SAND REPORT
SAND2004-0742
Unlimited Release
Printed March 2004

Expected Losses, Insurability, and Benefits from Reducing Vulnerability to Attacks

Rolf E. Carlson, Mark A. Turnquist, and Linda K. Nozick

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under Contract DE-AC04-94AL85000.

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865)576-8401
Facsimile: (865)576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: http://www.doe.gov/bridge

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd
Springfield, VA 22161

Telephone: (800)553-6847
Facsimile: (703)605-6900
E-Mail: orders@ntis.fedworld.gov
Online order: http://www.ntis.gov/ordering.htm

SAND2004-0742
Unlimited Release
Printed March 2004

Expected Losses, Insurability, and Benefits from Reducing Vulnerability to Attacks

Rolf E. Carlson
Advanced Information and Control Systems
Sandia National Laboratories
P.O. Box 5800
Albuquerque, NM 87185-1351

Mark A. Turnquist
Cornell University
Ithaca, NY

Linda K. Nozick
Cornell University
Ithaca, NY

Abstract

A model of malicious attacks against an infrastructure system is developed that uses a network representation of the system structure together with a Hidden Markov Model of an attack at a node of that system and a Markov Decision Process model of attacker strategy across the system as a whole. We use information systems as an illustration, but the analytic structure developed can also apply to attacks against physical facilities or other systems that provide services to customers. This structure provides an explicit mechanism to evaluate expected losses from malicious attacks, and to evaluate changes in those losses that would result from system hardening. Thus, we provide a basis for evaluating the benefits of system hardening. The model also allows investigation of the potential for the purchase of an insurance contract to cover the potential losses when safeguards are breached and the system fails.

ACKNOWLEDGEMENTS

Thanks to Dean Jones for his support of this effort.

AUTHOR CONTACTS

Rolf E. Carlson
Phone: 505-844-9476
E-mail: recarls@sandia.gov

Mark A. Turnquist
Phone 607-255-4796
E-mail: mat14@cornell.edu

Linda K. Nozick
Phone 607-255-6496
E-mail: lkn3@cornell.edu

I. INTRODUCTION

Interest is widespread in protection of infrastructure from malicious attack, and protection of computer and information systems is an important part of this overall concern. Most available literature on information system security focuses on tactical questions—How can intruders be best detected? What system changes can be implemented to reduce known vulnerabilities? How can vulnerabilities in new software be identified and fixed before release? These tactical questions are very important, but in this paper we focus on four questions at a more strategic level:

1. *For a given system, can we estimate the expected loss rate due to malicious attacks, as a function of some basic system characteristics and parameters?*

2. *If so, can we estimate the probable benefit of various types of "system hardening" as a basis for cost-benefit evaluation of potential system modifications and/or investments?*

3. *By putting uncertain losses from malicious attacks in an economic context, can we begin to understand the "insurability" of systems against such losses?*

4. *From a managerial perspective, what can such analysis tell us about how safe is safe enough as organizations (both public and private) struggle with the question of how to allocate resources for system security?*

This paper provides insight on strategies that might be followed by a system owner to reduce expected losses from adversarial attacks. By focusing on expected losses, we are also adopting a perspective that is consistent with the theory of insurance in the economics literature [1], and thus a secondary goal is to offer insights about insurability of systems against losses from deliberate attacks. Being able to balance system hardening against other forms of protection, such as insurance, is important for a complete cost/benefit analysis. Cost/benefit analyses of potential security upgrades are important for effective management of information systems and other types of infrastructure systems.

Our application context in this paper is information systems [with a particular interest in supervisory control and data acquisition (SCADA) systems], but the general approach is likely to be useful for assessing losses from malicious attacks in other kinds of systems as well. The basis for our analysis is a representation of the system of interest as a network of nodes and arcs. Nodes represent system assets, and arcs represent opportunities for attackers to move within the system.

Several previous authors have used graph-based methods to represent attackers or defenders in security analyses. Phillips and Swiler [2] introduced the concept of an "attack graph" to represent sets of system states and paths for an attacker to pursue an

objective in disrupting an information system. Dacier [3] created a related concept, termed a "privilege graph," to represent varying levels of privileges attained by an intruder on different processors in a computer network. Several subsequent papers (e.g., [4], [5], [6]) have extended these initial ideas.

A number of authors have used Markov models to represent uncertainties in system state in the face of attacks, especially from viruses, worms, and Trojan horses (e.g., [7], [8]). Soh and Dillon [9] used a Markov representation as the basis for a model of intrusion detection. This has been extended to more complex representations using Hidden Markov Models (HMM) that focus on intruder detection using indicators that indirectly reflect potential attacker activities (see, for example, [10], [11], [12]).

Jha et al. [5] and Sheyner et al. [6] introduce the idea of using Markov Decision Processes (MDP) for situations in which the attack path is probabilistic. We are also interested in using MDP tools for analyzing the strategy of system intruders; but our work is based on a different type of state representation for the system, and our objective is not to reflect detailed actions by either the attacker (often termed "atomic attacks") or the defender. In Section III, we show how our work can be connected to intruder detection analyses, but our focus is not on intruder detection, per se, but on the strategic questions of loss rates, insurability, and evaluation of investments in system security.

To address these questions, we first construct an HMM to represent an attack at a single node in a system. Then we develop an aggregated representation of that single-node model for inclusion in an MDP model of attacker strategy within a network representation of the entire system. Third, the MDP solution is used to compute expected losses from different classes of attackers, as a means of tying the analysis to the notion of "insurability." Finally, the sensitivity information from the MDP solution is used to indicate the parts of the system in which "hardening" against attacks may be most effective. To our knowledge, this is the first effort to use HMMs and MDPs in this way to evaluate economic losses from malicious attacks in systems and to assess potential benefits of hardening measures.

Section II offers a general system description and the notation used throughout the paper. In Section III, we describe the model for an attack at a single node in the system. In Section IV, we then use a connected set of abstractions of the single-node models to create a MDP model for the attacker's strategy across the system as a whole. This network-level model forms the basis for expected loss calculations in Section V, and a discussion of insurability in Section VI. The MDP also provides information on the relative sensitivity of expected losses to "hardening" of specific nodes, and we explore these implications in Section VII. Conclusions and suggestions for further research are discussed in Section VIII.

II. SYSTEM DESCRIPTION

A system, or target of evaluation (TOE), provides services to external users. The system can be an information network, physical facility, or both. Adversaries attack the system, seeking to degrade the services offered, and successful attacks result in economic loss for the system owner. The owner of the TOE may or may not sell services in an open market, but we will assume that losses can be measured in monetary terms.

In general, there may be several categories of attackers. We will define a set C of categories. Attackers in different categories may have different levels of skill and may pursue different strategies in attacking the system.

The TOE is represented as a directed graph $G = \{V, E\}$ with finite sets V of nodes and E of edges. We assume that there are real-valued functions θ_c defined on V representing the loss incurred by the system owner if a particular node is successfully attacked (i.e., breached) by an attacker of category $c \in C$. Any node $v \in V$ for which $\theta_c(v) > 0$ will be regarded as a system asset. Nodes may represent levels of privilege on a given processor (e.g., admin privileges), access to certain protected files, or the ability to initiate actions on some processor. The construct is intended to be quite general to allow application of the analysis tools developed here in different types of situations.

Our primary attention is on a class of adversaries that is rational and well informed. By "rational," we mean that the adversaries make decisions by weighing risks and benefits, that they will follow a strategy that maximizes the expected loss they can inflict on the system, and that they will quit when the risks of detection outweigh the incremental expected loss they can inflict. By "well informed," we mean that the adversaries know the values of system assets (i.e., they know $\theta_c(v)$ for all nodes $v \in V$), so they can direct their attacks to do the most damage.

The focus on adversaries who weigh risks and benefits of attacks in a rational way might be considered limiting, but this framework can actually account for quite a wide range of plausible behaviors. One might argue, for example, that a terrorist does not consider the risk of detection, and focuses only on the expected losses that can be inflicted. This is handled in a very straightforward way in the framework created here, by simply setting the cost of detection (as viewed by the attacker) to zero. Another extension of this basic model is to consider a utility function for the attacker that may be nonlinear in expected costs. This would allow representation of more general types of risk-prone or risk-averse behavior on the part of the attacker. The basic model structure described here can also accommodate that extension in a straightforward way.

Our focus on well-informed adversaries is useful because it leads to an estimate of expected losses that is likely to be an upper bound on losses from less well-informed attackers. At the end of Section IV, we discuss an extension of the analysis to consider

adversaries who know less about the system they are attacking, and describe one way to accomplish that extension in a relatively easy way. Further exploration of this topic is, however, an important area for additional work.

The adversaries mount their initial attacks at entry points to the system, and if an attack at a particular node is successful, they can traverse edges from the successfully breached node to other nodes in the network that are connected to the one breached. Traversing an edge entails a risk of detection. An attacker who has successfully breached a node may also choose to quit and exit the system. The adversary is assumed to make the decision that is most favorable to him/her.

The owner has an objective to maintain the integrity of the system. We interpret that objective as minimizing the expected losses incurred, but other representations of the owner's objective might be equally valid. Our focus on minimizing expected losses is consistent with the objectives of evaluating insurability of the system, but extension to consider other measures of system integrity is also very useful.

Throughout the paper, we use a simple remote monitoring system (Figure 1) as an example to illustrate the calculations. This example is built up in steps through the sections.

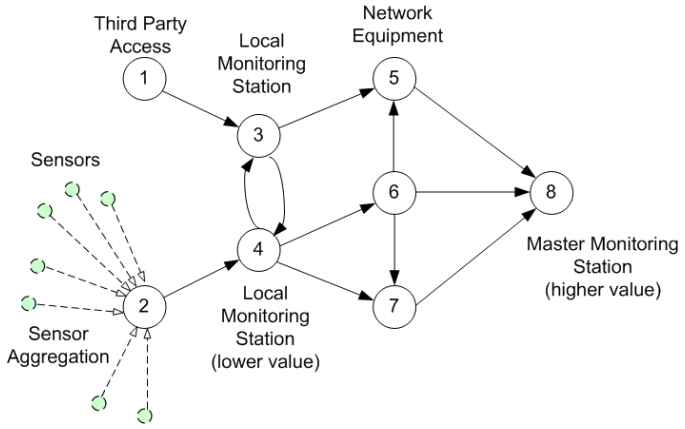

Figure 1. An example network for monitoring the state of an infrastructure system.

