examples of Mathematical Models That Could be described in NUSAM

Several categories of mathematical models may be used for Modeling and Simulation of Physical Security during effectiveness evaluations.

1. Consequence Models (Note: this is a list of standard models and software, along with references rather than detailed discussions)
   - Effects of Sabotage
   - Radiological Release Effects
   - Chemical Release Effects

2. Security Performance models
   - Sensing
   - Assessment
   - Communications // Communication and Propagation Model, Site Coverage Model
   - Delay
   - Response times, including detection times and force deployment time

3. Timeline Models: See the section on Path Analysis in Appendix E as well as Appendix O: Analysis Additional Technical Information on Performance Metrics for Timeline Analysis
   - Adversary timelines
   - Response timelines
   - Relation of path analysis to timeline analysis

4. Aspects of Weapon Performance and Effects
   - Direct fire weapons: use of Probability of Hit/Probability of Kill models, including factors involved such as distance and exposure
   - Treatment of area-kill weapons such as grenades
   - Other factors: Treatment of ammunition loadings, reload times, and firing rates as a function of distance
   - Blast effects for vehicle bombs, IEDs (ATF car bomb chart)
   - Stand-off weapons

5. Physical Aspects of the Site and Facility
   - Terrain, including vegetation
   - Modelling of Buildings and construction
   - Modelling of roads and infrastructure
   - Network models of the site, including Adversary Sequence Diagrams and polygon meshes

Note: The following categories of models are used in simulations

6. Vision
    - Line-of-Sight including treatment of terrain
    - Resolution: handling of detection, classification, identification (including friend/foe)
    - Daylight versus night-time conditions, including effects of lighting
    - Viewing objects in different states, such as
        - Combatants: standing, kneeling, prone positions
        - Vehicles: Viewed from different angles and traveling at different speeds

7. Hearing
    - Types of events that are heard
    - Factors including in the model, such as directionality

8. Movement
    - Speeds over different terrain for personnel and vehicles

9. Aspects of Entities
    - General
        - Exposure to weapon as a percentage of total exposure possible
    - Personnel
        - Types of states of health, including combat ineffectiveness
        - Treatment of fatigue
        - Suppression
        - Treatment of rules of engagement and force-continuum rules
    - Vehicles
        - States of health, including disablement
        - Ability to mount/dismount entities