



# Continuous Integration and Delivery Practices for Cyber-Physical Systems: An Interview-Based Study

IORELLA ZAMPETTI, University of Sannio, Italy

DAMIAN A. TAMBURRI, Eindhoven University of Technology and the Jheronimus Academy of Data Science (JADS), The Netherlands

SEBASTIANO PANICHELLA, Zurich University of Applied Sciences, Switzerland

ANNIBALE PANICHELLA, Technical University of Delft, The Netherlands

GERARDO CANFORA, University of Sannio, Italy

MASSIMILIANO DI PENTA, University of Sannio, Italy

Continuous Integration and Delivery (CI/CD) practices have shown several benefits for software development and operations, *e.g.*, faster release cycles and early discovery of defects. For Cyber-Physical System (CPS) development, CI/CD can help achieving required goals, such as high dependability, yet it may be challenging to apply. This paper empirically investigates challenges, barriers, and their mitigation occurring when applying CI/CD practices to develop CPSs in 10 organizations working in 8 different domains. The study has been conducted through semi-structured interviews, by applying an open card sorting procedure together with a member-checking survey within the same organizations, and by validating the results through a further survey involving 55 professional developers. The study reveals several peculiarities in the application of CI/CD to CPSs. These include the need for (i) combining continuous and periodic builds, while balancing the use of Hardware-in-the-Loop (HiL) and simulators; (ii) coping with difficulties in software deployment (iii) accounting for simulators and HiL differing in their behavior; and (vi) combining hardware/software expertise in the development team. Our findings open the road towards recommenders aimed at supporting the setting and evolution of CI/CD pipelines, as well as university curricula requiring interdisciplinarity, such as knowledge about hardware, software, and their interplay.

CCS Concepts: • **Software and its engineering** → **Maintaining software**.

Additional Key Words and Phrases: Continuous Integration and Delivery, Cyber-Physical Systems, Empirical Software Engineering

## 1 INTRODUCTION

Cyber-Physical Systems (CPSs) comprise heterogeneous software and hardware components interacting with each other. They aim at automating operations in different domains, such as automotive, aerospace, healthcare, or railways. As it happens for any software system, CPSs continuously evolve to cope with new customer

---

Authors' addresses: Fiorella Zampetti, [fiorellazampetti@gmail.com](mailto:fiorellazampetti@gmail.com), University of Sannio, Via Traiano, 9, Benevento, Italy, 82100; Damian A. Tamburri, [d.a.tamburri@tue.nl](mailto:d.a.tamburri@tue.nl), Eindhoven University of Technology and the Jheronimus Academy of Data Science (JADS), St. Janssingel 92, 's-Hertogenbosch, The Netherlands, 5211 DA; Sebastiano Panichella, [panc@zhaw.ch](mailto:panc@zhaw.ch), Zurich University of Applied Sciences, Steinberggasse 13, Winterthur, Switzerland, 8400; Annibale Panichella, [a.panichella@tudelft.nl](mailto:a.panichella@tudelft.nl), Technical University of Delft, Van Mourik Broekmanweg 6, Delft, The Netherlands, 2628 XE; Gerardo Canfora, [canfora@unisannio.it](mailto:canfora@unisannio.it), University of Sannio, Via Traiano, 9, Benevento, Italy, 82100; Massimiliano Di Penta, [dipenta@unisannio.it](mailto:dipenta@unisannio.it), University of Sannio, Via Traiano, 9, Benevento, Italy, 82100.

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© 2022 Association for Computing Machinery.

1049-331X/2022/11-ART \$15.00

<https://doi.org/10.1145/3571854>

requirements and technology changes. However, CPSs require a tailored development and operation (DevOps) process and are more challenging to evolve than conventional software [32, 51, 73, 74].

In such a context, adopting effective Continuous Integration and Delivery (CI/CD) practices off the DevOps menu is extremely relevant for setting the execution environment, *e.g.*, Hardware-in-the-Loop (HiL) or simulators. Even though CI/CD has been found effective in introducing several advantages in software development, *e.g.*, the reduction of release cycles and the early discovery of defects [75], its application implies overcoming barriers and challenges [9, 33].

When enacting CI/CD for CPSs, it is expected that further barriers and challenges will arise. In general, existing CI/CD technology cannot be applied to CPSs as is [39]. On the one hand, CPSs demand suitable Verification & Validation (V&V) techniques, and the interaction with HiL or the need to replace them with suitable mock-ups or simulators make the application of CI/CD challenging at best. On the other hand, while for conventional software systems good and bad practices for applying CI/CD have been defined [15, 85], for what concerns CPSs, the practice is still immature to be able to do so. Specifically, the combination of (very diversified and evolving) hardware devices and software, the complex execution scenarios, and the need for simulating hardware components during some build stages introduce new facets that must be considered when setting up a CPS development process, and in particular CI/CD pipelines for CPS development.

This paper aims to empirically investigate the challenges and barriers practitioners encounter while setting up and maintaining a CI/CD pipeline for CPSs, as well as the mitigation strategies adopted to deal with them. Specifically, the study has been conducted through (i) semi-structured interviews with 10 industrial practitioners involved in CPS development for 8 different domains, *i.e.*, aerospace, automotive, energy, healthcare, railways, robotics, identification technology (*i.e.*, Radio Frequency Identification - RFID), and acoustic sensors, (ii) by applying open coding [35] and card sorting [68] to the interview transcripts, (iii) by conducting a member-checking survey within the same organizations involved in the interviews aiming at corroborating the relations between challenges/barriers and mitigation strategies, and (iv) by assessing the relevance of the identified challenges/barriers and related mitigation through a survey involving 55 practitioners involved in CPS development for 9 different domains.

We start by characterizing the CI/CD practices of the interviewed organizations, focusing more on their build automation processes. In doing this, we target three aspects of CI/CD for CPSs, namely (i) the pipeline setting, (ii) the involved phases (*e.g.*, static analysis, testing, delivery, etc.), and (iii) the usage and configuration of simulators and/or HiL. After that, we look at challenges and barriers the organizations encounter, as well as mitigation strategies being adopted to deal with them.

The elicited set of challenges, barriers and mitigation strategies are impactful by providing insights to project leaders and developers, guiding them to configure CI/CD pipelines for CPSs, as well as to staff projects properly coordinate resources with different skills and expertise and acquire equipment. Furthermore, results highlight directions in which education for CPS development must be improved. This includes not only covering interdisciplinary topics between software development, measurements, and automated control, but also a proper introduction to software architectures and design principles, making CPS development flexible enough when switching between simulators and HiL. Last but not least, we identify areas where further research is required, among others domain-aware decision making, the integration of simulators and HiL in the pipeline, and further research in the area of test automation and flakiness detection/avoidance. The specificity of each CPS not only makes some lessons hard to generalize (each pipeline tends to be very different from others), but this also poses challenges when leveraging machine learning approaches upon developing recommender systems.

**Paper structure.** The rest of the paper is organized as follows. As basis for the study, Section 2 discusses the relevant literature. Section 3 describes the study methodology, while Section 4 reports and discusses the study results. Section 5 details the study implications, while threats to the study validity are discussed in Section 6. Finally, Section 7 concludes the paper and outlines future directions.

The study material (after redacting interview transcripts) is available online [84].

## 2 RELATED WORK

This section discusses the literature related to (i) CPS development leveraged for the inception of our study, (ii) CI/CD process, and (iii) CI/CD good and bad practices. Note that this is not an exhaustive systematic literature review on the topic, but rather it points out papers discussing challenges in CPS development and in CI/CD. Finally, it is important to highlight that, while challenges related to CPS development are already investigated from previous literature, to the best of our knowledge, there is very limited empirical evidence on how such challenges translate when setting a CI/CD pipeline for CPS development.

### 2.1 Development of Cyber-Physical Systems

CPSs are more complex and difficult to design, develop, test, and integrate than conventional software systems [32, 51, 73, 74]. Specifically, Törngren *et al.* [74] investigated how CPSs' engineering deals with the complexity of CPS design, and of the environment in which CPSs operate. In this context, it is of paramount importance to perform run-time verification of safety requirements [27], as well as testing encapsulating model-in-the-loop (MiL) [66], software-in-the-loop (SiL), and hardware-in-the-loop (HiL) [2]. With respect to previous studies, we investigate how CPS complexity impacts the setting of CI/CD pipelines, and how developers deal with such complexity.

Considering the costs, risks, and complexity of conducting system testing in a real environment [12, 40], simulation is becoming one of the cornerstones in developing and validating CPSs. CPS developers mainly rely on basic simulation models [29, 67], as well as rigid body [50, 86] and soft body simulation environments [25, 62]. The usage of CPS simulation environments enables automated test generation and execution [37, 54]. However, the limited budget allocated for testing activities and the virtually infinite testing space pose challenges for adequately exercising the CPS behavior [4, 20, 82]. We complement previous studies by looking at the challenges, barriers, and related mitigation strategies when integrating and combining simulators and HiL in CI/CD to support the development, V&V, and evolution of CPSs.

Related to DevOps applications in a CPS context, Park *et al.* [56] analyzed the use and challenges of the digital twin to enable DevOps approaches for cyber-physical production systems to continuously improve them. Specifically, Park *et al.* identified challenges related to (i) discrepancies between models and their physical counterparts, (ii) integration between heterogeneous models due to the complexity of CPSs, and (iii) security issues due to the tight coupling between the digital twin and the physical environment. Instead of only looking at automating the production process, we focus more on the CI/CD process for CPS development and evolution.

Finally, Mårtensson *et al.* [52] identified factors to consider for applying CI to software-intensive embedded systems, such as complexity of user scenarios, compliance to standards, long build times, security, and test environments. These factors represent real impediments for companies who want to adopt CI for embedded systems. While using different research methods, our study is wider than Mårtensson *et al.* (10 semi-structured interviews, plus an external survey with 55 participants vs. case studies with 2 companies), and considers the whole CI/CD process from development to delivery to the customer side. Finally, while we confirm findings from Mårtensson *et al.* [52], our study deepens the analysis of different CI/CD aspects (*e.g.*, setting, phases, and execution environment) for CPSs, and not only in relation to seven CI cornerstones.

Table 1 summarizes the main challenges during CPS development, as stated in previous literature, that are used to drive our study, although we do not focus on specific implementation details of simulators. We leverage the challenges identified by the aforementioned studies to devise the interview guide, in particular those related to (i) the complexity of the underlying environment, (ii) certification and compliance to standards, (iii) test automation, (iv) testing of safety requirements, and (v) MiL, HiL, and simulators.

Table 1. Challenges in CPS development from previous literature.

Ref.	CPS-related development challenges
[74]	Environment complexity, co-designing hardware and software
[27]	Test generation/automation, verification of safety requirements
[2]	Integration of MiL, SiL, and HiL
[12, 40]	Where testing is performed (HiL vs. simulators)
[25, 29, 50, 62, 67, 86]	Implementation of simulators
[4, 20, 82]	Simulator challenges/adequacy
[52]	Standards, long build, security, architecture, test environments of embedded systems
[56]	Digital twin adoption in manufacturing and related design challenges

## 2.2 CI/CD process

Hilton *et al.* [34] found that CI is becoming very popular in open source projects. The latter is also true in industry, even if Ståhl and Bosh [69, 70] found that there is not a uniform adoption of CI in industry. Furthermore, Vasilescu *et al.* [75] showed that CI practices improve developers' productivity without negatively impacting the overall code quality. Finally, Ståhl *et al.* [71], in a study involving three companies, found that the lack of traceability may prevent the application of CI in conventional software systems.

From a different perspective, Elazhary *et al.* [18] looked at the extent to which companies follow the CI practices by Fowler and Foemmel [21] through interviews. Their results emphasized differences among companies in terms of repository structure, testing automation, long build, and deployment challenges. While we share some goals with Elazhary *et al.*, our study, and the dimensions being investigated, relate to CI/CD application for CPS development. In a different study, Elazhary *et al.* [17] used grounded theory to investigate human factors in CI. Even if our study considers human factors, it is not focused on that.

Vassallo *et al.* [79] investigated, by surveying developers of a large financial organization, the adoption of the CI/CD pipeline during development activities, confirming what is known from existing literature (*e.g.*, the execution of automated tests to improve the quality of their product), or confuting them (*e.g.*, the usage of refactoring activities during normal development).

Finally, deepening the continuous delivery practice, Chen [10] analyzed four years of CD adoption in a multi-billion-euro company, and identified a list of challenges related to CD adoption. Savor *et al.* [64], instead, by analyzing the CD adoption in two industrial companies, found that it does not negatively impact developer productivity even when the project increases in terms of size and complexity.

Differently from previous studies, our goal is to shed light on the CI/CD process focusing on the peculiarities of CPS development.

## 2.3 CI/CD barriers and bad practices

Different authors studied barriers and/or challenges in adopting CI/CD. These were initially identified by Duvall *et al.* [13], and are related to the need for maintaining a fully automated build process, handling dependencies, having different levels of builds, and coping with different target environments.

Hilton *et al.* [33] studied barriers developers encounter when moving toward CI, *i.e.*, quality assurance, security, and flexibility. Olsson *et al.* [55], instead, looked at the challenges faced while migrating towards CD: the complexity of the deployment environment, the need to achieve timely delivery, and the lack of a complete overview of all the development projects.

Previous research also found that CI/CD may be wrongly applied, leading to bad practices. Specifically, CI/CD antipatterns have been defined by Duvall [15], and empirically elicited by Zampetti *et al.* [85] from interviews

and Stack Overflow posts. Our study is complementary to that although, where appropriate, we compare the practices observed in our context (CPS-specific) with bad practices recommendations from previous studies. Researchers have developed different kinds of tools to detect and remove antipatterns from CI configuration files [23, 78], analyzing the pipeline aging by observing its execution [77], skipping builds [3], or coping with security-related issues in infrastructure-as-code [60].

To the best of our knowledge, there is no such broad investigation on the application of CI/CD in CPS development and evolution, as well as the challenges and barriers faced together with mitigation strategies to overcome them.

### 3 EMPIRICAL STUDY DEFINITION AND PLANNING

The *goal* of this study is to investigate the CI/CD practices for CPS development, to identify challenges and barriers encountered in such practices, together with mitigation strategies adopted to overcome them. The *perspective* is of researchers interested to support developers in configuring CI/CD pipelines for CPSs, and practitioners setting, using and evolving CI/CD pipelines for CPS development. The *context* from which we have inferred the set of challenges and barriers with related mitigation strategies encountered when setting or evolving CI/CD pipeline for CPS development consists of 10 organizations developing CPSs for 8 different domains. To assess the identified set of challenges/barriers and related mitigation strategies, we have surveyed 55 practitioners (not involved in the first step of this study) developing CPSs for 9 different domains.

We start by creating organizational profiles by looking at the CI/CD practices adopted by the interviewed organizations, and in general all the practices the organizations are adopting to automate different stages of a build. Specifically, we look at the conditions that determine (i) the setting of the CI/CD pipeline, *e.g.*, whether an incremental build is used, when the build is triggered, or whether build matrices are used; (ii) the phases instantiated in the pipeline, *e.g.*, static analysis, various testing levels, or deployment; and (iii) the use of simulators and HiL in the context of the CI/CD pipeline.

The study addresses the following two research questions:

- **RQ<sub>1</sub>**: *What are the challenges and barriers respondents encounter, and how do developers deal with them?* After having characterized the CI/CD and build automation practices, we investigate the challenges (*e.g.*, the need to cope with a slow build or flakiness, or with phases not easy to automate) and barriers (*e.g.*, limited availability of software and/or hardware resources) encountered by the interviewed organizations when dealing with the setting or evolution of the CI/CD pipeline for CPS development. Furthermore, we highlight the strategies (*e.g.*, to adopt a pipeline that relies on both simulators and HiL in different build stages) adopted by the interviewed organizations to deal with challenges and barriers. We validated the relations between challenges/barriers and mitigation strategies through a member-checking survey with the same interviewees or practitioners belonging to the same team of the interviewees involved in the semi-structured interviews.
- **RQ<sub>2</sub>**: *How relevant are the identified CI/CD challenges/barriers and their mitigation for practitioners involved in CPS development?* While in the previous research question we identified a set of challenges and barriers with related mitigation strategies as experienced by our interviewees, this research question aims at performing an external validation by surveying practitioners involved in the setting, evolution or usage of a CI/CD pipeline for CPS development.

The study methodology used for addressing the research questions is depicted in Figure 1 and described in the following. After having recruited participants to be involved in the semi-structured interviews through personal knowledge, we conducted the interviews and transcribed their content. Note that, since this is an exploratory study, we prefer to rely on convenience sampling, as previously done in literature [18, 33]. This is because practitioners involved in CPS development represent a hidden population, therefore we did not have a sampling frame [8].

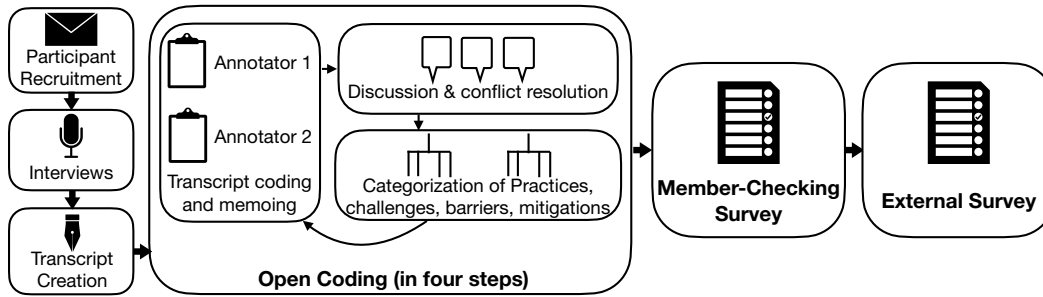


Fig. 1. Study methodology.

The latter helps us to conveniently reach a suitable number of study participants. After that, we performed an incremental (in four steps) open coding [35] of the transcripts, discussing the independent coding of multiple annotators, solving conflicts, and creating, through a card sorting strategy [68], categorizations for practices, as well as for barriers, challenges, and mitigation strategies. The relationships between challenges/barriers and mitigation strategies have been validated through a member-checking survey, and, finally we performed a further survey to validate our findings beyond the interview context.

### 3.1 Data Collection: Semi-structured interviews

We defined the interview structure through an iterative process, which started from the existing knowledge on the topic summarized in Table 1 (see Section 2). From such knowledge, all theoretical pending points were distilled and matched with interview structure areas and questions for each interview structure area. As summarized in Table 2, we start with demographics about the organization and the interviewee, and get a first glance at the development and lifecycle management practices [5] adopted in the context of interest. Then, we gather data about the pipeline structure and technology, paying particular attention to V&V and deployment. We then explore the usage of simulators and HiL. We also investigate the presence of any machine learning or (ML)-intensive components to be automated (e.g., trained/tuned) or executed by the pipeline over any CPSs software artifact, or, conversely, the use of ML and Artificial Intelligence (AI) for pipeline automation (e.g., as part of the testing oracle), i.e., AIOps [11]. After that, we investigated how the interviewees configure the overall CI/CD pipeline in terms of build triggering strategies and the possibility to handle different pipeline configurations, each one environment-specific. The interview ends with general questions about the main benefits achieved, barriers encountered, and challenges to tackle when configuring and evolving the CI/CD pipeline.

**Interview participant selection.** The interview participants have been selected based on personal knowledge, with the goal of identifying experienced practitioners over the theoretical constructs (CI/CD pipelines for CPS) under investigation. The resulting study size (10) is not particularly high, yet on the same order of magnitude as similar interview-based studies on CI/CD [18, 33] (although the study by Hilton *et al.* was followed by a larger survey). It has to be considered also that, differently from previous studies, we targeted a very specific development domain and technology (i.e., application of CI/CD for CPSs in industrial settings). After participants accepted our invitation, we gave them an overview of the questions to expect in the interview, to allow them to gather any additional information.

