SANDIA REPORT

SAND2016-9533
Unlimited Release
Printed September 2016

A Complex Systems Approach to More Resilient Multi-Layered Security Systems

Nathanael J. K. Brown, Katherine A. Jones, Alisa Bandlow, Linda K. Nozick, Lucas A. Waddell, and Drew Levin

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

Sandia National Laboratories is a multi-mission 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.

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

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

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

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

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.fedworld.gov
Online order: http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online

SAND2016-9533
Unlimited Release
Printed September 2016

A Complex Systems Approach to More Resilient Multi-Layered Security Systems

Nathanael J. K. Brown, Katherine A. Jones, Alisa Bandlow, Linda K. Nozick, Lucas Waddell, Drew Levin, and Jonathan Whetzel
Department 06131, Operations Research and Computational Analysis (ORCA),
Department 06133, Systems Readiness and Sustainment Technology
Sandia National Laboratories
P.O. Box 5800
Albuquerque, New Mexico 87185-MS1188

Abstract

In July 2012, protestors cut through security fences and gained access to the Y-12 National Security Complex. This was believed to be a highly reliable, multi-layered security system. This report documents the results of a Laboratory Directed Research and Development (LDRD) project that created a consistent, robust mathematical framework using complex systems analysis algorithms and techniques to better understand the emergent behavior, vulnerabilities and resiliency of multi-layered security systems subject to budget constraints and competing security priorities. Because there are several dimensions to security system performance and a range of attacks that might occur, the framework is multi-objective for a performance frontier to be estimated. This research explicitly uses probability of intruder interruption given detection (P_I) as the primary resilience metric. We demonstrate the utility of this framework with both notional as well as real-world examples of Physical Protection Systems (PPSs) and validate using a well-established force-on-force simulation tool, Umbra.

ACKNOWLEDGMENTS

The authors would like to thank the following Sandia National Laboratories staff for making various contributions to the project: Kristin L. Adair, John L. Russell, Mark K. Snell, Christopher J. Symonds, Sarah Walsh, and Jonathan H. Whetzel. Additionally, we would like to acknowledge the outstanding technical partnership with Cornell University and the student who made significant contributions to parts of this work, Ningxiong Xu.

This work was supported by Laboratory Directed Research and Development funding from Sandia National Laboratories.

CONTENTS

FIGURES

NOMENCLATURE

ASO	alarm station operator
CCTV	Closed-Circuit Television
DHS	Department of Homeland Security
DoD	Department of Defense
DOE	Department of Energy
FAR	false alarm rate
GA	genetic algorithm
HPC	high performance computing
IAEA	International Atomic Energy Agency
LDRD	Laboratory Directed Research & Development
LP	linear program
MILP	mixed integer linear program
MLS	multi-layered security
MORS	Military Operations Research Society
MVP	most vulnerable path
NAR	nuisance alarm rate
NNSA	National Nuclear Security Agency
P_D	probability of detection
P_I	probability of interruption given detection
PPS	physical protection system
RFT	response force time
RV	random variable
SNL	Sandia National Laboratories
TA-V	Tech Area V

1. INTRODUCTION

The configuration of layered security measures is at the center of efforts to protect a range of systems, from high-value facilities to large-scale infrastructures. Historically, analyses of security systems have been performed using directed graph and path analysis tools like Adversary Sequence Diagrams (ASD). However, there are many dimensions in the design space of a security system, including selection of technologies, alternative locations/configurations, different threats, and competing cost limitations. The dimensionality of this problem makes it effectively impossible to evaluate all permutations of potential system architectures using traditional methods. The experience of the individuals configuring the system drives the careful examination of a small subset of architectures.

The key goal of this Laboratory Directed Research and Development (LDRD) is the creation of a consistent, robust mathematical framework using complex systems analysis algorithms and techniques to better understand the emergent behavior, vulnerabilities and resiliency of multi-layered security systems subject to budget constraints and competing security priorities. Because there are several dimensions to security system performance and a range of attacks that might occur, the framework must be multi-objective for a performance frontier to be estimated. Since security measures can fail for a range of reasons, this research explicitly includes resiliency as a dimension of system performance.

The tools developed under this project can directly benefit the Department of Energy (DOE) and the National Nuclear Security Administration (NNSA) as well as other federal agencies such as the Department of Homeland Security (DHS) and the International Atomic Energy Agency (IAEA). All agencies have goals, objectives, or mandates related to protecting material (chemical, biological, radiological, and nuclear), facilities, or supply chains in a fiscally responsible manner. For example, the NNSA strategic plan specifically states an objective to "improve understanding of the interaction between risk and cost. The NNSA will execute programs at the lowest cost without sacrificing either critical mission elements or our commitment to operating in a safe, secure, and environmentally sound manner."

This document describes the capabilities developed over the span of the LDRD, including the associated applications. The first section of this document describes the key concepts and technical approach, but relies heavily on referencing the publications that we produced as a result of this effort. The second section describes our validation using Dante/Umbra. Since we were not able to gain access to alarm station operator (ASO) data, we used CCTV operators as a surrogate and conducted a separate literature review which appears in the appendix.

2. SUMMARY OF TECHNICAL APPROACH

The key goal of this LDRD was to create a mathematical framework to better understand the emergent behavior, vulnerabilities, and resiliency of multi-layered security systems subject to budget constraints and competing security priorities. This goal required that we:

- Create a mathematical representation of a multi-layered security system represented as a complex system.
- Provide insight into the trade-off between performance and cost.
- Generate investment strategies with resilience metrics that can be independently validated.

In order to achieve these goals, we had to create the following artifacts:

- A model of the security architecture of a physical protection system (PPS) as informed by the Physical Security Center of Excellence.
- A representation of the "human element" including intruder behavior and Alarm Station Operator (ASO) performance as impacted by nuisance and false alarm rates (NAR/FAR).
- An optimization to estimate the triple-objective trade-off frontier: probability of interruption (P_I, primary resilience metric), investment cost, and NAR/FAR.

The following sections describe the technologies we used to create these different artifacts.

2.1. Model of a Physical Protection System

We use a network of nodes and arcs to represent a PPS, where each node maps a physical location, and each arc defines the path between two nodes. Zero or more investments can be made on each arc and can include either detection elements (sensors) or delay elements (barriers such as fences or walls). Figure 1 shows a network representation of a simple PPS where an intruder can enter from any perimeter node and desires to reach the magenta node at the center. Once an intruder is detected, the response force is notified and must intercept (interrupt) the intruder before he/she reaches the target (a.k.a., "hands on target"). Our initial approach to computing the probability of interruption assumed a constant response force time (RFT) and constant intruder travel times across each link (though different links may have different travel times due to the addition of delay investments) and is described in [1]. Though a fairly standard approach, this method only includes the uncertainty of detection and not the uncertainty of the various travel times. To improve this model, we introduced uncertainty around the RFT in the form of exponentially distributed times as described in [2]. Unfortunately, the assumption of an exponentially distributed RFT forced us to use a simulation which was fairly compute-intensive. We improved this approach by assuming that all travel times (both intruder as well as response force) are Gaussian in nature, which allowed us to efficiently add time uncertainty to the P_I calculation without significant loss of generality. Additionally, we wanted to account for network impacts due to lighting and weather effects as well as intruders with enhanced skills, which would be seen as variations in performance for the various detection and delay elements. The addition these new sources of model uncertainty are described in [4].

