



Research Paper

The use of the Simplex architecture to enhance safety in deep-learning-powered autonomous systems[☆]

Federico Nesti^a,^{*}, Niko Salamini^a, Mauro Marinoni^a,^b, Giorgiomaria Cicero^{a,b}, Gabriele Serra^b, Alessandro Biondi^{a,b}, Giorgio Buttazzo^a

^a Department of Excellence in Robotics & AI, Scuola Superiore Sant'Anna, via Moruzzi 1, Pisa, Italy

^b Accelerat S.r.l., Via L. Alamanni Lotto D/2, 5A, San Giuliano Terme, Italy

ARTICLE INFO

Keywords:

Safe and secure machine learning
Reinforcement learning
Simplex architecture
Functional monitor

ABSTRACT

Recently, the outstanding performance reached by neural networks in many tasks has led to their deployment in autonomous systems, such as robots and vehicles. However, neural networks are not yet trustworthy, being prone to different types of misbehavior, such as anomalous samples, distribution shifts, adversarial attacks, and other threats. Furthermore, frameworks for accelerating the inference of neural networks typically run on rich operating systems that are less predictable in terms of timing behavior and present larger surfaces for cyber-attacks.

To address these issues, this paper presents a software architecture for enhancing safety, security, and predictability levels of learning-based autonomous systems. It leverages two isolated execution domains, one dedicated to the execution of neural networks under a rich operating system, which is deemed not trustworthy, and one responsible for running safety-critical functions, possibly under a different operating system capable of handling real-time constraints.

Both domains are hosted on the same computing platform and isolated through a type-1 real-time hypervisor enabling fast and predictable inter-domain communication to exchange real-time data. The two domains cooperate to provide a fail-safe mechanism based on a safety monitor, which oversees the state of the system and switches to a simpler but safer backup module, hosted in the safety-critical domain, whenever its behavior is considered untrustworthy.

The effectiveness of the proposed architecture is illustrated by a set of experiments performed on two control systems: a Furuta pendulum and a rover. The results confirm the utility of the fall-back mechanism in preventing faults due to the learning component.

1. Introduction

The recent explosion of artificial intelligence (AI) gave rise to the development of deep neural networks (DNNs) with super-human performance in many specific tasks, such as image classification (Masana et al., 2022), object detection (Zou et al., 2023), and control (Rybczak et al., 2024). These results attracted the interest of several industries, which started adopting deep learning methods in cyber-physical systems (CPSs), such as robots, drones, and vehicles, to perform complex tasks and increase their level of autonomy. As a consequence, a growing number of autonomous systems are being deployed, and self-driving cars are seemingly a few years away.

Nonetheless, the outstanding performance of AI algorithms is not free from drawbacks. One of the main obstacles to the deployment

of deep learning models in CPSs is their low level of trustworthiness with respect to the strict industrial-grade standards applied in critical fields. Trustworthiness typically refers to a set of desirable properties that an AI-based CPS should have, including safety, security, and time predictability.

Safety and fault-tolerance. Deep learning models are known to produce imperfect predictions for many reasons. A common cause is the distribution shift between train and test datasets that frequently occur when neural networks operate in the real world, e.g., when new unseen objects appear in the scene (Quiñonero-Candela et al., 2022). In these cases, such unreliable outputs should be detected and excluded by a fault-tolerant mechanism, which should also prevent the propagation of faults and ensure that the system can be kept in a safe state.

[☆] This article is part of a Special issue entitled: 'AI-Driven Innovations in Cyber-Physical Systems' published in Engineering Applications of Artificial Intelligence.

^{*} Corresponding author.

E-mail address: federico.nesti@santannapisa.it (F. Nesti).

Security. Differently from the software development practice adopted for safety-critical systems, the implementation of a neural network typically relies on sophisticated frameworks (e.g., PyTorch or TensorFlow) developed without following security- or safety-related programming practices, running on a rich operating system (e.g., Linux), and consisting of large code-bases that rely on several third-party libraries. These software components are all targets for cyber-attacks. For this reason, such components should be isolated from the safety-critical ones to protect the latter from possible attacks originated in non-trustworthy software. In addition to that, DNNs are prone to several types of adversarial attacks (Szegedy et al., 2014; Biggio and Roli, 2018), namely small malicious perturbations applied to the input that are imperceptible to humans, but can cause a neural network to produce a wrong output with a very high confidence score. Such adversarial attacks can also be performed in the physical world, i.e., without altering the digital representation of the DNN input, through dedicated patches or 3D-printed adversarial objects placed in the environment (Kurakin et al., 2017; Nesti et al., 2022).

Time predictability. Popular software and hardware platforms used to execute DNN inference tasks typically exhibit low predictability in the time domain, which causes model outputs to be delivered with highly variable delays. In fact, both deep learning frameworks and modern hardware acceleration technologies are designed to optimize the average-case behavior, rather than the worst-case one. Consequently, response times are subject to considerable variability that can degrade the overall system performance and compromise system stability (Casini et al., 2020). For instance, accelerating neural models on general-purpose graphics processing units (GPGPUs) significantly reduces the average response time, but introduces highly variable delays in their execution (Cavicchioli et al., 2017; Buttazzo, 2022). This phenomenon is exacerbated if multiple DNNs are to be executed on the same accelerator, due to the higher number of conflicts that may arise on the shared computational resources (e.g., caches, buses, memory banks, etc.).

1.1. This paper

To cope with these problems, this paper proposes an architectural solution that can address safety, security, and predictability issues in CPSs that include AI-based components, such as advanced sensory perception modules or neural controllers. This is obtained by (1) separating computations in two execution domains established by a type-1 real-time hypervisor and (2): building a Simplex architecture driven by a *safety monitor*.

The two domains are defined as follows: a *rich domain*, in charge of running all the AI-based components on a rich OS, and a *safe domain*, responsible for executing all the safety-critical functions in a more protected environment. In our experimental setting, the two domains are hosted with strong isolation on a single AMD Zynq Ultrascale+ system-on-chip and managed by CLARE, a type-1 real-time hypervisor that is part of the Clare Software Stack (Accelerat Srl, 2026). The two domains cooperate by means of inter-domain communication provided by the CLARE hypervisor, which guarantees bounded execution latency and jitter, as well as strong isolation mechanisms, such as secure cache partitioning and memory bandwidth reservation (Modica et al., 2018) to reduce inter-domain interference. Security hardening of the whole system is established at the hypervisor level by means of run-time security monitoring, address space layout randomization, and control-flow integrity.

To enhance safety, this work is inspired by the Simplex architecture (Sha et al., 1994; Bak et al., 2009; Mohan et al., 2013a), which relies on a *safety monitor* that analyzes the state of the system, and/or the behavior of DNNs, commanding a switch to a simpler but safer controller to recover from potentially unsafe states.

The effectiveness of the proposed architecture has been tested and evaluated using two control systems: a Furuta inverted pendulum and

an autonomous vehicle based on the AgileX Scout Mini rover equipped with a camera and a LiDAR. In both cases, the proposed architecture proved to be capable of managing the unsafe behavior of the AI-based components in a reasonable and effective way.

In summary, this work makes the following contributions:

- We propose a general architecture for AI-powered cyber-physical systems that runs on a single platform with two execution domains managed by a type-1 real-time hypervisor. While previous hypervisor-based architectures focused on specific applications without generalizing to diverse CPSs with AI components, our design ensures strong spatial and temporal isolation between domains while guaranteeing safety, security, and predictability for any CPS by leveraging a Simplex-inspired switching mechanism.
- We present two practical implementations of the architecture and its components that manage unsafe behaviors and faults in AI components. While prior works either provided only theoretical designs or validated them on limited systems, our implementations demonstrate the approach on both a Furuta inverted pendulum and an autonomous rover, highlighting the applicability across different CPS types.

2. Background and related work

The research concerning safe and secure architectures for AI-powered CPSs is wide and multi-faceted, as reported by several surveys (Pereira and Thomas, 2020; Olowononi et al., 2020). Given the scope of this paper, we focus on previous work using DNNs for autonomous systems.

A line of research within the safe and robust deep learning focuses on verifying and testing DNNs. Verification methods aim at formally proving that certain properties hold for a DNN, while testing methods seek to find corner cases that cause a DNN to produce incorrect outputs. Despite the broad interest in such topics, they are not addressed in this paper since the focus is on architectural aspects of safety and security, while neural networks are considered untrustworthy by nature. The interested reader may refer to Huang et al. (2020) for an in-depth discussion. The key concept of our work is that the AI components are considered untrustworthy by definition. Hence, it is necessary to introduce a dedicated architecture to isolate their unsafe behaviors.

2.1. Hypervisor-based architectures

As the use of DNNs in CPSs is a relatively new topic, there is still no established approach for developing safe, secure, and predictable AI-powered CPSs. Currently, several architectural solutions have been proposed to handle the untrustworthy behavior of DNNs.

A typical approach consists of executing the AI-based components and the safety-critical ones on different platforms (Reke et al., 2020; Liu et al., 2017; Gutiérrez et al., 2018), communicating via the Robot Operating System (ROS) or custom channels. Such design requires multiple platforms with redundant resources and hardware components, making it unsuited for autonomous systems that sometimes present stringent size, weight, and power requirements.

Although other single-platform architectures have been proposed (O’Kelly et al., 2019; Meier et al., 2015), adopting hypervisor technology, to the best of our knowledge, is becoming the de-facto solution for supporting mixed-criticality applications. In fact, type-1 hypervisors are favored in mixed-criticality systems due to their small code base, which reduces both attack surface and execution overhead, and their direct hardware control, which enhances security, safety, and time predictability. Typically, they also offer time-predictable latency, efficient inter-domain communication, and effective scheduling of virtual machines (Cittadini et al., 2023).

Several research efforts have integrated type-1 hypervisors in specific CPS applications. For instance, Klein et al. (2018) used seL4 to

separate trusted and untrusted software in UAVs, ensuring security and isolation. Almeida and Prochazka (2009) employed PikeOS for secure partitioning in spacecraft avionics. Craveiro et al. (2009) utilized ARINC-653 standard in the AIR operating system for aerospace systems, focusing on IA-32 and Sparc architectures. Perez and Gutiérrez (2017) implemented real-time communication in the Xstratum hypervisor, integrating ARINC-653 standard with support for DDS. Farrukh and West (2022) developed a low-overhead hypervisor solution by providing timing guarantees using SCHED_DEADLINE for Linux, though it has limitations for complex AI-based solutions. Biondi et al. (2021) proposed a hypervisor-based architecture for safety-critical systems, providing isolation, security, and real-time communication, while Cittadini et al. (2023) presented a dual-domain architecture applied to UAVs. These studies highlight the versatility and advantages of type-1 hypervisors in various CPS applications. Building on such previous work, our paper focuses on implementing an architecture for safe use of DNNs in the wild. The main strength and novelty of this paper is that the proposed architecture is general enough to serve as a blueprint for safe, secure, and predictable integration of AI-based components across diverse types of CPS. The hardware and software stack can be configured depending on the specific application requirements, as demonstrated by the two reported case studies (an inverted pendulum and an autonomous rover). Such two prototypes provide concrete guidelines that readers can adapt for their own applications.

2.2. The Simplex architecture

The Simplex architecture was originally proposed to safely test newly-designed and complex controllers on safety-critical systems (Sha et al., 1994) and perform safe online control system updates (Seto et al., 1998; Bak et al., 2011). Further developments employed the Simplex architecture to enhance the security (Mohan et al., 2013b) and safety (Vivekanandan et al., 2016) of CPSs. It was also recently used to adopt deep learning in safety-critical applications (Biondi et al., 2020), in a programming framework to ensure safety of ROS-based systems (Desai et al., 2019), and also to safely integrate reinforcement learning components in CPSs (Phan et al., 2020).

The main idea behind the Simplex architecture is the use of the following three interacting components:

- A *safe controller*: it is a simple controller designed to manage situations that might lead the system in an unstable or dangerous behavior. It typically performs simple halt procedures and/or safe operations defined by classical control techniques.
- A *high-performance controller*: it is an untrustworthy controller, typically showing higher performance with respect to the safe controller but without any guarantee of safety.
- A *safety monitor or runtime monitor*: it is a module responsible for deciding whether it is safe to let the high-performance controller regulate the system or it is better to switch to the safe controller.

The switching rule adopted by the safety monitor is tailored for the specific system and requires to be designed for each use case. This work uses the Simplex architecture to increase safety when an AI-based component is adopted in the system, presenting the implementation details for two specific use cases that can be useful for the practitioner interested in developing such systems.

