Proceedings of the 2015 Winter Simulation Conference
L. Yilmaz, W. K. V. Chan, I. Moon, T. M. K. Roeder, C. Macal, and M. D. Rossetti, eds.

MULTI-LAYERED SECURITY INVESTMENT OPTIMIZATION USING A SIMULATION EMBEDDED WITHIN A GENETIC ALGORITHM

Nathanael J. K. Brown
Katherine A. Jones

Sandia National Laboratories
Post Office Box 5800
Albuquerque, NM 87185, USA

Linda K. Nozick
Ningxiong Xu

School of Civil and Environmental Engineering
Cornell University
220 Hollister Hall
Ithaca, NY 14853, USA

ABSTRACT

The performance of a multi-layered security system, such as those protecting high-value facilities or critical infrastructures, is characterized using several different attributes including detection and interruption probabilities, costs, and false/nuisance alarm rates. The multitude of technology options, alternative locations and configurations for those technologies, threats to the system, and resource considerations that must be weighed make exhaustive evaluation of all possible architectures extremely difficult. This paper presents an optimization model and a computationally efficient solution procedure to identify an estimated frontier of system configuration options which represent the best design choices for the user when there is uncertainty in the response time of the security force, once an intrusion has been detected. A representative example is described.

1 INTRODUCTION

The analysis of multi-layered security systems, such as those protecting high-value facilities or critical infrastructures, involves the evaluation of many dimensions of the design space. There are many options to be considered during the selection of a design configuration, such as available technology alternatives, locations in the system where that technology could be placed and technology-specific configuration settings, the spectrum of potential threats that need to be protected against, and budget or resource limitations. Exhaustive evaluation of all possible configurations in a large security system is generally impossible, so automated identification of alternatives is extremely useful. This paper extends previous work, described in Brown et al. (2015), which develops an attacker-defender framework operationalized via a genetic algorithm (GA) by including uncertainty in the response time of the protective force as well as merging the GA with a domain-specific greedy method to seed the initial population.

This research is related to research focused on attack graphs. An attack graph is a network modeling technique used to represent path selection for an intruder. The core idea is to represent each discrete action on the part of an intruder (which yields a change in system state) as an edge in a graph; hence, collections of actions can be identified via path finding analyses. A numeric score can be associated with each edge, where these scores can be the probability of success or some type of benefit to cost computation. Phillips and Swiler (1998) first suggested the concept of an attack graph. Since then there have been many extensions to this modeling paradigm including Chen et al. (2009), Ou et al. (2006), and Sheyner et al. (2002). Generally, the automatic identification of attack graphs has been found to be quite difficult. The method developed in this paper uses the underlying physical network directly so that an attack graph does not need to be explicitly identified.

This research is also related to the extensive literature on attacker-defender models. Conceptually, our model is similar to models suggested by Romero et al. (2012) and Brown et al. (2013) among others. Our model is similarly focused on investment planning to thwart attacks, however, our model of intruder behavior focuses on the probability of interruption whereas these models focus on satisfying demands for services. This model is also similar to that suggested in Reilly et al. (2012) in that both address how an attacker might assess opportunities. However, Reilly et al. (2012) focuses on repetitive attack processes and therefore what parts of a transportation network should be made inaccessible to shipments of specific types of hazardous materials. We focus on understanding which paths are relatively easier to breach as a mechanism to understand which technologies should be deployed where.

2 METHODS

2.1 Mathematical Formulation

We represent the problem as an attacker-defender model, in which the attacker has perfect knowledge of the security measures in place. The attacker's goal is to reach a specific target that is being 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 interruption of the adversary. The probability of interruption is defined as 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. Increasing the probability of interruption is accomplished by adding detection and delay security measures (such as cameras and fences). Each investment has an associated cost and NAR/FAR.

These core ideas yield the following optimization model. Suppose there is a single location that is assumed to be the origin for the attacker, as well as a different single location that contains the item of interest (a.k.a. the target). Further, suppose there is a network that connects these two locations, and that network is composed of directed arcs and nodes. At each of those links, technologies can be located that either affect the travel time on the arc or the detection probability or both. The goal of the optimization is then to suggest which technologies to place at which locations, so as to identify an acceptable trade-off between the probability that the intruder is interrupted, investment costs, and NAR/FAR. This optimization can be expressed as follows.

$$\max_{I_{ij}^{y}} \left[\min_{z^r} \sum_r g^r z^r\right], \min_{I_{ij}^{y}} \left[\sum_{ij} \sum_y A_{ij}^{y} I_{ij}^{y}\right], \min_{I_{ij}^{y}} \left[\sum_{ij} \sum_y c_{ij}^{y} I_{ij}^{y}\right] \quad (1)$$

