Effective Defence Against Zero-Day Exploits Using Bayesian Networks

Tingting Li and Chris Hankin

Institute for Security Science and Technology,
Imperial College London, UK
{tingting.li, c.hankin}@imperial.ac.uk

Abstract. Industrial Control Systems (ICS) play a crucial role in controlling industrial processes. Unlike conventional IT systems or networks, cyber attacks against ICS can cause destructive physical damage. Zero-day exploits (i.e. unknown exploits) have demonstrated their essential contributions to causing such damage by Stuxnet. In this work, we investigate the possibility of improving the tolerance of a system against zero-day attacks by defending against known weaknesses of the system. We first propose a metric to measure the system tolerance against zero-day attacks, which is the minimum effort required by zero-day exploits to compromise a system. We then apply this metric to evaluate different defensive plans to decide the most effective one in maximising the system tolerance against zero-day attacks. A case study about ICS security management is demonstrated in this paper.

1 Introduction

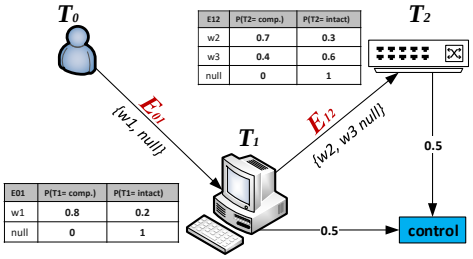
Cyber security of industrial control systems has increasingly become a severe and urgent problem, owing to the wide use of insecure-by-design legacy systems in ICS, and the potential physical damage of breached ICS to infrastructures, environment and even human health [19]. The rapid integration of ICS with modern ICT technology has further intensified the problem. Increasing attention has been drawn to this issue from various sectors such as industry, government and academia. Whilst most ICS vulnerabilities inherited from IT systems are well studied and can be defended by conventional security controls, very little effort can be made to combat zero-day exploits, because they are often unknown to the vendor and hence there is no patch available to fix them. One of the most famous cyber attacks against ICS is Stuxnet disclosed in 2010 [4], which was distributed by an infected USB flash drive, propagated across the corporate network, and eventually compromised the PLCs to disrupt the operation of industrial plants. Four zero-day vulnerabilities played crucial roles in gaining access to targets and propagating the malware. Until September 2010, there were about 100,000 hosts over 155 countries infected by Stuxnet [4]. The threat from zero-day exploits is still on the rise. In 2014, 245 incidents were reported to ICS-CERT and 38% of these incidents were identified as having an "unknown access vector", and ICS-CERT specifically mentioned the exploitation of zero-day vulnerabilities as one of the methods used by attackers [11]. In July 2015, a zero-day vulnerability in Adobe Flash Player has been acknowledged by ICS-CERT, which is able to gain access to critical infrastructure

networks via spear-phishing emails [10]. Later in August, ICS-CERT continuously released six advisories and six alerts about zero-day vulnerabilities on Siemens SIMATIC S7-1200 CPUs, Schneider Electric DTM, Rockwell Automation PLCs, etc.

It is extremely difficult to detect and defend against zero-day exploits. Sophisticated hackers are able to discover zero-day exploits before the vendors become aware of them. Since it is difficult to directly stop zero-day attacks, we consider the problem from a novel perspective, by seeking a way to make ICS sufficiently robust against zero-day attacks. We are able to reduce the risk of potential zero-day exploits to an acceptable level by strategically defending against the known attack vectors.

A typical APT attack targeting ICS has to exploit a chain of vulnerabilities at different hosts to eventually breach the control devices (e.g. PLCs). The involved exploits use either known or zero-day vulnerabilities to propagate across the network. Whilst we can hardly defend against the exploitation of zero-day vulnerabilities, we can alternatively deploy effective defences against the known vulnerabilities such that the risk of the whole attack chain being exploited can be overall reduced. A key attribute "exploitability" of weaknesses is borrowed from CWE [2] to reflect the sophistication of a zero-day weakness and the required attacking effort. Weaknesses with higher exploitability are likely to cause higher risk for the system. With regard to an acceptable level of risk, we define the tolerance against a zero-day weakness by the minimal required exploitability of the weakness to cause the system risk exceed the acceptable level. By using Bayesian Networks, we can prove that defending against known weaknesses is able to increase the tolerance, and find out the defence that maximizes the tolerance.

We express an acceptable level of risk by conceptual safety-related requirements of ICS [14] such as the integrity of monitoring data, the availability of control and reliable communication. Cyber attacks targeting ICS might violate such requirements to a certain degree. By modelling these requirements as nodes of a Bayesian Network, we can define the acceptable risk by the severity of a requirement being violated. Next we use a simple example to further illustrate the motivation of this work.



	tole. of T_1	tole. of T_2
no control	34%	10%
c_1 against w_1	42%	70%
c_2 against w_2	36%	17%
c_3 against w_3	36%	16%

Fig. 1: **Left**: a simple Stuxnet-like attack scenario. **Right**: Tolerance improvement

