

Understanding Misconfigurations in ROS: An Empirical Study and Current Approaches

Paulo Canelas
pacsantos@cmu.edu
School of Computer Science,
Carnegie Mellon University
Pittsburgh, PA, USA

Alcides Fonseca
amfonseca@fc.ul.pt
Faculdade de Ciências,
Universidade de Lisboa
Lisboa, Portugal

Bradley Schmerl
schmerl@cmu.edu
School of Computer Science,
Carnegie Mellon University
Pittsburgh, PA, USA

Christopher S. Timperley
ctimperley@cmu.edu
School of Computer Science,
Carnegie Mellon University
Pittsburgh, PA, USA

Abstract

The Robot Operating System (ROS) is a highly popular framework and ecosystem that allows developers to build robot software systems from reusable, off-the-shelf components. Systems are often built by customizing and connecting components into a working ensemble via configuration files. While reusable components allow developers to quickly prototype working robots in theory, in practice, ensuring that those components are configured and connected appropriately is fundamentally challenging, as evidenced by the large number of questions asked on developer forums. Developers must be careful to abide by the assumptions of individual components, which are often unchecked and unstated when building their systems. Failure to do so can result in misconfigurations that are only discovered once the robot is deployed in the field, at which point errors may lead to unpredictable and dangerous behavior. Despite misconfigurations having been studied in the broader context of software engineering, robotics software (and ROS in particular) poses domain-specific challenges with potentially disastrous consequences. To understand and improve the reliability of ROS projects, it is critical to identify the types of misconfigurations developers face. To that end, we study ROS Answers, a Q&A platform, to identify and categorize misconfigurations during ROS development. We then conduct a literature review to assess the coverage of these misconfigurations using existing detection techniques. In total, we find 12 high-level categories and 50 sub-categories of misconfigurations. Of these categories, 27 are not covered by existing techniques. To conclude, we discuss how to tackle those misconfigurations in future work.

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 the author(s) 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.
ISSTA 2024, 16-20 September, 2024, Vienna, Austria

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-XXXX-X/18/06
<https://doi.org/XXXXXXXX.XXXXXXX>

CCS Concepts

• **Computer systems organization** → **Robotics**; • **General and reference** → *Empirical studies*.

Keywords

Robotic Systems, Misconfigurations, Empirical Study, Literature Review

ACM Reference Format:

Paulo Canelas, Bradley Schmerl, Alcides Fonseca, and Christopher S. Timperley. 2024. Understanding Misconfigurations in ROS: An Empirical Study and Current Approaches. In *Proceedings of ACM SIGSOFT International Symposium on Software Testing and Analysis (ISSTA 2024)*. ACM, New York, NY, USA, 13 pages. <https://doi.org/XXXXXXXX.XXXXXXX>

1 Introduction

The Robot Operating System (ROS), known as the “Linux of Robotics”, is the de facto open-source framework for building robot software [39]. ROS’s package ecosystem provides developers with reusable, off-the-shelf components that implement common robot functions (e.g., perception, planning, localization, drivers) [14, 30]. In theory, ROS allows developers to quickly prototype robot software by integrating such components and adjusting their parameters via configuration files (e.g., Launch XML, ROS Param YAML) to match their intended application and environment.

Despite the relative ease of integrating components using configuration files in ROS, *correctly* configuring systems presents a considerable challenge. For instance, the software’s configuration depends on the robot’s hardware and operating environment (e.g., specific types of sensors and their placement in the robot).¹ When communicating with each other, components must make matching assumptions about their environment (e.g., when the camera is 10 cm above the wheels, a reference transformation is required).² Finally, timeliness properties must be considered when integrating components to avoid negatively impacting robot behavior (e.g., frames may be dropped when processing high-resolution image streams, leading to unstable and dangerous motion).³

¹<https://answers.ros.org/question/59087>

²<https://answers.ros.org/question/227092>

³<https://answers.ros.org/question/248656>

To provide plug-and-play functionality, ROS components are typically required to make assumptions about the context in which they are used (e.g., a topic should receive messages of the correct type at a certain frequency) that are neither checked nor documented [14]. Misconfigurations occur when one or more components make different, conflicting assumptions about the robot, leading to unintended and potentially dangerous behavior (e.g., property damage, human harm) during deployment. Given the importance of safety within this domain, it is vital to identify misconfigurations before the robot is deployed, and as early as possible. To that end, the robotics software engineering community has begun to develop tools to detect certain misconfigurations, such as those related to physical units [28], architecture [48], and reference frames [27].

To systematically tackle the misconfiguration problem, it is critical to understand the types of misconfigurations that occur in the wild and whether existing tools are designed to detect them. Based on our own experiences with ROS, we know that physical units, architectural, and reference frames are not ROS’s only categories of misconfiguration. Software misconfigurations have been thoroughly studied in different contexts of software development (e.g., security [13, 40], databases [33, 57], cloud computing [24, 50, 59] and networks [36]). However, misconfigurations within ROS are inherently different due to their cyber-physical nature.

In this work, we set out to **identify the broader set of misconfigurations that impact ROS systems and to determine which detection techniques address them and which misconfiguration types are going undetected**. This knowledge can guide future research in the robotics software engineering community in developing novel tools and techniques to address them.

We first derive a taxonomy of misconfigurations within ROS systems by conducting an empirical study of relevant questions posted to ROS Answers, a ROS-specific Q&A site similar to Stack Overflow. Secondly, after devising a taxonomy of misconfigurations, we determine the extent to which state-of-the-art analysis tools help to address those issues by conducting a systematic literature review of analysis papers published at several major software engineering, architecture, testing, and robotics conferences. Finally, as part of our analysis, we highlight misconfigurations unaddressed by existing techniques and further discuss research opportunities in developing new techniques for them.

Through our study, we make the following contributions:

- A taxonomy of misconfigurations in ROS, based on a qualitative study of a popular Q&A platform (Section 3);
- A literature review of the state-of-the-art approaches and how they cover the misconfigurations (Section 4);
- A dataset of misconfigurations and questions manually analyzed and categorized that can be used to guide future studies and develop novel techniques (<https://anonymous.4open.science/r/artifact-study-misconfigurations-5D6A>).

2 Background

In this section, we provide a high-level introduction to ROS as a middleware for building component-based robotics software and as an open-source ecosystem of reusable components.

Systems in ROS are built as a collection of independent processes, known as *nodes* or *components*, that are each responsible

for providing certain functions (e.g., perception, planning, control, driver interfacing). At its core, ROS’s responsibility is to provide the “plumbing” that facilitates communication between distributed components. The bulk of communication within ROS follows an anonymous publish/subscribe pattern [43]. At one end, components (e.g., a camera driver) publish messages to named topics (e.g., `/camera/color/image_raw`). On the other hand, components (e.g., object detection) subscribe to those same topics to receive messages. Neither the publishers nor the subscribers are directly aware of one another’s identities (i.e., they are spatially decoupled [15]). Moreover, communications are defined at run-time via calls to the ROS API (e.g., strings are used to state topics by name).

One of ROS’s biggest strengths is its rich ecosystem of generic open-source components that can be reused for common robot functions [30, 35]. For example, MoveIt! [8] provides motion planning and execution for manipulation, `ros_control` [7] provides implementations of low-level controllers (e.g., velocity, joints, effort), and `ros_localization` [37] provides filters for state estimation (e.g., by fusing GPS and IMU data). By reusing off-the-shelf components, developers can, in theory, reduce the cost and complexity of building robot software. However, because these components are generic, they must be configured to work in specific contexts.

ROS uses configuration files to customize and compose those components into a functioning ensemble. These include Launch XML (or Python) files, which launch and compose each component within the system and ROS Parameter YAML files are used by components at run-time to customize their behavior (e.g., specifying a color format and topic name for camera images). General-purpose components use these configuration files to tailor their behavior to a particular robot, application, or environment. Particularly complex and variable components and subsystems (e.g., Nav2, MoveIt!, `ros_control`) go beyond providing a fixed set of parameters and embed a limited domain-specific language within the ROS parameter system.

Together, these aspects of ROS are designed to enable rapid prototyping, encourage component reuse (e.g., components can be swapped without modifying the rest of the system), and reduce the cost and complexity of building robots. However, due to ROS’s dynamic, spatially decoupled architecture, developers must follow assumptions and conventions when integrating components into their system, which are neither enforced nor documented [1, 14]. Failing to do so leads to *misconfigurations*. Often, such misconfigurations do not produce meaningful errors and are only uncovered via manual and laborious debugging of erroneous behaviors at run time.

3 Study of Misconfigurations

To understand which misconfigurations current tools can address, we first need to know the different types that exist. Therefore, we ask the following research question:

RQ1: What kinds of misconfigurations do developers make when building robot software systems with ROS?

To answer this question, we perform an empirical study of ROS Answers, the, until recently, primary Q&A forum for ROS [14].⁴

⁴In August 2023, Open Robotics decided to move to the Robotics Stack Exchange (robotics.stackexchange.com).

