The SPIDER Cyber Security Investment Component (CIC)

Maria Tsiodra and
Michail Chronopoulos
City, University of London,
Cass Business School,
EC1Y 8TZ London, UK
michalis.chronopoulos@city.ac.uk

Matthias Ghering and
Eirini Karapistoli
CyberLens,
18 King William St,
London EC4N 7BP, UK
matthias.ghering@cyberlens.eu

Neofytos Gerosavva and
Nicolas Kylilis
8bells,
23 Agias Paraskevis,
Strovolos, Nicosia 2002, Cyprus
neofytos.gerosavva@8bellsresearch.com

***Abstract*—Recent security incidents worldwide demonstrate the increase in the complexity and severity of cyber security threats. The attackers become better organized and the attack vectors are using more advanced methods and tools. Therefore, within the currently evolving and complex 5G cyber security landscape, both businesses and end-users need to find ways to enhance their cyber security preparedness level in order to safeguard their infrastructures and assets. Additionally, modern organizations need to invest in cyber security technologies to proactively address the identified cyber risks, based on the specific individual characteristics of their infrastructures. For this reason, investing in cyber security constitutes nowadays an essential financial and operational decision aiming to reduce the financial risk that successful cyber-attacks entail. In this paper, we demonstrate how capital budgeting techniques for gauging the financial risk of cyber attacks may be integrated within an optimisation model for optimal selection of mitigation measures into a single unified decision-making framework.**

Keywords: Cyber risk assessment, Multi-phase attacks, Set cover problem, Optimal control selection

I. Introduction

The latest quarterly security forecast report by Canalys [1], estimates a global spending of $60.2 billion on security products and services in 2021, which reflects a 10% growth compared to 2020. In the same report's worst-case scenario, the perspective is for annual growth of 6.6% taking into account a deeper and longer economic impact due to COVID-19. Thus, cyber security will remain a high priority as the range of threats gets broader and new vulnerabilities emerge, while the frequency of attacks is unlikely to decrease. Furthermore, according to [2], [3], many boards of directors will formally request improved data and understanding of the returns after years of intensive investment in cyber security. This is mainly due to a growing spending in cyber security proportionately to the investment made in new technologies.

However, despite the continuous growth in cyber security investment during the last years, over 12 billion records, containing various types of personal identifiable information, were reportedly compromised in 2020, while the number of known ransomware attacks increased by nearly 60% [4]. This demonstrates how cyber adversaries have been improving their *modus operandi* and manage to stay one step ahead of those attempting to protect their networks [5]. Eventually, as a result of the constantly increasing risk, companies must become more sophisticated and upgrade their methods of securing their assets, as attackers tend to be more incentivized to compromise an organization's infrastructure in order to achieve a variety of goals [6]. Moreover, the complexity of the methodologies used by adversaries put the organizations in a position to prioritize their security strategies. Hence, each company should consider the threats to which it is most exposed, identify the associated system vulnerabilities and take measures to mitigate them.

Consequently, investing in cyber security constitutes an essential financial and operational decision. However, it is often not feasible from a cost-benefit standpoint, since patching most, if not all, of a firm's possible security vulnerabilities may result in over-investment.

Hence, identifying the financial impact of a cyber breach and choosing the best set of mitigation measures are two of the most important challenges that organizations must tackle, yet, doing so, requires the implementation of innovative methodologies that combine risk assessment and optimization techniques. In this paper, we demonstrate how a risk assessment framework based on the discounting cash flow (DCF) method can be combined with the optimal selection of mitigation measures, formulated as a set cover problem, into the Cyber security Investment Component (CIC) of the H2020 SPIDER platform[1].

We proceed in Section II by discussing some related work and then present the integration of the CIC within the general SPIDER platform in Section III. Section IV presents an overview of the risk assessment and optimisation framework and Section V concludes offering directions for further research.