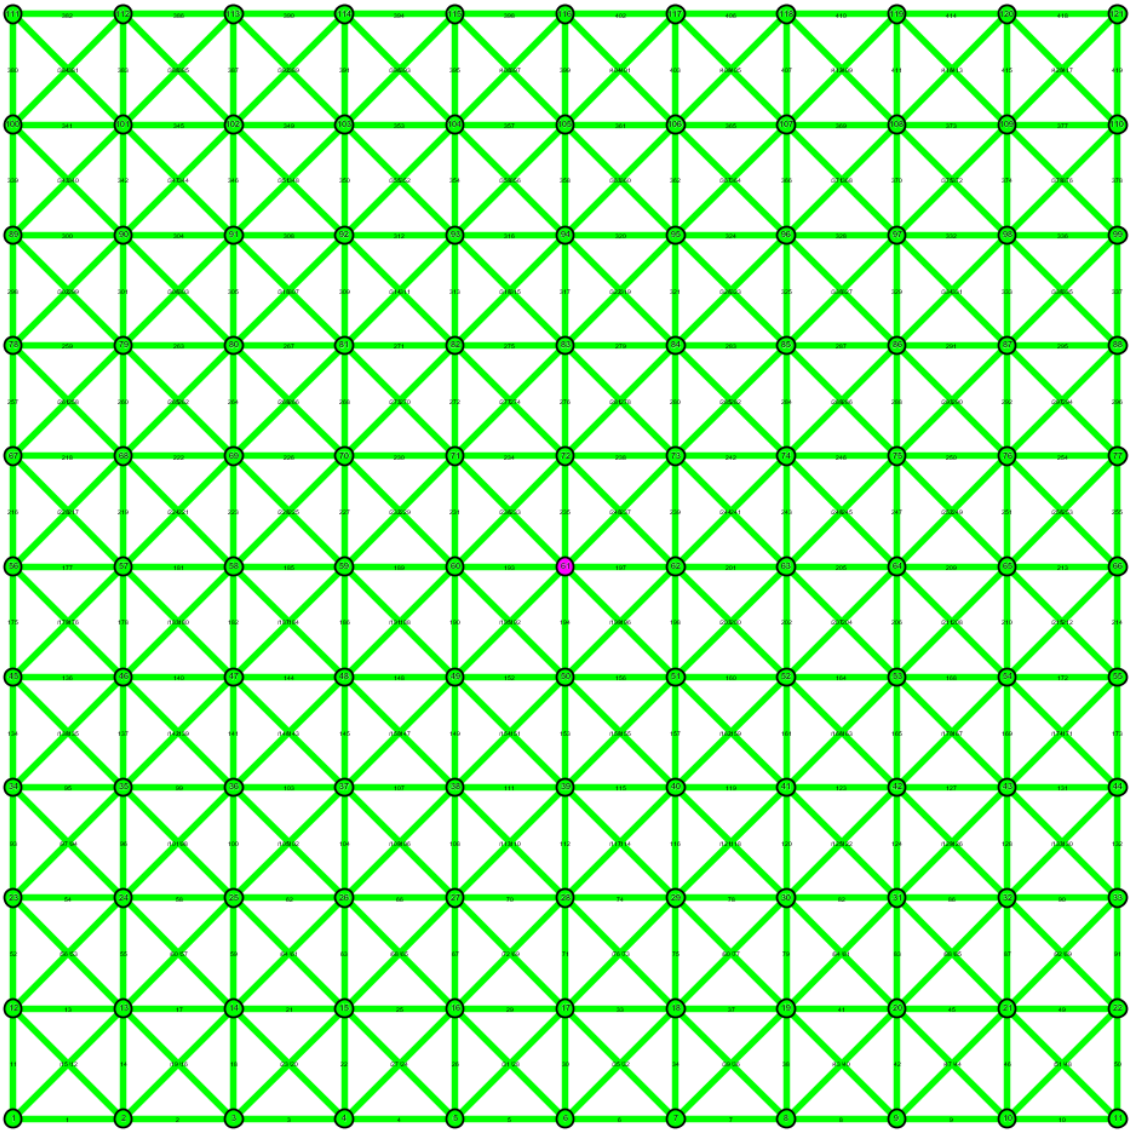

Figure 1: PPS Network Model with Single-Arc Investments

2.2. Representation of the "Human Element"

We addressed two different aspects of the human element in our model: the intruder and the ASO. For the intruder, we assume that they have perfect knowledge of the physical layout of the PPS that they wish to infiltrate. Given this knowledge, they will always take the path that provides them with the lowest P_I, often referred to as the most vulnerable path (MVP). The description of how this path is determined appears in [1] (no time or weather uncertainty), [2] (uncertainty in the RFT) and [4] (uncertainty in the RFT, intruder travel times, and lighting/weather effects). To enhance the intruder model, we included the possibility that some intruders have "enhanced" knowledge of the sensors and barriers such that they can degrade the system's performance as described in [4].

For the ASO, we wanted to create a representation such that they could be viewed as an integral part of the PPS that could be enhanced with greater investment. In [1], [2] and [5], it is assumed that there is a sufficient number of ASOs available such that every sensor detection is immediately assessed and reported to the response force. Our use of NAR/FAR minimization in the optimization was meant to ensure that the ASOs would not be overwhelmed by a high number of alarms. To improve upon this model, we incorporated three new elements as described in [3]: a minimum alarm assessment and communication time, the impact of alarm queuing due to elevated NAR/FAR, and the effect of trust lag time which can occur if the ASO views the system as unreliable (due to high NAR/FAR). With this improved model, an investment in more ASOs is reflected as a decrease in overall RFT which results in a higher P_I. Conversely, architectures which utilize a large number of sensors without a sufficient number of ASOs can be eliminated as candidates since the sheer number of false alarms can make it impossible to process an alarm in a timely manner (i.e., before the intruder gets hands on target).

Our standard queuing model assumes that all alarms are of equal importance, an assumption unlikely to be true in practice. In the context of a high consequence security system, it may be advantageous to assess specific alarm types ahead of others. A priority queue allows events to be serviced based upon an exogenous ranking of importance; events of higher priority will be serviced first. Further, priority queues can help avoid deadlock conditions by allowing high-priority events to be seen, even when the overall arrival rate of events is higher than the rate at which they can be serviced. For these reasons, we developed a priority queue model which establishes priority based on distance from target, where the sensors with the highest priority are those located closest to the target. In scenarios with high NAR/FAR, the use of a priority queue results in a higher P_I and lowers the required number of ASOs compared to a standard queue.

When developing a priority queue model, there are many choices that must be made. After talking to PPS experts, we chose to model the PPS using a *preemptive repeat* priority queue. *Preemptive* refers to the fact that high-priority sensors can preempt lower-priority ones, interrupting their service and sending them back to the front of the queue. When an interrupted sensor gets reassessed later, we assume that this assessment is *repeated* (not resumed), meaning that the process starts over from the beginning as if the sensor was being assessed for the first time. We also assume that the time between sensor activations is exponentially distributed corresponding to a Poisson arrival process and that the time taken by ASOs to service an alarm is normally distributed with a coefficient of variation of 0.1. These assumptions lead to non-parametric skewed-normal distributions of alarm waiting plus assessment times (see Figure 2, for example), precluding the use of an analytical approximation of the expected values.

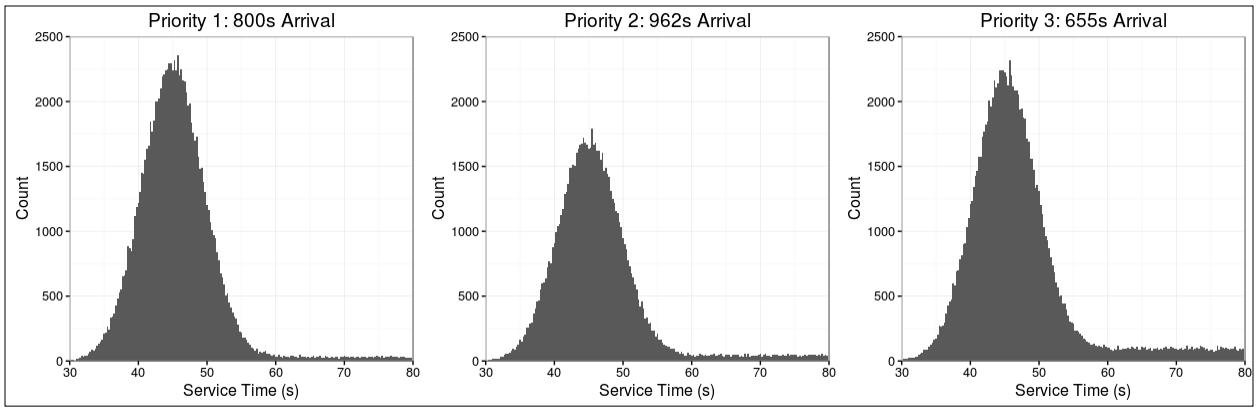

Figure 2: Distribution of Sampled Service (Wait + Assessment) Times for a Three-Level Priority Queue

Given a PPS architecture with known sensor placement, service times, and NAR/FAR, we calculate the expected value of P_I using simulation. In our simulation, sensor events are generated randomly according to their NAR/FAR, and the alarms are serviced based upon their arrival order and sensor priority. Queue wait plus assessment times, RFTs, and link travel times are sampled and converted to a binomial random variable that indicates whether or not an intruder successfully reaches the target given that a sensor with a specific priority level is triggered. The binomial random variables corresponding to each sensor priority are sampled until a predefined level of confidence of the expected value is obtained, giving us an estimate for $P_I^{(k)}$, the probability of interruption given that the intruder was detected by a sensor with priority level k.

Confidence is defined in terms of two parameters, α and ε, such that the simulation stops once the sampled binomial probabilities are within ε of the true expected value with a confidence probability of $1-\alpha$. We use two formulas to determine the current level of confidence. For the case where the sampled binomial random variable has produced events of only one type (either all successes or failures), we use the generalized form of the Rule of Three [7]:

$$\varepsilon = \frac{-\log \alpha}{n} \qquad (1)$$

where n refers to the number of samples taken. If samples of each type exist, we use the normal approximation interval:

$$\varepsilon = z_{1-\frac{\alpha}{2}} \sqrt{\frac{p(1-p)}{n}} \qquad (2)$$

where n refers to the number of samples taken, z_σ refers to the σ quantile of a Gaussian distribution, and p refers to the proportion of successes. In the case where low-priority sensor wait times cannot be sampled due to non-ergodic queue conditions (the low-priority events

are 'frozen out' by higher-priority events), $P_I^{(k)}$ for that specific sensor priority level is assumed to be zero.

Once $P_I^{(k)}$ is estimated for each priority level k, we can calculate the probability that an intruder is successful by taking the product (across all priority levels) of the probabilities that the intruder either 1) avoids detection by a sensor having the specific priority level, or 2) triggers the sensor but is able to reach the target before the sensor is assessed and responded to. The overall probability of interruption is simply the complement of this probability of intruder success:

$$P_I = 1 - \prod_k \left(\left[1 - P_D^{(k)} \right] + P_D^{(k)} \left[1 - P_I^{(k)} \right] \right) \quad (3)$$

where $P_D^{(k)}$ is the probability that the intruder is detected by a sensor with priority level k.

To test the general effectiveness of a priority queue model, we performed two experiments. First, we measured P_I for queues having one, two, and three priority levels as the NAR/FAR varied. For simplicity, the priority levels were assumed to be equidistant from the target, and each priority level was given an equal alarm arrival rate. The results demonstrate that splitting queues into multiple priority levels allows for more robust performance as NAR/FAR increases (see Figure 3). Using a standard queue with high NAR/FAR, it is likely that no sensor would ever be assessed quickly enough to interrupt the intruder. Using a priority queue, at least those sensors having the highest priority will be assessed in sufficient time to possibly lead to intruder interruption.

