

# Towards Compositional Verification for Modular Robotic Systems \*

Rafael C. Cardoso   Louise A. Dennis   Marie Farrell   Michael Fisher   Matt Luckcuck

Department of Computer Science  
The University of Manchester  
Manchester, United Kingdom

{rafael.cardoso, louise.dennis, marie.farrell, michael.fisher, matthew.luckcuck}@manchester.ac.uk

Software engineering of modular robotic systems is a challenging task, however, verifying that the developed components all behave as they should individually and as a whole presents its own unique set of challenges. In particular, distinct components in a modular robotic system often require different verification techniques to ensure that they behave as expected. Ensuring whole system consistency when individual components are verified using a variety of techniques and formalisms is difficult. This paper discusses how to use compositional verification to integrate the various verification techniques that are applied to modular robotic software, using a First-Order Logic (FOL) contract that captures each component's assumptions and guarantees. These contracts can then be used to guide the verification of the individual components, be it by testing or the use of a formal method. We provide an illustrative example of an autonomous robot used in remote inspection. We also discuss a way of defining confidence for the verification associated with each component.

## 1 Introduction

Autonomous robotic systems are being used more frequently in safety-critical scenarios. Examples include monitoring offshore structures [16], nuclear inspection and decommissioning [4], and space exploration [31]. Ensuring that the software which controls the robot behaves as it should is crucial, particularly as modern robotic systems become more modular, and are deployed alongside humans in both safety- and mission-critical scenarios.

Robotic software is often developed using a middleware to facilitate interoperability, such as ROS<sup>1</sup> or GenoM<sup>2</sup>. They share some abstract concepts [26], namely that systems are composed of communicating components. The approach that we describe in this paper is not restricted to any particular robotic middleware, instead it can be applied to a range of modular systems.

Distinct components in a modular system often require different verification techniques, ranging from software testing to formal methods. In fact, *integrating* (formal and non-formal) verification techniques is crucial particularly for robotic applications [14]. It is essential for the verification to be carried out using the most suitable technique or formalism for each component. However, performing compositional verification via linking heterogeneous verification results of individual components is difficult and the current state-of-the-art for robotic software development does not provide an easy way of achieving this.

Our approach is to construct a high-level First-Order Logic (FOL) contract specification of the system. The FOL contract describes the expected input, required output, and Assume-Guarantee [19] conditions for each component. This abstract specification can be seen as a logical prototype for individual

---

\*Work supported by UK Research and Innovation, and EPSRC Hubs for Robotics and AI in Hazardous Environments: EP/R026092 (FAIR-SPACE), EP/R026173 (ORCA), and EP/R026084 (RAIN).

<sup>1</sup><http://www.ros.org/>

<sup>2</sup><https://www.openrobots.org/wiki/genom>

components and, via our calculus for chaining individual specifications, the entire modular system. Once the FOL contracts have been checked for consistency, they can be used to guide the verification of each component. This involves ensuring that the verification encodes the assumptions and guarantees as, for example, test cases, assertions, or formal properties.

Our approach can be used in a Top-Down manner, to guide the system’s development from abstract specification to concrete implementation, via verification; or in a Bottom-Up fashion, to check the consistency of existing verification techniques. In this paper we present a Top-Down example. Our approach enables developers to choose the most suitable verification technique for each component, but also to link the conditions being verified across the whole system.

The next section presents background and related work. Section 3 describes our contribution for specifying modular systems using contracts written in FOL; then, we illustrate our approach via a simple remote inspection example. In Section 4 we discuss how the use of distinct forms of verification (e.g., testing, simulation, formal methods, etc.) affects our confidence in the verification results for the whole system. Finally, Section 5 concludes and outlines future directions.

## 2 Background and Related Work

Modular robotic systems that are used in safety- or mission-critical scenarios require robust verification techniques to ensure and certify that they behave as intended. These techniques encompass a range of software engineering tools and methodologies: from practical testing and simulation, through to formal verification [20, 14]. For example, to show that a property always holds: model-checkers [9] exhaustively explore the search space, and theorem provers [3] use mathematical proof. These techniques offer a means for proving that the software system is correct, which can be used as evidence to improve public trust or to gain certification, as needed. Formal verification can be applied to the implemented system; or to an abstract specification, which can then be further refined to program code.

