This article has been accepted for publication in IEEE Transactions on Control of Network Systems. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TCNS.2023.3272849

The Impact of Network Design Interventions on the Security of Interdependent Systems

Pradeep Sharma Oruganti, Parinaz Naghizadeh, and Qadeer Ahmed

***Abstract*— We study the problem of defending a Cyber-Physical System (CPS) consisting of interdependent components with heterogeneous sensitivity to investments. In addition to the optimal allocation of limited security resources, we analyze the impact of an orthogonal set of defense strategies in the form of network design interventions in the CPS to protect it against the attacker. We first propose an algorithm to simplify the CPS attack graph to an equivalent form which reduces the computational requirements for characterizing the defender's optimal security investments. We then evaluate four types of design interventions in the network in the form of adding nodes in the attack graph, interpreted as introducing additional safeguards, introducing structural redundancies, introducing functional redundancies, and introducing new functionalities. We identify scenarios in which interventions that strengthen internal components of the CPS may be more beneficial than traditional approaches such as perimeter defense. We showcase our proposed approach in two practical use cases: a remote attack on an industrial CPS and a remote attack on an automotive system. We highlight how our results closely match recommendations made by security organizations and discuss the implications of our findings for CPS design.**

***Index Terms*—Cybersecurity, cyber-physical systems, game theory, attack graphs.**

I. Introduction

Vulnerabilities in modern Cyber-Physical Systems (CPS) are increasingly exploited by attackers to launch sophisticated attacks on their safety-critical components. Automation, interdependence between assets in a network, and connectivity between different networks, all complicate the task of protecting the many assets within a CPS. Further, modern attacks are initiated and choreographed over multiple assets in the network, with the attackers remaining undetected for long stretches of time as they work their way to the most critical targets [1]–[3]. In response, CPS operators need to decide on an optimal allocation of their often limited security resources throughout a network by taking into account the functionality and security attributes of different components.

Given the conflicting goals of the attacker and the CPS operator, game-theoretic modeling and analysis can be used to provide insights and recommendations for the operators' optimal security decisions. In particular, there has been significant work on security games on networks for attack detection and improving network resilience [4]–[15]. Several of these works have used "attack graph" models to study attacks on interconnected CPS. The motivation for these models is that, to successfully compromise targets internal to the network, attackers generally initiate stepping-stone attacks from external nodes, and gradually work their way to the critical assets. As such, the nodes in the attack graph are used to represent the CPS assets, while the connectivity between them shows all the components that an attacker needs to (sequentially) compromise in order to reach the CPS's most critical assets. In this paper, we similarly use an attack graph model to analyze how a CPS defender can optimally deploy its security resources to best protect the CPS against an attacker.

A. Contributions and paper overview

We present two main extensions over existing works that have used an attack graph formalization to study CPS security: analyzing optimal security investments when assets have heterogeneous return-on-investment, and assessing the impacts of network design interventions. We detail each of these extensions, along with the main analytical and practical implications of considering them.

1) Modeling assets' return-on-investment: First, we extend the attack graph models studied in prior works (e.g. [7], [13], [15]) by introducing a *return-on-investment* feature, κ_i, for each asset v_i. This term, which is heterogeneous across assets, can be thought of as the rate of decrease in that asset's security risk per unit of investment. This captures realistic scenarios in which investing in some assets can provide better "bang for the buck". To the best of our knowledge, an attack graph with non-uniform node sensitivities has only been considered in [16] but with a primary focus on numerical experiments. Our work therefore extends this literature by introducing nodes' sensitivities to investments and providing an analytical study of the resulting games.

In particular, in Section III, we present an algorithm for transforming the attack graph of the resulting security game into an "equivalent" reduced form graph which considerably simplifies the computation load of identifying the optimal security investments and assessing the expected loss of the network (this is achieved by reducing the number of decision variables and constraints in the underlying minmax optimization problem). We further show that the resulting equilibrium investment strategies may recommend spreading investments

Pradeep Sharma Oruganti and Qadeer Ahmed are with the Department of Mechanical and Aerospace Engineering at The Ohio State University. Parinaz Naghizadeh is with the Departments of Integrated Systems Engineering and Electrical and Computer Engineering at The Ohio State University. Emails: {oruganti.6, ahmed.358, naghizadeh.1}@osu.edu.

P. Naghizadeh acknowledges support from NSF grant CCF-2144283.

Authorized licensed use limited to: The Ohio State University. Downloaded on June 08,2023 at 18:01:49 UTC from IEEE Xplore. Restrictions apply.
© 2023 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

This article has been accepted for publication in IEEE Transactions on Control of Network Systems. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TCNS.2023.3272849

on assets internal to the network; this is contrary to previous results in the homogeneous return-on-investment model which could only identify perimeter defense (as opposed to strengthening internal assets) and "min-cut" strategies (as opposed to spreading investments) as optimal for the defender [13], [15].

2) Assessing the impacts of design interventions: Our second contribution is to analyze an orthogonal set of defender actions in the form of *network design interventions*. In particular, in addition to optimally allocating her security budget, the defender can choose to modify the CPS by adding new nodes in the attack graph (as detailed shortly). To the best of our knowledge, the only other work considering network design interventions is [14] which looks at hiding or revealing edges of an attack graph to change an attacker's perception, while the original network is not modified.

Specifically, we focus on four possible re-design actions that can be taken by a CPS operator, which result in the introduction of additional nodes in the attack graph:

(a) Adding a node in series with existing nodes in the graph. Examples include adding an encryption device, or requiring stronger passwords.
(b) Adding a node in parallel with an existing node. Examples include adding an additional user to the CPS.
(c) A hybrid case of simultaneously adding a series and a parallel node to an existing node. Examples include adding an additional sensor to provide redundant information for anomaly detection.
(d) Adding additional input nodes. Examples include introducing an additional functionality in the system, such as adding Bluetooth connectivity to a device.

In Section IV, we consider each of these interventions when applied to a base network. We find the equilibrium outcomes of the security game on the modified attack graph, and compare the resulting expected network losses against that of the base network to elaborate on the security implications of each design intervention.

3) Numerical experiments and practical implications: In Section V, we illustrate both our attack graph reduction algorithm and our proposed design interventions in two (numerical) use cases: a remote attack on an industrial SCADA system and a remote attack on an automotive system. We also discuss the practical implications of our findings. In particular, for the SCADA system, we compare our recommended investment strategy against a perimeter defense strategy. The strategy recommended by our approach outperforms the perimeter defense strategy which matches current trends in industry practice [17]. Further, in our analysis of the automotive system which follows the penetration testing report [18], we find that our findings closely match the countermeasures recommended by security agencies. These observations indicate the potential value of our proposed framework as an analytical tool to help in strategic decision-making.

B. Related work

Our work is within the literature on using an attack graph formalization in the study of CPS security [5]–[7], [13]–[16], as detailed in Section I-A. Complimentary to these models, there exists a rich literature on *network interdiction games* as an alternative approach to the study of optimal resource allocation in networks (see [8] for a survey). In general, the attacker in a network interdiction game aims to identify the shortest path from the source nodes to the target assets, and the defender's security investments are aimed at "lengthening of the arcs" in the attack graph to thwart the attacker. One difference between network interdiction and attack graph formulations is that the former models do not tend to capture intermediate losses from the traversed assets in the attack path, i.e., a loss is incurred only when the attacker reaches the target asset through its selected (shortest) path. In contrast, an attack graph formulation allows us to model the loss from the intermediate assets, with a loss incurred even if the attacker only manages to partially progress through an attack path. We further compare our attack graph model parameters with those in network interdiction games in Section II-B.

Optimal cyber-risk management and security resource allocation has also been studied using concepts from Probabilistic Risk Analysis (PRA) in [19]–[21]. The networks considered in these works are different from stepping-stone attack graphs, in that attacks may be targeted at any individual node directly. Attack graphs models based on Bayesian graphs and Markov chains have also been used to study the overall vulnerability of IT systems in [6], [22], [23]; however, these works do not consider design interventions or optimal investment decisions.

An earlier version of our work appeared in [15] for the homogeneous return-on-investment model. We extend [15] by generalizing our reduction algorithm and design interventions to account for heterogeneous returns-on-investment. We show that the recommendations from our new model outperform the perimeter defense strategies recommended by [15] when asset sensitivities are taken into account, and use two new case studies to show that our findings are close to realistic manufacturer decisions and security agencies' recommendations.

II. The Security Game Framework

A. The attack graph

We consider a cyber-physical system (CPS) modeled as an acyclic directed attack graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$, where $\mathcal{V}$ represents the set of nodes and $\mathcal{E}$ represents the set of edges of the graph. A directed edge $(v_i, v_j) \in \mathcal{E}$ connecting node v_i to node v_j indicates that an attack on v_j can be launched once v_i is compromised. The attacks can be initiated from any of the outermost *entry* or *source* nodes of the graph, and are aiming to reach the *target* or *goal* asset. The set of entry nodes is represented as $\mathcal{V}_s \subseteq \mathcal{V}$ and the unique target node is $v_g \in \mathcal{V}$.

