SAND2016-2179C

The Impact of False and Nuisance Alarms on the Design Optimization of Physical Security Systems

Alisa Bandlow[1], Katherine A. Jones[1], Nathanael J.K. Brown[1], Linda K. Nozick[2]

[1] Sandia National Laboratories, P.O. Box 5800
Albuquerque, New Mexico 87185-1188 USA
{abandlo, kajones, njbrown}@sandia.gov
[2] Cornell University, 220 Hollister Hall
Ithaca, New York 14853 USA
lkn3@cornell.edu

Abstract. Despite the known degrading impact of high nuisance and false alarm rates (NAR/FAR) on operator performance, analyses of security systems often ignores operator performance. We developed a model to analyze the impact of nuisance alarm rates on operator performance and on overall system performance. The model demonstrates that current methods that do not account for operator performance produce optimistic estimates of system performance. As shown in our model, even low NAR/FAR levels and the associated alarm queueing effect can increase operator detect and response time, which in turn reduces the amount of time the response force has to interrupt the intruder. An illustrative analysis shows that alarm processing times can be higher than the assessment time due to queue wait times and that systems with only one or two operators can become overwhelmed as NAR increases, decreasing system performance.

Keywords: Human-systems Integration · Physical Protection Systems · Nuisance Alarms · Operator Performance

1 Introduction

Identifying an optimal design for a physical security system is critical to mission performance. Facility or system owners are sometimes willing to invest millions of dollars to increase intruder delay times by a few seconds. In most systems, a human operator assesses the alarmed sensors and then calls on the response force to investigate. With limited budgets, sites are searching for technologies that can reduce the number of staff employed, thereby reducing the overall cost of the security system. However, more technologies and sensors will usually increase the nuisance and false alarm rates (NAR/FAR), which in turn may require additional operators to respond to the increase in alarms.

Standard physical protection system (PPS) assessment methods include red team exercises [1], adversary sequence diagrams, design basis threat and fault tree analysis

[2]. Analyses of physical security systems primarily focus on the intruder delay times due to physical barriers, the reliability of sensors, alarm assessment, and the response times of the protective force [1, 2, 3, 4, 5]. These methods typically assume that the operator is ready to begin assessment as soon as the intruder alarm is generated and that the operator behavior never deviates from policy and training, such as ignoring alarms.

As the rate of nuisance and false alarms increases, a system's perceived reliability decreases, causing the operator to lose trust in the system. Moray, Inagaki, & Itoh [6] found that trust was impacted the most when system reliability fell below 90%. This loss of trust can result in delayed response to alarms (the "cry wolf" effect) [7, 8]; probability matching response rates [9, 10]; and, in extreme cases, failure to respond, ignoring or disabling alarms [11, 12].

Despite the known degrading impact of high of NAR/FAR on those monitoring the alarms, analyses of security systems often ignore operator performance. Thus, an operator who is slow to respond to an alarm or who simply ignores or disables the alarm can weaken security systems that are considered highly reliable. Without including a more realistic human response and assessment time, current system performance estimates may be overly optimistic.

We previously developed a model to optimize the design of a PPS [13]. Building on that model, we developed a new model to analyze the impact of nuisance alarm rates on operator performance and overall system performance. The NAR/FAR level and the associated alarm queueing effect for a proposed system design impact the speed at which the operators will respond to alarms, which in turn affects the amount of time the response force has to interrupt the intruder.

2 Physical Protection System (PPS)

The goal of a physical protection system (PPS) is to use detection, delay and response to prevent an adversary from reaching a target [2]. **Detection** is the discovery of an adversary when a sensor detects an abnormal event. A person assesses the alarm to determine if it is valid (an adversary is detected) or invalid (a nuisance or false alarm). **Delay** is the use of obstacles to increase the adversary task time. Obstacles can be passive barriers (e.g., locks, fences, and Jersey walls) or active barriers (e.g., engagement by the response force, pop-up vehicle barriers). **Response** is the actions taken by the response force to prevent adversary success.

