Towards a Modular Human-Robot Safety Control System Using Petri Nets

Philipp Kranz[a], Fabian Schirmer[b], Marian Daun[c] and Tobias Kaupp[d]
Center for Robotics (CERI), Technical University of Applied Sciences Würzburg-Schweinfurt, Schweinfurt, Germany

Keywords: Human-Robot Collaboration, Safety, Modularity, Petri Nets.

Abstract: In industrial human-robot collaboration, where humans and robots operate in a shared workspace, the paramount concern is the safety of the human operator. The prevailing safety practices evaluate safety based on the overall assembly sequence, with the most critical task within the sequence being the limiting factor for all other tasks. This approach often results in significant limitations and the potential exclusion of collaborative interaction. However, the integration of human and robotic capabilities can facilitate the automation of processes, enhancing overall flexibility. The modular safety control system presented in this work employs a decentralized approach using Petri nets to evaluate the safety of humans and robots on a task-basis. This enables bridging the gap between the current, static regulatory framework and the necessary adaptivity of modern production systems.

1 INTRODUCTION

The concept of human-robot collaboration (HRC) represents a production paradigm that is designed to enhance flexibility within the context of Industry 4.0 production landscapes (Krüger et al., 2009). However, the increased flexibility that HRC affords also gives rise to an elevated risk of safety incidents. In the case of conventional robot applications, the workspaces between humans and robots are separated by safety fences, which effectively eliminate the majority of safety risks. In contrast, the safety assessment process assumes a pivotal role in HRC, as humans and robots have to solve assembly tasks together and interact in the same workspace simultaneously (Manjunath et al., 2024).

Safety assessment is a critical issue and an ongoing research challenge in HRC (Arents et al., 2021). The objective of safety assessment is to eliminate the majority of potential risks at design time. While this approach works for traditional robot and automation solutions, it significantly constrains the dynamic aspects of human-robot collaboration, impeding the acceptance of such applications. HRC should therefore use passive measures wherever possible, but these should be complemented by active measures backed up by an intelligent monitoring strategy.

This is particularly important as most accidents in HRC occur during process and maintenance operations (Lee et al., 2021). The causes of these accidents are often malfunctions, untrained operators, improper methods, or operator fatigue. However, what is not considered in many safety assessment approaches is a system-wide view resulting from the interaction between human and robot (Berx et al., 2022). Human-robot interaction is therefore both a driver of the flexibility of HRC, but also introduces new risks that are difficult to detect with current safety control systems. Robot safety standards provide limited support for the implementation of such flexible automation solutions and do not take human factors and the trade-off between various industrial requirements and safety into account (Hanna et al., 2022).

This work presents a model-based approach that serves as a foundation for the development of a modular safety control system for HRC assembly. The system uses a fully decentralized adaptive control scheme in which independent sub-systems (human and robot) coordinate with each other. The control scheme is based on the four MAPE activities: Monitor (M), Analyze (A), Plan (P), and Execute (E). (Weyns et al., 2013). To this end, this work constructs a state-space model in the form of Petri nets, which is

[a] https://orcid.org/0000-0002-1057-4273
[b] https://orcid.org/0000-0002-7032-8242
[c] https://orcid.org/0000-0002-9156-9731
[d] https://orcid.org/0000-0003-3017-5816

Kranz, P., Schirmer, F., Daun, M. and Kaupp, T.
Towards a Modular Human-Robot Safety Control System Using Petri Nets.
DOI: 10.5220/0013011900003822

In *Proceedings of the 21st International Conference on Informatics in Control, Automation and Robotics (ICINCO 2024) - Volume 1*, pages 384-391
ISBN: 978-989-758-717-7; ISSN: 2184-2809

capable of evaluating the status of the human and the robot on an assembly task level and calculating the current hazard potential in the form of a Safety Risk Indicator (SRI).

The paper is structured as follows: Section 2 reviews the literature on HRC risk management and model-based approaches, including the use of Petri nets. Section 3 explains the proposed modular safety control system and how to calculate a Safety Risk Indicator. Section 4 details the Petri Net models for human, robot, and safety control systems. Section 5 evaluates the approach in an industrial use case, and Section 6 concludes with a discussion, a summary, and future research directions.

