

Dynamic Cyber Risk Assessment for Connected Medical Devices: the NEMECYS Approach

Gencer Erdogan^{1,*†}, Laura Carmichael^{3,†}, Steve Taylor^{3,†}, Simeon Tverdal^{4,†} and Andrea Neverdal Skytterholm^{2,†}

¹Sustainable Communication Technologies, SINTEF Digital, Oslo, Norway

²Software Engineering, Safety and Security, SINTEF Digital, Trondheim, Norway

³IT Innovation Centre, University of Southampton, Southampton, U.K.

Abstract

Connected Medical Devices (CMDs) face many critical cybersecurity challenges. Cybersecurity risk assessment is the industry de facto standard process to assess and mitigate potential cybersecurity risks. However, current cybersecurity risk assessments in the CMD domain are typically static, and many situations are highly dynamic involving changing circumstances of patient care priorities or new vulnerabilities detected at runtime. Thus, there is a clear need to support dynamic, runtime cybersecurity risk assessment where new events are reflected automatically in risk levels, and appropriate controls are recommended for unacceptable risks to return the residual risk to an acceptable level. In the EU project NEMECYS, we are developing an approach to dynamic cyber risk assessment for CMDs. The objective of this paper is to provide a high-level introduction to the NEMECYS project, and then explain in more detail our proposed dynamic cyber risk assessment approach for CMDs.

Keywords

cybersecurity, cyber risk, dynamic, risk assessment, connected medical devices

1. Introduction

Cybersecurity of medical devices connected to the internet (CMDs) faces three main critical challenges. These are: *A) complex and evolving standards*—The guidelines for medical device cybersecurity are complex, overly broad, and incomplete, struggling to keep pace with the rapidly evolving cyber threats and the need for integration in advanced, multi-institutional, and multi-device healthcare environments. *B) cost vs. care trade-off*—Implementing and maintaining cybersecurity measures is expensive, and an excessive focus on cybersecurity can detract from other aspects that are also critical to high-quality healthcare delivery, such as focus on operational efficiencies and providing person-centred care. *C) lifecycle and connectivity challenges*—Medical devices must be inherently secure, but connecting them introduces additional vulnerabilities, especially from other devices with unknown weaknesses, complicating the

Joint Proceedings of RCIS 2024 Workshops and Research Projects Track, May 14-17, 2024, Guimarães, Portugal

*Corresponding author.

†These authors contributed equally.

✉ gencer.erdogan@sintef.no (G. Erdogan); l.e.carmichael@soton.ac.uk (L. Carmichael); s.j.taylor@soton.ac.uk (S. Taylor); simeon.tverdal@sintef.no (S. Tverdal); andrea.skytterholm@sintef.no (A.N. Skytterholm)

🆔 0000-0001-9407-5748 (G. Erdogan); 0000-0001-9391-1310 (L. Carmichael); 0000-0002-9937-1762 (S. Taylor); 0000-0003-1660-4127 (S. Tverdal); 0000-0001-7507-6366 (A.N. Skytterholm)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

Table 1
NEMECYS project information

Project full name:	NEw MEDical CYbersecurity assessment and design Solutions
Acronym:	NEMECYS
Funding:	EU Horizon Europe
Start/end date:	1 January 2023 - 31 December 2025
URL:	https://nemecys.eu/

overall security landscape. In the EU project NEw MEDical CYbersecurity assessment and design Solutions (NEMECYS), we are addressing these three critical challenges. Table 1 shows the project information including funding, duration, and a URL for the project homepage.

The expected tangible outputs of the project are a set of toolboxes to assess the cybersecurity of CMDs in multi-stakeholder scenarios with different domains of control, as well as help relevant stakeholders to follow security by design best practices and meet regulations. Case studies in NEMECYS include home dialysis treatment, wearable devices for continuous monitoring of movement disorders, software as a medical device, and hospital based point-of-care testing. The toolboxes are grouped in the following three categories.

