ARL

US Army Research Laboratory

Modeling Cyber-Physical War-Gaming

by Edward J M Colbert, Alexander Kott, Lawrence P Knachel III, and Daniel T Sullivan

NOTICES

Disclaimers

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof.

Destroy this report when it is no longer needed. Do not return it to the originator.

ARL

US Army Research Laboratory

Modeling Cyber-Physical War-Gaming

by Edward J M Colbert and Alexander Kott
Computational and Information Sciences Directorate, ARL

Lawrence P Knachel III
ICF International, Columbia, MD

Daniel T Sullivan
Raytheon Company, Dulles, VA

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED (From - To)
August 2017	Technical Report	June 2016–June 2017

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
Modeling Cyber-Physical War-Gaming	
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)	5d. PROJECT NUMBER
Edward J M Colbert, Alexander Kott, Lawrence P Knachel III, and Daniel T Sullivan	
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	8. PERFORMING ORGANIZATION REPORT NUMBER
US Army Research Laboratory ATTN: RDRL-CIN-S Aberdeen Proving Ground, MD 21005-5067	ARL-TR-8079

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT

War games are simulations of what could happen in the real world. War games can serve multiple functions, such as training, testing a new system, testing existing processes, or discovering new knowledge. However, theoretic foundations and guidance for cyber war games are lacking. Here we illustrate how operational military war games and cyber war games share similar constructs. We also provide a game-theoretic approach to mathematically analyze attacker and defender strategies in cyber war games. Using a fairly realistic simulation, we empirically demonstrate applying game-theory models to quantify risks and benefits when assessing strategies.

15. SUBJECT TERMS

SCADA, Supervisory Control and Data Acquisition, ICS, industrial control system, CPS, cyber-physical system, cybersecurity, simulation, war game

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE	SAR	46	Edward Colbert
Unclassified	Unclassified	Unclassified			19b. TELEPHONE NUMBER (Include area code) 301-394-1674

Standard Form 298 (Rev. 8/98)
Prescribed by ANSI Std. Z39.18

Contents

List of Figures

List of Tables

Acknowledgments

We appreciate Dr Brian Rivera, Mr Jerry Clarke, and Mr Curtis Arnold for supporting cyber-physical system research at the US Army Research Laboratory. We are also grateful to the Red and Blue cell players who contributed to the tabletop war game for their enthusiastic participation.

1. Introduction

Cyber war-gaming is often used by commercial and public-sector organizations. Such war games typically involve employees of an organization who play roles in a human-based (and occasionally computer-based) simulation of a cyber attack and respond to the attack. A number of consulting organizations provide war-gaming design and facilitation as a service.[1] The growing prominence of cyber war games is hardly surprising. Cyber conflicts involve problems of an adversarial nature, not unlike those in military practice, and these are often solved by resorting to game-theoretic models, or by simulations.

However, theoretical foundations and guidance for cyber war-gaming are lacking. In this report, we offer a game-theoretic model of cyber attack and defense, compare the model-based analysis with experiences of an actual tabletop war game, and offer recommendations on using game-theoretic analysis for enhancing accuracy and value of such cyber war games.

1.1 Current Practices of Cyber War-Gaming

In spite of its growing popularity among corporate and government organizations (e.g., Casey and Willis[2]), the term "cyber war game" is not particularly well defined and may refer to many different forms of an exercise, test, simulation, or emulation event. Typically, unlike penetration testing in which "white hat" hackers seek to find the company's technical vulnerabilities, a corporate cyber war game often emphasizes a business scenario involving a cross-section of the company's business functions.[3,4,5] Modern cyber war-gaming in corporate and government scenarios[6,7] use the extensive wisdom and experience gained from military war-gaming.[8–10]

The war game is structured—often by a specialized consulting organization hired for this purpose (e.g., Hale[1])—to simulate experiences of a real cyber attack and realistic responses to it. Participants of the war game often comprise the company's employees from multiple functional areas: information security, application development, network operations, facilities management, customer service, production, marketing, legal and public affairs, financial, and distribution, for example. These players gather in one or several conference rooms for a duration of anywhere from 4 h to 3 days, and under the guidance of professional facilitators, proceed to enact the events of a cyber attack, usually developing a strong commitment and passion in the game.

Great diversity in the types and forms of cyber war games is found in current cyber exercises. Such diversity can be characterized along several dimensions:

- Breadth of business functions: The focus of a war game may range from strictly technical considerations of vulnerabilities, capabilities, and software and hardware activities to a broad coverage of business functions (e.g., financial, media, legal, and business operations aspects), where the technical cyber compromise is merely a starting point of the scenario. This may correlate with the seniority of decision-makers involved—broader scenarios may involve leaders higher on the corporate ladder.

- Scale of an entity under consideration: A war game may concern itself with an entity limited to a single web server to large-scale operations or a multinational corporation; it could be a single system or network, a site or an enterprise, or an international system of enterprises.

- Realism of the game: A war game may range from a tabletop exercise with little more than paper and pencil to computer-assisted simulations to the use of emulated cyber ranges, and even to attacks on operational systems.

The process of designing, implementing, and executing a cyber game normally involves most of the following steps (not necessarily sequential or in this order). They could be performed by a consulting organization in close collaboration with company personnel (e.g., Casey and Willis[2]):

- Defenses: Identify security mechanisms, tools, and personnel; their attributes and capabilities.

- Threats: Hypothesize suitable threats; their capabilities and likely tactics, techniques, and procedures; goals; limitations; access opportunities; skill levels; time; and resources available.

- Attacks: Formulate and select at least the following 2 attack scenarios: 1) most likely to be executed by the threat against this organization and 2) most dangerous—the one that may be less likely but would cause the greatest damage.

- Players: Recruit relevant participants/defenders of the company—who are relevant to the site or enterprise being war-gamed and who would realistically be engaged in defending against a given a threat scenario.

- Blue cell(s): Organize the participants/defenders into a team or teams responsible for planning and executing defensive actions; such a team is called a Blue cell. Preferably, participants are organized into teams (cells) that are reflective of the actual organizational structure of the company.

- White cell: Provide individuals who have experience in war games, who can document important information emerging during the war game, guide and facilitate the game, and adjudicate or arbitrate outcomes of individual actions taken by participants. This team is commonly called the White cell and is usually an outside consultant.

- Red cell: Provide or designate a team of individuals who play the role of the attacker; such a team is called the Red cell.

- Rooms and props: Prepare physical facilities, means of communication, and paper or computer-based products to conduct the game.

- Play: Assemble all cells and begin execution of the scenario; let the Red cell attack, Blue cell defend, and White cell declare the status as it evolves; and inject additional events to keep the game moving in the right direction.

- Payoff: Analyze the observations and results of the game, and formulate recommendations.

The outcomes and benefits of the cyber war game vary depending on the goals of the organization. These goals may include identifying hidden vulnerabilities, disproving incorrect assumptions, and offering additional procedures and training. A cyber war game may produce the following benefits:

- Identify poorly understood risks.

- Educate and entertain personnel.

- Obtain support of senior decision-makers.

- Explore the extent of potential disruption(s) to various business functions.

- Clarify the roles and responsibilities of cyber responders.

- Improve communications.

- Allow stakeholders to get to know one another and build relationships.

- Understand decision-making authorities.

- Highlight interactions with third-party business partners.

- Identify potential gaps in an organization's preparedness and response plans.

1.2 Military Practice of Course-of-Action Analysis

Much of the techniques used in cyber war-gaming appear to be influenced or directly borrowed from military war-gaming practices. The US Naval War College has provided extensive expertise and documented guidance for military war-gaming.[9,10] In the US Army, war-gaming practice is often called Course-of-Action (COA) Analysis.[11,12] This analysis is performed by military units of various sizes, from a small unit called a squad (4 to 10 Soldiers) to a very large organization called a corps (20,000+ Soldiers); we will use a unit called a brigade (1,500 to 3,200 Soldiers).

Somewhat comparable to a midsize corporation, a US Army brigade includes several thousand professional Soldiers and officers, hundreds of combat and support vehicles, helicopters, sophisticated intelligence and communication equipment and specialists, artillery and missiles, engineers, medical units, and repair shops. In a battle, these assets might perform hundreds of complex tasks (similar to corporate "business functions"): collecting intelligence; movements; direct and indirect fires; constructing roads, bridges, and obstacles; transporting and handling supplies; managing the civilian population; command and control, and so on. Unlike in cyber war-gaming, the threat (i.e., the enemy that the brigade fights against) tries to apply physical destruction to the brigade, although cyber attacks are often also a part of the threat's repertoire.

Detailed planning of a military operation requires an intensive effort of highly trained professionals, called the brigade planning staff. Typically, the group consists of 4 or 5 officers, ranging in rank from captain to lieutenant colonel, who perform this work with the support of a subordinate staff. The process normally takes from 2 to 8 h—not unlike a typical cyber war game. The physical environment often consists of a tent extended from the back of one or several Army trucks or armored command-and-control vehicles; folding tables and chairs, and—similar to a cyber war game—either computer screens or paper maps on which the officers draw unit symbols with their movement arrows.