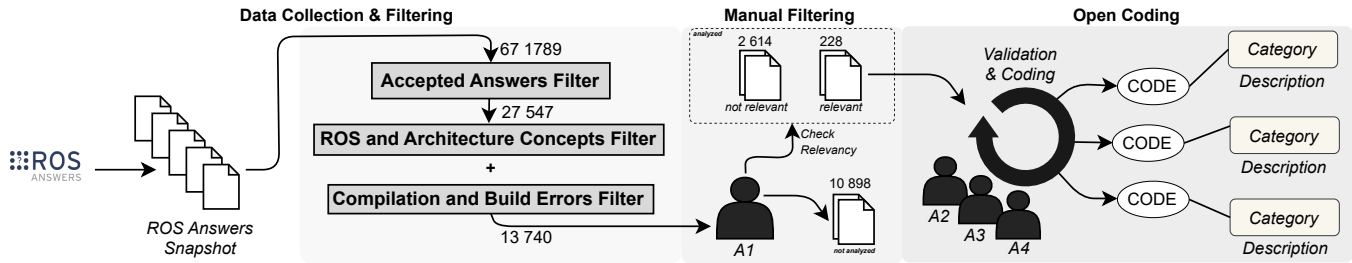


Figure 1: Three-step methodology for analyzing ROS Answers questions. Step 1 collects a snapshot of 67 189 questions. **Step 2** selects questions with accepted answers containing ROS or software architecture concepts and filters questions about compilation or building issues. **Step 3** produces a sample of 228 questions filtered based on their relevance. Questions are divided into stages, where their codings are iteratively improved.

Q&A platforms are designed for users to post their problems, including explaining the scenario where the problem occurs. While only some of the questions pertain to misconfiguration, we found many such instances in a prospective search. Q&A websites provide more detailed misconfiguration examples than social coding platforms (e.g., GitHub) since commit messages often lack important context, making it difficult or impossible to reliably identify misconfigurations. Furthermore, issue trackers often describe bugs in individual components rather than the difficulties of integrating those components into a working system. To that end, we perform a thematic analysis of questions posted to ROS Answers.

In the rest of this section, we describe our methodology (Section 3.1) and its associated threats to validity (Section 3.2) before presenting our taxonomy of ROS misconfigurations (Section 3.3).

3.1 Methodology

Figure 1 outlines our high-level methodology, which takes inspiration from studies of similar Q&A platforms [2, 46]. Below, we describe each step of our methodology.

Data Collection & Filtering. We first gathered all 67 189 questions posted to ROS Answers between January 1st, 2011, and November 20th, 2022. Figure 2 presents an example⁵ of a question and its accepted answer. We then filtered these questions to a set of 27 547 by selecting only those with an accepted answer as they offer an alternative perspective and accepted solution from the user.

During the second step, we narrowed the accepted questions to those referring to ROS and architectural concepts, expecting to obtain misconfiguration questions. Reducing the set of questions is a common practice in the literature [2, 4, 46]. We used common ROS concepts as defined in the ROS Wiki:⁶ node, subscribe topic, message, parameter, service, action, launch, publish, and subscribe. Since developers indirectly define their system’s architecture when changing configurations, we also identify questions related to software architecture errors [20]: architecture, mismatch, assumption, incompatibility, inconsistency, integration, and configuration. Subsequently, the set was further refined by removing questions containing keywords related to installation and build errors (e.g., `build error` or `compil*`). Build errors are comparatively easy

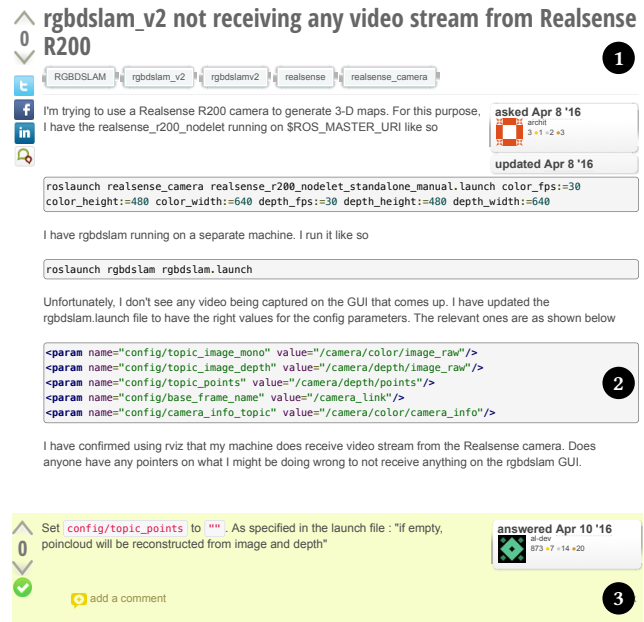


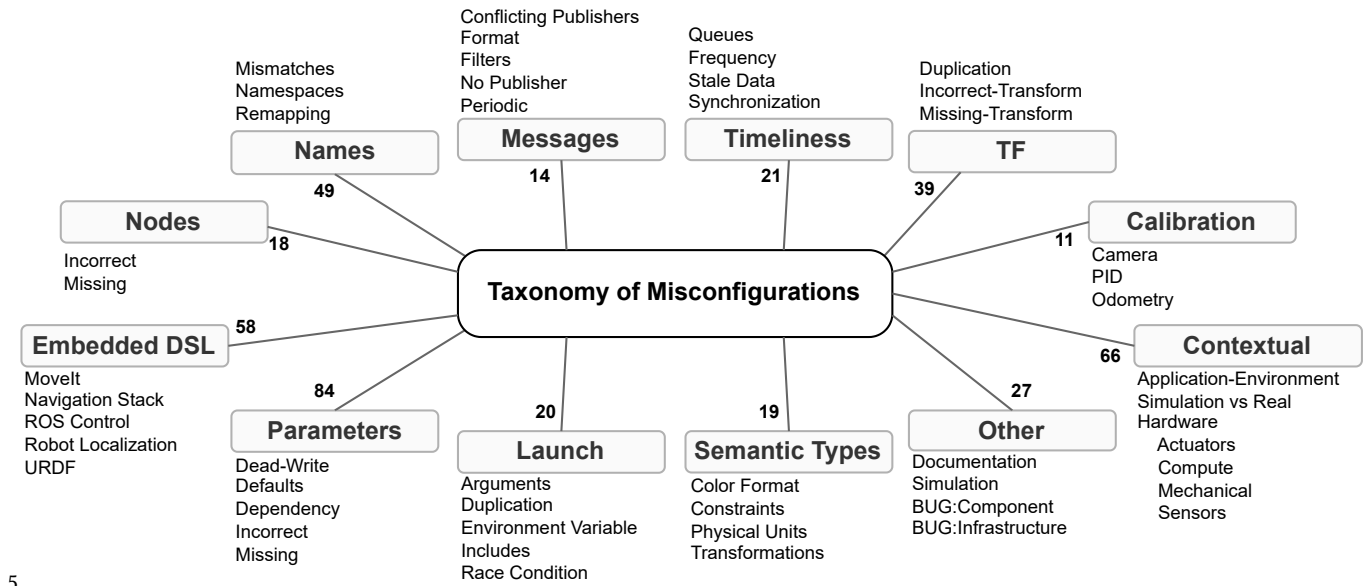
Figure 2: Example of a ROS Answers question corresponding to a misconfiguration where the developer incorrectly defined a parameter value. Each question contains a title (1), content with text and source code (2), and metadata about the author, date, and number of votes. Questions may include comments, answers, and an accepted answer (3).

to determine and diagnose, as developers can examine the error messages generated. This work focuses on undetected misconfigurations during the deployment process, subsequently impacting robots behaviors. These errors are significantly more challenging to detect and trace back to their source. This final filtering step provided a total of 13 740 accepted questions considered suitable for sampling.

Manual Filtering. The first author (A₁) randomly selected questions from the 13 740 questions and labeled them as relevant or irrelevant by analyzing their content to determine if it described

⁵<https://answers.ros.org/question/231458>

⁶<http://wiki.ros.org/ROS/Concepts>



5

Figure 3: Mindmap of the misconfigurations identified from the study presenting the 12 high-level categories of misconfigurations, and their 50 sub-categories level. Each misconfiguration contains the number of questions annotated with the code. Each question may refer to more than one misconfiguration.

a failed attempt to configure the system (i.e., interacting with the configurations or source code files).

Open Coding. The first author (A_1) iteratively provided sets of relevant questions to three other authors (A_2, A_3, A_4) who applied an open coding [16] approach to construct a taxonomy of misconfigurations. At each step, the authors individually proposed an updated set of codes for their sub-set of questions before discussing those codes and merging them into a revised taxonomy. The subsets were constructed such that each question was analyzed by three different authors, allowing a diversity of perspectives to be captured and reducing the error rate. This process continued until reaching the saturation point [21] (i.e., no further changes were made to the taxonomy after reaching the end of a step) after analyzing 228 questions, resulting in a final taxonomy of 12 categories and 50 subcategories of ROS misconfiguration.