1. Tools that dynamically balance the analysis of cyber risks and benefits to ensure the security provided is appropriate, considering both the clinical advantages and ethical considerations, without adding unnecessary and expensive security measures.
2. Tools that help manufacturers follow the best practice and meet regulations to make sure their connected medical devices comply with cybersecurity requirements set by the EU Medical Devices Regulation (MDR) and the EU In Vitro Diagnostic Medical Devices Regulation (IVDR), while minimising the cost.
3. Tools that help medical device manufacturers, CMD system integrators and operators (e.g. hospitals) to design, build, and deploy security into medical devices and connected medical device scenarios, in a cost-effective way.

This paper focuses mainly on Point 1 above, where we present the work related to dynamic cyber risk assessment for CMDs. The relevance of this work to the *Research Challenges in Information Science* conference topics are information security, risk management, and e-health, as well as model-driven engineering, web-based applications and services, cyber-physical systems, and (medical) internet of things.

From a research challenge perspective, we are addressing two main limitations in the state of the art. First, current cybersecurity risk assessments are typically static, and many situations are highly dynamic involving changing circumstances of patient care priorities or new vulnerabilities detected at runtime. There is a clear need to support dynamic, runtime cybersecurity risk management where new events are reflected automatically in risk levels, alarms raised when risk rises unacceptably, and appropriate controls recommended to return the residual risk to an acceptable level [1]. Second, existing cybersecurity risk analysis schemes are mainly concerned with controlling risks and do not consider benefits and trade-offs between cybersecurity risks and clinical benefits. Existing risk/benefit methodologies for medical devices, e.g. ISO 14971 [2] are focused on clinical risk/benefit, and only superficially concerned with cybersecurity risks.

The proposed dynamic cyber risk assessment approach for connected medical devices builds on our existing work CORAS [3, 4] and Spyderisk [5]. Section 2 briefly describes CORAS and explains how it is extended to support dynamic risk assessment for CMDs. Section 3 briefly describes Spyderisk and how it is extended to support automated risk assessment for CMDs. Section 4 explains how CORAS and Spyderisk will collaborate to provide dynamic and automated cyber risk assessment for connected medical devices. Section 5 relates our work to existing work, while Section 6 concludes the paper.

2. CORAS Risk Indicators

CORAS is a method for conducting security risk assessment [3] in line with ISO 27005 [6]. CORAS provides a customized language for threat and risk modelling and comes with detailed guidelines explaining how the language should be used to capture and model relevant information during the various stages of security risk assessment. The CORAS method provides a web-based tool [7] designed to support documenting, maintaining, and reporting assessment results through risk modelling. The tool is also available as open source [8]. The reader is referred to the aforementioned sources for further details about CORAS.

CORAS is extended and specialized for the CMD domain with respect to method, language, and tool in the NEMECYS project. From a method perspective, we introduce a step to identify domain-specific cyber-risk indicators, as part of risk assessment. From a language perspective, we introduce the concept of indicators in the conceptual model of CORAS and extend the CORAS modelling language with a graphical construct for indicators to capture CMD specific risk indicators. From a tool perspective, we extend the CORAS web-based modelling tool to support dynamic risk assessment based on indicator values obtained through other tools, such as monitoring and testing tools.

Figure 1 shows an excerpt CORAS threat model, featuring indicators developed within the NEMECYS project. This model pertains to a use case involving medical patches for monitoring hydration and daily bodily changes in patients. Before we explain the threat scenario captured in Figure 1, we need to clarify central concepts in CORAS [3].

- A *threat* is a potential cause of an unwanted incident. We distinguish between deliberate human threat, accidental human threat, and non-human threat such as malware or failure.
- A *threat scenario* is a chain or series of events that is initiated by a threat and that may lead to an unwanted incident.
- A *vulnerability* is a weakness, flaw or deficiency that opens for, or may be exploited by, a threat to cause harm to or reduce the value of an asset.
- An *unwanted incident* is an event that harms or reduces the value of an asset.
- An *asset* is something to which a party assigns value and hence for which the party requires protection.
- A *party* is an organisation, company, person, group or other body on whose behalf the risk assessment is conducted.
- An *indicator* is a piece of information that is relevant for assessing the risk level. An indicator may be assigned to any risk-model element.

