SafeNet of Unsafe Devices
Extending the Robot Safety in Collaborative Workspaces

Federico Vicentini, Nicola Pedrocchi and Lorenzo Molinari Tosatti
Institute of Industrial Technologies and Automation (ITIA), National Research Council (CNR), via Bassini 15 Milan, Italy

Keywords: Safety, Industrial Robots, Network Controlled System, System Engineering, Functional Safety, ISO 13489, ISO 10218.

Abstract: Collaborative workspaces represent the benchmark scenario of contemporary and future industrial robotics, where hybrid production systems and multimodal interactions among human operators and robots in cooperative tasks can foster the flexibility of robotic systems. Physical interactions together with dynamic workspace-sharing represent some reference applications in ISO 10218-2, where restrictive conditions for safety are posed at system level, eventually limiting the robot execution speed. With the aim of extending the use of industrial robots in shared environments and allowing the use of generically unsafe sensory and computational components for advanced applications, a methodology called SafeNet is presented. It considers the system as a device at large and applies the concept of functional safety (ISO 13489-1) with a set of architectural procedures and implementations. The safety aspects of structure, reliability and monitoring are addressed by a redundant system of computational nodes distributed over a network. SafeNet systems can be upgraded to candidate for safe Performance Levels.

1 INTRODUCTION

Collaborative workspaces (Fig. 1) are widely reckoned by both the industrial and the academic communities as one of the elective scenarios for the present-day industrial robotics. Safety, specifically, is one of the predominant functional aspects at both machine and system levels. Under this perspective, robots, as stand-alone machines, benefit from several technologies designed for a transparent and safe physical Human-Robot Interaction (pHRI) (De Santis et al., 2008; Alami et al., 2006). Such technologies support entirely new benchmarks for service robotics, as well as for many industrial applications. Examples include compliant actuation systems (Grebenstein et al., 2012; Bicchi et al., 2008; Zinn et al., 2004; Pratt and Williamson, 1995) and lightweight platforms (Kock et al., 2011; Albu-Schäffer et al., 2007a) that feature compliant behavior attained by mechanics and control (Albu-Schäffer et al., 2007b) and that display limited energy transfer in impacts (Haddadin et al., 2008; Haddadin et al., 2009). Together with internal or add-on sensing, *e.g.* tactile skins (Vogel et al., 2011), such compliant platforms represent a class of elective devices for shared environments. In such a context, safety issues are predominantly treated in terms

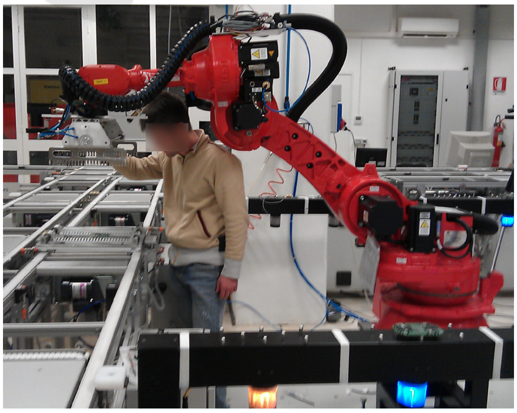

Figure 1: Paradigmatic scenario of a collaborative workspace for an industrial robot cell without fences.

of hazardous impacts or energy transfer, whose magnitudes and way of assessing are under discussion in ISO/TS 15066 (ISO, 2011c).

Physical HRI, however, is only a form of collaboration in shared workspaces. Paradigmatic workflows may, in fact, involve a mix of hand-guided procedures and contactless co-presence in the same safeguarded space. Such scenarios are particularly relevant for industrial robots, which as stand-alone devices have to comply with eventual stops or speed

Vicentini F., Pedrocchi N. and Molinari Tosatti L..
SafeNet of Unsafe Devices - Extending the Robot Safety in Collaborative Workspaces.
DOI: 10.5220/0004479502760283
In *Proceedings of the 10th International Conference on Informatics in Control, Automation and Robotics* (ICINCO-2013), pages 276-283
ISBN: 978-989-8565-71-6

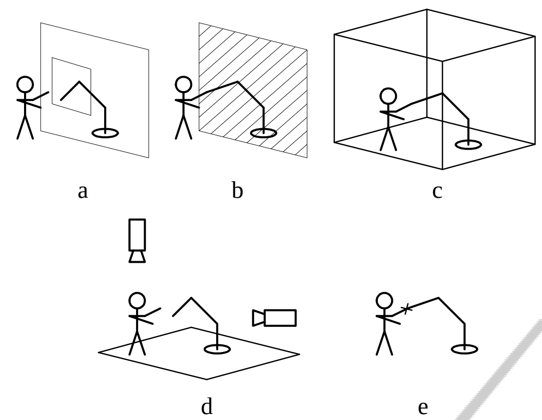

Figure 2: Classes of cooperative tasks as in ISO 10218-2:2011 Annex E (graphics copyright by ISO International Organisation for Standardisation): a) hand-over window, b) interface window, c) inspection, d) collaborative workspace, e) hand-guided robot.

limitations in such safeguarded spaces, as required by ISO 10218-1 (ISO, 2011a). Many optional safety packages in commercial controllers (KUKA Roboter, 2012; ABB Robotics, 2008) are, in fact, available for joints position safe checks at runtime, providing the basic information for a safe assessment of the robot configuration within a safeguarded space. This, in turn, represents the necessary condition for integrating safe application modes (Fig. 2) in dynamically shared environments as in ISO 10218-2 (ISO, 2011b).

