Proceedings of the 2024 Winter Simulation Conference
H. Lam, E. Azar, D. Batur, S. Gao, W. Xie, S. R. Hunter, and M. D. Rossetti, eds.

HOW HARD IS IT TO ESTIMATE SYSTEMIC ENTERPRISE CYBER RISK?

Ranjan Pal[1] , Rohan Xavier Sequeira[2] , and Sander Zeijlmaker[1]

[1]MIT Sloan School of Management, Massachusetts Institute of Technology, Cambridge, MA, USA
[2]Electrical and Computer Engineering, University of Southern California, Los Angeles, CA, USA

ABSTRACT

Systemic enterprise cyber risk typically arises when a single (software) vulnerability common across many enterprise computing devices across the globe is exploited by adversaries, and results in catastrophic aggregate cyber-loss consequences to be borne by CRM entities. Examples of such vulnerability exploitation incidents include the *Log4j* and *SolarWinds* cyber-attacks. The important question we ask here is: *how hard is it to discover these 'single vulnerabilities' in enterprise information systems?* We prove that answering this question is NP-hard. Alternatively, leave alone humans, even a computer via cyber-attack simulations might not (in the worst case) discover in finite time such vulnerabilities. Consequently, CRM entities can only expect and prepare for an inevitable catastrophic systemic cyber-incident in time rather than predict the likelihood of one. Likewise, we propose the policy implications of our research for the CRM market stakeholders and elucidate relevant action items for effective systemic enterprise CRM.

1 INTRODUCTION

We are in the pervasively digital age where societal service sectors are getting increasingly driven by IT and IoT technology. Consequently, with this rapidly growing cyber-terrain the likelihood of cyber-incidents (henceforth synonymous with cyber-attacks in this paper) increases manifold. A salient characteristic of ecosystems carved out by societal service sectors is that the sectors are interdependent upon each other. As an example, industries in the natural gas, electricity, healthcare, air transport, and postal/courier sectors rely upon one another for business continuity (see Figure 1). A cyber-incident directly impacting business continuity (BC) of enterprises in a given sector can simultaneously affect the BC in other sectors. To drive home this point, let us briefly work through *four* real-world examples of recent such cyber-attacks.

1.1 Examples of Potent Real World Cyber-Attacks on Societal Service Sectors

A single software vulnerability in the *Log4j* cyber-attack (launched by Chinese cybersecurity researchers in November, 2021 via exploiting *Log4Shell* as a zero-day vulnerability in the low-profile Log4j software utility embedded in billions of enterprise devices) could have caused widening ripples with catastrophic societal consequences. The Log4j vulnerability that affected approximately 40% of global enterprise devices allows attackers to remotely control and execute code on vulnerable machines after which they could cause adverse enterprise impact ranging from minimal to lack of business continuity for days if not weeks.

The *SolarWinds* cyber-incident of 2020 involved Russian hackers insert (in March 2020) a single malicious trojan code (*Sunburst*) into the software update of *Orion* - an IT performance management system widely used by private and government enterprises worldwide, to gain access to confidential business workflow information. More than 18,000 enterprises had applied the Sunburst update globally and this resulted in an accumulative and irreparable damage worth billions of US dollars.

The *Colonial Pipeline* cyber-attack was a ransomware cyber-incident launched by *DarskSide* in May 2021 that forced the company (the largest US pipeline system stretching to about 5500 miles on the East Coast) to close down business operations and freeze their IT systems. The attack vector has never been made public but it is certain that it is a single 'element' - either an unpatched vulnerability, or an

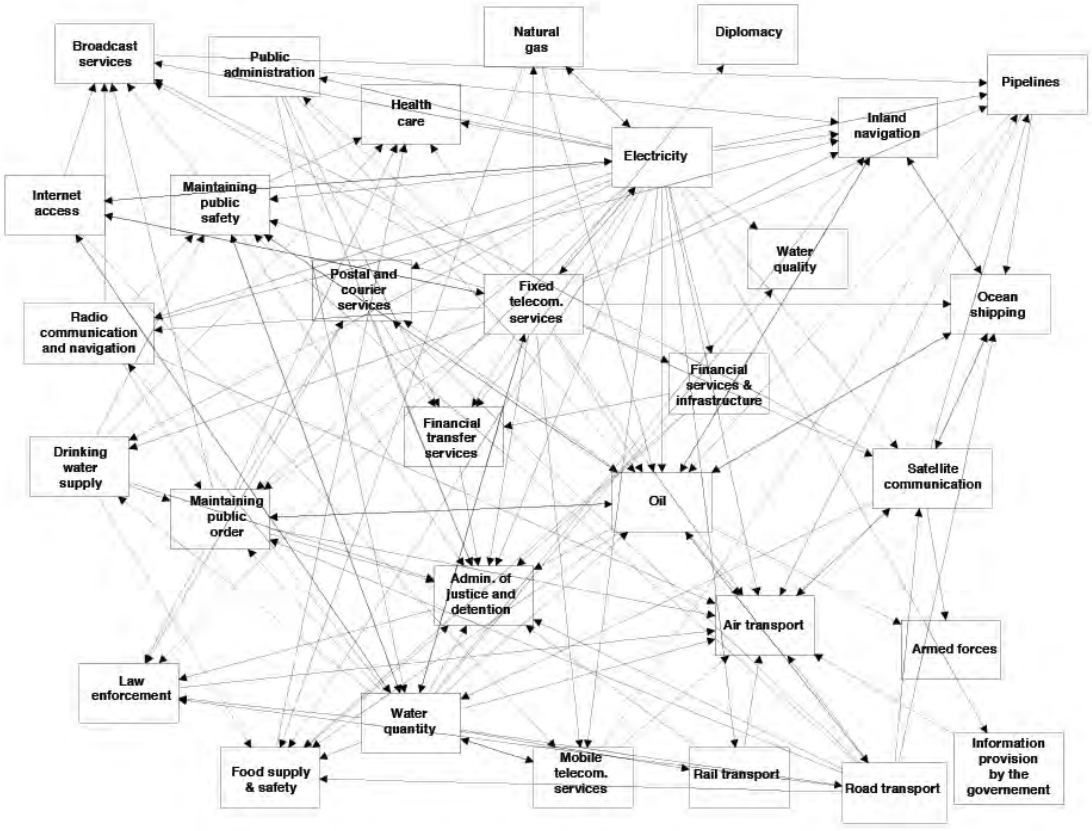

Figure 1: Showcasing the chart of complex service dependencies among networked enterprise sectors with critical cyber infrastructure. [*Source:* Netherlands Organization for Applied Scientific Research (TNO)].

employee getting phished, or an insider leaking access credentials to hackers. The business disruption caused consumer panic and many gas stations ran low or out of fuel for about a week hampering the functioning of multiple segments of the transportation industry (including American Airlines that ran short on fuel) Colonial Pipeline had to pay out USD 4.4M as ransom payment to resume services.

