

Defense trees for economic evaluation of security investments

Stefano Bistarelli* Fabio Fioravanti Pamela Peretti

Dipartimento di Scienze
Università degli Studi “G. d’Annunzio”
Pescara, Italy

E-mail: {bista, fioravanti, peretti}@sci.unich.it

Abstract

In this paper we present a mixed qualitative and quantitative approach for evaluation of Information Technology (IT) security investments.

For this purpose, we model security scenarios by using defense trees, an extension of attack trees with attack countermeasures and we use economic quantitative indexes for computing the defender’s return on security investment and the attacker’s return on attack.

We show how our approach can be used to evaluate effectiveness and economic profitability of countermeasures as well as their deterrent effect on attackers, thus providing decision makers with a useful tool for performing better evaluation of IT security investments during the risk management process.

1 Introduction

Security has become today a fundamental part of the enterprise investment. In fact, more and more cases are reported showing the importance of assuring an adequate level of protection to the enterprise’s assets.

In order to focus the real and concrete threat that could affect the enterprise’s assets, a risk management process is needed in order to identify, describe and analyze the possible vulnerabilities that must be eliminated or reduced. The final goal of the process is to make security managers aware of the possible risks, and to guide them toward the adoption of a set of countermeasures which can bring the overall risk under an acceptable level.

The determination of the acceptable risk level and the selection of the best countermeasure is unfortunately not an easy duty. There are no standard methodologies for the

process, and often security managers have to decide among too many alternatives. Usually, two possible approaches for the security risk management process can be followed: the qualitative and the quantitative ones. The qualitative approach is based on relative evaluation of risks, whilst the quantitative approach tries to give precise and objective measures of risk.

In this paper we define a methodology to mix the benefit of the two approaches. The qualitative approach will be used to depict security scenarios (via a modified version of attack trees [18, 1, 19, 16]), and quantitative indexes [10, 11] will be used to measure risk.

More in detail, we define *defense trees* by extending attacks trees with countermeasures. We label each node representing a specified vulnerability with a set of countermeasures which mitigate the damage of threats using such a vulnerability. Then, economic indexes are used as labels for the countermeasures.

The *Return on Investment* (ROI) [21, 20] index gives a measure of the efficacy of a specific security investment in a countermeasure w.r.t. a specific attack. The *Return on Attack* (ROA) [4] is instead an index that is aimed at measuring the convenience of attacks, by considering the impact of a security solution on attacker’s behavior.

The paper has the following structure. In Section 2 an introduction to the concepts of security risk management is given. In particular, in Section 2.1 a qualitative approach to scenario analysis based on attack trees is exemplified, and in Section 2.2 the quantitative indexes ROI and ROA are introduced.

The innovative part of the paper starts in Section 3 where we introduce defense trees as an extension with countermeasures of classical attack trees. Then, in Section 4 the defense trees are enriched with economic indicators (ROI in Section 4.1 and ROA in Section 4.2 respectively). After a brief discussion about the selection of the most promis-

*partially supported by: Istituto di Informatica e Telematica, C.N.R., Pisa, Italy. E-mail: stefano.bistarelli@iit.cnr.it

ing countermeasure when only one index (ROI or ROA) is available, a simple approach to compose the two indexes is provided in Section 4.3. Finally, Section 5 summarizes the paper and sketches some directions for future work.

2 Security risk management

The Risk Management process is a fundamental activity in an enterprise since it allows senior managers to make good decisions, thus protecting the organization and its ability to achieve its mission. Many risks can affect an organization's resources: risks related to the political and social environment where the organization works (*strategic risks*); risks related to the money market and interest rate (*financial risks*), and risks related to its business processes (*operative risks*).

In this paper we pay attention to the *Security Risk Management* process [21], that focuses on protecting an enterprise's assets from the *Information Technology Risk* (as part of the operative risk). The Information Technology Risk considers interruption of services, diffusion of reserved information or loss of data stored on IT systems. More precisely, the risk function can be defined as follows.

