A GAME ORIENTED APPROACH TO MINIMIZING CYBERSECURITY RISK

SCOTT MUSMAN & ANDREW J. TURNER
The MITRE Corporation, McLean, Va, USA.

ABSTRACT

Information and Communication Technology (ICT) systems are now ubiquitous in all aspects of our society. With an ability to create ICT incident effects via cyberspace, criminals can steal information or extort money, terrorists can disrupt society or cause loss of life, and the effectiveness of a military can be degraded. These threats have caused an imperative to maximize a system's cyber security resilience. Protecting systems that rely on ICT from cyber-attacks or reducing the impacts that cyber incidents cause is a topic of major importance. In this paper, we describe an approach to minimizing cybersecurity risks called Cyber Security Game (CSG), where CSG can be viewed as a form of model-based system security engineering. CSG is a method and supporting software that quantitatively identifies mission outcome focused cybersecurity risks and uses this metric to determine the optimal employment of security methods to use for any given investment level. CSG maximizes a system's ability to operate in today's contested cyber environment by minimizing its mission risk. The risk score is calculated by using a cyber mission impact assessment (CMIA) model to compute the consequences of cyber incidents, and by applying a threat model to a system topology model and defender model to estimate how likely attacks are to succeed. CSG takes into account the widespread interconnectedness of cyber systems, where defenders must defend all multi-step attack paths and an attacker only needs one to succeed. It employs a game theoretic solution using a game formulation that identifies defense strategies to minimize the maximum cyber risk (MiniMax), employing the defense methods defined in the defender model. This paper describes the approach and the models that CSG uses.

Keywords: cybersecurity, game theory, return on investment, risk assessment, risk management

1 INTRODUCTION

Information and communications technology (ICT) is now integral to almost every aspect of our daily activities. The detrimental aspect of introducing ICT dependencies, however, is that this makes us susceptible to impacts from cyber incidents. Criminals can steal and extort money or information, terrorists can disrupt society or cause loss of life, and the effectiveness of a military can be degraded, all because of an ability to create incident effects in cyberspace, and often without any requirement of physical proximity. Protecting ICT from cyber incident effects or reducing their impacts on operational activities has become of international importance. There is an escalating imperative to identify and minimize operational cyber risk.

In almost all circumstances, we are interested in achieving operational resilience: the ability for mission systems to continue to acceptably achieve their function in the face of incidents, while also doing what's best to fulfill other security requirements. Achieving such resilience almost always must be pursued in the face of resource limitations. Often we need to justify the costs and resources needed, so a problem with trying to achieve resilience is that we need a way to measure it, including how to identify whether efforts to improve it are successful or not. If operational resilience is defined in a qualitative rather than quantitative way very little prescriptive advice can be offered to increase it. Many of today's risk management methods provide only generic guidance, or recommended best practices, and produce only risk rankings. As described by [1] risk rankings lack the information required to support optimal allocation of resources to manage those risks, and more importantly: ranking doesn't contain information of how an attacker will modify their behavior in the face of defender actions.

ISSN: 2041-9031 (paper format), ISSN: 2041-904X (online), http://www.witpress.com/journals
DOI: 10.2495/SAFE-V8-N2-212-222

The difficulty with trying to manage one's cyber risk is that it requires a holistic view of how a system fulfils its intended purpose. This requires a knowledge of the purpose, the entire range of hazards that are possible, and a need to consider all attack avenues. Due to the interconnectedness of cyber systems attackers can exploit seemingly-non-critical cyber components to bypass security controls and other defenses. Examples include Stuxnet [2] and the Target data breach [3], where non-critical systems were compromised to reach the important ones. In other words, effective cybersecurity does not just involve defending high impact resources, but all inter-connected resources that provide a pathway to those that are mission-critical. Some have proposed techniques where a defender considers only the resources that directly cause impacts [4, 5]. But without an explicit representation of how the ICT components interconnect this approach fails to defend the seemingly non-critical resources that act as stepping stones and consequently fails to defend indirect attack paths an attacker might use to bypass defenses. An attacker can choose any attack method, ranging from those with a high expected value (payoff), to speculative probes that might reveal a weakness not immediately apparent to a defender obsessed only with "the crown jewels." To defend only against the most obvious attack paths simply allows an attacker to select the next most promising one, and so on, until some exploitable chink in the armor is discovered. Thus, without a holistic understanding of the system, one is likely to make the mistake of addressing only some of the hazards but not others, or to over-invest in addressing some hazards even though the risks from other hazards are higher. A comprehensive assessment is needed to be able to provide prescriptive advice to identify which defensive tools and methods are needed and where.

