

# MROS: Runtime Adaptation For Robot Control Architectures

Darko Bozhinoski, Mario Garzon Oviedo, Nadia Hammoudeh Garcia, Harshavardhan Deshpande, Gijs van der Hoorn, Jon Tjerngren, Andrzej Wąsowski and Carlos Hernandez Corbato

## ARTICLE HISTORY

Compiled August 30, 2021

## ABSTRACT

Known attempts to build autonomous robots rely on complex control architectures, usually implemented with the Robot Operating System (ROS). Runtime adaptation is needed in these systems, to cope with component failures and with contingencies arising from dynamic environments—otherwise these affect the reliability and quality of the mission execution. Existing proposals on how to build self-adaptive systems in robotics usually require a major re-design of the control architecture and rely on complex tools unfamiliar to the robotics community. Moreover they are hard to reuse across applications.

This paper presents MROS: a model-based framework for run-time adaptation of robot control architectures based on ROS. MROS uses a combination of domain specific languages to model architectural variants and capture mission quality concerns, and an ontology-based implementation of the MAPE-K and meta-control visions for runtime adaptation. The experiment results obtained applying MROS in two realistic ROS-based robotic demonstrators show the benefits of our approach in terms of the quality of the mission execution, and MROS's extensibility and re-usability across robotic applications.

## KEYWORDS

self-adaptive systems, models-at-runtime, variability, autonomous robots, control architecture, ontologies

## 1. Introduction

Robotic systems are designed by integrating and configuring individual robot capabilities in a governing control architecture, most frequently based on the Robot Operating System (ROS) [31]. However, static control architectures fall short when addressing context variability in open-ended, dynamic environments, where internal errors also compromise the quality and autonomy of mission execution. Self-adaptive Systems methods offer various solutions to deal with the context variability resulting from uncertainty and contingencies during execution [34]. However, existing approaches to apply these methods in robotic systems are marginal, because they require a major re-design of the control architecture, they use complex tools unfamiliar to the robotics community, and they are hard to reuse across applications. Therefore, there is a pressing demand for solutions to integrate self-adaptation in robot control architectures. This paper demonstrates how model-based self-adaptation can be easily integrated

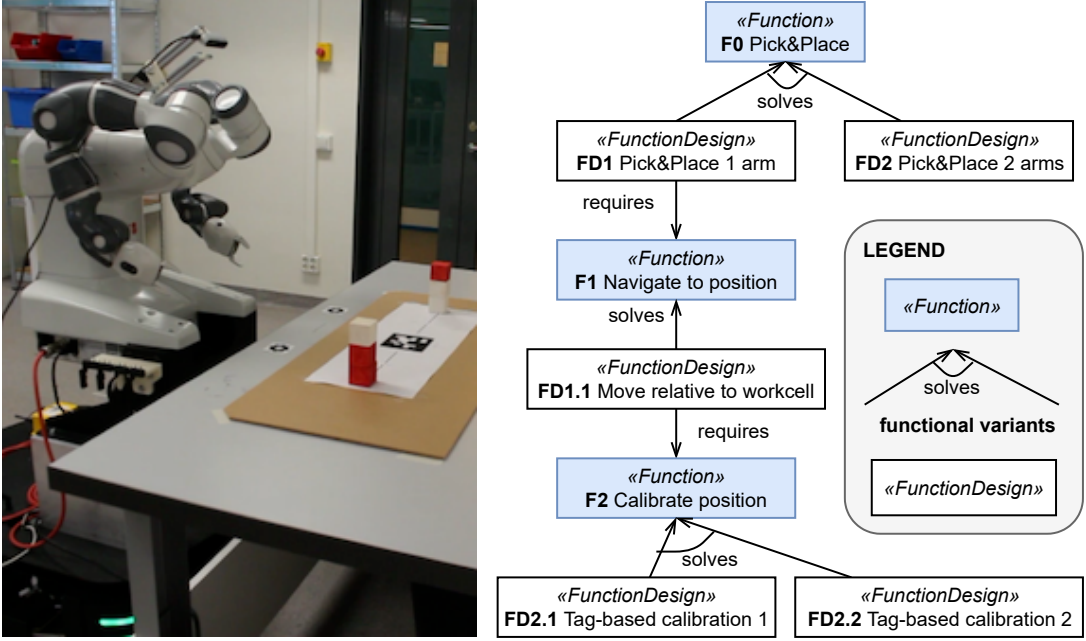


Figure 1.: Left: Demonstrator D1 based on an ABB YuMi manipulator. Right: the TOMASys model of its control architecture.

in robot control architectures to increase mission reliability and quality.

Robotics researchers have come up with sophisticated control architectures that are able to perform very well a specific mission [14, 33]. However, these architectures are mission and robot platform specific and are unable to address other missions and robots without undergoing major modifications [22].

An example of a robot performing a complex task is the mobile manipulator shown on Figure 1. The robot uses QR-tags to calibrate its position in the work cell, and the wooden blocks with which it builds a pyramid. The original ROS-based control architecture suffered from errors locating the QR-tag in varying light conditions, and from task failures whenever an arm bumps lightly into any obstacle (the internal safety function prevents any further motion of that arm). However, the system offers redundancy that can be exploited for self-adaptation. The challenge is to integrate self-adaptation into its architecture without major development effort, and maintaining extensibility to other tasks (e.g. navigation between the workstations), so that it can be reused for other robots.

We claim that using a model-based solution to integrate self-adaptation in robot control architectures addresses this challenge and results in increased mission quality and reliability at a lower cost. Modeling has been successfully applied in various domains to solve issues of architectural composition both statically (software product lines [30]) and at runtime (dynamic software product lines [11], models-at-runtime [5, 7]). Using models of the robot control architecture raises the abstraction level in the implementation of robot control systems, improving reuse from system to system, which is presently rare and often *ad hoc*. While work has been done in this direction [9], given the complexity of robotics systems, it is unclear how to proceed to maximize the benefits, and to avoid redesigning and re-implementing the adaptation layer for each robot from scratch.

In this work, we propose MROS: a model-based approach for self-adaptation of robot control architectures, contributing two key insights. First, the use of different Domain

Specific Languages to model different levels of abstraction allows for lean and extensible tools to design self-adaptation in robotics. Second, an explicit separation of concerns, between mission-specific logic and system management, and between general and mission-specific concerns, allows to reuse adaptation rules and artifacts across missions and robotic systems. For system design, MROS uses a combination of platform-specific Domain Specific Languages (DSLs) to model the robot architecture. These DSLs are close to the ROS ecosystem and effectively capture functional and task quality concerns of a robotic system for run-time adaptation. For run time adaptation, MROS realizes the *meta-control* vision [21] using a novel implementation of the MAPE-K loop based on ontological reasoning to drive the run-time adaptations.

A parallel, but equally important, goal of this paper is to advance experiment design for model-based self-adaptation, to push the community from anecdotal evidence towards data. From methodological perspective, we propose to evaluate self-adaptation frameworks at two levels: (i) how well, and at what computational cost, a given framework improves mission performance and reliability, and (ii) how reusable (vs idiosyncratic) a framework is; to how large class of systems it applies, and how much specialized development is needed. We evaluate the former in a quantitative experiment, and the latter in a qualitative assessment based on engineering experience from the project, and its architectural qualities.