Definition 2.1 (Security Risk [21]) *The Security Risk is a function of the likelihood of a given threat-source exercising a particular potential vulnerability, and the resulting impact of that adverse event on the organization's assets.*

At the beginning of the Security Risk Management process, the different *assets* that compose the IT system are identified and analyzed.

Definition 2.2 (Asset) *An asset is any tangible or intangible item owned by an organization that has a value for an enterprise and that needs protection.*

During the risk management process, the following phases are performed for each asset: risk assessment, mitigation and monitoring.

Risk Assessment The Risk Assessment phase look for identifying risks, determining the possible damages, quantifying the impact of potential threats and providing an economic balance between the economic impact of risk and the cost of risk mitigation. The output of the risk assessment phase is a report that describes *threats* and *vulnerabilities* that can harm a system, gives measures about the risk and provides recommendations toward the implementation of effective *countermeasures*. Following [9, 21],

- a *threat* is the potential for a threat-source to exercise (by accidental trigger or intentional exploit) a specific vulnerability;
- a *vulnerability* is a flaw or weakness in system security procedures, design, implementation, or internal controls that could be exercised (by accidental trigger or

intentional exploit) by an attack and result in a security breach or a violation of the systems security policy;

- a *countermeasure* is a control which should be implemented in order to reduce the ability for an attacker to leverage existing system vulnerabilities.

Risk Mitigation During the Risk Mitigation a systematic methodology is used by senior management to prioritize, evaluate and implement countermeasures recommended by the risk assessment process. Based on the risk level presented in the risk assessment report, the implementation actions are prioritized. Every alternative solution is analyzed and then the most appropriate and cost-effective ones are selected for actual implementation.

Monitoring The Monitoring phase is the last phase of the risk management process. During this phase actual effectiveness of implemented countermeasures is evaluated.

In this paper we pay attention on the security risk assessment phase where the security officer can follow two approaches: the *qualitative* and the *quantitative* one.

2.1 Qualitative approaches

The *qualitative approach* [6] evaluates the security risk level of an IT system by using a variety of polling, interview, and questionnaire techniques with the aim of comparatively ranking assets and threats according to their perceived criticality and likelihood, respectively. They usually adopt *scenario analysis*, which requires the construction of different scenarios of computer security compromise, in order to illustrate how vulnerable an organization is to information technology attacks [8].

A particular kind of instruments that can be used to conduct a scenario analysis are *attack trees* [18, 19]. Attack trees provide a formal and methodical way of describing how attacks against a system can be performed.

An attack scenario can be represented in a tree-based structure whose root is the attacker's *goal* and paths from leaf nodes to the root represent the different ways of achieving this goal. The root of the tree is associated with an asset of the IT system under consideration. Leaf nodes represent simple subgoals which lead the attacker to (partially) damage the asset by exploiting a single vulnerability. Non-leaf nodes (including the tree root) represent attack subgoals and can be of two different types: *or*-nodes and *and*-nodes. Subgoals associated with *or*-nodes can be achieved by achieving any of its child nodes, whilst *and*-nodes represent subgoals which can only be achieved by achieving all its child nodes.

Each path from leaf nodes to the root ending in an achieved subgoal represents a different attack strategy in the considered scenario. Below we provide examples of how attack trees can be used to model an attack scenario and to identify which vulnerabilities can be exploited in order to harm a system.

Example 1 An enterprise's server is used to store information about customers. We use attack trees to model two different attack scenarios for this asset: (1) theft of data stored on the server (Figure 1), and (2) theft of the server itself (Figure 2).




Figure 1: An example attack tree: theft of data stored on the server.




Figure 2: An example attack tree: theft of the server.

2.2 Quantitative approaches

The *quantitative approach* [15] assigns absolute numeric attribute values to assets, threats, vulnerabilities and countermeasures. The exact identification of risk and the cost/benefit justification of countermeasures are fundamental for constructing a good risk mitigation strategy.

