11TH INTERNATIONAL CONFERENCE
ON CRITICAL INFORMATION
INFRASTRUCTURES
SECURITY

10-12 October 2016
UIC HQ Paris

CRITIS
2016

# Effective Defence Against Zero-Day Exploits Using Bayesian Networks

**Tingting Li** and Chris Hankin

*Institute for Security Science and Technology*
*Imperial College London*

RITICS Imperial College London

# Research Institute in Trustworthy Industrial Control Systems

EPSRC Engineering and Physical Sciences Research Council | CPNI Centre for the Protection of National Infrastructure | CRITIS 2016

£2.4M programme, 5 coordinated projects.

Phase 1 (Directorship) awarded 01/01/14, Chris Hankin, Imperial College London.

Phase 2 awarded 01/10/14.

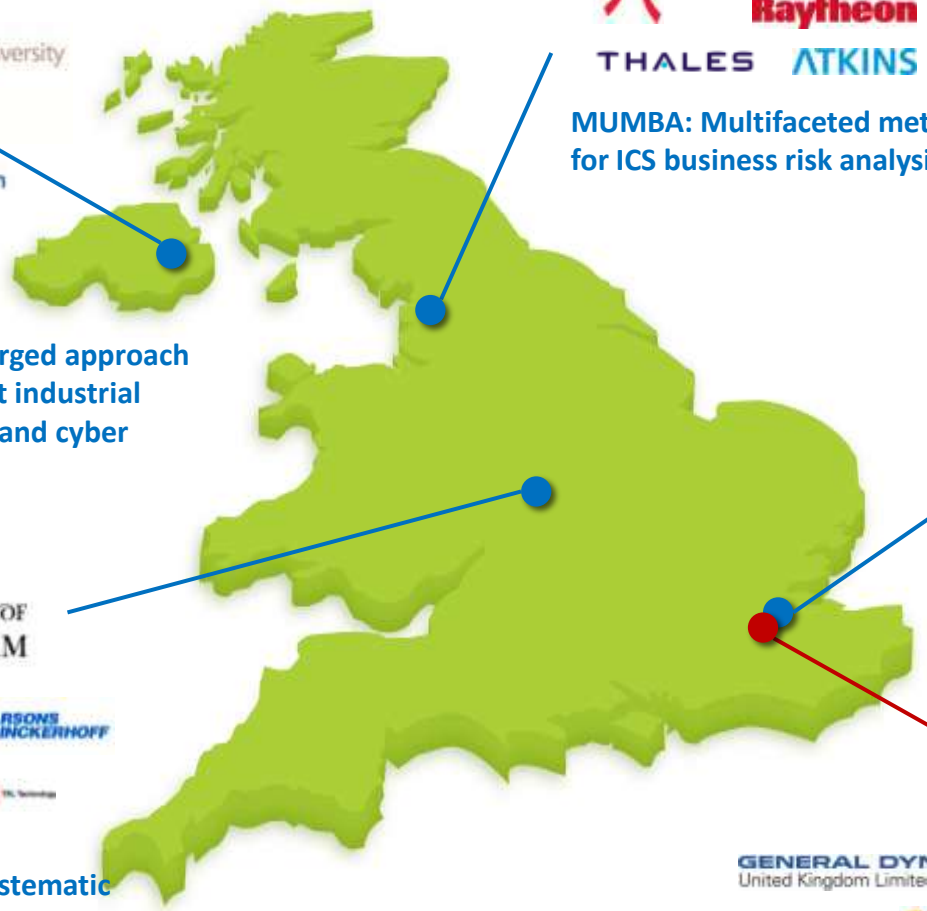

Lancaster University, Airbus, Raytheon, Thales, Atkins

**MUMBA: Multifaceted metrics for ICS business risk analysis**

Queen's University Belfast, Scottish and Southern Energy Power Distribution, Thales, Statnett

**CAPRICA: Converged approach towards resilient industrial control systems and cyber assurance**

City University London

**CEDRICS: Communicating and evaluating cyber risk and dependencies in ICS**

Alstom, Adelard

University of Birmingham, nationalgrid, Parsons Brinckerhoff, RSSB

**SCEPTICS: A systematic evaluation process for threats to ICS (incl. national grid and rail networks)**

Imperial College London

General Dynamics United Kingdom Limited, Atkins, Symantec.

