IEEE Access
Multidisciplinary ⋮ Rapid Review ⋮ Open Access Journal

Received January 8, 2018, accepted February 27, 2018, date of publication March 12, 2018, date of current version March 28, 2018.

Digital Object Identifier 10.1109/ACCESS.2018.2814481

On Multi-Phase and Multi-Stage Game-Theoretic Modeling of Advanced Persistent Threats

QUANYAN ZHU[1], (Member, IEEE), AND STEFAN RASS[2], (Member, IEEE)
[1]Department of Electrical and Computer Engineering, New York University Tandon School of Engineering, Brooklyn, NY 11201, USA
[2]System Security Group, University of Klagenfurt, 9020 Klagenfurt, Austria

Corresponding author: Stefan Rass (stefan.rass@aau.at)

ABSTRACT Advanced persistent threats (APT) are considered as a significant security threat today. Despite their diversity in nature and details, a common skeleton and sequence of phases can be identified that these attacks follow (in similar ways), which admits a game-theoretic description and analysis. This paper describes a general framework that divides a general APT into three major temporal phases, and fits an individual game model to each phase, connecting the games at the transition points between the phases (similarly to ''milestones'' accomplished during the launch of an APT). The theoretical description is derived from a running example. The benefit of this game-theoretic perspective is at least threefold, as it 1) helps to systematize the threat and respective mitigation actions (by turning them into pure strategies for the gameplay); 2) provides optimized actions for defense and attack, where the latter can be taken as a (non-unique) indication of neuralgic points; and 3) provides quantitative measures of resilience against an APT, in terms that can be defined freely by a security officer. We illustrate this approach with a numerical example.

INDEX TERMS K.6.5 security and protection, H.4.2.a decision support, I.2.1 applications and expert knowledge-intensive systems.

I. INTRODUCTION

Contemporary hacks of large-scale (critical or enterprise) infrastructures have been reportedly successful, due to their technical sophistication, paired with stealthiness. Indeed, ''cyber-weapons'' have long become a standard part of national defense agendas, with hacks and denial-of-service attacks on critical infrastructures today playing the same role as conventional weapon tests that have been made public to demonstrate a nation's capability of warfare. Independently and in parallel to the development of the defense mechanisms, an equally powerful arsenal of attack systems has evolved, partially mimicking legal online and web services up to cloud computing. The criminal counterpart of legal online services has nowadays become known under the term ''cybercrime-as-a-service'', describing a fully developed business sector secretly acting and growing in the internet, offering a huge lot of techniques, technology and services that support all sorts of cybercrime. From the honest party's perspective, a particular form of such cybercrime are advanced persistent threats, which summarize a diverse lot of highly sophisticated techniques to slowly penetrate a system without causing any harm at first, but waiting for the final strike until it is too late for the infected system to recover or prevent the attack. The process is indeed comparable to a normal virus infection like HIV, which after infection can remain stealthy for a while before the virus becomes active and the disease breaks out. Similarly, other viruses have an incubation phase after the infection and may show symptoms at a time when it is already too late for a cure. APTs follow a similar rationale, and here we decompose them into temporal phases of infection, stealthy infiltration, and causing damage; where the APT typically exposes itself visibly not before it is too late to fix the damage.

A. CONTRIBUTIONS

Our work presents a game-theoretic approach to capture the three previously sketched phases of an APT infection, where each phase is captured by its own specific game model. The overall APT model is then formed by connecting the three game models at the transitions between APT infection and stealthy takeover, and upon the transition to the phase where the APT actively harms the victim system.

The division into phases and separate game models helps to identify countermeasures in each phase systematically and comprehensively, and lets us find an optimal defense (cure) in each phase. To this end, game-theory is there to simultaneously compute the optimal behavior for both, the attacker pushing forward the APT, and the defender, fighting an invisible enemy, since one of the APT attacker's main goals is to go undetected for as long as possible.

Our game-theoretic model relies on the observation that an APT has to exploit multiple vulnerabilities and penetrate multiple layers of the system to inflict damage on the physical systems. Therefore, a multi-stage and multi-phase game-theoretic framework is proposed to capture the strategic interactions of an attacker with a sequence of agents in the system. At each phase, an APT launches different types of attacks including spearphishing, reconnaissance, exploitation, and command and control. Each phase is composed of multiple stages of local interactions in which an attacker has to take multiple actions to be stealthy and successful. The interaction at each stage and phase is modeled by a game. The structure of the game can vary across the phases and the stages. For example, the spearphishing game is modeled as a simplified Bayesian game. This is to capture the information asymmetry between the players and give a baseline framework that can be extended in follow-up work (we come back to this below). The multi-stage penetration of the networks is modeled using a sequence of nested games. The final-phase physical-layer infrastructure protection is modeled as a finite zero-sum game. The sequential nature of the play allows us to compose the heterogeneous games together into one large-scale game. The composed game provides an integrated view of APT and its interactions with the defending system. We propose *Gestalt Nash equilibrium* as a solution concept for the composed game, which provides a theoretic underpinning for a holistic risk assessment of APT through the analysis of Gestalt Nash equilibrium. Also, we use it as a unifying framework to design and deploy cyber mechanisms to reduce risks and achieve a guaranteed level of security.

The main finding in this work is that game-theoretic methods provide an elegant mean of capturing the heterogeneity and workflow of an APT and enable the protection design at every stage and phase of the system. Our main contribution can be summarized as follows:

- We propose a composable game-theoretic framework and the associated solution concept to capture the multi-stage and multi-phase stealthy behavior of APTs.
- We provide a holistic approach to assess long-term risks of an APT attack that consists of human, cyber and physical interactions with the critical infrastructure.
- We propose adaptive learning methods to compute the equilibrium and design automated and optimal defense across multiple layers of the system.

B. ORGANIZATION OF THE PAPER

The rest of this work is structured as follows: to properly embed our work into the existing landscape of APT threat mitigation, section II discusses a selection of related research on the topic. Section III describes the general structure of an APT, letting a more concrete example follow to derive and illustrate the game-theoretic description of each phase. Section IV-A puts the abstract model to work, showing which results can be obtained and how to interpret them. Conclusions are drawn in section V.

II. RELATED WORK

Our work is related to recent investigations on several prominent APT attacks, including Stuxnet [1], Duqu [2], Flame [3], and Aurora [4]. The detailed analysis of APT has been explored in [5] and [6]. Different types of defense mechanisms for APT-related attacks have been proposed at both cyber and physical layers of the system. The cyber solutions include moving target defense [7], [8], trust mechanisms [9], and defense-in-depth techniques [10]. The physical layer solutions include watermarking [11], adding redundancies [12] and resilient control mechanisms [13], [14]. Holistic cyber-physical solutions for APT attacks have been investigated in [15] and [16], which takes into account the coupling between the two layers of the system and proposes co-design defense mechanisms for APT. In this work, our objective is aligned with those in [15] and [16] to develop an integrated cyber defense solution that not only addresses the multi-stage cyber threats but also human and physical ones. Hence a multi-phase and multi-stage framework naturally arises from this pursuit.

Game theory is a useful tool for modeling attacks and assessing cyber risks [17]. The application of game theory to APT modeling has be explored in [18] and [19]. In [18], a FlipIt game is proposed as the framework for "Stealthy Takeover," in which players compete to control a shared resource. The attacker can periodically compromise a system completely, in the sense of learning its entire state, including its secret keys. In [19], a multi-stage modeling is used to capture explicitly the dynamics of the attack over multiple layers of the system. Moving target defense is used as a defense mechanism to mitigate the impact of the APT. Both works have designed adaptive and learning algorithms to adapt defense strategies based on online observations. Our work extends the idea of multi-stage modeling to a multi-stage and multi-phase framework which includes not only the cyber layer of the defense but also the human and physical layers. The social engineering at the human layer of the system is often phase one of the APT attack which exploits the vulnerabilities in human perception and cognition. The physical layer is the ultimate objective of the APT attack, and hence an explicit modeling of the physical defense is necessary to provide a holistic view for the APT.

III. ADVANCED PERSISTENT THREATS

While the dark count of APT incidents may be considerably larger, a reasonable picture of the general structure of an APT can be drawn based on reported incidents [20]–[23]. Following these reports, and abstracting from the details of each specific step, an APT usually undergoes a series of phases, with each phase being decomposed into several steps to achieve the respective (sub)goal. We condense the more fine-grained view on this "kill-chain" (highlighting the terms used in the taxonomy of [24] by italicization; cf. also [25] for an alternative discussion of the topic in similar terms) into three major phases that we model as games:

- Phase 1 (Initial penetration and establishment): this comprises a phase of information gathering (*reconnaissance*), design of a made to measure malware (*development*), preparing the trojan and droppers (*weaponization*), and *delivery* (transmission into the victim infrastructure, e.g., by a phishing email, or similar).
- Phase 2 (Learning and propagation): this is a repeated sequence of *exploitation* to get deeper into the system, and *installation* to leave artifacts and backdoors for an easy return later.
- Phase 3 (Damage): this includes methods to interact with the victim system's compromised resources via previously left artifacts (*command and control*), and *actions on the target* (causing the actual damage).

The concrete steps being taken along this structure are specific for each threat; see [26]–[28] to mention only a few instances. The respective tasks completed by the adversary are then tailored to the specific infrastructure at hand, and it is difficult to give a general "recipe" here. However, what can be said is that the human factor plays a crucial role in the initial infection phase. In general, common measures include, but are not limited to, (spear) phishing, bring-your-own device, and many other forms of social engineering. In light of strong network security widely found nowadays (firewalls and intrusion detection systems have become standard components), sending out malware via email has become a popular technique to begin an initial infection.

The infection itself can be logically divided in the malware obtained by clicking a malicious weblink or opening a virulent file (e.g., with macros in it, or similar), any of which causes a (drive-by) download of malware to the computer. This first occurrence of malware does not necessarily cause any damage by itself, but is in many cases a sleeping trojan (a.k.a. "dropper") whose sole purpose is to infect further machines, and to wait for the remotely triggered deployment of the actual malware (in Phase 3 of the above list). This remote control is usually up to a command-and-control (C&C) server, commonly itself being hidden behind a hierarchy of proxies (cf. [27]). Once access to the infrastructure has been gained, Phase 1 can be considered as complete.