Many formal methods exist, and most follow a simple paradigm: if the program is executed in a state satisfying a given property (the pre-condition), then it will terminate in a state that satisfies another property (the post-condition) [13]. It is up to the software engineer to specify these properties (sometimes referred to as assumptions and guarantees, respectively). The *Assume-Guarantee* (or *Rely-Guarantee*) approach [19] enables components to be assessed (verified) individually and then the separate assessments to be combined, potentially under some assumption about the underlying concurrency. To combine assessments, specifications of a component’s pre- and post-conditions are needed; this is similar to the notion of Hoare logic that underlies these techniques [17]. In this work, we refer to this combination of assumption/pre-condition and guarantee/post-condition as a *contract* that must be preserved.

Our work takes inspiration from Broy’s logical approach to systems engineering [5, 6] in which he distinguishes three kinds of artefacts: (1) system-level requirements, (2) functional system specification, and (3) logical subsystem architecture. Each of these artefacts is represented as a logical predicate in the form of assertions, with relationships defined between them which he then extends to assume/commitment contracts. His treatment of these contracts is purely logical [6], and, in this paper, we present a similar technique that is specialised to the software engineering of modular robotic systems.

In [21], a workflow is proposed for systematically verifying the design of models of a cyber-physical system using a combination of formal refinement and model-checking. Our work also deals with several different levels of abstraction, but we tackle the use of compositional verification in modular systems.

Recent work has analysed a portion of the literature and identified common patterns that appear in robotic missions (e.g. patrolling and obstacle avoidance) [22]. They provide verified LTL/CTL specifica-




Figure 1: We specify the contract for each component. Here we denote the assumption/pre-condition by  $\mathcal{A}(\bar{i})$  and the guarantee/post-condition by  $\mathcal{G}(\bar{o})$ . These contracts are used to guide the verification approach applied to each component, denoted by dashed lines, such as software testing for a black-box implementation (Component 1). The solid arrows represent data flow. Note that we use the bar notation  $(\bar{i}, \bar{o})$  to indicate a vector of variables.

tions for these commonly found missions which can then be reused in future developments. Further, it is not clear how their support for compositional verification can be extended to support heterogeneous components such as those in our example. Related work includes the identification of Event-B specification clones in cyber physical system specifications [15].

Other compositional approaches include OCRA [8], AGREE [10], CoCoSpec [7], SIMPAL [29], DRONA [12], and a methodology to decompose a system into assume-guarantee contracts that are then validated through simulation [27], however, none explicitly incorporates heterogeneous techniques.

### 3 A First-Order Logic Framework

No single verification approach suits every component in a modular robotic system [14]. Components such as hardware interfaces or planners may be amenable to formal verification, whereas, sensors may require software testing or simulation-based testing.

As illustrated in Fig. 1, we could use logical specifications (e.g. temporal logic), model-based specifications (e.g. Event-B [1] or Z [28]), or algebraic specifications (e.g. CSP [18] or CASL[24]) amongst others to specify the components in the system. Each of these formalisms offers its own range of benefits, and each tends to suit the verification of particular types of behaviour. Also, we may only have access to the black-box or white-box implementation of a component.

Our approach facilitates the use of compositional verification techniques for the components in modular robotic systems. We achieve this by specifying high-level contracts in FOL and we employ temporal logic for reasoning about the combination of these contracts. In this way, we attach the assumptions over the input ( $\mathcal{A}(\bar{i})$ ) and guarantees over the output ( $\mathcal{G}(\bar{o})$ ) to individual components (see Fig. 1).

#### 3.1 A Calculus for Chaining Component Specifications

Specifically, we use *typed* FOL, potentially with the addition of algebraic operators, to specify assumptions and guarantees. For each component,  $C$ , we specify  $\mathcal{A}_C(\bar{i}_C)$  and  $\mathcal{G}_C(\bar{o}_C)$  where  $\bar{i}_C$  is a variable representing the input to the component,  $\bar{o}_C$  is a variable representing the output from the component, and  $\mathcal{A}_C(\bar{i}_C)$  and  $\mathcal{G}_C(\bar{o}_C)$  are FOL formulae describing the assumptions and guarantees, respectively.