**RITICS@Imperial: produce models and tools in support of effective interventions.**

# Background

- Zero-day vulnerabilities are *unknown* or *undisclosed* and there is *no patch* available to fix them.
- Zero-day exploits have demonstrated their essential contributions in *Stuxnet*.
  - *Four different zero-day* vulnerabilities to gain access to targets and propagate. [Symantec 2011]

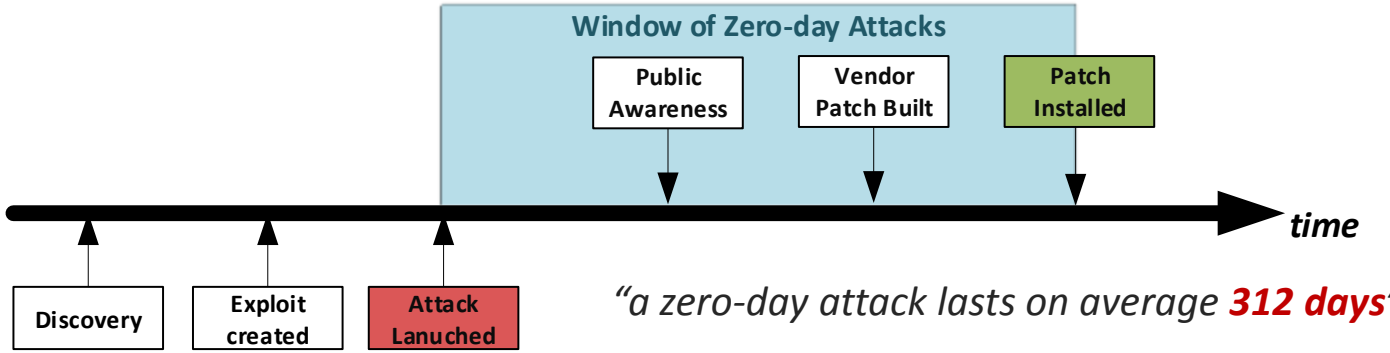


*"a zero-day attack lasts on average **312 days**"* [Bilge & Dumitras 2012]

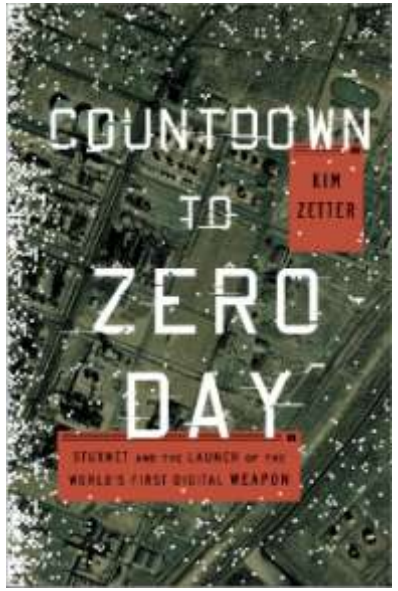


- The *threat* from zero-day exploits is still *on the rise*.
  - 38% of the 245 reported incidents to ICS-CERT have "unknown access vector". [ICS-CERT 2015]
  - In August 2015, ICS-CERT released six advisories & six alerts about zero-day vulnerabilities.
  - Average price up to $100,000. [Wikipedia]
- *Defending* against zero-day attacks is *a challenging task* [FireEye]
  - Traditional signature-based security measures are incapable of combating zero-day attacks.
  - Skilled hackers and their malware can go undetected for months, even years…

# Proposed Approach

- Difficult to directly stop zero-day. *Can we make ICS sufficiently robust against zero-day ?*
- A typical APT exploits *a chain of vulnerabilities*: either *known* or *zero-day* to propagate.
- Alternatively deploy defences strategically against the known vulnerabilities
  - the likelihood of the whole attack chain being exploited can be *overall reduced*.