Labeling. Finally, we labeled each of the 228 questions using the final taxonomy. Half of the questions were labeled by one pair of authors (A_1, A_2) and the other half by a different pair (A_3, A_4). We calculated the agreement by dividing the number of codes both authors agreed on by the total number of codes used. This first step led to an agreement of 84.12% and 85.5% for each pair of authors. Then, each pair compared the codes that differed by one code and adjusted their code if in agreement. Furthermore, to determine documentation-related questions, the authors collected all questions annotated with documentation, discussed random instances of these, and re-annotated all questions until they reached an agreement on using this code. Finally, the authors discussed 26 questions with initial disagreements. The authors who did not analyze a given question were arbiters during the discussion.

3.2 Threats to Validity

External Validity. We identify two primary external validity threats: the generality of our results to (a) different ROS versions and distributions and (b) expert users. The first threat relates to the possible predominance of ROS 1 over ROS 2 questions and the impact of ROS distributions. Given the relatively recent release of ROS 2 [34] in 2018, we expected more questions related to ROS 1. We found that the analyzed questions rarely specified ROS or distribution versions (37 out of 228), making it infeasible to determine version-specific misconfigurations. Furthermore, as there are few architectural differences between ROS versions, we believe that our findings generalize to both versions. The second threat concerns the applicability of our findings to real-world scenarios. By sampling data from a popular ROS Q&A platform, we are addressing real-world developer issues. Nevertheless, we recognize that industrial settings may present unique, undisclosed misconfigurations. Our taxonomy provides a basis for further studies in such contexts.

Internal Validity. We identify four main threats to internal validity: the initial sampling method, the generalizability of the sampled data, biases in question analysis, and potential misrepresentation of misconfiguration types in the sampled ROS Answers questions. The first threat regards the initial sampling and validation step with only one author, possibly introducing personal biases in the selection. To mitigate this step, we performed a preliminary study similar to the current one, in which all authors looked at a sample of questions and validated them as relevant or not relevant. To address the second threat to validity, we sampled relevant ROS Answers questions, inspecting, validating, and categorizing them until reaching theoretical saturation. This approach avoids introducing unaddressed categories of misconfigurations when analyzing new questions. We addressed the third by iteratively analyzing

questions, validating their relevance and the misconfiguration categories with at least two authors of the paper. The fourth threat relates to the sample’s representativeness. We focus on questions only related to ROS concepts and software architecture, which are more likely to contain misconfigurations. However, we acknowledge that other categories of misconfigurations may be overlooked in this study due to this restriction, but that does not invalidate the categories that we find.

3.3 Results

In this section, we describe each high-level category and sub-category of misconfigurations in detail. Figure 3 presents the mindmap of the taxonomy of misconfigurations. Below, we present examples of questions related to each Misconfiguration .

Messages. Most ROS communication occurs via messages exchanged between components over named topics. Through our analysis, we identified five misconfigurations related to messaging.

Format misconfigurations arise due to a mismatch between two or more components in the message format that are exchanged on a shared topic. For instance, the *initialpose* topic, representing the initial position and orientation of the robot in the map, accepts *geometry_msgs/Pose* messages. Both publishers and subscribers must respect the message format when exchanging messages on this topic. However, a misconfiguration occurs when either end breaks the “contract” and expects a different message.

Components may expect a specific number of publishers to the topics to which they subscribe. This assumption eases components’ expectations of the frequency of messages they receive.

No Publisher misconfigurations occur when a component subscribes to a topic no publisher sends messages to. The subscriber waits indefinitely for messages that never arrive. Figure 4 presents an example where the subscriber never receives messages due to a topic name mismatch. **Conflicting Publishers** misconfigurations appear when there is more than one simultaneous publisher to a topic that should only be accessed by a single publisher. This leads to messages that may provide opposite, conflicting instructions.

The `message_filters` API is used to filter incoming messages on a given topic (i.e., a callback is only triggered for messages that satisfy a given condition) and synchronize messages across multiple topics (e.g., invoke a callback once data has been received from multiple sensors). In particular, we observed difficulties in using this API (`Filters`) to synchronize two topics without losing messages. Finally, **Periodic** misconfigurations occur when the correct system execution relies on messages being periodically published at a specific frequency (e.g., camera, lidar, IMU data). Misconfigurations occur when a component stops publishing data continuously and other parts of the system continue to wait, indefinitely, to receive that data. For instance, a pedestrian detection node must consistently publish images regarding the pedestrian’s position estimates to function correctly.

Launch. Launch files are the primary means of orchestration within ROS, used to launch and glue together individual components with specific configurations into the system.

Developers often use the ability to import other launch files recursively to improve modularity and simplify reuse by writing individual launch files for separate sub-systems (i.e., a collection of

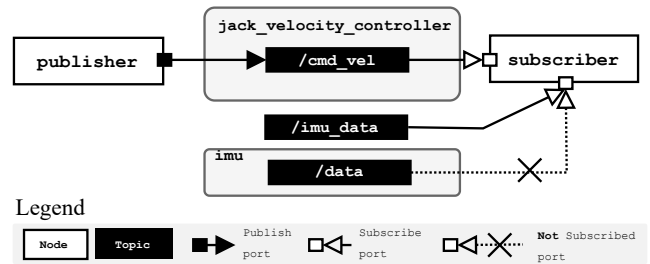


Figure 4: Example of a Name Mismatch from ROS Answers, where the developer mistyped the subscriber’s topic name. The subscriber expects data from `imu/data`, but due to the wrong connection to `imu_data`, no data is received.

components that work together to perform tasks such as perception, planning, or control). When writing a launch file for an entire system, developers introduce a layer of abstraction, requiring them to only reason about what launch files to include rather than worrying about configuring every individual node and subsystem. We observed cases where the developer either includes inappropriate launch files for the given context or fails to include crucial launch files necessary for the robotic system’s functioning.

We also observe node Duplication errors in launch files, where two or more nodes are instantiated with the same name (e.g., by accidentally launching the same node with the same name). In this case, ROS complains at run-time that the name is taken, and the second node crashes upon launch.

Since ROS architectures are defined at run-time, launch files are susceptible to Race Conditions . For instance, nodes may publish messages to topics before the subscriber finishes launching, causing those messages to be lost. Nodes can also be sensitive to the ordering of `<node>` and `<param>` tags within launch files: In ROS 1, a `<node>` may launch before the parameters are stored on the global parameter server, leading to parameter misconfigurations.

To allow components to be customized to a particular system, launch files support Arguments , whose values may be provided by the command line, a parent launch file, or a specified default value. Those values can be accessed via string interpolation (e.g., via `$(arg name-of-arg)`) within the launch file or through the command line, and are typically used to specify ROS parameter values, rename nodes, control the inclusion of particular nodes, or to set the system time. Arguments are prone to many of the same issues as parameters (e.g., dead writes and unintentional use of default values).

Environment variables are also used to customize the behavior of individual components in launch files. Misconfigurations can occur when necessary environmental variables for nodes are not specified or when incorrect values are assigned to those variables.

Parameters. In ROS, parameters are used to adjust the behavior of components to their intended deployment. Parameter values are typically provided by launch files, which are used to compose multiple components into a functioning ensemble. Listing 1 presents an example of a launch file with two parameter-related misconfigurations where the developer forgets to include a `tf_prefix`. This

```

1 <!-- robot urdf model -->
2 <param name="robot_description" command="cat_$(
3   find_urdf_pkg)/urdf/my_robot.urdf" />
4
5 <!-- robot state publisher node -->
6 <node pkg="robot_state_publisher"
7   type="state_publisher"
8   name="robot_state_publisher">
9
10   <param name="~tf_prefix"
11     value="robot_name"
12     type="str" />
13 </node>

```

Listing 1: Example where `tf_prefix` parameter required for operating with multiple robots is Missing, leading the system to use the Default Parameter value.

parameter is critical when working with multi-robot systems, to create separate transform trees for each robot. Since the parameter is not defined, the system, by omission, uses the default value, and a single transform tree is used for all the robot systems.

Since parameters are defined and used at run-time within ROS, components may unexpectedly crash when a required parameter is Missing, or behave in an unintended manner when the component falls back on a Default value for a missing parameter. Both of these types of parameter misconfiguration can also be caused by Dead Writes where the wrong name is used to specify a parameter (e.g., due to a typographical error or a name change refactoring). These cases can be hard to debug as there is no static checking and warnings may not be produced for missing or unused (i.e., dead-write) values.

Misconfigurations can occur when using Incorrect parameter values. Those values may be universally incorrect (e.g., out of bounds) or contextually incorrect for the given robot, environment, and application. When defining parameter values, developers also need to be cautious of potential Dependency issues, where the behavior of a given parameter is changed by the value of another parameter (e.g., a parameter that enables or disables a feature).