Such that

$$\sum_r z^r = 1 \quad (2)$$

$$z^r \in \{0,1\} \ \forall r \quad (3)$$

$$I_{ij}^{y} \in \{0,1\} \ \forall (i,j), y \quad (4)$$

where r is an index for the paths that connect the intruder origin with the target of interest, g^r is the probability of interruption for path r, z^r is a binary variable that takes on a value of 1 if path r is selected and zero otherwise, y is an index on the combinations of investments that can be added to each link (hence forth referred to as package y), I_{ij}^{y} is a binary variable that takes on a value of 1 if investment package y is placed on arc (i,j) and 0 otherwise, A_{ij}^{y} is the NAR/FAR for investment package y on link (i,j) and c_{ij}^{y} is the cost of investment package y on link (i,j). It is important to notice that we consider packages of technologies on a link so as to properly consider the interaction between different

technologies. For example, when two technologies are placed on arc, the probability of detection is likely to be substantially lower than the sum of the probabilities.

Equation (1) gives the objectives for the system owner. The first term expresses the goal of the defender to select technologies that yield the highest probability of interruption of the intruder when the intruder selects the "weakest" path available. That is, for the intruder, the key decision is the path to select, given the investments in security measures the defender has made. This is what yields the max-min structure for the first goal. The second objective given in Equation (1) for the system defender is to minimize NAR/FAR and the third objective is to minimize costs. Taken together, these three terms define a trade-off space for the defender in the selection of security technologies. Equation (2) requires that the intruder select a single path. Equations (3) and (4) give the binary restrictions on the decision variables.

2.2 Solution Procedure

We use a genetic algorithm to solve the formulation given in equations (1)-(4). As in all genetic algorithms (Deb, 2012), there are three key steps that are iteratively employed: the steps of computing the performance of each member of the population, crossover and mutation. We create most of the initial population via Monte Carlo simulation and employ a greedy algorithm to generate the remainder. Using these solutions, an initial estimate of the frontier is created by cleaning and decimating these initial solutions. The cleaning process examines each solution (which gives a complete investment strategy) and attempts to remove as many investments as possible without decreasing the probability of interruption. This process can lead to solutions which are significantly less expensive. The decimation process examines each of the cleaned solutions and randomly removes single investments, each of which are added to the population of solutions (and the efficient frontier, if appropriate) as long as the newly created solution has a non-zero probability of interruption.

Each new pair of individuals is created via genetic crossover from two parents that are randomly selected from the solution pool, proportionally to their fitness. For crossover, we use a region-based crossover procedure similar to that described in Cohoon and Pairs (1987). This is a 2D crossover strategy where all technologies in a contiguous region of the network are swapped. The region swapped is randomized between pairs of parents. This structure better represents the topology in a solution and therefore is more likely to preserve effective pieces of intermediate solutions. The advantages of this type of encoding has been demonstrated on related topology problems including electromagnetic topology optimization and circuit design (e.g. Im et al. 2003, Bui and Moon 1995, and Moon et al. 1998) .

To estimate the fitness of an individual, we use an evolving estimate of the efficient frontier and compute the distance of that individual to this estimate using Manhattan distances, as suggested in Krause (1987).

For mutation, we use a variable-rate mutation strategy based on the level of homogeneity across the population of solutions. A related variable-rate mutation strategy employed in Brown et al. (2013) is used to counter the tendency for crossover to produce homogenous populations as described in Sait (1999). Mutations are based on the level of consistency amongst security investment strategies across the population of solutions. The percentages T_H and 100-T_H are defined to represent thresholds for which each security investment is considered to be used or not "most of the time", respectively, where T_H was selected to be 25%. A similarity ratio is used to determine the percent of investments that fall into either one of these categories for the given population and provides a measure of the consistency of security investment strategies across the population. As this ratio increases, the mutation rate increases up to a maximum value of 5%.

The stopping criteria uses an upper bound specified by a fixed number of crossover-mutation iterations. If, however, more than three crossover-mutation iterations occur consecutively where the Pareto frontier does not improve, then it is assumed that a steady-state frontier has been achieved and the solution procedure terminates.

The remainder of this section gives the computations used to compute the probability of interruption and the greedy algorithm used to seed the initial population.

2.3 Computing the Probability of Interruption

Suppose the response time for the protective force is known with certainty. The computation of g^r can be illustrated as follows. Consider the three link path given in Figure 1. The intruder must proceed from A, to B, to C and finally to D. Assume the response time for the protective force is six minutes. The intruder could be detected on the first, second, or third link, or not at all. If they are detected on the third link (i.e., from C to D), there is insufficient time for the protective force to respond. But, if they are caught on the first or second links, there is sufficient time to respond. Hence, the coefficient g for this route is 0.13 + (1-0.13)*0.45 which equals 0.5215, which is the probability that they are interrupted given the probabilities of the two possible interruption scenarios: they are detected on link A-B (0.13) or they are detected on link B-C (0.45) and not detected on link A-B (1-0.13).