The input for the staff's effort comes usually from the unit commander as a high-level specification of the operation. With this input, the planning staff works as a team, called the Blue cell, to perform actual war-gaming, which includes the following:

- Predicting enemy actions or reactions. This is done by the Red cell, usually the officer who specializes in collecting and analyzing enemy intelligence.[13] The Red cell plays the role of the enemy to help the Blue cell understand possible actions and responses of the enemy. Similarly to the cyber war

games, the Red cell provides 2 cases of enemy actions: the most likely plan of enemy actions and the most dangerous (to the brigade) plan of enemy actions. The latter might be the same as the former but usually is different as it involves assumptions of greater capabilities on the part of the enemy.

- Planning and scheduling the detailed tasks required to accomplish the specified COA,[14] preventing or responding to the threat actions (like the Blue-cell defenders would do in a cyber war game), and allocating tasks to the diverse forces constituting the brigade (like elements of a corporate response to cyber attack).

- Estimating success or failure of friendly and enemy actions, and battle losses.[15] This is similar to the function performed by the White cell in cyber war-gaming.

The process of estimating enemy actions and friendly actions may repeat in several cycles until a convergence is achieved: the Red cell is unable to suggest any further improvements of the enemy actions, and the Blue cell is unable to suggest any further improvements of the friendly actions. This hints at reaching something akin to Nash equilibrium in game theoretic terms.

This war-gaming usually produces a plan/schedule in a synchronization matrix format, a type of Gantt chart (Fig. 1). The chart's columns represent time periods. The rows contain functional classes of actions, such as Maneuver (which in turn includes such subclasses as Main Effort and Security), Combat Service Support (e.g., logistics), Military Intelligence, and so on. This plan-schedule's content, recorded largely in the matrix cells, includes the tasks and actions of the friendly force's subunits and assets, their objectives and manner of execution, expected timing, dependencies and synchronization, routes and locations, availability of supplies, combat losses, enemy situation and actions, and so on.

	H+ 0:00	H+ 0:15	H+ 0:30	H+ 0:45	H+ 1:00	H+ 1:15	H+ 1:30	H+ 1:45	H+ 2:00
Security	Guard Advance	Unit 1 Tactical March		PL PL Atlanta	Assailable Flank				
	Unit 0 Unopposed Advance	Unit 1 Cross PL PL Blue			Unit 1 Move to Support Need				
	Unit 1 Tactical March				Unit 0 Maintain Contact with the Enemy				
					Unit 0 Isolate Enemy Force				
					Unit 1 Attack Enemy Alpha				
Main Effort	Unit 2 Movement to Contact Unit K	Unit 3 Tactical March	Unit 3 Tactical March	Unit 3 Tactical March	Unit 2 Movement to Contact Tactical March	Unit 3 Tactical March	Unit 3 Tactical March	Unit 3 Tactical March	Unit 3 Unopposed Advance
	Unit 2 Movement to Contact Tactical March	Unit 3 Uncoil	Unit 3 Cross PL PL Yellow	Unit 3 Cross PL PL Blue	Unit 3 Tactical March	Unit 3 Cross PL PL Boston	Unit 3 Cross PL PL Baltimore	Unit 3 Cross PL PL New York	Unit 3 Tactical March
	Unit 2 Main Force Advance		Unit 2 Tactical March		Unit 3 Cross PL PL Atlanta	Unit 2 Transition to Hasty Attack	Unit 2 Tactical March	Unit 2 Tactical March	
	Unit 2 Tactical March		Unit 2 Cross PL PL Yellow			Unit 2 Maneuver as Needed to Reach	Unit 2 Cross PL PL Blue	Unit 2 Cross PL PL Atlanta	

Fig. 1   Fragment of a war-gaming output in a synchronization matrix format. The horizontal axis is time, and the vertical axis describes specific war-game functions (adapted from Rasch et al.[14]).

Ultimately, the purpose of the military war game is for the Blue cell to consider and select a small (manageable, often on the order of 3) number of COAs that are seen as most advantageous to the Blue side. In doing so, the Blue cell has to make an assumption about the COA that would be adopted by the Red side. In military war-gaming, there are several ways to approach this difficult decision.

One approach is to consider the "most likely" COA of the opponent—that is, the Red COA that the Blue cell feels is most likely to be adopted by the Red side. This assessment of likelihood might be based on the Blue cell's knowledge of the Red side's preferences (e.g., the COAs that the Red side has adopted in previous battles). Alternatively, the Blue cell might decide that the most likely Red COA is the one that provides the Red side with the greatest advantage or greatest utility in the battle.

Another approach might be to consider the "most dangerous" or "most damaging" Red COA: the COA that would cause the greatest damage to the Blue side. Note that the "most likely" and "most damaging" Red COAs are often different.

Having selected the most likely Red COA and the most damaging Red COA, the Blue cell usually attempts to select a Blue COA that would perform sufficiently well against both of the Red COAs.

Besides creating this tangible set of potential strategies, other valuable outcomes of this war-gaming are similar to those of corporate cyber war games: identification

of hidden vulnerabilities, incorrect assumptions, risks and losses, education, and clarity of roles and responsibilities.

The remainder of this report is outlined as follows. We first propose a game-theoretic model of a contest between cyber attacker and cyber defender. Then we describe an actual cyber war game designed and led by one of this report's authors. We explore how our theoretical model may apply to the actual war game. Finally, we discuss the insights and benefits that the model brings to the cyber war game and offer a set of practical recommendations for designing and conducting cyber war games.

2. Game-Theoretic Method

In this section, we formulate an approach to analyze a cyber war game scenario that combines elements of game theoretic and risk analytic treatments. In a later section, we discuss an example of a real-world process where an approximation of this approach is used.

2.1 Mathematical Model

In the following formulation, we are partly inspired by Hausken.[16] Consider the cyber-physical system (CPS) depicted in Fig. 2.
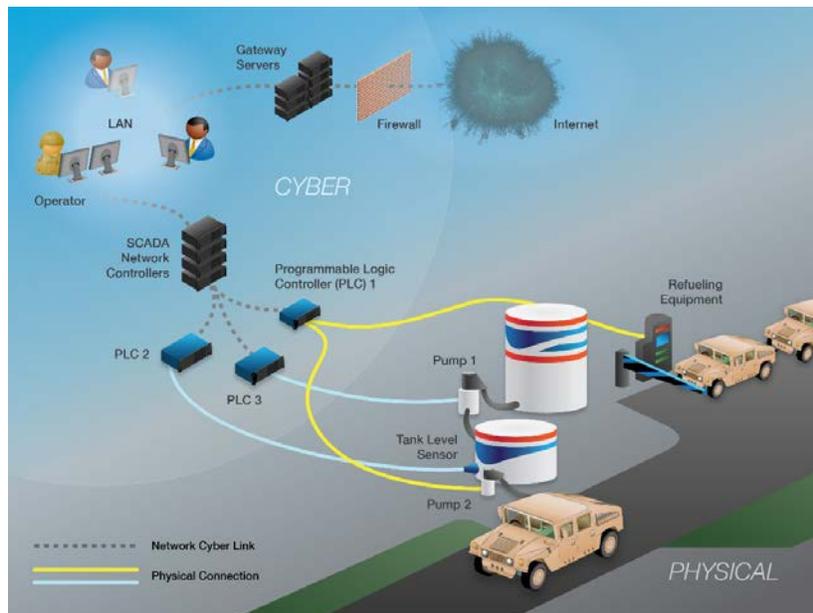

Fig. 2   Sample CPS connected to the Internet. Programmable logic controllers (PLCs) control fuel tanks and refueling equipment.

For this system, a cyber attacker desires to obtain a benefit b by accessing the system via the Internet and eventually obtaining control of the fuel system's programmable logic controllers (PLCs). In doing so, the attacker would have to penetrate defensive mechanisms and actions of the defenders in several layers of the CPS.

To make the discussion more general, in Fig. 3 we replace the details of Fig. 2 with an abstract scheme where the attacker enters the attack surface and penetrates a series of layers guarded by the defender before arriving at the target. The firewall in Fig. 2 is represented as Layer 1 in Fig. 3. Applying security best practices to securing the gateway servers in Fig. 2 is depicted as Layer 2 in Fig 3. Using access control lists (ACLs) on the Supervisory Control and Data Acquisition (SCADA) network controllers in Fig. 2 is illustrated as Layer 3 in Fig. 3.
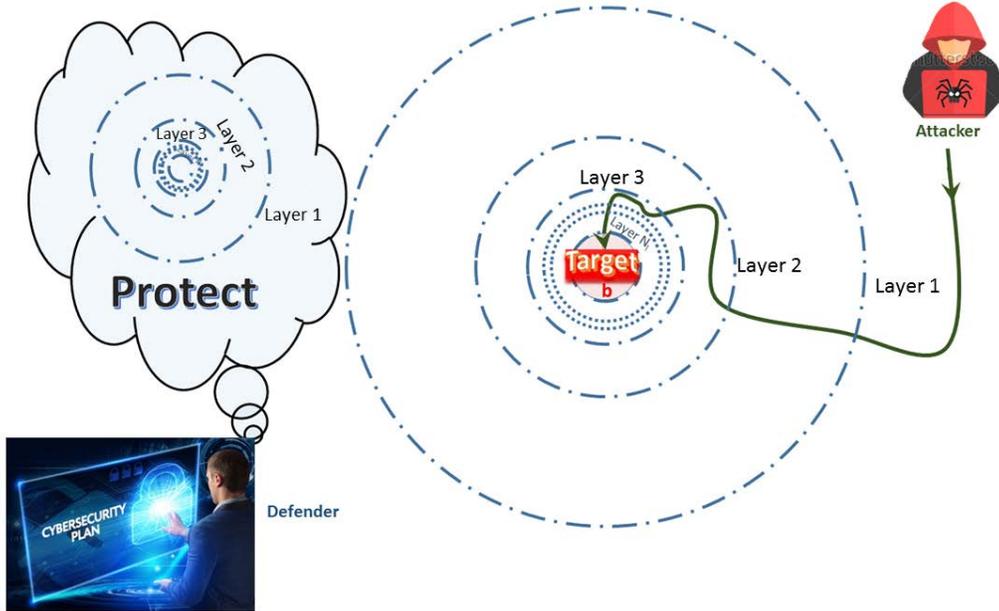
Fig. 3   Abstract illustration of an attacker's path through N_l cyber layers surrounding a central target, which is protected by a defender. The attacker receives benefit b after penetrating the final layer.