Fig.1 shows a Stuxnet-like attack path launched by an adversary T_0 to gain access to a workstation T_1 (by exploiting w_1), which is then used as a foothold to further compromise the PLC T_2 (by exploiting w_2 or w_3). Exploitability of each known weakness is given in the tables. Both T_1 and T_2 equally contribute to the satisfaction of the requirement about available control. The requirement about control is indicated by a box in the bottom-right of Fig.1. We set the acceptable risk as the probability of the control

requirement being violated must be lower than 27%. We also assume there is a zero-day weakness at T_1 or T_2. By using our approach based on Bayesian networks, we obtained results in the right table of Fig.1. Without any control deployed, either a zero-day exploit at T_1 with 34% exploitability or a zero-day exploit at T_2 with 10% exploitability is sufficient to bring the risk beyond the acceptable level. It can be found that by deploying different controls, the tolerance against zero-day exploits is generally increased. Some increments are rather small as in this demo example the effectiveness of controls is set to 10% only. Different controls bring about different improvement on the tolerances depending on their effectiveness, the exploitabilities of their combating weaknesses and their influence coverage over targets. Our approach is able to take all these factors into account and find out the most effective defence against zero-day attacks.

2 Modelling and Problem Representation

In this section we formally use Bayesian Networks (BN) to model ICS-targeted attacks with zero-day exploits involved and evaluate the risk. A discrete random variable is captured by a chance node in BN with a finite set of mutually exclusive states and a conditional probability distribution over the states. We further defined three types of chance nodes for different purposes: (i) *target nodes* indicate valuable assets in ICS with a set of known and zero-day weaknesses, (ii) *attack nodes* captures available attack methods between a pair of targets, and (iii) *requirement nodes* are designed to model particular objectives for evaluation. A *Bayesian Risk Network* is established based on the three types of nodes, where complete attack paths are modelled by target and attack nodes, and the damage of successful attacks are evaluated against requirement nodes.

Definition 1. *Let $\mathcal{T}$ be a set of **target nodes** $\mathcal{T} = \{T_1, \ldots, T_n\}$. Parent nodes of a target T_x is denoted by $T'_x \in pa(T_x)$. The domain states of a target node is $\Omega(T_x) = \{c, i\}$, representing the target being compromised or intact respectively. Let $\mathcal{R} = \{R_1, \ldots, R_m\}$ be a set of **requirement nodes** and the domain states of a requirement node is $\Omega(R_x) = \{c, v\}$, indicating respectively the requirement being complied or violated. Parents of a requirement node could be target nodes or other requirement nodes $pa(R_x) \subseteq \mathcal{T} \cup \mathcal{R}$.*

Target nodes represent valuable assets where zero-day weaknesses might be exploited in addition to the known weaknesses. Requirement nodes capture safety-related objectives, which are used to evaluate the impact of cyber attacks on the system safety. Detailed modelling and reasoning about these requirements can be found in [14].

Definition 2. *Let $W = \{w_1, \ldots, w_m\}$ be a set of weaknesses, $\omega : \mathcal{T} \times \mathcal{T} \rightarrow 2^W$ gives possible weaknesses that can be exploited from one target to another. Let $\mathcal{E} = \{E_{T'_1 T_1}, \ldots, E_{T'_n T_n}\}$ be a set of **attack nodes** connecting a target and its parents. The domain states $\Omega(E_{T_i T_j}) = \omega(T_i, T_j) \cup \{null\}$, including all weaknesses on T_j that can be exploited from T_i, or none of them is exploited (i.e. null). The **exploitability** e_{w_x} of a weakness w_x is the likelihood of w_x being successfully exploited.*

Definition 3. *Let w_z be a **zero-day exploit** with uncertain exploitability $e_{w_z} \in [0, 1]$, $w_z \in \Omega(E_{T'_x T_x})$ indicates there is a zero-day exploit at the target T_x.*

Unlike Bayesian Attack Graphs (BAG) [18] that are constructed based on states and attributes, we build a Bayesian network at the level of assets and model multiple weaknesses between a pair of assets by a single attack node, rather than multiple attack edges. Each attack node hence becomes a decision-making point for attackers to choose a (known or zero-day) weakness to proceed. Such Bayesian networks enable us to model zero-day exploits without knowing details about them (e.g. pre-requisites or post-conditions), but concentrate on analysing the risk caused by zero-day exploits.

Definition 4. *Let $\mathcal{C} = \{c_1, \ldots, c_k\}$ be a set of **defence controls** and $d(c_x) \in 2^W$ be weaknesses that can be defended by c_x. If $w_i \in d(c_j)$, then by deploying c_j, the exploitability of w_i is scaled by $\varepsilon \in [0, 1]$, where ε is the effectiveness of c_j.*

A defence control is able to reduce the exploitability of its combating weaknesses. If ε is set to 50%, then applying c_j reduces the exploitability of $w_i \in d(c_j)$ by 50%.

