

# Predictions of Network Attacks in Collaborative Environment

Martin Husák, Pavel Čeleda  
Institute of Computer Science,  
Masaryk University, Brno, Czech Republic  
husakm@ics.muni.cz, celeda@ics.muni.cz

**Abstract**—This paper is a digest of the thesis on predicting cyber attacks in a collaborative environment. While previous works mostly focused on predicting attacks as seen from a single observation point, we proposed taking advantage of collaboration and exchange of intrusion detection alerts among organizations and networks. Thus, we can observe the cyber attack on a large scale and predict the next action of an adversary and its target. The thesis follows the three levels of cyber situational awareness: perception, comprehension, and projection. In the perception phase, we discuss the improvements of intrusion detection systems that allow for sharing intrusion detection alerts and their correlation. In the comprehension phase, we employed data mining to discover frequent attack patterns. In the projection phase, we present the analytical framework for the predictive analysis of the alerts backed by data mining and contemporary data processing approaches. The results are shown from experimental evaluation in the security alert sharing platform SABU, where real-world alerts from Czech academic and commercial networks are shared. The thesis is accompanied by the implementation of the analytical framework and a dataset that provides a baseline for future work.

**Index Terms**—intrusion detection, alert correlation, information sharing, collaboration, prediction, situational awareness

## I. INTRODUCTION

Collaboration has emerged as an important topic in network security, namely in the form of information exchange [16, 17], collaborative intrusion detection systems [18, 19], collaborative correlation of intrusion detection alerts [20], and other areas [21]. The security teams incline to share the knowledge and experience, as well as timely information on current threats and attack, since the foundation of cyber security as we know it as of today. Sharing of timely information on current threats and attacks seems to be a promising form of early warning [22], which can be used to take preemptive measures to defend a network. It is no surprise that both researchers and practitioners have taken the initiative to create methods and build tools that would allow automated security information exchange and use them to increase the level of protection of the networks.

This paper is a digest of the thesis on predicting cyber attacks and the next step of an adversary in a collaborative environment [1], one of the promising use cases of security information exchange [19]. In the past, we have seen many attempts to predict the upcoming attacks or attack steps. However, a majority of such methods faced serious drawbacks that caused these methods never to reach technological maturity

and deployment in production. For example, they required a library of attack descriptions or an observation point with very detailed visibility into the network traffic and highly precise intrusion detection. It was, and still is, highly demanding to have a single observation point in the network that could provide every piece of information needed to detect and successfully predict the next event in an ongoing attack. Learning from the past to predict future attacks can also be problematic in the continuously evolving threat landscape. However, a collaborative approach seems to make the problem easier to approach. First, many attacks, e.g., malware infections, were observed to happen in the same fashion in close time and across large network areas. Thus, we can use observations from a friendly network to predict upcoming events in our network. The collaboration and information exchange may serve as an early warning system [22]. In addition, heterogeneity of intrusion detection systems may provide a complex view on an attack even though the individual intrusion detection systems do not have complete detection capabilities and, thus, may complement each other [18].

Predicting the next step of an adversary is an integral part of *cyber situational awareness* [23], an application of the concept of *situational awareness* into cyber security. One of the most widely used definitions of situational awareness is the one by Endsley [24]: “*Perception of the elements in the environment within a volume of time and space, the comprehension of their meaning and the projection of their status in the near future.*” The thesis [1] illustrates how to increase cyber situational awareness in its three levels: perception, comprehension, and projection [25], from collecting alerts from multiple sources to analyzing them and predicting the next move of an adversary. Research on predicting cyber attacks have not been conducted in a collaborative environment yet. Although such an environment would seem much more difficult to perceive and comprehend, and, in many aspects, it really is, it opens new possibilities for projecting cyber attacks. Not only we aim to predict when the attackers strike and how, but also which targets will they choose and where can we observe the associated events.

This paper is structured into seven sections. The research questions are stated in Section II. The background of the thesis is summarized in Section III. The contributions of the thesis are summarized in the following three sections. Section IV discusses improvements in the perception level, i.e., extending

capabilities of intrusion detection systems towards information sharing. Comprehension of the shared intrusion detection alerts via alert correlation and data mining is discussed in Section V. Section VI discusses attack prediction using shared alerts and mined predictive rules. Finally, Section VII concludes the paper with answers to research questions, and suggests future research.

