Using Cyber Defense Exercises to Obtain Additional Data for Attacker Profiling

Joel Brynielsson*†, Ulrik Franke‡, Muhammad Adnan Tariq*, Stefan Varga*§
*KTH Royal Institute of Technology, SE-100 44 Stockholm, Sweden
†FOI Swedish Defence Research Agency, SE-164 90 Stockholm, Sweden
‡SICS Swedish Institute of Computer Science, P.O. Box 1263, SE-164 29 Kista, Sweden
§Swedish Armed Forces Headquarters, SE-107 85 Stockholm, Sweden
Email: joel@kth.se, ulrik.franke@sics.se, tari@kth.se, stefan.varga@mil.se

***Abstract*—In order to be able to successfully defend an IT system it is useful to have an accurate appreciation of the cyber threat that goes beyond stereotypes. To effectively counter potentially decisive and skilled attackers it is necessary to understand, or at least model, their behavior. Although the real motives for untraceable anonymous attackers will remain a mystery, a thorough understanding of their observable actions can still help to create well-founded attacker profiles that can be used to design effective countermeasures and in other ways enhance cyber defense efforts. In recent work empirically founded attacker profiles, so-called attacker personas, have been used to assess the overall threat situation for an organization. In this paper we elaborate on 1) the use of attacker personas as a technique for attacker profiling, 2) the design of tailor-made cyber defense exercises for the purpose of obtaining the necessary empirical data for the construction of such attacker personas, and 3) how attacker personas can be used for enhancing the situational awareness within the cyber domain. The paper concludes by discussing the possibilities and limitations of using cyber defense exercises for data gathering, and what can and cannot be studied in such exercises.**

***Index Terms*—Cyber defense exercise; behavioral modeling; attacker persona; cyber situational awareness.**

I. Introduction

To "know your enemy and know yourself" was the key to winning battles in the era of Sun Tzu during the 6th century B.C., but this perpetual truth also holds for contemporary "cyber battles." In today's society, a large organization's computer infrastructure is often complex, distributed and under constant change, which makes "knowing yourself" a nontrivial task. Possibly even more challenging, though, is to "know your enemy" since today's attacks may be hard to detect, and even harder to attribute to a perpetrator. Additionally, attackers possess varying skills, have different motives, and may be organized in teams, etc. In this paper we assume that a fair view of the threat that attackers pose can help improve cyber defense. For this purpose we propose investigating ways to obtain additional empirical data from so-called cyber defense exercises (CDXs) [1] to create more elaborated attacker profiles. As mentioned, a key motive for the creation of these profiles is to build upon an informed awareness regarding the threat, as a necessary complement to a thorough understanding of one's own strengths and weaknesses.

In user-centered design, realistic profiling of the users of a system is done by creating so-called personas [2]. Recently, this methodology has been extended further into "attacker personas" [3], which is used for building more extensive profiles of the adversary, that can in turn be used for an in-depth discussion regarding the overall threat in terms of analyzing both one's own organization and the possible threats to the organization [4], [5]. Furthermore, attacker personas can be used to convey a sufficiently accurate sense of the cyber threat that is relatively easy to understand to other stakeholders, such as senior management and non-technical personnel, thereby viewing the persona also as a communications tool [2]. A crucially important aspect, however, is to empirically obtain the data to be used for creating relevant attacker personas. This is not an easy thing to accomplish: as commonly accepted within the user-centered design field, there is a discrepancy between what you do and what you say that you do, so therefore it is important to find the right methods for making the "attackers" actually generate the relevant, domain-specific information to be used for the profiling. That is, one must come up with relevant activities and means for making observations in order to obtain the data needed for the persona creation process.

Building on experiences from past CDXs, this paper discusses possible CDX setups, and a way to enrich attacker personas by using both qualitative and quantitative observations. The main idea is to use a CDX to enrich attacker profiles by designing a part of the exercise in such a way as to stimulate the attacking teams to perform attacks that are as close to assumed realistic attacker behavior as possible. This attacker behavior can then be observed and used for determining the "user" characteristics. In the long run, a number of such observations can turn an initial uninformed view of the threat, e.g., as being posed by the stereotypical "script kiddie profile," into a well-informed attacker persona description that can play an integral role for analyzing the overall organizational threat.

