Continuous Risk Management for Industrial IoT: a Methodological View

Carolina Adaros-Boye, Paul Kearney, and Mark Josephs

Birmingham City University, Birmingham, B4 7XG, UK
carolina.adarosboye@mail.bcu.ac.uk
paul.3.kearney@bcu.ac.uk
mark.josephs@bcu.ac.uk

Abstract. Emergent cyber-attacks and exploits targeting Operational Technologies (OT) call for a proactive risk management approach. The convergence between OT and the Internet-of-Things in industries introduces new opportunities for cyber-attacks that have the potential to disrupt time-critical and hazardous processes. This paper proposes a methodology to adapt traditional risk management standards to work in a continuous fashion. Monitoring of risk factors is based on incident and event management tools, and misbehaviour detection to address cyber-physical systems' security gaps. Another source of information that can enhance this approach is threat intelligence. Risks are calculated using Bayesian Networks.

Keywords: cybersecurity, risk monitoring, IoT, IIoT, ICS

1 Introduction

In automated and smart industries, a certain level of cyber-risk is usually accepted, especially if security controls have conflicts with safety, performance, or availability [4]. For example, security updates can compromise reliability of operations [8], defences against brute force attack can lock out valid users in the middle of a crisis, and control flow integrity checks can present throughput overheads [15]. Finding methods to monitor cybersecurity risks continuously is relevant to allow better preparedness in the case of cyber-threats, particularly in cases of risks with a low likelihood but a high impact.

Many companies are engaging in Industry 4.0, or connecting their existing Operational Technologies (OT) to their information systems. This paradigm, also known as the Industrial Internet of Things (IIoT), is used in various business domains including industries such as electrical, utilities, transport, manufacturing, and building management, among others. These systems differ from typical IT because their core functionalities are based on cyber-physical systems. Industrial operations have requirements that introduce challenges in the application of security controls and many legacy cyber-physical systems were not originally designed having security in mind [15][5]. Limitations on memory and processing capabilities plus real-time response requirements restrict having strong authentication or encryption. The attack surface is also larger than in traditional IT because of the additional communication and information processing layers. Multiple access points, and numerous nodes make difficult to monitor and control all the devices connected to the network and who accesses them. They also work with specific purpose communication protocols which are unknown to traditional network security tools [7]. Some widely used protocols, such as Modbus, Profinet, and Bacnet, among others, do not consider basic security controls such as authentication and data integrity checking.

Compensating controls such as rigorous physical security, network segmentation, and continuous monitoring can counteract the lack of built-in security in Industrial Control Systems (ICS) [2]. However, access points enabled for maintenance, configuration and support activities performed by external or internal personnel can also be a threat vector for an incident either malicious or unintended. However, even if rigorous security controls are in place, organisations still will need to assume a certain level of risk, since it is impossible to cover all the possible flanks of attack. Therefore, the main objective of our research is to find mechanisms to continuously monitor security risks and increase cyber-situational awareness in IIoT. Our key questions are: What information do you need to know in order to monitor security risks in ICS? How can that information be derived from what you can actually measure? How can existing cyber-risk management frameworks be adapted for a more dynamic risk monitoring? How can these modifications be introduced?

In a previous paper we presented a proposal for a continuous risk assessment method for ICS/IIoT [1]. Detection of anomalies or misbehaviour in the system's physical measurements is considered as a mean to fill gaps left by typical intrusion detection methods. The present paper proposes a process oriented view of this approach based on workflow and descriptions of activities with their expected inputs and outputs. Through this perspective, we identify which activities are covered by a traditional risk management process and which need to be defined. Using the ISO/IEC 27005 [11] standard as a reference, workflows are presented together with a blueprint of how to integrate standard cyber-risk management frameworks with a continuous risk assessment paradigm.

2 Why continuous risk management for ICS?

Risk management deals with the fact of not being able to control all aspects of a situation. Often decisions need to be made with incomplete information. Internal and external conditions are always changing, and so is the availability and accuracy of information. Thus, monitoring risks can help to validate assumptions and to check if a risk is becoming more likely to materialise than initially thought, or even transforming into an imminent issue. A cybersecurity programme addressing known vulnerabilities should help avoiding to become an easy target and deter opportunistic attackers, but not those who might go the extra mile to develop elaborate exploit mechanisms. Examples of this are targeted ICS malware such as Stuxnet, TRITON, and LockerGoga. In the case of ICS, a cyber-attack can cause material or environmental damage and even compromise human lives and safety. Abuse of privileges from an insider is also considered an important threat in ICS.

Monitoring risks is “maintaining ongoing awareness of an organisation’s risk environment, risk management program, and associated activities to support risk decisions” [3]. Typically this is done as a periodic and discrete activity that is not integrated with operational processes. We propose that risk monitoring should make use of near-real time operational data. Signs of an attack in IIoT can be discovered by monitoring not only network and software related variables, but also sensor’s data. Then, if traditional cybersecurity controls are bypassed, the status of physical variables can give signs that an unusual situation is happening.