## II. RESEARCH QUESTIONS

To formalize the scope of the thesis [1], we posed three research questions. Notice that the research questions follow the levels of situational awareness, on which this whole thesis is structured. The research questions are as follows:

(i) *How can we make use of security alert sharing for building cyber situational awareness?*

First, we aimed at investigating the capabilities of cybersecurity alert sharing platforms from technical and non-technical perspectives. To build cyber situational awareness using the shared data, we need to ensure, from the technical perspective, that the sharing platform is automated and delivers accurate and timely information about relevant cybersecurity events in the collaborating networks [16]. From the non-technical perspective, a sharing community needs to be established, so that the data are rich in heterogeneity and relevant for correlations [26]. Further, questions of trust and privacy issues may appear in a collaborative environment [19]. Fulfilling all the conditions and prerequisites would then enable us to receive relevant data for further research, which relates to the *perception* level of situational awareness.

(ii) *How to understand shared security alerts and effectively discover patterns in such data?*

Second, there was a need to understand the data and discover relations and repeating patterns in them. Simple statistical approaches were not sufficient in this case, given the nature of the data. There are many events related to cyber security that are expected to appear frequently in the data. For example, alerts of network scanning are expected to fold the majority of the shared alerts due to the omnipresent nature of network scans and their relatively easy detection [27]. What we were looking for, however, were the hidden attack patterns, that would not be noticed in plain sight. Data mining is a suitable approach to discovering such patterns, and its usage in cyber security is on the rise these days [28]. Understanding the data and knowing about frequent patterns in the data correspond to the *comprehension* level of situational awareness.

(iii) *How can we predict upcoming cybersecurity events, and how can we evaluate such predictions?*

Finally, when the relations and patterns in cybersecurity alerts are discovered, we may use them to predict future events. This task was studied in earlier works [2], but not in a collaborative environment, which promises a larger overview and new capabilities of attack predictions. The open question, however, is a task of evaluating the prediction. This task is feasible in a laboratory environment but is problematic in live networks as it would require the execution of an intrusion and access to the exploited hosts. Large-scale collaborative environment

makes this task even harder. Altogether, this research question corresponds to the *projection* level of situational awareness.

## III. BACKGROUND AND RELATED WORK

Background for this thesis is composed of two main topics. The first topic is collaboration and information exchange in cyber security. The second is the topic of attack projection, prediction, and forecasting.

The topics of collaboration and information exchange were always present in the field of cyber security, and are well documented in the technical [16, 17] and scientific [18, 19, 21] literature. Therefore, we focused on automated information exchange, such as in collaborative intrusion detection systems and alert sharing platforms. Non-technical aspects of information sharing, such as community management, legal obligations, and issues of trust among the sharing peers, were shown to be at least of the same interest as the technical implementations. Both technical and non-technical aspects of information sharing were illustrated on an example of an intrusion detection alert sharing platform SABU<sup>1</sup>, in which we conducted the experiments and measurements presented in the thesis. A schema of SABU is depicted in Figure 1, where we depicted the core functionality of the platform in the top half of the image. The central component of the platform is Warden, a hub where all the alerts are sent to and redistributed to recipients. The data are alerts provided by intrusion detection systems, honeypots, and other tools, such as third-party alert sharing platforms. The sources of data are deployed in heterogeneous networks of academic and commercial institutions. Warden sends the data to the recipients, such as incident handlers and analysts in security teams and automated active network defense devices. The bottom half of the figure depicts the AIDA framework, which is another outcome of the thesis and is discussed later in this paper. In short, the AIDA framework receives the alerts from Warden, processes them, makes predictions, and sends alerts of predicted events back to Warden.

