A Game Theoretic Analysis of Cyber Threats

Paul Tavolato[1], Robert Luh[1,2] and Sebastian Eresheim[1,2]
[1]Research Group Security and Privacy, University of Vienna, Kolingasse 14-16, A-1090 Vienna, Austria
[2]Department of Computer Science, UAS St. Pölten, A-3100 St. Pölten, Austria

Keywords: Cyber Threat Analysis, Game Theory.

Abstract: Cyber threat analysis is crucial to securing modern IT systems. In the ongoing project described here a strictly mathematical method for threat analysis is sketched. The threat landscape between an attacker (hacker) and a defender (system owner) is modeled along the formalisms of stochastic game theory, thus opening the way for a rigorous formal analysis. The key benefit of the project is its applicability to real-world situations. Therefore, the information about possible attack and defense actions is taken from several proven data sources resulting in a large number of actions (173 attack actions and 115 defense actions). We present an adaptation of the so-called Princess-and-Monster game to model the problem. Various problems with the formalization are discussed. To keep the model manageable despite the claim of practicality, it is applied only to specific scenarios mimicking real-world situations.

1 INTRODUCTION

Threat analysis, as part of cyber risk analysis, is a key factor in today's IT management. Risk analysis studies the probability of potential risks and what impact they would have if, in fact, they were to occur. Such analysis is an indispensable requirement for the planning of preventive and counteractive measures. As these measures are generally financially relevant, a quantitative approach to risk analysis is preferable. This implies that threat analysis, too, must rest on a strong quantitative foundation.

This paper focuses on the process of assessing the capabilities and activities of unknown intelligence entities or criminals aimed at an organization's IT system. In other words, we consider malicious acts that seek to disrupt proper operation of an IT system by violating one or more of the central cybersecurity properties: the CIA triad – confidentiality, integrity and accessibility as defined in (Keyser, 2018).

In the literature some methods for threat analysis have been proposed, including STRIDE and Pasta (see for example (Shostack, 2014), (Tarandach & Coles, 2020), or (Swiderski & Snyder, 2004)). These methods, albeit useful in practice, have an essential drawback: they are mainly of an informal or semi-formal nature and thus lack a strong mathematical background. We suggest using game theory as the mathematical framework for a rigorous treatment of threat analysis.

By formally defining the situation of an attacker and a defender as a game, one can analyze essential properties of the situation, giving valuable input regarding the costing and planning of a defense strategy together with the necessary measures to enforce this strategy with the aim to detect, mitigate or – at best – prevent attacks on the system.

The main objectives of the project are to develop a rigorous mathematical framework for analyzing real-world threat scenarios. The mathematical framework is based on game theory. A high degree of practical relevance is achieved by mimicking real-life situations in cyber security by taking information about attack and defense activities from several proven data sources. Another innovative element of the project is the incorporation of stochastic features in the game theoretic model: Each action by one of the players is assigned a success probability. Only if the action succeeds, the effects associated with that action are realized.

The current state of the project comprises a game theoretic model representing the threat landscape of a defender and an attacker as a game. The game is based on a so-called Princess-and-Monster game, which is adapted to fit cyber threat situations. Furthermore, a complete list of attack actions is compiled containing for each action a description, the

Tavolato, P., Luh, R. and Eresheim, S.
A Game Theoretic Analysis of Cyber Threats.
DOI: 10.5220/0011792700003405
In *Proceedings of the 9th International Conference on Information Systems Security and Privacy (ICISSP 2023)*, pages 706-713
ISBN: 978-989-758-624-8; ISSN: 2184-4356

skill requirements of the action, a success probability, and the damage effect in case of a successful execution of the action. A complete list of defense actions is compiled, too, containing the skill requirements of the action, a success probability, and eventually a mitigating ("healing") effect of the action.

2 RELATED WORK

Some formal aspects of threat and/or attack modeling are described in (Guelzim & Obaidat, 2015). A method widely used in threat analysis is based on attack-defense trees (ADT). Lately, semi-formal and formal treatments of ADTs were developed; see Wideł et al. (Wideł, Audinot, Fila, & Pinchinat, 2019) for an overview of the use of formal models in security.