With the addition of communication channels back to the field, this depiction could represent a Command and Control (C2) system, or a portion of a Supervisory Control and Data Acquisition (SCADA) system in the electric power grid. Potential unauthorized entry to the system is via nodes 1 and 2, and a variety of possible paths for an attacker present themselves as various nodes are breached. The sensors are portrayed to suggest that there may be network elements that exist but are not included in the analysis.

III. Adversary and Attack Characterization at a Node

An attack on a system node and the interaction between the attacker and the intrusion-detection system is modeled using a hidden Markov model (HMM). The general concept of such a model is represented in Figure 2. The attacker's actions (the lower portion of the diagram) are assumed to progress through a set of states as a Markov process. Occupancy of various states may result in emanations that are observable by the system operator (represented by the "signals" in Figure 2). The system operator can define some subset of emanations that, if observed, will cause the user to be placed on a watch list for monitoring. Some other set of emanations will cause the system to evict the user (correctly or incorrectly) under the premise that the user is an attacker. If the attacker reaches a set of states that we call "breach states" without being evicted, we say that the node has been breached, and no further emanations will cause the system to evict the attacker at that node.

The state space in the HMM is defined to encompass both the categories of attackers and the level of monitoring or concern that the system attaches to a user (potential attacker). Thus, for example, if we are representing two categories of attackers (nominally referred to as "weak" and "strong" attackers), the nodes representing the attack space are expanded into two layers, as shown in Figure 3. Attackers may gain strength in the course of their attacks and make transitions from the "weak" layer to the "strong" layer. In addition, if we define a set *M* of monitoring levels, the set of layers expands further to reflect attackers in various categories being monitored at different levels. A different HMM may be tailored for each different node in the system, allowing implicit characterization of the node, attack, attacker, and threat level.

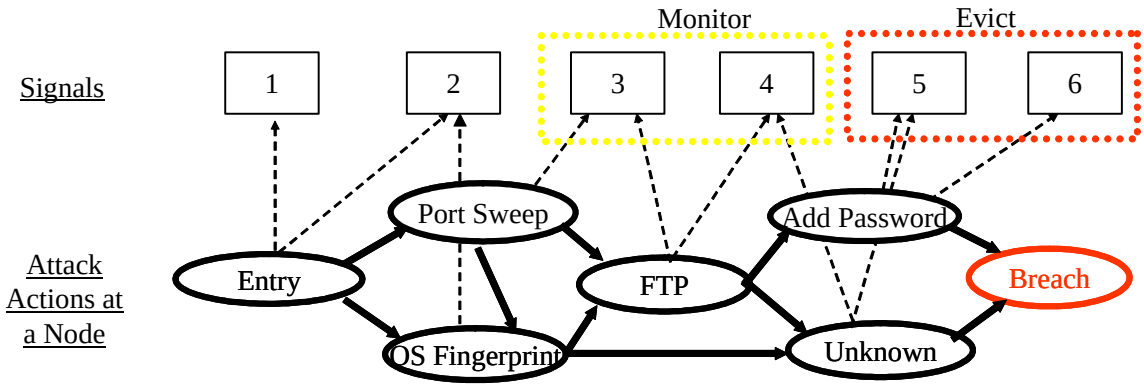

Figure 2. A Hidden Markov Model characterizing an attack at a system node.

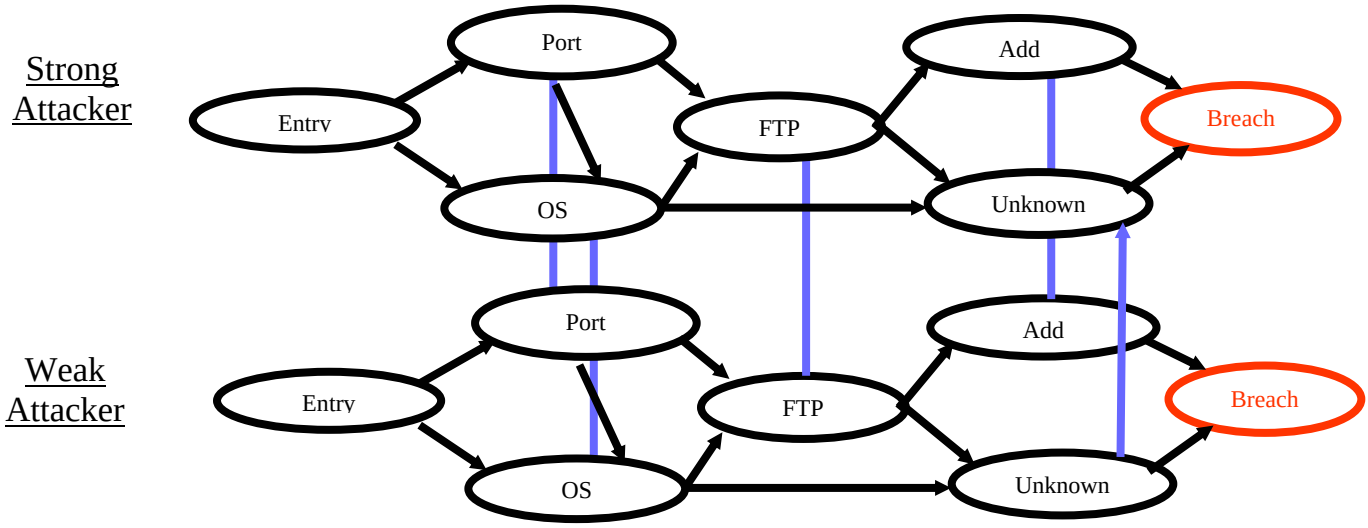

Figure 3. Layered state representation to reflect multiple categories of attackers.

We use a discrete-time, discrete-state HMM characterized by the following equations:

$$X_{n+1} = A^T X_n \tag{1}$$

$$Y_n = BX_n \tag{2}$$

for transition steps $n = 1, 2, \ldots, \infty$. The state of the system (i.e., presence of the attacker in some node in the lower portion of Figure 2) is represented by the (column) probability vector, X. The dynamics of the system are governed by (1), where A is a transition matrix (i.e., it satisfies the properties $a_{ij} \geq 0$ and $\sum_j a_{ij} = 1$.) The states of the system are not observed directly. The process Y is observed, which is a function of the state of the underlying Markov process, X. Each column of B specifies a conditional probability distribution over the possible observations, given that the underlying (hidden) system is in a particular state.

In intrusion-detection applications, HMMs are typically used as the basis for an estimation problem. That is, given a set of observations, $\{Y_1, Y_2, \ldots\}$, which may also include noise, it is desired to estimate the state of the Markov chain at some point in time. Often, this is cast as a "filtering" problem: given $\{Y_1, Y_2, \ldots, Y_k\}$, estimate X_k, assuming that A and B are known matrices. However, in many applications, it is assumed that A and B are unknown, and the problem may include estimating those matrices from the data as well as using the resulting estimates to estimate X_k. This is termed *adaptive* estimation. Several algorithms are available for constructing the estimates of A and B from sequences of observations (for example, see the discussion in [13]).

For our purposes, we assume that A and B are known (or have been estimated). We want to use the estimated HMMs at various nodes as the basis for a network-level model of attacker strategy. To do this, we will abstract the HMM at node *v* to a simpler representation, as shown in Figure 4, which reflects two categories of attackers, denoted as "weak" and "strong." An attacker in category *c* enters an "Attack *v*" state for that category. The attacker continues to occupy that state until the attack is detected (and the attacker is evicted), or the attack is successful and transitions to a "Breach *v*" state associated with a category *c'*. We adopt a convention that the attacker categories are ordered by increasing "strength" of the attacker, and we allow attackers to change categories (e.g., become "stronger") through a successful attack on a node (i.e., $c' \geq c$). The system can change its level of monitoring as a result of possible emanations from the attacker during the attack, and this change in monitoring level must be reflected in the transition probabilities shown in Figure 4. The simplified representation of the HMM at a node is characterized by $\frac{|C|(|C|+3)}{2}$ basic parameters – the transition probabilities $p_{cc'}(v)$ and $d_c(v)$ shown in Figure 4, representing successful attacks and detection (eviction), respectively.