To cover the background on the second topic, we conducted a thorough literature review of attack projection, prediction and forecasting methods [2]. The problem was set in a context of research on intrusion detection and cyber situational awareness. A taxonomy of methods was provided, and each category was described in detail and evaluated. Briefly, the traditional methods are based on discrete or continuous models. Discrete models, such as attack graphs and their stochastic variants, allow for modelling the attacks and predicting the next move of an adversary via matching the attacker's behavior to a known attack model [29]. Continuous models, such as time series, allow for predicting changes in the security situation and increase or decrease in attack rates. Novel methods based on machine learning and data mining allow for automated extraction of attack models for discrete methods [30] or directly training a neural network to predict attacks [31]. There

<sup>1</sup><https://sabu.cesnet.cz/en/start>

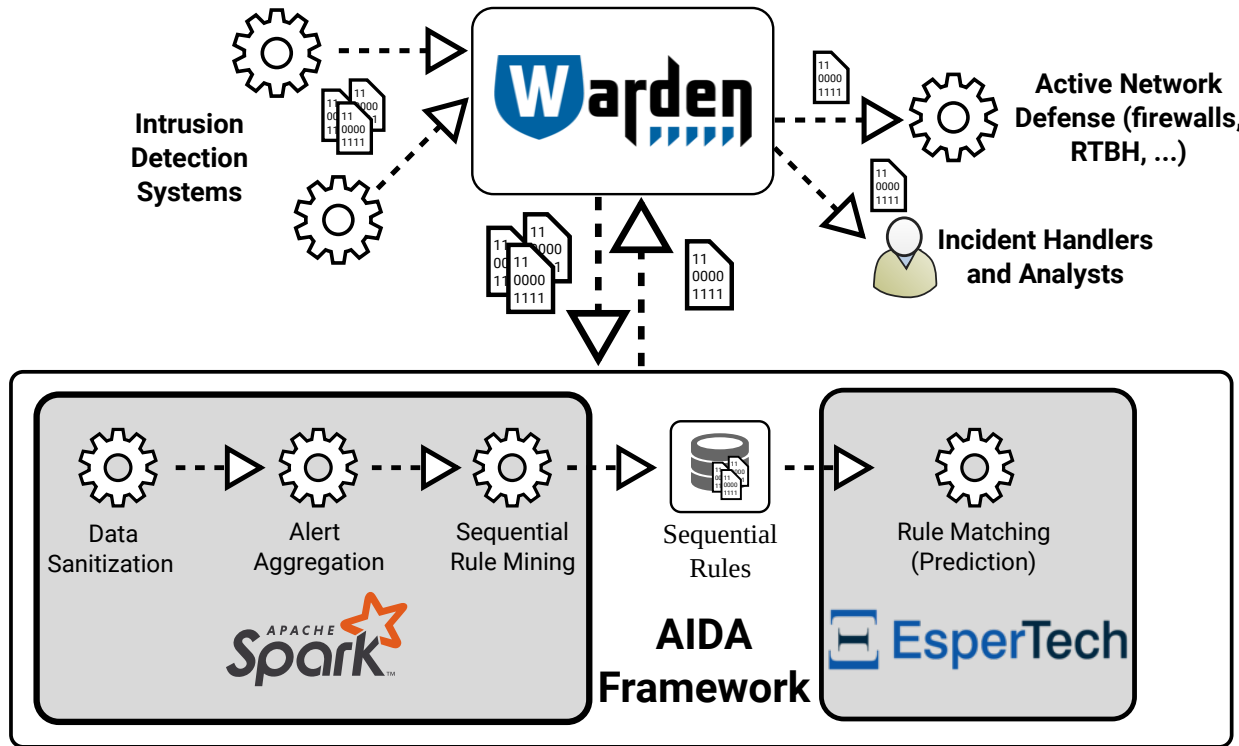


Fig. 1. The schema of AIDA, framework for correlation and analysis of intrusion detection alerts.

are also numerous other methods that are either very novel, not well known, or highly specialized.

#### IV. PERCEPTION LEVEL: INTRUSION DETECTION

To answer the first research question, we had to review what types of information are being shared among the cybersecurity community. The alerts raised by honeypots and intrusion detection systems and shared via SABU alert sharing platform proved to be valuable source of information for increasing the cyber situational awareness as they describe what is happening in the networks that share the information. Thus, we delved deeper into the alerts and sources of the data and discussed how the major intrusion detection systems contribute to the alert sharing platform, what alerts they provide, and how they could be improved or adjusted for the tasks discussed in the thesis. The two most important sources of intrusion detection alerts that contribute to the SABU alert sharing platform are network-based intrusion detection systems and honeypots.