In this scheme, the attacker has devised a set of N_a strategies $\{s_{a,i}\} \in S_a$, where S_a is the attack-strategy space. The attack strategies are identified by index i. Likewise, the defender has developed a set of N_d strategies $\{s_{d,j}\} \in S_d$, where S_d is the defense strategy space and j is the defense strategy index. We describe a simplified model in which both attacker and defender have complete knowledge of the system and can therefore determine each other's strategies. The defender's strategies are accomplished by selecting specific subsets of cyber-defense mitigations $\{m_{d,k}\} \in M_d$, where M_d is the set of all mitigations, and k is the mitigation index.

As shown in Fig. 3, there are N_l layers that the attacker needs to penetrate. We identify these layers by the index l.

There are costs for both the attacker and the defender for each specific strategy tuple $\{i, j\}$. That is, given an attack strategy i and a defense strategy j, the attacker suffers a cost $C_{a,ij}$ to accomplish his goal, and the defender spends a cost $C_{d,ij}$ to deploy his defense strategy.

Finally, given a strategy choice $\{i, j\}$, the attacker is assumed to penetrate layer l with probability $p_l(s_{a,i}, s_{d,j})$, or $p_{l,ij}$ in shorthand notation. As shown in Fig. 3, if the attacker penetrates all N_l layers, he reaches the target T and obtains a benefit b.

In the mathematical model below (Eqs. 1–3), we consider the scenario of when the attacker gains benefit b, the defender loses an equivalent value (b) of his assets. This constraint could easily be modified as the situation demands.

The attacker and defender expected utility u_a depends on the benefit b, the total probability of penetrating all layers $P_{ij}^T = \prod_{l=1}^{N_l} p_{l,ij}$, and expended costs $C_{a,ij}$. For the attacker:

$$u_a(s_{a,i}, s_{d,j}) = b \prod_{l=1}^{N_l} p_l(s_{a,i}, s_{d,j}) - C_a(s_{a,i}, s_{d,j}), \qquad (1a)$$

or, in shorthand notation

$$u_{a,ij} = b\, P_{ij}^T - C_{a,ij}. \qquad (1b)$$

Likewise, for the defender

$$u_d(s_{a,i}, s_{d,j}) = b[1 - \prod_{l=1}^{N_l} p_l(s_{a,i}, s_{d,j})] - C_d(s_{a,i}, s_{d,j}), \qquad (2a)$$

or, in shorthand notation

$$u_{d,ij} = b\,(1 - P_{ij}^T) - C_{d,ij}. \qquad (2b)$$

The challenge for each player is to select the strategy, s_a^* and s_d^* for attacker and defender, respectively. As will be described in Section 2.2, there are a number of methods the attacker and defender may use to select their strategies s_a^* and s_d^*. Upon selection of the player strategies s_a^* and s_d^*, the total probability of penetration can be computed:

$$P_T^* = \prod_{l=1}^{N_l} p_l(s_a^*, s_d^*). \qquad (3)$$

2.2 Strategy Selection

Once the attacker and defender costs and penetration probabilities are known or assumed, payoff matrices $u_{a,ij}$ and $u_{d,ij}$ can be computed and inspected by each player to determine their preferred strategy.

2.2.1 Pure Strategy Equilibria

We first describe the case in which the players choose a single unique strategy (pure strategy approach) as opposed to estimating a probabilistically weighted set of multiple strategies (a.k.a., mixed strategy approach).

The model outlined in Section 2.1 does not describe a zero-sum game (defined as $u_{a,ij} = -u_{d,ij}$), and a Nash equilibrium state for s_a^* and s_d^* may not exist. However, since the number of strategies considered will be low enough to be manageable by a human, manual methods can be used to search for a local saddle point or a Nash equilibrium.

That is, if we search over each defense strategy j for the preferred attack strategies $\{i_j\}$ of the attacker

$$\{i_j\} = s_{a,i}^{max,j} = \underset{s_{d,j} \in S_d}{\operatorname{argmax}} u_a(s_{a,i}, s_{d,j}), \tag{4a}$$

and, likewise, search over each attack strategy i for the preferred defense strategies $\{j_i\}$ of the defender

$$\{j_i\} = s_{d,j}^{max,i} = \underset{s_{d,j} \in S_d}{\operatorname{argmax}} u_d(s_{a,i}, s_{d,j}), \tag{4b}$$

and find saddle points in the payoff matrices, then one or more "most likely" equilibrium strategies may be found. If only one equilibrium point is found, then it is by definition a Nash equilibrium. We illustrate this method further in Section 3.4.

Equilibrium strategies such as these may be preferred by the players in the case when the payoff matrices are fully disclosed to both players, and both players must choose their strategy simultaneously.

2.2.2 Strong Stackelberg Strategy Equilibria

A common method for finding solutions to nonzero-sum games is to assume that one of the players has the opportunity to be first in selecting his strategy; that player chooses a mixed strategy solution, and the opponent follows with a pure strategy (i.e., a single strategy with 100% probability). The equilibrium strategies in this scenario are known as Strong Stackelberg Equilibria (SSE).[17,18] If the war game is

played in this manner, the leader and the follower could use one of the many methods for finding SSE solutions, such as those described in Korzhyk et al.[19] While we do not discuss SSE equilibria further in this work, we later consider an example where multiple strategies are considered to some extent.

2.2.3 Playing Against the Most Likely Strategy of the Opponent

As described in Section 1.2, a player may choose a strategy by first estimating the most likely strategy of the opponent based on any available information about the opponent. Then the player selects his own strategy based on how successful it will be against the most likely strategy of the opponent.

2.2.4 Playing Against the Most Damaging Strategy of the Opponent

Alternatively, as described in Section 1.2, a player may choose a strategy by first estimating the most damaging strategy of the opponent (i.e., the strategy in which the opponent would impose the most severe losses on the player). Then the player selects his own strategy based on how successful it will be against the most damaging strategy of the opponent.

2.3 Sample Calculation

To illustrate, we provide sample calculations for a simple scenario. Suppose a freelancing cyber-crime group is engaged by an anonymous third party to penetrate controls of a munitions plant. Two layers of the network need to be defeated: $l = 1$ and $l = 2$. If successful, the attacker will be paid a benefit b of \$50,000. The defender of the plant consistently uses a single strategy $s_d^* = s_{d,1}$.

Based on the preliminary reconnaissance of the plant's network, the attacker considers 2 possible strategies. The first strategy $s_{a,1}$ will use an available exploit that will rapidly defeat the defenses of Layer 1 with probability $p_{1,11} = 0.9$. However, the exploit is noisy and will likely alert the defenders, thereby reducing the probability of penetrating Layer 2 to $p_{2,11}$ to 0.1. The cost associated with this "fast and noisy" strategy $s_{a,1}$ is only $C_{a,11} = \$5,000$. The second strategy $s_{a,2}$ will require the development of a new, stealthy exploit for Layer 1 and will take longer to deploy. In this case, $p_{1,21} = 0.9$, $p_{2,21} = 0.8$, and the costs to the attacker are much higher, $C_{a,21} = \$15,000$.

Using Eq. 1a, we find that the expected attacker utilities for the first and second attacks are, respectively,

$$u_{a,11} = \$50,000\,(0.9 \cdot 0.1) - \$5,000 = -\$500$$
$$u_{a,21} = \$50,000\,(0.9 \cdot 0.8) - \$15,000 = \$21,000.$$

Clearly, the second attack has higher utility, and therefore the attacker selects $s_a^* = s_{a,2}$. The attacker's probability of success $P_T^* = P_{21}^T = 0.9 \cdot 0.8 = 0.72$.

2.4 Practical Considerations

To generalize, the overall process for calculating the game-theoretic quantities in Section 2.1 is as follows:

1. Collect information about S_a, S_d, $p_l(s_{a,i}, s_{d,j})$, $C_a(s_{a,i}, s_{d,j})$ and $C_d(s_{a,i}, s_{d,j})$ from empirical and experimental sources.

