SAND2009-4139C

Quantifying the Degree of Balance in Physical Protection Systems

Gregory D. Wyss
Sandia National Laboratories, New Mexico
P.O. Box 5800, MS-0757, Albuquerque, NM, 87185-0757

Abstract

"Balanced protection" is recognized as an important heuristic of an effective physical protection system. A physical protection system is viewed as balanced when every pathway by which an adversary might attempt to accomplish a goal presents a similar level of difficulty (e.g., similar resources required, similar probability of timely detection, similar likelihood of being neutralized, etc.). Historically, the degree of balance within the protection system has been assessed subjectively when the security analyst compares relevant security metrics for several of the most vulnerable attack pathways using expert judgment. This paper proposes an objective method to obtain a quantitative metric to represent the degree of balance within a security system. The method is based on the mathematics of exponentially weighted moving averages, wherein the performance of the security system for many of the most vulnerable attack pathways are aggregated to provide a single metric. By using an exponentially weighted average instead of a simple mean or a linear weighted average, the method automatically applies higher weighting to the most vulnerable attack pathways and minimal weighting to the less desirable pathways. The method is relatively insensitive to the number of attack pathways considered in the average as long as each attack pathway is developed to a comparable level of granularity. This characteristic is very different from a simple mean or a linear weighted average, which can be very sensitive to the number of terms selected for use in the averaging process. This method can be used to help inform risk management decisions within a single site and across multiple sites.

Introduction

The objective of a physical protection system is to protect a facility or system against potential attacks by malevolent adversaries. For high-consequence facilities, the key functions of the physical protection system are described by the philosophy "detect, delay, respond." The protection system must detect that the system or facility is under attack and then delay the adversary from achieving his or her attack objectives for long enough that an effective response can be initiated to interrupt and ultimately neutralize the attack. The ability of the physical protection system to accomplish these objectives is measured by two probabilities:

1. the probability that the attack is detected in time for an effective response to be mustered to interrupt the adversary's attack tasks (P_I, the probability of interruption, which is sometimes also called the probability of timely detection), and
2. the probability that the response will be able to successfully neutralize that adversary attack (P_N, the probability of neutralization).

To be successful, the protection system must succeed in both of these functions. Thus, the probability that the protection system is effective in defeating an adversary attack, P_E , is the probability that the attack is both interrupted and neutralized, so $P_E = P_I \cdot P_N$, and P_E is an estimate of the likelihood that an adversary will succeed given that they attempt a particular attack scenario. P_E is estimated for each potential attack scenario to which a facility or system is susceptible, and the scenarios with the lowest value of P_E are the most advantageous to the adversary. The adversary is assumed to be intelligent and knowledgeable about the protection system, so he or she will not likely attempt an attack scenario that is significantly suboptimal (i.e., the adversary will likely select from among the few scenarios with the lowest value P_E). For this reason, it is reasonable and convenient to rate the performance of a physical security system according to the lowest value of P_E for any attack path at the facility or system.

Another desirable characteristic for a physical protection system is that it provide *balanced protection*.[1] A physical protection system is viewed as balanced when every pathway by which an adversary might attempt to accomplish a goal presents a similar level of difficulty (e.g., similar resources required, similar probability of timely detection, similar likelihood of being neutralized, etc.). This has been described colloquially as "a facility should not have tightly locked doors and wide-open windows." Mathematically, this means that the value of P_E for the adversary's one or two most advantageous attack scenarios should not be dramatically lower than it is for the next several most attractive scenarios. This definition is simple in concept and can be applied as a heuristic when evaluating physical protection systems when an analyst visually scans the P_E values for the adversary's several most attractive scenarios. However, the definition of balance remains heuristic unless it can be defined according to a mathematical formula.

The purposes of this paper are (a) to describe the desirable characteristics that such a formula should exhibit and (b) to evaluate several such potential formulas against these characteristics.

Desirable Characteristics of a Balance Metric

The objective of a balance metric is to measure the degree to which the attractiveness of an adversary's "good" attack options are comparable to his best attack option. Since P_E is the most commonly used metric for scenario attractiveness, a balance metric should be computed on the basis of P_E, but should ideally support the use of other scenario attractiveness metrics if they should become widely used. In addition, it should consider the scenarios that would be attractive but discount the scenarios that would be unattractive to an adversary.

