Using Simulations to Determine Economic Cost of a Cyber Attack

Nina Sokol, Viktoria Kežman, Stjepan Groš*
Faculty of Electrical Engineering and Computing
University of Zagreb, Zagreb, Croatia
{nina.sokol,viktoria.kezman,stjepan.gros}@fer.hr
*Corresponding author

***Abstract*—Cyber attacks pose an increasing threat to organizations, with their growing frequency leading to significant financial consequences. Accurately estimating the costs of cyber security incidents remains a challenge for many organizations. To address this, in this paper we introduce a novel methodology for calculating the cost of a cyber attack based on summing costs at the tactical level. We experimented with this approach through a series of simulations using a novel simulator. Our goal was to determine if this approach is feasible, while also identifying shortcomings of the simulator in its current state. In these simulations, we studied how the cost of an incident changes with different defense tactics. Simulation results provide insights into different defense strategies and their impact on total costs. Our proposed methodology, utilizing the simulator for both simulation and cost determination, offers a unique approach to determining the costs incurred by cyber incidents. We also provide a detailed discussion on the potential shortcomings of this approach and research questions that still need to be answered for final conclusions to be drawn.**

***Index Terms*—cost calculations, simulations, financial damage, cyber attack, methodology for cost calculation, sources of economic damage**

I. Introduction

Cyber attacks pose a significant challenge for organizations due to the uncertainty they bring. This underscores the importance of conducting risk assessments to determine the potential costs involved. By evaluating these risks, organizations can make informed decisions about investing in security measures, such as acquiring cyber insurance. Many organizations struggle to accurately estimate the actual cost of cyber security incidents, which often differs from their estimates. Organizations require a clear understanding of potential losses in a cyber incident to determine the optimal level of investment in security measures.

This information is also valuable for both cyber security consultancies and insurance companies, as they often struggle to obtain consistent data on incidents which is then used to assess risks and set insurance premiums. Banks can also benefit from this information when they are assessing risk and calculating credit scores to approve bank loans. It also enables better cost comparisons across different types of organizations and a deeper understanding of the types of costs more common in certain sectors. Ultimately, this can change organizations' attitudes towards cyber security, improve their understanding of the cost of breaches, and justify additional investment, training, or increased cyber insurance coverage. Moreover, it will demonstrate the importance of cyber security to employees, encouraging them to be more careful and follow good practices.

There was a lack of uniformity in how costs were collected and calculated in the existing literature on this topic, making it difficult to directly compare figures from different sources. The estimates vary significantly from study to study. Researchers employ various methods to estimate the costs of cyber security breaches.

Bottom-up approaches involve identifying each cost separately and then adding them up to get a total cost estimate. Generally, industry experts help identify these costs, and estimates are made through surveys conducted at the firm level [1]. Anderson et al. offer a framework for systematically measuring cyber crime costs, distinguishing direct losses, indirect losses, and protection costs while separating cyber crimes from supporting infrastructure. However, this framework has not been utilized to develop a tool or methodology for directly measuring the costs of cyber crime [2]. Event studies assess how a cyber security breach affects a firm's value using statistical analysis. They employ the market model to estimate stock returns without the breach, considering market trends and security responsiveness [3]. Case studies investigate the costs of cyber security breaches through company interviews and secondary research [4].

This paper presents experiments conducted to evaluate the financial impact on organizations resulting from a cyber attack. Our goal was to test the feasibility of using the simulator to calculate the cost of a cyber attack and to examine how the incident cost varies with different defense tactics. The *Cyber Conflict Simulator* (CCS) is used for simulating attack and defense scenarios and determining potential damage. What makes it unique and innovative is that the simulator allows modeling different types of organizations, including different attack and defense scenarios.

The simulator allows organizations to comprehensively understand their structure, evaluate risks, and enhance skills in countering cyber threats without putting their actual systems and data at risk. This simulator enables interactive simulations in which every action initiated by participants has consequences. Participants take on roles as attackers and defenders, using a simulator that allows them to interact and observe the outcomes of their actions. This tactical-level tool was specifically designed for incident response training and incorporates a simple damage calculation mechanism. This study explored the feasibility of employing this mechanism for damage assessments. The goal is to build a model of the organization, identify the cost factors contributing to the total damage, and then use the simulations to obtain results.

In the simulations, the sources of damage were varied during the first part of the experiments, while in the second part, we varied the defense strategies. For each of the simulations, the total cost of the attack was calculated using the proposed method.

The paper is organized as follows: After the introduction section, Section II provides a brief survey of related work. Following that, Section III outlines different cost categories and describes the simulator used in the experiment. Section IV presents the methodology employed in this study. Section V describes the experiments conducted and the results obtained through the simulations. Section VI is dedicated to discussing these results and addressing the encountered challenges. Finally, Section VII presents the conclusions drawn from the study and outlines potential future research directions.

II. Related Work

In the realm of cyber security, several studies have addressed the issue of estimating financial damage due to cyber attacks.

The costing methodology used in the study by Ponemon Institute and IBM [5] is activity-based costing (ABC). The study included 507 organizations that experienced breaches in the last year, with 3,211 interviews conducted to gather relevant information. To collect data, participants estimated costs using their knowledge by rating them on a number line during interviews. The methodology involved identifying and assigning costs to activities based on actual usage.

Heyburn [6] employed interviews with organizations to develop a cost estimation tool for assessing the full costs of cyber security breaches. The methodology involved a mapping exercise to identify cost categories, leading to the creation of a comprehensive questionnaire. Riek's [7] study employs surveys to assess the costs of cyber crime on individual victims, using a mathematical model to compare cost categories. Lis and Mendel's [8] research focuses on assessing the effectiveness of cyber security measures, particularly for critical infrastructure, taking into account costs, benefits, and the specific indicator to measure the return on security investment.

In summary, literature on specific attack and defense simulations and the calculation of the damage resulting from such attacks is relatively scarce, often industry-specific [8] and non-generalizable [7], or exclusively focused on theoretical frameworks and cost calculations [7].

Kuhl [9] introduced a simulator designed to model different computer network configurations and generate attacks automatically based on the network setup. It operates at a technical level and focuses on simulating network-based attacks, but it wasn't used to assess the actual damage caused.

In the simulator used in this paper, technical details are not as crucial. It is important to simulate attack and defense scenarios to see where the damage can occur, involving people and their actions. The simulations reflect how these actions impact the resulting damage, providing a more realistic approach to assessing damage.

III. Background

Our motivation lies in the need for a thorough listing of all costs that may arise as a result of a cyber attack so that we can determine which ones we can quantify and which ones we cannot. Costs have been identified from the literature we have examined. After assembling the list, we analyzed the costs by categorizing them based on their type and time frame. Additionally, we will describe the simulator used in the experiment and its features that were used to calculate costs.