The first source of alerts is a plethora of network-based intrusion detection systems. Due to the interests of the research community centered around alert sharing platform SABU, we focused on network flow monitoring [32] and flow-based intrusion detection [33]. We used flow monitoring extended by HTTP parsing to enhance the capabilities of flow-based intrusion detection towards more fine-grained detection of network reconnaissance activities. Knowledge of the HTTP requests on a large-scale allowed us to detect application-level network scans and brute-force attacks against web services that were found to frequently complement each other and, thus,

displaying an interesting attack pattern [3]. Subsequently, we discussed preliminary approaches to detect attack patterns via means of complex event processing [4], which provided a solid technical background for later work. Further, we observed events, such as web crawling and accesses to phishing websites, that would otherwise be undetectable with plain network flow monitoring [5].

The second source of alerts are the honeypots that are ideal for the analysis of potentially malicious network traffic as any packet incoming to a honeypot is by nature suspicious. We discussed how to monitor the honeypots and how to use them as a source of alerts [6]. Due to observations of the behavior of honeypots during large-scale DDoS attack, we were able to identify problematic network traffic that causes false positive alerts raised by honeypots [7]. As we showed in a case study, honeypots are not so free of false-positives as they were believed to be and might become a burden in certain situations, such as during the reflected DDoS attacks (DRDoS), in which the honeypots may be abused as reflectors of malicious network traffic. In another study on using honeypots, we illustrated how to capture phishing emails to speed up the reaction to phishing campaigns [5]. Capturing phishing emails at honeypots and using network monitoring to detect accesses to phishing websites contributed to the list of interesting patterns in malicious activity.

The contributions of the first research direction of the thesis are mostly in finding ways to improve existing tools and approaches or finding compliance with various novel requirements, both technical and non-technical. Technical issues were

covered by an increase in the variety of the intrusion detection systems and their outputs, more fine-grained detection methods, and reducing false-positive detections. The non-technical issues were mostly related to assuring the quality of intrusion detection alerts and managing the sharing community. However, a novel issue emerged with novel legal frameworks, such as GDPR (General Data Protection Regulation), which raised uncertainty among the sharing community [8]. GDPR states that IP addresses and other identifiers, which are shared within the SABU platform, are personal data. Thus, we found a way to comply with the law by identifying risks to privacy [8] and conducting a balance test that showed that the benefits of information sharing in cyber security are higher than risks of harm to privacy [9].

The main outcomes of the first research directions are as follows. First, the network-based intrusion detection system can use novel approaches of extended flow monitoring to detect more fine-grained events that could be used later [3]. Second, even the honeypots can produce false-positive alerts, even though they were believed to not do so. Nevertheless, the information exchange can help in identifying the false positive alerts [7]. Outside of the technical perspective, one has to be careful about the legal issues regarding privacy, which became an important issue for security practitioners [9].

#### V. COMPREHENSION LEVEL: ALERT CORRELATION

The second research question aims at problems associated with the comprehension of security alerts and the effective discovery of attack patterns to achieve the comprehension level of situational awareness. First, the preliminary visual analytics [10] suggested several directions, in which we started investigating the alerts. Alert aggregation [11] emerged as an important topic to address. Subsequently, we discussed the use of data mining methods to analyze the alerts and discover attack patterns.

In order to extract knowledge from the data, we loosely followed the phases of alert correlation that were set in the fundamental related work [34]. However, in our situation, we had to deal with the alerts from heterogeneous sources distributed in the network address space as well as geographically and organizationally. Our first attempts involved attempts to capture and visualize the relations between various data sources and types of data they shared [10]. The preliminary visualizations helped us to identify the need for alert aggregation, for which we elaborated typical use cases and found out that over 85 % of the alerts can be aggregated [11]. In more detail, about 1 % of the alerts are duplicates that appear due to errors or sharing already shared information. More than 54 % of the alerts are continuations, i.e., alerts of events that were already reported but still continue, such as long-term scans. The remainder of the aggregable alerts are alerts from different data sources that report the same event, such as large-scale scans.