Within the network infrastructure, again various options are available for the malware to propagate; say sending further emails, using shared network drives, or similar. In any case, the goal of the learning and propagation phase (Phase 2) is to infect a maximal number of devices, or equivalently, infect a maximal portion of the infrastructure. This phase can comprise an intensive communication with the command and control server, even bidirectionally, in order to report the system details to the attacker, who in turn can deploy updates and hand-crafted specific attacks to help the malware infect more of the victim network. Once a "sufficient" number of machines has been infected, Phase 3 can be initiated.

In the final third stage of the APT, the actual malware is uploaded and deployed over the so-established network of bots. We stress that phase three may entail a relaunch of phase 2 (learning and propagation) inside the victim system. Since a full detailed discussion of phase repetitions within the model extend beyond the scope of this work, we leave this to future work.

As practical attacks demonstrated, the third phase is not necessarily a single powerful blow, but usually stealthy (such as in case of Stuxnet, which acted slowly to go undetected for a long time), and can be even a compound of several simultaneous attacks (such as was the case in the Ukranian power grid, where many power switches were opened, while a simultaneous denial-of-service attack was launched on the support hotlines of the energy provider).

Since these phases usually follow and depend on one another, it appears reasonable to model the whole APT as a sequential game, with (at least) three phases covering the above steps. The term "phase" is here used for a temporal period, as opposed to the spatial concept of a "stage", which describes the particularities of the network infra*structure*. Consequently, each of the three phases, sub-games in the overall sequential APT game, can be thought of its own (sequential) game, going through various stages in the network. Taking an email infection as an example, one stage can be gathering information for a forged identity, with the next stage being the acquisition of a forged email account, from which spam can be distributed in the company (third stage of Phase 1).

Likewise, in Phase 2, the stages may be more directly associated with the different networks in the company. That is, once the malware has been opened by a person with access to subnet X, the next stage could be considered as reached once the email has been forwarded and opened by a person with access to a (logically or physically disjoint) subnet Y, and so on.

The players (attacker and defender) can be different in each phase of the APT game. For example, the spammer causing the initial infection and having the dropper installed is not necessarily the same person that writes the actual damage program. Indeed, "cybercrime-as-a-service" is a contemporary term to describe the full-fledged supply-chains that have illegally evolved to offer the full spectrum of services required to mount a successful APT (ranging from spammers handling the initial infection, up to software engineers providing exploit kits that help with the propagation in Phase 2, etc.).

Thus, for a game-theoretic modeling, it is useful to think of the sequential game as a two-player game, with each player physically being a whole team of actors. This applies not only to the attacker but also to the defender since the counteractions against phishing (spam filters, awareness training, etc.) can be quite different from those on the technical network level (e.g., signature-based malware scanning, or similar).

However, the description allows to identify a few characteristics that help with setting up a game-theoretic model per stage. Specifically, Table 1 gives a high-level overview of strategies on either side (by either team), and the objective of

TABLE 1. Phase games and modeling: Each phase is composed of games specified with the set of players, the action sets (cf. e.g., [26] for further examples), the player objectives and payoff functions.

Phase	Attacker actions (player 2)	Defender's options (major actions available to player 1)	Game's objective	Model / payoff structure
1	social engineering, hacking, $\cdots$	raise security awareness, trainings, incident documentation and alerting, $\ldots$	initial infection (enter the next phases)	expected loss once the attacker is inside (player 1 is minimizing)
2	infect shared network drives, send infected email attachments/links, $\cdots$	virus/malware scans (spot-checking), re-installation of network devices, change system configurations, audits, penetration tests, $\cdots$	maximize the size of the infection (relative to the network size), get closer to the inner business assets	expected loss per stage (player 1 is minimizing)
3	misuse/damage of system	anomaly detection, logging, $\cdots$	maximal damage in the long run	expected damage in the long run (player 1 is minimizing)

the game in each phase. Note that the strategies listed there for player 1 (the defender) are cumulative, in the sense of not explicitly repeating strategies available in one stage to be also available in another phase. The displayed options are thus only major representatives of what can be done, and the table is in no way exhaustive.

While the overall game is sequential, each subgame can have its own and distinct structure. For example, the game modeling Phase 1 can be a simple repeated matrix game with 0/1-outcomes, since a failure of one infection attempt may just be repeated until success (we present a more sophisticated modeling next). Differently, Phase 2 may be modeled as a sequential game, since the move from one stage to the next (within the network) obviously depends on which past stages have been successfully conquered. The third game, in turn, may be a repeated game again, possibly involving an element of continuous time, since the action can be taken at any time, as long as it goes undetected.

A. MULTI-PHASE AND MULTI-STAGE APT-GAMES

We consider a phishing attack at Phase 1 aiming to steal the password of the administrator to get into the network. This phase will be followed by a stealthy attack (Phase 2 game) that targets the SCADA software (e.g., Siemens STEP 7 as in Stuxnet) of an industrial control system. The Phase 3 of the attack is to leverage the control of the software to launch a physical attack on the control system. Each phase can be modeled as a game.

In Phase 1, a popular way to make contact and initiate the intrusion is through social engineering. From the long list of techniques, which includes (spear)phishing emails, malicious gifts (USB gadgets or similar), baiting (seemingly lost USB sticks), waterholing, tailgating, etc., we will subsequently use *phishing* as a showcase model in Phase 1. The success of baiting, for example, is a different story than phishing and depends on whether or not foreign USB sticks or other similar malicious hardware can be connected to the system. Even if this is technically prevented, there is a residual risk of the USB stick being plugged in at the home computer, installing a key-logger to acquire login information for the remote enterprise network (e.g., when the user is working at home).

Phase 2 is put into more concrete terms in section III-B.2. It is as a multi-stage sequential game over N stages enumerated from 1 to N, where the N-th stage is the outermost one entered upon winning the Phase 1 game. The game played at stage 1 in Phase 2 lets the attacker enter the 0-th stage, in which the Phase 3 game is played.

At Phase 3, once the attacker has the control of the command and control of the system, it can launch any attack that serves its hidden agenda. In other words, the final consequence of the attack or the game payoff $I(0)$ will depend on the objective of the attacker at Phase 3.

Note that one reason to divide the attack into different phases is that APT can leverage different resources to achieve its attack objectives. The Phase 1 attack can be outsourced to botnets who specialize in the spearphishing. The Phase 3 attack requires the expert knowledge of the control systems.

B. DESCRIBING THE GAMES PER PHASE

Towards instantiating the model, we will use an example APT scenario upon whose description the games for each phase will be defined. Fig. 2 shows how the phase games that we concretely describe in the following are connected. We stress that this kind of modeling is not mandatory, but only a proposal that must be adapted to the particular situation at hand.

Most risk management is done in qualitative terms, meaning that there are hardly any precise figures available on likelihoods or impact (as we need to define our models). To capture this, we will define our games in a qualitative scoring $low/medium/high$, with each linguistic term having a fixed numerical representative that is specified below. Of course, the representative and scoring must be consistent throughout all phases and stages, for otherwise, the connections between the games (indicated in Fig. 1) would not be meaningful, as the parameters in one game would have different units to that in the next game. Moreover, we stress that representatives for the ranks in each game must be chosen with magnitudes of mutually compatible and meaningful interpretations in the application context (e.g., gain and cost should be measured in the same units and have similar numeric ranges. For example, if the gain is a million $, but the cost is only 300$, then the cost should be rated as *negligible*, relative to the gain being valued as *high*). The costs, gains, and likelihoods used in the following will thus be understood abstract and must be defined specifically w.r.t. the context of the model.

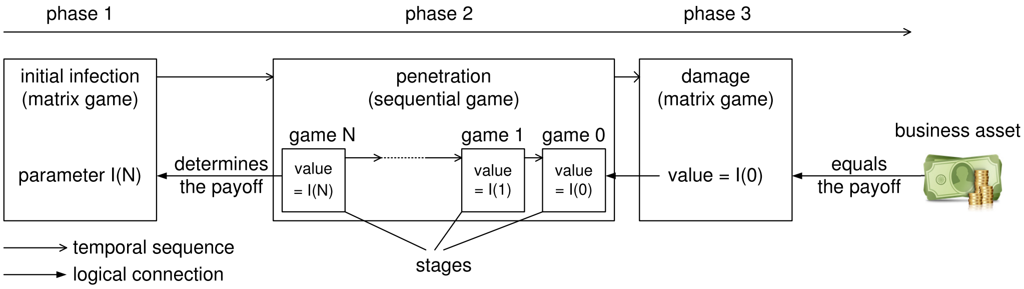

FIGURE 1. Multiple Phases of an APT Game: Phase 1 is the initial infection and Phase 2 is the penetration, followed by Phase 3 damage attacks. Phase 2 consists of multiple stages, each of which constitutes a game between the attacker and the system. The three phases are modeled by game models that are sequentially concatenated together.

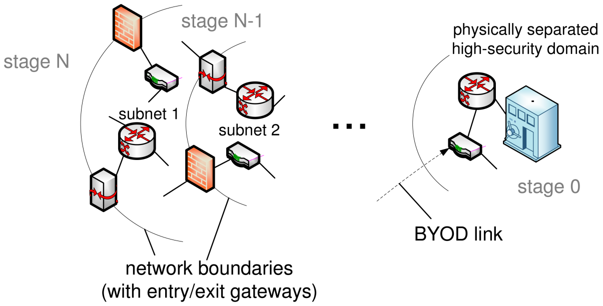
stage N
stage N-1
subnet 1
subnet 2
...
physically separated high-security domain
stage 0
BYOD link
network boundaries (with entry/exit gateways)

FIGURE 2. Example of Phase 2 game played over the infrastructure network.