A. Framework for analyzing the costs

In this section, we will describe the incident response life cycle and categorize different types of costs, relating them to different phases of the incident response life cycle. Furthermore, we'll discuss the time frames linked with these costs, such as short-term, medium-term, and long-term.

The NIST incident response life cycle breaks incident response down into four main phases: Preparation; Detection and Analysis; Containment, Eradication, and Recovery; and Post-Event Activity [10].

In this paper, regarding the type of costs, we distinguish between *organization's*, *non-organization's*, and *defense* costs.

Organization's costs arise directly within the organization, including everything that could be simulated and estimated using the simulator. *Non-organization's* costs are those whose calculation depends not only on the organization itself but also on external stakeholders. For example, if intellectual property is stolen, resulting in losses, assessing these losses relies not only on the organization but also on external stakeholders. *Defense* costs quantify preventive measures, which are associated with the Preparation phase. The Preparation phase covers the work an organization does to get ready for incident response. This can include employee training, risk analysis, insurance premiums, as well as preventive measures for critical infrastructure, i.e., physical security systems such as surveillance cameras, intrusion protection systems, and other measures to ensure the security of computers and equipment. This phase includes work done to prevent incidents from happening. However, these prevention costs are not included in our list in this study, as we focus specifically on expenses resulting from the incidents.

Regarding the time frame, we distinguish between *short-term*, *medium-term*, and *long-term* time frame.

The *short-term* time frame focuses on the Detection and Analysis as well as Containment, Eradication, and Recovery phase. Detection and Analysis phase includes accurately detecting and assessing the incident, while Containment, Eradication, and Recovery phase focuses on keeping the incident impact as small as possible and mitigating service disruptions.

The *medium-term* time frame refers to the Post-Event Activity phase, which includes costs incurred after the immediate period of the breach has passed. For example, fines imposed by regulatory bodies. The *long-term* time frame refers to a period after the incident has been resolved and its initial consequences have passed. For example, if stolen data is threatened to be published years after the incident has been resolved. We haven't set a strict time limit for defining short, medium, or long-term, as it varies depending on the type and duration of the breach.

B. Cost Definitions

In this section, we provide a list of all costs that may arise as a result of an incident. Each type of cost will be described and listed in Table I under the corresponding category. When an incident occurs, various stakeholders may suffer from the incident, including the affected organization, clients whose data may be compromised, investors, partners whose data may also be at risk, regulatory bodies responsible for investigating or imposing sanctions, media, and the public. Depending on the specifics of the incident and the industry, there may be other stakeholders involved as well. For each cost, it will be stated which stakeholders are affected by it.

The costs are defined as follows:

1) *External services* [6], [11], [12]. This cost includes all payments for services provided by external firms hired by the organization. These services typically include a range of activities such as log analysis, investigation of suspicious system components, identification of vulnerabilities leading to the attack, system restoration, computer forensics, reverse engineering, fixing vulnerabilities, travel expenses, and others. The primary objective of engaging these external services is to minimize the impact of the attack on the organization's operations and overall business continuity. Furthermore, legal fees and public relations (PR) expenses can also be included, as they contribute to the organization's efforts to mitigate the impact of the attack. As this can be simulated and represents a cost borne by the organization, this cost is categorized as *organization's* cost. External services can fall within both short-term and medium-term time frames.
2) *Extortion payments* [6]. This refers to the ransom paid by firms to restore access to services and data denied due to the incident. Extortion payments are a concern during the incident (short-term time frame), but even after it's resolved, there's a risk of being blackmailed for data release, with uncertain timing. This could happen right away or years later (medium-term time frame). This cost can also be simulated, and since it's a cost borne by the organization, we consider it as *organization's*. It's also worth noting that if third-party data is stolen, the cost may also be borne by that third party, but we do not take that into account here.
3) *Financial theft* [6]. Attackers can steal money from organizations through unauthorized transactions. This cost represents the exact amount of money stolen from the organization's accounts. It can be simulated and represents an *organization's* short-term cost.
4) *Card data theft* [2]. The bank covers the costs related to card number theft, such as card replacement, investigations, potential refunds, etc. This represents a *non-organization's* cost and it falls within short-term and medium-term time frame.
5) *Increased insurance premium* [6], [11], [13]. Compared to car insurance, if there's an accident, the premium usually goes up, suggesting someone might not have been careful. The same applies to incidents. If there's an incident, the insurance premium often increases. An increased insurance premium is *organization's* long-term cost as it is borne by the organization and persists beyond the immediate aftermath of the incident.
6) *Employee work time* [6], [11], [14]. This refers to the additional tasks incurred by employees following an incident, which are direct consequences of the incident. Each employee has their standard duties, but when an incident occurs, these additional tasks become necessary and are handled by IT department personnel. These additional tasks represent an *organization's* cost as they require resources to resolve. This cost falls within the short-term time frame. While these tasks arise due to the incident, it doesn't necessarily mean that regular business operations will stop. However, it can be challenging to model when they'll stop.
7) *Fines* [11], [12]. This cost includes all monetary fines or compensation that the organization may be obligated to pay as a result of a security incident. This includes fines imposed by regulatory bodies, compensation to victims, or costs of legal settlements. This cost falls within the medium-term time frame.
8) *Intellectual property theft* [11]. This cost refers to the financial loss incurred due to the theft or compromise of intellectual property. Assessing this loss depends not only on the organization but also on external stakeholders, so we consider this as *non-organization's* cost. IP theft falls within medium-term and long-term time frames due to the time needed for the individual to use the stolen intellectual property in their products or services and begin selling them.
9) *Notification and reporting cost* [6], [11], [12]. Organizations engage in stakeholder communication, determining whom to inform based on regulations that mandate notification of individuals whose data has been compromised or reporting the incident to relevant authorities. In case of a personal data breach, the organization must promptly notify the competent supervisory authority within 72 hours, unless the breach is unlikely to pose a risk to individuals' rights and freedoms [15]. This represents *organization's* short-term cost.
10) *Physical equipment damage* [6], [14]. As a result of a cyber attack, especially within SCADA systems, physical equipment may suffer damage [6]. The organization

bears the cost of repairing or replacing this equipment, making it an *organization's* short-term cost.

