

OPTIMIZING CYBER-RESILIENCE IN CRITICAL INFRASTRUCTURE NETWORKS

Ranjan Pal¹, Rohan Xavier Sequeira², Sander Zeijlmaker¹, and Michael Siegel¹

¹MIT Sloan School of Management, Massachusetts Institute of Technology, Cambridge, MA, USA

²Electrical and Computer Engineering, University of Southern California, Los Angeles, CA, USA

ABSTRACT

With the expanding cyber-risk terrain spanning business processes in digitally driven enterprises with critical infrastructure, it is inevitable in time that system process continuity (SPC) will be affected (e.g., via ransomware) for certain inter-dependent processes of such an enterprise, and hamper business continuity. We are interested in the question: *how should managers of such enterprises optimize cyber-resilience (i.e., the ability to maintain SPC via absorbing and adapting to an adverse cyber-incident) for any complex networked critical infrastructure (CI) (sub-)system with multiple process functionality components (PFCs)?* We prove via an algorithmic graph-theoretic approach that optimizing or approximately optimizing cyber-resilience within a pre-specified enterprise cyber-protection budget in any CI with networked and inter-dependent PFCs is NP-hard. Consequently, we propose a computationally tractable graph-based Monte-Carlo simulation framework to ‘optimize’ (boost) cyber-resilience within any PFC network by allocating a constrained cyber-protection budget among PFCs in accordance with their Katz centralities in the PFC network.

1 INTRODUCTION

The modern operational technology (OT) driven enterprise market is crucial to businesses spanning a wide range of (public and private) sectors that include energy, finance, retail, chemical, power, manufacturing, transportation, bio-technology, and other end-user verticals. Here, OT encompasses cyber-physical systems (CPSs) that control and monitor physical equipment and processes serving such businesses. It is quite fair to say that such businesses are supported atop IoT-driven critical infrastructure (CI) responsible for necessarily providing high quality of service (QoS) to (mission critical) societal applications and often in real-time. It is projected by the World Economic Forum (WEF) that the OT driven enterprise market will see a (non-)linear growth increase with the increase in the global Industrial Internet of Things market that currently stands at approximately 85.5 billion USD in 2023 to almost USD 169 billion by 2028.

The physical machinery underlying an enterprise CI supporting business processes in these sectors, that traditionally used to be dumb, are often embedded today with software-programmable IoT devices such as sensors, smart phones, actuators, programmable logic controllers (PLCs), programmable automation controllers (PACs), and other intelligent electronic devices (IEDs). Furthermore, these devices are legacy in nature; they can communicate with each other over a wireless network (e.g., WiFi, 5G) and/or the Internet through proprietary network protocols, and are often managed atop by a social (logical) network of enterprise employees who oversee both IT and OT systems.

This resulting human-managed, smart, and networked cyber-physical infrastructure with (IT) software and IoT driven operational technology controls is supporting increasingly new forms of enterprise business process applications built upon software stacks. The steady growth of such modern cyber-physical infrastructures supporting business process communication between and across C-suites, middle management, and bottom levels is primarily due to (a) rising cost of labor, (b) pressure on businesses to satisfy the two-fold constraint of meeting receding deadlines under increasing demand, (c) organizational push to improving quality control via real-time data driven decision making, and (d) mitigating human error in increasingly automated business processes. In summary, any communication network within cyber-physical




Figure 1: Illustration of cyber-attack anticipation and CI cyber-resilience dynamics post the attack event.

CI (CIs from here on in the paper) comes with the immense promise to generate significant economic and societal benefit through improved efficiency, productivity, and reliability by supporting a plethora of (critical) day-to-day modern business processes.

A CPS-driven enterprise network (EN) can necessarily sustain process continuity for OT system processes both, for itself and for other reliant system processes dependent on this EN (e.g., as in a supply chain environment internal/external to an enterprise) *only if* relevant (sub-)systems spanning the EN are cyber-resilient in the face of system-disrupting cyber-breach incidents (cyber-attacks) with financial consequences for the enterprise. As an example, according to a World Economic Forum report (Martel et al. 2019), a six-hour winter blackout in mainland France due to a non cyber-resilient EN (sub-)system within an electricity grid could result in financial damages of around 1.5 billion Euro (or USD 1.7 billion) to the mainland society. Hence, cyber-resilience of EN (sub-)systems is a critical and necessary condition for sustaining system process continuity (SPC) of individual enterprises and enterprise supply chain ecosystems spanned by them. Here, a ‘(sub-)system’ refers to any (sub-)network at the physical/logical abstractions (process levels) of an EN.

In this paper, in accordance with NIST and the US National Academy of Science, we define "cyber-resilience" as the *ability* of any EN (sub-)system to successfully absorb and adapt to such cyber-breach incidents (Kott and Linkov 2019) by providing at least a minimum acceptable level of a performance measure that sustains enterprise SPC. Alternatively, in accordance with the NIST jargon, cyber-resilience is the ability of cyber-incident response in any EN (sub-)system to any cyber-incident.

In other words, this definition showcases two things: (i) cyber-attacks are inevitable for an enterprise built upon a digital and Internet networked backbone (see Anderson (2020) for reasons sourced in factors such as security technology, security economics, adversarial evolution, and humans in the loop), and (ii) it is often a priority of an enterprise management to *always* sustain SPC within any EN (sub-)system, even if at degraded levels of system performance.

