Design and proof of concept of a prediction engine for decision support during cyber range attack simulations in the maritime domain

1st Markos Antonopoulos
Institute of Communication & Computer Systems (ICCS)
Athens, Greece
https://orcid.org/0000-0001-9052-0293

2nd Giorgos Drainakis
Institute of Communication & Computer Systems (ICCS)
Athens, Greece
https://orcid.org/0000-0003-2443-2783

3rd Eletherios Ouzounoglou
Institute of Communication & Computer Systems (ICCS)
Athens, Greece
https://orcid.org/0000-0002-5078-3248

4th Giorgos Papavassiliou
Institute of Communication & Computer Systems (ICCS)
Athens, Greece
https://orcid.org/0000-0002-5073-2395

5th Angelos Amditis
Institute of Communication & Computer Systems (ICCS)
Athens, Greece
https://orcid.org/0000-0002-4089-1990

***Abstract*—Globalization and the rapid increase in world trade have contributed to greater demand for international transport and logistics and, consequently, the expansion of the maritime industry, one of the oldest and most vital industries of the global economy, accounting for more than 80% of world trade. Over the last few years, the evolving digitization in the maritime industry has led to a significant increase in the number of cyber-attacks on ports and ships, and thus cyber risk management is considered as one of the main challenges for the sector. In this paper, we present the capabilities of a prediction engine that could be used as a decision support system for maritime cyber security personnel. By describing the integration of such engine with risk and econometric models, we draw perspectives for using such a tool into cyber ranges for training relevant port stakeholders in prioritizing their actions during a cyber-attack.**

Keywords— cyber security, maritime, cyber range, sequence modelling, event prediction, risk assessment, econometric model

I. Introduction

Despite the ever-growing adoption of networked technology in almost all areas of contemporary infrastructures, both a standardized technology plus a systematic methodology for testing and evaluation of such systems and related products from a cyber security perspective remain elusive. Pertinent traditional testing assumed that the behaviour of such networked systems is determined by the behaviour of their individual components; Therefore, traditional testing focused separately on each isolated component by emulating a static interface to the rest of the network. However, the behaviour of such systems is not determined solely from the behaviour of their individual components, but also by their highly complex interactions and interdependencies.

Apart from testing, the security of contemporary networked infrastructure requires trained professionals and specialized tools, able to detect network vulnerabilities and to respond to cyber threats in real-time, analyse them, plus monitor and maintain the network's integrity and secure function. Due to the ever-changing landscape of cyber threats, it is perceivable that cyber-security professionals require continuous, ever-lasting training. Additionally, the tools utilized by cyber-security experts require a robust, updatable design able to address the needs dictated by threats of a continuously evolving nature.

Cyber ranges are an emerging technology promising to provide solutions to the above questions. The National Institute of Standards and Technology defines cyber ranges as interactive and/or simulated representations of events of an organization's local network, system, tools and applications [1][1]. Simply put, a cyber range is a virtual simulation environment for networks, able to incorporate elements like actual or simulated devices, virtual machines, software, webpages, simulated traffic etc. Both security professionals and the development of respective tools can benefit from such simulation environments. Cyber ranges can serve as virtual playgrounds for real-time testing of hands-on skills and tools utility in simulated cyber-attacks.

In this work, we focus on cyber threats in the maritime domain. We propose a tool aiming to serve as a decision support system for maritime cyber security personnel during an ongoing simulated cyber-attack. The tool mainly serves as a prediction engine. The prediction engine aims at organizing the knowledge acquired from simulated cyber-range attack scenarios plus potentially any additional prior knowledge in a systematic statistical/probabilistic representation able to make predictions on the evolution of an ongoing or hypothetical cyber-attack. Combined with the output of infrastructure-specific risk-quantification and econometric models, such a representation is expected to provide a comprehensive, systematic and queryable model of the simulated attack scenarios, facilitating educated decisions during a real time cyber-attack.

978-1-6654-9952-1/22/$31.00 ©2022 IEEE

The paper is organized as follows: Section II presents known related work in the field. Section III describes the proposed tool and its capabilities for providing probabilistic predictions on future cyber security events, while Section IV presents the integration of maritime specific risk estimation and econometric models to the output of the engine. The deployment of the proposed tool into a cyber range environment in the context of the Cyber-MAR project is described in Section V. Finally, conclusions are drawn in Section VI.