A path P_{ij} between nodes v_i and v_j is a sequence of connected nodes $\{v_i, v_{i+1}, \dots, v_j\}$, i.e. $P_{ij} = \{v_i \to v_{i+1} \to \dots \to v_j\}$; let $\mathcal{P}_{ij}$ denote the set of all such paths. All the nodes that can be reached from a node $v \in \mathcal{V}$ (through one or more steps and including v itself) are denoted as $Post(v)$, and all the nodes from which v (including itself) can be reached are denoted $Pre(v)$. Each node v_i is endowed with a stand-alone loss (financial or functional) $L_i \geq 0$, incurred if the node is successfully compromised. We assume $L_g > 0$, where L_g is the loss associated with the target asset v_g. Table I in Appendix I summarizes our notation.

Authorized licensed use limited to: The Ohio State University. Downloaded on June 08,2023 at 18:01:49 UTC from IEEE Xplore. Restrictions apply.
© 2023 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

This article has been accepted for publication in IEEE Transactions on Control of Network Systems. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TCNS.2023.3272849

B. The security game

We study a Stackelberg game between an attacker and a defender. The defender acts first by deploying defense resources over the nodes of $\mathcal{G}$. Let $x_i \in \mathbb{R}_{\geq 0}$ denote the security investment on node v_i. We assume that given an investment x_i, the probability of successful attack on node v_i is given by:

$$p_i(x_i) = p_i^0 e^{-\kappa_i x_i} \tag{1}$$

Here, $p_i^0 \in (0, 1]$ denotes the default probability of compromise under no investment, and $\kappa_i \geq 1$ is a node's sensitivity to investments (with higher κ indicating higher marginal benefit-on-investment). Similar models have been considered in prior works [7], [13], [15] when $\kappa_i = 1, \forall\ i$. Our work extends these works by introducing node sensitivities and providing an analytical study of the resulting games.[1]

We consider a game of full information, i.e., the attacker and defender both have knowledge of the network topology, all node attributes, and each other's utility functions and action sets. The attacker's action consists of selecting one path $P_{sg} \in \mathcal{P}_{sg}$ to initiate a sequence of attacks starting from some $v_s \in \mathcal{V}_s$ with the objective to reach and compromise the target node v_g. Assuming a worst-case attacker, its goal is to identify the path to perform stepping-stone attacks that would lead to the maximum expected loss on the CPS. In response, the defender chooses an investment profile $\mathbf{x} = [x_1, x_2, \ldots, x_{|\mathcal{V}|}]$ to minimize the loss in face of such attacker.

Formally, the defender solves the following problem:

$$\begin{aligned} &\min_{\mathbf{x}} \max_{P_{sg} \in \mathcal{P}_{sg}} \sum_{v_i \in P_{sg}} L_i \prod_{v_j \in P_{sg} \cap Pre(v_j)} p_j(x_j) \\ &\text{s.t.} \quad \sum_{i=1}^{|\mathcal{V}|} x_i \leq B \ , \ \text{and} \quad x_i \geq 0, \ \forall i \in \mathcal{V} \ . \end{aligned} \tag{2}$$

Here, B is the security budget available to the defender. We use $\mathbf{x}^*$ and $\mathbf{L}^*$ to denote the optimal solution of (2) and the expected loss under this investment profile, respectively. The solution to (2) determines the Stackelberg equilibrium strategies for the defender. We note that the objective function is strictly convex, and the feasible region is non-empty and compact; therefore, a solution to (2) exists and is unique.

Throughout our analysis, we assume the defender can place investments on nodes preceding the target node to protect it, but not on the target itself (i.e. $x_g^* = 0$). This is a mild assumption, and resembles real-life scenarios where security investments on certain components cannot be made due to reasons such as conformance to standards, functional requirements, or ownership. Additionally, we assume that the defender has access to a *sufficient* budget B, formally stated below. The proof is provided in Appendix III-A.

Lemma 1 (Sufficient budget): For any given instance of the Stackelberg game on the attack graph, there exists a *sufficient budget* B such that if the optimal investment on a non-entry node v_j is $x_j^* = 0$ at the solution of (2) given budget B, then $x_j^* = 0$ at the solution of (2) under any $B' \geq B$.

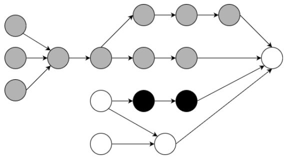
Fig. 1: An attack graph in its original form.

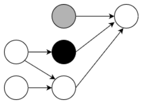
Fig. 2: Reduced attack graph of Fig. 1, with the gray/black nodes in Fig. 1 reduced to a single gray/black node.

This choice still allows us to evaluate how the defender prioritizes the expenditure of a limited budget, without considering cases in which some nodes are not attended to due to lack of resources. In other words, if a node receives zero investment in the optimal profile under sufficient budget, it will continue receiving no investment even if the defender procures more security budget. All following analysis is performed assuming sufficient budget is available.[2]

III. Attack Graph Reductions

We begin our analysis by showing that the attack graph $\mathcal{G}$ of the security game described in Section II can be transformed into an "equivalent" reduced form which considerably simplifies the optimization problem in (2). Formally, we define equivalence between attack graphs as follows.

Definition 1 (Equivalent graphs): Two attack graphs $\mathcal{G}_1$ and $\mathcal{G}_2$ are *equivalent* if they have the same expected loss $\mathbf{L}^*$ under their respective optimal strategies $\mathbf{x}_1^*$ and $\mathbf{x}_2^*$. We denote this by $\mathcal{G}_1 \equiv \mathcal{G}_2$.

Our motivation for proposing such attack graph reductions is two-fold. First, our reduction procedure leads to an equivalent attack graph with a reduced number of nodes, which simplifies problem (2) by reducing the number of decision variables. Moreover, we propose using this reduction algorithm in conjunction with our network re-design interventions presented in Section IV. Specifically, we are in general interested in evaluating whether a network re-design intervention can be effective by reducing the expected loss in the network. To this end, it is sufficient to compare the losses on the reduced forms of the attack graphs before and after the intervention. This will in turn reduce the computational requirement when assessing different candidate interventions. An illustration of our reduction procedure's outcome is shown in Figs. 1 and 2.

To see why the computational load of security assessment can be lowered by our approach, first note that through the addition of a variable, problem (2) can be converted into a minimization problem with $|\mathcal{P}_{sg}|$ inequality constraints where $\mathcal{P}_{sg}$ denotes the set of all the attack paths leading to v_g. The current optimization problem has a total of $m = n + p + 1$ constraints, were n indicates the number of assets, p

[1] Similar elements appear in shortest path network interdiction game formulations [8]. Specifically, an arc (v_i, v_j) in those models has length $c_{ij} + d_{ij} x_{ij}$ given the (typically binary) interdiction decision x_{ij}. The parameters p_i^0 and κ_i in our model are similar to c_{ij} and d_{ij} in such models.

[2] We discuss the insufficient budget case in Appendix IV.

Authorized licensed use limited to: The Ohio State University. Downloaded on June 08,2023 at 18:01:49 UTC from IEEE Xplore. Restrictions apply.
© 2023 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

This article has been accepted for publication in IEEE Transactions on Control of Network Systems. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TCNS.2023.3272849

indicates the number of paths from the source nodes to the target, and the one additional constraint indicates the budget constraint. In general, using interior-point methods, solving this optimization problem with n variables and m constraints has an overall computational complexity of $O(n^{1.5}m^3\log(\frac{1}{\epsilon}))$, where ϵ indicates the degree of accuracy within which the solution is obtained [24], [25]. Accordingly, a reduction in the number of variables (together with the resulting reduction in the number of attack paths) will reduce the computation time to solve (2). We provide numerical examples in Section V.

We present our reduction procedure as a number of subroutines to be applied to series paths, parallel paths, and input nodes in the attack graph. All proofs in this section are presented in the online appendix of the extended version of this paper [26], and show that the attained reduced graph following each subroutine leads to an equivalent graph (in the sense of Definition 1) to the original graph.

A. Series path reductions

We first present attack graph reductions which ultimately replace any series path $\{v_i \to v_{i+1} \to \cdots \to v_{i+n}\}$ with a single equivalent node. Formally, we say $\{v_i \to v_{i+1} \to \cdots \to v_{i+n}\}$ is a series path if $Pre(v_{k+1}) = \{v_i, v_{i+1}, \dots v_k, v_{k+1}\}$ and $Post(v_k) = \{v_k, v_{k+1}, \dots, v_{i+n}\}$ for all $k \in \{i, \dots, i+n-1\}$. For readability, the default loss p_i^0 for individual nodes is dropped in the remainder of this section. This is without loss of generality, as they can be subsumed in the stand-alone loss L_i of the node.

We begin by identifying nodes that will receive a zero investment at the optimal equilibrium profile, and show that these can be subsumed in their preceding nodes to obtain an equivalent attack graph.

Lemma 2 (Series zero investments): Consider a series link $\{v_i \to v_{i+1} \to \cdots v_{i+n-1} \to v_{i+n}\}$.

- Start at $m = i+n$. The pair of nodes $v_{m-1} \to v_m$ can be replaced with a node with $\kappa_{eq} = \kappa_{m-1}$ and $L_{eq} = L_{m-1} + L_m$, and $x_m^* = 0$, if and only if
$$L_{m-1} \geq (\frac{\kappa_m}{\kappa_{m-1}} - 1)L_m \ . \tag{3}$$
- If (3) is not satisfied, for $i < m \leq i+n-1$, the pair of nodes $v_{m-1} \to v_m$ can be replaced with a node with $\kappa_{eq} = \kappa_{m-1}$ and $L_{eq} = L_{m-1} + L_m$, and $x_m^* = 0$, if and only if
$$\kappa_{m-1} \geq \kappa_m \text{ or } L_{m-1} \geq \Bigg(\frac{\frac{\kappa_{m-1}}{\kappa_{m-2}} - 1}{1 - \frac{\kappa_{m-1}}{\kappa_m}}\Bigg)L_m \ . \tag{4}$$

Intuitively, under the conditions in the lemma, either the earlier node v_i provides higher marginal return on investment, or has a substantially higher stand-alone loss than the subsequent node v_j, and as such, the defender is better off adopting "perimeter defense" and investing all budget on the outer node v_i. Note that this finding is consistent with previous results in [13], [15], [16], which had studied the special case of $\kappa_i = \kappa_j$. Lemma 2 further extends these results as it identifies conditions under which the defender distributes her investments over inner nodes as well.