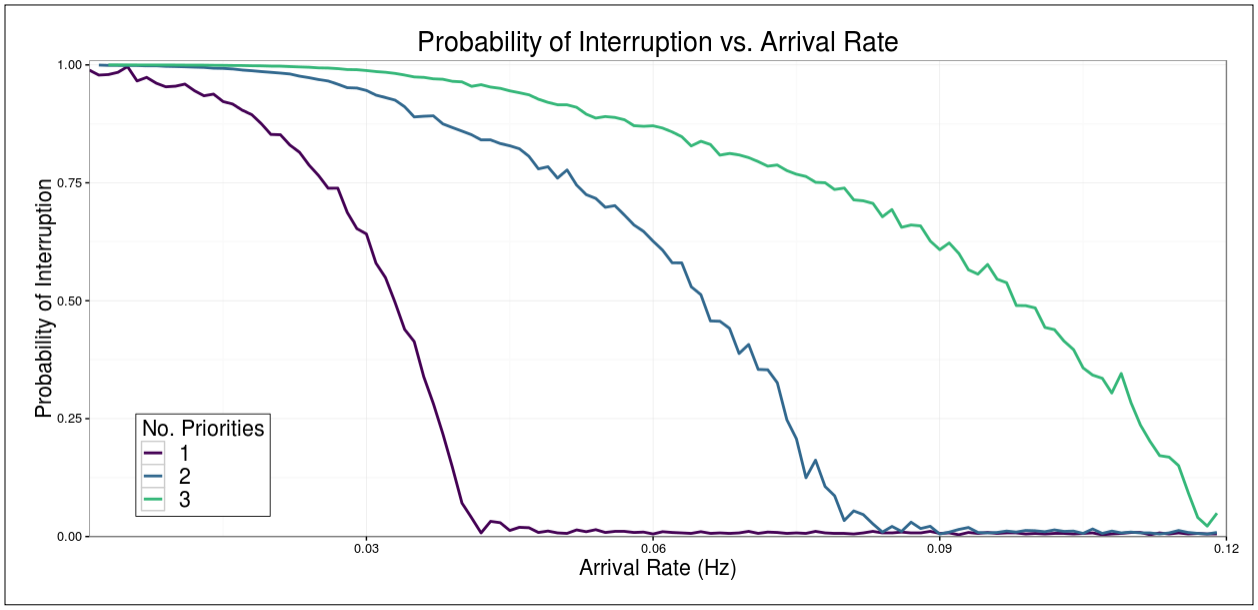

Figure 3: PI vs. NAR/FAR for Queues having One, Two, and Three Priority Levels

Next, we examined the effect that the number of ASOs has on queues with one, two, and

three priority levels, given high NAR/FAR. Similar to the previous experiment, NAR/FAR and distance from the target were assumed to be equal across priority levels. The results show that queues split into priority levels are able to operate successfully with fewer ASOs in situations where standard queues struggle (see Figure 4).

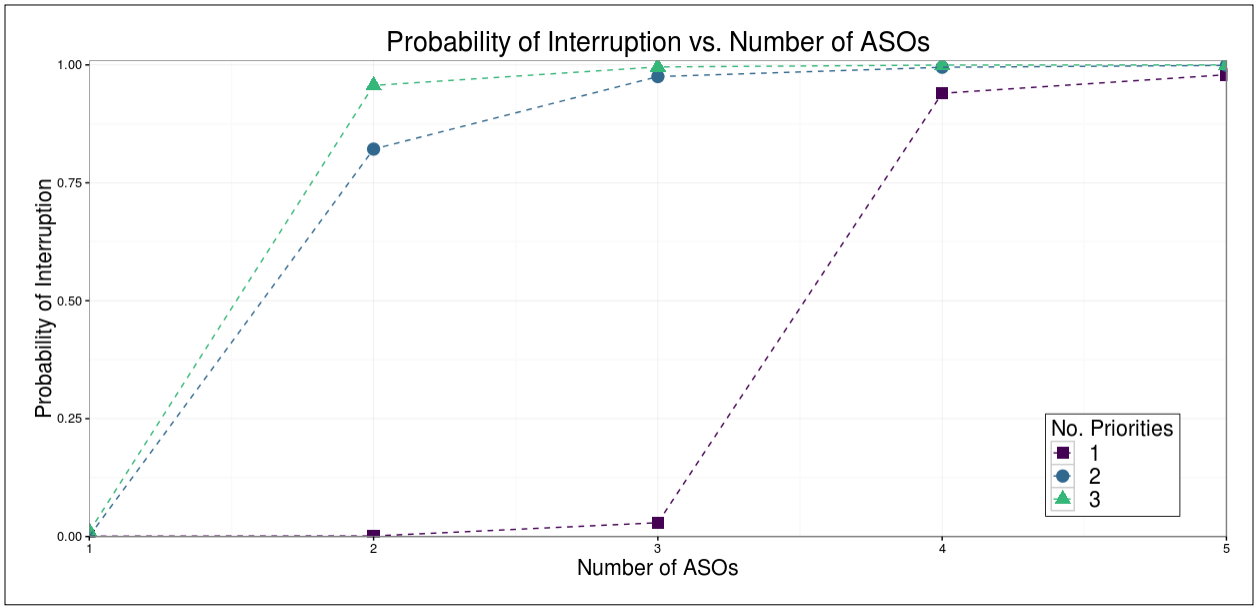

Figure 4: PI vs. No. of ASOs for Queues having One, Two, and Three Priority Levels

Based on these results, it is clear that priority queues possess a distinct advantage over standard queues in scenarios having high NAR/FAR. Priority queues are able to remain functional in situations where arrival rates are higher than service rates because they provide a mechanism for ASOs to ignore low-priority signals, guaranteeing quick service and response to higher-priority events. The GA optimizer is able to detect these scenarios and either remove the superfluous low-priority sensors to save money or increase the number of ASOs to improve performance.

We assigned priority levels to sensors based on a scheme where the sensors with the highest priority are those located closest to the target. Future work would examine alternative prioritization strategies. For example, while it is critical to detect intruders who have managed to get close to the target, it may be more beneficial to prioritize detecting them upon their initial arrival to give the response team more time to perform a successful interruption. Alternatively, sensors might be prioritized based upon their relative confidence and reliability; giving precedence to sensors with low NAR/FAR and high detection rates.

2.3. Investment Planning Optimization

At the heart of our investment planning optimization is a game-theoretic attacker-defender model. The main ideas behind this approach are:

- **Intruder goal**: Minimize the probability that the travel time remaining after detection will exceed the response time of the protective force (probability of interruption).
- **System owner goal**: Maximize the probability that the intruder will be interrupted given that the intruder can adapt to different investment strategies.
- **System owner decision**: What technologies and physical barriers to invest in and where to place them subject to budget and false alarm rate limits.

One approach to solving this problem could be to examine all possible architecture permutations and simulate the attacker-defender response for each architecture. To illustrate why this approach is not feasible, we refer once again to our network model of a PPS. For even the small PPS representation of Figure 1, there are 420 arcs. If we assume six different investment types, then there can be up to 64 different investment permutations per arc. The number of investment permutation for the entire network is then 64^{420} or about 10^{758} (greater than the number of particles in the universe!). By comparison, our "real world" example of SNL's Tech Area V has roughly 4800 arcs which results in approximately 10^{8670} permutations.

Since an optimization-via-simulation approach is not feasible with even a modestly-sized PPS, we considered two different mathematical approaches to solving this problem: Mixed-Integer Linear Program (MILP) and metaheuristic. The advantage of the MILP is that it provides an elegant, provably optimal, closed-form solution. Unfortunately, this approach does not scale well to large networks. The work on this approach was initiated in 2014, but has been refined and submitted as [4].

The metaheuristic approach uses a genetic algorithm (GA) to perform a multi-objective optimization trading off investment cost, NAR/FAR, and P_I. The advantages of the GA approach are that it has a more flexible implementation, scales well to large problems, and is HPC-compatible (due to its parallel implementation). The disadvantage of this approach is that there is no guarantee of optimality in the solutions produced. Our initial foray was presented at the 82nd MORS Symposium (2014) in Working Group 28 and was nominated for the Barchi Prize (best symposium paper). In response to this nomination, we generated a paper more fully describing the details of what was presented in [1]. Due to extenuating circumstances at MORS (possibly related to a vacancy in the editor position), this paper was never accepted for publication, but still represents a major milestone for this project.

A significant improvement to the optimization approach made in [4] over the previous approach is the use of investment layers instead of individual arc investments. In this context, an investment layer is a collection of same-type arc investments which form a closed perimeter around the target as shown in Figure 5. The advantage of this approach is that it reduces the number of permutations by several orders of magnitude while utilizing a reasonable investment strategy that would likely be employed in practice. The main disadvantage is that the investment placement is more restricted and may not allow for fully optimal solutions.