2. Compute cost-utility functions (payoff matrices) and total penetration probability matrix using Eqs. 1, 2, and 3.

3. Utilize the payoff matrices to determine the best strategy (see Section 2.2).

The most difficult step may be the first. Quantitative data, such as $p_{1,11} = 0.9$ and $p_{2,11} = 0.8$ in our illustrative example, come from subject matter experts and may be difficult to estimate. Obtaining such information, particularly from subject matter expects, may be a difficult and expensive process that can yield subjective, inconsistent, and/or unreliable results.

2.5 Relation to Military War-Gaming Practice

As described in Section 1.2, the normal process of military war-gaming (COA analysis) is as follows. Having collected and considered the relevant information, the Blue cell officers propose a friendly defense strategy $s_{d,1}^H$, where the superscript H indicates a strategy devised by a human in the military war game. The strategy $s_{d,1}^H$ denotes a sequence of activities and associated resources, time periods, and spatial locations. This strategy is recorded in the synchronization matrix (e.g., Fig. 1). Then the Red cell who plays the role of the enemy proposes enemy strategy $s_{a,1}^H$. Given $s_{d,1}^H$, the Blue cell produces a modified strategy $s_{d,2}^H$, and so on. In each iteration, and for each activity, the planning team uses its experience and doctrinal guides to determine whether the activity will succeed and the losses each side will encounter as the result of the activity. These iterations continue until equilibrium is reached, where the Red cell is unable to suggest any further improvements to their attack strategy, and the Blue cell is unable to suggest any further improvements to their defense strategy, or both cells are exhausted in an unsuccessful attempt to find such an equilibrium. Implicitly or explicitly, the losses are expected to be kept below some maximum allowable value. In the human war game activity, risk is often assessed qualitatively by the players by estimating the likelihood of accomplishing the military objective of the battle (e.g., capturing certain terrain or destroying the enemy's forces).

Our game-theoretic model and strategy selection methods offer an analytical tool for the Blue and Red cell to decide on their strategies $s_{d,1}^H$ and $s_{a,1}^H$. The payoff and penetration matrices can be used as general tools for considering all strategies as long as no targets in the layers have been compromised and the costs are not changed. While we do not offer a mathematical framework for determining strategies in the war game where targets in the layers have been compromised and costs need to be modified dynamically, such a framework could be constructed using the same general methodology.

3. Experimental Investigation of Cyber War-Gaming

In this section, we explore an empirical example of calculating the cost, utility, and probability of success for different strategies of the attacker and defender in a war game.

3.1 General

As a more elaborate example, we consider the application of our game-theoretical framework to a tabletop war-gaming activity conducted at the US Army Research Laboratory.[20] In this event, a fictitious AQUA CPS was designed and presented to 2 teams of human cyber experts—a Red cell and a Blue cell—representing the attacker and defender, respectively.

3.2 AQUA War Game Information Packet

A technical information packet[21] describing the AQUA system was provided to both teams before the exercise began. Highlights from the information packet follow.

The AQUA is a food-processing plant that produces packaged meals. The process map for our fictitious AQUA plant is shown in Fig. 4.

The plant executes 6 manufacturing processes. The meat and vegetables are cooked separately. Once they are cooked, the meals are prepared and packaged in a material suitable for high-pressure processing. Once the high-pressure process is completed, the meals are placed in boxes and stored in a warehouse.
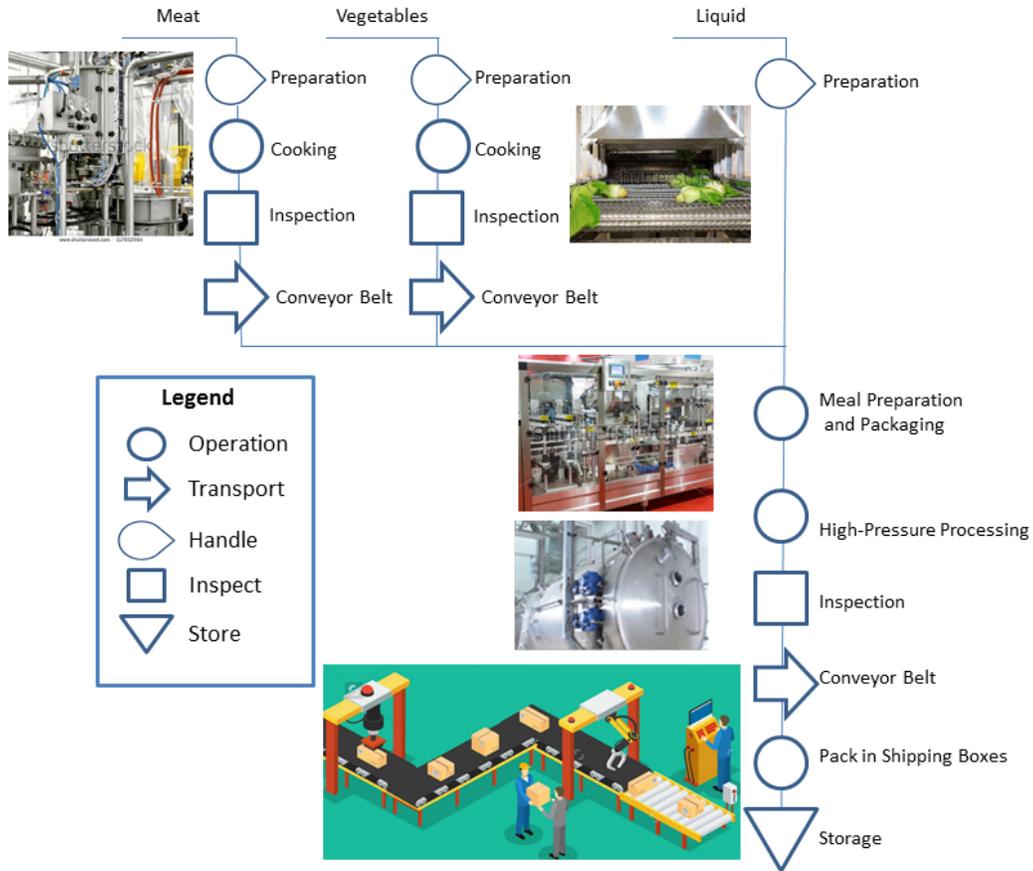
Fig. 4   Process map for the production line in the fictitious AQUA plant

The plant network used by AQUA is shown in Fig. 5. It consists of 6 PLCs, several workstations, a closed circuit television (CCTV) system, and a wireless network for tablet computers. Technicians use the tablet computers to access the human–machine interface (HMI) displays. The plant network is not connected to the corporate network or the Internet. All plant machinery is hard-wired to the input and output modules of the PLCs. The CCTV cameras are hard-wired to the digital video recorder.
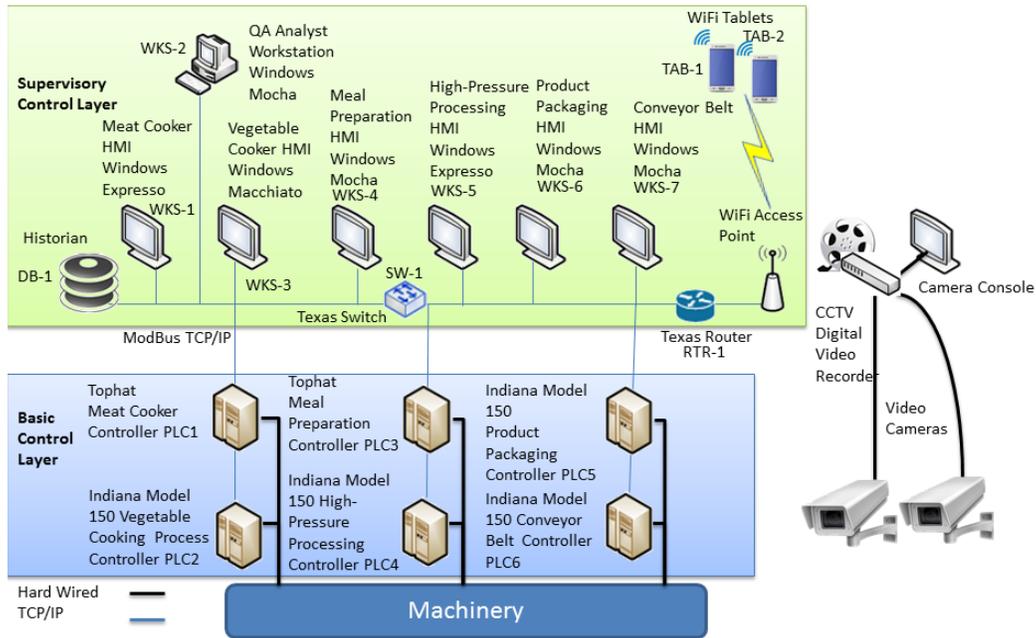
Fig. 5 **Plant network used to control AQUA production; wire fidelity (WiFi) incorporated**

3.3 Cyber-War-Game Method and Execution

Before the war game began, in addition to reviewing the information packet read-ahead document, both teams were allowed to ask the game facilitators specific technical questions. Answers were shared with both teams. During the exercise, the Red and Blue cells met separately to review all of the technical information and discuss strategies. Each cell documented its potential strategies. After this, the cells converged to share their findings. The Red cell shared its attack strategies with the Blue cell. Mitigations, counter-attacks, and counter-mitigations were discussed. Details (see Colbert et al.[20]) about the Red-cell strategies and Blue-cell mitigations from the table-top war game are presented in the Appendix. The representation of those results in terms of our game-theoretic framework (Section 2.1) is described in the following subsections.