Note that by repeated application of Lemma 2, working our way from the last node backwards to the first node, we can convert any series path $\{v_i \to v_{i+1} \to \cdots \to v_{i+n}\}$ to a reduced form in which all remaining nodes should have non-zero investments at the optimal investment profile. Following this, we conduct the remaining series reduction as detailed below, which will result in any series path being replaced by a single equivalent node.

Lemma 3 (Series reduction): Consider a series path $\{v_i \to v_{i+1} \to \cdots \to v_{i+n} \to v_t\}$. Assume that $x_j^* \neq 0$ at all of these nodes. Then, this path can be replaced by a single equivalent node v_{eq} with $\kappa_{eq} = \kappa_i$ and stand-alone loss
$$L_{eq} = \frac{\kappa_{i+1}L_i}{\kappa_{i+1} - \kappa_i}\prod_{j=i+1}^{i+n}\Big(\frac{\kappa_{j-1}}{\kappa_{j+1}}\Big(\frac{\kappa_{j+1} - \kappa_j}{\kappa_j - \kappa_{j-1}}\Big)\frac{L_{j-1}}{L_j}\Big)^{\frac{-\kappa_i}{\kappa_j}} \Big(\frac{\kappa_{i+n}}{\kappa_t - \kappa_{i+n}}\Big(\frac{L_{i+n}}{L_t}\Big)\Big)^{\frac{-\kappa_i}{\kappa_{i+n}}} \ . \tag{5}$$

B. Parallel path reductions

After performing the proposed series link reductions, the reduced attack graph can contain parallel paths of the form $\{(v_i \to v_{i+1} \to v_j), (v_i \to v_{i+2} \to v_j), \dots, (v_i \to v_{i+n} \to v_t)\}$, where $Pre(v_{i+j}) = \{v_i\}$ and $Post(v_{i+j}) = \{v_t\}$, for all $j \in I := \{1, 2, 3 \dots, n\}$. In this section, we identify scenarios under which parallel paths of this form can be replaced by a single equivalent node v_{eq}.

Similar to the series reduction case, we first identify cases in which we can determine, a priori, if one or more of the parallel nodes should receive zero investment under the optimal investment strategy, and can therefore remove them prior to solving for the optimal investment profile.

Lemma 4 (Parallel zero investments): Consider a set of parallel paths $\{(v_i \to v_{i+1} \to v_g), (v_i \to v_{i+2} \to v_g), \dots, (v_i \to v_{i+n} \to v_g)\}$ with $L_{i+1} < L_{i+2} < \cdots < L_{i+n}$. Let $\kappa_{par} := \sum_{k \in I}\frac{1}{\kappa_k}$. Then, $x_{i+1}^* = 0$ if and only if
$$L_i \geq \Big(\frac{1 - \kappa_i\kappa_{par}}{\kappa_i\kappa_{par}}\Big)(L_{i+1} + L_g) \tag{6}$$

Intuitively, the above lemma can be interpreted as follows. It may arise that due to the security attributes of the parallel nodes (which follow (6)), the optimal action is to equate the losses across all paths to that of $\mathbf{L}_l = L_i + L_l + L_g$ where v_l is the loss the parallel node with the least stand-alone loss L_l. This lemma can be applied repeatedly: with v_{i+1} receiving no investment, we can check Lemma 4 over the remaining nodes until no additional reductions of this type are possible. Note also that in the special case when $\kappa_i = \kappa, \forall i$, through repeated application of Lemma 4, the set of paths will be replaced with the single path containing the parallel node with the highest stand-alone loss; this matches our earlier results in the homogeneous κ model [15].

Following repeated application of Lemma 4, all remaining parallel nodes in sets of the form $\{(v_i \to v_{i+1} \to v_g), (v_i \to v_{i+2} \to v_g), \dots, (v_i \to v_{i+n} \to v_g)\}$ will receive non-zero investments. These sets can be further replaced by a single equivalent node, as shown in the following lemma.

Lemma 5 (Parallel reduction): Consider a set of parallel paths, $\{(v_i \to v_{i+1} \to v_g), (v_i \to v_{i+2} \to v_g), \dots, (v_i \to$

Authorized licensed use limited to: The Ohio State University. Downloaded on June 08,2023 at 18:01:49 UTC from IEEE Xplore. Restrictions apply.
© 2023 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

This article has been accepted for publication in IEEE Transactions on Control of Network Systems. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TCNS.2023.3272849

$v_{i+n} \to v_g)\}$ with $\kappa_{par} = \sum_{r=i+1}^{i+n} \frac{1}{\kappa_r}$ such that the conditions of Lemma 4 are not satisfied. Then $\mathcal{G}$ can be replaced with a single equivalent node v_{eq} with $k_{eq} = k_i$ and

$$L_{eq} = \frac{L_i}{1-\kappa_i\kappa_{par}} \prod_{j=i+1}^{i+n} \left(\frac{\kappa_i\kappa_{par}}{1-\kappa_i\kappa_{par}} \left(\frac{L_i}{L_j+L_g} \right) \right)^{\frac{-\kappa_i}{\kappa_j}} \quad (7)$$

It is to be noted that the outcome of these steps can lead to parallel paths being reduced to series links. In such a scenario, further reduction of the attack graph can be achieved by looping between the procedures in Lemmas 2-5, until no further series or parallel reduction is possible.

C. Input node reductions

Finally, we look at a possible reduction of multiple input nodes. Similar to the previous sections, we begin by providing a condition on input nodes given which one can a priori guarantee that their first successor node, v_t, will receives a zero investment in the optimal investment profile.

Lemma 6: Consider a set of paths $\mathcal{G} = \{(v_i \to v_t \to v_g), (v_{i+1} \to v_t \to v_g), \cdots, (v_{i+n} \to v_t \to v_g)\}$ with $L_j = L, \forall j \in \mathcal{I} = \{i, i+1, i+2, \ldots, i+n\}$. Let $\kappa_{par} = \sum_{r=i}^{i+n} \frac{1}{\kappa_r}$. Then, $x_t^* = 0$ if and only if $\kappa_t\kappa_{par} \leq 1$ or $L \geq (1-\kappa_t\kappa_{par})(L_t + L_g)$.

Similar to the result in [15], this lemma states that the defender is better off choosing a "perimeter defense" if the stand-alone loss v_t is substantially lower than the input nodes; it further extends that result by showing that the same is true if the inner node v_t has a relatively lower return-on-investment.

We now look at the case when the conditions of Lemma 6 are not satisfied (i.e., $x_t^* \neq 0$), and show that multiple input nodes can be replaced by a single equivalent node.

Lemma 7 (Source node reduction): Consider a set of input paths $\{(v_i \to v_t), (v_{i+1} \to v_t), \ldots, (v_{i+n} \to v_t)\}$ with equal stand-alone loss input nodes, i.e., $L_{i+1} = \cdots = L_{i+n} = L$. Assume the conditions of Lemma 6 are not met. Then, this set can be replaced with a single equivalent node v_{eq} such that

$$L_{eq} = \frac{L\kappa_i\kappa_{par}}{\kappa_i\kappa_{par}-1} \left(\frac{1}{\kappa_i\kappa_{par}-1} \left(\frac{L}{L_t} \right) \right)^{\frac{-1}{\kappa_i\kappa_{par}}} \quad (8)$$

where $k_{par} = \sum_{r=i}^{i+n} \frac{1}{k_r}$ and $k_{eq} = \frac{1}{k_{par}}$.

D. Reduction algorithm

We now present our proposed attack graph reduction procedure in Algorithm 1. The statement and proof of Proposition 1 are based on the sequence of Lemmas 2-7 presented earlier. This proposition generalizes our earlier work [15] as well as related results in prior works [13], [16].

Proposition 1: Given a sufficient budget B, Algorithm 1 leads to an equivalent reduced form $\mathcal{G}_r$ of attack graph $\mathcal{G}$.

In Appendix III, we further detail how the optimal investments obtained from the reduced graph $\mathcal{G}_r$ can be mapped back to the optimal investments on the original graph $\mathcal{G}$.