In the next stage of the analysis, we examined the data mining methods and identified methods that are suitable for the task of extracting attack models. The methods of sequential pattern and rule mining proved suitable for our tasks; the

sequences correspond to the series of actions of attackers as observed by the intrusion detection systems [12]. Namely, the top-k sequential rule mining [35] offers good performance and usable results that are directly translatable into predictive rules. The confidence value of the mined rules, i.e., a probability that a certain sequence of alerts will be followed by another as described in the rule, were often found to reach values of 0.8 and 0.9 on a scale from 0 to 1. Subsequently, we experimentally evaluated the use of sequential rules for predicting attacks [13]. The results showed that the outputs are stable in time and, thus, a mined rule can be used for prediction at least in the following days and weeks. Further, most of the rules predict events at least several minutes before they happen and, thus, leave enough time to respond to the predictions.

The main contributions of this research direction are the experimental evaluation of the proposed methods and a detailed description of hands-on experience with processing the shared intrusion detection alerts from heterogeneous sources. The important findings are that almost 85 % of the security alerts are duplicate entries and can be aggregated in favor of the 15 % of alerts representing unique events [11]. Subsequently, the methods of sequential pattern and rule mining were shown to extract usable models of attackers' activities. Namely, the top-k sequential rule mining offers good performance and results that are directly translatable into predictive rules with a promising confidence [12, 13]. Finally, aggregation of the alerts increased the usability of the mined rules because it prevented mining the rules that would describe repeated reporting of long-lasting events and other patterns frequently found in the data.

#### VI. PROJECTION LEVEL: ATTACK PREDICTION

In the third research direction, we discussed the implementation and evaluation of a framework for security event prediction in a collaborative environment. First, we introduced the AIDA framework and discussed the design considerations and deployment options. We outlined the experiences from the deployment of the AIDA framework in the alert sharing platform [14]. Subsequently, we introduced the stand-alone deployment mode of the AIDA framework that allowed for experimentation with datasets. Finally, we created our own dataset that can be used to evaluate the software and reproduce our results.

AIDA framework<sup>2</sup> [14] is a modular framework for the stream-based analysis of intrusion detection alerts using the concepts of big data processing, data mining, and complex event processing. A simplified schema of the AIDA framework and its deployment in the SABU platform is depicted in Figure 1. The framework receives alerts from Warden, the central component of the SABU platform. The alerts are distributed to processing components of the framework as a data stream using the Kafka message broker<sup>3</sup>. The data are first processed by three components based on Apache

<sup>2</sup><https://github.com/CSIRT-MU/AIDA-Framework>

<sup>3</sup><https://kafka.apache.org/>

Spark<sup>4</sup> framework for stream-based data processing. The first component performs sanitization and semantic filtering of the data, such as filtering alerts of low interest. The other two components perform aggregation and data mining, as described in previous works [11, 12, 13]. The rules, mined by the top-k sequential rule mining algorithm [35], are stored in a database, where they are accessed by the rule matching component. The rule matching component is based on Esper<sup>5</sup> and its Event Processing Language (ELP) that allows for translation of the sequential rules into SQL-like queries over the stream of alerts. If the first part of the rule is found using the ELP query, the remainder of the rule is predicted and reported as an alert that is sent back to Warden.

To allow for the reproducibility of the results and further experimentations with the shared data and the AIDA framework, we collected and published a dataset [15]. The dataset contains intrusion detection alerts obtained via an alert sharing platform SABU for one week. A plethora of heterogeneous intrusion detection systems and honeypots deployed across several organizations contributed to the sharing platform. The alerts are stored in the IDEA format<sup>6</sup> and categorized using the taxonomy of security events included in the IDEA definition. The network entities (IP addresses, hostnames, etc.) are anonymized. However, the list of interesting features (presence on blacklists, geolocation, etc.) of such entities at the time of data collection is provided.

The main contribution of this research direction is the development of the AIDA framework and achieving the final level of situational awareness by projecting the upcoming attack in operational deployment. Creating and publishing the dataset and publishing the AIDA framework as open-source software allows for the reproducibility of our research and may serve for our future work as well as for the future work of the research community.