II. Related Work

Network attack modelling has a long-standing history. Several approaches have been proposed, each one with its own merits and disadvantages. Attack graphs form a concrete class of such models, further categorized in rule-based [2] [3] and probability-based [4] approaches.

Rule-based attack graphs aim to "provide an efficient representation and algorithmic tools to identify the possible cases system vulnerabilities can be exploited in a network" [5]. The cited references describe several attempts following this approach; it is stressed, however, that such approaches rely on a comprehensive and accurate knowledge of both the network under consideration plus its vulnerabilities. Additionally, any changes and/or updates on the network should also be reflected thoroughly in the respective models. Both these requirements are rarely feasible in practice.

Probability-based attack graphs aim to map sensor (e.g. Intrusion Detection Systems, abbr. IDS) observable events to high level attack patterns using probabilistic modelling [6] [7] [8] [9] [10]. The approach seems not to rely so much on detailed knowledge of the network. However, the definition of high level attack patterns is not a trivial task and requires the involvement of domain experts. How much detail, or how to update these attack patterns remain open questions of active research.

An alternative approach is to model the attacker's behaviour in various terms, including intent, capability, opportunity etc [11] [12] [13]. Although the approach seems rational, its formalization is quite undeveloped. This is also evident on the development of corresponding methodologies which currently lack the maturity to provide robust models.

We note that in any of the aforementioned methodologies, the kind of the infrastructure (e.g. port, factory etc) the network of interest relies in is expected to play an additional role in the designed model. It is reasonable to expect that different cyber threat events will be observed in a port than in a factory. In what follows, we are interested in modelling sequences of cyberthreat events pertaining to a port infrastructure.

III. Description of the Engine

As mentioned in the previous section, each network attack modelling methodology comes with its own merits and disadvantages. In this work, we draw mainly on the methods described in [6], for the following reasons:

- The resulting prediction model does not rely on a detailed knowledge of the network of interest.
- The resulting model is built in a straightforward way from the raw event sequences captured by robust and mature tools during the cyber-attack. Additionally, in order to make real-time predictions during an on-going cyber attack, it requires solely the up-to-the-moment captured sequence of events from the same tools.
- Although training and using the resulting model does not require definition of higher-level attack patterns, the basic structure underlying the model (i.e. the suffix tree, which is described later in this section) consists in a compact representation of all acquired knowledge in terms of already observed (sub-) sequences of events. This enables a systematic consideration of various features of already observed event (sub-) sequences, facilitating the definition and detection of relevant higher-level attack patterns.

The core idea of the proposed tool is to systematically model past knowledge on cyber-attack behaviour patterns, where the cyberattacks are simulated in a cyber range, employed with IDS sensors and Security Information and Event Management (SIEM) tools. The corresponding event logs are captured and recorded by the IDS sensors and the SIEM output events. In this section, we give precise definitions and present the mathematical and algorithmic basis of the engine. First, we draw a distinction between the notions of cybersecurity event and cybersecurity incident.

In terms of the event logs provided by the IDS and SIEM modules, a cybersecurity event is a large and complex structure of information regarding a large variety of features of what can be defined as a cybersecurity event. These features include type of event, source and destination IPs, ports and protocols, among many others. A cybersecurity incident is a modelling term, which we use to refer to a subset of the information describing a cybersecurity event. The definition of this subset is made by modelling decisions. Including all available information on a cybersecurity event in this subset would be impractical for two reasons; First, it can be quite complicated to detect and model patterns in sequences of such rich and complex structures. Even if this is accomplished, it could be confusing to present all this information to a cybersecurity professional in a way that facilitates real time decision making. Second, it would most certainly require vast amounts of data to capture all the diversity in every (time-varying) feature. Thus, what the engine will perceive as an incident is a modelling decision and is to be defined according to what practical purpose it is planned to serve, plus the availability of pertinent data.

Useful alternative definitions of what the engine will perceive as an incident may be the following:

- a triplet of (IDS_event_type, Source_IP, Dest_IP), where IDS_event_type is an IDS indication e.g. dns, fileinfo, flow, http, alert etc.
- a triplet of (SIEM_event_type, Source_IP, Dest_IP), where SIEM_event_type is an SIEM indication e.g. malware, zerologon attack, modicon quantum attack etc.

A more generic model of an incident may also be the following, which takes into account subnetworks rather that specific IPs.