Figure 1 shows the adversary task timeline and the relationship to the three PPS functions. The total time for the adversary to accomplish their goal is labeled Adversary Task Time. The Adversary Task Time is impacted by the delay provided by the PPS. Any delay provided by the PPS before detection (the dotted line labeled Adversary Undetected) does not count towards system effectiveness. After the first alarm at time T_0, the alarm information is assessed to determine if it is valid or a false or nuisance alarm. If the alarm is assessed to be valid at time T_A, the alarm information is communicated to the response force. Additional time is required for the response force to deploy and respond to the adversary. The time at which the response force interrupts the adversary is T_I. If the adversary is not interrupted, they will complete their task at time T_C.

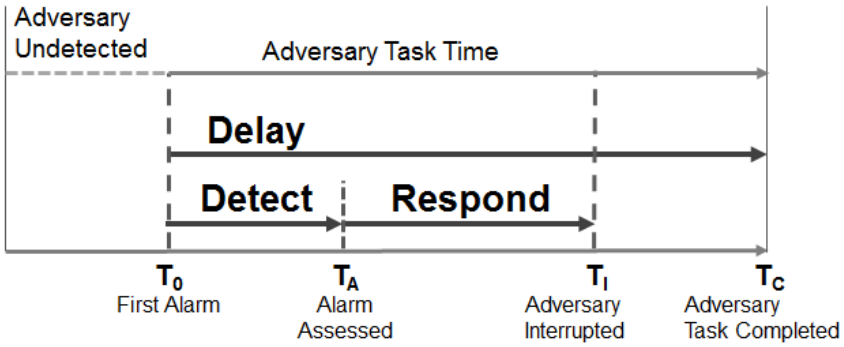

Fig. 1. Relationship between PPS functions (delay, detect, respond) and the adversary task timeline.

3 PPS Design Optimization Model

We previously developed a model to optimize the design of PPS [13]. We use probability of detection, probability of interruption, delay time, and response time as defined by [2] as elements in our model. We represent the problem as an attacker-defender model, in which the attacker (adversary) has perfect knowledge of the security measures in place. The attacker's goal is to reach a specific target which is protected by a physical security system. In this model, the defender is the designer and operator of the security system. The defender's goals are to minimize investment cost, minimize the nuisance alarm rate and false alarm rate (NAR/FAR), and maximize the probability of interrupting the adversary. The probability of interruption (P_I) is the probability that the travel time of the security response force will be less than the travel time remaining for the attacker once they have been detected, allowing interception before the target has been reached. We use the P_I given detection as the relevant measure of interest for the quality of the path from the perspective of the attacker.

The model is not a simulation of the attacker attempting to reach the target. Instead, the model performs an implicit enumeration of all attacker paths [4, 5]. For each PPS design solution, an algorithm explores the paths the attacker can take to calculate the worst (lowest) P_I [3]. This P_I is then assigned to the solution. Increasing the P_I is accomplished by adding detection and delay security measures (such as cameras and fences). Each technology investment has an associated cost and NAR/FAR.

The goal of the optimization is to suggest which technologies to place at which locations. The model only places barriers and sensors outside and on the exterior of buildings. The final output is a Pareto frontier that identifies a collection of efficient solutions. This allows a decision maker to identify an acceptable trade-off between the probability that the intruder is interrupted, investment costs, and NAR/FAR.

4 Operator Performance Model

We developed an operator performance model to begin including a more detailed representation of operator behavior with a goal of analyzing the impact of alarm rates on operator performance and on overall system performance. We chose to focus on two impacts of alarm rate on performance. First, dependent on the alarm assessment time, there is a maximum number of alarms a single operator can assess in a day. Second, the alarm rate can affect the operator's trust in the alarm system's reliability.