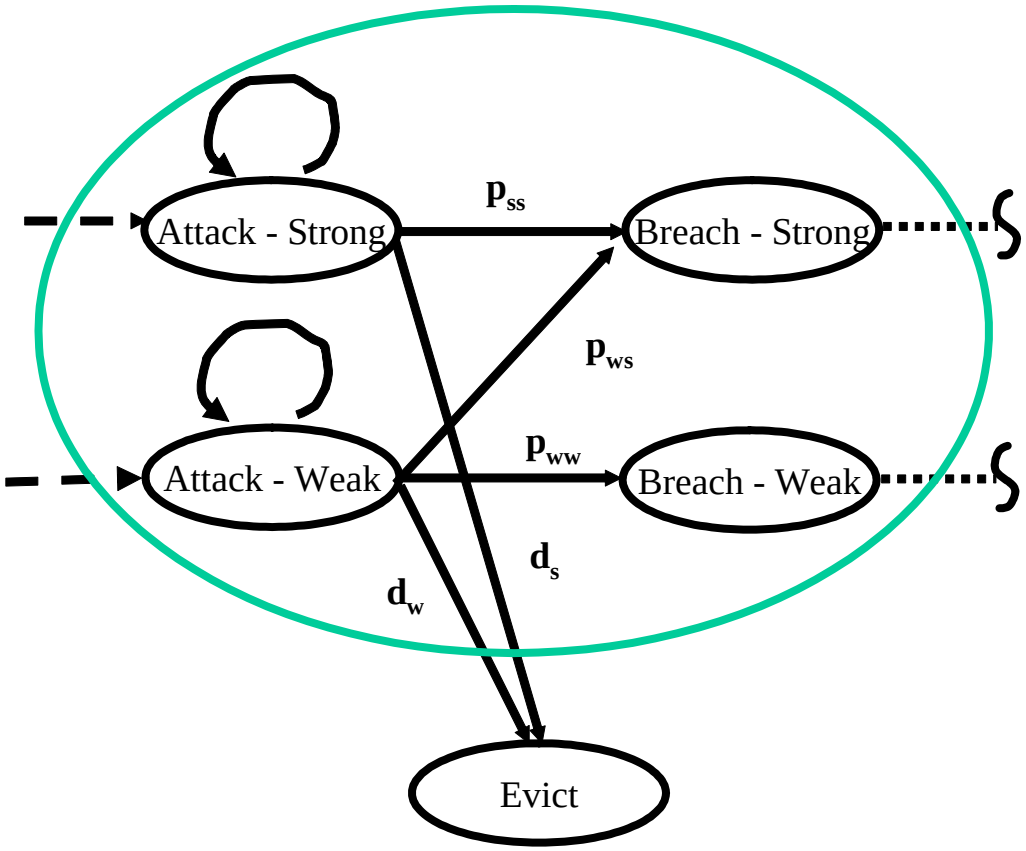

Figure 4. An aggregated abstraction of the HMM at a node.

To make the abstraction in Figure 4 useful, we must be able to derive the values of $p_{cc'}(v)$ and $d_c(v)$ from the underlying A and B matrices of the HMM. The attack states in Figure 4 are transient states, and the breach states and eviction state are absorbing. The transition probabilities $p_{cc'}(v)$ and $d_c(v)$ are specified so that the probabilities of absorption in the breach and evict states match those from the original HMM, and so that the expected number of transitions prior to absorption also matches the original HMM. To do this, we construct an augmented state space for the HMM by adding an eviction state. The transition probabilities to the eviction state are given by the b_{ij} values from the B matrix corresponding to emanations that are specified to cause eviction. Transition probabilities between levels of monitoring are also specified by b_{ij} values, for emanations i that correspond to causing increased monitoring. The original transition probabilities (in the A matrix) are adjusted to account for the probability of eviction and/or change in monitoring level. The resulting transition matrix for the augmented state space will be denoted as P.

Figure 5 is an example of the augmented Markov model corresponding to Figure 2, where there are two categories of attackers (weak and strong, abbreviated to W and S) and two levels of monitoring by the system (which we may think of as "normal" and "high," abbreviated as N and H). The nodes (states) in the model are designated as 1SN, 3WN, 4SH, etc., indicating the combination of operation, attacker category, and monitoring level.

To be more formal about the expansion of the state space illustrated in Figure 5, if the original set of states (e.g., as shown in Figure 2) contains K states and we define $|C|$ categories of attackers, the state vector X will contain $K|C|$ states (e.g., as shown in the example in Figure 3). The augmented state space, X', will then contain $K|C||M|$ +1 states (as illustrated in Figure 5).

We define a subset Ω_v of the observation states in Y (at node v) such that if an observation is recorded in Ω_v, the user is evicted from the system. We will also define a subset ψ_v of those states that corresponds to the system increasing its monitoring level. The sets Ω_v and ψ_v reflect the security protocols in place in a given system. Making these sets larger (especially Ω_v)allows faster detection of attackers, but also causes more false alarms and can preclude legitimate users from performing necessary functions. Making the sets smaller reduces the sensitivity of the detection system.

The subset of states in X' denoting "breach" states (for the various combinations of attacker category and monitoring level) plus the "evict" state are absorbing. Given the transition matrix P, we can group the absorbing states at the end of the list and partition P as follows:

$$P = \begin{bmatrix} Q & Z \\ 0 & H \end{bmatrix} \tag{3}$$

The submatrix *Q* represents transitions among the transient states, *Z* represents transitions from the transient states into the absorbing states, and *H* represents transitions within the set of absorbing states. In applications of interest here, *H* is usually an identity matrix.

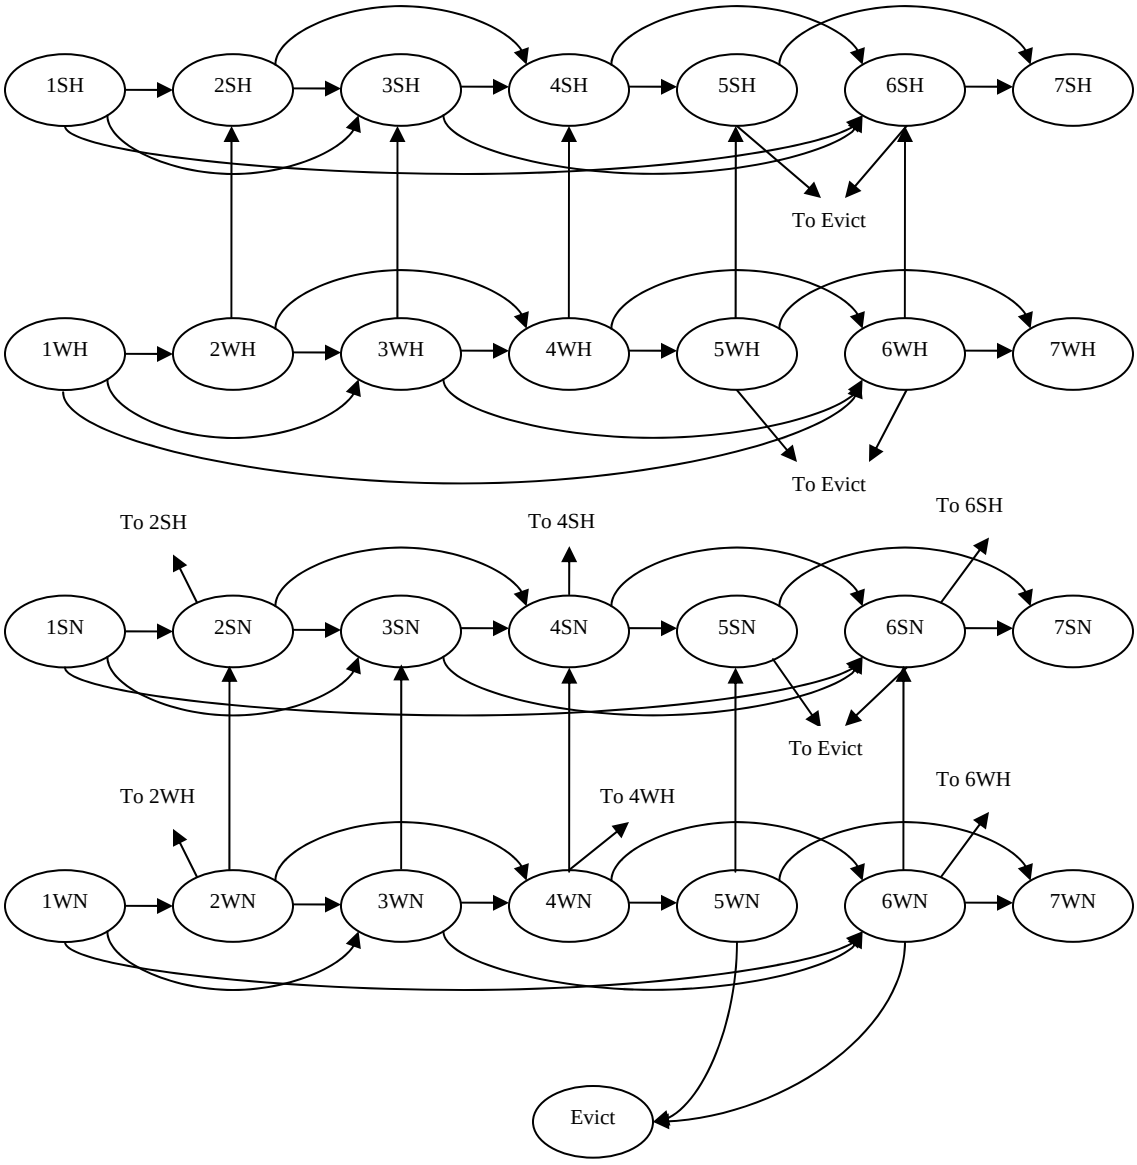

Figure 5. States and transitions for the expanded Markov model representing the network in Figure 2.

The Fundamental Matrix $\Phi = (I - Q)^{-1}$ contains elements ϕ_{ij}, interpreted as the expected number of visits to state *j* before absorption, given that the system started in state *i* (see, for example, [14]). The chain, *X'*, has |*C*||*M*| identifiable entry states (initiation of an attack by an attacker in category $c \in C$ who is being monitored at level $m \in M$). We will use $\Gamma_e(c)$ to denote the set of entry states in *X'* that are aggregated into the "attack" state

for category *c* in the reduced network shown in Figure 4. In general, this will include entry states corresponding to different system monitoring levels. We use $\Gamma_b(c)$ to denote the set of breach states in *X′* aggregated into the breach state for category *c* in the reduced network.