We present in the following games at different phases. Each game has a unique structure. For example, a Bayesian game for Phase 1, a nested game for Phase 2, and a finite matrix game for Phase 3. The games presented below could be generalized to an arbitrarily large action set and multiple players. For reader's convenience, we first present simple illustrations of the games below and then generalize them in Section IV.

1) PHASE 1 GAME: SPEARPHISHING

As is well known from sources like [29], various social engineering techniques are available to overcome the initial barrier, which is in most cases tricking a human being into opening a malicious content. Each such social engineering technique constitutes another (attack) strategy in the game, with defense strategies corresponding naturally (see [29]). Assuming that the adversary will repeat a failed attempt, the game is also repeated, with stochastically independent trials and a finite number of strategies for the two players (attacker and defender).

In a *phishing game*, the attacker's strategy is to construct a message that looks normal to the user, while the defender's strategy is to determine whether the email is legitimate. We can use a Bayesian game framework to model the scenario. Let θ be the type of the sender. $\theta = 1$ means that the sender is the phisher. $\theta = 0$ means that the sender is a legitimate sender. The attacker can act legitimate or send a spear-phishing email. Sending a legitimate email does not launch an attack, but it will consume the cognitive attention of the user and make the user less on alert. The attacker can plant malware or a link to it in an illegitimate message, which can spread over the networks once opened. The targeted user or the receiver can determine to open the email with or without cognitive detection (D or ND). The cognitive detection requires a cost c_d. This Bayesian game can be represented by two payoff matrices:

For $\theta = 1$, we have

$$\mathbf{A}_1 = \begin{array}{c|c|c} & \text{L} & \text{NL} \\ \hline \text{D} & c_d & c_d + p_I I(N) \\ \hline \text{ND} & 0 & I(N) \\ \hline \end{array}$$

For $\theta = 0$, we have

$$\mathbf{A}_0 = \begin{array}{c|c} & \text{L} \\ \hline \text{D} & c_d \\ \hline \text{ND} & 0 \\ \hline \end{array}$$

In the payoff matrices above, the row player is the receiver or the targeted user who aims to minimize the cost. The column player is the sender who aims to maximize the cost in $\mathbf{A}_1$ or merely choose to act legitimate in $\mathbf{A}_0$, depending on the type of the player. The value $I(N)$ is the expected payoff for the attacker once it gets inside the infrastructure (anticipating another N protective layers to be overcome until the inner business asset can be attacked). Thus, the value $I(N)$ measures the loss upon a successful initial intrusion, equivalently, if the defender loses the Phase 1 game.

The parameter p_I is the probability of misdetection, i.e., the error made by the human cognition. (Here we assume that the spam filter has failed to filter the email.) The parameter p_I can be assessed using human cognitive models and survey studies in the literature of cognition and perception. In general, p_I can be time-varying and determined by past experience (learning), and can change over time. Here, we view p_I as a given parameter which models the error rates of the human decision-making.

a: MORE GENERAL PHASE 1 MODELS

It is evident that the phishing game sketched above can be generalized in various ways, e.g., as Stackelberg games [30], or similar. In any case, it is important to bear in mind that the attacker can get into the system on many paths, including *indirect* ones and the action space on both sides is expectedly larger than the binary example that we used above.

For example, it is not always necessary to inject malware through phishing, if a supplier for the victim system is the easier target and vehicle to bring in malicious code (say, if the third party from which a software is obtained can be hacked and the malware travels into the system on the legitimate ways by installing so-far trusted software from a supplier company; see [28] for an example). Our presentation is intentionally simplifying matters here but practically needs to be done on a much more fine-grained level.

2) PHASE 2 GAME: LEARNING AND PENETRATION

Once the malware has been placed initially, i.e., game 1 has been completed with at least one win for the second player, an APT usually continues to establish a stealthy backdoor through which an exploit kit can be deployed. This exploit kit is similar to a vulnerability scanner, as it probes and investigates its surrounding and can continuously report back to the command and control server, i.e., the adversary. The adversary, player two in this game, can then deploy updates to the malware and have it install specific code to penetrate and infect further parts of the system.

For game-theoretic modeling, it is useful to think of the infrastructure as being logically divided into several sections (zones, or similar), which manifest themselves as subnetworks. For high-security domains, these may even be physically separated from the outer networks, in which case a bring-your-own-device "link" may be possible, or social engineering attacks like tailgaiting may be a strategy to smuggle in a hardware keylogger to secretly read access credentials.

In our case, the phase-2-game can be defined by partitioning the whole infrastructure graph into subnets, enumerating them in order of distance to the central asset of interest. The Phase 2 game is sequential, being composed of a sequence of N games, each of which models the transition from one stage (subnet) to the next, starting from stage N until stage 0, where the valuable business assets can be attacked (and Phase 3 starts). Fig. 2 shows an example.

Let us call $\mathbf{B}_{2,n}$ the game modeling the jump from stage $n \leq N$ to stage $n-1 \geq 0$. The function $I : \{0, 1, \ldots, N\} \to \mathbb{R}$ gives the saddle-point value of each stage game. It denotes the consequence of the APT attack if it becomes successful over the stages.

The game-play itself is defined by the different possibilities to overcome the entry/exit points between the networks. Thus, to set up the game strategies, we need to consider the physical network devices and the possibilities to penetrate them. Drawing from Table 1, such strategies can be: (i) stealing login credentials (e.g., for VPN tunnels), (ii) using shared network drives, and (iii) exploiting outdated patch levels (and induced vulnerabilities of firewalls, ...).

Let us simplify matters again by assuming that each subnet has two connections to its neighbors (for reasons of availability by redundancy), and call the network connection points "firewall 1" and "firewall 2". Again, abstracting from the various concrete scenarios that may be definable in practice, let us sketch our example using the strategies "penetrate firewall 1" (PF1) and "penetrate firewall 2" (PF2), not saying how this is done concretely, although we distinguish attacks in terms of the likelihood to succeed, and the costs associated with mounting them. A specific attack j has thus a chance of $p_{j,i}$ to succeed under defense action i, at an attack cost of c_i, irrespectively of whether the game at this stage was won or lost (note that as a generalization, the costs may be different in each stage; we keep it constant for simplicity here).

The defender, in turn, may randomly reconfigure the firewall or security systems (e.g., by redirecting traffic over different firewalls, patching systems, activating honeypots, etc.). Thus, we let the defender have two strategies to reconfigure either entry point 1, the firewall (RCEP1) or entry point 2, the gateway (RCEP2). Both strategies aim at thwarting the attacker's current activities by changing the configuration so that the malware in its current version can no longer use the exploits discovered so far.

The payoff structure of the sequential game (at stage n) takes the form (with success rates p_1, p_2 and costs c_1, c_2 for both attacks) is denoted as $\mathbf{B}_{2,n}$ and defined as

	PF1	PF2
RCEP1	$p_{1,1} \cdot I(n-1) - c_1$	$p_{1,2} \cdot I(n-1) - c_2$
RCEP2	$p_{2,1} \cdot I(n-1) - c_1$	$p_{2,2} \cdot I(n-1) - c_2$

with the rationale that the transition from stage n to stage $n-1$ goes undetected if the defense is done at just the wrong firewall in this moment. Note that it is convenient though not necessary to assume a circular payoff structure here, meaning that no dominating strategies exist (in which case the game would degenerate to a single optimal choice for both players). This induces a unique mixed equilibrium that can be computed in closed form (since the game is 2×2), and thus defines the value function $I(n)$ for all stages $0 \leq n \leq N$ recursively.

Once the innermost stage $n = 0$ is reached, the Phase 3 game starts, defining $I(0)$ as the saddle point value of the final Phase 3 game.

3) PHASE 3 GAME: CAUSING DAMAGE

In this game, the primary goal is to cause the maximal damage over repeated rounds of the game. For that sake, the malware will take all actions to remain undetected and act hiddenly by causing only little damage but doing so continuously (such as was the case for Stuxnet). One example technique is using covert channels embedded within DNS requests [26], but many more ways are imaginable, each of which may become its own strategy in the Phase 3 game model.

The game for our example will thus be a static game with independent repetitions, in which the defender's action is restricted to checking the current system state by all available means (it is advisable not to rely on sensor data only, if the malware could have manipulated it, such as was done by Stuxnet). Hence, the defense strategies are spot checks on the infrastructure regarding its configuration, while the attack actions depend on the specific possibilities of the infrastructure.

Completing our example, the Phase 3 game is again a matrix game with, say, two strategies per player. For the defender, this would mean *spot checking* two *critical points* (strategies SCP1, SCP2), corresponding to the spots where the attacker may cause (some) damage (attack strategies are thus *damaging P1* (DP1) or *damaging P2* (DP2)). The cost structure may thus be given as

$$\mathbf{C} = \begin{array}{c|c|c} & \text{DP1} & \text{DP2} \\ \hline \text{SCP1} & 0 & \ell_2 \\ \hline \text{SCP2} & \ell_1 & 0 \end{array}$$

The losses for player 1 are herein defined as ℓ_i if the business asset i has been damaged. Thus, player 1 minimizes the overall damage, the value of which is denoted as $I(0)$, which will serve as an input to stage 1 of the phase 2 game. There are several options to quantify the loss parameters ℓ_1 and ℓ_2. One pragmatic way to categorize the loss using indicators by choosing $\ell_i \in \{0, 1, 2, \cdots\}$ to inform the level of damages that could incur. Another way to quantify the loss is to use the physical models for the assessment of the physical consequences when components or subsystems of the infrastructure system fail. This approach has been adopted in the recent works of [15] and [31] in which control system models are used to assess the impact on the feedback loop.

Of course, this game can be generalized in many ways, say, by allowing the attacker to be successful even in case of directly hitting a defense measure (scenarios (SCP1, DP1) and (SCP2,DP2), both being rated with zero loss for the defender in the above instance of the game), or by making the game also sequential to account for cumulative losses across repetitions. We leave both generalizations for future work here, for the sole sake of simplicity of the upcoming examples.