11) *Investment loss* [6], [11]. This cost covers the financial loss an organization experiences due to the withdrawal of investors, donors, or reduced financial support following a security incident [6]. It can involve reduced capital inflow, contract cancellations, or diminished potential resources for ongoing operations and development. It's important to note that these losses occur primarily as a result of the incident, as current contributors are no longer willing to provide financial support due to its occurrence. This represents *organization's* medium-term cost.
12) *Business process interruption* [6], [11]. This cost covers financial losses due to business disruptions, including revenue loss from service unavailability and decreased sales, often caused by attackers or deliberate shutdowns. More precisely, they refer to costs incurred by non-IT personnel involved in mitigating or resolving the incident when business processes are disrupted. It can also occur if clients are lost, reducing service demand, or if there is a lowering of employee productivity due to the emotional impact of the attack. Here, we're referring to clients directly affected by the incident who may decide not to continue doing business with the organization that experienced that incident. Industrial organizations use SCADA systems for remote data management and control. When this system becomes unavailable or partially disabled in an attack, manual intervention is required, leading to reduced operational efficiency and additional expenses compared to automated control [14]. Business process interruption is an *organization's* cost because when a business process is disrupted, the organization calculates the lost earnings for that period.
13) *Increased cost of equity*. [16] This refers to the situation where a company may face higher demands from investors for returns on their invested capital after the cyber incident occurs. This means that investors will require a higher percentage of return to compensate for the increased risk resulting from security threats and vulnerabilities that the company is exposed to after the incident. This represents *organization's* medium-term cost.
14) *Stock prices* [6]. When a cyber attack becomes public, organizations often see a drop in their stock prices. This is due to a combination of factors, including not just the organization's reputation but also industry sector dynamics and the overall market climate when the incident was disclosed. Consequently, the owners bear the financial consequences. This represents *non-organization's* cost, which we have previously stated is excluded from our total cost calculation.
15) *Reputation damage* [11], [12]. Damage to a firm's reputation may result in lost business, loss of both existing and potential clients, decreased stock market value, etc. These are all consequences that translate into costs, which have already been listed and described. Therefore, we mention reputation damage here as the cause, while the consequences are the costs we have already identified. We further elaborate on this in Section VI.

We categorized costs in Table I based on their time of occurrence and feasibility of calculation and simulation during the incident. This categorization method worked best for us, allowing us to organize costs according to our needs and effectively simulate calculations. When calculating the total cost, we consider the *organization's* costs (those incurred solely within the organization). Specifically, we model the short-term costs using the simulator (Group 1), and then we add the medium-term costs (Group 2) to the total cost calculation. Accurately estimating *non-organization's* and long-term costs (Groups 3, 4, 5, 6) poses significant challenges as their calculation relies not only on the organization but also on external stakeholders. Therefore, we do not include them in the simulation or the calculation of total damage.

C. Simulator

The experiments were conducted using the *Cyber Conflict Simulator* (CCS) [17]. This tool stands out for its ability to abstract low-level technical details of an incident and simulate all events at the tactical level. For instance, log analysis, typically a technical task, is represented in the simulator as an action without getting into technical details. However, it still requires time to execute and provides results similar to real-world analyses conducted by people. This approach aligns well with our goal of calculating the cost of a breach because we don't need to know exact technical details to determine the cost.

In the simulator, the foundation for any simulation is referred to as the *cyber landscape*. The *cyber landscape* is a description of an organization's cyber environment that includes information and communication technologies in the organization, but also people and business processes, and their interdependencies. The *cyber landscape* is, in some way, a digital twin of the organization but includes more information, not only related to digital technologies. The *cyber landscape* allows models of multiple organizations, enabling simulations of interactions between them.

The *cyber landscape* is composed of various objects, such as computers, routers, firewalls, as well as business services, loss objects, and employees. In this paper, the primary focus is on the loss object, which stores data related to expressions used for cost calculation and the specific type of loss incurred. Loss objects don't have a real-world equivalent as objects like computers and routers do. They are a part of the landscape and represent containers for storing loss-related information. Each object within the landscape is characterized by a set of attributes that are crucial for simulation. For instance, the *Computer* object features the *Is Available* attribute, determining its accessibility over the network. Overall, the organization's structure can be modeled with varying levels of detail, ranging from a high-level overview that includes

TABLE I
COST BREAKDOWN BY TYPE AND TIME FRAME

	Short-term	Medium-term	Long-term
Organization's cost	External services Extortion payments Financial theft Employee work time Physical equipment damage Business process interruption Notification and reporting cost (Group 1)	External services Extortion payments Fines Investment loss Increased cost of equity (Group 2)	Increased insurance premium (Group 3)
Non-organization's cost	Card data theft Stock prices (Group 4)	Card data theft Intellectual property theft Stock prices (Group 5)	Intellectual property theft Stock prices (Group 6)

only essential components to a finely detailed representation encompassing every individual computer, person, service, and business process within the organization.

The simulator allows players to take on various roles during the simulation. These roles could include attacker, incident manager, IT team member, or management team member. Each player is assigned a specific role, which comes with a unique set of actions. Their role involves assessing the current situation, making decisions, and assigning tasks accordingly. These tasks are then delegated to individuals, referred to as *actors* who carry them out and may later provide feedback on the results. For instance, if the IT team manager decides that log analysis is necessary for gathering more information, they will delegate this task to a simulated individual (*actor*) within the IT team. The actor will then execute the task and report the findings. Also, each employee has their *Security Awareness* control set to a specific level. As the name suggests, this control represents the employee's awareness of security, enabling them, for example, to report the receipt of potentially malicious emails. For example, if the employee's control is set to 0.5, it means that they will report suspicious emails in approximately half of the cases and not report in the other half.

In the simulation, every decision and action made by players, whether they are playing the role of attackers or defenders, significantly affects the course of events and the final results. Players initiate actions in real time within the simulation, but the simulator also provides a mode where recorded actions can be replayed. An *attack sequence* comprises the actions undertaken by the attacker in the simulation. These actions closely mirror the MITRE ATT&CK Pattern [18] tactics and techniques, such as *Recon* or *Create Spearphishing Mail*. These actions allow the attacker to gather information about the targeted organization or create deceptive emails (spearphishing) aimed at specific employees within the organization. Conversely, a defense sequence encompasses all the steps taken by defenders to protect against such attacks.

The simulator aims to recreate the duration of each action as closely as possible to its real-world counterpart. This means that actions may take several hours or even days to complete, depending on the action type and selected parameters. However, since the simulation time is limited by the duration of the exercise itself, it is essential to be able to speed up or slow down the simulated time within the simulator. CCS provides this option, allowing each player to choose one of five execution speeds available at the top of their window. Players need to agree on the execution time because the simulator adjusts to the time of the slowest player.