Definition 5. *Let $\mathcal{B} = \langle \mathcal{N}, \mathcal{P}_T, \mathcal{P}_E, \mathcal{P}_R, P_{T_0} \rangle$ be a **Bayesian Risk Network** , where*

- *$\mathcal{N} = \mathcal{T} \cup \mathcal{E} \cup \mathcal{R}$, including target nodes, attack nodes and requirement nodes.*
- *$\mathcal{P}_T = \{P_{T_1}, \ldots, P_{T_n}\}$ includes conditional probabilities of all non-root target nodes given their parents such that P_{T_x} denotes $P(T_x \mid \bigcup_{T'_x \in pa(T_x)} E_{T'_x T_x})$, where $P(T_x \mid \bigcup_{T'_x \in pa(T_x)} E_{T'_x T_x})$=1$-\prod_{T'_x \in pa(T_x)}(1-P(T_x|E_{T'_x T_x}))$ by noisy-OR operator [17]. $P(T_x|E_{T'_x T_x})$ is the probability of T_x given the weakness used at $E_{T'_x T_x}$.*
- *$\mathcal{P}_E = \{P_{E_{T'_1 T_1}}, \ldots, P_{E_{T'_n T_n}}\}$ includes conditional probability distribution for all attack nodes such that $P_{E_{T'_x T_x}}$ denotes $P(E_{T'_x T_x} \mid T'_x)$.*
- *$\mathcal{P}_R = \{P_{R_1}, \ldots, P_{R_n}\}$ includes decomposition of all requirement nodes such that P_{R_x} denotes $P(R_x \mid pa(R_x))$, where $P(R_x \mid pa(R_x)) = \Sigma_{R'_x \in pa(R_x)} P(R_x|R'_x)$, and $P(R_x|R'_x)$ is the assigned proportion of R'_x in R_x.*
- *P_{T_0} is the prior probability distribution of the root node T_0.*

$P(T_x)$ is the unconditional probability of $T_x \in \mathcal{T}$, which can be obtained by:

$$P(T_x) = \begin{cases} \sum_{E_{T'_x T_x}} P_{T_x} \sum_{T'_x} P_{E_{T'_x T_x}} P(T'_x) & \text{if } w_z \notin \Omega(E_{T'_x T_x}) \\ \sum_{E_{T'_x T_x}} P_{T_x} \sum_{T'_x} P_{E_{T'_x T_x}} P(T'_x) + P(T_x|E_{T'_x T_x} = w_z) \sum_{T'_x} P_{E_{T'_x T_x}} P(T'_x) & otherwise \end{cases}$$

$P(T_x)$ is obtained by its parent node $P(T'_x)$ recursively until it hits the root T_0 whose probability distribution is known. $\sum_{E_{T'_x T_x}}$ denotes $E_{T'_x T_x}$ is marginalized. P_{T_x}, P_{R_x} and $P_{E_{T'_x T_x}}$ are given by $\mathcal{P}_T$, $\mathcal{P}_R$ and $\mathcal{P}_E$ respectively. $P(T_x|E_{T'_x T_x} = w_z)$ equals to the uncertain exploitability of the zero-day exploit w_z at T_x.

$P(R_x)$ denotes the unconditional probability of $R_x \in \mathcal{R}$ given its parents R'_x and $P(R_x) = \sum_{R'_x} P_{R_x} \prod_{R'_x \in pa(R_x)} P(R'_x)$, where $pa(R_x)$ are marginally independent.

$\mathcal{P}_T$ is given by conditional probability tables (CPT) for each target node. Each entry of the CPT is the probability of a target being compromised (resp. intact) when a weakness is chosen, which equals to the exploitability e_{w_x} (resp. $1 - e_{w_x}$) of the chosen weakness. Such a CPT is shown in the upper part of Fig.2(a). When w_1 is used,

the chance of T_1 being compromised $P(T_1 = c)$ is 0.8, equivalent to the exploitability of w_1. When a target node has multiple parent nodes, noisy-OR operator [17] is applied to calculate the joint probability of parents, as in BAG [15][18]. $P_{E_{T'_x T_x}}$ decides the chance of each weakness being used. Here we assume attackers choose uniformly from available weaknesses. As given in the lower CPT in Fig.2(a), when the parent target is intact ($T_2 = i$), no attack can be continued towards the next target (i.e. *null* is the only choice). When the parent target is compromised, the probability is equally distributed over the available weaknesses $\Omega(E_{T_2 T_4}) = \{w_4, w_5, null\}$. If a zero-day exploit exists at T_x, the extra contribution of w_z is added to $P(T_x)$. We make the same assumption as in [18][8] that such *Bayesian Risk Networks* are directed acyclic graphs.

Definition 6. *A Bayesian Risk Network $\mathcal{B}$ is constructed for a given system. The **tolerance against zero-day attacks** of the system is represented by (κ, Z), where*

- *κ is a defined acceptable level of risk, expressed by $\kappa := P(N_a = s) \leqslant L$, where the probability of a fixed node $N_a \in \mathcal{T} \cup \mathcal{R}$ being at a particular state $s \in \Omega(N_a)$ is used to define the risk and L is the upper bound of $P(N_a = s)$.*
- *$Z := \langle z_1, \ldots, z_n \rangle$ is a tolerance tuple with each element corresponding to the tolerance against a zero-day exploit at each target node. Thus the tolerance $z_i \in Z$ against a zero-day at an arbitrary target $T_i \in \mathcal{T}$ is obtained by :*

$$z_i = \underset{P(T_i | E_{T'_i T_i} = w_z)}{\operatorname{argmax}} \kappa := P(N_a = s) \leqslant L$$

$P(T_i | E_{T'_i T_i} = w_z)$ equals to the exploitability of the zero-day exploit w_z at T_i and z_i is the maximum exploitability of w_z subject to κ. $P(N_a)$ is the unconditional probability of a target or requirement node, which can be obtained by Definition 5.