3.4 Description of Game-Theoretic Variables

In this section, we illustrate how to derive the game-theoretic computations from the Red- and Blue-cell strategies.

3.4.1 Attack and Defense Strategies

The Red cell submitted 5 attack strategies $\{s_{a,i}\}$, $i = 1 \dots 5$, listed in Table 1.

Table 1 Red-cell attack strategies

Attack strategy index {i}	Attack strategy ($s_{a,i}$) description
1	Layer-2 attack
2	Subversion and Espionage Directed against the Army (SAEDA)
3	Rival employer attack
4	Jumping the airgap
5	Human Interface Device (HID) attack

Note: Details are in the Appendix.

The Blue cell proposed 15 specific mitigations for their defense against the proposed attack strategies. Since it is not feasible to consider all possible ($N_{d,max} = 2^{15} = 32,768$) defender strategies that could be constructed from the 15 mitigations, 4 defense strategies $\{s_{d,j}\}$, $j = 1 \ldots 4$ were identified by selecting specific mitigations. These 4 strategies offer potentially good security with reasonably acceptable costs. A fifth strategy, "No Action", was also noted and is identified here as *j=0* for reference. Defender mitigation costs were estimated manually by the Blue cell.

3.4.2 Attacker Computations

Four attack layers (cf. Fig. 3) were identified by the Red cell. We reference the attack layers with index $\{l\}$ and list descriptions of the layers in Table 2.

Table 2 Penetration layers for attack

Attack strategy index (l)	Layer description
1	Penetration into wireless Layer-2 network
2	Penetration of router RTR-1
3	Penetration of switch SW-1 into plant network
4	Access to PLC

Attacker costs were estimated by the Red cell for each attack strategy $s_{a,i}$. Costs were estimated generally as a function of Blue-cell mitigation. This differs slightly from our mathematical framework, which requires costs as a function of $\{ij\}$, so we must translate the Red-cell input appropriately. We separate the attacker costs according to whether the costs were dependent on defender strategies (mitigations) or not. Fixed costs, listed in Table 3, are independent of mitigations. Differential costs, listed in Table 4, are dependent on mitigations. Many of the fixed costs are associated with preliminary research, reconnaissance, and bribes, which occur before the attack begins (i.e., before penetrating Layer 1).

Table 3　Fixed attacker costs and success probabilities

Attack index {i}	Attacker's efforts, expenses	Layer index {l}	Cost (h or $)	Probability of success	Mitigations of interest
1	Research for WiFi attack	NA	120 h	NA	NA
1	Time to brute-force crack WiFi password	1	24 h	1.0	1
1	Time to brute-force crack RTR-1 password	2	48 h	1.0	8
1	Time to brute-force crack SW-1 password	3	48 h	1.0	8
1	Time to set up modification of PLC traffic	4	4 h	1.0	7,8,10
2	Research, bribe	NA	$65,000	1.0	NA
3	Research, bribe	NA	$320,000	1.0	NA
4	Research, setup	NA	$35,000	0.5	NA
5	Research, bribe	NA	$40,000	0.5	NA

Note: These costs were independent of mitigation or defense strategy. We assume labor costs of $1,000/h for Red-cell activities.

For the purposes of computing our game-theoretic model, we construct attacker fixed cost $C_{a,ij}^0$ and penetration probability $p_{i,j}^0$ matrices for attacker and defender strategies {i,j} using the fixed Red-cell input from Table 3:

$$C_{a,ij}^0 = \begin{bmatrix} 244 & 244 & 244 & 244 & 244 \\ 65 & 65 & 65 & 65 & 65 \\ 320 & 320 & 320 & 320 & 320 \\ 35 & 35 & 35 & 35 & 35 \\ 40 & 40 & 40 & 40 & 40 \end{bmatrix}, \text{ and } \quad p_{i,j}^0 = \begin{bmatrix} 1.0 & 1.0 & 1.0 & 1.0 & 1.0 \\ 1.0 & 1.0 & 1.0 & 1.0 & 1.0 \\ 1.0 & 1.0 & 1.0 & 1.0 & 1.0 \\ 0.5 & 0.5 & 0.5 & 0.5 & 0.5 \\ 0.5 & 0.5 & 0.5 & 0.5 & 0.5 \end{bmatrix}.$$

Table 4 Differential attacker costs

Attack Index {i}	Mitigation index {m}	Impact of mitigation	Layer index {l}	Cost (h)	Probability of success
1	1	Additional time needed for cracking WiFi	1	24	1.0
	7	Higher probability of detection	4	0	0.9
	8	Higher probability of detection	2	0	0.5
	8	Higher probability of detection	3	0	0.5
	8	Higher probability of detection	4	0	0.5
	10	Advanced research needed and higher probability of detection	4	4	0.7
2	2	Lower probability of successful subversion	4	0	0.2
	3	Wastes labor and higher probability of detection	4	20	0.7
	7	Higher probability of detection	4	0	0.8
	8	Higher probability of detection	2	0	0.5
	8	Higher probability of detection	3	0	0.5
	8	Higher probability of detection	4	0	0.5
	10	Advanced research needed and higher probability of detection	4	4	0.5
	11	Higher probability of detection	3	0	0.25
3	3	Wastes labor and higher probability of detection	4	20	0.7
	7	Higher probability of detection	4	0	0.9
	10	Advanced research needed and higher probability of detection	4	4	0.7
	12	Defeats the attack	3	0	0.0
	14	Lower probability of success	3	0	0.5
4	6	Lower probability of success	3	0	0.5
	10	Advanced research needed and higher probability of detection	4	4	0.5
	13	Malware may be detected	3	0	0.9
5	6	Lower probability of HID installation	3	0	0.5

Notes: These costs are dependent on Blue-cell mitigations. We again assume labor costs of $1,000/h for Red-cell activities.

Next, we use the Red-cell input from Table 4 to construct differential attacker cost $C_{a,lm}^{i}$ and penetration probability $p_{a,lm}^{i}$ matrices, one for each attack {i}. Since the Red-cell cost definitions are dependent on Blue-cell mitigations and the protections of each defensive layer, we index these matrices over layer l and mitigation m. While this information could be stored and computed using 3-D matrices $C_{a,ilm}$ and p_{ilm}, for the purposes of illustration, we use 5 separate 2-D matrices (one for each attack i) in the following discussion.

For example, for Attack 3, by examining Table 4, we can construct differential cost $C^3_{a,lm}$ and penetration probability matrices $p^3_{a,lm}$:

$$C^3_{a,lm} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 20 & 0 & 0 & 0 & 0 & 0 & 0 & 4 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}, \text{ and}$$

$$p^3_{a,lm} = \begin{bmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 0.0 & 1 & 0.5 & 1 \\ 1 & 1 & 0.7 & 1 & 1 & 1 & 0.9 & 1 & 1 & 0.7 & 1 & 1 & 1 & 1 & 1 \end{bmatrix}.$$

Since our game-theoretical model references defender strategy j instead of mitigation m, we recompute the differential attacker cost over $\{l,j\}$: $C^3_{a,lj} = C^3_{a,lm}M_{mj}$, where $M_{mj} = M^T_{jm}$ is the transpose of M_{mj} (see Eq. 5). The result for attack $i = 3$ is

$$C^3_{a,lj} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 4 & 4 & 4 & 24 \end{bmatrix}.$$

We next need to condense $\{l\}$ and construct a single differential cost matrix for all attacks $\{i\}$. By summing row vectors over $\{l\}$, we obtain $C^3_{a,j} = [0\ 4\ 4\ 4\ 24]$ for Attack 3 and perform similar calculations for the other 4 attacks. The differential cost matrix $C^D_{a,ij}$ over $\{i,j\}$ is then

$$C^D_{a,ij} = \begin{bmatrix} C^1_{a,j}; C^2_{a,j}; C^3_{a,j}; C^4_{a,j}; C^5_{a,j} \end{bmatrix} = \begin{bmatrix} 0 & 28 & 28 & 28 & 28 \\ 0 & 4 & 4 & 4 & 4 \\ 0 & 4 & 4 & 4 & 24 \\ 0 & 4 & 4 & 4 & 4 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}.$$

We can now compute the total attacker cost matrix $C_{a,ij}$ over $\{i,j\}$:

$$C_{a,ij} = C^0_{a,ij} + C^D_{a,ij} = \begin{bmatrix} 244 & 272 & 272 & 272 & 272 \\ 65 & 69 & 69 & 69 & 69 \\ 320 & 324 & 324 & 324 & 344 \\ 35 & 39 & 39 & 39 & 39 \\ 40 & 40 & 40 & 40 & 40 \end{bmatrix}.$$

The differential penetration probabilities for the mitigations are also computed from information in Table 4 by constructing the differential penetration probabilities p^D_{im} for each attack $\{i\}$ and each mitigation $\{m\}$:

$$
p_{im}^D = \begin{bmatrix}
1 & 1 & 1 & 1 & 1 & 1 & 0.9 & 0.125 & 1 & 0.7 & 1 & 1 & 1 & 1 & 1 \\
1 & 0.2 & 0.5 & 1 & 1 & 1 & 0.8 & 0.125 & 1 & 0.5 & 0.25 & 1 & 1 & 1 & 1 \\
1 & 1 & 0.7 & 1 & 1 & 1 & 0.9 & 1 & 1 & 0.7 & 1 & 0.0 & 1 & 0.5 & 1 \\
1 & 1 & 1 & 1 & 1 & 0.5 & 1 & 1 & 1 & 0.5 & 1 & 1 & 0.9 & 1 & 1 \\
1 & 1 & 1 & 1 & 1 & 0.5 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1
\end{bmatrix}.
$$

Each entry is computed by multiplying probabilities over layers $\{l\}$ for each $\{i,m\}$ combination. For example, as seen in Table 4, Attack 2's Mitigation 8 has penetration probabilities of 0.5 for Layers 2, 3, and 4, so that $p_{im}^D = (0.5)^3 = 0.125$.