2.3. Safety monitor

Many different runtime monitoring functions have been proposed in the literature. From a classical control perspective, a closed-loop system can be analyzed around a certain equilibrium with Lyapunov-based techniques to estimate the region of asymptotic attraction (ROA) (Zubov, 1965), a subset of the state space where the system is guaranteed to converge to the equilibrium. Estimating the ROA of the

system regulated by the safe controller is vital to extract the boundaries that the state shall not overcome when controlled with the high-performance controller, since it will probably become unrecoverable by the safe controller. When both the plant and the safe controller are simple enough, Lyapunov-based techniques might be the safest options, as they always provide subsets of the actual ROA; it is also possible to optimize the maximum-volume ellipsoid that fits inside the ROA (Seto and Sha, 1999). This approach has been used to safely control a Furuta inverted pendulum in Section 4. Alternatively, neural networks can be trained to provide an estimation of the ROA (Ferreira and Krogh, 1997, 1999), although it is advisable to somehow certify the behavior of the neural network before deployment in the safety monitor.

A different approach is based on what is called reachability analysis (Bak et al., 2014), which estimates the set of states that might be reached from the current state. This indicates whether the system is on the verge of instability, or, in general, close to an unsafe state set. However, such techniques can be resource-intensive and might not be feasible for the stringent memory and processing requirements of real-time systems. Moreover, they also present scalability limitations when dealing with complex high-dimensional systems.

In some complex cases, it is difficult to define equations of a system that are inclusive of the unsafe situations that the system might encounter. In these cases, it is possible to define a safety envelope, as we adopted to achieve safe control of the rover. The safety envelope can be configured to impose limitations to the maximum velocity, as well as to restrict the allowed position or distance from obstacles in close proximity. Additionally, it is possible to use monitoring techniques that analyze the behavior of the neural network: Uncertainty estimation (Abdar et al., 2021), out-of-distribution detection (Yang et al., 2024), or adversarial defenses (Rossolini et al., 2023) might be used in combination to find the situations where the neural network does not provide safe outputs. These techniques are not considered for this paper, but future work will investigate them.

2.4. Comparison

While several studies adopted hypervisor- or Simplex-based architectures for safe CPS operation, only a few works explicitly treated AI components as untrusted. For instance, Cittadini et al. (2023) presented a dual-domain hypervisor-based architecture for supporting AI-powered CPSs on heterogeneous platforms, but did not propose a general methodology for managing faults in the AI model.

In contrast, our work proposes a general architecture to explicitly manage AI faults across diverse CPS applications, demonstrating the effectiveness of the approach on two case studies, an inverted pendulum and an autonomous rover. Table 1 summarizes the key differences between prior architectural approaches and the proposed work.

3. Proposed architecture

This section provides a high-level description of the architecture introduced in Section 1, focusing on its main components. The implementation details specific to the two considered use cases are described in Sections 4 and 5, respectively.

The proposed architecture is illustrated in Fig. 1 and consists of two execution domains concurrently running in strong isolation on the same hardware platform. CLARE hypervisor is in charge of managing two isolated Virtual Machines (VMs) on the hardware platform that may host stand-alone software execution environments powered by rich OSes (e.g., Linux) as well as real-time OSes or firmware (e.g., FreeRTOS). Each domain has both private and shared resources, whose access is defined at configuration time and then controlled by CLARE hypervisor at run time. In particular, each domain may have (i) exclusive access to computing resources (e.g., CPUs and accelerators), devices (e.g., Ethernet interfaces, DMAs), and memory resources, and (ii) mutual access to

Table 1

High-level comparison between related architectural approaches and this work. The AI as *untrusted* dimension is decomposed into three subcategories.

Approach	Single platform	Hypervisor-based	AI components			Validated on CPS
			focus on AI	AI as untrusted	AI faults	
Multi-platform AI/CPS architectures (Reke et al., 2020; Liu et al., 2017; Gutiérrez et al., 2018)	No	No	No	No	No	Yes
Simplex-based AI control (Phan et al., 2020; Vivekanandan et al., 2016; Mohan et al., 2013b)	Yes	No	Yes	Yes	No	Yes
Hypervisor-based CPS (no AI focus) (Klein et al., 2018; Almeida and Prochazka, 2009; Craveiro et al., 2009; Perez and Gutiérrez, 2017; Farrukh and West, 2022)	Yes	Yes	No	No	No	Yes
Hypervisor-based CPS with AI components (Biondi et al., 2021, 2020; Cittadini et al., 2023)	Yes	Yes	Yes	Yes	Partial	No
Cittadini et al. (2023)	Yes	Yes	Yes	Partial	No	Yes
This work	Yes	Yes	Yes	Yes	Yes	Yes

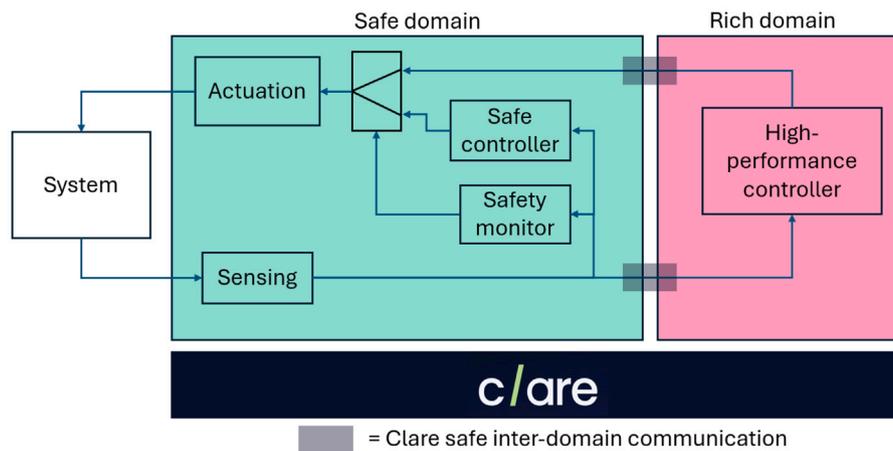


Fig. 1. Overall architecture of the system with functional blocks.

other physical and logical resources (e.g., shared-memory areas, inter-domain signals) that are mainly used for inter-domain communication mechanisms. The proposed architecture is composed of two isolated VMs that host the safe and rich domains, respectively. Each domain has private access to its dedicated resources required to implement its functionality, and to two inter-domain channels to implement full-duplex communication.

The two domains have different levels of criticality:

- The safe domain has a high criticality and is responsible for all the safety-critical tasks, such as sensing, actuation, communication, safe control, and safety monitoring. This domain is typically powered by a real-time operating system (RTOS).
- The rich domain has a low criticality and is responsible for all the high-performance computations. In the scheme illustrated in Fig. 1, the only AI-based component is the high-performance controller, but other controllers and low-criticality tasks might be present. This domain is deemed untrustworthy with respect to safety, security, and time predictability, and it is typically powered by a rich operating system such as Linux.

The default operational mode of the system uses the high-performance AI-based controller in the rich domain. In such an operational mode, the safety monitor constantly evaluates the state of the system to detect any anomaly or dangerous situation. Whenever this occurs, the AI-based controller is disconnected and the output of the safe controller is used instead. Please note that while this is not showed in Fig. 1 for simplicity, the safety monitor may be triggered also

by other conditions specific for each implementation. For instance, the safety monitor can read the output of the high-performance controller or its internal activations (in case a neural network is used). The safety monitor switches back to the AI-based controller whenever it is deemed safe, functionally acting as a state machine, which is general enough to implement the required switching strategy for the system at hand. For instance, one could use the dead-zone approach, hysteresis switching, or blending control (Zhu and Antsaklis, 2015).

CLARE hypervisor enhances system security through VM runtime monitoring. When anomalous execution behaviors are detected within a virtual machine, the hypervisor can be configured to shut down the affected VM and restore a trusted backup. Since the monitoring mechanism operates outside the VM, it remains isolated and cannot be compromised by malicious software running inside the VM. In addition, CLARE integrates address space layout randomization and control-flow integrity at the hypervisor level. All these features are natively supported by CLARE and are designed to introduce only minimal overhead, thereby preserving system performance while providing strong security guarantees. The details about the implementation of these security mechanisms are out of the scope of this study.

This architecture is general enough to be seamlessly adapted and configured in different ways to match the system requirements. The actual configuration of the architecture in terms of resource assignment depends on the target application and its criticality level. Specific configurations are discussed in detail in the following sections when addressing the two case studies.

The proposed architecture provides multiple benefits: (i) Possibility of executing a high-performance but untrustworthy controller on a

physical system, safely handling sporadic wrong outputs by switching to a simpler but safer controller; (ii) Capability of safely managing deadline misses caused by the unpredictable timing behavior of the rich OS by exploiting the switching mechanism; (iii) Capability of isolating all the safety-critical components in the high-criticality domain, protecting the physical system from possible cyber-attacks to the rich OS or adversarial attacks against the AI components; (iv) Run-time security monitoring at the hypervisor level; (v) Possibility to train a neural controller directly on a physical system to fine-tune the network to refine the model.

To achieve the benefits mentioned above, however, the two isolated domains have to be carefully managed by properly addressing the following potential issues:

- The use of two controllers in alternation could result in an unstable behavior if not properly handled.
- The end-to-end latency resulting from the cooperation of the two domains must be precisely characterized and taken into account in the design of software activities since it could jeopardize the system stability or result in poor performance due to the reduced utilization of the AI components.
- Typically, an embedded computing platform has several devices. Each device must be managed by a domain, potentially decreasing the performance of the system; hence, the design of the architecture must be decided carefully, depending on the specific application requirements.

All these aspects are analyzed in the context of two use cases presented in Sections 4 and 5, each discussing a specific architectural configuration, the design of the control loop period, and the switching rules of the safety monitor.

4. First case study: the Furuta pendulum

This section describes the specialization and implementation of the architecture presented in Section 3 to safely control a Quanser Furuta pendulum (Quanser, 2026). An overview of the hardware system and the software tasks is first provided and then followed by a description of the safety monitor and the controllers. Experiments are finally presented for evaluation purposes. The code implementation will be made available upon request.

The primary safety requirement considered in this basic system is that the pendulum's pole must always be kept in equilibrium. As the dynamics of the pendulum evolves rapidly, this system constitutes a good case study to assess the real-time capabilities of the proposed control architecture. Please note that, although DNN-based controllers are not strictly required to control an inverted pendulum, this system provides interesting and valuable insights useful for the implementations on more complex systems, such as the one described in Section 5.

The Furuta pendulum, technically a rotary inverted pendulum, is in fact one of the standard benchmarks for real-time control problems. Fig. 2 shows the pendulum used in this work. The controller acts on the DC motor voltage with a PWM signal. The motor axis is linked to the horizontal arm (from now on simply indicated as the arm), while the pendulum rotates freely at the end of the arm. Both angle values are acquired by an encoder. The angle of the arm is α , while the angle of the pendulum with respect to the vertical is θ . The state of the system consists of the arm and pendulum angles and the corresponding angular velocities, formally $\mathbf{x} = (\theta, \alpha, \dot{\theta}, \dot{\alpha})^T \in \mathcal{X} \subset \mathbb{R}^4$. The system is modeled in the standard form $\dot{\mathbf{x}} = f(\mathbf{x}, u)$, where $u \in \mathbb{R}$ is the control input and f is the non-linear dynamics. The equations of the system can be found in Cazzolato and Prime (2011).

The pole is brought in the upright position with a standard energy-based swing-up controller (Åström and Furuta, 2000), while the high-performance controller is used for pendulum stabilization when it is already up.

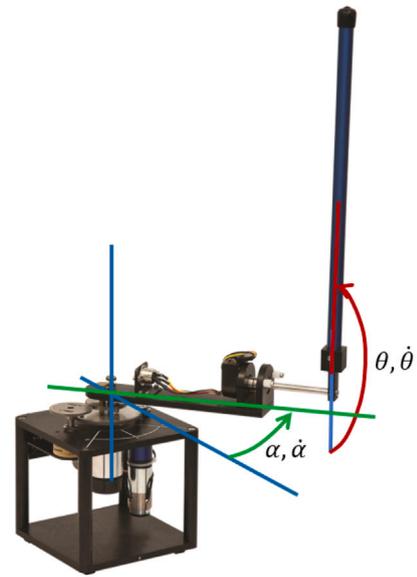


Fig. 2. Quanser's Furuta pendulum used in this demo.

