

Using a Real-Time Cybersecurity Exercise Case Study to Understand Temporal Characteristics of Cyberattacks

Aunshul Rege¹(✉), Zoran Obradovic², Nima Asadi², Edward Parker¹,
Nicholas Masceri¹, Brian Singer¹, and Rohan Pandit¹

¹ Department of Criminal Justice, Temple University, Philadelphia, USA
rege@temple.edu

² Computer and Information Sciences Department,
Temple University, Philadelphia, USA

Abstract. Anticipatory cyber defense requires understanding of how cyber adversaries make decisions and adapt as cyberattacks unfold. This paper uses a dataset of qualitative observations conducted at a force on force (“paintball”) exercise held at the 2015 North American International Cyber Summit (NAICS). By creating time series representations of the observed data, a broad range of data mining tools can be utilized to discover valuable verifiable knowledge about adversarial behavior. Two types of such analysis discussed in this work include clustering, which aims to find out what stages show similar temporal patterns, and peak detection for adaptation analysis. Collectively, this mixed methods approach contributes to understanding how adversaries progress through cyberattacks and adapt to any disruptions they encounter.

Keywords: Adaptive human behaviour · Dynamic decision making · Temporal analysis · Time series data · Clustering · Field research

1 Introduction

Today’s information networks and integrated systems are highly networked, thereby increasing the attack surface, resulting in greater cyberattacks [2]. Yet, conventional cyberattack management is reactionary and does not capture Advanced Persistent Threats (APTs), which increasingly target critical infrastructures and consistently circumvent traditional security measures, resulting in large and costly damages [1]. It is therefore essential that commercial and government organizations develop defenses which are able to respond rapidly to, or even foresee, new attack strategies and tactics [2]. While many important contributions in anticipatory/proactive cybersecurity have been made, they are technical in nature and downplay the relevance of the human agents behind the cyberattacks, and their decision-making processes and adaptation strategies [2].

This paper employs quantitative data science methods of time series analysis to assess the observed adversarial behavior at a force on force (“paintball”) exercise. Collectively, this mixed method contributes to understanding

how adversaries progress through cyberattacks and adapt to any disruptions they encounter. This paper is structured as follows. Section 2 outlines the mixed methodology of observations and time series analysis. Next, the computational results are discussed. Finally, this paper discusses relevant findings and possible implications for adversarial movement and adaptability.

2 Methodology

In the Criminological discipline, crime scripts provide a systematic understanding of the crime commission processes [3]. The applications of crime scripts to cyberattacks as they unfold remains understudied. In the technical domain, crime scripts appear as intrusion chain models that capture the step-by-step process of cyberattacks. While there are many models of adversarial intrusion chains, we use the 12-step cyber intrusion chain model in [1], as it offers detailed attack stages that allow for thorough data analysis.

The Merit Network and the Michigan Cyber Range provide a virtual platform called Alphaville, which is used for cybersecurity training exercises. Alphaville emulates a typical city and consists of five locations: a school, a library, a city hall, a small business, and a power company, each of which contains servers and firewalls with intentional vulnerabilities. During the 2015 North American International Cyber Summit (NAICS), the researchers observed a five-hour force on force “paintball” exercise, where teams battled to claim Alphaville’s network by controlling critical servers. Researchers observed one of the teams participating, which consisted of four members (henceforward referred to as Subjects S1, S2, S3, and S4).

Temporal analysis aims to extract and characterize the trends, patterns, and variations within a process over time using time series data. In order to create the time series, the timestamped observations of the team’s actions and their durations were utilized. In this work, each time point in the generated time series represents a one minute time span. For each time point, the value of each time series is the accumulated number of minutes spent by the entire team on its corresponding intrusion stage. After creating the time series representation of the data, we performed temporal analyses of the intrusion process through data mining methods, namely, clustering and peak detection. We performed clustering of the time series in order to achieve a verifiable measurement of co-activation and co-dependence of intrusion stages. Clustering allows similar time series (measured by comparing the amplitude of the time series, which is the total amount of time in minutes allocated to each intrusion stage during each minute of the exercise practice) to be placed in groups. A high similarity between the time series of the intrusion stages A and B is an indication that whenever intrusion stage A was performed within a time point, the possibility of performing stage B during that time point was higher than any other intrusion stage.

In this work, we use Agglomerative hierarchical clustering [4]. The reason behind choosing this clustering model is its power in providing the order and

similarity hierarchy of the clusters, and the fact that no a priori information about the number of clusters to be made is required.

To understand the team's adaptation measures when facing disruptions, the time series were then analyzed for detecting local peaks after these disruptions occurred. We employed a Peak-Valley detection algorithm [5] to detect the 'adaptation stages' by finding peak values that were above the global mean (average of the mean of all time series amplitudes), which were separated from those stages that the red team spent minimal time on (peak values below global mean).

3 Results

3.1 Observed Duration of Adversarial Intrusion Chain Stages

The observed data summarized at Fig. 1 suggest that the team spent approximately 49% (140 min) of the exercise time on entering the system, establishing foothold, and moving laterally to gain further control over systems. This was followed closely by Reconnaissance (stages 2, 3, 4, and 5), which took up roughly 44% of the exercise time or 125 min. The researchers did not find other intrusion stages during observations and so these stages are excluded from further analysis.