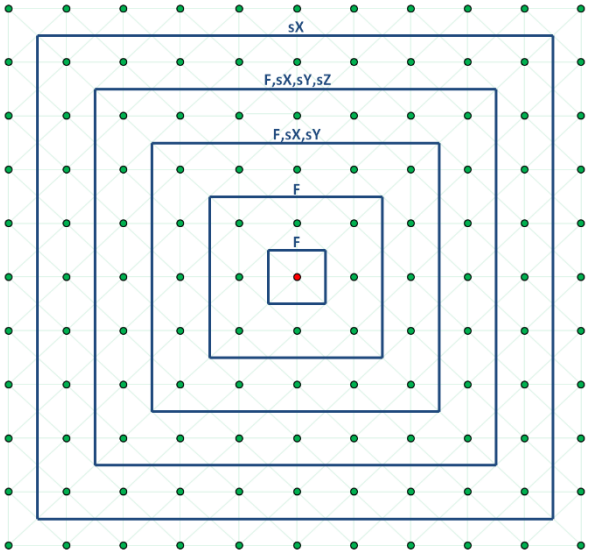

Figure 5: PPS Network Model with "Layered" Investments

As previously mentioned, the only source of uncertainty in our initial PPS model was the probability of detection (P_D) for each sensor. Our final GA-based investment optimization procedure (as described in [4]) includes the following sources of uncertainty:

- RFT (Gaussian RV)
- Arc travel time (Gaussian RV)
- ASO assessment and standard queue wait times (Gaussian RV)
- Environmental and lighting impacts (scenario-based)
- Intruder capabilities (scenario-based)

In order to address the last two sources, a stochastic optimization approach had to be taken. The stochastic optimization requires that we evaluate the architecture's P_I for each weather/intruder scenario (where the technology performance varies according to the scenario) and optimize against the worst case value across all scenarios. The advantage of this approach becomes clear when compared to two alternate approaches that might be chosen by a naïve designer:

- Use the average value P_I across all scenarios as the performance metric
- Create an "average" architecture which uses the average sensor/barrier values (across all scenarios) to conduct a single scenario optimization

For the first case study, we compare the worst-case P_I against the average P_I across all scenarios. The plot in Figure 6 illustrates that selecting an architecture based on the average P_I could leave the PPS vulnerable to the worst case where the P_I is always lower (often substantially). Each point in the plot represents a solution on the Pareto Frontier.

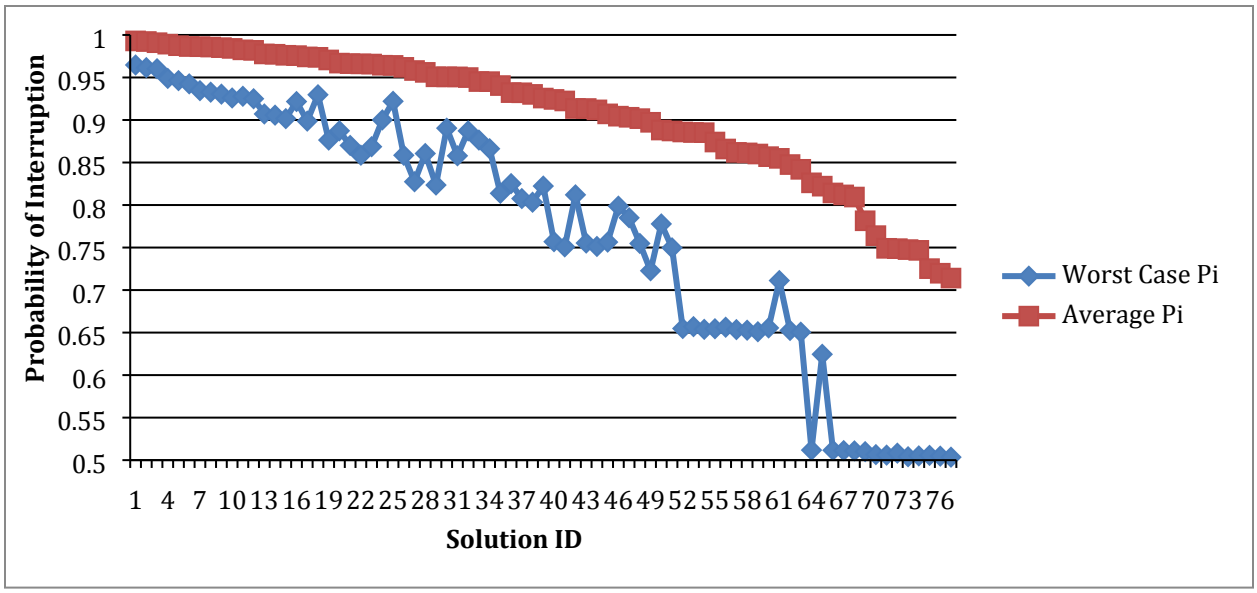

Figure 6: Worst Case vs. Average PI

For the second case study, we determine the average performance value for each technology investment across all scenarios. We can then perform a single scenario optimization which produces architectures based on these average performance values. When comparing the P_I from the average architecture to the true average across all scenarios (as shown in Figure 7), we see that the scenario-average P_I is much lower. When compared to the worst-case (WC) P_I across all scenarios, the differential is even more stark (as shown in Figure 8).