Real World Examples of Critical Infrastructure Cyber (Non-)Resilience - As a practical example of non cyber-resilient system behavior in the **energy sector**, the *Kyivoblenergo* regional electric distribution company in Ukraine faced a cyber-attack in the year 2015 that eventually lead to regional power blackouts for some hours (*hence exhibiting cyber non-resilience due to a lack of business continuity*) that not only disrupt consumer lifestyle and the local energy transmission business, but all other businesses (e.g., manufacturing, healthcare) that were (critically) dependent on power. This is an example of cyber non-resilience (representing the orange and red curves in Figure 1) on part of the electric distribution company due to the management’s lack of maintaining business (i.e., power supply) continuity for multiple hours. As another practical example of non cyber-resilient system behavior in the **healthcare sector**, the UK’s *National Health Service* (NHS) was attacked in mid-2017 by cyber-criminals using the *WannaCry* ransomware that

locked medical records from medical staff access, and the adversary demanded a USD 300 payment in Bitcoin that would double after three days. The attack disrupted health services in hospitals across Britain (*showcasing the inability of NHS to absorb and adapt to a cyber-attack*) as multiple patient appointments (around 19000) had to be canceled and emergency ambulances diverted to non-emergency locations. This is an example of cyber non-resilience (representing the orange and red curves in Figure 1) on part of NHS and certain hospitals in Britain due to the managements' lack of maintaining system process continuity (i.e., holding patient appointments, medical sessions etc.,) for days.

1.1 Research Motivation and Goal

We have the following research motivation and goals for this paper.

Motivation - Enterprise managers of CIs would want to always be in the blue and green zones of Figure 1 after the occurrence of an adverse cyber event. Failing to do so will result in service disruption of processes relying upon these OT systems (e.g., internal OT driven IT supply chains), in addition to the obvious disruption of OT system processes. The ability to provide the basic minimum quality of service (QoS) to service clients 24/7 is the first and foremost priority of every CI-driven service enterprise business, but is not guaranteed when the services are run atop OT systems whose security management quality is widely known to be moderate at best, as reported by the WEF (Martel et al. 2019). According to the standard NIST framework, this reflects the ability of any digital system infrastructure to absorb and adapt to the cyber-attack. We denote this ability as our notion of cyber-resilience in our paper. *Subsequently, we are motivated to optimize this ability for a general CI system exhibiting (a) a physical communication network among its CPS elements, and overlaid above by (b) logically networked and interdependent relationships between system functionality components - each component driven by CPS elements mentioned in (a).*

Goal - Our goal in this paper is to design a methodology that enables an enterprise management to optimize our proposed notion of cyber-resilience for any logical network of system functionality components as a function of a given enterprise allotted cyber-protection budget for CI system security.

1.2 Research Contributions

There is a lack of (formally-backed) research principles on how C-suites and CI system managers should optimize cyber-resilience for their business processes served atop complex and networked CPS environments having interdependent relationship between processes. To the best of our knowledge, there exists no systematic framework that specifically addresses cyber-resilience optimization in settings where system processes are interdependent and networked. In reality, most (if not all) real world OT driven ICSs are characterized by interdependent processes (Khan and Madnick 2021). *Hence, in this research we take a first pass at proposing a general formal framework to optimize cyber-resilience in any such aforementioned complex networked critical infrastructure.*

We formulate an **algorithmic graph-theoretic framework** (see Sections 3 and 4) aimed to enable a networked CI system management computationally achieving a desired level of cyber-resilience within a pre-specified enterprise cyber-protection budget for networked and interdependent process functionality components PFCs in any networked (sub-)system induced by the CI. The framework necessitates solving a constrained optimization problem that should ideally output a set of networked PFCs among which the pre-determined cyber-protection budget needs to be allocated. This allocation will achieve a desired optimal level of cyber-resilience (expressed as a formal condition), given failure thresholds of individual PFCs to be deemed dysfunctional. Evidently, there could be multiple output sets of PFCs characterizing a feasible solution and we are interested to find the set (if unique) consisting of the most critical PFCs. We term this set solution as an optimal cyber-resilience ensuring policy (OCEP). We show that the optimization problem under consideration to derive an OCEP is computationally intractable, i.e., NP-hard. Moreover, the budget-constrained optimization problem targeting achieving a mathematical approximation of OCEP on which the cyber-protection budget should be allocated for obtaining desired cyber-resilience levels is

also NP-hard. In other words, given arbitrary instances of an enterprise PFC cyber-protection budget, a CI (sub-)system network structure, and a pre-specified desired dysfunctionality conditions for PFCs, it is very difficult even for a computer (leave alone humans) to *always* determine a budget-constrained OCEP or even a budget-constrained approximate OCEP.

Consequently, we propose (in Section 5) a **graph-based heuristic framework** (validated via Monte Carlo simulations) that strategically allocates a cyber-protection budget among the enterprise PFCs in accordance to the Katz eigenvector centralities of PFCs in a network. Our proposed graph heuristic framework is computationally tractable in nature. Monte Carlo simulations of such a framework on practical CI topologies illustrate that our proposed Katz eigenvector centrality based heuristics result in *boosted PFC network cyber-resilience* when compared to a non-strategic budget allocation approach. This is because **crown jewel** PFCs are made strategically more resilient. Subsequently, we are the first to design a computationally tractable graph heuristic to resolve the hardness challenge of the aforementioned cyber-resilience optimization problem and formally boost CI system cyber-resilience. In all of our Monte Carlo simulations, we adopt randomized and scaled PFC network topologies motivated by the PFC network structure in the real-world electricity microgrid in Boston, Massachusetts, USA (Khan and Madnick 2021).

