DRAFT

Report #04-006

OPTIMAL RESOURCE ALLOCATION FOR SECURITY IN RELIABILITY SYSTEMS

Azaiez, N. & Bakir, N.

CREATE REPORT

Under FEMA Grant EMW-2004-GR-0112

November 25, 2005

Center for Risk and Economic Analysis of Terrorism Events
University of Southern California
Los Angeles, California

Optimal Resource Allocation for Security in Reliability Systems

CREATE Report

November 3, 2004

Naceur Azaiez

Industrial Engineering Department

King Saud University

Vicki M. Bier

Department of Industrial and Systems Engineering

University of Wisconsin-Madison

Abstract

Recent results have used game theory to explore the nature of optimal investments in the security of simple series and parallel systems. However, it is clearly important in practice to extend these simple security models to more complicated system structures with both parallel and series subsystems. The purpose of this paper is to begin to address this challenge. While achieving fully general results is likely to be difficult, and may require heuristic approaches, we are able to find closed-form results for systems with moderately general structures, under the assumption that the cost of an attack against component any given increases linearly in the amount of defensive investment in that component. These results have interesting and sometimes counterintuitive implications for the nature of optimal investments in security.

Acknowledgment

This material is based upon work supported in part by the U.S. Army Research Laboratory and the U.S. Army Research Office under grant number DAAD19-01-1-0502, by the U.S. National Science Foundation under grant number DMI-0228204, and by the United States Department of Homeland Security through the Center for Risk and Economic Analysis of Terrorism Events (CREATE) under grant number EMW-2004-GR-0112. Any opinions, findings, and conclusions or recommendations expressed herein are those of the author and do not necessarily reflect the views of the sponsors.

1. Introduction

Most past applications of game theory and similar approaches to defense against intentional threats to security have dealt either with components in isolation (Major, 2002; Woo, 2002, 2003; O'Hanlon et al, 2002), or with simple series and parallel systems (Bier and Abhichandani, 2003; Bier et al., 2004). In the reliability area, Levitin and colleagues have by now amassed a large body of work applying reliability analysis to problems of security; see for example Levitin (2002, 2003a, 2003b), Levitin and Lisnianski (2000, 2001, 2003), and Levitin et al. (2003). Much of this work combines reliability analysis with optimization, to identify the most cost-effective risk reduction strategies; however, the threat is usually assumed to be static, rather than responding in an adaptive way to the defenses that have been implemented.

In the real world, however, we will frequently be concerned about protecting the functionality of complex systems with arbitrary structures from adaptive threats. For example, we may be concerned about preserving the functionality of electricity transmission and distribution systems, about protecting nuclear power plants against terrorist attacks or sabotage, or about ensuring the existence of a viable transportation route from one major city to another, in situations where potential attackers may be able to observe some or all of our defenses and adapt their strategies accordingly. Thus, it is important to extend the existing results to address more complex situations.

There are in principle several ways to do this. At the most general, investments in the security of the various components of the system by the defender change the function giving the success probability of an attack on that component as a function of the level of effort expended by the attacker. In response to this function, the attacker then determines the level of effort to be expended on attacking each component (and hence the success probabilities of those attacks). However, a variety of simplifications to this general model are possible. For example, one could assume that the level of effort expended by the attacker on each component to be attacked is a constant, and hence investments by the defender change only the success probability of an attack on each component. Alternatively, one could hold constant the success probability of an attack on each component. In this case, defensive investments could be interpreted as increasing the cost or level of effort that the attacker would need to expend in order to achieve that probability of success.

Here, we adopt this latter approach, and assume that the defender attempts to deter attacks by making them as costly as possible to the attacker. For example, cost could be measured in terms of higher levels of technology needed in order to mount an attack, or an increased probability of attackers being captured. This particular approach to simplifying the problem discussed above clearly is not fully general. For example, there may be attack strategies (such as computer attacks) where the cost to the attacker is essentially negligible; it may be more plausible to model such situations by assuming that defensive investments reduce the success probabilities of attacks, rather than increasing their cost. Moreover, the assumption that the defender wishes to maximize the cost of a least-cost attack is only a proxy for the goal of deterring attacks, since it could for example yield wasteful solutions in which the attack costs resulting from the defender's investments vastly exceed the available resources of any possible attacker. However, we believe that the model of deterring attacks by maximizing the cost of launching an attack is nonetheless reasonable in some circumstances, especially where the defender does not have good information about the resources available to the attacker(s).

Section 2 of this paper presents some results from the literature on least-cost failure diagnosis, which will be extended and adapted to model least-cost attack strategies. Section 3 then extends those results to systems with more general structures than those discussed in the prior work. Section 4 uses those results to characterize optimal attack strategies, and the corresponding optimal defense strategies. Finally, Section 5 discusses the conclusions of our work, and presents some directions for future work.

2. Results of Prior Work

As stated above, the approach used here models optimal attack strategies by analogy with existing results for *least-expected-cost failure-state diagnosis* of reliability systems. In this section, we discuss the least-cost diagnosis problem, and summarize the existing results of interest to the current study.

Consider a reliability system, the components of which are to be tested sequentially in order to identify the state of the system (operating or failed). A cost is incurred for testing each component of the system. The initial failure probability of each component (before testing) is known, as well as the system configuration. The

problem is to determine the optimal inspection procedure for identifying the system state at minimal expected inspection cost.

Butterworth (1972) solved the problem for simple series and parallel systems, and established a sufficient condition for optimality of diagnosis strategies in *k-out-of-n* systems. Halpern (1974) developed an optimal sequential testing procedure for *k-out-of-n* systems with equal testing costs for all components, and Halpern (1977) gave results for *series-parallel* and *parallel-series* systems. A series-parallel system consists of several stages in series with each other, where each stage consists of one or more components in parallel; conversely for parallel-series systems. Ben-Dov (1981) developed an optimal procedure for *k-out-of-n* systems with general costs, and provided another proof of the result for *series-parallel* and *parallel-series* systems from Halpern (1977). Cox et al. (1989) extended the optimal sequential inspection problem to the *minimum-expected-cost classification problem*, by introducing a discrete-valued "classification function," which corresponds to the "structure function" of the system in the special case of the least-cost inspection problem. In particular, Cox et al. (1989) suggested three heuristic procedures for solving minimum-expected-cost classification problems, and showed that for some important special cases (*series-parallel*, *parallel-series*, and *k-out-of-n* systems), one or more of those heuristics produce an optimal solution. Finally, Cox et al. (1996) extended the optimal sequential inspection problem to include situations in which both the system structure and the component failure probabilities are uncertain, and are characterized by probability distributions.