Nevertheless, robots and robot systems, compulsorily featuring speed limitations in the safeguarded workspace (normative *status quo*), may conversely need higher task speeds and, additionally, may require the use of pervasive sensing and context awareness. This monitoring capability almost always needs distributed sensor equipments dedicated to the detection of the environment and users. Sensor processing and interpretation could, in turn, require significant computational power, so that collaborative workspaces would be, in a general sense, portrayed as distributed robotic systems[1] (Fig. 3). The resulting paradigmatic architecture is therefore a network of general-purpose devices, notably including unsafe nodes and where safe/unsafe controllers are parts of a wider set of data producers/consumers.

In this paper we discuss a methodology developed in fact to fullfil ISO 10218-2 safety requirements for a robotic system with unsafe nodes (robots included) through a set of architectural and procedural actions over the system. The two key concepts are that (i)

[1]specifically Network Controlled Systems (Gupta and Chow, 2010; Hespanha et al., 2007), when control actions proper are distributed among several nodes .

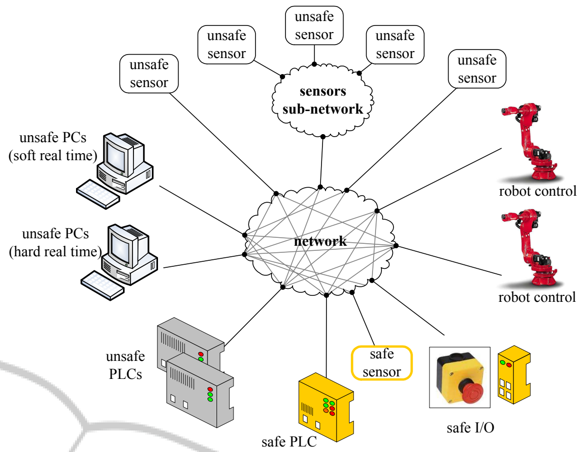

Figure 3: A robotic distributed system including both safe and unsafe nodes/devices.

the system at hand can be seen as a single (complex) device that (ii) has to display *functional safety* as a whole. Functional safety is the "part of safety relating to the Equipment Under Control (EUC) and the EUC control system that depends on the correct functioning of the Electric/Electronic/Programmable Electronic (E/E/PE) safety-related systems, other technology safety-related systems and external risk reduction facilities" (IEC, 2010). Since the system at hand can be considered a single EUC when used for interacting with and monitoring the collaborative workspace, it is required to be validated with respect to functional safety criteria as in ISO 13489-1 (ISO, 2006). Equivalently, components in a system are not required to be safe *per se* but, rather, the system functional safety depends on *to which extent the residual probabilities of failures in exchanged data can be limited.*

The core methodology here discussed aims at extending the functional safety of data flows before any usage of such data in the network. Applications eventually using such safe data in safety functions do not contribute to the preliminary safe assessment of data. Rather, being the *way* the nodes are safely checked relevant for the overall risk assessment, such network can freely integrate both safe and un-safe sensors/devices. This would make the exclusive use of individually safety-rated devices non-necessary for a safe system integration. A relative freedom in the integration of subsystems, remarkably computational nodes in PC-based robotic applications, is considered to be beneficial for the evolution of industrial robotic cells towards fully-collaborative fully-open environments. Such freedom of components choice, sometimes actually being the only choice because of required specific technologies that are not supported by safety-rated devices, reflects the concept of extension