Another important characteristic of a balance metric is that it be relatively independent of the level of detail employed in the physical protection system analysis. To appreciate the need for this characteristic, one must understand a key fact about attack scenarios: if one thinks about the system long enough, one can conceive of an almost infinite number of possible attack scenarios, most of which would not be attractive to the adversary when compared to the most attractive scenario. These unattractive scenarios should not contribute to the balance metric. That is, the balance metric should be relatively insensitive to the number of scenarios identified during the security analysis.

The Mathematics of Averages

The mathematics of averages can be used to develop a metric for the balance of a protection system. One can compute some sort of average P_E value for the scenarios that are attractive, but not most attractive, and compare that average to the P_E value for the most attractive scenario. If the two values differ little, then the physical protection system can be viewed as balanced, while a large difference between the values indicates an unbalanced protection system. This section examines several types of averages that may be appropriate for this purpose, while the next section considers the relative advantages and disadvantages of each as a possible balance metric.

This paper focuses specifically on the use of moving averages. A moving average (sometimes called a rolling average or a running average) is typically used to examine the characteristics of a system which exhibits noisy behavior but is evolving over time. The moving average smooths out short-term fluctuations in the data and highlights longer-term trends or cycles. Examples include financial data such as stock prices, trading volumes, employment, and gross domestic product, as well as environmental data such as long-term climate trends (e.g., distilling long-term trends from rapidly fluctuating daily and even annual temperature data). Thus, a moving average focuses attention on data for a particular subset of the overall data. For the purposes of measuring balance within a physical protection system, the subset of data upon which one wishes to focus is the set of scenarios that are attractive, but not most attractive, to an adversary. Thus, before applying a moving average to the data from a protection system, the analyst will sort the scenarios in order of increasing P_E so that those scenarios with the highest P_E values will receive the highest consideration in the moving average and, hence, exert the greatest influence on the balance metric.

Simple Moving Average

A simple average represents the mean of a data set, and is computed as the sum of all points in the data set divided by the total number of entries in the data set. The sum is unweighted; that is, all points in the data set have exactly the same contribution to the calculation, and no point is more important than any other in its contribution to the sum.

A simple moving average is similar to a simple average except that it considers only the previous *n* data points.[2] Like the simple average, this is an unweighted mean. When an *n*-day simple moving average is applied to a time series of data such as the daily closing price for a commodity, a simple moving average is simply the unweighted mean of the closing prices from the previous *n* days. When the next data point is added to the set (e.g., the closing commodity price for the next day), this new value comes into the sum, and an old value drops out as it is beyond the *n*-day window for the moving average.

The analyst must select a parameter in order to calculate a simple moving average: *n*, which is the number of data points to be included in the average or the length of the time window over which the moving average is calculated. The analyst selects this parameter value based on convention or expert judgment in order to emphasize the type of trend that is being investigated. For example, a 10-day window emphasizes short-term movement, while a 200-day window emphasizes long-term trends.

Linear Weighted Moving Average

A weighted average uses multiplying factors to give different weights to different data points within the data set. This can be done for several reasons. For example, the analyst may assign greater importance to more recent data because it is believed to be a better predictor of future trends than older data. The analyst may also be unsure of how many data points should be used in a trend analysis. In a simple moving average, the most recent or most important data point has exactly the same contribution to the metric as the most distant or least important data point. The application of a weight to each point in the moving average calculation can be used to reflect the analyst's belief that older or more distant data should be given lesser consideration.

The simplest weighted moving average applies a weighting factor that decreases linearly as one moves from the most important value toward less important values in the data set.[3] For data in a time series, the weights decrease toward zero linearly with each previous day in the time series. In an N-day weighted moving average, the most recent day has weight N, the second most recent $N–1$, and so forth until the N^{th} day in the average has a weight of 1, and all successive days have a weight of 0. The data value for each day is multiplied by its weight, and the resulting weighted values are summed. The denominator of the weighted moving average is simply the sum of the weighting factors, which is the triangle number $\frac{N\cdot(N+1)}{2}$.