IV. GAME SOLUTIONS AND ANALYSIS

In this section, we will formally define the multi-phase multi-stage (MPMS) game and characterize the Nash equilibrium of the game. Let $i = 1, 2, \cdots, P$ denote the number of phases of the APT. In Section III-A, we have seen that the MPMS game is composed of three stages, i.e., $P = 3$. Each phase is composed of different types of games. Phase 1 is represented by a one-stage Bayesian game. Phase 2 is modeled by a sequential dynamic game of N stages, while Phase 3 is captured by a one-stage matrix game. Let G_{ij} be the game at phase i and stage j, and S_i be the number of stages at phase i. In our example, we have $S_1 = 1, S_2 = N, S_3 = 1$. The main elements of each game $G_{i,j}$ can be represented by the triplet $\langle \mathscr{P}_{i,j}, \mathscr{A}_{i,j}, \mathscr{U}_{i,j} \rangle$, where $\mathscr{P}_{i,j}$ is the set of player of the game; $\mathscr{A}_{i,j} = \{A^p_{i,j}, p \in \mathscr{P}_{i,j}\}$ is the (pure) action sets of the players with player p's (pure) action set denoted by $A^p_{i,j}$; $\mathscr{U}_{i,j} = \{U^p_{i,j}, p \in \mathscr{P}_{i,j}\}$ is the set of utility functions or preferences of the players with player p's utility function denoted by $U^p_{i,j} : \Pi_{p \in \mathscr{P}_{i,j}} \mathscr{A}^p_{i,j} \to \mathbb{R}$. The mixed strategies of the players are denoted by $x^p_{i,j} \in \Delta(A^p_{i,j}), p \in \mathscr{P}_{i,j}$. The average payoffs under mixed strategies are denoted by $\bar{U}^p_{i,j} : \Pi_{p \in \mathscr{P}_{i,j}} \Delta(\mathscr{A}^p_{i,j}) \to \mathbb{R}$. We use $\mathscr{G}_i := \{G_{ij}, j = 0, \ldots, S_i\}$ to denote the Phase 1 game, and $\mathbf{G}(i, i')$ to denote the MPMS game from Phase i to Phase i', where $i \leq i', i, i' = 1, \cdots, P$.

The players of games can be different at distinct phases and stages. Game $G_{1,1}$ is the game between a sender (SD) and a targeted user (TU). Game $G_{2,j}, j = 0, 1, \cdots, N$, is a game between a network administrator (NA) and a propagating malware (PM). $G_{3,1}$ is a game between the critical infrastructure designer (CI) and an attacker (AT). To capture the distinct feature of the Bayesian game of $G_{1,1}$ and the dynamics of the Phase 2 game, we can also include the type space $\Theta = \{0, 1\}$, and the transition kernel $\mathscr{K}$, respectively. The type space Θ will extend the triplet representation of the Phase 1 game to a quadruplet $\langle \mathscr{P}_{i,j}, \mathscr{A}_{i,j}, \mathscr{U}_{i,j}, \Theta \rangle$ with the action sets and utility functions of the players being type-dependent, i.e., $\mathscr{A}_{i,j} = \{A^{p,\theta}_{i,j}, p \in \mathscr{P}_{i,j}, \theta \in \Theta\}$ and $\mathscr{U}_{i,j} = \{U^{p,\theta}_{i,j}, p \in \mathscr{P}_{i,j}, \theta \in \Theta\}$. The dynamic structure of the multi-stage Phase 2 game can be represented by the kernel of the dynamic game.

The composition of the games at three phases captures the human-cyber-physical nature of the APT attack on critical infrastructure, in which the adversary first makes use of human vulnerabilities and then propagate malware through cyber defense, and finally attacks the physical layer of the infrastructure. Note that the sequential composition of the games leads to a game of games. Games at each phase are of different types. Hence, an appropriate solution concept is needed to capture the heterogeneous nature of the game of games. A common solution concept for Bayesian games is Bayesian Nash equilibrium (BNE). The one for sequential games is sub-game perfect Nash equilibrium (SPNE), and the one for matrix is game is simply (mixed strategy) Nash equilibrium (NE). For the sequentially composed game G, we propose the notion of Gestalt Nash equilibrium (GNE) which allows the equilibrium to retain the solution feature of the decomposed games as well as have sub-game perfectness of the sequentially composed game. To formally define GNE, we let $x^p_{i,j}$ be the mixed strategy of a player $p \in \mathscr{P}_{i,j}$ at phase i and stage j. It is clear that for the game described in Section III-A, we have $\mathscr{P}_{1,1} = \{SD, TU\}$, $\mathscr{P}_{2,j} = \{NA, PM\}, j = 0, \cdots, N$.

Phase 1 Game: We start with the analysis of the phishing game at Phase 1. For $G_{1,1}$, the payoff to the user is in terms of cost while utility to the sender. The mixed-strategy of the sender is distinguished by $x^{SD,0}_{1,1}$ and $x^{SD,1}_{1,1}$ with the former for the legitimate sender and the latter for the illegitimate sender. The mixed-strategy profile $x^{TU*}_{1,1}$, $x^{SD*}_{1,1} = \{x^{SD,0*}_{1,1}, x^{SD,1*}_{1,1}\}$ constitutes a BNE of G_{11} if for all $x^{TU}_{1,1} \in \Delta(A^{TU}_{1,1}), x^{SD,\theta}_{1,1} \in \Delta(A^{SD,\theta}_{1,1}), \theta = \Theta$,

$$\bar{U}^{TU}_{1,1}(x^{TU*}_{1,1}, x^{SD*}_{1,1}) \leq \bar{U}^{TU}_{1,1}(x^{TU}_{1,1}, x^{SD*}_{1,1}), \tag{1}$$

$$\bar{U}^{SD,0}_{1,1}(x^{TU*}_{1,1}, x^{SD*}_{1,1}) \geq \bar{U}^{SD,0}_{1,1}(x^{TU*}_{1,1}, \{x^{SD,0}_{1,1}, x^{SD,1*}_{1,1}\}), \tag{2}$$

$$\bar{U}^{SD,1}_{1,1}(x^{TU*}_{1,1}, x^{SD*}_{1,1}) \geq \bar{U}^{SD,1}_{1,1}(x^{TU*}_{1,1}, \{x^{SD,0*}_{1,1}, x^{SD,1}_{1,1}\}), \tag{3}$$

In the context of the phishing game defined by the cost matrices $\mathbf{A}_0$ and $\mathbf{A}_1$, condition (1) can be rewritten as solving the following optimization problem

$$\min_{x_{1,1}^{TU}\in\Delta(A_{1,1}^{TU})} \bar{U}_{1,1}^{TU}(x_{1,1}^{TU}, x_{1,1}^{SD*}) := \sum_{i\in\Theta} p_{1,1}^{i}(x_{1,1}^{TU})^T \mathbf{A}_i x_{1,1}^{SD,i*}$$

where $p_{1,1}^i$ is the prior probability of a sender interacting with a player of type i. Similarly, conditions (2) and (3) can be also equivalently represented by the following optimization problems for $i \in \Theta$,

$$\max_{x_{1,1}^{SD,i}\in\Delta(A_{1,1}^{SD,i})} \bar{U}_{1,1}^{SD,i}(x_{1,1}^{TU*}, \{x_{1,1}^{SD,i}, x_{1,1}^{SD,-i*}\}) := (x_{1,1}^{TU})^T \mathbf{A}_i x_{1,1}^{SD,i*}$$

We denote the expected payoff to the user at the equilibrium as $\tilde{U}_1^{TU} = \bar{U}_{1,1}^{TU}(x_{1,1}^{TU*}, x_{1,1}^{SD*}) = \sum_{i\in\Theta} p_{1,1}^i (x_{1,1}^{TU*})^T \mathbf{A}_i x_{1,1}^{SD,i*}$.

Note that payoff matrices $\mathbf{A}_0$ and $\mathbf{A}_1$ of $G_{1,1}$ depend on the value $I(N)$, which is the value of the outermost stage of the Phase 2 game. Hence, the mixed strategies obtained using (1), (2) and (3) depend on the equilibrium outcome of the Phase 2 game.

Phase 2 Game: At Phase 2, the game is composed of a sequence of multi-stage penetration games. At stage n, the payoff function is represented by the a matrix $\mathbf{B}_{2,n}$, i.e.,

$$\bar{U}_{2,n}^{PM}(x_{2,n}^{NA}, x_{2,n}^{PM}) = (x_{2,n}^{NA})^T \mathbf{B}_{2,n}(x_{2,n}^{PM}) = -\bar{U}_{2,n}^{NA}(x_{2,n}^{NA}, x_{2,n}^{PM}).$$

At each stage, the saddle-point equilibrium of the zero-sum game $x_{2,n}^{NA*} \in \Delta(A_{2,n}^{NA})$, $x_{2,n}^{PM*} \in \Delta(A_{2,n}^{PM})$ satisfy the inequalities for every $j = 0, \cdots, N$, $x_{2,n}^{NA} \in \Delta(A_{2,n}^{NA})$, $x_{2,n}^{PM} \in \Delta(A_{2,n}^{PM})$,

$$\begin{aligned}(x_{2,n}^{NA})^T \mathbf{B}_{2,n}(x_{2,n}^{PM*}) &\geq (x_{2,n}^{NA*})^T \mathbf{B}_{2,n}(x_{2,n}^{PM*}) \\ &\geq (x_{2,n}^{NA*})^T \mathbf{B}_{2,n}(x_{2,n}^{PM}).\end{aligned}$$

The value of the game $I(n)$ is defined by the payoff matrix $\mathbf{B}_{2,n}$ at stage n as follows: $I(n) = (x_{2,n}^{NA*})^T \mathbf{B}_{2,n}(x_{2,n}^{PM*})$. Since the payoff matrix $\mathbf{B}_{2,n}$ depends on the outcome of the game at the following stage, the game value $I(n)$ is related to $I(n-1)$ through the following recursive relation

$$I(n) = val(\mathbf{B}_{2,n}(I(n-1))), \tag{4}$$

where val is the value operator that maps a payoff matrix to the value of the game. Here, the payoff matrix $\mathbf{B}_{2,n}$ explicitly depends on $I(n-1)$ (and possibly also $I(n)$ under a more generalized modeling).