We select a particular node N_a to define the risk κ, which could be a valuable target node or a critical requirement. Thus κ is defined by the likelihood of N_a being compromised or violated, e.g. the likelihood of a requirement being violated must be less than 30%. The presence of a zero-day exploit at any target is likely to increase the likelihood as its exploitability increases. Thus, we define the tolerance by the minimum required exploitability of a zero-day exploit at each target to violate κ, or alternatively the maximum exploitability of a zero-day exploit the system can tolerate subject to κ.

3 Case Study and Results

In this section, we present a hypothetical example to demonstrate our approach of finding effective defence against zero-day exploits. We start with the configuration of the example and then discuss the results by applying different defence controls.

3.1 Case Study Settings

A simple network is constructed in Fig.2(a) consisting of common types of assets in ICS – a *HMI*, a *workstation*, a *PLC* and a *RTU*. The four assets are modelled as four target nodes $\mathcal{T} = \{T_1, T_2, T_3, T_4\}$ of a Bayesian network. A special node *EXT*(denoted by T_0) represents the external environment of the network. We also select five common weaknesses $\{w_1, w_2, w_3, w_4, w_5\}$ from the *ICS Top 10 Threats and Countermeasures*

[1] and *Common Cybersecurity Vulnerabilities in ICS* [3]. These weaknesses are enumerated in Fig.2(b), which are attached to relevant attack nodes between a pair of targets. Exploiting different weaknesses yields different consequences depending on the exploitability of the chosen weakness. For instance, in order to compromise T_1, an attacker can choose to exploit w_1 or w_2, or keep hiding $null$. Currently we assume that attackers choose relevant weaknesses uniformly at each attack node. The chance of exploiting a node successfully is given in the relevant CPTs. An example CPT of T_1 is shown in Fig.2(a). When w_1 is chosen, the attacker has a priori 80% chance to compromise T_1. The exploitabilities of weaknesses are essential to construct such CPTs. In this case study, we consistently convert different levels of the CWE attribute "*Likelihood of Exploit*" [2] and the metric "*Exploitability*" from [1] into certain values. Weaknesses that are identified as *"Very High"* by CWE or *"Easy to Exploit"* in [1] are set to 0.8; Weaknesses with *"High"* level of exploitability are set to 0.7 and *"Moderate"* weaknesses have exploitabilities of 0.6. Thus, we derive the rightmost column of the table in Fig.2(b). The table in Fig2(c) lists a set of common defence controls [1] that are used in this case. We set a uniform effectiveness ε of all controls to 50%. Therefore the exploitability of w_1 becomes 0.4 after deploying c_1.

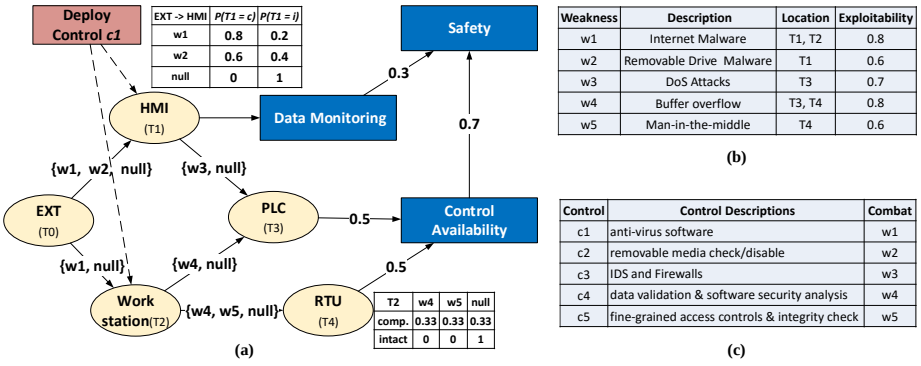


Fig. 2: (a) network; (b)selective common weaknesses; (c) selective common controls;

To model the cyber-physical effects of potential exploits, we consider three key requirements in the example. The target node T_1 has direct and dominated influence on the requirement about *data monitoring*. As two core field controllers, the *PLC* and *RTU* equally contribute to satisfying the requirement about *control availability*. The overall *safety* jointly relies on data monitoring (30%) and control availability (70%) . We make this particular configuration to reflect the common requirement of ICS that system availability generally outweighs the other aspects [13]. In the Fig.2(a), we use dashed lines to indicate the impact of deploying c_1 against w_1 at the target T_1 and T_2.