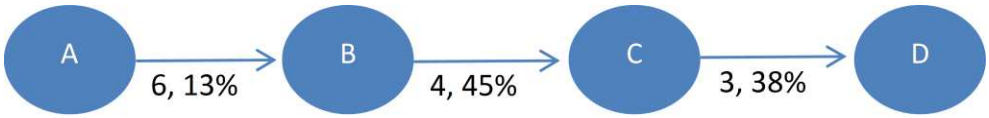

Figure 1: Example of computation g^r, with constant response force time.

This set of computations can be extended to networks via a label correcting algorithm (as given in Brown et al. 2015). Let $\{N, Z\}$ be the directed graph where N is the set of nodes and Z is the set of links. Let δ_u^- be the set of links $(i,u) \in Z$, T_{ij} be the travel time of link $(i,j) \in Z$ and D_{ij} be the detection probability on the link. We assume that the detection probabilities are independent. Also, let T be the required time to interrupt the intruder. Let a be the origin and b be the target. We define $I_i = (I_{i1}, I_{i2})$ for each $i \in N$, where I_{i1} represents the shortest travel time from node i to target b if the travel time is shorter than T, and I_{i2} represents the probability that the intruder is interrupted when they are in node i. The objective for the intruder is to minimize I_{a2}. The following algorithm is used to find the path that yields the lowest probability of interruption once detected.

1. Let $S \leftarrow b$, $I_b = (0,0)$, $C = \phi$, and $I_i = (inf, 1)$, $i \in N \backslash \{b\}$.
2. For $u \in S$, if $u = a$, then stop and report I_{a2}. Otherwise do the following:
    a. For each $v \in \delta_u^-$, If $I_{u1} + T_{vu} < I_{v1}$ and $I_{u1} + T_{vu} < T$, then let $I_{v1} = I_{u1} + T_{vu}$ and $I_{v2} = 0$.
    b. For each $v \in \delta_u^-$, if $I_{u1} + T_{vu} \geq T$ and $D_{vu} + (1 - D_{vu}) * I_{u2} < I_{v2}$, let $I_{v1} = I_{u1} + T_{vu}$ and $I_{v2} = D_{vu} + (1 - D_{vu}) * I_{u2}$.
    c. Find v^* that minimizes I_{v1} for $v \in N \backslash C \cup S$.
    d. If $I_{v^*1} < T$, then do the following:
        i. $S \leftarrow v^*$.
        ii. Go to step 3.
    e. If $I_{v^*1} \geq T$, do the following:
        i. Find v^{**} that minimizes I_{v2} for $v \in N \backslash C \cup S$.
        ii. $S \leftarrow v^{**}$.
        iii. Go to step 3.
3. Remove u from S, $C = C \cup \{u\}$ and go to step 2.

If we now consider an exponentially distributed force response time, the computation of g^r is somewhat more complicated and is illustrated as follows. Again, let =(N,A) be a directed graph with node set $N=\{1,2,\ldots n\}$ and arc set A. Associated with every arc is a two dimensional vector where the first

element in the travel time on the arc and the second is the probability of detection. We assume that these measures are non-negative. Let C(t,P_{OD}) be the label associated with a partial path from the origin o to the destination d of the partial path that is of length t and has a probability of interruption given detection of P_{OD}. For ease of discussion, assume that the nodes in the partial path are ordered sequentially from o to d with node o having a node number of one and the final node having a node number of one larger than the set of arcs. For simplicity, let's assume that this index is h with a maximum of H. The probability of interruption along a partial path is then as given below

$$D_{12}P\left(I|\sum_{ij}T_{ij}\right)+\sum_{i>1}^{H-1}D_{i,i+1}\,P\left(I|\sum_{j\geq i}^{H-1}T_{j,j+1}\right)\prod_{j<i}\left(1-D_{j,j+1}\right) \quad (5)$$

where $P\left(I|\sum_{ij}T_{ij}\right)$ is the probability of interruption when the accumulated travel time on the partial path is $\sum_{ij}T_{ij}$ and $D_{i,j}$ is the probability of detection on the link (i,j). Notice that this formula is equivalent to the following formula that can be applied recursively from the end of the path moving towards the origin, where we assume that we are currently at node j which is somewhere between nodes o and d.

$$D_{jk}P(I|\sum_{l\geq j}T_{l,l+1})+\left(1-D_{jk}\right)L_{kD} \quad (6)$$

where L_{kD} is accumulated probability of interruption given detection from node k, which is the successor node to j, to the end of the path.

We now demonstrate how an algorithm like that described above, which relies exclusively on the probability of interruption as the distance metric to conclude dominance in partial path labels, can fail.