2 RELATED WORK

This section summarizes the research underlying this work. It first reviews current safety regulations for HRC and highlights their limitations. Next, the benefits of a modular design for safety control systems are presented, followed by the use of Petri nets for modeling and their applications in HRC.

2.1 Safety Standards and Regulations for HRC

Under the European Union's legal framework, manufacturers, integrators, and end users must ensure robotic systems comply with the Machinery Directive's health and safety requirements. Compliance can be achieved by adhering to the directive or harmonized standards like ISO 10218 and ISO/TS 15066. While these standards are not legally binding, a mandatory risk assessment must identify hazards, estimate severity, analyze risks, monitor continuously, and implement risk mitigation measures (Hanna et al., 2022). Unlike the automotive sector, which follows ISO 26262 for safety analysis, no definitive procedures exist for robotics. Techniques such as Fault Tree Analysis (FTA) and Failure Mode and Effects Analysis (FMEA) are common but often oversimplify system behavior, leading to unreliable reliability assessments. There exists a notable gap between the existing static safety regulations and procedures and the necessary adaptability required by modern production systems (Hanna et al., 2022).

2.2 Modular Safety Control in HRC

The modular design of assembly systems allows for greater flexibility and versatility in their utilization (Siegert et al., 2021). However, the introduction of a new or modified application to such a system renders it a novel and potentially hazardous system from a safety perspective, necessitating a reassessment of the associated risks. A modular design of a safety control system would be advantageous for HRC, as it would allow for compensation in the event of deviations in the interaction between humans and robots, as well as enabling a more flexible response to dangerous situations. In their study, (Hillen et al., 2022) present an initial approach to addressing the discrepancy between the current static safety regulations and the requisite flexibility of modern production systems. The modular safety approach allows for comprehensive coverage of the design, development and runtime phases of risk assessment. In the runtime phase, a multitude of safety-related variables, including the present configuration and parameters of the systems, can be evaluated automatically.

2.3 Petri Nets for Safety Modeling

A Petri net (PN) is a formal model used to describe and analyze information flow and control systems, particularly for asynchronous and concurrent activities (Peterson, 1977). PNs have a long history in safety control, with Leveson and Stolzy using time PNs to model failures in safety-critical real-time software systems, focusing on fault tolerance and risk mitigation (Leveson and Stolzy, 1987). Recently, PNs have gained traction for safety, reliability, and risk assessment due to their formal, graphical, and mathematical nature, which enables realistic modeling of dynamic behaviors (Kabir and Papadopoulos, 2019). However, PNs can face state space explosion issues, limiting their use in large systems. Modularization has been proposed as a solution (Yevkin, 2011; Chiacchio et al., 2013). Given the complexity of human-robot collaboration (HRC) in assembly tasks, this work adopts a modular approach to enhance the scalability of the PNs used.

2.4 Petri Nets in HRC

Chao and Thomaz utilize time PNs to create a systematic model for reciprocal human-robot turn-taking (Chao and Thomaz, 2016). Their modular approach maps various resources and social actions, like speech, gaze, and gestures, enabling the robot to perform pick-and-place tasks while adapting to the human operator. This results in quicker reactions and improved task completion compared to simpler state-based methods. Similarly, Casalino et al. use time PNs to enhance scheduling in collaborative assembly

(Casalino et al., 2019). Their model allows for adaptive adjustments based on runtime data, responding to variations in human task durations and optimizing assembly activities. The principles of modularity and adaptivity can also inform the development of a safety control system for HRC. This work builds on the PNs in (Casalino et al., 2019), extending their approach to include a safety control loop.

3 MODULAR HRC SAFETY CONTROL SYSTEM

This section presents a modular safety control system for HRC. Firstly, the overall architecture based on MAPE principles is explained. The most common hazards for HRC assembly are then classified, and a SRI is calculated as a safety metric for task-based safety assessment.

3.1 Overall Architecture

The system architecture overview is depicted in Figure 1. In human-robot interaction during assembly, each step (1 to n) involves specific conditions related to components, actions, tools, and agents. These conditions, alongside environmental factors, can lead to safety hazards (accidents or losses) (Leveson, 2016). Hazards are categorized by severity, frequency, and likelihood of occurrence. To mitigate risks, we can use active controls that adjust system parameters or passive controls that raise awareness of potential hazards. Our safety control system employs a MAPE (Monitor, Analyze, Plan, and Execute) loop to continuously monitor conditions, analyze safety risks, and implement mitigation strategies, ensuring customized safety precautions for each assembly step (De Lemos et al., 2013).

