Proceedings of the 2023 Winter Simulation Conference
C. G. Corlu, S. R. Hunter, H. Lam, B. S. Onggo, J. Shortle and B. Biller, eds.

A NETWORK THEORY TO QUANTIFY AND BOUND CYBER-RISK IN IT/OT SYSTEMS

Ranjan Pal
Sander Zeijlemaker
Michael Siegel

Sloan School of Management
Massachusetts Institute of Technology
E94, 245 First Street
Cambridge, MA 02142, USA

Rohan Xavier Sequeira

Electrical and Computer Engineering
University of Southern California
3740 McClintock Ave
Los Angeles, CA 90089, USA

ABSTRACT

IT/OT driven industrial control systems (ICSs) such as water/power/transportation networks are increasingly meeting the daily functional needs of civilian society around the globe. This, alongside making societal businesses more automated, efficient, productive, and profitable. However, often poorly configured IoT security settings increase the chances of occurrence of (nation-sponsored) stealthy spread-based APT malware attacks in ICSs that might go undetected over a considerable period of time. The ICS enterprise management is often keen to get *apriori* statistical estimates of cyber-loss impact post any cyber-attack event such that it can plan ahead on its cyber-resilience budget. In this paper, we propose the first mathematical theory, based upon stochastic processes and concentration inequalities, to (a) statistically quantify apriori the cyber-loss impact (distribution) on an ICS infrastructure network post an APT cyber-attack event, and subsequently (b) bound the tail of such a cyber-risk distribution, for arbitrary impact distributions.

1 INTRODUCTION

Cloud and sensor network driven machine-to-machine (M2M) communication is triggering a paradigm shift in the way various industrial (decision-making) processes are conducted and managed. The aggressive adoption of this M2M communication for service applications in industrial control systems (ICSs) such as smart grids and across various (critical) industries/industry verticals such as automotive, utilities, home automation, healthcare, and security, is expected to rapidly accelerate the industrial Internet of Things (IIoT) market. To drive home our point, the global IIoT market, as of 2021, is worth USD 100+ billion dollars (projected to reach a trillion USD by 2028), with a steady yearly growth rate, i.e., a CAGR of $\approx$ 22.8%, according to a recent report by *Grand View Research, Inc.*

The effectiveness of an IIoT controlled ICS enterprise (e.g., a smart grid, a smart factory) solely relies upon the reliable and resilient functioning of networked IoT devices underlying an ICS infrastructure that operate collaboratively in collecting, transmitting, relaying, and intelligently processing application information over wireless/wired communication network (e.g., 5G, WiFi, Ethernet). However, the last half a decade has seen (nation-state sponsored) attacks with increasing frequency on IoT-driven industrial control systems (ICSs). The prime set of reasons (see (Xenofontos et al. 2021)) attributed to this trend are (a) the increase in the density and scale of sensor networks associated with an organization's underlying IoT network infrastructure forming a large cyber-attack surface that has sub-optimal OT/IT air-gaps, (b) hard-to-replace old but Internet exposed OT and control equipment modestly capable of being robust to cyber-attacks, (c) IoT-equipped machines (e.g., HMI computers, SCADAmaster computers, PLCs) from multiple vendors running a patchwork of proprietary, heterogeneous, and non-updatable software, (d) poor or absent (behavioral) organizational cyber-security practices (e.g., poorly configured default security IoT device

979-8-3503-6966-3/23/$31.00 ©2023 IEEE

settings), (e) organizational C-suites allocating insufficient budget to implement cyber-security awareness, monitoring, and prevention technology, (f) a rise in hybrid and in-secure remote work environments post the COVID outbreak, and (g) a significant rise in the number, type, and quality of cyber-attackers.

The above set of reasons increase the chances of occurrence of stealthy but popular *advanced persistent threat* (APT) attacks in IoT driven ICS network systems that (a) spear-phish IT administrators or other employees to gain persistent high level IT network access within the target industrial site via (but not limited to) user accounts, hardware and software assets, (b) consequentially, move laterally inside the site to find and exploit insecure devices and machines (e.g., botnets discovering open RDP ports using the *Shodan* search engine), (c) usually go undetected (due to weak security monitoring) over a considerable period of time (e.g., via a timed logic bomb in a malware that might delay activation of specified maximal adverse impact event(s)) - consequently not allowing the defender(s) to segment and isolate the network, and (d) result in cascading IoT device (and machine) failures that disrupt device and employee networked communication systems, and cause significant physical and/or service quality damage in the long-run. *The quantitative measure of such an adverse impact in an arbitrary IoT driven ICS network is usually a time-dependent and non-deterministic (random) variable over a family of APT cyber-attacks reflecting a loss (usually converted to tangible units w.r.t. a quality of service (QoS) metric) with financial consequences for the ICS driven enterprise.* As an example of an outcome of such a random variable, consider a modern cyber-inspired version of the Northeast power blackout of 2003 that was triggered by a transmission line failure and cascaded into a massive power failure affecting 55 million people in the northeast region of the USA, and resulted in an economic (first and third-party damage) of USD 6 billion (Minkel 2008).

1.1 Research Motivation and Goal

It is often the case in recent years, during the wake of major cyber-attacks in the past decade, that an ICS enterprise management (e.g., CEO, CISO, board) is interested to tangibly estimate *apriori* statistical metrics (e.g., mean, tail-risk) related to the cyber-loss impact post a cyber-attack (e.g., via an APT) event to budget cyber-resilience management within the enterprise. This is important to minimize tangible (e.g., monetary, stock value) and non-tangible (e.g., reputation) multi-party losses post an inevitable future cyber-attack event (as part of incident response) through investing effectively in cyber-protection mechanisms. Note that it is usual in practice for managers to find it difficult to estimate the hypothetical impact of an adverse cyber-incident (Butler 2002). This difficulty is aggravated by uncertainties in the knowledge of the cyber-risk terrain, system complexity, lack of cyber-incident data and cyber-loss impact metrics, and the inability to predict future cyber-incidents (Jalali et al. 2019; Komljenovic et al. 2016). Add to this is the role of cognitive biases that prevent even the most experienced of system managers to assess the impact of cyber-risk accurately enough (Tversky and Kahneman 1974; Madnick et al. 2016; Straub and Welke 1998). *In such environments, it is best that managers consider worst case statistical estimates of cyber-risk impact (the analysis of which needs to consider a family of parameters for a given cyber-attack type) into account rather than work with a faulty perception of exact cyber-risk impact.* As management guru Peter Drucker once famously said: *"if you cannot measure it, you cannot manage it"*. Our **research goal** in this paper is to develop a mathematical theory that (a) develops apriori, a generic closed form expression for the cyber-loss impact post the occurrence of a cyber-attack event drawn from a family of malware driven APT cyber-attacks, and (b) generates worst case apriori statistical estimates of cyber-loss (evaluated via the closed form expression) for these cyber-attacks to support cyber-resilience management planning.