Our main contributions include:

- *A reusable implementation of MROS meta-controller*, ready to integrate self-adaptation in ROS systems.
- *A set of design tools and languages* to model and deploy the MROS solution in ROS systems.
- *An evaluation of MROS in two case studies (mobile manipulator, factory floor navigation)* exceeding prior methodological standards. The experiments show performance and robustness improvements, and quantify the computational cost of the general architecture.
- *An analysis of benefits of model-driven methods for this problem.*

Design methodologies for autonomy and self-adaptation remain an open discussion topic in the robotics community. With this paper, developed in a joint effort of robotics and modeling researchers, we hope to get more members in the modeling community interested in the problem. We also contribute a validated approach and artifacts to the discussion of systematic design methods and reusable architectural components for robotics. We believe that the integration of Meta-control with ROS may bring models-at-runtime and system modeling methods to broader groups of practitioners in robotics.

The rest of the paper is organized as follows. Section 2 gives background information for model-based development and ROS, and self-adaptation using meta-control. Section 4 presents the two robotic demonstrators developed to validate our solution. Section 5 defines our experimental setup and results regarding the run time feasibility of MROS on both demonstrators. Section 6 discusses how MROS supports system extensibility and reusability, Section 7 discusses related work and, Section 8 presents the conclusions and future research direction.

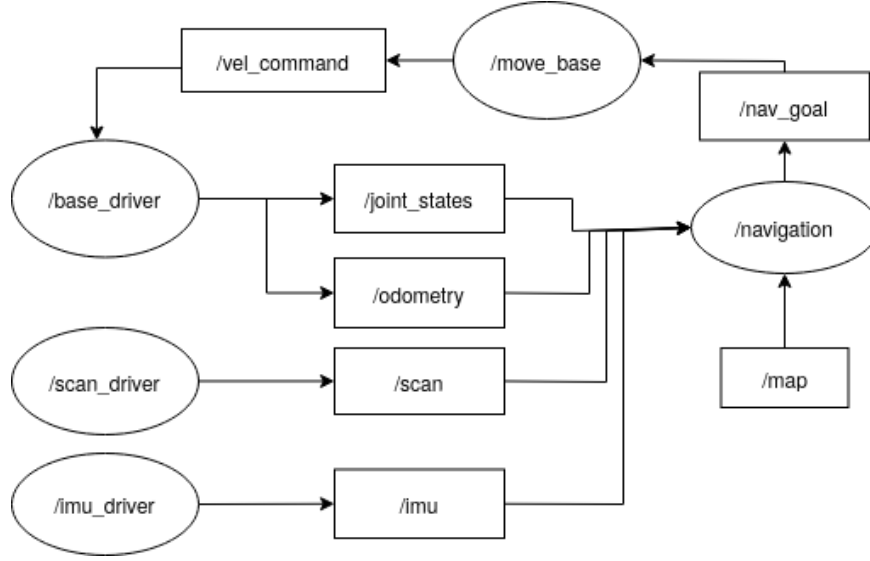


Figure 2.: ROS graph for a simple navigation application, extracted from a system at runtime by standard ROS tools. Ellipses denote nodes, boxes mark topics.

## 2. Background

### 2.1. ROS and the RosModel tooling

The Robot Operating System is a component-based, an open-source platform considered the *de facto* standard for the development of robotics systems, and the platform chosen to develop this work. A key to the success of ROS lies in imposing few architectural constraints and in offering many tailored components (robotics-specific functional components and hardware drivers). A typical ROS system consists of many small and mostly independent distributed programs called *nodes*. Nodes communicate via message passing, using a publish-subscribe mechanism, or service calls. Figure 2 shows an example of a runtime communication graph of one of the cases of this paper.

To facilitate model-based development with ROS, Hammoudeh Garcia and colleagues developed the *RosModel tooling* framework which contains a set of languages to formalize relevant properties of a ROS system architecture, in a manner compatible with model-based ecosystem [18]. Our paper uses two of these languages:

- *RosModel* model: specifies the communication interfaces of each node (the ports offered by a node) and the deployment information (the ROS package of a node and the packages it depends on)
- *RosSystem* model: reflects the composition of the instantiated nodes, how they are wired and configured.

Figure 3 shows the (static) *RosSystem* for the (runtime) example of Fig. 2, in the *RosModel tooling* graphical editor. Nodes are defined as components that contain a set of interfaces, with properties *name* and a reference to the definition of the instantiated interface. The format also captures connections between the interfaces. Associated tools validate all the connections (i.e. the match of the type of the message sent by the publisher and expected by the subscriber) and automatically generate a *launch file*, an XML script used in ROS to start, configure, and connect all the nodes that constitute the robotics system. The tooling offers static analyzers and runtime introspectors that

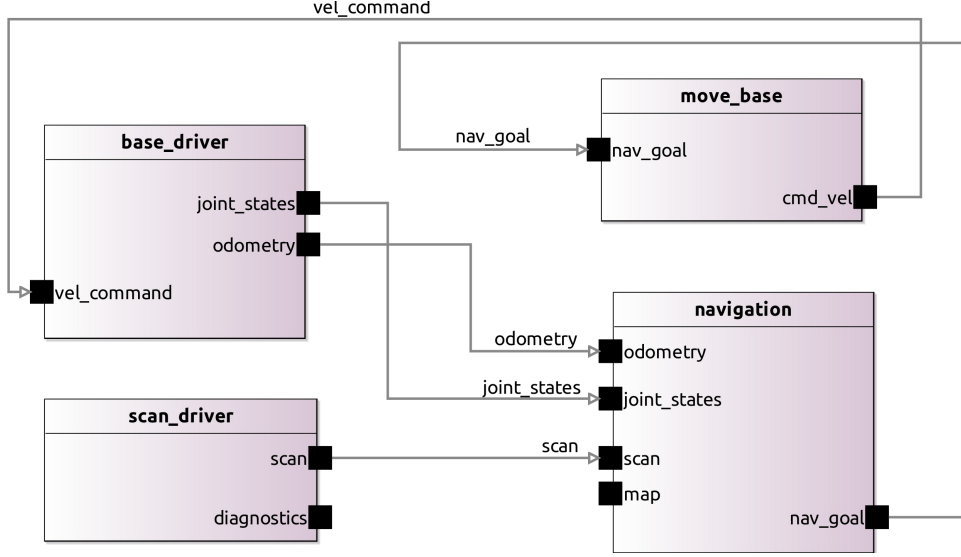


Figure 3.: Example RosSystem model for the system in Fig. 2

allow scaffolding the models from an existing software implementation.

## 2.2. Meta-control and the TOMASys Meta-model

Meta-control [3, 21] is a reference architecture for self-adaptive autonomous control systems inspired by MAPE-K (*Monitor-Analyze-Plan-Execute over a Knowledge base*, [26]) and supervisory control [8]. At its center is the *Teleological and Ontological Model for Autonomous Systems* (TOMASys) that allows to model the architecture of component-based systems, including architectural variations at both design-time and run-time [21]. Architectural models provide the appropriate level of abstraction to manage adaptation at runtime [17]. TOMASys captures the relation between the system’s requirements and their allocation to the system components through the concept of functional and physical architectures from systems engineering [16, 20]. In the Meta-control vision, a Meta-controller component implements runtime adaptation using the TOMASys model of the system and automatic reasoning for the evaluation, selection, and deployment of architectural variants.

In TOMASys, a *function* represents a capability that has been designed in the system, for example “Navigation” in a mobile robot. A *function design* is a possible realization of a function, an architectural variant able to deliver it. Internally, a function design prescribes a certain structure that delivers the function, i.e. it maps functionally to system structure, through specifications of components and their interconnection. TOMASys allows to model functional decomposition by designating *required* functions for a function design, capturing the dependencies between lower- and higher-level functionality. At runtime, the functional requirements are represented as instances called *objectives*, which are solved by instances of the *function designs*, the active function designs called *function groundings*.