In particular, the paper (Aslanyan, Nielson, & Parker, 2016) describes a formal system for analyzing quantified properties of attacks while in a paper by Gadyatskaya et al. (Gadyatskaya, Hansen, Larsen, Legay, Olesen, & Poulsen, 2016) Priced Timed Automata are used to formalize attack trees. However, defense actions are not included, and no stochastic aspects were used in this approach.

A stochastic analysis of attack trees is discussed in (Pekergin, Tan, & Fourneau, 2016). In (Buldas, Gadyatskaya, Lenin, Mauw, & Trujillo-Rasua, 2020) constraint programming is used to evaluate attack trees with incomplete information.

The first paper suggesting stochastic game theory (Shapley, 1953) in connection with cyber security was published in 2002 (Hamilton, Miller, Ott, & Saydjari, 2002). Kordy et al. (Kordy, Mauw, Melissen, & Schweitzer, 2010) formally proved the equivalence of attack-defense trees and game theory. There are several papers about the application of game theory to various aspects of security: a good overview, though restricted to cyber-physical systems, is given in (Etesami & Basar, 2019). Other game-theoretic approaches are discussed in (Bommannavar, Alpcan, & Bambos, 2011), (He, Zhuang, & Rao, 2012), (Luo, Szidarovszky, Al-Nashif, & Hariri, 2010), (Nguyen, Alpcan, & Başar, 2009), (Sallhammar, Helvik, & Knapskog, 2006), (Sajjan, Sankardas, & Dipankar, 2010), (Tabatabaei, 2016), (Zhang, Wang, & Zhuang, 2021), (Li, Peng, Zhu, & Basar, 2021), (Jie, Choo, Li, Chen, & Guo, 2019), and (Han, Niyato, Saad, & Başar, 2019).All the papers give some small examples to illustrate the viability of the method, but do not cover threat analysis for real-world situations.

The paper closest to our work is by Attiah et al. (Attiah, Chatterjee, & Zou, 2018); they study achievable mixed strategy Nash equilibria, but do not incorporate success probabilities.

3 GAME THEORY AND THREAT ANALYSIS

3.1 The Game

As a mathematical approach to rigorously model threat situations, game theory for non-cooperative games suggests itself. In cyber threat scenarios there are two opponents playing (i.e., fighting) against each other. Each opponent has an arsenal of actions, though there are different actions available to the two opponents (attack and defense actions, respectively).

The structure of the game discussed here is adopted from PenQuest (Luh, Eresheim, Großbacher, Petelin, Mayr, Tavolato & Schrittwieser, 2022), a digital cyber security role-playing game, where an attacker attempts to compromise an IT infrastructure and the defender tries to prevent or mitigate the threat. The attacker has a predefined goal (violating one part of the CIA triad), and the defender has a given infrastructure he wants to defend against attacks. The distinct strength of this game, which distinguishes it from others, is its closeness to real world situation: to render it as realistic as possible the attack and defense actions that are part of the game are taken from several proven data sources such as STIX (Structured Threat Information eXpression language) (MITRE Corporation, D), the APT kill chain by Hutchinson (Hutchins, Cloppert, & Amin, 2011), the CAPEC (Common Attack Pattern Enumeration and Classification) attack patterns (MITRE Corporation, A), the MITRE ATT&CK (Adversarial Tactics, Techniques & Common Knowledge) attack and mitigation patterns (MITRE Corporation, B), the NIST SP 800-53 Countermeasures (Joint Task Force Transformation Initiative, 2015), and MITRE D3FEND (MITRE Corporation, C). In the game mentioned above there are 173 different attack actions and 115 different defense actions defined.

Attack actions are subdivided into three categories along the attack stages: Reconnaissance, Initial Access, and Execution. A stage can only be entered if the preceding stage has been completed successfully. Each action requires a minimum skill level of the attacker (1, 2, 3, 4or 5) and has a success probability. Effects of a successful attack action may be twofold:

a) the damage achieved within the defender's system in the three dimensions Confidentiality, Integrity, Availability. Damage is measured with a number between 0 and 3, where 3 designates maximum damage and means that the attacker has reached his/her goal; damage accumulates throughout the game; and
b) a gain of insight into the defender's configuration, enabling a more specific choice of subsequent actions and thus resulting an increase in the success probabilities of those actions.