Within this approach several indexes can be used to estimate the effectiveness of an IT security investment.

Definition 2.3 (Single Loss Exposure [11]) The Single Loss Exposure (*SLE*) represents a measure of an organization's loss from a single threat event and can be computed by using the following formula:

$$SLE = AV \times EF$$

where, the Asset Value (*AV*) [11] is a synthetic measure of the cost of creation, development, support, replacement and ownership values of an asset, and the Exposure Factor (*EF*) [10] represents a measure of the magnitude of loss or impact on the value of an asset arising from a threat event, and is expressed as a percentage of the asset value.

Since not all the threats are equally likely to succeed the *SLE* value can be modified by considering the frequency of the given threat.

Definition 2.4 (Annualized Loss Expectancy [11]) The Annualized Loss Expectancy (*ALE*) is the annually expected financial loss of an organization which can be ascribed to a threat and can be computed by using the following formula:

$$ALE = SLE \times ARO$$

where the Annualized Rate of Occurrence (*ARO*) [11] is a number that represents the estimated number of annual occurrences of a threat.

It is important to notice that estimating the *ARO* could be very difficult. It is usually created upon the likelihood of the event and the number of attackers that could exploit the given vulnerability. For example a meteorite damaging the data center could be estimated to occur only once every 100,000 years and will have an *ARO* of 0.000001. In contrast, 100 data entry operators attempting an unauthorized access attempt could be estimated to occur six times a year per operator and will have an *ARO* of 600.

Summarizing the above indexes, *SLE* (and *EF*) gives a measure of the damage of a single threat; the *ARO* gives the likelihood of a threat to occur in a year; *ALE* tries to consider both the likelihood and the damage of each threat.

All of the above indexes do not consider the fact that the organization can try to build some defense for reducing the probability of vulnerability exploitation by attackers (e.g. implementing some firewall filtering), or reducing the damage of an attack (e.g. applying some backup strategies). We will now introduce two indexes able instead to consider also the presence of countermeasures: the Return on security Investment (*ROI*) and the Return on Attack (*ROA*).

The *ROI* can be used for providing an economic evaluation of an enterprise's expenditure in IT security. It can be used to compare alternative investment strategies and to evaluate whether an investment is financially justified, and can be computed using the following formula:

Definition 2.5 (Return on Investment [20]) The Return on Investment (*ROI*) index is defined as:

$$ROI = \frac{(ALE \times RM) - CSI}{CSI}$$

where *RM* is the risk mitigated by a countermeasure and represents the effectiveness of a countermeasure in mitigating the risk of loss deriving from exploiting a vulnerability (expressed as a numeric value in $[0,1]$), and *CSI* is the cost of security investment that an enterprise must sustain for implementing a given countermeasure.

If *ROI* is a positive number, the cost for the investment is financially justified. Otherwise, if *ROI* is zero or a negative number, the investment is not profitable.

Together with ROI we can consider also the Return on Attack index (ROA) proposed in [4], which is aimed at measuring the convenience of attacks considering the impact of a security solution on attacker’s behavior. In fact, ROI alone provides only a partial characterization of IT investments, because it lacks to explicitly consider attackers’ interests. Assuming that the organization’s loss is equal to the attacker gain is often a gross simplification. Also, the cost of an attack cannot be directly related to the cost of the security measure because different solutions at different costs might be perceived as equally expensive to break from the attacker’s viewpoint.

Definition 2.6 (Return on Attack [4]) *The Return on Attack (ROA) is the gain that an attacker expects from a successful attack over the losses that he sustains due to the adoption of security measure S by his target. Is defined as:*

$$ROA = \frac{GI}{cost\ before\ S + loss\ caused\ by\ S}$$

where GI is the expected gain from the successful attack on the specified target.

As shown in Section 4.3, a combined use of ROA and ROI indexes allows us to execute a more complete evaluation of a countermeasure, considering not only its effectiveness and profitability but also the deterrent effect produced on the attacker.