In Markov chains that have both transient and absorbing states, if *j* is one of the absorbing states, the probability that the system is absorbed in state *j*, given that the initial state was state *i*, is given by the ij^{th} element of the matrix ΦZ (for a proof of this result, see [15], page 157). We denote this conditional probability as $f_i(j)$:

$$f_i(j) = [\Phi Z]_{ij} \tag{4}$$

In the reduced state representation (Figure 4), equation (4) allows us to write two expressions:

$$\text{Prob(Breach in category c′ | Entry in category c)} = \frac{p_{cc'}(v)}{\sum_{c''\geq c} p_{cc''} + d_c(v)} \tag{5}$$

$$\text{Prob(Eviction | Entry in category c)} = \frac{d_c(v)}{\sum_{c''\geq c} p_{cc''} + d_c(v)} \tag{6}$$

In the full representation of the Markov chain (Figure 5), the entry in category c is represented by a set of entry states, $\Gamma_e(c)$. If the probability of entry in state *i* is π_i, the conditional probability of entry in state *i*, given that the attacker is of category *c*, is:

$$\frac{\pi_i}{\sum_{j\in\Gamma_e(c)} \pi_j}.$$

Then, in the full representation, we can write:

$$\text{Prob(Breach in category c′ | Entry in category c)} = \sum_{k\in\Gamma_b(c')} \sum_{i\in\Gamma_e(c)} f_i(k) \frac{\pi_i}{\sum_{j\in\Gamma_e(c)} \pi_j} \tag{7}$$

If the Eviction state in Figure 5 is the last state number (i.e., $K|C||M|+1$), then:

$$\text{Prob(Eviction | Entry in category c)} = \sum_{i\in\Gamma_e(c)} f_i(K|M||C|+1) \frac{\pi_i}{\sum_{j\in\Gamma_e(c)} \pi_j} \tag{8}$$

Equating the expressions in (5) and (7), and in (6) and (8), we have $\frac{|C|(|C|+3)}{2}$ equations in the unknowns $p_{cc'}(v)$ and $d_c(v)$. These equations do not allow unique solution for the unknowns, because there are $|C|$ linear dependencies in those equations (the conditional probabilities in (5) and (6) must sum to 1). However, it is also desirable that the expected number of transitions prior to absorption be equal for the full and reduced representations, and this provides $|C|$ additional equations. In general, the expected number of transitions prior to absorption for an attacker who enters in state *i* is:

$$n_i = \sum_j \phi_{ij} \tag{9}$$

For the reduced representation, the expected number of transitions prior to absorption for an attacker in category *c* is then:

$$n_i(c) = \frac{1}{\sum_{c' \geq c} p_{cc'}(v) + d_c(v)} \tag{10}$$

In the expanded representation, this conditional probability is:

$$n_i(c) = \sum_{i \in \Gamma_e(c)} n_i \frac{\pi_i}{\sum_{j \in \Gamma_e(c)} \pi_j} \tag{11}$$

Equating expressions (10) and (11), we can solve for the denominators in (5) and (6), so the final expressions for $p_{cc'}(v)$ and $d_c(v)$ are as follows:

$$p_{cc'}(v) = \frac{\sum_{k \in \Gamma_b(c')} \sum_{i \in \Gamma_e(c)} f_i(k) \frac{\pi_i}{\sum_{j \in \Gamma_e(c)} \pi_j}}{\sum_{i \in \Gamma_e(c)} n_i \frac{\pi_i}{\sum_{j \in \Gamma_e(c)} \pi_j}} \tag{12}$$

$$d_c(v) = \frac{\sum_{i \in \Gamma_e(c)} f_i(K|M||C|+1) \frac{\pi_i}{\sum_{j \in \Gamma_e(c)} \pi_j}}{\sum_{i \in \Gamma_e(c)} n_i \frac{\pi_i}{\sum_{j \in \Gamma_e(c)} \pi_j}} \tag{13}$$

The value of the simplified representation is that it allows us to construct a Markov Decision Process (MDP) of the attacker's strategy at the system level, without carrying along all the detail of states within the potential attacks at each node. This is the focus of the following section, and represents the third major step in our analysis.

Before proceeding to that discussion, the analysis is illustrated at a single node level, considering entry node 1 in Figure 1 and the HMM that might be constructed as part of an intrusion detection system there. Table 1 represents an example of a set of basic states that could represent various attacker actions, and the possible emanations that could result from attacker presence in these states. The definitions in Table 1 correspond to the diagram in Figure 2.

Table 1. Example attack states, emanation signals, and system operator actions.

State	Example Attack Operation	System Action if Detected	Possible Emanation "Signals"
1	Entry	None	None
2	Port Sweep	Monitor	y_1, y_2
3	Operating System Fingerprint	Monitor	y_1, y_3, y_5
4	FTP Connection	Evict	y_2, y_4
5	Password File Edit	Evict if Monitoring Level is High	y_3
6	Unknown	Unobservable	None
7	Breach	Unobservable	None

As a potential attack unfolds, the operator may observe unusual activity, or emanations from various attack operations. These emanations are represented by $y_1, \ldots, y_5$. To translate these emanations into a more concrete setting, consider an intrusion detection system based on a collection of pattern matching rules. If a suspicious pattern is recognized, then a rule is tripped, and the system is alerted to the activity. Invoking a rule produces the "emanation" that yields observability for the attack operation at that moment. If an emanation suggests a sufficient threat, the system might decide to evict the user. In Table 1, the operator has chosen to evict the user if emanations y_2 and y_4 are seen together (regardless of the current monitoring level attached to that user), or if y_3 is seen alone and the user is currently under suspicion (High monitoring level). However, if y_2 is seen together with y_1, or y_3 is seen together with y_1 and/or y_5, the user is placed on a watch list for monitoring but not evicted. Note that this structure includes the possibility of a "false positive" – a user in state 3 may cause an emanation of y_3 without either y_1 or y_5 and thus be evicted even though the system would not

normally evict a user for the action in state 3. State 6 represents an action that may not be catalogued by the system (e.g., a novel attack mode that does not leave traces detectable by the collection of pattern matching rules).

The collection of emanations listed in Table 1 allows us to construct 12 possible observable outcomes:

$$\{y_1, y_2, y_3, y_4, y_5, y_1y_2, y_1y_3, y_1y_5, y_2y_4, y_3y_5, y_1y_3y_5, none\}$$

To define the outcome space (Y) for the HMM, we expand each of these outcomes based on the current monitoring level (N or H), leading to 24 elements for the outcome space. Outcomes y_3H, y_2y_4L, and y_2y_4H constitute the set Ω_1, leading to eviction at node 1.

We assume that the intrusion-detection process has estimated the A and B matrices shown in Figures 6 and 7 for this node. From these matrices, we can construct estimates of $p_{WW}(1)$, $p_{WS}(1)$, $p_{SS}(1)$, $d_W(1)$, and $d_S(1)$, using equations (5)–(11). The resulting values are:

$p_{WW}(1) = .014$
$p_{WS}(1) = .016$
$p_{SS}(1) = .03$
$d_W(1) = .029$
$d_S(1) = .006$

This collection of five parameter values summarizes the HMM at node 1 for representation in the system-level model to be described in the following section.

from \ to	1WN	2WN	3WN	4WN	5WN	6WN	7WN	1SN	2SN	3SN	4SN	5SN	6SN	7SN	1WH	2WH	3WH	4WH	5WH	6WH	7WH	1SH	2SH	3SH	4SH	5SH	6SH	7SH	evict
1WN	0	0.7	0.25	0	0	0.05	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
2WN	0	0	0.179	0.179	0	0	0	0	0.012	0	0	0	0	0	0	0.63	0	0	0	0	0	0	0	0	0	0	0	0	0
3WN	0	0	0	0	0.836	0.054	0	0	0	0.014	0	0	0	0	0	0	0.096	0	0	0	0	0	0	0	0	0	0	0	0
4WN	0	0	0	0.764	0.08	0	0.01	0	0	0	0.016	0	0	0.004	0	0	0	0	0	0	0.004	0	0	0	0	0	0	0.002	0.12
5WN	0	0	0	0	0.932	0.01	0.02	0	0	0	0	0.018	0	0.008	0	0	0	0	0	0	0.008	0	0	0	0	0	0	0.004	0
6WN	0	0	0	0	0.05	0.91	0.01	0	0	0	0	0	0.02	0.004	0	0	0	0	0	0	0.004	0	0	0	0	0	0	0.002	0
7WN	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
1SN	0	0	0	0	0	0	0	0	0.3	0.2	0	0	0.5	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
2SN	0	0	0	0	0	0	0	0	0	0.272413	0.272413	0	0	0	0	0	0	0	0	0	0	0	0.455175	0	0	0	0	0	0
3SN	0	0	0	0	0	0	0	0	0	0	0	0.85	0.091044	0	0	0	0	0	0	0	0	0	0	0.058956	0	0	0	0	0
4SN	0	0	0	0	0	0	0	0	0	0	0.8133	0.08	0	0.014	0	0	0	0	0	0	0	0	0	0	0	0	0	0.006	0.0867
5SN	0	0	0	0	0	0	0	0	0	0	0	0.95	0.01	0.028	0	0	0	0	0	0	0	0	0	0	0	0	0	0.012	0
6SN	0	0	0	0	0	0	0	0	0	0	0	0.06	0.92	0.014	0	0	0	0	0	0	0	0	0	0	0	0	0	0.006	0
7SN	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
1WH	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0.7	0.25	0	0	0.05	0	0	0	0	0	0	0	0	0
2WH	0	0.000027	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0.493987	0.493987	0	0	0	0	0.012	0	0	0	0	0	0
3WH	0	0	0.000262	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0.739738	0.15	0	0	0	0.014	0	0	0	0	0.096
4WH	0	0	0	0.032768	0	0	0.004	0	0	0	0	0	0	0.002	0	0	0	0.731232	0.08	0	0.01	0	0	0	0.016	0	0	0.004	0.12
5WH	0	0	0	0	0.0016	0	0.0016	0	0	0	0	0	0	0.0008	0	0	0	0	0.1848	0.002	0.004	0	0	0	0	0.0036	0	0.0016	0.8
6WH	0	0	0	0	0	0.05	0.004	0	0	0	0	0	0	0.002	0	0	0	0	0.06	0.85	0.01	0	0	0	0	0	0.02	0.004	0
7WH	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0
1SH	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0.3	0.2	0	0	0.5	0	0
2SH	0	0	0	0	0	0	0	0	0.000862	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0.499569	0.499569	0	0	0	0
3SH	0	0	0	0	0	0	0	0	0	0.002204	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0.732989	0.129351	0	0.135456
4SH	0	0	0	0	0	0	0	0	0	0	0.06727	0	0	0.006	0	0	0	0	0	0	0	0	0	0	0.74603	0.08	0	0.014	0.0867
5SH	0	0	0	0	0	0	0	0	0	0	0	0.032768	0	0.012	0	0	0	0	0	0	0	0	0	0	0	0.237232	0.01	0.028	0.68
6SH	0	0	0	0	0	0	0	0	0	0	0	0	0	0.006	0	0	0	0	0	0	0	0	0	0	0	0	0.06	0.92	0.014	0
7SH	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0
Evict	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1