Table 3 summarizes demographic information about organizations and interviewees involved in the study. Five out of ten organizations are large (i.e., over 1,000 employees), one is medium (i.e., between 50 and 1,000 employees), two are small (between 10 and 20 employees), and two are micro (less than ten employees). Furthermore, the sample covers eight different domains: aerospace, automotive, energy, healthcare, railways, robotics, identification

Table 2. Interview structure

Section	Content
Overview	Company description, domain, programming languages Respondent background and role
CI/CD pipeline structure	Phases and steps Tools (versioning, build automation, CI/CD, use of containers) Verification and Validation approaches Deployment
Simulators and HiL	Simulator development/acquisition Simulator/HiL integration in the pipeline Simulators vs HiL tradeoffs
ML-based components	In the developed software In the pipeline
CI/CD pipeline configuration	Pipeline stability Build strategies Triggering
Conclusion	Challenges Barriers Expected benefits

Table 3. Participant Demographics

Org <sub>ID</sub>	Organization			CPS Exp. (Y)
	Domain	Size	Role	
O <sub>1</sub>	Aerospace	Small	R&D Manager	8
O <sub>2</sub>	Healthcare	Large	DevOps Architect	18
O <sub>3</sub>	Acoustic Sensors	Small	SW and HW Integrator	6
O <sub>4</sub>	Robotics	Medium	Team Leader	7
O <sub>5</sub>	Automotive	Large	R&D Manager	20
O <sub>6</sub>	Aerospace	Large	R&D Manager	20
O <sub>7</sub>	Railways	Large	SW and HW Integrator	10
O <sub>8</sub>	Railways	Micro	Team Leader	25
O <sub>9</sub>	Identification Technology	Micro	Software Engineer	3
O <sub>10</sub>	Energy	Large	Project Leader	5

technology (*i.e.*, RFID), and acoustic sensors. Finally, the participants' professional experience in the CPS field varies from 3 to 25 years, with varying job titles, and all of them are currently involved in the configuration of the CI/CD pipeline.

As it will be clearer later, we intentionally selected participants having different maturity levels in the implementation of a CI/CD pipeline for CPSs. That is, we also included organizations that, while having experience in setting CI/CD pipelines, only partially automated CPS builds, without having a full-fledged CI/CD pipeline. This allowed us to understand, in those cases, how they automated certain phases, as well as the reasons why they are still facing challenges in having a complete CI/CD pipeline.

**Conducting interviews.** Interviews were conducted using a videoconferencing system, by one researcher (with the support of one-two other researchers), following an order based on interviewees' availability. Before starting the interview, the interviewer recalled study goals and gathered consent for recording. The interview

structure was followed rigorously, varying only the level of detail over different areas of the interview based on the provided answers. For instance, if a participant mentioned the use of simulators, we asked deeper questions on the topic, while we skipped questions not applicable to a given participant. It is important to remark that, interviews are treated as independent from each other, meaning that questions were not adjusted over different interviews. This is because, as shown in Table 3, the involved organizations cover a broad range of domains, and the main goal was to achieve a similar understanding among those domains.

**Creating transcripts.** After interviews have been completed, a researcher transcribed the audio, creating a document organized into sections as Table 2. The transcripts contain a total of 15,329 words and 787 sentences.

### 3.2 Data Analysis from interview transcripts

Two authors, experts of the domain, (hereinafter referred to as “coders”), independently used online spreadsheets to assign codes (*i.e.*, open coding) to sentences in the transcripts. The coding was carried out following the approach illustrated by Hoover [35], *i.e.*, annotating a code near sentences of the transcript. A code is defined as a mnemonic label identifying a concept defined in a text fragment, *e.g.*, by applying the label ‘TEST’ to any part of text reflecting a software testing activity. Wherever appropriate, the coder added a memo that could be leveraged to better explain the observed phenomenon, as well as to identify possible relationships between codes dealing with different aspects of the CI/CD pipeline setting and evolution.

Open coding has been performed over four subsequent sessions by arranging the 10 interview transcripts into four groups. Each group included two, three, four, and one interview, respectively. After each coding session, the coders held a discussion meeting, in which similar codes created by multiple coders were merged, and conflicts were resolved. After each round, we computed the Krippendorff  $\alpha$  [45] to determine the achieved level of agreement. The obtained  $\alpha$  values for the four iterations were 0.65 (close to the minimum acceptability of  $\alpha = 0.66$  [45]), 0.71, 0.69, and 0.86 (substantial agreement). Starting from the second iteration, the coders could reuse, through a drop-down cell, codes created during previous iterations, or create new ones. Note that, to further limit agreement by chance, each code annotation was reviewed during the meetings, not just the disagreements.

During the discussion meetings, broad groups of codes were also defined. For instance, we distinguished codes belonging to the CI/CD pipeline from those related to the development process. Also, we started grouping codes belonging to different phases of the pipeline, and codes related to challenges, barriers, and mitigation strategies. Such a categorization started during the first discussion meeting and then was refined over the next ones. After the first two sessions of the open coding (after the first session the set of codes was too immature for this purpose), three researchers iteratively produced—by adopting a card sorting strategy [68]—the first version of a mind map grouping codes into categories. Such a mind map has been used as a support to ease the subsequent open coding phases and to evaluate the extent to which non-leaf nodes were saturated. Note that we do not expect a full saturation [63] in this study, due to the high diversity of the considered application domains. The mind map was then refined after each subsequent coding phase.

Overall, we identified a set of 179 codes, which led to the construction of a categorization of codes explaining the phenomenon, organized across 43 high-level categories.

Finally, the two coders performed three iterations over the transcripts, codes, and memos to derive relations between different codes. For instance, it is possible that process constraints (*e.g.*, the need to use a specific type of simulator or tool imposed by the domain, or to adopt certain coding standards), introduce challenges while setting the pipeline (*e.g.*, the need to cope with phases not easy to automate, or slow build and flakiness) that may be addressed by relying on a particular mitigation strategy (*e.g.*, push small changes when using incremental builds). As an example, when talking about flaky behavior experienced in the build process,  $O_4$  mentioned that: “*of course, we have some retry for network issues*”, while “*in case of resources problems we do not have retries, but the pipeline maintainers can open issues aimed at solving the problem*”. The outcome of this step consists of 90 relations



from 128 sentences. We will present how different codes relate to each other and are spread among different organizations through storytelling.

### 3.3 Member-checking survey to validate relations between challenges/barriers and mitigation strategies

To verify our understanding of how the interviewed organizations act to address the challenges and barriers encountered while setting and maintaining the CI/CD pipeline for CPS development, we conducted a member-checking survey by involving the interviewees themselves, or people working in the same team of the interviewees. Asking outside the same team, especially in large organizations, would have reached completely different projects or even different domains, even unrelated to CPS, making the member-checking worthless.

The survey has been designed by following guidelines for survey design and operation from social science [31] and software engineering [41–44, 57]. Specifically, the survey contains:

- An introduction explaining the study goals;
- A set of ten sections in which we validate the relations between the 10 challenges/barriers for which we found at least one mitigation strategy from the transcripts;
- A demographic section in which we asked the participant: the application domain, the role within the organization, the years of experience in CPS development, as well as information about the CI/CD pipeline (*i.e.*, (i) whether or not the organization has a CI/CD pipeline in place, (ii) years from its introduction, and (iii) how the participant interacts with it).

Since the main goal of the survey is to validate our correct understanding of the challenges/barriers and related mitigation strategies, we asked the participants to provide their personal contacts (among them the name of the organization) mainly for traceability purposes.

For each section in the survey, we start by asking whether or not the challenge/barrier has been faced at least once by the team they are working with. Specifically, instead of using a yes/no question, we added a third option aimed at highlighting those cases where the challenge/barrier cannot be encountered due to the development process adopted by the organization. For instance, if an organization does not use HiL in its development process, it will never experience problems due to the high cost or lack of scalability of the hardware devices. If the challenge/barrier has been encountered at least once within the organization, we list a set of questions, each one aimed at investigating the adoption of the identified mitigation strategy to overcome the previously presented challenge/barrier. Specifically, the respondent could choose between three different options: (i) yes, and we used it, (ii) yes, but we never used it, and (iii) no. Two out of 17 questions dealing with mitigation strategies provide only two options: (i) yes, it happened, and (ii) no, it never happened. At the end of each section, respondents could use an optional free comment field to provide additional mitigation strategies adopted for overcoming the related challenge/barrier.

The questionnaire was administrated through Survey Hero<sup>1</sup>, and we kept the questionnaire open for 12 weeks. Note that nobody reported having particular issues (*e.g.*, privacy issues) with the used survey administration tool. Furthermore, because of constraints imposed during the survey administration, we had to keep it anonymous.

After closing the survey, we obtained eleven responses from the 10 organizations involved in the semi-structured interviews. Specifically, for O<sub>5</sub> we obtained two different responses, even if one of them did not provide demographic information. Among the 10 respondents providing their personal contacts, four of them have also participated in the semi-structured interviews. Furthermore, 4 are R&D Managers, 3 are software and hardware integrators, 2 are DevOps Architects, and 1 is a DevOps QA Engineer. In terms of years of experience with CPS development, five respondents have between 1 and 5 years of experience, two respondents between 5 and 10, and the remaining two more than 10 years. Seven out of nine participants (the ones answering this specific question),

<sup>1</sup><https://www.surveymhero.com/>

declare that their organization already has in place a CI/CD pipeline used while developing CPSs (1 introduced it less than one year ago, 1 has a mature pipeline introduced more than five years ago, while 5 between one and five years ago), and in terms of the way they interact with the pipeline, among the six participants who answered this question, 1 only uses the CI/CD pipeline, 2 are involved in its setting and maintaining, and 3 set, maintain, and use it for their development tasks.

### 3.4 Evaluation through an external survey

To address RQ<sub>2</sub> we conducted a survey involving practitioners using (or trying to set up) a CI/CD pipeline for CPS development in their organization. To recruit participants we used two different sources:

- (1) *Snowball sampling* [30], *i.e.*, we shared the survey link to some personal contacts and encouraged them to indicate us further participants. This choice has been dictated because, while we had a relatively limited set of contacts reachable with our knowledge, snowballing could help us to reach the relevant people (those involved in CPS development by relying on a CI/CD pipeline).
- (2) *An infrastructure for recruiting survey participants*, namely *Prolific*<sup>2</sup>. This platform allows to reach additional participants, by paying a small fee. The platform has a participant screening facility (we required participants to have at least a bachelor's degree in computer science or similar, and knowledge about relevant software development technology, including versioning, monitoring, virtualization, and testing). Also, similarly to what was done in the member-checking survey, we collected further information about CI/CD competences to further filter participants. At the same time, we are aware that with *Prolific* we have less control over the participants' reliability than with snowballing. To mitigate this problem, our online package contains separate results belonging to the snowball sub-sample and the *Prolific* sub-sample.

The online survey presented to the participants has: (i) an introduction explaining the study goals, *i.e.*, to assess a catalog of challenges and barriers concerning the setting and maintaining of a Continuous Integration (CI) and Continuous Delivery (CD) pipeline for CPS development; (ii) 14 sections in which we validate challenges, barriers and mitigation strategies; and (iii) a demographic section similar to the one described in Section 3.3.

We started by asking, for each challenge/barrier (properly grouped in categories), whether they have ever encountered it as a factor preventing/limiting the setting up of a CI/CD pipeline, or, if the participant did not encounter it, whether she perceive the challenge/barrier as a real impediment. Specifically, the respondent could choose between four different options: (i) yes, it is relevant (and I encountered it), (ii) yes, it is relevant (but I never encountered it), (iii) no, I do not consider it as relevant, and (iv) does not apply to my context. If at least one of the challenges/barriers in the category was felt as relevant to the respondent, the survey shows a new section asking about the mitigation strategies used (or felt as relevant) to address the previously selected challenges/barriers by using a multiple choice answer. Specifically, the respondent could choose among the mitigation strategies we obtained as a result of RQ<sub>1</sub>, but could also add new (unseen) mitigation strategies. Finally, the survey contains an open-ended question aimed at collecting other challenges/barriers that did not apply to the 10 interviewed organizations.

Also in this case, the survey has been administrated through Survey Hero, and nobody reported having particular issues with this administration tool. For the snowball sample, the questionnaire has been left open for one month, and due to constraints imposed during the survey administration, we kept it anonymous. For what concerns *Prolific*, we obtained the requested responses within the same days the survey has been opened.

In the end, we obtained 19 responses from the snowball sampling, and 50 further responses from *Prolific*. However, through a screening of the participants' answers we discarded 14 responses from *Prolific*, *i.e.*, (i) it was difficult to infer whether or not the participant works for CPS development, *e.g.*, education or applications for cosmetic stores, and (ii) the participant declares to not have a CI/CD pipeline in place within the organization,

<sup>2</sup><https://www.prolific.co>




Fig. 2. CPS domains from the external validation survey

and at the same time declares that the CI/CD pipeline has been adopted only recently. As a result, we obtain a final set of 55 valid responses covering 9 different application domains (as shown in Figure 2).

Among the respondents providing demographic information (51), in terms of the role played in their organization, there are: 23 software and hardware integrators, 13 R&D Managers, 7 DevOps Architects, 5 software developers/testers, 1 Project Manager, and 1 CTO (Chief Technology Officer). In terms of years of experience with CPS development, 19 respondents have less than 1 year of experience, 27 between 1 and 5, two respondents between 5 and 10, and the remaining three more than 10 years. 47 out of 51 respondents declare that their organization already has in place a CI/CD pipeline used while developing CPS (19 introduced it less than one year ago, seven have a mature pipeline introduced more than five years ago, while 21 between one and five years ago.) Finally, in terms of the way our respondents interact with the pipeline, 31 only use the CI/CD pipeline, six are involved in its setting and maintaining, and the remaining 10 set, maintain, and use the pipeline for their development tasks. Finally, among the respondents who declare that their organization does not have a CI/CD process in place for CPS development, three declare being involved in setting it.

#### 4 STUDY RESULTS

In the following, we report and discuss the results addressing the RQs defined in Section 3. To properly contextualize challenges, barriers, and their mitigation strategies, it is important to summarize the development process of the interviewed organizations. Specifically, Section 4.1 briefly describes, for each organization participating in the semi-structured interviews, the CPS development process, focusing more on the adoption of CI/CD pipelines and, in general, on their level of build automation. The interested reader could find more details in the Appendix.

##### 4.1 Contextualization: Organization Profiles

Table 4 provides an overview of the main analyzed dimensions for the 10 organizations considered in our study. In the following we briefly describe them.

Table 4. Summary of the CPS development process adopted within the 10 interviewed organizations (*i.e.*,  $O_{ID}$ ). The ✓ (✗) occurs when the property (does not) apply to the organization; the — represents cases where the property is not applicable/available for the organization; ● means that the phase is automatized within the pipeline; ◐ means that the phase is automatized but not included in the pipeline; ○ means that the phase is done manually.

Property		Organizations									
		$O_1$	$O_2$	$O_3$	$O_4$	$O_5$	$O_6$	$O_7$	$O_8$	$O_9$	$O_{10}$
Prog. Language		C	C# C++	C C++	C++ Python	RTJ	C C++	C	C C++	C# Java	Java Python
Pipeline Maturity		< 1	[1, 5)	✗	—	< 1	—	[1, 5)	✗	✗	≥ 5
Phases	Static Analysis	●	●	✗	●	✗	●	○	○	✗	●
	Unit Test	●	●	○	●	●	●	●	○	◐	●
	Int. Test	✗	●	○	●	✗	●	●	○	◐	●
	System Test	✗	●	✗	●	✗	◐	◐	○	○	✗
	Non-Func. Test	●	◐	○	○	◐	○	◐	○	◐	●
	Deploy	✗	●	◐	●	●	●	●	◐	◐	●
Triggering	Continuous		✓	—	✓	✓	✓	✓	—	—	—
	Incremental		✓	—	—	—	—	—	—	—	✓
	Nightly	✓	✓	—	✓	✓	—	—	—	—	—
Pipeline Config.	Env.	Stable	Domain specific	—	Stable	—	Device specific	Device specific	—	—	Stable
	Staged Builds	✗	✓	—	✗	✗	✗	✓	—	—	✗
Mocking		✗	✗	—	✗	✗	✗	✗	—	—	✓
Simulators		Ext.	Int.	—	Ext.	Int.	Int.	Int.	Int.	Int.	Int.
HiL		✗	✓	✓	✓	✓	✓	✓	✓	✓	✓
Containerization	VMs	—	—	—	—	✓	—	✓	—	✗	✗
	Docker	—	—	—	Deploy and HiL	✗	✗	—	—	Deploy	HiL Deploy

4.1.1  $O_1$  (Aerospace).  $O_1$  is involved in verification and validation (V&V) tasks for aerospace software (*i.e.*, on-board software for satellites), hence their CI/CD pipeline is only for V&V and not for development. The standards in the aerospace domain enforce the adoption of conventional programming languages, *i.e.*, “We mainly use ANSI C-99 following the MISRA rules”, as well as the need for certifying software.

$O_1$  started to adopt CI/CD practices for CPSs less than one year ago. Due to the application domain and the related standards and certification constraints, the pipeline compiles the software provided and developed by the customer, relies on SonarQube for static code analysis checks, and executes unit and robustness tests to “check how the system behaves/reacts in the presence of unexpected inputs”. The triggering of the pipeline is almost manual, even if there are scheduled nightly builds for running test suites requiring a long time to complete.

Finally,  $O_1$  cannot involve HiL in the pipeline, as it would require a clean room not accessible from the outside. Instead, it relies on third-party simulators provided by the customer, reducing the costs/efforts needed to develop the simulators from scratch, as well as guaranteeing the trustworthiness of the outcome being produced.

4.1.2  $O_2$  (Healthcare).  $O_2$  is a large organization involved in the healthcare domain. It adopts conventional programming languages, *i.e.*, mainly C# and C++ during the development process.

O<sub>2</sub> has a CI/CD pipeline in place for CPS development that has been introduced 4 years ago, and they are still improving it. Furthermore, based on its application domain, O<sub>2</sub> is constrained to “*follow medical application frameworks providing a base set of rules in terms of how to build applications and how to integrate them*”.

O<sub>2</sub> adopts both incremental and nightly builds. While nightly builds leverage HiL and run three different types of testing, namely unit/component, sub-system, and system testing, incremental builds leverage self-developed simulators to provide developers fast feedback about the impact of their changes, *i.e.*, only a subset of the whole set of functional tests are executed. Furthermore, both incremental and nightly builds run static code analysis tools. Finally, nightly builds imply an automated deployment on a “real” Computed Tomography (CT) scanner, *i.e.*, “*physical systems that are equivalent to the real hardware in the CT Scanner but not connected to anything around it which has a simulator running on it*”.

4.1.3 O<sub>3</sub> (Acoustic Sensors). O<sub>3</sub> is involved in CPS innovation for the industry, among others, the development of the SPL Noise Meter Board, by using conventional programming languages, *i.e.*, Python for testing and C, C++ for micro-controllers development. Each team is composed of both software and hardware experts who work together.

O<sub>3</sub> does not have a CI/CD pipeline for CPS development, however, the deployment is fully automated, while the testing is manual, *i.e.*, impossibility to automatically test acoustic signals. Finally, at the moment, O<sub>3</sub> only uses real hardware devices, yet they wish to include simulators in their CI/CD process.

4.1.4 O<sub>4</sub> (Robotics). O<sub>4</sub> is involved in the development of autonomous robots, and similarly to O<sub>3</sub>, each team accounts for both hardware and software experts. In their development process, O<sub>4</sub> mainly adopts C++, together with Python for users’ interfaces and for interacting with the hardware devices.

O<sub>4</sub> has a fully containerized (using Docker) pipeline for CPS development. It relies on continuous and nightly builds for running regression testing activities on already packaged components and for deployment to the customers. Furthermore, continuous builds also execute static code analysis tools to inform developers about code quality degradation, and unit tests relying on simulators. The application domain does not introduce certification constraints, while it hinders the automation of non-functional testing within the pipeline. Finally, O<sub>4</sub> relies on third-party simulators and HiL into different stages of the whole CI/CD process.