We construct the corresponding *Bayesian Risk Network* in Fig.3, where the unconditional probability distribution over possible states of each node is computed. The node T_0 denotes the untrusted external environment where attackers can launch any attacks, and thus the probability of its compromised state is 100%. Fig.3 simulates the example ICS *without* any control deployed or any zero-day exploits, and the chance of the safety

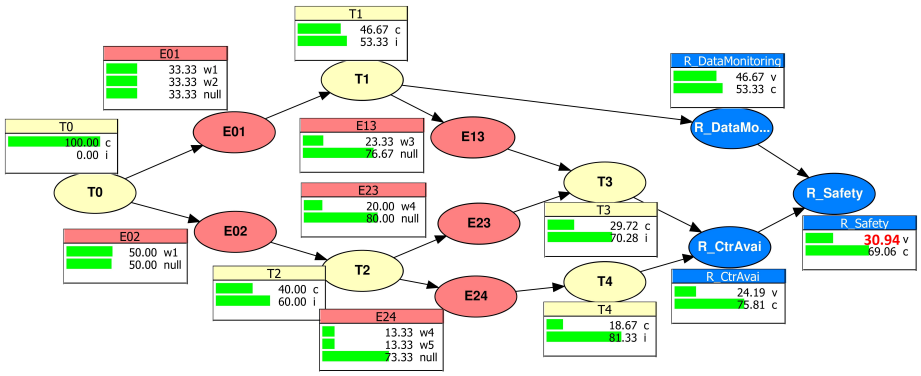


Fig. 3: posterior risk distribution with no control deployed. (by *Hugin Lite* [9])

being violated is about 30.94%. In the following parts of the paper, the *risk* of the system is referred to the probability of the safety requirement being violated $P(R_{\text{safety}} = v)$.

In the next sections, we add zero-day exploits to each target and deploy different controls, in order to evaluate the impact of controls on the system tolerance against those zero-day exploits. We first present the results with an individual control in Section 3.2 and further discuss the results with multiple controls deployed in Section 3.3.

3.2 Results – Deploying a Single Control

We run four trials of the experiment in each of which a zero-day exploit w_z is added to each target. In each trial, different defence controls are individually deployed and the updated risks over scaled exploitabilities of the zero-day exploit (e.g. 20%, 40%, 60% and 80%) are computed. In the four charts of Fig.4, the upper curve with markers illustrates the trend of the risk with *none* control by varying exploitabilities of w_z. This curve is used as the baseline to evaluate the mitigated risk by deploying each control, which are indicated by the coloured bars respectively.

We discover that the existence of zero-day exploits w_z increases the risk of the system. As shown in Fig.3, the *apriori* risk is about 30.94% without any zero-day exploit, which is raised to 34.23% with a w_z of 80% exploitability at T_1, and to 34.6% at T_2. It is worth noting that the risk caused by zero-day exploits starts to *exceed* the risk without zero-day exploits only when the zero-day exploit reaches a certain exploitability, which is seemingly counter-intuitive. At T_1, the risk exceeds the apriori risk (30.94%) when the w_z reaches a higher exploitability (49%). It is because the known weaknesses w_1 and w_2 at T_1 already have rather high exploitabilities, the presence of low-level zero-day exploits would actually reduce the overall chance of T_1 being compromised as we assume the attacking methods are chosen uniformly. Therefore it is possible that the risk *with* zero-day exploits is lower than the risk *without* zero-day exploits when the zero-day exploits are at very low exploitabilities. However, since zero-day exploits can be hardly detected, their exploitabilities tend to be very high in reality. From Fig.4, the zero-day exploit at T_2 is the most threatening one as it brings the greatest increment to the risk, while that at T_4 is the least threatening one. This is simply because T_2 influences more subsequent nodes than T_4.

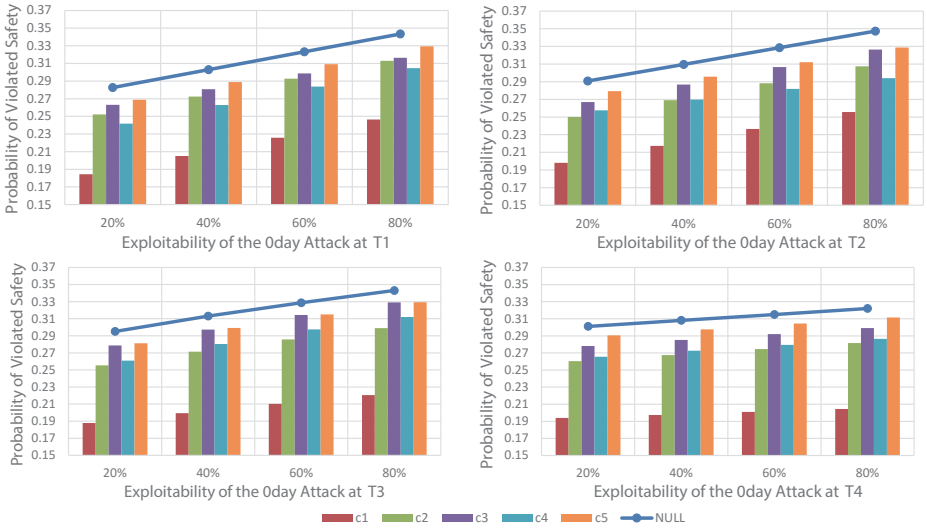


Fig. 4: risk distribution with different controls on each target with an zero-day exploit

It can be found that the control c_1 is the most effective one to reduce the risk such that the risk drops to 24.59% from 34.23% with a w_z of 80% exploitability at T_1, In the bottom-left chart of Fig.4, we notice similar risk mitigation of c_3 and c_5 to combat the w_z at T_3. These two controls mainly target for w_3 at T_3 and w_5 at T_4 respectively. The similar mitigation is probably due to the symmetric positions of T_3 and T_4 in the network and their equal contribution to satisfy the control availability requirement.