We now state the results for optimal inspection of series and parallel systems. The existing results for *series-parallel* and *parallel-series* systems are not given here, as they are special cases of the more general results to be established in the next section, and are not needed for the development of those results.

To begin, consider a series system of n independent components. Let the testing procedure be such that component $i+1$ is tested only if component i is found operational, for all components $i=1, 2\ldots n-1$. Also, assume that the testing cost of component i is c_i, and the failure probability of component i is q_i, and let $p_i = 1 - q_i$. Then, the following result holds; see for example Ben-Dov (1981).

Theorem 2.1

In a series system, testing components *i=1, 2…n* in sequential order is optimum (in the sense that it minimizes expected testing cost) if and only if:

$$c_1 / q_1 \leq c_2 / q_2 \leq \ldots \leq c_n / q_n \qquad (1)$$

In this case, the expected testing cost is given by

$$C = c_1 + \sum_{i=2}^{n} [\prod_{j=1}^{i-1} p_j] c_i \quad \square \qquad (2)$$

Now, consider a parallel system, and assume that the testing procedure is such that component *i+1* is tested only if component *i* is found failed, for all components *i=1, 2…n-1*. Then, using the same notation as above, the following result also holds; see again Ben-Dov (1981).

Theorem 2.2

In a parallel system, testing components *i=1, 2…n* in sequential order is optimum (in the sense that it minimizes expected testing cost) if and only if:

$$c_1 / p_1 \leq c_2 / p_2 \leq \ldots \leq c_n / p_n \qquad (3)$$

In this case, the expected testing cost is given by

$$C = c_1 + \sum_{i=2}^{n} [\prod_{j=1}^{i-1} q_j] c_i \quad \square \qquad (4)$$

3. Extension to Systems with More General Structures

Real-world systems for which reliability is important often involve complex combinations of series and parallel subsystems. Therefore, it is important to generalize the results given above to more general combined series/parallel systems. Here, we restrict our attention to systems of independent components that can be represented "without replications," in the sense defined by Azaiez and Bier (1995); that is, systems that can be represented using only AND/OR logic in such a way that each component appears only once. This implies, among other things, that all parallel subsystems must satisfy *one-out-of-n* success logic (i.e., that the success of any single parallel train in a subsystem must be sufficient for success of the entire subsystem). Thus, for example, *k-out-of-n* systems (for $1 < k < n$) will not be considered here, since in any representation of such systems using only AND/OR logic, at least one component will appear more than once.

Figure 1 below presents an example of a combined series/parallel system that can be represented with no replications. Note that the only operations involved in

constructing such combined series/parallel systems are placing subsystems and/or simple components in series and/or in parallel with each other. The *series-parallel* (respectively, *parallel-series*) systems discussed by Ben-Dov (1981) and Cox et al. (1996) are special cases of such systems. Before proceeding to the main result, some definitions similar to those in Azaiez (1993) are introduced.

3.1 *Definitions*

We now introduce the following definitions:

1. A subsystem S is called a series (parallel) subsystem with constituents $S_1 \ldots S_n$ (for $n > 1$) if S can be obtained by placing $S_1 \ldots S_n$ in series (in parallel).
2. A series (parallel) subsystem S is called a maximal series (parallel) subsystem if no other subsystems of the entire system can be obtained by placing additional components or subsystems in series (parallel) with S.
3. The constituents $S_1 \ldots S_n$ of a series (parallel) subsystem S are called the basic constituents of S if none of them is itself a series (parallel) subsystem. That is, each basic constituent S_i of a series subsystem must be either a simple component or a parallel subsystem, and conversely for the basic constituents of a parallel subsystem.

For instance, the system represented in Figure 1 is a maximal series subsystem whose basic constituents are subsystem S_3 and component 5. In turn, S_3 is a maximal parallel subsystem whose basic constituents are subsystem S_2 and component 4, and so on.

It follows that every series (parallel) subsystem has a unique set of basic constituents. Also, any system with more than one component must be either a maximal series or a maximal parallel subsystem. Next, we provide an "initialization algorithm" that will be used to derive the optimal testing policy of an arbitrary combined series/parallel system.

3.2 *Initialization Algorithm*

Consider a combined series/parallel system that can be represented with no replications, as discussed above. The following algorithm is used to order the basic constituents of all subsystems of such a system, prior to identifying the optimal inspection policy.

Step 1

Consider any maximal series subsystem S for which all the basic constituents $S_1 \ldots S_n$ are simple components. Let c_i be the testing cost of component S_i, and let p_i and q_i be the success and failure probabilities of component S_i, respectively, for all $i=1 \ldots n$. Then, do the following:

1. Reorder and re-label the components (if necessary) so that condition (1) above holds. We say that $S = (S_1 \ldots S_n)$ is now ordered.
2. Since the expected cost of testing the series subsystem $S = (S_1 \ldots S_n)$ is given by equation (2) above, set $C(S) = c_1 + \sum_{i=2}^{n} [\prod_{j=1}^{i-1} p_j] c_i$.
3. Set $P(S) = \prod_{i=1}^{n} p_i$ and $Q(S) = 1 - P(S)$ to be the success and failure probabilities of subsystem S, respectively.

Similarly, for any maximal parallel subsystem S for which all the basic constituents $S_1 \ldots S_n$ are simple components, and using the same notation as above, do the following:

4. Reorder and re-label the components (if necessary) so that condition (3) above holds. We say that $S = (S_1 \ldots S_n)$ is now ordered. Whenever S is to be tested, it should be tested sequentially according to the established order, such that component S_{i+1} is tested only if component S_i is found failed.
5. Since the expected cost of testing the parallel subsystem $S = (S_1 \ldots S_n)$ is given by equation (4) above, set $C(S) = c_1 + \sum_{i=2}^{n} [\prod_{j=1}^{i-1} q_j] c_i$.
6. Set $Q(S) = \prod_{i=1}^{n} q_i$ and $P(S) = 1 - Q(S)$ to be the failure and success probabilities of subsystem S, respectively.