Cyber attacks and the concomitant defensive actions can be viewed as game playing between two players [6]. One way this can be represented is as a zero-sum game (one where both attacker and defender assign the same value to gain or loss). In two player, finite, zero-sum games, the game theoretic solution concepts of Nash Equilibrium and Minimax produce the same solution – where the defender tries to minimize the maximum payoff for the attacker. When implementing a technique to use to solve this type of game, we need to be concerned with our ability to define the "game state" (i.e. what game board looks like), how attacker and defender moves change the game state, and to be able to evaluate each game state. Since each mission system is going to be different, and since operational objectives are whatever they happen to be, for cyber, this means the game board must be built for each system, and we need to be able to evaluate how the modeled system can fulfill its purpose in the face of cyber incidents. This is made possible by reasoning about cyber incident effects, rather than the incident instances that cause the effects [7], and by leveraging our past work on Cyber Mission Impact Assessment (CMIA) [8, 9] as a way of assessing the goodness of a "game state." Combining the "mission impact" measure with a system topology based [10, 11] threat likelihood model allows us to estimate mission risk.

We describe a program that allows us to play a cyber security game with an objective to minimize a mission system's cyber risk. We call our software CSG, the Cyber Security Game. Our approach is theoretically grounded in game theory, yet focuses on being practically useful. We describe the specifics of how the game is formulated, and each of the models of the system needed to play the game. A cyber mission impact assessment (CMIA) model is used to identify the consequences of cyber incidents. A model is used to describe the ICT topology, connectivity, access relationships, and asset type information. Models define the defensive measures that can be employed, and a default attacker model is provided. Each of these models is algorithmically linked to aspects of the risk calculation so that changes to each can directly alter the mission risk metric. CSG is a model-based solution, meaning that user effort is focused on developing the models of the system and its usage scenario, and CSG's algorithms then

operate on the models to compute mission risk and optimize the employment of security measures. The outcome of running CSG allows us to identify the optimal set of security tools and resilience techniques to reduce the mission systems' cyber risk for any given defender cost.

The next section describes our game, describes how we represent cyber mission systems, how we assess cyber threats, how we have implemented the game, and how we use it to perform a portfolio analysis of tools and methods that can reduce cyber risk. We then discuss aspects of our game program and how it solves the problem we have posed.

2 THE CYBER SECURITY GAME

CSG is an algorithmic method that relies on models that describe the system, its purpose, the threat environment, and the defender capabilities. It runs algorithms on those models to produce results. CSG algorithms automate several expert level capabilities, such as the combinatorics of possible incidents, attack path discovery, and portfolio analysis, so that analysts do not have do them manually. When aspects of the system, defenses, or threats change (i.e., new vulnerabilities are discovered), a defender merely updates the appropriate model and reruns CSG to assess the impact of the changes. An overview of CSG is provided here.

CSG's game formulation searches through the combinatorics of cyber incidents, attack paths, impacts, and defender method employments using Minimax. CSG is formulated as a two-player zero-sum game. As such it implements rational decision making, where both players work to best counteract each other's moves. CSG assumes that the attacker knows about the system they are attacking. This is reasonable because CSG's focus is on defense employment for the long term, assuming that attackers will be able to learn about and may pursue targets over the lifetime of a system. Methods for trying to deceive, delay, or deter an attacker over the short term would motivate a different game formulation. Examples of different formulations are shown in [6].