## VII. SUMMARY AND CONCLUSION

This paper provided a digest of a thesis on predicting network attacks in collaborative environment [1]. The thesis and this paper are loosely structured by the research questions that correspond to the three levels of situational awareness (perception, comprehension, and prediction). The research questions were stated and discussed in Section II and answered in Sections IV–VI. To summarize the whole thesis, we found out that the sharing of intrusion detection alerts allows for increasing the cyber situational awareness in all three of its levels. On the first level, *perception*, we discussed how to obtain the alerts from heterogeneous and distributed sources that allow for interesting observations, such as finding if an event is isolated or part of a larger event. Various sources of information may complement each other quite well, even in scenarios that are not so straightforward. On the second level, *comprehension*, we showed how to understand the data and discover interesting patterns in them using alert correlation

and data mining. An important observation is that the stability of the results in the time intervals of days and week allows for the meaningful usability of the predictive analyses over the alerts that complete the cyber situational awareness on the third and final level, *projection*.

The main contribution of the thesis is the combination of information sharing with predictive analytics [13, 14], a combination that was called for [19], yet not thoroughly investigated. Making use of shared intrusion detection alerts may also be considered as a contribution due to the heterogeneity of the data and other aspects that complicate processing the data [11]. Extracting patterns from the data [12] and using it to predict upcoming events served not only as an early warning for receivers of the predicted alerts but also as a significant data volume reduction process. Instead of processing millions of alerts per day by every receiving peer in the sharing platform as observed in SABU, the receivers may filter and receive only those alerts that are relevant for their networks and systems [14]. Thus, the predictions serve the purpose of personalized filtering of the shared data for the receivers, for which it is easier to process tens or hundreds of alerts. A major scientific contribution of this work is also a literature review of attack projection, prediction, and forecasting methods, that was published as a survey paper [2]. For the sake of reproducibility of our research and also for use by other researchers, we also crafted a dataset of intrusion detection alerts from heterogeneous sources and published an open-source implementation of the AIDA framework that was used for the experiments.

We believe there is a lot of potential for future work. As discussed in the thesis, modeling the attack patterns via graph models could potentially improve the performance of the pattern matching and predictions. Manipulating with the parameters of data mining would allow for more interesting patterns in security events, such as distributed and multi-stage attacks. Another interesting direction of future work would be the visualization of the data and the processes discussed in this thesis. Although we presented visual data analytics for the preliminary understanding of the shared intrusion detection alerts [10], it is only a small fraction of what could be visualized. Finally, a lot of research opportunities is in the further development of intrusion detection alert sharing platforms and other collaborative and information sharing tools. Alert correlation supported by data mining, attack projection and prediction, and evaluation of the overall security situation will probably stay as an interesting research topic as we are constantly facing novel threats, attack vectors, and attack patterns. Similarly, future cybersecurity operations may require the sharing of different sets of information, including the sharing of information on incident response and attack mitigation. Combining more and more heterogeneous information from distributed sources allows for a plethora of research opportunities.

<sup>4</sup><https://spark.apache.org/>

<sup>5</sup><http://www.esper.tech.com/esper/>

<sup>6</sup><https://idea.cesnet.cz/en/index>

## ACKNOWLEDGMENT

This research was supported by ERDF “CyberSecurity, CyberCrime and Critical Information Infrastructures Center of Excellence” (No. CZ.02.1.01/0.0/0.0/16\_019/0000822).