If the entire system is now ordered (i.e., if all maximal series or parallel subsystems are now ordered according to steps 1-3 or 4-6, as appropriate), then stop. Else, go to step 2.

Step 2

Consider each non-ordered maximal series (respectively, parallel) subsystems $S = (S_1 \ldots S_n)$ in which all basic constituents are either ordered subsystems or simple components. If any basic constituent S_i is a simple component, then let $C(S_i)$ be the testing cost of S_i, and $P(S_i)$ and $Q(S_i)$ be the success and failure probabilities of S_i, respectively.

For each maximal series subsystem $S = (S_1 \ldots S_n)$ in turn, do the following:

7. Reorder and re-label the basic constituents (if necessary) so that the following condition holds:

$$C(S_1)/Q(S_1) \le C(S_2)/Q(S_2) \le \ldots \le C(S_n)/Q(S_n) \quad (5)$$

We say that $S = (S_1 \ldots S_n)$ is now ordered.

8. Set $C(S) = C(S_1) + \sum_{i=2}^{n} [\prod_{j=1}^{i-1} P(S_j)] C(S_i)$. (6)

9. Set $P(S) = \prod_{i=1}^{n} P(S_i)$ and $Q(S) = 1 - P(S)$ to be the success and failure probabilities of S, respectively.

Similarly, for each maximal parallel subsystem $S = (S_1 \ldots S_n)$, do the following:

10. Reorder and re-label the basic constituents (if necessary) so that the following condition holds:

$$C(S_1)/P(S_1) \le C(S_2)/P(S_2) \le \ldots \le C(S_n)/P(S_n) \quad (7)$$

We say that $S = (S_1 \ldots S_n)$ is now ordered.

11. Set $C(S) = C(S_1) + \sum_{i=2}^{n} [\prod_{j=1}^{i-1} Q(S_j)] C(S_i)$. (8)

12. Set $Q(S) = \prod_{i=1}^{n} Q(S_i)$ and $P(S) = 1 - Q(S)$ to be the failure and success probabilities of S, respectively.

Repeat step 2 as needed until all subsystems have been ordered.

END.

3.3 *Initialization Example*

We now apply the above algorithm to the system given in Figure 1, with the following data. Let the inspection costs be given by: c_1=10; c_2=12; c_3=7; c_4=6; and

c_5=10. Also, let the success probabilities be given by: p_1=0.7; p_2=0.8; p_3=0.67; p_4=0.6; and p_5=0.9.

For step 1, we consider the maximal parallel subsystem S_1 formed by components *2* and *3*. The ratios of cost to success probability are 15 and 10.45 for components *2* and *3*, respectively. Based on step 1 of the initialization algorithm, we need to renumber the components so that constituent S_1^1 of subsystem *1* will be component *3*, and constituent S_2^1 of subsystem *1* will be component *2*. To distinguish the subsystem before and after ordering, in this example we will use $\vec{S}^1$ to denote the subsystem after ordering, so that the ordered subsystem will be given by $\vec{S}^1 = (3, 2)$. The corresponding expected cost is given by $C(\vec{S}^1) = c_3 + q_3 c_2 = 10.96$. Also, the failure probability is given by $Q(\vec{S}^1) = q_2 q_3 = 0.066$ and the success probability is therefore $P(\vec{S}^1) = 1 - Q(\vec{S}^1) = 0.934$.

In step 2, we begin by considering the series subsystem S^2 made up of component *1* and ordered subsystem $\vec{S}^1$; i.e., $S^2 = (1, \vec{S}^1)$. Set $C(1) = c_1 = 10$ and $P(1) = p_1 = 0.7$. The ratios of cost to failure probability are 33.3 and 166.1 for component *1* and ordered subsystem $\vec{S}^1$, respectively. Therefore, the ordered series subsystem $\vec{S}^2$ is given by $\vec{S}^2 = (1, \vec{S}^1)$. The corresponding expected cost of the ordered subsystem $\vec{S}^2$ is given by $C(\vec{S}^2) = C(1) + P(1)\, C(\vec{S}^1) = 17.67$. Also, the success and failure probabilities of $\vec{S}^2$ are 0.65 and 0.35, respectively.

We next consider the parallel subsystem $S^3 = (\vec{S}^2, 4)$ made up of ordered series subsystem $\vec{S}^2$ and component *4*. The ratios of cost to success probability are 28.4 and 10.0 for ordered subsystem $\vec{S}^2$ and component *4*, respectively. This yields the ordered subsystem $\vec{S}^3 = (4, \vec{S}^2)$. The failure probability of $\vec{S}^3$ is given by $Q(\vec{S}^3) = q_4\, Q(\vec{S}^2) = 0.14$, leading to a success probability of $P(\vec{S}^3) = 0.86$. The corresponding expected cost of the ordered subsystem $\vec{S}^3$ is given by $C(\vec{S}^3) = C(4) + Q(4)\, C(\vec{S}^2) = 13.07$.

Finally, we consider the entire system $S = (\vec{S}^3, 5)$. Set $C(5) = c_5 = 10$ and $P(5) = p_5 = 0.9$. The ratios of cost to failure probability are 94.7 and 100.0 for ordered subsystem $\vec{S}^3$ and component *5*, respectively. Therefore, *S* is already ordered. However, to distinguish between the initial version of system *S* and the

ordered one, we will as above denote the latter by $\vec{S}$. Moreover, the expected testing cost for the ordered system is given by $C(\vec{S}) = C(\vec{S}^3) + P(\vec{S}^3)\, c_5 = 13.07 + 0.86\ (10) = 21.67$. Also, the system success probability is $P(\vec{S}) = P(\vec{S}^3)\, p_5 = 0.78$. □

3.4 *Optimal Inspection Policy*

The following lemma, for which the proof is omitted, is a natural extension of the lemma stated in Ben-Dov (1981). Moreover, its proof relies on the same basic argument used in the induction proof given in Ben-Dov (1981).