The remainder of the paper is structured as follows. Section II introduces the concept of cyber security exercises in general, CDXs in particular, and discusses their usefulness for scientific research purposes. Section III provides the necessary background regarding personas as a means to portray users for whom to design, and attacker personas as a means to portray

978-1-5090-3865-7/16/$31.00

elusive "users" that may pose a threat to the information security of an organization. Then, Section IV discusses the importance of obtaining relevant empirical data to be used for both attacker persona validation and for obtaining a higher-level understanding of how attacks are carried out, and presents principles and best practices for CDX design for this purpose, e.g., the use of incentive structures. A discussion regarding the applicability of the described methods can be found in Section V. Finally, Section VI concludes the paper.

II. Cyber Defense Exercises

There are different types of cyber exercises. One type is a table-top exercise that deals with cyber security on a conceptual level. Other types take advantage of actual physical resources, e.g., computers and networks to provide more or less realistic environments for training purposes. Hoffman et al. [6] propose a general taxonomy with small scale internal exercises, national capture the flag (CTF) exercises, semester-long class exercises, and cyber defense exercises, CDXs. Sometimes an element of competition is involved. The goals of cyber exercises vary between the development of offensive (attack) skills that are typically the focus of CTF-exercises, and the development of defensive skills that are trained in CDXs.

A CDX typically involves computers that are interconnected and run a number of services that are supposed to be secure and available to legitimate users. The training infrastructure, sometimes referred to as a cyber range, may be set up predominantly with physical hardware, but also as a virtual environment consisting of virtual computers and network equipment. As mentioned, the exercises may be of different scale with respect to the number of participants, number of geographical sites, etc. When disperse geographical locations are used the participating sites are normally connected via presumably secure virtual private networks, VPNs, in order to preserve the integrity of the exercise and to prevent "leakage" of exercise activities to the public internet.

As noted by Mauer et al. [7], the first U.S. military CDX premiered in 2001. Since then a number of subsequent exercises have been carried out under the auspices of the U.S. military. The practice of conducting similar exercises, although of less bellicose nature, was later adopted and developed by academia as well [6], [8], initially in the United States. Due to the nature of computer security, where scientific insight benefits greatly by being coupled with applied skills [9], hands-on exercises are increasingly developed and used in contemporary university education endeavors [1], [10], [11], [12], [13] whilst already being the main focus in professional training activities [14].

The scope and purpose of CDXs are (to provide) "interactive learning opportunities in realistic scenarios" [7]. A possible theoretical explanation for the successes of CDX training may be derived from Kolb's pedagogical model of experiential learning [15]. Beside the technical training components of CDXs [16], Hoffman et al. [6] argue that they also provide training for other components, such as for legal, ethical and forensic competencies. By convention the participants are divided into different teams, typically at least one blue team that is assigned to defend resources, one red team that is tasked to attack, degrade or destroy resources, and one white team (exercise management) who are observers or judges [8], [17], [18].

In order to keep track of the progress of the participants there is a need for solid performance metrics that can be measured [19]. Patriciu and Furtuna [20] list metrics of two main types: the number of occurrences of some event, and the time elapsed between certain conditions, e.g., up-time and down-time [7].

There are a number of advantages with CDXs. Foremost they provide a controlled environment where it is possible to observe attacks in some detail, but also to positively attribute them to an attacker. As mentioned earlier, it is also possible to train non-technical skills, such as the improvement of teamwork and other interpersonal skills [7]. CDXs can also produce scientifically valuable data, labeled datasets [21], [22], that can be used for research. The rationale for this is that genuine data is hard to come by, and that the validity of artificially generated datasets may be questioned. In the widely used DARPA/MIT Lincoln Laboratory dataset from 1999, for example, it was discovered [23] that all malicious packets had a time to live (TTL) of 126 or 253, which obviously would be unrealistic in a real-life scenario. CDXs, at least, have real people performing attacks, even if some of them are carried out using simple standard tools [21].