II. Related Work

A strand of the cyber security literature draws on the theory of investment under uncertainty [7], with the main objective to derive the expected value of investment in cyber security controls along with the investment threshold price and the probability of investment within a given time horizon [8]. For example, Gordon *et al.* [9] show that information sharing regarding vulnerabilities can decrease uncertainty about risks, and, in turn, the value of deferment options. More recently, Benaroch [10] develops a real options model to cast the cyber security investment problem as one of selecting a subset of uncertainty-reducing mitigation measures, whose availability is controlled by decision-makers and their size is log-normally distributed. In the same line of work, Chronopoulos *et al.* [2] analyse how uncertainty over the cost of a cyber attack and the arrival of a control impacts the optimal time of investment in cyber security. Although this line of work has contributed significantly to the area of investment under uncertainty, it ignores the degree to managerial discretion hedges financial risk, which can be measured by its Value at Risk (VaR) and by its conditional VaR (CVaR). Such risk measures can be developed to gauge the financial risk exposure of an organisation following a security breach, however, applications within cyber security economics remain underdeveloped.

Examples of empirical models that focus on the development of risk measures within a cyber security context include Wang *et al.* [11], who develop a model of investment in information security and utilise VaR to evaluate different investment tradeoffs. Specifically, using data on daily activities from a large US financial institution, they measure the risk of daily losses an organisation faces due to security exploits and use extreme value analysis to simulate the distribution of the daily losses and estimate the VaR. Rakes *et al.* [12] present an integer programming model for determining optimal countermeasure selection based on threat likelihoods, under expected value and worst-case conditions. An extension of this line of work is presented in Sawik [13], who utilises the same source of data but applies VaR and CVaR within the integer programming model of [12]. Taking the perspective of a smart grid, Law & Alpcan [14] investigate the impact of false data injection attacks and present a game-theoretic approach to smart grid security by combining quantitative risk management techniques with decision making on protective measures. Results indicates that different risk measures may lead to different defence strategies, but the CVaR allows a decision maker to prioritise high-loss tail events.

Despite their novelty, the aforementioned models overlook key uncertainties, such as the time it takes to exploit a vulnerability and the cost a system incurs once a vulnerability is compromised. Such features are also ignored in models for optimal selection of mitigation measures. Indeed, while the latter have evolved considerably from standard to multi-objective, bi-level optimisation models, these have been developed mainly within a deterministic context. For example, the problem of optimal selection of mitigation measures is often cast as a set cover problem, motivated by the application potential of coverage models to the allocation of emergency response resources [15] and to homeland security, e.g. for optimally screening checked baggage on commercial aviation flights [16]. More pertinent to cyber security is Zheng et al. [17], who cast the problem of optimal selection of mitigation measures as a set cover problem, whereby they first solve a deterministic version to analyse the incentive to implement complementary mitigations to reduce supply chain vulnerabilities. Sub-

[1]SPIDER: a cyberSecurity Platform for vIrtualiseD 5G cybEr Range services

sequently, they extend the deterministic version to allow for limitations on the choice, as well as uncertainty over the efficacy of the different mitigation measures. Also, in the same line of work as [18], [19] and [20], [21] develop a game-theoretic framework, whereby the defender chooses a security plan seeking to minimise its security risk, while the attacker aims to maximise it via the most effective attack path. This is modelled as a min-max optimisation problem, where the maximisation problem is the attacker's, and the minimisation problem is the defender's, keeping in mind the reaction of the attacker.

In order to quantify the risk exposure that a security breach entails and propose a set of optimal mitigation measures, we will draw upon the aforementioned optimisation techniques and combine them with capital budgeting methods for the evaluation of serial projects within a cyber security context. Examples of the latter within the context of project scheduling include Creemers [22], who studies the Net Present Value (NPV) of a project with multiple phases that are executed in sequence. A cash flow may be incurred at the start of each phase and a payoff is obtained at the end of the project, while the duration of each phase is a random variable with a general distribution function. The novelty of this work is that it derives an exact closed-form expression for the moments of the NPV of a project as well as a closed-form approximation of the distribution of the project's NPV. This combination of capital budgeting and optimisation techniques reflects the main functionality of the CIC, which will implement the cyber economic models within the project context. In this paper, we will present the innovation aspects of this tool, its reference architecture and the process of suggesting security controls, the details of the sub-components composing the CIC, the investment decision support mechanisms implemented by the CIC, the interactions of the CIC with other SPIDER components and the way of visualising the results to the end user.