A TOMASys model example is shown in Fig. 1, using UML notation with stereotypes. A high-level function F0 (Pick&Place) represents the capability of the ABB mobile manipulator to pick and place small objects. The model defines two function designs named FD1 and FD2 that *solve* this function. FD2 represents an architectural variant

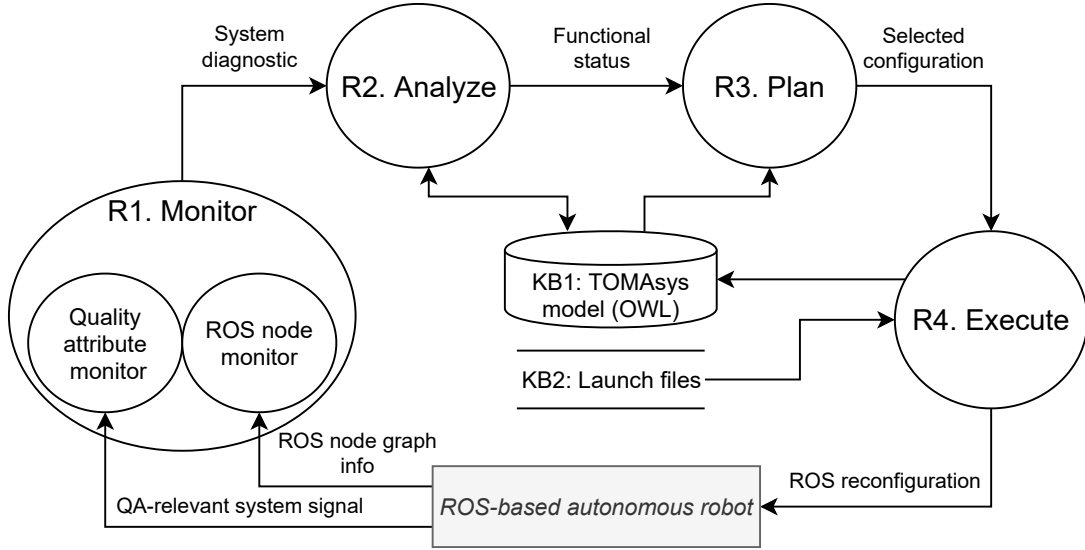


Figure 4.: Data flow diagram for the run-time operation of the MAPE-K loop implemented by the MROS Meta-controller.

that solves Pick&Place using both arms of the robot (which jointly can reach the entire workspace), and it is implemented as a monolithic program in the ABB programming language. FD1 represents a variant that solves the function using only one arm and the ability of the mobile base to move and approach objects if they are out of reach of the arm. It is implemented as a set of ROS nodes specifically parametrized to solve F0, and it depends on function F1 (Move position) implemented by function design FD1.1 (Move relative to work cell). FD1.1 in turn requires the function F2 (Calibrate position) delivered by two designs, FD2.1 and FD2.2, which detect the QR tag to calibrate the position of the robot in the work cell. FD2.1 and FD2.2 use different camera settings configured for different light conditions.

### 3. MROS Model-Based Meta-control

MROS is a model-based solution for meta-control that integrates easily in a typical ROS development workflow. It provides a general reusable and extensible meta-control node that extends any ROS system with runtime self-adaptation. Reuse is maximized by separating the general reasoning principles, the adaptation related to system configuration, and error handling from task-related adaptations. Extensibility is achieved by using ontological reasoning to implement the decision-making in the meta-controller. At design time, MROS greatly simplifies the development effort by extending the *RosModel tooling* with a series of plugins that allow ROS developers to model architectural variants and their quality attributes, and to automate the generation and deployment of the meta-controller.

#### 3.1. Runtime Adaptation

MROS implements a MAPE-K loop using a runtime model conforming to the TOMAsys meta-model serving as a knowledge base (KB1). The realization of MAPE-K with ontological reasoning is a core contribution of this work. It allows to separate: (i) the

TOMASys model of the robotic system, coded with the Ontology Web Language (OWL), (ii) the adaptation rules, coded with the Semantic Web Rule Language (SWRL), and (iii) the self-adaptation process—a ROS node `mros1_reasoner`. Figure 4 summarizes the main data processing steps in the MROS MAPE-K loop. We discuss them in order below.

### *Step R1. Monitor*

Two monitoring processes track the status of active ROS nodes and the degree of fulfillment of non-functional requirements using quality attributes as a proxy. The goal is to (i) detect and identify parameters misconfigured by developers in ways that lead to sub-optimal robot behavior, (ii) identify erroneous node behavior, and (iii) observe violation of system constraints (i.e. level of quality attributes). The monitors operate at a fixed frequency and produce diagnostics using the standard ROS diagnostic mechanism [1]. This means they easily integrate into an existing system, and might use or provide services already needed in the system, regardless of the presence of a meta-controller.

The *ROS node monitors* leverage *RosModel* tools. The existing `ros_graph_parser` node is used to create a *RosSystem* model of the running ROS-based system. A newly developed `rosgraph_monitor` compares that model with the desired *RosSystem* model of base application (cf. Table 1). It publishes a diagnostic error message in case a difference is detected. The *Quality attribute monitors* are implemented by a separate *observer* node per each quality attribute. They report normalized values for the quality attributes: zero means that the quality attribute is at the required level, one denotes maximum deviation from the required level. The skeleton of these nodes is automatically created from the *RosSystem* model, for the attributes and the system signals (ROS topics) specified by the system architect. Developers need to implement the logic in the observers to obtain the quality attribute level from the given signals. For the evaluation cases presented below, we implemented observers for two attributes: safety (risk of collision with humans, using proximity sensors) and energy (using the runtime battery consumption and a battery model).

### *Step R2. Analyze*

The system diagnostic data produced in Step R1 is used to infer the status of the functional architecture of the system. Any new facts, for instance a new component status when a node has reported a fatal error, or a change in the level of a quality attribute, are asserted in the TOMASys model (KB1). Then, ontological reasoning and application-independent rules (declared in SWRL, as proposed in [23]) are applied to the model to infer the current status of the functional architecture, including the status of the TOMASys objectives and the applicability of the alternative function designs. Finally, the status is updated in KB1.

### *Step R3. Plan*

If the meta-controller discovers in Step R2 that any objectives are violated, an architectural adaptation is proposed by searching for alternative function designs in the TOMASys model (KB1). If several available designs meet the required quality attribute levels, the one maximizing an application-specific utility function is selected. The `mros1_reasoner`, used in R2–R3, is implemented with the Owlready2 library [27], the off-the-shelf reasoner Pellet to infer the functional status, and a custom logic for



<i>Step</i>	<i>Input artifact</i>	<i>Design-time activity</i>	<i>Output artifact</i>
1	Application requirements	Application modeling (manual + reverse engineering tools)	<i>RosSystem</i> model of base application
2	<i>RosSystem</i> model of base application	Modeling architectural variants	<i>RosSystem</i> models
3	<i>RosSystem</i> model of base application	Generate meta-controller deployment configuration	(YAML)
4	<i>RosSystem</i> models	Generate deployment configuration for variants ( <b>automatic</b> )	ROS Launch files
5	<i>RosSystem</i> models	Generate observers of the quality attributes ( <b>automatic</b> )	Observers boilerplate code
6	<i>RosSystem</i> models	Generate runtime model for the meta-controller ( <b>automatic</b> )	<i>TOMASys model (OWL)</i>

Table 1.: Key steps in development of MROS meta-controller.

