Dissecting Cyberadversarial Intrusion Stages via Interdisciplinary Observations

Aunshul Rege
Criminal Justice
Temple University
Philadelphia PA USA
rege@temple.edu

Shanchieh Yang
Computer Engineering
Rochester Institute of Technology
Rochester NY USA
Jay.Yang@rit.edu

Alyssa Mendlein
Criminal Justice
Temple University
Philadelphia PA USA
alyssa.mendlein@temple.edu

Katorah Williams
Criminal Justice
Temple University
Philadelphia PA USA
katorah.williams@temple.edu

Shao-Hsusan Su
Computer Engineering
Rochester Institute of Technology
Rochester NY USA
ss9382@rit.edu

Stephen Moskal
Computer Engineering
Rochester Institute of Technology
Rochester NY USA
sfm5015@rit.edu

ABSTRACT

Advanced Persistent Threats (APTs) are professional, sophisticated threats that pose a serious concern to our technologically-dependent society. As these threats become more common, conventional response-driven cyberattack management needs to be substituted with anticipatory defense measures. Understanding adversarial behavior and movement is critical to improve our ability to proactively defend. This paper focuses on understanding adversarial movement and adaptation using a case study from a real-time cybersecurity exercise. Through multidisciplinary methodologies from social and hard sciences, this paper presents a mechanism to dissect cyberadversarial intrusion chains to unpack movement, and adaptations.

KEYWORDS

Advanced Persistent Threats (APTs); cyber adversary behavior; intrusion chain; qualitative field research; intrusion alerts

ACM Reference format:

Aunshul Rege, Shanchieh Yang, Alyssa Mendlein, Katorah Williams, Shao-Hsusan Su, and Stephen Moskal. 2020. Dissecting Cyberadversarial Intrusion Stages via Interdisciplinary Observations. In *Proceedings of the Sixth International Workshop on Security and Privacy Analytics (IWSPA '20)*, March 18, 2020, New Orleans, LA, USA. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3375708.3380317>

1. Introduction

Advanced Persistent Threats (APTs) are a real concern in today's technology-driven society, putting intellectual property, sensitive information, critical infrastructures, and national security all at

risk. These threats are defined by adversarial intent, opportunity, and capability [1]. APTs can use simple tools or escalate their Techniques, Tactics, and Procedures (TTPs) as they see necessary [2]. As suggested by their names, these attackers are persistent and often engage in repeated, coordinated attacks to achieve eventual success [3]. In addition, APTs are professional in their planning and are not deterred by obstacles, instead being capable of adapting their approach [2]. To make matters worse, APT trends for 2016 suggest that these threats are becoming increasingly common, dynamic, and deceptive [4].

Conventional cyber attack management, which is typically response-driven, has proved to be an inadequate approach for combating APTs [5]. Instead, there is an acknowledged need for anticipatory defense measures to enhance the capacity of defending forward [6, 7]. In order to develop these measures, though, there needs to be an understanding of the human element of APTs, and there has been little focus on this aspect of attacks in the open literature.

This paper focuses on a multidisciplinary methodological approach, weaving together human behavior from observations and technical logs, to better understand adversarial movement, and adaptation. The authors note that this research is exploratory (and thus preliminary) in nature as this mixed-methods approach offers a unique perspective in studying human behavior in cyberattacks. The next section of the paper details cyber intrusion chains, or how attackers progress through cyberattacks. The third section details the multidisciplinary data collection methodology at a real-time cybersecurity exercise, the Collegiate Penetration Testing Competition, and addresses assumptions and limitations of the data. In the fourth section, the temporal breakdown of the intrusion chains is provided using both observed and technical data, and highlight discrepancies that arise between these two data points. Finally, this paper offers recommendations for future research.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

IWSPA '20, March 18, 2020, New Orleans, LA, USA

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7115-5/20/03...\$15.00

<https://doi.org/10.1145/3375708.3380317>

2. Cyber Intrusion Chains

There have been many models of cyber attacks throughout the last decade, such as those by Lockheed Martin and SANS [8]. However, the model that the authors feel is most appropriate for the depth and complexities of cyber crime as a human action is DELL's Lifecycle of an Advanced Persistent Threat, shown in Figure 1 [2]. This model breaks the cyberattack process, or intrusion chain, into a series of 12 stages that are sequential in nature.