1.2 Research Contributions

We make the following theoretical research contributions in this paper.

1. We design a novel stealthy cyber-malware spreading framework for a parameterized family of APT-type cyber-attacks in IoT driven ICS networks, that captures the time-varying *attack-defense-impact* trio as a time-dependent *Markov-Feller* (MF) continuous stochastic process. This process

is ideally suited to capture the stealthy infection spread of a parameterized family of APTs across a network topology, and their impact launching time periods. Our proposed model extends a huge literature on *attack-defense* type models that omit providing a systematic framework to quantify the adverse impact on (wireless) organizational network assets when modeling cyber-malware spread as a time-dependent continuous stochastic process only (see Section 2).

2. Furthermore, for the MF malware spread process, we provide a closed-form analysis of the non-deterministic time-aggregate adverse impact of an APT attack on the entire ICS network. As a research novelty, this 'time-space' adverse impact analysis extends the *Factor Analysis of Information Risk* (FAIR) model (Freund and Jones 2014)(a standard industry model to quantify cyber-risk) to settings where network connections between system risk variables are explicitly accounted for, and showcases the likelihood of extreme impact events in the ICS network (see Section 3).
3. We derive tight error bounds of empirical Conditional-Value-at-Risk (CVaR), i.e., the measure of APT cyber-risk, that is usually obtained in practice, w.r.t. the true theoretical CVaR estimates of the time and space aggregate adverse impact distribution in an ICS network. Since cyber-risk managers will only have access to empirical estimates of CVaR, the latter's accuracy is of paramount importance to their business. To this end, we conduct the derivation of (theory-practice) error bounds using a rigorous analysis based on the *theory of large deviations* (TLD) in probability theory. Specifically, we derive a tight upper and lower bound of the CVaR (APT risk) estimation error using properties of the *Chernoff-Hoeffding* and the *McDiarmid* concentration inequalities from TLD, respectively. *The novelty of this analysis is in the tight (empirical) estimation error bounds obtained in theory as a function of the finite number of empirical data samples of the impact distribution that a cyber-risk manager might have access to in practice.* The theory also hints at the threshold number of empirical samples from the adverse impact distribution a cyber-risk manager (e.g., insurer) should demand of the ICS network manager to satisfy its error tolerance (see Section 4).

Related Work and Novelty - Deriving formal tail cyber-risk bounds with performance guarantees is completely new in the networked cyber-risk quantification literature - let alone such quantification in APT affected ICS network settings. In recent years, there has been a number of efforts modeling the spread of stealthy cyber-malware (e.g., APTs) in arbitrary communication networks using approaches seeded by the SIS methodology (Yang et al. 2015; Yang et al. 2017; Yang et al. 2017; Xu et al. 2011; Xu et al. 2012; Xu et al. 2014; Wu et al. 2017). In the same regard, the use of the network-based SIS model (also known as contact processes), as a continuous time Markov chain, to model an epidemic spread is a standard in computer epidemiology studies, and has been used and analyzed in multiple research efforts (Liggett 2012; Liggett 2013; Bezuidenhout and Grimmett 1990; Durrett and Liu 1988; Mountford et al. 2016; Bailey 1975; Van Mieghem et al. 2008; Van Mieghem and van de Bovenkamp 2015; Pastor-Satorras et al. 2015) - be it related to cyber-infections, or otherwise. *However, none of these above-mentioned works on the dynamics of spread processes account for the time and node aggregate cyber-loss impact (and their statistics) post a cyber-breach incident.* This is a basic necessity for enterprises to budget plan on cyber-resilience management activities prior to the occurrence of any cyber-attack.

2 THE APT MALWARE SPREAD PROCESS SPECIFICS

In this section, we propose models to (a) capture stealthy malware spread dynamics in an ICS network characterized by an APT attack, and (b) formulate the node and time aggregate adverse impact of the spread in the ICS network. *Wherever applicable, we complement modeling elements with real-world parallels borrowed from popular cyber-attacks conducted on ICS enterprises.*

2.1 The Spread Model

Network Model - We consider an ICS network of IoT and/or OT devices with sensors (henceforward, both being referred to as an IoT device) labeled $1,\cdots,N$ on a simple unweighted bidirectional graph.

Each IoT device (inside an IoT-driven ICS subnet and/or across subnets) is capable of getting infected (a) *indirectly*, i.e., by malware transmission (e.g., via emails, AUTORUN, open port exploitation through message broadcasting) from neighboring infected nodes (e.g., post the event when a DMZ inside an ICS is breached), and (b) *directly*, for example via it downloading malicious code from the Internet or the code being injected on them via a backdoor (e.g., the event when the DMZ is breached due to social engineering attacks or due to infected plug-able external devices). The network is represented by a symmetric adjacency matrix $A \in \{0,1\}^{N\times N}$, with $a_{ii} = 0$ for all i, and $a_{ij} = 1$ indicates a connection between network nodes i and j, and $a_{ij} = 0$ indicates otherwise.

As examples of a direct cyber-infection from popular cyber-attacks, we have (a) phishing-driven *BlackEnergy3* malware infecting IoT driven components (acting as nodes) of the Ukraine power grid (in 2015) via which login credentials for these components were obtained by hackers, (b) camera software vulnerabilities exploited by hackers to get entry into computers (both cameras and the computers acting as nodes) of the SCADA systems of a Turkish oil pipeline (in 2008), and (c) the *Stuxnet* worm utilizing four zero-day exploits to infiltrate the Supervisory Control and Data Acquisition (SCADA) systems controlling uranium centrifuges (acting as nodes among other CPS components that include PLC-controlled variable frequency drives (VFDs) and print spoolers) in an Iranian nuclear plant. As examples of corresponding indirect cyber-infection (post the direct infection) from the above-mentioned cyber-attacks, we have (a) hackers opening switches that distribute power to the Ukrainian power grid and overwriting switch-controlling firmware controlling serial-to-ethernet controllers, (b) causing the Turkish oil pipeline to become over-pressurized via control commands on the IoT-controlled SCADA computers, and (c) causing the uranium centrifuges to slow up and down, crossing through mechanical resonances, till their failure via compromised PLC controller controlled VFDs. More generally, the direct-indirect nature of cyber-attacks on ICSs have been studied in (Jalali et al. 2019; Martinez-Moyano et al. 2015; Kure and Islam 2019; Sepúlveda Estay 2021; Zeijlemaker et al. 2018).