4.1.5 O<sub>5</sub> (Automotive). O<sub>5</sub> is a large organization operating in the automotive domain working on the software-focused driving platform. This is the only organization in our study relying on real-time languages, *i.e.*, real-time Java to cope with scheduling requirements of embedded systems.

O<sub>5</sub> already has a CI/CD pipeline in place mainly for deployment purposes, even if it is working on improving it. However, unlike the others, O<sub>5</sub> relies on virtual machines instead of using Docker containers. Moreover, O<sub>5</sub> does not test all the developed modules together since it “*deploy[s] individual bundles to a platform*.”

Finally, since O<sub>5</sub> develops software for embedded entertainment in the automotive domain, the HiL is only available for a final validation on the customer’s side, so most of the work is done relying on virtual environments.

4.1.6 O<sub>6</sub> (Aerospace). O<sub>6</sub> operates in the aerospace domain, and it is mainly involved in developing and refining the routing algorithm for the Free Route Airspace (FRA). Similarly to O<sub>1</sub>, it relies on conventional languages: “*C and C++ [are] used for the back-end*.”

O<sub>6</sub> already has a CI/CD pipeline including static code analysis, unit testing, integration testing, and deployment. Similarly to O<sub>1</sub>, it is required that the developed code satisfies strict certification requirements that are mainly checked by relying on code coverage tools. However, differently from other organizations, O<sub>6</sub> does not rely on nightly builds, meaning that also time-intensive tasks are executed at each change: “*even the slow builds are continuously built*.” Finally, the pipeline provides a monitoring mechanism for what concerns aspects of the real-time operating system such as scheduling and memory that “*gives us the possibility to collect feedback/evidence that may help us in obtaining the certifications*.”

As regards HiL and simulators,  $O_6$  relies on both, however it does “*not have simulators and HiL in the same pipeline mostly for certification issues.*”

**4.1.7  $O_7$  (Railways).**  $O_7$  is involved in delivering software for railways, *i.e.*, Train Control Management System (TCMS). In terms of programming languages being used, the interviewee mentions the need of adapting the programming language to the device on which the software has to be executed.

$O_7$  already has a CI/CD pipeline in place for CPS development that, at the moment, is in a continuous improvement state. Based on the application domain,  $O_7$  adopts staged builds following the “green-build rule”. In the first stage, the build process is executed on a virtual machine, and in the presence of a green status, all the components are deployed together, enabling the execution on the virtual train. In the presence of a green status, it is possible to move to the next stage that relies on the hardware test track, “*where [there is] the whole set of devices and even some more that [are not] in the virtual train.*” Finally, in the presence of a green status it is possible to run the last stage relying on a real train. All the stages include functional testing, while the deployment is automated only for the first stage. Based on the previous statements, it is possible to conclude that  $O_7$  adopts both simulators and HiL in different stages of the build process, with the use of HiL occurring only in the last stage of the pipeline.

**4.1.8  $O_8$  (Railways).**  $O_8$  is involved in the railways domain, *i.e.*, the development of a specific component used for transmitting data between on-board and ground applications.  $O_8$  uses C and C++ (*i.e.*, conventional programming languages), and it has strict certification requirements, *e.g.*, compliance with the railway standards and specifications.

$O_8$  does not have a CI/CD pipeline in place for CPS development and it has, in general, little automation in the development process, *i.e.*, only the adherence to standards and specifications is automated. Finally, due to the high cost of the hardware devices in this particular domain,  $O_8$  mainly relies on simulators that are self-developed. However, once per week,  $O_8$  performs a testing session with a real “*train running in a real environment with real traffic*”.

**4.1.9  $O_9$  (Identification Technology).**  $O_9$  is involved in “*develop[ing] software relying on identification technologies such as RFID [(Radio Frequency IDentification),] Bluetooth low energy or bar codes*”, relying on conventional programming languages such as Java and C#.

Due to a lack of culture for setting a pipeline dealing with sensors and actuators,  $O_9$  does not have a CI/CD pipeline for CPS development. However, the testing phases are almost fully automated. For what concerns the deployment of CPS-related software,  $O_9$  relies on Docker for creating images that are manually deployed onto the servers. The development process also features a monitoring component for the internal development platform and customers’ devices, to notify about anomalies and errors, as soon as they occur. Finally, the development process considers both (self-developed) simulators and HiL.

**4.1.10  $O_{10}$  (Energy).**  $O_{10}$  is involved in the development of prototypes and proof of concepts for the energy domain. It has a mature (*i.e.*, introduced in 2016) pipeline for CPS development that uses conventional programming languages, mostly Java and Python.

Other than having a compilation phase, the CI/CD pipeline is aimed at executing static code analysis tools and linters, unit, and integration tests, followed by a deployment phase where the packaged version of the software is usually stored into an artifact repository as a docker image.  $O_{10}$  does not rely on nightly builds, while it only uses incremental builds.

$O_{10}$  does not need to run the software on embedded devices, implying that  $O_{10}$ , other than simulating the hardware when needed, mainly replaces it with mock-ups. Only when the real devices are available and it is safe to use them for testing,  $O_{10}$  uses Docker images for checking the correct behavior over the real devices.

Table 5. Process-related challenges

Category	ID	Challenge	Organizations
General	PRC <sub>1</sub>	Cycle-time reduction	O <sub>2</sub>
	PRC <sub>2</sub>	Onboard developers	O <sub>7</sub>
Culture	PRC <sub>3</sub>	Limited CI/CD culture	O <sub>1</sub>
	PRC <sub>4</sub>	Limited CI/CD culture for CPS development	O <sub>9</sub>
Environment	PRC <sub>5</sub>	Complexity of the environment	O <sub>2</sub> , O <sub>4</sub> , O <sub>5</sub> , O <sub>6</sub> , O <sub>8</sub>
	PRC <sub>6</sub>	Variability of the environment	O <sub>7</sub>
	PRC <sub>7</sub>	Lack of redundancy in the environment	O <sub>7</sub>
Testing	PRC <sub>8</sub>	Test cases manually derived	O <sub>5</sub> , O <sub>8</sub>
	PRC <sub>9</sub>	Test cases manually executed	O <sub>3</sub> , O <sub>9</sub>
	PRC <sub>10</sub>	Different interpretations for the same requirements	O <sub>7</sub>
	PRC <sub>11</sub>	Need a controlled environment for test automation	O <sub>3</sub> , O <sub>4</sub>
	PRC <sub>12</sub>	Complexity in oracle specification for test automation	O <sub>3</sub> , O <sub>5</sub> , O <sub>6</sub> , O <sub>8</sub> , O <sub>9</sub>
	PRC <sub>13</sub>	Complexity for deriving integration tests	O <sub>10</sub>
Deployment	PRC <sub>14</sub>	Complexity for deriving safety tests	O <sub>4</sub> , O <sub>5</sub> , O <sub>8</sub>
	PRC <sub>15</sub>	Late deployment	O <sub>2</sub>
Simulators	PRC <sub>16</sub>	Expensive deployment	O <sub>7</sub>
	PRC <sub>17</sub>	Lack of trustworthiness for simulators	O <sub>3</sub>
	PRC <sub>18</sub>	Complexity for oracle automation with simulators	O <sub>8</sub>

#### 4.2 RQ<sub>1</sub>: What are the challenges and barriers respondents encounter, and how do developers deal with them?

This research question describes barriers and challenges emerging from the semi-structured interviews. We start by describing the challenges related to the CPS development process in general. Then, we describe barriers and challenges encountered when setting and maintaining the CI/CD pipeline for CPS development, together with the related mitigation strategies. Note that we did not find mitigation strategies for all the barriers and pipeline-related challenges and, as described in Section 3.3, the member-checking survey only considers the barriers/challenges for which there was an explicit mitigation strategy reported by at least one of the interviewed organizations.

**4.2.1 Process-related challenges.** Table 5 reports the process-related challenges identified in our interviews, together with the traceability among which challenge has been encountered by which organization. It is important to remark that process-related challenges may not be specific to CI/CD, but are, more in general, challenges in the development process that, based on what was reported by the interview participants, have an impact on setting up and maintaining a CI/CD pipeline.

The challenges have been grouped into six different categories, *i.e.*, general, culture, environment, testing, deployment, and simulators. For each category, in the following, we provide a brief description of the challenges belonging to it, together with some examples.

**General.** This category accounts for two challenges, each one mentioned by only one out of ten organizations. One of the main benefits of adopting a CI/CD pipeline is related to the overall cycle time reduction (PRC<sub>1</sub>). However, even if O<sub>2</sub> has already invested effort and money in reducing the release time, it already sees space for reducing it: “*The biggest problem ... is cycle time. Three years ago, the cycle time was six weeks, while now we could do it every day. It is still not enough from a developer perspective because the feedback is not fast enough.*” While this challenge also applies to conventional software, when it comes to the CPS context, the challenge is

exacerbated mainly due to the need of interacting with both HiL and simulators. In this regard, O<sub>2</sub> mentioned that the cycle time cannot be easily reduced due to (i) the high costs for the infrastructure, and (ii) the translation of test strategies to hardware devices being very demanding”.

O<sub>7</sub> is facing problems when trying to onboard new developers (PRC<sub>2</sub>) mainly due to the complexity of the railways’ domain, as also found by Törngren *et al.* [74]. The interviewee stressed that in the railways’ domain it is crucial to follow specific standards that need to be known and properly understood by developers and testers.

**Culture.** This category groups two challenges related to the presence of a limited CI/CD culture in the development teams. This may limit the possibility of properly leveraging CI/CD facilities throughout the development process. O<sub>1</sub> reports the adoption of a pipeline that only includes tasks that are easy to automate mainly due to “*lack of knowledge*” (PRC<sub>3</sub>), as also found by Zampetti *et al.* [85]. Instead, while O<sub>9</sub> has already in place a pipeline for developing and deploying mobile apps to the app-store (*i.e.*, “*The setting of a CI/CD pipeline in the mobile context has been very easy*”), it does not have a pipeline for CPS development due to “*a lack of a deeper knowledge in the CI/CD context for CPS*”, in particular for what concerns the interaction between software and hardware components (PRC<sub>4</sub>). Specifically, there is a need for knowledge on how to properly account for the inclusion and setting of both HiL and simulators in the CI/CD pipeline configuration, as well as how to include a feedback mechanism to gather information directly from the field.

**Environment.** This category features three different challenges dealing with the characteristics of the physical environment in which the developed code has to be deployed.

Among them, only PRC<sub>5</sub>, *i.e.*, environment complexity, is mentioned by multiple organizations (five out of ten), while the remaining two only come from O<sub>7</sub>. The complexity of the environment impacts the execution environment being set (*i.e.*, simulators or HiL). The unavailability of third-party simulators (and the need for self-developing them) impacts the ability to simulate certain behaviors, or even in deviations between HiL and simulated environments. The consequence is that builds executed on simulators will have a different outcome when run on HiL. For instance, O<sub>4</sub> mentioned: “*Walking is not so easy to simulate so we need a real walking robot for spotting bugs*”, while O<sub>8</sub> stated: “*It could be difficult, demanding and expensive to have a one-to-one relationship between simulators and real systems*”. Our findings stress what is already known from previous literature in terms of relying on simulated environments, *i.e.*, the testing over simulators may fail to expose problems that would only manifest when running the system on the real hardware [52].

O<sub>7</sub> faces a problem related to the high environment variability (PRC<sub>6</sub>) [74], due to trains having different characteristics: “*We can rarely copy-paste software that has to run on different train architectures.*” At the same time, O<sub>7</sub> also faces a challenge due to the structure of its development process that is not cloud-based and has no redundancy (PRC<sub>7</sub>), implying that “*in the presence of network issues or server issues we are totally black and this is affecting everyone.*”

**Testing.** This category groups seven challenges. O<sub>5</sub> and O<sub>8</sub> mention as a challenge the substantial manual effort required for the test case specification process (PRC<sub>8</sub>). O<sub>3</sub> and O<sub>9</sub>, instead, felt the manual execution of testing activities to be challenging, *i.e.*, PRC<sub>9</sub> (*e.g.*, “*Another big barrier is related to the test case execution that, at the moment, we are doing manually since both the environment setting and the oracle definition require manual intervention*” for O<sub>9</sub>). Our findings confirm what is already pointed out by Mårtensson *et al.* [52] in terms of the presence of complex user scenarios implying the need of manual testing.

O<sub>7</sub> found it difficult to automate the test case specification mainly because the standards might be interpreted differently by different developers, and both might be correct (PRC<sub>10</sub>) — “*how do you read the standard? The standard is interpreted so the same requirement can be differently interpreted by different people (a challenge for automation).*” A different challenge experienced by O<sub>3</sub> and O<sub>4</sub> is related to the need for a controlled test environment (PRC<sub>11</sub>) impacting the execution environment to be used in the pipeline. For instance, O<sub>3</sub> mentioned: “*Since the output of the system is sound and the test should check the sound quality it is better to have it in a controlled environment that makes use of simulation.*”



Table 6. Pipeline-related barriers

Category	ID	Barrier	Organizations
Resources	B <sub>1</sub>	Limited human resources	O <sub>8</sub> , O <sub>9</sub>
	B <sub>2</sub>	Limited availability of software and/or hardware resources	O <sub>1</sub> , O <sub>2</sub> , O <sub>3</sub> , O <sub>4</sub> , O <sub>5</sub> , O <sub>6</sub> , O <sub>7</sub> , O <sub>8</sub> , O <sub>9</sub> , O <sub>10</sub>
Domain	B <sub>3</sub>	Complex non-functional requirements	O <sub>6</sub>
	B <sub>4</sub>	Security configuration prevents CD	O <sub>2</sub>
	B <sub>5</sub>	HiL not usable, e.g., for safety or security reasons	O <sub>1</sub> , O <sub>2</sub> , O <sub>8</sub> , O <sub>9</sub> , O <sub>10</sub>

Another test automation challenge is related to oracle specification (PRC<sub>12</sub>), as mentioned by five out of 10 organizations. The impossibility of specifying an automated oracle hinders what kinds of tests one can run in the pipeline. This may happen, for instance, when one needs to evaluate a signal received from a sensor, i.e., “*The main challenges for automatizing the test execution: a good way to model the test itself and have an oracle that can compare with the actual behavior.*” stated by O<sub>3</sub>. This aspect has already been mentioned by Mårtensson *et al.* [52], however, while they only talked about usability testing, we stress more the impediment in automatically determining and checking the test oracles, also for functional testing mainly due to outcome coming from real hardware devices working in a real environment with many external factors to control for, e.g., to check the quality of the acoustic signal coming from sensors (O<sub>3</sub>).

The remaining two challenges are related to difficulties encountered when specifying/deriving integration (PRC<sub>13</sub>) and safety (PRC<sub>14</sub>) tests. As regards the former, O<sub>10</sub> develops prototypes requiring the interconnection of many different sub-components. This makes it difficult to determine the expected system behavior: “*It is quite hard to derive integration test cases due to the complex combination of all different parts.*” As regards the specification of safety tests, in agreement to what indicated by Gautham *et al.* [27], O<sub>4</sub>, O<sub>5</sub>, and O<sub>8</sub> pointed out the complexity to identify situations “*that could never happen.*” or “*that you do not expect to happen.*” Checking for safety requirements is highly important, especially in those domains, such as aerospace and railways, where the safety integrity level (SIL) of the system must be equal to or higher than three.

**Deployment.** This category features two challenges occurring when deploying software on the customers’ side. Having deployment too late in the development process (PRC<sub>15</sub>) may result in installation issues (PC<sub>9</sub> in Table 7), as experienced by O<sub>2</sub>: “*we will not be able to run the software on the system because the installation even does not work on the system, because the update/upgrade does not work, or because the system behavior is not being considered in the early stages of development.*”

Then, there are cases where the deployment is expensive (PRC<sub>16</sub>) in terms of time and effort needed to complete it. This impacts both the type of execution environment adopted within the pipeline, as well as the build triggering strategy. As experienced by O<sub>8</sub> in the railways’ domain, the deployment on a test track requires “*one day with people involved in the testing and on a train a couple of days where many people need to be involved.*”

The late and expensive deployment is strictly related to the CPS nature. Indeed, as already highlighted in Section 4.1, the organizations deploy on real hardware devices only during the last stages of the overall CI/CD process, mainly due to the high costs of the hardware in specific domains such as railways and aerospace.

**Simulators.** The last category, among the process-related challenges, deals with the usage of simulators. O<sub>3</sub> pointed out the presence of scenarios where it is complex to trust the outcome provided by the simulators since there might be many external factors impacting the behavior of the system in a real environment (PRC<sub>17</sub>). Finally, as reported by O<sub>8</sub>, some scenarios cannot rely on simulators. Specifically, if it is complex for a human to specify the expected behavior for some scenarios, of course, it is not possible to rely on simulators that can emulate the same behavior (PRC<sub>18</sub>).

**4.2.2 Barriers for CI/CD pipeline setting and maintaining and related mitigation.** Table 6 summarizes the five barriers encountered by the ten organizations when applying CI/CD to CPSs. These barriers have been grouped into two categories, described in the following.

**Resources.** This category groups the barriers dealing with limited availability of human ( $B_1$ ) and software and/or hardware resources ( $B_2$ ), both influencing the type of execution environment adopted within the pipeline. While we are aware that those barriers can also apply to conventional software systems, the barriers worsen for CPS development, where it is mandatory (i) to rely on simulators, mostly self-developed where you need high expertise about the domain, and (ii) to use HiL that is very expensive in particular CPS domain such as railways and aerospace. For instance,  $O_8$  mostly relies on HiL due to limited availability of human resources having the skills needed to develop/configure simulators – “*given the needs and the budget of our company, it’s much better for more complex scenarios to rely on the hardware in the loop and only use simulations when whatever needs to be simulated is very simple.*”

All the interviewed organizations reported the limited availability of software and hardware resources. Specifically,  $O_6$  mentioned: “*Based on the fact that in the avionics domain the cost of the hardware is very expensive, we do most of the work in simulated environments*”, while  $O_7$  stated that “*Resources for the hardware devices (hardware test tracks and testbeds as real trains) represent an issue for us. We have a limited number of test tracks.*”

As reported in Table 8, the analysis of the interviews’ transcripts (see Table 8) has elicited two mitigation strategies: (i) prioritize and select the test cases to be included within the pipeline (*i.e.*, “*Some strategies rely on genetic algorithms to optimize the resources available for the testing execution environment*” from  $O_1$ ), and (ii) adopt incremental builds mainly relying on impact analysis, as reported by  $O_2$  – “*for what concerns rolling builds we try to limit the amount of testing being executed in them to be as fast as possible.*” The member-checking survey confirms the previous findings, and, as shown in Table 8, six out of ten organizations ( $O_1, O_2, O_5, O_6, O_7, O_{10}$ ) report to rely on test prioritization, while  $O_3, O_4, O_8$ , and  $O_9$  consider it useful while having never used it. As regards the adoption of incremental builds, instead,  $O_1, O_2, O_4, O_7$  and  $O_9$  mention its adoption, while  $O_8$  considers it a useful approach to deal with limited hardware/software resources.

Alternative solutions reported in the member-checking survey to cope with limited availability of resources are “*architectural changes with improved testing concepts*” ( $O_7$ ), and, unsurprisingly, “*platform virtualization*” ( $O_5$ ).

**Domain.** This category includes three different barriers, two of them highlighted by only one organization. Specifically,  $B_3$  and  $B_4$  are related to difficulties arising when automating certain phases in the CI/CD pipeline. For instance,  $O_6$  had to cope with the use of a real-time operating system which made task automation difficult, *i.e.*, “*the complexity of integrating within the pipeline the execution of nonfunctional testing and system testing*”, while  $O_2$  could not implement automated deployment due to security policies for the healthcare domain: “*We cannot deploy at the moment because a change in the security configuration of the software prevented our standard [deployment] process.*”

$B_5$  is related to coping with a complex execution environment. Specifically,  $O_{10}$  mentions that they could not integrate HiL in the CI/CD pipeline for safety reasons, and adopts simulation/mocking for the hardware devices to overcome it. As shown in Table 8, all the organizations facing this barrier used the same mitigation strategy to deal with it. Furthermore,  $O_2$  mentions the possibility to rely on “*digital twin hardware that avoids the safety issues (no moving parts, no radiation) but simulates the hardware to some much better*”.