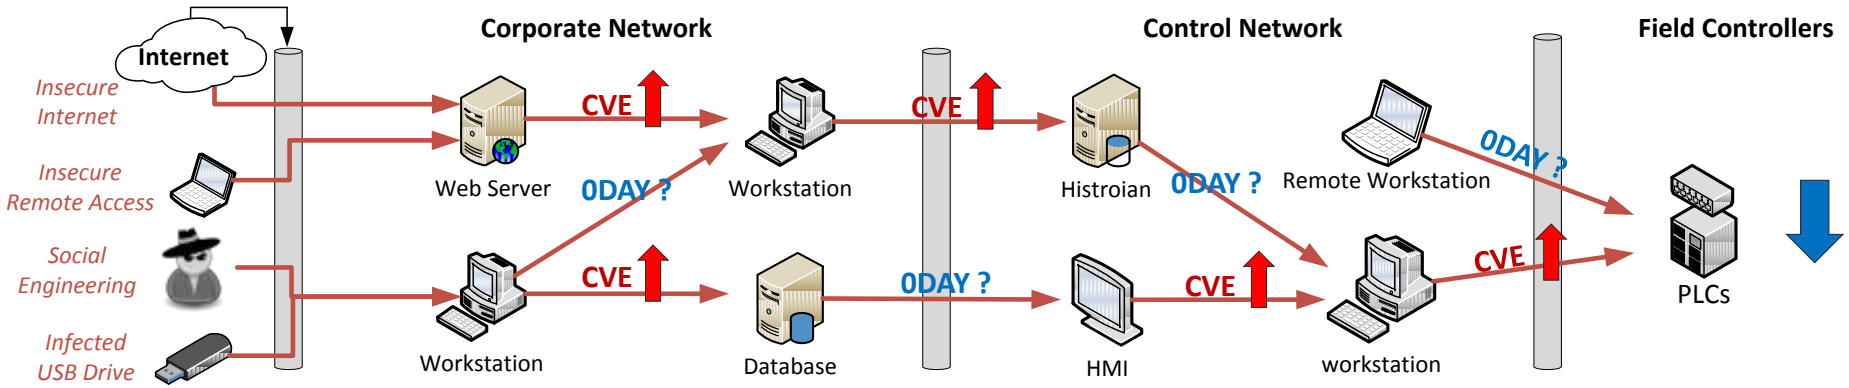


- The "*exploitability*" (from CWE) reflects *sophistication* and the *required attacking effort* of a 0day weakness.
- A security metric "*tolerance against zero-day exploits*" is defined by *the minimal required exploitability* of the zero-day exploit to cause the system risk *exceed the acceptable level*.
- By using *Bayesian Networks*, we can prove that defending against known weaknesses can
  - *increase* and *maximize* the tolerance.

# Modelling and Representation

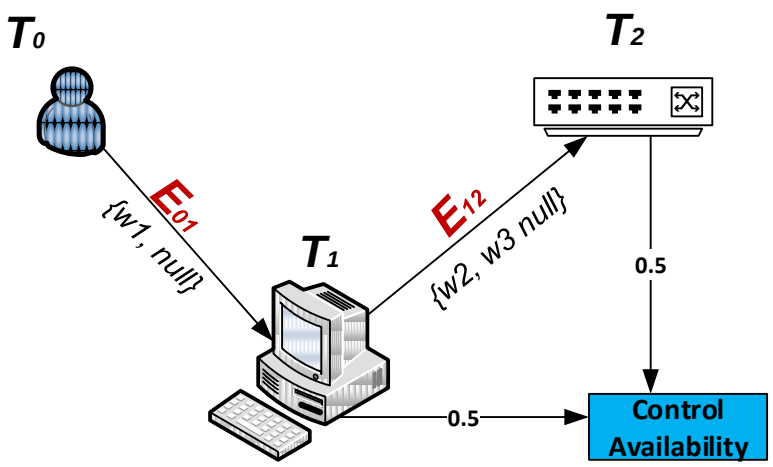


(a) Toy Case

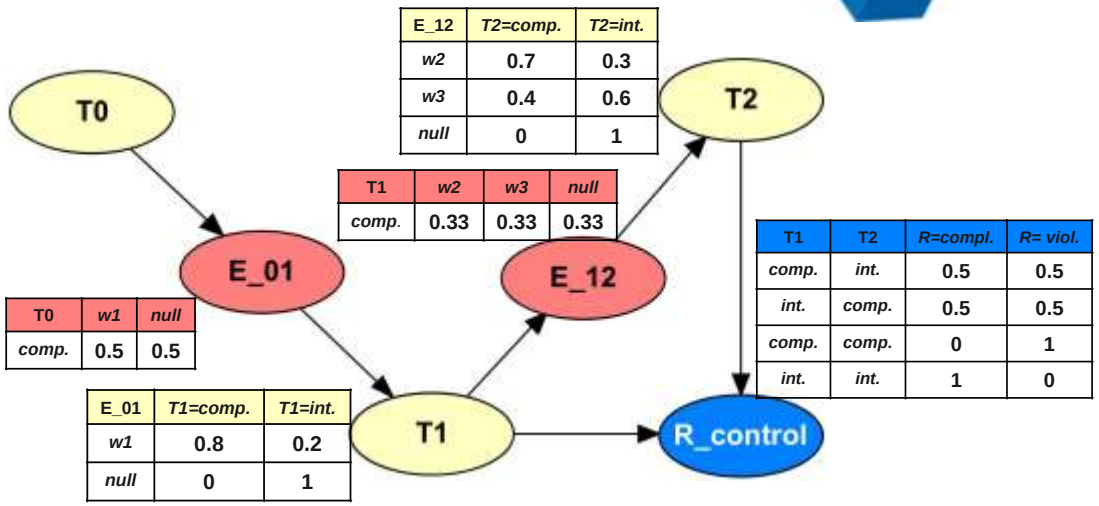


| E_12 | T2=comp. | T2=int. |
|---|---|---|
| w2 | 0.7 | 0.3 |
| w3 | 0.4 | 0.6 |
| null | 0 | 1 |

| T1 | w2 | w3 | null |
|---|---|---|---|
| comp. | 0.33 | 0.33 | 0.33 |

| T0 | w1 | null |
|---|---|---|
| comp. | 0.5 | 0.5 |

| E_01 | T1=comp. | T1=int. |
|---|---|---|
| w1 | 0.8 | 0.2 |
| null | 0 | 1 |

| T1 | T2 | R=compl. | R= viol. |
|---|---|---|---|
| comp. | int. | 0.5 | 0.5 |
| int. | comp. | 0.5 | 0.5 |
| comp. | comp. | 0 | 1 |
| int. | int. | 1 | 0 |