The respective saddle-point strategies at stage n can be obtained in the same way, but now depend on $I(n-1)$ (and through this, indirectly also on the saddle point strategies for the previous stages; however, a more direct calculation is the recommended choice in practice):

$$x_{2,n}^{NA*} \in \operatorname*{argmin}_{x_{2,n}^{NA}\in\Delta(A_{2,n}^{NA})} \left[\max_{x_{2,n}^{PM}\in\Delta(A_{2,n}^{PM})} \mathbf{B}_{2,n}(I(n))\right]$$

$$x_{2,n}^{PM*} \in \operatorname*{argmax}_{x_{2,n}^{PM}\in\Delta(A_{2,n}^{PM})} \left[\min_{x_{2,n}^{NA}\in\Delta(A_{2,n}^{NA})} \mathbf{B}_{2,n}(I(n))\right]$$

Phase 3 Game: At the final stage of the Phase 2 game, the payoff matrix $\mathbf{B}_{2,1}$ depends on the outcome of the game of the Phase 3 game, in which the attacker selects targets and the infrastructure chooses to invest resources on the protection. The value of the Phase 3 game is denoted by $I(0)$, given by

$$I(0) = (x_{3,1}^{CI*})^T \mathbf{C}(x_{3,1}^{AT*}),$$

where $\mathbf{C}$ is the payoff matrix of the Phase 3 game; $x_{3,1}^{CI*} \in \Delta(A_{3,1}^{CI})$, $x_{3,1}^{AT*} \in \Delta(A_{3,1}^{AT})$ are the saddle-point strategies of the Phase 3 game.

Gestalt Game Equilibrium: The computation of the GNE of the MSMP game $\mathbf{G}(1, P)$ will require a backward induction from the last stage of the last phase and propagate backward to the first stage of the first phase. The backward induction will yield a sub-game perfect GNE in mixed strategies given by the profile

$$x_{\mathbf{G}(1,P)}^* = (x_{i,j}^{p*}, p \in \mathscr{P}_{i,j}, i = 1, \ldots, P, j = 0, \ldots, S_i),$$

and the corresponding game values $v_{\mathbf{G}(1,P)}^* = (\tilde{U}_1^{TU}; I(n), n = 1, \cdots N; I(0))$. The strategy profile $x_{\mathbf{G}(1,P)}^*$ has sub-game perfectness properties. The truncated mixed strategy is $x_{\mathbf{G}(i,P)}^* = (x_{i',j}^{p*}, p \in \mathscr{P}_{i',j}, i' = i, \cdots, P, j = 0, \cdots, S_{i'})$ also constitutes a sub-game perfect GNE in mixed strategies of the truncated game $\mathbf{G}(i, P)$.

The existence of a GNE in mixed strategies can be guaranteed since each stage constitutes a finite game and the existence of its equilibrium is ensured. The game values $v_{\mathbf{G}(1,P)}^*$ are good indicators of the risk at a given stage. Due to the sequential features of the game, the risk at the first stage and the first phase $\tilde{U}_1^{TU}$ includes the risks at later stages and phases. Risk assessment using MSMP games gives a holistic view of the impact of APT on the entire cyber, physical and human layers of the infrastructure. The risk assessed at one stage endogenously integrates the long-term risk of the system, which is less myopic than independent stage-by-stage or phase-by-phase risk assessment. The risk estimated at Phase 1 will inform the targeted users of the potential consequences of their action, and reduce reckless decision-making of the users. The mixed strategies obtained at each stage $x_{\mathbf{G}(1,P)}^*$ gives optimal defense strategies against a strategic attacker. These strategies can be automated and built into the cyber defense mechanism of the infrastructure. For example, an automated alert system could be established to recommend responses to received messages. The firewalls and intrusion detections can be configured according to the mixed strategies to defend against the penetration of APTs. The physical control system can be further reconfigured to mitigate the loss at the final phase.

The analysis of the equilibrium provides not only a risk assessment benchmark but also a way to design and deploy cyber mechanisms to further reduce the risks and achieve a guaranteed level of security. The mechanism design can include the design the action space, payoffs and information structure. To achieve the desired security guarantee, we can introduce additional strategies in the action space by investing in new technologies. The payoff structures can be changed by introducing penalties and incentives for the users.

Information asymmetry can also be used to reduce the footprint of the system and the knowledge of the attacker.

A. INSTANTIATING THE MODEL AND A NUMERICAL EXAMPLE

As Fig. 1 implicitly suggests, instantiating the model is done *backwards* starting from the last phase. Picking up the APT modeling example, we would thus start from the inner most Phase 3 game and compute its equilibrium value $I(0)$.

1) PHASE THREE GAME

So, let us redefine the above payoff structure to protect an asset of high value, with the scale being $(low, medium, high) = (1, 2, 3)$. Each scale measures the worst-case outcome of an action in appropriate units. It gives the cost structure $\mathbf{C}$ for the Phase 3 game (with player 1 minimizing the losses $\ell_1 = \ell_2 = 3$ $(high)$) as

$$\mathbf{C} = \begin{array}{c|c|c} & \text{DP1} & \text{DP2} \\ \hline \text{SCP1} & 0 & 3 \\ \hline \text{SCP2} & 3 & 0 \\ \hline \end{array}$$

Its value is $val(\mathbf{C}) = 3/2$, which equals the starting point $I(0) = val(\mathbf{C}) = 3/2$ for the sequential Phase 2 game.

Solution and Interpretation: The expected damage in the Phase 3 game is 3/2, which is exactly between "low" and "medium" (whatever this may mean for the real application). The optimal defense action is to choose strategies SCP1 and SCP2 equiprobable (i.e., following the Nash-equilibrium strategy profile for the defender and the attacker)

$$(x_{3,1}^{CI*}, x_{3,1}^{AT*}) = ((1/2, 1/2), (1/2, 1/2)).$$

Likewise, the attacker is best off by the same strategy as the defender. If the defender sticks with this behavior, then the (implicit) assumption on this game to be zero-sum assures that the expected damage of 3/2 is *optimal* against all alternative adversarial behavior (i.e., cannot be increased by the attacker).

2) PHASE TWO GAME

Now, using the 2×2-game structure sketched in section III-B.2, we can plug concrete values into the game defining $I(n)$ to compute the sequence of equilibrium values per stage recursively. In the usual absence of exact values for probabilities, we will quantify the parameters in a similar discrete scale being $(low, medium, high) = (0.1, 0.5, 0.9)$, say, based on a subjective consensus among the experts concerned with the infrastructure.

We will use $p_{1,1} = 0.1$ (low), reflecting that it is quite unlikely to easily overcome firewall 1 if its configuration is unknown. Firewall 2 is assumed to be weaker, admitting a medium chance of penetration even if the configuration has been changed shortly before; $p_{2,2} = 0.5$ $(medium)$. Otherwise, if the configurations are known, the attacker has a good chance to break through, making $p_{1,2} = p_{2,1} = 0.9$ $(high)$. For the costs, assume that it is more expensive to attack firewall 1 ($c_1 = 2$ $(medium)$) than firewall 2 ($c_2 = 1$ (low)).

TABLE 2. Equilibria for Phase 2 game.

player →	Defender's equil. strategy $x_{2,n}^{NA*}$			Attacker's equil. strategy $x_{2,n}^{PM*}$		
	stage			stage		
strategy ↓	1	2	3	1	2	3
RCEP1	0,89	0	1	1/3	0	0
RCEP2	0,11	1	0	2/3	1	1

This choice of parameters endows $\mathbf{B}_{2,n}$ with a unique equilibrium in mixed strategies for (at least) the stages $n = 1, 2, 3$, along with the respective values.

Using this setting, we get the following sequence up to stage no. $N = 3$ as $(I(n))_{n=0}^{3} \approx (3/2, 2.28, -0.14, 1.13)$, so that $I(N) \approx 1.13$ can go into the Phase 1 game as a cost parameter.

The equilibrium strategies per stage are obtained for both players (network administrator (NA) and propagating malware (PM)) by minimizing the loss for player 1, i.e., maximizing the gain for player 2. The results are listed in Table 2. Interestingly, and not shown in the table, continuing the recursive computation of equilibria along (4), it is revealed that the pure strategy $x_{2,n}^{NA*} = (0, 1)$ remains optimal for $n > 3$ (at least up to $n = 17$), and the respective value becomes stationary around $\approx 2/3$. This means that – in this particular example – it would not make any sense to add more stages (assuming that they look the same as the three that we have) since this adds nothing to the overall security. This observation also states an important fact that the defense-in-depth strategies in cyber defense is effective for protection up to some depth. It is shown by the recursive formula (4) that there exists a limit point as the number of stages increases.

Solution and Interpretation: The values $I(1), I(2), \ldots$ represent the worst-case damages expected per stage, provided that the defender in each stage acts according to the Nash equilibrium in the respective stage. The equilibrium is found exactly as for a conventional matrix game, only using the values $I(n)$ and $I(n-1)$ in the respective payoff matrix in the n-th stage. Concretely, we find the equilibria listed in Table 2, where the equilibria in stage 2 and 3 are pure strategies for the defender (as opposed to a unique mixed equilibrium in stage 1).