III. The CIC Integration in the SPIDER Architecture

The CIC takes into account and integrates uncertainties regarding the time that an adversary needs to exploit a vulnerability and the associated cost, with the aim of assessing the estimated value of the cost that the 5G network will bring to its owners once its weaknesses have been exploited. Then, by implementing the necessary mitigation measures (controls), the CIC focuses on improving the coverage of the vulnerabilities in each asset. The CIC's outputs will be fed into the SPIDER dashboard enabling the 5G system administrators to reach the best investment decisions, taking into account any required resource constraints. This is achieved by enabling the 5G infrastructure risk auditors and investment decision support managers to communicate actively with the SPIDER platform in order to provide preferences, rules, policies, recommendations, and risk priorities, which will be used afterwards to instantiate the SPIDER cyber economic models. The ultimate goal is to produce a CIC whose outcomes would be interpretable and adaptable to risk and monetary changes and constraints. Hence, the novelty of the CIC is twofold:

i. The economic framework facilitates a thorough assessment of the organisation's risk exposure, taking into account the sequential nature of a cyber attacks and key associated uncertainties.
ii. The results from the economic models are used by optimisation functions to determine the optimal sets of measures for mitigating cyber risk subject to a budget constraint and risk preferences.

The CIC incorporates user expectations, rules, policies, suggestions and risk priorities generated by the SPIDER dashboard, as well as system data provided by the Continuous Risk Assessment Engine (CRAE) and the SPIDER platform in order to provide personalized real-time investment suggestions. Figure 1 illustrates how the CIC outputs could be displayed in the SPIDER dashboard. Each rectangle in the diagram represents a distinct page with specific functionality and information. The first page of Figure 1 provides a summary of vulnerability statistics allowing the Risk Auditor to quickly assess the system's status. The Asset Overview section presents a network of all the assets inside the system; the presence of vulnerabilities in assets can be indicated by colour coding them. Page 3 presents a complete list of vulnerabilities in the system. Each of the vulnerabilities is represented by their respective identifier. Through the vulnerability list section, the Risk Auditor can also proceed to the Vulnerability details page. An overview of the vulnerability, remediation recommendations, and a

list of affected assets can be found on this section. The fifth page shown in Figure 1 allows the user to select optional constraints and preferences to be used in the control optimisation. The suggested optimal controls can be found on the last page of the UI.

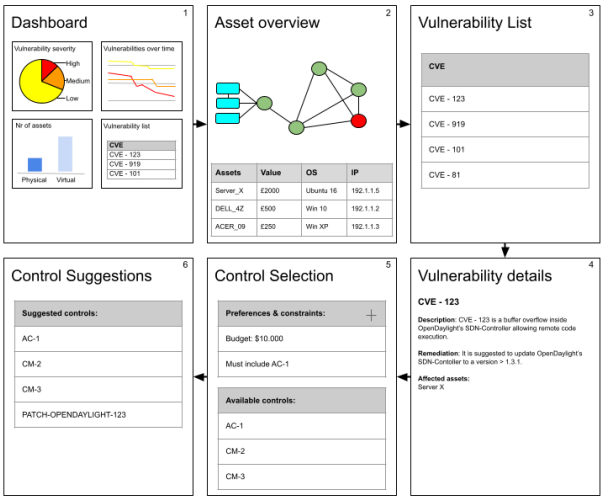

Fig. 1. CIC related pages of the SPIDER dashboard.

The CIC, which is also referred to as Decision Support System (DSS), is highlighted in red in the partial view of the SPIDER reference architecture presented in Figure 2. As shown in the SPIDER platform's reference architecture the DSS receives the measured risk from the Risk Calculation Engine (the part of SPIDER that calculates risks based on given assets relationships, vulnerabilities, controls, and threat appetites) and sends the suggested mitigation actions/security controls to the Operational Dashboard for Emulated Scenarios as well as to the Operational Dashboard for Simulated Scenarios.

Also, as illustrated in the Figure 3 the CIC consists of:

i. The CRAE feeds the CIC with the identified vulnerabilities and their respective successful exploitability probability of the 5G infrastructure, the 5G deployment assets as well as their respective connections. Also, it provides a set of controls that can be used in order to mitigate the risk of the potential exploitation of these vulnerabilities. Since the CRAE is able to continuously assess the risks related to the 5G infrastructure, the CIC must wait until it has obtained adequate amount of information in order to calculate an optimal decision, and will have to automatically recalculate its decision once new data has arrived.
ii. The SPIDER Dashboard, feeds the CIC with a variety of user data, including budgeting constraints, regulations and additional user preferences.
iii. The Kafka client helps the CIC to interact with other SPIDER components including the CRAE and the SPIDER Dashboard. It parses the data coming from the CRAE as well as the input coming from the SPIDER Dashboard and stores it within the User/System database.
iv. The User & System DB stores the data collected by the Kafka client for later use by the Economic Models. Some of this data can be fed back to the Kafka client in the form of statistics, such that they can be visualised by the SPIDER Dashboard.
v. The economic models use the data stored in the CIC's database to derive the valuation of a serial cyber security breach.
vi. The valuation produced by the Economic Models will be optimised by the Control Optimisation. This process results in a set of optimal controls given constraints provided by the user.