However useful CDXs may be, there are also a number of drawbacks. The training environment is still an artificial milieu with an unproportionate amount of malicious activities compared to the everyday situation [21]. The duration of a CDX is limited which in turn increases the workload per time unit of the participants. Furthermore, the absence of "background internet noise" [21] or "internet background radiation" [24], such as scans, diverse automated attacks, malformed packets, flooding backscatter, etc., may also contribute to distort both the psychological reality of the exercise and the validity when using exercise data to, e.g., tune intrusion detection systems [25]. Psychological factors such as stress levels may also be hard to recreate in exercises.

To summarize, it appears that CDXs are suitable for the training of cyber security specialists, but also for the collection of labeled datasets that can be used for research.

III. Personas and Attacker Personas

The term persona was first introduced by Cooper [2], and has become a critical part of user-centered design. Emphasizing the use of fictitious characters for representing the user in the design phase of a system, Cooper argues that the term user is elastic in nature, because the user will "bend and stretch" and behave differently under different circumstances. Referring to the resulting risk that the term "user" is perceived differently by different people within the design team, Cooper proposes personas as a tool to bridge this gap.

A persona is a collection of common attributes derived from a group of users who share common goals and usage patterns represented into a single personification. The personas are typically created by using collected data about the anticipated larger user community [26].

Each persona comprises attributes such as motivations, attitudes, skills, activities and goals, a picture, etc., and other domain relevant attributes related to the purpose of the persona. Note that Cooper, in the context of user-centered design, refers to skills principally as the ability to use/operate software. The idea behind developing personas is to make way for a system design that is user-centered, and to provide a methodology that minimizes design biases resulting from different interpretations of the term "user" that may result in an undesired designer-centered or technology-centered design of the system. The main attributes in the persona descriptions should reflect relevant properties of the target group in mind, but additional verbose descriptions that are not completely researched are also allowed in order to make the descriptions more life-like and personal.

However, the persona methodology has also been criticized where, e.g., Portigal [27] argues that the persona can be perceived as a stereotypical representation of the user, and thus does not represent the real user. Another potential pitfall is using too many attributes for the persona, in which case Chapman et al. [28] claim that it narrows down the number of people it really identifies considerably, maybe to the point of irrelevance.

Steele and Jia [29] proposed that the personas technique can also be of use for the design of secure systems. To denote personas with these adversarial aspects in mind, they proposed the usage of the "anti" prefix, e.g., anti-scenarios, anti-use-cases, and anti-personas. Later the term "attacker persona" was proposed [3] to specifically encompass the behavior of archetypical attackers of systems.

The critical question that we are addressing in this paper is how we can develop these attacker personas further, since real attackers in action tend to take precautions to cover their tracks, thus making observation difficult.

Atzeni et al. [3] have tried to answer this question by presenting a methodology which relies on using existing IT security literature as the major source of information. The proposed methodology mainly relies on data derived from open source threat taxonomies followed by affinity diagrams to rationalize their assumptions. However, we argue that merely relying on the literature can potentially lead to attacker personas that are stereotypical which, in turn, can result in attacker personas that are either exaggerating or underestimating the attacker threat. Thus, we conclude that attacker personas can and should be further developed by combining *empirical* data with data extracted from the literature. Such empirical data can to some extent be gathered from the (friendly) defending side through interviews [30], but to gather data from the (non-cooperative) attacking side calls for more elaborate techniques.

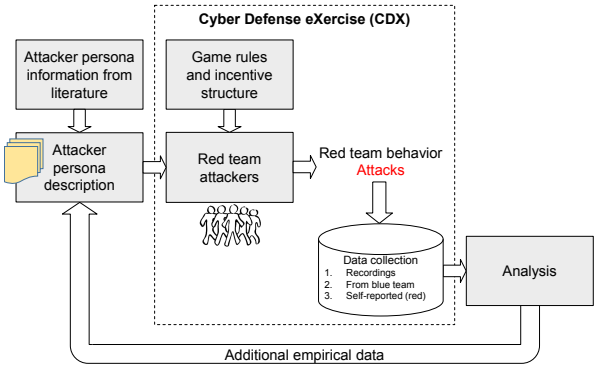

Fig. 1. The scheme envisions the process of using a CDX for collection of empirical data to be used for attacker persona refinement.