To enhance modularity, we implemented a decentralized safety control system based on the MAPE loop, as shown in Figure 2. In this fully decentralized approach, each host operates its own MAPE loop (De Lemos et al., 2013). The local M, A, P, and E components coordinate with corresponding peer components from other hosts, facilitating flexible information exchange on system status and analysis results. The local P components manage adjustment actions that activate their local E components to execute actions on the sub-systems. In HRC assembly, these sub-systems correspond to the agents involved, which include both the human and the robot.

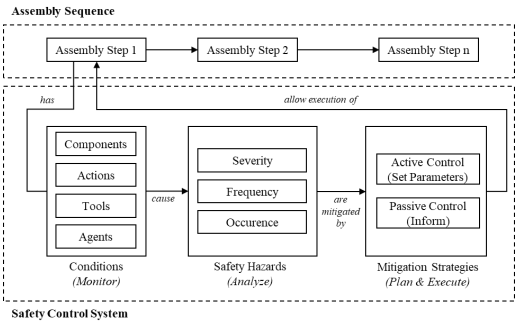

Figure 1: Overview of the proposed system with the two main components, the assembly sequence and the safety control system. Each assembly step is associated with specific conditions that give rise to different safety hazards. These hazards are analyzed to apply appropriate mitigation strategies, which legitimize execution of the assembly step.

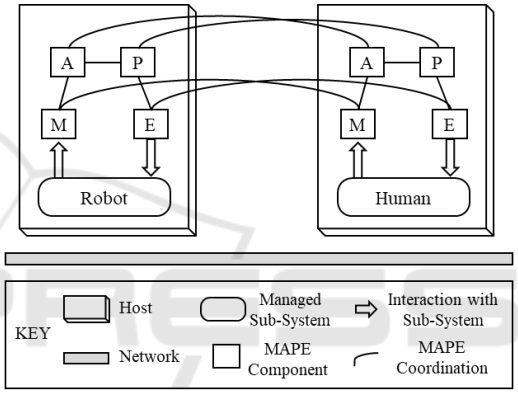

Figure 2: Decentralized pattern of a collaborative workspace, consisting of two sub-systems, the human and the robot, adapted from (De Lemos et al., 2013). The MAPE components of the human sub-system communicate with the MAPE components of the robot sub-system, making the overall system modular and adaptable to future changes.

3.2 Hazard Sources in HRC Assembly

In order to conduct a safety assessment for Human-Robot Collaboration (HRC), it is essential to identify potential risk factors. Most accidents in HRC occur during the execution of tasks (Lee et al., 2021). Each task involves components, actions, tools, and agents, all of which contribute to safety hazards. In adverse scenarios, a combination of these elements and environmental conditions may lead to safety risks (Leveson, 2016). Therefore, hazards are attributed to these conditions, forming the basis for calculating a risk indicator for specific tasks. This study focuses on task-specific safety hazards, assuming that the robot and workspace are inherently safe.

In assembly, components are the individual parts that are combined, while sub-assemblies consist of at

least two components. The final product represents the highest level, incorporating both sub-assemblies and components. Handling components, especially by robots, can create various hazards, such as sharp edges posing risks to workers and difficult-to-grip geometries that increase the likelihood of dropping heavy components. An ordinal rating system categorizes components as low, medium, or high risk based on their characteristics.

Actions are operations necessary for assembling components into sub-assemblies or final products, performed by either robots or humans. This study focuses on the following actions: 1) **Pick**: the robot or human grasps a component, 2) **Place**: the robot or human moves and releases a component at a specific point, 3) **Join**: the robot or human connects components or sub-assemblies with the required force, 4) **Hold**: the robot or human fixes a component in position, often with the robot acting as a third hand in collaborative tasks, and 5) **Screw**: the robot or human drives a screw using a screwdriver, with the appropriate tool.

Tools are objects required for specific actions, used by either humans (e.g., screwdrivers, hammers) or robots (as end-effectors or gripped tools). However, tools can also introduce safety risks, influenced by their geometry (e.g., piercing by a screwdriver bit) or function (e.g., crushing or entanglement).