NotPetya was one of the most potent cyber-attacks (launched by the Russian hacking group *Sandworm* in June 2017) in history arising out of cyber-warfare between Russia and Ukraine. NotPetya was a modified version of Petya, that took advantage of two known exploits in older Windows versions: *EternalBlue* and *Mimikatz* to embed malware into a Kiev-based tax filing software application used by multiple enterprises. The specialty of NotPetya was that it did not resort to user action as is the case in social engineering attacks, and was fast in its task of rapidly traveling from system to system accessing admin credentials. Banks in Ukraine and their transit hubs went non-operational within a span of a minute. On a global scale, business activities of (a) the global shipping giant Maersk, (b) pharmaceutical giant Merck, and (c) Cadbury manufacturer Mondelez and (d) courier giant TNT Express, resulted in shutdowns and took days to resume normal functionality. The total estimated economic impact of the NotPetya cyber-attack has been estimated to be atleast USD 10 billion.

1.2 The Nature, Structure, and Agency View of Systemic Cyber Risks

Cyber-attacks such as the ones mentioned above fall in the category of systemic cyber risk. These risks, a subset of the broader category of systemic risk usually arise when a single cyber-trigger event such as injection of malicious code inside a computer network or a tiny set of vulnerabilities (be it technology or

human driven) when exploited by adversaries, result in widespread failures (of the type of business disruption or enterprise security compromise) spanning corporate organizations (enterprises), critical infrastructure, and/or nations. This is the basic nature of systemic cyber risk.

In terms of structure, systemic cyber risks broadly take two forms - a chain reaction (a vertical failure) and contagion effect (horizontal failure) (Forscey et al. 2022). As examples of each, hackers can take down a single Internet exchange point that will prevent it from providing critical services to multiple end-users and businesses at the same time in a given geographical locality triggering a cascading chain reaction effect of global business disruptions reliant on these local businesses. Likewise, a malicious vulnerability exploit in an operating system (OS) affecting multiple computers (including those in enterprises using the software) across the globe via a contagion effect can adversely impact the core operations for all users using the OS. Take note that a contagion effect can lead to a chain reaction.

Finally, what is more important is the size of the systemic cyber risk terrain as viewed by agencies managing and governing systemic cyber risk. In this regard, there have been multiple (overlapping) notions of systemic cyber risk put forward by agencies such as Cybersecurity and Infrastructure Security Agency (CISA), European Systemic Risk Board (ESRB), and big multinational insurance companies (Forscey et al. 2022). While the CISA focuses more on systemic risks on critical infrastructure that span within a nation, ESRB complements CISA's notion on systemic risk to include national and international economies that are formed from enterprise business activities. Whereas, multinational insurance agencies (like Munich Re) characterize systemic risks to be those that are non-insurable due to accumulative effects of correlated individual cyber risk spread across clients in different societal service sectors and geographical regions.

1.3 The Causal Factors of Systemic Cyber Risk

One has to understand that a cyber-network of interdependent businesses across societal service sectors is large and significantly more densely connected than many other societal networks, which amplifies systemic risk. In a broad sense, there are five main causal factors of systemic cyber risk (Forscey et al. 2022).

1. *Risk Concentration*, wherein common technologies such as software, OSs; and third-party service providers create *"common vulnerabilities, data bottlenecks and single points of system failure."* that amplify systemic cyber risk and its consequential social and economic impacts.
2. *Complexity* wherein intricate mutual dependencies on technical, contractual, and financial dimensions of an interdependent societal network leads to poor visibility of single points of failure.
3. *Opacity* wherein there is a lack of transparency between managers of enterprise information systems and the enterprise clients on component specifics of privately owned and operated technology systems.
4. *Scale*, wherein large scale attacks can be launched by adversaries at a very low marginal cost due to the fast spreading pathways in dense interconnected networks, and
5. *Intelligent Adversaries*, wherein the bad actors high in computational, cognitive, and motivational, and technical resource availability (when compared to some enterprise cyber risk management teams) can discover simple yet undisclosed vulnerabilities that can cause a chain reaction of cyber-breach and/or business disruption events in an interdependent societal network.

Subsequently, *it is imperative to focus research on discovering vulnerabilities contributing to systemic cyber risk.* To drive home the importance of this point, in the World Economic Forum Annual Meeting 2023 in Davos, Over 93% of cybersecurity experts and 86% of business leaders (surveyed across 300 experts) opined that *"a far-reaching, catastrophic cyber event is likely in the next two years"*.

1.4 Estimating Systemic Cyber Risk: Data Collection and Simulation Challenges

The most important step to estimate systemic cyber risk involves detecting causal elements (i.e., cyber vulnerabilities) that can generate systemic cyber risk. There are two broad approaches to this detection task: data collection, and penetration testing simulations.

The approach of *data collection* is loaded with obstacles. Many enterprises do not have a software bill of materials (SBOMs) and/or hardware inventories for security managers to locate systemic-level points of failure. This is the intra-organizational opacity between enterprise information system managers and the software/hardware driven information systems used within the enterprise. Even if one (e.g., cyber-insurance companies) considers collecting information on the most important software/hardware components within an enterprise's information systems, the information set obtained from managers often misses out on some components generated through external network scans - such is the complexity of the dependencies between software and hardware components. Add to this, enterprises make a conscious effort not to report dependencies between information system components, vulnerabilities (unless not mandated by law that is prevalent in very few countries around the world) to cyber-insurers or in public in the fear of losing competitive advantage or inviting legal action, and also to remain evasive of exploitative malicious actors.