**4.2.3 Pipeline-related challenges and related mitigation.** Table 7 summarizes the pipeline-related challenges faced by the 10 organizations. The challenges have been grouped into five categories, each one related to a specific aspect of the CI/CD pipeline setting and evolution, *i.e.*, pipeline properties, thoroughness, simulators, HiL, and flaky behavior. In the following, we discuss each identified challenge, together with some examples from the study participants’ experiences, and related mitigation strategies.

Table 7. Pipeline-related challenges

Category	ID	Challenge	Organizations
Pipeline Properties	PC <sub>1</sub>	Long build execution time	O <sub>1</sub> , O <sub>2</sub> , O <sub>4</sub> , O <sub>5</sub> , O <sub>6</sub> , O <sub>7</sub> , O <sub>8</sub>
	PC <sub>2</sub>	Build time estimation	O <sub>9</sub>
	PC <sub>3</sub>	Static code analysis tools configuration	O <sub>3</sub> , O <sub>7</sub>
	PC <sub>4</sub>	Lack to access the production code from the pipeline	O <sub>1</sub>
	PC <sub>5</sub>	CI/CD configuration highly coupled with the environment	O <sub>2</sub> , O <sub>5</sub>
	PC <sub>6</sub>	Reusability of build artifacts	O <sub>2</sub>
Thoroughness	PC <sub>7</sub>	Development environment detached from the execution environment	O <sub>1</sub>
	PC <sub>8</sub>	Detecting deployment-related errors	O <sub>2</sub> , O <sub>6</sub>
	PC <sub>9</sub>	Continuous installation	O <sub>2</sub> , O <sub>3</sub> , O <sub>4</sub> , O <sub>5</sub> , O <sub>7</sub> , O <sub>8</sub>
	PC <sub>10</sub>	Closing the loop introduces performance degradation	O <sub>5</sub>
	PC <sub>11</sub>	Complexity in closing the loop due to uncontrollable factors	O <sub>4</sub> , O <sub>9</sub>
	PC <sub>12</sub>	Complexity in closing the loop due to data collection from the field	O <sub>5</sub>
Simulators	PC <sub>13</sub>	Limited in their functionality	O <sub>1</sub> , O <sub>2</sub> , O <sub>4</sub> , O <sub>5</sub> , O <sub>7</sub> , O <sub>8</sub> , O <sub>9</sub> , O <sub>10</sub>
	PC <sub>14</sub>	Functional correctness	O <sub>5</sub> , O <sub>6</sub> , O <sub>7</sub> , O <sub>10</sub>
	PC <sub>15</sub>	Deal with real-time properties	O <sub>5</sub> , O <sub>9</sub>
	PC <sub>16</sub>	Interaction with the environment	O <sub>2</sub> , O <sub>3</sub> , O <sub>4</sub> , O <sub>5</sub> , O <sub>6</sub> , O <sub>7</sub> , O <sub>8</sub> , O <sub>9</sub>
	PC <sub>17</sub>	Accessibility	O <sub>1</sub> , O <sub>5</sub> , O <sub>7</sub> , O <sub>9</sub> , O <sub>10</sub>
HiL	PC <sub>18</sub>	Availability	O <sub>10</sub>
	PC <sub>19</sub>	Automated deployment on HiL	O <sub>7</sub> , O <sub>8</sub> , O <sub>9</sub>
	PC <sub>20</sub>	Test Automation on HiL	O <sub>2</sub> , O <sub>4</sub> , O <sub>6</sub> , O <sub>7</sub> , O <sub>9</sub>
	PC <sub>21</sub>	Costs and scalability	O <sub>1</sub> , O <sub>2</sub> , O <sub>3</sub> , O <sub>4</sub> , O <sub>5</sub> , O <sub>7</sub> , O <sub>8</sub> , O <sub>9</sub>
	PC <sub>22</sub>	Dependency installation	O <sub>4</sub>
Flaky Behavior	PC <sub>23</sub>	Features' interaction	O <sub>2</sub>
	PC <sub>24</sub>	HiL availability	O <sub>10</sub>
	PC <sub>25</sub>	HiL inputs	O <sub>5</sub> , O <sub>10</sub>
	PC <sub>26</sub>	Lack of control over resources	O <sub>2</sub> , O <sub>4</sub> , O <sub>5</sub> , O <sub>6</sub> , O <sub>7</sub> , O <sub>9</sub> , O <sub>10</sub>
	PC <sub>27</sub>	Network issues	O <sub>1</sub> , O <sub>2</sub> , O <sub>4</sub> , O <sub>5</sub> , O <sub>6</sub> , O <sub>7</sub> , O <sub>9</sub> , O <sub>10</sub>
	PC <sub>28</sub>	Timing issues	O <sub>4</sub> , O <sub>10</sub>

**Pipeline Properties.** This category accounts for six different challenges, two of which deal with the build execution time (PC<sub>1</sub> and PC<sub>2</sub>), while the remaining four are related to the overall pipeline configuration. Four out of 10 organizations faced long build execution time, influencing the type of tasks automatized within the pipeline. For example, O<sub>6</sub> mentioned: “*Slow builds hinder the inclusion of running non-functional testing in the pipeline.*” While this is also considered a relevant challenge for conventional applications [14, 77, 85], for CPSs the problem can be further exacerbated when deploying and executing software on simulators or HiL. The latter confirms what is already found by Mårtensson *et al.* [52] highlighting how working with a highly integrated (tightly coupled) system, a small delivery to the main track may cause building and linking of a large part of the system resulting in long build times. The latter has been also mentioned by O<sub>2</sub> where there is a single integration branch where the components developed by their 70 teams are integrated into a single join point: *i.e.*, “*each component has a test service so running unit tests is very fast but we have a huge amount of high-level testing that is easy to write but kills us in terms of execution time.*” By looking at the result of the survey (see Table 8), the interviewed organizations mentioned a wide set of actions to deal with the above challenge. One possibility is to prioritize and select only a subset of test cases in the test suite to be executed (used also by O<sub>1</sub>, O<sub>2</sub>, and O<sub>7</sub>, and considered a useful action by O<sub>4</sub> and O<sub>8</sub>). A different approach, highlighted by O<sub>2</sub>, deals with the introduction of parallelization within the overall build process, *i.e.*, “*We have 20 test machines in parallel for managing the overall test size, especially for nightly builds.*” The latter is also used by O<sub>4</sub>, O<sub>5</sub>, O<sub>6</sub>, and O<sub>7</sub>, while O<sub>8</sub> only felt it as useful. It is also possible to run the whole build process only within nightly builds, even if this may be controversial since it defeats the CI/CD purpose [13]. However, this is considered acceptable for O<sub>1</sub>, as its pipeline is limited in scope, *i.e.*, used only for V&V purposes. Also O<sub>2</sub>, O<sub>5</sub>, O<sub>6</sub>, and O<sub>7</sub> rely on nightly builds to execute time-intensive

Table 8. Relations between challenges/barriers and mitigation strategies as seen from the semi-structured interviews and the member-checking survey. In bold there are the organizations that do not rely on the mitigation while considering it a useful solution.

Challenge/Barrier	Mitigation	Organizations
B <sub>2</sub> : Limited hw/sw resources	Test Prioritization	O <sub>1</sub> , O <sub>2</sub> , <b>O<sub>3</sub></b> , <b>O<sub>4</sub></b> , O <sub>5</sub> , O <sub>6</sub> , O <sub>7</sub> , <b>O<sub>8</sub></b> , <b>O<sub>9</sub></b> , O <sub>10</sub>
	Incremental Builds	O <sub>1</sub> , O <sub>2</sub> , O <sub>4</sub> , O <sub>7</sub> , <b>O<sub>8</sub></b> , O <sub>9</sub>
B <sub>5</sub> : Domain hinders HiL	Rely on sim./mock-up	O <sub>1</sub> , O <sub>2</sub> , O <sub>8</sub> , O <sub>9</sub> , <b>O<sub>10</sub></b>
PC <sub>1</sub> : Long build	Test Prioritization	O <sub>1</sub> , O <sub>2</sub> , <b>O<sub>4</sub></b> , O <sub>7</sub> , <b>O<sub>8</sub></b>
	Adopt Parallelization	O <sub>2</sub> , O <sub>4</sub> , O <sub>5</sub> , O <sub>6</sub> , O <sub>7</sub> , <b>O<sub>8</sub></b>
	Nightly Builds	O <sub>1</sub> , O <sub>2</sub> , O <sub>4</sub> , O <sub>5</sub> , O <sub>6</sub> , O <sub>7</sub> , <b>O<sub>8</sub></b>
	Incremental Builds	O <sub>2</sub> , <b>O<sub>4</sub></b> , O <sub>5</sub> , <b>O<sub>6</sub></b> , O <sub>7</sub> , <b>O<sub>8</sub></b>
PC <sub>9</sub> : Continuous installation	Containerization	<b>O<sub>3</sub></b> , O <sub>4</sub> , O <sub>5</sub> , <b>O<sub>7</sub></b> , <b>O<sub>8</sub></b>
PC <sub>13</sub> : Sim. limited func.	Combine sim. and HiL	O <sub>4</sub> , O <sub>7</sub> , O <sub>8</sub> , O <sub>9</sub> , O <sub>10</sub>
PC <sub>16</sub> : Sim. coupled with env.	Combine sim. and HiL	O <sub>2</sub> , O <sub>3</sub> , O <sub>4</sub> , O <sub>8</sub> , O <sub>9</sub>
PC <sub>17</sub> : Sim. accessibility	Timeout	O <sub>1</sub> , O <sub>5</sub> , <b>O<sub>7</sub></b> , <b>O<sub>10</sub></b>
PC <sub>21</sub> : HiL costs and scalability	Combine sim. and HiL	<b>O<sub>1</sub></b> , O <sub>2</sub> , <b>O<sub>3</sub></b> , O <sub>4</sub> , <b>O<sub>5</sub></b> , O <sub>7</sub> , <b>O<sub>8</sub></b> , O <sub>9</sub>
	Green-build rule	O <sub>2</sub> , O <sub>3</sub> , O <sub>4</sub> , O <sub>7</sub>
PC <sub>26</sub> : No resources' control	Fix the code	O <sub>4</sub> , O <sub>6</sub> , O <sub>7</sub> , O <sub>9</sub> , O <sub>10</sub>
	Fix pipeline config.	O <sub>6</sub> , O <sub>7</sub> , O <sub>9</sub>
PC <sub>27</sub> : Network issues	Retry	O <sub>2</sub> , O <sub>4</sub> , O <sub>5</sub> , O <sub>6</sub> , O <sub>7</sub> , O <sub>10</sub>

tasks, while adopting incremental builds during working hours (O<sub>2</sub>, O<sub>5</sub>, and O<sub>10</sub>). The latter is also used by O<sub>7</sub> and O<sub>8</sub>, while O<sub>4</sub> and O<sub>6</sub> consider the mitigation useful even if they have never adopted it.

A different challenge, experienced by O<sub>9</sub>, that can also apply to conventional systems, while it is more critical for CPSs, is related to the build time variability (PC<sub>2</sub>), due to the adopted infrastructure “*since our platform works in the cloud we need to know how much time it is required to acquire and elaborate a huge amount of data points*”.

Moving to the overall pipeline configuration, in the absence of clear coding standards or guidelines, the adoption of code style checking tools becomes problematic, if not unfeasible (PC<sub>3</sub>). In this scenario, approaches for coding style inference may be desirable [61, 83]. Similar considerations apply to bug-finding tools, sometimes inapplicable to CPSs for automating code review, as experienced by O<sub>7</sub>: “*we need expertise on the developers' side for determining whether or not a train is behaving in the expected way.*” The latter is strictly related to PRC<sub>2</sub> where, in the presence of safety-critical systems, like the ones in the aerospace and railways domains, it is very difficult to find skilled experts in the domain from both the hardware and software viewpoints.

The lack of access to production code (as experienced by O<sub>1</sub>) limits the ability to properly set static analysis or testing tools (PC<sub>4</sub>) — “*One big challenge is that we need to guarantee the protection of the source code: How to*

*test a component without having its production code?*” The latter is a specialization of the restricted access to information due to security aspects impediment found by Mårtensson *et al.* [52]. On the same line, there is a challenge (PC<sub>5</sub>) related to the extent to which technology restrictions, or restrictions coming from the application domain, may impact the pipeline setting. For instance, O<sub>2</sub> mentioned that “*the Windows situation does not help us with dockerization*”, and at the same time, they are having trouble in properly configuring the CI/CD pipeline for CPS since “[*they*] need to follow medical application frameworks providing a base set of rules in terms of how to build applications and how to integrate them.” The latter results in the last challenge related to the impossibility to reuse previously built artifacts (PC<sub>6</sub>) in the integration branch (*i.e.*, O<sub>2</sub> mentioned: “*It’s a huge pain that we do not reuse artifacts*”), mainly due to constraints imposed by the domain.

**Thoroughness.** This category groups six challenges related to (i) ensuring the overall accuracy and completeness of the CI/CD pipeline (PC<sub>7</sub>, PC<sub>8</sub>, PC<sub>9</sub>) scattered across eight organizations, and (ii) closing the DevOps loop by gathering data from the hardware, *i.e.*, PC<sub>10</sub>, PC<sub>11</sub> and PC<sub>12</sub> experienced by three out of 10 organizations.

O<sub>1</sub> faces a challenge related to having a development environment detached from the execution environment (PC<sub>7</sub>). Another challenge (PC<sub>8</sub> experienced by O<sub>2</sub> and O<sub>6</sub>) occurs in the presence of incremental deployment, which makes it difficult to detect and isolate deployment errors. Furthermore, O<sub>6</sub> reported how this even makes it necessary to reconfigure the entire pipeline — “*you deploy blocks, if there is an error in one of the blocks detecting it and reconfigure and reset the pipeline is a problem.*” Finally, continuous installation (PC<sub>9</sub>) cannot be achieved due to the late deployment strategy (PRC<sub>17</sub>). This is because changes to the environment impact the pipeline configuration, which needs to be adapted every time. For what concerns continuous installation problems, O<sub>2</sub>, O<sub>3</sub>, O<sub>4</sub>, O<sub>5</sub>, O<sub>7</sub> and O<sub>8</sub> have encountered them, with O<sub>4</sub> pointing out that using containerization it is possible to facilitate the switching between software versions to deploy, meaning that it will be possible to handle the variability of the environment in terms of dependencies. As shown in Table 8, containerization is also used by O<sub>5</sub>, while O<sub>3</sub>, O<sub>7</sub> and O<sub>8</sub> consider it a viable solution.

Moving on to the need for closing the DevOps loop, the interviews indicated three different challenges hindering the acquisition of data from the physical environment (or hardware device). Working in a CPS context implies having a tight interaction with multiple hardware devices, *i.e.*, sensors and actuators, in which gathering data from them could be problematic due to the presence of many external environmental factors that must be taken into account, as well as the need for having invasive measurement instruments directly in the field. Specifically, O<sub>5</sub> stressed the introduction of performance degradation (PC<sub>10</sub>) due to invasive measurement instruments: “*The challenge is that monitoring becomes invasive with respect to the system performance.*”, as well as the presence of noise in the collected data (PC<sub>12</sub>): “*There are architectural ways to deal with that so that if some sensor does not update on time, you still can make a relatively informed decision. But even then, you have to make sure that the drift is not over a certain size because then you cannot make reasonable decisions anymore.*” O<sub>4</sub> and O<sub>9</sub> highlighted the presence of uncontrollable factors in a CPS execution environment, making it challenging to close the DevOps loop. For instance, O<sub>4</sub> reported: “*Differently from other software applications, there is data that we cannot control such as the presence of something on the floor that the robot is not able to perceive so it will fail. You have to analyze the video data and this is very hard.*”

**Simulators.** This category groups five challenges related to simulators’ issues and limitations stressed more in the CPS domain due to the high environment complexity [74], which very often results in having scenarios that cannot be emulated, such as in the presence of many external environmental factors to be controlled. Specifically, the need to develop them in-house or the lack of specific skills may lead to simulators that are limited in their functionality (PC<sub>13</sub>). For instance, O<sub>8</sub> stated: “*we prefer to spend time in testing on real hardware instead of spending time in developing complex simulators*”, while O<sub>4</sub> reported: “*Walking is not so easy to simulate, so we need a real walking robot for spotting bugs.*” As shown in Table 8, it is a common habit to adopt a pipeline that relies on both simulators and HiL in different build stages to overcome the above challenge. A clear example of this happens in

O<sub>7</sub>, where there is a build process made up of three different build stages, each one adopting a specific execution environment (see Section 4.2).

A lack of knowledge about the device/system to simulate can lead to wrong assumptions, affecting the simulator's correctness (PC<sub>14</sub>) as experienced within O<sub>10</sub>: *"This happens more at the beginning of a project when you are not too familiar with the device and you make assumptions on how it works."* These problems might have an impact on the whole CI/CD pipeline setting and trustworthiness, because it is possible to have deviations of the monitored system behavior between the real hardware and simulators.

As experienced by O<sub>5</sub>, the limited capability to simulate real-time properties (PC<sub>15</sub>) hinders the applicability of simulators or at least raises the need for further tests on HiL. The latter is also confirmed by O<sub>9</sub>: *"for what concerns the simulation for the RFID we think that the simulation will not give us any benefits due to their unpredictable behavior."*

Likewise PC<sub>13</sub>, the high level of interaction between different components (PC<sub>16</sub>) forces organizations to directly test feature interaction by using real devices, instead of simulating them. Indeed, when using simulators for CPSs it is important to remark that they have to interact with a too complex environment that must be simulated as well. As an example, O<sub>6</sub> mentions problems faced when simulating a car behavior *"for the CAN data, what do you want to wish to happen here? If you are driving around something you need to know how fast the wheels are turning, as well as what the engine revolutions are together with other sensitive data you might pick up over the canvas. There are a lot of details that are very application dependent."* Also in this case, as shown in Table 8, organizations rely on pipeline configurations including different execution environments, *i.e.*, five out of eight organizations facing the challenge declare that this is a useful mitigation strategy (O<sub>2</sub>, O<sub>3</sub>, O<sub>4</sub>, O<sub>8</sub>, O<sub>9</sub>).

If an organization has to test third-party software, as in the case of O<sub>1</sub>, there may be the need to run the simulated environment on a remote machine which may turn out problematic to be properly integrated into a local pipeline (PC<sub>17</sub>), due to network security restrictions. Such a scenario typically occurs in the development of safety-critical systems (which very often are CPSs), because the software needs to be tested by somebody different from the development organization. To deal with this problem, O<sub>1</sub> mentions the usage of "timeout" within the pipeline. As shown in Table 8, O<sub>1</sub> and O<sub>5</sub> handle external simulator unavailability through timeouts, while O<sub>7</sub> and O<sub>10</sub> consider this useful yet they do not use it. O<sub>1</sub> also mentions they often *"request some customization at the customer side of their simulators. Sometimes it is accepted, most of the times not."*

**HiL.** This category groups four challenges related to issues and limitations of using HiL in the CI/CD pipeline. As shown in Table 7, three out of four challenges in this category are experienced by multiple organizations, while PC<sub>18</sub> is organization-dependent. Specifically, O<sub>10</sub> faces problems with checking hardware availability before running tests (PC<sub>18</sub>): *"One of the biggest problems, when any particular hardware is involved, is that the hardware may either not be available, or it may be switched off"*.

From a different perspective, as experienced by O<sub>7</sub>, O<sub>8</sub>, and O<sub>9</sub>, deployment on HiL may be challenging (PC<sub>19</sub>). Specifically in O<sub>8</sub> *"remote installation cannot be used with real systems"*, while in O<sub>9</sub> *"The other challenge is related to having a fully automated deployment over the customers' server in which it is possible to have full control on what is going on and try to identify, as soon as possible, failures/errors occurring during the deployment."*

Testing on HiL (PC<sub>20</sub>) is considered very demanding to achieve. O<sub>2</sub> reports: *"If you translate test strategies to the hardware it is very demanding."* and this is mostly a consequence of limited human resources being available. However, there are cases where testing on HiL is constrained by the high cost and lack of scalability (PC<sub>21</sub>) of the hardware devices/systems *"This costs and does not scale"* for O<sub>2</sub>, or *"it is very costly to test on trains"* for O<sub>7</sub>.