2 RELATED WORK

We briefly state the related work in relation to two relevant themes: quantifying cyber-resilience in CI networks and optimizing cyber-resilience on the same.

Quantifying Cyber-Resilience - Most well-known system cyber-resilience metrics introduced in the research literature are engineering focused, and either model cyber-resilience as (a) a rebound of the system from cyber-shock to reach the usual state of equilibrium level of performance at which the system usually performs, or (b) a synonym for robustness allowing the system to function at degraded but acceptable levels of performance post a rebound from a cyber-shock (Francis and Bekera 2014; Linkov et al. 2013; Clark and Zonouz 2017; Woods 2015; Hosseini et al. 2016; Arghandeh et al. 2016; Gholami et al. 2018; Venkataramanan et al. 2019; Venkataramanan et al. 2019; Zuloaga et al. 2019; DiMase et al. 2015; Sterbenz et al. 2011; Chaves et al. 2017; Haque et al. 2018). Despite a highly application-dependent overloading in the definition of cyber-resilience across these works, the common aspect among these metrics is that they are derived using mathematical frameworks that all account for the cyber-vulnerability dynamics of each (sub-)system component or a network (Haque et al. 2018), alongside some accounting for an adversarial input to model the cyber-vulnerability dynamics.

However, a common drawback to all these metrics is the fact that none of them account for the extent of liabilities between CI networked PFCs - a salient complex system property, irrespective of whether the cyber-resilience measure is network dependent or not. More specifically, the degree of liability between (sub-)system components creates negative service degradation externalities that (non-linearly) percolate throughout an CI network when individual PFCs experience a cyber-shock. These percolating externalities, that directly influence the ability of components to absorb and adapt, go unaccounted for in the calculation of existing cyber-resilience metrics. In a prior work Pal et al. (2024), we alleviate the aforementioned pitfall by proposing a formal amalgamated methodology that accounts for the percolation of the negative externalities throughout a liability-driven CI network in determining a quantitative measure of cyber-resilience.

Optimizing Cyber-Resilience - Optimizing cyber-resilience over a quantified metric in complex networked engineering systems has been relatively less explored in the literature. This is because there have been relatively few efforts to quantify systems resilience in CI systems, let alone networked CI systems. Haque et al. (2018) and Haque et al. (2021) use the *Technique for Order Preference to Ideal Solution* (TOPSIS) methodology from operations research to identify the critical components of networked CI systems in ranking order. The ranking order can then be used by system managers to proportionally invest a cyber-security protection budget across the ranked nodes to boost or optimize cyber-resilience. *However, unlike in our work, Haque et al. (2018) and Haque et al. (2021) do not extend their methodology to design algorithms to boost cyber-resilience.* In addition, unlike us, *there exists no research to the best*

of our knowledge that formally classifies the computational complexity of the cyber-resilience optimization problem in CI systems with networked and interdependent components. Fang et al. (2016), Zio and Piccinelli (2010) and Barker et al. (2013) optimize systems resilience in critical infrastructure system networks with interdependent components. However they assume that network edges are brittle, i.e., they either function at full capacity or does not. In contrast system components are often non-brittle to function at partial capacities if adversely impacted by a cyber-attack event. In addition, these works do not classify the computational complexity of the cyber-resilience optimization problem in systems with networked and interdependent components.

3 SYSTEM MODEL

Every enterprise would ideally want to make the most out of investments it makes in a particular venture (e.g., cyber-resilience in our case). In this section, we set up the system model to analyze whether an operational technology driven CI enterprise management can achieve cyber-resilience in an optimal fashion given a budget constraint, and performance thresholds (below which, PFC dysfunctionality is assumed) of individual CI PFCs. Our formal model (and notations) for budget-constrained cyber-resilience optimization is adapted from an orthogonal setting of financial resilience proposed in Klages-Mundt and Minca (2022).

3.1 Formalizing Notations

We define a networked and inter-dependent CI PFC $(C, D, \beta, \theta, \mathbf{p})$ network as follows:

- n nodes, each representing a PFC inside a CI PFC network.
- m operational CI functionalities spread over the n PFCs, with each functionality possibly handled by more than one PFC.
- $\mathbf{p} = m \times 1$ vector of per-unit system management costs for operating each of the m CI functionalities.
- $D = n \times m$ matrix with $D_{ik} \geq 0$ the matrix (the rows spanning the multiple CI PFCs) of the fraction of system management efforts for ICS functionality k (spanning the columns of the matrix) exhausted by PFC i (adding to 1).
- $C = n \times n$ matrix with $C_{ij} \geq 0$ the fraction of operational (i.e., functional) dependency of CI PFC j upon CI PFC i and 0 along the diagonals.
- $\hat{C} = n \times n$ diagonal matrix with $\hat{C}_{ii} = 1 - \sum_j C_{ji}$ the fraction of the operational (i.e., functional) space of CI PFC i that is not functionally dependent on any another PFC $j \neq i$ in the CI (PFC i self-generates resources to service the functional space $\hat{C}_{ii}$).
- $\theta = n \times 1$ vector of failure thresholds for each CI PFC i - a lower bound of degraded PFC performance below (ideally much below normal/best performance) which the PFC is deemed dysfunctional.
- $\beta = n \times n$ diagonal matrix of (optional) additional system management costs (e.g., costs spent on maintenance agencies) incurred by the CI to when the respective CI PFC becomes dysfunctional.