<pre> 1 RosSystem { Name 'system_a' RosComponents ( 2   ComponentInterface { name move_base 3   RosParameters{ 4     RosParameter 'max_vel_x' { value 0.5 }, 5     RosParameter 'max_vel_y' { value 0.5 }, 6     RosParameter 'inflation_radius' { value 0.5 }, 7     RosParameter 'observation_sources' { value scan }}} 8 Parameters { 9   Parameter { name 'qa_safety' type Double value 0.41 }, 10  Parameter { name 'qa_energy' type Double value 0.48 }}} </pre>	<pre> 1 RosSystem { Name 'system_b' RosComponents ( 2   ComponentInterface { name move_base 3   RosParameters{ 4     RosParameter 'max_vel_x' { value 0.3 }, 5     RosParameter 'max_vel_y' { value 0.3 }, 6     RosParameter 'inflation_radius' { value 0.8 }, 7     RosParameter 'observation_sources' { value point_cloud }}} 8 Parameters { 9   Parameter { name 'qa_safety' type Double value 0.71 }, 10  Parameter { name 'qa_energy' type Double value 0.33 }}} </pre>
--	---

Figure 5.: Fragments of *RosSystem* models for two variants of a navigation system, specifying configurations of the `move_base` node and expected values for the quality attributes. Left: base variant using a laser scanner, Right: alternative using a camera.

the search and selection of system configuration in the Plan step.

#### *Step R4. Execute*

Enforce the adaptation selected in Step R3—stop nodes no longer needed, start new nodes needed, reconfigure the nodes remaining. It is implemented by a general `rosgraph_manipulator`, which follows an adaptation tactic suitable for ROS1, re-deploying the nodes by using launch files created by the *RosModel* tools.

### *3.2. Design Time Activities*

The application developer uses the MROS extension of the *RosModel* tools to describe the architecture of the domain application by specifying the set of possible configurations of components and their quality attributes. Table 1 summarizes the steps involved.

#### *Step 1. Model the application*

We first create a *RosSystem* model of the base application, starting with individual nodes. Two tools for automatic model generation from source code are available from the *RosModel* project [19]: a built-in plugin that generates the models for ROS packages



stored locally, and a web interface, available at <http://ros-model.seronet-project.de/>, for ROS packages publicly available. Once the *ROSModel* models describing single nodes are obtained, the developer can compose them, using either the graphical or the textual editor, to define the *RosSystem* model of the entire base application. Optionally, if MROS is used for an existing ROS application, the *RosModel* tools support automatic analysis of the complete existing system. This step results in a *RosSystem* model (a `.rossystem` file) which is validated by our tools and used to automatically generate executable artifacts (launch files to start the ROS system).

#### *Step 2. Model architectural variants*

Next, the developer models the adaptation possibilities—the architectural variants. Most adaptations can be sufficiently captured by parameter variations, but larger architectural variations are also possible like switching to a different hardware module). Figure 5 shows the representation of two different configurations of the same system. In this case, we have two different configurations of the node `move_base`. This node is in charge of interpreting the actions from the navigation module and translating them into commands to the base. Notice, how the maximum speed, inflation radius, and the type of sensor data used to adapt to the environment varies. The base variant `system_a` uses the output of a laser proximity scanner, while the alternative variant `system_b` uses the point cloud obtained from a 3D camera. Additionally global parameters are added to define the value of the quality attributes (abbreviated `qa`), in the case of this example for safety and energy.

#### *Step 3. Generate meta-controller deployment configuration*

We provide a *RosModel* plugin that automatically generates the application-specific configuration for the `mros1_reasoner` node to be integrated on top of the base application.

#### *Step 4. Generate deployment configurations for variants*

We automatically generate executable artifacts (KB2: ROS launch files in Figure 4), which can be called to execute runtime adaptations when a new configuration is selected by the reasoner (R4 in Figure 4).

#### *Step 5. Generate observers of quality attributes*

Another *RosModel* wizard is used to generate the scripts to observe the system. The GUI allows to select the quality attributes to be monitored. MROS generates a description of the ROS node as a *RosModel* and the skeleton of the Python code for monitoring.

#### *Step 6. Generate runtime model for the meta-controller*

Finally, MROS provides a script that given the *RosSystem* models of different adaptation strategies, automatically generates the TOMASys model that feeds the meta-controller at runtime.

In summary, the provided tools offer a holistic model-driven solution for the adaptation problem, exploiting models and automation, with a high degree of reuse both at design time *and* at runtime. Since many artifacts are needed, generating all of them offers not only productivity gains: the tools ensure consistency of the input artifacts and guarantee consistency of all the generated artifacts.

## 4. MROS Demonstrators

We use two different robotic systems to validate MROS and the developed artifacts.<sup>1</sup> Both are based on realistic applications for autonomous robots, concretely manipulation and navigation [2]. Mobile manipulation robots are designed to operate in semi-structured environments, such as a warehouse, a factory floor, and a healthcare facility. These environments change dynamically (e.g. the floor plan can change from day to day, some areas may be shared with humans and with other robots, etc). This context variability demands varying configurations of the robot control architecture and appropriate methods to support successful mission completion.

### 4.1. *Dual-Arm Manipulator (D1)*

In D1, the robot is a YuMi dual-arm manipulator standing on a Clearpath Ridgeback mobile platform and equipped with an Intel Realsense RGBD camera (Fig. 1). Its task is to build a pyramid with wooden blocks at the work station. The camera and QR-tags are used to calibrate the robot’s position in the work cell, and to locate the bricks and the target position of the pyramid. This information and prior knowledge about the blocks sizes, is used to compute the required manipulation motions. To improve robustness, we use a YuMi dual-arm manipulator. Its two arms provide architectural variability: the task can be performed using both or one arm. Each arm has limited range of operation, so when working with one arm, the robot uses the mobility of the Ridgeback platform to reach all the bricks.

### 4.2. *Mobile Robot Navigating in a Factory (D2)*

The mobile robot demonstrator (Fig. 6a) is concerned with navigation on a factory floor, where the robot moves between workstations to perform different manipulation tasks. The robot is the Ridgeback platform, equipped with two Hokuyo lasers, odometry and an inertial measurement unit, and is augmented for this demonstrator with an indoors localization system. The baseline control architecture for navigation is based on the ROS1 navigation stack [29], which uses off-the-shelf control and AI algorithms for planning and reactive obstacle avoidance. The robot navigates to a target location using odometry and sensor data and sending velocity twist commands to the Ridgeback.

The context variability includes: (i) new obstacles on the factory floor affecting the quality attribute values attained in the default configuration of the navigation stack (ii) variations in power consumption, resulting in changing battery depletion.

## 5. Feasibility of Runtime Adaptation

We analyze the implementation of MROS with respect to its management of component failures and runtime quality attribute levels, by using the two demonstrators and asking the following research questions:

**RQ1.** *How does MROS improve system robustness in the presence of component failures?*

---

<sup>1</sup>The source code of the meta-controller and the experiment scripts is released under the open source license and can be found in the following repos: [https://github.com/rosin-project/metacontrol\\_sim](https://github.com/rosin-project/metacontrol_sim) (Metacontrol) and [https://github.com/rosin-project/metacontrol\\_experiments](https://github.com/rosin-project/metacontrol_experiments) (experimental scripts)

**RQ2.** *How does MROS affect system and mission performance in the presence of component failures and system quality concerns?*

With **RQ1**, we demonstrate *qualitatively* how MROS addresses component failures on a manipulation mission. **RQ2** investigates *quantitatively* how the MROS reconfiguration affects the task and system performance of a robot by addressing functional and non-functional concerns in a navigation mission.