The system metric in CSG is a system risk score. This score comes from computing the different impacts that an attacker can cause combined conditioned on how difficult the system architecture and defenses makes it to cause those impacts (described later in this paper). In the game, the defender tries to minimize this value. A state in the game represents a configuration of the system, either with no defenses employed, or some status quo version of the system that one wants to improve. To compute the risk score for that state, for the attacker player generates an attack tree that looks multiple attack steps ahead to identify possible impacts. The leaf nodes of the attack tree provide an expected value (EV) computed from the impacts and the probabilities (as shown in Fig. 5). The defender then employs defenses to reduce expected value of the impacts the attacker can cause, and the game proceeds in typical game tree fashion. Starting from the initial state it employs MiniMax, to assess move pairs of a defender action (min), and the attackers revised attack tree given the defender changes (max). The resultant game tree assesses how each defender method can best reduce the risk score. Since the employment of each defender method incurs a cost, the game is over when either the game identifies the optimal set of defense methods to use when the defender has spent the amount of money they have allocated, or when a complete portfolio analysis is performed to compute the Pareto frontier for each price point.

Figure 1 demonstrates the output of CSG, where the cost and performance of each portfolio option is plotted. The red dots show the Pareto optimal portfolios. In this example, 55,296 defense portfolio options were considered. With no defenses the risk score was 8492934. With all the security methods applied, the risk was reduced to 2038408, for a cost of $250k. However, spending ~$39k, the risk can be reduced to 2779440. This represents 89% of the risk reduction for only 16% of the cost.

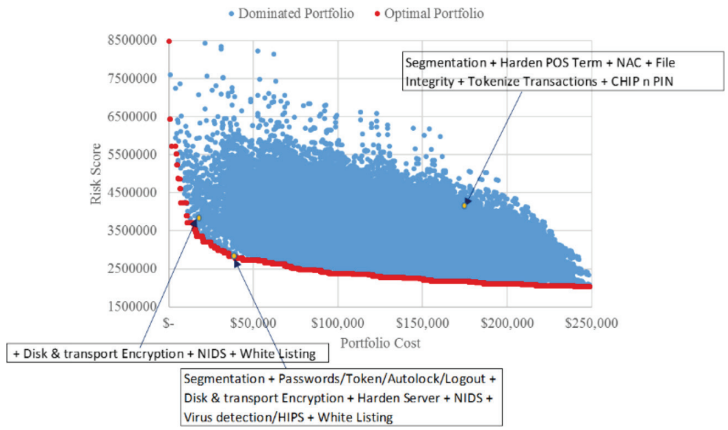

Figure 1: Security Portfolio Effectiveness vs. Cost.

CSG requires four system models. A CMIA process model computes impacts. A system topology model describes the component interconnectivity, type and trust relationships. It contains a built-in (default) attacker model. Lastly, models of the defender methods are needed. Each of these models is described below.

2.1 Modeling Incident Impacts

CSG uses CMIA process models to determine the consequence (losses) incurred from cyber incidents, and the CMIA software is embedded in CSG. CMIA makes it possible to capture mission details as an executable simulation. It models mission activities, ICT activities, activity durations, activity dependencies, ICT resources, temporal constraints, data, and control flows. ICT resources in the process model can be affected by cyber incident effects represented by the DIMFUI cyber incident effect taxonomy described in [7] and relates the occurrence of an effect to mission outcomes in the form of mission impacts. Figure 2 shows

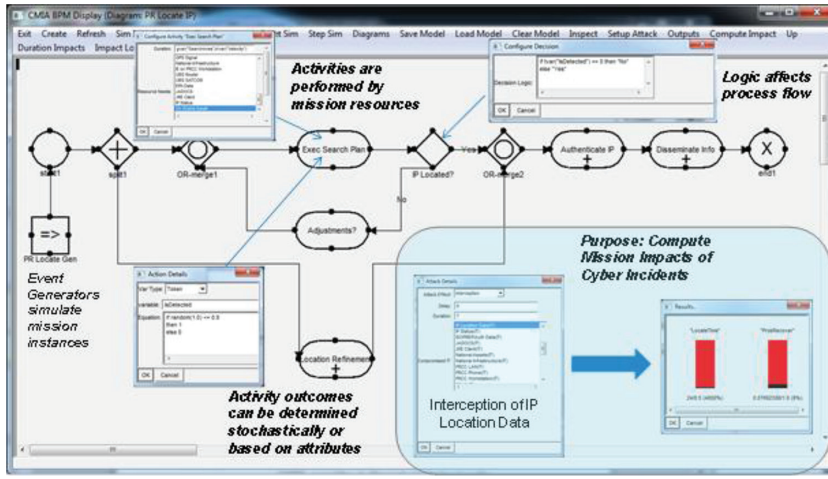

Figure 2: Process Model Showing Mission Activity Order and Control Flow Dependencies.

the CMIA tool displaying a Business Process Modeling Notation (BPMN) diagram, and shows some of the underlying details that turn the model into an executable simulation. It is outside the scope of this paper to describe CMIA in detail but for more information on the CMIA method and tool see [9, 12].