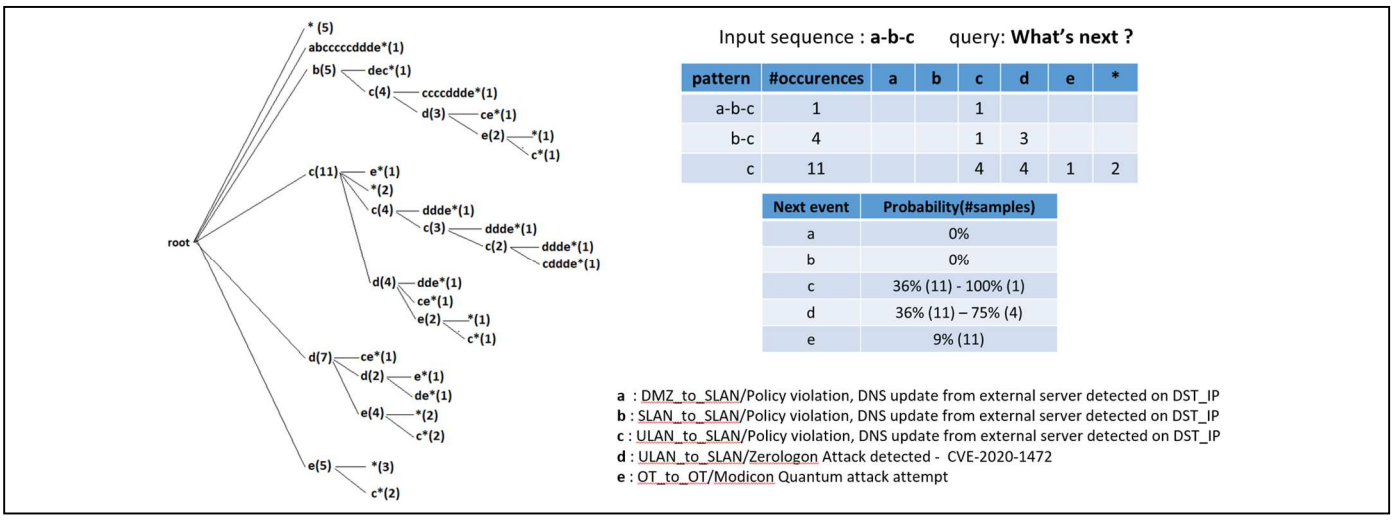

Fig. 1. Single step inference

- a triplet of (event_type, Source_subnetwork, Destination_subnetwork), where event_type is an SIEM or indication as above and Source_subnetwork, Destination_subnetwork may be the any subnetwork including the Operational Technology (OT) subnetwork, the ServersLAN, the Users_LAN, the DMZ etc.

We note that the approach is general enough to enable consideration of a wide variety of possible definitions of the notion of an incident; apart from IDS and/or SIEM event-based definitions, any type of discrete event concerning the port infrastructure maybe included and/or combined with others.

Once the notion of an incident has been defined in a useful and practical way, each incident will be matched to a code word, i.e., a string of characters. This enables the encoding of a sequence of incidents into a sequence of strings. The set of simulated attack sequences can then be represented compactly in a suffix tree structure [14], properly extended to handle sequences of strings instead of sequences of single characters. In what follows, we will demonstrate that such a representation models the knowledge that can be extracted in terms of attack behaviour patterns; these patterns essentially consist of incident (string) sub-sequences observed in the entire set of complete incident (string) sequences obtained by simulated cyber-attack scenarios.

Apart from a compact representation, the suffix tree structure consists of a complete and computationally efficient (i.e. searching a subsequence of length n is O(n)) implementation of a Variable Length Markov Model (VLMM); Given a discrete set E of possible events, an attack is modelled by an ordered finite sequence $\{e_i\}_{i=1}^{N}$. A VLMM considers all n-th order Markov models, i.e. the probabilities

$$P_n\{X_{t+1} = e | X_t = e_0, \dots, X_{t-(n-1)} = e_{n-1}\},$$

$$\text{where } e_i \in E \; \forall \; i = 0, \dots, n-1, n \geq 1 \qquad (1)$$

For more details and intuitive visualizations of the connection between VLMMs and Suffix Trees we refer to [6] [14]. Training a VLMM requires simply event sequences from simulated attack and is equivalent to finding the aforementioned corresponding Suffix Tree [6]. VLMMs have been shown to display good performance in predicting next attack steps [4] [6]. In our work, we have extended the classic suffix tree construction algorithm to be able to handle entire strings instead of single characters. This enlarges the number of different cybersecurity incidents the engine can take into consideration. Apart from inferring probabilities of possible next events as in [6], we have extended the inference algorithm to be able to answer more general queries considering multiple time steps into the future. We showcase the inner-workings of the engine in the following two examples