Threat Model - We consider cyber-threats that are representative of the malware-induced advanced persistent threat (APT) family (e.g, *WannaCry, NotPetya*) popularly affecting many ICS networks today. The *initial stage* of an APT, i.e., the spread of cyber-infection (malware such as the *BlackEnergy3*) through an ICS network (e.g., by open port scanning of vulnerable IoT devices) post initial malicious code injection on a set of devices, is (often) dynamically modelled using the seminal susceptible-infected-susceptible (SIS) methodology (Pastor-Satorras, Castellano, Van Mieghem, and Vespignani 2015). WLOG, we adopt the SIS model in our work. The rationale being that (a) it is intractable to plug all security deficiencies in a computer device, leave alone a system of devices (Pfleeger and Cunningham 2010), and (b) consequently, *IoT device i in the network is never immune*, i.e., *always eventually susceptible to infection in the cyber-world, despite measures taken via technology and/or human efforts to repair it post attack or prevent it from being attacked* (Anderson and Moore 2009). The latter point is because IIoT network security is primarily about the use of IDSs and security-monitoring driven alarm systems that are merely detective measures, and not preventive measures (though all existing preventive measures are inevitably imperfect).

2.2 The Spread Dynamics

State Evolutions - The state of node i at time t is denoted by $X_i(t)$, where $X_i(t) = 1$ indicates that i is infected at time t, and $X_i(t) = 0$ indicates that is susceptible. Each node can be infected by its neighbors, but is cured independently of all other nodes in the network. Each node in the IoT network is endowed with an independent exponential clock and changes its state when the exponential clock rings (from (Fahrenwaldt et al. 2018)). The rate of state changes by node i is given as a Markov chain:

$$\begin{aligned} X_i &: 0 \to 1, \quad \text{with rate } \left(\frac{1}{\alpha} + \beta \sum_{j=1}^{N} a_{ij} X_j(t)\right) \\ X_i &: 1 \to 0, \quad \text{with rate } \delta, \end{aligned} \tag{1}$$

for $\alpha, \beta, \delta > 0$ (from (Fahrenwaldt et al. 2018)). Here (a) α is the probability that a node becomes infected directly, e.g., by downloading malicious code from the Internet (Xu et al. 2012); (b) $\beta = \frac{1}{\gamma i_u(t)}$, where γ is the probability of a susceptible node v being infected by an infected neighbor u, i.e., $a_{uv} = 1$; $i_v(t)$ is the probability that node v is infected at time t; and (c) δ is the rate at which a node becomes susceptible from the infected state (Xu et al. 2011). We assume, for the purpose of analytical tractability and w. l. o. g, that α, β, and γ are uniform for all the nodes in the network. We rationalize this uniformity for another practical reason: in any ICS, the attacker and the defender both improve their strategies over time hand in hand (Furnell and Thomson 2009; Liginlal et al. 2009; Pattinson et al. 2012; Lewis 2003; Thalheimer 2006). Thus, at any point of time we find it rational to assume that the aforementioned parameters are uniform. We also assume that parameters $\alpha, \beta, \gamma, \delta$ are common knowledge to a network administrator. While the knowledge of β, γ, δ can be estimated from internal measurements and observations (and furthermore depends on cyber-protection quality that the administrator deploys) through network telemetry, we admit that it is difficult to accurately estimate α (apart from a naive history-based estimate) - as knowing it would severely constrain the ability of a strategic attacker. The latter point is only relevant if a system manager is trying to evaluate the loss impact post a cyber-attack event - not the goal of our paper where we pre-analyze potential loss impact over an entire parameterized family (w.r.t. $\alpha, \beta, \gamma, \delta$) of cyber-attacks if they were to happen. The state evolution logic adheres to reality that IDS and security-monitoring alarm-triggered incident response are imperfect, and takes time within which cyber-adversaries take network control.

The Underlying Stochastic Process - We first provide an intuition for the general audience of the N-dimensional Markov process X stitched out by the above-mentioned one-dimensional Markov chain, where N is the number of nodes in the ICS network. More specifically, at any time instant t, $X(t)$ is the vector of random variables $X_i(t)$'s - capturing the I/S state of each node in the ICS network. Hence, $X(t)$ is an N-dimensional random function over time, where each random instance of the single-dimensional $X_i(t)$ over t evolves according to one-dimensional Markov chain X_i. Geometrically, at each time instant t, the N-dimensional function represents a random (I/S) configuration of the entire IIoT network. Each such random instance of the N-dimensional function $X(t)$ evolves according to a special N-dimensional Markov chain (contributing to a special stochastic process known as a Feller process (Rogers and Williams 1994)) that is decomposable using the one-dimensional Markov chains X_i.

This intuition can be formally represented as follows. Let $(\Omega, \mathscr{F}, P)$ be a probability space with filtration $\mathbf{F} = (\mathscr{F}_t)_{t \geq 0}$, where $\mathbf{F}$ is right continuous and $\mathscr{F}_0$ contains all the P-null sets (Potter 2004). Ω is the sample space, and P is the probability function mapping each event in F to $[0,1]$. Intuitively, $\mathscr{F}_t$ can be thought of as the family of all N-dimensional random functions $X(t)$ charted out till time t, with $\mathscr{F}_0$ consisting of all functions whose probability measure of occurrence is zero (e.g., a point function in N dimensions, where the points are random initial (I/S) configurations of the N-node ICS network). The process X is a Markov process with state space $E = \{0,1\}^N$ with $X_0 = x \in E$. We assume that X is a Feller process with generator $G : C(E) \to \mathbf{R}$, notated by $G(f(x))_{|f \in C(E), x \in E}$ and expressed as:

$$\sum_{i=1}^{N} \left((1 - x_i)\frac{1}{\alpha} + \beta(1 - x_i) \sum_{j=1}^{N} a_{ij} x_j + x_i \delta \right) (f(x^i) - f(x)),$$

where state transitions $x^i_j = x_j$ for $i \neq j$ and $x^i_i = 1 - x_i$ (from (Fahrenwaldt et al. 2018)). The family $C(E)$ consists of all cadlag functions on E (these are N-dimensional random functions admitting jumps, i.e., discontinuities, in the stochastic process $\{X_t\}$), but form a Skorohod space on which probability measures are always defined (Liggett 2012). The concept of the jump is relevant in the N-dimensional setting because the (I/S) network configuration changes over time may have abrupt changes. The cadlag functions lying on the Skorohod space just enables us to be always able to quantify the probability mass of a collection of random functions. The Markov process X has exponential waiting times between jumps, with an exponential state space of cardinality 2^N. Note that for each $\alpha, \beta, \gamma, \delta$ instance, the Markov-Feller process captures all possible malware spread variations and resulting (I/S) node states in the N-node ICS network. Hence our choice of the stochastic process sets the necessary foundations.