The simulator calculates the cost of the attack based on the cost-calculating expressions. These expressions are defined as attributes within loss objects before the simulation begins. It is possible to define expressions that accumulate the cost based on changes in the values of attributes of another object in the landscape. It is easy to apply the same cost calculation expression to different simulations and *cyber landscapes* as cost-calculating expressions are hard-coded in the simulation once defined. The part of the cost calculating expressions that are meant to be changeable are constants that represent, for example, the organization's cost of downtime or cost of employee time. The organization has the option to provide this data in the preparation stages of the exercise so the simulation is better tailored to the organization. Alternatively, the organization may choose not to disclose this information and use average values that are already in the simulator as defaults, but this may result in larger error margins.

$$\begin{gathered} Accumulator(LossIT.Loss, (1 - PathExists \\ (SKAPC01, LANSKA, time)) * 1.388889, time) \end{gathered} \tag{1}$$

The cost calculation expression is illustrated in Equation 1 representing an attribute of the *LossIT* object. When forming this expression, we can use elements from the *cyber landscape* along with predefined expressions, basic math operations, constants, and parentheses. Predefined expressions, like *Accumulator* and *PathExists*, are functions themselves. The *Accumulator* function calculates the new value by adding the product of the *val* value and the time passed to the accumulated attribute. The *PathExists* function checks if a path exists in the *cyber landscape* between the start and end nodes (objects in the landscape), returning a boolean value of true or false.

The formula for calculating costs includes the use of objects (*Loss IT, SKAPC01, LANSKA*) and their attributes (*Loss*), as well as constants. Now, let's examine the meaning of the formulated expression. Essentially, this expression is checking if the computer *SKAPC01* is connected to the *LAN SKA* network. If the *PathExists* function returns false, it means that the computer is not connected to the network. In such a scenario, a cost of $ 5000 per hour, which is equivalent to $ 1.388889 per second, as specified in the formula, would be incurred. The *Accumulator* function would then add this amount, multiplied by the time the computer is disconnected from the network to the value of the *Loss* attribute in the *LossIT* object.

After each simulation, the simulator generates a *Simulation log*. This log offers a detailed overview of the timelines for every action taken by employees in defending against the attack, as well as the number of employees engaged in these activities. It provides details on the specific action, including its start and end times, the actor in the *cyber landscape* who performed the action, and the player who initiated it during the simulation. Additionally, it includes any other relevant parameters associated with the action. An example of this type of log is shown in Table II. The *Simulation log* also includes records related to costs, as illustrated in Table III. It shows how the cost was accumulated in the loss objects over time, so it is easy to compare the timestamp with action logs and determine which actions in the simulation affected the cost.

IV. Methodology

In this section, we describe our methodology to calculate costs during and after the incident. To determine the overall cost of an incident, we aggregate the costs associated with each tactical step taken by defenders, alongside assessing the damage resulting from the tactical steps executed by an attacker targeting the organization.

To calculate the cost, the list of potential costs described in Section III-B is reviewed and the relevant costs regarding the type of the attack are selected. For example, if there's no ransomware involved or no ransom payments made, we don't include extortion payments in the calculation. The selected costs can be further divided into two groups - those that are calculated during the simulation using the simulator and those that are added after the simulation has ended. The first group of costs corresponds to the costs in Group 1 in Table I, and these represent the costs that the organization incurs up until the Recovery phase of the incident, which is when the simulation ends. The second group of costs is not simulated and corresponds to costs in Group 2 in Table I.

Our cost calculations consist of two phases. The first phase of cost calculation is executed within the simulator during the simulation using the expressions defined in the loss objects, as described in III-C. Cost-calculation expressions start to calculate the cost as objects' attributes change values because of the actions taken during the simulation. This change can be continuous if the cost-calculating expression is time-dependent or a one-time addition to the total cost if the cost-calculating expression is not time-dependent.

The second phase of cost calculation occurs after the simulation ends. Since the simulator is still under development, it is currently not possible to calculate all the costs that the organization incurs up to the Recovery phase (Group 1 in Table I) while the simulation is ongoing. These costs must be added by analyzing the simulation log. Lastly, costs that can be estimated from Group 2 in Table I, such as fines and the cost of hiring external services, are added to the total cost.

As we simulated one attack sequence described in Section V-B, we applied the methodology as follows. The first phase of cost calculation within the simulator involves:

- The simulator checks how much of the organization's network is offline, and accumulates the cost of employees not being able to access the organization's network and the Internet. Since employees can do a certain percentage of their work offline [19], the cost is lower than if the computers were completely shut down. The simulator takes this into account. This is a cost that is continuously changing while the incident is being resolved.
- When data is stolen and later published or sold, the simulator multiplies the cost of one stolen record with the number of these records. This is a one-time cost that is added to the total cost when this event occurs.

The second phase in cost calculation involves analyzing the *Simulation log* and adding any selected costs that can be estimated but are not simulated:

- The simulator captures the start and end times of each action, enabling the calculation of costs related to employees performing incident response tasks. This is done by multiplying the time spent on each task by the employee's hourly wage [11].
- Furthermore, every object in the simulator is assigned a location. So, when an actor relocates to perform a task, travel expenses and allowances are added to the total cost based on the distance traveled and time spent traveling.
- Fixed costs, like fines and payments to attackers, are also included. This includes all legal costs arising from analyzing the company policy and notifying regulators, as well as all legal disputes with clients affected by the attack.
- Furthermore, any impact on employees' usual work is factored in.

V. Experiments

In this section, we present the results of the experiments we conducted. We conducted two experiments. One in which we had one simulation with a fixed attack and defense sequence. The second experiment had three simulations. In every one of those, the attack sequence was fixed and the defense sequences were varied.

TABLE II
SIMULATION LOG - ACTION LOG

Type	Beginning	Name/Sender	Ending/Title	Actor/Receiver	Player	Parameters/Message
Action	9.7.20220:45	MalwareScan	9.7.20220:55	IT Log&Backup Admin 01 (senior)	IncidentManager	Actor: IT Log&Backup Admin 01(senior)Target: BISPC02Mode: Delete malware

TABLE III
SIMULATION LOG - LOSS LOG

Time/Date	Financial Loss -Key Account Losses (€)	Financial Loss - IT (€)	Financial Loss - NDC (€)
10.07 01:30	0	26862	39651
10.07 02:00	140000	26862	39651
10.07 02:30	195953	26862	39651
10.07 03:00	335907	26862	39651

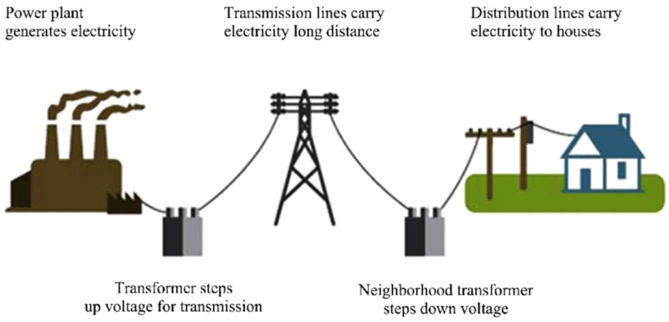

Fig. 1. Electric power system [20]