Semantic Types. Even when components correctly make assumptions about the message format shared on a given topic, they can still make incorrect assumptions about the message content (i.e., their semantic types). For instance, two components may correctly exchange a `sensor_msgs/Image`, but the publisher sends color images while the subscriber expects grayscale images.

Components implicitly assume that messages satisfy specific Constraints over their contents. For example, values are within certain bounds (e.g., positions, velocities, motor values), specific coordinate frames are used, or the range of laser scan measurements is respected. Type-related misconfigurations may stem from the improper use of images and point clouds, such as mismatched assumptions between components about the Color Format of the image or point cloud (e.g., grayscale vs. color images) or an incorrect assumption that all of the images and point clouds that are shared on a given topic have been subject to specific Pointcloud Transformations (e.g., resizing, compression, or color conversion) or Image Transformations.

Finally, misconfigurations can occur due to mismatched assumptions on the Physical Units of data that are exchanged between components. As robots interact with the real world, ensuring that the component’s physical units match is critical. For instance, a publisher describes rotational velocity using radians per second, but the subscriber expects the same quantity in degrees per second.

Names. Every node, topic, service, action server, and parameter within ROS has an associated name specified at run-time either as a field or property within a Launch XML file or as a string in the source code of a component.⁷ Since names are resolved at run-time, it is easy to introduce Mismatches between two or more components in the name of a resource (e.g., topic, parameter, service, or embedded DSL configurations). Figure 4 illustrates an example of one of these errors from our dataset, where the developer incorrectly defined the topic’s name in the subscriber. Since the name is incorrectly configured, the subscriber receives no messages as there are no publishers for that topic. Detecting this misconfiguration is challenging, as the topic names only differ by a specific character and are usually only detected during execution when the system does not behave as intended. Another instance of a name mismatch relates to typos when writing the configurations for the robot localization. For instance, the developer incorrectly defines the name of the `world_map`, a parameter specifying the frame treated as a fixed reference frame.

ROS implements a hierarchical naming structure to promote encapsulation. By convention, this system groups related resources and allows multiple instances of the same node to be used simultaneously (e.g., one node for each camera). ROS’s hierarchical naming structure is implemented via namespaces: Every resource belongs to a namespace, denoted by a forward slash in the name of that resource (e.g., `/right_camera/raw_image` belongs to `/right_camera`). Furthermore, namespaces may be stacked (e.g., `/vehicle_a/right_camera/raw_image`) where `/` denotes the root of the hierarchy, known as the global namespace. We observe that Namespaces misconfigurations occur most commonly when developers forget to use namespaces or otherwise use the wrong namespace, leading to naming collisions and causing the system to have an unintended architecture.

In addition to namespaces, ROS also relies on name Remapping to help encapsulation and reuse: ROS launch files allow the mapping of the name of a certain resource (e.g., topic, service) onto a different name within the context of a particular component. This ability is used to wire a general-purpose component into the necessary configuration for a specific system. For instance, when using the `image_proc/resize` component⁸ to resize camera images, users must remap the subscribed `image` topic (i.e., incoming camera images) and published `~image` topic (i.e., resized image) onto appropriate source and destination topics (e.g., `/right_camera/image_color` and the topic target named `/right_camera/image_color_resized`). Incorrect or missing remappings can lead to message loss and unintended behavior.

Nodes. Nodes (i.e., components) are the processes within ROS that collectively form a working robot by performing computation, exchanging information, and interacting with the robot hardware.

⁷<http://wiki.ros.org/Names>

⁸http://wiki.ros.org/image_proc

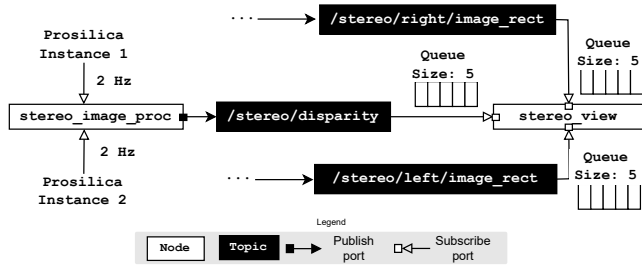


Figure 5: Example of a misconfiguration where the developer gets images from stereo_image_proc at 2Hz, converting them to disparity images. The stereo_view needs matching left, right, and disparity images, but the slow processing speed of disparity images causes left and right images to arrive faster and fill the queue, dropping messages. The stereo_view skips these until a matching of images is available.

In our study, we observe cases where either crucial components are **Missing** or an **Incorrect** component is used for the particular robot, environment, or intended application. For example, the developer incorrectly uses a planner, which leads the robot to navigate erroneously, possibly against walls. A special case for missing a component is a missing nodelet manager. In this case, developers do not define the qualified manager name, allowing all nodelets to be assigned to that manager.

Timeliness. Robots are real-time systems: messages must be sent between components (e.g., control signals, state estimates) by a certain deadline to prevent unintended and dangerous behavior. Timeliness misconfigurations occur when the timing assumptions of interacting components are mismatched.

Both publishers and subscribers within ROS 1 have associated **Queues**, which are used to buffer either outgoing or incoming messages. Determining an appropriate size for this queue is crucial to ensure both timely and correct behavior: A queue size that is too small can lead to message loss, whereas an overly long queue can lead to excessive compute resource usage and message delays. For instance, Figure 5 presents a queue misconfiguration within a robot with stereo vision. The stereo_view node requires three images: left, right, and the disparity image (i.e., the difference between left and right). In the proposed architecture,⁹ the disparity image is computed using the other images. However, computing the disparity image takes time and computational resources. Due to the desynchronization of the sensors, when the disparity image is ready, the original left and right images have already been overwritten in their corresponding queues. To fix this misconfiguration, the developer can increase the queue size to avoid overwriting or throttle the publishing rate of the cameras to account for time taken to produce the disparity image.

This example is an instance of **Frequency** misconfiguration between publishers and subscribers, where subscribers expect to receive messages at a given frequency to operate correctly.

To function safely, certain components rely on an uninterrupted stream of data that accurately describes the state of the robot and

its environment. For example, motion planning in a dynamic environment requires timely and accurate estimates of the position of both the robot and potential obstacles in the scene. Unintended and unsafe behavior can occur when **Stale Data** no longer accurately represent the current state of the robot and its environment, as messages are not published at a high enough frequency.

Finally, **Synchronization** of certain messages is essential for correct system operation. For instance, in example to showing the queue misconfiguration, Figure 5 presents an example where the sensors used are sources of multiple misconfigurations. One source of misconfiguration is the lack of synchronization between camera frequencies. There is a mismatch when the stereo_view processor receives the information from each camera. Synchronization of the receiving data is required to ensure that all images are consumed simultaneously, keeping the original images available before consuming the corresponding disparity image.

TF. Robot systems typically rely on a large number of 3D coordinate frames to reason about the relative position and orientation of the robot, its physical parts, and its environment. tf ¹⁰ is a core ROS library that uses a special tree structure to allow users to transform between coordinate frames (e.g., to determine the position and orientation of the robot’s gripper relative to the robot’s base).

We observe three major types of TF-related misconfigurations:

Incorrect Transform, either due to a typo, a misunderstanding of the transform tree semantics, or a mistake about the geometry of the robot. **Missing Transform**, where the developer forgets to provide a transformation between a parent and child frame. Finally,

Duplication of transforms, where the developer publishes the same TF transform from multiple conflicting sources. (i.e., there should be a single source of truth).

Embedded DSL. ROS, its associated toolchain, and some of its most popular (and general purpose) packages rely on their own custom configuration formats. Some ROS packages embed their configuration formats inside of ROS’s parameter system, effectively forming an embedded DSL. In other cases, a standalone file is used (e.g., URDF). We observed issues related to the use and configuration of the Navigation Stack, *ros_control*, *robot_localization*, MoveIt!, and URDF. These misconfigurations contain more specific types of issues (e.g., parameter issues). This high-level category presents the types of configuration files related to these misconfigurations.

Navigation Stack provides mobile robots with the ability to use odometry and sensor values to localize their position and navigate within a 2D plane by sending velocity commands to the mobile base.¹¹ Specifically, we saw issues concerning the correct definition of motion planning parameters. For instance, developers may select an inappropriate planner or fail to adapt the parameters of that planner to the robot and environment.