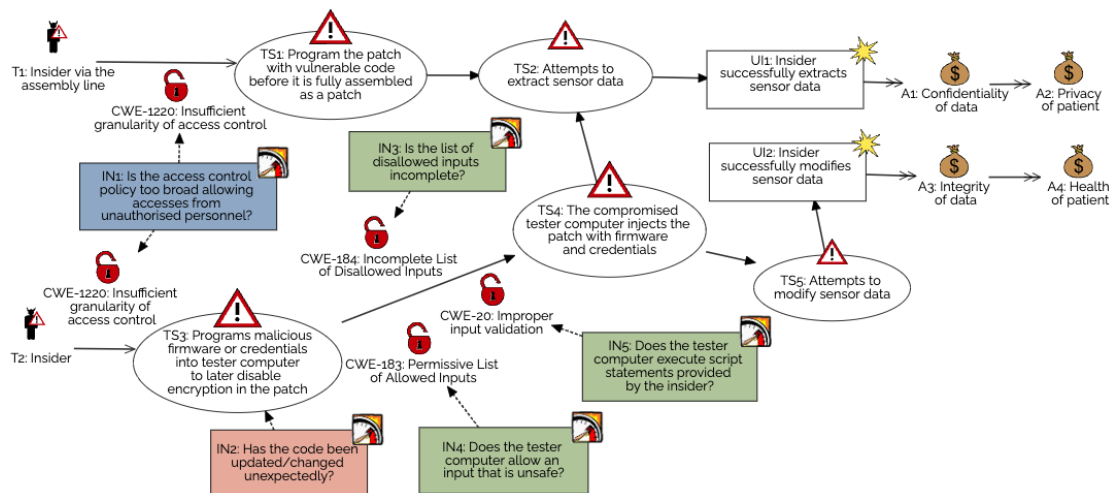


Figure 1: CORAS threat model with indicators.

Figure 1 illustrates cyber risks caused by the two deliberate human threats *T1: Insider via the assembly line* (third party) and *T2: Insider* in the organization. Threat *T1* exploits the weakness (that can become a vulnerability) *CWE-1220: Insufficient granularity of access control* [9] and initiates the threat scenario *TS1: Program the patch with vulnerable code before it is fully assembled as a patch*. One potential motivation of the attacker is to extract data from many patches, thus, the threat scenario *TS2: Attempts to extract sensor data*, which in turn may lead to the unwanted incident *UI1: Insider successfully extracts sensor data*. This unwanted incident may in turn cause harm to the assets *A1: Confidentiality of data* and thus harm asset *A2: Privacy of patient*.

Threat *T2* also exploits the vulnerability *CWE-1220* to initiate the threat scenario *TS3: Programs malicious firmware or credentials into tester computer to later disable encryption in the patch*. The tester computer is a computer developers use to test the patches. Next, the insider exploits three vulnerabilities that lead to the threat scenario *TS4: The compromised tester computer injects the patch with firmware and credentials*. The exploited vulnerabilities are *CWE-183: Permissive list of allowed inputs*, *CWE-184: Incomplete list of disallowed inputs*, and *CWE-20: Improper input validation*. Threat scenario *TS4* may then lead to the threat scenarios *TS2* or *TS5: Attempts to modify sensor data*, where the former may lead to unwanted incident *UI1*, and the latter may lead to unwanted incident *UI2*. As mentioned above, unwanted incident *UI1* may cause harm to the assets *A1* and *A2*, while unwanted incident *UI2* may cause harm to assets *A3: Integrity of data* and thus harm asset *A4: Health of patient*.