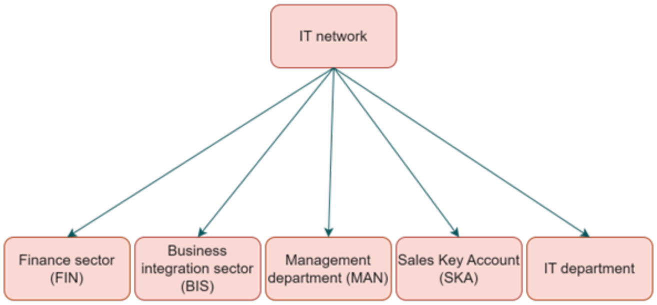

Fig. 2. The topology of the IT network in TSO Enterprise n

Our primary objective was to determine the feasibility of using the simulator as a tool to calculate the total cost of an incident. Additionally, we aimed to identify the potential limitations and shortcomings of the approach. Our secondary objective was to study how the cost of an incident changes concerning different defense tactics.

Firstly, the topology of the organization under attack will be described. Then, different defense sequences that were used to mitigate the attack will be examined, along with how the total cost of the attack is influenced by various defense tactics. Finally, the results of the experiments will be presented, emphasizing the primary sources of damage contributing to the total cost of the attack and the impact of different defense tactics on the overall cost.

A. Simulation environment

The organization selected for the experiments is designed to resemble the structure of the Croatian Transmission System Operator Inc (HOPS) [21]. The transmission system, which is managed by the Transmission System Operator (TSO), is responsible for transporting electrical energy from the power plants and transmitting it over long distances. It is one of the three key components of the electrical power system, alongside power plants and the distribution system as can be seen in Figure 1.

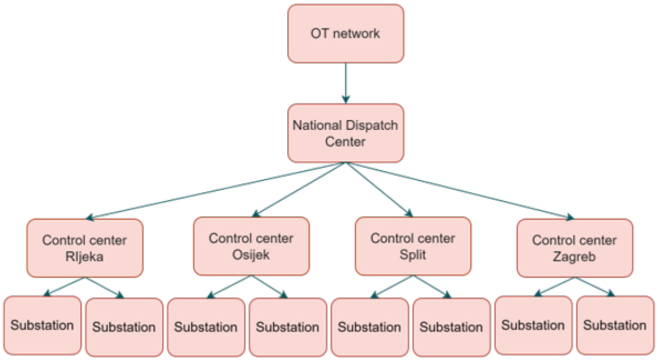

Fig. 3. The topology of the OT network in TSO Enterprise

The TSO can be divided into two main parts. The first part consists of four departments that use the corporate information technology (IT) network (Figure 2): Finance sector (FIN), Business integration sector (BIS), Management department (MAN), Sales key Account sector (SKA). The top priority in the IT network is keeping data confidential.

The second part of TSO is called the operational technology (OT) network (Figure 3). This network includes the National Dispatch Center in Zagreb and control centers in Region 1, Region 2, Region 3, and Region 4 [22]. These centers are responsible for managing substations across the country. A substation mainly consists of a transformer and equipment for control and protection. The secure transmission and distribution of electrical energy are ensured by the control centers while maintaining production and consumption in balance. A Supervisory Control And Data Acquisition (SCADA) system is used for continuous monitoring [23]. SCADA systems provide supervisory control of the geographically distributed parts of the system, in this case, substations. They collect data and send it back to the control centers. In case of emergencies or system attacks, specific measures can be taken, including disconnecting substations if necessary. Since the OT network controls physical components, ensuring that services are always available to users is the top priority. Availability means that the services managed by the control system are consistently accessible.

In addition to this organization, referred to as *TSO Enterprise*, there are also other companies, including a *Rapid Response Team Company* and the *Attacker Organization* within the *cyber landscape*.

B. Simulations

For the experiments, we used a model of the organization *TSO Enterprise*, as detailed in Section V-A. Throughout the experiments, we conducted four simulations, each featuring the same attack sequence. In every simulation, the attackers consistently executed the same attack, while the defensive strategies varied.

In all simulations, the cost calculation expressions are pre-defined within the simulator before the initiation of the simulation, as described in Section III-C. These expressions helped us calculate the cost of business downtime when parts of the network were offline or shut down. We also factored in costs related to stolen data after its sale or publication by calculating the cost for each record.

The sequence of actions taken by the attacker against the organization was as follows.

The attacker takes control of a web server. Upon gaining control of the web server, the attacker injected malicious code into the application that employees use to access their emails. The compromised web server serves as a delivery server for the malware and as a command and control server. The command and control server is used for attack coordination, facilitating activities such as sending spear phishing emails, redirecting to the infected website, and transferring stolen data, serving as an intermediary to obscure the attacker's actions. The attacker then performs recon to find out as much information as possible about the target TSO which will allow him to compromise the organization. Using various techniques, the attacker reveals employees' identities and their corresponding email addresses. The attacker crafts and sends personalized spear phishing emails to targeted individuals. These emails are meant for specific targets and contain deceptive content to lure them into opening and clicking on a link that will lead them to the infected web application. Targeted individuals open the received messages, and depending on their *security awareness* they either read the received message or click on the link, resulting in the activation of malicious code on their computers. When specific individuals click on the link, the malicious code downloads and runs on their computers, granting the attacker remote access and control. The malicious code provided the attacker with internal network access. The keylogging functionality of the malware captures authentication data for unauthorized access.

Ultimately, the attacker successfully steals documents related to the SCADA configuration as well as employment contracts. The SCADA configuration files contain data related to system settings and parameters, including information about system architecture, security settings, identification data, and other technical details relevant to SCADA systems. On the other hand, employment contracts typically contain information about employees, their roles, responsibilities, salaries, contractual obligations, and personal identification information such as names, addresses, payment details, and similar information.

The first simulation (simulation A) was used to determine how much different sources of damage participate in the total incident cost in our experiment, while the other three simulations (simulation B, C, and D) were used to assess how various defense tactics impact the overall incident cost in our experiment. The total cost of an incident is defined as the sum of various cost components, as explained in Section III-B.

We will now explain simulations A, B, C, and D in more detail.

Simulation A used a defense sequence that changed the sources of damage (Assumption sets A1, A2, A3, A4) to calculate the total cost, as shown in Table IV. These sources of damage were varied based on assumptions regarding the organization's insurance coverage against cyber attacks, the duration of business interruption during incident resolution, and whether the organization had an existing contract with Rapid Response Company (RRC) or would be engaging with them for the first time during an incident.