IV. Cyber Defense Exercise Design

We propose that a CDX, in addition to its overarching goals and objectives, can also be specifically tailored for the purpose of refining attacker personas according to Fig. 1. A medium- to large-scale CDX is typically organized with defenders (blue teams), attackers (red teams) and a number of support teams. All CDX participants are given a scenario which ideally is aligned with the overarching exercise objectives. The attackers are instructed to conduct attacks against various IT resources according to some rationale that is given to them, while the defenders are instructed to uphold the defenses for their assigned resources. Personas, according to Section III, can be used to set the mood and emulate realistic behavior for both attackers and defenders.

A. Incentive Structures

While it is clear that the actions taken by attackers (and defenders) depend on their incentives, this has traditionally been considered difficult to measure in CDXs. For example, Sommestad and Hallberg [22], note that "the incentives that real attackers or defenders act upon" appear difficult to assess in exercises or competitions.

However, while the *full* complexity of real-life incentives cannot be reflected in a CDX, this does not mean that *some* incentives cannot be deliberately included. In particular, the incentives of the game can be set to reflect different relevant motivations. It is useful to consider a few examples of how such incentive structures within the game might look:

Espionage. In this game, points are *awarded* for obtaining secret information stored in the opposing team's network. Points are *deducted* (possibly more than are awarded) if the opponents discover that information has been stolen. This setup does not explain *why* an attacker commits to espionage, but it does indicate measures prone to be taken by an attacker who is looking for information while simultaneously covering his/her tracks.

Insider. In this game, a member of the team is subverted by the exercise directors to act as an insider, assisting the opposing team regarding their objectives. Various setups are possible, ranging from those where the insider is

exposed after the end of the exercise, to those where his/her actions are kept secret even to the team he/she is secretly helping, and where he/she is given an additional secret payment afterwards by the exercise directors. This setup does not explain *why* someone decides to become an insider, but it offers the chance to study insider modus operandi in realistic circumstances.

Ideological attack. In this game, points are *awarded* for wreaking havoc on the opponent, and *nothing is deducted* for being caught or clumsy. This setup does not explain *why* an attacker has a certain ideological conviction, but it offers a chance to study attackers who do not care about caution or reprisal.

B. Data Collection

As discussed in the preceding sections, it is critical to collect empirical data for developing the attacker personas so that they are not just assumptions but also based on facts to the largest extent possible. Tariq et al. [5] argue that there are approaches to collect empirical data for the purpose of developing attacker personas. As discussed in previous sections, CDXs provide a unique opportunity to examine attack behavior in detail as attacks are executed by the red team. Drawing upon the taxonomy for cyber incidents proposed by Howard and Longstaff [31], factors such as tools, vulnerabilities, attacker actions, targets (e.g., CDX target systems) and unauthorized results (e.g., successful outcomes of attacks), are conveniently collectable.

However, not all types of data needed for the development of attacker personas are readily available in CDXs. As identified [31], relevant data about *actual* attackers and their objectives are not directly available because they do not participate in the CDX. As exemplified within many domains [32], [33], directly interviewing stakeholders is indeed possible but is also challenging and requires significant efforts. Another way is to use data collected from self-reported attackers. However, the question of the validity of such information remains, as it is possible that the respondents either underestimate or overestimate their own skills and knowledge, which in turn leads to a misleading picture [34].

It should be noted that the use of the attribute "skills," differs from its meaning in the context of user-centered design as developed by Cooper. For the context of attacker personas the term denotes an individual's ability to perform actual attacks, whereas Cooper refers to the skills with regards to operate the software of the IT system.