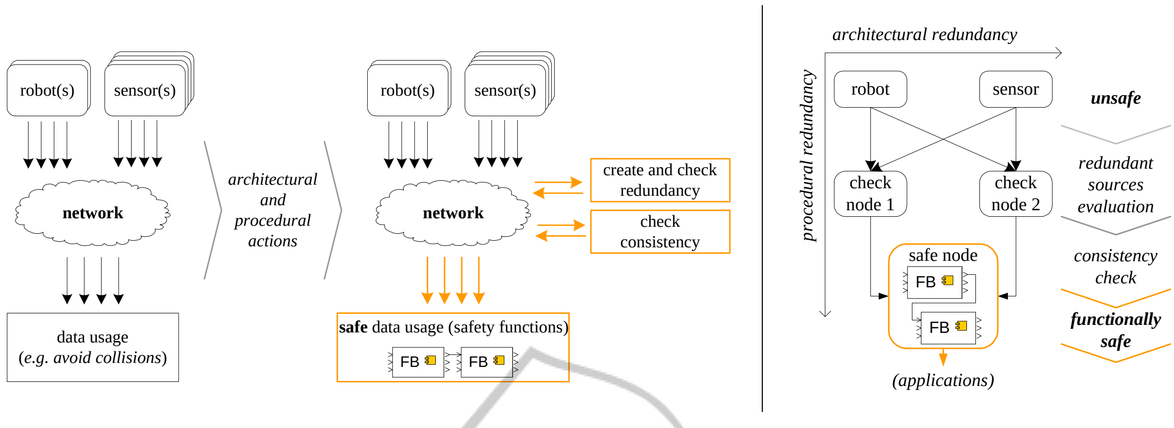

(a) (b)

Figure 4: Methodological framework for a SafeNet: in (a) a common network of mixed safe/unsafe devices turns into a safe network through a set of architectural and procedural actions, *i.e.* involving redundancy and consistency check over the network, making the following usage of data *functionally* safe. Such actions are detailed in (b) where different sources (robots and sensors) are procedurally evaluated in two (architecturally redundant) check nodes. Upon data procedural evaluation, data consistency between the two nodes is verified by a safe node/layer through safety functions (FB), then used.

of safety features in networks rather than restricting the usage of few rated protocols, as reckoned in the progressive introduction/standardization of safe protocols into the main families of automation fieldbuses (Moyne and Tilbury, 2007; Decotignie, 2005; Felser, 2005).

On top of the methodology here discussed, the system integration has nevertheless to provide a general assessment, evaluation and mitigation of risks according to ISO 12100 (ISO, 2010) wrapping guidelines, which are out of scope in this work. In the next section, instead, the procedural and architectural aspects of the extension of safety to data flows are discussed.

2 SAFENET FRAMEWORK

Functional safety is a key element of system design based on (i) well-tried components and methods and on (ii) the application of the principles of redundancy, diversity, monitoring. Functional safety is expressed as a ISO 13489-1:2006 Performance Level (PL), or equivalent IEC 61508 SIL level, which encapsulate the rate of reliability, failure-detectability and readyness to recovery of a component/system. Specifically for a robotic system, the required safety-rate is (ISO, 2011b):

$$\begin{cases} PL_r = d \quad \text{i.e.} \;\; PFH_d \in [10^{-7}, 10^{6}) \\ Cat.3 \text{ Designated Architecture} \end{cases}$$

where PL_r is the *required* performance level, PFH_d is the Probability of dangerous Failures per Hour and $Cat.3$ is one of the two safest ISO 13489-1:2006 categories of Safety Related Parts of a Control System (SRP/CS) using double channels (see details in section 2.2). Such functional safety rate is the aim of the actions (Fig. 4-a) that transform a network of unsafe devices in a *SafeNet* and that mainly involve a double set of data validation and cross-checking. The purpose is to reduce the probability of failing in detecting occasional inconsistencies in data processed by different nodes. Distributed systems, in particular, are likely to include sensors used for environmental monitoring that are eventually available for tracking the robot(s) motion inside a shared workspace as well. The possible configurations of sensors and controllers are very diverse, only optionally including native safety packages in robot control. External motion tracking information are, instead, rarely matching the safety-rate standards.

The above listed principles of *redundancy*, *diversity* and *monitoring* are therefore applied to the verification of such tracking information by a double independent elaboration of a single target information, obtained along both a procedural and an architectural dimension (Fig. 4-b). The procedural redundancy corresponds to the plain use of data from both the tracked (unsafe) robot and the tracking unsafe sensors, verifying that values match within given tolerances. The architectural redundancy, complementarily, is obtained distributing robots and sensors data in doubled flows for independent procedural evaluation. Then, both comparative units (check nodes) are verified for consistency by a final safe unit/layer, *i.e.* data are fed to safety functions coded according to IEC

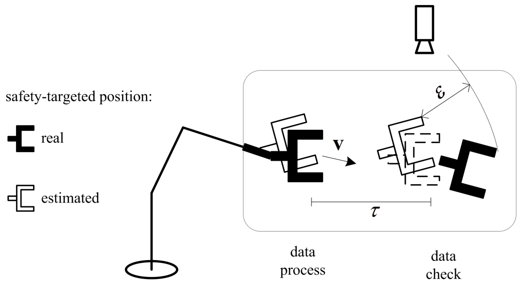

Figure 5: Threshold inaccuracy ε in motion tracking due to speed $\mathbf{v}$ and time latency τ.

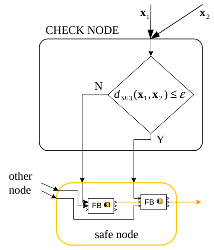

Figure 6: In each ckeck node evaluation, data differences are evaluated w.r.t. the task-dependent threshold ε. Evaluated data are then bound in safety functions (FB).