The value calculated for a linear weighted moving average is much less sensitive to the exact value of the parameter N (the window for the moving average) than is the simple moving average. It can be shown that a single additional value that is added to the linear weighted moving average (i.e., adding one day to the window, or increasing N by 1) will only contribute $\frac{2}{(N+1)(N+2)}$ as a fraction of the newly computed weighted average value. Still, a larger change in N causes a very significant change in the relative importance of specific values in the weighted average value. Consider a situation where N is doubled. In the original weighted average calculation, the N^{th} value was only $^{1}/_{N}$ as important to the computed weighted average value as the first value. When N is doubled, the N^{th} value becomes about ½ as important as the first value. This sensitivity may be important when one is highly uncertain as to the correct value for N.

Exponentially Weighted Moving Average

Whereas the linear weighted moving average applies weighting factors that decrease linearly toward zero as one moves toward less important or more distant data, the exponentially weighted moving average, as its name implies, uses weighting factors that decrease exponentially in the form of a power series. This influences the characteristics of the moving average in two ways. First, since the exponential weighting function decreases more rapidly than a comparable linear weighting function, the exponentially weighted moving average gives even more weight to the more recent or more important data than does its linear counterpart. Second, since the exponential weighting function never actually reaches zero no matter how many terms are applied, every member of the data set contributes at least something to the exponentially weighted moving average.

The characteristics of the exponentially weighted moving average are set through the selection of the weighting factor α, a value between 0 and 1, which determines the rate of decay for the exponential weighting function. If one begins with the most important or most recent point in the data set, then the value of the weighting factor for the k^{th} term in the weighted average is $\alpha (1 - \alpha)^{(k-1)}$. Thus, the first point contributes α as a fraction of the total exponentially weighted average, and the contribution of each successive term to the total is reduced by a factor of (1– α).

The value for α is sometimes expressed in terms of a period N such that $\alpha = \frac{2}{N+1}$. When expressed in this way, the first N terms in the weighted average represent approximately 86% of the total weight in the calculation, and the first $2N$ terms represent more than 98% of the total weight. The half-life of the weights (the interval over which the weights decrease by a factor of two) is approximately $\frac{N}{2.885}$ if $N > 5$. The value of α goes down by almost a factor of two when N is doubled.

In the exponentially weighted moving average calculation, the relative contribution of each term to the overall weighted sum is far more consistent than was observed for the linear weighted average. Recall that when the width of the window was doubled, the importance of the N^{th} value was increased from $^{1}/_{N}$ to about ½ as important as the first value. In other words, its importance to the average increased by a factor of about $^{N}/_{2}$, which is important if N is greater than about 5, and becomes enormous as N increases. By contrast, a comparison of the importance of the N^{th} value to that of the first value for the exponentially weighted average shows that the importance of the N^{th} value is increased by a factor that never exceeds 2.8, regardless of the value of N (or α). Even the importance of the $2N^{th}$ value (at which the sum of the weighting factors exceeds 98% of the total weight) never increases by a factor of more than 7.4 as N is doubled, regardless of the value of N (or α). Thus, there is far greater consistency in the relative weighting among terms in the exponentially weighted moving average across different values of N than exists in the linear weighted moving average.

Application of Averages as Balance Metrics

A physical protection system is viewed as balanced when every pathway by which an adversary might attempt to accomplish a goal presents a similar level of difficulty (e.g., similar resources required, similar probability of timely detection, similar likelihood of being neutralized, etc.). This definition can be applied as a heuristic when an analyst visually scans the P_E values for the adversary's most attractive scenarios. By definition, an assessment of balanced protection requires the analyst to consider multiple scenarios. One way to measure this is to count the number of scenarios that have P_E values that are comparable the most attractive scenario. The problem with this approach is that it is difficult to set a cutoff value that adequately expresses what it means to be "comparable." Even if such a criterion could be developed, the accuracy of such a scenario count could be inaccurate if several scenarios had P_E values that were just beyond the cutoff.