The indicators are defined as yes/no questions, for example, *IN1: Is the access control policy too broad allowing access from unauthorized personnel?* Indicator values may be obtained by different means. In some cases it is sufficient to base the indicator value on expert knowledge provided by someone who knows the target system well to decide the value, while in other situations it may be necessary to carry out security tests or monitor the application layer or the network layer in order to derive indicator values based on continuous monitoring.

Figure 1 illustrates five indicators: One indicator based on expert knowledge (*IN1*), one

indicator based on application layer monitoring (*IN2*), and three indicators based on security tests (*IN3*, *IN4*, and *IN5*). We also have an indicator category based on network layer monitoring and an indicator category based on unconventional data sources (not shown in the threat model in Figure 1). The latter category is to be further explored in the NEMECYS project and is intended to collect data from unconventional data sources, such as data from marketplaces, forums, blogs, and social media [10].

CORAS threat models with indicators, such as the example in Figure 1, facilitates dynamic risk assessment for CMDs in the following manner. First, we translate the model into an executable Bayesian network or multi-attribute decision tree script because CORAS diagrams are directed acyclic graphs and can be schematically translated into these formalisms [4]. Second, we use the indicator values (yes/no) as input parameters to the script and execute the script to obtain a propagated likelihood value, via the chain of threat scenarios, for an unwanted incident. Third, we rerun the script each time the indicator values are updated and thereby dynamically calculate new likelihood levels. We may, of course, input continuous indicator values to the scripts and define logic that returns more fine-grained likelihood values. An unwanted incident that harms one asset represents one risk. The consequence value for a risk is expected to be provided by the party. The risk level is calculated with respect to the likelihood value of an unwanted incident and its consequence value on an asset.

3. Spyderisk Automated Risk Assessment

Spyderisk is an automated risk assessment tool providing a systematic cause-and-effect modelling of threats [5]. Spyderisk follows the ISO 27005 standard for information security, cybersecurity and privacy protection [6]. The Spyderisk System Modeller provides a web-based graphical interface that can be used to create a model of assets and relations for the target system under test [11]. The knowledge base is then used to identify threats and their consequences, and threat likelihoods and risk levels are calculated. The end user can then decide what security controls to add (from those proposed) to the model to reduce the likelihood of threats, and re-calculate risks. After this, the end user “can choose to add or change the controls”, “re-design the system”, or “accept the system design” [11]. Spyderisk is also an open project [12]. The reader is referred to the aforementioned sources for further information about Spyderisk.

As part of the NEMECYS project, the Spyderisk knowledge base is being extended for the CMD domain. Such domain-specific knowledge is being acquired from multiple sources, including desk-based research, and project discussions driven by threat modelling in CORAS and system modelling in Spyderisk related to CMD use cases. Given that cybersecurity of CMD is a “patient safety concern” [13], one key area of focus for such extensions is understanding the various types of patient harm that may result from cybersecurity risks. For instance, cybersecurity breaches have the potential to affect the “safe operation” of CMDs that could cause injuries to patients and in some cases “threaten human life” [14]. As a further example, consider how loss of integrity and loss of availability of data and systems can “indirectly lead to a serious deterioration of health” through “indirect harm”, such as “a misdiagnosis”, “a delayed diagnosis”, “delayed treatment”, “inappropriate treatment”, “absence of treatment”, and “transfusion of inappropriate materials” [15]. Cybersecurity breaches may also compromise confidentiality of

data related to CMDs.

Of course, cyber risk assessment for CMDs also needs to take account of the “tensions” [16] between “patient safety”, “security design”, and “usability” of CMDs [17]. For example, in some cases, specified control strategies aiming to mitigate cybersecurity risks could compromise the usability of a CMD and increase patient safety risks [17], such as in “emergency” situations or as part of other types of “clinical workflow” [16].