The approach of *penetration testing* involves a cyber-attack simulation targeted at discovering and checking potential vulnerabilities so that they could be patched before real-life adversaries can exploit them. This approach includes ethical hacking activities. The target range of penetration testing activities span attempts to breach any number of endpoints or applications, from application protocol interfaces (APIs) to backend servers of an enterprise. The types of penetration testing include web application testing, network security testing, social engineering testing, and cloud security testing. Typical market products include the *Powershell Suite*, *Wireshark*, *Metasploit*, *MobSF*, and *Apktool*. However the penetration testing (pen-testing) approach is not without its share of considerable drawbacks. Pen-testing is inconsistent in the sense that it is performed by experts whose skill sets and strengths vary substantially. Pen-testing is a costly and resource intensive effort that is currently deployed by enterprises periodically but only for a constrained time duration. However, the cyber-threat environment and enterprise security posture is very dynamic that is likely to leave important undetected vulnerabilities from the pen-testing efforts. Alternative approaches such as deploying vulnerability scanning systems might alleviate this drawback but these systems don't incorporate context, and might result in noisy output that includes lot of issues that are low-risk and should not be given resolution priority and few issues that are high-risk and should necessarily be given resolution priority but those that escape the attention of the enterprise management.

It is evident that identifying vulnerabilities that cause systemic cyber risk, either via modern data collection or by pen-testing simulations and vulnerability scanning approaches are imperfect to result in zero chances of a systemic cyber risk event.

Goal - Our goal in this paper is to study the fundamental question: *how difficult (as a formal guarantee) is it to computationally find/search vulnerabilities that result in negligible or low chances of systemic cyber risk?* Note that computational approaches supersede manual approaches in terms of search effectiveness. Hence, a computationally difficult task would directly imply manual difficulty but not vice-versa.

1.5 Research Contributions

We make the following research contribution in this paper.

- We first prove using algebraic graph theory that for a set of enterprises represented as nodes connected through an inter-dependent service network (graph), and each node having a diverse portfolio of service liabilities spread across its neighbors to whom it provides service, even a small decrease to the portfolio (due to security mishaps post vulnerability exploitation in the enterprise's information systems) on any network edge (i.e., the liability guarantees are affected due to a cyber-attack on the enterprise node) could significantly impact the service reliability in the network (see Section 3).
- We subsequently prove that, even in the most ideal setting when a cyber risk manager (e.g., a cyber (re-)insurer) has complete knowledge of every client node's service liability portfolio across graph neighbors and client security investment portfolios for assets inside their enterprises, it is computationally intractable, i.e., NP-hard, to find out the number of enterprise nodes that are systemically affected post a cyber-breach event - *a necessary pre-requisite to optimal systemic*

cyber risk estimation and management at both, an individual enterprise (node), and the ecosystem (network) level. In other words, leave alone humans, even the most advanced vulnerability detection simulator software will not be able to detect with 100% accuracy in a finite amount of time, all vulnerabilities that might result in a systemic cyber-incident (unless P = NP). Here, the notion of an enterprise being *systemically affected* implies it being impacted indirectly by a direct cyber-attack on some upstream enterprise (in the graph), on the lines of compromise by malicious entities and/or business disruption to provide a given service (or providing service at a significantly low QoS) to any of its inter-dependent clients downstream. We prove the NP-hardness property by reducing the *Balanced Complete Bipartite Subgraph* (BCBS) problem - already known from theoretical computer science literature to be a NP-hard problem (Garey and Johnson 1979), to our task (see Section 4).

The established computational hardness of optimal systemic cyber risk assessment (and subsequent estimation) task does not in anyway undermine the existing (steady) success, or periodic volatility (see Section 4 for details) of cyber-insurance markets. It just establishes, in a formal and foundational manner for the first time, the inevitable limitations of preventing a systemic cyber-attack in society. Throughout the rest of the paper, we use the term 'dysfunctional' and 'systemically affected' interchangeably for a network node that is adversely impacted via a systemic cyber-attack event.

2 RELATED WORK

We review related work in this section in a concise and brief manner in the interest of space.

First and foremost, *ours is the first work of its kind to investigate into the computational tractability aspects of optimal systemic cyber risk estimation in an inter-dependent service network post a cyber-breach event - a necessary pre-requisite for effective CRM.* In recent related efforts Pal et al. (2021), Pal et al. (2023), the authors show the optimal cyber risk diversification problem (a related variant of the optimal systemic cyber risk estimation problem as optimal risk diversification relies upon optimal systemic cyber risk estimation as a problem instance) in service networks to be NP-hard for residual cyber risk managers. The diversification problem assumes that accurate knowledge of security investment portfolios of organizations is an information asymmetry (IA) challenge, and show that it is NP-hard to design optimal cyber (re-)insurance contracts under the IA challenge. In contrast, we in this paper show that even in the best case, when a cyber risk manager (e.g., (re-)insurer) has complete knowledge of every client (organization) node's service liability and security investment portfolios, it is computationally intractable, i.e., NP-hard, to find out the number of dysfunctional (systemically affected) organizational nodes after a cyber-incident. Irrespective of the hardness of estimating systemic cyber risk, policies in practice to manage cyber risk, systemic or otherwise, through cyber (re-)insurance solutions has been active since the last decade and a half. The proven potential of cyber-insurance to improve cybersecurity has been mathematically shown in seminal papers Shetty et al. (2010), Hofmann (2007), Pal and Golubchik (2010), Pal et al. (2014), Naghizadeh and Liu (2014), Pal et al. (2018), Pal et al. (2011), though without reaching market efficiency. However, this has not completely discouraged cyber-insurance providers from increasing their supply of solution products, that is steadily seeing an increase over the years. The current advent of cyber re-insurance solutions is fuelled (since 2017) by the recent massive cyber-attack impacts caused by large-scale DoS (Mirai) and ransomware (WannaCry, Petya) attacks that have led to cascading and aggregate supply-chain organizational claims upon insurers. To this end, *recent theoretical efforts have zoomed in to the statistical nature of loss impact distributions, and their influence on the feasibility (if not optimal profitability) of cyber re-insurance markets.* More specifically, in a series of efforts Pal et al. (2020), Pal et al. (2020), Pal et al. (2020), Pal et al. (2021), Pal et al. (2020), the authors have proved that spreading *catastrophic* heavy-tailed cyber risks that are identical and independently distributed (i.i.d.), i.e., not tail-dependent, *is not* an effective practice for cyber re-insurers, whereas spreading i.i.d. heavy-tailed cyber risks that are *not catastrophic* is. While this latter point has long been believed and empirically validated in the cyber-insurance research literature, the former point is a surprising new facet that the authors unravel