Figure 6. Estimated A matrix for the example analysis.

signals	1WL	2WL	3WL	4WL	5WL	6WL	7WL	1SL	2SL	3SL	4SL	5SL	6SL	7SL	1WH	2WH	3WH	4WH	5WH	6WH	7WH	1SH	2SH	3SH	4SH	5SH	6SH	7SH
f1,L	0	0.27	0.256	0	0	0	0	0	0.309825	0.276556	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
f2,L	0	0.07	0	0.48	0	0	0	0	0.139825	0	0.4233	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
f3,L	0	0	0.096	0	0.8	0	0	0	0	0.135456	0	0.68	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
f4,L	0	0	0	0.08	0	0	0	0	0	0	0.0833	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
f5,L	0	0	0.016	0	0	0	0	0	0	0.026656	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
f1,f2,L	0	0.63	0	0	0	0	0	0	0.455175	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
f1,f3,L	0	0	0.384	0	0	0	0	0	0	0.287844	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
f1,f5,L	0	0	0.064	0	0	0	0	0	0	0.056644	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
f2,f4,L	0	0	0	0.12	0	0	0	0	0	0	0.0867	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
f3,f5,L	0	0	0.024	0	0	0	0	0	0	0.027744	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
f1,f3,f5,L	0	0	0.096	0	0	0	0	0	0	0.058956	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
none,L	1	0.03	0.064	0.32	0.2	1	1	1	0.095175	0.130144	0.4067	0.32	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0
f1,H	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0.27	0.256	0	0	0	0	0	0.309825	0.276556	0	0	0	0
f2,H	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0.07	0	0.48	0	0	0	0	0.139825	0	0.4233	0	0	0
f3,H	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0.096	0	0.8	0	0	0	0	0.135456	0	0.68	0	0
f4,H	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0.08	0	0	0	0	0	0	0.0833	0	0	0
f5,H	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0.016	0	0	0	0	0	0	0.026656	0	0	0	0
f1,f2,H	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0.63	0	0	0	0	0	0	0.455175	0	0	0	0	0
f1,f3,H	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0.384	0	0	0	0	0	0	0.287844	0	0	0	0
f1,f5,H	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0.064	0	0	0	0	0	0	0.056644	0	0	0	0
f2,f4,H	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0.12	0	0	0	0	0	0	0.0867	0	0	0
f3,f5,H	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0.024	0	0	0	0	0	0	0.027744	0	0	0	0
f1,f3,f5,H	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0.096	0	0	0	0	0	0	0.058956	0	0	0	0
none,H	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0.03	0.064	0.32	0.2	1	1	1	0.095175	0.130144	0.4067	0.32	1	1

Figure 7. Estimated B matrix for the example analysis.

IV. EXPANDING TO THE SYSTEM LEVEL

At the system level, we represent the network as shown in Figure 1, but with each node expanded using a representation like the one in Figure 4. There may be several potential entry nodes to the system, and we denote by e_v the probability that a given attack is initiated at node v. Although data are being collected on the topic of information intrusion and attack [16], there is little publicly available analysis on the arrival distribution of these attacks or a characterization of the arrival process. We assume that attackers arrive as a Poisson process with overall rate λ. This assumption is consistent with a premise of uncoordinated attacks being mounted by individuals out of a large population and is similar to assumptions made in other adversarial situations (e.g., [17], [18]). The arrival rate of attacks at entry node v is then $e_v\lambda$. Each of these node-specific arrival rates is then further broken down into arrival rates by attackers in category $c \in C$.

We can create the basic building blocks of an MDP model of the attacker's strategy, using the state diagram shown in Figure 4. Attackers arrive at the "attack v" state at some rate ($e_v\lambda$ if node v is an "entry" node and via transitions from other nodes if node v is an "internal" node). At each transition, there are probabilities $p_{cc'}(v)$ that the attack is successful (i.e., transition to a "Breach v" state) and probabilities $d_c(v)$ that the attacker will be evicted. If neither success nor eviction occurs, the attacker remains in the "attack v" state, continuing his/her attack. The "Eviction" state is an absorbing state, and there is an associated cost to the attacker, which we denote as $\xi_c(v)$. If a "Breach v" state is reached, the attacker inflicts some loss, $\theta_c(v)$, on the system owner and then has choices about what to do next.

Breaching node v generally offers an opportunity to attack another node, i, in the network. This opportunity is represented by links from the "Breach v" states to "Attack i" states for a corresponding attacker category. These links are assumed to have an immediate cost to the attacker denoted by $s_c(v,i)$. This immediate cost represents the risk of detection associated with that transition. The attacker also has the choice of quitting (exiting from the system), presumably with an immediate cost $s_c(v,j) = 0$ (where j denotes an exit state).

We can pose the problem of finding the optimal attack strategy as an MDP over an infinite horizon. We define the expected reward to the attacker to be the expected loss that can be inflicted upon the system operator minus the expected costs to the attacker (resulting from risk of detection and associated penalties). As mentioned in Section II, there may be some categories of attackers (e.g., terrorists) for whom we would set the perceived expected costs of detection to be zero. Such an attacker would not voluntarily quit his/her attack, but is subject to the same mechanisms of detection and eviction. Also, as mentioned in Section II, a utility function that reflects risk-prone or risk-averse behavior can be substituted for the expected reward calculation in a straightforward way.

We assume that the objective of the attacker is to maximize his/her expected reward (or utility), and we examine the problem of finding the optimal attack strategy for this objective. Solving this problem positions us to adopt the perspective of the system operator and consider the actions that can have the largest impact on reducing the expected losses resulting from such attacks. We can also consider the potential for insuring the system against financial losses from such intentional attacks.

If the attacker is in state i and chooses action a_i, we denote the expected value of the future stream of rewards by $w(i,a_i)$. Each possible action a_i implies an immediate reward value $R_i(a_i)$ and a change in the transition probabilities that govern the process. We denote the elements of the transition matrix resulting from choosing action a_i as $P_{ij}(a_i)$. The MDP is positive bounded. At each breach node, there is always a possible decision with non-negative expected total return because the attacker can always choose to quit. It is bounded because the process is absorbed into one of two states ("Quit" or "Evict") that have $R_i(a_i) = 0$. Thus, the sum of future expected rewards is bounded from above. The absorption into states with $R_i(a_i) = 0$ also ensures that an optimal stationary deterministic policy exists [19], and that it is conserving. As a result, we can find the optimal policy through either policy iteration or linear programming.

From a computational standpoint, policy iteration is generally preferable to linear programming for finding solutions, but the linear programming formulation can yield an insight that is significant for our current purposes, so we proceed along that line. Puterman [19] describes the linear programming formulation for positive bounded expected total reward models. The formulation seeks the decision policy (choice of a_i) that maximizes the expected value of the reward stream, $w(i,a_i)$. We denote the resulting optimal expected value as $w^*(i)$.

As Puterman [19] describes in detail, the set of $w^*(i)$ is the smallest set of values of $w(i)$ for which the following inequalities hold for all states, i:

$$w(i) \geq R_i(a_i) + \sum_j P_{ij}(a_i) w(j) \tag{14}$$

If we then introduce an arbitrary set of positive scalars, β_i , with the requirement that $\sum_i \beta_i = 1$, the linear program can be written as follows:

$$\min \sum_i \beta_i w(i) \tag{15}$$

$$\text{subject to:} \qquad w(i) - \sum_j P_{ij}(a_i) w(j) \geq R_i(a_i) \qquad \forall\, i,\, a_i \tag{16}$$

$$w(i) \geq 0 \qquad \forall\, i \tag{17}$$

This linear program has a dual that can be expressed as follows:

$$\max \sum_i \sum_{a_i} R_i(a_i) x_i(a_i) \tag{18}$$

$$\text{subject to:} \quad \sum_{a_i} x_i(a_i) - \sum_j \sum_{a_i} P_{ij}(a_i) x_i(a_i) \le \beta_i \qquad \forall i \tag{19}$$

$$x_i(a_i) \ge 0 \quad \forall i, a_i \tag{20}$$

The primal linear program has many more constraints than variables, so it is more effective to solve the dual problem. In addition, it can be shown (see [19]) that in an optimal solution to the dual problem (18)–(20), there is no more than one non-zero $x_i(a_i)$ for each state i. The a_i for which $x_i(a_i)$ is non-zero indicates the optimal action a_i^* for each i.

The optimal values of the non-zero dual variables, $x_i(a_i)$ for each state i, indicate the "shadow prices" for affecting the rewards, $R_i(a_i)$. This provides useful information about the relative value of different "hardening" strategies that might be applied in the system. This is discussed further in Section VII.

In the application context of interest here, the state space for the MDP is a collection of "attack" and "breach" states (as shown in Figure 4) for the nodes in the TOE. For the "attack" states, there is a single action possible (i.e., no decisions are made). The immediate reward for this action is $R_i(a_i) = -d_c(v)\xi_c(v)$, where state i refers to an attack at node v by an attacker of category c. For the "breach" states, there are several possible decisions (i.e., attack other nodes k, or quit). For a decision to attack node k, the immediate reward is $R_i(a_i) = \theta_c(v) - s_c(v,k)$, that is, the value of the loss inflicted at node v (which the attacker has just breached) less the expected cost associated with moving from node v to node k. For a decision to quit, the immediate reward is $\theta_c(v)$. The change in the transition probabilities that accompany any decision is simply that the probability of transition to the chosen state ("attack k" or "quit") becomes 1.0.