IV. Example Scenario

A. Architecture

The CIC will determine the optimal set of controls for an emulated 5G architecture. Here, we will introduce a simple example of a 5G architecture that can be emulated by the SPIDER Platform. An overview of the architecture is presented in Figure 4 and consists out of the following components:

i. Open-Source MANO (OSM) is in charge of the orchestration of various network functions across all the computing domains. This makes the OSM a valuable target.
ii. The computing domains are represented by the (VIM #1 and VIM #2). These domains are virtual machines that can be hosted in a variety of (physical) places from datacentres to small edge computing devices, such as a server rack near a 5G antenna. As the name suggests, computing domains facilitate computing resources to a number

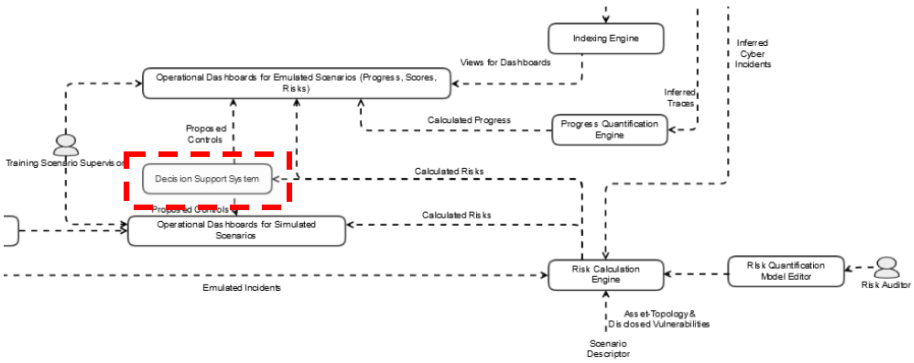

Fig. 2. Partial view of the SPIDER reference architecture.

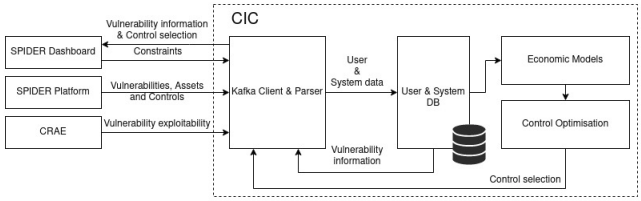

Fig. 3. Overview of the CIC architecture.

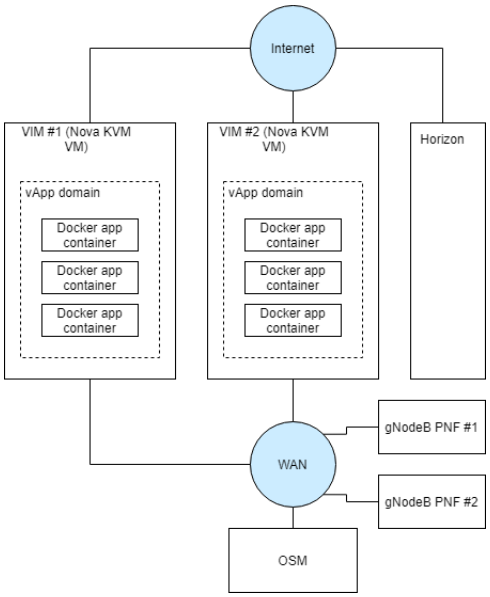

Fig. 4. A simplistic architecture overview of an artificial 5G infrastructure.

of applications and network functions. These applications can be owned by the telecommunication company itself or by third parties. Examples of such third parties would be streaming companies caching high demand movies closer to the end user, or autonomous vehicle companies offloading some of the vehicle's computations to nearby (low-latency) computing domains.

iii. To ensure that applications can't affect each other, they are run in separate isolated docker containers.

iv. The Horizon dashboard is used by the system administrators to preform maintenance on the computing domains. A system administrator with the appropriate permissions can use Horizon to add, remove or modify computing domains.

v. There is a Wide Area Network connecting the OSM, gNodeB base stations, and the computing domains.