3 Defense trees: adding countermeasures to attack trees

Attack trees can be used as a tool to easily provide a visual representation of an attack scenario, and can be used for scenario evaluation when enriched with attacker’s attributes (e.g. attacker’s competencies, costs, ...) [18, 19]. However, they do not take into account countermeasures which can be implemented by the defending organization and the costs sustained for such security investments.

For this reason we enrich standard attack trees by decorating every leaf node with a set of countermeasures. Each countermeasure associated with a leaf node represents a possible risk mitigation of the scenario showing the use of the specific vulnerability. We call such attack trees decorated with countermeasures *defense trees*.

Definition 3.1 (Defense Tree) *A Defense Tree is built by adding a set of countermeasures to the leaves of an attack tree.*

An example defense tree is presented below.

Example 2 *The attack trees used in Example 1 can be enriched with the possible countermeasures that can be introduced to protect the organization’s server as follows.*

Figure 3 and Figure 4 show some of the countermeasures which can be implemented to reduce the risk of data theft and the risk of server theft.




Figure 3: A defense tree for the attack tree of Figure 1.

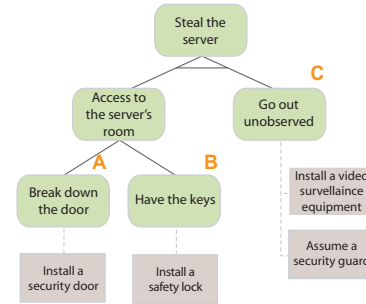


Figure 4: A defense tree for the attack tree of Figure 2.

In order to ease the process of identifying all possible attacks even in deep defense trees, we transform defense trees to defense trees in *disjunctive normal form* (DNF) where the and-nodes are moved towards the leaves of the tree.

In the example provided below, we show how to represent a defense tree in disjunctive normal form.

Example 3 *Figure 5 shows how the same scenario of Figure 4 can be represented by using the disjunctive normal form.*

The equivalence of the two representation derives from the similar equivalence of the logical formula

$$((A \text{ or } B) \text{ and } C) = (A \text{ and } C) \text{ or } (B \text{ and } C)$$

representing exactly the situation in Figure 4 and in Figure 5. A detailed discussion of the equivalent transformation of attack trees can be found in [14]. In the following we will always consider a defense trees in disjunctive normal form.




Figure 5: The defense tree of Figure 4 transformed in DNF.

4 Economic evaluation of threats

In order to obtain a more precise evaluation of attack/defense scenarios we enrich the defense tree modeling the considered scenario by using economic quantitative indexes (like ROI and ROA). Also, we can use this information to make a more informed decision in the selection of the countermeasures that to be implemented for protecting the system. In this way we combine the advantages of attack trees (ease of use, visual modeling of attack scenarios), with the advantages of quantitative approaches (the use of indexes).

The result is a decorated defense tree that can be used to evaluate the security investment that an organization needs to support. By considering both ROI and ROA indexes, we can consider two different points of view: the organization’s view and the attacker’s view. Analyzing the same tree we can consider, at the same time, how a person can attack a system and how an organization can protect it, and we can compare costs for the attacker to achieve his goals and costs for the organization to protect its own systems.

Besides analysing the considered attack scenario from the enterprise’s point of view, using both ROI and ROA, we can use attack trees for providing a scenario evaluation also from an attacker’s perspective. Looking at an attack scenario from the defender’s point of view, we can use ROI to determine what countermeasures are cost effective. On the other hand, by using ROA, we can see the same attack scenario from the attacker’s point of view and determine which are the best attack strategies.

In the following we show, by means of an example, how to label a defense tree in disjunctive normal form with the economic indexes presented in the previous sections.