Consider the five-node network given in Figure 2 which has two possible paths from A to D: A-B-C-D and A-B-C′-D. Assume the response time for the protective force is exponentially distributed with a mean of six minutes. Hence the probability that the response time will not exceed a given path travel time x, is $1-\exp\left(-\frac{x}{6}\right)$.

For path 1 (A-B-C-D), the intruder could be detected on the first, second or third link or not at all. If they are detected on link C-D, the probability they are interrupted is 39.35%, which is the probability that the response time will be no greater than 3 minutes. However there is only a 38% chance that they are detected on that link, which brings the probability of interruption on that link to 0.3935*0.38 which is about 15%. The probability of interruption on link B-C is calculated as (0.45*0.6886)+ (1-0.45)*(0.38)*(0.3935) which equals 0.3921, where the probability that the response time will be no greater than 7 minutes is 0.6886. Hence, the value of g for path 1 is then (0.13*0.8854)+(1-0.13)*[(0.45*0.6886)+ (1-0.45)*(0.38)*(0.3935)] which equals 0.4562. For path 2, the probability of interruption on link B-C′ is calculated as (0.2*0.9179)+(1-0.2)*(0.35)*(0.7364) which equals 0.3898. Notice that this value is less than the probability of interruption on link B-C, which would cause the path via node C to be pruned when using a Dijkstra-type algorithm. Hence the value of g for path 2 is then assigned to node A as (0.13*0.9698)+(1-0.13)*[(0.2*0.9179)+(1-0.2)*(0.35)*(0.7364)] which equals 0.4652. Notice that this value is greater than the value (0.4562) achieved by traversing path A-B-C-D, which is incorrect.

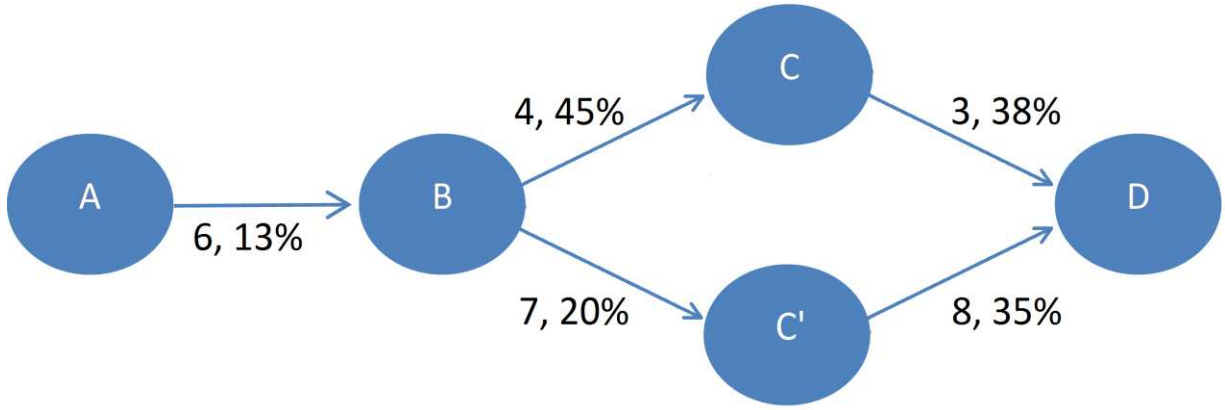

Figure 2: Example of computation g^r, with exponentially distributed response force time.

When the response time for the protective force is uncertain, path enumeration is required to identify the optimal path for the intruder. Rather than pursue this computational path, we randomly generate a collection of response times for the protective force from an exponential distribution but truncate the tail at the 95th percentile to avoid excessively long response times. For each solution (i.e. collection of investments across all arcs) we find the optimal path for the intruder assuming that the intruder has perfect knowledge as to what the response time will be. Notice that this is a conservative assumption because the intruder is assumed to possess this knowledge when they select their path.

2.4 Use of a Greedy Algorithm to Seed the Genetic Algorithm

In order to improve the quality of the initial population, a greedy algorithm is used to produce a single initial solution. That solution is then manipulated into a family of additional solutions which are combined with a population of randomly selected solutions. This initial collection of solutions is also used to normalize the different objectives so that the fitness function, which is the Euclidean distance from the Pareto Frontier to each new solution, is not heavily biased towards any single objective (due to the intrinsic differences in the magnitudes in which each objective is measured). This greedy algorithm adds investments by iteratively identifying the most desirable path to the intruder: the path that minimizes the probability of interruption given detection. It assumes there is a minimum desirable path length (which is a multiple of the response time) and a minimum desirable path detection rate (which is a plausibly high value). The greedy algorithm is described below.