We assume w.l.o.g. that matrix C is column sub-stochastic without which $\hat{C}^{-1}$ is not well-defined. Consequently, note that the matrix $I - C$ is invertible since the spectral radius $\rho(C) < 1$.

3.2 Formalizing ICS Component Performance

Subject to the above notations, the PFCs' inter-component dependency-induced performance level values are given by the following matrix: $\mathbf{V} = \mathbf{C}\mathbf{V} + \mathbf{D}\mathbf{p} - \beta \mathbf{1}_{\{\mathbf{v} < \theta\}}$, where $\mathbf{1}_S$ is the 1-0 valued vector indicating the entries of set S . This 'cost-benefit' equation directly follows from the meaning of the C , $\mathbf{p}$, and D notations with respect to achieving a threshold level of performance θ across the PFCs. These values can also be represented via the following equation in vector form for the ease of analysis: $\mathbf{v} = \hat{C}\mathbf{V} = \hat{C}(\mathbf{I} - \mathbf{C})^{-1}(\mathbf{D}\mathbf{p} - \beta \mathbf{1}_{\{\mathbf{v} < \theta\}})$. It follows from Elliott et al. (2014) (an adaptation from component performance in financial networks) that (i) $\hat{C}(\mathbf{I} - \mathbf{C})^{-1}$ is column-stochastic, and (ii) there always exists a solution for $\mathbf{v}$ via the application of Tarski's

fixed point theorem (Tarski 1955) on the complete lattice of solutions for $\mathbf{v}$. In this paper, we assume a scalar real quantity (for model tractability benefits) for individual component, i.e., PFC, performance - however, the insights obtained post analysis is applicable to settings with heterogeneous units spanning real and discrete measures.

3.3 Formalizing Cyber-Protection Portfolio to Achieve PFC Performance Thresholds

One of the salient system manager roles within a CI is to invest in PFC cyber-protection to ensure that ideally all PFCs always perform above a threshold level of performance - even post a cyber-attack event. To capture this concept, we add a vector of component-wise (i.e., for each PFC) cyber-protection investments $\gamma \geq 0$, which affect the dysfunctional status of PFCs post a cyber-attack event. *This vector is representative of a cyber-protection investment portfolio across CI PFCs.* Consequently, given an investment portfolio γ , component i becomes dysfunctional if $V_i + \gamma_i < [\hat{C}^{-1}\theta]_i$. This leads to post-investment component performance values

$$\tilde{\mathbf{v}} = \hat{C}(I - C)^{-1}(D\mathbf{p} - \beta\mathbf{I}_{\mathbf{v}+\gamma < \hat{C}^{-1}\theta}).$$

In other words, the formalism states that it is not always possible under cyber-protection budget constraints for every PFC to be functional post the adverse CI system impact due to a cyber-attack. However, cyber-protection investments made via this mechanism lowers the dysfunctional threshold of components.

4 CAN WE ACHIEVE OPTIMAL BUDGET CONSTRAINED CYBER-RESILIENCE?

In this section, we analyze whether an operational technology driven CI enterprise management can achieve cyber-resilience in an optimal fashion given a budget constraint, and performance thresholds (below which, PFC dysfunctionality is assumed) of individual CI PFCs. It is evident from the ‘defense in depth’ paradigm in cyber-security that certain ‘**crown jewels**’ (PFCs) will be given strategic cyber-protection investment importance because them being adversely affected can negatively impact a large part of a CI in the event of a cyber-attack, when compared to the other ‘jewels’. Here the term ‘optimal’ in the context of cyber-resilience implies allocating a given cyber-protection budget among *a set of CI PFCs* (nodes) in the interdependent CI PFC network that generates the maximum impact in terms of the positive allocation externalities on PFC network wide cyber-resilience. We structure this section in *two* parts: the *first* part formalizes the budget-constrained optimization problem; and the *second* part formally investigates, and provides managerial implications, on how hard it is to compute the solution to this problem in practically reasonable amount of time. The second part is extremely relevant in the context of CI system managers knowing timely enough on which PFCs (**crown jewels**) to invest a constrained cyber-protection budget.

4.1 The Optimization Problem

We formulate our budget-constrained cyber-resilience optimization problem over a cyber-protection investment portfolio γ as follows:

$$\begin{aligned} \max_{\gamma \geq 0} \quad & w(S) \\ \text{s.t.} \quad & \mathbf{1}^T \gamma \leq b \end{aligned}$$

where $\mathbf{1}$ is the all-ones vector.

Here, $f(S)$ - a cyber-protection investment impact function that is a set function outputting the $|U|$ -sized vector of influence exerted by the investment in cyber-protection on PFC set $S \subseteq U$ on each node in U (i.e., $f_u(S)$ = influence exerted on PFC node $u \in U$). We assume $f(\emptyset) = 0$. $w(S)$ - a scalar weighted sum function that outputs an importance weighting of PFC set S . In the simplest setting, each PFC is weighted by 1, i.e., each PFC given equal importance. An example of this scalar weight function is $w(S) = \sum_{i \in S} a_i v_i$ - the weighted sum of performance levels of each CI PFC in S , where $a_i > 0$, and $a_i = 1$ for all i representing the simplest (and often idealistic) setting. In practice, $a_i \neq 1$ for all i . θ is the vector of additional performance

levels for each PFC $u \in U$, induced by corresponding cyber-protection investments, that is needed to make sure that each PFC performs at or above a given threshold level of performance. b is the organizational cyber-protection budget to be allocated among PFC set S . γ_u is the individual cyber-protection investment spent on PFC $u \in U$.