Table 1 shows some examples of attack actions listing the action name, the attack stage, the necessary skill requirement of the attacker, the success probability, and the damage caused by the successful action.

Table 1: 7 Examples out of the 173 attack actions

Action	Stage	Skill req.	Success prob.	C/I/A damage
Vulnerability Scan	Recon.	2	0,6	0/0/0
Brute Force	Access	1	0,4	0/0/1
Spearphishing	Access	3	0,5	1/1/0
Manipulate Website	Execution	2	0,3	0/3/2
Buffer Overflow	Execution	3	0,5	1/2/3
Wipe Disk	Execution	3	0,3	0/0/3
Man in the Middle	Execution	2	0,4	1/1/1

Defense actions are subdivided into three categories as well: Prevention, Detection, and Response. Response actions are only effective when an attack has already been detected and therefore are available only in such a situation. Each action requires a minimum skill level of the defender (1, 2, 3, 4 or 5) and has a success probability. Effects of a successful defense action may be twofold:

a) the "healing" effect triggered by the successful action in the dimensions Confidentiality, Integrity, and Availability, thus undoing some or all of the damage already inflicted on the defender's system, which means that the damage number is decreased again. Such healing can only happen as a result of a successful response action.
b) a decrease of the success probability of certain attacker actions (e.g., awareness training reduces the success probability of the attacker action "phishing") and/or an increase in the success probabilities of other defense actions.

Table 2 shows some examples of defense actions including the action name, the action category, the necessary skill requirement of the defender, the success probability of the action, and the healing effects with respect to Confidentiality, Integrity and Availability triggered by the (response) action.

Table 2: 7 Examples out of the 115 defense actions.

Action	Action type	Skill req.	Success prob.	C/I/A healing
2-Factor-Authentication	Prev.	2	0,5	0/0/0
Encrypt Transmission	Prev.	2	0,7	0/0/0
Check Driver Integrity	Prev.	2	0,5	0/0/0
Analyze Network Protocol	Det.	3	0,5	0/0/0
Analyze File Access	Det.	2	0,3	0/0/0
Terminate Connection	Resp.	2	0,7	-1/-1/-1
Restore Configuration	Resp.	2	0,6	0/-1/-1

As game theory is applicable to a wide variety of quite different situations, one must specify the characteristics of the game under consideration:

- It is a **non-cooperative** game (quite obvious).
- It is a **two-player game** between an attacker and a defender.
- It is an **asymmetric** game. The two players have different actions to choose from.
- It is a game of **perfect recall**. Each player never forgets what s/he has done so far.
- Whether it is a **zero-sum** game or not depends on the details of the utility function. However, there is always a winner and a loser, there are no ties: either the attacker achieves her/his goal and wins, or the defender wins, if the attacker gives up without reaching her/his goal (time-out).
- It is a **non-deterministic** game: each player can in most situations choose between a number of possible actions non-deterministically, restricted only by some predefined constraints.
- It is a **stochastic** game: the success of a player's action is defined by a probability distribution.
- It is a game with **imperfect information**: the current state of the game and its history is mostly unknown to the players. In the beginning of the game, for example, the defender does not even know that the game has already started, and an attack is well under-way; and the attacker has no knowledge of the defender's configuration. In addition, neither player knows which actions are available to the opponent. Moreover, the defender does not know which actions the attacker has already

carried out successfully and the attacker does not know what defensive actions are already in place. It is only during the course of the game that parts of this information may gradually become unveiled. For example, the attacker may gain insight into the defender's configuration by conducting successful reconnaissance actions and the defender may somewhen find out that he or she is under attack.