2.3 The Adverse Impact Generating Stochastic Process Model

The *end goal* of an APT is to cause system-wide damage (e.g., revenue loss via business disruption) in the ICS network, after it has stealthily infected (a subset of) nodes in its initial stage, where infection does not imply node/device damage that is left for later. More specifically, at certain times post its infection stage, the APT attack decides to launch damage on certain infected nodes making them dysfunctional (incapable of providing QoS). These times, unknown to a system manager, are denoted by $(T_n)_{n\in\mathbf{N}}$. The corresponding number of cumulative damage launches by these time instants are counted through a stochastic process $M = (M(t))_{t\geq 0}$ (from (Fahrenwaldt et al. 2018)). The size of the negative/adverse impact on infected nodes becoming dysfunctional at any time instant t is modeled by another N-dimensional stochastic process $L = (L(t))_{t\geq 0}$, where $L(t) = (L_1(t),\ldots,L_N(t))^{\mathrm{T}}$, where each dimension represents the adverse impact on a given node in the ICS network. The value $L_i(t)$ for node i is zero if it is either not infected at time t or the APT attack does not launch damage on i at t.

Note that M is a left-continuous counting stochastic process with a deterministic and non-deterministic component (according to the seminal Doob-Meyer decomposition theorem (Brémaud 1981)), and adapted to the filtration $\mathbf{F}$, i.e., random functions spanned upto $M(t)$ is inside the family of functions $\mathbf{F}_t$. One of the components is its stochastic intensity $(\lambda(t))_{t\geq 0}$, where λ is a non-negative $\mathbf{F}$-predictable process (because it is a continuous time-adapted process that is left-continuous). The other component, $M(t\wedge T_n) - \int_0^{t\wedge T_n}\lambda(s)ds$, is the martingale, i.e., the non-deterministic component for all $n\in\mathbf{N}$. The martingale property ensures that the conditional expectation, $M'(t+\delta t) - \mathbf{E}[M'(t+\delta t)]|\mathbf{F}_t$, equals $M'(t)$ for a stochastic process $M'(t)$. In our setting, since M is a counting process, it is a sub-martingale by the Doob-Meyer decomposition theorem. In practical jargon, the sub-martingale property simply and evidently implies that the expected number of launched damages by the APT at time $t+\Delta t$ given time history is greater than the number of launched damages by time t. L is predictable (again due to it being a time-adapted process that is left-continuous) and non-negative. Both M and L are independent from the Markovian infection spread process X.

3 NETWORK AND TIME AGGREGATE ADVERSE CYBER-LOSS IMPACT

We (e.g., cyber-loss managers) are concerned with the expected network aggregate adverse impact of the APT attack over a fixed time window $[0,T]$ *with* $T>0$, where we assume that the tangible impact units are the same for all the nodes.

3.1 A Formal Model of the Expected Space (Network) Time Aggregate Cyber-Loss Impact

In practice, such units can be mapped to the loss in quality of experience (QoE) derived from dysfunctional nodes (e.g., dysfunctional IoT-embedded CNC industrial machines) (Suryanegara et al. 2019; Minovski et al. 2020; Serral-Gracià et al. 2010). Consider a measurable function $f(\cdot;\cdot):\mathbf{R}_+\times\mathbf{R}_+^N\to\mathbf{R}_+$, where the first argument refers to the time, and the second - to the N-dimensional (for N nodes) adverse impact (i.e., FAIR for networks) generated by an APT cyber-attack (ref. (Fahrenwaldt, Weber, and Weske 2018)). The *expected* aggregate impact (adopted from (Fahrenwaldt et al. 2018)) incurred over the period $[0,T]$, given by $\mathbb{E}\Big[\int_0^T f(t;L(t)\circ X(t))dM(t)\Big]$, obey the following equivalent equalities:

$$\mathbb{E}\Big[\int_0^T f(t;L(t)\circ X(t))dM(t)\Big] = \mathbb{E}\Big[\int_0^T f(t;L(t)\circ X(t-))dM(t)\Big],$$

$$\mathbb{E}\Big[\int_0^T f(t;L(t)\circ X(t))dM(t)\Big] = \mathbb{E}\Big[\int_0^T f(t;L(t)\circ X(t-))\lambda(t)dt\Big],$$

$$\mathbb{E}\Big[\int_0^T f(t;L(t)\circ X(t))dM(t)\Big] = \mathbb{E}\Big[\int_0^T f(t;L(t)\circ X(t))\lambda(t)dt\Big].$$

Here, the $\circ$ operator denotes the component-wise Hadamard product of vectors. The first equality is due to the fact that X and M are independent and never jump at the same time with probability 1 (in practical terms an ICS node is not infected and launched an attack upon at the same time). The second equality follows from the $\mathbf{F}$-predictability of the integrand (Brémaud 1981) (because M is left-continuous), and the third equality holds since the paths of X, i.e., the N-dimensional random functions, possess at most countably many jumps in $[0,T]$ and constitute a Lebesgue null set, i.e., a set with a probability measure zero, for each path (because in practice there are countably many malware spread paths and the probability of the individual occurrence of each is zero amongst an infinite continuum of possible paths). *Specific practical forms of $f(\cdot)$ will be discussed in Section 6.* Of equal importance as the mean, are the *tail properties* of the distribution $\int_0^T f(t;L(t)\circ X(t))\lambda(t)dt$ that reflects the statistical spread of the adverse impact (risk) distribution (the probability that network aggregate adverse impacts of an APT attack exceeds a certain percentile). The spread metric will characterize the notion of APT risk that we model in our work with the widely popular Conditional-Value-at-Risk (CVaR) risk measure of $\int_0^T f(t;L(t)\circ X(t))\lambda(t)dt$ from risk theory (McNeil et al. 2015). A detailed analysis of this CVaR metric characterizing APT risk is deferred till Section 4.

One could argue the feasibility of quantifying the adverse impact over time and space if the attack type and parameter space is (partially) unknown. *However, the aim of this paper is not to quantify this impact post the occurrence of an (APT) attack event.* We want to quantify (for cyber-protection budget planning CISOs and CEOs) how much cyber-risk (induced by statistics of the adverse impact) an ICS can potentially be subject to if it were to be affected by a family (characterized by all feasible $\alpha,\beta,\gamma,\delta,t$) of malware-spreading APT cyber-attacks. To complement our contribution rationale, note that it is usual in practice for managers to find it difficult to estimate the hypothetical impact of an adverse cyber-incident (Butler 2002). This difficulty is aggravated via uncertainties in the knowledge of the cyber-risk terrain, system complexity, lack of cyber-incident data and cyber-loss impact metrics, and the inability to predict future cyber-incidents (Jalali et al. 2019; Komljenovic et al. 2016). Add to this is the role of cognitive biases that prevent even the most experienced of system managers to assess the impact of cyber-risk accurately enough (Tversky and Kahneman 1974; Madnick et al. 2016; Straub and Welke 1998). In such environments, it is best that managers consider worst case cyber-risk impact (this analysis needs to consider a family of parameters) into account rather than have a faulty perception of the exact cyber-risk impact.