The objective function reflects the weighted impact of the number of functional PFCs in the interdependent CI network, and our goal is to maximize this impact that is directly proportional to the number of functional PFCs after the NIST ensorsed absorb and adapt phase of a cyber-attack event. Here, investment portfolio γ upper bounded by organizational cyber-protection budget b , reverses the dysfunctionality of a seed set of protection-invested PFCs $S_0 \subseteq T$ for which $\gamma_{u|u \in S_0} \geq \tilde{\theta}_u$. Furthermore, one could iteratively construct sets $S_i \subseteq T$ of PFCs whose dysfunctionality is reversed by propagating the positive externality effects arising from S_{i-1} via the following expression:

$$f_u(S_{i-1}) + \gamma_u \geq \tilde{\theta}_u.$$

It is evident that the CI network is fully cyber-resilient if the latter expression holds for all components in the CI. *Note that, w.l.o.g., we (in)equate an investment metric, γ , with an additional performance level metric, $\tilde{\theta}$, for the sake of analytical tractability with the realistic assumption that investments directly translate to performance improvements.*

Relevance of Optimization Problem to Cyber-Resilience Management - The cyber-resilience optimization problem can be interpreted as the following: given an impending PFC dysfunctionality cascade following a cyber-attack event on an CPS driven CI network, how do we find an optimal cyber-protection investment portfolio, i.e., *a budget-constrained cyber-resilience ensuring policy* (CEP) γ , to limit the number of PFC dysfunctions. The CEP, apart from maximizing $w(S)$ - a positive impact measure of the total number of functional and cyber-protected PFCs post the adapt and absorb phase of a cyber-attack, will obey the relation, $f_u(S_{i-1}) + \gamma_u \geq \tilde{\theta}_u$, for all u . *This will consequently minimize the number of dysfunctional PFCs within an ICS post a cyber-attack event.*

To achieve this goal, let T be the set of CI PFCs that would become dysfunctional without cyber-protection investment. Now magnify the view to only look at effects on the PFC nodes in T , while preserving the entire network structure. In particular, define the following: (i) I_T = diagonal matrix with $I_{uu} = 1$ for $u \in T$ and 0 otherwise, and (ii) $\Psi(T)$ maps to a system on the non-zero diagonal coordinates of I_T . Essentially, $\Psi(T)$ is the $|T| \times |U|$ matrix obtained by dropping zero rows of I_T . We can apply the above map to transform the system to look at

$$\bar{\mathbf{v}} := \Psi \hat{C} (I - C)^{-1} (D\mathbf{p} - \beta \mathbf{1}_{\mathbf{v} < \theta}).$$

This mathematical transformation removes CI PFCs that don't become dysfunctional without cyber-protection, while preserving the networked connections. The idea is that among the PFCs that would be dysfunctional without cyber-protection, some of them will escape dysfunctionality through cyber-protection invested in them, in addition to the positive externalities from other cyber-protected PFCs. Their performance value would then go above the threshold value below which dysfunctionality results. In a reverse causal relation of dysfunctionality, other CI PFCs would be indirectly saved from dysfunctionality (via positive externality effects from protection-invested PFCs) because their performance would increase.

4.2 How Hard is it to Compute a CEP?

Ensuring optimal cyber-resilience within an interdependent CI PFC network necessitates computing a CEP that will be used by a CI C-suite to optimally distribute a cyber-protection budget among CI PFCs. After all, CI managers will automate the CEP evaluation process. However, it need not always be the case (thanks to theoretical computer science) that for any given PFC network and cyber-protection budget constraints, a CEP automation algorithm can output one. We have the following result stating the computational hardness of computing a CEP for our budget-constrained cyber-resilience optimization problem, given arbitrary instance of a CI network topology and an organizational cyber-protection budget.

Theorem 1 Let $(C, D, \beta, \theta, \mathbf{p})$ be an interdependent industrial control system PFC network with n PFCs, and deterministic dysfunctionality performance thresholds θ across the PFCs. Let $0 \leq \ell < \alpha \leq 1$. Suppose αn PFCs become dysfunctional (directly or indirectly) in the event of a cyber-attack. Then it is NP-hard to determine whether there exists a CEP $\gamma_i \geq 0, \forall i$, with, $\|\gamma\|_1 \leq b$ in the form of a cyber-protection investment portfolio, such that at most ℓn PFCs become dysfunctional post the investment. In other words, optimal cyber-resilience in interdependent CI PFC networks is NP-hard. The optimal approximate (CEP) solution to the budget constrained optimization problem is also NP-hard. That is, an approximately optimal CEP is inapproximable to within a constant factor in polynomial time for some worst case instance $(C, D, \beta, \theta, \mathbf{p})$.