(b) Toy Case- Bayesian Risk Network (aprior )

**Definition 1.** 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 $\mathcal{N} = \mathcal{T} \cup \mathcal{E} \cup \mathcal{R}$:

- **Target Nodes** $\mathcal{T} = \{T_0, \ldots\ T_n\}$ for valuable assets;
  $\mathcal{P}_T$ – distribution for non-root targets over states $\Omega(T_x) = \{c, i\}$.
- **Attack Nodes** $\mathcal{E} = \{E_{01}, \ldots\ E_{mn}\}$ for weaknesses between targets;
  $\mathcal{P}_E$ – uniform distribution over $\Omega(E_{T_i T_j}) = \omega(T_i, T_j) \cup \{null\}$.
- **Requirement Nodes** $\mathcal{R} = \{R_0, \ldots\ R_k\}$ for safety-related requirments.
  $\mathcal{P}_R$ – distribution for all requirment nodes over states $\Omega(R_x) = \{c, v\}$.
- $P_{T_0}$ is the prior probability distribution of the root node $T_0$

**Definition 2.** Let $W = \{w_1, \ldots, w_m\}$ be a set of weaknesses, $\omega : \mathcal{T} \times \mathcal{T} \to 2^W$ gives possible weaknesses that can be exploited from one target to another. The **exploitability** $e_{w_x} \in [0, 1]$ of a weakness $w_x$ is the likelihood of $w_x$ being successfully exploited.

# Modelling and Representation

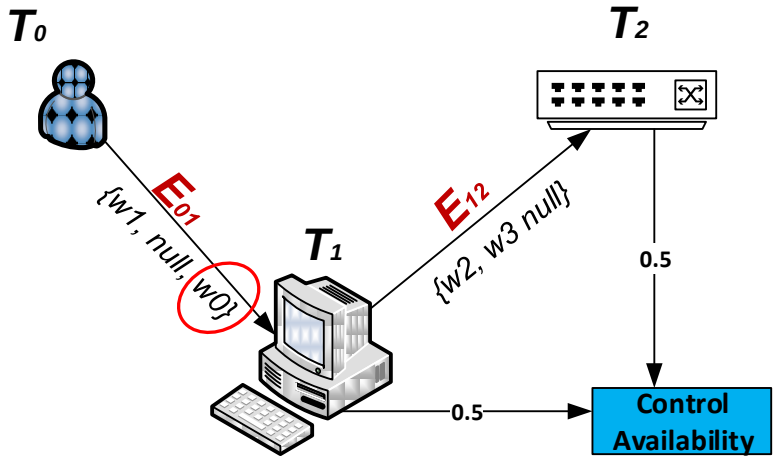


**(a) Toy Case with *a 0day at T1***

- *Define the risk* as the likelihood of a node being compromised/violated.
- The presence of a *zero-day exploit* would *increase the risk* as its exploitability increases.
- The *tolerance* of a zero-day exploit is the *minimum required exploitability* to reach the risk boundary *L*.
- Deploying controls can *reduce the exploitability* of known weaknesses to *increase the tolerance*.

**Definition 3.** A Bayesian Risk Network $\mathcal{B}$ is constructed for a given system. The **tolerance against zero-day attacks** of the system is represented by $(\kappa, Z)$, where

- $\kappa := P(N_a = s) \leqslant L$, an acceptable level of risk, $N_a \in \mathcal{T} \cup \mathcal{R}$, $s \in \Omega(N_a)$, $L \in [0, 1]$.
- $Z := \langle z_1, \ldots, z_n \rangle$, each element $z_i \in Z$ is the tolerance against the zero-day exploit at an arbitrary target $T_i \in \mathcal{T}$ and obtained by :

$$z_i = \underset{P(T_i | E_{T_i' T_i} = w_z)}{\operatorname{argmax}} \quad \kappa := P(N_a = s) \leqslant L$$

$z_i$ is the maximum exploitability of $w_z$ subject to $\kappa$.