In order to develop attacker personas, the following data can be collected:

Measured data from recordings. Behavioral characteristics of the red team attackers. Data about used tools, exploited vulnerabilities, targeted systems, attacker actions including details of attack execution as well as results of the attack activities (unauthorized access) or failures.

Data from observations. Information from the blue team defenders about the red team attackers. Since attacks often are carried out in multiple steps, there will sometimes be a dynamic interplay with competent defenders. Additional information about how the blue team defenders perceive attacker activities can therefore be collected in a systematic way.

Self-reported data. Details about the red team attackers. Biographical data about CDX participants acting as attackers can be collected. Data such as gender, age, marital status, etc., as well as information about their skill level is of interest, including current work, formal education, and other pertinent skills.

For the first type of data collection for persona and scenario development, which measures the skills, systems that record participants' actions on both the network level and the system level are needed. This can be accomplished by using network monitoring tools at the network level, and by setting up a sandbox environment at the system level. The cyber range mentioned in Section II is an example of a system supporting this kind of data gathering. The recorded information can then be related with the data collected from the other data in order to obtain further insight regarding how the relevant persona/attacker performs attacks.

For the second type of data collection, in which skills are observed, the blue team defenders, who presumably are the IT security experts, can provide insights based on *their observations* regarding the attackers' skills, behavior traits, goals, suspected motivations, and so forth. This data can be collected both qualitatively and quantitatively, and can then be used to further refine the personas.

The third type of gathered data—skills reported in questionnaires given to the red team prior to, or in conjunction with, the CDX—can be used to, if possible, infer additional insights by comparing their given "profiles" with their actual accomplishments in their attacking endeavors, to further refine the attacker personas.

The totality of the collected data can thus be used to validate the assumptions about the initial attacker personas used, which were created using data gathered from the literature similar to the example persona depicted in Fig. 2.

To get hold of the attributes mentioned above, it becomes pivotal that the tasks are designed in a way that they enable a certain level of diversity. Neither the scenarios, nor the tasks can therefore be too specific. That is, the idea of using attacker personas is to distinguish between categories of attackers that perform attacks in different ways. This difference in attack pattern might, e.g., be based on the motivations and the goals of an attacker which an attacker persona is well suited to describe. To further elaborate on the issue, let us consider an ideologically motivated attacker who has a particular goal, perhaps to deface a popular website. In this particular case, the attacker might not care about hiding his/her footprints since he/she wants to show his/her presence. On the other hand, an attacker motivated by espionage might attack the target in a more cautious way since the goal, by definition, is to steal critical information without getting caught. The approach taken by this latter kind of attacker will likely be more thoroughly planned and rather than attacking the system

Fig. 2. The persona in the figure depicts the ideologically motivated attacker "Martin." Martin has been attacking organizations for ideological reasons, primarily triggered by his beliefs regarding freedom of information.

directly the espionage attacker will probably take measures to conceal his/her actions. Similarly, a financially motivated attacker cares about the financial gain rather than the target itself, and will try to find a system which is easy to attack relative to the financial gain. Hence, this kind of attacker can be expected to try to scan as many different systems as possible.

V. Discussion

In this paper we have argued that attacker personas can be improved by using empirical data collected from CDXs and similar types of games or exercises. It is important, however, to consider both the limitations and the strengths of this claim. Following Raser [35], we distinguish between four criteria for the validity of gaming as a research tool:

1) psychological reality,
2) structural validity,
3) process validity, and
4) predictive validity.

For some research questions, these criteria are relatively easy to meet. If the objective is to find the success rate of remote code execution attacks (as described by Holm et al. [36]), then the exercise environment can be set up accordingly, and whenever a remote code execution attack is performed by the red team, the simulation environment ensures structural validity (operating systems, communication protocols, etc., all work just like in reality), process validity (finding vulnerabilities, using exploits, obtaining privilege escalation, etc., all work just like in reality), and predictive validity (what works in the simulated environment works in reality—if the real systems are configured just like the simulated ones). As for psychological reality, this research question requires only that participants, once in a while, actually attempt to perform a remote code execution attack.