CMIA process models can be probabilistic and stochastic, allowing one to bound the uncertainties associated with the model. Running combinatorics on the set of cyber incident effects allows the criticality of cyber resources to be estimated [12] or validated by mission subject matter experts. CMIA also supports the assessment of multiple (simultaneous) incidents, allowing CSG to run through combinations of the possible cyber incident effects that will be assessed using MiniMax.

2.2 Modeling the Attacker

CSG, provides a default attacker model that defines the probability that attacks will succeed given the topological constraints that the system imparts on the attacker. For the attacker to affect ICT resources that can cause impacts, the attacker must find a pathway to access them. The attack model conditions the probability of an attack succeeding with the following characteristics:

- Whether the attacker is trying to compromise a component they can directly connect to, or whether it requires crossing a network trust boundary to access
- Whether that component is the same type as one of the components that the attacker has already demonstrated they can compromise
- Whether a component is known to be vulnerable to known exploits
- Whether the component is a server that contains one or more network services
- Whether user roles, who have access to each resource, can leverage those roles to access other components in the network to create impact

Figure 3 shows the probabilities of attacker success given topology relationships. If they first try and attack host Win 7-2 before attacking something else, the diagram shows that is it possible to attack the host from the outside with probability of success $P(S \mid OIC)$. It also shows there is a probability in becoming an inside user or subverting an insider account, where the number of accounts N, will affect the risk $N \times P(S \mid IA)$. Once a host is compromised an attacker can compromise other resources on the devices with $P(S \mid HA)$.

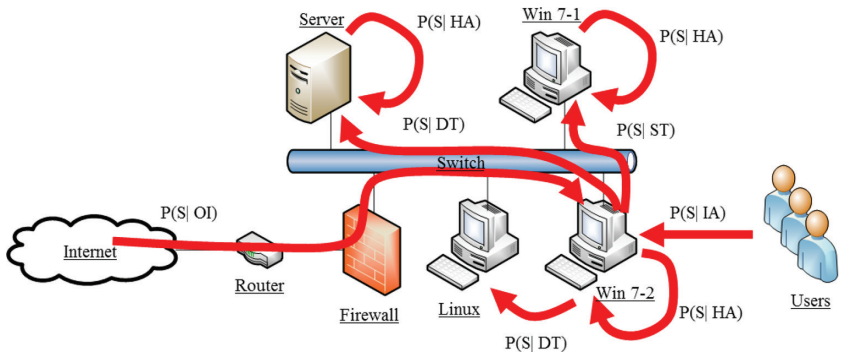

Figure 3: Default Attack Model.

Once in, the attacker can try to attack other cyber resources of the same type as one of the components already compromised then the probability is P(S|STC). If it's of a different type, then the attacker needs a new zero-day attack and has the probability $P(S \mid DTC)$. Throughout, whether the host is a client or a server will affect the probabilities. The probability of successfully navigating the network to compromise a cyber resource is computed using the chain rule shown in Equation 1. This attack model captures the basic security properties of segmentation, diversity, and least privilege and can be composed across multiple networks, trust, and segmentation boundaries. It mathematically credits a defender reducing the amount of access to components that can cause significant impacts, for diversifying components that can cause significant impacts when attacked in combination, it credits a defender for making it harder to reach the components that cause significant impacts

$$P(A_1, A_2, A_3, \ldots, A_n) = P(A_1 \mid A_2, A_3, \ldots, A_n) P(A_2, \mid A_3, \ldots, A_n) .. P(A_{n-1}, \mid A_n) P(A_n).$$