B. Risk Assessment

Figure 5 illustrates an example scenario, where the attack consists of 3 steps. Hence, in this case, $i = 1, 2, 3$ denotes the assets and each one has $j = 1, 2, 3, \ldots, m_i$ vulnerabilities, i.e. $\mathcal{V}_i = \{v_{i1}, v_{i2}, \ldots, v_{im_i}\}$. For example:

- $\mathcal{V}_1$ = {A Malicious Tenant uses his legitimately obtained position, An adversary leverages a Remote Code Execution (RCE) based on the CVE-2019-8943 vulnerability, An adversary leverages a Remote Code Execution (RCE) based on the CVE-2018-13415 vulnerability, Brute forcing Horizon administration passwords}
- $\mathcal{V}_2$ ={Exploiting a docker containerisation vulnerability CVE-2019-14271 to escape the docker

container, Gaining access to an existing VM by exploiting the CVE-2020-12689 vulnerability, Using Horizon privileges to create a new VM}

- $\mathcal{V}_3 = \{$Compromising the OSM$\}$.

In the first step, the attack must exploit one of the vulnerabilities in $\mathcal{V}_1$ in order to gain entry to the system. Once the first step is carried out, the attack moves to the second phase, where the attacker must exploit one of the vulnerabilities in $\mathcal{V}_2$. Note that after completing the second phase there is an optional step to move laterally to another virtual machine. Finally, the last step would be to compromise the OSM by exploiting the vulnerability in $\mathcal{V}_3$.

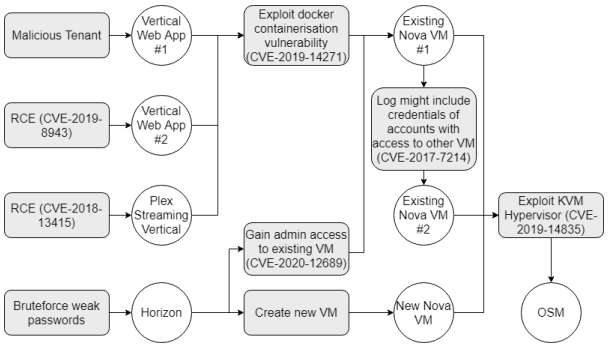

Fig. 5. Attack graph representing the attacker actions (edges/ovals) and resulting compromised states (nodes/circles) in a vulnerable artificial 5G infrastructure.

Each one of the three stages of the attack entails a financial impact for the organisation, which, as a function of the random exploitation time, is itself also a random variable. Indeed, the impact of the attack on any of the assets of the network can be expressed as:

Impact = (Asset Value) x (likelihood of being attacked) x (probability of being compromised)

Since the attacker may require a substantial amount of time to exploit a vulnerability [23], risk assessment should consider the present value of this impact. In turn, this introduces the need to estimate the distribution of the exploitation time, and, subsequently, the notion of discounting within the estimation of the impact of a cyber attack. Therefore, the risk assessment functionality of the CIC entails the calculation of the probability distribution of the expected impact taking into account various underlying uncertainties. The robustness and novelty of the CIC is reflected on the calculation of the expected impact based on input from the CRAE. The latter combines different sources of information, such as the business profile of organisations and cybersecurity information collected by CERTs and/or CSIRTs, and carries out risk analysis based on real-time monitoring of target infrastructures simulated in cyber ranges during training/preparedness sessions, thereby enabling a real-time analysis of cyber risks, threats and vulnerabilities of target systems.

Once the distribution $f_i(\cdot)$ is determined, the CIC will produce specific risk measures, e.g. VaR_α or CVaR_α, to gauge the financial risk exposure of the cyber attack, as shown in Figure 6. Note that VaR is the minimum project value for a given confidence level, α, during a specified time horizon, and CVaR is the expected value of the project given that it is less than the VaR.