URDF (Universal Robot Description Format) files describe robots in terms of their links, joints, transmissions, sensing capabilities, collision geometry, and physical properties (e.g., inertia, contact coefficients, joint dynamics).¹² URDF files are used for visualization, simulation, and motion planning. For instance, when developers

⁹<https://answers.ros.org/question/9108>

¹⁰<http://wiki.ros.org/tf>

¹¹<https://wiki.ros.org/navigation>

¹²<https://wiki.ros.org/urdf>

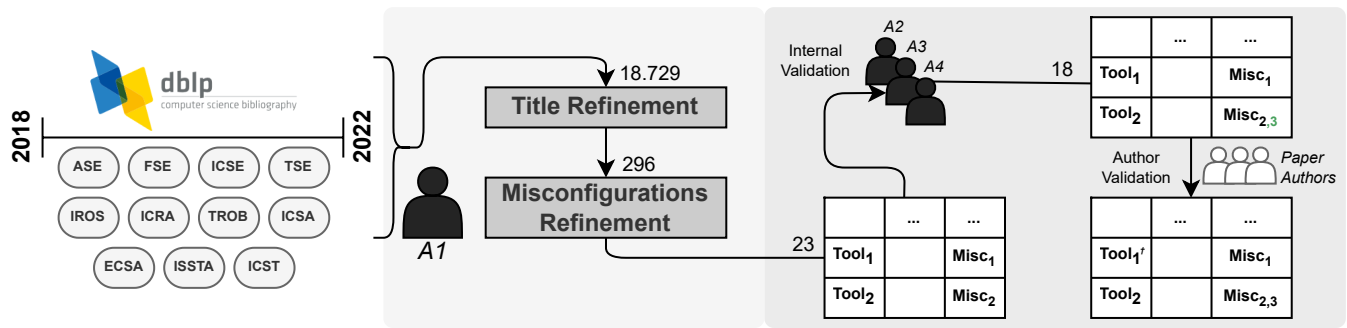


Figure 6: Methodology of the literature review. We collected prior work from 2018 to 2022 from top conferences and journals in software engineering, robotics, architecture, and testing. We refined the search by looking at the paper titles, then reading the papers and matching them to misconfigurations. We perform an internal validation with three authors and author validation with the authors of each technique to validate the proposed tool for misconfiguration matching.

forget to specify the joints between two links or incorrectly provide the robot description of two systems in the same file.

MoveIt is a platform for building manipulators using ROS that incorporates algorithms for motion planning, manipulation, kinematics, control, and navigation [8]. Successfully configuring it for a robot in a particular environment relies on careful configuration of numerous parameters (e.g., planner density, and padding offsets).

ROS Control allows developers to integrate and compose multiple off-the-shelf control algorithms into their system. To behave safely and operate as intended, each controller needs to be adapted to the specifics of each robot, which, when done incorrectly, can result in a misconfiguration.

Finally, **Robot Localization** provides a node collection for performing state estimation (i.e., Kalman filters) and integrating GPS data: they fuse data from multiple sensor sources (e.g., IMUs, GPSs, odometry) to obtain a robust estimate of the robot’s state (i.e., position, rotation, velocity). We observe issues that stem from missing or incorrect transforms, incorrect units, and sensor mismatches (e.g., attempting sensor fusion without wheel or visual odometry).

Calibration. Robots rely on a suite of sensors to perceive their environment. To ensure that the robot’s understanding of its environment is and remains accurate, those sensors must be calibrated. In our study, we observed misconfigurations due to the miscalibration of **Cameras**, **Odometry**, and **PID** controller parameters, all of which required a manual change. In all of these cases, the mistake was either (a) forgetting to calibrate entirely (e.g., camera intrinsics and extrinsics), (b) relying on default parameters that were inappropriate for the robot (e.g., PID defaults), or (c) using incorrect values (e.g., wheel radius).

Contextual. Misconfigurations can also occur when the system’s configuration is tweaked according to a specific context. A simple example is developing the robot’s software within a simulation context and deploying it to a physical robot or vice versa. Some configurations and parameters need to be different

when changing contexts, leading to **Simulation vs. Real** misconfigurations when the behavior does not match (typically by not adjusting the necessary configurations). Furthermore, we encountered **Application-Environment** misconfigurations where the component’s configuration depends on the type of application of the system and the surrounding environment. For instance, the configurations of the components inside and outside a warehouse can be different due to weather, lighting conditions, and surface grip.

The robot’s physical configuration (i.e., its sensors, actuators, mechanical components) naturally imposes restrictions over the set of plausibly correct software configurations; configuring the robot’s software therefore requires an understanding of the robot’s physical configuration. We identified four sub-categories where **Hardware** components require careful software configuration: (1) **Actuators**, where the configuration of the parameters for a component in the system depends upon the exact actuators that are used, (2) **Sensors**, where the type of sensors and their positioning on the system restricts parameter configuration, (3) **Mechanical**, where the robot’s mechanical hardware imposes restrictions of the space of meaningful and safe parameters, and (4) **Compute** where the configuration impacts resource usage. Furthermore, robots may have specific hardware limitations that must be addressed via software configuration. For instance, one source of misconfiguration in Figure 5 is resource-related, where components take longer to process overly large camera images.

Other Challenges. During the analysis, we encountered questions not related to misconfigurations. **Documentation** questions occur due to missing or outdated documentation in ROS. This category is not considered a misconfiguration as the developer could not progress with the configuration and reach a misconfiguration due to the lack of documentation. We highlight this category as part of our taxonomy since outdated or lack of documentation may cause developers to introduce misconfigurations. Developers may not know what components to use or how to use them correctly in their systems. We also encountered instances where the issues were related to the component having a **BUG:Component** rather than being a configuration issue, or the **BUG:Infrastructure** in which the system is running contained an error (e.g., **RViz**). Finally, we

identified cases where the misuse of Simulation prevents the correct functioning of the system. Developers can use simulation time to playback previously recorded data (e.g., sensor readings) in a time-synchronized manner (e.g., during testing and debugging). However, developers unfamiliar with the concept of simulation time may forget to define it, resulting in incorrect system behavior (e.g., different times in the receiving of data).

4 Study of Existing Tools

After categorizing the different types of misconfigurations in ROS-based systems, it is essential to identify which categories are overlooked when using existing state-of-the-art analysis techniques. To that end, we address the following question:

RQ2: To what extent do current techniques address these categories of misconfiguration?

Answering this question reveals the gaps in current tooling and helps guide the development of new tools to increase the coverage of misconfigurations that can be detected earlier in the development process, avoiding errors in deployment. To address this objective, we conducted a literature review following best practices outlined by Snyder et al. [44]. Our search strategy, influenced by Albonico et al. [3], focused on sourcing works from the software engineering, software architecture, software testing, and robotics communities, which are the likely places for such techniques to have been reported. Figure 6 illustrates our methodology for performing the literature review, which we describe below.

4.1 Methodology

Our literature review methodology consists of three stages:

Stage 1. Collection: Using DBLP,¹³ we collected all papers of journals and conferences and associated co-located events between 2018–2022 inclusive from the major software engineering (ICSE, FSE, ASE, TSE), software testing (ICST, ISSTA), software architecture (ICSA, ECSA), and robotics (ICRA, IROS, TROB) venues. We gathered 18 729 paper links, titles, venues, and respective years.

Stage 2. Refinement & Collection: The first author started by manually inspecting paper titles and searching for keywords that describe any of the subcategories of misconfigurations. Furthermore, for non-robotics venues, we searched in the titles for robotics-related topics. In contrast, in the robotics venues, we searched for concepts related to verification, testing, and repairing misconfigurations. Whenever a prior work seemed relevant, we manually inspected it by looking at the abstract, followed by the introduction, approach, and conclusion.

Then, we annotated the relevant papers with the misconfigurations they cover and their type of analysis (static or dynamic). We consider a technique static if it performs verification without executing the system and dynamic otherwise. Although general-purpose testing techniques potentially cover all the identified misconfigurations in theory, we only considered those that explicitly cover them in their problem statement or examples provided in their evaluation. This refinement considered 18 relevant papers.

Stage 3. Internal & Authors Validation: We performed a two-step validation of the matching each technique and the misconfiguration. We first validated our findings internally with the

other authors of this study (A_2, A_3, A_4), experts in software architecture, software engineering, and robotics, to mitigate against missing or incorrect categorizations. We then asked the authors of each technique to validate our findings externally. Since ROS Answers questions do not contain executable examples to test each technique’s ability to detect a particular kind of misconfiguration, the external author validation helps us validate our assumptions of each tool’s ability. To perform an external author validation, we gathered the contact details for the authors of each technique and emailed them to asking if our classification is correct and what other misconfigurations, if any, their technique addresses.¹⁴ We analyzed the answers received by the papers’ authors. When in disagreement, we compare both mappings and re-analyze the paper. From the proposed misconfiguration mapping, 3 out of the 10 authors of each technique proposed an update by adding only 1 extra category. For all but one of these addition, we agreed with the correction and updated our categorization. The sole case where we disagreed with the authors’ additional categories was Santos et al. [42], which allows developers to synthesize run-time monitors and discover bugs by writing properties using the HAROS Property Language (HPL) [41]. Similar to general-purpose test cases, these properties act as tests that may detect specific instances of misconfiguration but do not cover the overall *category* of misconfiguration. While this extensibility is essential, as is writing integration tests in general, we consider this and other similar general-purpose testing techniques outside the scope of this literature review as, with the appropriate instrumentation and scaffolding (i.e., test inputs), testing tools can theoretically identify any misconfiguration, even if it is seldom practical.