Simulation A starts when some employees report suspicious emails, at which point incident becomes visible to the company. The organization examines the content of these emails and checks other computers that accessed the same source, uncovering more harmful content. They enlist the help of a Rapid Response Company, which provides further details about the malicious code. Using this information, they conduct a more thorough inspection of the system, leading to the discovery of additional infected computers. All employees are warned about potential phishing emails and encouraged to report any suspicious messages. Harmful files are shared with an antivirus company for future detection. Vulnerable software and systems receive updates. At the end of the simulation, the attacker profits from one of the stolen documents by selling the sensitive data it contains, while the information from the other document is published, causing further damage to the organization. It is important to note that although the attacker may benefit from these actions, the potential financial damage to the organization is far greater due to various consequences

TABLE IV
SIMULATION ASSUMPTIONS IN SIMULATION A

Assumtion set	Cyber insurance	Interruption to business as usual	Prior contact with the RRC
A1	no	no	yes
A2	no	yes	yes
A3	yes	yes	yes
A4	no	yes	no

such as reputational damage, loss of customers, and legal fines.

Three defense strategies, in simulations B, C, and D shown in Table V, were compared to assess their impact on the overall incident cost.

The defense strategy in simulation B begins by examining suspicious emails that have been reported, identifying the source of these emails, and finding other employees who have also received malicious content in their emails. All computers that receive such emails are thoroughly checked, revealing malicious code on some of them. An external company is hired to analyze this malicious code. With the help of the external company, malicious code is removed from the organization's computers, security updates are installed, compromised user accounts are restored, and employees are warned about phishing. Again, the attacker sells some of the stolen sensitive information and publishes the rest.

In the defense sequence in simulation C, the organization only uses internal resources, and the defense sequence only differs a little compared to the first. All analyses and updates on the software are made by the employees of the organization. Also, compromised parts of the system were shut down to prevent the malicious code from spreading to other parts of the system.

In the defense sequence in simulation D, most of the defense work was transferred over to the external company. That company performed malicious code analysis and reverse engineering, digital forensics, and installation of the latest patches. Additionally, a PR company was hired.

C. Results

During each simulation, the simulator calculated part of the total incident cost. To incorporate the unaccounted costs, we obtained a simulation log after each simulation. Based on the log, we added missing costs to obtain the final value.

We would like to clarify that the numerical data and calculations provided are preliminary results obtained from simulations and have not been validated. As an indicator, we have provided ratios instead of actual numbers to give an approximate idea of the cost relationship that we obtained through simulations. At this stage, we trust the ratios more than actual numbers. Our main focus during this phase of our research was not to obtain specific numbers that would represent the total cost. Instead, our focus was on applying the described methodology in the simulator and identifying potential improvements to the methodology and the simulator used.

In simulation A, where the defense sequence was employed to manipulate the sources of damages, four different costs were generated based on a set of assumptions (Assumption set A1, A2, A3, A4) established before their calculation.

The first set of assumptions (A1) assumes that during the incident response, business-as-usual activities remain uninterrupted for employees. The organization doesn't have cyber insurance. The organization lacks cyber insurance but already holds a contract with a Rapid Response Company, thereby incurring no additional costs for their services.

The second set of assumptions (A2) differs from the first in that it assumes an interruption to the company's normal operations while employees respond to the attack, resulting in additional costs. This means that while employees focus on resolving the incident, their usual tasks are put on hold, leading to potential damages over time. For example, the system administrator needs to thoroughly check the system to try and find all the malicious code in the system instead of doing their usual work such as updating the system or installing upgrades. This interruption resulted in an approximately 13% increase in the total attack cost.

The third set of assumptions (A3) assumes that the organization has cyber insurance that pays out the insurance policy. The rest of the assumptions are the same as in the second set. The insurance policy decreased the cost of the attack by 33% compared to the second set.

The fourth and final set of assumptions (A4) considers the worst-case scenario that results in the highest total cost. In this case, there was an interruption to the business-as-usual activities while the employees were responding to the attack. Additionally, the organization lacked a contract with a Rapid Response Company (RRC) and had to hire one during the incident. Furthermore, the organization did not have cyber insurance. Compared to the first set of assumptions, the cost increased by approximately 13%. There was no significant increase compared to the second set, where the only difference was the assumption that the organization had a contract with the RRC, indicating that this cost is relatively low compared to others. However, the most notable difference is observed when compared to the third set of assumptions, with a cost increase of 50%.

Looking at all the sources of damages that were considered

TABLE V
SIMULATION ASSUMPTIONS IN SIMULATIONS B, C AND D

Simulation	Cyber insurance	Interruption to business as usual	RRC hired	PR company hired
Simulation B	no	yes	yes	no
Simulation C	no	yes	no	no
Simulation D	yes	yes	yes	yes

in calculations, we can also analyze which sources have the greatest impact on the total cost of the attack in our simulations. Costs connected to publishing and selling the stolen data are by far the greatest, at 64%. Unavailability of services and computers or business interruption costs make up about 13% of the total cost followed by the damages from an interruption to the business-as-usual activities which comes up to 11% of the total cost. Fines and legal costs make up about 10% of the total cost. All other costs sum up to the remaining 2%.

In the defense sequences that were used to compare how different defense strategies affect the total cost of the cyber incident, three total costs were calculated, one for each of the simulations B, C, and D.

In the simulation B, a Rapid Response Company was hired. However, they did not possess a cyber insurance policy. Most of the cost, around 90%, came from business losses and the value of stolen data.

In simulation C, the organization handles the attack using only its own resources. They shut down a large part of their network, causing extra costs. They didn't report the attack to regulators, so they got fined for losing important documents. With other expenses being low, the fine made up about half of the total cost in this case.

In simulation D, the organization prepared for the cyber attack and had contracts with the Rapid Response Company and a PR Company. Also, the organization recognized the importance of having cyber insurance. The most significant sources of damages in this case are again, business loss and costs connected to the stolen data, which make up about 75% of the total cost.

VI. DISCUSSION

In this section, we'll address the problems and limitations we encountered while preparing and conducting experiments. We'll start by discussing the limitations of the simulator, and then touch upon the categorization of costs. Next, we'll look into the investments and gains for the attacker, before concluding with potential paths for further research.

A. Challenges in validating cost calculation methodology

This is perhaps the most important and significant issue we encountered, and therefore we address it first. Validating our cost calculation methodology presents a significant challenge. Ideally, we would have some reference, for example, exact data about incidents and costs. This data could then be replicated with our method and results compared. However, such data of suitable accuracy is unavailable, and the best we could hope for is to compare orders of magnitude. Yet, even this is questionable as the context of the simulation and the one for the data have to be the same. Frequently, there isn't enough data to describe the context and transfer or replicate it within the simulator.