- Whether the game is one with **incomplete information** is a more complex question. Usually, in game theory "incomplete information" means that the players don't know each other's utility or pay-off function. In principle this information is clear on both sides: the defender knows that an attacker wants to violate at least one of the C-I-A dimensions and the attacker knows that the defender wants to prevent just that. In other words: the pay-off function is comprised of the numbers that represent the damage in the respective three dimensions. The defender wants these numbers to stay at zero and the attacker wants to increase at least one of these numbers to three. But there is some hidden information in the pay-off functions, too: the opponents do not know each other's skill level, the defender does not know about the attackers initiative - that is the time (and budget) limit of the attacker which determines when they will give up and stop attacking.
- We restrict the game to being turn-based: only one player can make a move at a time. Moreover, the moves are executed alternately. This property makes the game an **extensive form** game: the sequence of moves can be structured in a tree-like manner.

3.2 Game Theoretic Model

The game consists of the following information:

- The players: attacker and defender.
- The actions available to the respective players.
- A success probability for each action.
- A pay-off function that assigns rewards or deals/heals damage points to the players.

Given this information we can view attacker-defender scenarios as stochastic 2-player games:

$\Gamma_{PQ} = (\Pi, A, sp, S, U)$

$\Pi = \{att, def\}$	the players: attacker and defender
$A = A_{att} \cup A_{def}$	a finite set of actions of the players (different actions for attacker and defender)
$sp: A \rightarrow [0,1]$	the (dynamic) success probability of an action
$S = S_{att} \cup S_{def}$	the finite set of strategies for the attacker respectively the defender
$U: S_{att} \times S_{def} \rightarrow \mathbb{R}^2$	the utility function for the players that assign pay-offs to strategy combinations

When going through the various types of games introduced in the literature, we chose some aspects of the so-called "Princess and Monster Game" (Isaacs, 1965) and adapted it to the cyber threat domain. In the original Princess-and-Monster game, two players, the Princess and the Monster, are in a completely dark room with defined boundaries. They don't see each other and move around in the dark. Usually, the moving speed of the monster is defined as 1 and the moving speed of the princess is $\omega < 1$. The Monster wins when it catches the princess in time (before an initially defined time limit), that is when both are at the same spot (or close enough to each other) in the room. The utility function is the time it takes the Monster to catch the Princess. The Monster tries to minimize the catching time, the Princess tries to maximize it (to survive as long as possible). There are many variations of the game depending on the size of the room, its metric, and the possible moves of the players in it.

In our cyber-threat version of the game the Princess is the defender (together with the system to defend). The Monster – the attacker – is not in the room but can remotely control the Princess's moves. The room is a three-dimensional cube with a discrete net of size 4x4x4. The position of the defender is hence a point (x,y,z) with three integer coordinates, each in the range of 0 to 3. The coordinates correspond to the three damage categories Confidentiality, Integrity, Availability (see Figure 1). The defender starts at point (0,0,0) and the game ends when s/he reaches the target border of the room. In other words, the attacker is victorious when the target category defined at the beginning of the game reaches the value 3. The main effect of a successful control action of the attacker is to move the defender closer to the target border (to increase the respective value), while the defender can move farther away from it when s/he successfully executes an appropriate action with a healing effect. The utility function is the time as in the original game, thus making it a zero-sum game. But it has a predefined limit: if it takes the

attacker too long to move the defender to the border, he gives up: the game is over and the defender wins.

Besides the increase or decrease of the coordinates the respective moves may have additional effects such as dynamically changing success probabilities (some defensive actions can reduce the success probability of certain attack actions). In some cases, the success probability could even be reduced to 0.

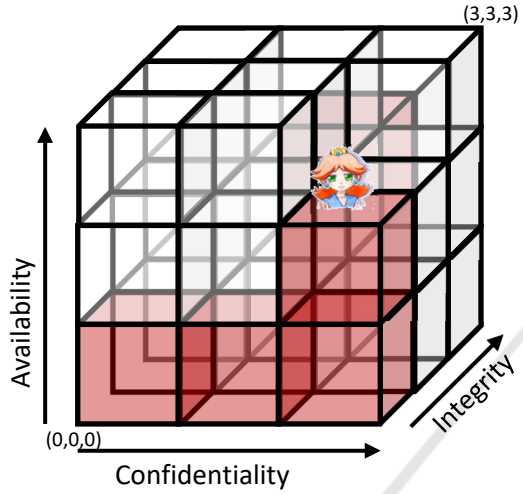

Figure 1: The princess is in position (2,1,2) implying that she is already in great danger in dimensions Confidentiality and Availability.