Lemma 3.1

Consider any ordered series or parallel subsystem $S = (S_1 \ldots S_n)$. Then in order to minimize the expected testing cost, testing of any basic constituent S_i must be performed to completion before moving on to testing of another basic constituent with a subscript higher than i. □

We are now ready to establish the main result of this section.

Theorem 3.1

Consider a combined series/parallel system S, ordered according to the *initialization algorithm*. Then, the optimal testing policy that minimizes the expected testing cost is to follow the orderings specified in the initialization algorithm. Moreover, if a basic constituent S_i^j of subsystem $S^j = (S_1^j \ \ldots \ S_n^j)$ is to be tested, then it should be tested to completion before moving on to testing of basic constituent S_{i+1}^j of that subsystem (or testing of some other subsystem), if needed. In this case, the optimal expected testing cost of the system will equal $C(S)$, as computed in the above algorithm.

Proof

The proof is established by induction on the number of simple components (i.e., the cardinality of S), denoted $|S|$. If $|S| = 2$, then S is either a series or a parallel system of two components. The result holds for any series or parallel system by Theorems 2.1 and 2.2. Assume now that the result holds if $|S| = k\text{-}1$. Let $|S| = k$, and

let $S = (S_1 \ldots S_n)$ be the representation of S in terms of its basic constituents. If $n = k$ (i.e., all basic constituents are simple components, and therefore S is either a series or a parallel system), then the result holds for the reason given above. Otherwise, some basic constituent (say, S_i) is not a simple component. Certainly, we know that $|S_i| < k$. Therefore, if S_i is to be tested, then by the induction hypothesis it should be tested to completion, and in the order specified in Theorem 3.1.

Now, consider a virtual system $\tilde{S} = (S_1, S_2 \ldots S_{i-1}, \tilde{s}, S_{i+1} \ldots S_n)$ that consists of the same basic constituents as S, except that subsystem S_i has been replaced by a simple component $\tilde{s}$ with testing cost $c = C(S_i)$ and success probability $p = P(S_i)$ (computed according to the initialization algorithm). From the lemma, testing S optimally is equivalent to testing $\tilde{S}$ optimally. However, we know that $|\tilde{S}| < k$. Therefore, the induction hypothesis specifies that $\tilde{S}$ should be tested as in the algorithm. If $\tilde{s}$ requires testing, then this is equivalent to testing S_i to completion in the order specified in the algorithm. It follows that S will also be tested optimally as in the algorithm, and consequently the optimal testing cost will equal $C(S)$ as computed in the algorithm. □

Note that the series-parallel and parallel-series systems given in Ben-Dov (1981) are special cases of the combined series/parallel systems considered here.

3.5 *Optimal Inspection Policy for Example*

The optimal testing procedure for the system discussed in Section 3.3 and shown in Figure 1 is as follows:

- Test $\vec{S}^3$ first. If it is found to be failed, then conclude that the system S is failed. Otherwise, test component *5*. If component *5* is failed, then S is failed; otherwise, S is operating.
  - To test $\vec{S}^3$, start by testing component *4*. If it is found to be operating, then conclude that $\vec{S}^3$ is operating. Otherwise, test $\vec{S}^2$.
    - To test $\vec{S}^2$, first test component *1*. If it is failed, then conclude that $\vec{S}^2$ (and therefore $\vec{S}^3$ and the entire system S) are failed. Otherwise, test $\vec{S}^1$.

- To test $\vec{S}^1$, test component *3* first. If it is found to be operating, then $\vec{S}^1$, $\vec{S}^2$, and $\vec{S}^3$ are also operating. Otherwise, test component *2*. If failed, then $\vec{S}^1$, $\vec{S}^2$, $\vec{S}^3$, and $\vec{S}$ are failed; otherwise, $\vec{S}^1$, $\vec{S}^2$, and $\vec{S}^3$ are operating.

One can easily show using the *initialization algorithm* that the expected testing cost of the above procedure is 21.67.

4. Optimal Attack/Defense Strategies

We now shift our discussion to the primary question of interest in this paper; namely, identification of optimal attack strategies, and corresponding optimal defenses. We therefore replace the terminology of inspection policies by the terminology of attack strategies. In the context of attack strategies, the costs will be the costs to the attacker of launching attacks on the various components of a targeted system, and the failure (respectively, success) probabilities will refer to failure (respectively, survival) of those components after being attacked.

We assume that each component can be attacked at most once. This might be a reasonable assumption, for example, if a component that survives an initial attack is extremely likely to survive all subsequent attacks; in that case, after a failed attack on one component, an attacker would prefer to attack other components, even if they were initially less attractive than the first one. This assumption is not as restrictive as it would initially appear. For example, the cost c_i and success probability p_i of an attack on component i could be specified to reflect not the cost and success probability of a single attack, but rather the cumulative cost and cumulative success probability of a large number of repeated attacks against the same component. (This approach to specifying the costs and success probabilities of attacks to reflect the effects of multiple attacks is still less than fully general, however. That is because the structure of our model essentially assumes that all attacks against any given component occur consecutively. Thus, our model would not allow an attacker to target component *1*, move on to try an attack against component *2* if the first attack fails, and then come back to again target component *1* if the attack against component *2* also fails.)

We define a feasible attack policy to be one that continues until either the system is disabled, or the attacker discovers (via failed attempts on components in all minimal cut sets of the system) that it will be impossible to disable the system. We assume that the attacker objective is to determine the feasible attack policy with minimum expected cost. In this case, the initial optimal attack strategy (before any defensive investments have been undertaken) will be analogous to the optimal inspection strategy given in Theorem 3.1, under the same assumptions about the system; i.e., that the system can be represented in a combined series/parallel configuration with no replications (as explained above), and that attacks on each component succeed or fail independently of the results of attacks on other components. (Note that the above assumptions will hold throughout the remainder of this paper.)