Proof Sketch - The proof of the theorem follows from a modified adaptation of the computational hardness proof of the financial resilience optimization problem in (Klages-Mundt and Minca 2022) for financial networks. The proof involves first reducing the Independent Set problem (Kleinberg and Tardos 2006) in theoretical computer science to an instance of the cyber-resilience optimization problem in interdependent OT PFC networks. The reduction involves a reduction gadget that involves constructing a directed bipartite graph with uniform weights (corresponding to performance thresholds and cyber-protection investment interventions) on the nodes. The next step involves reducing this gadget to a cyber-resilience optimization problem instance. This is done by solving for $\tilde{\theta}_u$ in $\tilde{\mathbf{v}} = (I + C)(D\mathbf{p} - \beta \mathbf{1}_{U_1 \cup U_2}) = 0$, where U_1 and U_2 are the bipartite node sets. The solution exists as a bipartite graph is a two layer directed acyclic graph (DAG) where the ‘cross-holdings’ are the fraction of PFC (node) functionality liabilities in the first layer that are held by the PFC nodes in the second layer.

Practical Implications of the Theorem - The theorem states that it is difficult even for a computer to always guarantee that an outputted CEP for arbitrary problem input instances optimally (i.e., maximizing $w(S)$) or even sub-optimally (within a fraction of the optimal) allocates cyber-protection investments across CI PFCs to ensure that post cyber-protection investment, a maximum of $\ell \cdot n$ CI PFCs are dysfunctional when compared to αn (with $\alpha > \ell$) PFCs prior to the investment.

This is not a negative result as many might perceive it to be. It simply says that for certain restrictive input instances of n , b , and α that cannot be known apriori, achieving a CEP will take the lifetime of the universe for a computer to output a CEP as an answer to the budget-constrained optimization problem. Usually, in most of the practice space of problem input instances, the CEP is obtained fast enough by a computer to the benefit of PFC network managers. As an example, when the cyber-protection budget is zero CEP computation is provably not NP-hard. However, the challenge is that there is no way to know apriori which are the miniscule number of problem instances for which CEP computation becomes hard even for a computer. Alternatively, *ICS managers - instead of focusing on having the best cyber-resilience management solution to the optimization problem, should invest in effective heuristics that can solve the cyber-resilience problem in interdependent PFC networks ‘sub-optimally’*. One might argue that there have been related efforts on classifying the computational difficulty of optimal resource (in our case cyber-protection investments) allocation in networked settings appearing in Kempe et al. (2003), Günneç et al. (2020), Demaine et al. (2014) that directly transfer to our problem setting. However, these works do not consider the percolation of externalities within the network, hence every problem instance in Kempe et al. (2003), Günneç et al. (2020), Demaine et al. (2014) cannot be mapped to an instance of our problem.

5 GRAPH HEURISTIC SIMULATIONS TO OPTIMIZE CI NETWORK CYBER-RESILIENCE

We showed that optimizing cyber-resilience, even if approximately so, in CI with networked and interdependent PFCs is a computationally difficult problem. Hence, we have to resort to computationally tractable heuristics to achieve ‘optimal’ cyber-resilience in such networks. In this section, we perform large scale Monte Carlo simulations (10K sample path runs per setting configuration) of $\alpha_n(\mathbf{E}_n, \gamma_n)$ ($1 - \alpha_n(\mathbf{E}_n, \gamma_n)$) - the cyber-resilience metric proposed by us in Pal et al. (2024) under the influence of a strategic graph-based cyber-protection budget allocation heuristic, for random graph settings where each graph represents interconnected and interdependent components of an enterprise infrastructure.




Figure 2: A real world CPS CI (located in Boston, USA) to generate random simulation topologies.

Our simulation setting on graph topologies is exactly the same as in Pal et al. (2024). The random graphs are adapted from the underlying real-world CPS network of the thermal plant CI based in Boston, Massachusetts, USA. In other words, we abstract out a graph from the real world thermal plant architecture (see Figure 2) and randomize the structure to generate multiple graphs of different shapes and sizes that might represent other similar real world thermal plants. As essential structural elements, we sample the in-degrees and out-degrees of graph nodes (PFCs), each, from both a heavy-tailed distribution (Pareto) and a light-tailed (Normal) distribution for the sake of ensuring completeness to generating random non-tree graph topologies. The power parameter of the Pareto distribution (for the plots shown in the paper) are taken WLOG to be from a Pareto distribution having a shape parameter of 1 and a scale parameter of 10 to capture practical heavy-tailed topologies. Likewise, in the case of light-tailed topologies, the graph in and out degrees are sampled WLOG from a plot-representative Normal distribution with mean and standard deviation of 20 and 4, respectively. In addition, we simulate $\alpha_n(\mathbf{E}_n, \gamma_n)$ and the number of simulation clock time steps until the total number dysfunctional nodes stabilize to below a certain threshold for two contagion settings: one where each PFC is resilient (does not become dysfunctional w.p. 1) post cyber-attack, and one where each PFC is brittle and fails immediately w.p. 1 upon a cyber-attack.

Our proposed heuristic allocates an enterprise cyber-resilience budget strategically among PFCs according to the **Katz** centrality measure (Landherr et al. 2010; Zhan et al. 2017) of the PFCs in their network - **crown jewels** in the PFC network getting greater share of the budget. Note that the Katz centrality measure is a specific type of eigenvector centrality measure that accounts for both, the indirect PFC contacts in an interdependent PFC network (as usual of a traditional eigenvector centrality measure) and the local influence of cyber-protection as characterized by the non-eigenvector based degree centrality measure. We study how (much) strategic cyber-protection investments improve PFC network cyber-resilience when compared to a non-strategic approach where a given amount of cyber-protection investment is not topology-driven.