So far, we have made an additional restriction: we assume the game to be turn-based. The players execute their actions alternatingly one after the other. This is necessary to precisely localize the effects of actions in time, as these effects might have consequences on actions executed thereafter.

3.3 Strategies

A prerequisite for game theoretic analysis is the definition of strategies and strategy combinations. A strategy is a sequence of actions chosen by a player; a strategy combination is the combination of the respective strategies of the two players. A pure strategy defines a fixed course of actions that a player will stick to. In many cases, especially in games with imperfect and/or incomplete information, it is better for a player not to use the same strategy all the time, but to mix strategies according to some probability distribution. A mixed strategy of a player p is a probability distribution q_P over the set of his/her pure strategies:

$$q_P : S_P \rightarrow [0,1] \quad \text{where } P \in \Pi = \{att, def\}$$

All the probabilities of the strategies of one player sum up to 1. In a game of perfect recall mixed strategies are equivalent to behavioral strategies which define a probability distribution over actions at each decision point of a player. We will use behavioral strategies as they are easier to understand in the context of this game. Furthermore, we assume that the number of strategies is finite. To guarantee this we do not allow a player to repeat an action within a strategy.

Nevertheless, the number of possible strategies is huge, taking into account the number of 173 attack actions and 115 defense actions. This huge number of strategies makes a straightforward game theoretic analysis leading to a Nash equilibrium in mixed strategies (which according to Nash's famous theorem always exists in a finite game) practically infeasible. To reduce the number of strategies the following aspects could be used:

- The actions are assigned a skill requirement (1, 2, 3 or 4) necessary to play this action. Given a predefined type of attacker and defender some actions can be excluded for the scenario under consideration.
- Strategies for an attacker must consist of three consecutive stages: reconnaissance, initial access, and execution. All actions are assigned to one of these stages. Entering the next stage requires the successful completion of the previous stage.
- Defense actions, too, have a type: prevention, detection, response. The use of certain actions is restricted by the progress of the game so far.

One problem to solve is the combination of the probability distribution over the actions at a certain point in the game (as defined in the behavioral strategy) and the success probabilities of these actions. An action is chosen due to the probability distribution; but to realize the effects of this action it must be successful, too. If the action chosen is not successful, it has no effects whatsoever: the pay-off does not change, and other effects are not realized either (the only exception is the attacker's initiative which decreases with unsuccessful attacker actions as well). To accomplish this, we multiply the probabilities governing the choice of actions with the success probabilities of the respective actions and then normalize the results for the interval [0,1] such that the probabilities sum up to 1 (at least one action must be chosen). This gives the adapted probability distribution for a successful behavioral strategy.

Even with the simplifications mentioned above, a game theoretic analysis is only possible for clearly

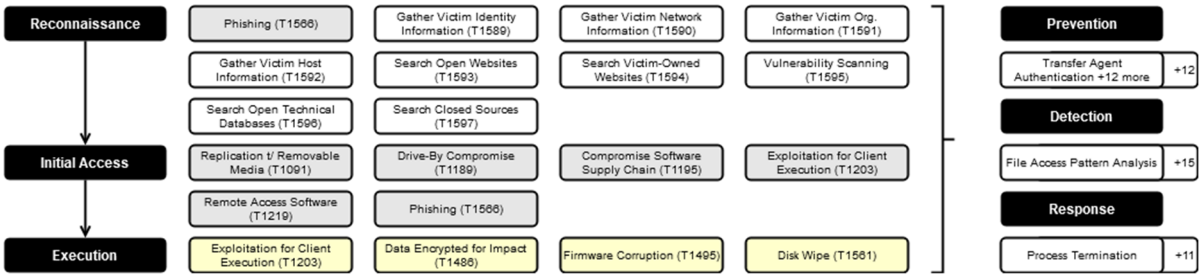

Figure 2: Set of attack actions available in a simplified ransomware scenario. Actions highlighted in grey deal only minor damage (+1), while yellow actions (Execution stage) have a high damage potential in the game (+2). Defense actions are exemplified on the right.

defined scenarios. Such a scenario defines the following characteristics of a game and thereby reduces the number of possible actions and strategies of the players:

- The skill level of the attacker
- The skill level of the defender
- The attacker's goal: which part of the CIA triad s/he wants to attack
- The configuration of the defender

3.4 Example Scenario

Let's have a short look at one such scenario, a ransomware attack on a conventional IT network configuration: a hacker with skill level 2 wants to infiltrate the defender's system in order to encrypt substantial data and make the system inoperable. Hence, the goal of the attack is the availability of the system. Within the game, this means that the attacker wants to drive the defender towards the third border of the game room – s/he wants to increase the third dimension of the defender's position to a value of 3. In our setting, the defender has the same skill level of 2. Figure 2 shows the available actions independent of skill level. The strategies available to the attacker in our example are composed of a subset of these attack actions and are structured in three stages:

Table 3 shows the possible attack actions in the reconnaissance stage, named after the corresponding MITRE ATT&CK technique: "Phishing", "Gather various victim Info", "Search various info", "Vulnerability Scanning". Table 4 shows the possible attack actions in the initial access stage: "Compromise Software Supply Chain" (i.e. provide a malicious update) and "Remote Access Software" (i.e. asking the user for a remote desktop connection); others, like "Use removable media", "Drive by compromise" or "Exploitation for client execution" require a higher skill level than 2. Table 5 shows the possible attack actions in the execution stage: "Data Encrypted for Impact" and "Disk Wipe". Again, the action "Firmware corruption" is not mentioned as it requires a too high skill level.

For the defender side there are more options available. In the prevention stage, out of more than 40 possible actions we can remove 16, since they require a higher skill level than 2. Some of the remaining actions would not prevent the ransomware attack, so they could be removed from the scenario, too, leaving 13 actions in total. In the detection stage we have a similar picture: out of 48 actions we can remove 15 because of a too high skill bar; and some more actions could be removed in view of the details of the configuration. In the response stage, there are 20 actions that fulfill the skill requirement.

Table 3: Reconnaissance actions in the example.

Action	Stage	Skill req.	Success prob.	C/I/A damage
Phishing	Recon.	2	0,5	1/1/0
Gather victim id info	Recon.	1	0,7	0/0/0
Gather victim network info	Recon.	1	0,5	0/0/0
Gather victim org. info	Recon.	1	0,7	0/0/0
Gather victim host info	Recon.	1	0,7	0/0/0
Search open websites	Recon.	1	0,7	0/0/0
Search victim websites	Recon.	1	0,7	0/0/0
Search open technical DB	Recon.	1	0,7	0/0/0
Search closed sources	Recon.	1	0,5	0/0/0
Vulnerability scanning	Recon.	2	0,6	0/0/0

Table 4: Initial access actions in the example.

Action	Stage	Skill req.	Success prob.	C/I/A damage
Compromise supply chain	Access	2	0,5	1/1/1
Remote access software	Access	2	0,1	0/1/0

Table 5: Execution actions in the example.

Action	Stage	Skill req.	Success prob.	C/I/A damage
Data encryption for impact	Exec.	2	0,5	0/1/2
Disk wipe	Exec.	2	0,3	0/0/3

Only a path of attack actions, which results in a damage value of 3 will render the attack successful. However, the defender has the opportunity to mitigate a current damage level by taking appropriate defensive actions, which could lower the damage level if successfully executed.

The analysis of the scenario should then yield information on optimal strategies (for the defender). Due to the success probabilities attached to the actions there will not be a single best strategy, but rather a set of probabilities denoting the success probability of a specific attack and how this probability would change due to specific defense actions.

4 CHALLENGES AND FUTURE WORK

The formalism for modeling attack-defense scenarios as described in this paper makes threat analysis possible in a strictly mathematical way. The main innovation here is that concepts of mathematical game theory are applied to tactical threat analysis for real-world situations. To keep the analysis as close to practice as possible we collected attack and defense actions from accepted data sources and vocabularies. As mentioned earlier the model so far contains 173 different attack actions and 115 defense actions that can be used in modeling attack and defense strategies. Subsuming all these aspects into one single model would clearly lead to too large a model, and hence render a reasonable mathematical treatment impossible. To keep the model in a manageable size we break down the model into scenarios along the lines of kill-chains (Hutchins, Cloppert, & Amin, 2011).