Figure 1: DELL's Intrusion Chain Model [2]

In stage 1, adversaries identify and select their targets. Stage 2 is marked by the assembly of a team with complementary or supplementary skills. As adversaries transition into stage 3, they begin the acquisition of the necessary tools and begin to build the vectors to be used in the attack. By stage 4, the adversarial team is working to identify infrastructure weaknesses and deploy social engineering tactics to obtain entry. This stage can be marked by the adversaries conducting a series of scans to find structural vulnerabilities. Stage 5 moves from a focus on pathways into the system to the examination of the target's security and defense systems. Identifying these systems is essential to infiltrating the system and being as discreet as possible. Additionally, examining these defense systems enables adversaries to create the necessary evasion and response plans [2]. Stage 6 is marked by a strengthening of the adversaries' foothold in the target's system. This is achieved by using the earlier developed attack vectors, or other skills. Once the adversaries' foothold is strong in the target's system, they can move to deploy and install malware into the environment in stage 7. Adversaries establish additional access points in the targeted environment in stage 8, while obtaining credentials to gain greater system access and control in stage 9. During stage 10, adversaries strengthen their presence by expanding laterally and deeper into the targeted environment. This expansion wider and further into the target's system allows the adversaries to have control over multiple parts of the target's system, thus improving their ability extract data. In stage 11, adversaries exfiltrate data or disrupt functionality to accomplish

their objectives. Finally, in stage 12, adversaries clean up and remove traces of their presence in the targeted environment. Note that the 'usage' of each stage may vary from one adversarial team to another and when the attack progresses further. This is expected since the intrusion chain describes a sequence of sub-objectives, together forming a potentially orchestrated overall attack plan.

3. Using Cybersecurity Exercises to Study Human Behavior

Cyber defense competitions were first developed in certain divisions of U.S. military service academies as a way to test the network defense skills of their students; shortly thereafter, other sectors realized the benefits of these competitions for all types of university students [9]. In the United States, there are four critical infrastructure cybersecurity exercises: US Army Research Lab, ICE-CERT/INL, US Cyber Storm, and Alphaville [10]. These exercises are ideal to train and educate current and future operators, owners, and users of critical infrastructure on how these systems are subjected to cyberattacks; how to defend these systems in real time; how to manage limited employee and monetary resources during and after cyberattacks; and how to better manage system confidentiality, integrity, and availability (CIA) [10].

Additionally, other cybersecurity exercises have developed to serve student populations. The Collegiate Cyber Defense Competition was the first of its kind among non-military students, and this type of competition has been growing in popularity ever since its creation [11]. These competitions have been shown to provide many benefits for students, such as bolstering skills and techniques for defending a network and learning to adapt and work in groups [9, 11].

The Collegiate Penetration Testing Competition (CPTC) consists of regional and nationwide competitions that seek to develop the skills required to "effectively discover, triage, and mitigate critical security vulnerabilities" by mimicking the cyber security testing done for real-world organizations [12]. This "pentesting" involves "simulating real attacks to assess the risk associated with potential security breaches" by understanding system vulnerabilities that could be exploited and what attackers might gain through successful exploitation [13, p. 1]. The CPTC structure is unlike other collegiate cybersecurity competitions, as each student team professionally penetrates an instantiation of the same network and reports the findings from these vulnerabilities back to the organizers, who pretend to be the company executives [12]. Regional competitions are held at universities across the nation and are two days, with penetration testing on the first day and reporting and results on the second day [12]. The top team from each region and the highest-ranked teams overall advance to the national competition, which is held in Rochester, NY. The three-day national final competition starts with entrance meetings for fictitious leadership of the organization on day one, followed by penetration testing on day two and reporting and exit meetings on day three [12].

Interestingly, such competitions are great platforms for researching adversarial behavior, group dynamics, decision-making and adaptability, and movement along intrusion chains. Studying actual APTs is problematic. First, their activity is covert in nature and researchers do not typically have access to their organizational and operational dynamics. Second, researching APTs directly raises ethical concerns and is risky. Accessing APTs is difficult as they are part of an underground culture that may be unknown or inaccessible. For these reasons, cybersecurity exercises, such as the CPTC, are ideal settings to study behaviors and activities that otherwise remain out of reach. While still not representative of reality, as noted in section 4, these competitions still offer a setting to start understanding adversarial behavior and, equally importantly, how to study it in a meaningful manner.