4.1. Hardware and on-board software

The reference computing platform is an AMD Zynq Ultrascale+ MPSoC (mounted on a ZCU102 board), a heterogeneous SoC that includes four Cortex-A53 cores, two Cortex-R5 cores, an FPGA-based Programmable Logic (PL), and several I/O peripherals (Xilinx, 2026). Thanks to its level of heterogeneity, this platform is suitable for different types of applications and it is also a good candidate for hosting multiple applications with mixed levels of criticality. The platform is managed by the CLARE Software Stack to establish two execution domains:

- a *Rich VM*, powered by Petalinux 2018.2 (an AMD/Xilinx Linux distribution), using three virtual CPUs with 2 GB of RAM and hosting the Caffe framework to perform the inference of neural networks;
- a *Safe VM*, powered by the Erika3 RTOS (Evidence Srl, 0000), using one virtual CPU with 256 MB of RAM, which hosts the safe controller, the system monitor, and a logging and telemetry component.

The sensing part is composed of two position encoders, one for the arm and one for the pendulum, which are sampled at 16 kHz through a specific logic synthesized on the FPGA. Please note that, although the FPGA is currently used only for encoder reading, its capability to support additional tasks (e.g., neural network acceleration) makes the chosen platform suitable for more complex applications. An Erika3 task controls the PWM actuation by driving a dedicated timer. Another Erika3 task in charge of telemetry sends the variables of interest through an UART port with a baudrate of 115200. The data are then sent to a web-based application running on a PC for monitoring and logging purposes.

As shown in Fig. 3, the Rich VM has no direct access to any hardware device involved in the plant control loop. Every access is mediated by the Safe VM, which communicates with the Rich VM through the inter-VM communication mechanism provided by CLARE. In particular, Erika3 tasks sample the position of both the arm and pendulum, computes the state of the system, and sends it to the Rich VM. The latter continuously monitors this port through a polling task, waiting for a new message. If the port contains a new message, a Linux application reads the state and performs a neural network inference

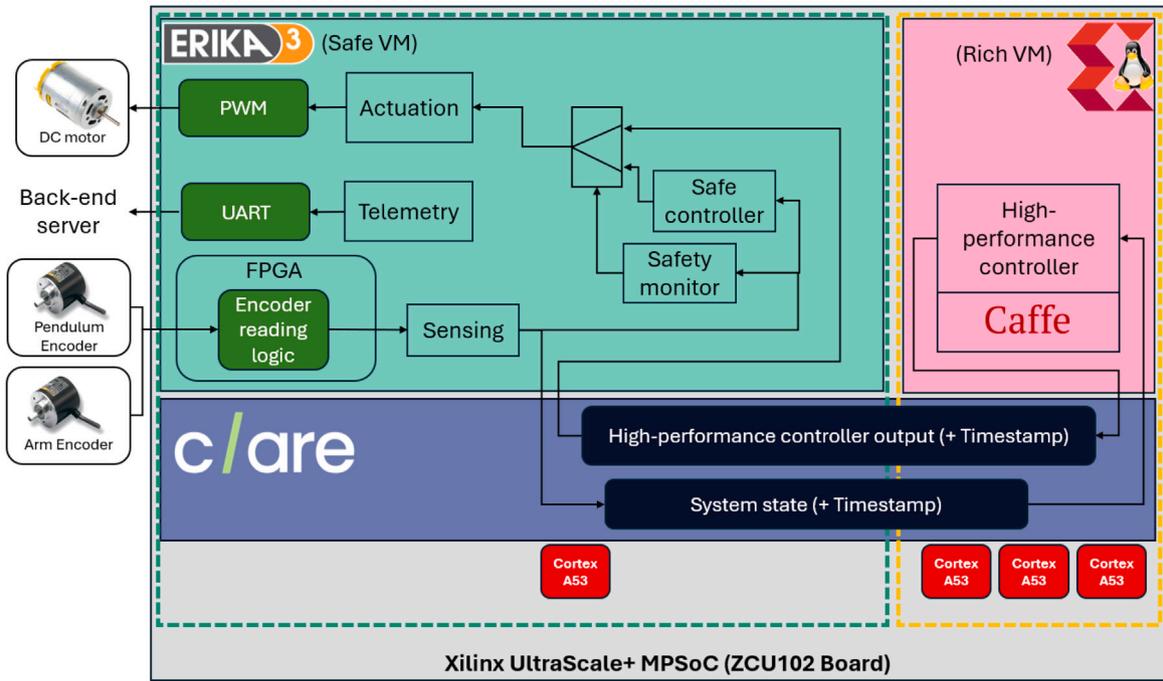


Fig. 3. Scheme of the hardware/software architecture managed by the CLARE hypervisor for safe and secure control of the Quanser rotary inverted pendulum. Please note that the blocks represent functional components and might not have a direct correspondence with the task set.

Table 2
Period, deadline and priority of the task set of the Safe VM (Erika3).

Task	Period [ms]	Deadline [ms]	Priority
SafetyMonitor	2	2	4
Control	4	4	3
Stream	10	10	2
Interface	20	20	1

using the Caffe framework. The output of this neural network is sent to the high-performance controller through another port and the loop is repeated. In this way, the Safe VM can safely read the output of the high-performance controller from this dedicated port. The sampling ports have a 1-to-N semantic and can contain one sample only. In this case, the channel is configured as a 1-to-1 port and always contains the freshest sample.

The software running on Erika3 combines auto-generated code (from Matlab/Simulink) and hand-written code. The Erika3 task set is shown in Table 2. In the priority column, higher values indicate higher priorities.

The **SafetyMonitor** and the **Interface** tasks host the auto-generated code to handle the finite state machine required for the safety monitor switching mechanism. The safety monitor (i) checks the risk of instability according to the system state, and (ii) checks whether the output port of the high-performance controller contains fresh messages, meaning that Linux is not missing any deadline. The safety monitor algorithm and the other controller involved are described in Section 4.2.

The **Control** task is responsible for:

- sensing (acquiring measurements from encoders);
- sending the measurements to the state port (towards the Rich VM);
- executing the safe controller algorithm;
- reading the advanced controller output from Linux through the dedicated port; and
- sending the correct actuation value to the PWM, according to the decision of the safety monitor.

Finally, the **Stream** task is responsible for periodically sending the system status to a web-based application through the UART, for visualization purposes.

4.2. Controllers and safety monitor

Two controllers are employed: the swing-up controller and the safe controller. A finite state machine decides which controller must be activated, depending on the state of the system. The state machine runs on the safe domain, together with all the safety-critical tasks (i.e., the two controllers mentioned above and the safety monitor).

At the system start-up, the pendulum is assumed to stand in the down-hanging position $\mathbf{x} = (0, 0, 0, 0)$, from which a standard energy-based swing-up controller (Åström and Furuta, 2000) is used to bring the pole close to the upright position $\mathbf{x} = (0, \pi, 0, 0)$. The safe controller used for this application is a standard proportional controller $u = Ke$, where e is the error between a certain reference point and the current state \mathbf{x} , and $K \in \mathbb{R}^{1 \times 4}$ is computed by standard linearization and pole-placing techniques to guarantee a stable behavior.

The high-performance controller is described in Section 4.3, but its behavior does not affect the switching mechanism of the safety monitor, which solely relies on the definition of safe controller and the notion of Region of Asymptotic Attraction (ROA).

Given an equilibrium point \mathbf{x}_0 in the state space (it can be chosen $\mathbf{x}_0 = \mathbf{0}$ without loss of generality), the ROA is the subset of the state space where the system is asymptotically attracted to the equilibrium, i.e., the system is stable. In this implementation on the Furuta pendulum, we are interested in estimating the ROA of the system controlled by the safe controller, since we want to know which states can be safely recovered to the equilibrium.

Formally, the ROA of the controlled system is defined as

$$ROA = \{ \mathbf{x} \in \mathcal{X} \text{ s.t. } \lim_{t \rightarrow \infty} \mathbf{x}(t) = \mathbf{0}, \mathbf{x}(0) = \mathbf{x} \}, \quad (1)$$

where $\mathbf{x}(t)$ is the state indexed in time t .

The ROA of a non-linear system cannot be computed in closed form, and must be estimated numerically. Section 2 provided a brief review of the methods proposed in the literature to estimate the ROA. The ROA

estimation technique used in this paper is based on approximating the ROA as an ellipsoid centered in the equilibrium (i.e., the origin), having a quadratic form $\mathbf{x}^T P \mathbf{x} < 1$, where $P \in \mathbb{R}^{4 \times 4}$. The value of the symmetric positive-definite matrix P can be found with an optimization process that aims at maximizing the volume of the ellipsoid, as explained in Seto and Sha (1999).

By the ROA definition, it is possible to define a subset \mathcal{R} of the state space that contains only the states that can be recovered safely to the equilibrium by the safe controller, that is

$$\mathcal{R} = \{\mathbf{x} \mid \mathbf{x}^T P \mathbf{x} < 1\} \subset ROA. \quad (2)$$

If the current state \mathbf{x} can be checked to be in \mathcal{R} , then we can ensure that \mathbf{x} is recoverable by the safe controller and the high-performance controller can be safely used (i.e., without the risk of instability). As soon as the current state exits \mathcal{R} , the safety monitor switches to the safe controller that is able to bring the state back to the equilibrium, thus preventing risky conditions that could possibly be caused by the high-performance controller. This kind of switching between the high-performance and a safe controller is typical of the Simplex architecture (Sha et al., 1994), as it allows exploiting the higher performance of non-certified controllers safely.

To prevent frequent switches between the two controllers, the commutation to the high-performance controller occurs when a safe subset $S \subset \mathcal{R}$ is reached by the state. In particular, the safe subset can be defined as

$$S(\theta_s) = \{\mathbf{x} \mid \mathbf{x}^T P \mathbf{x} < \theta_s\}, \quad (3)$$

where $\theta_s < 1$ is a user-selected threshold that defines how close to the equilibrium the safe subset should be.

Hence, the scalar quantity $\mathbf{x}^T P \mathbf{x}$ represents a sort of *instability index*. When such an index is close to 0, it means that the state is close to the equilibrium. Conversely, when the instability index is close to 1, the system is close to an unrecoverable instability. We empirically choose the threshold $\theta_s = 0.1$ to allow the safety controller to bring the state close to equilibrium before allowing the high-performance controller to intervene again. A simplified visualization of the stability regions used by the switching-based state machine is illustrated in Fig. 4, which shows the actual ROA of recoverable state (gray area), which is usually unknown, the estimated ROA \mathcal{R} (yellow area) and the safe region S (green area).

It is worth noting that to guarantee the stability of the overall system, a one-step-ahead simulation of the state is required. This is because, in general, a wrong control input at time t issued by the high-performance controller could move the state $\mathbf{x}(t+1)$ not only outside the ROA approximation \mathcal{R} , but also outside the actual, unknown ROA. To prevent this to happen, the safety monitor must work on a prediction $\hat{\mathbf{x}}(t+1)$ of the system state rather than on the actual state $\mathbf{x}(t)$. However, because of (i) the strict timing and resource constraints of the system, and (ii) the fact that \mathcal{R} is typically a coarse under-approximation of the ROA, this possibility was ignored, also because never occurred during our extensive experiments. Nevertheless, it should be taken into account in the design of such an architecture.

It is also worth remarking that all the ROA estimation methods might suffer from false negatives. Compared with false positives, false negatives are the most dangerous threat for the safety monitor. In fact, classifying an unsafe state as safe might lead to catastrophic failures. There is no closed-form algorithm to understand whether false positives exist. Simulation can however be used to sample the boundary of the estimated ROA to count the number of false positives (which might become unfeasible for highly-dimensional systems). However, the simulated model includes both parametric uncertainty and non-linear modeling inaccuracies with respect to the real-world version. Hence, in our specific case, the best option is to sample the estimated ROA boundary multiple times, each time varying the numerical parameters of the model. With a 5% variation and a grid size of 0.05 in the state space we found no false positives, which confirms that the

Table 3
Min and max times.

	Min [μ s]	Max [μ s]
Sensing	0.789	2.387
Reading hp controller output	1.694	3.179
Compute instability index	0.174	0.382
Finite state machine	0.113	3.124
Writing hp controller input	1.697	3.693
Linux2Erika end-to-end	288.579	3975.657

ellipsoid approximation is a coarse under-approximation of the actual ROA (although such a sampling cannot provide an absolute guarantee). If false positives were to be found, it might be useful to model a variable $\theta_U \in (0, 1)$ and define a parametric version of the estimated ROA: $\mathcal{R}(\theta_U) = \{\mathbf{x} \mid \mathbf{x}^T P \mathbf{x} < \theta_U\} \subset ROA$.