| | 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% |

$\kappa := P(R_{\text{control}=viol.}) \leq 27\%$
$\text{eff}(c_x) = 0.1$

**(b) Toy Case- *Tolerance without/with controls***

# Case Study – ICS Security Management

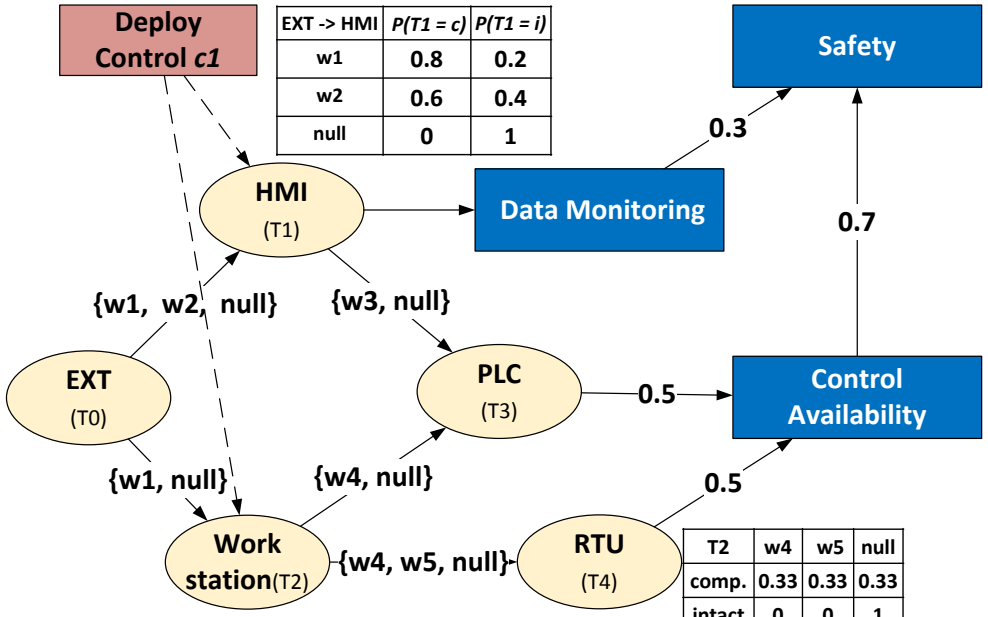


| EXT -> HMI | P(T1 = c) | P(T1 = i) |
|---|---|---|
| w1 | 0.8 | 0.2 |
| w2 | 0.6 | 0.4 |
| null | 0 | 1 |

| T2 | w4 | w5 | null |
|---|---|---|---|
| comp. | 0.33 | 0.33 | 0.33 |
| intact | 0 | 0 | 1 |

(a) network

| Weakness | Description | Location | Exploitability |
|---|---|---|---|
| w1 | Internet Malware | T1, T2 | 0.8 |
| w2 | Removable Drive Malware | T1 | 0.6 |
| w3 | DoS Attacks | T3 | 0.7 |
| w4 | Buffer overflow | T3, T4 | 0.8 |
| w5 | Man-in-the-middle | T4 | 0.6 |

(b) selective common weaknesses

| Control | Control Descriptions | Combat |
|---|---|---|
| c1 | anti-virus software | w1 |
| c2 | removable media check/disable | w2 |
| c3 | IDS and Firewalls | w3 |
| c4 | data validation & software security analysis | w4 |
| c5 | fine-grained access controls & integrity check | w5 |

(c) selective common controls

Covert CWE attribute *"Likelihood of Exploit"* to "*exploitability*" :

- Very High → 0.8
- High → 0.7
- Moderate → 0.6

- **Target nodes:** External, HMI, Workstation, PLC and RTU.
- Select five common ***weaknesses*** and ***countermeasures*** from *ICS Top 10 Threats and Countermeasures [1]* and *Common Cybersecurity Vulnerabilities in ICS [2]*.
- Safety jointly relies on data monitoring (30%) and control availability (70%).
- *PLC* and *RTU* equally contribute to the requirement on control availability.