### 5.1. Manipulation Mission (RQ1)

#### *Experiment setup*

To address **RQ1**, we conducted a study on D1 dual-arm manipulator (Section 4.1) in two scenarios:

- (1) System does not detect the workstation tag.
- (2) System cannot use both arms.

The first scenario aims to demonstrate the MROS ability to recover a function that requires multiple components (increased mission robustness). The second scenario demonstrates the MROS controller’s ability to recover from an error by realising an alternative set of functions to perform the task goal with degraded performance (increased system robustness).

#### *Results and discussion*

In the first scenario, the QR-tag for calibration is not detected, due to the light conditions. The meta-controller uses the TOMASys model in Fig. 1 to switch between the alternative configurations of the camera parameters FD2.1 and FD2.2 to detect the tag.

In the second scenario, we inject a failure by blocking one of the robot’s arms. MROS identifies the component error and uses the functional architecture model to find an alternative variant. It adapts the architecture to a configuration that uses the arm not blocked. Since the task cannot be completed with the remaining arm (the wooden blocks are out of range), it also switches the control algorithm to change positions of the mobile platform.

We conclude that the MROS meta-control augments the reliability of the overall system, delivering the expected behavior in real-time in the context of component failures.

### 5.2. Navigation Mission (RQ2)

#### *Experiment setup*

To address **RQ2**, we systematically compared a system with the *MROS* meta-controller to a benchmark system (*Baseline*) for mission success, component failure, and quality concerns in the navigation scenario of Section 4.2.<sup>2</sup> The robot is placed at a starting location and it navigates to a goal location. An example of the route during one run is shown in Fig. 6b; the starting position is at the bottom-left of the map and the goal is at the top-right. We considered six pairs of initial and goal positions.

---

<sup>2</sup>The data from this experiment can be found in <https://doi.org/10.5281/zenodo.5340886> released under the Creative Commons Attribution 4.0 International license.



(a) Simulation of the mobile manipulator in a factory.



(b) An example route of the mobile manipulator navigating.

Figure 6.: Mobile Manipulator in a factory

### *Baseline*

We developed 30 variants of a control architecture, with different quality expected performing the navigation mission. Each of these variants constitutes the *Baseline* system for the corresponding experiment run. The variants are obtained by selecting between a laser or a 3D camera as sensor input for obstacle avoidance, and different values of three parameters that govern the behavior of the `move_base` node that generates the motion commands: `inflation_radius` (minimal distance allowed to the closest obstacle), `max_vel` (maximum velocity allowed for the robot), `controller_frequency` (the frequency of the control loop).

### *MROS-D2*

In *MROS-D2*, the meta-controller runs on top of one of the architectural variants, performing run-time analysis on component failures and quality violations, and making decisions on when and how to switch to a different navigation configuration. *MROS-D2* calculates safety at run-time by measuring the braking distance of the robot before hitting an obstacle `d_break` [13] and the distance from the current location of the robot to the closest obstacle in its direction  $P$ . We define runtime safety to be 1 if the robot is at least `d_break` distance from the closest obstacle. As the robot approaches the closest obstacle, safety at run-time is computed as  $P/d\_break$ . Furthermore, *MROS-D2* computes the energy at run-time based on an energy model that simulates the power consumption of a robot. The energy model computes the instantaneous `power_load` of the robot using current controller frequency, velocity, acceleration, etc. We specify 0.6 as a required threshold for safety and 0.62 for energy. If the robot breaches these thresholds, *MROS-D2* initiates self-adaptation.

The independent variables in our experiment are: (i) the architectural variant initially deployed for the control architecture, and (ii) the contingencies during the mission: For each run, we compare the MROS and *Baseline* systems in terms of violation of system qualities thresholds, component failures handling, and mission performance. A run consists of an initial *system configuration*  $C$  and a *perturbation*  $P$ . We selected nine initial configurations (out of 30) that show the best-expected performance in a

set of preliminary tests. A perturbation models a sequence of contingencies within a mission. Each perturbation  $P$  is represented by a tuple  $\langle iv_1, iv_2, iv_3 \rangle$ , where each element represents a value for each contingency type:

- (1) *Cluttered environment*: the environment contains a set of new obstacles unknown to the robot (not in the map used by the navigation) in randomized positions. We quantify the level of clutter depending on the number of new obstacles as *low* (7), *medium* (14) and *high* (20).
- (2) *Unexpected increase in power consumption*: This could be the result of a sticky surface, a slope, some problem with the wheels, or activation of a new sensor that consumes more power. We quantify this contingency in three levels: 20%, 40%, and 60% increase in power consumption.
- (3) *Component failure*: We inject a failure in the laser sensor by killing its driver ROS node at run-time.


To estimate the effects of MROS on adaptation and mission performance we measure the following dependent variables on both systems:

- *The percentage of time the mission is above the threshold for safety risk and energy consumption.*
- *Mission success.* We defined a demanding metric for mission success that combines reaching the goal with performance and safety requirements. A mission is considered successful if (i) the robot reaches its goal, (ii) the average safety risk level for a run is  $< 0.4$ , (iii) the safety risk is below its threshold for more than 95% of the mission time; (iv) the energy-consumption is below its threshold level for more than 90% of the time.
- *Time to complete the mission.* We measure the time the robot needs to complete a single run.


### Results and Discussion

We collected data points for 6727 runs of both systems. Figures 7a and 7b illustrate the management of safety and energy at run-time during a single run for the Baseline system and for MROS-D2. The red line represents the threshold for each quality, while the highlighted area is a time period when contingencies are introduced. For the *Baseline* system (Fig. 7a), the total run-time is 39 seconds, the safety quality attribute value has an average of 0.47 and is above the defined limit for 4 seconds (10% of the run-time), and the QA Energy has an average of 0.38 and is above the required limit for 3 seconds (8% of the run-time). For the *MROS* system (Fig. 7b), the total run-time is 56 seconds. In this case, the average of the safety quality attribute value is 0.37 and it is above the defined limit for only 1 second (2% of the run-time). The average for QA Energy is 0.325 and it is above the required limit for 3 seconds (5% of the run-time). Moreover, when the MROS system is used, two reconfigurations occur, which can be clearly seen in Figure 7b where, in the highlighted areas, both QA Energy and Safety values drop down immediately after the limit is crossed and then they restart with a different trend. Each reconfiguration takes about 4.3 seconds, contributing to the MROS run being longer than the *Baseline*.

Table 2 shows the average percentage the *Baseline* and *MROS-D2* are above the safety threshold during the mission. On average *MROS-D2* is 2.6 times more effective in reducing the safety violation from 2.5% to 0.96% on average in a single mission. *MROS-D2* always outperforms the *Baseline*, independently of the level of clutter (number of obstacles). Furthermore, *MROS-D2* manages to keep the safety violation below 1% for



(a) Baseline without MROS meta-control.



(b) The system with MROS Meta-control.

Figure 7.: Safety (top) and Energy (bottom) quality attribute values for a single run. The red line indicates the maximum allowed value in each case. The highlighted areas mark the contingencies where a reconfiguration occurs.

<i>How cluttered the environment is</i>	<i>MROS</i>	<i>Baseline</i>
no clutterness contingency	0.99%	2.60%
low clutterness	0.98%	2.84%
medium clutterness	0.88%	1.73%
high clutterness	1.00%	2.70%
<b>Total</b>	<b>0.96%</b>	<b>2.50%</b>

The

Table 2.: Average time a mission is above the safety threshold.