As shown in Table 8, the study participants identified two possible strategies to deal with these cost and scalability problems: (i) relying on a mixed pipeline where continuous builds run on simulators and some periodic builds on HiL (used by O<sub>2</sub>, O<sub>4</sub>, and O<sub>9</sub>, and considered useful by O<sub>1</sub>, O<sub>3</sub>, O<sub>5</sub>, and O<sub>8</sub>), or (ii) adopting the green build rule when transitioning between simulators and HiL [15], as highlighted by O<sub>7</sub>: *"Only when the tests in the"*

*virtual train are green can we move to the next step.*”, and also used by O<sub>2</sub>, O<sub>3</sub> and O<sub>4</sub>. The alternative would be, as pointed out by the O<sub>5</sub> survey respondent, *“working with virtual devices instead of real hardware devices.”*

**Flaky behavior.** This category accounts for seven different root causes that may lead to non-determinism in the build execution used for CPS development. Flakiness related to non-determinism during test execution [89] has been largely studied [16, 46, 48, 58, 90] and approaches to detect and cope with it have been proposed [47, 49, 59, 65, 87]. While similar to conventional software, dependency installation within the pipeline (PC<sub>22</sub>) may result in pipelines having a flaky behavior, e.g., for O<sub>4</sub> *“ROS uses GitHub repositories for dependency resolution so when GitHub or the repositories are down our build jobs will fail due to the impossibility of resolving dependencies”*, or else little control over external resources (PC<sub>26</sub>), e.g., *“the most important root cause we experienced is related to the load on the server-side”*, the root causes behind flaky behavior in CPSs may be different from conventional software. Specifically, a CI/CD pipeline for CPSs can suffer from flakiness due to:

- The complex interacting environment (PC<sub>23</sub>), i.e., CPSs are systems of systems with tight interactions among different components, e.g., for O<sub>2</sub> *“the complexity of [the] subsystems whose features interact across many indirections may lead to non-deterministic behaviors”*;
- HiL unavailability (PC<sub>24</sub>), where without a proper check of the availability of hardware, the build outcome might fail intermittently since the pipeline was not able to properly communicate with the device, i.e., O<sub>10</sub> reported: *“We experienced flakiness in terms of non-deterministic behavior mainly due to hardware not being available”*. In this specific scenario, it is important to properly discriminate between intermittent failures caused by communication issues with the HiL from failures due to wrongly implemented functionality;
- Presence of noise in the measurements (PC<sub>25</sub>) when using HiL i.e., difficulty in removing the effect of external environmental factors from the data read from the sensors, as experienced by O<sub>5</sub> and O<sub>10</sub>. Specifically, for O<sub>10</sub> *“Other times the charge level that you read out would go a little bit higher or there is noise in the measurements”*, while for O<sub>5</sub> *“you need to understand what your sensors are sensing and what the acceptable range of inputs are”*;
- Network issues (PC<sub>27</sub>) where, for instance, glitches in the network lead to a connections being lost as reported by O<sub>10</sub>, stressed more in the CPS domain where you need to control among the communication occurring across a huge number of different hardware devices operating in a complex environment;
- Simulators not coping with timing issues (PC<sub>28</sub>), e.g., O<sub>10</sub> stated: *“the last problem is related to multi-threaded programming”*.

For what concerns flakiness mitigation, as highlighted in Table 8, when the problem is related to the lack of control over resources (PC<sub>26</sub>), the solutions adopted are (i) to change and fix the pipeline configuration, i.e., O<sub>7</sub> stated: *“The misbehavior is reported back to the integration team responsible for the Jenkins configuration to find a solution.”*), as well as (ii) to fix the root cause of the flaky behavior within the code: *“to not experience it anymore in the system”* from O<sub>2</sub>. When the root cause of the flaky behavior is in the networking (PC<sub>27</sub>), the organizations leverage the “usual” retries (O<sub>2</sub>, O<sub>4</sub>, O<sub>5</sub>, O<sub>7</sub>, O<sub>10</sub>): e.g., *“of course we have some retry for network issues”* for O<sub>4</sub>, or *“For what concerns flaky connections, you have to be concerned about missed messages and retries”* for O<sub>5</sub>, O<sub>6</sub>, instead only considers it a viable solution. Furthermore, the respondent belonging to O<sub>2</sub> mentioned as an alternative solution the *“introduction of quarantine builds together with an appropriate process of how to deal with these tests”*.

#### 4.3 RQ<sub>2</sub>: How relevant are the identified CI/CD challenges/barriers and their mitigation for practitioners involved in CPS development?

This research question describes the results of the evaluation of the findings in RQ<sub>1</sub> made through an external survey leveraging practitioners that have not been involved in the semi-structured interviews. Note that we have only validated the barriers and the pipeline-related challenges together with their associated mitigation strategies.

As regards the five barriers encountered when trying to configure a CI/CD pipeline for CPS development, by looking at the results in Figure 3, we found that among the participants who answered each question, the limited number of human and software/hardware resources together with the presence of complex non-functional requirements to be checked within the pipeline are the ones felt as more relevant ( $> 72\%$ ). Furthermore, while 30 out of 55 respondents still consider as relevant the barriers dealing with security aspects hindering the inclusion of HiL in the CI/CD process, 31% do not consider such barriers as a real impediment. All the three mitigation previously identified were considered relevant by the survey participants. Specifically, the adoption of test case prioritization techniques is predominant (31 out of 55 respondents), followed by the usage of simulators or mock-ups (28), and the usage of incremental builds (17).



Fig. 3. Results of barriers perception

Figure 4 shows the results of the survey in terms of the six challenges belonging to the Pipeline Properties category. Unsurprisingly, 47 out of 55 respondents consider the long build execution time as a relevant challenge. Also, while from the semi-structured interviews the remaining five challenges were experienced by one or at most two different organizations, the survey indicates how some of such challenges are felt as relevant by more than 69% of our participants. These are (i) the need to properly estimate the build time before timing out the CI/CD process, (ii) the difficulty in properly configuring static code analysis tools, and (iii) the presence of a CI/CD configuration highly coupled with the environment. Regarding the impossibility of having access to the production code, if we do not consider the seven participants reporting that this challenge cannot apply to their context,  $\approx 31\%$  of the respondents do not consider it as a relevant challenge.

Moving onto the mitigation strategies, more than half of our respondents (30) rely on test case prioritization techniques, 28 rely on parallelization, 26 rely on nightly builds for time-intensive tasks, while 18 consider useful the adoption of incremental builds during normal working hours. Finally, one participant reported a new mitigation strategy dealing with long builds where “we simulate faster than in the reality where possible”, however, the same participant also points out the drawback of this mitigation, i.e., having different build outcomes when using simulators and HiL—“this can introduce subtle timing differences in the test results”.

For what concerns pipeline-related challenges in the Thoroughness category dealing with ensuring the overall accuracy and completeness of the CI/CD pipeline (see PC<sub>7</sub>, PC<sub>8</sub> and PC<sub>9</sub> in Figure 5), differently from the RQ<sub>1</sub> results, more than half of our survey respondents consider the presence of a development environment detached



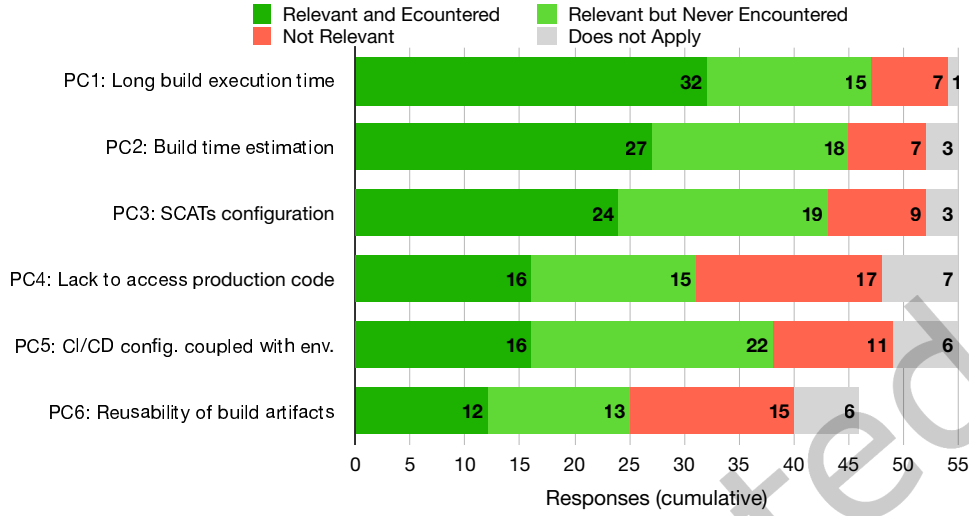


Fig. 4. Results of pipeline challenges perception: Pipeline Properties

from the execution environment as a real impediment to set up a CI/CD process for CPSs. This is also true for the difficulties in detecting deployment-related errors (39 respondents). The above differences stress the impossibility to have a “standardized” CI/CD configuration that can be applied to almost all the CPS domains. While all participants considered the adoption of containerization a viable solution to overcome these challenges, one new mitigation strategy comes up from a survey participant which, for PC<sub>8</sub>, suggests the possibility of developing and adopting static analysis tools able to analyze (and detect errors from) deployment scripts.

Moving to the three challenges related to closing the DevOps loop by gathering data from the hardware, *i.e.*, real environment, by looking at the bottom part of Figure 5 it is possible to state that more than 65% of the respondents consider them as relevant, with the presence of uncontrollable factors to account for having the highest percentage ( $\approx 71\%$ ). While from RQ<sub>1</sub> we did not find any mitigation strategy for these challenges, we obtained some feedback from eight survey respondents. First of all, it could be possible to continuously analyze the logs also after the operation has started. At the same time, one survey respondent points out the possibility to make the monitoring less impactful on performance by “*disabl[ing] invasive logging methods*”. For what concerns the presence of uncontrollable factors, one respondent pointed out how using continuous testing allows to “*better overcome the problem of the uncontrollable factors in real life systems and usually diminish the future costs and improve efficiency*”. There are also mitigation strategies dealing with the overall CI/CD process. Specifically, one respondent mentions the possibility to use parallel DataOps observability pipeline, *i.e.*, “*we use ELK, but it is still under debate/migration*”. While a different respondent highlights as possible mitigation the presence of “*cross-functional teams which bring in more collaborations and ideas*”.

Figure 6 shows the results of the external validation survey for the challenges dealing with the inclusion of simulators in the CI/CD process. As can be seen from the figure, at least 36 out of 55 respondents considered such challenges relevant. The only exception is the challenge of dealing with the impossibility of accessing the third-party simulators adopted in the pipeline (PC<sub>17</sub>). In this case, 13 respondents mention that it does not apply to their context, meaning that the simulators are mainly self-developed within the organization they belong to. In terms of mitigation, instead, (i) 28 respondents use a CI/CD process made up of both simulators and HiL for overcoming the presence of limited functionality, (ii) 24 use both simulator and HiL for overcoming the

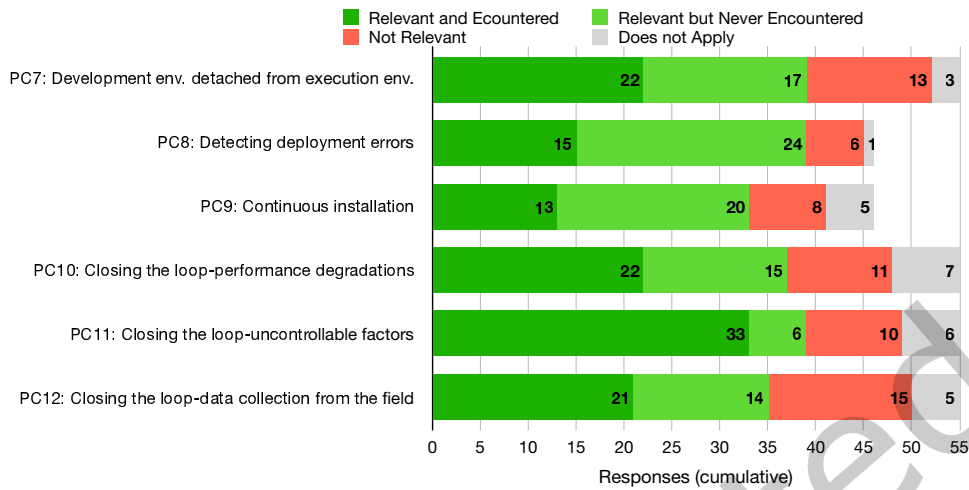


Fig. 5. Results of pipeline challenges perception: Thoroughness

complexity due to a tight interaction among different components and the environment, and (iii) among the 31 respondents struggling with the simulators' accessibility, 14 adopt the "timeout" feature.




Fig. 6. Results of pipeline challenges perception: Simulators

As regards the four challenges dealing with the inclusion of HiL in the CI/CD process, as shown in Figure 7, the costs and scalability challenge is the predominant one (49 out of 55 respondents), followed by the need to check for HiL availability (41), and the complexity for automating both deployment and testing activities on HiL (42 and 40 respondents for PC<sub>19</sub> and PC<sub>20</sub> respectively). No new mitigation strategy comes up from the survey results. However, 28 respondents confirm that the adoption of simulators and HiL in different build stages can help to deal with costs and scalability issues, while 19 adopt the "green-build" rule, *i.e.*, HiL can only be considered when the CI/CD process relying on simulators has a green status.

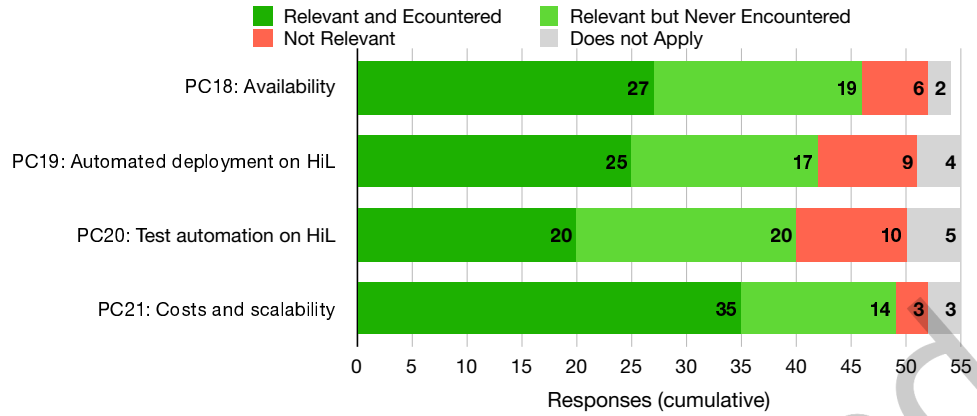


Fig. 7. Results of pipeline challenges perception: HiL

The last category of pipeline-related challenges being validated through the external survey considers the root cause for flaky behavior experienced in the CI/CD process. As shown in Figure 8, for each challenge we have that more than half of our respondents consider it relevant for CPSs. Moreover, 49 out of 55 respondents consider challenging to deal with HiL availability and simulators not coping with timing issues. Unsurprisingly, the two challenges being not specific to the CPS development, *i.e.*, PC<sub>22</sub> and PC<sub>26</sub>, are the ones where several respondents (13 and 12 respectively) mentioned that it is not relevant. Fixing the pipeline configuration is the most frequent mitigation strategy, as indicated by 31 respondents, while no further mitigation strategies are suggested.

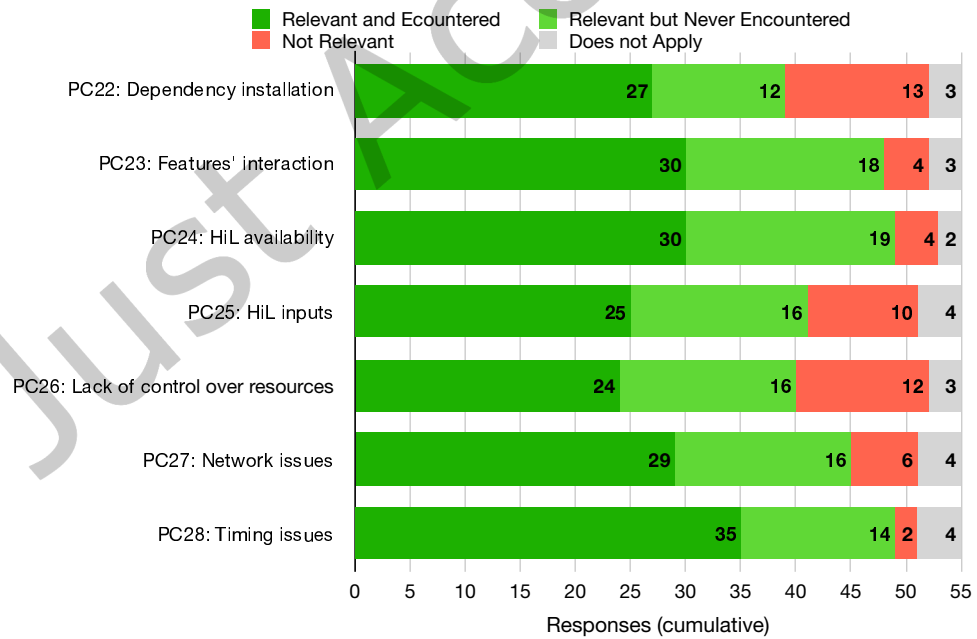


Fig. 8. Results of pipeline challenges perception: Flaky Behavior

Finally, by looking at the 15 answers to the open-ended question aimed at eliciting other challenges that we did not encounter in the semi-structured interviews, we gathered the following, additional challenges:

- (1) *Guaranteeing the supply chain security* (three respondents);
- (2) *The impossibility to use simulated environments* unless the quality of specific data types is ensured (one respondent);
- (3) *The need to have field tests included in the DevOps cycle even with a lower frequency* as both simulators and HiL can only cover a fraction of what usually happens during real field testing activities (one respondent);
- (4) *The difficulty to implement a “quick-retry” feature* in the CI/CD process, to selectively rollback at specific stages mainly because this is highly dependent on the infrastructure language (one respondent); and
- (5) *The difficulty to reduce the build execution time when dealing with HiL due to the need for checking the HiL availability, i.e., “very long hardware boot times”* (one respondent).

## 5 DISCUSSION AND IMPLICATIONS

This section summarizes the main findings and implications of our study. We divide the section into implications for (i) developers, (ii) educators, and (iii) researchers.

### 5.1 Implications for developers

We start by discussing what, based on insights learned from this study, developers must consider when trying to set up and evolve a CI/CD pipeline for CPS development.

**Simulators are necessary to achieve continuous builds on CI/CD pipelines.** Performing CI/CD on real hardware is often unfeasible, for different reasons. The automated deployment may be complicated, or the hardware may not be available on-site. Also, organizations doing V&V tasks only may have limited/no access to hardware, simulators, or even to the production code. Therefore, simulation is often the only choice available. However, having a reliable simulator is challenging for many CPS developers. While in some cases simulators come from hardware producers, in other circumstances the only option is to develop them in-house. This requires the allocation of suitable skills and efforts in the development process. Failing to do so would have severe consequences on the ability to setup not only CI/CD, but even simple test automation without relying on the hardware directly, when this is possible.

**Balancing the use of simulators and HiL in the pipeline.** Deploying and running CPSs on HiL at every change could be troublesome, expensive, and may result in slow feedback. At the same time, for the reasons mentioned before, it is unlikely that developers could fully trust a quality assessment performed solely on simulators. Therefore, it is highly desirable to configure staged builds relying on different execution environments, namely (i) continuous builds on simulators, aimed at providing fast feedback to developers (e.g., about the outcome of static checks, or possible integration issues discovered by tests); and (ii) periodic (e.g., nightly) builds on HiL, to verify whether the assumptions made on simulators are still valid, checking properties that often cannot be verified on simulators (e.g., response time properties), testing the system in scenarios that cannot be easily simulated, or verifying the compatibility of the software against hardware variants not fully reproduced by the simulators.