3) PHASE ONE GAME

Since this is a Bayesian game, player 1 (the defender) does not know against whom it is playing, so nature chooses the type of the email sender ($\theta = 0$ for an honest sender, and $\theta = 1$ for a phisher). In this Bayesian game, player 2 gets told the type, while player 1 (the remail recipient) remains unaware. Given this setting, we compute the Bayesian Nash equilibria as follows:

In matrix $\mathbf{A}_1$, let us plug in the value $I(N)$ obtained from the sequential Phase 2 game, and assume that the cost for cognitive detection to be "medium" (which may reasonably reflect that efforts need to be spent on checking an email for originality before opening any of its attachment or clicking on weblinks in it). This amounts to setting $c_d = 2$. Furthermore, let us take $p_I = 0.9$ to express that once the phishing email has passed the spam filter and perhaps was sent by a spear phisher, there is a high chance of the person opening the email (say, if it believes that the technology would have warned it about any danger). Under this setting, we instantiate the payoff matrices $\mathbf{A}_1$ and $\mathbf{A_0}$ to be

$$\mathbf{A}_1 = \begin{pmatrix} 2 & 3.01 \\ 0 & 1.13 \end{pmatrix}, \quad \mathbf{A}_0 = \begin{pmatrix} 2 \\ 0 \end{pmatrix}$$

Since player 2 knows its own type (honest or phisher), it has best responses in either game, being

$$x_{1,1}^{SD,1*} = (0,1), \quad \text{and } x_{1,1}^{SD,0*} = (1)$$

with the latter strategy being no real choice, but given here for the sake of completeness. As these are the only equilibria and they are in pure strategies, the payoff perceived by player 1 will be determined by the likelihood θ (slightly abusing the symbol θ for the type in Θ to also denote its probability of occurrence) of which game matrix ($\mathbf{A}_0$ or $\mathbf{A}_1$) is being used. Thus, the overall payoff for the first player (the receiver of the email who does not know which type the sender is; phisher or legitimate), comes to

$$U_{1,1}^{TU} = \theta \cdot \mathbf{A}_1 + (1-\theta) \cdot \mathbf{A}_0 = \begin{pmatrix} c_d + p_I \cdot \theta \cdot I(N) \\ \theta \cdot I(N) \end{pmatrix}$$

The optimal behavior for the receiver of the email thus depends on which cost is lower among $\{c_d + p_I \cdot \theta \cdot I(N), \theta \cdot I(N)\}$. Using our values from before, and assuming a spear phishing to be not very likely ($\theta = 0.1$ (*low*)), we obtain $U_{1,1}^{TU} = (2.91, 1.01)^T$ with a pure strategy equilibrium (for the minimizing player 1) as $x_{1,1}^{TU*} = (0,1)$, giving a value $val(G_{1,1}) \approx 1.01$.

Solution and Interpretation: The equilibrium payoff in the Phase 1 game represents the average ''entry hurdle'' that an attacker has to overcome to break into the infrastructure at first. Expressed in our qualitative scale *low* < *medium* < *high*, we would rate the value 1.01 slightly above low, so the protection is indeed quite good.

4) GESTALT NASH EQUILIBRIUM FOR MPMS GAMES

The three-phase game analysis leads to a GNE for the MPMS game

$$x^*_{G(1,P)} = \Big(\Big\{x_{1,1}^{SD,0*}, x_{1,1}^{SD,1*}, x_{1,1}^{TU*}\Big\}, \\ \Big\{x_{2,n}^{NA*}, x_{2,n}^{NA*}\Big\}_{n=1}^{3}, \Big\{x_{3,1}^{CI*}, x_{3,1}^{AT*}\Big\}\Big),$$

each component of which has been analyzed in Sections IV-A.1, IV-A.2, IV-A.3. The associated value of the game is

$$v^*_{G(1,P)} = (1.01, -1.13, 0.14, -2.28, 3/2),$$

being interpreted as the *loss* for the defender (overall minimizing player) on a *qualitative* scale. That is, the values point at the category within or the categories between which the real loss will be. Observe that a negative loss for the defender is actually a gain, or equivalently, *cost* for the attacker. In the numerical example, if the cost for an attack would exceed the value of the inner business asset (in our case 3, as it is rated ''high''), then attacking the system is simply not economical. The negative values suggest that the strategic attack of the system at the corresponding stages is costly.

The game value 1.01 at the first stage is an important risk assessment for thwarting the attack at an earlier stage. By taking into account the effectiveness of the defense of the ensuing stages and phases, a target user can make a more comprehensive and informed decision to avoid reckless actions. The user can increase its cognitive detection by investing more attention in checking the content of the email or reducing its cognitive load or cost using assistive software tools. We can see that as we reduce the cognitive cost c_d to less than 0.1 (e.g., by extending the qualitative scale with another level ''negligible'' represented by a value < 0.1), the user will adopt mixed strategy $(1, 0)$ at Phase 1, indicating that the users will pay attention when the impact of inattention (on future stages) is known to be high. Similarly, the GNE informs the earlier stage of defense to be aware of its consequence on future stages. This approach avoids myopic defense decisions and improves the efficiency of defense-in-depth mechanisms.

To make this assessment in our case concrete, let us compute the accumulated costs for the attacker in each phase. It is given by adding the averages along all stages, which amounts to computing $v^*_{G(1,P)} \cdot \mathbf{1}^T \approx 0.75$. This means that the attack is indeed worth mounting since the reward to be expected is at least 0.75 (though not the full value 3 of the business asset). The same calculation tells the defender that the protection is indeed quite good, as the expected losses are less than *low* on the qualitative scale.

B. PRACTICALLY COMPUTING EQUILIBRIA AND LEARNING

In this section, we will introduce learning mechanisms that allow the computation of the mixed-strategy GNE and the learning of the optimal strategies online. The mechanism is composed of multi-stage distributed and asynchronous fictitious-play learning algorithm in which a player at each stage and phase updates his mixed strategies based on his local observations [32], [33]. Let $f_{i,j,k}^p = [f_{i,j,k}^p(a_{i,j}^p)]_{a_{i,j}^p \in A_{i,j}^p} \in \Delta(A_{i,j}^p)$ be the mixed strategy of player p at stage j of phase i, with each entry denoting the empirical probability that player p chooses an action $a_{i,j}^p$. At each iteration, the player update their strategies $f_{i,j,k}^p$ for every $i = 1, \cdots, P$ and $j = 1, \cdots, S_i$,

$$f_{i,j,k+1}^p = f_{i,j,k}^p + \frac{1}{k+1}(e_{i,j,k}^p - f_{i,j,k}^p), \tag{5}$$

where $e_{i,j,k}^p = [e_{i,j,k}^p(a_{i,j}^p)]_{a_{i,j}^p \in A_{i,j}^p}$ is a vector with $e_{i,j,k}^p(a_{i,j}^p) = 1$ if at time k, player p chooses action $a_{i,j}^p$ and otherwise $e_{i,j,k}^p(a_{i,j}^p) = 0$. At iteration k, each player chooses the best action in response to the empirical strategies of the

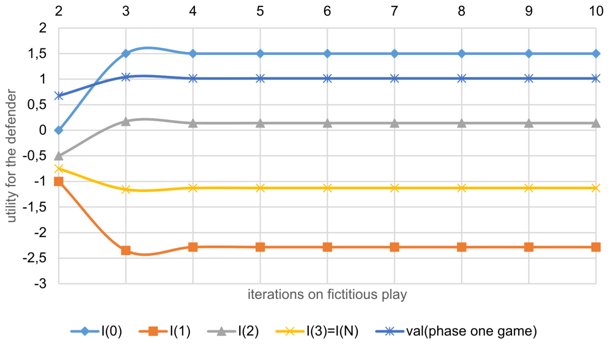

FIGURE 3. Speed of "learning" an optimal protection (GNE) using fictitious play.

other players, i.e.,

$$d^p_{i,j} \in \arg\,\mathrm{opt}_{d^p_{i,j}\in\mathscr{A}^p_{i,j}} \bar{U}^p_{i,j}(d^p_{i,j}, f^{-p}_{i,j,k}), \tag{6}$$

where $f^{-p}_{i,j,k} = \{f^{p'}_{i,j,k}, p' \in \mathscr{P}_{i,j}\backslash\{p\}\}$ and opt stands for either max or min depending on the player. The same algorithms (5) and (6) will be applied at every stage and phase. The action of the players at one stage can affect the payoffs of the actions at other phases or stages. Note that the convergence of the fictitious play algorithm is guaranteed for two-person two-action zero-sum games [34]. Hence the zero-sum game at penultimate stage will converge after the empirical mixed strategies at the final stage of the last phase converge. Using the argument of backward induction, we can show the learning mechanism will finally converge to a GNE. The algorithm provides an adaptive scheme for the users to change their configurations of defense strategies in response to either changes in the system and the additional knowledge of the attacker. Moreover, the algorithm provides a scalable and distributed way to compute the GNE in which the learning algorithm is run in parallel at every stage and phase of the system. The learning mechanism can be also used as a modeling tool to understand human behaviors at the first phase where the user attempts to learn from his past observations. In addition, the mechanism for Phase 2 and 3 can be implemented as a built-in protocol to defend against APT across the stages.

In Fig. 3, we have applied the fictitious play (FP) algorithms to compute the equilibria per phase/stage individually, and then put them together into an (approximate) GNE. The appeal of FP here is it reflecting the learning and experience of the actors (TU, NA and CI), in responding optimally to whatever they have learned about the attacker's action in the past. Fig. 3 indicates that already little experience (at least 4 iterations) suffices to get a quite accurate approximate GNE.

C. USING APT GAMES FOR SECURITY DESIGN

Having found that the values (individual risk assessments) are unsatisfying, we are free to alter the games in various ways, such as parameters, but also action spaces and payoffs. For example, there is no theoretical limit preventing us from working with an $m \times n$-matrix ($m, n > 2$) in all three stages, or in adopting a more fine-grained view on the costs (say, letting them be continuous rather than discrete variables). Actually, such an adaption is a mandatory part of higher level security management processes like ISO27000 anyway, which prescribe a cycle of *planning*, *doing*, *checking* and *acting* (PDCA-cycle). Our work fits into this cycle in providing a tool for three out of these four phases, as we can do planning (corresponds to game modeling), acting (corresponds to enforcing the equilibrium defense strategies in the daily business), and checking (amounting to deciding if the obtained values are satisfying or not).

Another aspect of the security design is related to the human factor. As we have seen in the GNE of the preceding numerical example, the cost of cognitive detection can be reduced by investing in automated detection and alert systems or removing the cognitive load of the users. The cognitive design of the users will be useful to move the equilibrium of the Phase 1 game toward a desirable target.