via theory. In addition, spreading *catastrophic* and *curtailed* heavy-tailed cyber risks that are (non) i.i.d., i.e., not tail-dependent, *is not* an effective practice for cyber-reinsurers. *Orthogonal to investigating the statistical (and economic) sustainability of cyber re-insurance markets, we investigate on the computational feasibility of underwriting optimal cyber-insurance contracts for managing systemic cyber risk.*

3 SENSITIVITY OF SERVICE NETWORK RELIABILITY TO A SYSTEMIC CYBER-ATTACK

Even before we answer the question on how hard it is to computationally estimate systemic cyber risk, we should first provably demonstrate the importance of this problem by emphasizing to what degree a cyber-attack of the least systemic potency on a single enterprise node affects the reliability of the entire interdependent service network ecosystem. *This is a sensitivity analysis task in the jargon of operations research.* Subsequently, we will first mathematically demonstrate that even a small decrease in the liability portfolio (due to security mishaps) between any two enterprises in a service network can have a *large* externality-induced impact on the service reliability in the overall network. We will then prove in Section 4 that the systemic cyber risk estimation problem is NP-hard.

3.1 System Model

We assume n organizations/enterprises in a smart city setting and networked together via service relationships. Each organization invests in m security-enhancing instruments (SEIs) [e.g., antivirus, firewalls, security software updates], that contribute to the overall security strength (SS), i.e., reflects the cyber-security posture, of the organization. Each SEI instrument k has a weight of ss_k indicating the security strength/effectiveness of the instrument. As a conservative assumption (to strengthen the reality behind the theory claims we will make in this paper), we let ss_k to be the same for all organizations for any given k. This could pathologically happen if for example every organization buys the same commercial antivirus package, or all their IT systems use the same OS whose manufacturer releases a common security patch at a given time. Also note here that by analysing this pathological example we are implying that if the sensitivity is high for this case, then a major direct cyber-attack on a source enterprise that hits the QoS of the latter significantly can have a massive systemic impact on the networked enterprise ecosystem.

Each organization i is assumed to allocate a normalized portfolio proportion of $D_{ik} \in [0,1]$ to SEI k. A D_{ik} value of 1 indicates i harnessing the full power of SEI k, and a D_{ik} value of 0 either indicates i not buying SEI k, or not using/deactivating it post purchase/download. The $n \times m$ matrix is denoted by $D = (D_{ik})$. We define $C = (C_{ij})$ to be the $n \times n$ matrix indicating the portfolio of strength of service liabilities between organizations in a service network. Thus, organization i is liable upon organization j for C_{ij} fraction of its service needs. As an example j could be a public cloud provider like Amazon (e.g., AWS) selling compute and storage instances, and i could be a start-up relying on a significant amount of such resources for their daily operations. C carves out a directed network (graph) with n nodes representing the organizations, and an edge from organization i to j of weight $C_{ij} > 0$ ($C_{ii} = 0$ for all i). Now, $\sum_j C_{ij}$ is the fraction of organization i's service needs that are met by organizations (enterprises) external to i. The remainder needs of i, denoted by $\hat{C}_{ii} = 1 - \sum_j C_{ij}$, is met through its self-owned resources. The matrix $\hat{C}$ is a diagonal matrix with $\hat{C}_{ii}$ on the diagonal. The valuation of the total security strength gained by organization i (as being part of the service network) is then given by

$$V_i = \underbrace{\sum_k D_{ik} ss_k}_{\text{Valuation due to SEIs of } i} + \underbrace{\sum_j C_{ij} V_j}_{\text{Valuation due to SEIs of nodes 'feeding' } i}, \tag{1}$$

where V_i is the sum of the security strength gained by organization i on its SEIs, together with the liability proportion induced sum of the security strength gained by organizations on which i depends for its service (the out-degree neighbors of i in the service graph) for providing service to their customers. *Without loss of generality we assume that V_i is linearly separable for analytical tractability purposes.* The rationale behind

this leveraged formulation (in line with Hemenway and Khanna (2016)) is simply the communication-theoretic fact that cyber-security strength is both, a function of the service provider side and also the service receiver side (illustrating the network externality effect).

Moreover, the sum total of the security strength gained by organizations as being part of the service network is significantly higher than the sum of the security strength of the SEI components, i.e., $||V||_1 \geq ||\vec{ss}||_1$ they individually invest in, due to positive network externalities. In matrix notation, Equation (1) becomes $V = D\vec{ss} + CV$ which implies $V = (I-C)^{-1}D\vec{ss}$. The matrix $I-C$ is invertible because we assume that $\hat{C}_{jj} > 0$, hence the column sums of C are all strictly less than one. In fact, the matrix $I-C$ is an *M-Matrix* (Poole and Boullion 1974), and consequently $(I-C)^{-1}$ is an inverse M-Matrix (Willoughby 1977; Johnson 1982). In order to evaluate an organization i's non-leveraged security quotient with respect to V_i - its leveraged security strength gained through network externality effects, we must scale the latter quantity by the fraction of security requirements i *does not* offload (as in SECaaS business models) to other organizations, and takes it upon itself by investing in SEIs. The resultant quantity is i's contribution to the total amount of positive externality generated in the service network through SEI investments made by all organizations. Specifically, the individual security quotient of organization i w.r.t. V is $v_i = \hat{C}_{ii}V_i$. The vector of individual security quotient values are the solution to the following system of equations:

$$v = \hat{C}V = \hat{C}(I-C)^{-1}D\vec{ss}. \tag{2}$$

The matrix C is column sub-stochastic because column i sums to $1-\hat{C}_{ii}$. *Thus, in strongly inter-sector coupled smart societies, the SEI-induced security quotient of organizations is heavily dependent on the weights C in the service network.*

3.2 What is the Reliability Cost of a Service Edge Going 'Down' on a Cyber-Attack?

In this section we will study the impact of a small hit on the security strength of a single enterprise (organization) to the overall hit in security strength of the interdependent service network. In high-level terms we term this impact to be the reliability cost (to be formally explained below).