# Case Study – Posterior Risk Distribution

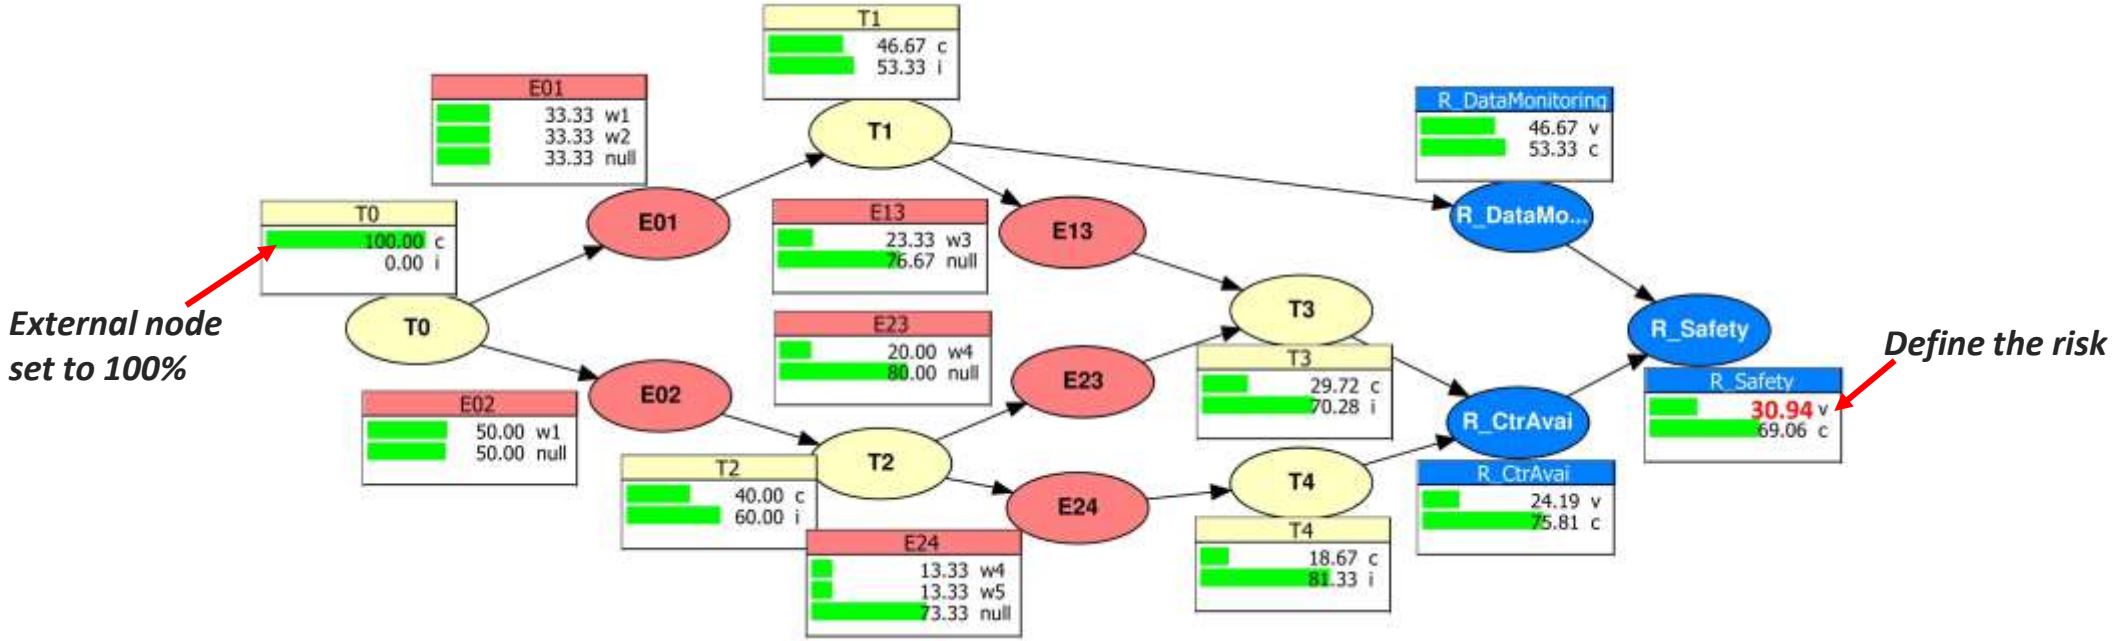


**(a) posterior risk distribution with no control deployed.**

- T0 is the untrusted external environment where attackers can launch any attacks (aprior set to 100%).
- The risk is defined by the chance of the safety requirement (*R_Safety*) being violated.
- Without any control deployed or any zero-day exploits, the current posterior risk is about 30.94%.