Fig. 1. Total time spent by the red team on each intrusion stage through the entire exercise

3.2 Time Series Generation and Clustering

Figure 2 shows the time series created for each intrusion stage. The clustering results are provided in Fig. 4, and an example of temporal pattern similarities is provided in Fig. 3. In Fig. 4, the vertical axis corresponds to the Euclidean distance of time series pairs. The clustering threshold, which determines the stages that are grouped together, was selected at the middle of the largest distance, which results in the red threshold line in Fig. 4. The results indicate the temporal similarities among intrusion stages; for instance, the occurrence of intrusion stage 3 (a peak in its time series), is more likely to be accompanied by the occurrence of the stages 4, 5, and 7 than any other intrusion stages.




Fig. 2. Time series representation of the observational data. The arrows at the bottom show the disruptions corresponding to Table 1. (Color figure online)




Fig. 3. An example of the similarities among time series; the peak/valley patterns happen more concurrently between intrusion stage pairs 3 and 5 compared to pairs 3 and 9, or pairs 5 and 9.

3.3 Analysis of the Adaptation Process

The local peaks of the time series and the global mean depicted by the horizontal red threshold (global mean of 0.908 min total engagement per one minute interval) can be observed in Fig. 2. We can observe that within the 10 min time frame after a disruptive event, the amount of time allocated to certain intrusion stages was above this threshold at multiple intervals, indicating that the red team focused more on these stages in response to that disruption. For instance, in Fig. 2, we observe a spike in stage 2 (Find/Organize Accomplices) after the 8:40 access failure disruption (detailed in Table 1). Possible explanations for




Fig. 4. Hierarchical clustering of the time series where each number corresponds to the intrusion stage number (Color figure online)

Table 1. Possible explanation for time spent on certain stages post disruptions

Time	Player	Hurdle	Disruption details (S)	Spiked stage (Mins Spent)	Stage sequence	Possible explanations for spike in the stage
8:33	S2	L	S3 Kills S2 attack chain	5(2), 3(1)	Con current	To test the targeted system's intrusion detection measure (spike in stage 5), the team was deciding which tools to use (spike in stage 3)
8:40	S2	L	S3: why do I keep losing my shell?	2(3), 3(1)	2, 3	Team member lost access, so may have sought help from other members (spike in stage 2) about which tools to use (spike in stage 3)
8:50	S3	S	S3 has a failed login attempt	7(1), 3(1)	7, 3	Team was possibly in stage 7 (spike in stage 7), moving laterally to strengthen foothold, but to gain access, may have tried different tools (spike in stage 3)
9:44	S2	L	S2 tries to get into the system	2(1.5), 3(1), 5(1)	2, (3, 5 concurrent)	Team member may be unsuccessfully trying to gauge target's defense measures (spike in stage 5) and hence may have sought help from other team members (spike in stage 2) about which tools to use (spike in stage 3)

disruptions and responses are provided in Tables 1, but these cannot be conclusive as they are based solely on observations, and as such, cannot account for the team's decision-making processes and dynamics.

4 Conclusion

There are some unavoidable limitations to this research such as generalizability and the fact that the case study is not representative of real cyberattacks.

However, the authors make the case that this paper is exploratory, methodologically unique, and based on one of the most reputable force on force (“paintball”) exercises in the United States.

The time series analysis offers some interesting findings about the adversarial intrusion chains:

“Dispersed Spikes May Indicate Nonsequential Progression of Intrusion Stages”. The greatest cumulative spike occurred for stage 2, but these spikes occurred at different times (Table 1, 8.40 and 9.44). This suggests that adversaries exhibit complex back and forth movement when they face disruptions.

“Parallel Stages and Stage 3 (Build/Acquire Tools)”. After each disruption the team focused on multiple stages at either the same time (concurrent) or with a slight temporal lead (Table 1), suggesting that stages occur in parallel rather than in sequence. Also after each disruption, Stage 3 always occurred in parallel with other stages (Table 1), which suggests that building/acquiring tools may be a relevant stage during most adaptations.

Accessing Systems is Key across Multiple Stages. Most disruptions (Table 1: 8.40, 8.50, and 9.44) were related to difficulties in gaining or maintaining access to target systems, which was an issue at multiple stages.

Acknowledgements. This material is supported by the National Science Foundation (NSF) CAREER Award No. 1446574 and partially by NSF CPS Award No. 1453040. The authors thank the Merit Network and the Michigan Cyber Range for allowing data collection at their 2015 NAICS event.

References

1. Cloppert, M.: Attacking the cyber kill chain. <http://digital-forensics.sans.org/blog/2009/10/14/security-intelligence-attacking-the-kill-chain>. Accessed 2 Feb 2014
2. Colbaugh, R., Glass, K.: Proactive Defense for Evolving Cyber Threats. Sandia National Laboratories [SAND2012-10177] (2012). <https://fas.org/irp/eprint/proactive.pdf>. Accessed 15 Feb 2017
3. Leclerc, B.: Crime scripts. In: Wortley, R., Townsley, M. (eds.) *Environmental Criminology and Crime Analysis*. Routledge (2016)
4. Rokach, L., Maimon, O.: Clustering methods. In: Maimon, O., Rokach, L. (eds.) *Data Mining and Knowledge Discovery Handbook*, pp. 321–352. Springer, New York (2005)
5. Schneider, R.: Survey of Peaks/Valleys identification in Time Series. University of Zurich, Department of Informatics, Switzerland (2011)