**Late delivery is the crux of CPS development.** For the reasons explained above, CPS software tends to reach target production hardware very late in the development. This has several negative side effects, including the late discovery of defects that could not be identified through simulation, but also having a system that reaches the end user very late. Allocating sufficient effort, resources, and competences to enable automated delivery is therefore highly desirable.

**Having hardware experts on-board may be a plus.** Based on what we learned from this study, it is clear how CPS development may highly benefit from the availability of both software and hardware experts so that it

is much easier to self-develop simulators whose behavior is as much as possible like the one of the real device. This would help to reduce the differences that might be observed in terms of build outcome, *e.g.*, the number and type of failing tests, when running the process on simulators and HiL. On the one hand, the presence of hardware experts in a team, when available, has been found useful by our interviewees. On the other hand, a context in which both hardware and software evolve makes tasks such as change impact analysis more challenging to handle, *e.g.*, to determine whether and to what extent a hardware change would impact some software components, or some software evolution would hinder the integration of certain pieces of hardware.

## 5.2 Implications for educators

In the following, we discuss what, based on insights learned from this study, would be expected for what concerns the creation (or enhancement) of curricula related to CPS development.

**Blended curricula with hardware and software competencies.** Any effective DevOps organizational setting or management of a CI/CD pipeline likely requires software engineering expertise other than what is currently taught in regular graduate-level courses, *e.g.*, knowledge about the hardware, software-hardware interplay, and domain standards expertise. On the one hand, university curricula shall strive to include such aspects in their teaching. On the other hand, practitioners should more actively engage in standardizing CPS application lifecycle management practices, patterns, and tools to enable the aforementioned educational augmentation exercise.

**Specialized courses on simulator development.** In a context for which CPS specific curricula are highly desirable, one competence assumes paramount importance, and this is the development of simulators. The latter requires combining knowledge from physics, automated control (*e.g.*, system dynamics, discrete systems), and virtual reality (many simulators leverage 3D or even virtual reality environments, similar to those used in video games).

**Teaching CI/CD in complex, heterogeneous environments.** CI/CD is oftentimes taught in the context of conventional system development. To favor the adoption of CI/CD for complex systems, and in particular for CPSs, courses on CI/CD should touch on topics related to (i) coping with complex hardware or simulators attached to the pipeline, and (ii) pondering fast builds with the need for testing a CPS on multiple devices (or simulators), where this is appropriate. Also, while conventional CI/CD literature advocates “building at every change” [13], CPS developers need to face with reality, and therefore such a common wisdom need to be revisited. Similarly, we found that for large and complex CPSs “retest all” does not work, and therefore incremental builds are a widely adopted practice.

**Software architectures for CPSs.** CPSs heavily interact with HiL interfaces (and sometimes multiple HiL, having different characteristics and varying APIs) and, during the development process, with simulators. The latter may be updated or even replaced by better ones. From an educational perspective, it is desirable that courses related to software architectures properly treat such scenarios, discussing the proper architectural choices or design choices allowing an easy (even at run-time) replacement of different kinds of HiL and simulators in the software systems. Software components of CPSs may need to be deployed on, or interact with, multiple types of devices (*e.g.*, a control software may be deployed on different car models). This requires that developers must have suitable knowledge of product line engineering and follow related practices when designing CPSs. Furthermore, it is desirable to teach prospective CPS developers about how to design a CPS architecture to make a system scalable, but also secure, and easy to be monitored and tested.

## 5.3 Implications for researchers

Implications for researchers aim at developing approaches and tools to support developers in setting up, maintaining, and using CI/CD pipelines for CPSs.

**The target environment of CPS is multifaceted and diversified, making CI/CD pipelines complex and expensive.** Often a CPS may target multiple devices, as well as both HiL and simulators. This may entail a build matrix against which the pipeline must be run, *i.e.*, the matrix describes different combinations of parameters (*e.g.*, simulator models, HiL instances, other settings) for the build. That being said, it is possible that, while some changes may entail different behavior on different matrix instances, other changes do not, and therefore running the build on all possible configurations would be a waste of resources. On the one hand, this stimulates research towards approaches aimed at recommending the creation of a suitable build matrix system based on similar systems, and in general systems targeting similar devices. Also, these kinds of recommenders should be able to point out the need for maintaining build matrices by learning “from the crowd”, *e.g.*, the need to prune out obsolete environments and add new ones. On the other hand, proper approaches should be developed to trigger builds on different matrix instances based on the changes performed.

**Coping with multiple root causes for flaky behavior.** The complex technological stack, the behavior of simulators and HiL, their (sometimes uncontrollable) unavailability or lack of accessibility, and the mechanisms used to collect test outputs (*e.g.*, sensors or video cameras) require not only to better monitor all possible elements causing flakiness, but also to combine and enhance various mitigation approaches, including checking the status of HiL/simulators, and leveraging the “usual” retries. As indicated by the participants, flaky behaviors in CPSs are often due to the complex interacting environment (*e.g.*, lack of complete control on the hardware status) rather than on the order with which the tests are executed. Hence, flaky test detectors that target flaky tests considering their ordering [24, 88] are not effective for environment-dependent flakiness. CPS-specific detectors could be inspired to those for undetermined specifications [87], or based on machine learning models [59] but trained on CPSs data and encompassing CPS-specific features, *e.g.*, changes to simulators or HiL configurations, as well as their build logs. To improve CI/CD infrastructures for CPS, it may be useful to develop recommenders, integrated into the pipeline, able to support developers in the identification of flakiness behavior, and identify its root causes.

**Challenges in automated test execution.** We found that one of the reasons that impede full CI/CD automation for CPS is the difficulty to automate test execution, especially when the system is deployed on the hardware. That is, the system receives inputs from sensors and interacts with actuators. Full test automation requires (i) tools, such as scenario generators or record replay tools able to seed inputs to the CPS, and (ii) the capability of CI/CD infrastructure to support the execution of such tools. Very often these tools are GUI-oriented and not particularly well-suited to be integrated in a CI/CD pipeline.

**Challenges in automated oracle creation.** The CPS execution environment (*e.g.*, simulators or HiL) drastically complicates the definition and automatic check of oracles. The latter requires to ponder several factors: (i) the test scenario (or requirement to assess); (ii) the accepted level of realism in simulations; (iii) the readiness level or maturity of the hardware proxies used in the pipeline; and (iv) the output sources, *e.g.*, based on actual sensors’ data or mocking/synthetic data. Besides, the oracles consist of value ranges (*e.g.*, time intervals) instead of scalars, or they may be signals that need to be properly processed, as highlighted in existing studies on testing for CPS [6, 53]. It may be important to account for non-functional properties, including timing ones [80]. Also, to cope with inputs originating from sensors or even from a multimedia recording of the CPS execution (as pointed out by  $O_3$  and  $O_4$ ), it is desirable to develop approaches for pattern recognition [28, 38, 72].

**Need for specific fault models.** Looking more broadly at configuring V&V phases within the CPS pipeline, respondents would like to early discover some defects through static analysis. This requires a clear fault modeling in the CPS context (as the ones for autonomous cars [26] and unmanned vehicles [81]), but also to develop CPS-specific linters, which can be integrated into the CI/CD pipelines to allow early detection of build failures, hence avoid to perform expensive testing activities, and hence long builds, which constitute a major problem for CPS developers according to our study results. Moreover, CPS-specific fault models can be useful for other purposes, not only to facilitate root-cause analysis [26], but also to create domain-specific mutation testing strategies, as it has been done in other cases such as deep learning [36] or mobile development [19, 76].

## 6 THREATS TO VALIDITY

Threats to *construct validity* concern the relationship between theory and observations. The interview participants might have misinterpreted our questions, or they might have reported their personal (and biased) views of the phenomenon. While this is typical for interview-based studies [17, 33], we mitigated the threat by using semi-structured interviews and following up with clarifications every time we realized this was needed.

There could be threats to construct validity related to how survey respondents interpreted the survey questions and provided their answers. We have mitigated this threat by providing a self-explanatory description of the challenges, barriers, and mitigation strategies. Also, we left them the possibility to provide open comments to also point out cases of misunderstanding. However, based on the provided answers, we had no evidence of cases where respondents had difficulties in understanding the posed questions. In addition, for the external survey, we leveraged demographic information to filter out responses where the information provided made it evident that a participant had not the required knowledge.

Threats to *internal validity* concern confounding factors that could have influenced our results. To limit subjectivity in our coding, we employed multiple coders, computed inter-rater reliability, and used follow-up discussions not only to resolve cases of inconsistent coding, but to review any single coding. We elicited codes and relations only based on explicit occurrences of words in the transcripts. However, we could not exclude imprecision due to our interpretation of the participants' answers.

Another threat could be the low representativeness of the respondents in the semi-structured interviews and, to some extent, in the external survey. In the first case, participants were obtained through personal contacts, as we need people available to participate in a relatively long interview. However, such participants cover a relatively diversified set of domains (8). As for the survey, the use of snowballing and especially the use of *Prolific* allowed us to mitigate a possible bias due to the direct personal contacts.

Threats to *reliability validity* relate to the extent to which results can be reproduced. To achieve this goal, we (i) have described the data collection and analysis process in detail, and (ii) provide in our replication package the detailed outcome of the coding phases.

Finally, threats to *external validity* concern the generalizability of our findings. The interview-based study has been conducted involving 10 organizations developing CPS for 8 different domains. We are aware that the obtained findings may not generalize to different organizations and domains. Indeed, from the performed interviews, we found that CI/CD pipelines were extremely different from case to case. Therefore, as in other interview studies conducted within a limited set of organizations, and also considering the study topic, the generalizability is relatively limited. To mitigate this threat, when addressing RQ<sub>2</sub> we have validated the findings collected in RQ<sub>1</sub> through an external survey with practitioners different from the ones involved in our semi-structured interviews, and belonging to 9 domains. Still, it is possible that, also in this case, as Figure 2 shows, some domains are better covered than others, as well as some are still not covered at all.

## 7 CONCLUSION AND FUTURE WORK

In this paper, we investigated the adoption, usage, and evolution of CI/CD pipelines for CPS development, by focusing on challenges and barriers that DevOps teams face when setting up or evolving CI/CD processes for CPS development highlighting that the configuration is highly dependent on the domain. The study is based on interviews from 10 organizations developing CPSs in 8 different domains, followed by a member-checking survey within the same development teams, and an external validation survey involving 55 participants from 9 domains. By performing an open coding on the interview results, we have elicited a set of challenges/barriers, along with their mitigation strategies.

The obtained findings are a first step towards supporting DevOps teams in properly using and configuring CI/CD for CPSs. Also, they have implications on how to enhance education/training for CPS developers, and

trigger future research. Based on that, future work aims at triangulating this study through other channels, e.g., in-field observations, and at investigating bad practices in applying and maintaining CI/CD for CPSs. In particular, our goal will be to automatically detect, by analyzing CI/CD pipeline configurations and run-time data, problematic situations (“smells”) that would require an intervention on the DevOps side, and, for what possible, automatically suggest repairs.

## ACKNOWLEDGMENTS

We gratefully acknowledge the Horizon 2020 (EU Commission) support for the project *COSMOS* (DevOps for Complex Cyber-physical Systems), Project No. 957254-COSMOS. Also, we would like to thank the study participants for taking their time to participate in the study and for providing insightful feedback.

## REFERENCES

- [1] [n.d.]. MISRA:2012 Guidelines for the use of the C language in critical systems. [https://www.academia.edu/40301277/MISRA\\_C\\_2\\_012\\_Guidelines\\_for\\_the\\_use\\_of\\_the\\_C\\_language\\_in\\_critical\\_systems.pdf](https://www.academia.edu/40301277/MISRA_C_2_012_Guidelines_for_the_use_of_the_C_language_in_critical_systems.pdf). Accessed: 2021-08-30.
- [2] Sara Abbaspour Asadollah, Rafia Inam, and Hans Hansson. 2015. A Survey on Testing for Cyber Physical System. In *Testing Software and Systems*. Springer Intern. Publishing, Cham, 194–207.
- [3] Rabe Abdalkareem, Suhaib Mujahid, Emad Shihab, and Juergen Rilling. 2021. Which Commits Can Be CI Skipped? *IEEE Trans. Software Eng.* 47, 3 (2021), 448–463.
- [4] Raja Ben Abdesslem, Annibale Panichella, Shiva Nejati, Lionel C Briand, and Thomas Stifter. 2018. Testing autonomous cars for feature interaction failures using many-objective search. In *IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 143–154.
- [5] Zuleyha Akgun, Murat Yilmaz, and Paul M. Clarke. 2020. Assessing Application Lifecycle Management (ALM) Potentials from an Industrial Perspective. In *27th European Conference on Systems, Software and Services Process Improvement, Düsseldorf, Germany*. 326–338. [https://doi.org/10.1007/978-3-030-56441-4\\_24](https://doi.org/10.1007/978-3-030-56441-4_24)
- [6] Aitor Arrieta, Shuai Wang, Ainhoa Arruabarrena, Urtzi Markiegi, Goiuria Sagardui, and Leire Etxeberria. 2018. Multi-objective black-box test case selection for cost-effectively testing simulation models. In *Proceedings of the Genetic and Evolutionary Computation Conference*. 1411–1418.
- [7] Roberto Bagnara, Abramo Bagnara, and Patricia M Hill. 2018. The MISRA C coding standard and its role in the development and analysis of safety-and security-critical embedded software. In *International Static Analysis Symposium*. Springer, 5–23.
- [8] Sebastian Baltes and Paul Ralph. 2022. Sampling in software engineering research: A critical review and guidelines. *Empirical Software Engineering* 27, 4 (2022), 1–31.
- [9] L. Chen. 2015. Continuous Delivery: Huge Benefits, but Challenges Too. *IEEE Software* 32, 2 (2015), 50–54.
- [10] Lianping Chen. 2017. Continuous Delivery: Overcoming adoption challenges. *Journal of Systems and Software* 128 (2017), 72 – 86.
- [11] Yingnong Dang, Qingwei Lin, and Peng Huang. 2019. AIOps: real-world challenges and research innovations. In *Proceedings of the 41st International Conference on Software Engineering: Companion Proceedings, ICSE 2019, Montreal, QC, Canada, May 25-31, 2019*. 4–5. <https://doi.org/10.1109/ICSE-Companion.2019.00023>
- [12] Alexey Dosovitskiy, Germán Ros, Felipe Codevilla, Antonio M. López, and Vladlen Koltun. 2017. CARLA: An Open Urban Driving Simulator. In *1st Annual Conference on Robot Learning, CoRL 2017 (Proceedings of Machine Learning Research, Vol. 78)*. PMLR, 1–16.
- [13] Paul Duvall, Stephen M. Matyas, and Andrew Glover. 2007. *Continuous Integration: Improving Software Quality and Reducing Risk*. Addison-Wesley.
- [14] Paul M. Duvall. 2010. Continuous integration. Patterns and Antipatterns. *DZone refcard #84* (2010). <http://bit.ly/l8rfVS>
- [15] Paul M. Duvall. 2011. Continuous Delivery: Patterns and Antipatterns in the Software Life Cycle. *DZone refcard #145* (2011). <https://dzone.com/refcardz/continuous-delivery-patterns>
- [16] Moritz Eck, Fabio Palomba, Marco Castelluccio, and Alberto Bacchelli. 2019. Understanding flaky tests: the developer’s perspective. In *Proceedings of the ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering, ESEC/SIGSOFT FSE 2019, Tallinn, Estonia, August 26-30, 2019*. 830–840.
- [17] Omar Elazhary, Margaret-Anne D. Storey, Neil A. Ernst, and Elise Paradis. 2021. ADEPT: A Socio-Technical Theory of Continuous Integration. In *43rd IEEE/ACM International Conference on Software Engineering: New Ideas and Emerging Results, ICSE (NIER) 2021, Madrid, Spain, May 25-28, 2021*. 26–30.
- [18] Omar Elazhary, Colin Werner, Ze Shi Li, Derek Lowlind, Neil A Ernst, and Margaret-Anne Storey. 2021. Uncovering the benefits and challenges of continuous integration practices. *IEEE Transactions on Software Engineering* (2021).



- [19] Camilo Escobar-Velásquez, Mario Linares-Vásquez, Gabriele Bavota, Michele Tufano, Kevin Moran, Massimiliano Di Penta, Christopher Vendome, Carlos Bernal-Cárdenas, and Denys Poshyvanyk. 2022. Enabling Mutant Generation for Open- and Closed-Source Android Apps. *IEEE Trans. Software Eng.* 48, 2 (2022), 186–208. <https://doi.org/10.1109/TSE.2020.2982638>
- [20] Erik Flores-García, Goo-Young Kim, Jinho Yang, Magnus Wiktorsson, and Sang Do Noh. 2020. Analyzing the Characteristics of Digital Twin and Discrete Event Simulation in Cyber Physical Systems. In *Advances in Production Management Systems. Towards Smart and Digital Manufacturing (IFIP Advances in Information and Communication Technology, Vol. 592)*. Springer, 238–244.
- [21] Martin Fowler and Matthew Foemmel. [n.d.]. Continuous integration. <https://web.archive.org/web/20200522100521/https://martinfowler.com/articles/originalContinuousIntegration.html>. Accessed: 2021-11-21.
- [22] Heinz Gall. 2008. Functional safety IEC 61508 / IEC 61511 the impact to certification and the user. In *2008 IEEE/ACS International Conference on Computer Systems and Applications*. 1027–1031.
- [23] Keheliya Gallaba and Shane McIntosh. 2018. Use and misuse of continuous integration features: An empirical study of projects that (mis) use Travis CI. *IEEE Trans. Software Eng.* 46, 1 (2018), 33–50.
- [24] Alessio Gambi, Jonathan Bell, and Andreas Zeller. 2018. Practical test dependency detection. In *2018 IEEE 11th International Conference on Software Testing, Verification and Validation (ICST)*. IEEE, 1–11.
- [25] Alessio Gambi, Tri Huynh, and Gordon Fraser. 2019. Generating effective test cases for self-driving cars from police reports. In *Proceedings of the ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering, ESEC/SIGSOFT FSE 2019, Tallinn, Estonia, August 26-30, 2019*. ACM, 257–267.
- [26] Joshua Garcia, Yang Feng, Junjie Shen, Sumaya Almanee, Yuan Xia, and Qi Alfred Chen. 2020. A comprehensive study of autonomous vehicle bugs. In *ICSE '20: 42nd International Conference on Software Engineering, Seoul, South Korea, 27 June - 19 July, 2020*. 385–396.
- [27] Smitha Gautham, Athira Varma Jayakumar, Abhi Rajagopala, and Carl Elks. 2021. Realization of a Model-Based DevOps Process for Industrial Safety Critical Cyber Physical Systems. In *2021 4th IEEE International Conference on Industrial Cyber-Physical Systems (ICPS)*. 597–604.
- [28] Farshad Gholami, Niousha Attar, Hassan Haghghi, Mojtaba Vahidi Asl, Meysam Valueian, and Saina Mohamadyari. 2018. A classifier-based test oracle for embedded software. In *2018 Real-Time and Embedded Systems and Technologies (RTEST)*. 104–111.
- [29] Carlos A. González, Mojtaba Varmazyar, Shiva Nejati, Lionel C. Briand, and Yago Isasi. 2018. Enabling Model Testing of Cyber-Physical Systems. In *Proceedings of the 21th ACM/IEEE International Conference on Model Driven Engineering Languages and Systems (Copenhagen, Denmark) (MODELS '18)*. Association for Computing Machinery, New York, NY, USA, 176–186.
- [30] Leo A Goodman. 1961. Snowball sampling. *The annals of mathematical statistics* (1961), 148–170.
- [31] Robert M. Groves, Floyd J. Jr. Fowler, Mick P. Couypper, James M. Lepkowski, Eleanor Singer, and Roger Tourangeau. 2009. *Survey Methodology, 2nd edition*. Wiley.
- [32] Philipp Helle, Wladimir Schamai, and Carsten Strobel. 2016. Testing of Autonomous Systems - Challenges and Current State-of-the-Art. *INCOSE International Symposium* (2016), 571–584.
- [33] Michael Hilton, Nicholas Nelson, Timothy Tunnell, Darko Marinov, and Danny Dig. 2017. Trade-Offs in Continuous Integration: Assurance, Security, and Flexibility. In *Proceedings of the 25th ACM SIGSOFT International Symposium on Foundations of Software Engineering, FSE 2017*.
- [34] Michael Hilton, Timothy Tunnell, Kai Huang, Darko Marinov, and Danny Dig. 2016. Usage, costs, and benefits of continuous integration in open-source projects. In *Proceedings of the 31st IEEE/ACM International Conference on Automated Software Engineering (ASE)*.
- [35] David L. Hoover. 2016. Textual Variation, Text-Randomization, and Microanalysis. In *DH. Alliance of Digital Humanities Organizations (ADHO)*, 223–225. <http://dblp.uni-trier.de/db/conf/dihu/dh2016.html#Hoover16a>
- [36] Nargiz Humbatova, Gunel Jahangirova, and Paolo Tonella. 2021. DeepCrime: mutation testing of deep learning systems based on real faults. In *ISSTA '21: 30th ACM SIGSOFT International Symposium on Software Testing and Analysis, Virtual Event, Denmark, July 11-17, 2021*. 67–78. <https://doi.org/10.1145/3460319.3464825>
- [37] Gunel Jahangirova, Andrea Stocco, and Paolo Tonella. 2021. Quality metrics and oracles for autonomous vehicles testing. In *2021 14th IEEE Conference on Software Testing, Verification and Validation (ICST)*. IEEE, 194–204.
- [38] Gunel Jahangirova, Andrea Stocco, and Paolo Tonella. 2021. Quality Metrics and Oracles for Autonomous Vehicles Testing. In *14th IEEE Conference on Software Testing, Verification and Validation, ICST 2021, Porto de Galinhas, Brazil, April 12-16, 2021*. IEEE, 194–204. <https://doi.org/10.1109/ICST49551.2021.00030>
- [39] Suzette Johnson, Harry Koehneman, Diane LaFortune, Dean Leffingwell, Stephen Magill, Steve Mayner, Avigail Ofer, Robert Stroud, Anders Wallgren, and Robin Yeman. 2017. *Industrial DevOps - Applying DevOps and Continuous Delivery to Significant Cyber-Physical Systems*. National Academies Press. <https://itrevolution.com/book/industrial-devops/>
- [40] Nidhi Kalra and Susan Paddock. 2016. Driving to safety: How many miles of driving would it take to demonstrate autonomous vehicle reliability? *Transportation Research Part A: Policy and Practice* 94 (12 2016), 182–193.
- [41] Barbara A. Kitchenham and Shari Lawrence Pfleeger. 2002. Principles of survey research part 2: designing a survey. *ACM SIGSOFT Software Engineering Notes* 27, 1 (2002), 18–20.