To see how these ideas are applied in the example shown in Figure 1 (with two categories of attackers and two levels of monitoring at each node, as in the previous section), we create the state diagram shown in Figure 8. (Some of the arcs are not shown completely to help prevent confusion in the diagram among arcs that appear to cross one another.)

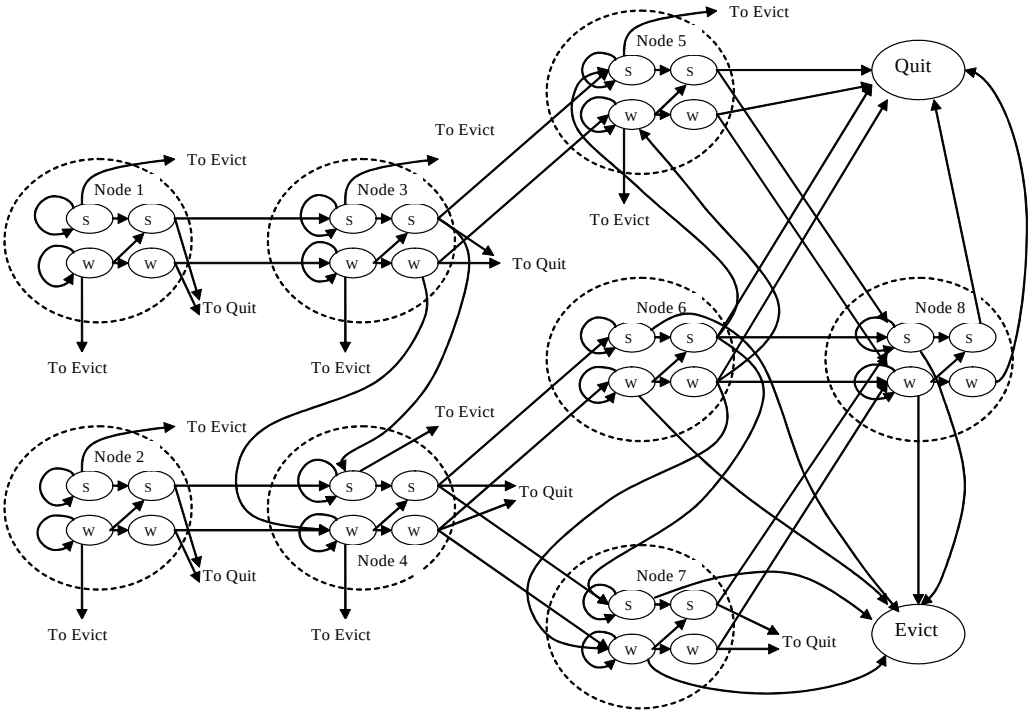

Figure 8. MDP state network and transitions.

The transition probabilities for this Markov chain come from the computations of $p_{WW}(v)$, $p_{WS}(v)$, $p_{SS}(v)$, $d_W(v)$, and $d_S(v)$ for nodes 1 through 8 in Figure 1. A summary of these values is given in Table 2. Note that the values shown for node 1 are the values we computed in Section III. Table 2 also provides information on the penalty to the attacker for detection at each node and the losses to the system operator if a given node is breached. Both the losses and potential penalties increase as the attacker proceeds "deeper" into the system, and the detection probabilities are also larger at the deeper nodes, reflecting somewhat tighter security at those locations.

Table 3 specifies the probability of detection, $r_c(v, k)$ associated with various moves that the attacker might make within the system. These probabilities are used to compute the expected costs to the attacker in category *c* associated with making a specific move from node *v* to node *k*, $s_c(v, k)$. For this example, we compute those expected costs as:

$$s_c(v, k) = r_c(v, k)\, \xi_c(k) \tag{21}$$

That is, the expected cost is the probability of detection during the move times the penalty cost associated with detection at the intended destination node. Other methods of specifying the $s_c(v, k)$ costs could also be used, but this is a reasonable and simple way of determining them.

In this example (see Table 2), we are making no distinction in the penalties, $\xi_c(k)$, among the categories of attackers, *c*. However, the formulation is general, and would allow such distinctions if desired.

The solution is summarized in Table 4. If node 3 or node 4 is breached, the optimal decision depends on the attacker category. A weak attacker will try to penetrate the system through nodes 5 and 7, while a strong attacker will attack through nodes 4 and 6 before going to node 7. This strategy reflects an opportunity for a strong attacker to inflict more potential damage on the system by proceeding through nodes 4 and 6, and the increased risk to the strong attacker is less than the increase in expected losses inflicted. The solution to the linear program representing the MDP provides a strategy that reflects an optimal policy for each category of attacker from any position in the network.

The existence of this strategy does not mean that all attackers will always proceed in exactly the way indicated. It does mean that if all attackers were rational and well informed (in the sense described at the beginning of the paper), this would be a strategy through which they could inflict the greatest amount of expected damage to the system. We can compute expected losses to the system in a conservative way by assuming that the system operator is always facing rational well-informed attackers who are optimizing their attack strategy. In reality, the overall pattern of attacks is likely to be less damaging than this because many attackers will have less information and will not necessarily optimize their strategy.

One very direct way to incorporate imperfect information on the part of the attackers in the analysis is to embed the MDP model in a simulation where uncertainty in the perceptions of the losses, $\theta_c(v)$, is reflected. This is one type of limitation on the information assumed to be available to the attackers. We might assert that an attacker in category c bases his/her strategy on a perception of $\theta_c(v)$ that may or may not be correct. Variations in the perceptions of the losses to be inflicted on the system operator by breaching specific nodes can lead to different attack strategies for different attackers in the same category, and the effect (from the system operator's perspective) is that the overall attacks appear to be following a mixed (or randomized) strategy. This form of simulation is a step in the general direction of considering the system to be a partially observable Markov decision process (POMDP) from the perspective of the attacker.

Suppose that a given attacker in category *c* perceives the loss at node *v* to be a Normal random variable $\theta'_c(v)$, with $E[\theta'_c(v)] = \theta_c(v)$, and standard deviation $\sigma_c(v)$. If we sample each of the $\theta'_c(v)$ distributions, we have a set of perceived losses, and we can then solve the MDP for the optimal strategy under that sample of perceptions. By repeating the process for many samples and recording the strategy for each category of attackers in each sample, we can construct an estimate of the probability of a given strategy for each attacker category.

Table 2. Example data for network nodes.

Node (see Figure 1)	Prob. of Success $p_{WW}(v)$	Prob. of Success $p_{WS}(v)$	Prob. of Success $p_{SS}(v)$	Prob. of Detection $d_W(v)$	Prob. Of Detection $d_S(v)$	Penalty for Detection $\xi_W(v)$	Penalty for Detection $\xi_S(v)$	Loss for Breach $\theta_W(v)$	Loss for Breach $\theta_S(v)$
1	.014	.016	.03	.029	.006	100	100	100	200
2	.009	.005	.011	.01	.006	100	100	100	200
3	.008	.004	.01	.008	.005	500	500	200	400
4	.007	.004	.009	.008	.005	500	500	200	400
5	.007	.003	.008	.01	.007	1000	1000	500	1000
6	.006	.003	.007	.01	.007	1000	1000	500	1000
7	.005	.003	.006	.01	.006	1000	1000	500	1000
8	.005	.002	.006	.1	.04	5000	5000	1000	2000

Table 3. Probability of detection for possible moves.

Arc	Prob. of Detection (r_{vk}) Weak Attacker	Prob. of Detection (r_{vk}) Strong Attacker
1-3	.01	.005
2-4	.02	.01
3-4	.05	.025
3-5	.1	.05
4-6	.07	.035
4-7	.07	.035
5-8	.1	.05
6-5	.02	.01
6-7	.03	.015
6-8	.04	.02
7-8	.1	.05

Table 4. Optimal decisions from MDP solution.

If this node has just been breached:	And the Attacker Category is:	Then do this next:
1	Weak	Attack 3
1	Strong	Attack 3
2	Weak	Attack 4
2	Strong	Attack 4
3	Weak	Attack 5
3	Strong	Attack 4
4	Weak	Attack 7
4	Strong	Attack 6
5	Weak	Attack 8
5	Strong	Attack 8
6	Weak	Attack 5
6	Strong	Attack 5
7	Weak	Attack 8
7	Strong	Attack 8

For example, in the test network shown in Figure 8, if we assume that weak attackers have a perception of losses that has a coefficient of variation of 0.3 at each node, and strong attackers have perception distributions with a coefficient of variation of 0.2 at each node, we can construct the solution summarized in Table 5 via this type of simulation.

Table 5. Optimal decisions from simulation of MDP solution with imperfect information on the part of attackers.

If this node has just been breached:	And the Attacker Category is:	Then do this next:	With this probability:
1	Weak	Attack 3	1.0
1	Strong	Attack 3	1.0
2	Weak	Attack 4	1.0
2	Strong	Attack 4	1.0
3	Weak	Attack 4	0.33
3	Weak	Attack 5	0.67
3	Strong	Attack 4	0.7
3	Strong	Attack 5	0.3
4	Weak	Attack 7	0.63
4	Strong	Attack 6	0.67
4	Strong	Attack 7	0.33
5	Weak	Attack 8	1.0
5	Strong	Attack 8	1.0
6	Weak	Attack 5	1.0
6	Strong	Attack 5	0.7
6	Strong	Attack 7	0.3
7	Weak	Attack 8	1.0
7	Strong	Attack 8	1.0

In this example, the introduction of imperfect information creates mixed strategies for both categories of attackers, but not at the same subsets of nodes. In each case, the strategy identified in the deterministic analysis is the most likely, but there are significant variations in attack strategies at nodes 3, 4, and 6.