<i>Power consumption</i>	<i>MROS</i>	<i>Baseline</i>
no power contingency	1.26%	2.63%
power increase 20%	1.16%	2.22%
power increase 40%	2.47%	2.63%
power increase 60%	2.63%	4.57%
<b>Total</b>	<b>1.86%</b>	<b>2.98%</b>

The

Table 3.: Average time a mission is above the energy threshold.

all contingency levels related to a cluttered environment. Considering these contingency levels, the *MROS-D2* meta-controller re-configures the system once every 2-3 runs.

Table 3 shows the average percentage the *Baseline* and *MROS-D2* are above the energy threshold during the mission. Again, *MROS-D2* outperforms the *Baseline* for all contingency cases related to the increase in power consumption. The biggest difference is the 60% energy increase contingency, where the result shows that MROS reduces the average violation of the energy constraint from 4.57% to 2.63% during the mission.

Table 4 shows that *MROS-D2* outperforms the *Baseline* for all contingency types. *MROS-D2* improves the mission success by 13.3% on average. In this calculation, we decided not to include the data for component failure, because the *Baseline* will always

<i>Contingency Type</i>	<i>MROS</i>	<i>Baseline</i>
no clutter contingency	78.16%	67.03%
low clutterness	77.92%	62.03%
medium clutterness	79.13%	70.16%
high clutterness	77.77%	62.06%
no power contingency	81.63%	59.51%
power increase 20%	79.38%	68.48%
power increase 40%	76.47%	71.01%
power increase 60%	76.35%	60.18%
component failure	73.00%	0%
<b>Total</b>	<b>78.50%</b>	<b>65.20%</b>

Table 4.: Mission success rates under various contingencies.

fail the mission when a laser failure is introduced (so the 13.3% reported is biased against *MROS-D2*). In contrast, *MROS-D2* effectively adapts when laser failure is introduced, completing the mission in 73% of the mission runs. Beyond the component failure contingency, the biggest difference in mission success can be seen in the case when there is not any increase in power consumption. In this case, the *Baseline* has 59.51%, while *MROS-D2* has 81.63% mission success.

Finally, MROS increases the mission completion time by 7.5, as a result a result of the Meta-controller choosing a slower configuration to satisfy the energy and safety quality concerns. As expected in MROS-D2, adaptations were triggered more often for higher contingency levels in terms of clutter and energy peaks.

It is important to note that one can engineer a more resilient robot, with more redundant hardware and control strategies, and then the reported numbers for *MROS* will improve further, as they are indeed a function of a more adaptive strategy, which the baseline system does not have. The experiment demonstrates however that the model-based MROS meta-controller is capable to exploit the hardware and software redundancy to increase robustness. Thus we conclude:

Model-based intelligent reconfiguration does improve mission and system performance for robot systems, effectively managing component failures and quality concerns.

## 6. Qualitative Analysis of MROS Design

MROS contributes two tools: the *RosModel* MROS plugins for ROS developers at design time, and the Meta-controller ROS packages to deploy it in a ROS system. In order to understand the generality of our contribution, we discuss the re-usability and extensibility of these artifacts.



### 6.1. Reusability

We prototyped the MAPE-K loop of the meta-controller for the D1 dual-arm manipulator demo without generative programming. Following a common practice in the field, we then generalized it for D2 using model-driven development.

For D1, we designed the OWL implementation of the TOMASys meta-model (`tomasys.owl`) to support the specification of MROS models, and the ontological reasoning for Steps R2 and R3 (`mros1_reasoner.py`). The TOMASys model was manually developed. We first analyzed the existing software for the system following the *ISE&PPOOA* MBSE methodology [16], obtaining the conceptual representation of its functional architecture and possible variants (Fig. 1). Then we modeled it in OWL, using `tomasys.owl` and the Protege editor. For D1, we used application specific solutions for steps R1 and R4 at runtime.

To exploit the commonality in the domain, we consequently settled to use the model-driven methods and the *RosModel* bridge tools. The resulting MROS tools offer:

- *Modeling languages and tools for ROS architectures:* (a) linking ROS and TOMASys via *RosSystem* (b) with consistency checking and syntax feedback
- *Extracting models from code*, allowing to add meta-control to existing projects and using new nodes with MROS.
- *Automatic generation of the ontological knowledge base* (otherwise a tedious and low-level task).
- *A reusable monitor of node status* `rosgraph_monitor` (used in Step R1)
- *A generator of boilerplate code for observers* (step R1).
- *A reusable meta-controller*, including tactics to detect and address violations of quality attributes and contingencies (steps R2 and R3)
- *A reusable navigation ontology* based on `tomasys.owl` including minimal domain-specific model of safety, energy and performance quality attributes in navigation, along with the associated SWRL adaptation rules (steps R2, R3).
- *Automatic generation of executable reconfiguration actions for the system architecture* (launch files) used by the `rosgraph_manipulator` to reconfigure the system (Step R4)

For demonstrator D2, we used this infrastructure to model the architectural variants with 30 different *RosSystem* models of the possible configurations of the navigation stack.

It is a well known fact that effective deployments of model-driven engineering require a very good domain implementation. It is somewhat surprising thus that often model-driven technology is demonstrated in isolation, or on small examples and use cases. The Robot Operating System provides an excellent platform, a rich and highly customizable implementation of the robotics domain. ROS with model-driven engineering is a “match made in heaven.” Using ROS was an enabling factor for this research: it opened access to standard formats, architectures, and interfaces. We benefited from the rich tooling and experience ecosystem, which allowed us to efficiently execute rich and realistic experiments, rarely seen in model-based self-adaptation research. Furthermore, using generative model-driven solutions with ROS meant that the meta-controller technology is readily accessible to unprecedented number engineers and researchers working on self-adaptation in robotics and software engineering using ROS.

## 6.2. System Extensibility

We analyze how the MROS framework supports system extensibility in the context of a mobile robot navigating in a factory scenario with respect to which aspects should one extend in the application: (i) to capture new components added to the system; (ii) to capture new quality attributes to be monitored.

### *Adding new components*

A 3D camera component was added so the mobile robot would be able to continue the task even if the laser fails. To add a new component we only needed to add the new configurations in the *RosSystem* models. We extended the number of system variants by adding new configurations that contain the camera component. We used the *RosModel* tools to model 3 new configurations which corresponds to 30 different variants in total. For monitoring, we did not need to add new observer models, because the current tooling allows component failures to be captured. For analysis and planning, we used the model-to-model transformation to automatically generate the updated knowledge base. Finally, the *RosModel* tools automatically generate the new reconfiguration actions.

### *Adding new quality attributes*

We extended the application scenario D2 with energy-efficiency. For monitoring, we developed an observer for energy-efficiency that calculates energy at run-time. Energy level was computed using an energy model that simulates the power consumption of a robot using controller frequency, velocity, acceleration, etc. For analysis (R2) and planning (R3), we extended the knowledge base with new information about the energy-efficiency of each system configuration. Manual work is unavoidable, but the artifacts are modular and reusable across robotic systems with the same quality attributes. Again as a final step, the *RosModel* tools automatically generate the reconfiguration actions.

In conclusion, in both cases the cost of development of the new features of the robot, greatly outweighs the cost if incorporating it into the meta-control system, which is mostly automatic, barring modeling the new configurations (easy) and new quality attributes (reusable).