4.1 Alarms as a Queue

Standard PPS assessment methods typically assume that the alarm station operator is ready to begin assessment as soon as the intruder alarm is generated. However, the operator may be busy assessing other alarms, thus delaying detection. There is also a maximum number of alarms that can be assessed in a day. When the alarm rate is high enough, the operator(s) may not be able to assess all of the alarms in the queue, which can lead to missed detections.

To model the arrival and assessment time of alarms more realistically, we use a queueing model. Our approach is to shift the mean of the response time distribution by adding the steady state results of a queuing model which assumes there are k independent servers and that all alarms generated are examined in the order received. We assume that the alarms are independent of one another and arrive via a Poisson process (M). We also assume that the service time for an alarm is arbitrarily distributed (G) and that there are k operators examining the alarms. In queueing notation, this implies an $M/G/k$ queue.

Gans [14] gives a classical result that a reasonable approximation for the average waiting time per alarm in the queue is as follows.

$$E\left[W^{M/G/k}\right] = \frac{C^2+1}{2} E\left[W^{M/M/k}\right] \quad (1)$$

where C is the coefficient of variation of the service time distribution, and the $E\left[W^{M/M/k}\right]$ is as follows.

$$\frac{(1-P_0)}{k\mu-\lambda}. \quad (2)$$

where $\rho = {}^{\lambda}/_{k\mu}$, λ is the arrival rate of the Poisson process, k is the number of operators, and μ is the mean of the service rate distribution.

P_0 is the probability that there are zero alarms in the system and is computed as follows.

$$P_0 = \frac{1}{\left[\sum_{i=0}^{k-1} \frac{(k\rho)^i}{i!}\right] + \frac{(k\rho)^k}{k!\,(1-\rho)}}. \quad (3)$$

Generally, this approximation works well for operators in the tens, which we assume to be a plausible range for this application.

Strengths of this approach are that it is tractable and allows us to represent investment in the operators that examine alarms. The weakness is that there may be priorities among the alarms, which this approach ignores.

4.2 Operator Trust in the Alarm System

In the standard PPS assessment methods, the operator performance can be included in the probability of detection (P_D) as follows.

$$P_D = P_S \text{ x } P_A. \quad (4)$$

where P_S is the probability that a sensor detects the abnormal event, and P_A is the probability that the cause of the alarm is accurately assessed by the operator. P_A does not include the operator's response rate, so these methods do not account for operator behavior that deviates from policy and training, such as ignoring alarms.

To better represent operator behavior, we allow the system NAR/FAR to effect operator response time. First, since the original PPS design model focuses on the optimization of technology investments, operators are added as an investment item with an annual cost per operator. More operators can deal with higher alarm rates, but this will increase investment costs. Second, system NAR/FAR rate affects operator response times. Operator distrust in the system is quantified as a delay in response time. Since the original model uses P_I as a measure of quality for solutions, we include delay in response time in our calculation of P_I in order to better anticipate the true performance of a design in practice.

The operator performance model requires the following inputs, to be provided by the site physical security expert. Each sensor type has an alarm rate (which includes correct detections, NAR and FAR) per day. The sensor alarm rates are aggregated to obtain the system alarm rate per day. The site will have a maximum number of operators that they are willing to hire *k*. The operator has an average assessment time for a single alarm *AT*. Many alarm station operators perform a primary task in addition to monitoring alarms [2, 15], so there will be an average lag time, *LT*, when the operator has to switch tasks to respond to an alarm.

As the sensors' P_S increases, the NAR also increases, and the operator's trust in the system decreases. This causes an increase in response times [10, 16, 17, 18], which we call the trust delay time *TDT*. We bin the alarm rates into categories of low, medium and high, based on acceptability levels in industry standards EEMUA Publication 191 [19] and ANSI/ISA-18.2 [20], which result in proportional trust delay times. We define the alarm rate levels for a single operator. The low category signifies the maximum alarm rate deemed to be acceptable, with high operator trust. The medium category signifies the maximum alarm rate deemed to be manageable. The high category contains alarm rates above the medium category and is deemed to be over-demanding

The operator's total response and assessment time *OAT* for a single alarm is calculated as follows:

$$OAT = AT + LT + TDT. \quad (5)$$

where the TDT value is selected based on the system NAR/FAR category. OAT is used to calculate μ in the queueing model.