Agents in HRC include humans, robots, or both, and the interaction modes depend on the task. For tasks performed exclusively by humans or robots, only their respective risk indicators are considered. In collaborative tasks, the interactions must be evaluated, as safety risks can escalate with certain interaction types. This work recognizes five interaction modalities as proposed by (Bauer et al., 2016): 1) **Human-Robot Cell**: Humans and robots work in separate areas, with the robot fenced off, preventing contact, 2) **Human-Robot Coexistence**: Humans and robots work in separate areas without a safety fence, with no direct contact, 3) **Human-Robot Synchronized**: Humans and robots share a workspace, working on the same component but at different times, 4) **Human-Robot Cooperation**: They work on different components in the same space simultaneously, with possible but unnecessary contact, and 5) **Human-Robot Collaboration**: Both work on the same component in the same space at the same time, with necessary contact.

3.3 Safety Metrics in HRC

In human-robot collaboration, safety hazards are usually assessed at design time, not at runtime, as it is proposed in this work. One of the most commonly used metrics for design time safety assessment, the risk priority number, is typically used to assess hazards, prioritize them and identify the most critical ones so that targeted mitigation strategies can be found (Afefy, 2015). The risk priority number is made up of the severity of a hazard, its occurrence probability and its detectability. As it is not possible to objectively measure detectability at runtime, the metric is adjusted accordingly and the frequency at which the actor is exposed to a hazard is used instead. The three components of our Safety Risk Indicator (SRI) are then defined as follows: 1) **Severity (S):** The potential impact or seriousness of a failure on the system is evaluated. Higher severity means more significant consequences, 2) **Frequency (F):** This measure the number of exposures to a situation where a potential hazard can occur, and 3) **Occurrence (O):** This assesses how likely it is for a particular failure to happen. A higher occurrence rate indicates that the failure is more common.

The defined conditions of an assembly task affect the SRI variables: Components and tools impact severity, the number of actions affects frequency, and the agents and their interaction modalities influence occurrence. All variables are scaled as low, medium, or high and weighted as follows: severity (S) [3], occurrence (O) [2], and frequency (F) [1]. The SRI is calculated as: $SRI = 3\ x\ S + 2\ x\ O + 1\ x\ F$

4 HRC SAFETY CONTROL MODEL

The modular HRC safety control system is modeled using PNs. The MAPE loops of the human and robot are integrated into a PN to model the control of a human, a robot and a human-robot task (Figure 3). The PN models in this paper are based on the work of (Casalino et al., 2019).

4.1 Model of a Robot Control System

To control a robot task, only the upper two components of the PN are used: the assembly task and the robot control loop. The robot status is represented by a blue token linked to the assigned task. Before starting the robot task, a condition check is initiated, providing all relevant information (components, tools, etc.). When T0 (Figure 3) is triggered, the safety control process begins.

The token first moves to ”Check Environment,” where current parameters of the robot and its environment are assessed. Using these values, an SRI value

is calculated to establish a task-specific safety metric, as detailed in Section 3.3. Robot parameters are then adjusted to meet the SRI safety requirements. If the requirements are satisfied, T8 fires and the robot proceeds to the assembly step and returns to "Wait/Idle" afterward. If not, T7 fires and the token moves to "Check Environment," and the safety control loop re-evaluates the conditions for potential task execution.

The robot token consists of both unchangeable and changeable values. Unchangeable values include the robot type and the end-effector or tool, which may vary depending on configuration or if a gripper change system is used. Changeable values involve current robot parameters (speed and force) and task-specific conditions, such as 1) components to be handled for the assembly step, 2) required actions, and 3) the interaction modality with the human operator. These variable values can be modified throughout the safety control cycle. Initially, task-specific conditions are recorded during the monitoring phase, followed by SRI calculation and adjustments to robot parameters. Thus, the value composition for the robot's place type includes:

$$Robot = Type \times CurrentSpeed \times CurrentForce \times Components \times Actions \times Tools \times Interaction \times SRI$$

4.2 Model of a Human Control System

The safety control system for the human operates similarly to that of the robot (Figure 3, lower part). After T0 is fired, the control loop starts, and the red human status token moves to "Check Environment." Depth cameras assess the human's availability and position at the collaborative workplace, while specific assembly step information is gathered. An SRI value for the human operator is calculated, mirroring the robot's process.