The main challenges that remain are on one side the definitions of more elaborate scenarios considering detailed configurations of the system under attack and on the other side a detailed game theoretic analysis of these scenarios. Such analysis must be of stochastic nature as one of the main analysis goals is the investigation of the influence of specific defense actions on the success probabilities of various attack strategies. Due to the huge number of possible actions in a realistic scenario a traditional game theoretic analysis might be very difficult, if not infeasible. The application of formal methods, such as stochastic model checking, could potentially be considered as a viable alternative.

Future work will consider additional parameters in the game, especially budgets: if costs are attached to defense actions, what is the relation between an increase in the defender's budget and the overall success probability of an attack strategy? Another option for future work would be giving up the turn-based property of the game at least partially: attacker and defender could execute more than one action in a row before the other player makes a move, especially in the beginning of the game. So far, the consequences of this have not yet been analyzed.

The main goal of the project is strategy synthesis for defenders: Given a configuration with real-world circumstances such as a given system configuration, budgets, skill levels and the like, what is the optimal defense strategy for anticipated threats?

ACKNOWLEDGEMENT

This work is part of the project INODES funded by the Austrian Science Fund (FWF), whose support is gratefully acknowledged.

REFERENCES

Aslanyan, Z., Nielson, F., & Parker, D. (2016). Quantitative Verification and Synthesis of Attack-Defence Scenarios. IEEE 29th Computer Security Foundations Symposium, CSF 2016, (S. 105-119). doi:10.1109/CSF.2016.15

Attiah, A., Chatterjee, M., & Zou, C. (2018). A Game Theoretic Approach to Model Cyber Attack and Defense Strategies. IEEE International Conference on Communications (ICC).

Bommannavar, P., Alpcan, T., & Bambos, N. (2011). Security risk management via dynamic games with learning. IEEE International Conference on Communications (ICC). doi:10.1109/icc.2011.5963330

Buldas, A., Gadyatskaya, O., Lenin, A., Mauw, S., & Trujillo-Rasua, R. (2020). Attribute Evaluation on Attack Trees with Incomplete Information. Computers and Security 88/101630.

Etesami, S. R., & Basar, T. (2019). Dynamic Games in Cyber-Physical Security: An Overview. Dynamic Games and Applications, S. 884–913. doi:doi.org/10.1007/s13235-018-00291-y

Gadyatskaya, O., Hansen, R. R., Larsen, K. G., Legay, A., Olesen, M., & Poulsen, D. B. (2016). Modelling Attack-defense Trees Using Timed Automata. In M.

Fränzle, & N. Markey (Hrsg.), FORMATS 2016, LNCS 9884. 9884, S. 35–50. Springer LNCS. doi:DOI: 10.1007/978-3-319-44878-7 3

Guelzim, T., & Obaidat, M. S. (2015). Formal methods of attack modeling and detection. In Modeling and Simulation of Computer Networks and Systems - Methodologies and Applications (S. 841-860). Elsivier.

Hamilton, S., Miller, W., Ott, A., & Saydjari, O. (2002). The role of game theory in information warfare. 4th Information Survivability Workshop, (S. 1-4).

Han, Z., Niyato, D., Saad, W., & Başar, T. (2019). Game Theory for Next Generation Wireless and Communication Networks - Modeling, Analysis, and Design. Cambridge University Press.

He, F., Zhuang, J., & Rao, N. (2012). Game Theoretic Analysis of Attack and Defense in Cyber-Physical Network Infrastructures. Proceedings of the Industrial and Systems Engineering Research Conference. doi:10.1.1.719.2652

Hutchins, E., Cloppert, M., & Amin, R. (2011). Intelligence-driven computer network defense informed by analysis of adversary campaigns and intrusion kill chains. Lead. Issues Inf. Warf. Secur. Res. 1/80.

Isaacs, R. (1965). Differential Games. New York: John Wiley and Sons.