Using an average as a measure of balance is desirable because it considers multiple scenarios. The averaging process minimizes the impact of the discrete cutoff values described above because it allows the actual data to speak for itself without regard for potentially artificial limits. The Department of Energy (DOE) makes use of this property of averages in its Graded Security Protection (GSP) Policy.[4] Each site is rated in regard for compliance with security requirements on the basis of the average P_E values over six significantly different and attractive scenarios.* This approach is essentially a simple moving average with $N = 6$. Thus, by comparing the lowest P_E value with the averaged P_E value calculated to satisfy the requirements of the GSP, one could estimate the balance of the protection system at the site. Note, however, that the value of this balance estimate is highly sensitive to the selection of N, the size of the group of scenarios included in the calculation. The value of the average is equally sensitive to the value of P_E for the 2nd most attractive scenario, the 6th most attractive scenario, and the 20th most attractive scenario (if N were at least that large).

This paper has shown how the average P_E calculation could be made significantly less sensitive to the value selected for N through the use of either a linear weighted or exponentially weighted average. This would, however, change the way that the average behaves, making it much more sensitive to the value of P_E for the most attractive scenario. Decision makers would need to evaluate whether this change would enhance or detract from their policy objectives. They may find it desirable to develop a separate balance metric that makes use of a weighted average and a larger value for N in order to quantify the degree to which the physical protection systems for which they are responsible exhibit the desirable characteristic of "balanced protection."

While averaging P_E provides a reasonable balance metric, a more robust metric can be generated by averaging the adversary's probability of success during an attack P_S, which is $(1–P_E)$. As one moves through the list of scenarios from most attractive to least attractive, the value of P_E increases towards 1.0 while the weighting factor used in both of the weighted averages decreases towards zero. By using P_S instead of P_E, one eliminates the competition between increasing metrics and decreasing weights, and introduces a compounding effect wherein less attractive scenarios have both lower metrics (P_S approaches zero) and lower weights. This makes the balance metric even less sensitive to the value selected for N, resulting in more robust performance of the balance metric.

Summary and Conclusions

An important principle of physical protection system design is that it should provide balanced protection. A physical protection system is viewed as balanced when every pathway by which an adversary might attempt to accomplish a goal presents a similar level of difficulty. Historically, the degree of balance within the protection system has been assessed subjectively as the security analyst compares P_E for several of the most vulnerable attack pathways using expert judgment. This paper describes how weighted averages can be used to effectively and objectively measure the degree of balance within a physical protection system. In so doing, the performance of the physical security system for many of the most vulnerable attack pathways are aggregated to provide a single balance metric. Linear and exponentially weighted averages automatically apply higher weighting to the

---

* It has been noted that averaging of P_E values across scenarios can also obscure information that is important to the overall security analysis. The smoothing effect of the averaging process can obscure the existence of a single high-vulnerability scenario when its P_E value is averaged with those of other less vulnerable scenarios. If one assumes that adversaries will actively seek out and exploit the best scenarios that they can identify, then the existence of a high-vulnerability scenario must be taken seriously and addressed. The DOE GSP addresses this issue administratively. We should note that the situation just described reflects a physical protection system that exhibits a serious lack of balance.

most vulnerable attack pathways and minimal weighting to the less desirable pathways. In particular, the exponentially weighted average is the least sensitive to the number of attack pathways considered.

By comparing P_E for each site's most attractive scenario with its averaged P_E as calculated according to the DOE GSP, a decision maker can gain insights into whether the site provides balanced protection, although this metric is very sensitive to the number of scenarios considered in the averaging calculation. It can be made less sensitive through the use of one of the weighted averaging methods described in this paper. A robust balance metric can be a valuable tool to inform risk management decisions both within a single site and across multiple sites.

Acknowledgments

This work was performed at Sandia National Laboratories. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. This paper has been approved for unlimited release as SAND2009-XXXXC.

References

[1] Garcia, M.L., *The Design and Evaluation of Physical Protection Systems*, Butterworth-Heinemann (Elsevier), Burlington, MA, 2001.

[2] Iman, R.L. and Connover, W.J., *Modern Business Statistics, 2nd Edition*, John Wiley & Sons, New York, 1989.

[3] Wikipedia, "Moving Average," http://en.wikipedia.org/wiki/Moving_average , viewed November 18, 2008.

[4] US Department of Energy, "Graded Security Protection Policy (U)," DOE O 470.3B, US Department of Energy, Washington, DC, August 12, 2008.