Due to the non-deterministic nature of human behavior, we can only passively control the human by informing them of the calculated SRI and associated risks. If T12 is satisfied, the task execution is confirmed, and the human receives assembly instructions. The token then moves to "Human Task" during execution and returns to "Wait/Idle" afterward. If T14 is fired, the token moves to "Check Environment", and the safety control loop is re-triggered.

The human token consists of both changeable and immutable values. Immutable values include the employee ID, job description, and relevant training, which help identify whether the individual is a trained assembly or maintenance operator or a non-specialist, the latter posing a higher safety risk. Like robots, human operators are assigned appropriate working conditions based on the task. However, the tool used is a variable value, while other conditions align with the robot's task-specific parameters. Since humans lack directly adjustable parameters, only passive mitigation strategies can be employed. Thus, the value composition for the human's place type includes:

$$Human = EmployeID \times Jobtitle \times Trainings \times Components \times Actions \times Tools \times Interaction \times SRI$$

4.3 Model of a Human-Robot Control System

In collaborative assembly situations, both, human and robot execute an assembly step together (Figure 3). For this purpose, both safety control systems work in parallel. When T0 fires, both systems are triggered and start their respective control loops, like they do when the safety for a pure robot or a pure human step is checked. As intended for the decentralized system and as described in Section 3.1, the corresponding MAPE components of the human and the robot control system are in close interaction to cover the overall safety hazards regarding the shared assembly step. Transition T1 to start the actual Human-Robot Task can only be fired, if both, the robot parameters are met and the residual risk for the human is tolerable.

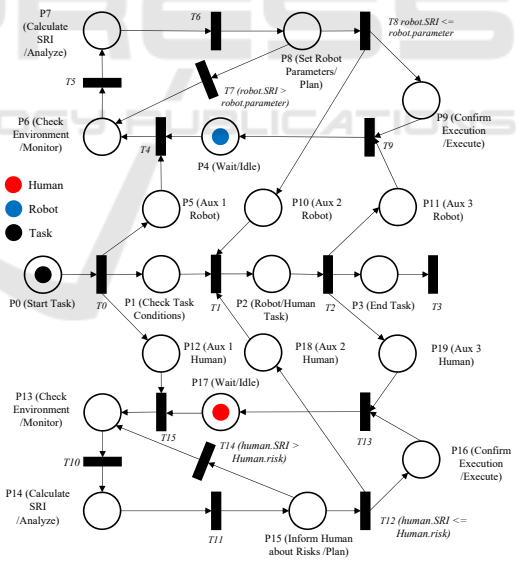

Figure 3: Petri net for the control structure of a human-robot task. Top: Robot safety control loop. Middle: HRC assembly step. Bottom: Human safety control loop.

In the human-robot tasks, the token values associated with both the robot and the human are integrated to compute the SRI. The aforementioned values comprise the employee ID, job title and trainings for the human, and the type, current speed and current force

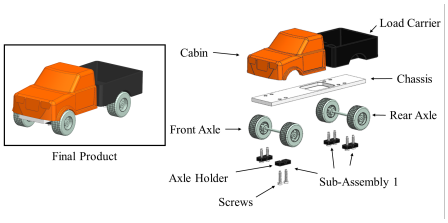

Figure 4: Exploded view of the product consisting of a base (cabin, load carrier, and chassis), a front axle, a rear axle, and four sub-assembly 1 (axle holder and two screws) (Schirmer et al., 2024).

for the robot. Task-specific values, including the components, actions and tools, are combined for both actors, thereby creating a new interaction modality. The calculated SRI is then applied to the human and the robot. In conclusion, the value composition for the place type of human-robot task is as follows:

$$Human\text{-}Robot = Type \times CurrentSpeed \times CurrentForce \times EmployeID \times Jobtitle \times Trainings \times Components \times Actions \times Tools \times Interaction \times SRI$$

5 APPLICATION

The safety control system, as previously created, is now applied to an industrial use case: the collaborative assembly of a toy pick-up truck. This paper demonstrates how the interaction between the MAPE components of the robot and human affect the SRI score, and how this can be manipulated by adjusting the human and robot parameters.