Algorithm 1 Reduction of attack graph $\mathcal{G}$ to an equivalent $\mathcal{G}_r$

```
Input: An attack graph $\mathcal{G}$
Output: An equivalent reduced attack graph $\mathcal{G}_r$
while series or parallel paths reducible do
    Series Paths Reduction Step:
    Gather all series paths in graph $\mathcal{G}$
    Lemma 2: remove series nodes with no investment
    Lemma 3: reduce remaining series links to one node
    Update the set of series paths in $\mathcal{G}$ accordingly
    Parallel Paths Reduction Step:
    Gather all parallel paths in graph $\mathcal{G}$
    Lemma 4: remove (some) parallel paths
    Lemma 5: reduce (some) parallel paths to one node
    Update the set of parallel paths in $\mathcal{G}$ accordingly
end while
Input Node Reduction Step:
Consider all input nodes in graph $\mathcal{G}$
Apply Lemma 6 if possible, else Lemma 7, to remove
  (some) of the input nodes
Update the set of input nodes in $\mathcal{G}$ accordingly
return reduced graph $\mathcal{G}_r$
```

IV. Network Design Interventions

While the attacker-defender games of Section II have been studied in a number of prior works in the homogeneous return-on-investment case (e.g., [7], [13]), their focus, similar to the analysis presented in Section III, has been on the study of the optimal investment strategy *given* a fixed network. In addition to extending these models by considering heterogeneous return-on-investments κ, this paper further evaluates the use of an orthogonal set of defender actions, in the form of *network design interventions*.

To illustrate the main ideas, we consider a minimal base network and four re-design actions as illustrated in Fig. 4: (a) adding a node in series; (b) adding a node in parallel; (c) a combination of series and parallel additions; and (d) adding an new entry node. We compare the overall loss on the networks obtained through these actions against those of the base network, and provide (intuitive) interpretations for the potential effects of each type of intervention.

These four types of interventions can be made in any general CPS. As mentioned earlier, our reduction approach in Section III can be used to simplify the task of comparing the expected losses following these interventions. In Section V, we will elaborate on the effect of these interventions in more general networks, and show that they match the intuitions obtained from the analysis of the base network in this section, using numerical examples motivated by applications in industrial cyber-physical systems.

The Base network: Consider the minimal attack graph shown in Fig. 3. This attack graph is minimal in the sense that the target node v_g is an interior node of the network, accessible only through a stepping-stone attack by compromising the entry node (v_1) as well as an intermediate node (v_2).

For simplicity, we let $p_i^0 = p$, $\forall i$ for the analytical results in this section; in the numerical illustrations, we additionally set

Authorized licensed use limited to: The Ohio State University. Downloaded on June 08,2023 at 18:01:49 UTC from IEEE Xplore. Restrictions apply.
© 2023 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

This article has been accepted for publication in IEEE Transactions on Control of Network Systems. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TCNS.2023.3272849

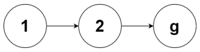

Fig. 3: A minimal base network.

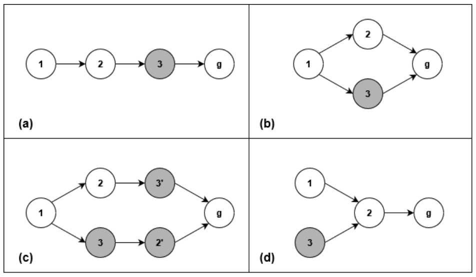

Fig. 4: Network design intervention alternatives.

$p = 1$ and highlight the impact of other problem parameters. We also assume the problem parameters are such that the conditions of Lemma 2 are not met, so that both nodes v_1 and v_2 receive non-zero investments at equilibrium. The expected loss for the base network in this case is:

$$\mathbf{L}^*_{\text{b}} = \frac{L_1\kappa_2 p}{\kappa_2 - \kappa_1}\left(\frac{\kappa_1}{\kappa_2 - \kappa_1}\left(\frac{L_1 p}{L_2p^2 + L_g p^3}\right)\right)^{\frac{-\kappa_1}{\kappa_2}} e^{-k_1 B}$$

We will next assess which design interventions can help lower this expected loss.

1) Series connection: increased endurance: The first intervention we consider is that of an addition of a node in series, as illustrated in Fig. 4(a). Security interventions in the form of adding encryption devices or requiring (stronger) passwords can be represented as this type of network redesign. Intuitively, we might expect that the addition of a series node will increase the endurance of the system as the attacker now has to compromise an extra node to get to the target.

We consider two possibilities for such interventions: closer to the target node (strengthening the core of the network) vs. closer to the entry node (strengthening perimeter defenses). In the former case when v_3 is added after v_2, if the conditions of Lemma 2 apply on v_3, then $x^*_3 = 0$. The loss of the modified network then will be similar to $\mathbf{L}^*_{\text{b}}$, but with L_2 replaced by $L_2 + L_3$ and L_g multiplied with p^4. As $\mathbf{L}^*_{\text{b}}$ is increasing in L_2 (since $\frac{\partial \mathbf{L}^*_{\text{b}}}{\partial L_2} > 0$) this means that while the series addition does reduce the probability of attack on the downstream node (v_g here), there may arise scenarios were this may not offset the increase in loss of the node immediately upstream (v_2 here). This means that, perhaps counter intuitively, attempts at "strengthening" the core of the network with components with lower return-on-investment or higher safety criticality (higher L) tends to backfire and increase the total loss.

Next, we look at the case when the conditions of Lemma 2 for v_3 are not met, which implies $x^*_3 \neq 0$. Depending on whether v_3 is introduced before (pre) or after ($post$) v_2, the equilibrium expected losses obtained are:

$$\mathbf{L}^*_{\text{srs, post}} = \frac{L_1\kappa_2 p}{\kappa_2 - \kappa_1}\left(\frac{\kappa_1}{\kappa_3}\left(\frac{\kappa_3 - \kappa_2}{\kappa_2 - \kappa_1}\right)\frac{L_1 p}{L_2 p^2}\right)^{\frac{-\kappa_1}{\kappa_2}}$$
$$\left(\left(\frac{\kappa_2}{\kappa_3 - \kappa_2}\right)\frac{L_2 p^2}{L_3 p^3 + L_g p^4}\right)^{\frac{-\kappa_1}{\kappa_3}} e^{-k_1 B}$$

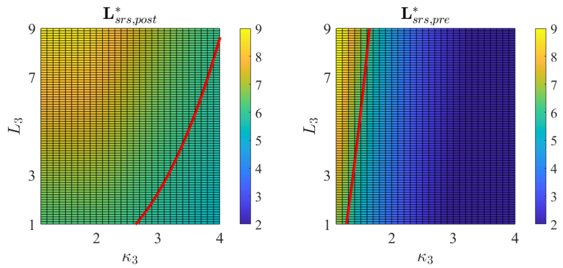

Fig. 5: Adding v_3 after (left) and before (right) v_2 in the base network. The red line indicates the level curve of the base expected loss $\mathbf{L}^*_b = 6$. All values to the left of the curve indicate a higher loss compared to the base.

$$\mathbf{L}^*_{\text{srs, pre}} = \frac{L_1\kappa_3 p}{\kappa_3 - \kappa_1}\left(\frac{\kappa_1}{\kappa_2}\left(\frac{\kappa_2 - \kappa_3}{\kappa_3 - \kappa_1}\right)\frac{L_1 p}{L_3 p^2}\right)^{\frac{-\kappa_1}{\kappa_3}}$$
$$\left(\left(\frac{\kappa_3}{\kappa_2 - \kappa_3}\right)\frac{L_3 p^2}{L_2 p^3 + L_g p^4}\right)^{\frac{-\kappa_1}{\kappa_2}} e^{-k_1 B}$$

We compare these expected losses against that of the base network numerically. In Fig. 5 we fix the values of L_1, L_2, and L_g and compare the resulting loss L^*_b (indicated using the red curve) against the losses $\mathbf{L}^*_{\text{srs, pre}}$ and $\mathbf{L}^*_{\text{srs, post}}$ as a function of L_3 and κ_3 of the added node. First, we note that in both cases, the total expected loss increases with L_3, meaning that the added node should itself have a low stand-alone loss for the expected loss to decrease relative to the base network. Further, it can be seen that for a v_3 added downstream (closer to the target) to lower the expected loss relative to the base network, it has to be a node with a relatively *high* return-on-investment. Adding the same node upstream would lead to a decrease in expected loss at lower κ_3.

Impacts of interventions on expected loss vs. attack probability on individual nodes: We note that there may exist a trade-off between minimizing total expected loss and the probability of attack on a given node. Figure 6a illustrates a scenario where the attack probability on v_2 (given by $p^*_2 = p^0_1 p^0_2 e^{-\kappa_1 x^*_1 - \kappa_2 x^*_2}$) changes depending on the security attributes of a node v_3 added immediately downstream. It can be seen that as κ_3 increases, higher security investments on v_3 lead to lower total expected losses, while the reduced investment on v_2 leads to a higher probability of attack on that node. That is, the designer opts to make interventions that decrease overall loss, despite the (negative) impacts it may have on some of the individual assets.

In summary, we observe that the addition of a series node v_3 can help strengthen the network if the added node has sufficiently low stand-alone loss L_3 and sufficiently high return-on-investment κ_3, with the benefits being higher if it is feasible to add the node closer to the perimeter of the network.

2) Parallel connection: structural/physical redundancy: Next we study the addition of a parallel connection, as illustrated in Fig. 4(b). This intervention can be seen as improving the number of redundant components in the system to improve its tolerance to physical failures; examples include, adding redundant communication lines, back-up generators, etc.