The tolerance of the system against zero-day exploits has been improved by deploying controls. In the top-right chart of Fig.4, at least a zero-day exploit with exploitability 31% is needed at T_2 to produce the risk 30%. With the help of c_2 , a zero-day exploit with much higher exploitability 74% at T_2 is required to reach the same level of risk.

3.3 Results – Deploying Combined Controls

A defence plan consists of multiple controls. There are five controls in the example with 2^5 combinations of them. We use bit vectors to represent including or excluding a control in a plan. For the controls $\{c_1, c_2, c_3, c_4, c_5\}$, a defence plan "10011" indicates to apply c_1, c_4 and c_5. We define $|d|$ as the number of controls included in a plan. Each row of Table 1 shows the result of deploying a defence plan, with the maximal risk $max(r)$ when the zero-day exploit at each target reaches its maximal exploitability 100%, and the mean risk reduction $\bar{\Delta}_x$ by deploying the plan. The mean risk reduction over the four targets is given by $\bar{\Delta}$. The acceptable risk κ is set to $P(R_{\text{safety}} = v) \leqslant 20\%$. The rightmost column shows the resulting tolerance against zero-day exploits. The symbol min indicates the risk already exceeds the acceptable level regardless of the existence of a zero-day exploit, while max denotes the system is fully tolerant to a zero-day exploit at the target, i.e. the acceptable level κ can never be violated even if the zero-day exploit reaches its maximal exploitability.

Table 1: Results of Selective Defence Plans

Plan	T_1		T_2		T_3		T_4		$\bar{\Delta}$	**Tolerance**
	$max(r)$	$\bar{\Delta}_1$	$max(r)$	$\bar{\Delta}_2$	$max(r)$	$\bar{\Delta}_3$	$max(r)$	$\bar{\Delta}_4$		
00000	0.362	0	0.365	0	0.357	0	0.328	0	0	min, min, min, min
10000	0.267	0.097	0.274	0.092	0.231	0.113	0.208	0.112	**0.414**	0.36, 0.23, 0.44, 0.57
11000	0.236	0.128	0.234	0.132	0.183	0.156	0.166	0.153	0.569	**0.66, 0.65, max, max**
10001	0.260	0.104	0.258	0.102	0.224	0.120	0.202	0.117	0.443	0.43, 0.31, 0.56, 0.87
10100	0.239	0.118	0.258	0.109	0.219	0.126	0.189	0.130	0.483	0.57, 0.41, 0.66, max
10010	0.247	0.118	0.225	0.123	0.214	0.131	0.189	0.130	0.502	0.56, 0.59, 0.75, max
00101	0.319	0.037	0.324	0.038	0.329	0.030	0.295	0.033	0.138	min, min, min, min
11100	0.212	0.145	0.223	0.145	0.175	0.165	0.153	0.166	0.620	0.87, 0.77, max, max
01101	0.293	0.064	0.289	0.073	0.287	0.068	0.259	0.069	**0.273**	min, min, min, min
11010	0.215	0.149	0.184	0.165	0.166	0.174	0.147	0.172	0.660	0.86, max, max, max
01111	0.253	0.105	0.227	0.118	0.253	0.103	0.222	0.106	0.430	0.42, 0.50, 0.32, 0.37
11011	0.208	0.156	0.167	0.175	0.159	0.180	0.142	0.178	0.689	0.92, max, max, max
11110	0.192	0.166	0.173	0.177	0.158	0.182	0.134	0.185	0.710	max, max, max, max
11111	0.185	0.173	0.156	0.187	0.151	0.189	0.129	0.190	0.740	max,max, max, max

The first row 00000 with no control deployed is still used as the baseline. In terms of the risk reduction, 10000 is the most effective one when $|d| = 1$, while the plans 11000, 11010 and 11110 are the most effective choice if we can implement 2, 3 or 4 controls respectively. We discover that implementing more controls does not always produce stronger defence. For instance, deploying c_2, c_3 and c_5 (i.e. the plan 01101) has risk reduction 0.273, which is lower than the reduction 0.414 by deploying c_1 only. Each control combats different weaknesses that are distributed over different nodes. Defending against more widespread weaknesses would generally produce more risk reduction across the network. Besides, weaknesses near the attack origin (i.e. the node T_0 in this case) tend to have greater impact on the risk of all subsequent nodes, and hence applying defences against *earlier* attacks are relatively more effective. The control c_1 combats a common weakness w_1 at both T_1 and T_2, and w_1 provides the initial access to the system for the adversary to induce further attacks.

Looking at the tolerance against zero-day attacks, implementing no control 00000 is obviously one of the worst cases. The control c_1 yields a tolerance $\langle 0.36, 0.23, 0.44, 0.57\rangle$, indicating certain sophistication of each zero-day exploits is required to individually violate κ. Deploying c_1 and c_5 further enhances the tolerance to $\langle 0.43, 0.31, 0.56, 0.87\rangle$. 11000 makes the system be fully tolerant of a zero-day at T_4 or T_5 (least threatening ones). At least 11110 is needed for the system to be tolerant of a zero-day at any target.