We did not include the trust delay time or allow variability in P_S or P_A in the original model. Since each solution on the Pareto frontier represents a unique PPS design solution, the inclusion of additional variables would make it difficult to quantify the impact of each variable on the probability of interruption. Instead, the operator performance analysis begins with a single PPS design solution generated by the original model. The number of operators is varied from 1 to *k*. The P_S for each sensor type and the P_A for operators are varied across a range provided by the site physical security expert. The alarm rate is calculated as a function of P_S in lieu of actual performance data.

The goal of this second model is to identify the acceptable trade-off between the system performance P_I, cost of employing more operators, P_S, and the NAR.

5 Illustrative Analysis

For our analysis, we assume a site size of 400m x400m, with a single target at the center of the site. The baseline architectures are generated using the PPS design optimization model, using a single barrier type (fence), a single sensor type with $P_S = 0.9$, and $P_A = 1.0$. We selected the PPS design option with the highest P_I.

Since we did not have access to actual operator performance data, we looked for data in the literature. We found a small set of studies that looked at changes in response time due to system reliability [10, 16, 17, 18, 21], but the target identification tasks were not analogous to alarm station operator tasks. Due to lack of data, the values used are notional (see Table 1).

We calculated the NAR/FAR categories based on an operator assessment time of 45 seconds, which gives a maximum possible rate of 1,920 assessments performed per day. The queueing model showed this value to be too high, even with no trust delay. We set this as the medium category threshold and set the low category threshold to half of that value. Increasing sensory sensitivity increases the rate of nuisance alarms [2], so we created the sensor alarm rate function (Table 1) to generate a range of alarms based on the P_S, from 1 alarm at $P_S = 0.3$ to 10 alarms at $P_S = 0.9$.

We performed two experiments: one which included trust delay in the assessment time and one that did not. First, the highest P_I PPS architecture with no trust delay was generated (Figure 2, a) and was used to compare the impact of no trust delay versus trust delay on system performance. Second, , the highest P_I PPS architecture with trust delay was generated (Figure 2, b) and compared to a second design with fewer sensors (Figure 2, c) to analyze the trade-offs in system performance P_I, ten-year cost, number of operators, P_S, and the NAR.

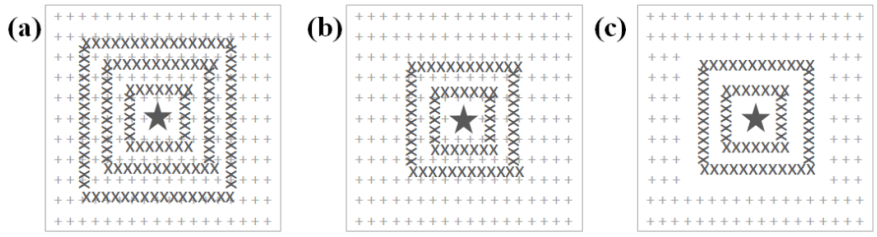

Fig. 2. The PPS architectures for Trust/No Trust Delay (a) and Low Sensor (b)/High Sensor (c) experiment conditions. The target is protected by fences and sensors.

Table 1. Notional values used in all of the operator performance model experiments.