Although V_i for any enterprise i is useful information for an audit agency partnering an under-writing residual cyber risk management (RCRM) firm (e.g, cyber-insurer, cyber re-insurer), what matters most in assessing systemic cyber risk and the subsequent pricing of RCRM contracts is the security quotient of i. Higher the quotient (signaling increased positive intent to take more control of services security, than to offload it) higher the pricing confidence of under-writers to assess systemic cyber risk and fairly price i.
Consider an example simplistic scenario where a single organization (e.g., a cloud provider) that serves 10% of the business needs of a start-up (the customer) i is hit by a cyber-attack (e.g., type of Sunburst APT hacks that enable hackers to slip malware into software updates of SolarWinds' Orion software widely used by organizations in multiple service sectors) that allows it to only service 7% of i's needs till breach-management is successful. Consequently, $\varepsilon = (10-7) = 3\%$ is the negative service liability impact on i due to the cyber-attack on the cloud provider (in reality the cyber-attack will simultaneously affect many organizations reliant on the cloud provider with different magnitudes of ε). In addition, assume in the above example that the customer is small enough to contribute only $r_i = 5\%$ of its business and security needs through its own resources (e.g., virtual machines, SEIs), and relies on other organizations (as in SECaaS business models) for the remaining 95%.

Clearly, organizations reliant on i for service availability such as in the above example will be vulnerable to low QoS despite the positive externalities in relation to security enhancement i may receive from its 'feeder' peer organizations. After all, it only takes to compromise the 'feeder' organizations - occurrences that are frequently on the rise, especially in the COVID and work-from-home (WFH) era. Hence with the thought of systemic cyber risk in mind, i is likely an unattractive candidate to sell marketable RCRM policies (e.g., inexpensive premiums, low deductibles), despite the demand from the latter. On the contrary, RCRM contract sellers would not want to incur high opportunity costs in missing out on providing service to the pervasive IT-reliant small and medium businesses (SMBs), as the latter presents a huge market scope.

In order to resolve this dilemma, a question RCRM contract sellers would ask is: *what is the maximum negative impact of a particular ε incurred upon a* single *outgoing edge from any given i on v_i (derived from Equation (2))?* - customer *i*'s SEI-induced security quotient. For the remainder of the paper, we term this impact as the **reliability cost** of a service edge going 'down' (hit by a cyber-attack and incapable of working at full strength). *The reason - a drop in the security quotient of a organization node i will proportionally adversely affect the reliability of downstream services to other organizations sourced at i.*

The answer to the question has significant CRM pricing and regulatory importance simply because (a) appropriately scaling this impact showcases a loose bound on the aggregate system reliability costs incurred in the entire network if multiple service edges were impacted, and (b) the effect of liability-induced negative network externalities (post cyber-attack events) emanating from the 'central' service providers (e.g., a cloud provider) or SMBs on their own customer nodes (*i* being only one of them) could decide the seller investment portfolios, premium pricing, and deductible structures in the RCRM business. We have the following result inspired by existing research Hemenway and Khanna (2016) in financial service networks characterizing the reliability cost when a service edge in a network goes 'down' due to a cyber-attack.

Theorem 1 For any given service network inducing a directed graph (that might include cyclic liabilities dependencies) if $\|C-\tilde{C}\|_1 < \varepsilon$, then $\|\vec{v}-\tilde{\vec{v}}\|_1 < \frac{\varepsilon}{r}\|D\vec{s}\vec{s}\|_1$, where ε reflects the liability performance hit on a given service edge post a cyber-attack; $\|\vec{v}-\tilde{\vec{v}}\|_1$ is the reliability cost due to the ε hit; and $r = \min_i(\tilde{\hat{C}}_{ii}, \hat{C}_{ii})$ is the minimum value of r_i (fraction of non-offloaded business/security needs) for any enterprise *i*. In addition, it always holds that $\|\vec{v}-\tilde{\vec{v}}\|_1 \leq 2\|D\vec{s}\vec{s}\|_1$, thus resulting in

$$\frac{\|\vec{v}-\tilde{\vec{v}}\|_1}{\|D\vec{s}\vec{s}\|_1} \leq \min\{\frac{\varepsilon}{r}, 2\}.$$

Moreover, if the service network is acyclic, and if $||C-\tilde{C}||_1 < \varepsilon$, then $||\vec{v}-\tilde{\vec{v}}||_1 < \varepsilon||\vec{s}\vec{s}||_1$.

Proof Sketch - Using the principle of mathematical induction on the individual levels (for acyclic graphs), we see that the sum of the security strength values on the outgoing edges from level *i* in the graph is at most $||\vec{s}\vec{s}||_1$. Since each organization *i*'s V_i value is at most $||\vec{s}\vec{s}||_1$, an ε change in the QoS of any edge corresponds to an absolute change of at most $\varepsilon||\vec{s}\vec{s}||_1$ - thereby proving the theorem. For the case of cyclic graphs, the application of the triangle inequality via algebraic manipulation results in the theorem proof.