Once p_{im}^D is constructed, we then use the translation matrix M_{mj} to calculate and redistribute the probabilities as a function of defender strategy j. For example, p_{ij} for attack $i=3$ and defense $j=2$ includes Mitigations 1, 6, 9, 10, and 12, which have probabilities of success of 1.0, 1.0, 1.0, 0.7, and 0.0, respectively. These 5 values multiply to $p_{ij=32}^D = 0.0$. Continuing this process, we find the resulting differential probabilities of success for all layers $\{l\}$ as a function of attack and defense strategies $\{i,j\}$ is

$$
p_{ij}^D = \begin{bmatrix}
1 & 0.7 & 0.63 & 0.63 & 0.07875 \\
1 & 0.5 & 0.08 & 0.02 & 0.00125 \\
1 & 0.0 & 0.0 & 0.0 & 0.0 \\
1 & 0.25 & 0.225 & 0.225 & 0.225 \\
1 & 0.5 & 0.5 & 0.5 & 0.5
\end{bmatrix}.
$$

The total probability matrix can then be computed by taking the element-by-element product with the fixed probability matrix: $P_{T,ij} = p_{ij}^0 \, p_{ij}^D$:

$$
P_{ij}^T = \begin{bmatrix}
1 & 0.7 & 0.63 & 0.63 & 0.07875 \\
1 & 0.5 & 0.08 & 0.02 & 0.00125 \\
1 & 0.0 & 0.0 & 0.0 & 0.0 \\
0.5 & 0.125 & 0.1125 & 0.1125 & 0.1125 \\
0.5 & 0.25 & 0.25 & 0.25 & 0.25
\end{bmatrix}.
$$

The utility cost of the attacker $u_{a,ij}$ is computed from Eq. 1b. For our tabletop war game, the benefit b was \$100,000 so that

$$
u_{a,ij} = b\, P_{ij}^T - C_{a,ij}^T = \begin{bmatrix}
756 & 428 & 358 & 358 & -193.25 \\
935 & 431 & 11 & -49 & -67.75 \\
680 & -324 & -324 & -324 & -344 \\
465 & 86 & 73.5 & 73.5 & 73.5 \\
460 & 210 & 210 & 210 & 210
\end{bmatrix}.
$$

3.4.3 Defender Computations

To perform the game-theoretical computations, we first construct a translation matrix M_{jm} between mitigations and defense strategies, based on Blue-cell input shown in Table 5.

Table 5 Blue-cell defense strategies and mitigations. Estimated costs are shown for individual mitigations and for the 5 strategies associated with the mitigations.

Mitigation			Defender strategy				
{m}	Description	Cost ($)	No action (j=0)	Basic security (j=1)	IDS+ (j=2)	IDS enhanced (j=3)	IDS+ and physical security (j=4)
1	Upgrade WiFi security	10,000	…	X	X	X	X
2	Install hardware firewalls	15,000	…	…	X	X	X
3	Install network honeypots	30,000	…	…	…	…	X
4	Configure VLANs	4,000	…	…	X	X	X
5	Install clear conduit	8,000	…	…	…	X	
6	Safeguard plant documents	10,000	…	X	X	X	X
7	Install IDS	40,000	…	…	X	X	X
8	Install layer-2 IDS sensor feed	40,000	…	…	…	…	X
9	Apply STIG controls	25,000	…	X	X	X	X
10	Upgrade training methods	15,000	…	X	X	X	X
11	Monitor ports on devices	30,000	…	…	…	X	X
12	Disallow USB media installs	5,000	…	X	X	X	X
13	Upgrade scanning station	15,000	…	…	X	X	X
14	Lock BIOS on devices	5,000	…	…	X	X	X
15	Upgrade physical security for PLCs	20,000	…	…	…	…	X
Defender strategy costs $C_{d,ij}$ ($) (independent of attack {i})		0	65,000	144,000	182,000	264,000

$$M_{jm} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 & 1 & 0 & 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 1 & 0 & 1 & 1 & 0 & 1 & 1 & 0 & 1 & 1 & 1 & 0 \\ 1 & 1 & 0 & 1 & 1 & 1 & 1 & 0 & 1 & 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 0 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \end{bmatrix}. \quad (5)$$

The vector defining defender costs by mitigation m, $C_{d,m} = [10 \ 15 \ 30 \ 4 \ 8 \ 10 \ 40 \ 40 \ 25 \ 15 \ 30 \ 5 \ 15 \ 5 \ 20]$, is also constructed (see Table 5). We use the translation matrix M_{jm} to compute defender strategy costs as a function of defender j: $C_{d,j} = M_{jm} C_{d,m} =$

$[0 \ 65 \ 144 \ 182 \ 264]$, $j = 0 \ldots 4$, where $M_{jm}C_{d,m}$ implies summation (matrix multiplication) over m. Note that defender costs are independent of attacker strategies i in this tabletop war-game example. This is a generalized case of our game-theoretic framework described in Section 2.1.

Although the defender costs are not dependent on the attacker strategies *{i}*, we construct the defender cost matrix over *{ij}*

$$C_{d,ij} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 65 & 65 & 65 & 65 & 65 \\ 144 & 144 & 144 & 144 & 144 \\ 182 & 182 & 182 & 182 & 182 \\ 264 & 264 & 264 & 264 & 264 \end{bmatrix}$$

and use it to compute the defender utility. Here, we have assumed that the defender strives to keep the portion of the defender's assets b that would have otherwise been forfeited to the attacker as a benefit. The probability that those assets are not forfeited is $p = 1 - P_{ij}^T$.

We can then compute the defender cost utility as

$$u_{d,ij} = b\left(1 - P_{ij}^T\right) - C_{d,ij} = \begin{bmatrix} 0 & 235 & 226 & 188 & 657.25 \\ 0 & 435 & 776 & 798 & 734.75 \\ 0 & 935 & 856 & 818 & 736 \\ 500 & 810 & 743.5 & 705.5 & 623.5 \\ 500 & 685 & 606 & 568 & 486 \end{bmatrix},$$

where P_{ij}^T is computed in Section 3.4.2.

4. Conclusions

In this report we have shown how game-theoretic models can be used to calculate payoff matrices, which are effective tools for advising offensive and defensive players in cyber war games. Without such analytical tools, players in cyber war games often have difficulties making rational, explainable strategy selections.

5. References

1. Hale J. Game theory: cyber preparedness. SC Media; 2014 Sept 2 [accessed 2017 Jun 20]. http://www.scmagazine.com/game-theory-cyber-preparedness /article/366363/.

2. Casey T, Willis B. War games: serious play that tests enterprise security assumptions. Intel; 2008 [accessed 2017 Jun 20]. https://www.sbs.ox.ac.uk/cybersecurity-capacity/system/files/Intel%20- %20Wargames-%20Serious%20Play%20that%20Tests%20Enterprise %20Security %20Assumptions.pdf.

3. Soloman D. The role of cyber war games in developing advanced cyber defence. SC Media; 2014 Jun 9 [accessed 2017 Jun 20]. http://www.scmagazineuk.com/the-role-of-cyber-war-games-in-developing-advanced-cyber-defence/article/354670/.

4. Bailey T, Kaplan J, Weinberg A. Playing war games to prepare for a cyberattack. McKinsey and Co.; 2012 [accessed 2017 Jun 20]. http://www.mckinsey.com/~/media/mckinsey/dotcom/client_service/BTO /PDF/MOBT26_Cyber_war_gaming.ashx.

5. Peters S. Cyber war games: top 3 lessons learned about incident responses. DARKReading; 2015 [accessed 2017 Jun 20]. http://www.darkreading.com/risk/cyber-war-games-top-3-lessons-learned-about-incident-response/d/d-id/1319813.

6. Perla PP. Now hear this – improving war-gaming is worthwhile – and smart. Proceedings. 2016;142(1):1,355 [accessed 2017 Jun 20]. http://www.usni.org /magazines/proceedings/2016-01/now-hear-improving-war-gaming-worthwhile%E2%80%94and-smart.