D. INSIGHTS FROM THE MPMS GAME MODEL

The practical appeal of dividing an APT into phases and modeling each phase as its own game, with interdependencies across the entire sequence of games has manifold advantages. First, focusing on the costs (thus minimizing efforts for the defender) enforces to adopt an economic viewpoint in the protection. Indeed, security is not about gaining something, but about avoiding loss, so minimizing losses (or costs) is a natural way to think about what the model helps us with.

Second, the interconnection between the games as described here helps us stepwise extend our view from a purely local risk assessment to one that covers the entire infrastructure components and their interplay. This systematizes the approach of reasoning about the whole system based on local security assessments.

Third, the games can be defined over scores that may be tailored to the specific context. This is an advantage and a drawback at the same time, since it necessarily makes the results subjective to some extent (indeed, other choices of representative values for the nominal scale *low* < *medium* < *high* are equally admissible, but fixing a common understanding of qualitative risk levels in numeric terms (insofar applicable) is a standard prerequisite in a risk management process),[1] but it greatly eases *risk communication* and aids *risk awareness*. In fact, a core requirement of a good risk management process is to formulate risks in terms being understandable for decision makers. The beauty of game-theoretic modeling lies in the freedom to let the decision maker define the scale a priori in which the results are being presented a posteriori. It must be noted, however, that the numbers coming out of the game have no meaning by the absolute values, but can only express *rankings* in terms of the risk being, say "above medium" (for a computed value of,

[1] Many national risk assessments specify concrete figures; for example Sweden [35], the Netherlands [36], Switzerland [37], Germany [38], etc.

TABLE 3. Different model parameterizations.

Var.	Meaning	Parameter set		
		1	2	3
$p_{1,1}$	likelihood to penetrate (even against active defense)	0.1	0.5	0.1
$p_{2,1}$		0.9	0.9	0.5
$p_{1,2}$		0.9	0.9	0.5
$p_{2,2}$		0.5	0.1	0.1
c_1	cost to penetrate firewall 1	2	2	1
c_2	cost to penetrate firewall 2	1	1	2
c_d	effort for a human user to detect a phishing email	2	1	3
p_I	likelihood to suspect phishing at all	0.9	0.5	0.9
θ	likelihood to get an email from a phisher	0.1	0.5	0.5
ℓ_1	loss if business asset 1 is damaged	3	3	3
ℓ_2	loss if business asset 1 is damaged	3	3	3

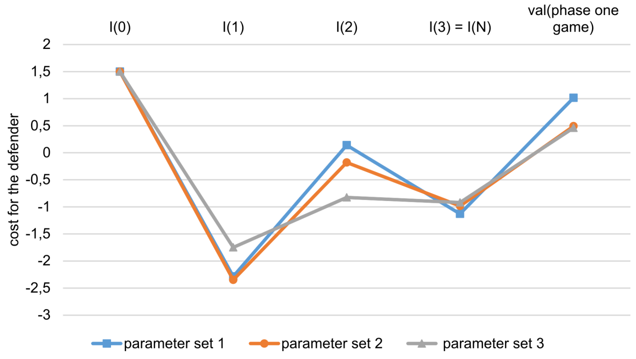

FIGURE 4. GNE equilibrium values per stage for the parameter sets in Table 3.

TABLE 4. Different subjective scales.

	probability			cost		
	low	med.	high	low	med.	high
scale 1	0.1	0.5	0.9	1	2	3
scale 2	0.01	0.5	0.7	5	10	30
scale 3	0.01	0.3	0.5	1	5	10
	$p_{1,1},\theta$	$p_{2,2}$	$p_{2,1},p_{1,2}$	c_2	c_1,c_d	ℓ_1,ℓ_2

say 2.3 > *medium* = 2 or "very high", if the game delivered a value > 3 if "high" is represented by the value 3).

To see how the results depend on the (subjective) choices of parameters and even the scale, we repeat the numerical example of section IV-A. In addition to the original parameter set, we use two more and different parameterizations of the model, which are listed in Table 3.

Solving the games (as described in section IV-A and plotting the results shows the different curves displayed in Fig. 4. Unsurprisingly, the numerical results are different, however, quite interestingly, the *shape* of the curves is retained over variations of the parameters. This indicates an "objective" lesson to be learned here, which is the following:

In entering **Phase 1**, the equilibrium being pure reflects what we observe in real life applications, namely the considerable success of phishing, due to people tending to be unaware and opening malicious mail. In our modeling, this effect is represented by the cost c_d, which a user needs to invest in order to detect the phishing. Whenever possible, a user may prefer not to invest too much effort into manual email checking, which our results corroborate.

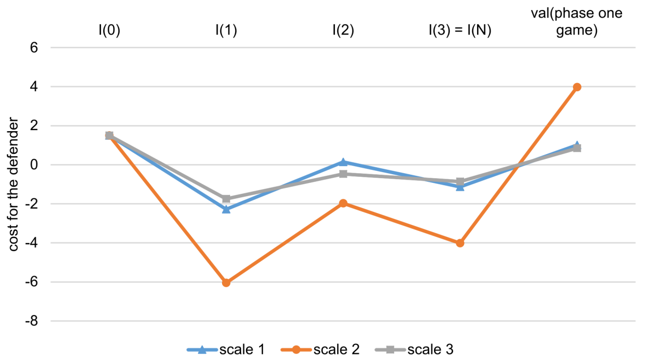

FIGURE 5. GNE equilibrium values under different scales as listed in Table 3.

During **Phase 2**, the zig-zag behavior of the value sequence, in our view, reflects the common effect of concentrating a defense on the weakest spot, while a change of the defense focus is only done upon knowing that the defense has been suboptimal before. That is, a person will be reluctant in changing a working protection; however, across stages, efforts may be distributed to protecting both firewalls evenly much (resulting in the pure strategies obtained for the defender in Table 2), and leading to the respective optimal attack strategies in turn. Our sequential game model reproduces this natural (intuitively expected) human behavior, but by varying its parameters, allows to optimize the defense under the natural human way of implementing it.

In Phase 3, the attacker's goal is obviously to avoid detection while causing maximal damage. Some of these attempts may be successful, while some others may be not, in which case the attacker will certainly repeat its trial (using a different technique then). So, the overall damage will certainly accumulate but can be bounded on average in the long run. The value obtained from the Phase 3 game is thus a measure of security risk of the physical infrastructure, which is used to assess the risks at the cyber layer.

Likewise, we may ask for the change in the results if different interpretations of the qualitative scale low/medium/high exist. Fixing the model parameters to the qualitative choices of Section IV-A (parameter set 1 in Table 3) but letting each value be with a different representative now, we can re-calculate the GNE and its value to visualize the effect of the subjective scale on the outcome. Table 4 shows the particular choices along with the scales used (including the scale from section IV-A as the first) and the parameters being set to the particular value. Fig. 5 plots the results. As with the different parameterization, the shape of the curve is retained, although the numerical values are different. Thus, the overall "trend" and risk picture across the infrastructure remains somewhat robust against the subjective choices in the parameters.

V. CONCLUSION

Despite the heterogeneity of concrete advanced persistent threats, they all seem to follow a similar pattern that can be divided into phases (temporal) and stages (spatial) domains. Each of these has its individual characteristics and needs different forms of protection. Game theory offers a convenient way of capturing the specific attack vectors being exploited throughout an APT, while letting us elegantly connect the individual models into an overall APT model that accurately models the heterogeneous environment in which an APT is mounted. In designing optimal defenses per phase and stage, we end up with an overall optimal APT defense, whose practical appeal is manifold: first, it can be defined in risk quantification terms that can be tailored to the specific application at hand, and directly relates to the set of possible counteractions (through the game's defense strategies). Second, the defense does involve a series of actors on both sides, which an accurate model should reflect. Our model does so by allowing different players in each phase of the APT. Third, the model reproduces what we expect intuitively, but goes beyond this confirmation in exhibiting interesting phenomena like the convergence of the risk along the Phase 2 game, or showing a general risk pattern across the infrastructure as has been visualized in section IV-D.

Practically, it is easy to compute optimal defenses using numerical techniques like fictitious play. The appeal of the latter is its natural resemblance to "learning" to play optimally as times goes by. This also mimics human behavior to some extent, and as our experiments showed, the learning can rapidly lead to an almost optimal defense behavior. Although APTs may themselves be too rare to take an algorithm like fictitious play to convergence (this algorithm is a mere computational vehicle anyway, and practically, the equilibrium should not be learned from experience only), the effect of adapting one's defense is positive in the reverse direction: since changes of configurations may occur more frequently than APTs, this in turn induces an element of uncertainty for the attacker from which security can be gained. Nonetheless, repeated defense actions such as awareness training for employees can be adapted on information collected from APT incidents having occurred elsewhere, so the learning from the past is, here, to be understood more broadly than in the narrow sense of just recording one's own and the direct opponent's past actions.

The models presented here are generic and open to generalizations in various ways, such as improved accounts for uncertainty (say, lifting the games to abstract spaces of probability distributions [39]), changing the models from matrix games into more general competitions (like signalling games or Stackelberg games), or simply by inducing further parameters to the models to gain flexibility. Our work shall thus stipulate further steps towards mathematically optimized designs that are resilient against advanced persistent threats. Since attackers do optimize their strategies for the victim (and have done so in the past, which culminated in the contemporary APT issue), adopting the same approach on the defender's side seems more than natural.

REFERENCES

[1] J. P. Farwell and R. Rohozinski, "Stuxnet and the future of cyber war," *Survival*, vol. 53, no. 1, pp. 23–40, 2011.

[2] B. Bencsáth, G. Pék, L. Buttyán, and M. Félegyházi, "The Cousins of Stuxnet: Duqu, flame, and gauss," *Future Internet*, vol. 4, no. 4, pp. 971–1003, 2012.

[3] K. Munro, "Deconstructing Flame: The limitations of traditional defences," *Comput. Fraud Secur.*, vol. 2012, no. 10, pp. 8–11, 2012.