4. Integrated Dynamic Risk Assessment

As illustrated in Figure 2, some of the results from CORAS will be used by Spyderisk, and some results from Spyderisk will be used by CORAS. One principal focus for collaboration is that of knowledge acquisition. For instance, the information being captured as part of the CORAS threat models for CMDs relates to key concepts in Spyderisk, such as threats, vulnerabilities, assets and consequences. Information generated through CORAS therefore can be used as an input to Spyderisk domain modelling activities, helping to inform extensions of the Spyderisk knowledge base for the CMD domain, and contributing to the wider Spyderisk Open Project. Another important area of collaboration is enriching threat models for CMDs. The Spyderisk System Modeller provides models of the target system and can generate numerous potential attack paths with respect to the target system model. This information can be used to identify new threat scenarios and indicators for CMDs in the CORAS threat models.

In addition to risk assessment tools, the NEMECYS project develops other tools, such as vulnerability scanners, firmware fuzzing, and procurement and compliance tools, which all will be used to provide values to the relevant indicators identified in the CORAS threat models. This information will help both CORAS and Spyderisk to calculate risk estimates. Together, CORAS and Spyderisk will produce a cyber risk assessment report for the stakeholders of CMDs.




Figure 2: CORAS and Spyderisk integrated dynamic risk assessment.

5. Related Work

To the best of our knowledge, there are only two closely related approaches that have a specific focus on dynamic risk assessment for connected medical devices. Leite et al. [18] introduce an approach based on dynamic risk assessment and control for medical cyber-physical systems (MCPS). Their approach gathers and uses safety data, patient characteristics, and other relevant information in runtime to dynamically and continuously optimize the system performance and

maintain safety. In their later work, Leite et al. expanded their method to better identify risk factors and build risk assessment models using Bayesian networks [19]. This model accounts for both changes in the patient's health and the system's ability to detect and respond to these changes.

Li et al. [20] propose a dynamic medical risk assessment model, for capturing the impacts of factors on the occurrence of adverse events. The authors use static fault trees to show risk scenarios and use Dynamic Bayesian network and Bayesian inference to analyze the operations of medical devices, in consideration of their failures, repairs, and human errors over time.

While these approaches assess safety risks caused by operational errors in the target system, our approach assesses cybersecurity risks that may compromise the security of medical devices, which in turn may harm patient treatment. Moreover, our approach assesses cybersecurity risks of the context in which connected medical devices are deployed and operate, such as the corruption of data monitored and gathered by medical devices used for treatment, the unavailability of critical services, and the compromise of private patient data.

6. Conclusion

There is a clear need to support dynamic and runtime cybersecurity risk management for connected medical devices (CMDs), and help existing risk/benefit methodologies for medical devices to explicitly address cybersecurity risks. In the EU project NEMECYS, we are addressing these challenges. We propose a dynamic cyber risk assessment approach for connected medical devices by combining risk indicator capabilities of CORAS [3, 4] and automatic risk assessment capabilities of Spyderisk [5]. The approach helps stakeholders including manufacturers, operators (hospitals) and integrators of CMDs to dynamically and automatically assess and mitigate cybersecurity risks while at the same time understanding the various types of patient harm that may result from cybersecurity risks. In the next phase of the NEMECYS project, our proposed approach will be validated in case studies provided by the stakeholder partners in the project.

Acknowledgments

This work has been conducted as part of the NEMECYS project, which is co-funded by the European Union (101094323), by UK Research and Innovation (10065802, 10050933 and 10061304), and by the Swiss State Secretariat for Education, Research and Innovation.

References

- [1] R. M. Czekster, P. Grace, C. Marcon, F. Hessel, S. C. Cazella, Challenges and Opportunities for Conducting Dynamic Risk Assessments in Medical IoT, *Applied Sciences* 13 (2023). doi:10.3390/app13137406.
- [2] ISO14971:2019, ISO 14971:2019 – Medical devices – Application of risk management to medical devices, 2019. URL: <https://www.iso.org/standard/72704.html>.
- [3] M. Lund, B. Solhaug, K. Stølen, *Model-Driven Risk Analysis: The CORAS Approach*, Springer Berlin, Heidelberg, 2011. doi:10.1007/978-3-642-12323-8.