4.2 Threats to Validity

External Validity. We identify the generalizability to other venues and the time frame selected as external threats. This literature review focused on searching for relevant work at top conferences and journals in software architecture, software engineering, software testing, and robotics. However, our findings may miss relevant tools by focusing on specific conferences and a particular time frame (2018–2022). Extending the search for further years and conferences and journals presents a challenging task as the intersection of multiple research areas quickly increases the prior work required for manual inspection.

Internal Validity. We identify the manual inspection of prior work and the subjectivity in categorization as internal threats. The initial paper refinement was done by only one of the authors, and there is a threat when selecting relevant prior work based on the paper titles where the author may overlook relevant papers. By searching for terms related to misconfigurations and the verification, testing, and repair of robotic systems, we expect to mitigate this threat. There is also a threat when performing the mapping of each technique to a misconfiguration. To mitigate this mapping, each technique was analyzed by two authors, and we tried to validate the mapping with the original paper authors.

¹³<https://dblp.org>

¹⁴At the time of submission, 10 of the 18 techniques were validated by the authors of those techniques.

Table 1: Overview of each technique, the type of analysis (Dynamic or Static), and the sub-categories of misconfigurations each addresses. The main categories of misconfigurations are as follows: (Ca) Calibration, (Co) Contextual, (M) Messages, (N) Names, (O) Other, (P) Parameter, (T) Semantic Types, † Author validated, * Authors Disagreement.

Reference	Venue	Year	Analysis	Misconfigurations
Kate et al. [28] [†]	FSE	2018	S	(T) Physical Units
Burgueno et al. [5] [†]	ICSE-RoSE	2018	S/D	(T) Physical Units (T) Constraints
Witte et al. [53]	ICSE-RoSE	2018	S/D	(N) Mismatches (M) No Publisher
Wuest et al. [54]	ICRA	2019	D	(P) Missing
Cramariuc et al. [9]	ICRA	2020	D	(Ca) Camera (Co) Sensors
Carvalho et al. [6] [†]	IROS	2020	D	(T) Constraints (M) Format (M) No Publisher
Wigand et al. [52]	IROS	2020	S	(Co) Actuators (Co) Application-Environment (Co) Mechanical
Kate et al. [27]	FSE	2021	S	(TF) Incorrect Transform (TF) Missing Transform
Jung et al. [26] [†]	FSE	2021	D	(P) Incorrect (P) Dependency (P) Defaults (Co) Application-Environment
Kortik et al. [31]	ICRA	2021	S/D	(M) Format (M) No Publisher (M) Conflicting Publishers
Santos et al. [42] ^{†*}	ICSE-RoSE	2021	S/D	(N) Mismatches (M) No Publisher
DeVries et al. [11] [†]	ICSE-SEAMS	2021	S/D	(Co) Application-Environment (Co) Mechanical (Co) Sim-vs-Real
Taylor et al. [45] [†]	ASE	2022	S	(T) Physical Units (TF) Incorrect Transform (TF) Missing Transform
Kim et al. [29] [†]	FSE	2022	D	(DSL) URDF (P) Incorrect (Co) Sim vs Real (Co) Sensors (Co) Actuators (BUG) Infrastructure
Das et al. [10]	ICRA	2022	D	(Co) Sensors (Ca) Camera
Heiden et al. [23] [†]	ICRA	2022	D	(O) Simulation
Timperley et al. [48] [†]	ICSA	2022	S	(M) Conflicting Publishers (M) No Publisher (M) Format (N) Mismatches (N) Remapping (P) Dead Write (P) Incorrect
Han et al. [22]	ICSE	2022	D	(P) Incorrect (P) Dependency

4.3 Results

Table 1 presents the mapping between the techniques and the misconfigurations. We identify the misconfigurations each technique addresses, the type of analysis it performs, static or dynamic, and the venue. We identified techniques that statically or dynamically detect misconfigurations, although some tools are static techniques whose verification is optionally extended with dynamic analysis. For instance, Burgueno et al. [5] statically analyze physical units described by a modeling language and generate model invariants checked during the execution of the system.

We also identified techniques that do not detect misconfigurations but rather automatically infer or optimize configuration parameters. The inference of the configurations helps prevent misconfigurations, as developers do not configure the system manually. For example, Wuest et al. [54] automatically infers geometric and inertia parameters, and Heiden et al. [23] probabilistically infers simulation parameters.

For each technique, we describe the misconfigurations it addresses with an important caveat. Although a technique addresses a specific misconfiguration, it does not mean it is solved. Some techniques address particular misconfigurations for specific robotic systems (e.g., Swarmlab) and contain limitations. For instance, Phys [28] and SA4U [45] both detect physical unit misconfigurations. While Phys performs static analysis, allowing it to detect physical unit errors before execution, SA4U requires execution information to detect the misconfigurations.

Overall, we identified 18 related works that address 23 of 50 sub-categories of misconfigurations. Parameters, Messages, and Contextual are the misconfigurations most addressed by current

techniques. On the other hand, no misconfigurations related to the Timeliness, Nodes, and Launch categories are currently addressed, and within the Embedded DSL category, only the URDF dialect is considered in current techniques.

5 Related Work

In this work, we studied the types of misconfigurations that developers face and what techniques can address them. As the presented misconfigurations are not specific to any ROS version or distribution, we expect our findings to generalize to the many ROS systems.

Prior work studied types of bugs in robotic systems and autonomous vehicles. For instance, 27.25% of the bugs in autonomous vehicles (AV) software detected are misconfigurations [18]. Instances of the misconfigurations we encountered were also found in Unmanned Aerial Vehicles (UAV) [51]. 19.6% of the bugs encountered in a study of two UAVs, PX4 and Ardupilot, are misconfigurations such as parameter misuse and missing, parameter limits related to hardware, and inconsistencies related to sensors and libraries. Our taxonomy of misconfigurations not only addresses ROS-specific issues but also other types of misconfigurations related to the cyber-physical nature of these systems.

Similar to the misconfiguration we encountered, the simulation to real-world transition is challenging in the services robotics domain [19]. Missing dependent components are also described in the literature [17] and presented in our taxonomy through particular instances of missing nodes and missing nodelet manager. Physical unit misconfigurations within ROS have also been quantitatively studied through projects on GitHub [38]. Issues related to URDF

files presented in this study, through the scope of Xacro XML language, are also a source of misconfigurations in prior work [2]. ROSDiscover [48] and ROSInfer [12] identify misconfigurations related to the structural composition and behavioral interaction between components, covering certain misconfigurations within our Naming, Parameters and Messages, and Timing categories. Finally, the ROBUST project [47, 49], a large-scale study of bugs in ROS, evidenced different types of misconfigurations related to missing runtime dependencies (e.g., nodes and configuration files), dangerous defaults (e.g., missing setting a parameter value), namespaces misconfigurations, and name mismatches. Unlike prior work, we focus on how developers misconfigure their systems and which techniques are available to address them.

Configuration errors are not specific to robotic systems, and prior surveys found these in other configurable systems [56]. For instance, some open source storage systems are prone to inconsistency errors of parameter values and value inconsistencies [58], similar to the incorrect parameter and dependent parameter errors we encountered, respectively. Similarly, Android manifests are also a source of incorrect attribute names and values [25]. Finally, configuration errors in databases arise according to the data types used and the ranges of accepted values for the configuration parameters [55].

6 Discussion

Misconfigurations are a critical concern in robotic systems, as these lead to unintended and potentially dangerous behavior. Our study and literature review identified a gap in the ability of state-of-the-art analysis tools to cover the space of ROS misconfigurations. In this section, we discuss some of the requirements that future analysis tools need to satisfy to improve the detection of the different categories of misconfigurations.

Misconfiguration analysis must work with ROS’s domain-specific languages and dialects. Building ROS Systems requires changing configuration elements distributed in multiple different configuration formats (e.g., Navigation Stack, MoveIt!, URDF). Through our study, we observe that these formats are a source of misconfiguration (Embedded DSL). As manually tracking many component configurations across different file formats and ensuring that they are consistent is challenging, automated techniques must detect misconfigurations spread throughout these files. Through our literature review, we observe that only one technique explicitly treats these DSLs as first class entities as part of its verification [29].

One avenue to address this concern is considering the DSL configurations in analysis tools. Current analysis tools do not explicitly consider the semantics of the configurations embedded within the different embedded DSLs. A future direction in improving the detection of misconfigurations is to incorporate this knowledge into analysis tools to enhance their capabilities. Alternatively, these separate configuration formats and files could be merged into a single analysis, verifying that configurations are correctly integrated across DSLs.