4. Integrated Methodology

The results presented in this paper are based on observations of one team's activity during the penetration testing phase of the 2017 CPTC national competition, during which teams had to professionally attack and identify vulnerabilities on a network supporting voting stations and election tallies. The activities and behaviors of the participating team, which had six members, were observed through two complementary collection methods for the duration of the 10-hour competition. First, the researchers directly observed the actions of the team. Second, the researchers deployed Suricata intrusion detection sensors to record the technical observables of the team's intrusion activities. This complementary approach allowed the researchers to compare the actions conducted from the attacking team's perspective to the adversarial actions detected in the targeted network. To the best of our knowledge, no prior study has employed such an integrated interdisciplinary approach to investigate attacker behaviors.

Observations are a qualitative method used in the social science to study human behaviors. In the context of cyberadversarial behavior, qualitative methods are particularly useful to unpack the underlying processes and mechanisms of human interactions, group dynamics, and adversarial intrusion chains. The research team observed the team in real-time during the entire 10-hour competition and then later compiled into a time-stamped 'adversarial actions' transcript. The observed data were analyzed to extract three main components. First, the team member's actions were categorized into the various stages of the intrusion chain identified above. Second, the team's adaptations to hurdles were explored. Finally, the researchers analyzed the observations to understand group dynamics and divisions of labor.

The Suricata alerts recorded are interpretations of the malicious activities transmitted over the network by the participating team members. The alerts are first collected through industry standard Splunk indexing system and statistically analyzed by attempting to map them into the previously described intrusion chain. Note that there is no commonly accepted mapping from intrusion alerts to intrusion chains. In fact, there are several similar intrusion chains proposed by different organizations over time. A

common denominator of intrusion chain stages may be qualified as Macro-attack Stages as in Reconnaissance, Exploitation, and Exfiltration. Not all intrusion stages can be observed through intrusion detection systems, such as Suricata. For example, Stage 2 in Dell's Intrusion Chain – Find and Organize Accomplices – may not be a stage that can be observed through network traffic. For this work, we consider Stage 4 for the Dell Intrusion as Reconnaissance, Stages 6~10 as Exploitation, and Stage 11 as Exfiltration, and map Suricata alerts into the three Macro-attack Stages.

By comparing the Suricata alerts and the qualitative observations mentioned earlier, this work aims at finding complementary information that can enrich our understanding of cyberadversary behaviors and their limitations.

4.1 Limitations

As with any research, this one has two main limitations: exercise-related and methodology-related. The CPTC has certain limitations that are inherent to its exercise design and logistics.

First, the exercise is not representative of reality; it is expedited in nature. Real cyberattacks are very likely to occur over extended time periods. Second, the participants are students and, as such, may not serve as an ideal stand-in for real, experienced cybercriminals. Research has indicated cybercriminals as having a high level of sophistication, intelligence, adaptation, and persistence [1, 2, 4, 6, 7]. Cybersecurity exercise participants, such as those in CPTC, may thus not represent adversaries who are as sophisticated and adaptive in their operations.

One methodological limitation is related to the single instance of real-time physical, in-person observations. These conclusions are based on the actions and behaviors of only one participating university team, and different groups may have exhibited different behaviors, thus leading to different findings. In addition, the real-time observations may not have captured all the relevant behaviors exhibited by the group. The Suricata intrusion alerts can only reflect the observable actions transmitted over the network. In other words, they do not explain why the attacker perform the action, and they do not reflect the team member discussions and activities on the local computer that are legitimate activities that do not trigger intrusion alerts.

Another important methodological limitation to note is that intrusion chain models are meant to depict the progression of cyberattacks, and are not made for mapping to technical observables such as Suricata intrusion alerts. For the intrusion chain discussed above [2], the Suricata alerts can only be approximately mapped into three categories: Reconnaissance (Research Targets), Exploitation (Initial Intrusion, Outbound Connection Initiated, Expand Access & Obtain Credentials, and Strengthen Foothold), and Exfiltration (Exfiltrate Data).