Figure 1 proposes an architecture to support a continuous risk monitoring and re-assessment methodology. A risk calculation engine updates key risk indicators based on different sorts of inputs, some of them are updated in a continuous stream and others on a batch basis. The effect of different events on the risk scores is determined by a Bayesian Network which is defined in a setup process establishing relationship between different events. The data capture and processing and alert generation are performed by a SIEM (Security Information and Event Management) tool which is fed by different sources of data. The approach also applies anomaly and misbehaviour detection techniques to physical variables in order to provide independent evidence of possible security issues.

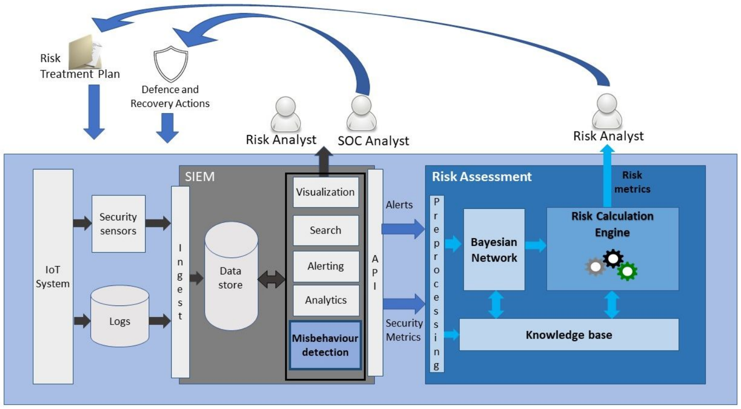

Fig. 1. Architecture of the approached proposed.

A combination of expert’s knowledge and machine learning techniques is used to model the system’s behaviour. For example, an operations expert should provide business rules and functional requirements, including forbidden states and safety considerations, a control and automation expert should know particularities of the programming language of controllers, their memory handling mechanisms, and communication protocols. This would enable the definition of what is normal and abnormal for parameters in different components of a system. Suspicious events in the IT components such as unusual logs, brute force attack access, and malicious code can be identified by a cybersecurity operations expert and monitored by regular off the shelf tools. Continuous updates aim for a better integration of operational processes, risk management, and security processes. This should allow shortening response times. A more detailed description of our ideas on continuous risk assessment for ICS can be found on our previous paper [1].

3 Users, relevant stakeholders, and related processes

The methodology is designed for risk analysts. However, as shown in Figure 1, the results can be shared with the Security Operations Centre (SOC) or equivalent area. Therefore, some of the alerts generated by the continuous risk assessment should be forwarded to them. Any suspicious event or behaviour that changes normal trends or any abrupt changes should be also escalated to the appropriate stakeholders, such as process owners.

Risk management processes in practice are often defined and performed by a separate area from the one dealing with security operations. This has different reasons including that often risk analysts work at a more strategic level and security analysts at a tactical level. Nevertheless, cybersecurity standards regard risk management as an integral part of a cybersecurity management system [10][8][11]. Lack of integration between security risks and operations management can take away the purpose of doing a risk assessment in the first place. Figure 2 represents how cybersecurity risk management should not just overlap but contain security operations management. Updated security metrics can reduce the levels of uncertainty involved in a risk analysis. In our approach, we define three categories of indicators depending on the degree of confidence they can give to predict an attack:Indicators of Risk (IoR), which can modify the estimation of likelihood of an attack, Indicators of Compromise (IoC), which are a type of IoR that can reveal a possible breach in any of its stages, and alerts, that indicate with a high degree of certainty an imminent issue.

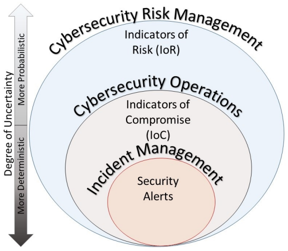

Fig. 2. Indicators hierarchy in a Cybersecurity Management System

Frequent updates of risks can also support decision-makers to make informed and rational choices. Think on a manufacturing company where plant A uses legacy systems, and plant B uses new generation IIoT with built-in security. Plant B is more efficient so they are evaluating to use the same technology in plant A, but the capital expenditure is high. Quantitative information about reduction of cybersecurity risks can contribute in supporting the business case for the migration by increasing the return on the investment calculation.

4 Related work

Given the need of having a better visibility in ICS, it is not surprising to see work proposing continuous assessment methods for cyber-physical security risks. The work done in [14] and [9] propose real-time risk analysis methods based on the human immune system. However, their focus is mostly on network security rather than addressing all the layers of the system, which is required in order to consider all possible cybersecurity risks.

Original approaches have been proposed, as well, to detect anomalies in IoT systems based on their behaviour by analysing sensor data and correlating events [13] [4] and detecting invalid or “prohibited” states of the system [16]. Their work is highly related to the ideas we are proposing and confirms that monitoring risk factors in cyber-physical systems (IoT, IIoT, ICS) in operational or near real time is not a misbegotten idea, despite the practical challenges it may offer.

Among other related research that we have taken into account is the work done in [7] and [5] which highlights challenges and requirements and describe techniques for intrusion detection in industrial cyber-physical systems. Work has been done in this regard by proposing methods to generate dynamic security metrics, including risk metrics in order to help deciding or automatically choose among alternatives for countermeasures [6]. Some of these even go further and suggest ways to deal with unknown threats [12]. This is also a consideration taken in our methodology. Some of these approaches can provide potentially useful techniques and methods to capture real-time security metrics which will be key inputs for our risk calculations. An example of this is the use of Bayesian Networks [18] and fuzzy logic [12] [18]. Nonetheless, many of them cannot be considered as an holistic cyber-risk assessment either by their limited scope [14] [9] or lack of consideration of the context and business impact [18].