The simulation approach can also be used to analyze other types of imperfect information on the part of attackers – for example, imperfect knowledge of what arcs exist in the network for movement among nodes, or even imperfect information of what nodes exist. We have not pursued analysis of all these possibilities for the example network, but the process is straightforward.

V. Expected Losses to the System Operator

The solution to the MDP (either deterministically or using simulation to reflect imperfect information) represents a strategy that an attacker could follow to maximize his/her expected reward. Because the attacker's reward is directly related to the expected losses that can be inflicted on the system operator, this strategy represents a reasonable basis for estimating the magnitude of those losses. The solution to the MDP provides the transition matrix, P^*, whose elements are $P_{ij}(a_i^*)$. The Markov model that underlies the attacker strategy is a transient model; eventually, any attacker ends up in one of the two absorbing states: "eviction" or "quit." If we add to the model an "Entry" state that delivers attackers to the entry nodes of the TOE at some rate, we can analyze the expected loss rate to the system operator.

We can partition the system states into entry states, interior states, and exit states; the transition matrix can be partitioned as:

$$P^* = \begin{bmatrix} 0 & D^* & 0 \\ 0 & Q^* & Z^* \\ 0 & 0 & I \end{bmatrix} \tag{22}$$

In (22), the submatrix D^* represents transition probabilities from the entry state to interior states; Q^* represents transitions among interior states; Z^* represents transitions from interior states to exit (absorbing) states; and I (the identity matrix) reflects the absorption in the exit states.

Using this notation, the expected number of visits to an interior state j, given that the process started in state i, is given by the ij element of $(I-Q^*)^{-1}$. Then, if the overall arrival rate of attackers is λ and the vector D^* contains the probabilities of system entry at various nodes by the various categories of attackers, the expected total visits per time period to the interior (transient) states are the entries in the row-vector $\lambda D^* \left(I - Q^*\right)^{-1}$. If we then create a loss vector, L (assumed to be a column vector), whose elements are $\theta_c(v)$ for the "Breach v" states in the set of transient states and zero otherwise, the expected loss per time period, $E(\Theta)$, is:

$$E(\Theta) = \lambda D^* \left(I - Q^*\right)^{-1} L \tag{23}$$

For the simple example described in Section IV, if the average arrival rate of attacks (λ) is three per day, 90% of the attackers are weak upon entering, and all attackers follow the optimal attack strategy shown in Table 4, the expected loss calculated using (23) is $1324 per day. Said another way, the average attacker costs the system $441.

With the calculation of expected loss in (23), we are now in a position to address three very important questions:

1. Is the system insurable against these losses, using the economic concepts of insurance theory?
2. How might the system operator act to reduce the expected loss (e.g., by changing security parameters, detection thresholds, etc.)?
3. Can we estimate the marginal value of specific measures to "harden" the system in terms of changes in expected losses as a basis for determining how much hardening is worthwhile?

In the following sections, we explore these questions.

VI. INSURING THE SYSTEM

Suppose that we have an insurance company that is willing to write a contract to reimburse the TOE owner for damages an adversary may inflict on system assets. This may also be a self-insurance structure, in which the TOE owner sets aside some amount of money per period into a reserve fund to cover losses sustained. Equation (23) in the previous section specifies the expected loss rate. This can be interpreted as either the expected cost to the insurance company per period or the size of the set-aside amount that the TOE owner would have to contribute each period to maintain the self-insurance.

The insurance underwriter faces some risk, however, since the losses in a given period may not always correspond to expectation. The aggregate loss in any time period is the sum of losses from many individual attackers, so using the Central Limit Theorem, the aggregate loss per time period is approximately Normally distributed, with mean given by (23). To determine the variance, we recognize that the per-period loss is the sum of a random number of i.i.d. random variables (losses from a random number of attackers, each of whom inflicts a random loss on the system). If we use N to denote the random number of attackers and ζ_i to denote the random loss from attacker i, the general form of such a variance is (see, for example, [14]):

$$\begin{aligned} Var(\Theta) &= Var\left(\sum_{i=1}^{N} \zeta_i\right) \\ &= E(N)Var(\zeta) + \left[E(\zeta)\right]^2 Var(N) \end{aligned} \tag{24}$$

The assumption of Poisson arrivals with a constant arrival rate allows us to argue that the number of attackers, N, is Poisson, and under this condition (24) simplifies to:

$$Var(\Theta) = \lambda E[\zeta^2] \tag{25}$$

where $\lambda = E(N)$.

To find $E[\varsigma^2]$, the expectation of the square of the loss from a single attacker, we reconsider the matrix $\left(I - Q^*\right)^{-1}$, whose elements we denote as ϕ_{ij}^*. The probability that a transient Markov chain ever makes a transition into state *j*, given that it started in state *i*, is (see, for example, [14]):

$$f_{ij} = \frac{\phi_{ij}^* - \delta_{ij}}{\phi_{jj}^*} \tag{26}$$

where δ_{ij} is 1 if $i = j$ and zero otherwise. Since we know that the Markov chain starts in the "entry" state and we can arbitrarily denote that as state 1, we are interested in the values of:

$$f_{1j} = \frac{\phi_{ij}^*}{\phi_{jj}^*} \tag{27}$$

If we let L_j denote the loss if state *j* is entered and a single attacker will inflict a given loss at most once (i.e., the Markov chain for a single attacker is acyclic), the probabilities in (27) determine a probability distribution of losses from a single attacker.

We can then compute the expected squared loss from a single attacker as follows:

$$E[\varsigma^2] = \sum_j f_{1j} L_j^2 \tag{28}$$

Substituting (27) and (28) into (25), we have the result that the variance of loss per time period is:

$$Var(\Theta) = \lambda E[\varsigma^2] = \lambda \sum_j f_{1j} L_j^2 = \lambda \sum_j \frac{\phi_{1j}^*}{\phi_{jj}^*} L_j^2 \tag{29}$$

Equations (23) and (29) then characterize the Normal distribution of loss per time period. Suppose the insurance company (or the TOE owner) had set aside an initial reserve of *U* dollars, and the per-period contributions (or insurance premiums) are π dollars. Claims Θ_t are paid each period *t*, and if we assume that losses are independent of one another, the sequence of per-period changes in available assets, $\pi - \Theta_t$, is a Wiener process. For such a process, as long as $\pi > E\left(\Theta\right)$, the ultimate ruin probability is (see, for example, [20]):

$$\begin{aligned} \rho(U) &= \lim_{t \to \infty} \Pr\left(U + \pi t - \sum_{i=1}^{t} \Theta_i < 0 \right) \\ &= e^{\frac{2\left(E(\Theta) - \pi\right)U}{Var(\Theta)}} \end{aligned} \tag{30}$$

The value of $\rho(U)$ can be used to determine the appropriate per-period premium π to cover a return on the reserves, *U*, and to ensure that the risk of ruin is sufficiently small.

It is also important to explore whether it is better to insure the system to cover expected losses, or to invest money in system hardening to reduce those losses. To begin such an evaluation, it is important to identify the places in the system that have the greatest leverage for reducing expected losses. This is the topic of the next section.

VII. Sensitivity Analysis of Expected Losses

The solution to the MDP for the attackers (found through the dual linear programming formulation in equations (18)–(20)) specifies a set of non-zero dual variables, $x_i(a_i)$, one for each state *i*. These dual variables indicate the "shadow prices" for affecting the rewards, $R_i(a_i)$, that appear on the right-hand side of the constraints in the primal linear programming problem. This provides useful information about the relative value of different strategies that might be applied to reduce expected losses.

The state space for the MDP is a collection of "attack" and "breach" states (as shown in Figure 4) for the nodes in the TOE. For the "attack" states, the immediate reward is $R_i(a_i) = -d_c(v)\xi_c(v)$, where state *i* refers to an attack at node *v* by an attacker of category *c*. For the "breach" states, a decision at node *v* to attack another node *k* has immediate reward $R_i(a_i) = \theta_c(v) - s_c(v,k)$, that is, the value of the loss inflicted at node *v* (which the attacker has just breached) less the expected cost associated with moving from node *v* to node *k*. In the example of interest here, we have computed $s_c(v, k) = r_c(v, k)\,\xi_c(k)$. Thus, the mechanisms available to reduce $R_i(a_i)$ are:

a) increase the detection probability, $d_c(v)$;
b) increase the penalty for detection, $\xi_c(v)$;
c) reduce the potential node loss, $\theta_c(v)$; or
d) increase the movement detection probability, $r_c(v, k)$.

By rank ordering the $x_i(a_i)$ values from largest to smallest, we can develop a list of places in the system where changes can have the greatest benefit and then examine our list of mechanisms for reducing $R_i(a_i)$ to identify the most effective action at that location.

For example, for the system analyzed in Section IV, there are seven $x_i(a_i)$ values greater than 1.0 at the states shown in Table 6. It is quite clear that the major losses to the system are because of the "strong" attackers, even though they constitute only 10% of the total attacks at entry. It is also clear that nodes 4, 5, and 6 are the most important places to focus attention on system hardening against these strong attackers. At the "attack" nodes, the mechanisms of interest are (a) and (b) from the list above, so we have a clear indication from the analysis of what strategies are likely to be most effective for system hardening, and where they should be focused.

Table 6. Identification of largest dual variables.

State (i)	Value of $x_i(a_i)$
Strong Attack, Node 5	15.19
Strong Attack, Node 4	12.72
Strong Attack, Node 6	11.96
Strong Attack, Node 3	6.19
Strong Attack, Node 7	2.31
Strong Attack, Node 2	1.79
Weak Attack, Node 2	1.76

In general, the detection probability, $d_c(v)$, can be increased by expanding the sets of observations that result in increased monitoring attention (ψ_v) and in eviction of users (Ω_v), as discussed in Section III. The sets Ω_v and ψ_v reflect the security protocols in place at a given system node, v. Making these sets larger (especially Ω_v) increases the probability of detection of attackers.