A. Example 1-Single step inference

Using a cyber range with integrated IDS and SIEM modules, we have simulated a simple attack on a virtual port network including PLCs. IDS sensed events were fed into the SIEM module, which correlated them and produced higher-level alarm events. We repeated 5 dry-runs of the attacks varying the first affected asset and recorded the resulting SIEM alarms sequence. SIEM alarm incidents were encoded by the triplet (SIEM_event_type, Source_subnetwork, Destination_sub network) as described above.

The resulting incident encodings and corresponding sequences are given in TABLE I.

The resulting suffix tree and inference on an exemplary input are depicted in Fig. 1.

The suffix tree encodes systematically all the raw knowledge on attack patterns (i.e. incident subsequences) from the 5 simulated scenarios. For example, one can see that the subsequence…-c-d-… has been observed 4 times in past attacks by simply traversing the tree as the subsequence dictates; The children of the node reached by traversing the tree in that way also shows that out for the 4 times …-c-d-… has been observed, 1 time was followed by a 'd', 1 time by a 'c' and 2 times by an 'e'. Note that each of these characters corresponds to an incident; In this way, we can use the engine as a real-time tool fed with the up-to-the-moment sequence of observed events and

predicting possible next events based on previous knowledge. Apart from that, the engine may be used for off-line studies on patterns of observed incident (sub-) sequences.

TABLE I. INCIDENT ENCODING AND SEQUENCES

Incident Encoding	'a' : DMZ_to_SLAN/Policy violation, DNS update from external server detected on DST_IP 'b' : SLAN_to_SLAN/Policy violation, DNS update from external server detected on DST_IP 'c' : ULAN_to_SLAN/Policy violation, DNS update from external server detected on DST_IP 'd' : ULAN_to_SLAN/Zerologon Attack detected 'e' : OT_to_OT/Modicon Quantum attack attempt
Encoded Incident (alarm) Sequences from XL-SIEM from 5 simulated attacks[a]	a-b-c-c-c-c-c-d-d-d-e-* b-c-d-e-* b-d-e-c-* b-c-d-c-e-* b-c-d-e-c-*
Abbreviations: *SLAN : Servers LAN* *ULAN : Users LAN* *OT: Operational Technology assets (PLCs, Power Transformers etc.)*	

[a.] (*) denotes end of sequence

B. *Example 2-Multistep inference*

The second example illustrates how the engine performs a more general type of inference on the same data. This time the query under consideration is of the type "what is the probability of observing event x in the following n events?"

To answer such a query, again we traverse the tree as dictated by the input sequence. From the reached node, we start a breadth-first-traversal of the tree which extends all possible following incident sequences until they reach a length equal to n. The estimated probability is the fraction of these sequences in which x is present. Fig. 2 presents this multi-step inference.