4.3. Results

To carry out experiments on this case study, the high-performance controller is implemented by a shallow neural network (one hidden layer, with tanh activation) taking the state \mathbf{x} as input and returning the scalar output u required to control the system. The hidden layer has $H = 128$ neurons, and the network is trained by the Covariance Matrix Adaptation (CMA-ES) (Auger and Hansen, 2012) evolutionary algorithm. The optimization process is performed in episodes, as in a standard reinforcement learning framework. An episode ends either after 1000 steps (with sample time 4 ms) or by being interrupted as soon as the pole angle α gets 20 degrees off the vertical or the arm angle θ drifts 45 degrees off the starting point. At each episode instant, the reward is $1 - 0.1|u|$, which is summed until termination of the episode, similarly to what happens in the CartPole OpenAI-gym (Brockman et al., 2016) environment. The term proportional to $|u|$ serves to enforce regularization. During training, random perturbations are injected to make the agents learn to withstand small pushes. The optimization is performed in a simulated environment, specifically a custom OpenAI-Gym environment that replicates the discrete dynamics of the Furuta pendulum.

CMA-ES considers the set of parameters of the network as the chromosomes to be evolved. Since the neural network has 4 inputs, the hidden layer has $4 \times H = 512$ weights and 128 biases, whereas the second layer consists of a single neuron with $H = 128$ weights and one bias. Hence, the total number of parameters is 769. The optimization is performed for 250 generations, with a population of 50 individuals, and a survival rate of 50%. To have a smoother control on the real system, the action is averaged in a moving window of two samples. Fig. 5 shows the results of the training process, where it is clear that the best agents are consistently performing well after just 100 generations.

Fig. 6 illustrates a situation in which the safety monitor decided to switch to the safe controller as a consequence of the injection of a disturbance (i.e., a push). A video¹ is also available online to view the actual behavior.

Table 3 shows the minimum and maximum time required for the execution of the different portions of the Control and Safety monitor tasks in 10,000 samples. The end-to-end latency of the control task is also showed (Linux2Erika end-to-end). This row in particular takes into account also the neural network inference time and the communication overhead between the Safe VM and Rich VM and back. It is possible to note that the maximum time required is less than the control task period (4 ms).

¹ <https://www.youtube.com/watch?v=pcrC1NZ12L0&t=191s>.

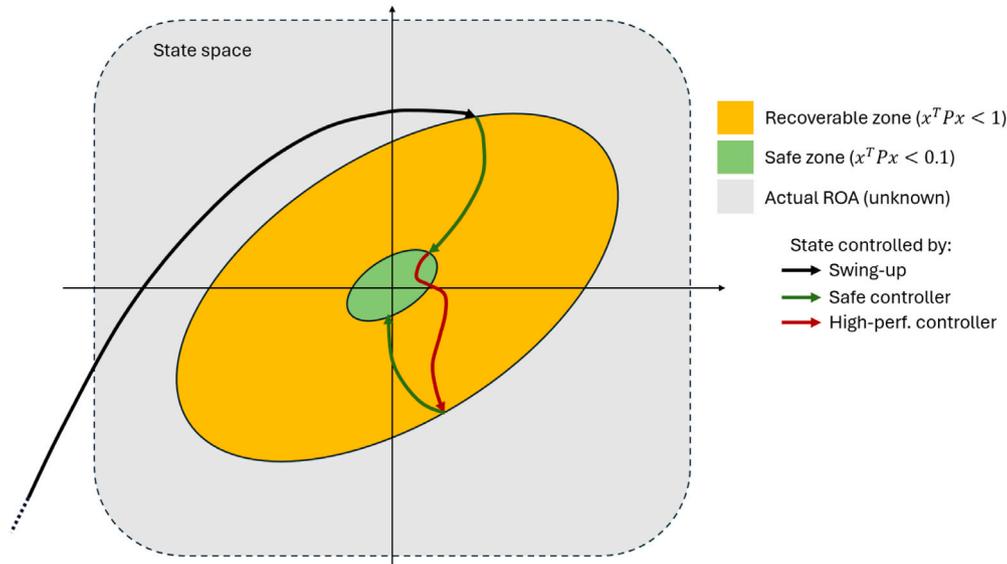


Fig. 4. Simplified 2D visualization of the state of the system and the various regions that define the behavior of the state machine. The actual region of asymptotic attraction (gray area) including all the states that are recoverable by the safe controller is actually unknown. At system startup, the swing-up controller (black arrow) brings the state close to the equilibrium. As soon as the state enters \mathcal{R} (yellow - states that can be surely recovered by the safe controller), the safe controller (green arrow) is activated to get the state close to the equilibrium and into S (green), where the high-performance controller can be activated safely. The state machine from this point switches between the high-performance and the safe controller depending on the instability index $x^T P x$.

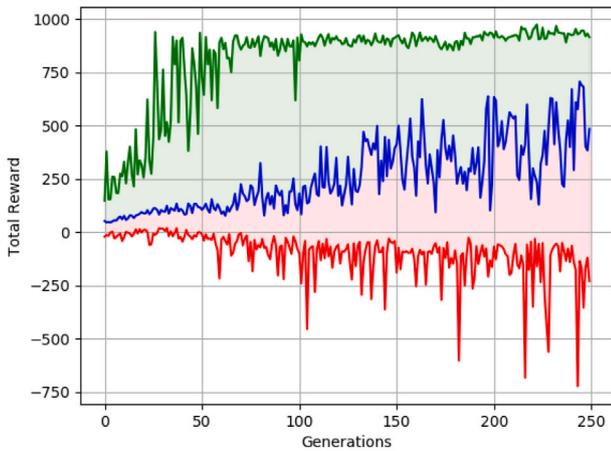


Fig. 5. Results of the training process for the population using CMA-ES. Green is the best for each generation, blue is the median, red is the worst; shaded areas include all the individuals in between.

5. Second case study: the AgileX Scout Mini rover

This section describes the implementation details of the architecture presented in Section 3 to control an AgileX Scout Mini rover (AgileX, 2026) for the task of LiDAR- and camera-based navigation and obstacle avoidance. Section 5.1 illustrates the rover and its features, Section 5.2 presents the specific architecture designed for this system, Section 5.3 specifies the controllers and the safety monitor used for the selected task, and Section 5.4 reports the results of the experiments performed on the physical system. The code implementation will be made available upon request.

5.1. The system

The AgileX Scout Mini rover is an agile and powerful platform, ideal for testing software for autonomous systems. Powered by a large

battery, the rover is equipped with both a LiDAR and a camera. Its design allows users to control it with built-in nodes that implement differential drive, making it behave like a unicycle.

The unicycle is a standard dynamical system whose state is described by $\mathbf{x} = (x, y, \theta)$, where (x, y) represents its position on the plane and θ its orientation. Being a non-linear system, its dynamics can be described as $\dot{\mathbf{x}} = f(\mathbf{x}, \mathbf{u})$. However, differently from the pendulum, where u is a single scalar, here $\mathbf{u} = (v, \omega)$ is a two-dimensional vector where v and ω are the linear and angular velocities (i.e., the yaw rate) of the rover, respectively.

The task carried out for this test case is navigation and obstacle avoidance, performed on a corridor using a LiDAR and a camera, as depicted in Fig. 7. Specifically, the rover starts at a random point of the corridor and must arrive at the end point without colliding with the corridor walls. Additionally, a few obstacles are placed throughout the route and a QR code is placed at the end of the corridor to signal the end of the course. The safety zones referred in the caption are introduced in Section 5.3.

5.2. Hardware and on-board software

We setup the rover with a Velodyne VLP-16 LiDAR, a Logitech 920 camera, and an AMD Kria KR260 board, equipped with an AMD Zynq UltraScale+ MPSoC. We chose the Kria KR260 board for two reasons:

- Compact footprint: mounting the board on the AgileX Scout MiniRover required a board with a limited size.
- Low power consumption: definitely lower compared to other alternatives such as Nvidia boards (Lizano and Westerlund, 2024).

The selected board includes an FPGA device that is not utilized in this implementation; however, it can be exploited for additional tasks, such as offloading and accelerating DNN inference. This flexibility makes the board an excellent choice for future extensibility, which will be investigated in future implementations.

Fig. 8 shows the hardware/software architecture of the system. As in the previous setup, platform virtualization is managed by CLARE. Since the application tasks present different criticality levels, two separate VMs running Ubuntu 22.04 LTS are utilized. They will be referred to

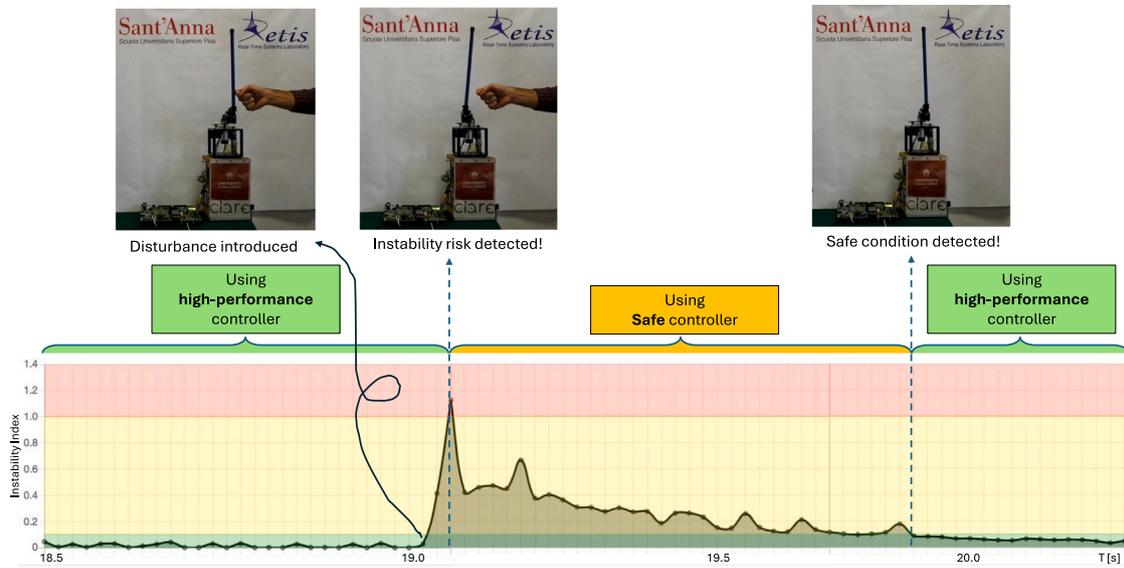


Fig. 6. A sample of the switching behavior of the overall architecture. When a strong disturbance is introduced, the high-performance controller brings the state of the system towards instability. This condition is detected by the safety monitor that promptly switches to the safe controller, which brings the state back to a safe neighborhood. The high-performance controller is then reactivated.

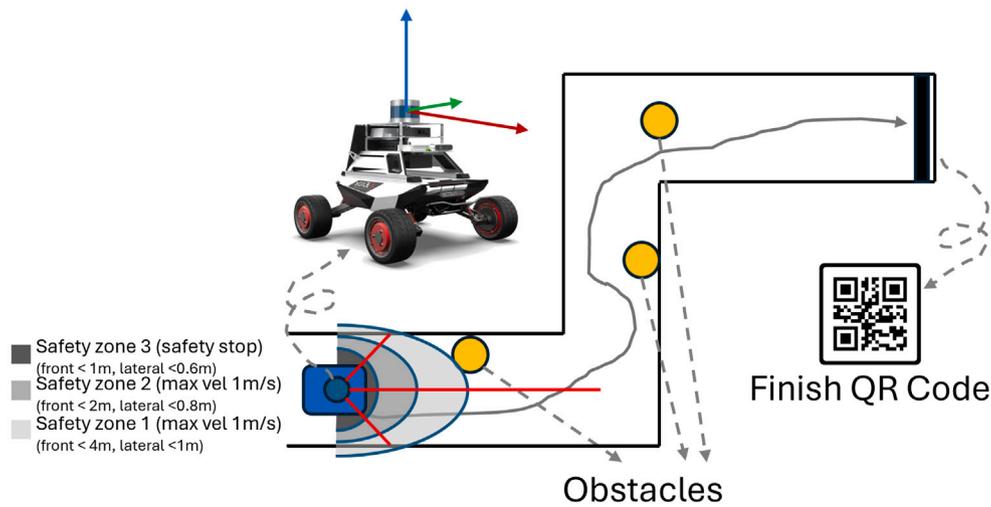


Fig. 7. The AgileX Scout Mini rover used for the second use case and a scheme of the corridor navigation task performed. The picture also illustrates the LiDAR reference frame on the rover, the three LiDAR distances that constitute the output of the LiDAR pre-processing algorithm, and the three safety zones defined by the safety monitor.

as `Safe_VM` and `Rich_VM`. The hypervisor is configured to provide static partitioning of the hardware resources among the two VMs. In particular, CPUs and memory-mapped areas (DRAM and device registers) are exclusively assigned to the VMs. This strategy minimizes resource contention and hypervisor intervention in virtualization of the resource, thus also reducing the power consumption.