Under these assumptions, and by analogy with Theorem 3.1, the optimal least-cost attack policy for an ordered series (parallel) path will consist of attacking basic constituent S_{i+1} only if an attack on basic constituent S_i fails (succeeds). In an ordered system, basic constituent S_i will be attacked before any basic constituent S_j with $j>i$, so we will say that S_i is "more attractive" to the attacker than S_j. This concept can also be generalized to components and/or subsystems not necessarily belonging to the same series or parallel path. In particular, if in an optimal attack strategy one component or subsystem will be attacked before another component or subsystem, we will say that it is "more attractive." In this context, "ordered" will mean from most attractive to least attractive.

Note that in a series subsystem, "more attractive" means "more fragile" (holding the attack costs equal). However, in a parallel subsystem, "more attractive" means "more robust" (again holding the attack costs equal). The intuition behind this is that if it will be impossible to disable a particular subsystem, the attacker would like to find that out before a lot of resources have been invested in attempting to disable the individual constituents of that subsystem. Therefore, in a parallel subsystem, the first basic constituent to be attacked should be either the strongest constituent of the subsystem (if the attack costs of all constituents are equal), or more generally the constituent with the lowest ratio of attack cost to probability of surviving the attack. This avoids wasting resources on constituents that can be disabled with high probability, if the attacker is unlikely to be able to disable the subsystem as a whole. (Another way to think about this is that in a parallel subsystem, an attack on a

constituent with a low probability of being disabled provides more "information" to the attacker on the likelihood of being able to disable the subsystem than attacks on more vulnerable constituents.)

We assume that the objective of the defender is to maximize the minimum expected cost of a feasible attack. We recognize that the cost of an attack to the attackers may not actually be a high priority to the defender in this context. However, the general idea of this defensive strategy is that by maximizing the minimum expected cost of a feasible attack, such an attack may become beyond the capabilities of some or all potential attackers (for example, by increasing the technological sophistication needed to achieve a particular success probability).

We assume that defensive investments in any given component increase the cost of attacking that component, but do not decrease the probability of an attack succeeding. (Another way of stating this is to assume that attackers respond to defensive investments by increasing their efforts enough to hold the probability of success constant, until that is no longer within their capabilities.) Moreover, we assume that the attacker is aware of any changes in the system (i.e., defensive investments) before launching an attack (the case of perfect knowledge), and selects the optimal attack strategy accordingly.

In this formulation, the defender's decision variables consist of the resources to devote to increasing the cost (to the attacker) of attacks on the various components, subject to a budget constraint B limiting the total defensive investments. We assume that the defender wishes to maximize the expected cost of a least-cost feasible attack, in the hopes that an attack will then be beyond the (unknown) capabilities of the attacker(s). The problem thus is to determine the optimal allocation of the total defensive budget B over the various components in order to maximize the expected cost of an optimal attack. We will assume that the cost of attacking a component is monotonically increasing in the defensive resources allocated to that component. This means that the optimum defensive investment will occur at the boundary of the feasible set, or in other words that the optimal defensive strategy is to spend the entire available budget.

We focus here on the case in which the cost of an attack against component i increases linearly in the amount of defensive investment in that component, x_i. In other words, if the initial cost of attacking component or subsystem S_i (before any defensive resources have been expended to increase this cost) is c_i, then a defensive

investment of x_i is assumed to increase the attack cost from c_i to $c_i + a_i x_i$. The assumption that the attack cost increases linearly in the defensive investment is restrictive, but may be realistic for some types of security improvements (e.g., installing additional protective devices in a facility), if the attacker must remove or disable each of these devices one after the other. The assumption of linearity may also be reasonable for other types of defensive investments, provided that the defensive budget is sufficiently small for the attacker cost to be well-approximated by a linear function within the feasible region.

4.1 *Series System*

Consider an ordered series system S of n components, for which the initial cost of an attack on component S_i is c_i, the probability of the component resisting an attack is q_i, and $p_i = 1 - q_i$ for $i=1…n$. Since the system is assumed to be ordered, relationship (1) holds. The problem is to determine the optimal allocation of defensive resources to maximize the expected cost of an optimal attack. (It should be clear here that optimality for the attacker is considered to be minimizing the expected cost of an attack. Moreover, as before, any feasible attack is assumed to be continued until either the system is disabled or the attackers have exhausted their options for disabling the system.)

Let $C(0, 0…0)$ be the initial cost of an optimal attack (before any defensive investments have been undertaken), and let $C(x_1, x_2…x_n)$ be the expected cost of an optimal attack after an investment of $(x_1, x_2…x_n)$ in components $(S_1, S_2…S_n)$, respectively. Note that $C(0, 0…0)$ is given by relationship (2) above. Then, the optimal defensive investment will be the solution to the following optimization problem:

$$\begin{aligned} &\max_{x_1, x_2 \ldots x_n} C(x_1, x_2 \ldots x_n) \\ &s.t. \quad \sum_{i=1}^{n} x_i \le B \\ &\qquad x_i \ge 0, i = 1 \ldots n \end{aligned} \tag{9}$$

Note that the feasible region of (9) is a compact set, and the objective function is increasing in each argument. Moreover, it is possible to show that the objective

function is also continuous (although not in general differentiable). Therefore, the following lemma holds.

Lemma 4.1

An optimal solution of optimization problem (9) exists, and occurs at the boundary of the feasible region $\{x_1, x_2 \ldots x_n \geq 0 \mid \sum_{i=1}^{n} x_i = B\}$.

If the components are ordered in terms of their attractiveness, then the minimum expected cost of a feasible attack would be given by

$$C(x_1, x_2 \ldots x_n) = c_1 + a_1 x_1 + \sum_{i=2}^{n} \prod_{j=1}^{i-1} q_j (c_i + a_i x_i) \qquad (10)$$

Note, however, that the budget allocation can change the order of attractiveness of the various components. This would also change the objective function of the problem. In particular, if after some defensive investment $(x_1, x_2 \ldots x_n)$ the components are ordered according to $(\pi(1) \ldots \pi(n))$, where π is a permutation of $(1, 2 \ldots n)$, then the objective function would become

$$C(x_1, x_2 \ldots x_n) = c_{\pi(1)} + a_{\pi(1)} x_{\pi(1)} + \sum_{i=2}^{n} \prod_{j=1}^{i-1} q_{\pi(j)} (c_{\pi(i)} + a_i x_{\pi(i)}) \qquad (11)$$

Thus, optimization problem (9) is not a standard optimization problem, since while the objective function can always be written as a linear function of the decision variables x_i, the specific form of that linear combination will vary depending on the values of the decision variables. Optimization problem (9) could still be solved by decomposing it into $n!$ linear programs (some of which may not be feasible, if the budget is not large enough to achieve some of the $n!$ possible orderings of the components), solve all of these linear programs individually, and then choose the sub-problem whose optimal solution gives the largest minimum expected cost to the attacker.