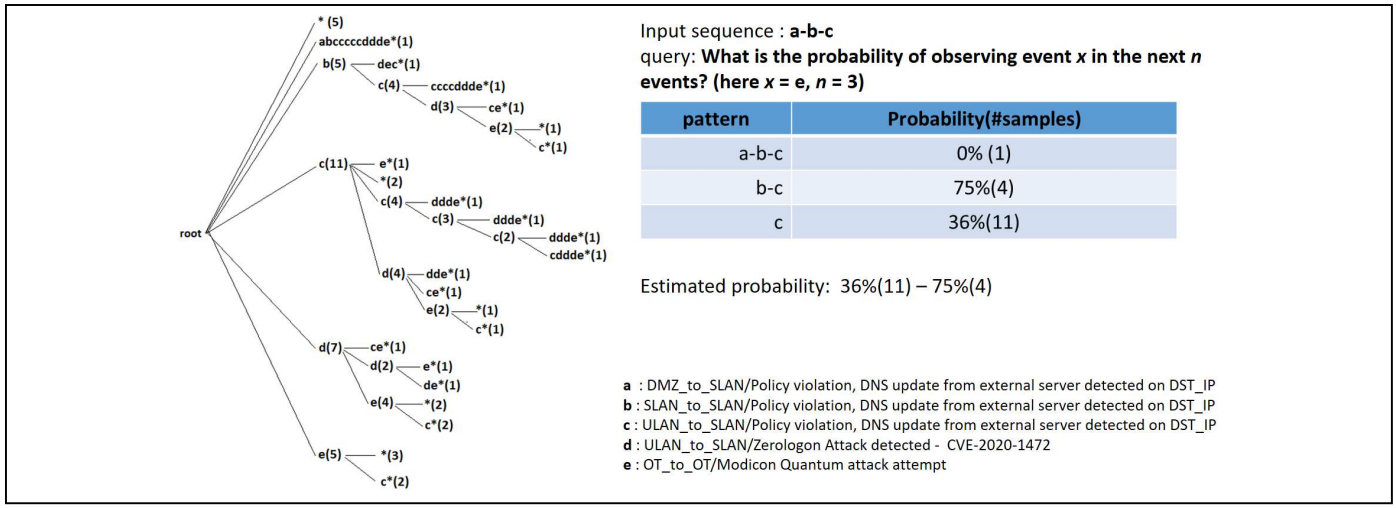

Fig. 2. Multistep inference

IV. RISK AND ECONOMETRIC MODELS INTEGRATION

Up to now, the presented method deals only with information regarding events at the network level and can be utilized to study attack patterns on a variety of networks. Although the cyber range simulated network considers port infrastructure and pertinent network events, the engine outputs only probabilistic predictions on events at the network level; For example, based on the attack patterns observed so far, the engine may infer that there is a high probability of observing a Denial of Service event in a PLC. Since we focus on port infrastructures, this is not enough for a cybersecurity professional studying the attack or making real-time decisions on how to prioritize his/her responses. Additional information, concerning the functional importance of a device in the port operations pipeline plus the econometric impact of inflicted downtimes is needed. Towards this end, we have integrated the output of maritime-specific risk estimation and econometric models to the output of the engine.

Timestamp	Prediction origin (IDS/SIEM)	Description	Source Subnetwork	Destination Subnetwork	est. Probability of occurrence	Num Of Delayed Vessels	Average Delay/Vessel (days)	Downtime (days)
03/22/2022, 10:43:03	IDS	ET POLICY DNS Update From External net	OT	SLAN	0.5 - 1	48	1.45	1.5
03/22/2022, 10:43:03	IDS	ET POLICY DNS Update From External net	ULAN	SLAN	0.5	48	1.45	2.2
03/22/2022, 10:43:03	SIEM	Policy violation, DNS update from external server detected on DST_IP	ULAN	SLAN	0.14 - 0.21	48	1.45	4.2
03/22/2022, 10:43:03	SIEM	Modicon Quantum attack attempt - STOP	OT	OT	0.43 -0.5	23	0.2	8.2
03/22/2022, 10:43:03	SIEM	Zerologon Attack detected	ULAN	SLAN	0.29 - 0.43	48	1.45	8.2

Fig. 3. Output of the Engine

The Maritime Cyber Risk Analysis (MaCRA) [15] framework models the port infrastructure plus the list of cyber-effects associated with any particular IT/OT asset in the port. This list of cyber-effects refers to the types of operational disruptions that could be inflicted by a successful cyber-attack (e.g. denial of service). The resulting risk estimation model accepts as inputs the proportion of available port IT and port OT infrastructures, i.e. two numbers from 0 to 1, and estimates risk –associated information including the number of delayed vessels, the average delay per vessel, equivalent downtime and an indicative recovery scenario.

Fusing and displaying risk estimations with the predictions of the engine as shown in Fig. 3 provides a much more comprehensive consideration of probable future incidents plus their associated implications for the port operations pipeline. Thus, the output of the engine provides the network defender with a means to make educated decisions on how to prioritize his/her actions.

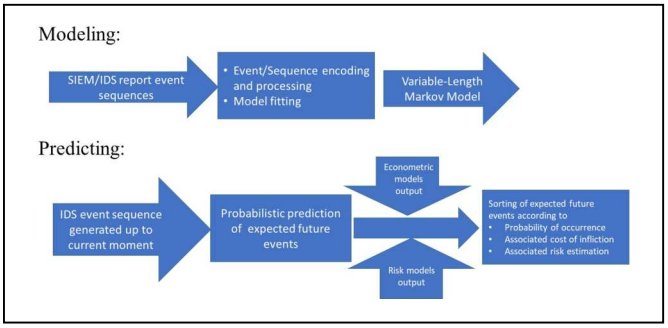

Fig. 4. The entire egine logic

Apart from this, the estimations of the risk model are used to make higher level econometric predictions regarding the economic impact of the downtime of a specific port on other ports, for specific products and involved companies. Each line of the columns Num. of delayed vessels, Average Delay/Vessel and Downtime can be utilized by the considered econometric model and provide estimations of the economic impact of the disruption of the port under attack to other ports for any specified product type [16].