5.1 Experimental Setup

The collaborative assembly of a toy pick-up truck, shown in Figure 4, is employed as a case study (Schirmer et al., 2024). The assembly sequence comprises tasks that are performed exclusively by humans, exclusively by the robot, or by an interaction of the two actors.

In the initial assembly phase, the robot prepares the truck base, consisting of the load carrier, cabin, and chassis, by positioning them upside-down in an assembly bracket. The human operator prepares sub-assembly 1 by inserting two screws into each axle holder. They then collaborate to fix the front axle: the robot positions it, and the operator secures it with two sub-assembly 1 using an electric screwdriver. This process is repeated for the rear axle.

The SRI of the truck assembly tasks was initially determined using the evaluation criteria set out in Section 3 (Figure 6 top). The severity is determined by the severity of the components and the severity of the tools used in the task. The number of actions is translated into the frequency, and the type of interaction between the agents is translated into the occurrence.

5.2 Safety Control System of the Toy Pick-Up Truck

Figure 5 shows the safety control of the toy truck modeled by using the PN from Section 4. The three main components are color-coded: the assembly sequence in black, the robot's safety control in blue, and the human's in red. Once the assembly starts and T0 is activated, the tasks 'Prepare Base' and 'Prepare Sub-Assembly 1' are initiated simultaneously, engaging both agents' safety control systems. As they progress through their MAPE cycles, their monitor components update to recognize that both agents are working in the same collaborative space, effectively merging the tasks. This results in an increased SRI from 7 (Prepare Base) and 5 (Prepare Sub-Assembly 1) to 11 (Figure 6 bottom).

The 'Fix Front Axle' and 'Fix Rear Axle' tasks require interaction, initially planned as collaborative. After T0 fires, both safety control systems activate and go through their MAPE loops. Their identical initial SRI of 14 (Figure 6 top) is too high for execution. To lower the SRI, the interaction can be adjusted to synchronization (SRI 10) or cooperation (SRI 12) (Figure 6 bottom). The synchronization option is chosen for execution due to its lower SRI, though other factors like execution time or ergonomics could change this decision.

6 DISCUSSION AND CONCLUSION

The presented safety control system can conduct safety assessments for collaborative assembly sequences on a task basis, allowing for more dynamic evaluations compared to existing EU-standard systems, which require full assembly assessments. This task-based approach promotes modular HRC systems and enables the reuse of safety assessments across different sequences. However, the current system has a limitation in its discrete assessment of safety hazards, which cannot account for safety-critical factors occurring during task execution.

In summary, this paper introduces a modular safety control system based on the MAPE cycle, enabling decentralized safety evaluations for both robot and human agents. By using Petri nets to represent the system, we demonstrate the behavior of the MAPE