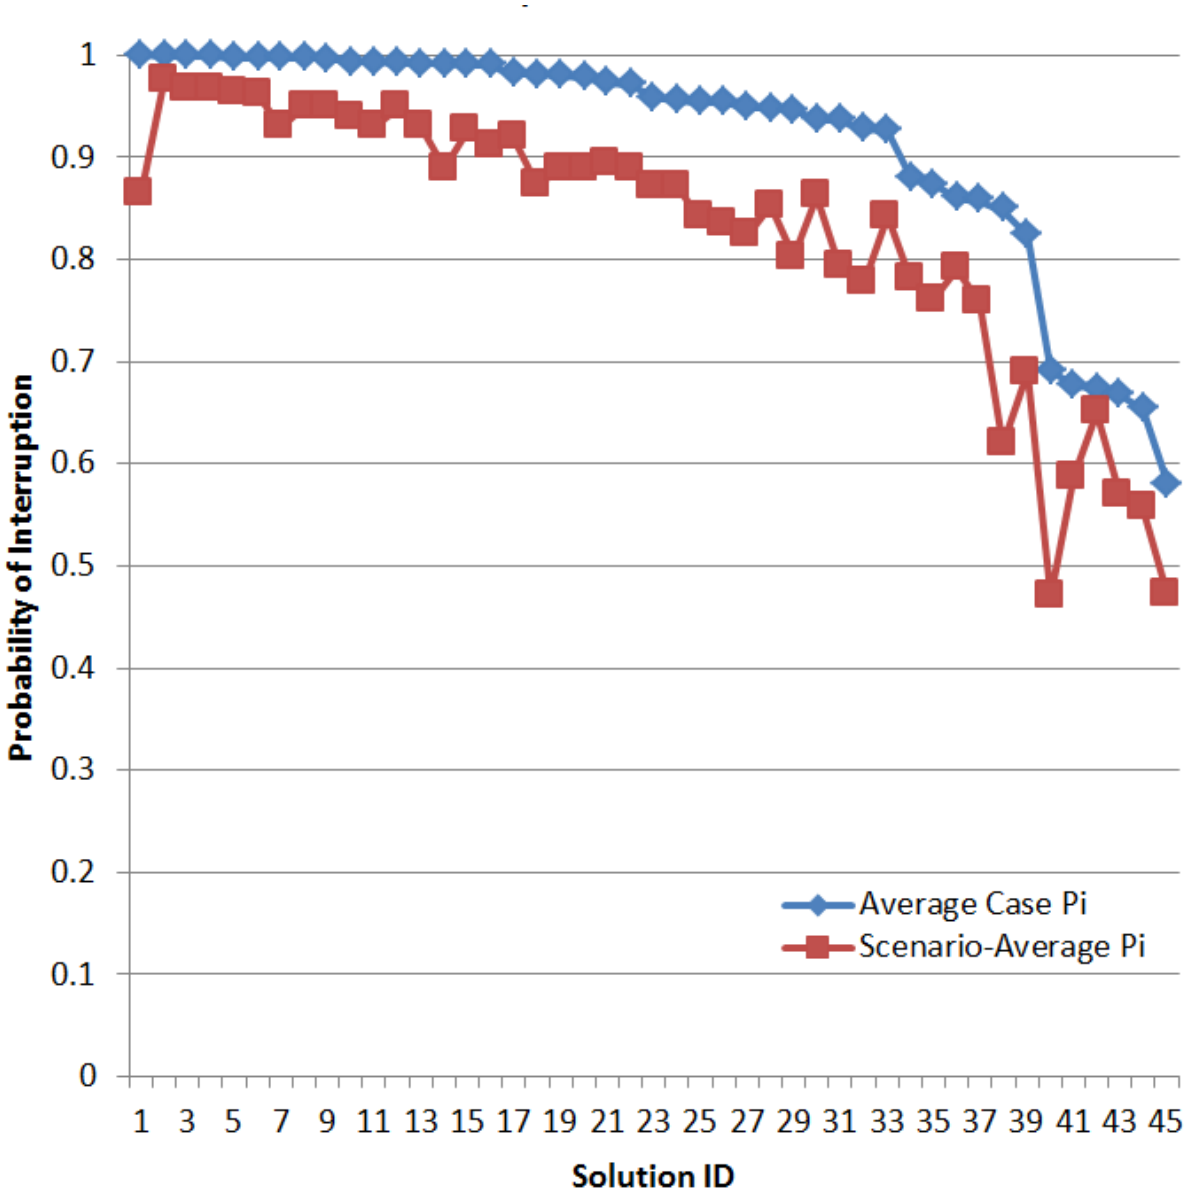

Figure 7: Average Architecture PI vs Scenario-Average PI

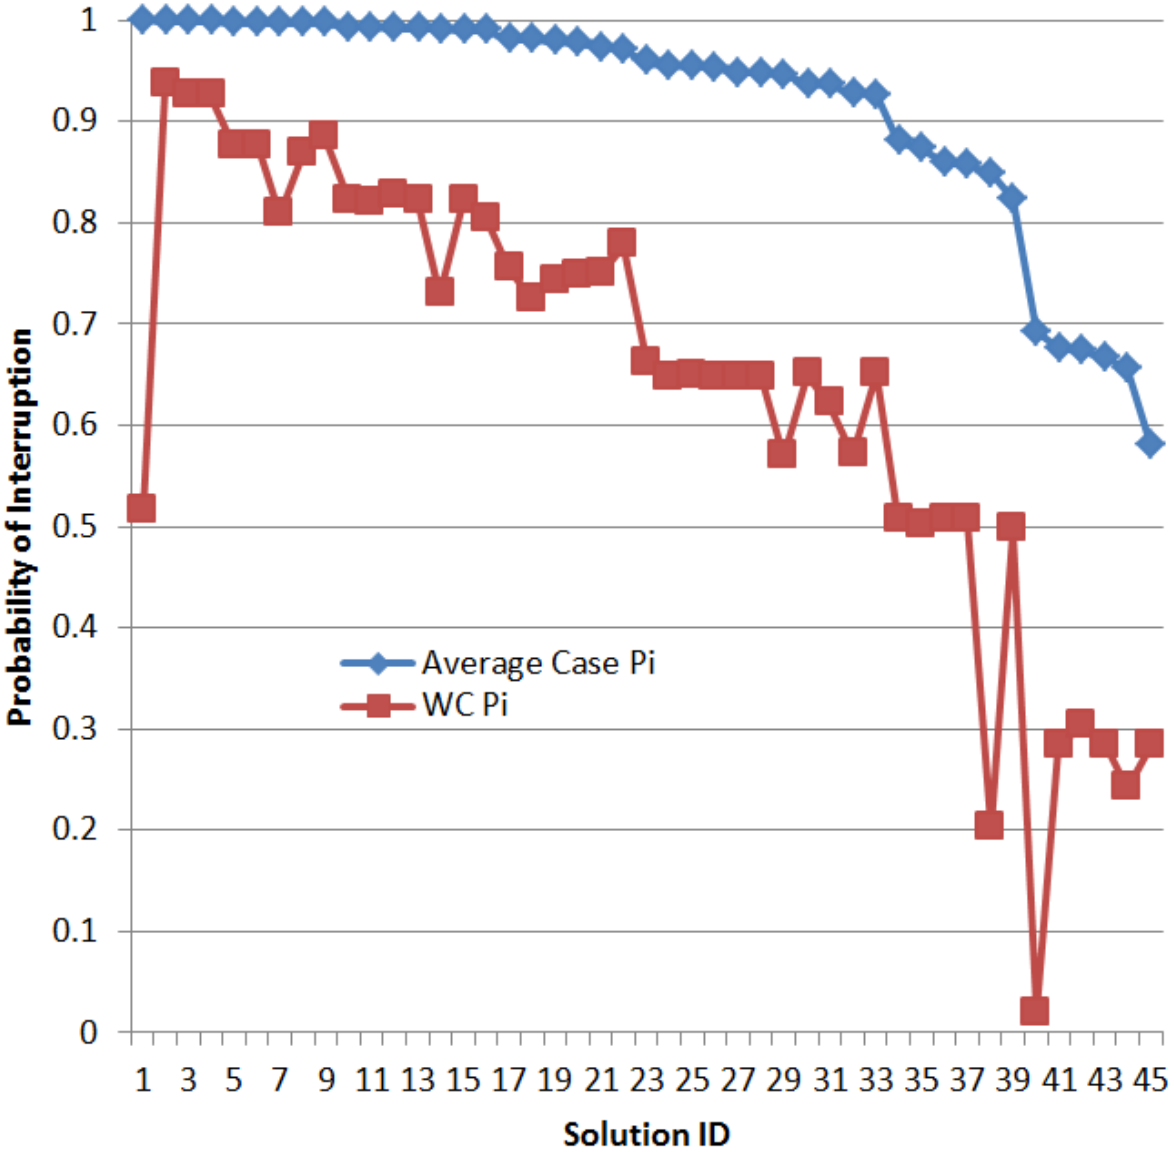

Figure 8: Average Architecture PI vs Worst Case PI

3. VALIDATION USING DANTE/UMBRA

In order to validate our methodology for calculating P_I for a given PPS architecture, we used the Dante Scenario Editor, which is built upon a well-established force-on-force simulation tool, the Sandia-developed Umbra Simulation Framework [6]. This validation process consists of two main steps: *scenario creation* and *batch analysis*. In order to illustrate the process, we will examine a PPS based on SNL's Tech Area V (TA-V).

3.1. Dante Scenario Creation

To create a PPS scenario, Dante first loads in 3D terrain data for the site of interest, and then produces a 2D grid overlay with a resolution that matches the one used by the MLS optimization engine. Dante then removes nodes (colored red) that overlap with buildings and other terrain objects that a person could not pass through. Figure 9 shows the loaded TA-V terrain data with an overlay grid resolution of 10 meters.

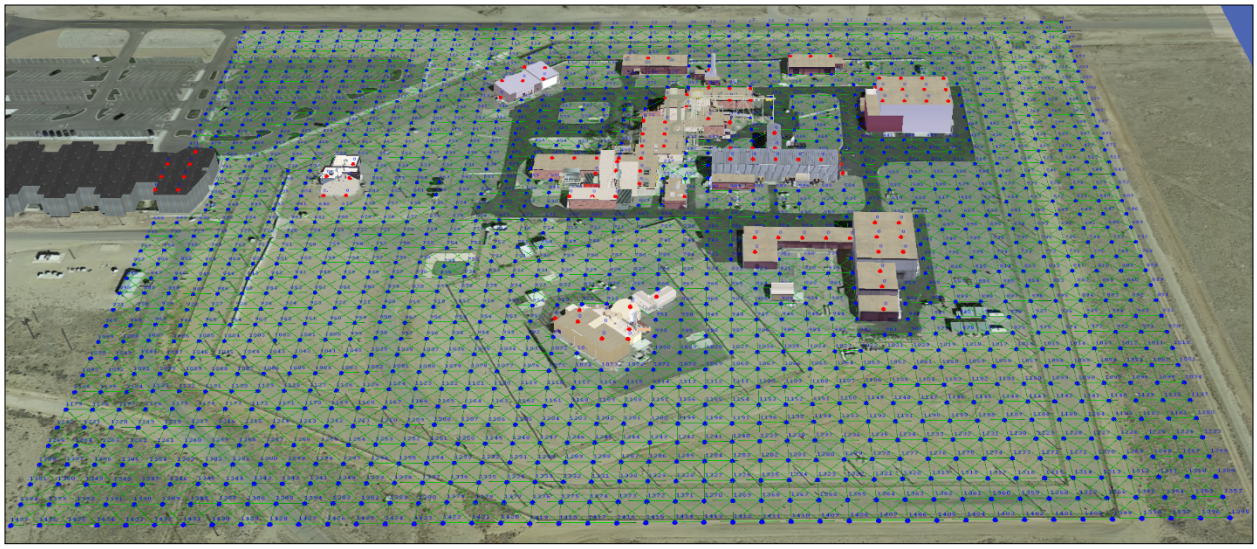

Figure 9: TA-V 3D Terrain Data with 2D Grid Overlay

Next, Dante imports a solution from the MLS optimization engine. This solution defines the PPS investments made, such as numbers and locations of fences and sensors. Each solution asset is mapped to a simulated sensor within the Dante library, and each sensor's coverage area is defined in such a way to make the architecture match the MLS engine's "layered" design philosophy (see Figure 5). Several adjustments to Dante's scenario building algorithm (e.g., extending fences and expanding buildings) were needed to ensure that the resulting design matched that of the MLS engine, with no unwanted gaps between investment assets and buildings. Figure 10 shows the Dante editor after MLS Solution A has been imported. This particular solution consists of four layers of fences, located at radii of 20, 30, 40, and 50 meters from the target, and three layers of sensors, located at radii 40, 50, and 60 meters from the sensor. The fences are represented by white arc investments and the sensors are

represented by blue arc investments. The Dante UI gives the user the opportunity to interactively explore a variety of additional information about each investment.