We split our analysis into two scenarios: 1) when the condition of Lemma 4 is satisfied, and 2) when it is not. In

Authorized licensed use limited to: The Ohio State University. Downloaded on June 08,2023 at 18:01:49 UTC from IEEE Xplore. Restrictions apply.
© 2023 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

This article has been accepted for publication in IEEE Transactions on Control of Network Systems. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TCNS.2023.3272849

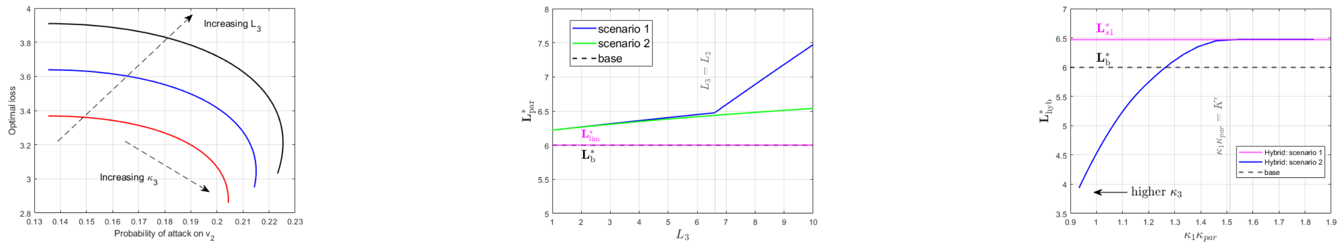

(a) Trade-off between p_2^* and $\mathbf{L}^*_{\text{srs,post}}$ after a series intervention

(b) Expected loss vs. L_3; parallel intervention (Fig. 4(b))

(c) Expected loss vs. κ_3; hybrid intervention (Fig. 4(c))

Fig. 6: Expected loss in different numerical scenarios studied

scenario 1, the node with the lower stand-alone loss between v_2 and v_3 receives no investment. The total expected loss here depends on the relative value of L_2 and L_3:

$$\mathbf{L}^*_{\text{par}} = \begin{cases} (L_1p + L_3p^2 + L_gp^3)\left(\frac{L_3+L_gp}{L_2+L_gp}\right)^{\frac{-\kappa_1}{\kappa_2}} e^{-\kappa_1 B}, \\ \text{if } L_2 \geq L_3; \\ (L_1p + L_2p^2 + L_gp^3)\left(\frac{L_2+L_gp}{L_3+L_gp}\right)^{\frac{-\kappa_1}{\kappa_3}} e^{-\kappa_1 B}, \\ \text{if } L_2 < L_3. \end{cases}$$

In this scenario, since $\kappa_2 > \kappa_1$ (assumption from the base case), even for low L_3, $\mathbf{L}^*_{\text{par}} > \mathbf{L}^*_{\text{b}}$. Additionally, since $\frac{\partial \mathbf{L}^*_{\text{par}}}{\partial L_3} > 0$, the loss increases with increasing L_3. Increasing κ_3 (until the condition of Lemma 4 satisfied) has no effect when $L_2 \geq L_3$, and a limited effect in reducing loss when $L_2 < L_3$. This means that overall, the addition of v_3 in this scenario tends to increase the expected loss relative to the base case.

In scenario 2, letting $\kappa_{par} = \frac{1}{\kappa_2} + \frac{1}{\kappa_3}$:

$$\mathbf{L}^*_{\text{par}} = \frac{L_1p}{1-\kappa_1\kappa_{par}}\left(\frac{\kappa_1\kappa_{par}}{1-\kappa_1\kappa_{par}}\left(\frac{L_1p}{L_2p^2+L_gp^3}\right)\right)^{\frac{-\kappa_1}{\kappa_2}} \left(\frac{\kappa_1\kappa_{par}}{1-\kappa_1\kappa_{par}}\left(\frac{L_1p}{L_3p^2+L_gp^3}\right)\right)^{\frac{-\kappa_1}{\kappa_3}} e^{-\kappa_1 B}.$$

Figure 6b numerically illustrates the effect of varying L_3 on $\mathbf{L}^*_{par}$ in both scenarios, with the other nodes fixed at the same attributes as in Fig. 5, and κ_3 is chosen such that it is the lowest value at which both v_2 and v_3 receive investment in scenario 2. Firstly, we observe that the parallel case has a higher total expected loss than the base, irrespective of the scenario. This is to be expected (in general) as the number of paths to the target have increased and the designer has to potentially split investments over multiple paths. Similar to the series case, it can be shown that the total expected loss keeps increasing with increasing L_3 and decreases with increasing κ_3. As κ_3 increases, the investment on v_3 decreases as it gives more return-on-investment. In the limiting case (very high κ_3), v_3 receives very little investment with all the budget going on v_2 to balance losses across both paths. Hence the total expected loss in the limiting case would be $\mathbf{L}^*_{\text{limit}} = L_1p + (L_2p^2 + L_gp^3)e^{-\kappa_2 B}$, so that $\mathbf{L}^*_{\text{limit}} = \mathbf{L}^*_{\text{b}}$.

In summary, adding a structurally redundant node v_3 does not in general improve the security posture compared to the base architecture. At best, with (very) high return-on-investment κ_3 or low stand-alone loss L_3, the total expected loss remains close to the base case. Overall, adding such redundancies can improve operational reliability, but increases the attack surface (and hence expected loss) in the CPS.

3) Hybrid connection: functional/informational redundancy: Our next intervention assesses the impact of introducing functional redundancy in a system. A functionally redundant component can be used to carry out the same tasks as an existing node; for instance an additional sensor can be added to attain signals for health monitoring, or anomaly detection and isolation. While functioning independently, information from such components is generally used in unison for decision making. As a result, a successful attacker would need to compromise *both* nodes (as least to some extent) to proceed in the stepping stone attack towards v_g. To capture this, we add the nodes $v_{2'}$ and $v_{3'}$ after v_3 and v_2, respectively.

Considering the similarity in the type of nodes v_2 and v_3, we assume $L_2 = L_3 = L$, and let the auxiliary nodes have zero loss, i.e., $L_{2'} = L_{3'} = 0$. We again denote $\kappa_{par} = \frac{1}{\kappa_2} + \frac{1}{\kappa_3}$. The defender equalizes the expected losses over both paths, leading to $x_2^*\kappa_2 = x_3^*\kappa_3$. Here, $x_2^* \neq 0$ and $x_3^* \neq 0$ when

$$L_1 < \left(\frac{1}{\kappa_1\kappa_{par}} - 1\right)Lp + \left(\frac{2}{\kappa_1\kappa_{par}} - 1\right)L_gp^3 .$$

Similar to the other sections, we divide our analysis into two scenarios: 1) when $\kappa_1\kappa_{par} > 2$ and 2) when $\kappa_1\kappa_{par} < 2$. In scenario 1, it can be seen that the above condition fails irrespective of the other security attributes leading to $x_2^* = x_3^* = 0$. Similar to the previous sections, if the stand-alone loss of the root of the paths is significantly more than that of the other nodes, all investment is made on the root. Since all of the parallel nodes have the same stand-alone loss, all of them receive zero investment. In this case, we get the total expected loss to be $\mathbf{L}^*_{\text{s1}} = (L_1p + Lp^2 + L_gp^4)e^{-\kappa_1 B}$ which does not depend on κ_3. Note that unlike the parallel case, L_g is multiplied with p^4 instead of p^3. This reflects the additional step that the attacker must perform to compromise the target. As a result, the addition of such functional redundancies (hybrid nodes) can *reduce* loss compared to the base network.

We next illustrate the effect of varying κ_3 numerically, in both scenarios, in Figure 6c. Decreasing $\kappa_1\kappa_{par}$ (increasing κ_3) beyond $K' = \frac{Lp+2L_gp^3}{L_1+Lp+L_gp^3}$ leads to decreasing $\mathbf{L}^*_{\text{hyb}}$. This means that unlike the parallel case, the hybrid architecture allows a significant reduction in total expected loss at higher κ_3. This is expected: although new paths to the target are included and it may seem that the the attack surface has

Authorized licensed use limited to: The Ohio State University. Downloaded on June 08,2023 at 18:01:49 UTC from IEEE Xplore. Restrictions apply.
© 2023 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

This article has been accepted for publication in IEEE Transactions on Control of Network Systems. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TCNS.2023.3272849

increased, each path is more robust and harder to compromise, since the information from all paths is fused at the target.

In summary, the addition of functionally/informationally redundant nodes can decrease the expected loss. This is because the probability of compromise *at* the target node is reduced with the addition of the hybrid node due to the additional series components included in these nodes. Further improvements can be achieved if the return-on-investment of the additional node v_3 also has high return-on-investment κ_3.

4) Additional input node: new features: We finally look at the effect of adding entry nodes to the network, illustrated in Figure 4(d). This intervention represents scenarios involving adding additional features or functionalities; e.g., adding Bluetooth, wireless connectivity, etc., to improve user experience. Considering equal stand-alone losses L for the input nodes, we first consider the case when the conditions of Lemma 6 are satisfied. Informally, this happens when the return-on-investment or the stand-alone loss of the entry nodes is sufficiently higher than the next node v_2 and hence $x_2^* = 0$. In this case, the total expected loss is

$$\mathbf{L}_{\text{inp}}^* = (Lp + L_2p^2 + L_gp^3)e^{\frac{-B}{\kappa_{inp}}}$$

where $\kappa_{inp} = \frac{1}{\kappa_1} + \frac{1}{\kappa_3}$. When the conditions of Lemma 6 are not met, the total expected loss is

$$\mathbf{L}_{\text{inp}}^* = \frac{L\kappa_2\kappa_{inp}}{\kappa_2\kappa_{inp}-1}\left(\frac{1}{\kappa_2\kappa_{inp}-1}\left(\frac{Lp}{L_2p^2+L_gp^3}\right)\right)^{\frac{-1}{\kappa_2\kappa_{inp}}} e^{\frac{-B}{\kappa_{inp}}}.$$

It can be shown (analytically and numerically) that both effective losses are higher compared to the base case. This follows intuition, since adding more functionalities only leads to a larger attack surface, without providing any downstream benefits as all prior attack paths remain unaffected.