The above approach should be computationally feasible at least when n is relatively small. In fact, while the number of linear programs to be solved may be quite large, any individual linear program would be quite simple, and would most likely require at most a few seconds of computational time using standard optimization software. Much the same approach could also be applied to parallel systems and general combined series/parallel systems, following the procedures given

in Section 3. However, this approach does not provide much insight into the qualitative properties of the optimal solution.

In order to investigate the qualitative properties of the optimal solution, we will assume that the cost-effectiveness parameters for investments in the various components are all equal; i.e., $a_i = a$ for all i. This is certainly a restrictive assumption, but allows us to fully characterize the optimal solution in this special case.

We now present results for the optimal budget allocation for defense of a series system as described above, where now we set $a_i = a$ for all $i=1, 2…n$. We begin by stating some preliminary results.

Proposition 4.1

If we have $(c_1 + aB)/q_1 \le c_2/q_2$, then the optimal allocation policy will be given by $(B, 0…0)$.

Proposition 4.1 states that if after spending the entire budget on the component that is initially most attractive it is still the most attractive, then the optimal policy will be to allocate the entire budget to that component.

Proof

From the hypothesis, component 1 will be the most attractive component for any feasible defensive investment $(x_1, x_2…x_n)$. Therefore, the objective function will be of the form:

$$C(x_1, x_2 \ldots x_n) = c_1 + ax_1 + \sum_{i=2}^{n} [\prod_{j=1}^{i-1} p_{\pi(j)}](c_{\pi(i)} + ax_{\pi(i)}) \qquad (12)$$

where π is a permutation of the components $(S_2…S_n)$ such that $(S_1, S_{\pi(2)}…S_{\pi(n)})$ is an ordered series system after investment of $(x_1, x_2…x_n)$. It follows (using Lemma 4.1) that

$$C(x_1, x_2 \ldots x_n) \le C_0 + a\{x_1 + \sum_{i=2}^{n} [\prod_{j=1}^{i-1} p_{\pi(j)}](x_{\pi(i)})\} \le C_0 + a\sum_{i=1}^{n} x_i = C_0 + aB$$

However, the right-hand side of the last inequality is simply $C(B, 0…0)$. □

Proposition 4.2

If all components are equally attractive initially, then under the optimal defensive investment they will still be equally attractive, the optimal resource allocation will be given by

$$x_i^* = Bq_i / \left(\sum_{j=1}^{n} q_j\right) \forall i = 1...n, \qquad (13)$$

and we will have

$$(c_i + ax_i^*)/q_i = (c_j + ax_j^*)/q_j \forall i, j = 1...n \qquad (14)$$

Proof

First, it should be clear that if $(S_1 \ldots S_n)$ is the ordered series system after an optimal allocation of $(x_1, x_2 \ldots x_n)$, then the expected cost of an attack, $C(x_1, x_2 \ldots x_n)$, can be written in the form $C(x_1,...,x_n) = C_0 + \sum_{i=1}^{n} \eta_i x_i$, where the coefficients η_i are non-increasing, and if $(c_i + ax_i)/q_i < (c_{i+1} + ax_{i+1})/q_{i+1}$, then $\eta_i > \eta_{i+1}$. Assume now that (14) does not hold. Then, equation (14) will be replaced by a strict inequality for some pair of i and $j=i+1$. In that case, it would be possible to improve the value of the objective function by increasing x_i and reducing x_{i+1} while holding their sum constant, because $\eta_i > \eta_{i+1}$. This allows the defender to further increase the expected cost of an attack, contradicting the hypothesis that the initial allocation was optimal. From (14), (13) immediately follows. □

From Propositions 4.1 and 4.2 (keeping in mind the arguments used in their proofs), the following result holds.

Corollary 4.1

If the first $k < n$ ordered components are equally attractive, and we have $(c_i / q_i) + aB / \left(\sum_{j=1}^{k} q_j\right) \le c_{k+1} / q_{k+1}$ for $i=1 \ldots k$, then the optimal allocation policy is given by $x_i^* = Bq_i / \left(\sum_{j=1}^{k} q_j\right), i \le k$, and $x_i^* = 0$ otherwise.

The corollary states that if the budget is not sufficient to reduce the attractiveness of the most attractive ones to the same level as that of the next most attractive component(s), then the less attractive components should not receive any

investment at optimality, and the optimal policy should protect all of the most attractive components evenly (i.e., keeping them equally attractive). Moreover, the allocation of defensive resources to any component that is among the most attractive should be proportional to the probability of that component failing in an attack (which is intuitively reasonable). We are now ready to state the general result for series systems.

Theorem 4.1 (Optimal allocation policy)

Consider an ordered series system $(S_1 \ldots S_n)$ satisfying (1), with available defensive budget B. If investing an amount x_i in component S_i will increase the attacker cost by ax_i for some positive value a, then the optimal investment policy that maximizes the minimum expected cost to the attacker is as follows:

1. Invest in protecting component S_1 until either the total budget is depleted or S_1 becomes only as attractive as S_2, whichever occurs first.
2. If the first k components ($1 < k < n$) are equally attractive and the budget is not yet depleted, then allocate the remaining budget among components $1 \ldots k$ while keeping them equally attractive until either the total budget is depleted or they become only as attractive as component S_{k+1}, whichever occurs first.
3. If all n components are equally attractive and the budget is not yet depleted, then allocate the remaining budget among all components while keeping them equally attractive.