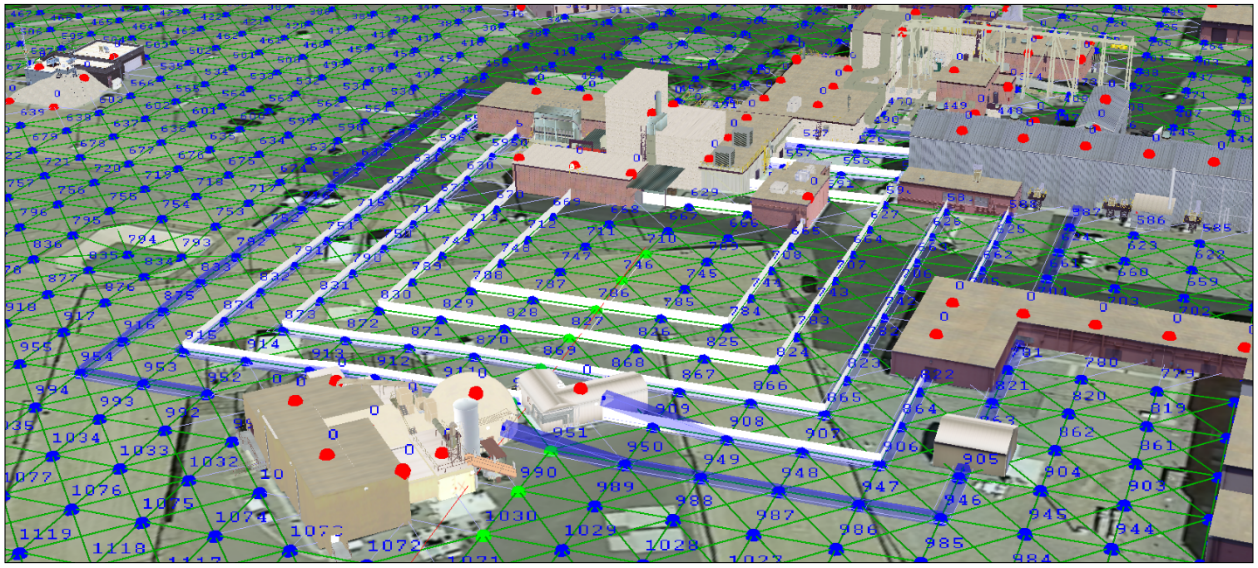

Figure 10: Solution A Investments

Once a solution is selected, the user is given the ability to set the parameters for a scenario via the dialogue box shown in Figure 11. Here, the user can input values for fence breach time, intruder speed, RFT, and sensor assessment time (which includes average queue waiting time for the MLS models that incorporate the single queue strategy). If the "Use Variable Time/Speed" box is checked, the quantity will be given a Gaussian distribution with the entered standard deviation; otherwise, the quantity will be constant.

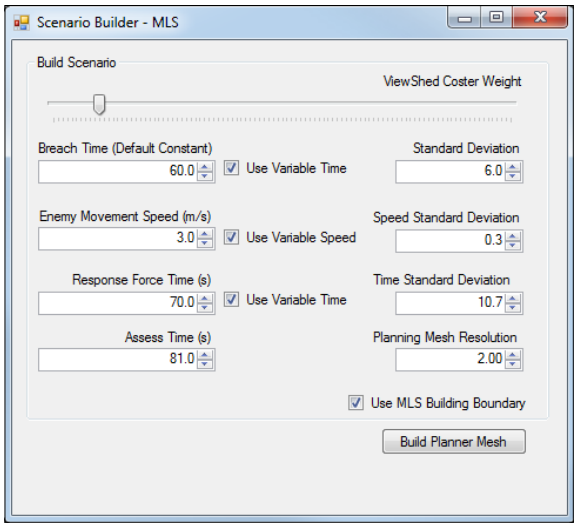

Figure 11: Dante Scenario Parameters

The Planning Mesh Resolution gives the resolution of the network on which the simulated intruder will be allowed to travel. Figure 12 shows Solution A with part of the intruder's navigation mesh highlighted.

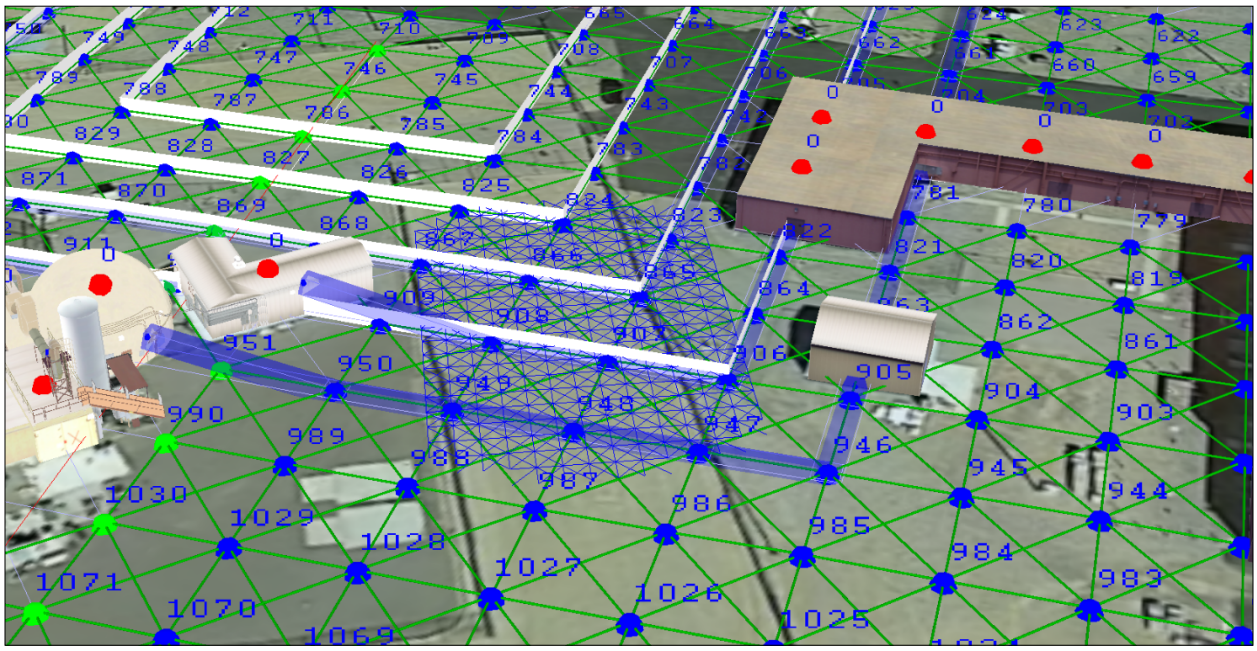

Figure 12: Solution A with Intruder Navigation Mesh

3.2. Dante Batch Analysis

Dante's batch analysis feature allows the simulation engine to run a scenario multiple times with a random start position for the intruder along the navigation mesh perimeter. This randomization is in addition to the possibly random parameters shown in Figure 11.

One of the main differences between the MLS optimization engine and Dante is the way that the intruder chooses his path to the target. In MLS, the intruder is assumed to have perfect knowledge of the PPS, and therefore always chooses the MVP. In Dante, the intruder uses the A* search algorithm for determining the best path from his random start location to the target. The A* algorithm is a heuristic where each link in the path is given a cost based on a multiplier of its length and a cost value. This cost value is based on geometric "costers" that increase the cost value for navigation links that cross a sensor's detection area (cost based on the ViewShed Coster Weight slider seen in Figure 11) or a fence (cost based on the fence breach time). The algorithm then performs a greedy search for the path with the lowest cost.

When an intruder encounters a sensor, the sensor detects them with a probability imported from the MLS engine's solution data. Each sensor possesses a timer that begins to run if the sensor reports the detection of an intruder, where the length of each timer is randomized based on the response force and assessment times provided for the scenario. The simulation run ends in a loss for the intruder if 1) any sensor's timer runs to completion before the intruder reaches the target, or 2) the intruder's path planning algorithm cannot find a possible

route to the target. Otherwise, the simulation ends in a win for the intruder. If an intruder encounters both a sensor and a fence in the same layer, the sensor will have a chance to detect the intruder *before* the intruder is delayed by the fence.

The results from a batch of scenario runs highlight how an intruder was successful/unsuccessful across the various runs. Figure 13 shows example results from a batch of 500 runs of a TA-V scenario. The intruder lost (blue) 498 out of the 500 simulation runs, or 99.6% of the time. Dante not only reports how many times the intruder wins and loses, but also gives the exact reason why each loss occurred. This information could be extremely valuable in determining which sensors and sensor types are critical to the PPS.

Overall Statistics

Condition	Count	Percentage	Average Time	Minimum Time	Maximum Time
Blue	498	99.6	219.9	0.8	391.4
->sensor75-11149 detected the red agent.	40	8.0	220.8	202.8	248.4
->sensor60-16444 detected the red agent.	36	7.2	223.4	197.4	243.2
->sensor60-16784 detected the red agent.	35	7.0	201.4	181.6	228.4
->sensor75-11489 detected the red agent.	34	6.8	203.2	185.6	245.0
->sensor60-18072 detected the red agent.	23	4.6	190.9	174.4	208.8
->sensor75-11157 detected the red agent.	22	4.4	210.7	163.0	245.0
->sensor60-16538 detected the red agent.	22	4.4	202.1	181.0	226.0
->sensor60-16721 detected the red agent.	22	4.4	211.9	185.8	233.2
->sensor60-17698 detected the red agent.	21	4.2	198.5	181.6	216.8
->sensor75-11481 detected the red agent.	20	4.0	203.1	173.0	237.6
->sensor60-16452 detected the red agent.	19	3.8	215.6	194.0	250.8
->sensor60-16500 detected the red agent.	17	3.4	215.2	180.6	238.4
->sensor75-11243 detected the red agent.	16	3.2	202.6	185.0	219.8
->sensor75-12777 detected the red agent.	15	3.0	195.3	178.8	211.0
->sensor75-11205 detected the red agent.	15	3.0	211.8	186.8	225.4
->sensor75-11346 detected the red agent.	13	2.6	340.8	322.0	361.2
->sensor60-18080 detected the red agent.	12	2.4	190.6	174.0	210.6
->sensor90-6074 detected the red agent.	12	2.4	355.3	331.4	382.8
->sensor75-11369 detected the red agent.	12	2.4	348.6	328.0	391.4
->sensor75-11426 detected the red agent.	12	2.4	206.5	184.0	244.0
->sensor75-12403 detected the red agent.	9	1.8	195.3	179.4	209.6
->sensor75-11256 detected the red agent.	8	1.6	347.0	306.2	365.2
->sensor60-16776 detected the red agent.	8	1.6	204.3	196.8	210.6
->Attacker failed to reach goal.	8	1.6	0.8	0.8	0.8
->sensor60-16654 detected the red agent.	7	1.4	205.3	187.6	218.0
->sensor75-11359 detected the red agent.	7	1.4	203.5	186.2	224.6
->sensor90-5961 detected the red agent.	7	1.4	351.4	325.8	375.0
->sensor90-6051 detected the red agent.	7	1.4	342.9	324.0	358.8
->sensor75-12785 detected the red agent.	6	1.2	180.1	155.4	198.4
->sensor75-11287 detected the red agent.	4	0.8	204.2	201.0	207.6
->sensor60-16582 detected the red agent.	4	0.8	192.4	186.8	196.8
->sensor75-12786 detected the red agent.	3	0.6	198.5	197.0	200.6
->sensor60-18081 detected the red agent.	1	0.2	181.4	181.4	181.4
->sensor60-16571 detected the red agent.	1	0.2	208.0	208.0	208.0
Red	2	0.4	405.5	375.6	435.4
->Attacker reached goal.	2	0.4	405.5	375.6	435.4