V. Applications

In this section, we provide numerical experiments to illustrate our previous analysis in two applications: a remote attack on an industrial SCADA system, where the goal of the attacker is to maliciously control physical actuators, and an attack on an automotive system, where the goal of the attacker is to remotely access the Controller Area Network (CAN) to send malicious commands. All computations are done using the CasADi optimization toolbox [27] with the IPOPT solver on a generic laptop using an Intel i7 CPU @ 2.8 GHz.

A. Remote attack on a SCADA system

We look at an attack on an industrial CPS discussed in [28]. Here the attacker tries to obtain the control of actuators by performing the following steps: 1) gaining administrative privileges; 2) bypassing DMZ firewalls; and 3) gaining access to an industrial PLC. A simplified version of the attack graph is shown in Fig. 7. The descriptions and security attributes of each node are given in Table II in Appendix II. We utilize the attack scoring mechanisms, the Common Vulnerability Scoring System (CVSS) [29] to quantify p^0. We quantify the respective stand-alone losses L for each node subjectively based on relative critically to the safety of the system.[3]

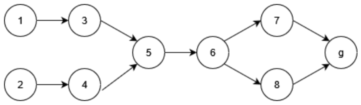

Fig. 7: Attack graph for a remote attack on an industrial CPS

We first solve (2) for this attack graph. This problem has 9 variables, 12 inequality constraints, and 1 equality constraint, and it took around 0.030s to compute the optimal solution. With a budget of $B = 5$ units, the optimal investment strategy is $\mathbf{x}_{\text{scada}}^* = \{1.4689, 1.4689, 0, 0, 2.0447, 0, 0, 0.0174\}$ and the expected loss is $\mathbf{L}_{\text{scada}}^* = 586.67$.

Next, we apply our reduction algorithm from Proposition 1 to first reduce the attack graph, and then find the optimal solution. The reduction steps are detailed below:

- First, we identify the two series links $\{v_1 \to v_3\}$ and $\{v_2 \to v_4\}$. We observe that $\kappa_3 = \kappa_1$ and $\kappa_4 = \kappa_2$. From Lemma 2, we immediately obtain $x_3^* = x_4^* = 0$. Nodes $\{v_1 \to v_3\}$ and $\{v_2 \to v_4\}$ can be replaced with equivalent nodes v_{13} and v_{24}.
- Next, we identify the parallel links $\{v_6 \to v_7 \to v_g\}$ and $\{v_6 \to v_8 \to v_g\}$. Since $\kappa_6\kappa_{78} > 1$, where $\kappa_{78} = \frac{1}{\kappa_7} + \frac{1}{\kappa_8}$, from Lemma 4 we obtain $x_7^* = 0$.
- Next we identify the series link $\{v_5 \to v_6 \to v_8\}$ and observe that $\kappa_5 = \kappa_6$. From Lemma 2, we get $x_6^* = 0$ followed by a reduction of this link to a single node v_{56}.
- We finally look at the input nodes $\{v_{13} \to v_5\}$ and $\{v_{24} \to v_5\}$. Let $L_{13} = L_{24} = L_{12}$ and $\kappa_{12} = \frac{1}{\kappa_1} + \frac{1}{\kappa_2}$. We observe that both the conditions from Lemma 6 are not satisfied, i.e. $\kappa_5\kappa_{12} > 1$ and $L_{12} > (\kappa_5\kappa_{12} - 1)L_{56}$. From Lemma 6 we get $x_5^* \neq 0$ and from Lemma 7, nodes $\{(v_{13} \to v_{56}), (v_{24} \to v_{56})\}$ can be replaced with a single equivalent node v_{in}.

Using this procedure, the attack graph in Fig. 7 is reduced to $v_{in} \to v_8 \to v_g$. The optimal loss on this network can again be found to be $\mathbf{L}_{\text{scada}}^* = 586.67$. This optimization problem now has only 3 variables, 3 inequality constraints, and 1 equality constraint, and took around 0.02s to solve. This is a 50% reduction in the computation time. While the absolute speed improvement seems minimal in this small problem, for much larger systems (for example, large-scale electric grids), we will get a significant improvement in absolute runtime as well.

Additionally, we look at how our obtained optimal defense strategy compares with a baseline "perimeter defense" strategy. Such a perimeter defense strategy is similar to the ones recommended by [15] and [13] (where such strategies are shown to be optimal in the homogeneous κ case and with rational defenders). Under such a strategy, only input nodes v_1 and v_2 would receive an investment of 2.25 each. The total expected loss in this case would be 2.09×10^4, which is much larger than $\mathbf{L}^*$. This highlights how accounting for the heterogeneity of assets in their return-on-investment can substantially impact optimal defense strategies.

B. Remote attack on an automotive system

1) System overview and quantification: Next, we consider the remote attack on an automotive Controller Area Network

[3] Note that security quantification generally involves significant expert judgement; the values here are chosen to showcase the proposed methodology.

Authorized licensed use limited to: The Ohio State University. Downloaded on June 08,2023 at 18:01:49 UTC from IEEE Xplore. Restrictions apply.
© 2023 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

This article has been accepted for publication in IEEE Transactions on Control of Network Systems. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TCNS.2023.3272849

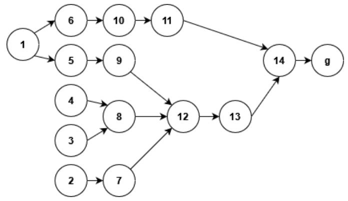

Fig. 8: Attack graph for an attack on an automotive system

(CAN), similar to the ones reported in [30] and [18]. Modern automotive systems provide many connectivity features such as Wi-Fi, Bluetooth (BT), cellular connectivity (CELL), and physical ports such as USB ports on the infotainment module. OBD-II ports (legally mandated for vehicle diagnostics) allow for a physical connection to the vehicle. In the worst-case, vulnerabilities in the communication protocols, firmware of the Electronic Control Units (ECUs), and software vulnerabilities on the infotainment module may allow access to the CAN and control of certain safety-critical actuators.

Once notified of a vulnerability or a flaw, the manufacturer is faced with making a decision on (optimal) resource allocation for immediate as well as long-term security of the vehicle [18]. This decision is complicated further by: 1) the number of vulnerabilities with varying impacts and complexity, and 2) fleet exposure based on hardware and software version combinations. We apply our proposed approach on this problem to showcase its benefits and possible insights in such scenarios. The simplified attack graph from the attacks described in [18] is shown in Fig. 8 and the descriptions of the nodes are provided in Table III in Appendix II.

We have set the numerical values of the nodes' attributes as follows. For simplicity, we classify the base probability of successful attack p^0 using a $\{$high, medium, low, very low$\}$ ordinal scale with probabilities of 1, 0.75, 0.5, and 0.25, respectively. The remaning security attributes for each node are derived subjectively following discussions provided in [18]. For quantifying κ, we have taken into account a few factors such as user interaction with these units, fleet exposure, and individual criticality. For example, Node-10 which represents the firmware on the telematics unit itself, has been given a low κ as it may require user interaction, and as such we assume that significant security improvement on this unit is harder as it may degrade customer experience. Node-6 on the other hand, represents remote communication procedures between backend manufacturer servers and the telematics unit to perform diagnostic operations without user interaction, and hence we have set it as having a higher κ.

2) Optimal investment decisions: For this system, due to standardization and being external to the manufacturer, we assume that no investment is possible on Node-1 ($x_1^* = 0$). Assuming a budget $B = 5$ units, we obtain $\mathbf{x}^*_{\text{veh}} = [0, 0, 0, 0, 0.8179, 2.8317, 0.0956, 0.3820, 0, 0.5796, 0, 0.2932, 0, 0]$ and $\mathbf{L}^*_{\text{veh}} = 1.8837$.

Since the numbers depend heavily on the quantification of the security attributes, we only interpret the *ranking* of the investments. It can be seen that Nodes 5, 6, and 10 receive the highest investment. This indicates that immediate resources, whether that be time, number of personnel, or monetary resources, should be used for securing the nodes responsible for remote access; in fact, according to [18], the manufacturer prioritized Nodes 6 and 10 in its defense strategy.

3) Design interventions: When the product (nodes) are in post-production, direct investment on these nodes may not be possible, and the manufacturer may instead perform a network redesign. The chosen countermeasure by the manufacturer in reality was to forward the remote request from the backend server (received from the customer) to another server which updates the configuration on the vehicle *only* by HTTPS allowing for a higher payload and security [18]. To perform the intended command, the requested configuration must match with the updated configuration. This can be seen as setting up a hybrid communication completely in the backend with minimal user interaction [31]. To capture this, we introduce the hybrid node Node-15 after Node-6 and in parallel with Node-10 with attributes $L_{15} = L_{10}$, $p^0_{15} = 0.25$, and $\kappa_{15} = \kappa_{10}$. With this redesign, we obtain $L^*_{veh} = 1.7625$, which is lower than the original base loss, indicating the effectiveness of the countermeasure in improving the vehicle's security, and matching the intervention analysis discussed in Section IV.