Theorem Implications in Practice - The theorem implies that the reliability cost, i.e., the maximum amount of hit to the security quotient, induced by a liability performance degradation on even a single edge could be very high if the service network graph is cyclic, but is bounded by $\varepsilon||\vec{s}\vec{s}||_1$ if the network is acyclic. Simply put, *negative externalities amplify in cyclic graphs.* For the example above, when r_i = 5%, the reliability cost could shoot upto atmost 20ε for a given ε value if the underlying service graph is allowed to have cycles - leave alone how aggregate reliability costs would be incurred if liability guarantees on multiple edges degraded simultaneously. In a service network design context, inter-dependency cycles should be avoided if and when possible. On the other hand for acyclic service graphs the reliability cost can go as low as 5ε for the above example. From a policy design viewpoint, it is fair to say that measures to modify graphical liability structures between organizations (e.g., promoting acyclic liability graphs) to reduce reliability costs is a distant reality - however, information disclosure regulations could be made stronger for organizations to be able to report the ε's, the *r* values (the ε's and the *r* enabling a public accountability of organizational cyber-security posture along with the degree of self-liability undertaken), and their liability relationships for RCRM contract providers to effectively estimate systemic cyber risk and subsequently price CRM contracts appropriately. In addition, there should be regulations in effect that mandate the existence of threshold security controls to be enforced upon IoT device manufacturers prior to packaging for sale. This would reduces chances of hits to security quotient values.

4 SYSTEMIC CYBER RISK ESTIMATION IS COMPUTATIONALLY INTRACTABLE

Thus far, we showcased the importance of organizational reliability costs to RCRM solution providers and showed that they could be very high in a service network under the effect of externalities. In this section, we show (along with a discussion on practical and policy implications) that even in the ideal (information symmetric) scenario where there is perfect knowledge to a RCRM contract underwriter of (a) reliability costs, (b) the service network, and (c) the individual organizational investments in SEIs, it is computationally hard, post a cyber-breach event(s) in any part of the network, to estimate the number of dysfunctional organizations who either cannot provide a given service to their customers or provide them at very low QoS. *In other words it is computationally hard to accurately estimate the systemic effects of a cyber-incident.* For the purpose of this paper, we have termed this number as a proxy but proportional estimate of systemic cyber risk post a cyber-breach event(s) and is a very important input in the interests of scalable RCRM businesses. More specifically, we address the following question under an ideal information symmetry scenario: *given an inter-dependent service network, if the total security strength of the SEIs drop (either on a single instrument or in combination) by some small amount d (due to budget-induced poor cyber-hygiene) for any organization and it becomes a successful target to a cyber-security breach w.r.t. a given service, what is the maximum number of organization nodes that will become dysfunctional w.r.t. to any corresponding service, due to the inter-dependent nature of the network?*

This is one of the most fundamental questions in systemic cyber risk management when it comes to judging the feasibility of cyber risk aggregation by third-party risk managers, simply because an 'unaccounted for' dysfunctional node of central importance in the service network could contribute to a heavy-tailed aggregate cyber risk distribution that might render RCRM infeasible (see Pal et al. (2020), Pal et al. (2021), Pal et al. (2020), Pal et al. (2020) for a detailed theory). In addition, it is often the case in real practice that many IT and cyber-physical driven organizations, irrespective of their sizes and reputation, do cut their costs with respect to investing in appropriate amounts of back-end cyber-security in favor of boosting/re-aligning their investments in other front-end ventures (Blau 2017). The primary reason behind this enigmatic trend is that the digital threat landscape changes constantly, and it's very difficult for the C-suite of enterprises to know the probability of any given cyber-attack succeeding — or how big the potential losses might be. Consequently, this adversely affects their judgement of the ROI on cyber-security investments, i.e., investments in SEIs. Insights from behavioral economics and psychology show that human judgment is often biased in predictably problematic ways. In the case of cybersecurity, some decision makers use the wrong mental models to help them determine how much investment is necessary and where to invest. As an example, some CEOs think that SEI instruments that create a fortified castle is all that's needed to keep a company safe, and cut costs on critical SEI components such as hiring specialist vulnerability testing engineers. As a result, the goals of a financial decision maker will always be oriented toward cyber risk mitigation instead of cyber risk management.

The fact that we show that answering the aforementioned fundamental question is NP-hard even for ideal information symmetric scenarios, renders the question NP-hard even for practical information asymmetric scenarios. *Our hardness result is based on the computational hardness of the Balanced Complete Bipartite Subgraph (BCBS) problem* that involves finding a maximum balanced clique in a bipartite graph. The problem is formally defined as follows (see Garey and Johnson (1979)).

Definition 1 (BCBS) Given a bipartite graph $G = (V_1, V_2, E)$ with $|V_1| = |V_2| = n$, find the largest integer K such that there exists sets $C_1 \subset V_1$ and $C_2 \subset V_2$ with the properties that $|C_1| = |C_2| = K$, and the induced graph on $C_1 \cup C_2$ is a complete bipartite subgraph of G.

It is well known from the theoretical computer science literature that the BCBS problem is NP-hard (Garey and Johnson 1979). In practical terms, this implies that there is no scalable algorithm (unless P = NP) that can compute the size of the maximum balanced clique in a bipartite graph in a reasonable (polynomial in n) amount of time. As a matter of fact, the problem is so computationally difficult that even *approximating* the size of the largest balanced clique is hard. More specifically, according to seminal

results in Feige (2002), Feige and Kogan (2004) for some $\delta > 0$ (i) it is Random 3-SAT hard to approximate BCBS to within a factor of n^{δ}, and (ii) provided the optimal solution to BCBS can be approximated to within a factor of $2^{(\log n)^{\delta}}$ for every $\delta > 0$, 3-SAT can be solved in time $2^{n^{3/4}+\varepsilon}$ for every $\varepsilon > 0$, and (iii) there exists no polytime algorithm that will approximate BCBS to within a factor of n^{δ} (unless P = NP).

We now state our main result, adapted for cyber-settings from financial network settings (Hemenway and Khanna 2016), elucidating the computational complexity of estimating systemic cyber risk that in our paper is equivalent to estimating the maximum number of dysfunctional (with respect to a given service) organization (enterprise) nodes in an inter-dependent service network post a cyber-breach event on a single node due to a C-suite induced reduction of d units of investment (information unknown to an RCRM solution provider) in cyber-hygiene by the said organization.