[4] G. Kurtz. (2010). *Operation 'Aurora,' Hit Google & Others*. [Online]. Available: http://siblog.mcafee.com/cto/operation-%E2

[5] N. Falliere, L. O. Murchu, and E. Chien, "W32. stuxnet dossier," Symantec Corp., Mountain View, CA, USA, White Paper Version 1.4, Feb. 2011.

[6] C. Tankard, "Advanced persistent threats and how to monitor and deter them," *Netw. Secur.*, vol. 2011, no. 8, pp. 16–19, 2011.

[7] S. Jajodia, A. K. Ghosh, V. Swarup, C. Wang, and X. S. Wang, Eds., *Moving Target Defense—Creating Asymmetric Uncertainty for Cyber Threats* (Advances in Information Security), vol. 54. New York, NY, USA: Springer, 2011.

[8] S. Jajodia, A. K. Ghosh, V. S. Subrahmanian, V. Swarup, C. Wang, and X. S. Wang, Eds., *Moving Target Defense II—Application of Game Theory and Adversarial Modeling*, (Advances in Information Security), vol. 100. New York, NY, USA: Springer, 2013.

[9] R. Uzal, N. C. Debnath, D. Riesco, and G. Montejano, "Trust in cyberspace: New information security paradigm," in *Managing Trust in Cyberspace*. Paris, France: Atlantis Press, 2013.

[10] J. Pawlick, S. Farhang, and Q. Zhu, "Flip the cloud: Cyber-physical signaling games in the presence of advanced persistent threats," in *Decision and Game Theory for Security*. Cham, Switzerland: Springer, 2015, pp. 289–308.

[11] B. Satchidanandan and P. R. Kumar. (Jun. 2016). "Dynamic watermarking: Active defense of networked cyber-physical systems." [Online]. Available: https://arxiv.org/abs/1606.08741

[12] S. Sridhar, A. Hahn, and M. Govindarasu, "Cyber–physical system security for the electric power grid," *Proc. IEEE*, vol. 100, no. 1, pp. 210–224, Jan. 2012.

[13] Q. Zhu and T. Başar, "Robust and resilient control design for cyber-physical systems with an application to power systems," in *Proc. 50th IEEE Conf. Decision Control Eur. Control Conf.*, Dec. 2011, pp. 4066–4071.

[14] Q. Zhu, L. Bushnell, and T. Başar, "Resilient distributed control of multi-agent cyber-physical systems," in *Control of Cyber-Physical Systems*. Heidelberg, Germany: Springer, 2013, pp. 301–316.

[15] Q. Zhu and T. Basar, "Game-theoretic methods for robustness, security, and resilience of cyberphysical control systems: Games-in-games principle for optimal cross-layer resilient control systems," *IEEE Control Syst.*, vol. 35, no. 1, pp. 46–65, Feb. 2015.

[16] Z. Xu and Q. Zhu, "Secure and resilient control design for cloud enabled networked control systems," in *Proc. 1st ACM Workshop Cyber-Phys. Syst.-Secur. PrivaCy*, 2015, pp. 31–42.

[17] M. Manshaei, Q. Zhu, T. Alpcan, T. Başar, and J.-P. Hubaux, "Game theory meets network security and privacy," *ACM Comput. Survey*, vol. 45, no. 3, Jun. 2013.

[18] M. van Dijk, A. Juels, A. Oprea, and R. L. Rivest, "FlipIt: The game of 'stealthy takeover,'" *J. Cryptol.*, vol. 26, no. 4, pp. 655–713, 2013.

[19] Q. Zhu and T. Başar, "Feedback-driven multi-stage moving target defense," in *Proc. Conf. Decision Game Theory Secur. (GameSec)*, 2013, pp. 246–263.

[20] K. Zetter. (2016). *Inside the Cunning, Unprecedented Hack of Ukraine'S Power Grid*. [Online]. Available: https://www.wired.com/2016/03/inside-cunning-unprecedented-hack-ukraines-power-grid/

[21] S. Gosk, T. Winter, and T. Connor. (2015). *Iranian Hackers Claim Cyber Attack on New York Dam*. [Online]. Available: http://www.nbcnews.com/news/us-news/iranian-hackers-claim-cyber-attack-new-york-dam-n484611

[22] C. Francescani. (2016). *U.S. Infrastructure Can be Hacked With Google, Simple Passwords*. [Online]. Available: http://www.nbcnews.com/news/us-news/u-s-infrastructure-can-be-hacked-google-simple-passwords-n548661

[23] D. Kushner. (2013). *The Real Story of Stuxnet*. Accessed: Apr. 11, 2016. [Online]. Available: http://spectrum.ieee.org/telecom/security/the-real-story-of-stuxnet/

[24] D. SecureWorks. (2014). *Advanced Threat Protection With Dell Secureworks*. Accessed: Jul. 27, 2016. [Online]. Available: https://www.secureworks.com/resources/sb-advanced-threat-protection-with-dell-secureworks

[25] M. J. Assante and R. M. Lee, "The industrial control system cyber kill chain," SANS Inst., North Bethesda, MD, USA, Tech. Rep. 36297, Oct. 2015, accessed: Oct. 21, 2016. [Online]. Available: https://www:sans:org/reading-room/whitepapers/ICS/industrial-control-system-cyber-kill-chain-36297

[26] M. Daly. (2009). *The Advanced Persistent Threat (or Informationized Force Operations)*. Accessed: Jul. 27, 2016. [Online]. Available: https://www.usenix.org/legacy/event/lisa09/tech/slides/daly.pdf

[27] Kaspersky Lab's Global Research&Analysis Team. (2013). *Red October, Diplomatic Cyber Attacks Investigation*. Accessed: Sep. 30, 2016. [Online]. Available: https://securelist:com/analysis/publications/36740/red-october-diplomatic-cyberattacks-investigation/

[28] MSS Global Threat Response. (2014). *Emerging Threat: Dragonfly/Energetic Bear—APT Group 28*. Accessed: Oct. 17, 2016. [Online]. Available: https://www:symantec:com/connect/blogs/emergingthreat-dragonfly-energetic-bear-apt-group

[29] K. D. Mitnick, *The Art of Intrusion: The Real Stories Behind the Exploits of Hackers, Intruders and Deceivers*. Hoboken, NJ, USA: Wiley, 2005.

[30] M. Zhao, B. An, and C. Kiekintveld, "Optimizing personalized email filtering thresholds to mitigate sequential spear phishing attacks," in *Proc. 20th AAAI Conf. Artif. Intell.*, 2016, pp. 658–665.

[31] A. Clark, Q. Zhu, R. Poovendran, and T. Başar, "An impact-aware defense against stuxnet," in *Proc. IEEE Amer. Control Conf.*, Jun. 2013, pp. 4140–4147.

[32] J. Robinson, "An iterative method of solving a game," *Ann. Math.*, vol. 54, no. 2, pp. 296–301, 1951.

[33] D. Fudenberg and D. K. Levine, *The Theory of Learning in Games*, 1st ed. Cambridge, MA, USA: MIT Press, May 1998.

[34] Q. Zhu, H. Tembine, and T. Başar, "Hybrid learning in stochastic games and its application in network security," in *Reinforcement Learning and Approximate Dynamic Programming for Feedback Control* (Computational Intelligence Series), F. L. Lewis and D. Liu, Eds. Piscataway, NJ, USA: IEEE Press, 2012, ch. 14, pp. 305–329.

[35] Swedish Civil Contingencies Agency (MSB). (2012). *Swedish National Risk Assessment*. Accessed: Oct. 17, 2016. [Online]. Available: https://www:msb:se/RibData/Filer/pdf/26621

[36] *National Risk Assessment 2011*, Netw. Analysts Nat. Secur., Nat. Inst. Public Health Environ., Bilthoven, The Netherlands, 2011.

[37] *Methode zur Risikoanalyse von Katastrophen und Notlagen für die Schweiz*, Bundesamt für Bevölkerungsschutz, Bern, Switzerland, 2013.

[38] D. Bundestag, "Unterrichtung durch die Bundesregierung: Bericht über die Methode zur Risikoanalyse im Bevölkerungsschutz 2010," in *Verhandlungen des Deutschen Bundestages: Drucksachen*. Berlin, Germany: Deutscher Bundestag, Mar. 2018, p. 4178. [Online]. Available: http://dipbt.bundestag.de/doc/btd/17/041/1704178.pdf

[39] S. Rass, S. König, and S. Schauer, "Uncertainty in games: Using probability-distributions as payoffs," in *Decision and Game Theory for Security* (Lecture Notes in Computer Science), vol. 9406. Cham, Switzerland: Springer, 2015.

QUANYAN ZHU received the bachelor's degree (Hons.) in electrical engineering from McGill University in 2006, and the master's degree from the University of Toronto in 2008. He is currently pursuing the Ph.D. degree with the Department of Electrical and Computer Engineering and also with the Coordinated Science Laboratory, University of Illinois at Urbana–Champaign. His research interests are complex large-scale dynamical systems, cyber-physical systems, security and resilience issues in communication networks, biological systems, smart grids, and control systems. He was a recipient of NSERC Canada Graduate Scholarship, the University of Toronto Fellowship, Ernest A. Reid Fellowship, and Mavis Future Faculty Fellowships. He was the TPC Chair (Smart Grid Track) of the 2012 First INFOCOM workshop on communications and control on sustainable energy systems and the General Chair of the 2016 Conference on Decision and Game Theory for Security.

STEFAN RASS received the double master's degree in mathematics and computer science from the University of Klagenfurt in 2005, the Ph.D. degree in mathematics in 2009, and the Habilitation degree on applied computer science and system security in 2014. He is currently an Associate Professor with the University of Klagenfurt, teaching courses on theoretical computer science, complexity theory, security, and cryptography. His research interests include applied system security, and complexity theory, statistics, decision theory, and game-theory. He has authored numerous papers related to security and applied statistics and decision theory in security. Closely related to the project he has co-authored the book *Cryptography for Security and Privacy in Cloud Computing*, (Artech House). He participated in various nationally and internationally funded research projects.