For other research questions, such as the incentives of real attackers and defenders, it might seem that the requirement for psychological reality becomes prohibitive, requiring the participants to actually *be*, say, ideologically or financially motivated. (Indeed, not even economic incentives for the participants are certain to make them financially motivated: "Subjects may make competitive choices not because they want to maximize their point totals, but because they want to beat the other person," as noted by Schlenker and Bonoma [37].) However, as argued in Section IV-A, there is a middle ground. By deliberately designing the incentives of the game, some knowledge about behavior under different incentives, corresponding to different attacker types, can probably be extracted, as illustrated in the examples given above.

These examples also shed some additional light on the interplay between personas and CDXs. The qualitative information making up personas is required for proper incentive structures in exercises to be set up. The results of the exercises can then serve to enrich the personas with realistic courses of action for different attacker types, operationalized by exercise incentives. Such behavioral information can be both qualitative, e.g., common modi operandi for espionage, and quantitative, e.g., the relative detection rates of ideological attackers compared to insiders.

However, it is also possible to use the data collected from CDXs to create completely new personas, rather than improve existing ones. Given enough behavioral data from the exercise, or several exercises, exploratory factor analysis can be carried out to create data-driven personas. This methodology is developed in greater detail by McGinn and Kotamraju [38].

VI. Conclusions

In this paper, we have elaborated on attacker profiling, and the possibility to create more realistic profiles using techniques employed within the user-centered design community. The paper has presented an approach to complementing the established practices of developing attacker personas. It has been argued that attacker personas that are solely based on literature studies and interviews run the risk of being stereotypical, and must be complemented with empirical knowledge originating from a more realistic environment. Further, a technique for complementing existing knowledge by obtaining additional empirical data from CDXs has been presented. The idea is to compare the reported skills of individual attackers or groups, that were acquired through questionnaires, with what they actually accomplish in a CDX in order to create more comprehensive attacker personas, that can be used to convey a sense of the cyber threat, and to ultimately aid in decision-making about cyber security.

The major contribution of the paper is the presented techniques for using CDXs to obtain empirical data to increase knowledge of attacker behavior, thereby allowing attacker personas of better quality to be developed.

Further work involves the refinement of the technical setup of CDXs in cyber ranges, to properly enable extensive data collection on a suitable abstraction level. In the future we will work in parallel with other kinds of empirical data gathering methods (literature studies, interviews, etc.) in order to further refine a set of attacker personas to ultimately come up with best practices for how to use CDXs as a means to enrich the attacker persona descriptions.

REFERENCES

[1] B. E. Mullins, T. H. Lacey, R. F. Mills, J. M. Trechter, and S. D. Bass, "How the cyber defense exercise shaped an information-assurance curriculum," *IEEE Security & Privacy*, vol. 5, no. 5, pp. 40–49, Sep.–Oct. 2007.

[2] A. Cooper, *The Inmates Are Running the Asylum: Why High-Tech Products Drive Us Crazy and How to Restore the Sanity*, 2nd ed. Indianapolis, Indiana: Sams Publishing, 2004.

[3] A. Atzeni, C. Cameroni, S. Faily, J. Lyle, and I. Fléchais, "Here's Johnny: a methodology for developing attacker personas," in *Proceedings of the Sixth International Conference on Availability, Reliability and Security (ARES 2011)*, Vienna, Austria, Aug. 2011, pp. 722–727.

[4] M. A. Tariq, J. Brynielsson, and H. Artman, "Storytelling for tackling organized cybercrime," in *Proceedings of the 26th BCS Conference on Human Computer Interaction*, Birmingham, United Kingdom, Sep. 2012.

[5] M. A. Tariq, J. Brynielsson, and H. Artman, "Framing the attacker in organized cybercrime," in *Proceedings of the 2012 European Intelligence and Security Informatics Conference (EISIC 2012)*, Odense, Denmark, Aug. 2012, pp. 30–37.

[6] L. J. Hoffman, T. Rosenberg, R. Dodge, and D. Ragsdale, "Exploring a national cybersecurity exercise for universities," *IEEE Security & Privacy*, vol. 3, no. 5, pp. 27–33, Sep.–Oct. 2005.

