Improving Security Technology Selections with Decision Theory

Shawn A. Butler
Carnegie Mellon University
Computer Science Department
Pittsburgh PA. 15217
1-412-683-6555
shawnb@cs.cmu.edu

ABSTRACT

Software engineers, using decision theory to make design choices, can improve their selections if they rely on outside expertise to reduce the uncertainty in their objectives. However, integrating specialized expertise presents another set of problems. This paper presents two challenges to using decision theory in selecting security technologies.

1. Introduction

The primary challenge in using decision theory in design problems is deciding how to apply decision theory techniques. I have been using decision theory techniques to develop a selection method for the security technology selection problem [1]. The application of decision theory in the development of this method led to several research design considerations. In this paper I describe two decisions that a researcher may have to address when adopting decision theory techniques in making design choices.

The first consideration was whether the goal of the method was to replicate the security manager's security selections or improve them. A second consideration was deciding the level and detail of information necessary to make reasonable selections. Both considerations impact the structure of the selection method and the source of information used to make security technology decisions. Although this paper addresses decision theory and security technology decisions, these considerations will likely be important when applying decision theory to other types of design decision problems.

2. Decision Theory

Decision theory is attractive in security technology selection problems because it provides a methodology to deal with the uncertainty and multi-objective nature of these decisions. In decision theory, risky decisions are those whose consequences are uncertain. Risky decisions can also have multiple objectives. Each objective has an attribute that is the degree to which a given decision objective has been attained [2]. Probability distributions can be associated with each attribute to reflect the expectations or uncertainties of decision makers. The power of decision theory is that it provides a systematic way to consider tradeoffs among attributes, which can be used to make decisions.

When an additive value model is valid[1], the value of each alternative is computed and ranked [3]. The value of the alternative is based on the objective attributes. Sensitivity analysis can be conducted to see how sensitive the rankings are to model assumptions. Models that rely heavily on decision maker expectations could produce better recommendations if the decision is sensitive to the expectations and the uncertainties (or variance) associated with the expectations could be reduced.

[1] The criteria necessary for the additive value model are beyond the scope of this paper.

3. Security Technology Selection Problem

The security technologies selection problem is the task of selecting the best set of security countermeasures for an information system. Inherently, the challenge is to quantify the benefits of security countermeasures and the consequences or outcomes of successful attacks. The benefit of a security countermeasure depends on how well it stops an attack, or mitigates the consequences of a successful attack.

There are three key elements of uncertainty in security technology selection problem. The first element of uncertainty is the attack. Most security engineers have little data concerning the frequency of attacks, however, the appropriate selection of countermeasures depends on which attacks will likely occur. The second element of uncertainty is the outcome of a successful attack. For example, once an attacker has access to the system, there are many potential paths. An attacker could do nothing or cause considerable damage depending on his motivation. Finally, the third element of uncertainty is the benefit from countermeasures. The effectiveness of a countermeasure in protecting or detecting an attack can only be estimated.

A security engineer must also balance multiple objectives when selecting security technologies. Consequences of successful attacks must be balanced with performance constraints, budget limitations and other design considerations. Each attack may result in a similar outcome, but with different attributes. For example, a denial of service attack and a virus could result in different levels of lost productivity. The fact that each attack can have different levels of outcomes will add complexity to the application of decision theory techniques.

4. Security Technology Selection Method

The security technology selection method takes advantage of the power of decision theory techniques. The possible outcomes from successful attacks can be thought of as objectives. For example, three possible objectives are 1) loss of life, 2) loss of revenue, and 3) loss of productivity. The attributes of these objectives will be the level of lives, revenue, and productivity lost. Usually, attributes are determined based on the security engineer's expectations about the attacks, the consequences of the attacks and the effectiveness of countermeasures. Reasonable alternatives can be generated from probability distributions that reflect these expectations. Figure 1 shows a typical expected revenue loss distribution of a virus attack.

5. Improved Decisions

One of the most important considerations in using decision theory is to decide whether decision theory is being used to replicate a decision or improve a decision. Of course, sometimes just the application of decision theory techniques in the design decision process may result in an improved decision, but it could just as easily result in a poor decision. When decision theory is used to replicate a decision then there may be an underlying assumption that there is sufficient expertise among the participants. If decision theory is going to improve the decision process, then outside expertise can be integrated into the decision process.

The security technology selection method relies on countermeasure expertise to improve the selection of countermeasures. Countermeasure expertise is used to more accurately represent the mitigation impact of a countermeasure given an attack in the security technology selection problem. The system security engineer provides the risk analysis consisting of the likelihood of attacks and their potential outcomes. The risk analysis phase of the selection method develops an outcome distribution for each relevant attack. If a security engineer uses anti-virus detection software, then either the frequency of successful virus attacks is reduced or the amount of revenue lost is reduced. Figures 2 and 3 show how the outcome distribution curves might shift when an anti-virus technology is used.

Countermeasure expertise may be able to more accurately determine the magnitude of the changes in the outcome distribution curves than security engineers. A security engineer may not have the experience or knowledge to accurately evaluate countermeasure effectiveness; therefore

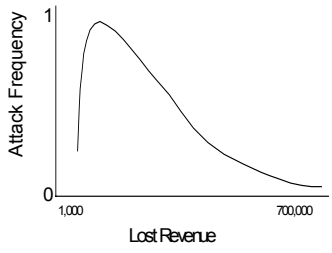

Figure 1

expert advice could result in a better countermeasure selection. Although the countermeasure specialist can provide relative shifts in the distribution curves, other factors such as a security administrator's skill level can influence the final distribution.