Alternatively, one might consider using the results from previously published papers that attempt to assess the costs of cyber incidents as a reference point. Yet, this is almost impossible because there's not enough information in them to be able to run simulations and compare the results.

In conclusion, this is a very hard question we'll have to tackle somehow in order to progress the field.

B. Problems and limitations encountered while using the simulator

There are several sources of errors that impact the accuracy of cost calculations, which need to be researched in more detail. Some of these were known to us before starting the research, while others became apparent during the course of our work.

Firstly, when introducing the simulator in subsection III-C, we emphasized its ability to model organizations with varying levels of detail, ranging from a broad overview to a very detailed representation. It is evident that the level of detail significantly impacts how costs are estimated and, consequently, their accuracy. This raises questions about the correlation between the level of detail in an organization's structure and the accuracy of the cost assessment of an attack. This remains an open question that we aim to explore in future research.

Another source of errors in our results arises from the fact that the simulator must calculate the duration of actions. However, these durations themselves are estimations and, as such, also contain errors. For instance, the duration of digital forensics in the simulation depends on the complexity of the task and the proficiency of the individual conducting the forensics. Characterizing these factors is challenging and requires further research. One potential approach we are considering is to use cyber ranges [24] to assess individuals' skills and, based on these assessments, create models of individuals within the simulation. This is also something that requires further research.

Some of the errors are caused by the simulator itself, i.e., its imperfection in damage calculations. When we started this research, one of the goals was to assess the shortcomings of the simulator and to identify areas for improvement that, hopefully, will be implemented in future versions. Some of those shortcomings we managed to circumvent by manually calculating and adding costs that were estimated from the simulation log. Obviously, a much better solution would be for the simulator to take into account as many sources of damages as possible.

C. Categorization of costs

The next issue we find important to emphasize is that there are multiple possible categorizations of costs. When creating Table I, we decided to group the costs by type and time period in which the cost is incurred. This decision was influenced by our specific case, namely, having a simulator that covers only the incident response period, and not post-incident or long-term periods. As a result, each source of cost has only one parameter attached to it, indicating *when* it is incurred, but the duration is only for the period in which it occurred. That's the reason why some costs, like *Card data theft*, occur several times. One variation to the categorization we have would be to add one more parameter to each cost, which defines its duration. This approach would prevent elements from repeating in the table but would add more information to the table, making it less readable.

D. Difficulties in determining cost start time and impact duration

One interesting, but probably hard-to-answer question, is when to start calculating certain costs during a simulation. In our experiment, the cost calculation started when the organization detected the first signs of an incident, specifically, when employees reported phishing emails. However, costs actually started to accumulate earlier, when the attacker breached the organization. Yet, it is a bit harder to pinpoint the exact moment. For example, if the attacker breaches the network but doesn't immediately cause damage, and the organization continues to function normally for a while, when does the damage start? Similarly, if the attacker steals data, when does the organization suffer damage?

E. Lack of precise definitions of damage

On a different note, the literature review reveals a lack of precise definitions of what constitutes damage resulting from a cyber incident. If there's no consensus on what generates costs, then calculating them becomes more challenging, perhaps even impossible. We explicitly excluded some cost categories from the simulations, assuming they're not incident-related. For instance, post-incident costs may include employee training or system updates, upgrades, and repairs [6]. Our perspective is that all these actions should be taken before an incident occurs, making them investments rather than damages. Therefore, we don't include these *preventive* costs in the calculations, as stated in subsection III-A.

In our calculations, we factored in certain medium-term expenses, such as fines and extortion payments, which could be estimated relatively accurately. However, quantifying losses from investments proved to be challenging. As the simulator primarily concentrates on the detection, response, and recovery phase of an incident, we manually incorporated these costs later.

F. The role of incident handlers in cost calculation

Furthermore, we faced a challenge in our cost calculation regarding how and whether to incorporate the contribution of individuals handling incidents into the cost calculation. Our approach to calculating the cost involved multiplying the hourly wage of such individuals by the hours dedicated to handling the incident. We justified this by recognizing that, during this time, they were unavailable for their regular tasks, resulting in damage to the organization.

However, even if the individual is not engaged in their regular tasks, it doesn't necessarily result in immediate damage. For example, let's consider a system administrator. If they are not performing their usual duties, it doesn't mean that the servers they maintain will suddenly stop functioning or encounter malfunctions. In fact, it's probable that operations will continue smoothly for a period without their direct involvement. However, over time, the risk of breakdowns increases, potentially leading to accumulated damages. This highlights the importance of further investigation. Specifically, it prompts the question: When does a business process begin to falter due to the absence of individuals responsible for maintaining the IT infrastructure upon which the process relies?

G. Incident response teams

When considering the internal incident response team (the one within the organization), the question arises of where and how to include their payment in the cost calculation, and whether to include it at all. An argument might be made that since they are paid by the organization regardless, their work shouldn't be added to the total cost of an incident.

To clarify this question, first, we believe that not many organizations have full-time internal incident response teams. Those teams cost a negligible amount of money, and if there are sporadic incidents or many incidents of low complexity, then it doesn't pay off to keep people on salary when they are not actively engaged. Instead, most organizations hire external incident response teams when a significant incident occurs, which is much more efficient. Now, when an organization hires an external incident response team, it is obvious that this cost should be added to the total cost of the incident. We argue that the same applies to internal teams as well; that is, the time they spend working on an incident should be included in the total cost. To illustrate this point, consider the case where an organization doesn't have an internal incident response team and instead hires external teams on an as-needed basis, but they constantly experience significant incidents. In that case, it is evident that we need to add the cost of the incident team to the total cost, but it is also clear that this scenario isn't much, if at all, different from the case where the incident response team is internal.

H. Investment and gains for the attacker

Another consideration is the investment required by attackers to execute an attack and the potential gains they stand to make. This is crucial for determining the likelihood of an attack, as rational attackers aim to minimize losses and their gains may not necessarily equal the damages incurred by the defense. It's important to note that we are specifically considering financially motivated attackers in this case, as APTs, hacktivists, and other non-financially motivated attackers present a different scenario.

Attackers operate in an environment of uncertainty, lacking perfect information about the costs and potential gains of their attacks. Nonetheless, they attempt to assess the profitability of an attack. According to leaked Conti chat messages [25], attackers conduct reconnaissance on their targets and select victims based on the sector and size of the organization. Conti, for instance, tends to target organizations believed to have the financial means to pay ransoms. This indicates that attackers face uncertainties in their planning and execution of attacks, and they mitigate these uncertainties by gathering information about their targets through Open-Source Intelligence (OSINT). Additionally, given the vast pool of potential victims, attackers may adopt strategies that target multiple victims simultaneously, focusing on those with weaker security measures and lower investment costs. In conclusion, this is a very interesting area of research that we intend to pursue further.