4.1 Computing ROI: the defender’s point of view

Given a defense tree in DNF, we describe the defender’s point of view by enriching the given tree using economic quantitative labels that can help determine countermeasures to be selected for implementation taking into account the organization’s return on investment. For each asset we want to protect, we proceed as follows:

- first, we consider - for each possible attack represented as a path in the tree - the Exposure Factor (EF) and the Annualized Rate of Occurrence (ARO); this part is very important because using incorrect data could lead to unexpected results;
- then, we estimate the Asset Value (AV) and we compute the Single Loss Exposure to a threat (SLE) and the Annualized Loss Expectancy (ALE) for each leaf in the DNF defense tree. In particular, when the last node before the leaves is an and-node, the computation of SLE and ALE is performed considering the EF and ARO of the and-node (since all the leaf vulnerability have to be exploited and not only one);
- then, we consider the cost of each countermeasure (CSI) and the percentage of Risk Mitigated (RM);
- finally, we compute the Return on security Investment (ROI) for each countermeasure.

We can use ROI to compare economic profitability of the different countermeasures that an organization can use to protect its own systems. For example, for each attack node, we could select the countermeasure which maximizes ROI among all countermeasures which are associated to its vulnerability nodes.

As an example consider the defense tree depicted in Figure 5 reflecting attacks to the server (the asset) and relative mitigation countermeasures. In the example we consider the value of the server estimated in 100.000 €, and the EF and the ARO of each attack as showed in Table 1. In the following we use (when available) statistics collected in [15] that are the results of combining the information from two surveys: a magazine survey in Information Week (October 1996) that asked “What Security Problems have resulted in financial losses?”, and another magazine survey, in InfoSecurity News May 1997 that asked “In the past 12 months, which of the following breaches have you experienced?”. We need now to compute SLE and ALE for each of the pos-

| Attack | EF | ARO |
|---|-----|------|
| break down the door and go out unobserved | 90% | 0,10 |
| open the door with keys and go out unobserved | 93% | 0,10 |

Table 1: Exposure Factor (EF) and Annualized Rate of Occurrence (ARO) for the tree in Figure 5

sible attacks. Considering the first attack of Figure 5 we can notice that for a successful attack we need both to break down the door and to go out unobserved. So, the EF and ARO are associated to the pair of actions (and not to the leaf). Similarly for the second attack.

We need now to compute SLE and ALE. For the first attack, we have $SLE=AV \times EF=100.000 \text{ €} \times 0.9=90.000 \text{ €}$ and $ALE=SLE \times ARO=90.000 \text{ €} \times 0.1=9.000 \text{ €}$. In a similar manner we can compute $ALE=9.300 \text{ €}$ for the second attack.

| Countermeasure | RM | CSI |
|--|-----|----------|
| Install a security door | 70% | 1.500 € |
| Install a video surveillance equipment | 10% | 3.000 € |
| Employ a security guard | 50% | 12.000 € |
| Install a security lock | 20% | 300 € |

Table 2: Possible countermeasure for the attack tree in Figure 5.




Figure 6: The defense tree of Figure 5 decorated with ROIs.

As a last step, by considering the countermeasure cost (CSI) and the amount of risk mitigated (RM) of Table 2, we can associate to each countermeasure the corresponding ROI. For the first countermeasure (installing a security door to mitigate the threat of breaking down a door), we have $ROI = \frac{(ALE \times RM) - CSI}{CSI} = \frac{(9.000 \text{ €} \times 0.7) - 1.500 \text{ €}}{1.500 \text{ €}} = 3.20$.

The resulting defense tree showing ROI for each countermeasures is depicted in Figure 6.

From the defense tree of Figure 6 the security manager can already make some considerations. To mitigate all the attacks, at least one countermeasure for path has to be selected. For each path, the countermeasure with highest ROI is selected (in fact, the higher the ROI the better the investment). So, for the first attack of the example the best countermeasure seems to be the installation of a security door with ROI=3,20. Similarly, for the second attack the best countermeasure is the installation of a safety lock with ROI=5.20.