Figure 13: Dante Batch Analysis Results

Dante also provides a path display result that shows the routes taken by the intruder over all scenario runs, color-coded to show which were wins (red) and which were losses (blue) for the intruder. Win paths are useful in that they indicate areas with potentially weak defenses. Figure 14 displays a portion of this result for the batch of simulation runs from Figure 13. Here, a circle represents the beginning of a path, a square represents the end of a path, and red indicates that the intruder successfully reached the target.

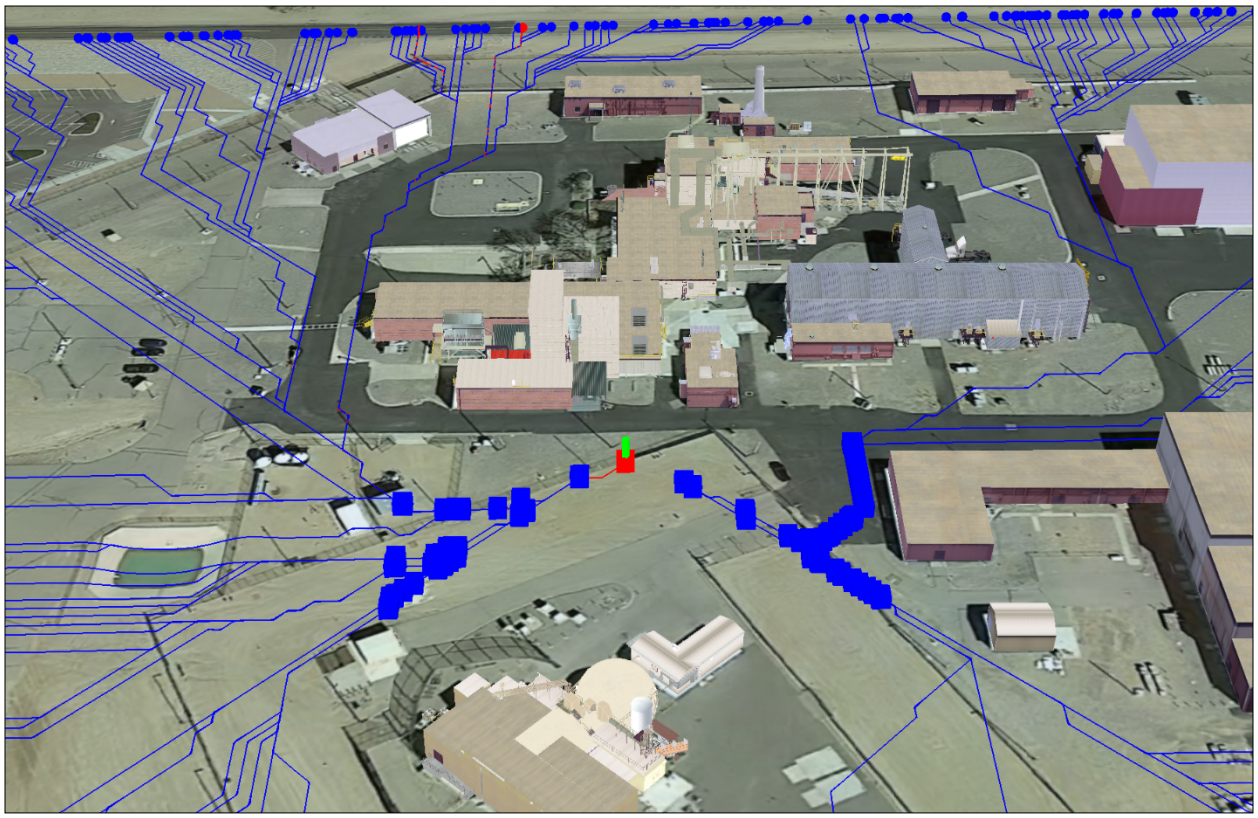

Figure 14: Dante Batch Analysis Path Display Result

These results verify the MLS optimization engine calculation, which reported a P_I of 0.997 for this solution. We saw similar consistency between MLS P_I and Dante batch analysis results for every solution that we tested.

4. SUMMARY AND CONCLUSIONS

The overarching goals of this LDRD project were achieved by creating an analysis framework and software implementation that:

- Automatically generates optimized and robust investment options for new and existing PPS designs
- Optimizes across multiple environmental and adversary scenarios to create a family of solutions that trade off performance and cost
- Models the impacts of system NAR/FAR on alarm station operators

To achieve these goals, we integrated the following capabilities into a rigorous analytic framework: PPS architecture modeling, modeling of intruder and ASO behavior, and game theory. These models were combined with a heuristic optimization engine to create an investment planning optimization which was validated using a real-world example simulated in a realistic 3D environment. Numerous artifacts were produced as a result of this effort including: four publications submitted to peer-reviewed journals, four technical advances, and a substantial library of software components which compose the framework. The external publications and conference presentations increase Sandia's visibility in the areas of:

- Advanced PPS investment planning optimization
- Multi-objective stochastic optimization
- Measuring the impacts of high NAR/FAR on PPS performance

This research is also in direct support of the Resilience in Complex Systems Research Challenge, which focuses on:

- Quantification of resilience metrics
- Techniques for addressing system uncertainty
- Methods for validation and verification

Next steps in this research for the analytic 2D model could proceed in at least the following three areas. First, this model assumed a single attacker identifying the weakest path during the most vulnerable weather and visibility conditions. In practice there may be locations which, if attacked, render other defenses less potent, such as a control center for video surveillance feed, for example. This opens up the possibility that teams of attackers working collectively with different goals may be able to create more potent attacks. Second, integrating additional priority queuing ideas into the modeling, including and extending beyond what was discussed in this report, would also be valuable. Finally, the analysis performed in this LDRD assumed that each sensor had a fixed NAR/FAR. For many sensors, it is likely that the weather conditions and visibility impact the NAR/FAR. This modification is easy to incorporate but also very likely to suggest that ASO staffing could fluctuate based on the environmental conditions.

Going beyond our current model, we could investigate adding 3D and real-time aspects. The addition of a third dimension in and of itself would add substantial

complexity to the processing but also significantly increase the realism. By moving to real-time analysis, the system could be used to dynamically assess threats and predict likely attack vectors in order to assist an ASO in determining the best course of action given an active threat environment. The latter extension would support the next generation physical security goals of moving beyond the current detect/delay/respond paradigm, reducing response times and improving situational awareness.

5. REFERENCES

1. Brown, N., K. Jones, L. Nozick, N. Xu, A. Bandlow, K. Adair, J. Gearhart, J. Russell, "Multi-layered Security Systems for Force Protection" (U), SAND2015-20767 C, Sandia National Laboratories, Albuquerque, NM, January 2015.
2. Brown, N., K. Jones, L. Nozick, and N. Xu, 2015. "Multi-Layered Security Investment Optimization Using Simulation Embedded Within a Genetic Algorithm." *Proceedings of the 2015 Winter Simulation Conference*. Edited by L. Yilmaz, W. K. V. Chan, I. Moon, T. M. K. Roeder, C. Macal, and M. D. Rossetti.
3. Bandlow, A., K. Jones, N. Brown and L. Nozick, 2016. "The Impact of False and Nuisance Alarms on the Design Optimization of Physical Security Systems." *Proceedings of the 7th International Conference on Applied Human Factors and Ergonomics, Advances in Human Factors and System Interactions*.
4. Brown, N., K. Jones, A. Bandlow, L. Waddell, and L. Nozick, 2016. "A Stochastic Programming Approach to the Design Optimization of Layered Physical Protection Systems." *2017 50th Hawaii International Conference on System Sciences (HICSS)* (accepted, available after January 2017).
5. Brown, N., K. Jones, L. Nozick, and N. Xu, 2016. "Optimizing the Configuration of Sensor Networks to Detect Intruders." *IEEE Systems Journal* (submitted).
6. Gottlieb, E., Harrigan, R., McDonald, M., Oppel, F., & Xavier, P. (2001). The Umbra simulation framework. *SAND2001-1533 (unlimited release), Sandia National Laboratory.*
7. Jovanovic, B., and P. Levy. "A look at the rule of three." *The American Statistician* 51.2 (1997): 137-139.

APPENDIX A: CCTV OPERATOR LITERATURE REVIEW

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, but the target identification tasks were not analogous to ASO tasks. Due to lack of data, we have used notional values in our work.