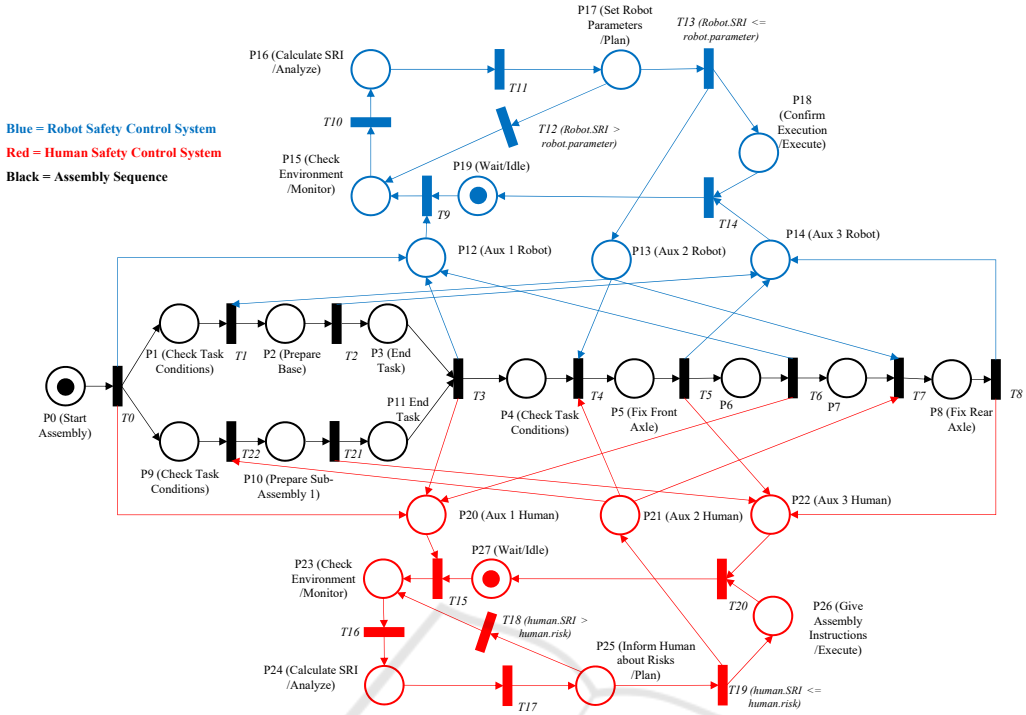

Figure 5: Petri net for the toy truck assembly. Different colors are used to distinguish between the three major parts of the PN: black for the assembly task, blue for the safety control system of the robot, and red for the safety control system of the human.

Initial Calculated Values

Task	Component	S (Comp)	Tools	S (Tool)	S	Actions	#	F	Agent	O	SRI
Prepare Base	Load Carrier	low	none		medium	Pick	3	low	Robot	none	7
	Cabin	low				Place	1				
	Chassis	medium				Join	2				
Prepare Sub-Assembly 1	Screws	low	none		low	Pick	2	medium	Human	none	5
	Axle Holder	low				Place	2				
						Join	8				
Fix Front Axle	Axle	medium	Screwdriver	low	medium	Pick	3	medium	Collaboration	high	14
	Sub-Assembly 1	low				Place	3				
						Hold	1				
						Screw	4				
Fix Rear Axle	Axle	medium	Screwdriver	low	medium	Pick	3	medium	Collaboration	high	14
	Sub-Assembly 1	low				Place	3				
						Hold	1				
						Screw	4				

Safety System Calculated Values

Task	Component	S (Comp)	Tools	S (Tool)	S	Actions	#	F	Agent	O	SRI
Prepare Base + Prepare Sub-Assembly 1	Load Carrier	low	none		medium	Pick	3	high	Coexistence	low	11
	Cabin	low				Place	1				
	Chassis	medium				Join	2				
	Screws	low				Pick	2				
	Axle Holder	low				Place	2				
						Join	8				
Fix Front Axle	Axle	medium	Screwdriver	low	medium	Pick	3	medium	Synchronization	low	10
	Sub-Assembly 1	low				Place	3		Cooperation	medium	12
						Screw	4				
Fix Rear Axle	Axle	medium	Screwdriver	low	medium	Pick	3	medium	Synchronization	low	10
	Sub-Assembly 1	low				Place	3		Cooperation	medium	12
						Screw	4				

Figure 6: Calculated SRI values for the truck use case. Top: Initially calculated values. Bottom: Calculated values from the safety control systems. Changes from the initial values are marked in blue.

components and their connection to the assembly sequence. Evaluations in an industrial context show the system's ability to discern various interaction modalities and adjust parameters to keep safe robot interactions within acceptable ranges.

In future work, in addition to a general risk assessment of the current task, we intend to integrate the detection and mitigation of more specific safety risks into our system. These will be continually reviewed to overcome our current limitation of discrete safety assessment, which can only identify risks before, but not during, task execution.

ACKNOWLEDGEMENTS

This research was partly funded by the Bayerische Forschungsstiftung under grant no. AZ-1512-21. We thank our industry partners Fresenius Medical Care, Wittenstein SE, Uhlmann und Zacher, DE software & control and Universal Robots.

REFERENCES

Afefy, I. H. (2015). Hazard analysis and risk assessments for industrial processes using fmea and bow-tie methodologies. *Industrial Engineering and Management Systems*, 14(4):379–391.

Arents, J., Abolins, V., Judvaitis, J., Vismanis, O., Oraby, A., and Ozols, K. (2021). Human–robot collaboration trends and safety aspects: A systematic review. *Journal of Sensor and Actuator Networks*, 10(3):48.