# Case Study – Deploying a Single Control

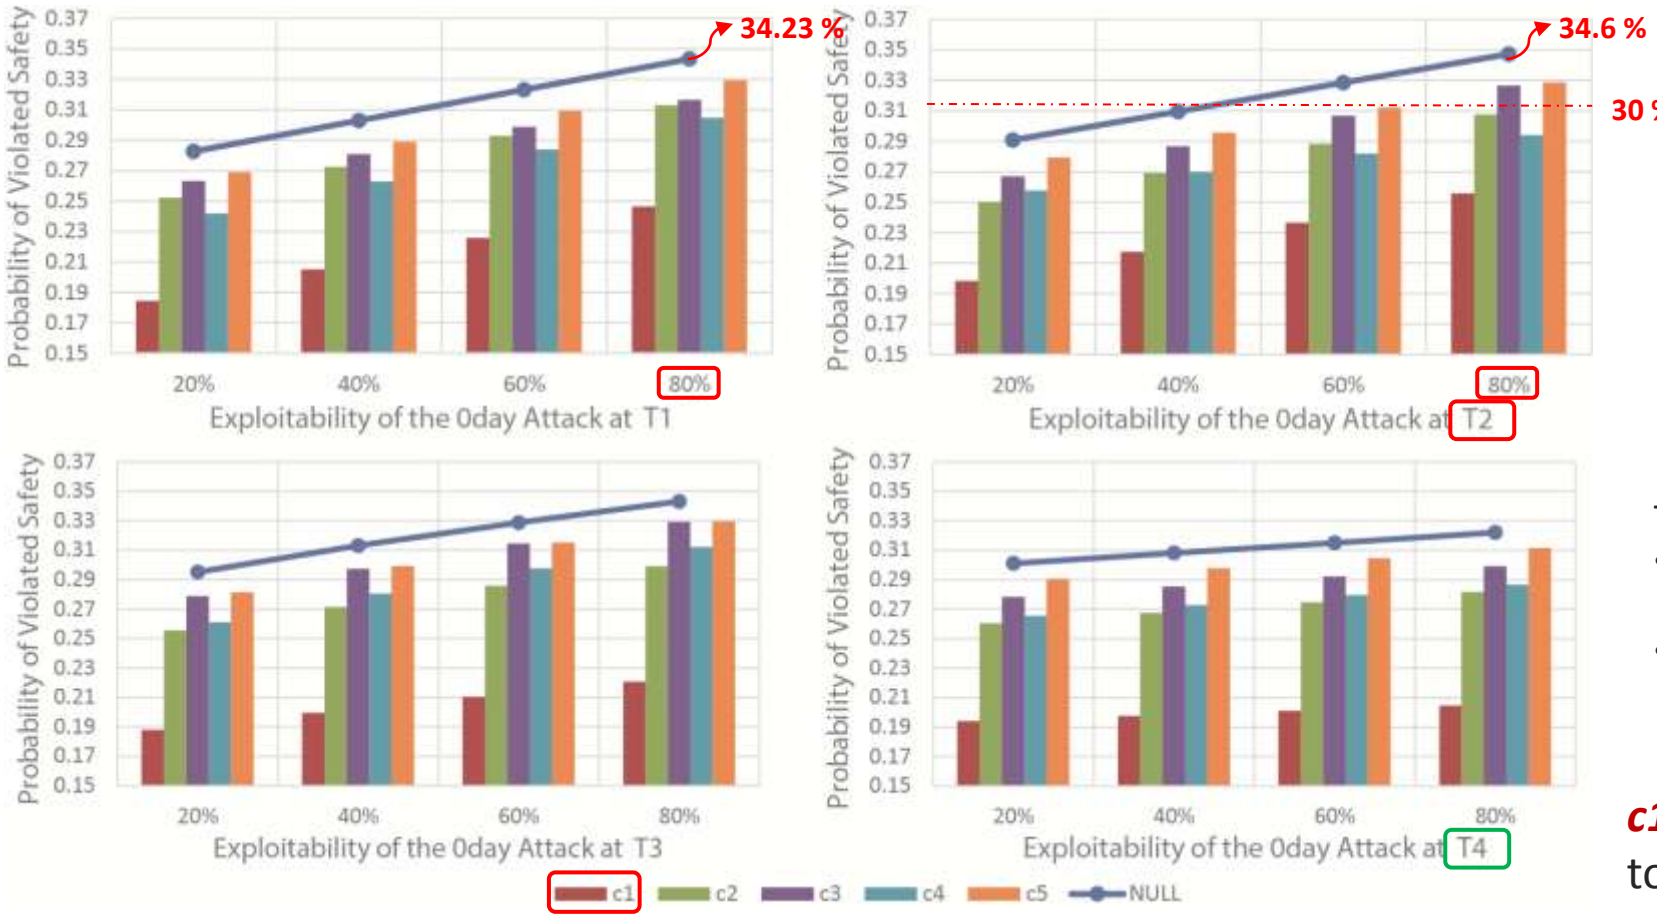


The zero-day exploits generally increase the risk of the system

- the risk (30.94%) is raised to 34.23% with a zero-day of 80% exploitability at T1.
- the risk (30.94%) is raised to 34.6% with a zero-day of 80% exploitability at T2.

Tolerance has been improved:

- With no control, the zero-day at T2 with > 31% is needed.
- With control c2, the zero-day at T2 with much higher >74% is needed.

***c1*** is the most effective control to mitigate the risk of zero-day

- Four experiments with a zero-day added at each target in respective experiment.
- Deploying a single control updated the risks over scaled exploitabilities of the zero-day exploit (i.e. 20%, 40%, 60%, 80%).
- The tolerance against zero-day has been improved (subject to the risk measured by $P(R_{\text{safety}} = violated) \leq 30\%$)

# Case Study – Deploy Combined Controls

*Mean risk reduction at the target*

*Mean risk reduction over all targets*

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 |

*Risk already exceeds the level regardless of any 0day.*

*Max risk incurred by the 0day*

*Fully tolerant of a 0day at the target*

*d = {~~c1~~, c2, c3, ~~c4~~, c5}*

- Five controls give 32 different defence plans. Bit vectors represent to include or not a particular control in a plan.
- The risk acceptable level is given by $P(R_{\text{safety}=v}) \leq 20\%$.

# Case Study – Deploy Combined Controls

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 |

$|d| = 1$ → 10000

$|d| = 2$ → 11000

$|d| = 3$ → 11010

$|d| = 4$ → 11110

Baseline with no control applied

Most effective plan with different # of controls

More controls do not always produce stronger defence (01101 vs. 10000)

- **11000:** fully tolerant of a 0day at T4 or T5,
- at least **11110** is needed to be resistant to 0day at any target.