Jie, Y., Choo, K.-K. R., Li, M., Chen, L., & Guo, C. (2019). Tradeoff gain and loss optimization against man-in-the-middle attacks based on game theoretic model. Future Generation Computer Systems, Volume 101, S. 169-179. doi:doi.org/10.1016/j.future.2019.05.078

Joint Task Force Transformation Initiative. (2015). SP 800-53 rev. 4. Recommended Security Controls for Federal Information Systems and Organizations. Gaithersburg.

Keyser, T. (2018). Security policy. In The Information Governance Toolkit (S. 57-62). CRC Press.

Kordy, B., Mauw, S., Melissen, M., & Schweitzer, P. (2010). Attack–Defense Trees and Two-Player Binary Zero-Sum Extensive Form Games Are Equivalent. In T. Alpcan, L. Buttyán, & J. Baras (Hrsg.), Decision and Game Theory for Security. GameSec 2010. . Lecture Notes in Computer Science, vol 6442. Springer. doi:10.1007/978-3-642-17197-0_17

Li, T., Peng, G., Zhu, Q., & Basar, T. (2021). The Confluence of Networks, Games and Learning - A game-theoretic framework for multi-agent decision making over networks. IEEE control system magazine, special issue on Distributed Nash Equilibrium Seeking over Networks.

Luh, R., Eresheim S., Großbacher S., Petelin T., Mayr F., Tavolato P., & Schrittwieser S. (2022). PenQuest Reloaded: A Digital Cyber Defense Game for Technical Education. EDUCON2022 – IEEE Global Engineering Education Conference, 2022. doi:10.1109/EDUCON52537.2022.9766700

Luo, Y., Szidarovszky, F., Al-Nashif, Y., & Hariri, S. (2010). Game Theory Based Network Security. Journal of Information Security 1/1, S. 41-44. doi:10.4236/jis.2010.11005

MITRE Corporation. (A). CAPEC—Common Attack Pattern Enumeration and Classification. https://capec.mitre.org/

MITRE Corporation. (B). MITRE ATT&CK. https://attack.mitre.org/

MITRE Corporation. (C). MITRE D3FEND. https://d3fend.mitre.org/resources/D3FEND.pdf

MITRE Corporation. (D). STIX—Structured Threat Information Expression | STIX Project Documentation. https://oasis-open.github.io/cti-documentation/

Nguyen, K. C., Alpcan, T., & Başar, T. (2009). Stochastic games for security in networks with interdependent nodes. Proceedings of the 2009 International Conference on Game Theory for Networks, GameNets '09.

Pekergin, N., Tan, S., & Fourneau, J.-M. (2016). Quantitative Attack Tree Analysis: Stochastic Bounds and Numerical Analysis. International Workshop on Graphical Models for Security, GraMSec 2016. doi:DOI:10.1007/978-3-319-46263-9_8

Sajjan, S., Sankardas, R., & Dipankar, D. (2010). Game theory for cyber security. CSIIRW '10: Proceedings of the Sixth Annual Workshop on Cyber Security and Information Intelligence Research. doi:doi.org/10.1145/1852666.1852704

Sallhammar, K., Helvik, B. E., & Knapskog, S. J. (2006). On stochastic modeling for integrated security and dependability evaluation. Journal of Networks 1/5, S. 31-42. doi:10.4304/jnw.1.5.31-42

Shapley, L. S. (1953). Stochastic games. PNAS 39/10, S. 1095-1100.

Shostack, A. (2014). Threat Modeling: Designing for Security. Wiley.

Swiderski, F., & Snyder, W. (2004). Threat Modeling. Microsoft Press.

Tabatabaei, M. (2016). Games and Strategies in Analysis of Security Properties. Dissertation. Université du Luxembourg.

Tarandach, I., & Coles, M. J. (2020). Threat Modeling: A Practical Guide for Development Teams. O'Reilly.

Wideł, W., Audinot, M., Fila, B., & Pinchinat, S. (2019). Beyond 2014: Formal Methods for Attack tree-based Security Modeling. ACM Computing Surveys. 52/4, Article 75. doi:https://doi.org/10.1145/3331524

Zhang, J., Wang, Y., & Zhuang, J. (2021). Modeling multi-target defender-attacker games with quantal response attack strategies. Reliability Engineering & System Safety, Volume 205.