Theorem 2 For every bipartite graph G on $2n$ nodes, and every $\varepsilon > 0$, there is an acyclic inter-dependent service network with $\Omega(n)$ organizations (enterprises), and a $d > 0$ such that computing the maximum number of organizations that could become dysfunctional with respect to a given service (as a proxy to computing an estimate of systemic cyber risk), following a cyber-breach event on a given organization due to a reduction of $d\varepsilon$ units of investment in SEIs, is as hard as solving the BCBS problem in G. In other words, estimating systemic cyber risk in interdependent service networks is NP-hard.

Proof Sketch - The proof sketch involves a brief explanation of a reduction and construction logic. Let $l > 0$ be any integer. Note that for an $n \times n$ balanced bipartite graph G, it is a computationally hard problem (i.e., NP-hard problem) to decide whether the largest balanced bipartite clique size in G is at least $K \times K$ or at most $\frac{K}{g} \times \frac{K}{g}$ for some arbitrary gap function g. Given G, we will construct a inter-dependent service network with $(2+l)n$ enterprises in a manner such that if G has a balanced bipartite subgraph of size k, then a drop in security investments on SEI by a networked organization (enterprise) by an amount $K\varepsilon$ can cause at least $(2+l)K$ enterprises to become dysfunctional in a systemic fashion. On the other hand, if the largest balanced bipartite subgraph of G is of size $\frac{K}{g}$, a drop in the security investments on SEI instruments by K can cause at most $K + \frac{K}{g}(l+1)$ organizations to become dysfunctional in a systemic fashion. The implication here is that estimating the maximum number of dysfunctional organization nodes (w. r. t. to a given service) induced by a fixed reduction in SEI investments is at least as hard as estimating the size of the maximum balanced bipartite clique. Now when $g = n^{\delta}$, choosing l = poly(n) gives us a gap of $((l+2)n)^{\delta'}$ for $\delta' < \delta$ and with the help of additional algebraic manipulations, we prove Theorem 2.

Theorem Implications in Practice - Apart from the evident fact that systemic cyber risk estimation in an inter-dependent service network is NP-hard, the theorem showcases multiple practical insights.

First, the result is true for acyclic bipartite interdependent service networks - in practice, service networks are likely to be cyclic making systemic cyber risk estimation in such networks even more computationally challenging. *Second*, in the presence of information asymmetry about reliability costs, network topology, vulnerability information sharing in public, etc., cyber risk terrain estimation is computationally hard. *Third*, it is computationally intractable to have an accurate estimate of cyber risk aggregated at an organizational node post a cyber-breach event - a strong counter-force to the foundational principle of RCRM solution providers who seek to attempt to identify the important "sources of aggregation" which if exploited or disrupted, have the potential to negatively impact many organizations, endpoint devices or individual persons. In practice, systemic cyber risk estimation is one of the most difficult aspects to managing cyber risk for any insurer or reinsurer. because of the "limitless" that could be imagined, and these constantly evolve over time with advancements in technology and threat actor capabilities. *Fourth*, either, the RCRM business will be difficult to scale for smart societies, or there should be significant capital influx from regulatory/financial authorities to RCRM solution providers in addition to the establishment of strong information disclosure laws by policy makers enforced upon RCRM clients. *Finally*, self-insurance will and should remain a major investment by enterprises for two reasons: (a) to 'insure' for risk categories that are either usually excluded by RCRM solution providers or are weakly covered, and (b) to sustain a threshold level of cyber risk management during periods of volatile RCRM markets. As an example of

(a), the current cyber-insurance market has excluded "infrastructure failure" (e.g., failure of energy and telecommunication networks) and "cyber war" between nation states as coverage areas due to them being failures that lead to "unmanageable aggregation risk". This is a big exclusion as such cyber-attacks are on the rise and leading organizations to rely on self-insurance and technology solutions to manage cyber risk.

5 POLICY RECOMMENDATIONS TO IMPROVE SYSTEMIC CYBER RISK ESTIMATION

We layout *three* policy making ideas that can improve the effectiveness of estimating systemic cyber risk.

Boost Capability to Recognize Software Dependencies - There should be significant effort by cyber risk managers of liability networked enterprises to understand software supply chains, given a recent published list by the US National Institute of Standards and Technology (NIST) identifying critical software categories such as operating systems and web browsers that are common across global enterprises. A way forward in this direction is the sharing of Biden-administration backed NIST-recommended Software Bill of Materials (SBOMs) that will allow vendors to transparently communicate contents of their software with enterprise management. Other complementary solutions include identifying elements of open source code that can be systemically critical (e.g., solutions innovated by the Linux Foundation and Google).

Identify Systemically Critical Enterprises in an Interdependent Network - Like post the case of the financial crisis of 2008 where the Financial Stability Board (FSB) designates global banks and insurance companies in financial networks as systemically important, a similar thing needs to be done by a regulatory agency for interdependent societal networks of enterprises. As precedence, a 2021 bill in US House of Representatives allows CISA to designate subset of systemically important critical enterprises. These are enterprises that are at most risk of triggering, propagating and/or suffering the adverse impact of systemic cyber risk events. Regulated platforms should be there in place for such enterprises to share among themselves systemic catalysis pointers. Agencies like the International Telecommunication Union (ITU) should set aside geopolitical considerations and work with tech companies to mitigate challenges to effective systemic cyber risk management in enterprise ecosystems.

Facilitate Collaborative and Modular Systemic Cyber Risk Assessment - It is necessary to convene global and disparate multi-stakeholder working groups where each stakeholder can focus on particular concentrated areas of systemic cyber risk (e.g., cloud, open source software), and then come together to share and combine insights that result in non-optimal but very effective estimates of systemic cyber risk across an interdependent societal network of enterprises. The Internet Security Research Group (ISRG) and the US National Security Telecommunications Advisory Committee (NSTAC) are examples of such working groups that can work with cyber-insurance companies for effective systemic cyber risk estimation.

6 SUMMARY

We proved that *optimal* systemic cyber risk management is computationally intractable, i.e., NP-hard, signifying the fact that optimal RCRM is utopic to achieve for worst case threat environments but should not be a deterrent to scaling cyber (re-)insurance markets. We proposed policy actions for improving management of space-time adverse impact of inevitable systemic (and worst case) cyber risk events.

ACKNOWLEDGMENTS

This study has been supported by funding from Cybersecurity at MIT Sloan (CAMS).