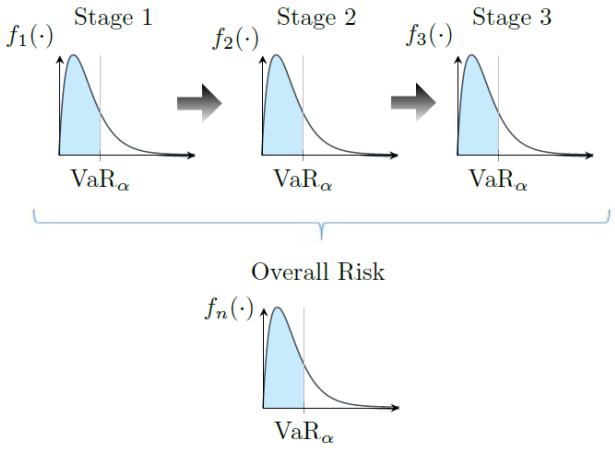

Fig. 6. Risk assessment of serial attack.

C. Optimisation of Controls

In its basic implementation, the optimisation framework within the CIC is formulated as a deterministic set cover problem but can be extended to a stochastic variant as in [17]. The optimisation objective is to reduce the anticipated cost of a security breach by applying patches on the affected areas of the system based on the required level of security coverage and patch efficacy. Thus, the CIC aspect would strive to reduce the cost of a cyber-attack while balancing financial and efficacy limitations. Hence, within the SPIDER context, the optimisation goal is constrained by: (a) budget restrictions due to the limited availability of financial capital B that affects

investments on mitigation measures; and (b) limitations relating to the organization's desired degree of efficacy. We denote by $\mathcal{C} = \{C_1, C_2, \ldots, C_\ell\}$ the set of available controls and E_{ijl} the efficacy of control C_l against vulnerability v_{ij}, where $l \in \{1, 2, \ldots, \ell\}$. Intuitively, E_{ijl} reflects the degree of protection offered by control C_l for a vulnerability v_{ij} for all assets $i = 1, 2, \ldots, n$ of the network. Also, x_l denotes whether a control is selected and y_l is the associated cost.

$$\min \sum_{l=1}^{\ell} x_l \tag{1}$$

$s.t.$

$$\sum_{l:v_{ij} \in C_l} x_l \geq 1, \quad \forall\ v_{ij} \tag{2}$$

$$\sum_{l=1}^{\ell} x_l y_l \leq B \tag{3}$$

$$x_l \in \{0, 1\} \tag{4}$$

Note that the solution to (1)-(4) reflects the minimum number of controls that offer a baseline coverage. Consequently, this formulation does not provide information about the residual risk following the implementation of the controls, and, therefore, it should be extended to include the risk measures from Section IV.B either in the objective function or in the constraints.

V. Conclusion

Efficient cybersecurity risk management relies on managerial strategies that are responsive to the various uncertainties associated with cyber attacks. The need for such strategies becomes particularly pronounced considering the critical impact that cyber attacks may have on organisations and the often very limited time to make executive decisions. In this paper, we take into account the serial nature of a cyber attack as well as key underlying uncertainties and develop an analytical framework to: i. evaluate the risk exposure of an organisation; and ii. propose an optimal set of mitigation measures. Thus, the contribution of our framework is that it extends the traditional DCF approach beyond a static context in order to demonstrate its application potential within a more complex setting that combines asset valuation, risk management and optimisation. To demonstrate the novelty of our model, we analyse the economic implications of a cyber attack by developing a case study based on a 5G network.

Directions for further research may include the extension of the proposed model for optimal selection of mitigation measures by casting it as a knapsack problem. This will not only address the limited scope of set cover problem, but, in addition, it will facilitate comparisons regarding the efficiency of different methods in terms of mitigating cyber risk. Furthermore, this approach will facilitate the direct integration of different risk measures within the objective function. Finally, game-theoretic cnsiderations as in [21] may also be including within the same framework.

Acknowledgements: This work was supported by the EU, H2020 and Project number 833685.

References

[1] Canalys, "Cybersecurity investment to grow 10% in 2021," https://www.canalys.com/newsroom/canalys-cybersecurity-2021-forecast?ctid=1912-da25dd360453873bd54b4bead2e63f3c, 2021.

[2] M. Chronopoulos, E. Panaousis, and J. Grossklags, "An options approach to cybersecurity investment," *IEEE Access*, vol. 6, pp. 12 175–12 186, 2017.