In addition to the lack of available data on operator assessment times, we were also unable to find data on operator assessment performance, probability of correct assessment P_A. An analogous domain is closed-circuit television (CCTV) operators in city surveillance. CCTV operators monitor multiple CCTVs to detect events and work in similar control room environments. The main difference between the operator domains is the monitoring task. CCTV operator performance research assumes that the operator is continuously monitoring CCTVs for events. ASOs tend not to perform continuous monitoring events and instead are notified via alerts to assess an event.

Vigilance Decrement

Vigilance, or sustained attention, is defined as the ability of observers to maintain their focus of awareness and remain alert to stimuli in the environment over prolonged periods of time (Davies and Parasuraman, 1982). Vigilance tasks, or monitoring tasks, involve maintaining attention over long periods of time to detect small changes in the information presented. Mackworth (1950) discovered the vigilance decrement, which is a decline in detection performance over time. Mackworth and Sawin and Scerbo (1995) found a decrease in performance during the first half hour of a monitoring period, which then stabilizes at a lower level. During a 90-minute monitoring task, Donald (2014) found that 23% of study participants lost concentration in the first 30 minutes, 62% in the second 30 minutes, and 50% in the final 30 minutes. However, a third of the participants showed no disengagement over the entire period.

Many factors contribute to the vigilance decrement, and SAND2014-17929 provides a good overview on this topic.

Detection Performance

Research shows that 100% target detection is difficult to achieve. Targets are also more difficult to detect in complex, “busy” footage.

Unsurprisingly, events that are expected and easily visible are more likely to be detected (Wells et al., 2006). Yet ASOs operate in a domain where real events are rare, and intruders attempt to be stealthy. In a study with 16 monitors of prison scenes with little movement and conspicuous, detection rates varied from 85%-97% (Tickner et al, 1972). However, when similar prison scenes were used showing a lot of movement, detection rates decreased dramatically (32%-100% across experimental conditions). Perfect detection rates could only be achieved when only one display was observed at a time. Donald (2011) found that no participant was able to identify all visible targets and only 12% detected 75% or more targets.

Studies show a decline in performance as the number of CCTV monitors increase. Wallace and Diffley (1998) found that the majority of operators reported confidence in monitoring up to 16 screens, with over half stating that they were comfortable monitoring a maximum of 4 screens. The operator-to-monitor ratio was 1:16 or higher in control rooms surveyed. In target detection tasks, as the number of monitors increased to 4, 9 and 16, detection performance measures gave accuracy of 83%, 84% and 64% respectively (Tickner and Poulton, 1973). Neil (2007) also found a decrease in target detection rates as the number of monitors increased from 1 (97%) to 6 (94%). He inferred that as the surveillance video became more complex, the number of screens monitored should be reduced to maintain the same detection rate. Rankin et al. (2012) agree that simple detection tasks produce better performance than more complex scenarios.

Operators can also suffer from change blindness. Change blindness occurs when people miss distinct changes in a visual scene (Dadashi et al, 2013). Becker and Pashler (2002) found that study participants failed to perceive changes in numbers in an array at which they had just looked and identified. Silverman and Mack (2006) found that change detection performance improved as the number of changed letters in a row increased.

Video Quality

Video quality can impact detection performance and increase the vigilance decrement. Video quality can be naturally degraded due to environmental factors such as lighting conditions, rain and fog. van Voortjuijsen (2005) degraded video quality by decreasing the frame rate and by increasing brightness. Study participant detection rates decreased as the video quality was degraded. In a comparison of target detection between high and low video quality, vigilance decrement was observed in the low quality video experiment where none was observed with the high quality video (Parasuraman et al., 2009). Participants also had a high false alarm rate and lower detection rate when observing low quality video.

Another aspect is how to determine the minimum video resolution required for detection. The Johnson Criteria (Johnson, 1985) is a widely used method for calculating the probability of target detection of an object imaged by an optical system. It has been widely used in the design of military systems. It can be used to determine the minimum resolution required to detect a target under various conditions. Sjaardema et al. (2015) found that while the criteria has been updated and improved over the years, it does not accurately predict target detection in all weather conditions and lacks modeling of the human element within a detection system.

The Rotakin test was developed by Jim Aldridge in 1989 to ensure the performance of CCTVs. The test is aimed at quantitatively determining deficiencies in the CCTV coverage by detecting blind spots and the limitations of operators in detecting camouflaged intruders in varying weather and lighting conditions. The Rotakin is designed to analyze three major performance parameters: viewing areas for cameras, size of images on monitors, and necessary observer response time. However, this research is specific to analog systems.

The most recent CCTV standards in the 2009 Operational Requirements (Cohen et al., 2009) do not make mention of the Rotakin test. Instead it outlines a two-level process for operators. Level 1 is a practice in planning where thought is given to the threat itself rather than the CCTV system. At this level statement of problem, stakeholders, risk assessment, and success criteria are outlined. Level 2 provides height based levels of detail which outline general requirements of figure size in order to: monitor, detect, recognize, and identify threats. The manual acknowledges that these minimum images to screen ratios may be smaller for CCTV digital systems that have higher resolution but also may be larger when cameras are in areas with low lighting or poor angles of view.

References

1. Aldridge, J., 1989. "The Rotakin-A Test Target for CCTV Security Systems." *Home Office Scientific Research and Development Branch.*
2. Becker, M.W. and H. Pashler, 2002. "Volatile visual representations: Failing to detect changes in recently processed information." *Psychonomic Bulletin & Review*, *9*(4), pp.744-750.
3. Cohen, N., J. Gattuso, and K. MacLennan-Brown, 2009. "CCTV Operational requirements manual 2009." *Home Office Scientific Developement Branch.*
4. Dadashi, N., A.W. Stedmon, and T.P.Pridmore, 2013. "Semi-automated CCTV surveillance: The effects of system confidence, system accuracy and task complexity on operator vigilance, reliance and workload." *Applied ergonomics*, *44*(5), pp.730-738.
5. Donald, F.M., 2011. "Optimal characteristics of inserted graphic objects in stimulating CCTV operator vigilance and performance." *Doctoral dissertation, Faculty of Humanities, University of the Witwatersrand, Johannesburg*.
6. Donald, F.M. and C.H. Donald, 2015. "Task disengagement and implications for vigilance performance in CCTV surveillance." *Cognition, Technology & Work*, *17*(1), pp.121-130.
7. Johnson, J., 1985. "Analysis of image forming systems." In *Selected papers on infrared design. Part I and II* (Vol. 513, p. 761).
8. Keval, H. and M.A. Sasse, 2006. "Man or gorilla? Performance issues with CCTV technology in security control rooms." In *16th World Congress on Ergonomics Conference, International Ergonomics Association* (pp. 10-14).
9. Neil, D., N. Thomas, and B. Baker, 2007. "Threat image projection in CCTV." In *Optics/Photonics in Security and Defence* (pp. 674102-674102). International Society for Optics and Photonics.
10. Parasuraman, R., E. de Visser, E. Clarke, W.R. McGarry, E. Hussey, T. Shaw, and J.C. Thompson, 2009. "Detecting threat-related intentional actions of others: effects of image quality, response mode, and target cuing on vigilance." *Journal of Experimental Psychology: Applied*, *15*(4), p.275.

11. Rankin, S., N. Cohen, K. MacLennan-Brown, and K. Sage, 2012. “CCTV operator performance benchmarking.” In *Security Technology (ICCST), 2012 IEEE International Carnahan Conference on* (pp. 325-330). IEEE.
12. Sawin, D.A. and M.W. Scerbo, 1995. “Effects of instruction type and boredom proneness in vigilance: Implications for boredom and workload.” *Human Factors: The Journal of the Human Factors and Ergonomics Society*, *37*(4), pp.752-765.
13. See, J., 2014. “Vigilance: A Review of the Literature and Applications to Sentry Duty. SAND2014-17929.” *Sandia National Laboratories*.
14. Silverman, M.E. and A. Mack, 2006. “Change blindness and priming: When it does and does not occur.” *Consciousness and cognition*, *15*(2), pp.409-422.
15. Sjaardema, T.A., C.S. Smith, and G.C. Birch, 2015. “History and Evolution of the Johnson Criteria. SAND2016-6368.” *Sandia National Laboratories*.
16. Tickner, A.H., E.C. Poulton, A.K. Copeman, and D.C.V. Simmonds, 1972. “Monitoring 16 television screens showing little movement.” *Ergonomics*, *15*(3), pp.279-291.
17. Tickner, A.H. and E.C. Poulton, 1973. “Monitoring up to 16 synthetic television pictures showing a great deal of movement.” *Ergonomics*, *16*(4), pp.381-401.
18. van Voorthuijsen, G., H.A.J.M. van Hoof, M. Klima, K. Roubik, M. Bernas, and P. Pata, 2005. “CCTV effectiveness study.” In *Proceedings 39th Annual 2005 International Carnahan Conference on Security Technology* (pp. 105-108). IEEE.
19. Wallace, E. and C. Diffley, 1998. “CCTV: Making it work.” *CCTV control room ergonomics. Police Scientific Development Branch Publication*, *14*, p.98.
20. Wells, H., T. Allard, and P. Wilson, 2006. “Crime and CCTV in Australia: Understanding the relationship.” *Centre for Applied Psychology and Criminology: Bond University, Australia.*

DISTRIBUTION

1	MS1188	Marcey L. Hoover	6130
1	MS1188	Dean A. Jones	6131
1	MS1188	Patrick Finley	6131
1	MS0899	Technical Library	9536 (electronic copy)
1	MS0359	D. Chavez, LDRD Office	1911


Sandia National Laboratories