First, the path that leads to the smallest probability of interruption given detection is identified. Investments are added to this path in order of largest incremental benefit with respect to path length until the path length goal is achieved. This incremental benefit is computed as the ratio of the change in path length to the investment cost. Once the path length is elongated sufficiently to meet the goal, the investments that focus on improving the probability of detection are added to the path based on their incremental benefit (where the incremental benefit is computed as the ratio of the improvement in the detection probability to the investment cost). Once the goal for the probability of detection is reached, a new path is identified and the process is repeated. This process concludes when there are no paths for which the length is shorter than the minimum desirable path length and there is no path with a detection probability that is less than the minimum desirable path detection rate.

At the conclusion of this algorithm, there are two final steps. First, any investment that does not impact the probability of interruption given detection is removed. This is accomplished by removing each investment in turn and computing the probability of interruption given detection. If this value does not decline, the investment is removed. Second, new solutions are built from this initial solution by randomly

removing investments. The solutions that stem from the greedy are only about 0.1% of the initial population but constitute 100% of the initial Pareto frontier (less the no-investment option).

3 REPRESENTATIVE EXAMPLE

3.1 Notional Security Investments

The following example uses notional data to demonstrate the application of the formulation and solution procedure. It assumes that an entirely new system is to be designed, rather than improving upon an existing design. The target is some high value asset that the intruder wants to access. There are two buildings on site, which are treated as barriers that the intruder cannot pass through. For the purposes of this analysis, there is a single intruder, and the force response time is exponentially distributed with a mean of 45 seconds. We randomly generate a collection of 50 response times for the protective force from an exponential distribution and use these values for each security investment strategy evaluation.

The security investments under consideration for this example include radar (R), fence (F), buried cable (BC), magnetic (Mag) sensor, microwave (Mic) sensor, and security camera (SC). The fence, buried cable, microwave sensor, and magnetic sensor investments can be applied on a per link basis. If radar is installed, it will cover a four link by four link (200 foot by 200 foot) area. Security cameras cover a two link by two link (100 foot by 100 foot) area directed away from the investment node location, to the north, northeast, east, or southeast of the node.

Table 1 lists the attributes of each of the possible investments. It shows the cost, NAR/FAR, probability of detection, and delay time expected for each of the investments. Only the fence is considered to be a barrier that can increase delay time on affected links. The other investments are sensors and only impact the probability of detection on the link. For the purpose of this analysis, we allow any combination of investment types to appear at the same location. When multiple detection technologies are added to a single link, it is assumed that the sensors are complementary and the total probability of detection is slightly increased. The composite probability of detection is set to (N-1)*3% higher than the maximum of the detection values, where N is the number of link sensors and the maximum probability of detection cannot exceed 99%.

Table 1: Investment options.

Investment	10-Year Cost (Thousands)	NAR/FAR	Probability of Detection on Affected Links	Delay (Seconds) on Affected Links
Radar (R)	500	2	0.75	-
Fence (F)	3	-	-	30
Buried Cable (BC)	200	4	0.9	-
Magnetic (Mag) Sensor	20	2	0.8	-
Microwave (Mic) Sensor	30	2	0.9	-
Security Camera (SC)	90	2	0.8	-

3.2 Results

This example demonstrates that by using the greedy algorithm to seed the GA versus strictly random initialization, a better Pareto Frontier (PF) is generated. When the initial population of 10,000 solutions is generated, an initial PF is also created to use for determining solution fitness. As can be seen in Table 2, the initial PF created via the greedy initialization is more sparse but has much lower cost solutions with lower NAR/FAR values. The highlighted solution (with 0.96 probability of interruption) is likely heavily

influenced by the initial greedy solution, which appears in the table as having a 0.99 probability of interruption. We shall focus on the highlighted solution to see how it is improved from the perspective of NAR/FAR and investment cost after being processed by the GA.

Table 2: Initial Pareto frontier for greedy versus random initialization.

Greedy Initialization				Random Initialization		
10-Year Cost (Thousands)	NAR/FAR	Probability of Interruption		10-Year Cost (Thousands)	NAR/FAR	Probability of Interruption
1110	172	0.189		27911	1804	0.4137
1338	186	0.192		29064	1670	0.4513
1370	188	0.256		26128	1806	0.4986
1382	188	0.301		25782	1902	0.5765
1644	196	0.304		26285	1956	0.5889
2146	240	0.336		28327	1800	0.8304
2149	240	0.36		27231	1810	0.8305
2337	262	0.684		29132	1758	0.8368
2346	262	0.738		27051	1822	0.8383
2349	262	0.81		29079	1788	0.839
2402	266	0.864		27573	1904	0.8391
2438	268	0.9408		27107	1990	0.8392
2441	268	0.96		27272	1966	0.8394
2471	270	0.99		27826	1866	0.8398
				27830	1888	0.84
				28998	1820	0.9278
				27596	1940	0.9304
				29186	2060	0.9336
				29782	1940	0.937
				29627	1986	0.9394
				31075	2032	0.9528
				31357	2030	0.9542
				29450	2096	0.9567
				36229	2480	0.9574