Observations and Analysis - We observe from Figure 3 that cyber-resilience in non-brittle networks is significantly higher from their brittle counterparts when PFC network topologies exhibit a Normal degree distribution, i.e., a light-tailed distribution. An opposite trend holds when PFC network topologies exhibit

Percentage of Initial Infected Nodes (IINs) - Katz Centrality




Figure 3: Cyber-resilience coefficient (RC) and total time to recovery ratio (TTRR) with and without heuristic intervention by varying (i) # of PFC nodes in brittle (BM) and non-brittle (non-BM) modes, (ii) PFC degree distribution, and (iii) % of IINs for Gaussian (GD) and Pareto (PD) PFC degree distributions.

a Pareto, i.e., heavy-tailed degree distribution. That is cyber-resilience in brittle networks is quite similar (and high) to that in non-brittle networks. The reason for this latter trend is the fact that differences in degree centrality among PFC nodes is more pronounced in networks with a heavy-tailed distribution than that in networks with light-tailed distributions. Hence graph heuristic driven proportionate cyber-protection allocations in networks with heavy-tailed distributions better PFC network cyber-resilience in both brittle and non-brittle networks. Such a pronounced effect is not visible for light-tailed degree distributional networks where the heuristic over-invests and under-invests over the PFC node space compared to the externalities they generate. *In other words, our proposed cyber-protection investment allocation heuristic among the PFC network nodes precisely targets those nodes whose dysfunctionality can lead to a cascading dysfunctionality effect within the PFC network post a cyber-incident.* When it comes to the median time units to adapt and absorb a cyber-incident, light-tailed degree distributional networks exhibit an oscillating performance with increase in the size of the PFC network. We observe that the performance increases with increasing node size, and then again decreases for further increased node sizes, and so on. The reason for this oscillating trend is that for certain arbitrary node sizes (varying across simulation instances) a non-strategic externality oblivious graph heuristic sufficiently 'mis-matches' the investment amount compared to the externality the PFC node generates, and this pattern repeats as we increase the size of the PFC network. On the contrary, for PFC networks with heavy-tailed degree distributions, the performance is improving with

increased node sizes. The reason for this is the fact that differences in Katz centrality among PFC nodes is more pronounced in networks with a heavy-tailed distribution than that in networks with light-tailed distributions, and the former network type escapes this oscillating performance effect.

6 SUMMARY

In this paper, we were interested in the question: *how should managers formally optimize cyber-resilience of CPS-driven critical infrastructure having networked and multiple inter-dependent process functionality components (PFCs)?* We proved via an algorithmic graph-theoretic framework that this task is NP-hard. Consequently, we proposed a computationally tractable and practical graph heuristic framework to ‘optimize’ (boost) cyber-resilience within any PFC network by strategically allocating a cyber-protection budget among PFCs in accordance with their Katz centralities in the PFC network. Alternatively, we proportionately allocate an enterprise cyber-protection budget preferring the CI **crown jewel** PFCs. We validated the efficacy of our strategic graph heuristic framework with extensive Monte Carlo simulations on real-world adapted CI topologies based upon a real-world electricity microgrid PFC network in Boston, USA.

ACKNOWLEDGEMENTS

This study has been supported by funding from Cybersecurity at MIT Sloan (CAMS) and the Agentur fÅNur Innovation in der Cybersicherheit GmbH (Cyberagentur). The Agentur fÅNur Innovation in der Cybersicherheit GmbH did not interfere in the research process and its results.

REFERENCES

- Anderson, R. 2020. *Security Engineering: A Guide to Building Dependable Distributed Systems*. Hoboken, NJ: John Wiley & Sons.
- Arghandeh, R., A. Von Meier, L. Mehrmanesh, and L. Mili. 2016. “On the Definition of Cyber-Physical Resilience in Power Systems”. *Renewable and Sustainable Energy Reviews* 58:1060–1069.
- Barker, K., J. E. Ramirez-Marquez, and C. M. Rocco. 2013. “Resilience-Based Network Component Importance Measures”. *Reliability Engineering & System Safety* 117:89–97.
- Chaves, A., M. Rice, S. Dunlap, and J. Pecarina. 2017. “Improving the Cyber Resilience of Industrial Control Systems”. *International Journal of Critical Infrastructure Protection* 17:30–48.
- Clark, A. and S. Zonouz. 2017. “Cyber-Physical Resilience: Definition and Assessment Metric”. *IEEE Transactions on Smart Grid* 10(2):1671–1684.
- Demaine, E. D., M. Hajiaghayi, H. Mahini, D. L. Malec, S. Raghavan, A. Sawant *et al.* 2014, April 7–11. “How to Influence People with Partial Incentives”. In *Proceedings of the 23rd international conference on World wide web*, 937–948. Seoul, Republic of Korea.
- DiMase, D., Z. A. Collier, K. Heffner, and I. Linkov. 2015. “Systems Engineering Framework for Cyber Physical Security and Resilience”. *Environment Systems and Decisions* 35(2):291–300.
- Elliott, M., B. Golub, and M. O. Jackson. 2014. “Financial Networks and Contagion”. *American Economic Review* 104(10):3115–3153.
- Fang, Y.-P., N. Pedroni, and E. Zio. 2016. “Resilience-Based Component Importance Measures for Critical Infrastructure Network Systems”. *IEEE Transactions on Reliability* 65(2):502–512.
- Francis, R. and B. Bekera. 2014. “A Metric and Frameworks for Resilience Analysis of Engineered and Infrastructure Systems”. *Reliability Engineering & System Safety* 121:90–103.
- Gholami, A., T. Shekari, M. H. Amirioun, F. Aminifar, M. H. Amini and A. Sargolzaei. 2018. “Toward a Consensus on the Definition and Taxonomy of Power System Resilience”. *IEEE Access* 6:32035–32053.
- Günneç, D., S. Raghavan, and R. Zhang. 2020. “Least-Cost Influence Maximization on Social Networks”. *INFORMS Journal on Computing* 32(2):289–302.
- Haque, M. A., G. K. De Teyou, S. Shetty, and B. Krishnappa. 2018. “Cyber Resilience Framework for Industrial Control Systems: Concepts, Metrics, and Insights”. In *International Conference on Intelligence and Security Informatics (ISI)*, 25–30. Miami, FL, USA. October 9–11.
- Haque, M. A., S. Shetty, K. Gold, and B. Krishnappa. 2021. “Realizing Cyber-Physical Systems Resilience Frameworks and Security Practices”. In *Security in Cyber-Physical Systems: Foundations and Applications*, edited by S. Ali, L. Khan, and Y. Jararweh, 1–37. Cham, Switzerland: Springer.