Linux is chosen as the operating system for this dual-VM configuration due to its flexibility in supporting the execution of modern applications and for its wide support for third-party software, with particular reference to device drivers for cameras and LiDARs. Furthermore, Linux enables the execution of the ROS2 (Macenski et al., 2022) middleware, which is particularly useful to implement the tasks to be performed by the rover. ROS2 is used both as a communication infrastructure and to leverage pre-compiled sensing and actuation nodes controlling the velocity of the rover and collecting the point cloud data from the LiDAR. The Linux version running in the `Safe_VM` is restricted to a minimal configuration, stripped of Internet access

and all the redundant services that might interfere with safety-critical tasks. Conversely, the Linux version running in the `Rich_VM` is kept with a stock configuration, including Internet access to enable remote access to non-critical services, which is also useful for simulating a cyber-attack (see Section 5.4).

Both the `Safe_VM` or the `Rich_VM` are exclusively assigned one of the two Ethernet ports (`Eth0` and `Eth1`). The same applies for the two USB 3.0 ports (`USB0` and `USB1`). The responsibilities of the two domains, corresponding to the two VMs, are summarized below:

- Safe domain (`Safe_VM`):
 - Capture the 3D point cloud from the LiDAR device through the `Eth0` interface and preprocess it to obtain distance measures from the LiDAR.
 - Send the LiDAR distance measures to the `Rich_VM` and receive the corresponding neural network output.

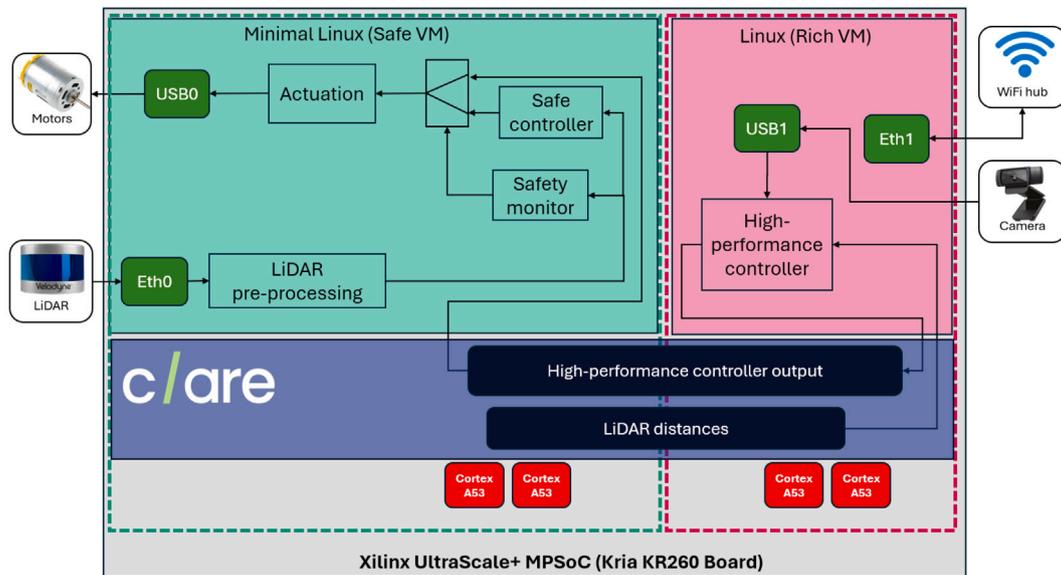


Fig. 8. The hardware/software architecture managed by the CLARE hypervisor for safe and secure control of the AgileX rover.

- Run the safety monitor, which enables actuation via either the safe or the high-performance controller output from the Rich_VM.
- Manage actuation through an USB2CAN interface using the USB0 interface connected to the CAN interface of the motors.
- Rich domain (Rich_VM):
 - Read the LiDAR distance measures from the Safe_VM.
 - Capture the image data from the Camera connected to the USB1 interface.
 - Execute the high-performance controller based on LiDAR distance measures and the RGB camera data and send the control to the Safe_VM.
 - Implement a wireless communication using the Eth1 interface.

The CLARE software stack provides a component, named CLARE-Middleware, that enables inter-domain communication using Cyclic Asynchronous Buffers (CABs) implemented in shared memory. A CAB is a data structure that operates using two pointers. The write pointer indicates where the data element will be written, while the read pointer shows where the next element will be read. Domains can write or read from these buffers according to the access control defined on the hypervisor configuration. To separate safe and unsafe data, the hypervisor is configured with two distinct buffers, referred to as Safe2Rich and Rich2Safe, to enable unidirectional communication between the domains. The Safe2Rich buffer is write-only for the Safe_VM and read-only for the Rich_VM, while the Rich2Safe buffer has the opposite access permissions.

The two domains run distinct instances of the ROS2 Humble framework. The tasks are therefore organized in nodes and exchange information via ROS2 topics. Fig. 9 illustrates how the nodes and topics interact. The `velodyne_node`, the `USB2CAN_node`, and the `camera_node` are pre-compiled and provided by the vendors. The `safe_node` implements all the functionality of the Safe_VM: LiDAR pre-processing, safety monitor, and actuation. Conversely, the `hp_controller_node` runs the high-performance controller, consisting of a neural network implemented in C++ and performs the

inference on the pre-processed the LiDAR data. The node also implements a QR code reader. The communication between the `safe_node` and the `hp_controller_node` is managed by CLARE-Middleware. The `safe_node` invokes the CLARE-Middleware API to write the pre-processed LiDAR data to the Safe2Rich buffer. The `hp_controller_node` polls the buffer for the data, processes it using the neural network, and sends the output to the `safe_node` via the Rich2Safe buffer. The `safe_node` polls the Rich2Safe buffer, executes the safety monitor, and applies the motor actuation (details in Section 5.3). The polling is supervised by a timer to guarantee that the task deadline is respected.

From a temporal standpoint, the LiDAR acquisition is the bottleneck, with a typical rate of 10 Hz. Hence, all the nodes are configured to run with a period of 100 ms.

Note that both the safe controller and the safety monitor are executed in the safe domain only and do not rely on camera inputs, but only on LiDAR data. In fact, images often suffer from several limitations, such as lighting conditions, perspective distortion, or other environmental factors. Furthermore, distance estimation from cameras, which is essential for obstacle detection and avoidance, is subject to higher errors with respect to LiDARs, even when considering stereo vision systems. Conversely, LiDAR data is much more structured and directly provide precise distance measurements. Furthermore, certified LiDAR technologies are also available and do not necessarily require learning-enabled components, for which it is hard to achieve a high level of assurance, to process the data they produce. For these reasons, camera data are processed by the high-performance controller only, while LiDAR data are processed by the safe one to enable trustworthy perception for safety-related purposes.

The raw LiDAR data consists of a 3D point cloud, whose size depends on the type of LiDAR and the number of channels (i.e., vertical beams) it generates. For instance, the Velodyne VLP-16 returns up to 300.000 points every 100 ms. However, since the specific obstacle avoidance application does not require such a massive amount of spatial data, the point cloud is first pre-processed to discard the points that are above the height of the rover (0.5 m). The selected points are then converted to distances from the LiDAR and grouped in bins depending on the angular position. The output of the LiDAR pre-processing is then a 3-dimensional vector d , where each element corresponds to the distance at -45 , 0 , and 45 degrees with respect to the x-axis of the LiDAR. This is because, in our case study, the rover is allowed to move forward only.

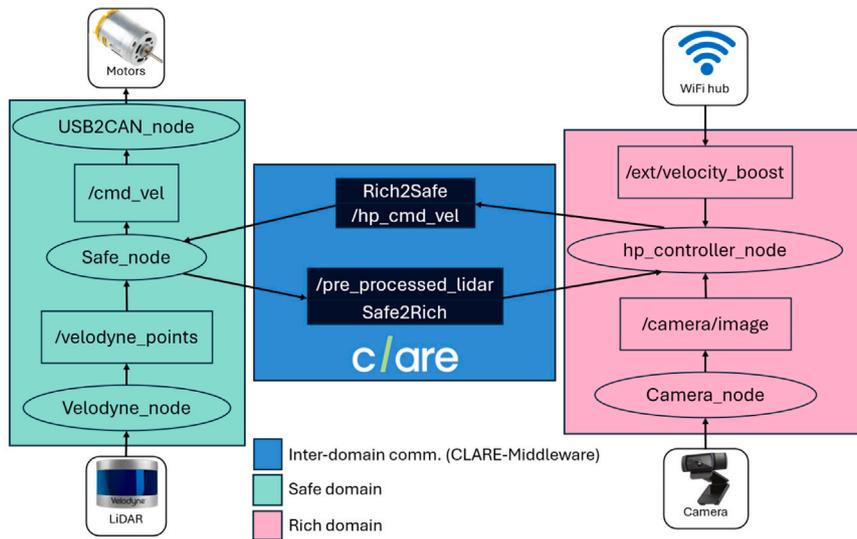


Fig. 9. Graph of ROS2 nodes and topics. The inter-domain communication is implemented with a custom transport.

5.3. Controllers and safety monitor

The safe controller is designed to limit the velocity of the rover to avoid collisions. The velocity limits are imposed according to the risk evaluated by the safety monitor. The safety monitor defines three different safety zones with increasingly stringent speed limits. Outside of these zones there is no risk of collision, and the output of the high-performance controller is used directly. The first (Z1), second (Z2), and third (Z3) safety zones are defined according to the distances returned by the front and lateral rays, as depicted in Fig. 7. In Z1 and Z2, the safe controller saturates the velocity command to 1.5 m/s and 1 m/s, respectively. These numerical values depend on the rover braking dynamics and are devised experimentally because of the high slip and slide motion given by the low friction between the rover wheels and the corridor pavement. The effectiveness of the zones is validated in the experimental subsection to avoid collisions when the rover is commanded to travel at its maximum velocity, which is about 3 m/s. Inside zone Z3, a safety stop is triggered by the safe controller and the human operator must acknowledge it before re-enabling the high-performance controller.

The high-performance controller is a 4-layer neural network. The input layer consists of an array of three elements encoding the distances coming from the LiDAR preprocessing d . The two hidden layers are fully connected, have 100 neurons each, and use a ReLU-activation function, while the output layer has two fully connected neurons using the \tanh activation; the two output neurons form a vector $A(d)$ whose elements represent the rover control variables (v, ω) . During regular operation, the output of the network is used directly, limiting the maximum velocity between -1 m/s and 1 m/s. However, we also included the possibility for a “simulated” attacker to boost the velocity of the high-performance controller by multiplying its output by a factor η defined in the topic `/ext/v_boost` depicted in Fig. 9. The network is trained as the actor of a deep RL agent, specifically using Deep Deterministic Policy Gradients (DDPG) (Lillicrap, 2015). All the hyperparameters are initialized with the typical values, as found in <https://github.com/ggliu/pytorch-ddpg>. The actor is trained on a custom environment that simulates LiDAR-based corridor navigation with randomly-generated paths of maximum length of 20 m. The reward at each step r is set to $r = D - 1 + E$, where D is the forward distance towards the goal traveled by the rover and E is the term added at the end of the episode defined as follows: $E = -100$ if a collision occurs (which terminates the episode right away); $E = 300$ if the rover reaches the goal (i.e., end of the corridor) before the maximum number of steps

in the episode (1000 steps, sampling time 100 ms); $E = 0$ otherwise. The -1 term is useful in the initial training phase to induce the rover to move forward rather than staying still to avoid collisions. Please note that the reward is not optimized for an optimal behavior of the rover, since the main contribution of the paper is the architectural configuration to enhance safety. The high-performance controller is seamlessly tested and integrated on the real rover with no need for fine-tuning. The camera is used to detect a specific QR code placed at the end of the corridor. When that QR code is detected, the rover stops and starts rotating on itself, signaling the completed course.

5.4. Results

The experiments are performed in the corridors of our laboratory. The selected course is about 20 m long with a right turn and a left turn similar to the one shown in Fig. 7. The QR code is printed on an A4 sheet and placed so that it is visible given the height of the camera. All the data from the experiments have been recorded as *rosbags* that capture the topics of interest.