[7] B. Mauer, W. Stackpole, and D. Johnson, "Developing small team-based cyber security exercises," in *Proceedings of the 2012 International Conference on Security and Management (SAM'12)*, Las Vegas, Nevada, Jul. 2012, pp. 213–217.

[8] P. Sroufe, S. Tate, R. Dantu, and E. Çankaya Celikel, "Experiences during a collegiate cyber defense competition," *Journal of Applied Security Research*, vol. 5, no. 3, pp. 382–396, 2010.

[9] M. Bishop, "Computer security education: Training, scholarship, and research," *IEEE Computer*, vol. 35, no. 4, pp. 30–32, Apr. 2002.

[10] J. Brynielsson, "An information assurance curriculum for commanding officers using hands-on experiments," *ACM SIGCSE Bulletin*, vol. 41, no. 1, pp. 236–240, Mar. 2009.

[11] D. Jacobson, "Teaching information warfare with lab experiments via the Internet," in *Proceedings of the 34th ASEE/IEEE Frontiers in Education Conference*, Savannah, Georgia, Oct. 2004, pp. T3C/7–12.

[12] S. K. Sharma and J. Sefchek, "Teaching information systems security courses: A hands-on approach," *Computers & Security*, vol. 26, no. 4, pp. 290–299, Jun. 2007.

[13] J. Hill, C. Carver, J. Humphries, and U. Pooch, "Using an isolated network laboratory to teach advanced networks and security," in *Proceedings of the 32nd ACM SIGCSE Technical Symposium on Computer Science Education*, Charlotte, North Carolina, Feb. 2001, pp. 36–40.

[14] S. Cooper, C. Nickell, V. Piotrowski, B. Oldfield, A. Abdallah, M. Bishop, B. Caelli, M. Dark, E. K. Hawthorne, L. J. Hoffman, L. C. Pérez, C. Pfleeger, R. Raines, C. Schou, and J. Brynielsson, "An exploration of the current state of information assurance education," *ACM SIGCSE Bulletin*, vol. 41, no. 4, pp. 109–125, Dec. 2009.

[15] D. A. Kolb, *Experiential Learning: Experience as the Source of Learning and Development*. Englewood Cliffs, New Jersey: Prentice-Hall, Inc., 1984.

[16] S. Glumich and B. Kropa, "DefEX: Hands-on cyber defense exercises for undergraduate students," in *Proceedings of the 2011 International Conference on Security and Management (SAM'11)*, Las Vegas, Nevada, Jul. 2011, pp. 487–493.

[17] J. Kick, "Cyber exercise playbook," MITRE Corporation, Wiesbaden, Germany, Tech. Rep. MP140714, Nov. 2014.

[18] N. Wilhelmson and T. Svensson, *Handbook for planning, running and evaluating information technology and cyber security exercises*. Stockholm, Sweden: National Defence College, 2014.

[19] J. Brynielsson, U. Franke, and S. Varga, "Cyber situational awareness testing," in *Combatting Cybercrime and Cyberterrorism: Challenges, Trends and Priorities*, B. Akhgar and B. Brewster, Eds. Springer International Publishing, 2016, ch. 12, pp. 209–233.

[20] V.-V. Patriciu and A. C. Furtuna, "Guide for designing cyber security exercises," in *Proceedings of the 8th WSEAS International Conference on E-Activities and Information Security and Privacy*, Puerto de la Cruz, Tenerife, Canary Islands, Spain, Dec. 2009, pp. 172–177.

[21] B. Sangster, T. J. O'Connor, T. Cook, R. Fanelli, E. Dean, W. J. Adams, C. Morrell, and G. Conti, "Toward instrumenting network warfare competitions to generate labeled datasets," in *Proceedings of the 2nd Workshop on Cyber Security Experimentation and Test (CSET'09)*, Montreal, Canada, Aug. 2009.