Equation 1: General Form of the Chain Rule

2.3 Modeling System Topology and Applying the Attack Model

CSG computes impacts and then uses the attacker model to estimate the probability that the impacts will occur given the constraints of the system topology. The topology model includes cyber components, applications, data, user account groups, and access controls that enforce trust relationships. Model elements include single ICT resources as well as ICT resource pools that represent functionally identical groupings of resources of the same type. The model also requires resource type information for each ICT resource. This makes it possible to know when the same attacker exploit from an earlier step might be reused, and when it can't. The existence of connections, firewall rules, and the access of user roles define connectivity capabilities and restrictions between ICT resources. An example of a topology model is shown in Fig. 4. User group populations are also included in the model, but the diagram does

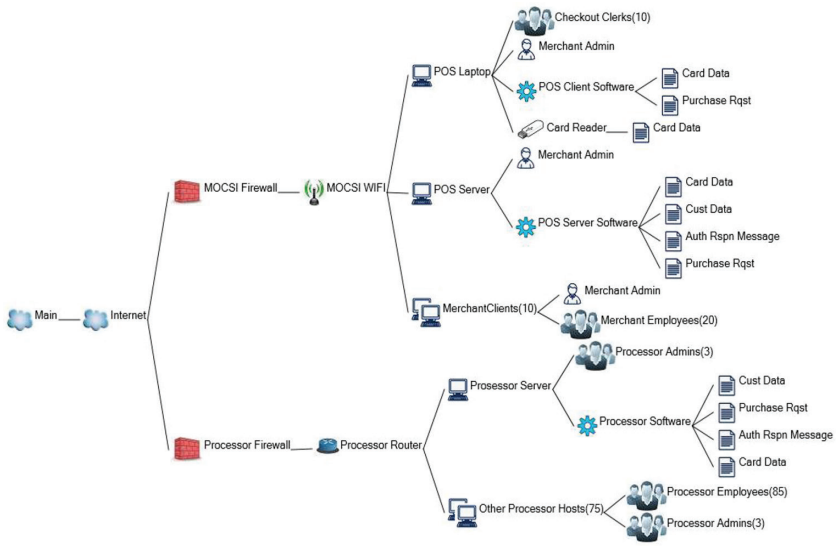

Figure 4: A Simple Topology Model for a Point of Sale System.

not show the trust relationships of who can communicate with who, that are also captured in the model.

2.4 Computing the Risk Score

The topology model makes it possible to compute attack trees. This process is illustrated in Fig. 5. The figure illustrates the complexity of analyzing even a simple four host system. All the mission impacts in the figure would be computed from a CMIA model of the system. The impacts shown include impacts from when multiple components are compromised. Even though the model in the figure contains only two subnets, the model represents a trust relationship from S1 and S2 to S4. The attack tree is then generated, using the attacker model to estimate the probability of attack steps succeeding.

In Fig. 5 only the attack steps from the internet are part of the final tree. This is possible for two reasons. First, rational attackers will always use the optimal path to compromise a component in the system. Second, we are evaluating risks in the context of MiniMax. This allows suboptimal paths to be pruned from the tree. Even so, this small four node example produces a more expansive attack tree than we can fully illustrate in the figure (only some pathways are shown). At the end of each path in the attack tree an expected value is computed (based on the worst-case – the max in MiniMax). To do this it is necessary to consider different pathways to achieve the same impacts. This is necessary because trust relationships (i.e., as imposed by firewall rules) are not necessarily directionally symmetric, so each pathway has a different probability of success. The sum of EV's at the end of each branch in the attack tree represents the risk score for the system. An alternative risk score for risk averse organizations (such as critical infrastructure components) could be the maximum EV that exists in the tree. In the next section we will discuss how defensive measures that reduce these EV's are modeled.