To the best of our knowledge, this is the first paper to present a methodological view which, as well adapts the traditional risk management process to continuous operation. This does not contradict the fact that, overall, our work gathers ideas already presented in different publications that have been reviewed as part of the state of the art and that were mentioned in this section.

5 The Continuous Risk Assessment Methodology

Industry standards do not define processes for a continuous risk assessment. However, we found that our approach does not have fundamental contradictions with the contents of frameworks such as IEC 624433, ISO 270005, NIST 800-37, or OG86. Most Risk Management models describe risk assessments as a PDCA cycle (Plan-Do-Check-Act) where risk monitoring (check) has the purpose of identifying changes in risk factors in an early stage [11]. For this, an organisation should identify the reassessment frequency and triggering criteria for cyber-risks where the period reflects the fast-changing nature of cyber security [10]. Rather than re-defining the traditional risk management process we propose to extend it to support methods that work on a continuous basis.

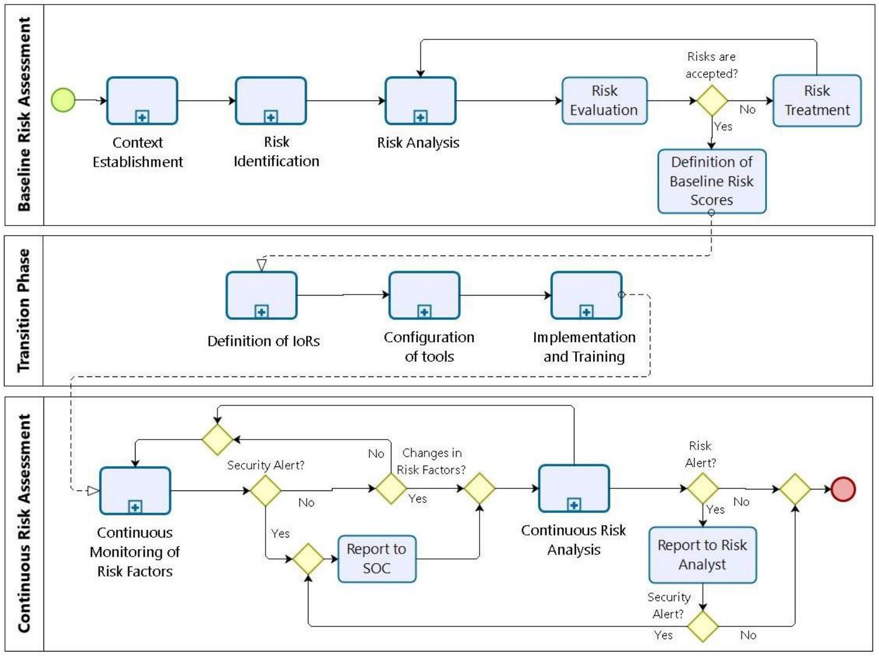

Fig. 3. Macro-process of the methodology.

Our approach proposes three phases: the baseline (initial) risk assessment phase, the transition phase, and the continuous risk assessment phase, performed during normal operation. The baseline risk assessment follows a standard process, as defined by the ISO/IEC 27005 standard. However, additional activities and work products are added as a preparation for the continuous risk assessment. In particular, information needed to model the system's functional behaviour under normal operation is gathered during the Context Establishment. In the transition phase the metrics that will be gathered during the continuous risk assessment are defined and the supporting tools for the continuous risk assessment are configured. In the continuous risk assessment phase the information from the baseline risk analysis gets updated based on operational data. Risk scores are modified when a significant change is detected. Figure 3 shows a workflow of the “macro-process” of the methodology. A modified version of BPMN, which is a widely-used notation for workflow modelling, is used to represent in separate lanes which activities are covered by the ISO/IEC 27005 standard and which are added by us. This is done just for modelling and representation purposes with the understanding that lanes are originally meant to separate roles. In the following sub-sections different sub-process are described for each phase.

5.1 Baseline Risk Assessment

In this phase a risk analysis is undertaken to generate baseline risk scores. Figure 4 shows the different sub-processes involved. In the context establishment all the information relevant to risk management is gathered [11]. This is critical for the success of any risk assessment since it sets the priorities, methods, and risk tolerance level. As it is not feasible to analyse all the possible risks, a scope needs to be defined to focus in the most critical assets, also known as “crown jewels” [2]. Involvement of experts is crucial to define the variables of the system that can be monitored in the continuous phase, and to model their behaviour. The availability of operational data of the system under normal conditions is also important to model the system’s behaviour. Characteristics of the infrastructure, network architecture, business rules, and configuration of the controllers need to be known. Outputs of the Context Establishment added in our approach are a list of sources of data to identify risk factors including online and batch automatic input feeds, as well as manual ones. This information will be used later to define the IoRs.