Validation of the safety zones. The mechanism of the safety zones is the key concept that allows keeping the overall system safe, i.e., without collisions. The threshold of each zone must be carefully selected to account for the braking space required for a complete stop. Hence, it is crucial that the safety zones keep the rover away from colliding with obstacles also in the worst case, i.e., when the rover is traveling at its maximum velocity, which is about 3 m/s. The validation of the safety zones for the frontal distance is done by setting the high-performance controller to output a constant velocity close to the maximum velocity of the rover when facing a wall or an obstacle. If the safety zones are designed correctly, there should be no collision, since the progressively stringent safety zones will decrease the maximum velocity requested to the rover. Fig. 10 illustrates this process showing how a certain distance margin must be considered to take into account the rover braking space. This configuration is sufficient for our purposes.

The same considerations hold for the lateral thresholds as well. However, for the scope of this paper, they are intentionally kept smaller, since the corridor is about two meters wide and more conservative safety zones would have triggered many false-positive safety stops. Also, considering that the high-performance controller’s neural network has a \tanh activation, its output will always be between -1 and 1 m/s, which is slow enough for the system to be able to avoid collision with these safety zones.

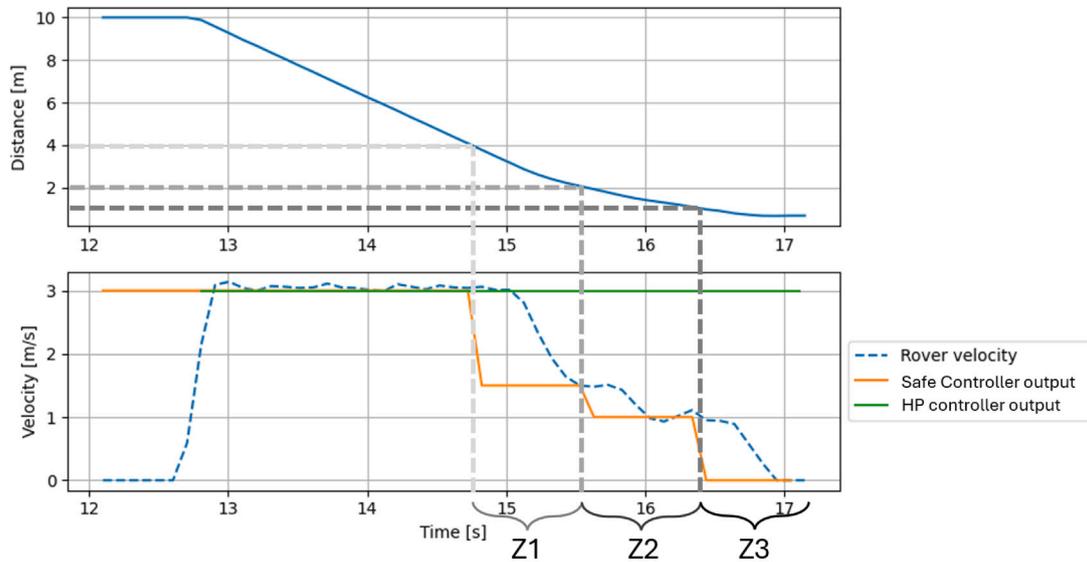


Fig. 10. Validation of the safety zones. The graph above shows the distance to the front wall when the rover is controlled at the maximum velocity by the high-performance controller. The graph below illustrates how the requested velocity is regulated according to the safety zones. The rover measured velocity is also showed (dashed blue) to visualize the braking dynamics that clearly requires a certain margin to avoid collision.

With this setup, false positives and false negatives primarily depend on the accuracy and precision of the LiDAR sensor. A false positive may be triggered by a spurious measurement detected too close to the rover, causing a safety stop. While this degrades performance, it does not compromise safety. Conversely, a false negative would require the LiDAR points corresponding to an obstacle to be erroneously estimated as farther than the safety threshold, a condition that would typically imply sensor spoofing or a malfunction rather than a transient measurement error. Although no false positive nor false negatives were recorded during our experiments, it is wise to include a fault detection module as a safety monitor for sensors, as this is a common practice in safety-critical systems. Since the focus of this paper is not on detecting sensor faults, this aspect is not further addressed in this work.

Sample run. Fig. 11 shows the output of the high-performance controller, which controls the velocity and the position of the rover in the corridor. In this run, no safety stop is triggered and the QR code is reached seamlessly. The safety zones are not visualized, as the velocity limitations are not affecting the high-performance controller: since the maximum velocity that the neural network can output with no external velocity boost is 1 m/s in standard condition, even when in Z2 the safe controller cannot limit the velocity unless a safety stop is triggered.

Since the system consistently presents a good behavior in normal conditions (including the avoidance of small obstacles placed close to the walls), we introduce two different dangerous situations to stress the systems in terms of safety and security: (i) sudden obstacles, and (ii) a simulated attack that modifies the output of the network.

Sudden obstacles. This experiment is meant to validate the behavior of the rover in the presence of obstacles that might suddenly appear and could not be avoided by going around. Specifically, a cardboard box is thrown in front of the rover while navigating the corridor.

Fig. 12 shows the frontal distance extracted from the LiDAR and the corresponding actuation command. As soon as the obstacle is detected, a safety stop is triggered within the 100 ms LiDAR acquisition period.

Simulated attack. The Safe_VM does not have internet access and cannot be interfered with. Conversely, the Rich_VM can access internet for external visualization. This channel can be used to simulate a possible cyber-attack on the Rich_VM. This test shows that the isolation property enforced by the hypervisor on the two domains and

the safety zones mechanism prevent (or at least limit) unsafe behaviors from the high-performance controller.

Specifically, the high-performance controller is subscribed to a float32 topic that acts as a multiplicative factor for the output of the neural network. In this way, when an external user publishes a factor η in the topic, the output of the neural network will no longer be limited between -1 and 1 , but between $-\eta$ and η . Although this kind of attack might not be a proper cyber-attack, it is useful to illustrate how the safe controller takes care of unsafe velocity inputs from the high-performance controller.

For this experiment, the rover is tasked to complete the same course with $\eta = 2$. The results reported in Fig. 13 show that the rover is able to reach the QR code without safety stops. However, since now the maximum velocity from the high-performance controller is 2 m/s, the safety zones are entered and hence the actual velocity command is limited accordingly. The safety zones are triggered during the left and right curves of the corridor; also, on the straight part, the rover shows a much more “wavy” behavior that triggers Z1 (and Z2 for a brief moment) because the walls were getting too close. This behavior of the neural network is likely arising because the agent has never traveled at such high velocities, while at low velocities the wave pattern is less visible.

Temporal analysis. Fig. 14(a) shows the distribution of 10.000 measurements of the overhead introduced by the round-trip communication of a 64-byte package (which is more than enough for our purposes) between the safe and the rich domain and back. From the histogram it is clear that the communication between the two domains introduces a negligible overhead, with more than 99% of the data passed in 10 μ s or less.

Fig. 14(b) shows the distribution of execution times for the safe_node (3000 samples). This includes the time required to send LiDAR distances to the Rich_VM, perform inference using the high-performance controller’s neural network, and receive its output. From this distribution, we can infer that this implementation consistently meets the safe domain deadline of 100 ms.

6. Discussion, lessons learned, and conclusions

Discussion. This paper presented an architecture for CPSs designed to enhance safety, security, and predictability levels of systems controlled with learning-based components. Specifically, the main concept

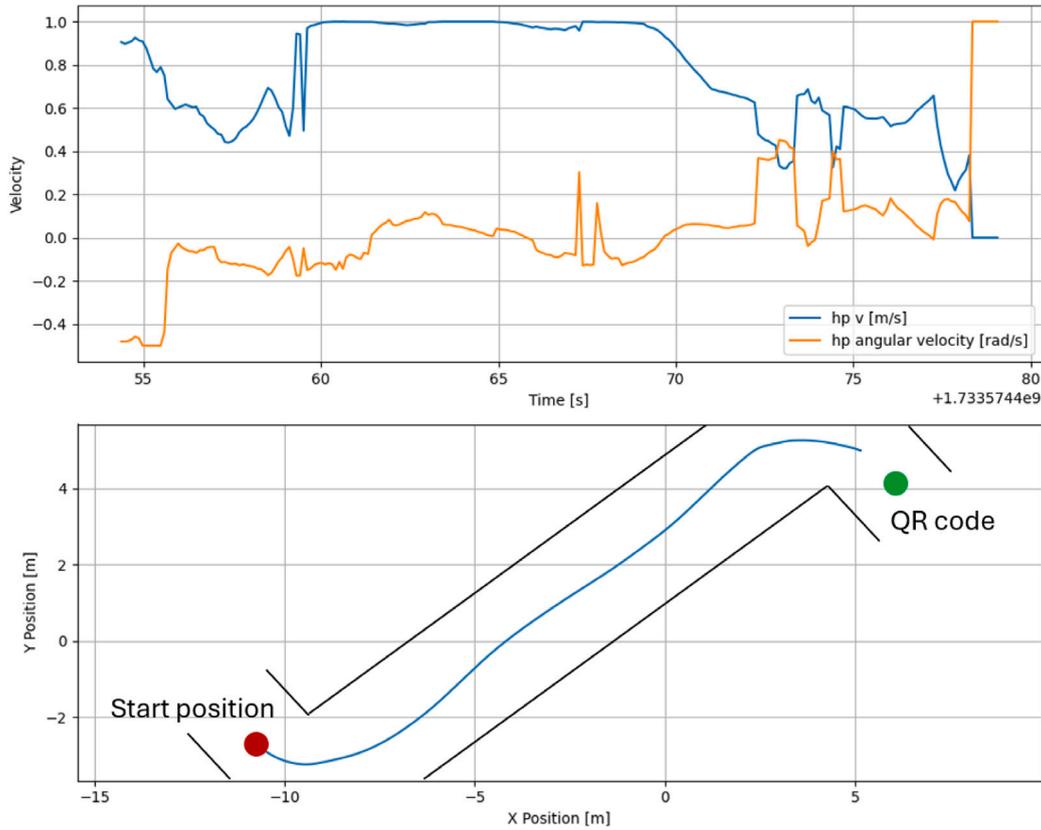


Fig. 11. Velocity command from the high-performance controller (above) and position in the corridor (below) during a sample run with no safety stops triggered. The rover reaches the QR code without issues. When the QR code is detected, the rover is spun on itself at 1 rad/s.

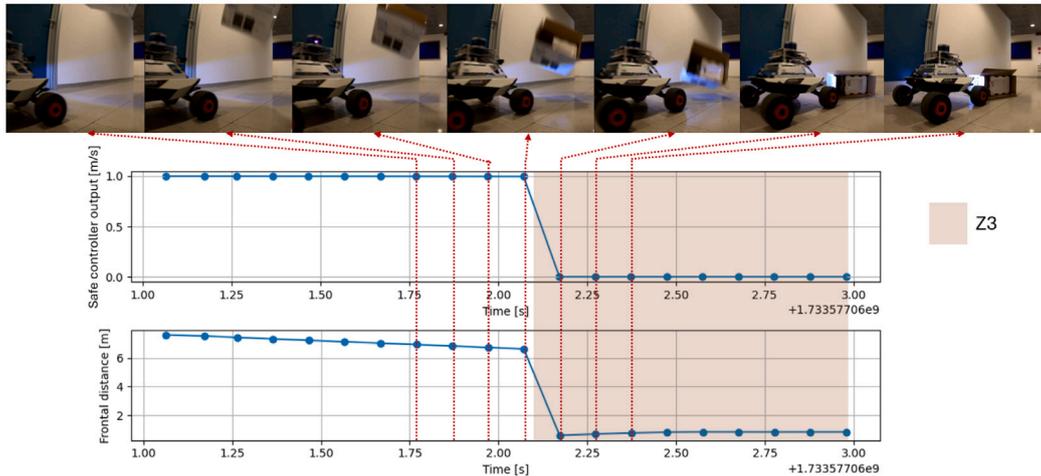


Fig. 12. Velocity command (above) and frontal distance to the closest obstacle (below) during a run with a sudden obstacle appearance. The rover reacts within its 100 ms period to trigger a safety stop.

behind such an architecture is to separate the AI components from the safety-critical ones by using advanced virtualization technologies (provided by the CLARE hypervisor). In fact, AI components are always deemed untrustworthy, and a Simplex-like switching mechanism is used to exclude the learning-based component upon a fault detection and replace it with a simpler but more reliable module.

The paper also presented two fully functional implementations of the proposed architecture: an inverted pendulum, used to highlight

the main features of the architecture, and a rover obstacle avoidance application, representing a much more complex and complete use-case that showcases the main advantages of the proposed approach. The proposed architecture, however, is general enough to be adapted to other common CPS applications, such as drones, robots, and even self-driving cars. Table 4 summarizes the two different hardware/software setups, highlighting how different applications can be implemented with the same architectural approach.