Because an attacker may have to compromise multiple ICT resources to cause significant impacts there is a need to look enough steps ahead to identify those cases. In CSG this number of attack steps to explore is a parameter that can be set. Usually, unless a system is

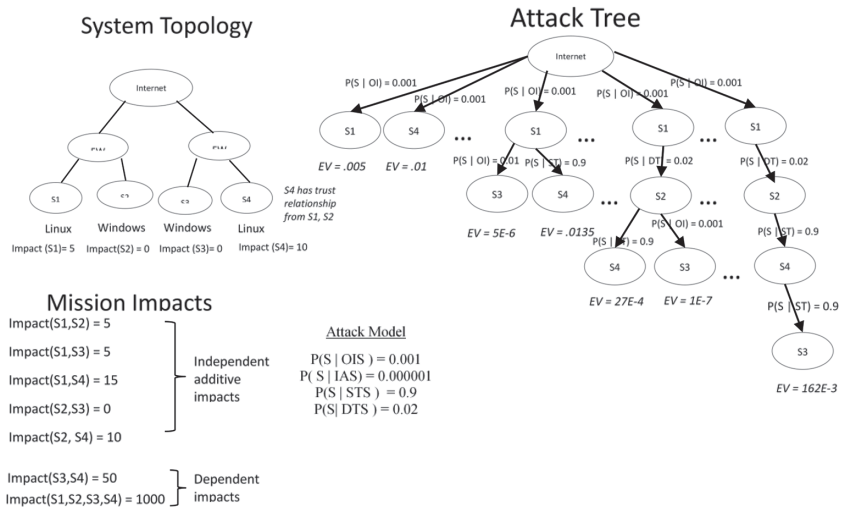

Figure 5: Topological Attack Graph Calculation for Risk Scoring.

composed of particularly layered defensive boundaries, looking 3-4 steps ahead should be sufficient.

2.5 Modeling Defender Methods

To assess defender choices CSG requires models of the defensive measures. These are shown in Fig 6. Some of the defense measures reduce the likelihood that attacks will succeed. This is typically accomplished in one of two ways. One way applies protections to the cyber resources, and the other is to change access. Each defense measure requires an assessment of how well the method is expected to work. Measures that reduce the chance of attack success are shown in the table in Fig. 6. The table lists: the measure name, cost estimates for employing the measure, and measure effectiveness assessment against each attack effect. This effectiveness assessment is defined as a score from 0 to 100. A score of 0 means that the measure is of no use to prevent the effect occurring, and a score of 100 means that the measure completely stops all the attacks that would cause that effect. An interpretation for this score is that a value of 40 should imply that 40% of the attacker exploits that we expect or anticipate as possible would no longer succeed.

Other defender measures can be represented by changes to the CMIA process model. The bottom left of Fig. 6 shows how a redundant server can be represented in the process model. Including a redundant path in the process illustrates the presence of a redundant server. If an incident only affects the original server then no impact occurs. However, since CSG looks multiple attacker steps ahead it will eventually identify the attack scenario that causes incidents on both servers.

Other defensive measures can be represented in the topology model. This is accomplished by copying the topology model and modifying it to reflect any proposed changes. The likelihood that attacks will succeed can be reduced by changing the network topology or changing