VI. Conclusion and future work

We studied the effect of network design interventions on a network of interdependent assets with varying sensitivity to investment. We first proposed an algorithm to convert large attack graphs to a reduced attack graph with fewer variables and constraints, allowing for more computationally efficient system level loss analysis and comparison. We then considered four potential types of design interventions by a defender of this system: a series node addition, a parallel node addition, a hybrid node addition, and an additional input node. We find that added endurance (series) or informationally redundant (hybrid) components can help decrease the expected loss, while adding new features (input) or physically redundant (parallel) components in general increases expected loss, and can only be justified if there is additional stand-alone benefit to these additions. We further showcased the usability of the proposed approach by applying it on two use cases. The results recommended by this approach are close to realistic decisions taken by the manufacturers and to outcomes from studies performed by security agencies. Future work includes, generalizing the reduction algorithm to consider scenarios with insufficient budget, and extending our proposed framework with Reinforcement Learning based techniques to study both optimal security investments and network design interventions in multi-stage attacker-defender games.

References

[1] N. Falliere, L. O. Murchu, and E. Chien, "W32. stuxnet dossier," *White paper, Symantec Corp., Security Response*, vol. 5, no. 6, p. 29, 2011.

[2] A. Greenberg, "Hackers remotely kill a jeep on the highwaywith me in it," *Wired*, vol. 7, p. 21, 2015.

[3] J. Slay and M. Miller, "Lessons learned from the maroochy water breach," in *International conference on critical infrastructure protection*. Springer, 2007, pp. 73–82.

[4] S. Amin, G. A. Schwartz, and S. S. Sastry, "Security of interdependent and identical networked control systems," *Automatica*, vol. 49, no. 1, pp. 186–192, 2013.

Authorized licensed use limited to: The Ohio State University. Downloaded on June 08,2023 at 18:01:49 UTC from IEEE Xplore. Restrictions apply.
© 2023 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

This article has been accepted for publication in IEEE Transactions on Control of Network Systems. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TCNS.2023.3272849

[5] T. H. Nguyen, M. Wright, M. P. Wellman, and S. Baveja, "Multi-stage attack graph security games: Heuristic strategies, with empirical game-theoretic analysis," in *Proceedings of the 2017 Workshop on Moving Target Defense*, 2017, pp. 87–97.

[6] N. Liu, J. Zhang, H. Zhang, and W. Liu, "Security assessment for communication networks of power control systems using attack graph and mcdm," *IEEE Transactions on Power Delivery*, vol. 25, no. 3, pp. 1492–1500, 2010.

[7] A. R. Hota, A. A. Clements, S. Bagchi, and S. Sundaram, "A game-theoretic framework for securing interdependent assets in networks," in *Game theory for security and risk management*. Springer, 2018, pp. 157–184.

[8] J. C. Smith and Y. Song, "A survey of network interdiction models and algorithms," *European Journal of Operational Research*, vol. 283, no. 3, pp. 797–811, 2020.

[9] J. Zeng, S. Wu, Y. Chen, R. Zeng, and C. Wu, "Survey of attack graph analysis methods from the perspective of data and knowledge processing," *Security and Communication Networks*, vol. 2019, 2019.

[10] J. Milošević, M. Dahan, S. Amin, and H. Sandberg, "A network monitoring game with heterogeneous component criticality levels," in *2019 IEEE 58th Conference on Decision and Control (CDC)*. IEEE, 2019, pp. 4379–4384.

[11] M. Pirani, E. Nekouei, H. Sandberg, and K. H. Johansson, "A game-theoretic framework for the security-aware sensor placement problem in networked control systems," *IEEE Transactions on Automatic Control*, 2021.

[12] M. Pirani, J. A. Taylor, and B. Sinopoli, "Strategic sensor placement on graphs," *Systems & Control Letters*, vol. 148, p. 104855, 2021.

[13] M. Abdallah, P. Naghizadeh, A. R. Hota, T. Cason, S. Bagchi, and S. Sundaram, "Behavioral and game-theoretic security investments in interdependent systems modeled by attack graphs," *IEEE Transactions on Control of Network Systems*, vol. 7, no. 4, pp. 1585–1596, 2020.

[14] S. Milani, W. Shen, K. S. Chan, S. Venkatesan, N. O. Leslie, C. Kamhoua, and F. Fang, "Harnessing the power of deception in attack graph-based security games," in *International Conference on Decision and Game Theory for Security*. Springer, 2020, pp. 147–167.

[15] P. Sharma Oruganti, P. Naghizadeh, and Q. Ahmed, "The impact of network design interventions on cps security," in *2021 60th IEEE Conference on Decision and Control (CDC)*, 2021, pp. 3486–3492.

[16] M. Abdallah, D. Woods, P. Naghizadeh, I. Khalil, T. Cason, S. Sundaram, and S. Bagchi, "Morshed: Guiding behavioral decision-makers towards better security investment in interdependent systems," in *Proceedings of the 2021 ACM Asia Conference on Computer and Communications Security*, 2021, pp. 378–392.

[17] K. Bissell, R. M. Lasalle, and P. Dal Cin, "The cost of cybercrimeninth annual cost of cybercrime study," *Ponemon Institute and Accenture Security. https://www. accenture. com/_acnmedia/PDF-96/Accenture-2019-Cost-of-Cybercrime-Study-Final. pdf# zoom*, vol. 50, 2019.

[18] Z. Cai, A. Wang, W. Zhang, M. Gruffke, and H. Schweppe, "0-days & mitigations: roadways to exploit and secure connected bmw cars," *Black Hat USA*, vol. 2019, p. 39, 2019.

[19] M.-E. Paté-Cornell and M. A. Kuypers, "A probabilistic analysis of cyber risks," *IEEE Transactions on Engineering Management*, 2021.

[20] M. D. Smith and M. E. Paté-Cornell, "Cyber risk analysis for a smart grid: How smart is smart enough? a multiarmed bandit approach to cyber security investment," *IEEE Transactions on Engineering Management*, vol. 65, no. 3, pp. 434–447, 2018.

[21] M.-E. Paté-Cornell, M. Kuypers, M. Smith, and P. Keller, "Cyber risk management for critical infrastructure: a risk analysis model and three case studies," *Risk Analysis*, vol. 38, no. 2, pp. 226–241, 2018.

[22] P. Xie, J. H. Li, X. Ou, P. Liu, and R. Levy, "Using bayesian networks for cyber security analysis," in *2010 IEEE/IFIP International Conference on Dependable Systems & Networks (DSN)*. IEEE, 2010, pp. 211–220.

[23] S. Abraham and S. Nair, "Exploitability analysis using predictive cybersecurity framework," in *2015 IEEE 2nd International Conference on Cybernetics (CYBCONF)*. IEEE, 2015, pp. 317–323.

[24] D. S. Hochbaum, "Complexity and algorithms for nonlinear optimization problems," *Annals of Operations Research*, vol. 153, no. 1, pp. 257–296, 2007.

[25] I. Pólik and T. Terlaky, "Interior point methods for nonlinear optimization," in *Nonlinear optimization*. Springer, 2010, pp. 215–276.

[26] P. S. Oruganti, P. Naghizadeh, and Q. Ahmed, "The impact of network design interventions on the security of interdependent systems," 2023. [Online]. Available: https://arxiv.org/abs/2302.05411

[27] J. A. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, "Casadi: a software framework for nonlinear optimization and optimal control," *Mathematical Programming Computation*, vol. 11, no. 1, pp. 1–36, 2019.

[28] C. Few, J. Thompson, K. Awuson-David, and T. Al-Hadhrami, "A case study in the use of attack graphs for predicting the security of cyber-physical systems," in *2021 International Congress of Advanced Technology and Engineering (ICOTEN)*. IEEE, 2021, pp. 1–7.

[29] Common vulnerability scoring system v3.1: Specification document. [Online]. Available: https://www.first.org/cvss/v3.1/specification-document

[30] C. Miller and C. Valasek, "Remote exploitation of an unaltered passenger vehicle," *Black Hat USA*, vol. 2015, no. S 91, 2015.

[31] Black hat presentation, 0-days & mitigations: Roadways to exploit and secure connected bmw cars. [Online]. Available: https://www.youtube.com/watch?v=HS8MoFa0TNs

Pradeep Sharma Oruganti received his B.Tech. in Mechanical Engineering from the Indian Institute of Technology, Hyderabad, India in 2012 and M.Sc in Mechanical Engineering from The Ohio State University (OSU), Columbus, OH, USA in 2017. He is currently pursing a Ph.D in Mechanical Engineering at OSU. His research interests include control theory, game theory and system dynamics with applications in safety and security assurance in cyber-physical systems.

Parinaz Naghizadeh received her Ph.D. degree in electrical engineering from the University of Michigan, Ann Arbor, in 2016. She is an Assistant Professor with the Department of Integrated Systems Engineering and the Department of Electrical and Computer Engineering at the Ohio State University. Her research interests include network economics, game theory, algorithmic economics, and reinforcement learning. She is a recipient of the 2022 NSF CAREER Award.

Qadeer Ahmed received Ph.D. degree in Control Systems from Mohammad Ali Jinnah University, Islamabad, Pakistan 2011. Now, he is Research Associate Professor with Mechanical and Aerospace Engineering Department and Center for Automotive Research, The Ohio State University, Columbus, OH, USA. His research includes control and diagnostics of automotive systems with focus on their efficiency, safety, and security.