Input	Variable	Notional Values for Experiments
Max number of operators	k	5
Operator assessment time	AT	45 seconds
Coefficient of variation in assessment time	C	0.1 (assume standard deviation is 10% of a mean with Gaussian distribution)
Probability of assessment	P_A	range from 0.6 to 1.0 in 0.1 steps
Lag time for task switching	LT	1 second
Trust delay time categories	TDT	low: 0 seconds (high trust, immediate response) medium: 10 seconds high: 20 seconds
Alarm rate categories		low threshold: 960 alarms/day or 40 alarms/hour medium threshold: 1,920 alarms/day or 80 alarms/hour high threshold: > 1,920 alarms/day
Sensor alarm rate		$12 * (P_S)^2$
Sensor probability of detection	P_S	range from 0.3 to 0.9 in 0.05 steps
Response force time		70 seconds (assume standard deviation is 10% of a mean with Gaussian distribution)

5.1 No Trust Delay Versus Trust Delay

The baseline architecture for this experiment is three fences surrounding the target with 272 sensors (NAR/FAR range 272-2720 alarms/day) across the site (Figure 2). The best $P_I = 0.9992$ with 3operators. The operator performance model varies the number of operators, P_S, and P_A. The step effect seen in the graph in Figure 3 is due to increases in P_S that are not significant enough to raise the NAR per the sensor alarm rate function (see Table 1).

No Trust Delay Condition. This experiment excluded the trust delay time from assessment times to establish a baseline and to show the impact of the queueing model. Figure 3 shows the average processing time for a single alarm, which includes the queue wait time and the operator assessment time. The higher times indicate a longer queue. The queueing model shows that even at the low alarm rates, the average processing time is higher than the operator assessment time due to queue wait times. With the notional values used, one to two operators are most sensitive to increasing P_S (which increases the NAR), with minimal impact on operator groups of size three to five.

The higher alarm processing times negatively affect system performance in terms of P_I. Higher P_S lead to higher P_I for all P_A, as seen in Figure 4, a and Figure 5, a for three to five operators. However, the increased alarm processing times for one to two operators reduce the maximum P_I that can be achieved. When P_I drops to zero, the NAR level overwhelms the operator, who can no longer clear the queue.

Trust Delay Condition. The effects of the trust delay are seen for one to two operators (Figure 3). For one operator, the processing time jumps when the medium trust delay time is added. For two operators, the processing time jumps when the high trust delay time is added. Trust delay has no impact on three to five operators, which have the same processing times as the no trust delay condition.

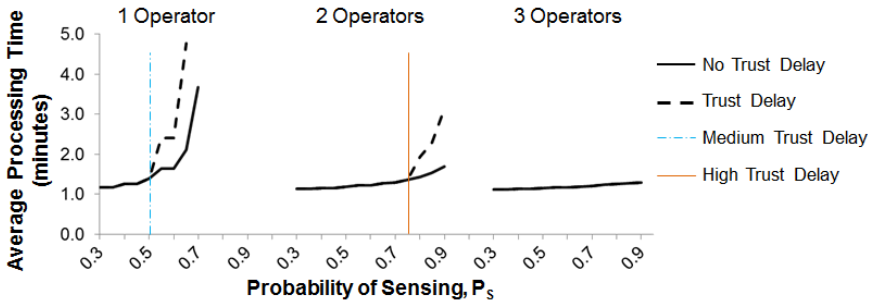

Fig. 3. Average processing times for a single alarm (in minutes) for the no trust/trust delay conditions. No trust/trust delay have the same values for the three operator group. Four and five operator groups have similar results to the three operators, so they are not shown.

As in the no trust delay condition, the increase in alarm processing times impacts the maximum P_I that can be achieved. The dip in the one operator P_I curve in Figure 4, b is caused by the addition of the medium trust delay and occurs for all P_A. The decline in the P_I curve for two operators is caused by the addition of the high trust delay and is most visible for the lower $P_A = 0.6$ and 0.7. Figure 5, b shows that for the highest NAR level ($P_S = 0.9$), the longer processing times for two operators (Figure 3) greatly reduce the system's P_I for all P_A.