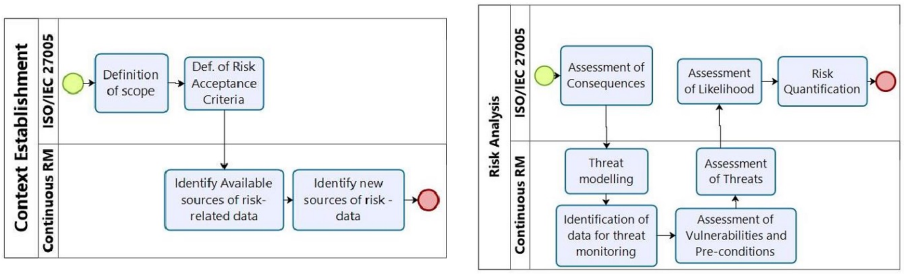
Context Establishment
ISO/IEC 27005
Continuous RM
Definition of scope
Def. of Risk Acceptance Criteria
Identify Available sources of risk-related data
Identify new sources of risk - data
Risk Analysis
ISO/IEC 27005
Continuous RM
Assessment of Consequences
Assessment of Likelihood
Risk Quantification
Threat modelling
Assessment of Threats
Identification of data for threat monitoring
Assessment of Vulnerabilities and Pre-conditions

Fig. 4. Sub-Processes of the Baseline Risk Assessment.

The following step is the risk identification, which is not shown in Figure 4, since it is exactly as in ISO/IEC27005, which means it encompasses the identification of assets, threats, security controls, vulnerabilities, and impacts. In the Risk Analysis process, a threat model should be developed to describe different possible attacks. An attack taxonomy specific to cyber-physical systems needs to be considered, such as the one proposed in [17]. This taxonomy differentiates a ”change” which is an effect at a system level from an ”impact”, which is the implication for the business. Changes that relate to each threat are used to define IoRs, IoCs and alerts. The likelihood of a risk is estimated based on available data such as pentesting results and CVSS scores, for quantification of the vulnerability level, and threat intelligence sources to estimate levels of threat. A Bayesian Network is used to link the dependencies between different factors and their likelihood of occurrence given a series of conditions.

In the Risk Evaluation process, risks are prioritised and categorised. The risks above the acceptable level must be considered for further review and treatment. Risk acceptance should only happen when risks are below a defined threshold. When a risk is reduced to an acceptable level instead of being eliminated it is known as a "residual" risk. In exceptional cases, the organisation might decide to accept a risk that is higher than their accepted level. It must be noted that to acknowledge a high level of risk and accepting it is totally different, in principle, from underestimating a risk. Optimism bias is common when risks are not properly analysed, and can lead to an organisation to be unprepared to deal with security issues when they arise. On the other hand, accepting a risk implies that there is awareness about its potential consequences, but a decision to accept it is formally made. In these cases, risk monitoring becomes a compensating control that allows actions to be taken promptly if there is evidence that the risk is rising to become an issue. In the cases where the risk treatment plan, or parts of it, requires more time to be put in place or to show results, the risk scores will remain over the threshold until the controls are effectively working on reducing it. The Risk Evaluation and Risk treatment processes that we define do not have major variations from the ISO/IEC 27005 standard. The definition of baseline risk scores, provides a benchmark that is based on the available information of the condition of the system in a certain point of time. For this reason, this methodology considers the results of the initial risk assessment not a definitive evaluation but just the setting of initial scores for the continuous risk assessment.

5.2 Transition Phase

In this phase the variables that will be monitored in continuous risk assessment are linked to the risks analysed in the previous stage. The attack model done during the risk analysis should provide a record of the means of detection or inferring possible changes that different elements of the system can suffer in different stages of an attack, as well as the pre-conditions that can allow a certain attack method to be executed. Each change and precondition can be related to a risk metric that termed an IoR.

For indicators related to physical variables, a misbehaviour detection model needs to be defined to address specific needs of ICS. Commonly used means of detection will usually not be sufficient to alert of changes in physical variables. Analysing data from sensors and actuators can allow to detect physical effects of cyber attacks. This is highly relevant, since insiders present important threats in ICS [2] and, as these actors will have valid access, they can bypass traditional security controls. Tools need to be configured, including firewalls, Intrusion Detection Systems (IDS), malware detection, network monitoring, log monitoring, as well as misbehaviour detection. A SIEM will be used to process information and handle alerts. The modules that need to be setup and configured for the next phase are the misbehaviour detection module, the pre-processing of risk metrics and calculation of IoRs.

The transition phase can be complex considering the amount of tools, variables, and methods that the methodology comprises. To overcome this, it is possible to approach it as a project where incremental changes are applied at different stages. For example, start with a reduced scope (e.g. only critical processes and systems). A training period will be necessary for users, as well as for detection algorithms. Tools settings and parameters will need to be adjusted, and calibrated in order to maximise accuracy, and minimise false positives.

5.3 Continuous Risk Assessment

The main processes that interact in this phase are the Continuous Monitoring of Risk Factors, the Continuous Risk Analysis and the generation and report of security and risk alerts. The difference between a security alert and a risk alert is that while the first requires immediate action, the other might not. Figure 5 shows the continuous risk monitoring process where inputs are monitored in order to identify significant changes. At this stage there might be conditions that trigger immediately a security alert that is reported to the SOC, prior to the risk analysis as shown in Figure 3.