Misconfiguration analysis require information about the robot’s physical environment, hardware, and intended application to reliably detect misconfigurations. As cyber-physical

systems that interact with the real world, correctly defining robot configurations depends on the context in which the system is used. For instance, in Section 3.3: Contextual , we identified misconfigurations arising from the lack of knowledge when changing configurations that depend on the environment, hardware, and type of application. For instance, the positioning of sensors in the hardware, indoor and outdoor environments, the frequency, quality, and size of images provided by the sensors, and the type of robotic system are all factors found in this work that impact the configuration of software components. If analysis tools intend to improve upon their verification, they must consider this contextual information. However, as the physical environment and hardware information is often missing, our literature review found that only 6 of 18 consider this information to optimize the configuration values.

Future work can provide application, physical environment, and system hardware information to analysis tools to improve their verification through two possible approaches: domain-specific languages and artifact mining. Introducing a domain-specific language allows developers to specify properties regarding the context in which the system is executed (e.g., whether it is executed indoors or outdoors). When existing, contextual information can be obtained by analyzing existing artifacts (e.g., Phys [28], ROSDiscover [48], SA4U [45]), and inspecting the information these dialects provide.

Static analysis is not sufficient to detect all misconfigurations. Tools must be able to analyze run-time behavior. As executing the system is expensive, time-consuming, and possibly dangerous, it is ideal to detect these misconfigurations prior to system execution using static analysis tools. While static analysis techniques perform great work in reducing the cost of detecting misconfigurations, these are still bounded by the limited context understanding, which is not provided in ROS-based systems, and have scalability issues, being challenging to analyze large codebases [32]. Detecting some misconfigurations requires complex runtime behavior information not available to static analysis tools. For instance, a set of parameter values may need to be corrected according to the system’s execution (Incorrect Parameters),¹⁵ the incorrect calibration of the system is only detectable when executing the cameras (Calibration),¹⁶ or compute issues may arise when hardware actively interacts with the real world (Contextual).¹⁷

Future analysis tools can improve their detection of misconfigurations by augmenting the static checking with properties. Future analysis tools for robotics could use runtime behavior information to generate test cases, monitor runtime configurations, or use machine learning techniques to predict potential misconfigurations based on contextual information and system’s execution. These different approaches for analysis are currently used in techniques observed in our literature review and present a promising research direction. For instance, HAROS [41] generates monitors to track properties during runtime, SA4U [45] instruments the source code to obtain runtime information to help detect physical unit misconfigurations and Swarmbug [26] which performs multiple executions of the system while removing environment configuration variables to detect configurations responsible for the buggy behaviors.

¹⁵<https://answers.ros.org/question/30235/>

¹⁶<https://answers.ros.org/question/10975/>

¹⁷<https://answers.ros.org/question/195186/>

Furthermore, future research can also focus on system properties defined using specification languages. Properties defined in these domain-specific languages can help monitor misconfigurations only detectable dynamically while interacting with the real world.

7 Concluding Remarks

In this work, we conduct an empirical study to categorize the types of misconfigurations that occur within ROS and to determine the extent to which those misconfigurations are covered by existing analysis tools. We find 50 categories of misconfiguration, of which 27 are found not to be addressed and 23 are partially addressed by existing tools. Through this study, we identify promising areas for future research, discuss some of the requirements for future analysis tools, and, through our taxonomy, identify where detailed specific datasets are needed to drive the development of tools to tackle particular kinds of misconfiguration.

Acknowledgments

This work was supported by Fundação para a Ciência e Tecnologia (FCT) in the LASIGE Research Unit under the ref. (UIDB/00408/2020, UIDP/00408/2020 and EXPL/CCI-COM/1306/2021), the CMU Portugal Dual PhD program (SFRH/BD/151469/2021), and NSF-USDA-NIFA #2021-67021-33451. The authors would like to thank Bogdan Vasilescu and the Squareslab group for their feedback on this work.

References

- [1] Afsoon Afzal, Claire Le Goues, Michael Hilton, and Christopher Steven Timperley. 2020. A Study on Challenges of Testing Robotic Systems. In *International Conference on Software Testing, Validation and Verification*. 96–107.
- [2] Nicholas Albergro, Vivek Rathi, and John-Paul Ore. 2022. Understanding Xacro Misunderstandings. In *International Conference on Robotics and Automation*. 6247–6252.
- [3] Michel Albonico, Milica Dordevic, Engel Hamer, and Ivano Malavolta. 2023. Software engineering research on the Robot Operating System: A systematic mapping study. *Journal of Systems and Software* 197 (2023), 111574.
- [4] Michel Albonico, Ivano Malavolta, Gustavo Pinto, Emitza Guzman, Katerina Chinnappan, and Patricia Lago. 2021. Mining energy-related practices in robotics software. In *International Conference on Mining Software Repositories*. 483–494.
- [5] Loli Burgueño, Tanja Mayerhofer, Manuel Wimmer, and Antonio Vallecillo. 2018. Using physical quantities in robot software models. In *International Workshop on Robotics Software Engineering*. 23–28.
- [6] Renato Carvalho, Alcino Cunha, Nuno Macedo, and André Santos. 2020. Verification of system-wide safety properties of ROS applications. In *International Conference on Intelligent Robots and Systems*. 7249–7254.
- [7] Sachin Chitta, Eitan Marder-Eppstein, Wim Meeussen, Vijay Pradeep, Adolfo Rodríguez Tsouroukdissian, Jonathan Bohren, David Coleman, Bence Magyar, Genaro Raiola, Mathias Lüdtke, and Enrique Fernández Perdomo. 2017. ros_control: A generic and simple control framework for ROS. *The Journal of Open Source Software* 2 (2017), 456.
- [8] David Coleman, Ioan A. Şucan, Sachin Chitta, and Nikolaus Correll. 2014. Reducing the Barrier to Entry of Complex Robotic Software: a MoveIt! Case Study. *Journal of Software Engineering for Robotics* 5, 1 (May 2014), 3–16.
- [9] Andrei Cramariuc, Aleksandar Petrov, Rohit Suri, Mayank Mittal, Roland Siegwart, and Cesar Cadena. 2020. Learning Camera Miscalibration Detection. In *International Conference on Robotics and Automation*. 4997–5003.
- [10] Sandipan Das, Navid Mahabadi, Addi Djikic, Cesar Nassir, Saikat Chatterjee, and Maurice F. Fallon. 2022. Extrinsic Calibration and Verification of Multiple Non-overlapping Field of View Lidar Sensors. In *International Conference on Robotics and Automation*. 919–925.
- [11] Byron DeVries, Erik M. Fredericks, and Betty H. C. Cheng. 2021. Analysis and Monitoring of Cyber-Physical Systems via Environmental Domain Knowledge & Modeling. In *International Symposium on Software Engineering for Adaptive and Self-Managing Systems*. 11–17.
- [12] Tobias Dürschmid, Christopher S. Timperley, David Garlan, and Claire Le Goues. 2024. ROSInfer: Statically Inferring Behavioral Component Models for ROS-based Robotics Systems. In *International Conference on Software Engineering*.
- [13] Birhanu Eshete, Adolfo Villafiorita, and Komminist Weldemariam. 2011. Early Detection of Security Misconfiguration Vulnerabilities in Web Applications. In *International Conference on Availability, Reliability and Security*. 169–174.
- [14] Pablo Estefo, Jocelyn Simmonds, Romain Robbes, and Johan Fabry. 2019. The Robot Operating System: Package reuse and community dynamics. *Journal of Systems and Software* (2019), 226–242.
- [15] Patrick Th. Eugster, Pascal A. Felber, Rachid Guerraoui, and Anne-Marie Kermarrec. 2003. The many faces of publish/subscribe. *Computing Surveys* 35, 2 (Jun 2003), 114–131.
- [16] Douglas Ezzy. 2013. *Qualitative Analysis*. Routledge.
- [17] Anders Fischer-Nielsen, Zhoulai Fu, Ting Su, and Andrzej Wasowski. 2020. The forgotten case of the dependency bugs: on the example of the robot operating system. In *International Conference on Software Engineering, Software Engineering in Practice*. 21–30.
- [18] Joshua Garcia, Yang Feng, Junjie Shen, Sumaya Almane, Yuan Xia, and Qi Alfred Chen. 2020. A comprehensive study of autonomous vehicle bugs. In *International Conference on Software Engineering*. 385–396.
- [19] Sergio García, Daniel Strüber, Davide Brugali, Thorsten Berger, and Patrizio Pelliccione. 2020. Robotics software engineering: a perspective from the service robotics domain. In *European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 593–604.
- [20] David Garlan, Robert Allen, and John Ockerbloom. 2009. Architectural Mismatch: Why Reuse Is Still So Hard. *IEEE Software* 26, 4 (2009), 66–69.
- [21] Barney G Glaser and Anselm L Strauss. 2017. *Discovery of grounded theory: Strategies for qualitative research*. Aldine.
- [22] Ruidong Han, Chao Yang, Siqi Ma, Jianfeng Ma, Cong Sun, Juanru Li, and Elisa Bertino. 2022. Control Parameters Considered Harmful: Detecting Range Specification Bugs in Drone Configuration Modules via Learning-Guided Search. In *International Conference on Software Engineering*. 462–473.
- [23] Eric Heiden, Christopher E. Denniston, David Millard, Fabio Ramos, and Gaurav S. Sukhatme. 2022. Probabilistic Inference of Simulation Parameters via Parallel Differentiable Simulation. In *International Conference on Robotics and Automation*. 3638–3645.
- [24] Peng Huang, William J. Bolosky, Abhishek Singh, and Yuanyuan Zhou. 2015. ConfValley: a systematic configuration validation framework for cloud services. In *European Conference on Computer Systems*. 19:1–19:16.
- [25] Ajay Kumar Jha, Sunghee Lee, and Woo Jin Lee. 2017. Developer mistakes in writing Android manifests: an empirical study of configuration errors. In *International Conference on Mining Software Repositories*. 25–36.
- [26] Chijung Jung, Ali Ahad, Jinho Jung, Sebastian G. Elbaum, and Yonghui Kwon. 2021. Swarmbug: debugging configuration bugs in swarm robotics. In *European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 868–880.
- [27] Sayali Kate, Michael Chinn, Hongjun Choi, Xiangyu Zhang, and Sebastian G. Elbaum. 2021. PHYFRAME: type checking physical frames of reference for robotic systems. In *European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 45–56.
- [28] Sayali Kate, John-Paul Ore, Xiangyu Zhang, Sebastian G. Elbaum, and Zhaogui Xu. 2018. Phys: probabilistic physical unit assignment and inconsistency detection. In *European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 563–573.
- [29] Seulbae Kim and Taesoo Kim. 2022. RoboFuzz: fuzzing robotic systems over robot operating system (ROS) for finding correctness bugs. In *European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 447–458.
- [30] Sophia Kolak, Afsoon Afzal, Michael Hilton, Claire Le Goues, and Christopher S Timperley. 2020. It Takes a Village To Build a Robot: An Empirical Study of the ROS Ecosystem. In *International Conference on Software Maintenance and Evolution*. 430–440.
- [31] Sitar Kortik and Tejas Kumar Shastha. 2021. Formal Verification of ROS Based Systems Using a Linear Logic Theorem Prover. In *International Conference on Robotics and Automation*. 9368–9374.
- [32] William Landi. 1992. Undecidability of static analysis. *Letters Programming Language Systems* 1, 4 (dec 1992), 323–337.
- [33] Xiangke Liao, Shulin Zhou, Shanshan Li, Zhouyang Jia, Xiaodong Liu, and Haochen He. 2018. Do You Really Know How to Configure Your Software? Configuration Constraints in Source Code May Help. *Transactions on Reliability* 67, 3 (2018), 832–846.
- [34] Steven Macenski, Tully Foote, Brian Gerkey, Chris Lalancette, and William Woodall. 2022. Robot Operating System 2: Design, architecture, and uses in the wild. *Science Robotics* 7, 66 (2022), eabm6074.
- [35] Steve Macenski, Tom Moore, David V. Lu, Alexey Merzlyakov, and Michael Ferguson. 2023. From the desks of ROS maintainers: A survey of modern & capable mobile robotics algorithms in the robot operating system 2. *Robotics Autonomous Systems* 168 (2023), 104493.
- [36] Ratul Mahajan, David Wetherall, and Thomas E. Anderson. 2002. Understanding BGP misconfiguration. In *Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication*. 3–16.