A third methodological limitation is that it is not guaranteed that each source IP represents one team member, as observations

suggested that team members often moved about during the exercise and used each other's machines. A member could thus use zero, one, or multiple computers during the exercise.

While these limitations are valid and limit the generalizability of the findings, they are still useful to (i) study human behavior, decision-making, adaptation, and group dynamics [8], (ii) explore how different methodologies might complement and supplement each other, and (iii) identify how each discipline-specific methodology can learn from the other.

5. Findings

5.1 Intrusion Chain Analysis via In-Person Observations

There were three main analyses for the team's performance along the intrusion chain: (i) overall time spent on the various intrusion chain stages by the entire team, (ii) temporal distribution of emphasis at different times during the competition, and (iii) time spent on the various intrusion chain stage by each team member.

5.1.1 Overall time spent by entire team on various intrusion chain stages. Overall, the temporal analysis for the observed team indicates that the most of the team's time (49%) was spent in the planning stages of organizing accomplices (stage 2), acquiring tools (stage 3) and researching the target (stage 4), as seen in Figure 2. In particular, just researching the target took up 37% of the team's time overall.

It is not possible to tell from this analysis, though, whether these were all initial research, or research done once later stages had been reached and the group realized it had insufficient information. The middle stages of the intrusion chain made up for 35% of the time, including deployment (stage 6), initial intrusions (stage 7) and expanding access (stage 9). Only 11% of the team's time was spent exfiltrating data (stage 11).

There were a few stages in which the team spent very little time; the team spent 2% of their time testing for detection (stage 5), about 1% of their time initiating outbound connection (stage 8) and less than 1% was spent covering tracks (stage 12). This may be related to the structure of the exercise rather than the amount of time a real-world hacker would spend conducting these tasks. Participants may not be as concerned with being detected as an actual cyber hacker would be, because the team's goal is to just uncover all the vulnerabilities and understand what information would be exploited [13].

5.1.2 Temporal distribution of the entire team's emphasis at different times during the competition. Figure 2 presents the team's time spent in different stages at different parts of the day. There is more time spent in earlier stages in the morning, and more time spent in later stages in the afternoon.

A significant amount of time in the morning was spent in the planning stages (73%). Only 23% of the morning time is spent on the middle stages, stages 6 through 9. In the afternoon, this temporal stage allocation shifts to 48% of the team's time being

focused on the middle stages. However, the planning stages still encompass 25% of the time, and exfiltrating data (stage 11) takes up 21%.

Figure 2: Time Spent by Observed Team on Intrusion Chain Stages Overall and in the Morning (Yellow) and Afternoon (Blue)

5.1.3 Time spent on the various intrusion chain stage by each team member. The group's division of labor in intrusion stage minutes can be seen in Figure 3. In terms of researching targets (stage 4), where most of the team's time was spent overall, each member spent the same amount of time on this stage. However, in other stages, such as 2, 9 and 11, it is clear that certain group members spent more time with these tasks than others.

Figure 3: Time Spent (in Minutes) by Each Member of Observed Team on Intrusion Chain Stages

Team members 0 and 5 used most of their time in the planning stages (2 through 4), while member 2 spent a majority of his time researching targets (stage 4), expanding access (stage 9) and exfiltrating data (stage 11). Member 3 used the least time overall to complete tasks. This temporal distribution could be related to the skill sets of, and actions performed by, the various team members as observed by the researchers as noted in Table 1.

The results provided above are interesting. Forty-nine percent of the team's time was spent on reconnaissance and planning, which aligns with previous study temporal allocations from other cybersecurity exercises [8]. This finding reflects that real cybercriminals do indeed spend considerable amounts of time planning and conducting reconnaissance [8].

In addition, the division of labor results might be helpful in understanding how groups operate together. In the case of the observed team, members worked together to research targets (stage 4) and expand access (stage 9), and both of these stages combined accounted for a majority of the team's time. For less time-consuming stages, the team divided the labor less evenly, and appear to have specialized in certain areas.