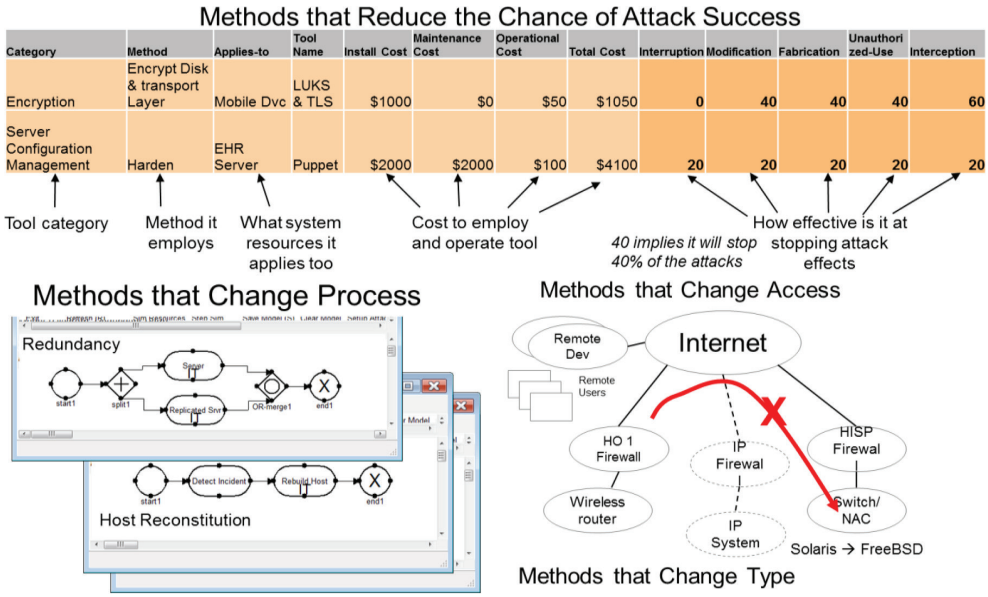

Figure 6: Modeling Defender Methods.

access controls. Risks can also be reduced by the principle of least privilege. This involves reducing the number of pooled resources (i.e., client workstations and/or users) with access. Lastly, the risks can be reduced by diversifying the components in the system, forcing the attacker to require multiple different attacks to create an impact [11].

3 SUMMARY

In this paper, we have described the Cyber Security Game (CSG). CSG is both a method and software that implements the method. CSG implements an approach that represents a quantitative, assessment of a mission system's cyber risk. CSG formalizes the information gathering activities that occur during a traditional risk assessment into computable models, artifacts that describe the system (i.e., a system topology model and impact model). CSG also algorithmically encapsulates expert capabilities in software that leverages these computable artifacts to implement a cyber risk assessment that is a consistent and comprehensive process. CSG artifacts also make it practical to keep risk assessments up to date as the system or missions change. A defender need only update the appropriate CSG model(s) to reflect changes, and rerun CSG to update the assessment.

CSG explores the possible attack paths for as-is or as-may-be system topologies. It uses MiniMax search to investigate how attackers react to defender actions. Finally, it performs a portfolio analysis to identify optimal risk reduction portfolios. Such computations are challenging for non-experts to perform consistently, especially without assisting tools. Implementing these in a tool form can help "level the playing field" for organizations that perform risk assessments, by reducing the need for risk assessment specialists and by producing artifacts that clearly capture the system, mission, threat and defense and assumptions that the assessment used.

When remediating risks it is easy to over invest in defending one portion of the system at the expense of underinvesting in defending others. CSG combats this by looking across the entire set of incidents, and outcomes. It apportions defenses to reduce the risks systematically. Each step in the game attempts to mitigate the largest risks existing at that point in the game. At the same time, it looks some steps ahead to avoid making a locally greedy decision. Each defender move may cause attacking some other component(s) to have the largest payoff for an attacker. Therefore, CSG can answer important cybersecurity questions like, "how much diversity is enough?" CSG answers this by being able to diversify components until the best attacker move becomes one that attacks other components not being protected by diversity. At that point, in a game theoretic sense, additional diversity is wasted given that the attacker already has a higher payoff elsewhere. By exploring the game-tree, CSG portfolios are always balanced to reduce overall risk.

When used with a decision target (i.e., reduce the risk by 70%) or a cost threshold (i.e., how much money one wants to spend), CSG can be used to determine whether the available set of defense methods can achieve some risk reduction target and identify the optimal set of defense methods needed to reach that target. If there is no a priori decision target, CSG can also be used to perform an entire portfolio analysis. This makes it possible to identify the Pareto frontier allowing the user to identify the optimal set of defense methods to use at each given price point, and to understand if and when there is a diminishing return on their investment.

THIS PAPER HAS BEEN APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED. CASE NUMBERS 15-3140 AND 16-3240

REFERENCES

[1] Cox, A., Some limitations of risk = threat × vulnerability × consequence for risk analysis of terrorist attacks. *Risk Analysis*, **28**(6), pp. 1749–1761, 2008. https://doi.org/10.1111/j.1539-6924.2008.01142.x

[2] Lagner, R., *To Kill a Centrifuge*. The Langner Group, 2013.

[3] United States Senate Committee on Commerce, Science, and Transportation, *A "Kill Chain" Analysis of the 2013 Target Data Breach*. United States Senate , Washington, D.C., USA, 2014.

[4] Garvey, P.R. & Patel, S.H., *Analytical frameworks to assess the effectiveness and economic-returns of cybersecurity investments*. Military Communications Conference (MILCOM), 2014 IEEE, Baltimore, MD, USA, 2014.