Significant changes reported during the continuous risk monitoring process will lead to new values of IoRs and IoCs, and to trigger alerts, if applicable. Adjustments can be either because previous assumptions have been proven inaccurate or biased, or because of internal and external changes. Awareness about relevant changes will be provided by continuous processing and analysis of data from different sources including security sensors and logs, and threat and vulnerability intelligence sources. The use of a SIEM will allow an important amount of the data used for monitoring is supplied in near real time, or in a batch modality, but automatically. However, manual inputs, such as uploading data files and manually changing parameters in absence of the adequate tools, can also be considered. For example, when there is information from a zero-day vulnerability coming from another source this can make it necessary to make manual modifications to some parameters.

Figure 6 shows the continuous risk analysis, where indicators are mapped with their corresponding risk scores whose values are consequently updated. The risk calculation engine recalculates risk scores every time that a condition in the system requires an IoR to change its value. As conditions are monitored continuously, risk scores can be updated at any moment. If any updated risk score exceeds the acceptable level ”risk alert” is triggered. As it is not possible to capture all the possible attack mechanisms in the threat modelling, a special risk score will be defined to represent unknown or zero-day attacks. Combinations of IoRs that cannot be linked to a known risk will be linked to this score.

All alerts get reported to the risk analyst to be analysed in order to consider the implementation of additional security and risk mitigation measures. In the

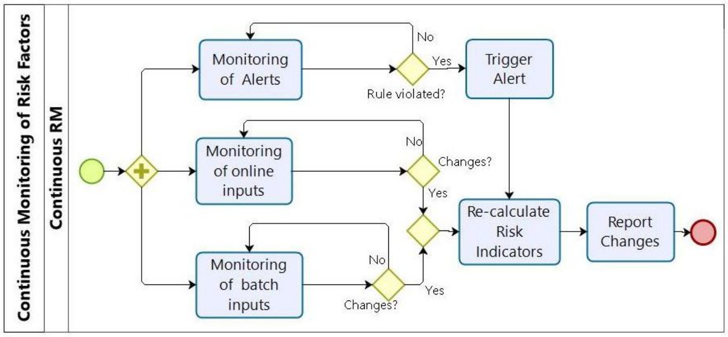
Continuous Monitoring of Risk Factors
Continuous RM
Monitoring of Alerts
Rule violated?
No
Yes
Trigger Alert
Monitoring of online inputs
Changes?
No
Yes
Re-calculate Risk Indicators
Report Changes
Monitoring of batch inputs
Changes?
No
Yes

Fig. 5. Continuous Risk Monitoring

case of security alerts and risk alerts that are judged to require some sort of immediate reaction the SOC analyst is informed. Security alerts can be raised automatically by rules managed by the SIEM, by the results of the risk calculations or by the risk analyst. While this approach does not deal with any of the processes related to the SOC operations, it contributes on giving an instance of collaboration between risk management and security operations.

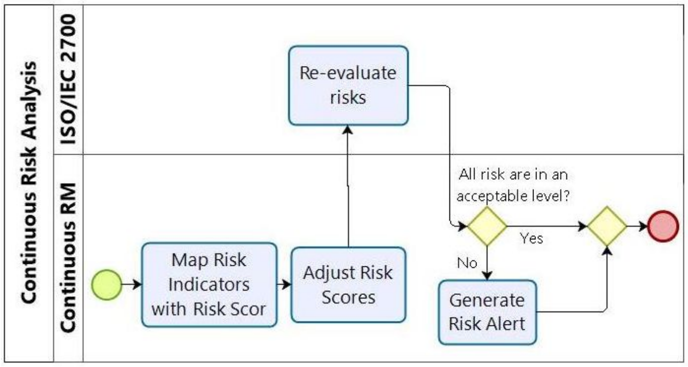

Fig. 6. Continuous Risk Analysis

6 Using Bayesian Networks in risk estimations

The initial risk quantification is based on Bayesian Network whose nodes represent events that can lead to one or more risks. An incidence matrix indicates whether or not there is a relationship between nodes. In the graphical representation of the network the nodes are linked by arrows. For each existing relationship a table of conditional probabilities is defined to calculate the likelihood of an event ”A” given the occurrence of an event ”B”. The overall likelihood of a risk is estimated based on a chain of events. An analogous process is done to link the IoRs to each action allowing to update likelihoods in the presence of IoRs during the continuous risk analysis.

In [1] an example of a temperature control in a Data Centre, supervised and controlled by a Building Management System (BMS) is used as use case to explain continuous risk assessment. A BMS controls environmental variables such as temperature, pressure, and humidity, and it could also monitor variables related to the use of the building such as perimeter security and access control, and utilities consumption, among others. In some contexts the environmental control of the building can be critical, for example in a chemical or nuclear plant. In data centres the environmental control system is used to keep temperature and humidity within levels that allow the servers to perform at their best capacity, and to prevent overheating and other unsafe conditions. In the case of temperature this is usually 25 degrees Celsius.