Thus, the interested stakeholder, e.g. the port involved or a specific company can utilize the combined output of predictions and associated risk assessments to consider, analyse and correlate attack patterns from a much broader perspective. We note however, that the database quantifying the econometric impact of each pertinent risk assessment and its integration with the engine is still under development. The entire engine pipeline for modelling, training and inference is depicted in Fig. 4.

V. Technical Details

The Prediction Engine (PE) along with its supporting modules, i.e., the IDS and the SIEM are instantiated as self-contained virtual machines (VMs), all of which are imported into DIATEAM's Hybrid Network Simulation (HNS) cyber range platform [17], as depicted in Fig. 5. HNS is a commercial platform for simulation-based cyber training and testing. Access to each VM is provided via a remote-viewer functionality. In our topology, the Prediction Engine, the IDS and the SIEM are co-located in the same network (Fig. 5).

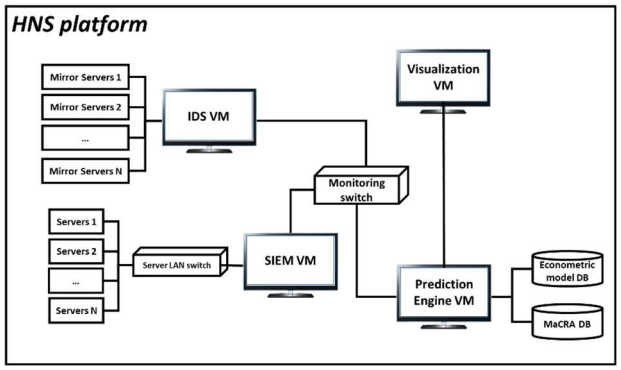

Fig. 5. Integration of the Prediction Engine along with its accompanying modules into an actual cyber range solution (HNS platform)

The above-mentioned modules are configured with static IPs and communicate via a monitoring switch. Additionally, the topology includes a Visualization module and the MaCRA and Econometric models, linked to the PE (Fig. 5). Both the IDS and the SIEM continuously monitor the system's assets e.g. a local area network (LAN) of servers and send any IDS alerts and SIEM detected incidents to the PE, via Syslog-ng mechanism [18] and RabbitMQ [19], respectively. The PE processes the detected alerts thereafter, using the algorithms described in Section III, to extract future events. The predicted future events are then combined with the information of MaCRA (downtime) and Econometric model (economic impact), to provide the holistic PE's output i.e., the total risk estimation and its associated economic impact. Both the MaCRA and the Econometric model are implemented as static databases (DBs). As such, they are incorporated in the PE's VM, where they can be queried accordingly. A more dynamic approach, where both models will be configurable by the end-user (parameterization) is left for future work. The final step of the prediction pipeline includes the propagation of the PE's output to the Visualization module (again via Syslog-ng mechanism), where the predicted results can be presented in a user-friendly manner (dashboard, graphs, etc.).

VI. Conclusion

We have developed a tool aiding the cyber security professional in dealing with cyber-attacks targeting port infrastructures simulated in a virtual environment (i.e. a cyber range). The tool draws on and extends promising modelling approaches and provides probabilistic predictions on future events combined with associated risk and econometric assessments. It remains an open question, how this tool may be further developed for usage in a real-world setting. VLMMs have been shown to provide good predictions, but it is perceivable that complex real-world environments resulting in far more complicated and (intentional or not) obfuscated event sequences need further analysis and more refined approaches on how the aforementioned notion of an incident of interest may be defined.

Acknowledgment

This paper is part of the research efforts under the Cyber-MAR project, which has received funding from the European Union's Horizon 2020 research and innovation program under grant agreement No. 833389. Content reflects only the authors' views, and the European Commission is not responsible for any use that may be made of the information it contains. The authors would like to thank all participants of the Cyber-MAR H2020 project for their continuous and rich involvement in the realization of the results described in this work.