Two radar charts are shown in Fig. 5 and 6 to provide an intuitive way to visualise the tolerance at the four different targets. The 100% coverage corresponds to the symbol max in tolerance tuples. Fig. 5 shows that deploying more controls does not always guarantee a larger tolerance coverage. The defence plans with c_1 involved tend to be most effective ones in terms of risk reduction and tolerance coverage. 10100 is able to fully protect the system from the zero-day exploit at T_4 because c_1 defends both T_1 and T_2 (where all attacks have to pass through), and c_3 further defends T_3, which greatly limits the damage the zero-day at T_4 can cause to their subsequent sharing requirement node. The tolerance of four effective plans (in terms of $\bar{\Delta}$) is drawn in Fig. 6. The

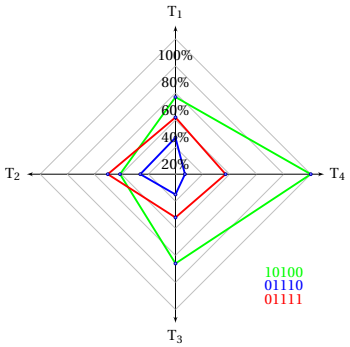


Fig. 5: tolerance coverage on each target

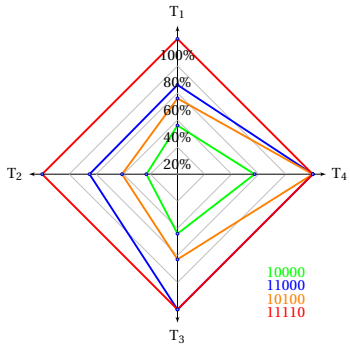


Fig. 6: comparing plans of high $\bar{\Delta}$

coverage against four targets are expanded at various rates. The zero-day exploit at T_4 seems to be the easiest one to be defended, while T_1 and T_2 are the most difficult ones. Three out of the four plans in Fig. 6 make the system immune from the zero-day exploit at T_4, but only 11110 can protect the system from the zero-day exploits at T_1 and T_2.

4 Related Work

Bayesian Networks (BN) have been widely applied in complex systems modelling, risk assessment and diagnostics [22]. BN offer a graphical model to combine and capture complex causal relations between various factors. BN are also able to reason about possible effects and update beliefs in the light of emerging events (e.g. deploying controls), and finally produce decisions that are visible and auditable. *Bayesian Attack Graphs* (BAG) was introduced by Liu and Man [15] by combining attack paths and Bayesian inference methods for probabilistic analysis. Poolsappasit et al. [18] introduced static risk analysis and dynamic risk analysis based on BAG in order to find the most cost-efficient security plans. Muñoz-González et al. [16] further improved the work by providing an efficient probabilistic inference mechanism based on *Belief Propagation* and *Junction Tree* to compute unconditional probabilities. BN were used in [12] to study interdependencies between safety and security for CPS, where the main focus is on analysing the impact of different factors on safety and security. By contrast, we explicitly modelled all possible attack paths by exploiting a chain of known or unknown weaknesses, and evaluated the damage of such cyber attacks against key safety-related requirements.

Combating zero-day attacks has attracted an increasing attention. Wang et al. [20] present a novel security metric *k-zero day safety* to count the minimum number of zero-day vulnerabilities required for compromising network assets. The following work in [21] evaluated the robustness of networks against zero-day attacks in terms of network diversity. Particularly network diversity was formally defined as a security metric by the number of distinct resources, the least and average attacking effort. With regard to the most effective defence for ICS, our work focused on the impact of deployed controls on mitigating the risk from zero-day attacks. Fielder et al. [7] compared three key decision making techniques (i.e. game theory, combinatorial optimisation and a hybrid of the two) to find effective defence for security managers. The work [5] provided a

co-evolutionary agent-based simulation to find optimal defences for ICS, and then [6] considered the cost-effectiveness of defences in various zones of ICS.

5 Conclusion and Future Work

In this paper we studied the possibility of improving the tolerance of ICS against zero-day attacks by means of defending against known weaknesses. We first formally defined the tolerance as a metric by the minimum required exploitability of a zero-day exploit to bring the system into a critical state. Such a metric captures the required zero-day attacking effort, and hence higher tolerance indicates more effort should be invested by an adversary to discover a more sophisticated zero-day flaw. Tolerance against the zero-day exploits at different assets is diverse, depending on the topological position and known weaknesses of an asset. We further built a simulation based on Bayesian Networks to analyse the zero-day threat propagation across ICS. Attackers are able to choose a known or a zero-day (if there is one) weakness at each step, to propagate the risk from one target to the next. Depending on the exploitability of the chosen weakness and its previous exploited targets, the probability of success can be computed. A complete attack path needs to successfully exploit a chain of such weaknesses to reach the final valuable targets of ICS. Deploying security controls combating known weaknesses at each step could actually reduce the chance of the whole attack path being breached. In this case, higher exploitability of zero-day weaknesses is required to reach the same risk level, which means the tolerance of the system against zero-day exploits has been improved. Our approach is able to find the most effective combination of available defence controls to maximize the tolerance and the zero-day attacking effort. A case study about security management of ICS was also demonstrated in this paper.