Bauer, W., Bender, M., Braun, M., Rally, P., and Scholtz, O. (2016). Lightweight robots in manual assembly–best to start simply. *Frauenhofer-Institut für Arbeitswirtschaft und Organisation IAO, Stuttgart*, 1.

Berx, N., Decré, W., Morag, I., Chemweno, P., and Pintelon, L. (2022). Identification and classification of risk factors for human-robot collaboration from a system-wide perspective. *Computers & Industrial Engineering*, 163:107827.

Casalino, A., Zanchettin, A. M., Piroddi, L., and Rocco, P. (2019). Optimal scheduling of human–robot collaborative assembly operations with time petri nets. *IEEE Transactions on Automation Science and Engineering*, 18(1):70–84.

Chao, C. and Thomaz, A. (2016). Timed petri nets for fluent turn-taking over multimodal interaction resources in human-robot collaboration. *The Int. Journal of Robotics Research*, 35(11):1330–1353.

Chiacchio, F., Cacioppo, M., D'Urso, D., Manno, G., Trapani, N., and Compagno, L. (2013). A weibull-based compositional approach for hierarchical dynamic fault trees. *Reliability Engineering & System Safety*, 109:45–52.

De Lemos, R., Giese, H., Müller, H. A., Shaw, M., Andersson, J., Litoiu, M., Schmerl, B., Tamura, G., Villegas, N. M., Vogel, T., et al. (2013). Software engineering for self-adaptive systems: A second research roadmap. In *Software Engineering for Self-Adaptive Systems II*, pages 1–32. Springer.

Hanna, A., Larsson, S., Götvall, P.-L., and Bengtsson, K. (2022). Deliberative safety for industrial intelligent human–robot collaboration: Regulatory challenges and solutions for taking the next step towards industry 4.0. *Robotics and Computer-Integrated Manufacturing*, 78:102386.

Hillen, D., Huck, T. P., Laxman, N., Ledermann, C., Reich, J., Schlosser, P., Schmidt, A., Schneider, D., and Uecker, D. (2022). Plug-and-produce... safely! end-to-end model-based safety assurance for reconfigurable industry 4.0. In *Int. Symposium on Model-Based Safety and Assessment*, pages 83–97. Springer.

Kabir, S. and Papadopoulos, Y. (2019). Applications of bayesian networks and petri nets in safety, reliability, and risk assessments: A review. *Safety science*, 115:154–175.

Krüger, J., Lien, T., and Verl, A. (2009). Cooperation of human and machines in assembly lines. *CIRP Annals*, 58(2):628–646.

Lee, K., Shin, J., and Lim, J.-Y. (2021). Critical hazard factors in the risk assessments of industrial robots: causal analysis and case studies. *Safety and health at work*, 12(4):496–504.

Leveson, N. and Stolzy, J. (1987). Safety analysis using petri nets. *IEEE Transactions on Software Engineering*, SE-13(3):386–397.

Leveson, N. G. (2016). *Engineering a safer world: Systems thinking applied to safety*. The MIT Press.

Manjunath, M., Raja, J. J., and Daun, M. (2024). Early model-based safety analysis for collaborative robotic systems. *IEEE Transactions on Automation Science and Engineering*.

Peterson, J. L. (1977). Petri nets. *ACM Comput. Surv.*, 9(3):223–252.

Schirmer, F., Kranz, P., Rose, C. G., Schmitt, J., and Kaupp, T. (2024). Holistic assembly planning framework for dynamic human-robot collaboration. In *Intelligent Autonomous Systems 18*, pages 215–227, Cham. Springer Nature Switzerland.

Siegert, J., Krispin, L., Ramez, A., El-Shamouty, M., Schlegel, T., Zarco, L., Roth, F., and Mannuß, O. (2021). Model-based approach for the automation and acceleration of the ce-conformity process for modular production systems: Future requirements and potentials. *ESSN: 2701-6277*.

Weyns, D., Schmerl, B., Grassi, V., Malek, S., Mirandola, R., Prehofer, C., Wuttke, J., Andersson, J., Giese, H., and Göschka, K. M. (2013). On patterns for decentralized control in self-adaptive systems. In *Software Engineering for Self-Adaptive Systems II*, pages 76–107. Springer.

Yevkin, O. (2011). An improved modular approach for dynamic fault tree analysis. In *Annual Reliability and Maintainability Symposium*, pages 1–5.