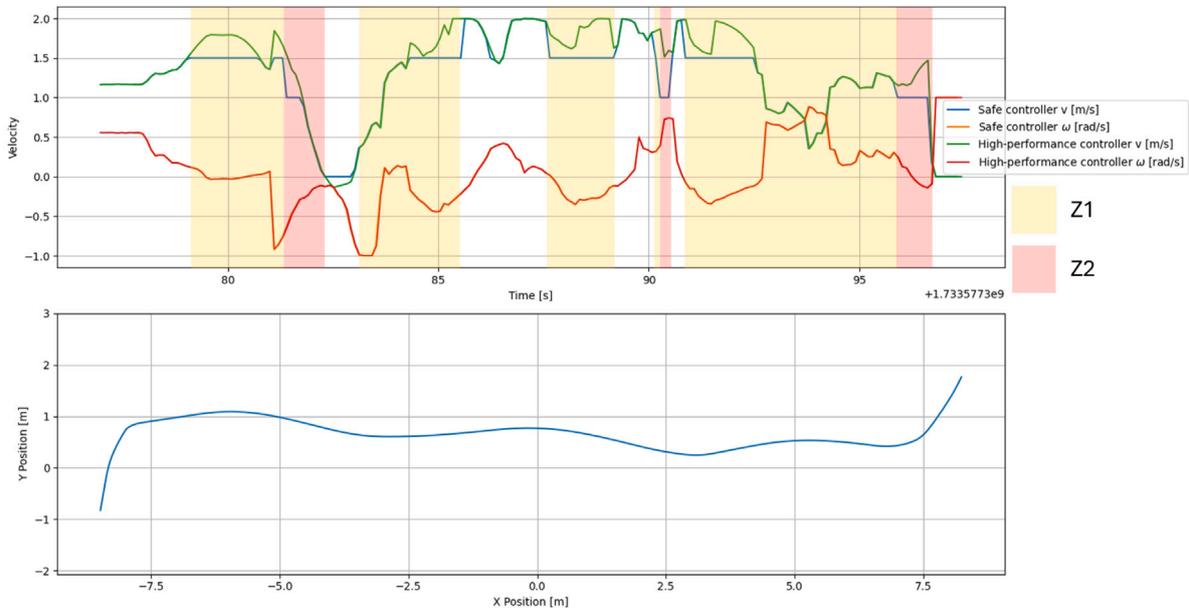


Fig. 13. Velocity command from the high-performance controller and actual actuation value (above) and position in the corridor (below) during a run with $\eta = 2$. The rover reaches the QR code, but there are a few situations where the safety zones limit the high-performance controller command, but never triggering a safety stop.

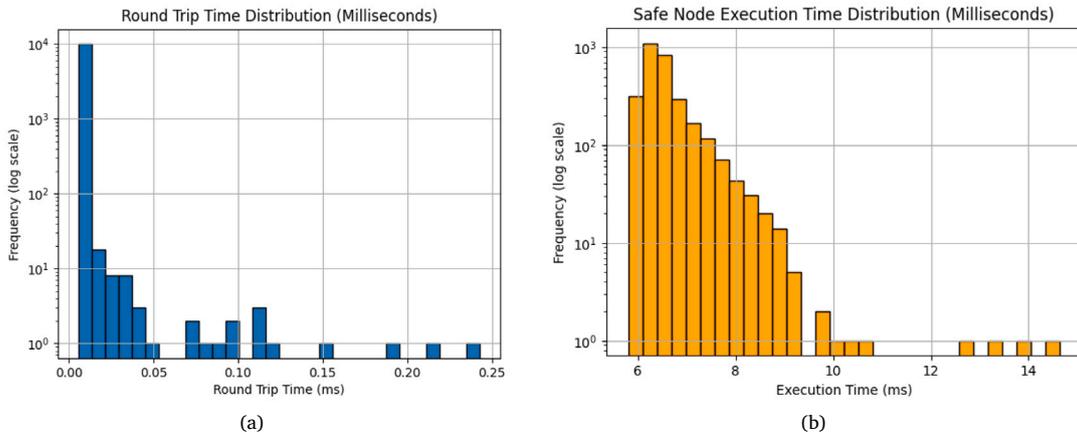


Fig. 14. (a) Round-trip communication times between the two domains for a 64-byte package (10.000 runs). (b) Distribution of the execution times of the `safe_node` (3000 runs).

Table 4

Comparison between the Quanser Furuta pendulum and AgileX Scout Mini rover case studies.

Feature	Quanser Furuta pendulum	AgileX Scout Mini Rover
Platform	Xilinx UltraScale+ (ZCU102)	Xilinx UltraScale+ (Kria KR260)
OS configuration	Rich VM: Petalinux 2018.2 (with Caffee) Safe VM: Erika3 RTOS	Rich VM: Ubuntu 22.04 Safe VM: Ubuntu 22.04 (minimal config)
CPU allocation	Rich VM: 3 CPUs Safe VM: 1 CPU	Rich VM: 2 CPUs Safe VM: 2 CPUs
Sensors	2 encoders, 16 kHz via FPGA	LiDAR (VLP-16, 10 Hz), RGB camera (Logitech 920, 10 FPS)
Safety monitor challenges	Fast and highly non-linear dynamics, strict real-time, ROA estimation	Obstacle avoidance, tuning of the safety zones.
Task	Swing-up and stabilization	Corridor navigation and obstacle avoidance with QR code detection
Control	Safe: linear controller; Rich: RL-based controller	Safe: LiDAR-based velocity limiter + safety zones; Rich: RL-based controller on LiDAR inputs
Temporal constraints	4 ms period	100 ms period (LiDAR acquisition bottleneck)

Limitations of the approach. Despite the advantages discussed throughout the paper, the proposed architecture introduces non-negligible complexity both in the hardware and software design. Specifically, the introduction of hypervisor technology requires specific know-how, a dedicated configuration, and programming tools; moreover, the use of different operating systems on the same platform further constrains the resources available to both domains.

While for the pendulum the configuration is rather simple and device assignment did not cause any issue, the rover application was much more complex. In fact, each sensor had to be connected through a specific port (Ethernet for the LiDAR, USB for the motors) and, since these ports were also used by the rich domain for camera and internet access, these devices had to be partitioned between the two domains, complicating both the design and testing/debugging phases. Moreover, for the pendulum, a single CPU core was enough to execute the safe domain tasks, allowing the use of three cores for the high-performance controller; for the rover, the safe domain was much more complex, requiring two cores to run ROS2 and all the assigned tasks on Linux. Device assignment required an iterative process to properly dimension the system. Additional complexity is required to manage the inter-domain dedicated communication, which also introduces a certain latency. Such a latency is practically negligible in the case of the rover (for which the deadline is 100 ms), but may take up a significant fraction of the period for high-rate control systems such as the pendulum (4 ms). For this reason, the periods of these tasks also required to be carefully engineered. The use of ROS2 is another important addition in the case of the rover with respect to the case of the pendulum. ROS2 allows using pre-compiled packages to read the sensors and control the motion of the rover; on the other hand, it introduces non-negligible complexity, and, sometimes, unexpected behavior, probably due to the intricate code-base of the middleware. This was found to be especially complicated when dealing with the custom transport between the safe and rich domains, which required several iterations to work properly and with low latency. The case of the pendulum was simpler in this sense, while it required custom-designed FPGA logic to sample the encoders. Another point worth discussing is that controlling a system with two different, possibly alternating controllers requires specific precautions such as the introduction of a finite state machine that keeps the system in a safe state.

From a functional point of view, the safety monitors designed for the two cases are based on different concepts and are selected according to the definition of safety for the task at hand: the pendulum is monitored with a model-based technique built on a strong theoretical support, while the rover is monitored using LiDAR data that define a safety envelope to avoid crashes. It is worth noting how the Simplex architecture is adaptable to different safety monitoring methods, which might also include additional logic conditions, such as cyber-security warning systems that detect whether the rich domain is attacked. All these downsides might slow down the design and development process, but the final tuned architecture significantly enhances the safety and security levels of autonomous CPSs, as demonstrated by our experiments.

Future Work. Future work will investigate several different directions, including:

- **The role of heterogeneous hardware for neural network acceleration.** FPGA is a solid alternative to the GPU for accelerating deep learning models. For both case studies, shallow networks were sufficient to showcase the proposed architecture, hence they were executed on the CPU without temporal issues. However, more complex visual tasks rely on deeper neural networks that require acceleration to work in real time. For this reason, future work will consider FPGA-accelerated DNNs in the implementation (Bao et al., 2025; Chen and Yu, 2025).

- **Implementing vehicle-in-the-loop simulation.** Autonomous vehicles require extensive testing. However, it is often impossible to cover all the possible critical scenarios to validate the entire software stack. Vehicle-in-the-loop simulation tackles this problem by creating a digital twin that replicates the movements of the physical vehicle in a virtual world. The synthetic sensor data can be sent in real-time to the rich domain, which would act in the physical world based on simulated sensory inputs.
- **Evaluating the robustness of visual perception models.** Since DNNs are prone to adversarial attacks, the proposed architecture can be used to evaluate the robustness of the models to physical adversarial attacks. Also, detection algorithms can be used as triggers for the safety monitor.
- **Benchmarking the real-time and power properties of the system.** All the previous points (especially the vehicle-in-the-loop simulation) will require a careful analysis of the latencies introduced by the isolation. Including larger (accelerated) models will certainly bring more complexity into the design, which should always be backed by a comprehensive timing analysis. Furthermore, the power consumption overhead introduced by this approach will be investigated.
- **Adding novel recovery and security features.** The hypervisor can be used in conjunction with an anomaly detection system (Borgioli et al., 2023) deployed within the VMs to identify anomalies and restore the affected VM from a clean backup.

CRediT authorship contribution statement

Federico Nesti: Writing – original draft, Software, Methodology, Data curation, Conceptualization. **Niko Salamini:** Software. **Mauro Marinoni:** Software, Methodology. **Giorgiomaria Cicero:** Software, Conceptualization. **Gabriele Serra:** Software. **Alessandro Biondi:** Writing – review & editing, Supervision. **Giorgio Buttazzo:** Writing – review & editing, Supervision.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Federico Nesti reports financial support was provided by Fondazione SERICS. Giorgiomaria Cicero reports a relationship with Accelerat srl that includes: board membership and equity or stocks. Gabriele Serra reports a relationship with Accelerat srl that includes: employment and equity or stocks. Alessandro Biondi reports a relationship with Accelerat srl that includes: board membership and equity or stocks. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work has been partially supported by project SERICS (PE00000014) under the NRRP MUR program funded by the EU - NGEU and project RETICULATE of the PRIN 2022 program.

Data availability

Data will be made available on request.