## AUTHORED PUBLICATIONS

- [1] Martin Husák. “Prediction of Network Attacks in Collaborative Environment”. Doctoral thesis. Masaryk University, Faculty of Informatics, Brno, 2019. URL: <https://is.muni.cz/th/dmpga/>.
- [2] Martin Husák, Jana Komárková, Elias Bou-Harb, and Pavel Čeleda. “Survey of Attack Projection, Prediction, and Forecasting in Cyber Security”. In: *IEEE Communications Surveys & Tutorials* 21.1 (Firstquarter 2019), pp. 640–660.
- [3] Martin Husák, Petr Velan, and Jan Vykopal. “Security Monitoring of HTTP Traffic Using Extended Flows”. In: *2015 10th International Conference on Availability, Reliability and Security*. Toulouse: IEEE, 2015, pp. 258–265.
- [4] Petr Velan, Martin Husák, and Daniel Tovarník. “Rapid prototyping of flow-based detection methods using complex event processing”. In: *NOMS 2018 - 2018 IEEE/IFIP Network Operations and Management Symposium*. April 2018.
- [5] Martin Husák and Jakub Čegan. “PhiGARo: Automatic Phishing Detection and Incident Response Framework”. In: *Availability, Reliability and Security (ARES), 2014 Ninth International Conference on*. Fribourg, Switzerland: IEEE, 2014, pp. 295–302.
- [6] Martin Husák and Martin Drašar. “Flow-based Monitoring of Honey-pots”. In: *Security and Protection of Information 2013*. Brno: University of Defence, 2013, pp. 63–70.
- [7] Martin Husák and Martin Vizváry. “POSTER: Reflected attacks abusing honeypots”. In: *Proceedings of the 2013 ACM SIGSAC conference on Computer & Communications Security*. Berlin, Germany: ACM, 2013, pp. 1449–1452.
- [8] Martin Horák, Václav Stupka, and Martin Husák. “GDPR Compliance in Cybersecurity Software: A Case Study of DPIA in Information Sharing Platform”. In: *Proceedings of the 14th International Conference on Availability, Reliability and Security*. ARES '19. Canterbury, CA, United Kingdom: ACM, 2019, 36:1–36:8.
- [9] Václav Stupka, Martin Horák, and Martin Husák. “Protection of personal data in security alert sharing platforms”. In: *Proceedings of the 12th International Conference on Availability, Reliability and Security*. Reggio Calabria: ACM, 2017, “65:1–65:8”.
- [10] Martin Husák and Milan Čermák. “A Graph-based Representation of Relations in Network Security Alert Sharing Platforms”. eng. In: *2017 IFIP/IEEE Symposium on Integrated Network and Service Management (IM)*. Lisbon: IEEE, 2017, pp. 891–892.
- [11] Martin Husák, Milan Čermák, Martin Laštovička, and Jan Vykopal. “Exchanging Security Events: Which And How Many Alerts Can We Aggregate?” eng. In: *2017 IFIP/IEEE Symposium on Integrated Network and Service Management (IM)*. Lisbon: IEEE, 2017, pp. 604–607.
- [12] Martin Husák, Jaroslav Kašpar, Elias Bou-Harb, and Pavel Čeleda. “On the Sequential Pattern and Rule Mining in the Analysis of Cyber Security Alerts”. In: *Proceedings of the 12th International Conference on Availability, Reliability and Security*. Reggio Calabria: ACM, 2017, “22:1–22:10”.
- [13] Martin Husák and Jaroslav Kašpar. “Towards Predicting Cyber Attacks Using Information Exchange and Data Mining”. In: *2018 14th International Wireless Communications Mobile Computing Conference (IWCMC)*. June 2018, pp. 536–541.
- [14] Martin Husák and Jaroslav Kašpar. “AIDA Framework: Real-Time Correlation and Prediction of Intrusion Detection Alerts”. In: *Proceedings of the 14th International Conference on Availability, Reliability and Security*. ARES '19. Canterbury, CA, United Kingdom: ACM, 2019, 81:1–81:8.
- [15] Martin Husák, Martin Žádník, Václav Bartoš, and Pavol Sokol. *Dataset of intrusion detection alerts from a sharing platform*. June 2019. URL: <http://dx.doi.org/10.17632/p6tym3fghz.1>.