3.2 A Cyber-Risk Manager's Closed-Form Expression for Aggregate Cyber-Loss

A closed form expression is often a necessary first step for cyber-risk managers to accurately estimate space-time aggregate cyber-loss in any part of the ICS network. To this end, we first propose a framework to derive a closed form expression for the node-aggregate cyber-loss impact, $\int_0^T f(t;L(t)\circ X(t))\lambda(t)dt$, in a given time period for an APT compromised IIoT network. More specifically, we will first propose a method for a cyber-risk manager to accurately approximate any general f in closed form as a polynomial function for analysis tractability (for ease of taking an integral), and follow that up with theoretically bounding the function approximation error.

Assumptions - Though X is a Markov process with cadlag paths, we assume that all $f\in C(E)$ are continuous functions over time, i.e., $f(t;L(t)\circ X(t)$ is continuous. However, *it is a challenging problem to compute aggregate (over time and space) loss values directly on top of general continuous loss functions.* Thankfully, the applicability of the *Stone-Weierstrass* theorem (Brosowski and Deutsch 1981) allows us to work with polynomial loss functions that are outcomes of uniformly approximating arbitrary continuous functions defined over a compact (closed and bounded) set in $\mathbf{R}^n$ - as in our case, with the Hadamard product lying in a compact set. We also assume that time-aggregated loss function f will map into a one-dimensional tangible scalar value, as is usual in practice. For a multi-variate f, we convert its range to a single dimensional output, i.e., $f:\mathbb{R}_+^N\to\mathbb{R}_+$, via the transformation:

$$f(x_1,\ldots,x_N)=g\left(\Lambda(x_1,\ldots,x_N)\right),$$

where we assume $\Lambda:\mathbb{R}_+^N\to\mathbb{R}_+$ to be a linear increasing aggregation function, and $g:\mathbb{R}_+\to\mathbb{R}_+$ to be continuous and increasing. g is also assumed to be bounded on $[0,\|\Lambda(L)\|_\infty)$, where $\|\cdot\|_\infty$ denotes the

L^{∞} norm. An example of Λ is $\Lambda(x_1,\ldots,x_N) = \sum_{i=1}^{N} \alpha_i x_i,\ \alpha_i \geq 0$. g, for example, could be of the form $g\left(\Lambda\left(x_1,\ldots,x_N\right)\right) = \Lambda\left(x_1,\ldots,x_N\right)$.

Polynomial Approximation of a Continuous f - As mentioned above, it can be quite cumbersome for a cyber-risk manager to accurately evaluate any general f (and consequently the aggregate space-time cyber-loss within an ICS network) for a given ICS network. It would be good if any given f could be approximated through a polynomial function that is much amenable to computational and analysis ease. *In view of recent developments in function approximation theory following (Davidson and Donsig 2009), a polynomial closed form approximation to a general f can indeed be constructed via the following steps:*

1. Choose (a) $d \in \mathbb{N}$ - a pre-specified choice for the degree of the polynomial, and (b) a bound $\varepsilon > 0$.
2. Select a constant $u \in \mathbb{R}_+$ (based on prior network and cyber-loss impact knowledge) such that the node-aggregated cyber-loss impact is bounded as per the following relation:

$$\mathbb{P}(\Lambda(L) > u) \leq \varepsilon.$$

3. From the space of all degree-d polynomials, choose the best uniform approximation $p_d(x) := \sum_{\ell=0}^{d} a_\ell x^\ell\ (a_0, a_1, \ldots, a_d \in \mathbb{R})$ (Davidson and Donsig 2009) of g on the compact interval $[0,u]$. The subsequent function approximation error is denoted by $e_d(g)$, and given as

$$e_d(g) =: \max_{x \in [0,u]} |g(x) - p_d(x)| = \|g - p_d\|_{\infty,[0,u]},$$

where the L_∞ norm is used to extract the maximum possible error.

How Much is the Approximation Error for a Cyber-Risk Manager? - We now quantify the 'cost' (error) of the approximation borne by the cyber-risk manager as an outcome of approximating f. The optimal degree-d polynomial approximation of $f(L \circ X)$ obtained from Steps 1-3 above is given by

$$\bar{f}_d(L \circ X) := \begin{cases} p_d(\Lambda(L \circ X)), & \text{if } \Lambda(L) \leq u \\ 0, & \text{if } \Lambda(L) > u \end{cases}$$

We then have the following result following developments in (Davidson and Donsig 2009; Fahrenwaldt et al. 2018) in relation to the approximation error induced by the polynomial $\bar{f}_d$.

Theorem 1 *The function approximation error for $f = f(x_1,\ldots,x_N)$ (obtained via the polynomial approximation method in (Davidson and Donsig 2009; Fahrenwaldt et al. 2018)) incurred by the cyber-risk manager is defined through the following inequation $\left\|f(Z) - \bar{f}_d(Z)\right\|_{L_1} \leq e_d(g) + m \cdot \varepsilon$, where $Z = L \circ X$, and m is a real number that satisfies*

$$|f(l \circ X)| = |g(\Lambda(L \circ X))| \leq |g(\Lambda(L))| \leq m,$$

for all possible realizations of L, where the L_1 prevents amplifying outlier effects during approximation.

Proof Sketch - The theorem is easily validated through algebraic simplification on the following inequality:

$$\|\left[f(Z) - p_d(\Lambda(Z))\right] \cdot \nVdash_{[0,u]}(\Lambda(L)) + f(Z) \cdot \nVdash_{(u,\infty)}(\Lambda(L))\|_{L_1} \leq \mathbb{E}\left[|g(\Lambda(Z)) - p_d(\Lambda(Z))| \cdot \nVdash_{[0,u]}(\Lambda(L))\right].$$

Implications for Cyber-Risk Management - The theorem implies that there always exists a closed form expression for approximate tangible space-time aggregate cyber-loss impact post an APT cyber-attack event on an ICS network guaranteed through polynomial approximation mathematics. The L_1 norm of the function approximation error induced by $\bar{f}_d$ is (a) tightly bounded from above by $e_d(g) + m \cdot \varepsilon$, and (b) is very low for small enough ε. From the viewpoint of cyber-risk management, the theorem results are useful as they simply denote that it is within the scope of the ICS cyber-risk manager to influence the quality of its approximate for aggregate space-time cyber-loss by controlling its (enterprise budget driven) risk appetite (via ε). *Once the cyber-risk manager gets hold of a parametric (approximate) expression for space-time aggregate cyber-loss, it could apriori estimate (compute) the mean value of such an expression over a statistical distribution of the parameters.* This loss amount reflects the impact of business discontinuity.