Team member	Skills/actions observed
0	scanning, reverse engineering, bruteforcing, escalating privileges, network mapping, and authentication schemes
1	Scanning, general networking, reverse engineering, bruteforcing, and exploiting web applications
2	reverse engineering, scanning, and bruteforcing
3	Scanning
4 & 5	linux, databases, hashing, scanning, bruteforcing, advanced exploit techniques, reverse engineering, and exploiting web applications.

Table 1: Team Member and Skills/Actions Observed

5.2 Mapping Intrusion Alerts to Macro-Level Intrusion Chain Stages

A total of 47,876 non-duplicate intrusion alerts were collected for all ten participating teams during the national final. Alerts that have exactly identical attributes were consolidated into a single one to prevent exaggeration of intrusion activities. Note that duplicate alerts can happen due to multiple recordings of the same activity. Also note that a single intrusion activity can also trigger multiple different alerts, but this study does not perform complex alert aggregation, to avoid over-simplification of technically observed activities. Out of the 47,876 alerts, 45,472 (94.88%) are estimated to be Reconnaissance (Stage 4), 2,298 (4.8%) for Exploitation (Stages 6~10), and 151 (0.32%) for Exfiltration (Stage 11). For the observed team, there were a total of 4,272 alerts, out of which 4,229 (98.99%), 20 (0.47%), and 23 (0.54%) were for Reconnaissance (Stages 2-4), Exploitation (Stages 6-10), and Exfiltration (Stage 11), respectively.

The observed team's exploitation activities (Stages 6-10) were observed relatively less by Suricata when comparing to the overall percentage spread for all teams. However, this is clearly not the case as compared to the 'usage' in minutes reported in Section 5.1.

According to the in-person observations, the team members spent approximately the same amount of time in Stages 6-10 as in Stage 4, at the 35% level. Even for Stage 11, the in-person observations

accounts for 11% of time spent, but there were less than 1% of intrusion alerts corresponding to such activities. Note that this does not mean either of the two observation methods is incorrect. One possible explanation is that the significant human effort spent in the Exploitation and Exfiltration stages captured during observations may only lead to a small number of actual malicious actions being executed by the team, and even fewer of these actions may have triggered the intrusion alerts.

Figure 4 shows a breakdown of intrusion alerts collected for the 6

Figure 4: Number of Suricata Alerts Collected for the Participating Team, Broken Down by Source IPs Over Time

source IPs used by the participating team members over time. Note that the students in this competition were not trying to 'cover their tracks' and thus no random or fictitious source IP tactics were used.

One may see that there are significantly more alerts in the morning of the competition (until Hour 4) while there were comparable amounts of time spent between morning and afternoon as shown in Figure 3. This difference is because while it was observed that the team spent more time on the later stages of the intrusion chain, there were fewer actual malicious exploitation and exfiltration activities being transmitted and observed over the network. Also note that the SrcIP 1 triggers more alerts (mostly Reconnaissance types) early on, but SrcIP 2 was detected more for its Reconnaissance activities later on.

To have a clear comparison on how each SrcIP is observed by Suricata for the different macro-stages, Figure 5 shows the alerts normalized within each stage in each hour for each SrcIP. The total value in each hour can be 1.0, 2.0, and 3.0 depending on whether 1, 2, or all 3 macro-stages were observed. For the

Figure 5 : Sum of Alerts Observed for the Participating Team, normalized within each Stage and Broken Down by Source IPs Over Time

observed team, no Exploit activities were observed by Suricata after Hour 5, and Exfiltration activities were only observed in Hour 4 and 7. Reconnaissance activities were observed in all 10 hours.

While SrcIP 1 and SrcIP 2 were mostly used for Reconnaissance, SrcIP 3 accounted for 12 out of 20 Exploitation alerts in the first 5 hours. This suggests that SrcIP 3 was used to advance along the intrusion chain rather quickly. Referencing to Figure 4, this may suggest that SrcIP 3 was used by Competitor 2 or 4 (recalling Stages 6-10 are part of the Exploit Macro-stage). Meanwhile, SrcIP 6 accounted for 21 out of 23 alerts for Exfiltration. Interestingly, multiple competitors were observed (in-person) to perform Stages 11 and 12, so it is possible that SrcIP 6 was used to exfiltrate either by multiple competitors (0, 1, 2, and 5) or the actions were executed by one competitor with inputs from others.