- Hosseini, S., K. Barker, and J. E. Ramirez-Marquez. 2016. "A Review of Definitions and Measures of System Resilience". *Reliability Engineering & System Safety* 145:47–61.
- Kempe, D., J. Kleinberg, and É. Tardos. 2003. "Maximizing the Spread of Influence Through a Social Network". In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 137–146. Washington, D.C., USA: ACM. Held August 24–27, 2003.
- Khan, S. and S. E. Madnick. 2021. "Cybersafety: A System-theoretic Approach to Identify Cyber-vulnerabilities & Mitigation Requirements in Industrial Control Systems". *IEEE Transactions on Dependable and Secure Computing* 19(5):3312–3328.
- Klages-Mundt, A. and A. Minca. 2022. "Optimal Intervention in Economic Networks using Influence Maximization Methods". *European Journal of Operational Research* 300(3):1136–1148.
- Kleinberg, J. and E. Tardos. 2006. *Algorithm Design*. Delhi, India: Pearson Education India.
- Kott, A. and I. Linkov. 2019. *Cyber Resilience of Systems and Networks*. 1st ed. Cham, Switzerland: Springer.
- Landherr, A., B. Friedl, and J. Heidemann. 2010. "A Critical Review of Centrality Measures in Social Networks". *Wirtschaftsinformatik* 52:367–382.
- Linkov, I., D. A. Eisenberg, K. Plourde, T. P. Seager, J. Allen and A. Kott. 2013. "Resilience Metrics for Cyber Systems". *Environment Systems and Decisions* 33(4):471–476.
- Martel, E., R. Kariger, and P. Graf. 2019. "Cyber Resilience in the Electricity Ecosystem: Principles and Guidance for Boards". Report, World Economic Forum, Geneva, Switzerland.
- Pal, R., R. X. Sequeira, and M. Siegel. 2024. "A Mathematical Theory to Quantify Cyber-Resilience in IT/OT Networks". In *2024 Winter Simulation Conference (WSC)*, 624–635 <https://doi.org/https://dl.acm.org/doi/10.5555/3643142.3643194>.
- Sterbenz, J. P., E. K. Cetinkaya, M. A. Hameed, A. Jabbar and J. P. Rohrer. 2011. "Modelling and Analysis of Network Resilience". In *Proceedings of the 2011 Third International Conference on Communication Systems and Networks (COMSNETS)*, 1–10. Bangalore, India: IEEE. Held January 4–8, 2011.
- Tarski, A. 1955. "A Lattice-Theoretical Fixpoint Theorem and its Applications". *Pacific Journal of Mathematics* 5(2):285–309.
- Venkataramanan, V., A. Hahn, and A. Srivastava. 2019. "CP-SAM: Cyber-Physical Security Assessment Metric for Monitoring Microgrid Resiliency". *IEEE Transactions on Smart Grid* 11(2):1055–1065.
- Venkataramanan, V., A. K. Srivastava, A. Hahn, and S. Zonouz. 2019. "Measuring and Enhancing Microgrid Resiliency Against Cyber Threats". *IEEE Transactions on Industry Applications* 55(6):6303–6312.
- Woods, D. D. 2015. "Four Concepts for Resilience and the Implications for the Future of Resilience Engineering". *Reliability Engineering & System Safety* 141:5–9.
- Zhan, J., S. Gurung, and S. P. K. Parsa. 2017. "Identification of Top-K Nodes in Large Networks using Katz Centrality". *Journal of Big Data* 4(1):1–19.
- Zio, E. and R. Piccinelli. 2010. "Randomized Flow Model and Centrality Measure for Electrical Power Transmission Network Analysis". *Reliability Engineering & System Safety* 95(4):379–385.
- Zuloaga, S., P. Khatavkar, L. Mays, and V. Vittal. 2019. "Resilience of Cyber-Enabled Electrical Energy and Water Distribution Systems considering Infrastructural Robustness under Conditions of Limited Water and/or Energy Availability". *IEEE Transactions on Engineering Management* 66(4):554–566.

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.

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 in 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 szejil@mit.edu.

MICHAEL SIEGEL is a Principal Research Scientist with the MIT Sloan School of Management. His primary research interest lies in cyber-security management of information systems. He is the founding co-Director of the Cybersecurity at MIT Sloan (CAMS) center within the MIT Sloan School of Management. His email is msiegel@mit.edu.