Notice however that sometimes a countermeasure can mitigate more than one attack (as is the case of the employment of a security guard in the defense tree of Figure 5). In this case a more detailed analysis has to be performed, and an overall ROI considering all the attacks and all the countermeasures of the tree has to be computed. Another consideration is about the ROI of a specific countermeasure. From the defense tree of Figure 6 we can see that the same countermeasure (for instance the employment of a security guard), can have a different ROI in different attacks (ROI=0.62 and ROI=0.61, respectively). This happens because the level of risk mitigation (RM) of a countermeasure could be strictly depending from the specific attack, and the ALE of the attacks could be completely different. We leave

the development of solutions to these problems as future work (see Section 5).

4.2 Computing ROA: the attacker's point of view

Given a defense tree in DNF, also the attacker's point of view can be considered by using ROA as a countermeasure label. We proceed as follows:

- we consider for each tree the expected gain deriving from a successful attack (GI);
- then, we estimated the attack cost to be sustained by an attacker to succeed when no countermeasure is present (Cost) and the added cost when the countermeasure is implemented (Loss);
- finally, with the above data the Return on Attack (ROA) is computed and used as a label for each countermeasure.

As an example, consider again the defense tree depicted in Figure 5. This time the tree is analysed from an attacker perspective. Let us suppose that the attacker has an advantage that can be economically quantified as 30.000 € for a successful attack to the server. By using the data of Table 3 we can compute the ROA for each countermeasure. The first four lines of Table 3 describe the countermeasures for the first attack, while the last four are related to the second attack. Notice that the cost an attacker has to

| Countermeasure for attack 1 | Cost | Loss |
|--|---------|---------|
| None | 4.000 € | 0 € |
| Install a security door | 4.000 € | 2.000 € |
| Install a video surveillance equipment | 4.000 € | 1.000 € |
| Employ a security guard | 4.000 € | 1.500 € |
| Countermeasure for attack 2 | Cost | Loss |
| None | 4.200 € | 0 € |
| Install a safety lock | 4.200 € | 200 € |
| Install a video surveillance equipment | 4.200 € | 1.000 € |
| Employ a security guard | 4.200 € | 1.500 € |

Table 3: Estimated cost and loss for two attacks.

pay depends on the attack and not on the countermeasure installed. In Table 3, for instance, the cost to be sustained by the attacker from stealing the server is different (4.000 € or 4.200 €): the loss instead depends on the specific countermeasure (2.000 € when encountering a security door vs 1.000 € for a video surveillance installation).

The data in the table are used to compute ROA for all the countermeasures in the tree. So, for instance when installing a security door we can obtain a $ROA = \frac{GI}{\text{cost before } S + \text{loss caused by } S} = \frac{30.000 \text{ €}}{4.000 \text{ €} + 2.000 \text{ €}} = 5,00$. In a similar manner we can compute ROA for all the other countermeasures as shown in Figure 7.

The defense tree of Figure 7 can be analyzed by the security manager in a similar manner as already described above




Figure 7: The defense tree of Figure 5 decorated with ROA.




Figure 8: The defense tree of Figure 5 decorated with ROI and ROA indexes.

for the ROI. This time the lower the ROA the lower the incentive for an attacker to try the specific attack. So, for the first attack of the example the best countermeasure seems to be the installation of a security door with a ROA=5.00. Similarly, for the second attack the best countermeasure is the employment of a security guard with ROA=5.26 (that could be used also as a countermeasure for the other attack!).

4.3 Putting together the evaluations

When the process of labeling the considered defense tree with ROI and ROA is complete, we can put together those indexes and perform a synthetic evaluation so as to determine the security investment that provides the best return on investment and that best discourages attacks.

The risk management process team should ideally select a countermeasure maximizing ROI and minimizing ROA. When such a countermeasure does not exist a countermeasure should be selected that either:

- maximizes ROI or minimizes ROA,
- is any Pareto-optimal countermeasure, or
- maximizes a user-defined function of ROI and ROA.