7. Peck M. What's the best way to war game cyberwarfare? GovTechWorks; 2016 [accessed 2017 Jun 20]. https://www.govtechworks.com/whats-the-best-way-to-war game-cyberwarfare/#gs.Ct26IrY.

8. Perla PP, McGradey E. Why war-gaming works. Nav War Coll Rev. Summer 2011;64(3):111.

9. Burns S, editor. War gamers' handbook: guide for professional war gamers. Newport (RI): US Naval War College; 2015 [accessed 2017 Jun 20]. https://www.usnwc.edu/getattachment/Research---Gaming/War-Gaming/WGD-HB---Complete-2.pdf.aspx.

10. McHugh FJ. Fundamentals of war gaming. 3rd ed. Newport (RI): US Naval War College; 1966 Mar [accessed 2017 Jun 20]. http://www.dtic.mil /cgi-bin/GetTRDoc?Location=U2&doc=GetTRDoc.pdf&AD=AD0686108.

11. Headquarters, Department of the Army. Staff organization and operations. Washington (DC): Headquarters, Department of the Army; 1997. Field Manual No.: FM 101-5.

12. Joint Publication 5-0. Joint operation planning. Arlington (VA): Joint Chiefs of Staff; 2011 Aug 11 [accessed 2017 Jun 20]. http://www.dtic.mil/doctrine /new_pub/jp5_0.pdf.

13. Kott A, Singh R, McEneaney W, Milks W. Hypothesis-driven information fusion in adversarial, deceptive environments. Info Fus. 2011;12(2):131–144.

14. Rasch R, Kott A, Forbus K. Incorporating AI into military decision making: an experiment. IEEE Intel Sys. 2003;18.4:18–26.

15. Kott A, Ground L, Langston J. Estimation of battlefield attrition in a course of action analysis decision support system. Military Operations Research Society Workshop on Land and Expeditionary Warfare. 1999 June.

16. Hausken K. Probababilistic risk analysis and game theory. Risk Analysis. 2002;22(1):17.

17. Leitmann G. On generalized Stackelberg strategies. Optim Theory App. 1978;26(4):637.

18. von Stengel B, Zamir S. Leadership games with convex strategy sets. Games Econ Beh. 2010;21(1–2):282.

19. Korzhyk D, Yin Z, Kiekintveld C, Conitzer V, Tambe M. Stackelberg vs. Nash in security games: an extended investigation of interchangeability, equivalence, and uniqueness. J Artif Intell Res (JAIR). 2011;41:297

20. Colbert E, Sullivan DT, Wong KW, Smith SC, Stephenson S, Sfakianoudis V, Ritchey RP, Parker TW, Knachel LP, Hutchinson SE, et al. RED and BLUE teaming of a US Army SCADA system: table-top exercise final report. Aberdeen Proving Ground (MD): Army Research Laboratory (US); 2015 Oct. Report No.: ARL-TR-7497.

21. Colbert E, Sullivan D, Wong K, Smith S. Table-top exercise: intrusion detection capabilities for us army SCADA systems, information packet. Aberdeen Proving Ground (MD): Army Research Laboratory (US); 2015 Oct. Report No.: ARL-TR-7498.

Appendix. AQUA War-Game Attacker and Defender Strategy Details

The contents of this appendix are taken directly from a previously published report, Colbert et al.[1] This appendix summarizes the Red cell's attack plans and the Blue cell's mitigations of a tabletop cyber war game held at the US Army Research Laboratory. See Section 3.1 of this report for a description and diagrams of the fictional AQUA cyber-physical system.

A.1 RED Team

A.1.1 Attack 1: Strategy $s_{a,1}$ – Layer-2 Attack

The first attack proposed by the Red cell against the AQUA Industrial Control System gains access by cracking the password of the wireless network. The authentication credentials of the wireless tablets are then spoofed to penetrate the rest of the plant network. Complete control of the PLC traffic is gained, which is used to undercook the chicken in the Meal Ready-to-Eat (MRE) production line. Details of the assumptions and the attack methodology are given below, as well as discussion of some countermeasures by the Blue cell.

A.1.1.1 Assumptions

For this attack, the Red cell assumed they would have physical proximity to the plant (i.e., they could position themselves within range of the wireless access point). Furthermore, they assumed that the wireless access point's pre-shared key (PSK) was not overly complex (i.e., methods other than full brute-force could crack the key).

A.1.1.2 Attack Methodology

A.1.1.2.1 Obtain WiFi protected access 2 (WPA2) key. By sniffing the wireless access point, the Red cell would discover Media Access Control (MAC) addresses of the tablets, specifically those of the wireless tablets. By spoofing de-authenticate packets, the Red cell would force the wireless tablet to re-authenticate with the access point. The Red cell would capture the re-authentication packets and use a password cracker (e.g., on the Amazon Elastic Cloud Compute [EC2]) to obtain the pre-shared WPA2 key. Alternatively, rainbow tables could be used to crack the key.

---

[1] Colbert E, Sullivan DT, Wong KW, Smith SC, Stephenson S, Sfakianoudis V, Ritchey RP, Parker TW, Knachel LP, Hutchinson SE, et al. RED and BLUE teaming of a US Army SCADA system: table-top exercise final report. Aberdeen Proving Ground (MD): Army Research Laboratory (US); 2015 Oct. Report No.: ARL-TR-7497.

A.1.1.2.2   Spoof and authenticate. Once the Red cell has the pre-shared key, they would spoof the MAC address of one of the plant's wireless tablets and authenticate to the WPA2 access point.

A.1.1.2.3   Connect to Texas router. After gaining access to the plant network, the Red cell would connect to the Texas router and authenticate using default credentials, which are assumed to be present on the router.

A.1.1.2.4   Covert layer-2 tunnel. The Red cell would then establish a Layer-2 Tunneling Protocol (L2TP) tunnel between their rogue wireless device and the Texas router in order to evade detection by Layer-3 devices.

A.1.1.2.5   Open the Switch and/or router. The Red cell would then use the Maintenance Operations Protocol–Remote Console (MOPRC) to connect to the Texas switch, using default credentials or no credentials. MOPRC is a Layer-2 protocol that is enabled by default on Texas devices. At this point, the Red cell can either change the default credentials or remove the need for credentials by updating the configuration. The router or switch configuration changes will be accomplished in memory, so no trace of tampering will be found if the router is rebooted.

A.1.1.2.6   Create new VLANs. Gaining access to the switch, the Red cell would create 3 separate virtual local area networks (VLANs): one VLAN for the PLCs, one VLAN for the HMI workstations, and the last VLAN for the Historian. The Red cell will configure the switch so all VLAN traffic is first sent to their rogue wireless device, allowing modification of the traffic. The switch configuration changes will be accomplished in memory, so no trace of tampering will be found if the switch is rebooted.

A.1.1.2.7   Spoof the switched port analyzer (SPAN) traffic. After isolating the SPAN port on the Texas switch, the Red cell would send "normal" traffic to the SPAN port to avoid detection by any network sensors that may be present.

A.1.1.2.8   Meals NOT fit to eat. Since the rogue wireless device now intercepts all Modbus traffic between the HMIs and the PLCs via the VLANs, the Red cell can modify traffic in each direction. They would impersonate the PLCs and send "correct" Modbus responses back to the HMIs and the Historian. Likewise, they would send "incorrect" Modbus messages to the meat cooker PLC to cause the chicken to be undercooked. The goal is for the chicken to pass the visual quality-assurance test, but it will not be fit to eat.

A.1.1.2.9 Randomize the attack pattern. The Red cell would launch this attack at randomized times for only 5–10% of the MREs to make this problem difficult to troubleshoot.

A.1.2 Attack 2: Strategy $s_{a,2}$ – SAEDA

In this attack, the Red cell coerces a plant employee to install a rogue wireless device or malware into the plant network. The Red cell, via remote access to the wireless device or via the malware, would misdirect the PLCs in such a manner as to disrupt or destroy plant MRE processes.

A.1.2.1 Assumptions

It is assumed that the subverted employee could successfully perform malicious actions without being noticed directly, by video surveillance, by perimeter guards, or by network monitoring.

A.1.2.2 Attack Methodology

A.1.2.2.1 Identify and befriend plant employees. The Red cell would identify plant employees using physical observation and surveying job web sites such as Monster, LinkedIn, and Dice. The Red cell would then establish a personal relationship with the employee by frequenting businesses or evening hangouts of the employee, and befriending him or her at those locations.

A.1.2.2.2 Subversion of the employee. Once plant employees are identified, the Red cell would try to subvert the employee using blackmail, bribes, or other techniques like Chinese water torture.

A.1.2.2.3 Implantation of rogue wireless device. The Red cell would give the coerced plant employee a Raspberry Pi-based device (for example) to plug into unused switch ports or workstation USB ports. This Raspberry Pi-based device would setup a rogue wireless access point using WiFi.

For an additional level of covertness, the WiFi radio access could be replaced with a cellular radio, and the Raspberry Pi connection to the plant network could be made with a vampire tap rather than through a direct USB or Ethernet port.

A.1.2.2.4 Disrupt or destroy MRE process. Once the Red cell has their rogue Raspberry Pi device connected, the device would flood the PLCs with packets, or flash the PLC with defective firmware or program code to destroy the MRE

production process. The rogue Raspberry Pi device can also carry out a Layer-2 attack.