61131-3/61508 in natively safety-rated logic devices or mapped in safe I/Os distributed over safe protocols. Such final step is compactly represented in Fig. 4-b by a safe node that acts as the safety gate between the safety functions domain and all the general purpose CPUs or unsafe sensors.

As a result, such architecture is equivalent to a SRP/CS distributed in two components and a safe node, suitable to fulfill the preliminary conditions for a *PL d* implementation, *i.e.* the dual structure and the availability of monitoring coverage.

Proceeding with the tracking configuration, the procedural and architectural aspects of the SRP/CS are discussed in the following subsections: robots and sensors data-check (procedure, Section 2.1) mainly involve kinematic and accuracy considerations, while the data-flows dispatching (architecture, Section 2.2) are considered according to ISO 13849-1:2006 guidelines.

2.1 Procedural Data Check

Considering a basic configuration with a robot moving along a joint trajectory $[\mathbf{q}, \dot{\mathbf{q}}]$, with speed $\mathbf{v}$ and tracked by a set of sensors (Fig. 5), each unit verifies that motion data from robots and sensors correspond, *i.e.* whose difference remain in a given safe interval. The motion data difference d_{SE3} can be evaluated in each check node (Fig. 6) according to any Lie-algebra consistent metrics[2]. Intervals and/or allowed regions for motion data verifications depend on the system and the application, *e.g.* largely on speed $\mathbf{v}$ and on the networking that may affect the data exchange. Measurement inaccuracies depend, in fact, on several factors, either spatial or temporal:

- errors in calibration that usually depend on the position inside the workspace due to the anisotropy of the calibration procedures

$$\varepsilon_{calib} = \varepsilon_{calib}(\mathbf{q}) = \varepsilon_{calib|sens} + \varepsilon_{calib|frames}$$

where $\varepsilon_{calib|sens}$ is the intrinsic precision of the sensor and $\varepsilon_{calib|frames}$ is the accuracy of the hand-eye calibration (Tsai, 1987). In case of calibration procedures based on same sensors used during the tracking, the inaccuracy propagates from the sensor precision to the errors in reference frame alignments;

- tracking errors of the manipulators,

$$\varepsilon_{dyn} = \varepsilon_{dyn}(\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}}, \mathbf{f}_{ext}) \simeq \varepsilon_{dyn}(\dot{\mathbf{q}}, \mathbf{f}_{ext}) \ll \varepsilon_{calib}$$

usually negligible in presence of accurate modeling of the residual flexibilities at both link and joint levels and proper compensatory control strategies;

- temporal misalignment (τ in Fig. 5) between the sampled poses and the real pose

$$\varepsilon_{lat} = \varepsilon_{lat}(\Delta T_{sens})$$

where (see Fig. 7)

$$\Delta T_{sens} = T_{offset} + \Delta T_{proc} + \Delta T_{tx} + \sum_k jit_k \geq 0$$

is the cumulative time delay due to channels asynchronicity (T_{offset}), sensor information processing time (T_{proc}), protocol-dependent transfer latency (T_{tx}) and related jitters, that ends up into a blind time-of-flight for the robot.

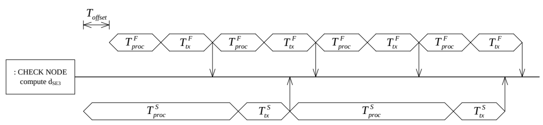

Figure 7: Time delay over network.

The overall inaccuracy $\varepsilon = \sum \varepsilon_k \geq \varepsilon_{min} > 0$, $\forall k$ sources listed above, introduces a non-null risk of erroneous tracking of the robot (risk factor, $RF \geq RF_0$) that increases more than linearly with ε (Fig. 8).

[2] rototranslations as in (Strobl and Hirzinger, 2006), rotations as in (Moakher, 2002) or plain Euclidean distances in $\mathbb{R}^3$

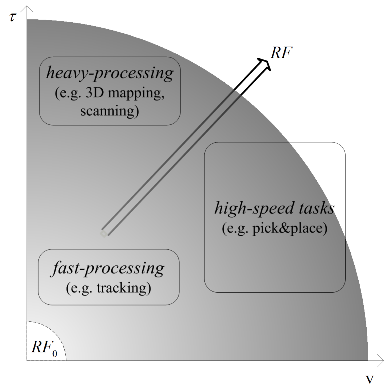

Figure 8: Risk factor map: types of robot applications and sensing routines in a time vs. speed domain. Gray gradient represents the superlinear increase of risk factor *RF*.