The PF is updated by executing ten crossover-mutation iterations of the genetic algorithm. In general, the PF generated by the greedy initialization performs better at the upper end with solutions that have lower cost and lower NAR/FAR but with higher probability of interruption. Since solutions with a minimum of 90% probability of interruption are likely to be required, this characteristic is more important than having better low-end solutions as generated by the random initialization technique. The highlighted solution from the initial PF table (Table 2) can be seen to have improved in the final PF table (Table 3) with values that are significantly lower for both the investment cost and NAR/FAR. The investment strategy associated with the highlighted solution appears in Figure 3.

Table 3: Final Pareto frontier for greedy versus random initialization.

Greedy Initialization			Random Initialization		
10-Year Cost (Thousands)	NAR/FAR	Probability of Interruption	10-Year Cost (Thousands)	NAR/FAR	Probability of Interruption
1050	148	0.18	90	16	0.064
1040	156	0.18	630	82	0.16
1073	158	0.189	510	102	0.192
1106	172	0.192	750	100	0.192
1115	172	0.256	820	114	0.272
1121	172	0.304	1481	158	0.7232
1411	180	0.336	1681	162	0.7291
1426	178	0.36	2218	262	0.8292
1432	178	0.384	4156	386	0.838
1435	178	0.464	3297	406	0.8392
1438	178	0.54	4383	356	0.8392
1462	180	0.684	3417	424	0.9244
1489	180	0.81	4800	410	0.9278
1581	224	0.83	4003	512	0.9321
1646	212	0.864	4713	476	0.9334
1781	228	0.96	4950	482	0.9444
2471	270	0.99	5757	488	0.9527
			5670	534	0.954
			5828	484	0.9565
			7178	630	0.9573
			35737	2464	0.9574

In Figure 3, the green links between nodes can be traversed by an intruder trying to reach target node 55. The fences are represented by the dark blue lines that are orthogonal to the node paths. The links are labeled with the assigned investments which can be any of those listed in Table 1. The resultant investment strategy is relatively symmetric, which is the ideal case since an intruder may attempt access from any of the perimeter nodes. The lowest probability of interruption path is indicated by the red links and follows the node path 6-17-28-37-46-55. Note that when using the sampling procedure, the paths from nodes 6 to 55, 60 to 55 and 106 to 55 all produce an equal probability of interruption as long as that value is less than the 95th percentile of the response time distribution (which is 135). This characteristic is due to the placement of the detection technologies at the perimeter of the system and the delay technologies in the interior, as specified by the suggested architecture, giving the protective force substantial time to respond.