References

- Åström, K., Furuta, K., 2000. Swinging up a pendulum by energy control. *Automatica* 36 (2), 287–295. [http://dx.doi.org/10.1016/S0005-1098\(99\)00140-5](http://dx.doi.org/10.1016/S0005-1098(99)00140-5), URL <http://www.sciencedirect.com/science/article/pii/S0005109899001405>.
- Abdar, M., Pourpanah, F., Hussain, S., Rezazadegan, D., Liu, L., Ghavamzadeh, M., Fieguth, P., Cao, X., Khosravi, A., Acharya, U.R., et al., 2021. A review of uncertainty quantification in deep learning: Techniques, applications and challenges. *Inf. Fusion* 76, 243–297.
- Accelerat Srl, 2026. The CLARE software stack, URL <https://accelerat.eu/clare>.
- AgileX, 2026. AgileX Scout Mini, URL <https://global.agilex.ai/products/scout-mini>.
- Almeida, J., Prochazka, M., 2009. Safe and secure partitioning with PikeOS: towards integrated modular avionics in space. In: *DASIA 2009-Data Systems in Aerospace*. Vol. 669, p. 27.
- Auger, A., Hansen, N., 2012. Tutorial CMA-ES: evolution strategies and covariance matrix adaptation. In: *Proceedings of the 14th Annual Conference Companion on Genetic and Evolutionary Computation*. pp. 827–848.
- Bak, S., Chivukula, D.K., Adekunle, O., Sun, M., Caccamo, M., Sha, L., 2009. The system-level simplex architecture for improved real-time embedded system safety. In: *Proceedings of the 15th IEEE Real-Time and Embedded Technology and Applications Symposium. RTAS'09, San Francisco, CA, USA*.
- Bak, S., Johnson, T.T., Caccamo, M., Sha, L., 2014. Real-time reachability for verified simplex design. In: *2014 IEEE Real-Time Systems Symposium*. pp. 138–148. <http://dx.doi.org/10.1109/RTSS.2014.21>.
- Bak, S., Manamcheri, K., Mitra, S., Caccamo, M., 2011. Sandboxing controllers for cyber-physical systems. In: *2011 IEEE/ACM Second International Conference on Cyber-Physical Systems*. pp. 3–12. <http://dx.doi.org/10.1109/ICCP.2011.25>.
- Bao, Z., Li, H., Zhang, W., 2025. A vision transformer inference accelerator for KR260. In: *Proceedings of the 2024 8th International Conference on Computer Science and Artificial Intelligence. CSAI '24, Association for Computing Machinery, New York, NY, USA*, pp. 245–251. <http://dx.doi.org/10.1145/3709026.3709107>.
- Biggio, B., Roli, F., 2018. Wild patterns: Ten years after the rise of adversarial machine learning. *Pattern Recognit.* 84, 317–331.
- Biondi, A., Casini, D., Cicero, G., Borgioli, N., Buttazzo, G., Patti, G., Leonardi, L., Bello, L.L., Solieri, M., Burgio, P., et al., 2021. SPHERE: A multi-SoC architecture for next-generation cyber-physical systems based on heterogeneous platforms. *IEEE Access* 9, 75446–75459.
- Biondi, A., Nesti, F., Cicero, G., Casini, D., Buttazzo, G., 2020. A safe, secure, and predictable software architecture for deep learning in safety-critical systems. *IEEE Embed. Syst. Lett.* 12 (3), 78–82.
- Borgioli, N., Thi Xuan Phan, L., Aromolo, F., Biondi, A., Buttazzo, G., 2023. Real-time packet-based intrusion detection on edge devices. In: *Proceedings of Cyber-Physical Systems and Internet of Things Week 2023*. In: *CPS-IoT Week '23, Association for Computing Machinery, New York, NY, USA*, pp. 234–240. <http://dx.doi.org/10.1145/3576914.3587551>.
- Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., Zaremba, W., 2016. OpenAI Gym. <http://dx.doi.org/10.48550/arXiv.1606.01540>, ArXiv E-Prints, arXiv:1606.01540, arXiv:1606.01540.
- Buttazzo, G., 2022. Can we trust AI-powered real-time embedded systems? In: *OpenAccess Series in Informatics (OASIS)*, Vol. 98, Proc. of the HIPEAC Workshop on Next Generation Real-Time Embedded Systems. NG-RES 2022, Budapest, Hungary.
- Casini, D., Biondi, A., Buttazzo, G., 2020. Timing Isolation and Improved Scheduling of Deep Neural Networks for Real-Time Systems. *Softw.: Pract. Exp.* 50 (9), 1760–1777.
- Cavicchioli, R., Capodieci, N., Bertogna, M., 2017. Memory Interference Characterization Between CPU Cores and Integrated GPUs in Mixed-Criticality Platforms. In: *Proc. of the 22nd IEEE International Conference on Emerging Technologies and Factory Automation. ETFA 2017, Limassol, Cyprus*.
- Cazzolato, B.S., Prime, Z., 2011. On the dynamics of the furuta pendulum. *J. Control Sci. Eng.* 2011 (1), 528341.
- Chen, Y., Yu, L., 2025. FPGA-based MCS-YOLOv5 lightweight ship detection model-tion on KR260. In: *2025 5th International Conference on Artificial Intelligence and Industrial Technology Applications. AIITA*, pp. 1764–1768. <http://dx.doi.org/10.1109/AIITA65135.2025.11047633>.
- Cittadini, E., Marinoni, M., Biondi, A., Cicero, G., Buttazzo, G., 2023. Supporting AI-powered real-time cyber-physical systems on heterogeneous platforms via hypervisor technology. *Real-Time Syst.* 59 (4), 609–635.
- Craveiro, J., Rufino, J., Schoofs, T., Windsor, J., 2009. Flexible operating system integration in partitioned aerospace systems. In: *Actas do INForum-Simposio de Informatica*. pp. 49–60.
- Desai, A., Ghosh, S., Seshia, S.A., Shankar, N., Tiwari, A., 2019. SOTER: a runtime assurance framework for programming safe robotics systems. In: *2019 49th Annual IEEE/IFIP International Conference on Dependable Systems and Networks. DSN, IEEE*, pp. 138–150.
- Evidence Srl, Erika enterprise RTOS, URL <https://www.erika-enterprise.com/>.
- Farrukh, A., West, R., 2022. FLYOS: Integrated modular avionics for autonomous multicopters. In: *2022 IEEE 28th Real-Time and Embedded Technology and Applications Symposium. RTAS, IEEE*, pp. 68–81.
- Ferreira, E., Krogh, B., 1997. Using neural networks to estimate regions of stability. In: *Proceedings of 1997 American Control Conference*. Vol. 3, pp. 1989–1993.
- Ferreira, E., Krogh, B., 1999. Switching controllers based on neural network estimates of stability regions and controller performance. http://dx.doi.org/10.1007/3-540-64358-3_36.
- Gutiérrez, C.S.V., Juan, L.U.S., Ugarte, I.Z., Vilches, V.M., 2018. Towards a distributed and real-time framework for robots: Evaluation of ROS 2.0 communications for real-time robotic applications. arXiv preprint arXiv:1809.02595.
- Huang, X., Kroening, D., Ruan, W., Sharp, J., Sun, Y., Thamo, E., Wu, M., Yi, X., 2020. A survey of safety and trustworthiness of deep neural networks: Verification, testing, adversarial attack and defence, and interpretability. *Comput. Sci. Rev.* 37, 100270.
- Klein, G., Andronick, J., Fernandez, M., Kuz, I., Murray, T., Heiser, G., 2018. Formally verified software in the real world. *Commun. ACM* 61 (10), 68–77.
- Kurakin, A., Goodfellow, I.J., Bengio, S., 2017. Adversarial examples in the physical world. In: *Proc. of the 5th Int. Conference on Learning Representation. ICLR, Toulon, France*.
- Lillicrap, T., 2015. Continuous control with deep reinforcement learning. arXiv preprint arXiv:1509.02971.
- Liu, M., Niu, J., Wang, X., 2017. An autopilot system based on ROS distributed architecture and deep learning. In: *2017 IEEE 15th International Conference on Industrial Informatics. INDIN, IEEE*, pp. 1229–1234.
- Lizano, S.A., Westerlund, T., 2024. Comparison of edge computing platforms for hardware acceleration of AI: Kria KV260, Jetson Nano and RTX 3060. *J. Edge Comput.* 15 (3), 123–135.
- Macenski, S., Foote, T., Gerkey, B., Lalancette, C., Woodall, W., 2022. Robot operating system 2: Design, architecture, and uses in the wild. *Sci. Robot.* 7 (66), eabm6074. <http://dx.doi.org/10.1126/scirobotics.abm6074>, URL <https://www.science.org/doi/abs/10.1126/scirobotics.abm6074>.
- Masana, M., Liu, X., Twardowski, B., Menta, M., Bagdanov, A.D., Van De Weijer, J., 2022. Class-incremental learning: survey and performance evaluation on image classification. *IEEE Trans. Pattern Anal. Mach. Intell.* 45 (5), 5513–5533.
- Meier, L., Honegger, D., Pollefeys, M., 2015. PX4: A node-based multithreaded open source robotics framework for deeply embedded platforms. In: *2015 IEEE International Conference on Robotics and Automation. ICRA, IEEE*, pp. 6235–6240.
- Modica, P., Biondi, A., Buttazzo, G., Patel, A., 2018. Supporting Temporal and Spatial Isolation in a Hypervisor for ARM Multicore Platforms. In: *Proceedings of the 18th IEEE International Conference on Industrial Technology. ICIT 2018, Lyon, France*.
- Mohan, S., Bak, S., Betti, E., Yun, H., Sha, L., Caccamo, M., 2013a. S3A: Secure System Simplex Architecture for Enhanced Security and Robustness of Cyber-Physical Systems. In: *Proceedings of the ACM Int. Conference on High Confidence Networked Systems. HiCoNS, Philadelphia, PA, USA*.
- Mohan, S., Bak, S., Betti, E., Yun, H., Sha, L., Caccamo, M., 2013b. S3A: Secure system simplex architecture for enhanced security and robustness of cyber-physical systems. In: *Proceedings of the 2nd ACM International Conference on High Confidence Networked Systems*. pp. 65–74.
- Nesti, F., Rossolini, G., Nair, S., Biondi, A., Buttazzo, G., 2022. Evaluating the robustness of semantic segmentation for autonomous driving against real-world adversarial patch attacks. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. WACV*, pp. 2280–2289.
- O'Kelly, M., Sukhil, V., Abbas, H., Harkins, J., Kao, C., Pant, Y.V., Mangharam, R., Agarwal, D., Behl, M., Burgio, P., et al., 2019. F1/10: an open-source autonomous cyber-physical platform. arXiv preprint arXiv:1901.08567.
- Olowononi, F.O., Rawat, D.B., Liu, C., 2020. Resilient machine learning for networked cyber physical systems: A survey for machine learning security to securing machine learning for CPS. *IEEE Commun. Surv. Tutor.* 23 (1), 524–552.
- Pereira, A., Thomas, C., 2020. Challenges of machine learning applied to safety-critical cyber-physical systems. *Mach. Learn. Knowl. Extr.* 2 (4), 579–602.
- Perez, H., Gutiérrez, J.J., 2017. Handling heterogeneous partitioned systems through ARINC-653 and DDS. *Comput. Stand. Interfaces* 50, 258–268.
- Phan, D.T., Grosu, R., Jansen, N., Paoletti, N., Smolka, S.A., Stoller, S.D., 2020. Neural simplex architecture. In: *NASA Formal Methods: 12th International Symposium, NFM 2020, Moffett Field, CA, USA, May 11–15, 2020, Proceedings*. vol. 12, Springer, pp. 97–114.
- Quanser, 2026. Quanser rotary inverted pendulum, URL <https://www.quanser.com/products/rotary-inverted-pendulum/>.
- Quiñonero-Candela, J., Sugiyama, M., Schwaighofer, A., Lawrence, N.D., 2022. *Dataset Shift in Machine Learning*. MIT Press.
- Reke, M., Peter, D., Schulte-Tiggas, J., Schiffer, S., Ferreira, A., Walter, T., Mathes, D., 2020. A self-driving car architecture in ROS2. In: *2020 International SAUPEC/RobMech/PRASA Conference. IEEE*, pp. 1–6.
- Rossolini, G., Nesti, F., Brau, F., Biondi, A., Buttazzo, G., 2023. Defending from physically-realizable adversarial attacks through internal over-activation analysis. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 37, pp. 15064–15072.
- Rybczak, M., Popowniak, N., Lazarowska, A., 2024. A survey of machine learning approaches for mobile robot control. *Robotics* 13 (1), 12.

- Seto, D., Krogh, B., Sha, L., Chutinan, A., 1998. The simplex architecture for safe online control system upgrades. In: Proceedings of the 1998 American Control Conference. ACC (IEEE Cat. No.98CH36207). Vol. 6, pp. 3504–3508 vol.6. <http://dx.doi.org/10.1109/ACC.1998.703255>.
- Seto, D., Sha, L., 1999. A case study on analytical analysis of the inverted pendulum real-time control system.
- Sha, L., Rajkumar, R., Gagliardi, M., 1994. The simplex architecture: An approach to build evolving industrial computing systems. In: Proceedings of the ISSAT Int. Conference: Reliability and Quality in Design. Seattle, Washington, USA.
- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., Fergus, R., 2014. Intriguing properties of neural networks. In: Proc. of the 2nd International Conference on Learning Representations. ICLR 2014, Banff, AB, Canada.
- Vivekanandan, K., et al., 2016. A simplex architecture for intelligent and safe unmanned aerial vehicles. In: IEEE International Conference on Embedded and Real-Time Computing Systems and Applications.
- Xilinx, 2026. Zynq UltraScale+ device technical reference manual-UG1085, URL <https://docs.amd.com/r/en-US/ug1085-zynq-ultrascale-trm/Zynq-UltraScale-Device-Technical-Reference-Manual>.
- Yang, J., Zhou, K., Li, Y., Liu, Z., 2024. Generalized out-of-distribution detection: A survey. *Int. J. Comput. Vis.* 1–28.
- Zhu, F., Antsaklis, P.J., 2015. Optimal control of hybrid switched systems: A brief survey. *Discrete Event Dyn. Syst.* 25, 345–364.
- Zou, Z., Chen, K., Shi, Z., Guo, Y., Ye, J., 2023. Object detection in 20 years: A survey. *Proc. IEEE* 111 (3), 257–276.
- Zubov, V.I., 1965. Methods of A. M. Lyapunov and their application. *Am. Math. Mon.*