[3] ENISA, "Enisa threat landscape - emerging trends," https://www.enisa.europa.eu/publications/emerging-trends, 2020.

[4] Help Net Security, "Cybersecurity investments will increase up to 10% in 2021," https://www.helpnetsecurity.com/2021/01/26/cybersecurity-investments-2021/, 2021.

[5] ENISA, "2020 it spending: Cybersecurity remains an investment priority despite overall it budget cuts, kaspersky found," https://www.businesswire.com/news/home/20200930005611/en/2020-IT-Spending-Cybersecurity-Remains-an-Investment-Priority-Despite-Overall-IT-Budget-Cuts-Kaspersky-Found, 2020.

[6] L. A. Gordon and M. P. Loeb, "The economics of information security investment," *Economics of Information Security*, pp. 105–125, 2004.

[7] A. K. Dixit, R. K. Dixit, and R. S. Pindyck, *Investment under uncertainty*. Princeton university press, 1994.

[8] A. Etheridge and M. Baxter, *A course in financial calculus*. Cambridge University Press, 2002.

[9] L. A. Gordon, M. P. Loeb, and W. Lucyshyn, "The impact of information sharing on cybersecurity underinvestment: A real options perspective," *Journal of Accounting and Public Policy*, vol. 34, no. 2, pp. 509–519, 2015.

[10] M. Benaroch, "Real options models for proactive uncertainty-reducing mitigations and applications in cybersecurity investment decision making," *Information Systems Research*, vol. 29, no. 2, pp. 315–340, 2018.

[11] J. Wang, A. Chaudhury, and H. R. Rao, "Research note—a value-at-risk approach to information security investment," *Information Systems Research*, vol. 19, no. 1, pp. 106–120, 2008.

[12] T. R. Rakes, J. K. Deane, and L. P. Rees, "It security planning under uncertainty for high-impact events," *Omega*, vol. 40, no. 1, pp. 79–88, 2012.

[13] T. Sawik, "Selection of optimal countermeasure portfolio in it security planning," *Decision Support Systems*, vol. 55, no. 1, pp. 156–164, 2013.

[14] Y. W. Law, T. Alpcan, and M. Palaniswami, "Security games for risk minimization in automatic generation control," *IEEE Transactions on Power Systems*, vol. 30, no. 1, pp. 223–232, 2015.

[15] L. Brotcorne, G. Laporte, and F. Semet, "Ambulance location and relocation models," *European Journal of Operational Research*, vol. 147, no. 3, pp. 451–463, 2003.

[16] S. H. Jacobson, L. A. McLay, J. E. Kobza, and J. M. Bowman, "Modeling and analyzing multiple station baggage screening security system performance," *European Journal of Operational Research*, vol. 52, no. 1, pp. 30–45, 2005.

[17] K. Zheng, L. A. Albert, J. R. Luedtke, and E. Towle, "A budgeted maximum multiple coverage model for cybersecurity planning and management," *IISE Transactions*, vol. 51, no. 12, pp. 1303–1317, 2019.

[18] H. M. Almohri, L. T. Watson, D. Yao, and X. Ou, "Security optimization of dynamic networks with probabilistic graph modeling and linear programming," *IEEE Transactions on Dependable and Secure Computing*, vol. 13, no. 4, pp. 474 – 487, 2016.

[19] A. Fielder, E. Panaousis, P. Malacaria, C. Hankin, and F. Smeraldi, "Game theory meets information security management," in *IFIP International Information Security Conference*. Springer, 2014, pp. 15–29.

[20] V. Viduto, C. Maple, W. Huang, and L.-P. David, "A novel risk assessment and optimisation model for a multi-objective network security countermeasure selection problem," *Decision Support Systems*, vol. 53, no. 3, pp. 599–610, 2012.

[21] M. Khouzani, Z. Liu, and P. Malacaria, "Scalable min-max multi-objective cyber-security optimisation over probabilistic attack graphs," *European Journal of Operational Research*, vol. 278, no. 3, pp. 894–903, 2019.

[22] S. Creemers, "Moments and distribution of the net present value of a serial project," *European Journal of Operational Research*, vol. 267, no. 3, pp. 835–848, 2018.

[23] A. Marotta, F. Martinelli, S. Nanni, A. Orlando, and A. Yautsiukhin, "Cyber-insurance survey," *Computer Science Review*, vol. 24, pp. 35–61, 2017.