- [4] G. Erdogan, A. Gonzalez, A. Refsdal, F. Seehusen, A method for developing algorithms for assessing cyber-risk cost, in: 2017 IEEE International Conference on Software Quality, Reliability and Security (QRS'17), IEEE, 2017, pp. 192–199. doi:10.1109/QRS.2017.29.
- [5] S. Phillips, S. Taylor, M. Boniface, M. Surridge, Automated knowledge-based cybersecurity risk assessment of cyber-physical systems, *Authorea Preprints* (2023) 1–23. doi:10.36227/techrxiv.24061590.v1.
- [6] ISO27005:2022, ISO/IEC 27005:2022 – Information security, cybersecurity and privacy protection – Guidance on managing information security risks, 2022. URL: <https://www.iso.org/standard/80585.html>.
- [7] CORAS, CORAS Web Application Tool, 2024. URL: <https://coras.tools/>.
- [8] CORAS, CORAS GitHub, 2024. URL: <https://github.com/stverdal/CORAS-The-Explorer>.
- [9] CWE, Common Weaknesses Enumeration (CWE). MITRE. CWE-1220: Insufficient Granularity of Access Control., 2024. URL: <https://cwe.mitre.org/data/definitions/1220.html>.
- [10] P. H. Meland, S. Tokas, G. Erdogan, K. Bernsmed, A. Omerovic, A systematic mapping study on cyber security indicator data, *Electronics* 10 (2021). doi:10.3390/electronics10091092.
- [11] Spyderisk, Spyderisk System Modeller Documentation, 2023. URL: <https://spyderisk.org/documentation/modeller/latest/>.
- [12] Spyderisk GitHub, Spyderisk GitHub repository, 2024. URL: <https://github.com/Spyderisk>.
- [13] N. O'Brien, S. Ghafur, M. Durkin, Cybersecurity in health is an urgent patient safety concern: We can learn from existing patient safety improvement strategies to address it, *Journal of Patient Safety and Risk Management* 26 (2021) 5–10. doi:10.1177/2516043520975926.
- [14] L. Coventry, D. Branley, Cybersecurity in healthcare: A narrative review of trends, threats and ways forward, *Maturitas* 113 (2018) 48–52. doi:10.1016/j.maturitas.2018.04.008.
- [15] MDCG, MDCG 2023-3, Questions and Answers on vigilance terms and concepts as outlined in the Regulation (EU) 2017/745 on medical devices, 2023. URL: https://health.ec.europa.eu/system/files/2023-02/mdcg_2023-3_en_0.pdf.
- [16] P. A. H. Williams, A. J. Woodward, Cybersecurity vulnerabilities in medical devices: a complex environment and multifaceted problem, *Medical Devices: Evidence and Research* 8 (2015) 305–316. doi:10.2147/MDER.S50048.
- [17] A. Ray, *Cybersecurity for Connected Medical Devices*, Elsevier Inc., 2021. doi:10.1016/C2018-0-03213-2.
- [18] F. L. Leite, D. Schneider, R. Adler, Dynamic risk management for cooperative autonomous medical cyber-physical systems, in: *Computer Safety, Reliability, and Security*, Springer International Publishing, Cham, 2018, pp. 126–138. doi:10.1007/978-3-319-99229-7_12.
- [19] F. L. Leite, D. Schneider, R. Adler, Dynamic risk assessment enabling automated interventions for medical cyber-physical systems, in: *Computer Safety, Reliability, and Security*, Springer International Publishing, Cham, 2019, pp. 216–231. doi:10.1007/978-3-030-26601-1_15.
- [20] M. Li, Z. Liu, X. Li, Y. Liu, Dynamic risk assessment in healthcare based on Bayesian approach, *Reliability Engineering & System Safety* 189 (2019) 327–334. doi:10.1016/j.res.2019.04.040.