- [42] Barbara A. Kitchenham and Shari Lawrence Pfleeger. 2002. Principles of survey research: part 3: constructing a survey instrument. *ACM SIGSOFT Software Engineering Notes* 27, 2 (2002), 20–24.
- [43] Barbara A. Kitchenham and Shari Lawrence Pfleeger. 2002. Principles of survey research part 4: questionnaire evaluation. *ACM SIGSOFT Software Engineering Notes* 27, 3 (2002), 20–23.
- [44] Barbara A. Kitchenham and Shari Lawrence Pfleeger. 2002. Principles of survey research: part 5: populations and samples. *ACM SIGSOFT Software Engineering Notes* 27, 5 (2002), 17–20.
- [45] Klaus Krippendorff. 2004. Reliability in Content Analysis: Some common Misconceptions and Recommendations. *Journal of the Royal Statistical Society. Series B (Methodological)* 30, 3 (2004), 411–433.
- [46] Wing Lam, Patrice Godefroid, Suman Nath, Anirudh Santhiar, and Suresh Thummalapenta. 2019. Root Causing Flaky Tests in a Large-Scale Industrial Setting. In *Proceedings of the 28th ACM SIGSOFT International Symposium on Software Testing and Analysis (Beijing, China) (ISSTA 2019)*. ACM, 101–111.
- [47] Wing Lam, Stefan Winter, Angello Astorga, Victoria Stodden, and Darko Marinov. 2020. Understanding Reproducibility and Characteristics of Flaky Tests Through Test Reruns in Java Projects. In *31st IEEE International Symposium on Software Reliability Engineering, ISSRE, Coimbra, Portugal, October 12-15, 2020*. 403–413.
- [48] Wing Lam, Stefan Winter, Anjiang Wei, Tao Xie, Darko Marinov, and Jonathan Bell. 2020. A large-scale longitudinal study of flaky tests. *Proc. ACM Program. Lang.* 4, OOPSLA (2020), 202:1–202:29. <https://doi.org/10.1145/3428270>
- [49] Johannes Lampel, Sascha Just, Sven Apel, and Andreas Zeller. 2021. When life gives you oranges: detecting and diagnosing intermittent job failures at Mozilla. In *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 1381–1392.
- [50] Antonio Loquercio, Elia Kaufmann, René Ranftl, Alexey Dosovitskiy, Vladlen Koltun, and Davide Scaramuzza. 2019. Deep drone racing: From simulation to reality with domain randomization. *IEEE Transactions on Robotics* 36, 1 (2019), 1–14.
- [51] Ivano Malavolta, Grace Lewis, Bradley Schmerl, Patricia Lago, and David Garlan. 2020. How Do You Architect Your Robots? State of the Practice and Guidelines for ROS-Based Systems. In *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering: Software Engineering in Practice (Seoul, South Korea) (ICSE-SEIP '20)*. ACM, New York, NY, USA, 31–40.
- [52] Torvald Mårtensson, Daniel Ståhl, and Jan Bosch. 2016. Continuous integration applied to software-intensive embedded systems—problems and experiences. In *International Conference on Product-Focused Software Process Improvement*. Springer, 448–457.
- [53] Claudio Menghi, Shiva Nejati, Lionel Briand, and Yago Isasi Parache. 2020. Approximation-refinement testing of compute-intensive cyber-physical models: An approach based on system identification. In *2020 IEEE/ACM 42nd International Conference on Software Engineering (ICSE)*. IEEE, 372–384.
- [54] Vuong Nguyen, Stefan Huber, and Alessio Gambi. 2021. SALVO: Automated Generation of Diversified Tests for Self-driving Cars from Existing Maps. In *2021 IEEE International Conference on Artificial Intelligence Testing (AITest)*. IEEE, 128–135.
- [55] Helena Holmstrom Olsson, Hiva Alahyari, and Jan Bosch. 2012. Climbing the “Stairway to Heaven” – A Multiple-Case Study Exploring Barriers in the Transition from Agile Development Towards Continuous Deployment of Software. In *Proceedings of the 2012 38th Euromicro Conference on Software Engineering and Advanced Applications (SEAA '12)*. IEEE Computer Society, Washington, DC, USA, 392–399.
- [56] Heejong Park, Arvind Easwaran, and Sidharta Andalām. 2019. Challenges in Digital Twin Development for Cyber-Physical Production Systems. In *Cyber Physical Systems. Model-Based Design*. Springer International Publishing, Cham, 28–48.
- [57] Shari Lawrence Pfleeger and Barbara A. Kitchenham. 2001. Principles of survey research: part 1: turning lemons into lemonade. *ACM SIGSOFT Software Engineering Notes* 26, 6 (2001), 16–18.
- [58] Gustavo Pinto, Breno Miranda, Supun Dissanayake, Marcelo d’Amorim, Christoph Treude, and Antonia Bertolino. 2020. What is the Vocabulary of Flaky Tests?. In *MSR '20: 17th International Conference on Mining Software Repositories, Seoul, Republic of Korea, 29-30 June, 2020*. 492–502.
- [59] Gustavo Pinto, Breno Miranda, Supun Dissanayake, Marcelo d’Amorim, Christoph Treude, and Antonia Bertolino. 2020. What is the vocabulary of flaky tests?. In *Proceedings of the 17th International Conference on Mining Software Repositories*. 492–502.
- [60] Akond Rahman, Chris Parmin, and Laurie Williams. 2019. The Seven Sins: Security Smells in Infrastructure As Code Scripts. In *Proceedings of the 41st International Conference on Software Engineering (Montreal, Quebec, Canada) (ICSE '19)*. IEEE Press, 164–175.
- [61] Steven P. Reiss. 2007. Automatic code stylizing. In *22nd IEEE/ACM International Conference on Automated Software Engineering (ASE 2007), November 5-9, 2007, Atlanta, Georgia, USA*. 74–83.
- [62] Vincenzo Riccio and Paolo Tonella. 2020. Model-based Exploration of the Frontier of Behaviours for Deep Learning System Testing. In *Proceedings of the ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE '20)*. Association for Computing Machinery.
- [63] Terry Rowlands, Neal Waddell, and Bernard McKenna. 2016. Are We There Yet? A Technique to Determine Theoretical Saturation. *J. Comput. Inf. Syst.* 56, 1 (2016), 40–47. <http://dblp.uni-trier.de/db/journals/jcis/jcis56.html#RowlandsWM16>
- [64] Tony Savor, Mitchell Douglas, Michael Gentili, Laurie Williams, Kent Beck, and Michael Stumm. 2016. Continuous deployment at Facebook and OANDA. In *Companion proceedings of the 38th International Conference on Software Engineering (ICSE Companion)*. 21–30.

- [65] August Shi, Jonathan Bell, and Darko Marinov. 2019. Mitigating the effects of flaky tests on mutation testing. In *Proceedings of the 28th ACM SIGSOFT International Symposium on Software Testing and Analysis, ISSTA*. 112–122.
- [66] Hesham Shokry and Mike Hinchey. 2009. Model-based verification of embedded software. (2009).
- [67] Sebastian Sontges and Matthias Althoff. 2018. Computing the Drivable Area of Autonomous Road Vehicles in Dynamic Road Scenes. *IEEE Trans. Intell. Transp. Syst.* 19, 6 (2018), 1855–1866.
- [68] Donna Spencer. 2009. *Card sorting: Designing usable categories*. Rosenfeld Media.
- [69] Daniel Ståhl and Jan Bosch. 2014. Automated software integration flows in industry: a multiple-case study. In *Companion Proceedings of the 36th International Conference on Software Engineering*. ACM, 54–63.
- [70] Daniel Ståhl and Jan Bosch. 2014. Modeling continuous integration practice differences in industry software development. *J. Syst. Softw.* 87 (2014), 48–59.
- [71] Daniel Ståhl, Kristofer Hallén, and Jan Bosch. 2016. Continuous Integration and Delivery Traceability in Industry: Needs and Practices. In *42th Euromicro Conference on Software Engineering and Advanced Applications, SEAA 2016, Limassol, Cyprus, August 31 - Sept. 2, 2016*. IEEE Computer Society, 68–72.
- [72] Andrea Stocco, Michael Weiss, Marco Calzana, and Paolo Tonella. 2020. Misbehaviour prediction for autonomous driving systems. In *ICSE '20: 42nd International Conference on Software Engineering, Seoul, South Korea*. 359–371.
- [73] Sirasak Tepjit, Imre Horváth, and Zoltán Rusák. 2019. The state of framework development for implementing reasoning mechanisms in smart cyber-physical systems: A literature review. *Journal of Computational Design and Engineering* 6, 4 (04 2019), 527–541. arXiv:<https://academic.oup.com/jcde/article-pdf/6/4/527/33135124/jcde.2019.04.002.pdf>
- [74] Martin Törngren and Ulf Sellgren. 2018. *Complexity Challenges in Development of Cyber-Physical Systems*. Springer International Publishing, Cham, 478–503.
- [75] Bogdan Vasilescu, Yue Yu, Huaimin Wang, Premkumar T. Devanbu, and Vladimir Filkov. 2015. Quality and productivity outcomes relating to continuous integration in GitHub. In *ESEC/SIGSOFT FSE*. ACM, 805–816.
- [76] Mario Linares Vásquez, Gabriele Bavota, Michele Tufano, Kevin Moran, Massimiliano Di Penta, Christopher Vendome, Carlos Bernal-Cárdenas, and Denys Poshyvanyk. 2017. Enabling mutation testing for Android apps. In *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering, ESEC/FSE 2017, Paderborn, Germany, September 4-8, 2017*. 233–244. <https://doi.org/10.1145/3106237.3106275>
- [77] Carmine Vassallo, Sebastian Proksch, Harald Gall, and Massimiliano Di Penta. 2019. Automated Reporting of Anti-Patterns and Decay in Continuous Integration. In *Proceedings of the 41st International Conference on Software Engineering, ICSE 2019, Montreal, Canada, May 25 - 31, 2019*. IEEE.
- [78] Carmine Vassallo, Sebastian Proksch, Anna Jancso, Harald C Gall, and Massimiliano Di Penta. 2020. Configuration smells in continuous delivery pipelines: a linter and a six-month study on GitLab. In *Proceedings of the 28th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 327–337.
- [79] Carmine Vassallo, Fiorella Zampetti, Daniele Romano, Moritz Beller, Annibale Panichella, Massimiliano Di Penta, and Andy Zaidman. 2016. Continuous Delivery Practices in a Large Financial Organization. In *32nd IEEE International Conference on Software Maintenance and Evolution (ICSME)*. 41–50.
- [80] Chunhui Wang, Fabrizio Pastore, and Lionel C. Briand. 2019. Oracles for Testing Software Timeliness with Uncertainty. *ACM Trans. Softw. Eng. Methodol.* 28, 1 (2019), 1:1–1:30. <https://doi.org/10.1145/3280987>
- [81] Dinghua Wang, Shuqing Li, Guanping Xiao, Yepang Liu, and Yulei Sui. 2021. An exploratory study of autopilot software bugs in unmanned aerial vehicles. In *ESEC/FSE '21: 29th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 20–31.
- [82] Rajaa Vikhram Yohanandhan, Rajvikram Madurai Elavarasan, Premkumar Manoharan, and Lucian Mihet-Popa. 2020. Cyber-Physical Power System (CPPS): A Review on Modeling, Simulation, and Analysis With Cyber Security Applications. *IEEE Access* 8 (2020), 151019–151064.
- [83] Fiorella Zampetti, Saghan Mudbhari, Venera Arnaudova, Massimiliano Di Penta, Sebastiano Panichella, and Giuliano Antoniol. 2022. Using code reviews to automatically configure static analysis tools. *Empir. Softw. Eng.* 27, 1 (2022), 28. <https://doi.org/10.1007/s10664-021-10076-4>
- [84] Fiorella Zampetti, Damian A. Tamburri, Sebastiano Panichella, Annibale Panichella, Gerardo Canfora, and Massimiliano Di Penta. 2022. *Continuous Integration and Delivery Practices for Cyber-Physical Systems: An Interview-Based Study - Online Dataset*. <https://doi.org/10.5281/zenodo.6337040>
- [85] Fiorella Zampetti, Carmine Vassallo, Sebastiano Panichella, Gerardo Canfora, Harald C. Gall, and Massimiliano Di Penta. 2020. An empirical characterization of bad practices in continuous integration. *Empirical Software Engineering* 25, 2 (2020), 1095–1135.
- [86] Eleni Zapridou, Ezio Bartocci, and Panagiotis Katsaros. 2020. Runtime Verification of Autonomous Driving Systems in CARLA. In *Runtime Verification*. Springer International Publishing, Cham, 172–183.
- [87] Peilun Zhang, Yanjie Jiang, Anjiang Wei, Victoria Stodden, Darko Marinov, and August Shi. 2021. Domain-Specific Fixes for Flaky Tests with Wrong Assumptions on Underdetermined Specifications. In *43rd IEEE/ACM International Conference on Software Engineering, ICSE*

2021, Madrid, Spain, 22-30 May 2021. 50–61.

- [88] Sai Zhang, Darioush Jalali, Jochen Wuttke, Kivanç Muşlu, Wing Lam, Michael D Ernst, and David Notkin. 2014. Empirically revisiting the test independence assumption. In *Proceedings of the 2014 International Symposium on Software Testing and Analysis*. 385–396.
- [89] Celal Ziftci and Diego Cavalcanti. 2020. De-Flake Your Tests : Automatically Locating Root Causes of Flaky Tests in Code At Google. In *2020 IEEE International Conference on Software Maintenance and Evolution (ICSME)*. 736–745.
- [90] Behrouz Zolfaghari, Reza M. Parizi, Gautam Srivastava, and Yoseph Hailemariam. 2021. Root causing, detecting, and fixing flaky tests: State of the art and future roadmap. *Softw. Pract. Exp.* 51, 5 (2021), 851–867. <https://doi.org/10.1002/spe.2929>

Just Accepted

## APPENDIX - DETAILS ABOUT THE INTERVIEW PARTICIPANTS

The interviews' transcripts together with the labeling procedure helped us in forming organization profiles, focusing more on how the interviewed companies set and maintain a pipeline for CPS development.

In the following, we detail the development process of the interviewed organizations, focusing more on the status of their CI/CD pipeline in terms of (i) build triggering strategies (e.g., continuous or periodic), (ii) co-existence of multiple pipeline configurations (e.g., for different devices) and the frequency of changes occurring to them, (iii) the phases being automated/executed within the pipeline, and (iv) the usage and setting of HiL and simulators within the pipeline (e.g., a pipeline can be a mix of different environments used in different circumstances). In the following, for each organization participating in the interviews, we describe—by leveraging the codes elicited during the interviews' transcripts analysis phase—the CPS development process and especially the adoption of CI/CD pipelines and, in general, of build automation<sup>3</sup>.

### O<sub>1</sub> (Aerospace)

CONTEXT: O<sub>1</sub> is involved in verification and validation (V&V) tasks for aerospace software (i.e., on-board software for satellites), hence their CI/CD pipeline is only for V&V and not for development. This is because, due to the safety integrity level (SIL) of the CPS, the development and the V&V teams and pipelines must be kept distinct [22]. O<sub>1</sub> relies on conventional programming languages dictated by standards in the aerospace domain (“*We mainly use ANSI C-99 following the MISRA rules*”). This implies the need for certifying software (i.e., following the Motor Industry Software Reliability Association (MISRA) standards [1, 7]).

PIPELINE STATUS: O<sub>1</sub> has started adopting CI/CD practices less than one year ago, mainly due to a limited culture within the team about CI/CD principles. Moreover, it does not have a strict separation of roles for what concerns the type of interaction with the pipeline (“*a developer who needs to customize a CI/CD pipeline by simply using yaml files can customize it directly*”). O<sub>1</sub> does not rely on build matrices with jobs related to different environment variants since “*the pipeline does not have to change/evolve based on the changes in the technologies being used (version for compilers and or programming languages), the aerospace domain follows the waterfall process. So everything is [frozen]: no changes may occur later on in the process.*”

AUTOMATED TASKS: Due to the application domain and the related standards and certification constraints, the pipeline compiles the software provided and developed by the customer, relies on SonarQube for (i) checking the fulfillment of the MISRA rules for certification, (ii) identifying maintainability problems (i.e., “*we also have non-functional requirements expressed in terms of rules available in SonarQube*”) mainly related to the presence of duplicated code, and (iii) identifying bugs as soon as they are introduced, and executing unit and robustness tests to “*check how the system behaves/reacts in the presence of unexpected inputs ([e.g., ] inputs having values out of the admissible range)*”. Furthermore, considering the overall scope of the pipeline, its triggering is manual, even if there are also nightly builds used for running test suites requiring a long time to complete. It is important to note that the testing criteria to derive the test cases to include within the pipeline are expressed from the customer as non-functional requirements (“*i.e., use MC/DC for deriving the test suite*”). Finally, O<sub>1</sub> has to consider time constraints for the pipeline setting to deal with possible issues that may arise when launching the simulator (e.g., memory leaks or impossibility to access the simulator).