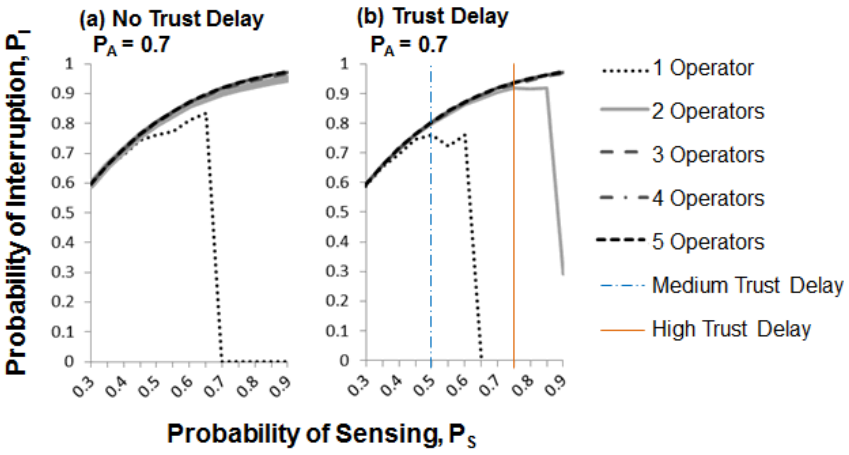

Fig. 4. System performance for operator P_A = 0.7 improves as the P_S increases. The higher NAR levels have a negative impact on one to two operators. In both conditions, three to five operators have the same values.

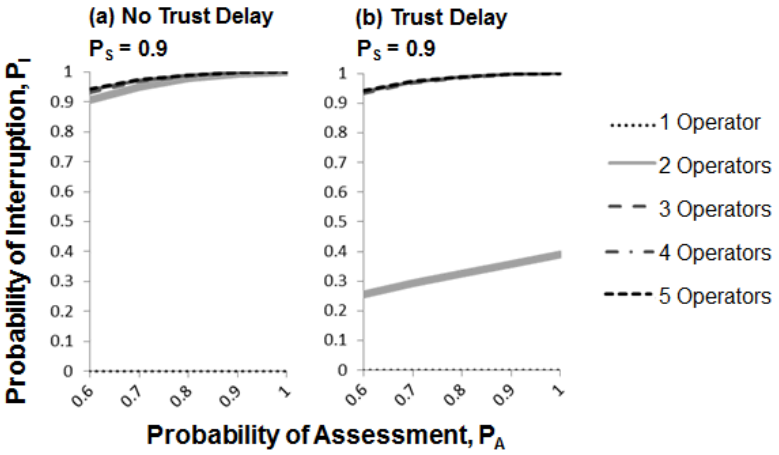

Fig. 5. Comparison of system performance P_I when P_S = 0.9 for the no trust/trust delay conditions as operator P_A increases. P_I = 0 for one operator. In both conditions, three to five operators have the same values.

5.2 Low Versus High Number of Sensors

The baseline architecture for this experiment is two fences surrounding the target (Figure 2 b, c). The high sensor architecture has 324 sensors (NAR range of 324-3240

alarms/day) across the site and has a ten-year cost of $115M. The best $P_I = 0.9995$ with five operators. The low sensor architecture has a sensor gap around and within the fences to simulate a work area that is not alarmed in order to avoid nuisance alarms. The best $P_I = 0.9898$ with three operators, 184 sensors (NAR range of 184-1840 alarms/day) and has a ten-year cost of $67M.

Let us assume a goal of $P_I \geq 0.8$. Initially, we assume an operator $P_A = 1.0$. We exclude the single operator (Figure 6, a) because he can only achieve the target P_I for a narrow band of P_S. In the low sensor solution, two operators can achieve $P_I \geq 0.8$ (Figure 6, b). The high sensor solution achieves better system performance at lower P_S, but eventually the two operators' P_I declines to zero. For the low sensor with two and three operators, the performance curves are almost identical. Therefore, the low sensor solution with two operators appears to be the better choice since it is significantly cheaper and has lower NAR than any high sensor solution.

What if operator $P_A = 0.6$? The low sensor solution can never achieve $P_I \geq 0.8$. With three operators, the high sensor solution meets the target, but P_I eventually drops below the target (Figure 6, c). With four operators, the high sensor solution meets the target (Figure 6, d). The high sensor solution with four operators achieves the best system performance, but costs $105M. The system owner will need to decide if the increased performance warrants the increased cost and higher NAR.