While spatial and control inaccuracies may be considered negligible in most of practical cases, with $\varepsilon_{calib} + \varepsilon_{dyn} \leq 5mm$, the latency-dominated inaccuracy ε_{lat} plays a significant role in building the overall blind time-of-flight τ along which the robot moves without any chance of detection (Fig. 5). The latency component, in fact, assumes the dominant role in evaluating data from sensors. Considering in fact a group S of slower devices w.r.t. the group F of faster devices - *e.g.* robots - with sampling frequency $[1-5)ms \ni T^F_{samp} < T^S_{samp} \in [5-20]ms$, and timing reasonably being

$$T_{offset} \leq T^F_{proc}$$
$$T^S_{tx} \simeq T^F_{tx} \in [1-5]ms$$
$$T^S_{proc} = T^S_{samp} + T^S_{comput} \gg T^F_{proc},$$

the sensor processing happens to be the prominent contributor to the overall time misalignment in data checking, *i.e.*

$$\tau \simeq T^S_{proc}$$
$$\varepsilon \simeq \varepsilon_{lat}.$$

As a result, demanding applications, *e.g.* fast robot motion - which is currently not allowed in standards (ISO, 2011b) - and time-expensive environmental monitoring, happen to require larger tolerances or larger uncertainty regions (*e.g.* larger risk factor Fig. 8) where each check node enters a safe state.

From a SafeNet procedural stand point, the monitoring principle would benefit from a reduction of such *RF* or, correspondingly, an improvement in quality of the sensor channels. The monitoring of channels, and their quality at large, tend to limit the number of failures (*i.e.* $d_{SE3} > \varepsilon$) per time unit, significantly contributing to the improvement of the system reliability, which in turn is one of the steering parameters in ISO 13489-1:2006 PL assessment.

2.2 Architectural Designation and Performance Level

The set of architectural actions (Fig. 4), aimed at differentiating and doubling the data flow evaluation, provide the necessary *structure* of a *PL d* class of functional safety. Architecturally this is equivalent to distributing a SRP/CS over 3 components, being able to cross-monitor the double data channels. In a ISO 13489-1:2006 *Cat*.3 architecture with dual channel I-L-O (input-logic-output) modules (Fig. 9), all monitoring functions are, in fact, performed by the safety functions in the safe node.

From an implementation point of view, this can be achieved by embodying the check nodes and the safe node in 3 separate PLCs (Fig. 10) networked through any suitable protocol (chiefly Ethernet-based) to the system and mutually through a *safe* protocol. On

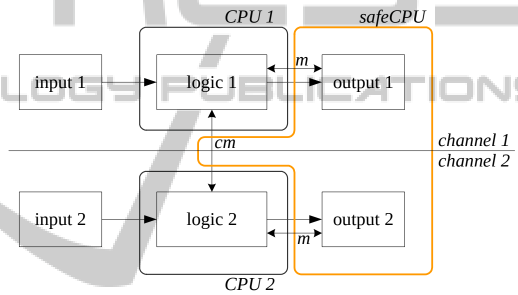

Figure 9: ISO 13489-1:2006 model for *Cat*.3 designated architecture with deployed SRP/CS components (CPU1, CPU2 and safeCPU) outlining their logical domain. *m* are monitored safe state tasks execution. *cm* is th ecross-monitoring of both channels.

top of dual channel architectures (mandatory for *PL d* and *PL e* rating), the actual Performance Level is determined by the degree of *reliability*, in the form of mean time to dangerous failures ($MTTF_d$), together with the *monitoring capability*, in the form of Diagnostic Coverage (*DC*) of a system. The $MTTF_d$ at system level is not discussed in detail in this work, while some considerations are due to highlight that the *DC* level of a *Cat*.3 architecture cannot be null, *i.e.* $DC < 60\%$. A minimum requirement *DC* can be achieved, for instance for inputs in SRP/CS, through “monitoring some characteristics of the sensor (response time, [...])” (ISO, 2006), with the possibility of improving the *DC* through cyclic and/or parallel methods for monitoring the sensor lines (*DC* up to 99%). *DC* rate is normally evaluated for all I-L-O modules. As an implementation example, a low range $60\% \leq DC < 90\%$ can be supported by watchdog components (Fig. 10) in each channel in charge of monitoring the availability and timing of (i) the