## REFERENCES

- [16] ENISA. *Detect, SHARE, Protect – Solutions for Improving Threat Data Exchange among CERTs*. [https://www.enisa.europa.eu/activities/cert/support/information-sharing/detect-share-protect-solutions-for-improving-threat-data-exchange-among-certs/at\\_download/fullReport](https://www.enisa.europa.eu/activities/cert/support/information-sharing/detect-share-protect-solutions-for-improving-threat-data-exchange-among-certs/at_download/fullReport). October 2013.
- [17] ENISA. *Standards and tools for exchange and processing of actionable information*. [https://www.enisa.europa.eu/publications/standards-and-tools-for-exchange-and-processing-of-actionable-information/at\\_download/fullReport](https://www.enisa.europa.eu/publications/standards-and-tools-for-exchange-and-processing-of-actionable-information/at_download/fullReport). November 2014.
- [18] Carol Fung and Raouf Boutaba. *Intrusion Detection Networks: A Key to Collaborative Security*. CRC Press, 2013. ISBN: 9781138198890.
- [19] Emmanouil Vasilomanolakis, Shankar Karuppayah, Max Mühlhäuser, and Mathias Fischer. “Taxonomy and Survey of Collaborative Intrusion Detection”. In: *ACM Computing Surveys* 47.4 (May 2015), 55:1–55:33.
- [20] Huwaida Tagelsir Elshoush and Izzeldin Mohamed Osman. “Alert correlation in collaborative intelligent intrusion detection systems – A survey”. In: *Applied Soft Computing* 11.7 (2011), pp. 4349–4365.
- [21] Guozhu Meng, Yang Liu, Jie Zhang, Alexander Pokluda, and Raouf Boutaba. “Collaborative Security: A Survey and Taxonomy”. In: *ACM Computing Surveys* 48.1 (July 2015), 1:1–1:42.
- [22] Ramaki Ali Ahmadian and Atani Reza Ebrahimi. “A survey of IT early warning systems: architectures, challenges, and solutions”. In: *Security and Communication Networks* 9.17 (2016), pp. 4751–4776.
- [23] Alexander Kott, Cliff Wang, and Robert F. Erbacher. *Cyber defense and situational awareness*. Vol. 62. Springer, 2014.
- [24] Mica R. Endsley. “Situation awareness global assessment technique (SAGAT)”. In: *Aerospace and Electronics Conference, 1988. NAECON 1988., Proceedings of the IEEE 1988 National*. IEEE, 1988, pp. 789–795.
- [25] Mica R. Endsley. “Toward a Theory of Situation Awareness in Dynamic Systems”. In: *Human Factors* 37.1 (1995), pp. 32–64.
- [26] Florian Skopik. *Collaborative cyber threat intelligence: detecting and responding to advanced cyber attacks at the national level*. CRC Press, 2018. ISBN: 9781138031821.
- [27] Elias Bou-Harb, Mourad Debbabi, and Chadi Assi. “Cyber Scanning: A Comprehensive Survey”. In: *IEEE Communications Surveys & Tutorials* 16.3 (2013), pp. 1496–1519.
- [28] A. L. Buczak and E. Guven. “A Survey of Data Mining and Machine Learning Methods for Cyber Security Intrusion Detection”. In: *IEEE Communications Surveys & Tutorials* 18.2 (Secondquarter 2016), pp. 1153–1176.
- [29] Xinzhou Qin and Wenke Lee. “Attack plan recognition and prediction using causal networks”. In: *Computer Security Applications Conference, 2004. 20th Annual*. December 2004, pp. 370–379.
- [30] Hamid Farhadi, Maryam AmirHaeri, and Mohammad Khansari. “Alert Correlation and Prediction Using Data Mining and HMM”. In: *ISecure* 3.2 (2011).
- [31] Xing Fang, Maochao Xu, Shouhuai Xu, and Peng Zhao. “A deep learning framework for predicting cyber attacks rates”. In: *EURASIP Journal on Information Security* 2019.1 (May 2019), p. 5.
- [32] R. Hofstede, P. Čeleda, B. Trammell, I. Drago, R. Sadre, A. Sperotto, and A. Pras. “Flow Monitoring Explained: From Packet Capture to Data Analysis with NetFlow and IPFIX”. In: *IEEE Communications Surveys & Tutorials* 16.4 (Fourthquarter 2014), pp. 2037–2064.
- [33] A Sperotto, G. Schaffrath, R. Sadre, C. Morariu, A Pras, and B. Stiller. “An Overview of IP Flow-Based Intrusion Detection”. In: *IEEE Communications Surveys & Tutorials* 12.3 (Thirdquarter 2010), pp. 343–356.
- [34] Fredrik Valeur, Giovanni Vigna, Christopher Kruegel, and Richard A. Kemmerer. “Comprehensive approach to intrusion detection alert correlation”. In: *IEEE Transactions on Dependable and Secure Computing* 1.3 (July 2004), pp. 146–169.
- [35] Philippe Fournier-Viger and Vincent S. Tseng. “Mining top-k sequential rules”. In: *International Conference on Advanced Data Mining and Applications*. Springer, 2011, pp. 180–194.