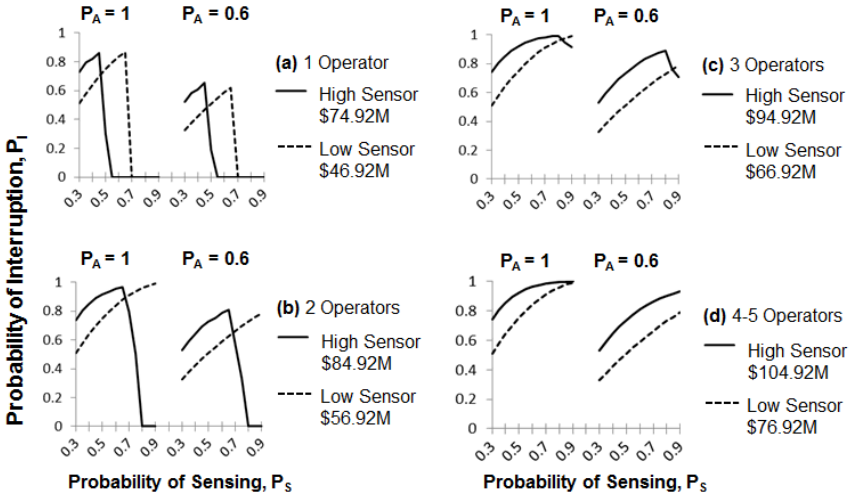

Fig. 6. Comparison of system performance P_I for the low/high sensor solutions when operator $P_A = 0.6$ and 1.0. In both conditions, five operators have similar performance to four operators.

6 Conclusion and Future Work

With limited budgets, sites are searching for technologies that can reduce the number of staff employed, thereby reducing the overall cost of the security system. However, more technologies and sensors will increase the nuisance and false alarm rates (NAR/FAR), which in turn will require more operators to maintain current system performance levels. While analyses of security systems often ignore operator performance, its inclusion is important to improving the accuracy of system performance estimates. Our illustrative analysis demonstrates that current methods that do not account for operator performance produce optimistic estimates.

First, the queueing model shows that even at the low alarm rates, the average alarm processing time can be higher than the operator assessment time due to queue wait times. In our illustrative analysis, the lowest processing time was well above the assessment time, even for the larger groups of operators.

Second, system owners need to consider the current alarm rate generated by their system and the additional alarms generated by new technologies. There is a maximum the number of alarms that can be assessed during a day, which can lead to missed detections when the system alarm rate exceeds this threshold. Our model can help system owners understand where these thresholds occur.

Third, operator assessment performance can have a major impact on system performance. If we assume an overly optimistic assessment performance when analyzing architectures, then the system performance will also be lower than anticipated.

Future work will include adding alarm priorities in the queueing model and adding response rate to the probability of assessment in our model, which adds the scenario where operators silence alarms without assessment. There are many more details that can be added, but adding more realism to the operator's probability of assessment requires more studies. There is limited data on how well alarm station operators can detect a target against various backgrounds on a video monitor. Studies are needed to obtain a baseline of operator assessment performance, including the impact of factors such as the vigilance decrement, the work environment, and the psychophysical characteristics of the assessment tasks.

Acknowledgements

The authors would like to thank John L. Russell and Judi See of Sandia National Laboratories. Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. SAND 2016-XXXX

References

1. Duggan, R.A., Wood, B.J.: Red Teaming of Advanced Information Assurance Concepts. In: Proceedings of DISCEX'00, pp. 112--118. IEEE, (2000)
2. Garcia, M.: The Design and Evaluation of Physical Protection Systems. Butterworth-Heinemann, Oxford (2007)