References

[1] National initiative for Cybersecurity Education, Cyber Ranges, National Institute of Standards and Technology (NIST), US Department of Commerce, 2017.

[2] L. Wang, A. Liu, and S. Jajodia, "Using attack graphs for correlating, hypothesizing, and predicting intrusion alerts". Computer Communications, 29(15):2917–2933, 2006

[3] S. Noel and S. Jajodia, "Advanced vulnerability analysis and intrusion detection through predictive attack graphs". Critical Issues in C4I, Armed

Forces Communications and Electronics Association (AFCEA) Solutions Series, International Journal of Command and Control, 2009.

[4] X. Qin and W. Lee, “Attack plan recognition and prediction using causal networks”, in Proceedings of 20th Annual Computer Security Applications Conference, pages 370–379, IEEE, December 2004.

[5] S.J. Yang, H. Du, J. Holsopple, and M. Sudit, “Attack Projection”, in: A. Kott, C. Wang and R. Erbacher R. Eds, Cyber Defense and Situational Awareness, Advances in Information Security, vol 62. Springer, Cham, 2014, https://doi.org/10.1007/978-3-319-11391-3_12

[6] D.S. Fava, S. R. Byers, and S. J. Yang, “Projecting cyberattacks through variable-length markov models”, IEEE Transactions on Information Forensics and Security, 3(3):359–369, September 2008.

[7] H. Du and S. J. Yang, “Characterizing transition behaviors in internet attack sequences”, in Proceedings of the 20th International Conference on Computer Communications and Networks (ICCCN), Maui HI, USA, August 1–4 2011.

[8] C. Cipriano, A. Zand, A. Houmansadr, C. Kruegel, and G. Vigna, “Nexat: A history-based approach to predict attacker actions”, in Proceedings of the 27th Annual Computer Security Applications Conference, pages 383–392. ACM, 2011.

[9] B.-C. Cheng, G.-T. Liao, C.-C. Huang, and M.-T. Yu, “A novel probabilistic matching algorithm for multi-stage attack forecasts”. IEEE Transactions on Selected Areas in Communications, 29(7):1438–1448, 2011.

[10] F. Soldo, A. Le, and A. Markopoulou, “Blacklisting Recommendation System: Using Spatio-Temporal Patterns to Predict Future Attacks”. IEEE Journal on Selected Areas in Communications, 29(7):1423–1437, August 2011.

[11] A. Steinberg, “Open interaction network model for recognizing and predicting threat events”, in Proceedings of Information, Decision and Control (IDC) ’07, pages 285–290, Febuary 2007.

[12] J. Holsopple, M. Sudit, M. Nusinov, D. Liu, H. Du, and S. Yang, “Enhancing Situation Awareness via Automated Situation Assessment”. IEEE Communications Magazine, pages 146–152, March 2010.

[13] H. Du, D.F. Liu, J. Holsopple, and S.J. Yang, “Toward Ensemble Characterization and Projection of Multistage Cyber Attacks,” in Proceedings of the 19th International Conference on Computer Communications and Networks (ICCCN), Zurich, Switzerland, August 2–5 2010. IEEE.

[14] T. C. Bell, J. G. Cleary, and I. H. Witten, Text Compression. Englewood Cliffs, NJ: Prentice-Hall, 1990.

[15] K. Tam and K. Jones, “MaCRA: a model-based framework for maritime cyber-risk assessment,” WMUJ. Marit. Aff., vol. 18, no. 1, pp. 129–163, 2019.

[16] O. Jacq *et al*., "The Cyber-MAR Project: First Results and Perspectives on the Use of Hybrid Cyber Ranges for Port Cyber Risk Assessment," *2021 IEEE International Conference on Cyber Security and Resilience (CSR)*, 2021, pp. 409-414, doi: 10.1109/CSR51186.2021.952796

[17] "What is Cyber Range - Definition - DIATEAM", DIATEAM Cyber Range & Cyber Solutions, 2022. [Online]. Available: https://www.diateam.net/what-is-a-cyber-range/. [Accessed: 08- Apr- 2022].

[18] "syslog-ng - Log Management Solutions", Syslog-ng.com, 2022. [Online]. Available: https://www.syslog-ng.com/. [Accessed: 08- Apr- 2022].

[19] "Messaging that just works — RabbitMQ", Rabbitmq.com, 2022. [Online]. Available: https://www.rabbitmq.com/. [Accessed: 08- Apr- 2022].