5.3 Discussion on Integrated Multidisciplinary Methods: Intersections and Gaps

This paper highlights a unique integrated multidisciplinary methodology to study adversarial behavior. As noted in Section 4, the in-person observation data were analyzed to understand (i) the team member's actions in relation to intrusion chains, (ii) any team adaptations, and (iii) group dynamics. Sections 5.1-5.3 focused only on the first of these three components: mapping behaviors and actions to intrusion chains.

Comparing the in-person observations and intrusion alerts shows discrepancy in efforts observed versus volume of alerts collected. Both methods allow interpretation of how different attack stages were used by each team member or SrcIP. In a way, the multidisciplinary analysis shows great potential that was not viable via only the individual analysis. For example, the mappings between the competitors and the SrcIPs have revealed previously invisible collaboration between the competitors. With additional detailed analysis and data to be collected, one can potentially explain the team dynamics or even adaptation by cross-analyzing the two types of data.

The current work through CPTC'17 has its limitations. In terms of (ii) team adaptations, an in-depth temporal analysis of the in-person observations can reveal the key points in the competition when team had turning points (tried something new or adapted to roadblocks). Such rich insights into the struggles and hurdles experienced by the team could be used to 'zoom in' to those timeframes in the Suricata logs to further study the team's actions. Thus, the observations could help analyze the Suricata logs more efficiently, which is further discussed in Section 6.2.

In terms of (iii) team dynamics, the in-person observations can give unique insights into how tasks were divided amongst team members, how decisions were made, how sub-teams were created to handle 'mini' tasks, how team members helped each other when they were stuck or experienced failures. All these aspects of human interaction occur before and as the team executes actions (type commands at computer terminals). While such interactions

are not explicitly captured in intrusion alerts, comparing the competitor IDs and the SrcIP can shed light on the exact actions performed as a result of the discussions and decisions between the competitors. A detailed and systematic analysis to connect the SrcIPs and the observed competitor activities will help enable more conclusive team dynamic analysis.

6. Directions for Future Research

This paper has put forth the idea of understanding the complexity of cyber attacker decision-making through integrated-methodologies. By understanding where team members spent their time and how they functioned, more effective strategies can be created to defend forward by preventing the progress along intrusion chains. This case study raises several lines of further inquiry.

6.1 Identifying and Measuring Intrusion Chains and Behavior

During a cyberattack, it is very likely that cyber adversaries use multiple intrusion chains simultaneously. However, both observations and technical logs cannot clearly identify which member actions correspond to various intrusion chains, which impacts our understanding of temporal stage allocation along the intrusion chain. An intrusion chain that reflects both adversary behavior as well as the technical observability will be helpful to further the study of cyber attack progression behavior.

This case study primarily used the metric of time to measure human actions and intrusion chains. This metrics, while useful, may not reflect the actual malicious activities or the technically detectable actions. Comparing the technical detectable intrusion activities with the human observables can inform the mismatch and, thus, offer a more comprehensive understanding of the cyberattack behavior and their effects. Several other metrics can be used, such as effort, specific objectives and motivation, and personalities and culture, to complement technical observables.

6.2 Adversarial Adaptations

Observations suggested that team members were able to adapt in some situations during the exercise, which may be indicative of their skill level and knowledge base, or their overall role in the exercise.

For instance, member 4 tried several ways to get reverse shells and member 5 attempted to bruteforce and get remote code executions on various occasions. Interestingly the same two competitors along with member 0 also tried new targets, which is suggestive of their skill level and knowledge. Members 1, 2, and 3 did not show any clear indications of adapting during the exercise, with the latter being a novice who looked to other team members for guidance.

One useful recommendation for future research would be to identify how adaptations can be identified in technical logs, and how these, in turn, could be aligned with the observed behavior data. This might further shed light on the amount of overlap in

the two data points. Furthermore, this research might demonstrate how much of the decision-making process to pursue adaptation occurs in both the observed and log data.

6.3 Injecting Variations into Exercises

During this exercise, the participants were not working against a rival or defender; would their behavior and movement along the intrusion chain be different if they were working against another competing team or a defending team? Additionally, how would the behavior change if the competitors were working on extended time frames similar to a real attack? How would their behavior and the outcome of the event change if they had a different group dynamic? By placing various constraints, environmental changes, and altering the structure of the group, more in-depth observations could be made to study how these groups operate in various environments.