VII. Conclusions

In this paper, we proposed a novel approach for calculating the cost of a cyber attack based on tactical level simulation. We also identified potential applications of this approach in risk management, cyber insurance, and determining the real cost of incidents that happened.

Specifically, based on a literature survey, we defined different types of costs and categorized them based on their timeline and tangibility. We used a novel tool to simulate attack and defense sequences on a TSO organization. During simulations, the simulator added costs we identified as they occurred. Due to the technological shortcomings of the simulator, we had to resort to some manual calculations based on the detailed simulation log.

We ran a number of simulations in our experiment in which we varied the sources of damages and analyzed their impact on the total cost of the attack. We also analyzed how different defense tactics impact the total cost. These simulations, are very likely of questionable accuracy and validity but are only a first step toward having a method to determine cyber incident damages that could be applied to a specific organization. Our goal was to do initial research into the feasibility of such an approach but also to identify shortcomings of the simulator for it to be improved.

Finally, we discussed the drawbacks of our approach and we identified a number of potential issues that warrant further research. Some of those issues might be fundamental, while others only introduce errors in results obtained using our method. A very important question is also if it is possible to assess the magnitude of errors, as very likely those errors would not be possible to eliminate.

We intend to continue with the research described in this paper, specifically, we plan to address identified research questions while at the same time using improved versions of the simulator to make the whole process as lean as possible. To obtain the best possible results, the feedback from the scientific and professional community is invaluable, which was a primary motivation to write this paper.

References

[1] Kaspersky. Measuring financial impact of it security on businesses: It security risks report 2016. Technical report, Kaspersky Lab, 2016. https://media.kaspersky.com/pdf/b2b/kaspersky-it-security-risks-report-2016.pdf Accessed on Oct. 10, 2023.

[2] R. Anderson, C. Barton, R. Böhme, R. Clayton, C. Ganán, T. Grasso, M. Levi, T. Moore, and M. Vasek. Measuring the changing cost of cybercrime. In *Workshop on the Economics of Information Security*, 2019.

[3] trumpwhitehouse.archives.gov. CEA Report: The Cost of Malicious Cyber Activity to the U.S. Economy. https://trumpwhitehouse.archives.gov/articles/cea-report-cost-malicious-cyber-activity-u-s-economy/, 2018. Accessed on Oct. 12, 2023.

[4] E. Mossburg, J. Gelinne, and H. Calzada. Beneath the surface of a cyberattack: A deeper look at business impacts. *Deloitte*, 2016.

[5] Ponemon I. 2019 cost of a data breach report. Technical report, Ponemon Institute and IBM Security, 2019. https://www.ibm.com/downloads/cas/RDEQK07R Accessed on Oct. 10, 2023.

[6] H. Heyburn, A. Whitehead, L. Zanobetti, J.N. Shah, and S. Furnell. Analysis of the full costs of cyber security breaches. *Ipsos MORI Report*, 2020.

[7] M. Riek, R. Böhme, M. Ciere, C. Ganan, and M. van Eeten. Estimating the costs of consumer-facing cybercrime: A tailored instrument and representative data for six eu countries. In *Workshop on the Economics of Information Security*, volume 2, 2016.

[8] P. Lis and J. Mendel. Cyberattacks on critical infrastructure: An economic perspective. *Economics and Business Review*, 5(2):24–47, 2019.

[9] M. E. Kuhl, M. Sudit, J. Kistner, and K. Costantini. Cyber attack modeling and simulation for network security analysis. In *2007 Winter Simulation Conference*, pages 1180–1188, 2007.

[10] Atlassian. Get to know the incident response lifecycle. https://www.atlassian.com/incident-management/incident-response/lifecycle. Accessed on Oct. 12, 2023.

[11] R. Layton and P. A. Watters. A methodology for estimating the tangible cost of data breaches. *Journal of Information Security and Applications*, 19(6):321–330, 2014.

[12] T. Caldwell. The true cost of being hacked. *Computer Fraud & Security*, 2014(6):8–13, 2014.

[13] S. Romanosky, L. Ablon, A. Kuehn, and T. Jones. Content analysis of cyber insurance policies: how do carriers price cyber risk? *Journal of Cybersecurity*, 5(1):tyz002, 02 2019.

[14] S. Patel and J. Zaveri. A risk-assessment model for cyber attacks on information systems. *Journal of Computers*, 5(3):352–359, 2010.

[15] GDPR. Art. 33 GDPR – Notification of a personal data breach to the supervisory authority - General Data Protection Regulation (GDPR) — gdpr-info.eu. https://gdpr-info.eu/art-33-gdpr/. Accessed on Mar. 18, 2024.

[16] D Malliouris and Andrew C Simpson. Underlying and consequential costs of cyber security breaches: Changes in systematic risk. In *Workshop on the Economics of Information Security*, 2020.

[17] Utilis d. o. o. Cyber conflict simulator - utilis. https://ccs.utilis.biz/, 2021. Accessed on Dec. 7, 2023.

[18] MITRE. MITRE ATT&CK™. https://attack.mitre.org/. Accessed on Dec. 7, 2023.

[19] Deloitte. The economic impact of disruptions to internet connectivity, a report for facebook. Technical report, Deloitte LLP, 2016. https://www.deloitte.com/an/en/Industries/tmt/perspectives/the-economic-impact-of-disruptions-to-internet-connectivity-report-for-facebook.html Accessed on Oct. 10, 2023.

[20] Scientific Figure on ResearchGate. Evaluation of solar energy potential for the red sea project, kingdom of saudi arabia. https://www.researchgate.net/figure/Power-Distribution-Model-Electricity-generation-transmission-and-distribution-Source_fig1_332525724. Accessed on Nov. 12, 2023.

[21] HOPS d. d. About us. https://www.hops.hr/en/about-us. Accessed on Oct. 10, 2023.

[22] HOPS d. d. System operations. https://www.hops.hr/en/system-operation. Accessed on Oct. 10, 2023.

[23] A. Daneels and W. Salter. What is SCADA? *International Conference on Accelerator and Large Experimental Physics Control Systems*, 1999.

[24] N. Chouliaras, G. Kittes, I. Kantzavelou, L. Maglaras, G. Pantziou, and M. A. Ferrag. Cyber ranges and testbeds for education, training, and research. *Applied Sciences*, 11(4), 2021.

[25] I. W. Gray, J. Cable, B. Brown, V. Cuiujuclu, and D. McCoy. Money over morals: A business analysis of conti ransomware. In *2022 APWG Symposium on Electronic Crime Research (eCrime)*, pages 1–12, 2022.