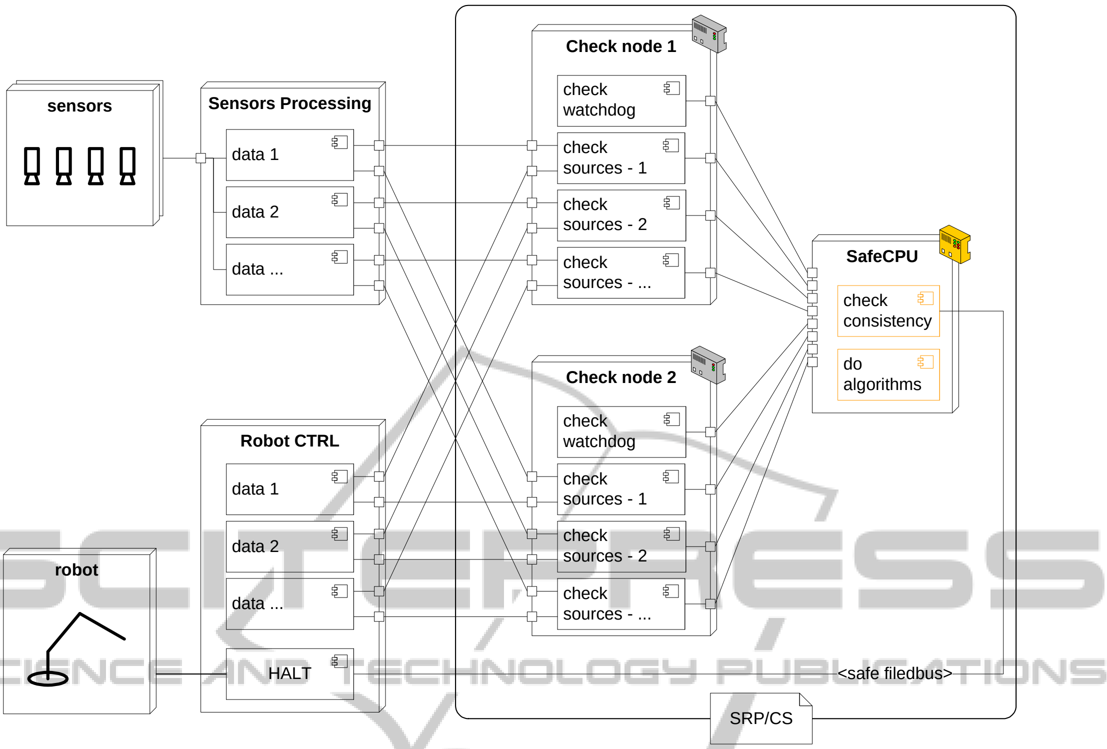

Figure 10: UML deployment diagram of a system made of a robot and a sensor set connected to the SRP/CS made of 2 standard PLCs (check node 1 and 2) and one safePLC (safe CPU). "double to double" connections from/to the SRP/CS are visible for all data feeds in each data source (data1, data 2, data ...). Watchdogs are present for DC purpose and *check consistency* component in SafeCPU node is in charge of handling the safe state. *do algorithms* component in SafeCPU node represents the data usage in a functionally safe mode, *i.e.* through safety functions.

data transfer protocol (*e.g.* port access) and (ii) the data sampling/processing routines. In particular, the L modules in both channels are directly connected to the safePLC (Fig. 11) through safe protocols, ensuring a supervised output for each channel (*m* in Fig. 9). The same apply at inter-logic level (*cm* in Fig. 9).

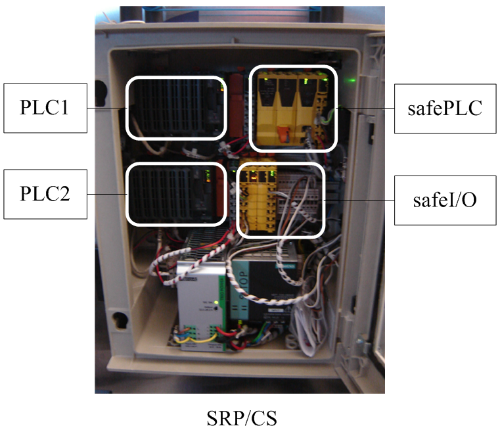

Figure 11: Deployed SRP/CS featuring actual hardware components (PLC1, PLC2 and safePLC).

Finally, ISO 13489-1:2006 requirements for functional safety include also the use of (application-dependent) well-tried procedures, components and methods in system development in form of a review of measures for avoiding the common causes of failures (CCF) that have to gain a minimum score of 65 according to quantification in Tab. 1.

Table 1: Measures against CCF (common causes of failures) scores.

measure	max score
Separation between the safety circuits	15
Diversity in design and technologies	20
Draft / Application / Experience in applying well-tried procedures	20
Assessment / Analysis	5
Competence / Training of developers	5
Environmental influences EMC and others	35

3 CONCLUSIONS

A methodology has been outlined discussing procedural and architectural actions aiming at qualifying a robotic system with a functional safety rate equal (at least) to *PL d* , as requested by ISO 10218-2:2011, in the case of entire/partial presence of unsafe nodes (Fig. 12).

The core concept introduced in such a methodology (SafeNet of unsafe devices) considers the system as a device at large, which has to display functional safety in its parts and nodes. Required level of functional safety has been reviewed to be formulated on the basis of system-level reliability and monitoring ($MTTF_d$ and *DC*), to require well-tried and consistent practices (CCF counter-measures), and, most importantly, to stand on a class of dual channel monitored architectures where the SRP/CS is able to consistently check the availability and validity of data feed and component behaviors. Such structural feature is obtained through the main characteristic of the SafeNet that involves the creation of procedural and architectural redundancies over the network, variously interconnecting robots and sensors. In this way, general systems of designated architectures $Cat.B/1/2$ can be upgraded to $Cat.3$ and can provide necessary conditions for *PL d* rate achievement (Fig. 13). The safety rate upgrade is mainly in charge of a SRP/CS distributed on 3 components that provide the *structure* for the designated cathegory as well as the *reliability* and *diagnostic coverage*.