4 ANALYZING TAIL APT RISK IN AN ICS NETWORK

Thus far, we derived analytical and tractable closed-form solutions to the first moment, i.e., expectation, of the node-aggregate cyber-loss distribution in an ICS network. While this metric is useful to a cyber-risk manager, in practice they are interested in both expected cyber-loss estimates, as well as *in the knowledge of the tail, i.e., loss spread, of a cyber-loss distribution*. An industry standard cyber-risk metric to measure this tail is the Conditional-Value-at-Risk (CVaR) metric. *In our paper, we synonymously term this metric as the APT risk, and provide accuracy guarantees of empirically estimating the APT risk, when compared to the theoretical ground-truth.* To do so, we propose a rigorous framework based on the theory of large deviations (TLD) that (a) first uses *concentration inequalities* (Boucheron et al. 2013) to analyse the deviations of empirical estimates of the APT risk from the empirical mean of the cyber-loss impact, and (b) subsequently provides upper and lower bounds of the deviations of the empirical estimates of the APT risk from theoretical ground truth. Concentration inequalities in the TLD deal with deviations of functions of independent random variables from their expectation.

4.1 Background for Analyzing Empirical CVaR

Here, we provide a background on empirically analysing the sample-driven CVaR metric of any risk distribution simply because this is what a cyber-risk manager will only have access to. The CVaR at level $\alpha \in (0,1]$ of a random variable X is

$$\mathrm{CVaR}_\alpha(X) \triangleq \inf_v \left\{ v + \frac{1}{\alpha}\mathbb{E}\left[(X-v)^+\right] \right\}.$$

It is well known from (Rockafellar et al. 2000) that, when X has a continuous distribution, that $\mathrm{CVaR}_\alpha(X) = \mathbb{E}\left[X \mid X \geqslant \mathrm{VaR}_\alpha(X)\right]$, where

$$\mathrm{VaR}_\alpha(X) \triangleq \sup_x \{x \mid \mathbb{P}(X \geqslant x) \geqslant \alpha\}$$

is the α-quantile (or VaR) of X. While this relation does not necessarily hold if X has a discontinuous distribution, *we can nonetheless roughly interpret* $\mathrm{CVaR}_\alpha(X)$ as the expected loss over the $\alpha\%$ worst cases. We now make the following assumption necessary for deriving tight bounds of finite sample CVaR on aggregate network cyber-loss impact.

Assumption 1. The random variable X satisfies $\mathrm{supp}(X) \subseteq [0,U]$, and its samples $X_1,\ldots,X_n$ are independent.

Suppose one were to map the outcome of a stochastic process modeling the aggregate cyber-loss over time and space in an ICS network, the values obtained on multiple sample paths of the stochastic process would be independent. This would closely map empirical estimates of such aggregate loss estimates from the real network at different points in time - thereby justifying the ‘independence’ aspect of the assumption. An upper bound criterion (as stated in the assumption) is common to the application of many concentration inequalities, and without loss of generality we use supp $(X) \subseteq [0,U]$, to reflect the fact that the minimum loss impact value is zero.

For the case of CVaR_α, we denote the simple empirical CVaR estimator by $\widehat{\mathrm{CVaR}}_\alpha$, and express it as:

$$\widehat{\mathrm{CVaR}}_\alpha\left(X_1,\ldots,X_n\right) \triangleq \inf_v \left\{ v + \frac{1}{n\alpha}\sum_{i=1}^{n}\left(X_i - v\right)^+ \right\}, \tag{2}$$

where $(y)^+ = \max(y,0)$. This empirical estimator is an intuitive and typically popular one (Rockafellar, Uryasev, et al. 2000) based on the method of moments. Such estimators are *efficiently computed* for large n and most convex loss functions - the CVaR function, CVaR_α, being one having a piece-wise linear loss function (Ben-Tal and Teboulle 2007). *In the analysis to follow, we denote $X = X_T$ to the random variable instantiating the node-aggregate cyber-loss upto a pre-specified time period T in an ICS network.*

4.2 The Lower Bound of (Empirical APT Risk - True APT Risk)

We have the following result stating the tight lower bound of tail space-time aggregate APT cyber-risk.

Theorem 2 *Consider a cyber-risk manager having access to* $X = (X_1, \ldots, X_n)$ *- a vector of n samples from the true node and time aggregate cyber-loss impact distribution in an ICS network. Then*

$$\mathbb{P}\left(\widehat{\mathrm{CVaR}}_{\alpha}(X) \geqslant \mathrm{CVaR}_{\alpha}(X) + \varepsilon\right) \leqslant \mathrm{e}^{-2\frac{\varepsilon^2\alpha^2}{U^2}\cdot n},$$

where $\mathrm{CVaR}_{\alpha}(X)$ *equals* $\mathrm{APTRisk}_{\alpha}(X)$*, and* $\widehat{\mathrm{CVaR}}_{\alpha}(X)$ *equals* $\widehat{\mathrm{APTRisk}}_{\alpha}(X)$.

Proof Sketch - We use the celebrated McDiarmid's concentration inequality (McDiarmid 1989) that says we need on the order of $n_H \triangleq (\sum_{i=1}^{n} c_i^2/\varepsilon^2)\log(1/\delta)$ samples to estimate the sample mean within a precision of ε with probability at least $1-\delta$. This combined with the fact that $\mathbb{E}\left[\widehat{\mathrm{CVaR}}_{\alpha}(X_1, \ldots, X_n))\right] \leqslant \mathrm{CVaR}_{\alpha}(X_1, \ldots, X_n)$ proves the theorem.

Implications to Cyber-Risk Management - The result provides a closed form expression of the tight lower bound of the deviation in empirical APT risk (measured as empirical CVaR) with the theoretical true value of the APT risk, as a function of sample count. *In practice, a higher sample count reduces the deviation*, and hence a cyber-risk manager (e.g., enterprise cyber-risk officer, cyber-insurer) should get access to sufficient number of samples to reduce tail cyber-risk estimation uncertainty.

4.3 The Upper Bound of (Empirical APT Risk - True APT Risk)

We have the following result stating the tight upper bound of tail space-time aggregate APT cyber-risk.