Figure 8 shows the defense tree labeled with both ROI and ROA. The double labeling gives the security manager the complete view of the scenario. The first step is the elimination of countermeasures dominated by some other and concentrate on the Pareto-optimal [17].

To do this we build the graphs of Figure 9 where for each attack, all countermeasures are compared. Figure 9.1 shows the countermeasures of the first attack. We can see




Figure 9: Comparing countermeasures of the defense tree of Figure 8.

how countermeasure $c1$ is better than countermeasures $c2$ and $c3$. In fact, $c1$ has both greater ROI and smaller ROA than $c2$ and $c3$. For this attack the security manager can easily choose the best solution. Figure 9.2 shows instead the countermeasures of the second attack. In this case only countermeasure $c5$ can be easily discarded.

At this step some countermeasures can be discarded (those with white background in Figure 8). For the first attack countermeasure $c1$ (install a security door) is selected. For the second attack, if the security manager wants to maximize ROI, then countermeasure $c4$ (install a safety lock) will be selected, whilst if he/she prefers to minimize ROA, the selected countermeasure will be $c6$ (the employment of a security guard).

5 Conclusions and Future Work

In this paper we presented our proposal for extending attack trees, a qualitative instrument used for modeling attack scenarios, with countermeasures and economic quantitative indexes. This extension allows us to evaluate effectiveness and profitability of countermeasures as well as their deterrent effect on attackers.

A related approach is proposed in [7] where the authors presents an economic model for determining the optimal security investment for protecting a system from a single threat. They consider three parameters: the monetary loss produced by an occurring breach (λ), the probability of a threat (t), and the probability that an attack would be successful (v) (that correspond respectively to our SLE , ARO and EF). The expected benefit of an IT investment is modeled as a function of the security investment ($EBIS(z)$). By assuming that "as the investment in security increases, the information is made more secure, but at a decreasing rate", the optimal amount of investment is determined by maximizing the relative difference between benefits and costs. Another economic-based framework is proposed in [13] where a game-theoretic approach is used for inferring the attacker's intents, objectives and strategies which are modeled using economic incentives and utilities.

The methodology presented in this paper provides a basis for future work along several research directions.

We are interested in extending our work so as to provide a solution to the problem of selecting a *set* of effective and profitable countermeasures which mitigate the risk deriving from all attacks in an attack tree. While it may seem obvious to compute the solution cost of a set $C = \{c_1, c_2\}$ of countermeasures as the sum $CSI_C = CSI_{c_1} + CSI_{c_2}$ of costs of the countermeasures in C , it should be noticed that the total cost of implementing a set of countermeasures could realistically be less than CSI_C (e.g. discounted price of bundled security solutions) or greater than CSI_C (e.g. when countermeasures must be managed by different employees, due to the existence of separation of duty constraints [3]).

On the other hand, it is not clear how to compute the value of the Risk Mitigated attribute for C , as any value between $\max(RM_{c_1}, RM_{c_2})$ (one countermeasure strictly entails the other) and $(RM_{c_1} + RM_{c_2})$ (completely independent countermeasures) appears to be acceptable depending on the type and nature of countermeasures and the asset being protected.

The problem of selecting a set of countermeasures becomes even more challenging under the realistic assumption that a single countermeasure can be used to mitigate risk associated with multiple vulnerabilities.

We also plan to investigate how to leverage existing results on constraint semirings [2] and their use in attack trees rewriting [14] for computing attribute values of and/or nodes as functions of attribute values of their children in the considered defense tree. Results borrowed from probability [12] and possibility theory [5, 22] can also be useful for estimating frequency and likelihood of attacks from frequency and likelihood of vulnerabilities used in the attack.

The annual rate of occurrence (ARO) of attacks can be difficult to estimate, because organizations are typically reluctant to make attack data publicly available due to the negative influence this may have on their reputation. Thus, another interesting direction of research may consist in exploring how Return On Attack (ROA) and other information about the attacker, like, for example, non-economic motivation, risk attitude and type of attackers (which can range from script-kiddies to organized crime and cyberterrorist), can influence the annual rate of occurrence of attacks, also from a game theoretical perspective.