REFERENCES

Blau, A. 2017. "The Behavioral Economics of why Executives Underinvest in Cybersecurity". *Harvard Business Review* 95(3):22–25.

Feige, U. 2002. "Relations between Average Case Complexity and Approximation Complexity". In *Proceedings of the Thiry-Fourth Annual ACM Symposium on Theory of Computing*, STOC '02, 534–543. New York, NY, USA: Association for Computing Machinery.

Feige, U. and S. Kogan. 2004. "Hardness of Approximation of the Balanced Complete Bipartite Subgraph Problem". Technical report, Technical Report MCS04-04, Department of Computer Science and Applied Math.

Forscey, D., J. Bateman, N. Beecroft, and B. Woods. 2022. *Systemic Cyber Risk: A Primer*. Washington, D.C.: Carnegie Endowment for International Peace.

Garey, M. R. and D. S. Johnson. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness*, Volume 174. San Francisco, CA, USA: W. H. Freeman and Company.

Hemenway, B. and S. Khanna. 2016. "Sensitivity and Computational Complexity in Financial Networks". *Algorithmic Finance* 5(3-4):95–110.

Hofmann, A. 2007. "Internalizing Externalities of Loss Prevention through Insurance Monopoly: An Analysis of Interdependent Risks". *The Geneva Risk and Insurance Review* 32:91–111.

Johnson, C. R. 1982. "Inverse M-matrices". *Linear Algebra and its Applications* 47:195–216.

Naghizadeh, P. and M. Liu. 2014, June 23-24. "Voluntary Participation in Cyber-Insurance Markets". In *Workshop on the Economics of Information Security (WEIS)*, 23–24. State College, PA, USA.

Pal, R. and L. Golubchik. 2010. "Analyzing Self-Defense Investments in Internet Security under Cyber-Insurance Coverage". In *IEEE 30th International Conference on Distributed Computing Systems*, 339–347. Genoa, Italy: IEEE.

Pal, R., L. Golubchik, and K. Psounis. 2011. "Aegis A Novel Cyber-Insurance Model". In *Decision and Game Theory for Security*, edited by J. S. Baras, J. Katz, and E. Altman, 131–150. Berlin, Heidelberg: Springer Berlin Heidelberg.

Pal, R., L. Golubchik, K. Psounis, and P. Hui. 2014, April 27-May 2. "Will Cyber-Insurance Improve Network Security? A Market Analysis". In *IEEE INFOCOM 2014 - IEEE Conference on Computer Communications*, 235–243. Toronto, Canada.

Pal, R., L. Golubchik, K. Psounis, and P. Hui. 2018. "Improving Cyber-Security via Profitable Insurance Markets". *ACM SIGMETRICS Performance Evaluation Review* 45(4):7–15.

Pal, R., Z. Huang, S. Lototsky, X. Yin, M. Liu, J. Crowcroft *et al*. 2021. "Will Catastrophic Cyber-Risk Aggregation Thrive in the IoT Age? A Cautionary Economics Tale for (Re-) Insurers and Likes". *ACM Transactions on Management Information Systems (TMIS)* 12(2):1–36.

Pal, R., Z. Huang, X. Yin, M. Liu, S. Lototsky and J. Crowcroft. 2020. "Sustainable Catastrophic Cyber-Risk Management in IoT Societies". In *2020 Winter Simulation Conference (WSC)*, 3105–3116 https://doi.org/10.1109/WSC48552.2020.9384051.

Pal, R., Z. Huang, X. Yin, S. Lototsky, S. De, S. Tarkoma *et al*. 2020. "Aggregate Cyber-Risk Management in the IoT Age: Cautionary Statistics for (Re) Insurers and Likes". *IEEE Internet of Things Journal* 8(9):7360–7371.

Pal, R., P. Liu, T. Lu, and E. Hua. 2023. "How Hard is Cyber-Risk Management in IT/OT Systems? A Theory to Classify and Conquer Hardness of Insuring ICSs". *ACM Transactions on Cyber-Physical Systems (TCPS)* 6(4):1–31.

Pal, R., P. Liu, T. Lu, and X. Yin. 2021. "Cyber Re-Insurance Policy Writing is NP-Hard in IoT Societies". In *2021 Winter Simulation Conference (WSC)*, 1–12 https://doi.org/10.1109/WSC52266.2021.9715524.

Pal, R., K. Psounis, J. Crowcroft, P. Hui, S. Tarkoma, A. Kumar *et al*. 2020. "When are Cyber Blackouts in Modern Service Networks Likely? A Network Oblivious Theory on Cyber (Re) Insurance Feasibility". *ACM Transactions on Management Information Systems (TMIS)* 11(2):1–38.

Poole, G. and T. Boullion. 1974. "A Survey on M-matrices". *SIAM Review* 16(4):419–427.

Shetty, N., G. Schwartz, M. Felegyhazi, and J. Walrand. 2010. "Competitive Cyber-Insurance and Internet Security". In *Economics of Information Security and Privacy*, edited by T. Moore, D. Pym, and C. Ioannidis, 229–247. Boston, MA: Springer US.

Willoughby, R. A. 1977. "The Inverse M-matrix Problem". *Linear Algebra and its Applications* 18(1):75–94.

AUTHOR BIOGRAPHIES

RANJAN PAL is a Research Scientist with the MIT Sloan School of Management, and an invited working group member of the World Economic Forum. His primary research interest lies in developing interdisciplinary cyber risk/resilience management solutions. He serves as an Associate Editor of the ACM Transactions on MIS. His email address is ranjanp@mit.edu. Ranjan has a PhD in computer science from the University of Southern California and was a postdoc at the University of Cambridge.

ROHAN XAVIER SEQUEIRA is a PhD student and Annenberg Fellow in the department of electrical and computer engineering (ECE) at the University of Southern California. His research interest lies is cyber risk management, privacy, and distributed systems. His email address is rsequeir@usc.edu. Rohan got his MS in ECE from the University of Michigan Ann Arbor.

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. He is the President of the Security, Stability, and Resilience Special Interest Group of the System Dynamics Society. His email is szeijl@mit.edu.