In this example it is assumed that the temperature setting may be changed from a small number of privileged accounts, which may be accessed remotely. ICS tend to have restricted remote access or to be isolated from corporate and public access networks either physically or through the configuration of a DMZ (demilitarised zone). Nevertheless, it can be the case that remote access enabled for a system’s administrator [2]. It is also often that cyber-hygiene measures that are considered to be basic in IT systems are not be implemented in ICS because they conflict with the availability, integrity and even safety of an operation. For example, a typical defence against brute force attacks is locking access to the system after a number of failed login attempts. However, in an emergency situation which can be time-critical, locking an ICS system because of human mistake such as a forgotten password could be riskier than exposing the system to a brute force attack.

Figure 7 shows a simplified representation of a Bayesian Network to quantify cyber-risks for this case study. Two possible attacks on the temperature control allow an attacker to achieve the goal of disturbing the normal operation of the data centre. The first attack is based on changing the temperature setting to a considerably higher value than the acceptable limit, which is 30 degrees Celsius. The second attack consists of disabling the temperature control. All nodes have states that depend on the achievement of certain goals by the attacker. In each node, the conditional probabilities of an attacker’s goal been achieved is registered in a table given each of the possible combination of states of the previous nodes. When a change is observed, then the likelihood of a possible estate can increase which will means that the estimation of the probabilities that certain attacker’s goal can be achieved increases, as well. This will be computed and, if applicable, presented to the user as an IoR. An attack goal can be associated to more than one IoR and each IoR can be associated with more than one attack goal. Hence, an IoR by itself is not conclusive to detect a certain type of attack, but by having a combination of them the level of uncertainty can be reduced.

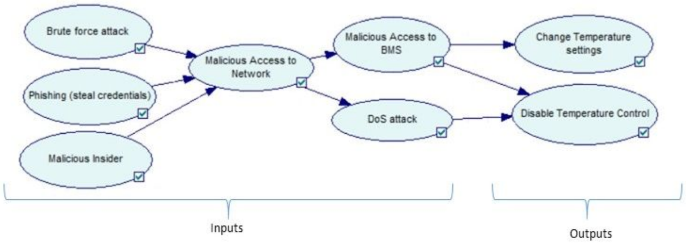
Brute force attack
Phishing (steal credentials)
Malicious Insider
Malicious Access to Network
Malicious Access to BMS
DoS attack
Change Temperature settings
Disable Temperature Control
Inputs
Outputs

Fig. 7. Example of Bayesian Network

In Figure 7 the nodes represent different attack goals and the figure represents a combination of possible kill-chains. The arrows indicate that the probability that the destination action has been performed is directly influenced by whether the source action has been performed. In Figure 8, the nodes on the left hand side of each diagram are associated with actions (as in Figure 7), but the nodes in the right are associated with observations (termed IoR). The arrows denote that the probability that the relevant phenomenon has been observed depends directly on whether the action has been performed. Inferencing is "forward" in Figure 7, and "backward" in the two diagrams of Figure 8, so that in the initial risk analysis inputs nodes are the actions and output nodes the likelihood of a risk and in the continuous risk analysis the input nodes are the IoRs which allow calculating the likelihood of particular actions in a kill chain taking place as well as the associated risks.

Not all means of detection are deterministic and Bayesian Networks allow calculation of the probability of an action being detected. An IoR is not necessarily related to a single action, but can be a symptom of several causes. A Bayesian Network can help inferring a set of possible malicious actions and the likelihood of each one of them taking place. If we take as an example "traffic from an unusual location" as an IoR, from the diagram at the left in Figure 8 it can be observed that it is related to three nodes. So if traffic from an unusual location is detected, this could be evidence of a brute force attack to gain access, or of malicious access, and also could be indicative of a Denial of Service attack (DoS) in its early stages. Nevertheless, when this indicator is correlated with other indicators it can give a more precise information. For example, if there are also multiple failed access attempts followed by a successful one this increases the probability that unauthorised access was gained through a successful brute force attack.

In the case of an insider attack, a different approach needs to be taken since in this scenario an agent that can be either malicious, or coerced to act maliciously, will have valid access credentials. This means that there will be no anomalies at a network or software level. For these cases, it is necessary to count on IoRs that can detect a functional misbehaviour in the system. The diagram at the right in Figure 8 shows a Bayesian network for different IoRs that are calculated through misbehaviour detection. In a sophisticated attack it could be

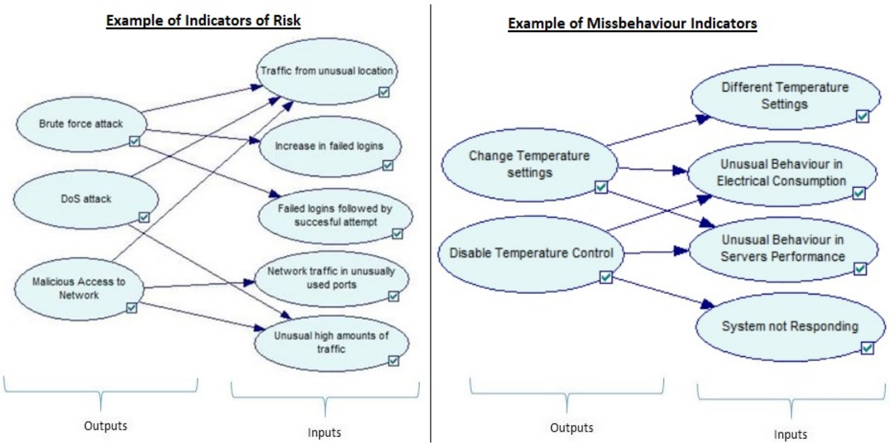
Example of Indicators of Risk
Brute force attack
DoS attack
Malicious Access to Network
Traffic from unusual location
Increase in failed logins
Failed logins followed by succesful attempt
Network traffic in unusually used ports
Unusual high amounts of traffic
Outputs
Inputs
Example of Missbehaviour Indicators
Change Temperature settings
Disable Temperature Control
Different Temperature Settings
Unusual Behaviour in Electrical Consumption
Unusual Behaviour in Servers Performance
System not Responding
Outputs
Inputs