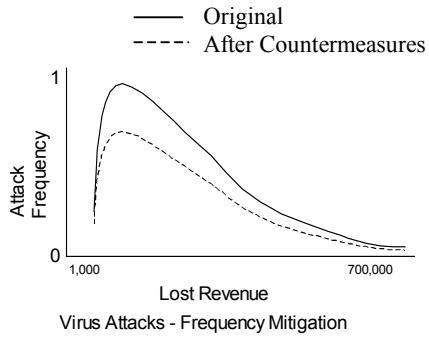

Figure 2.

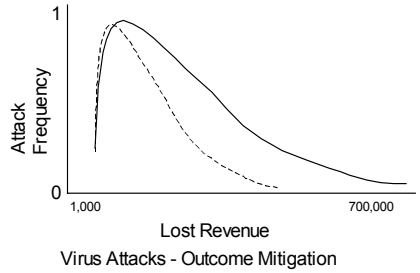

Figure 3.

Currently, my security technology selection method integrates only countermeasure expertise. Since so few organizations report incidents there is little reliable statistical data about attack frequencies. If threat data becomes available it would be possible to use statistically analyzed attack frequencies, adjusted for unique system characteristics, instead of the security engineer's opinion. If expertise is available then decision theory techniques can be used to improve design decisions.

6. Too Much Information

Another important consideration in the application of decision theory to design problems is the level and detail of the information used in the decision process. Two issues arose in the development of the security technology selection method. The first was the mismatch in attack information between the security engineers and countermeasure experts. The second issue was developing a process for combining or rolling-up the information to the final selection. Again, these problems are not unique to the security technology selection method.

The primary mismatch between security engineers and countermeasure experts is that they operate with different levels of information. The process of eliciting the risk analysis from security engineers requires that they validate a list of attacks. Security engineers are asked to add to the list of attacks provided or further define attacks that may be too general. Four risk analyses resulted in the elimination of two attacks and the addition of one attack. The list of attacks appeared to represent most security engineer's concerns.

The process of eliciting countermeasure expertise requires that the expert select from an extensive list of countermeasures those that are most appropriate for each attack. In my research, the initial difficulty was that the experts required a refinement of the attacks before they could suggest appropriate countermeasures. For example, countermeasure experts refine the Denial of Service attack into those that target mail servers and those that attack specific applications, such as mail servers. This required additional refinement from the security engineer, which is not always possible. The security engineer may not know enough about the difference among variations of general attacks to be able to estimate attack frequencies.

The fact that countermeasure experts distinguish among attacks with finer detail than security engineers argues for integrating outside expertise, if possible. In the security technologies selection method, differences are resolved on a case-by-case basis. The application of decision theory techniques to other types of design decisions may also result in a mismatch between expert information and engineering information.

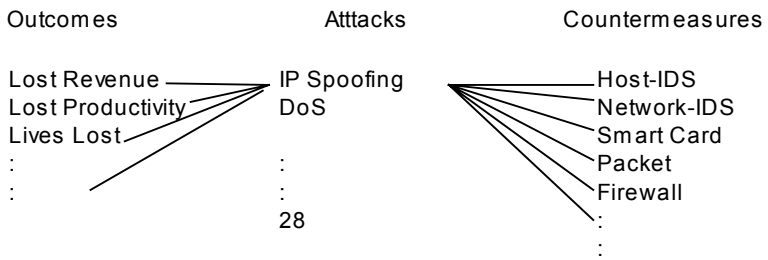

Figure 4.

Initially, the security technology selection problem appears to have an overwhelming amount of data making it difficult to re-assemble the data into meaningful recommendations. In the current security technologies selection method, there are 28 different types of attacks, over 40 countermeasures, and at least half a dozen possible outcomes. Figure 4. shows how many different attributes could be associated with each class of outcome. A distribution curve must be established for each outcome. For each attack, the security engineer is asked to provide three possible values, low, high, and expected, for each outcome. An important step in the development of the security technology selection method was to determine a way to reduce the possible combinations.

In order to reduce the number of combinations that are actually analyzed, the security engineer provides information about the top 3 or 4 outcomes and the countermeasure expert focuses on the countermeasures that provide a moderate level of protection or mitigation. The most efficient way to reduce the possible combinations is to focus on the most important. This technique eliminates aspects of the decision process that don't significantly contribute to the final selection.

Once the data is collected from a case study and the countermeasure expert has provided mitigation information, the challenge is to combine the information so that countermeasures can be recommended. Each case study will result in approximately 28 different outcome distributions, each distribution adjusted for several countermeasures. Although not completed, I expect that the data will allow countermeasures to be weighted based on their overall contribution to mitigating outcomes.

7. Conclusion

I have described a few issues that must be considered when applying decision theory techniques to design decision problems. There are other research issues that make application of these techniques challenging, such as developing techniques to turn expertise into quantifiable information. Although the security technology selection method takes one approach in using decision theory, there are many alternatives that may result in even better design processes. I am encouraged by my initial results that show decision theory can be a significant tool in software engineering practice.

ACKNOWLEDGEMENTS

I would like to acknowledge the Software Engineering Institute for their support and countermeasure technology expertise. I would also like to thank the Department of the Army for their participation as my case study.

REFERENCES

[1] Butler S., Chalasani, P., Jha S., Shaw M. *The Potential of Portfolio Analysis in Guiding Software Decisions*, 14 June 2000, Position Paper for First Workshop on Economics-Driven Software Engineering Research (EDSER-1), affiliated with the 21st International Conference on Software Engineering (ICSE'99), May 1999.

[2] Fischer, G. and Fischbeck, P. Multiattribute Preference Models: A Brief Overview. Working Draft, undated.

[3] Keeney, R. and Raiffa, H. Decisions with Multiple Objectives: Preferences and Value Tradeoffs. *Cambridge University Press,* 1999