Proof

Using the fact that the objective function of each optimization sub-problem in the decomposition approach mentioned above is linear, one can allocate the budget sequentially (i.e., in a greedy manner) without affecting optimality. Thus, by decomposing the budget and allocating it according to the three steps in Theorem 4.1 (applying first Proposition 4.1, then Corollary 4.1 as appropriate, and finally Proposition 4.2 if needed), the result follows. □

4.2 *Parallel Systems*

Consider an ordered parallel system $(S_1 \ldots S_n)$, for which the initial costs of attacking its components satisfy (3), and an available budget B for protecting the system. If we replace q_i by $p_i = 1 - q_i$, $i=1 \ldots n$, where p_i is the probability that component S_i fails given an attack, we can obtain results analogous to those in Section 4.1. More precisely, we have the following result, the proof of which is analogous to that of Theorem 4.1.

Theorem 4.2 (Optimal allocation policy)

Consider an ordered parallel system $(S_1 \ldots S_n)$ satisfying (3), with available defensive budget B. If investing an amount x_i in component S_i will increase the attacker cost by ax_i for some positive value a, then the optimal investment policy that maximizes the minimum expected cost to the attacker is as follows:

1. Invest in protecting component S_1 until either the total budget is depleted or S_1 becomes only as attractive as S_2, whichever occurs first.
2. If the first k components ($1 < k < n$) are equally attractive and the budget is not yet depleted, then allocate the remaining budget among components $1 \ldots k$ (while keeping them equally attractive) until either the budget is depleted or they become as attractive as component S_{k+1}, whichever occurs first.
3. If all n components are equally attractive and the budget is not yet depleted, then allocate the remaining budget among all components while keeping them equally attractive. □

This implies in particular that in a parallel configuration, if the initial costs of attacking the various components are all equal, then the optimal defensive strategy would further protect the most robust components and subsystems before considering the fragile ones. Of course, when the cost of attacking the most robust components gets sufficiently high, they will become less attractive, and more fragile components or subsystems will then have priority for further protection.

4.3 *General Combined Series/Parallel Systems*

We consider now the general case of combined series/parallel systems that can be represented without replications, in which attacks against the various components succeed or fail independently, as assumed above. The following lemma, for which the proof is omitted, is simply an adaptation of Lemma 3.1 from the context of optimal (least-cost) testing strategies to optimal attack strategies.

Lemma 4.2

Consider any ordered series or parallel subsystem $S = (S_1 \ldots S_n)$. Then in order to minimize the expected cost of an attack, an attack against any basic constituent S_i should be completed before moving on to attempt an attack against basic constituent S_{i+1}. □

This lemma states that an optimal attack strategy will not involve attacking some components in S_i, then others in S_{i+1}, and then still others in S_i. Therefore, any basic constituent S_i can be treated as if it were a simple component. This lemma and Theorems 4.1 and 4.2 now allow us to state the general result for an arbitrary series/ parallel system under the restrictions of independence and no replications. The proof is omitted, as it suffices to follow the steps of the *initialization algorithm* and then apply the above results as appropriate.

We consider an arbitrary combined series/parallel system S with no replications, ordered as in the *initialization algorithm* (adapted as appropriate to apply to optimal attack strategies rather than optimal testing strategies). As before, the defender's objective function is taken to be maximizing the expected cost of an optimal (i.e., least-cost) attack, and investing an amount x_i in component i is assumed to increase the cost of an attack against that component by ax_i for some positive value a. Then, the following result holds.

Theorem 4.3 (General combined series/parallel system)

Under the assumptions above, the optimal defensive strategy is as follows:

1. If the system S is a maximal series subsystem consisting of basic constituents $(S_1 \ldots S_n)$, then determine the optimal allocation of defensive investments to these constituents by applying Theorem 4.1 to $(S_1 \ldots S_n)$, with the exceptions that the ordering of the basic constituents

satisfies (5) instead of (1), and the wording "component(s)" is everywhere replaced by "basic constituent(s)" (to reflect the fact that the basic constituents of S may be subsystems instead of simple components).

2. If the system S is a maximal parallel subsystem consisting of basic constituents $(S_1 \ldots S_n)$, then determine the optimal allocation of defensive investments to these constituents by applying Theorem 4.2 to $(S_1 \ldots S_m)$, with the exceptions that the ordering of the basic constituents satisfies (7) instead of (3), and the wording "component(s)" is everywhere replaced by "basic constituent(s)."
3. Determine the optimal allocation of defensive investments among the basic constituents of each subsystem considered in steps 1 and 2 above, by again applying either step 1 or step 2 above as appropriate.
4. Repeat step 3 until decisions have been made regarding the optimal defensive investment in all simple components of the system. □

This result basically states that the most attractive components or basic constituents should be protected first, before considering the next most attractive. This is intuitive from the fact that the "improvement rate" (as given by the coefficient a) is the same for all components. Note, however, that the components most attractive to the attacker need not be defended first if different components have different coefficients a_i. In fact, if some moderately attractive components can be significantly improved with only a minor investment, while the most attractive ones require much larger investments to achieve similar improvements, then a limited budget would not necessarily be allocated only to the most attractive components. In that case, the attractiveness of the various components to the defender (in allocating defensive investments) need not be the same as their attractiveness to the attacker, since attractiveness to the defender depends on the cost-effectiveness of defensive investments in the various components, not only the cost and success probability of attacks against the various components. (Note, however, that if the attacker does not eventually target all components, but rather can afford to target only a single minimal cut set of the system, then it may no longer be optimal to defend any but the most attractive components, regardless of the cost-effectiveness of defensive investments in less attractive components.)

The results stated above are for the case of "perfect knowledge," where the attacker is assumed to be fully aware of any improvements made to the system by the defender prior to launching an attack. However, it should be clear that the optimal defensive strategy derived under the assumption of perfect knowledge is conservative for the defender, in the sense that the expected attack cost in the case of perfect information is a lower bound on the expected attack cost in the more general case of "imperfect knowledge." Moreover, the result in Theorem 4.3 depends on the assumption that the attacker will eventually target all minimal cut sets of the system, if not already successful in an earlier attack. However, the algorithm in Theorem 4.3 still yields a conservative strategy for the defender if the attacker cannot afford to target all minimal cut sets, but the defender does not know which minimal cut sets the attacker plans to target.