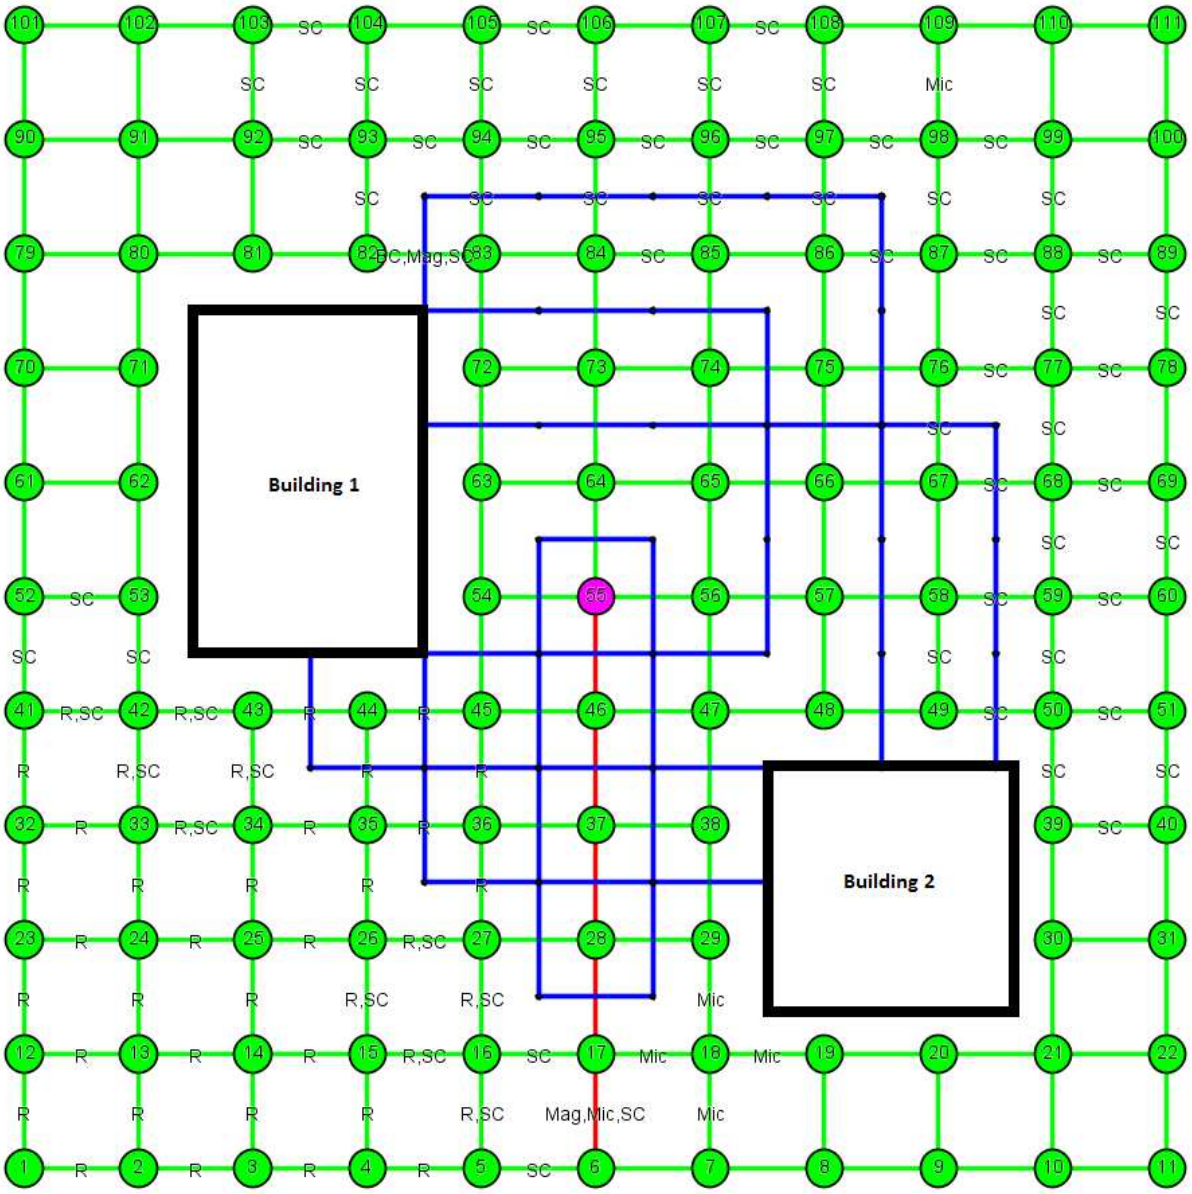

Figure 3: Investment strategy for 0.96 probability of interruption solution.

4 CONCLUSION

There are many dimensions in the design space of a security system, including technology options, alternative placement configurations, diverse threats, and budget considerations, making it effectively impossible to evaluate all potential system architectures. We have developed a game-theoretic model to optimize the design of security systems and demonstrated it on a realistic but notional problem instance. The model takes into account different configurations of both sensor technologies for detection and security barriers for delay, while also considering the uncertainty associated with response force times. The model also includes the ability to consider budget limitations and the impact of false alarms on system performance. In contrast to standard genetic algorithms, which typically use a collection of randomly generated solutions for the initial population, we used a greedy algorithm to form an initial

Pareto Frontier and augment the collection of random solutions. The example demonstrated that use of the greedy algorithm to seed the GA versus strictly random initialization resulted in generation of an improved Pareto Frontier of solutions.

There are at least three extensions to this model that would be useful. First, the capabilities, composition, and motivation of the attack force should have more variability. This model assumes a single attack force on a single path with a specific goal of reaching the target. It is likely useful to include some representation of the capability and composition of the attack force, such that some attackers are able to decrease the effectiveness of the security system via specialized capabilities or by launching coordinated attacks. Attacker motivation would also be useful to represent, since some attackers might want to escape the system after target acquisition, potentially changing their selected path. Second, investment in the defending response force is not considered by this model. The model could be extended to consider investments, such as an increased number of security response personnel, which would decrease force response time or ensure accurate assessment of alarms. Finally, the performance of many security technologies varies based on weather conditions or other environmental considerations. For this reason, some sensors may be complementary rather than substitutes (as they might appear in this analysis). One method to include this in the model would be to create a stochastic program with a collection of scenarios representing different weather and lighting conditions. For this approach, it may be important to maximize the minimum effectiveness of the system, where the effectiveness is evaluated across the different weather and lighting conditions.

ACKNOWLEDGMENTS

The authors would like to thank John L. Russell, Kristin Adair, and Alisa Bandlow 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 2015-2594 C.