Other interesting extensions to the work presented in this paper include considering how vulnerabilities can be used for attacking multiple assets of an organization, how to replace fixed attribute values with constraints (e.g. intervals), and how to use fuzzy logic techniques to define functions combining ROI and ROA indexes.

We hope our work can help encourage research and experimentation with use of economic indexes and combined development of attacker/defender perspectives during evaluation of alternative security investments.

References

- [1] D. Balzarotti, M. Monga, and S. Sicari. Assessing the risk of using vulnerable components. In *Proceedings of the First Workshop on QoP*, Milan, Italy, September 2005.
- [2] S. Bistarelli. *Semirings for Soft Constraint Solving and Programming*, volume 2969 of *LNCS*. Springer, 2004.
- [3] D. Clark and D. Wilson. A comparison of commercial and military computer security policies. In *IEEE Symposium on Computer Security and Privacy*, April 1987.
- [4] M. Cremonini and P. Martini. Evaluating information security investments from attackers perspective: the Return-On-Attack (ROA). In *Fourth Workshop on the Economics of Information Security*, June 2005.
- [5] D. Dubois and H. Prade. *Possibility theory: An Approach to the Computerized Processing of Uncertainty*. Plenum Press, 1988.
- [6] M. Gilbert. Disaster recovery planning: Conducting a risk analysis. white paper 11, Hill Associates, 2003.
- [7] L. A. Gordon and M. P. Loeb. The economics of information security investment. *ACM Trans. Inf. Syst. Secur.*, 5(4):438–457, 2002.
- [8] K. J. S. Hoo. How much is enough: a risk management approach to computer security. In *Workshop on Economics and Information Security*, 2002.
- [9] B. D. Jenkins. Security risk analysis and management. white paper, Norman Data Defense Systems, Inc., 1998.
- [10] M. Krause and H. F. Tipton. *Handbook of Information Security Management*. Auerbach Publications, 1999.
- [11] R. L. Krutz, R. D. Vines, and E. M. Stroz. *The CISSP Prep Guide: Mastering the Ten Domains of Computer Security*. Wiley, August 2001.
- [12] D. V. Lindley. *Making Decisions*. Wiley and Sons, 1985.
- [13] P. Liu and W. Zang. Incentive-based modeling and inference of attacker intent, objectives, and strategies. In *Proc. CCS'03*, pages 179–189. ACM Press, 2003.
- [14] S. Mauw and M. Oostdijk. Foundations of attack trees. In *Eighth Annual International Conference on Information Security and Cryptology*, LNCS. Springer, 2005.
- [15] J. W. Meritt. A method for quantitative risk analysis. In *Proceedings of the 22nd National Information Systems Security Conference*, October 1999.
- [16] A. Moore, R. Ellison, and R. Linger. Attack modeling for information security and survivability. Technical report, Software Engineering Institute CMU/SEI-2001-TN-001, 2001.
- [17] V. Pareto. *Manual of Political Economy*. Augustus M. Kelley, 1971. orig. (1960).
- [18] B. Schneier. Attack trees: Modeling security threats. *Dr. Dobbs's Journal*, December 1999.
- [19] B. Schneier. *Secrets & Lies: Digital Security in a Networked World*. John Wiley & Sons, 2000.
- [20] W. Sonnenreich, J. Albanese, and B. Stout. Return On Security Investment (ROSI): A practical quantitative model. In *Proc. 3rd Int. Workshop on Security in Information Systems, WOSIS 2005, In conjunction with ICEIS2005*, pages 239–252. INSTICC Press, 2005.
- [21] G. Stoneburner, A. Goguen, and A. Feringa. Risk management guide for information technology systems. Nist special publication 800–30, NIST, 2002.
- [22] L. A. Zadeh. Fuzzy sets as a basis for a theory of possibility. *Fuzzy Sets and Systems*, 1:3–28, 1978.