5. Conclusions and Directions for Further Work

The results presented here represent an initial attempt to extend existing results for defense of individual components (e.g., Major, 2002; Woo, 2002, 2003; O'Hanlon, 2002) or simple series and parallel systems (e.g., Bier and Abhichandani, 2003; Bier et al., in press) to combined series/parallel systems of more realistic complexity. As such, it yields interesting and sometimes counterintuitive insights, such as the observation that defending the stronger train(s) in a parallel subsystem can actually impose greater burdens on prospective attackers than hardening the weaker train(s).

As a preliminary analysis of a complex problem, there are obviously numerous possible extensions, refinements, and alternative model formulations that might be worth exploring. For example, we have indicated briefly how our approach could be applied in the case where the cost-effectiveness of defensive investment in the various components differs. However, our results so far are limited to the case where the cost of attacks increases linearly in the defensive investment. Since this may hold only for a limited range of defensive investments, it would be of interest to characterize the nature of the optimal defensive investment in the more realistic case where the cost of an attack is a concave function of the defensive investment, even if this can be done only for series systems rather than more general system structures.

Similarly, our current model has the disadvantage that it can lead to wasteful levels of expenditure by defenders, in cases where the maximum expected attack cost

achievable by the defender vastly exceeds the resources of likely attackers. In practice, it seems likely that defenders will be significantly resource-constrained, in which case this may not be a significant limitation. However, to enhance the applicability of our model, it would be interesting to extend our results to the case where the defender does not have a firm budget constraint, and instead chooses the total level of investment based on the value of the system being protected, and some (perhaps uncertain) assessment of the attacker budget.

Models of imperfect attacker information might also be of interest. In this type of formulation, failed attacks against particular components might serve to update the attacker's assessment of their attractiveness (e.g., in a Bayesian manner). This could serve to model "probing" attacks, whose purpose is in part to gather information about system defenses. A model of this form might have the feature that attackers can come back to re-target components that had previously been attacked unsuccessfully, if subsequent attacks revealed other components to be even more difficult to disable.

Additionally, it might be worthwhile to model the situation in which defensive investment reduces the success probability of an attack, rather than increasing its cost. Such a model, while perhaps mathematically less tractable, would be more intuitively appealing, since the defensive investments would more clearly be defending the system, rather than merely attempting to deter attackers. A revised model of this sort would also apply in situations where the attacker is not realistically budget-constrained, such as some types of extremely low-cost computer attacks, which can be extremely damaging but are within the resource constraints of enterprising teenagers.

Finally, in a more fully general model, defensive investments in the various components of a system would change the function relating the success probability of an attack to the level of effort expended by the attacker. In response to this function, the attacker would then simultaneously determine both the levels of effort to be expended on attacking the various components, and the success probabilities of those attacks. Such a model might then make it possible to assess the merits of differing types of defensive strategies, such as increasing the cost of attacks (which might deter some attackers, but would not necessarily reduce the success probabilities of those attackers with the most resources), reducing the success probabilities of attacks (without necessarily deterring any attackers), etc.

References

Azaiez, M.N., 1993. Perfect aggregation in reliability models with Bayesian updating. PhD dissertation, University of Wisconsin-Madison.

Azaiez, M.N., Bier V.M., 1995. Perfect aggregation in general reliability models with Bayesian updating, Applied Mathematics and Computation 73 281-302.

Ben-Dov, Y., 1981. Optimal testing procedures for special structures of coherent systems, Management Science 27 (12) 1410-1420.

Bier, V.M., Abhichandani, V., 2003. Optimal allocation of resources for defense of simple series and parallel systems from determined adversaries. In: Risk-Based Decisionmaking in Water Resources X, American Society of Civil Engineers, Reston, VA, pp. 59-76.

Bier, V.M., Nagaraj, A., Abhichandani, V., in press. Protection of simple series and parallel systems with components of different values, Reliability Engineering and System Safety.

Butterworth, R.W., 1972. Some reliability fault-testing models, Operations Research 20 335-343.

Cox L., Qiu, Y., Kuehner, W., 1989. Heuristic least-cost computation of discrete classification functions with uncertain argument values, Annals of Operations Research 21 1-30.

Cox, L., Chiu, S., Sun, X., 1996. Least-cost failure diagnosis in uncertain reliability systems, Reliability Engineering and System Safety 54 203-216.

Halpern, J., 1974. Fault-testing of a k-out-of-n system, Operations Research 22 1267-1271.

Levitin, G., 2002. Maximizing survivability of acyclic transmission networks with multi-state retransmitters and vulnerable nodes, Reliability Engineering and System Safety 77 189-199.

Levitin, G., 2003a. Optimal multilevel protection in series-parallel systems, Reliability Engineering and System Safety 81 93-102.

Levitin, G., 2003b. Optimal allocation of multi-state elements in linear consecutively connected systems with vulnerable nodes, European Journal of Operational Research 150 406-419.

Levitin, G., Lisnianski, A., 2000. Survivability maximization for vulnerable multi-state systems with bridge topology, Reliability Engineering and System Safety 70 125-140.

Levitin, G., Lisnianski, A., 2001. Optimal separation of elements in vulnerable multi-state systems, Reliability Engineering and System Safety 73 55-66.

Levitin, G., Lisnianski, A., 2003. Optimizing survivability of vulnerable series-parallel multi-state systems, Reliability Engineering and System Safety 79 319-331.

Levitin, G., Dai, Y., Xie, M., Poh, K.L., 2003. Optimizing survivability of multi-state systems with multi-level protection by multi-processor genetic algorithm, Reliability Engineering and System Safety 82 93-104.

Major, J., 2002. Advanced techniques for modeling terrorism risk, Journal of Risk Finance 4 (1) 15-24.

O'Hanlon, M., Orszag, P., Daalder, I., Destler, M., Gunter, D., Litan, R., Steinberg, J., 2002. Protecting the American Homeland, Brookings Institution, Washington, DC.

Woo, G., 2002. Quantitative terrorism risk assessment, Journal of Risk Finance 4 (1) 7-14.

Woo, G., 2003. Insuring against Al-Qaeda, Insurance Project Workshop, National Bureau of Economic Research, Inc. (Downloadable from website http://www.nber.org/~confer/2003/insurance03/woo.pdf).

Figure 1 Sample system *S* that can be represented with no replications