There are several promising lines of research following this work: (i) we currently considered only the individual zero-day weakness at different targets, and we will explore the consequence of combining multiple zero-day exploits. (ii) intelligent adversarial models are needed to decide the likelihood of different attack paths. (iii) as addressed in [14], security controls might have negative impact on the other criteria of ICS. We will look for possible extensions to model those criteria into the simulation. The cost of deploying controls is also an essential factor to decide the most effective defence, which will be considered in our future work.

Acknowledgement

This work is funded by the EPSRC project RITICS: Trustworthy Industrial Control Systems (EP/L021013/1).

References

1. BSI. Industrial control system security top 10 threats and countermeasures 2014, mar 2014. "www.allianz-fuer-cybersicherheit.de/ACS/DE/_downloads/techniker/hardware/BSI-CS_005E.pdf".
2. S. Christey, R. Glenn, et al. Common weakness enumeration, 2013.

3. U. S. Department of Homeland Security. Common cybersecurity vulnerabilities in industrial control systems, 2011. "www.ics-cert.us-cert.gov/sites/default/files/documents/DHS_Common_Cybersecurity_Vulnerabilities_ICS_20110523.pdf".
4. N. Falliere, L. O. Murchu, and E. Chien. W32. stuxnet dossier. *White paper, Symantec Corp., Security Response*, 5, 2011.
5. A. Fielder, T. Li, and C. Hankin. Defense-in-depth vs. critical component defense for industrial control systems. In *Proceedings of the 4rd International Symposium for ICS & SCADA Cyber Security Research*. British Computer Society, 2016.
6. A. Fielder, T. Li, and C. Hankin. Modelling cost-effectiveness of defenses in industrial control systems. In *International Conference on Computer Safety, Reliability, and Security*, pages 187–200. Springer, 2016.
7. A. Fielder, E. Panaousis, P. Malacaria, C. Hankin, and F. Smeraldi. Decision support approaches for cyber security investment. *Decision Support Systems*, 86:13–23, 2016.
8. M. Frigault and L. Wang. Measuring network security using bayesian network-based attack graphs. In *2008 32nd Annual IEEE International Computer Software and Applications Conference*, pages 698–703, July 2008.
9. Hugin Expert A/S. Hugin lite 8.3. *http://www.hugin.com*, 2016.
10. ICS-CERT. Incident response activity july 2015 - august 2015, 2015. "https://ics-cert.us-cert.gov/monitors/ICS-MM201508".
11. ICS-CERT. Incident response activity september 2014 - february 2015, 2015. "www.ics-cert.us-cert.gov/monitors/ICS-MM201502".
12. A. J. Kornecki, N. Subramanian, and J. Zalewski. Studying interrelationships of safety and security for software assurance in cyber-physical systems: Approach based on bayesian belief networks. In *Computer Science and Information Systems (FedCSIS), 2013 Federated Conference on*, pages 1393–1399. IEEE, 2013.
13. R. Langer. *Robust Control System Networks-How to achieve reliable control after Stuxnet*. New York, New York, United States: Momentum Press, 2012.
14. T. Li and C. Hankin. A model-based approach to interdependency between safety and security in ics. In *Proceedings of the 3rd International Symposium for ICS & SCADA Cyber Security Research*, pages 31–41. British Computer Society, 2015.
15. Y. Liu and H. Man. Network vulnerability assessment using bayesian networks. In *Defense and Security*, pages 61–71. International Society for Optics and Photonics, 2005.
16. L. Muñoz-González, D. Sgandurra, M. Barrère, and E. Lupu. Exact inference techniques for the dynamic analysis of attack graphs. *arXiv preprint arXiv:1510.02427*, 2015.
17. J. Pearl. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann, 2014.
18. N. Poolsappasit, R. Dewri, and I. Ray. Dynamic security risk management using bayesian attack graphs. *IEEE Transactions on Dependable and Secure Computing*, 9(1):61–74, 2012.
19. K. Stouffer, J. Falco, and K. Scarfone. Guide to industrial control systems (ics) security. *NIST special publication*, 2011. "http://csrc.nist.gov/publications/nistpubs/800-82/SP800-82-final.pdf".
20. L. Wang, S. Jajodia, A. Singhal, P. Cheng, and S. Noel. k-zero day safety: A network security metric for measuring the risk of unknown vulnerabilities. *Dependable and Secure Computing, IEEE Transactions on*, 11(1):30–44, 2014.
21. L. Wang, M. Zhang, S. Jajodia, A. Singhal, and M. Albanese. Modeling network diversity for evaluating the robustness of networks against zero-day attacks. In *Computer Security-ESORICS 2014*, pages 494–511. Springer, 2014.
22. P. Weber, G. Medina-Oliva, C. Simon, and B. Iung. Overview on bayesian networks applications for dependability, risk analysis and maintenance areas. *Engineering Applications of Artificial Intelligence*, 25(4):671–682, 2012.