- [37] T. Moore and D. Stouch. 2014. A Generalized Extended Kalman Filter Implementation for the Robot Operating System. In *International Conference on Intelligent Autonomous Systems*. 335–348.
- [38] John-Paul Ore, Sebastian G. Elbaum, and Carrick Detweiler. 2017. Dimensional inconsistencies in code and ROS messages: A study of 5.9M lines of code. In *International Conference on Intelligent Robots and Systems*. 712–718.
- [39] Morgan Quigley, Ken Conle, Brian Gerkey, Josh Faust, Tully Foote, Jeremy Leibs, Rob Wheeler, and Andrew Ng. 2009. ROS: an open-source Robot Operating System. *ICRA Workshop on Open Source Software* 3, 3.2 (01 2009), 1–6.
- [40] Akond Rahman, Shazibul Islam Shamim, Dibyendu Brinto Bose, and Rahul Pandita. 2023. Security Misconfigurations in Open Source Kubernetes Manifests: An Empirical Study. *Transactions on Software Engineering and Methodology* 32, 4 (2023), 99:1–99:36.
- [41] André Santos, Alcino Cunha, and Nuno Macedo. 2018. Property-based testing for the robot operating system. In *International Workshop on Automating TEST Case Design, Selection, and Evaluation*. 56–62.
- [42] André Santos, Alcino Cunha, and Nuno Macedo. 2021. The High-Assurance ROS Framework. In *International Workshop on Robotics Software Engineering*. 37–40.
- [43] André Santos, Alcino Cunha, Nuno Macedo, Rafael Arrais, and Filipe Neves dos Santos. 2017. Mining the usage patterns of ROS primitives. In *International Conference on Intelligent Robots and Systems*. 3855–3860.
- [44] Hannah Snyder. 2019. Literature review as a research methodology: An overview and guidelines. *Journal of Business Research* 104 (2019), 333–339.
- [45] Max Taylor, Johnathon Aurand, Feng Qin, Xiaorui Wang, Brandon Henry, and Xiangyu Zhang. 2023. SA4U: Practical Static Analysis for Unit Type Error Detection. In *International Conference on Automated Software Engineering*.
- [46] Fangchao Tian, Peng Liang, and Muhammad Ali Babar. 2019. How Developers Discuss Architecture Smells? An Exploratory Study on Stack Overflow. In *International Conference on Software Architecture*. 91–100.
- [47] Christopher Timperley and A Wasowski. 2019. 188 ROS bugs later: Where do we go from here? *ROSCON'19* (2019). https://roscon.ros.org/2019/talks/roscon2019_188_bugs_later.pdf
- [48] Christopher Steven Timperley, Tobias Dürschmid, Bradley R. Schmerl, David Garlan, and Claire Le Goues. 2022. ROSDiscover: Statically Detecting Run-Time Architecture Misconfigurations in Robotics Systems. In *International Conference on Software Architecture*. 112–123.
- [49] Christopher S. Timperley, Gijs van der Hoorn, André Santos, Harshvardhan Deshpande, and Andrzej Wasowski. 2024. ROBUST: 221 bugs in the Robot Operating System. *Empirical Software Engineering* 29, 3 (2024), 57.
- [50] Tetsuya Uchiumi, Shinji Kikuchi, and Yasuhide Matsumoto. 2012. Misconfiguration detection for cloud datacenters using decision tree analysis. In *Asia-Pacific Network Operations and Management Symposium*. 1–4.
- [51] Dinghua Wang, Shuqing Li, Guanping Xiao, Yepang Liu, and Yulei Sui. 2021. An exploratory study of autopilot software bugs in unmanned aerial vehicles. In *European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 20–31.
- [52] Dennis Leroy Wigand, Niels Dehio, and Sebastian Wrede. 2020. Model-Based Specification of Control Architectures for Compliant Interaction with the Environment. In *International Conference on Intelligent Robots and Systems*. 7241–7248.
- [53] Thomas Witte and Matthias Tichy. 2018. Checking consistency of robot software architectures in ROS. In *Workshop on Robotics Software Engineering*. 1–8.
- [54] Valentin Wüest, Vijay Kumar, and Giuseppe Loianno. 2019. Online Estimation of Geometric and Inertia Parameters for Multirotor Aerial Vehicles. In *International Conference on Robotics and Automation*. 1884–1890.
- [55] Tianyin Xu, Jiaqi Zhang, Peng Huang, Jing Zheng and† Tianwei Sheng, Ding Yuan, Yuanyuan Zhou, and Shankar Pasupathy. 2013. Do not blame users for misconfigurations. In *Symposium on Operating Systems Principles*. 244–259.
- [56] Tianyin Xu and Yuanyuan Zhou. 2015. Systems Approaches to Tackling Configuration Errors: A Survey. *Computing Surveys* 47, 4 (2015), 70:1–70:41.
- [57] Zuoning Yin, Xiao Ma, Jing Zheng, Yuanyuan Zhou, Lakshmi N. Bairavasundaram, and Shankar Pasupathy. 2011. An empirical study on configuration errors in commercial and open source systems. In *Symposium on Operating Systems Principles*. 159–172.
- [58] Zuoning Yin, Xiao Ma, Jing Zheng, Yuanyuan Zhou, Lakshmi N. Bairavasundaram, and Shankar Pasupathy. 2011. An empirical study on configuration errors in commercial and open source systems. In *Symposium on Operating Systems Principles*. 159–172.
- [59] Jiaqi Zhang, Lakshminarayanan Renganarayana, Xiaolan Zhang, Niyu Ge, Vasanth Bala, Tianyin Xu, and Yuanyuan Zhou. 2014. EnCore: exploiting system environment and correlation information for misconfiguration detection. In *Architectural Support for Programming Languages and Operating Systems*. 687–700.