## 7. Related Work

Dynamic architectures and self-adaptation have been an active topic of research in model-driven software engineering for a decade. The so called *models-at-runtime* paradigm [6, 7] has been exploring the use of abstractions to reduce complexity and cope with increasingly complex runtime adaptation problems. One group of prior works related to the use of models in robot software is based on variability management and feature models that can represent the functionalities provided by a software system symbolically [15, 25]. The previous efforts have focused on the solution space of architecture design, that is the possible configurations of components. The discrete variants in the architecture are modeled separately from properties and quality attributes (the problem space). However, formal semantics, a common meta-model, for automated reasoning about the relation between both the solution and the problem space is missing. We address this shortcoming by extending the TOMASys meta-model to represent quality attributes and implement it as an ontology for automated reasoning.

Non-functional requirements, such as performance, reliability or safety, have been

considered in the dynamic reconfiguration of robot software architecture [28]. In a newer work, Brugali et al. have integrated feature models, standard modelling languages (UML-MARTE) with queuing network models to support reconfiguration in order to maintain an adequate level of performance [9]. These solutions require intensive modelling efforts, MDE skills and completely new toolsets. In contrast, we chose a less invasive path, in that we exploit existing codebase and practices in the ROS ecosystem.

Iftikhar and Weyns [24] propose Active FOrmal Models for Self-adaptation (ActivFORMS). The adaptation goals that are verified offline are guaranteed to be observed at runtime. The method uses timed automata for modeling the behavior of the multi-robot system, a method unlikely to be accepted by a broader robotics community. Its expressiveness is lower than of linear hybrid control systems, which means they require careful modeling and abstraction for any typical robotics problem. Our ontology-based approach presented offers more natural models of trade-offs and requirements, including non-linearities, thanks to the use of description logics.

Aldrich et al. [4] leverage predictive data models to enable automated robot adaptation to changes in the environment at run-time. While the approach clearly depicts the benefits of using models by capturing high-level artefacts (combining application and system logic), it makes it extremely challenging for a developer to make use of them in robotic scenarios because: (i) it does not introduce models that can be reused for a different application; (ii) it does not give insights how someone can build similar models; (iii) it does not provide automated infrastructure to leverage those models.

The BRICS [10] and RobMoSys projects [32] have proposed model-based approaches to develop robotic systems. Our work is inspired by the meta-models resulting from those efforts, but goes beyond realizing their proposed future roadmap to use the models at run time for adaptation and autonomy.

Cheng et al. propose [12] a framework that uses GSN assurance case models to manage run-time adaptations for ROS-based systems. While the framework clearly shows the value of systematically integrating assurance information from GSN models to ROS specific information to guide runtime monitoring and adaptation, it has couple of drawbacks: (i) it uses custom developed libraries specific to the approach, rather than standard libraries in ROS (such as ROS Diagnostics); (ii) it does not reuse preexisting practises developed in the ROS ecosystem raising the entry barrier for ROS developers to effectively use it. Also, thanks to the tight integration of the ROS framework and model-driven automation, we are able to execute extensive experiments, beyond simple demos. We quantify the improvement of the self-adaptation, and its computational cost.

As discussed in section 2, our work exploits of two existing streams of contributions. Hernández et al. [21, 23] proposed principles for knowledge-based self-adaptation and the concept meta-control based on off-the-shelf reasoners, but did not provide a solution to obtain the architectural models or to realize their approach on robotic systems. Hammoudeh et al. [18] created the model-based RosModel tooling to support development ROS systems, but it could not model architectural variants or quality concerns. Our work leverages both works to create a complete solution to develop self-adaptation in ROS systems.

## 8. Conclusions and Future Work

This paper introduces the MROS framework to support the integration of self-adaptation in robot control architectures based on ROS. MROS promotes separation of system management decision making from the mission logic, using ontological reasoning and

design models.

MROS uses platform specific DSLs to represent the architectural models of variants of the robotic system. We have developed extensions of the *RosModel* tools that lower the modelling effort for developers and allow to automate the implementation and deployment of the self-adaptation mechanisms. The use of an ontology-based knowledge based and reasoning for the MROS MAPE-K loop facilitates reusing MROS adaptation rules across robotic applications, and extend them to address new quality concerns beyond safety and energy, which have been addressed in this work.

We developed two MROS tools: *RosModel* tooling extensions to support the development phase, and a ROS package to deploy the MROS meta-controller for run time adaptation, and applied in two different complex robotic systems for validation. Our results from systematic experiments in the mobile robot demonstrator show how MROS can be used to increase mission reliability and safety and energy concerns. The implementation also illustrates how the MROS DSL approach at design time and ontological approach at run time allowed to easily reuse the adaptation rules from component failure, and extend them to safety concerns and then to energy.

The architecture of MROS provides a framework to explore the integration of methods from robotics, software engineering and artificial intelligence. For example, currently we are investigating deep learning methods for quality prediction, and the integration of MROS meta-controller with different solutions to manage the mission logic, such as behaviour trees or planning methods. Moreover, we are working on a version of MROS for ROS2, which offers better support for component and system management than ROS1, and more industrial traction. This will open new opportunities and challenges for self adaptive systems and modelling in robotics.

## References

- [1] “ROS Diagnostics Package.” [Online]. Available: <http://wiki.ros.org/diagnostics>
- [2] ABB, “Abb demonstrates concept of mobile laboratory robot for hospital of the future,” <https://new.abb.com/news/detail/37279/hospital-of-the-future>, Oct 2019.
- [3] E. Aguado, Z. Milosevic, C. Hernández, R. Sanz, M. Garzon, D. Bozhinoski, and C. Rossi, “Functional self-awareness and metacontrol for underwater robot autonomy,” *Sensors*, vol. 21, no. 4, 2021. [Online]. Available: <https://www.mdpi.com/1424-8220/21/4/1210>
- [4] J. Aldrich, D. Garlan, C. Kästner, C. Le Goues, A. Mohseni-Kabir, I. Ruchkin, S. Samuel, B. Schmerl, C. S. Timperley, M. Veloso *et al.*, “Model-based adaptation for robotics software,” *IEEE Software*, vol. 36, no. 2, pp. 83–90, 2019.
- [5] N. Bencomo, R. B. France, B. H. C. Cheng, and U. Aßmann, Eds., *Models@run.time — Foundations, Applications, and Roadmaps*, ser. Lecture Notes in Computer Science, vol. 8378. Springer, 2014.
- [6] N. Bencomo, S. Götz, and H. Song, “Models@run.time: a guided tour of the state of the art and research challenges,” *Software & Systems Modeling*, Jan 2019. [Online]. Available: <https://doi.org/10.1007/s10270-018-00712-x>
- [7] G. Blair, N. Bencomo, and R. B. France, “Models@ run.time,” *Computer*, vol. 42, no. 10, pp. 22–27, Oct 2009.
- [8] M. Blanke, M. Kinnaert, J. Lunze, and M. Staroswiecki, *Diagnosis and Fault-Tolerant Control*. Springer-Verlag Berlin, 2006.
- [9] D. Brugali, R. Capilla, R. Mirandola, and C. Trubiani, “Model-based development of qos-aware reconfigurable autonomous robotic systems,” in *2018 Second IEEE*