6.4 Multidisciplinary Integrated Methods

Despite quantitative and experimental methods being the preferred strategy for causal investigations [14, 15, 16], this research shows the contributions that qualitative methods can provide in isolation, and in conjunction with technical methods. Qualitative research answers questions about the underlying “why?” or “how?” of cyberattack processes [14, 15, 16]. In doing so, qualitative methods are ideal for better understanding the behavior and movement of adversaries in cyberattacks [16], by unpacking mechanisms and offering insights into relationships that cannot be explored otherwise [14].

Other research should also combine this multidisciplinary approach to answer questions about adversarial decision-making in cyber attacks. This approach could ensure that important insights are not overlooked or separated within disciplines or methods of analysis. It is essential to understand adversarial decision-making if we are going to create anticipatory responses that improve our ability to defend forward.

ACKNOWLEDGMENTS

This work is supported through National Science Foundation EAGER Awards # 1742747 and #1742789 and partially through

CAREER Award #1453040. The authors thank the 2017 Collegiate Penetration Testing Competition for letting them collect observation and log data. The authors also thank Temple University’s ethics board for reviewing the human subject aspects of this research.

REFERENCES

- [1] Ingoldsby, Terrance R. "Attack tree-based threat risk analysis." *Amenaza Technologies Limited* (2010): 3-9.
- [2] Dell. (2012) "Lifecycle of an Advanced Persistent Threat." Retrieved from <http://www.findwhitepapers.com/force-download.php?id=27730> (2012).
- [3] RSA. (2012). Stalking the kill chain. Retrieved from <http://www.emc.com/rsa>.
- [4] F-Secure. (2016). 5 Advanced Persistent Threat trends to expect in 2016. Retrieved from <https://business.f-secure.com/5-advanced-persistent-threat-trends-to-expect-in-2016>.
- [5] Cloppert, Mike. "Security Intelligence: Attacking the Cyber Kill Chain." *SANS Computer Forensics* (2009).
- [6] Hutchins, Eric M., Michael J. Cloppert, and Rohan M. Amin. "Intelligence-driven computer network defense informed by analysis of adversary campaigns and intrusion kill chains." *Leading Issues in Information Warfare & Security Research* 1, no. 1 (2011): 80.
- [7] Barnum, Sean. "Standardizing cyber threat intelligence information with the structured threat information expression (stix)." *Mitre Corporation* 11 (2012): 1-22.
- [8] Rege, A., E. Parker, B. Singer, and N. Masceri. "A qualitative exploration of adversarial adaptability, group dynamics, and cyber-intrusion chains." *Journal of Information Warfare* 16, no. 3 (2017): 1-16.
- [9] Carlin, Anna, Daniel P. Manson, and Jake Zhu. "Developing the Cyber Defenders of Tomorrow with Regional Collegiate Cyber Defense Competitions (CCDC)." *Information Systems Education Journal* 8, no. 14 (2010): n14.
- [10] Rege, A. & Adams, J. "The need for more sophisticated cyber-physical systems war gaming exercises". (2019). *Proceedings of the 18th European Conference on Cyber Warfare and Security*.
- [11] Bei, Yan, Robert Kesterson, Kyle Gwinnup, and Carol Taylor. "Cyber defense competition: a tale of two teams." *Journal of Computing Sciences in Colleges* 27, no. 1 (2011): 171-177.
- [12] CPTC. (n.d). Collegiate Penetration Testing Competition. Retrieved from <https://nationalcptc.org/>
- [13] Weidman, Georgia. *Penetration testing: a hands-on introduction to hacking*. No Starch Press, 2014.
- [14] Maxwell, Joseph A. "Causal explanation, qualitative research, and scientific inquiry in education." *Educational researcher* 33, no. 2 (2004): 3-11.
- [15] Howe, Kenneth R. "A critique of experimentalism." *Qualitative inquiry* 10, no. 1 (2004): 42-61.
- [16] Prowse, Martin, and Laura Camfield. "Improving the quality of development assistance: What role for qualitative methods in randomized experiments?." *Progress in Development Studies* 13, no. 1 (2013): 51-61.