Theorem 3 *Consider a cyber-risk manager having access to* $X = (X_1, \ldots, X_n)$ *be a vector of n samples from the true node and time aggregate cyber-loss impact distribution in an ICS network. Then for any* $\varepsilon \leqslant 0$,

$$\mathbb{P}\left(\widehat{\mathrm{CVaR}}_{\alpha}(X) \leqslant \mathrm{CVaR}_{\alpha}(X) - \varepsilon\right) \leqslant 3\mathrm{e}^{(-\frac{1}{5})\alpha(\frac{\varepsilon}{U})^2\cdot n},$$

where $\mathrm{CVaR}_{\alpha}(X)$ *equals* $\mathrm{APTRisk}_{\alpha}(X)$*, and* $\widehat{\mathrm{CVaR}}_{\alpha}(X)$ *equals* $\widehat{\mathrm{APTRisk}}_{\alpha}(X)$.

Proof Sketch - We use the celebrated Hoeffding's inequality (Hoeffding 1994) that says we need on the order of $n_H \triangleq (U/\varepsilon)^2\log(1/\delta)$ samples to estimate the sample mean within a precision of ε with probability at least $1-\delta$. This combined with the fact that $\widehat{\mathrm{CVaR}}_{\alpha}(X_1, \ldots, X_n) \geqslant \frac{1}{n\alpha}\sum_{i=1}^{\lfloor n\alpha \rfloor} X_{(i)}$ holds, where $X_{(i)}$ are the decreasing order statistics of X_i, proves the theorem.

Implications to Cyber-Risk Management - The result provides a closed form expression of the upper bound of the deviation in empirical APT risk (measured as empirical CVaR) with the theoretical true value of the APT risk, as a function of sample count. *In practice, a higher sample count reduces the deviation*, and hence a cyber-risk manager should ideally get access to sufficient number of samples to reduce tail cyber-risk estimation uncertainty.

5 SUMMARY

In this paper, we proposed the first theoretical framework for ICS enterprise managers to accurately estimate *apriori* and tightly bound APT risk in general IoT driven ICS networks for a parametric family of stealthy spread-based APT cyber-attacks. The importance of this task lies in planning a budget for enterprise cyber-resilience management activities before the occurrence of cyber-attacks. We first modeled the time-varying *attack-defense-impact* trio pertaining to our threat model as a continuous Markov process. Subsequently, we provided a closed form expression for the node and time aggregate cyber-loss impact in an ICS network due to a spread-based APT cyber-attack. Finally, we proposed a rigorous analysis motivated by concentration inequalities from probability theory, to derive tight non-asymptotic bounds of the difference between the empirical estimation of the APT risk measure and the true (ground truth) value of the APT risk.

ACKNOWLEDGMENTS

This work is sponsored by Cybersecurity at MIT Sloan (CAMS) at the Massachusetts Institute of Technology.

REFERENCES

Anderson, R., and T. Moore. 2009. "Information Security: Where Computer Science, Economics and Psychology Meet". *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 367:2717–2727.

Bailey, N. T. 1975. *The Mathematical Theory of Infectious Diseases and its Applications*. Charles Griffin & Company Ltd, 5a Crendon Street, High Wycombe, Bucks HP13 6LE.

Ben-Tal, A., and M. Teboulle. 2007. "An Old-New Concept of Convex Risk Measures: The Optimized Certainty Equivalent". *Mathematical Finance* 17(3):449–476.

Bezuidenhout, C., and G. Grimmett. 1990. "The Critical Contact Process Dies Out". *The Annals of Probability*:1462–1482.

Boucheron, S., G. Lugosi, and P. Massart. 2013. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press.

Brémaud, P. 1981. *Point Processes and Queues: Martingale Dynamics*, Volume 50. Springer.

Brosowski, B., and F. Deutsch. 1981. "An Elementary Proof of the Stone-Weierstrass Theorem". *Proceedings of the American Mathematical Society*:89–92.

Butler, S. A. 2002. "Security Attribute Evaluation Method: A Cost-Benefit Approach". In *Proceedings of the 24th International Conference on Software Engineering*, 232–240.

Davidson, K. R., and A. P. Donsig. 2009. *Real Analysis and Applications: Theory In Practice*. Springer Science & Business Media.

Durrett, R., and X.-F. Liu. 1988. "The Contact Process on a Finite Set". *The Annals of Probability*:1158–1173.

Fahrenwaldt, M. A., S. Weber, and K. Weske. 2018. "Pricing of Cyber Insurance Contracts in a Network Model". *ASTIN Bulletin: The Journal of the IAA* 48(3):1175–1218.

Freund, J., and J. Jones. 2014. *Measuring and Managing Information Risk: A FAIR Approach*. Butterworth-Heinemann.

Furnell, S., and K.-L. Thomson. 2009. "Recognising and Addressing 'Security Fatigue"'. *Computer Fraud & Security* 2009:7–11.

Hoeffding, W. 1994. "Probability Inequalities for Sums of Bounded Random Variables". In *The Collected Works of Wassily Hoeffding*, 409–426. Springer.

Jalali, M. S., M. Siegel, and S. Madnick. 2019. "Decision-Making and Biases in Cybersecurity Capability Development: Evidence from a Simulation Game Experiment". *The Journal of Strategic Information Systems* 28(1):66–82.

Komljenovic, D., M. Gaha, G. Abdul-Nour, C. Langheit, and M. Bourgeois. 2016. "Risks of Extreme and Rare Events in Asset Management". *Safety Science* 88:129–145.

Kure, H. I., and S. Islam. 2019. "Assets Focus Risk Management Framework for Critical Infrastructure Cybersecurity Risk Management". *IET Cyber-Physical Systems: Theory & Applications* 4(4):332–340.

Lewis, J. 2003. "Cyber Terror: Missing in Action". *Knowledge, Technology & Policy* 16(2):34–41.

Liggett, T. M. 2012. *Interacting Particle Systems*, Volume 276. Springer Science & Business Media.

Liggett, T. M. 2013. *Stochastic Interacting Systems: Contact, Voter and Exclusion Processes*, Volume 324. Springer Science & Business Media.

Liginlal, D., I. Sim, and L. Khansa. 2009. "How Significant is Human Error as a Cause of Privacy Breaches? An Empirical Study and A Framework for Error Management". *Computers & Security* 28(3-4):215–228.

Madnick, S., M. S. Jalali, M. Siegel, Y. Lee, D. Strong, R. Wang, W. H. Ang, V. Deng, and D. Mistree. 2016. "Measuring Stakeholders' Perceptions of Cybersecurity for Renewable Energy Systems". In *International Workshop on Data Analytics for Renewable Energy Integration*, 67–77. Springer.

Martinez-Moyano, I. J., R. Oliva, D. Morrison, and D. Sallach. 2015. "Modeling Adversarial Dynamics". In *2015 Winter Simulation Conference (WSC)*, edited by L. Yilmaz, W. K. V. Chan, I. Moon, T. M. K. Roeder, C. Macal, and M. D. Rossetti, 2412–2423. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

McDiarmid, C. 1989. "On the Method of Bounded Differences". *Surveys in Combinatorics* 141(1):148–188.