Authorized licensed use limited to: The Ohio State University. Downloaded on June 08,2023 at 18:01:49 UTC from IEEE Xplore. Restrictions apply.
© 2023 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

This article has been accepted for publication in IEEE Transactions on Control of Network Systems. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TCNS.2023.3272849

Appendix I
Summary of Notation

TABLE I: Summary of notation

Symbol	Description
$\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$	Directed acyclic graph with nodes $\mathcal{V}$ and edges $\mathcal{E}$
v_g	Unique target node
$Post(v)$	Set of all nodes that can be reached from v
$Pre(v)$	Set of all the nodes from where v can be reached
P_{ij}	An attack path from v_i to v_j
$\mathcal{P}_{ij}$	Set of all paths from v_i to v_j
x_i	Security investment on node v_i
$p_i(x_i)$	Probability of successful attack on v_i given an investment of x_i on it
κ_i	Sensitivity of node v_i to investments
p_i^0	Default probability of successful attack on v_i
L_i	Stand-alone loss of v_i if compromised
B	Security budget of the defender

Appendix II
Description and Security Attributes for the Numerical Experiments in Section V

TABLE II: Description and security attributes for Fig. 7

Node	Description	$\mathbf{p^0}$	**L** (x10^6)	κ
1	Email host	0.18	0.01	1
2	Web app	0.18	0.01	1
3	Privilege escalation	0.09	0.02	1
4	Web app host	0.09	0.02	1
5	Defeat DMZ access control	0.09	20	3
6	Connect to firewall	0.13	0.2	3
7	Gain access to engg. workstation	0.08	1000	5
8	Gain access to SCADA controls	0.08	2000	5
g	Manipulate PLC	1	10000	-
9	Manipulate PLC	0.07	50	5
g	Malicious actuation	-	100	-

TABLE III: Description and security attributes for of the nodes and their characteristics in the attack graph illustrated in Fig. 8.

Node	Description	p^0	L	κ
1	Cellular connection	1	1	1
2	OBD-II	0.25	1	1
3	Wi-Fi	0.5	1	1
4	USB	0.25	1	1
5	Connected services	0.75	5	3
6	Connection to vehicle	0.75	10	1
7	Internal network diagnostics	0.75	5	3
8	Head Unit internal arch.	0.75	5	3
9	Connected services comm.	0.75	5	3
10	Remote diagnostic comm.	0.75	20	2
11	Telematics Control Unit	0.25	20	2
12	Head unit	0.5	5	2
13	CAN tx/rx	0.25	5	1
14	Central Gateway	1	20	1
g	Vehicle CAN	1	50	-

Appendix III
Obtaining Optimal Investments on the Original Graph and Proof for Lemma 1

While Algorithm 1 provides the final reduced graph and the optimal expected loss. The resulting nodes may lose their physical meaning through the reduction process. It is essential for the defender to query the optimal investments on the nodes of the original attack graph. These are obtained through the reduction procedure and can be stored during each iteration of Algorithm 1.

From Lemma 3, given some budget T invested over a series link $\{v_i \rightarrow v_{i+1} \rightarrow \ldots \rightarrow v_{i+n} \rightarrow v_t\}$, the optimal investment on each node is given by:

$$x_t^* = \frac{-1}{\kappa_t} \log\Big(\frac{L_{i+n}\kappa_{i+n}}{L_t(\kappa_t - \kappa_{i+n})}\Big),$$

$$x_j^* = \frac{-1}{\kappa_j} \log\Big(\frac{L_{j-1}\kappa_{j-1}}{L_j\kappa_{j+1}}\Big(\frac{\kappa_{j+1} - \kappa_j}{\kappa_j - \kappa_{j-1}}\Big)\Big),$$

$$x_i^* = T - x_t^* - \sum_{j=i+1}^{i+n} x_j^*, \ \forall j \in I = \{i+1, i+2, \ldots, i+n\}.$$

Similarly, from Lemma 5, the optimal investments on a parallel network $\{(v_i \rightarrow v_{i+1} \rightarrow v_g), (v_i \rightarrow v_{i+2} \rightarrow v_g), \ldots, (v_i \rightarrow v_{i+n} \rightarrow v_g)\}$ is given by:

$$x_j^* = \tfrac{-1}{\kappa_r} \log\Big(\frac{\kappa_i\kappa_{par}}{1 - \kappa_i\kappa_{par}}(\frac{L_i}{L_j + L_g})\Big), \forall r = \{i+1, \ldots, i+n\}$$

$$x_i^* = T - \sum_{r=i+1}^{i+n} x_j^*, \ \forall j \in I.$$

And from Lemma 7, for a graph $\{(v_i \rightarrow v_t), (v_{i+1} \rightarrow v_t), \ldots, (v_{i+n} \rightarrow v_t)\}$ with $L_j = L, \ \forall j \in I$, the optimal investments on the nodes are given by:

$$x_t^* = \frac{-1}{\kappa_t} \log\Big(\frac{L}{L_t(\kappa_t\kappa_{par} - 1)}\Big)$$

$$x_j^* = \frac{1}{\kappa_j\kappa_{par}}\Big(T - x_t^*\Big), \forall j \in I.$$

We see that other than the input nodes in each link, the optimal investments on the other nodes do not depend on the budget T. Hence, any increase in budget would not change the investment on these nodes. Additionally, a sufficient budget T over each link would then be the minimum budget T such that:

$$T - \sum_{r=i+1}^{i+n} x_r^* \geq 0$$

A sequence of steps that the defender could follow to assess the utility of a design intervention would be:

1) Perform a network design intervention on the original graph
2) Run Algorithm 1 to obtain the optimal expected loss after the intervention while storing the optimal investments on the reduced nodes during each iteration.
3) Output the optimal investments on every node.

Authorized licensed use limited to: The Ohio State University. Downloaded on June 08,2023 at 18:01:49 UTC from IEEE Xplore. Restrictions apply.
© 2023 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

This article has been accepted for publication in IEEE Transactions on Control of Network Systems. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TCNS.2023.3272849

A. Proof of Lemma 1

Proof: The proof is by construction, and follows directly from the above arguments on finding the optimal investment profiles in each series or parallel subnetwork. ■

Appendix IV
Series reduction under insufficient budget

Under an insufficient budget, the nodes closer to the end of the link receive investment first. Investment is made sequentially, starting from the target node and moving upstream, each node receiving their optimal investments following Lemma 3 until the budget runs out. This is proved in the following lemma:

Lemma 8: Consider an attack graph $\mathcal{G}$ containing the series of nodes $\{v_i \rightarrow v_{i+1} \rightarrow \cdots \rightarrow v_{i+n} \rightarrow v_{i+n+1}\}$ and a budget T spent over these nodes. Under a optimal investment strategy $\mathbf{x}^*$, there exits a non-zero optimal investment on $v_j, j \in \{i, i+1, i+2, \ldots i+n\}$, when:

$$\sum_{v_t \in Post(v_j)} x_t^* < T$$

Proof: Consider a series path $\mathcal{G}_t = \{v_0 \rightarrow v_1 \rightarrow v_2 \cdots \rightarrow v_t\}$ such that $\sum_{j=1}^{t} x_j^* \geq T$. We prove the lemma by induction. For the base case, consider a node v_{-1} added upstream to get $\mathcal{G}_1 = \{v_{-1} \rightarrow v_0 \rightarrow v_1 \rightarrow v_2 \cdots \rightarrow v_t\}$. From the discussion in Appendix III we know that the investments on the t downstream nodes do not change. From Lemma 3, $x_{-1}^* = T - \sum_{j=0}^{t} x_j^*$. Additionally, depending on the node attributes of v_0 and v_{-1}, $x_0^* \geq 0$ which implies $\sum_{j=0}^{t} x_j^* \geq T$ and $x_{-1}^* \leq 0$. Hence, within the feasible region the minimum occurs at $x_{-1}^* = 0$, proving the base case. Next for the inductive step, consider the path $\{v_{1-n} \rightarrow \cdots \rightarrow v_{-1} \rightarrow v_0 \rightarrow v_1 \rightarrow v_2 \cdots \rightarrow v_t\}$. We assume the statement is true for this sequence of $n+t$ nodes. With this assumption, the upstream n nodes can be replaced with single node with equivalent loss $L_{eq} = \sum_{j=1-n}^{0} L_j$. Finally, applying the same steps as the induction base, it is straightforward to show that the statement holds for $n+1+t$ nodes with only the t nodes downstream receiving non-zero investments. ■

Conjuncture 1: By viewing the entire attack graph as a series of sub-networks (series or parallel), following Lemma 8, under an insufficient budget, the sub-networks closer to the target receive their respective optimal investments until the budget is depleted.

Conjuncture 2: If the budget T over a parallel link $\{(v_i \rightarrow v_{i+1} \rightarrow v_g), (v_i \rightarrow v_{i+2} \rightarrow v_g), \ldots, (v_i \rightarrow v_{i+n} \rightarrow v_g)\}$ is insufficient, the *root node* v_i does not receive any investment and the rest of the split across the parallel nodes v_j, $j \in I = \{i+1, i+2, \ldots, i+n\}$. The optimal investments on each node then being:

$$x_r^* = \frac{1}{k_r k_{par}} \Big(T + \sum_{j=I\setminus\{r\}|} \frac{1}{k_r} \log\Big(\frac{L_j + L_g}{L_r + L_g}\Big)\Big), r \in I$$

The remaining budget may or may not equalize expected losses across all paths across all parallel paths, with certain nodes receiving no investment.

Authorized licensed use limited to: The Ohio State University. Downloaded on June 08,2023 at 18:01:49 UTC from IEEE Xplore. Restrictions apply.

© 2023 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.