HiL AND SIMULATORS: O<sub>1</sub> cannot involve HiL in the pipeline, as it would require a clean room not accessible from the outside. Instead, it relies on third-party simulators dictated by the customer. This is because the customer follows “*a framework for simulation aimed at hosting different simulators for different satellite models for the digital twin of the satellite*”. Relying on third-party simulators helps in reducing the costs/efforts needed to develop the simulators from scratch, as well as, helps in guaranteeing the trustworthiness of the outcome being produced and

<sup>3</sup>The interested reader could find the code mind maps in the replication package [84].

provided to the customer. Of course, the level of trustworthiness increases for those cases where the simulator is provided by the same vendor of the hardware device that must be simulated.

### O<sub>2</sub> (Healthcare)

CONTEXT: O<sub>2</sub> is a large organization involved in the healthcare domain: it provides Computed Tomography (CT) scanners for clinical use. As regards the development process, O<sub>2</sub> has a team for each component being developed—around 17 different teams working on 70 branches— together with an integration branch where all the other branches are integrated into a “*single joined point*”. Furthermore, each team adopts conventional programming languages, *i.e.*, mainly C# and C++.

PIPELINE STATUS: O<sub>2</sub> already has a CI/CD pipeline in place for CPS development that has been introduced 4 years ago, and they are still improving it. Furthermore, based on its application domain, O<sub>2</sub> is constrained to “*follow medical application frameworks providing a base set of rules in terms of how to build applications and how to integrate them*”. The latter requires the adoption of processes aimed at verifying whether or not the overall development process adheres to the regulatory standards for developing medical applications.

AUTOMATED TASKS: O<sub>2</sub> adopts both incremental and nightly builds. Of course, the tasks involved in the different types of builds, as well as, the execution environment involved in them vary. Specifically, nightly builds leverage HiL, and run three different types of testing, namely unit/component, sub-system, and system testing. To provide developers fast feedback about the impact of their changes, O<sub>2</sub> relies on incremental builds executing only a subset of the whole set of functional tests— by doing “*impact based testing to figure out the impact of the changes and select the tests to be executed based on the impact*.” To control the overall build execution time, O<sub>2</sub> encourages developers to push small changes leading to “*small sets of tests to be executed*.” Finally, both incremental and nightly builds run static code analysis tools mainly aimed at identifying maintainability and security flows in the code.

There is a specific type of build aimed at checking performance requirements like “*test whether each component (some components) stays within the resource limits they are assigned to*”. The outcome of such a build is compared over time to identify and monitor possible performance degradation within the whole system. Moreover, O<sub>2</sub> has a specific DevOps team for checking the fulfillment of security requirements, even if this is not done continuously while only “*near the finalization of the product*”, and it is not automated.

HiL AND SIMULATORS: As explained above, both the triggering strategies adopted by O<sub>2</sub> and the tasks being automatized within each type of build influence the choice between using simulators and/or HiL. Nightly builds have an automated deployment on a “*real*” CT Scanner “*without reusing existing artifacts while building all of them from scratch in a clean environment*”, for executing the whole test suite in a real production environment. Note that, when talking about “*real*” CT Scanner, O<sub>2</sub> refers to “*physical systems that are equivalent to the real hardware in the CT Scanner but not connected to anything around it which has a simulator running on it*”.

For what concerns simulators, O<sub>2</sub> relies on self-developed simulators— there are suitable knowledge and skills to properly develop simulators, *i.e.*, O<sub>2</sub> develops both the software and the hardware. However, at the moment, O<sub>2</sub> does not use simulators (“*mainly used for functional testing only*”) for checking non-functional (*i.e.*, performance) requirements.

### O<sub>3</sub> (Acoustic Sensors)

CONTEXT: O<sub>3</sub> is involved in CPS innovation for the industry, among others the development of the SPL Noise Meter Board, *i.e.*, a low-cost, high quality, electronic sensing board capable of measuring noise of the environment. It does not have any separation of roles between the members of the team (“*The team is the company*.”), however, the team is composed of both software and hardware experts who work together, simplifying the overall development process, in particular for those activities requiring the integration and communication between

software and hardware components (*“Useful for sensors’ integration... [it] help[s] having knowledge about the hardware components, how they work and how it is possible to communicate with them.”*)

O<sub>3</sub> adopts a pull-request (PR) development process with one branch per feature (*i.e.*, *“Several branches for maintaining and developing different features.”*) Furthermore, even if it does not have strict guidelines in terms of coding standards, O<sub>3</sub> attempts to adopt similar coding styles within each branch. Finally, it relies on conventional programming languages, *i.e.*, Python for testing and C, C++ for micro-controllers development.

PIPELINE STATUS: At the moment, O<sub>3</sub> does not have a CI/CD pipeline for CPS development.

AUTOMATED TASKS: Even if O<sub>3</sub> does not have a CI/CD pipeline in place, the deployment is fully-automated, while the testing is still manual mainly due to the impossibility of automating the oracle specification, in particular for testing acoustic signals. Furthermore, even if O<sub>3</sub> does not have certification constraints for the developed code, they need to cope with certification constraints *“for the acoustic signals.”*

HiL AND SIMULATORS: O<sub>3</sub> only uses real hardware devices, even if within the organization there is the wish of including simulators in the process to test the acoustic signal (*i.e.*, the main outcome of their product) in a controlled environment, *i.e.*, *“removing noise from the surrounded environment.”*

#### O<sub>4</sub> (Robotics)

CONTEXT: O<sub>4</sub> is involved in the development of autonomous robots, and is made up of several development teams where each team accounts for both hardware and software developers. Furthermore, it adopts a PR development process with one branch per feature (*i.e.*, *“We have a branch for each feature that needs to be implemented and/or improved and we use PRs to merge the work in the stable release branch.”*) Based on the application domain, it mainly adopts C++, together with Python for users’ interfaces and for interacting with the hardware devices.

PIPELINE STATUS: O<sub>4</sub> has a fully containerized (using Docker) pipeline for CPS development. It relies on continuous and nightly builds, even if they are not used for running time-intensive tasks, *i.e.*, *“(that is not so expensive in terms of execution time)”*, while for running regression testing activities on already packaged components and for deployment to the customers. Furthermore, the CI/CD configuration is pretty stable meaning that, even if each branch may rely on a customized CI/CD process, the configuration does not have to change over time.

AUTOMATED TASKS: Our interviewee mentions the execution of static code analysis tools to inform developers about code quality degradation, and unit tests relying on simulators. Only when a PR is peer-reviewed and there are no failures in the entailed CI/CD process, it is possible to merge the change on the stable repository and enact a release process for shipping the product to the customers. O<sub>4</sub> monitors the overall quality of the development process in terms of static analysis metrics and code coverage from the unit test execution. The application domain does not introduce certification constraints, *i.e.*, *“If you want to sell a robot you do not need to have a certified robot”*, while it hinders the automation of non-functional testing within the pipeline. Specifically, our interviewee mentions the manual execution of reliability and safety tests *“running the robot a long time with a guy supervising the test execution to understand when and why the robot starts to not work anymore”*, or *“guaranteeing that once pressing the stop button the robot actually shuts down”*.

HiL AND SIMULATORS: O<sub>4</sub> relies on third-party simulators and HiL. One point raised by our interviewee is related to the partial usage of Docker on the hardware so that it is possible to run the robot in a privileged mode and switch between software versions quite easily: *“each one may choose the version of the software that has to be run over the robot.”*

#### O<sub>5</sub> (Automotive)

CONTEXT: O<sub>5</sub> is a large company operating in the automotive domain working on the software-focused driving platform. The development process is organized into three different teams each one with a specific goal: *“one*

*working on virtual machines, one working on web services because we provide DevOps solutions for embedded systems, and finally, we have a small team working on customer delivery with the goal of adapting our tools to the customers' needs.*" This is the only organization in our study relying on real-time languages, *i.e.*, real-time Java to cope with scheduling requirements of embedded systems.

**PIPELINE STATUS:** O<sub>5</sub> already has a CI/CD pipeline in place mainly for deployment purposes (*"we are able to support updating software on devices on the fly"*), even if it is working on improving it — *"it is still kind of an infancy we are still working on improving"*.

**AUTOMATED TASKS:** O<sub>5</sub> mainly uses the pipeline for deployment. However, differently from other organizations, O<sub>5</sub> relies on virtual machines instead of using containers for several reasons: (i) better control over the resources (*i.e.*, *"the ability to enforce our resource usage inside the virtual machine while you do not have quite the same extent with a container"*), (ii) versioning capability, *i.e.*, *"when a new service comes in, you register it so it is easy to start a new version of this service, run down the old one and switch over the new one during run-time"*, and (iii) memory safety guarantee, *i.e.*, *"by looking at a recent post by both Google and Microsoft we found that around 70% of the security violations are due to failures of memory safety. So by using a garbage-collected environment, we can prevent those issues from occurring."* Going deeper into how the deployment process works, O<sub>5</sub> first creates the virtual machine, (*i.e.*, emulating the virtual environment), then the OSGi infrastructure, and finally it tests individual modules. The latter means that O<sub>5</sub> does not test all the developed modules together since it *"deploy[s] individual bundles to a platform."*

For what concerns the verification of non-functional requirements, O<sub>5</sub> performs security and performance testing, even if they are not included in the pipeline. Specifically, for real-time systems it is important to monitor the impact of each change on performance properties to be able to identify, as soon as possible, the change introducing performance degradation, *i.e.*, *"we have various performance tests that we run regularly to track our performance as the system evolves."*

**HiL AND SIMULATORS:** Since O<sub>5</sub> develops software for embedded entertainment in the automotive domain, the HiL is only available for a final validation on the customer's side: *"then employ our customer for the last mile"*, so most of the work is done relying on virtual environments.

## O<sub>6</sub> (Aerospace)

**CONTEXT:** Similar to O<sub>1</sub>, O<sub>6</sub> operates in the aerospace domain, and is mainly involved in the development and refining of the routing algorithm for the Free Route Airspace (FRA). For what concerns the programming language being adopted, O<sub>6</sub> relies on conventional languages: *"C and C++ [are] used for the back-end."*

**PIPELINE STATUS:** O<sub>6</sub> already has a CI/CD pipeline mainly for deployment and testing purposes, that is under continuous improvements. Moreover, our interviewee mentions that the pipeline is more an MLOps than a simple DevOps pipeline.

**AUTOMATED TASKS:** Among the phases being automated there are: (i) static code analysis for identifying maintainability flows and spotting bugs as soon as they are introduced, (ii) unit testing, (iii) integration testing, and (iv) deployment. Furthermore, the execution of non-functional testing activities is mainly carried out manually and outside the pipeline, due to the high complexity of the real-time operating system under development. Similarly to O<sub>1</sub>, it is required that the developed code satisfies strict certification requirements that are mainly checked by relying on code coverage tools. Differently from other organizations, O<sub>6</sub> does not rely on nightly builds, meaning that also time-intensive tasks are executed at each change, *i.e.*, *"even the slow builds are continuously built."* O<sub>6</sub> recommends developers to use private builds before pushing their changes on the stable release branch, at least for what concerns the execution of unit testing. Finally, the pipeline provides a monitoring mechanism for what concerns aspects of the real-time operating system such as scheduling and memory that *"gives us the possibility to collect feedback/evidence that may help us in obtaining the certifications."*



**HiL AND SIMULATORS:** O<sub>6</sub> relies on both simulators and HiL, however it does *“not have simulators and HiL in the same pipeline mostly for certification issues.”* Specifically, it is possible to rely on real devices only when there is enough trustworthiness about the software in terms of correct behavior, as well as the absence of crashes gained by relying on self-developed simulators.

## O<sub>7</sub> (Railways)

**CONTEXT:** O<sub>7</sub> is involved in delivering software for railways, *i.e.*, Train Control Management System (TCMS), and similarly to what is reported for the aerospace domain, due to the safety integrity level of the software under development, developers and testers must be different (*i.e.*, *“Testers and Developers are in separate teams in presence of new functionality to be implemented both start together to implement and write test cases.”*)

**PIPELINE STATUS:** O<sub>7</sub> already has a CI/CD pipeline in place for CPS development – *“introduced two years ago”* – that, at the moment, is in a continuous improvement state since it does not automatize the whole development process (*i.e.*, *“The deployment on the real train or on the hardware test track is not automated at the moment even if we are working on making it automatic.”*) In terms of programming language, the interviewee mentions the need of adapting the programming language to the device on which the software has to be executed, however, they mainly rely on conventional languages.

Based on the application domain, O<sub>7</sub> adopts staged builds following the “green-build rule”. In the first stage, the build process is executed on a virtual machine, *i.e.*, *“[a] virtual train, software running on a PC that should behave like it does on a real train”*. Once a change occurs on a specific component the related build process is enacted and, in presence of a green status, all the components are deployed together so that it is possible to enable the execution using the virtual train (*“the devices are run in some kind of containers and we have frameworks building and connecting the whole set of devices and components.”*). If the build process ends with a successful state, it is possible to move to the next stage that relies on the hardware test track – *i.e.*, *“where we have the whole set of devices and even some more that we do not have in the virtual train.”* Finally, if the build process for the second stage ends with a green status, it is possible to run the last stage relying on a real train. Note that, each device/component has a proper CI/CD configuration.

**AUTOMATED TASKS:** O<sub>7</sub> uses the pipeline to automatically test basic functionality (*i.e.*, *“We test specific train functionality such as whether we should activate the train in [a specific] mode”*), as well as, the interaction between different components/devices (*i.e.*, *“we have a long sequence of events for each test that involves different devices and components so we are mainly doing integration testing.”*) The test suites used in different stages of the build process may be different since *“for some test cases, we are not allowed to rely on the virtual environment while we must consider the hardware track or a real train.”* At the moment, O<sub>7</sub> has automated deployment within the pipeline only for the first stage, *i.e.*, relying on the virtual train (for which the overall build execution is *“around one hour and [a] half”*), while it is done manually for what concerns the other two stages: hardware test track and a real train. Testing against non-functional requirements is also done manually, because of the high variability and complexity of the environment.

Going deeper into how developers interact with the CI/CD pipeline, O<sub>7</sub> enforces developers to run private builds before pushing their changes on the main stable repository. The private builds are aimed at executing the same test suite later executed on the CI/CD servers – *“for the moment we cannot configure the number and type of tests to be executed locally”*. Furthermore, the “green-build rule” is used for determining the development tasks: *“in presence of a failure all developers are stopped until the build becomes green again.”*

**HiL AND SIMULATORS:** O<sub>7</sub> adopts both simulators and HiL in different stages of the build process, with the use of HiL occurring only in the last stage of the pipeline.

### O<sub>8</sub> (Railways)

CONTEXT: The application domain of O<sub>8</sub> is railways, and in particular the development of a specific component used for transmitting data between on-board and ground applications. O<sub>8</sub> uses C and C++ (*i.e.*, conventional programming languages), and it has strict constraints for what concerns the production code that has to satisfy strict certification requirements, as well as the compliance with the railway standards and specifications. The latter is mainly checked by relying on *“a specific complex tool that can be configured based on specifications and standards.”*

PIPELINE STATUS: Due to the limited availability of human resources together with the complexity and the safety integrity level of the application domain, O<sub>8</sub> does not have a CI/CD pipeline in place for CPS development and it has, in general, little automation in the development process.

AUTOMATED TASKS: Only the adherence to standards and specifications is automated, while functional *“tests are written manually starting from requirements and system specification[s] but also their execution requires a manual effort”*. This is also the case for non-functional and integration testing (*i.e.*, *“we have a set of testers in front of a screen who monitor and check for the presence of any discrepancies about what is expected and what is instead observed while running the system.”*) Due to the effort and time needed to manually verify the reliability of the software under test, functional tests are executed at every change, while integration tests are only executed when the change impacts the *“interfaces with other modules/components.”*

HIL AND SIMULATORS: Due to the high cost of the hardware devices involved in this particular application domain, O<sub>8</sub> mainly relies on simulators that are self-developed (*“we do not rely on third party very expensive simulators”*). However, once per week, O<sub>8</sub> performs a testing session with a real *“train running in a real environment with real traffic [and] possibly without people.”*

### O<sub>9</sub> (Identification Technology)

CONTEXT: O<sub>9</sub> is involved in *“develop[ing] software relying on identification technologies such as RFID [(Radio Frequency IDentification),] Bluetooth low energy or bar codes*, other than mobile app development for which there is a CI/CD pipeline used for testing and automated deployment on the play stores. For what concerns the CPS development, O<sub>9</sub> relies on conventional programming languages such as C# and Java.

PIPELINE STATUS: The limited availability of human resources, together with a lack of culture for setting a pipeline dealing with sensors and actuators, results in not having a CI/CD pipeline in place for CPS development.

AUTOMATED TASKS: The testing phases are almost fully automated. Specifically, there are *“RFID-readers connected to a network”* on which it is possible to execute unit and integration testing activities automatically. For what concerns integration testing, it is important to remark that there are cases requiring the manual intervention of the tester (*i.e.*, *“For instance, when we need to test a transfer of tags between different antennas we cannot use automation”*), as well as cases where it is required to interact with the hardware devices that cannot be simulated. Of course, in this specific setting, it is not possible to guarantee the overall reproducibility of the results of the test, however, *“the reproducibility of the test in this context is not required.”* Furthermore, O<sub>9</sub> does not run the whole test suite at each change. Instead, they manually select some test cases based on impact analysis: *“select what are the test cases that are impacted by the change that, consequently, need to be executed”*. Other than having unit and integration testing activities, O<sub>9</sub> also executes, from time to time, performance testing.

For what concerns the deployment of CPS-related software, O<sub>9</sub> relies on Docker for creating images that are manually deployed onto the servers.

Finally, the development process also features a monitoring component for the internal development platform and customers' devices, to notify about anomalies and errors, as soon as they occur.

**HiL AND SIMULATORS:** The development process adopted by  $O_9$  relies on both (self-developed) simulators and HiL. Simulators are developed based on specific organization needs and use case scenarios, implying that they are limited in their functionality.

### $O_{10}$ (Energy)

**CONTEXT:**  $O_{10}$  is involved in the development of prototypes and proof of concepts for the energy domain. The development of prototypes rather than real products represents concrete facilitation, since there may be less stringent constraints in terms of pipeline setting and evolution.

**PIPELINE STATUS:** What is mentioned above justifies the presence of a mature (*i.e.*, introduced in 2016) pipeline adopted within the organization for CPS development that uses conventional programming languages, mostly Java and Python. The pipeline configuration is pretty stable probably due to the development of prototyping solutions that do not need to be shipped to real environments.

**AUTOMATED TASKS:** Other than having a compilation phase, the CI/CD pipeline is aimed at executing unit and integration tests (*"Our pipeline is mostly for unit testing (80%) but there is also some integration testing."*), followed by a deployment phase where the packaged version of the software is usually stored into an artifact repository as a docker image. Furthermore, safety requirements, such as checking that a battery is not charged more than a certain rate, are specified and checked through unit test cases that do not involve the real devices. Thanks to the need for developing prototyping solutions, the pipeline accounts for static code analysis tools and linters that are mainly used for checking maintainability issues only (*i.e.*, *"They are not used for checking out bugs, but mostly for making sure that the code is easy to read for other colleagues and for maintainability purposes."*) Moving the attention on the triggering strategies,  $O_{10}$  does not rely on nightly builds. It only uses incremental builds so that each build execution time does not overcome the 10 minutes rule.

**HiL AND SIMULATORS:** Looking at the execution environment,  $O_{10}$  does not need to run the software on embedded devices meaning that it *"tr[ies] to find devices having interfaces to communicate with. So basically we run our software on a traditional machine and it just communicates with the hardware.* So, differently from other organizations,  $O_{10}$ , other than simulating the hardware when needed (*i.e.*, *"we simulate the battery for testing the charging protocol"*), mainly replaces it with mock-ups (*i.e.*, *"it is very easy to mock a client just to see if our software sends the right commands or does not use any register twice"*). Only when the real devices are available and it is safe to use them for testing,  $O_{10}$  uses Docker images for checking the correct behavior over the real devices, as well.