Fig. 8. Link between Changes and Indicators of Risk (IoR)

possible that the attacker tampers with the temperature values displayed in the BMS preventing alerts being triggered directly by the change in the temperature values or settings. An example of this is the case of Stuxnet where normal operational values were replayed in the system while the frequency of the motor drives of the nuclear centrifuges was changed. To address this, anomalies in other variables, such as unusual behaviour in electrical power consumption or in server's performance, could raise an alert concerning the temperature.

A key challenge of risk management is trying to quantify what you do not know. Sometimes there is enough information to identify and analyse a risk by making plausible assumptions. However, at other times there is no prior information at all. This is the case of zero-day attacks. As it is not possible to identify all possible risks in the baseline risk assessment, it is important to consider a likelihood of an undesirable event with undefined characteristics. To address this, every IoR will have among its possible causes an unknown or undefined cause as shown in Figure 9. This has not been included explicitly in previous figures for simplicity of presentation. When a combination of IoRs cannot be mapped to any known risk, the likelihood of an unknown risk unfolding is increased. One or more nodes could also be defined to represent non-cybersecurity related triggers for an IoR.

7 Future work

As the present paper provides just a conceptual approach, considerable work needs to be done to prove specific methods can work under the proposed framework. Future work will focus on further demonstration of the methodology using different use cases related to ICS and IIoT. One of our goals is to explore the correlation of data from physical variables (sensor and actuator data) to reveal

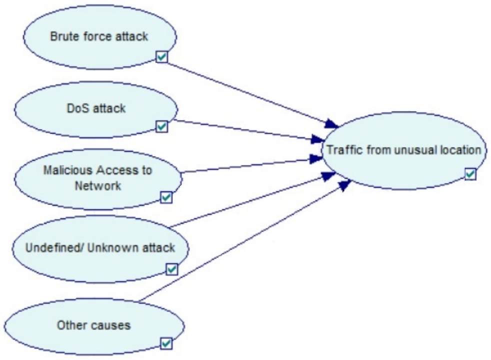
Brute force attack
DoS attack
Malicious Access to Network
Undefined/ Unknown attack
Other causes
Traffic from unusual location

Fig. 9. Addressing Unknown attacks

information about the physical system that is not directly observable and that an attacker might be trying to obfuscate.

Due to the breadth and complexity of our proposal, several challenges need to be addressed including proving the potential scalability of the approach. Selecting appropriate methods, rules, and training mechanisms to detect misbehaviour and avoid false positives is another challenge for which a good level of understanding about the expected normal behaviour of the system is required. It is proposed to formalise expert knowledge during the initial risk assessment where context and business rules should be captured. This should include specifications about normal behaviour of key processes of the system, a list of relevant variables to be monitored and availability of monitoring mechanisms. Possible correlations between physical variables should be identified and analysed, such as temperature and pressure. Additionally, machine learning techniques can be used to model patterns of behaviour that cannot be defined mathematically by experts. It must be noted that when data is gathered to model the normal behaviour of the system, this is done under the assumption that it has not yet been compromised.

Minimising cyber-security risks requires being up to date with new threats and vulnerabilities, as well as with new cyber-security tools and methods for prevention, detection, response, and recovery. As the methodology aims to be technology agnostic, it should be possible to introduce new monitoring and detection tools after the initial implementation.

In many industries cyber-security in ICS and IIoT systems has been neglected and there is lack of personnel equipped to deal with it. The continuous monitoring of cyber-risks should help increase cyber-security awareness and identify cybersecurity weaknesses. This also should include processes, people and training, since not all risks are technical since cybersecurity vulnerabilities can also be related to processes and procedures.

An more advanced approach will be to define the likelihood in terms of a distribution rather than a single point estimation, using Monte Carlo techniques.

This can give a more realistic view of the uncertainty level than single point probabilities. The specific technique and algorithm used to calculate the baseline risk scores should be defined during the context establishment. This methodology is general enough to work with different emerging approaches and ideas. Examples of this are deep learning and artificial intelligence, threat intelligence, and Security Orchestration, Automation and Response (SOAR).

8 Conclusions

Methods to maintain continuous cyber-risk awareness can support rational and well informed decisions as well as improve times and effectiveness of reactions in the event of an incident. The methodology proposed links detection mechanisms with risk assessment by using security metrics to calculate risk indicators in real time. Through this, we postulate the idea that integrating risk management with the security operations should enable a better prioritisation of security resources and more timely reactions. Current availability of off-the-shelf tools for real time data analytics also can constitute a powerful resource to make this approach feasible.