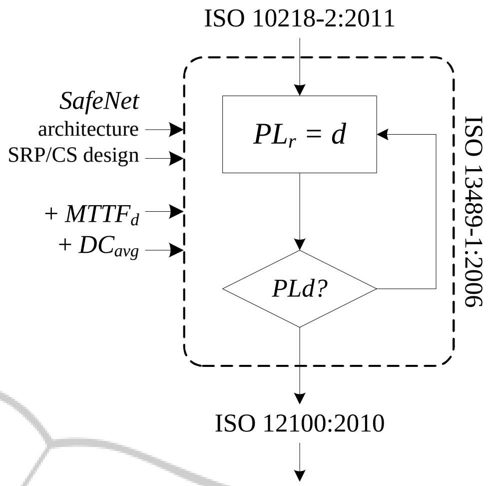

Figure 12: Design and evaluation process for robot system safety: application and scenario requirements from ISO 10218-2:2011 are elaborated according to functional safety procedures (ISO 13489-1:2006) until validation, before proceeding to general risk assessment in ISO 12100:2010.

ACKNOWLEDGEMENTS

This work has been partially suported by CNR Flagship Program "Factory of the Future", FdF-SP1-T3.1, Project FACTOry Technologies for HUMans Safety.

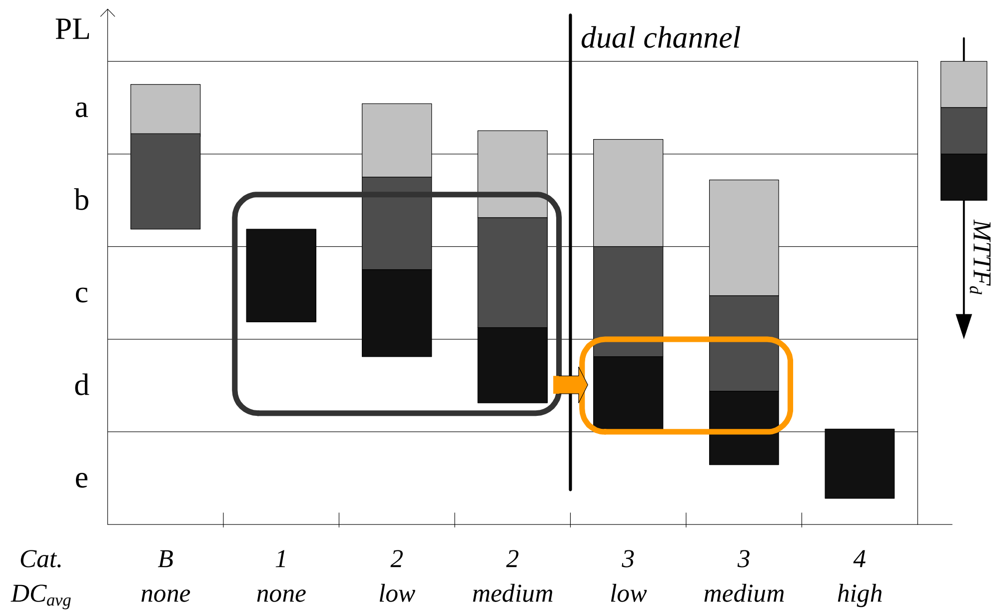

Figure 13: Effect of SafeNet methodology on the PL assessment: a generic network of unsafe devices is outlined in gray box, while applying redundancy architectural measures it is possible to move to *PL d* class (outlined in dark yellow).

REFERENCES

ABB Robotics (2008). EPS and SafeMove White Paper WHPEPS-20.

Alami, R., Albu-Schäffer, A., Bicchi, A., Bischoff, R., Chatila, R., De Luca, A., De Santis, A., Giralt, G., Guiochet, J., Hirzinger, G., Ingrand, F., Lippiello, V., Mattone, R., Powell, D., Sen, S., Siciliano, B., Tonietti, G., and Villani, L. (2006). Safe and Dependable Physical Human-Robot Interaction in Anthropic Domains: State of the Art and Challenges. In Bicchi, A. and De Luca, A., editors, *Procceedings IROS Workshop on pHRI - Physical Human-Robot Interaction in Anthropic Domains*.

Albu-Schäffer, A., Haddadin, S., Ott, C., Stemmer, A., Wimböck, and Hirzinger, G. (2007a). The dlr lightweight robot: design and control concepts for robots in human environments. *Industrial Robot: An International Journal*, 34:376–385.