REFERENCES

Brown, N. J. K., K. A. Jones, L. K. Nozick, N. Xu, A. Bandlow, K. A. Adair, J. L. Gearhart, and J. L. Russell. 2014. "Multi-layered Security System for Force Protection." MORS Journal (submitted).

Brown, N., J. Gearhart, D. Jones, L. Nozick, N. Romero, and N. Xu. 2013. "Multi-Objective Optimization for Bridge Retrofit to Address Earthquake Hazards." In Proceedings of the 2013 Winter Simulation Conference, edited by R. Pasupathy, S.-H. Kim, A. Tolk, R. Hill, and M. E. Kuhl, 2475-2486. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Bui, T. N., and R. Moon. 1995. "On Multi-Dimensional Encoding/Crossover." In Proceedings of the 6th International Conference on Genetic Algorithms, edited by L. J. Eshelman, 49-56. Pittsburg, PA: Morgan Kaufmann Publishers Inc.

Chen, F., J. Su, and Z. Yi. 2009. "A Scalable Approach to Full Attack Graph Generation." In Lecture Notes in Computer Science: Proceedings of the First International Symposium on Engineering Secure Software and Systems, 5429, 150-163.

Cohoon, J. P., and W. D. Pairs. 1987. "Genetic Placement." IEEE Transactions on Computer Aided Design, 6, 956-964.

Deb, K. 2012. "Advances in Evolutionary Multi-objective Optimization." In SSBSE 2012, edited by G. Fraser. Springer-Verlag, Berlin, Heidelberg.

Im, C., H. Jung, and Y. Kim. 2003. "Hybrid Genetic Algorithms for Electromagnetic Topology Optimization." IEEE Transactions on Magnetics, 39(5), 2163-2169.

Krause, Eugene F. 1987. Taxicab Geometry. New York, NY: Dover Publications, Inc. ISBN 0-486-25202-7.

Moon, B., R., Y. S. Lee, and C. K. Kim. 1998. "GEORG: VLSI Circuit Partitioner with a New Genetic Framework." Journal of Intelligent Manufacturing, 9, 401-412.

Ou, X., W. Boyer, and M. McQueen. 2006. "A Scalable Approach to Attack Graph Generation." In Proceedings of the 13th Annual ACM Conference on Computer and Communication Security, 336-345. New York, NY: Association for Computing Machinery Inc.

Philips, C. A., and L. P. Swiler. 1998. "A Graph-Based System for Network Vulnerability Analysis." In Proceedings of the 1998 New Security Paradigms Workshop, Association for Computing Machinery, 71-81. New York, NY: Association for Computing Machinery Inc.

Reilly, A., L. Nozick, N. Xu, and D. Jones. 2012. "Game Theory-based Identification of Facility Use Restrictions for the Movement of Hazardous Materials Under Terrorist Threat." Transportation Research Part E: Logistics and Transportation Review, 48(1), 115-131.

Romero, N., N. Xu, I. Dobon, and D. Jones. 2012. "Investment Planning for Electric Power Systems Under Terrorist Threat." IEEE Transactions on Power Systems, 27(1), 108-116.

Sait, S. M., and H. Youssef. 1999. Iterative Computer Algorithms with Applications in Engineering: Solving Combinatorial Optimization Problems. Los Alamitos, CA: IEEE Computer Society.

Sheyner, O., J. Haines, S. Jha, R. Lippmann, and J. M. Wing. 2002. "Automated Generation and Analysis of Attack Graphs." In Proceedings of the IEEE Computer Society Symposium on Research in Security and Privacy, 273-284.

AUTHOR BIOGRAPHIES

NATHANAEL J. K. BROWN is a Software/Electrical Engineer at Sandia National Laboratories in Albuquerque, New Mexico. He has a Masters in Electrical Engineering from Purdue University. The main focus of his career has been on the design and development of advanced software algorithms ranging from handwriting recognition to multi-objective optimization using heuristic techniques. His e-mail address is njbrown@sandia.gov.

KATHERINE A. JONES is an Operations Research Analyst at Sandia National Laboratories in Albuquerque, New Mexico. She has a Master's degree in Systems Engineering from the University of Virginia. Her work interests are in developing simulation and optimization models for complex systems. Her e-mail address is kajones@sandia.gov.

LINDA K. NOZICK is a Professor of Civil and Environmental Engineering at Cornell University. Her Ph.D. is in Systems Engineering from the University of Pennsylvania. Her research interests are in the modeling of complex systems with a particular focus on infrastructure systems and resiliency to natural and man-made hazards. Her email address in lkn3@cornell.edu.

NINGXIONG XU is a Research Associate in the School of Civil and Environmental Engineering at Cornell University. He earned a Ph.D. in Operations Research from Stanford University. His research interests are in large scale optimization under uncertainty and the application of those tools to transportation systems. His e-mail address is nx22@cornell.edu.