[5] Carin, L., Cybenko, G. & Hughes, J., Cybersecurity strategies: the QuERIES methodology. *Computer,* **41**(8), pp. 20–26, 2008. https://doi.org/10.1109/mc.2008.295

[6] Roy, S., Ellis, C., Shiva, S., Dasgupta, D., Shandilya, V. & Wu, Q., *A Survey of Game Theory as Applied to Network Security*. 43rd Hawaii International Conference on System Sciences (HICSS), Koloa, HI, USA, 2010.

[7] Temin, A. & Musman, S., *A Language for Capturing Cyber Impact Effects. MITRE Technical Report MTR-10344*. MITRE Corporation, Washington DC, 2010.

[8] Musman, S., Temin, A., Tanner, M., Fox, F. & Pridemore, B., *Evaluating the Impact of Cyber Attacks on Complex Missions*. 5th International Conference on Information Warfare and Security, Dayton, OH, USA, 2010.

[9] Musman, S.T.A., *A Cyber Mission Impact Assessment Tool*. Homeland Security Technologies Conference, Boston, MA, 2015.

[10] Jajodia, N.S., Topological vulnerability analysis. *Cyber Situational Awareness, Advances in Information Security,* 2010.

[11] Wang, L., Jajodia, S. & Noel, S., *k-zero day safety: Measuring the security risk of networks against unknown attacks*. European Symposium on Research in Computer Security, Athens, Greece, 2010.

[12] Musman, S., Tanner, M., Elsaesser, E. & Lewis, L., *A Systems Engineering Approach to Crown Jewels Estimation and Mission Assurance Decision Making*, Proceedings of the IEEE Symposium on Computational Intelligence in Cyber Security, Paris, France, 2011.

[13] Musman, S., Tanner, M., Temin, A. & Elsaesser, E., *Computing the Impact of Cyber Attacks on Missions*. 2011 IEEE International Systems Conference (SysCon), Montreal, QC, 2011.

[14] Musman, S. & Agbolosu-Amison, S., *A Measurable Definition of Resiliency Using "Mission Risk" as a Metric*, MITRE Corp, McLean, VA, USA, 2014.

[15] Dhanjani, N., Rios, B. & Hardin, B., *Hacking: The Next Generation*, O'Reiley Media Inc., Sebastopol, CA, 2009.

[16] Noel, S., Ludwig, J., Jain, P., Jhonson, D., Thomas, R., McFarland, F., King, B., Webster, S. & Tello, B., *Analyzing Mission Impacts of Cyber Actions (AMICA).* ATO IST-128 Workshop on Cyber Attack Detection, Forensics and Attribution for Assessment of Mission Impact, Istanbul, Turkey, 2015.

[17] Nguyen, N., Alpcan, T. & Basar, T., *Stochastic games for security in networks with interdependent nodes*. International Conference on Game Theory for Networks, Istanbul, Turkey, 2009.

[18] Jormakka, J. & Jolsa, J., Modelling information warfare as a game. *Journal of Information Warfare,* **4**(2), pp. 12–25, 2005.

[19] Sallhammar, K. & Knapskog, S., *Using Game Theory in Stochastic Models for Quantifying Security*. Proceedings of the 9th Nordic Workshop on Secure IT-systems, Espoo, Finland, 2004.

[20] MSM, "Making Security Measurable," [Online], available at https://makingsecuritymeasurable.mitre.org/, 8 July 2013.

[21] Alpcan, T. & Basar, T., *A game theoretic analysis of intrusion detection in access control systems*. 43rd IEEE Conference on Decision and Control, Nassau, Bahamas, 2004.

[22] Watters, J., Morrissey, S., Bodeau, D. & Powers, S., *The Risk-to-Mission Assessment Process (RiskMAP): A Sensitivity Analysis and an Extension to Treat Confidentiality Issues*. MITRE, McLean, VA, USA, 2009.

[23] S. Noel, J. S, O. B. & J. M., *Efficient minimum-cost network hardening via exploit dependency graphs*. Proceedings of the 19th Annual Computer Security Applications Conference *(ACSAC'03)*, 2003.