Albu-Schäffer, A., Ott, C., and Hirzinger, G. (2007b). A unified passivity-based control framework for position, torque and impedance control of flexible joint robots. *The International Journal of Robotics Research*, 26(1):23–39.

Bicchi, A., Peshkin, M. A., and Colgate, J. E. (2008). Safety for physical human-robot interaction. In *Springer Handbook of Robotics*, pages 1335–1348. Springer Berlin / Heidelberg.

De Santis, A., Siciliano, B., De Luca, A., and Bicchi, A. (2008). An atlas of physical human-robot interaction. *Mechanism and Machine Theory*, 43(3):253–270.

Decotignie, J.-D. (2005). Ethernet-based real-time and industrial communications. *Proceedings of the IEEE*, 93(6):1102 –1117.

Felser, M. (2005). Real-time ethernet - industry prospective. *Proceedings of the IEEE*, 93(6):1118 –1129.

Grebenstein, M., Chalon, M., Friedl, W., Haddadin, S., Wimböck, T., Hirzinger, G., and Siegwart, R. (2012). The hand of the dlr hand arm system: Designed for interaction. *The International Journal of Robotics Research*, 31(13):1531–1555.

Gupta, R. and Chow, M.-Y. (2010). Networked control system: Overview and research trends. *Industrial Electronics, IEEE Transactions on*, 57(7):2527 –2535.

Haddadin, S., Albu-Schäffer, A., De Luca, A., and Hirzinger, G. (2008). Collision detection and reaction: A contribution to safe physical human-robot interaction. In *Intelligent Robots and Systems, 2008. IROS 2008. IEEE/RSJ International Conference on*, pages 3356–3363.

Haddadin, S., Albu-Schäffer, A., and Hirzinger, G. (2009). Requirements for safe robots: Measurements, analysis and new insights. *The International Journal of Robotics Research*, 28:1507–1527.

Hespanha, J., Naghshtabrizi, P., and Xu, Y. (2007). A survey of recent results in networked control systems. *Proceedings of the IEEE*, 95(1):138 –162.

IEC (2010). *IEC 61508:2010: Functional safety of electrical/electronic/ programmable electronic safety-related systems*. International Electrotechnical Commission, Geneva, Switzerland.

ISO (2006). *ISO 13849-1:2006: Safety of machinery – Safety-related parts of control systems – Part 1: General principles for design*. International Organization for Standardization, Geneva, Switzerland.

ISO (2010). *ISO 12100:2010: Safety of machinery – General principles for design – Risk assessment and risk reduction*. International Organization for Standardization, Geneva, Switzerland.

ISO (2011a). *ISO 10218-1:2011: Robots and robotic devices – Safety requirements for industrial robots – Part 1: Robots*. International Organization for Standardization, Geneva, Switzerland.

ISO (2011b). *ISO 10218-2:2011: Robots and robotic devices – Safety requirements for industrial robots – Part 2: Robot systems and integration*. International Organization for Standardization, Geneva, Switzerland.

ISO (2011c). *ISO/TS 15066:2011: Robots and robotic devices Collaborative robots*. International Organization for Standardization, Geneva, Switzerland.

Kock, S., Vittor, T., Matthias, B., Jerregard, H., Kallman, M., Lundberg, I., Mellander, R., and Hedelind, M. (2011). Robot concept for scalable, flexible assembly automation: A technology study on a harmless dual-armed robot. In *Assembly and Manufacturing (ISAM), 2011 IEEE International Symposium on*, pages 1 –5.

KUKA Roboter (2012). KUKA.SafeOperation product catalog.

Moakher, M. (2002). Means and averaging in the group of rotations. *SIAM Journal on Matrix Analysis and Applications*, 24(1):1–16.

Moyne, J. and Tilbury, D. (2007). The emergence of industrial control networks for manufacturing control, diagnostics, and safety data. *Proceedings of the IEEE*, 95(1):29 –47.

Pratt, G. and Williamson, M. (1995). Series Elastic Actuators. In *Intelligent Robots and Systems 95. 'Human Robot Interaction and Cooperative Robots', Proceedings. 1995 IEEE/RSJ International Conference on*, volume 1, pages 399–406.

Strobl, K. and Hirzinger, G. (2006). Optimal hand-eye calibration. In *Intelligent Robots and Systems, 2006 IEEE/RSJ International Conference on*, pages 4647–4653.

Tsai, R. (1987). A versatile camera calibration technique for high-accuracy 3D machine vision metrology using off-the-shelf tv cameras and lenses. *Robotics and Automation, IEEE Journal of*, 3(4):323–344.

Vogel, C., Poggendorf, M., Walter, C., and Elkmann, N. (2011). Towards safe physical human-robot collaboration: A projection-based safety system. In *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on*, pages 3355 –3360.

Zinn, M., Khatib, O., Roth, B., and Salisbury, J. (2004). Playing it safe [human-friendly robots]. *Robotics Automation Magazine, IEEE*, 11(2):12–21.