McNeil, A. J., R. Frey, and P. Embrechts. 2015. *Quantitative Risk Management: Concepts, Techniques and Tools-Revised Edition*. Princeton University Press.

Minkel, J. 2008. "The 2003 Northeast Blackout–Five Years Later". *Scientific American* 13.

Minovski, D., C. Åhlund, K. Mitra, and R. Zhohov. 2020. "Defining Quality of Experience for the Internet of Things". *IT Professional* 22(5):62–70.

Mountford, T., J.-C. Mourrat, D. Valesin, and Q. Yao. 2016. "Exponential Extinction Time of the Contact Process on Finite Graphs". *Stochastic Processes and their Applications* 126(7):1974–2013.

Pastor-Satorras, R., C. Castellano, P. Van Mieghem, and A. Vespignani. 2015. "Epidemic Processes in Complex Networks". *Reviews of Modern Physics* 87(3):925.

Pattinson, M., C. Jerram, K. Parsons, A. McCormac, and M. Butavicius. 2012. “Why do Some People Manage Phishing E-Mails Better Than Others?”. *Information Management & Computer Security*.

Pfleeger, S., and R. Cunningham. 2010. “Why Measuring Security is Hard”. *IEEE Security & Privacy* 8(4):46–54.

Potter, P. 2004. “Stochastic Integration and Differential Equation”. *Stochastic Modeling and Applied Probability* 21.

Rockafellar, R. T., S. Uryasev et al. 2000. “Optimization of Conditional Value-at-Risk”. *Journal of Risk* 2:21–42.

Rogers, L. C. G., and D. Williams. 1994. “Diffusions, Markov Processes and Martingales, Volume 1: Foundations”. *John Wiley & Sons, Ltd., Chichester* 7.

Sepúlveda Estay, D. A. 2021. “A System Dynamics, Epidemiological Approach for High-Level Cyber-Resilience to Zero-Day Vulnerabilities”. *Journal of Simulation*:1–16.

Serral-Gracià, R., E. Cerqueira, M. Curado, M. Yannuzzi, E. Monteiro, and X. Masip-Bruin. 2010. “An Overview of Quality of Experience Measurement Challenges for Video Applications in IP Networks”. In *International Conference on Wired/Wireless Internet Communications*, 252–263. Springer.

Straub, D. W., and R. J. Welke. 1998. “Coping With Systems Risk: Security Planning Models for Management Decision Making”. *MIS Quarterly*:441–469.

Suryanegara, M., D. A. Prasetyo, F. Andriyanto, and N. Hayati. 2019. “A 5-Step Framework for Measuring the Quality of Experience (QoE) of Internet of Things (IoT) Services”. *IEEE Access* 7:175779–175792.

Thalheimer, W. 2006. “Spacing Learning Events Over Time: What the Research Says”. *Retrieved March* 21:2007.

Tversky, A., and D. Kahneman. 1974. “Judgment Under Uncertainty: Heuristics and Biases: Biases in Judgments Reveal Some Heuristics of Thinking Under Uncertainty.”. *American Association for the Advancement of Science* 185(4157):1124–1131.

Van Mieghem, P., J. Omic, and R. Kooij. 2008. “Virus Spread in Networks”. *IEEE/ACM Transactions On Networking* 17(1):1–14.

Van Mieghem, P., and R. van de Bovenkamp. 2015. “Accuracy Criterion for the Mean-Field Approximation in Susceptible-Infected-Susceptible Epidemics on Networks”. *Physical Review E* 91(3):032812.

Wu, Y., P. Li, L.-X. Yang, X. Yang, and Y. Y. Tang. 2017. “A Theoretical Method for Assessing Disruptive Computer Viruses”. *Physica A: Statistical Mechanics and its Applications* 482:325–336.

Xenofontos, C., I. Zografopoulos, C. Konstantinou, A. Jolfaei, M. K. Khan, and K.-K. R. Choo. 2021. “Consumer, Commercial and Industrial IoT (in) Security: Attack Taxonomy and Case Studies”. *IEEE Internet of Things Journal*.

Xu, S., W. Lu, and L. Xu. 2012. “Push-and Pull-Based Epidemic Spreading in Networks: Thresholds and Deeper Insights”. *ACM Transactions on Autonomous and Adaptive Systems (TAAS)* 7(3):1–26.

Xu, S., W. Lu, L. Xu, and Z. Zhan. 2014. “Adaptive Epidemic Dynamics in Networks: Thresholds and Control”. *ACM Transactions on Autonomous and Adaptive Systems (TAAS)* 8(4):1–19.

Xu, S., W. Lu, and Z. Zhan. 2011. “A Stochastic Model of Multivirus Dynamics”. *IEEE Transactions on Dependable and Secure Computing* 9(1):30–45.

Yang, L., M. Draief, and X. Yang. 2017. “Heterogeneous Virus Propagation in Networks: A Theoretical Study”. *Mathematical Methods in the Applied Sciences* 40(5):1396–1413.

Yang, L.-X., M. Draief, and X. Yang. 2015. “The Impact of the Network Topology on the Viral Prevalence: A Node-Based Approach”. *Public Library of Science* 10(7):e0134507.

Yang, L.-X., P. Li, X. Yang, and Y. Y. Tang. 2017. “Distributed Interaction Between Computer Virus and Patch: A Modeling Study”. *arXiv Pre-Print arXiv:1705.04818*.

Zeijlemaker, S., J. Uriega, and G. Pasaoglu Kilanc. 2018. “Malware Dynamics: How to Develop a Successful Anti-Malware Defense Reference Architecture Policy”. In *Proceedings of the 36th International Conference of the System Dynamics Society, Reykjavik, Iceland*.

AUTHOR BIOGRAPHIES

RANJAN PAL is a Research Scientist with the MIT Sloan School of Management, USA. His primary research interest lies in developing interdisciplinary cyber risk/resilience management solutions. His email address is ranjanp@mit.edu.

SANDER ZEIJLEMAKER is a Research Affiliate with the MIT Sloan School of Management, USA. His primary research interest lies in developing cyber risk governance solutions based upon system dynamics. His email is szeijl@mit.edu.

MICHAEL SIEGEL is a Principal Research Scientist with the MIT Sloan School of Management, USA. His primary research interest lies in developing interdisciplinary cyber risk/resilience management solutions. His email is msiegel@mit.edu.

ROHAN XAVIER SEQUEIRA is a PhD student and Annenberg Fellow in the department of Electrical and Computer Engineering at the University of Southern California, USA. He has a MS in ECE from the University of Michigan, Ann Arbor. His research interest lies is cyber-risk management and distributed systems. His email address is rsequeir@usc.edu.