3. Bennett, H.A.: The "EASI" Approach to Physical Security Evaluation (NUREG-760145). U.S. Nuclear Regulatory Commission, Washington, D.C. (1977)
4. Matter, J.C.: SAVI: A PC-Based Vulnerability Assessment Program (SAND88-1279). Sandia National Laboratories, Albuquerque (1988)
5. Winblad, A., Snell, M., Jordan, S.E., Key, B., Bingham, B., Walker, S.: The ASSESS Outsider Analysis Module (SAND89-1602C). In: 30th Annual Meeting of INMM, pp. 420--425. (1989)
6. Moray, N., Inagaki, T., Itoh, M.: Adaptive Automation, Trust, and Self-Confidence in Fault Management of Time-Critical Tasks. Journal of Experimental Psychology: Applied, vol. 6, no. 1, pp. 44--58 (2000)
7. Bliss, J.P.: An Investigation of Alarm Related Incidents and Incidents in Aviation. International Journal of Aviation Psych., vol. 13, no. 3, pp. 249--268 (2003)
8. Breznitz, S.: Cry Wolf: The Psychology of False Alarms. Lawrence Erlbaum Associates, Hillsdale (1984)
9. Bliss, J.P., Dunn, M.C.: Behavioural Implications of Alarm Mistrust as a Function of Task Workload. Ergonomics, vol. 43, no. 9, pp. 1283--1300 (2000)
10. Bliss, J.P., Gilson, R.D., Deaton, J.E.: Human Probability Matching Behavior in Response to Alarms of Varying Reliability. Ergonomics, vol. 38, no. 11, pp. 2300--2312 (1995)
11. Sorkin, R.D.: Why Are People Turning Off our Alarms? Journal of the Acoustical Society of America, vol. 84, no. 3, pp. 1107--1108 (1988)
12. Xiao, Y., Seagull, F.J., Nieves-Khouw, F., Barczak, N., Perkins, S.: Organizational-Historical Analysis of the "Failure to Respond to Alarm" Problems. IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans, vol. 34, no. 6, pp. 772—778 (2004)
13. Brown, N.J.K., Jones, K.A., Nozick, L.K., Xu, N.: Multi-Layered Security Investment Optimization Using a Simulation Embedded Within a Genetic Algorithm. In: Proceedings of 2015 WinterSim, pp.2424--2435. IEEE Press (2015)
14. Gans, N., Koole, G., Mandelbaum, A.: Telephone Call Centers: Tutorial, Review, and Research Prospects. Manufacturing & Service Operations Management, vol. 5, no. 2, pp. 79--141 (2003)
15. Bransby, M.L., Jenkinson, J.: The Management of Alarm Systems: A Review of Current Practice in the Procurement, Design and Management of Alarm Systems in the Chemical and Power Industries. Health and Safety Executive Research Report CRR 166 (1998)
16. Dixon, S.R., Wickens, C.D., Chang, D.: Unmanned Aerial Vehicle Flight Control: False Alarms Versus Misses. In: Proceedings of HFES, pp. 152--156. SAGE Publications (2004)
17. Bustamante, E.A., Bliss, J.P., Anderson, B.L.: Effects of Varying the Threshold of Alarm Systems and Workload on Human Performance. Ergonomics, vol. 50, no. 7, pp. 1127--1147 (2007)
18. Chancey, E.T., Bliss, J.P., Proaps, A.B., Madhavan, P.: The Role of Trust as a Mediator Between System Characteristics and Response Behaviors. Human Factors, vol. 57, no. 6, pp. 947--958 (2015)
19. EEMUA: Alarm Systems: A Guide to Design, Management and Procurement (Third Edition) Publication 191. The Engineering and Materials Users Association, London (2013)
20. ISA: ANSI/ISA-18.2-2009 Management of Alarm Systems for the Process Industries. International Society of Automation, Research Triangle Park (2009)
21. Dixon, S.R., Wickens, C.D., McCarley, J.S.: On the Independence of Compliance and Reliance: Are Automation False Alarms Worse Than Misses? Human Factors, vol. 49, no. 4, pp. 564--572 (2007)