Using the analysis framework created in this paper, we can trace the effects of a specific change in the protocols at nodes 4, 5, and 6, for example, through to a change in expected losses for the system operator, and thus we have a quantitative (and perhaps even monetary) measure of the effectiveness of the change. As an illustration, suppose that the system operator made changes at nodes 4, 5, and 6 so that the detection probability for strong attackers at those nodes increased by 25% in the system as studied in Section IV. Tracing this change through the MDP solution shows that it would reduce the expected losses (as calculated in Section V) to $1145 per day, a reduction of $179 per day (or about 13.5%) from the original conditions.

We must note that the benefit of reduced losses achieved by making the sets Ω_v larger at nodes 4, 5, and 6 must be balanced against the cost of increased "false alarms." As noted in Section III, increasing the set of signals that will cause eviction of a user is also likely to make it more difficult for legitimate users to accomplish their work. Although it is often difficult to assign a cost to a false alarm, it may be possible to estimate the cost of lost work in some situations and adjust the estimated benefits accordingly. This is an area of potential further work.

VIII. CONCLUSION

We have created a framework for analyzing expected losses from malicious attacks against an infrastructure system. This framework also allows the system operator (and a potential insurer) to evaluate the insurability of the system against such losses and to identify where the most effective changes in the system can be made to "harden" it against attacks. We have included an example of a SCADA-related information system to illustrate the ideas, but the framework should have much broader applicability, to both information systems and physical facility infrastructure.

The basis for our analysis is a representation of the system of interest as a network of nodes and arcs. Nodes represent system assets, and arcs represent opportunities for attackers to move within the system. We construct an HMM to represent an attack at a single node in the system. Then we develop an aggregated representation of that single-node model for inclusion in an MDP model of attacker strategy within a network representation of the entire system. The MDP solution is used to compute expected losses from different classes of attackers, as a means of tying the analysis to the notion of "insurability." Finally, the sensitivity information from the MDP solution is used to indicate the parts of the system in which "hardening" against attacks may be most effective. To our knowledge, this is the first effort to use HMMs and MDPs in this way to evaluate economic losses from malicious attacks in systems and to assess potential benefits of hardening measures.

The framework we have created offers several important directions for further research. Certainly one of these directions is further work on analyzing attacks by adversaries with imperfect information. We have described one approach, using simulation, to incorporate some aspects of imperfect information in the analysis, but there is much additional work in that direction that is likely to be very useful.

One potentially useful approach is to treat the system (from the perspective of the attacker) as a partially observable Markov decision process (POMDP). While the HMM provides partial observability of the system into the attack process, the POMDP provides partial observability of the attacker into the system. We could then represent the uncertainty with which the attacker and system owner face each other. Perfect observability of the system by the attacker represents the threat level posed by an insider, such as a former system administrator, while partial observability can provide for a more realistic representation of any number of attacker classes consisting of system outsiders, such as hackers, who have varying, but imperfect levels of system knowledge.

The level of detail that is needed to accurately represent the uncertainty faced by the attacker through the POMDP will need to be considered carefully however, because the model has the potential to become intractable [21]. Use of parallel computing architectures to enable large simulations might be employed to model and understand the effects of an attack against a region of the United States, for example. Extensions of this

work to develop an optimal strategy for multi-agent attackers under uncertainty is another promising potential outcome of more detailed investigation of the POMDP solution structure.

A second direction for enhancing the characterization of attackers is in relaxing the assumption that all attackers seek to maximize their expected total reward (or maximize the expected damage they can do). Treating some attackers as risk-averse and others as risk-prone is likely to make the analysis richer in providing insight on loss characteristics.

Third, there is much room for further exploration of defensive strategies, following the lines laid out in Section VII. We have identified four basic types of changes that can be implemented to reduce the expected rewards to attackers (and hence, losses to the system): (1) increase the detection probability, $d_c(v)$; (2) increase the penalty for detection, $\xi_c(v)$; (3) reduce the potential node loss, $\theta_c(v)$; or (4) increase the movement detection probability, $r_c(v, k)$. In particular, strategies that have combined effects on more than one of these basic parameters are likely to be effective. For example, an operator strategy that distributes the valuable pieces of the systems (as measured by $\theta_c(v)$), thus spreading out the assets and reducing the potential losses, combined with adding more barriers, or transitions, before reaching a particular valued node, could be an effective means to defend the system.

Finally, there is a need for significant empirical work to verify and validate, or to uncover significant weaknesses in, the analysis framework developed here. It is important to develop additional experience in estimating parameters for the types of models described here and to determine the scalability of the approach for very large networks. In general, good data for analyses of system security are hard to collect, and the model described here requires significant supporting data. Collection of appropriate data will require some effort, but the overall approach appears to be a promising direction for understanding and improving security of networked infrastructure, and thorough empirical testing is an important next step.

REFERENCES

[1] Albrecher, H. *Dependent Risks and Ruin Probabilities in Insurance*. Technical Interim Report IR-98-072, International Institute for Applied Systems Analysis, Laxenburg, Austria, September 1998. <http://www.iiasa.ac.at/cgi-bin/pubsrch?IR98072>

[2] Phillips, C.A., and Swiler, L.P. A Graph-Based System for Network-Vulnerability Analysis. In ACM *Proceedings for the 1998 New Security Paradigms Workshop*, 1998, pp. 71-81.

[3] Dacier, M., *Towards a Quantitative Evaluation of Computer Security*. PhD Thesis, Institut National Polytechnique de Toulouse, Toulouse, France, 1994.

[4] Swiler, L.P., Phillips, C.A., Ellis, D., and Chakerian, S. Computer-Attack Graph Generation Tool. *Proceedings of the 2nd DARPA Information Survivability Conference and Exposition*, 2001, vol 2, pp. 307-321.

[5] Jha, S., Sheyner, O., and Wing, J.M. *Minimization and Reliability of Attack Graphs*. Research Report CMU-CS-02-109, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, 2002. <http://reports-archive.adm.cs.cmu.edu/anon/2002/CMU-CS-02-109.pdf>

[6] Sheyner, O., Haines, J., Jha, S., Lippmann, R., and Wing, J.M. Automated Generation and Analysis of Attack Graphs. *Proceedings of the IEEE Computer Society Symposium on Research in Security and Privacy*, Berkeley, CA, May 2002, pp. 273-284.

[7] Katsikas, S.K., Gritzalis, D., and Spirakis, P. Attack Modelling in Open Network Environments. In *Communications and Multimedia Security II*, 1996, pp. 268-277.

[8] Katsikas, S., Spyrou, T., Gritzalis, D., and Darzentas, J. Model for Network Behaviour under Viral Attack. *Computer Communications*, February 1996, vol. 19, no. 2, pp. 124-132.

[9] Soh, B.C., and Dillon, T.S. Setting Optimal Intrusion-Detection Thresholds. *Computers & Security*, 1995, vol. 14, no. 7, pp. 621-631.

[10] Warrender, C., Forrest, S. and Pearlmutter, B. Detecting Intrusions Using System Calls: Alternative Data Models. In *Proceedings of the 1999 IEEE Symposium on Security and Privacy*, May 1999, pp. 133-45.

[11] Ourston, D., Matzner, S. Stump, W., Hopkins, B., and Richards, K. Identifying Coordinated Internet Attacks. In *Proceedings of the Second SSGRR Conference*, Rome, 2001, paper 109.

[12] Ourston, D., Matzner, S. Stump, W., and Hopkins, B. Evaluating Network Intrusion Detection Algorithm Performance as Attack Complexity Increases. In *Proceedings of the Third SSGRR Conference*, L'Aquila, 2002, paper 74.

[13] Elliott, R.J., Aggoun, L., and Moore, J.B. *Hidden Markov Models: Estimation and Control*. Springer, New York, 1995.

[14] Ross, S. *Introduction to Probability Models, 7th edition*. Harcourt – Academic Press, San Diego, 2000.

[15] Bremaud, P. *Markov Chains: Gibbs Fields, Monte Carlo Simulation and Queues*, Springer, New York, 1999.

[16] Howard, J., An Analysis of Security Incidents on the Internet 1989 - 1995, PhD Thesis, Carnegie Mellon University, April 1997. <http://www.cert.org/research/JHThesis/Start.html>

[17] Borodin, A., Kleinberg, J. Raghavan, P., Sudan, M., and Williamson, D. Adversarial Queuing Theory, In *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing*, 22-24 May 1996, pp. 376-385.

[18] Gaver, D., and Jacobs, P. *Stochastic and Deterministic Model of Targeting with Dynamic and Error-Prone BDA*. Technical Report NPS-OR-97-018, Accession Number: ADA331679, Naval Postgraduate School, Institute for Joint Warfare Analysis, Monterey, CA, September 1997.

[19] Puterman, M.L. *Markov Decision Processes*. Wiley, New York, 1994.

[20] Bohman, H. *Risk Theory and Wiener Processes*. ASTIN Bulletin, Vol. VII, Part I, December 1972, pp. 96-99. <http://www.casact.org/library/astin/vol 7, no 1/>

[21] Mundhenk, M., Goldsmith, J., Lusena, C., and Allender, E. Complexity of finite-horizon Markov decision process problems. *Journal of the ACM*, 47(4):681-720, 2000.

DISTRIBUTION

1	MS0451	Sam Varnado, 5500
1	MS0451	Dean Jones, 6221
1	MS0455	Reynold Tamashiro, 5517
1	MS0455	Laurence Phillips, 5517
1	MS0455	Steven Goldsmith, 5517
1	MS0784	Mike Skroch, 5512
1	MS0784	Jennifer Depoy, 5512
1	MS0785	Robert Hutchinson, 5516
1	MS0785	Timothy S. McDonald, 5514
2	MS1110	Cynthia Phillips, 9215
1	MS1351	Juan Torres, 5517
10	MS1351	Rolf Carlson, 5517
1	MS1371	Bruce Varnado, 4143
1	MS0161	Patent and Licensing Office, 11500
1	MS9018	Central Technical Files, 8945-1
2	MS0899	Technical Library, 9616