- International Conference on Robotic Computing (IRC)*, Jan 2018, pp. 129–136.
- [10] H. Bruyninckx, M. Klotzbücher, N. Hochgeschwender, G. Kraetzschmar, L. Gherardi, and D. Brugali, “The BRICS Component Model: A Model-based Development Paradigm for Complex Robotics Software Systems,” in *Proceedings of the 28th Annual ACM Symposium on Applied Computing*, ser. SAC ’13. New York, NY, USA: ACM, 2013, pp. 1758–1764. [Online]. Available: <http://doi.acm.org/10.1145/2480362.2480693>
  - [11] R. Capilla, J. Bosch, P. Trinidad, A. R. Cortés, and M. Hinchey, “An overview of dynamic software product line architectures and techniques: Observations from research and industry,” *J. Syst. Softw.*, vol. 91, pp. 3–23, 2014. [Online]. Available: <https://doi.org/10.1016/j.jss.2013.12.038>
  - [12] B. H. Cheng, R. J. Clark, J. E. Fleck, M. A. Langford, and P. K. McKinley, “AC-ROS: Assurance Case driven Adaptation for the Robot Operating System,” in *Proceedings of the 23rd ACM/IEEE International Conference on Model Driven Engineering Languages and Systems*, 2020, pp. 102–113.
  - [13] W. Chung, S. Kim, M. Choi, J. Choi, H. Kim, C.-b. Moon, and J.-B. Song, “Safe navigation of a mobile robot considering visibility of environment,” *IEEE Transactions on Industrial Electronics*, vol. 56, no. 10, pp. 3941–3950, 2009.
  - [14] N. Correll, K. E. Bekris, D. Berenson, O. Brock, A. Causo, K. Hauser, K. Okada, A. Rodriguez, J. M. Romano, and P. R. Wurman, “Lessons from the amazon picking challenge,” *CoRR*, vol. abs/1601.05484, 2016. [Online]. Available: <http://arxiv.org/abs/1601.05484>
  - [15] K. Czarnecki, P. Grünbacher, R. Rabiser, K. Schmid, and A. Wasowski, “Cool features and tough decisions: a comparison of variability modeling approaches,” in *Sixth International Workshop on Variability Modelling of Software-Intensive Systems, Leipzig, Germany, January 25-27, 2012. Proceedings*, U. W. Eisenecker, S. Apel, and S. Gnesi, Eds. ACM, 2012, pp. 173–182. [Online]. Available: <https://doi.org/10.1145/2110147.2110167>
  - [16] J. L. Fernandez-Sánchez and C. Hernández, *Practical Model Based Systems Engineering*. Artech House, in print 2019.
  - [17] D. Garlan, “Software architecture: A travelogue,” in *Proceedings of the on Future of Software Engineering*, ser. FOSE 2014. New York, NY, USA: ACM, 2014, pp. 29–39. [Online]. Available: <http://doi.acm.org/10.1145/2593882.2593886>
  - [18] N. Hammoudeh Garcia, M. Lüdtke, S. Kortik, B. Kahl, and M. Bordignon, “Bootstrapping mde development from ros manual code - part 1: Metamodeling,” in *2019 Third IEEE International Conference on Robotic Computing (IRC)*, Feb 2019, pp. 329–336.
  - [19] N. Hammoudeh García, H. Deshpande, A. Santos, B. Kahl, and M. Bordignon, “Bootstrapping mde development from ros manual code: Part 2—model generation and leveraging models at runtime,” *Software and Systems Modeling*, Apr 2021. [Online]. Available: <https://doi.org/10.1007/s10270-021-00873-2>
  - [20] C. Hernandez and J. L. Fernandez-Sanchez, “Model-based systems engineering to design collaborative robotics applications,” in *2017 IEEE International Systems Engineering Symposium (ISSE)*, Oct 2017, pp. 1–6.
  - [21] C. Hernández, J. Bermejo-Alonso, and R. Sanz, “A self-adaptation framework based on functional knowledge for augmented autonomy in robots,” *Integrated Computer-Aided Engineering*, vol. 25, no. 2, pp. 157–172, 2018.
  - [22] C. Hernandez Corbato, M. Bharatheesha, J. van Egmond, J. Ju, and M. Wisse, “Integrating different levels of automation: Lessons from winning the amazon robotics challenge 2016,” *IEEE Transactions on Industrial Informatics*, vol. PP,



- no. 99, pp. 1–1, 2018.
- [23] C. Hernandez Corbato, Z. Milosevic, C. Olivares, G. Rodriguez, and C. Rossi, “Meta-control and self awareness for the ux-1 autonomous underwater robot,” in *Robot 2019: Fourth Iberian Robotics Conference*, M. F. Silva, J. Luís Lima, L. P. Reis, A. Sanfeliu, and D. Tardioli, Eds. Springer International Publishing, 2020, pp. 404–415.
  - [24] M. U. Iftikhar and D. Weyns, “Activforms: Active formal models for self-adaptation,” in *Proceedings of the 9th International Symposium on Software Engineering for Adaptive and Self-Managing Systems*, 2014, pp. 125–134.
  - [25] Kang et al., “Feature-oriented domain analysis (foda) feasibility study,” CMU/SEI-90-TR-21, Tech. Rep., 1990.
  - [26] J. Kephart and D. Chess, “The vision of autonomic computing,” *Computer*, vol. 36, no. 1, pp. 41 – 50, January 2003.
  - [27] J.-B. Lamy, “Owlready: Ontology-oriented programming in python with automatic classification and high level constructs for biomedical ontologies,” *Artificial Intelligence in Medicine*, vol. 80, pp. 11 – 28, 2017. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0933365717300271>
  - [28] A. Lotz, J. F. Inglés-Romero, C. Vicente-Chicote, and C. Schlegel, “Managing run-time variability in robotics software by modeling functional and non-functional behavior,” in *Enterprise, Business-Process and Information Systems Modeling*, S. Nurcan, H. A. Proper, P. Soffer, J. Krogstie, R. Schmidt, T. Halpin, and I. Bider, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 441–455.
  - [29] E. Marder-Eppstein, E. Berger, T. Foote, B. Gerkey, and K. Konolige, “The office marathon: Robust navigation in an indoor office environment,” in *International Conference on Robotics and Automation*, 2010.
  - [30] K. Pohl, G. Böckle, and F. van der Linden, *Software Product Line Engineering*. Springer Verlag, 2005.
  - [31] M. Quigley, K. Conley, B. P. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng, “ROS: an open-source Robot Operating System,” in *ICRA Workshop on Open Source Software*, vol. 3, 2009, pp. 1–6.
  - [32] RobMoSys Project, “RobMoSys - Composable Models and Software.” [Online]. Available: <https://robmosys.eu/>
  - [33] S. Thrun, M. Montemerlo, H. Dahlkamp, D. Stavens, A. Aron, J. Diebel, P. Fong, J. Gale, M. Halpenny, G. Hoffmann, K. Lau, C. Oakley, M. Palatucci, V. Pratt, P. Stang, S. Strohband, C. Dupont, L.-E. Jendrossek, C. Koelen, C. Markey, C. Rummel, J. van Niekirk, E. Jensen, P. Alessandrini, G. Bradski, B. Davies, S. Ettinger, A. Kaehler, A. Nefian, and P. Mahoney, “Winning the darpa grand challenge,” *Journal of Field Robotics*, 2006, accepted for publication.
  - [34] D. Weyns, N. Bencomo, R. Calinescu, J. Camara, C. Ghezzi, V. Grassi, L. Grunske, P. Inverardi, J.-M. Jezequel, S. Malek, R. Mirandola, M. Mori, and G. Tamburrelli, “Perpetual assurances for self-adaptive systems,” in *Software Engineering for Self-Adaptive Systems III. Assurances*, R. de Lemos, D. Garlan, C. Ghezzi, and H. Giese, Eds. Cham: Springer International Publishing, 2017, pp. 31–63.