- Five controls give 32 different defence plans. Bit vectors represent to include or not a particular control in a plan.
- The risk acceptable level is given by $P(R_{\text{safety}=v}) \leq 20\%$.

# Case Study – Tolerance Coverage

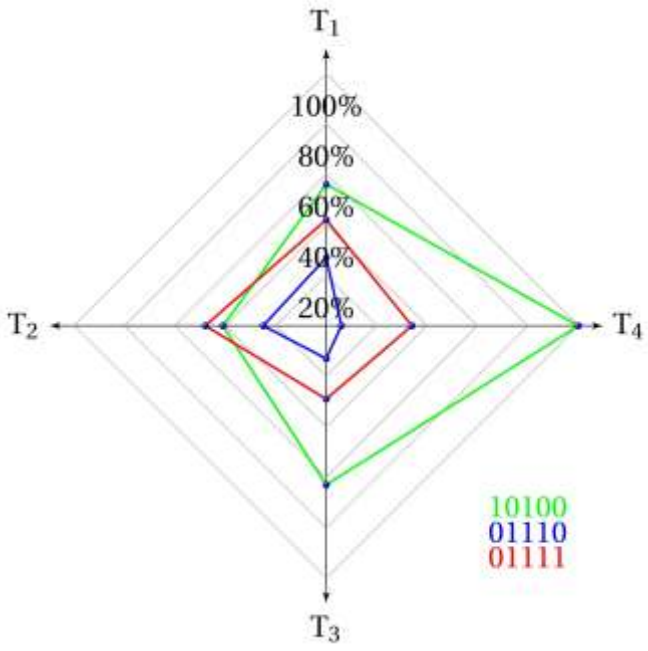


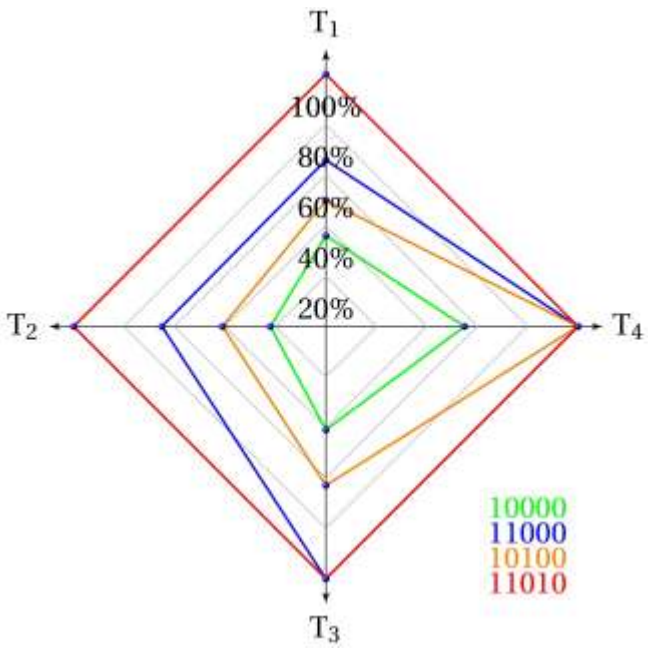


- Map *tolerance* tuples to the *coverage percentage* of a radar chart.
- The *max* in tolerance tuples corresponds to *100% coverage*.
- The coverage against the 0day at four different targets are expanded at various rates.
- **Left:** more controls do not always guarantee a larger tolerance coverage.
- **Right**: the 0day at *T4* seems to be the easiest one to be defended, while the 0day at *T1&T2* are the most difficult ones.

# Conclusion

- Improved the tolerance against zero-day attacks by defending against known weaknesses.
- Defined the tolerance as a metric by the minimum required exploitability of a zero-day exploit to bring the system into a critical state.
- Higher tolerance required more attacking effort to discover more sophisticated zero-day flaws.
- A simulation based on Bayesian Networks was built to model the risk propagation of zero-day attacks, and the reduced risk by deploying difference defence plans.
- Found the effective combination of available defence controls that maximizes the tolerance.

# Future Work

- *Ongoing*: Combating Zero-day Exploits by Moving Target Defence.
  - MTD covers *Diversity*, Shuffle and Redundancy [Hong & King 2016].
  - Attacker: less effort is needed to exploit a vulnerability already encountered
  - Defender: diversifying a target consumes limited budget and increases maintenance cost.
  - Question: most cost-efficient way to increase zero-day tolerance.
- *Ongoing*: Attack Simulation with Zero-day Life-cycle Model
  - Discovery rate, Exploit creation rate, Patch development rate and Patch application rate.
  - Different sizes of attack window, and different decision-making for defense.

- Optimization to find the way to maximize the tolerance.
- Adversarial modelling to better mimic the most likely paths for different attackers.
- More defence might not give stronger defence (defence-in-depth and other alternatives).
- Call for realistic cases and examples!!

Thank you very much !