[22] T. Sommestad and J. Hallberg, "Cyber security exercises and competitions as a platform for cyber security experiments," in *Proceedings of the 17th Nordic Conference on Secure IT Systems (NordSec 2012)*, Karlskrona, Sweden, Oct.–Nov. 2012, pp. 47–60.

[23] M. V. Mahoney and P. K. Chan, "An analysis of the 1999 DARPA/Lincoln Laboratory evaluation data for network anomaly detection," in *Proceedings of the 6th International Symposium on Recent Advances in Intrusion Detection (RAID 2003)*, Pittsburgh, Pennsylvania, Sep. 2003, pp. 220–237.

[24] R. Pang, V. Yegneswaran, P. Barford, V. Paxson, and L. Peterson, "Characteristics of internet background radiation," in *Proceedings of the 2004 ACM SIGCOMM Internet Measurement Conference (IMC 2004)*, Taormina, Sicily, Italy, Oct. 2004, pp. 27–40.

[25] T. Sommestad and U. Franke, "A test of intrusion alert filtering based on network information," *Security and Communication Networks*, vol. 8, no. 13, pp. 2291–2301, Sep. 2015.

[26] J. Pruitt and J. Grudin, "Personas: Practice and theory," in *Proceedings of the 2003 conference on Designing for User eXperiences (DUX'03)*, San Francisco, California, Jun. 2003, pp. 1–15.

[27] S. Portigal, "True tales: Persona non grata," *interactions*, vol. 15, no. 1, pp. 72–73, Jan.–Feb. 2008.

[28] C. N. Chapman, E. Love, R. P. Milham, P. ElRif, and J. L. Alford, "Quantitative evaluation of personas as information," in *Proceedings of the Human Factors and Ergonomics Society 52nd Annual Meeting*, New York, Sep. 2008, pp. 1107–1111.

[29] A. Steele and X. Jia, "Adversary centered design: Threat modeling using anti-scenarios, anti-use cases and anti-personas," in *Proceedings of the 2008 International Conference on Information & Knowledge Engineering (IKE 2008)*, 2008, pp. 367–370.

[30] M. A. Tariq, J. Brynielsson, and H. Artman, "The security awareness paradox: A case study," in *Proceedings of the 2014 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM 2014)*, Beijing, China, Aug. 2014, pp. 704–711.

[31] J. D. Howard and T. A. Longstaff, "A common language for computer security incidents," Sandia National Laboratories, Livermore, California, Tech. Rep. SAND98-8667, Oct. 1998.

[32] E. Eriksson, H. Artman, and A. Swartling, "The secret life of a persona: When the personal becomes private," in *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI'13)*, Paris, France, Apr.–May 2013.

[33] J. Horgan, "Interviewing terrorists: A case for primary research," in *Terrorism Informatics: Knowledge Management and Data Mining for Homeland Security*, ser. Integrated Series in Information Systems, H. Chen, E. Reid, J. Sinai, A. Silke, and B. Ganor, Eds. Springer US, 2008, vol. 18, ch. 4, pp. 73–99.

[34] J. S. Giboney, J. G. Proudfoot, S. Goel, and J. S. Valacich, "Measuring hacking ability using a conceptual expertise task," in *Proceedings of the 2015 ADFSL Conference on Digital Forensics, Security and Law*, Daytona Beach, Florida, May 2015, pp. 123–133.

[35] J. R. Raser, *Simulation and Society: An Exploration of Scientific Gaming*. Boston, Massachusetts: Allyn and Bacon, Inc., 1969.

[36] H. Holm, T. Sommestad, U. Franke, and M. Ekstedt, "Success rate of remote code execution attacks: Expert assessments and observations," *Journal of Universal Computer Science*, vol. 18, no. 6, pp. 732–749, Mar. 2012.

[37] B. R. Schlenker and T. V. Bonoma, "Fun and games: The validity of games for the study of conflict," *Journal of Conflict Resolution*, vol. 22, no. 1, pp. 7–38, Mar. 1978.

[38] J. McGinn and N. Kotamraju, "Data-driven persona development," in *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI'08)*, Florence, Italy, Apr. 2008, pp. 1521–1524.