A.1.2.2.5 Alternative access method. If the Red cell insider cannot install the rogue Raspberry Pi-device, the Red cell would create malware. The Red cell would provide instructions for the insider to manually install the malware on a machine in the plant network. The malware would damage the PLCs by causing rapid power-cycling or cause circuit breakers to trip, thus disrupting plant productivity.

A.1.3 Attack 3: Strategy $s_{a,3}$ – Rival Employer Attack

Social engineering is used in this attack to gain initial access to the plant network. Posing as a rival employer searching for new hires, the Red cell acquires vital plant information by interviewing plant employees. When a job opening becomes available at the AQUA plant, an insider sponsored by the Red cell takes the job and consequently performs malicious activities to destroy or damage the AQUA plant processes.

A.1.3.1 Assumptions

In order for this attack to work, the AQUA plant must have job openings and the Red cell must find qualified applicants for those jobs.

A.1.3.2 Attack Methodology

A.1.3.2.1 Fake Competitor. The Red cell began this scenario by masquerading as a competitor business to the AQUA plant. The fictitious business will contact plant employees and invite them for interviews. During the interviews, the Red cell would ask probing questions about the plant processes and acquire critical reconnaissance information.

A.1.3.2.2 Hiring of Insider into Plant. Assuming the plant has job openings, using the information gleaned from the interviews, the Red cell would try to get one of their own members hired into an AQUA plant vacancy.

A.1.3.2.3 Install Dial-up Access. Some of the older HMI workstations may have modems installed. Using that modem, or a covertly concealed modem, the Red cell insider would connect a HMI workstation modem into a telephone port to give the Red cell dial-in access to the plant network.

A.1.3.2.4   Access to the Plant Network. With access to the plant network, the Red cell could perform many of the actions described earlier. Some examples are:

- Disable protective controls on the plant machinery and then run a malicious program to overpressure the high-pressure processor, potentially damaging or destroying plant machinery.

- Introduce an implant on the Texas switch and proceed to attack the PLC devices with malicious Modbus commands, as before.

- Wipe the Historian using a SQL injection attack or use malware to send Modbus messages to the Historian with erroneous data. Either attack to the Historian would raise questions about quality of the MREs, and the plant will discard all food in production, resulting in financial losses.

A.1.4   Attack 4: Strategy $s_{a,4}$ – Jumping the Airgap

In this attack, the Red cell assumes network access to the plant network could be obtained, using a method similar to those proposed in earlier Attacks 1–3. The target plant machine in this attack is the machine that writes the media for patching the plant equipment. In a real scenario, this could be an isolated system. For AQUA, the machine is on the corporate network. Once that machine is owned, malware with capabilities for disrupting/destroying and Command and Control (C2) is installed and executed strategically.

A.1.4.1   Assumptions

The Red cell assumed that network access to the plant network has been successful, either through a network or insider attack.

A.1.4.2   Attack Methodology

A.1.4.2.1   Access method. Similar to previous attacks, the Red cell would need to access the plant network. For this attack, access would be achieved using one of the following attacks, or a similar method to those described in Attacks 1–3:

- Spear phishing

- Web site "Watering hole" attack

- Attacks against the wireless network

- Insider attack

A.1.4.2.2 Malicious patching workstation. The Red cell would then compromise the corporate "patching" workstation, which creates software patches and antivirus (AV) updates for plant devices. Each patch (CD or other media) that is created on the patching workstation would also include additional malware that will execute on the machine being patched. The malware will be launched using AutoRun and will include methods to escalate priviliges and achieve persistence.

A.1.4.2.3 Disable/Destroy Plant Process. As described earlier, once access is achieved to the plant network devices, many options are available for disabling or destroying the plant processes.

A.1.5 Attack 5: Strategy $s_{a,5}$ – Human Interface Device (HID) Attack

In this attack, a HID, such as a USB mouse, infected with malware launches additional malware into the plant network, ultimately disrupting or destroying the plant processes.

A.1.5.1 Assumptions

The malicious HID can be made to appear legitimate.

A.1.5.2 Attack Methodology

A.1.5.2.1 Malicious HID. A Red cell insider would introduce a malicious HID device into the plant. Alternatively, malicious code would be introduced into a HID using supply-chain techniques.

A.1.5.2.2 Connection of the malicious HID. The malicious HID would then be connected to one of the plant workstations, either by the insider or by a plant operator.

A.1.5.2.3 Launch of malware from the HID. Malware is launched automatically from the HID without human intervention, once the HID is connected to the host workstation. Malicious code, generated by the HID, then runs on the workstation, sending malicious Modbus commands to the PLCs.

A.1.5.2.4 Disable/destroy plant process. As described earlier, once access is achieved to the plant network devices, many options are available for disabling or destroying the plant processes.

A.2 Blue-Cell Mitigations

The first Blue-cell countermeasure, based on Attack 1, was to remove the WiFi network, since it presents a significant attack surface for access to the plant network. This would eliminate the access vector for this attack completely.

If the AQUA plant must keep the WiFi network, the wireless network would be put in its own subnet, firewalled, and/or isolated from the plant network. Intrusion-detection methods would be used to detect unauthorized access in the wireless network segment and the other segments of the plant network.

The Blue cell also recommended using machine learning algorithms to develop a profile of normal network activity and tracking the use of each MAC address with that algorithm. The spoofed wireless tablet traffic would then stand out and be detected.

A lively discussion between Red and Blue cell members developed with regard to mitigations and counter-mitigations that could be applied to the Attack 2. The net result is summarized as a list of Blue-cell countermeasures:

- The Blue cell decided they will require all system administrators to undergo background investigations, and will require all employees to undergo security awareness training.

- An IEEE 802.1X port security policy of Deny All Permit by Exception (DAPE) would be implemented, and all USB ports would be locked down so that new devices cannot be added. All workstations and network elements would be hardened using the Department of Defense (DoD) Security Technical Implementation Guide (STIG) guidelines.

- AutoRun would be disabled for CDs which would prevent automatic loading of malware placed on the CD.

- The Blue cell also decided to install internal honeypots to detect suspicious activity and to stand up WiFi jammers to prevent a Red-cell remote connection to and from the Raspberry-Pi device.

- A network management system would be installed on the plant network to capture and detect unauthorized Internet Protocol (IP) addresses connecting with plant equipment and domain name system (DNS) queries. This will enable detection of communications to and from the Raspberry-Pi device.

- The AirTight wireless monitoring system will be installed to detect if third generation (3G) Code-Division Multiple Access (CDMA) or WiFi are broadcast within the plant.

- All plant Ethernet cabling would be placed within clear conduit tubing and monitored by guards to prevent taps being introduced.

- Insider threat is obviously a significant problem and is very difficult to mitigate completely.

AQUA uses plain old telephone system (POTS) telephony. In response to Attack 3, the Blue cell mitigated the ability to use a modem in this attack by establishing a policy to physically disable RJ-11 modem connectors in workstations not requiring modem capabilities.

To prevent the Red cell from using "war-dialing", a POTS security system would be installed so that an incoming call would be disconnected immediately. Automatic callback would be established only if the remote modem phone number is on an authorized list.

User permissions would be limited for plant equipment by establishing role-based access controls for all employees.

Network access to the network switches and routers would be restricted by installing ACLs. Physical access to the network equipment would be highly restricted to the network elements. Switches and routers would be patched appropriately according to the STIG.

To mitigate Attack 4, the Blue cell banned all USB media from the plant. Patches and AV updates can only be applied via CDs. In addition, a separate scanning workstation would scan all CDs to check for malware. The scanning workstation would calculate and verify Message Digest 5 (MD5) and SHA hashes of all patches and AV updates on the CDs. AutoRun would be disabled on all workstations.

Finally, the response to Attack 5 included implementing access restrictions and disabling all unneeded USB ports in accordance with the DOD STIG.

INTENTIONALLY LEFT BLANK.

List of Symbols, Abbreviations, and Acronyms

2-D	2-dimensional
3-D	3-dimensional
3G	third generation
ACL	access control list
AV	anti-virus
CCTV	closed circuit television
CDMA	Code-Division Multiple Access
COA	course of action
CPS	cyber-physical system
DAPE	Deny All Permit by Exception
DNS	domain name system
DoD	Department of Defense
EC2	Elastic Cloud Compute
HID	Human Interface Device
HMI	human–machine interface
IDS	Intrusion Detection System
IP	Internet Protocol
L2TP	Layer 2 Tunneling Protocol
MAC	Media Access Control
MD5	Message Digest 5
MOPRC	Maintenance Operations Protocol–Remote Console
MRE	Meal Ready-to-Eat
POTS	plain old telephone system
PLC	programmable logic controller
PSK	pre-shared key

SAEDA   Subversion and Espionage Directed against the Army

SCADA   Supervisory Control and Data Acquisition

SPAN   switched port analyzer

SSE   Strong Stackelberg Equilibria

STIG   Security Technical Implementation Guide

USB   universal serial bus

VLAN   virtual local area network

WiFi   wireless fidelity

WPA2   WiFi protected access 2

INTENTIONALLY LEFT BLANK.