This paper regards continuous risk management as an extension of ISO/IEC 27005 in alignment with IEC 62443 and other frameworks rather than defining it as a totally different process. Because each ICS will have its own requirements, it is not possible to generate a one-size-fits-all solution. However, defining steps for each organisation to develop their own continuous risk assessment strategy can serve as a guideline that can be used broadly across different systems. Expert knowledge and contextual information are captured during initial assessment and transition phases. It is not possible to provide a general answer to some of our questions since our risk analysis approach needs to be tailored to the particularities of each system.

As in many cases the data gathered from security controls and IT elements of the system does not tell the whole story, and may itself be compromised by an attacker. Anomaly and misbehaviour detection techniques based on physical variables, such as temperature and power consumption, address this gap which can help to overcome different challenges in ICS cybersecurity, such as detection of physical attacks, detection of malicious insiders, and detection of unknown threats and zero day attacks.

Cybersecurity and OT operations, including safety management, have been developed on totally separate tracks. However, much of the knowledge in the field of control engineering and safety can be useful in the implementation of cybersecurity controls, as well as in the development of a future continuous risk assessment paradigm.

References

1. Adaros Boye, C., Kearney, P., Josephs, M.: Cyber-Risks in the Industrial Internet of Things (IIoT): Towards a Method for Continuous Assessment. In: International Conference on Information Security. pp. 502–519. Springer (2018)

2. Cyber-X Labs: 2019 Global ICS and IIoT Risk. A data-driven analysis of vulnerabilities in our industrial and critical infrastructure. Tech. rep., Cyber-X Labs (2018)
3. Dempsey, K., Chawla, N.S., Johnson, A., Johnston, R., Jones, A.C., Orebaugh, A., Scholl, M., Stine, K.: Information Security Continuous Monitoring (ISCM) for Federal Information Systems and Organizations: National Institute of Standards and Technology (NIST) Special Publication 800-137 (2012)
4. Desnitsky, V., Kotenko, I., Nogin, S.: Detection of anomalies in data for monitoring of security components in the internet of things. In: 2015 XVIII International Conference on Soft Computing and Measurements (SCM). pp. 189–192. IEEE (2015)
5. Ding, D., Han, Q.L., Xiang, Y., Ge, X., Zhang, X.M.: A survey on security control and attack detection for industrial cyber-physical systems. Neurocomputing 275, 1674–1683 (2018)
6. Gonzalez-Granadillo, G., Dubus, S., Motzek, A., Garcia-Alfaro, J., Alvarez, E., Merialdo, M., Papillon, S., Debar, H.: Dynamic risk management response system to handle cyber threats. Future Generation Computer Systems 83, 535–552 (2018)
7. Han, S., Xie, M., Chen, H.H., Ling, Y.: Intrusion detection in cyber-physical systems: Techniques and challenges. IEEE systems journal 8(4), 1052–1062 (2014)
8. Health and Safety Executive (HSE): Cyber Security for Industrial Automation and Control Systems (IACS) OG86. Tech. rep., UK Government (2018)
9. Huang, H., Xie, D.: Real-time network risk evaluation paradigm-inspired by immune. In: 2015 11th International Conference on Natural Computation (ICNC). pp. 786–790. IEEE (2015)
10. International Electrotechnical Commission: IEC 62443 2-1: Establishing an industrial automation and control system security program (2011)
11. ISO/IEC: ISO/IEC 27005:2011. Information security risk management (2011)
12. Kotenko, I., Saenko, I., Ageev, S.: Countermeasure security risks management in the internet of things based on fuzzy logic inference. In: 2015 IEEE Trustcom/BigDataSE/ISPA. vol. 1, pp. 654–659. IEEE (2015)
13. Kotenko, I.V., Levshun, D.S., Chechulin, A.A.: Event correlation in the integrated cyber-physical security system. In: 2016 XIX IEEE International Conference on Soft Computing and Measurements (SCM). pp. 484–486. IEEE (2016)
14. Liu, C., Zhang, Y., Zeng, J., Peng, L., Chen, R.: Research on Dynamical Security Risk Assessment for the Internet of Things inspired by immunology. In: 2012 8th International Conference on Natural Computation. pp. 874–878. IEEE (2012)
15. McLaughlin, S., Konstantinou, C., Wang, X., Davi, L., Sadeghi, A.R., Maniatakos, M., Karri, R.: The cybersecurity landscape in industrial control systems. Proceedings of the IEEE 104(5), 1039–1057 (2016)
16. Sicard, F., Zamai, E., Flaus, J.M.: Filters based Approach with Temporal and Combinational Constraints for Cybersecurity of Industrial Control Systems. IFAC-PapersOnLine 51(24), 96–103 (2018)
17. Yampolskiy, M., Horváth, P., Koutsoukos, X.D., Xue, Y., Sztipanovits, J.: A language for describing attacks on cyber-physical systems. International Journal of Critical Infrastructure Protection 8, 40–52 (2015)
18. Zhang, Q., Zhou, C., Tian, Y.C., Xiong, N., Qin, Y., Hu, B.: A fuzzy probability bayesian network approach for dynamic cybersecurity risk assessment in industrial control systems. IEEE Transactions on Industrial Informatics 14(6), 2497–2506 (2018)