Each individual component,  $C$ , obeys the following implication:

$$\forall \bar{i}_C, \bar{o}_C. \mathcal{A}_C(\bar{i}_C)[C] \Rightarrow \diamond \mathcal{G}_C(\bar{o}_C)$$

where  $\bar{i}_C$  and  $\bar{o}_C$  are the inputs and outputs, respectively, of  $C$ ;  $\mathcal{A}_C(\bar{i}_C)[C]$  represents the execution of  $C$  with the assumption  $\mathcal{A}_C(\bar{i}_C)$ ; and ‘ $\diamond$ ’ is Linear-time Temporal Logic (LTL)’s [25] “eventually” operator.

So, this implication means that if the assumptions,  $\mathcal{A}_C(\bar{i}_C)$ , hold in the specification or program code of  $C$ , then upon execution of the component *eventually* the guarantee,  $\mathcal{G}_C(\bar{o}_C)$ , will hold. Note that our use of temporal operators here is motivated by the real-time nature of robotic systems and could be of use in later extensions of this calculus for larger, more complex systems.

Components in a modular robotic architecture are ‘chained’ together so long as their types/requirements match. Similarly, we can compose the contracts of individual components in a number of ways, the simplest being a sequential composition. The basic way to describe such structures is to first have the specification capture *all* of the input and output streams and then to describe how these are combined in the appropriate inference rules. The proof rule, PR1, for such linkage is as expected:

$$\frac{\begin{array}{l} \forall \bar{i}_1, \bar{o}_1. \mathcal{A}_1(\bar{i}_1) \Rightarrow \diamond \mathcal{G}_1(\bar{o}_1) \\ \forall \bar{i}_2, \bar{o}_2. \mathcal{A}_2(\bar{i}_2) \Rightarrow \diamond \mathcal{G}_2(\bar{o}_2) \\ \bar{o}_1 = \bar{i}_2 \\ \vdash \forall \bar{i}_2, \bar{o}_1. \mathcal{G}_1(\bar{o}_1) \Rightarrow \mathcal{A}_2(\bar{i}_2) \end{array}}{\forall \bar{i}_1, \bar{o}_2. \mathcal{A}_1(\bar{i}_1) \Rightarrow \diamond \mathcal{G}_2(\bar{o}_2)} \quad (\text{PR1})$$

Intuitively, this states that if two components are sequentially composed with the output of the first equal to the input of the second and the guarantee of the first implies the assumption of the second then, we can deduce that the assumption of the first component implies that the guarantee of the second will “eventually” hold.

The illustrative example that we present in the next section contains a neat, linear chain of components and so it is easy to see how this proof rule works. But a realistic robotic system could be much more complex than this. For example, a component might have multiple, branching output streams that are each used as input to other distinct components. Thus, we propose PR2 to account for branching:

$$\frac{\begin{array}{l} \forall \bar{i}_1, \bar{o}_1. \mathcal{A}_1(\bar{i}_1) \Rightarrow \diamond \mathcal{G}_1(\bar{o}_1) \\ \forall \bar{i}_2, \bar{o}_2. \mathcal{A}_2(\bar{i}_2) \Rightarrow \diamond \mathcal{G}_2(\bar{o}_2) \\ \bar{i}_2 \subseteq \bar{o}_1 \\ \vdash \forall \bar{i}_2, \bar{o}_1. \mathcal{G}_1(\bar{o}_1) \Rightarrow \mathcal{A}_2(\bar{i}_2) \end{array}}{\forall \bar{i}_1, \bar{o}_2. \mathcal{A}_1(\bar{i}_1) \Rightarrow \diamond \mathcal{G}_2(\bar{o}_2)} \quad (\text{PR2})$$

Here, the input to the second component,  $\bar{i}_2$ , is a subset of the output,  $\bar{o}_1$ , of the first. We have also devised other proof rules which are omitted here for brevity.

### 3.2 Remote Inspection Case Study

The various middleware often used to develop robotic systems enable the combination of subsystems, potentially developed independently of one another. For example, an autonomous rover might have a navigation subsystem that is responsible for traversing a planet’s surface, and a subsystem for sampling and analysing rock and soil samples.

In this paper, we use the illustrative example of a simplified navigation subsystem for an autonomous rover, whose goal is to perform remote inspection of particular targets, while avoiding any obstacles, on a 2D grid map that has been previously generated. This navigation subsystem is composed of components for *Detection*, a *Planner* and a cognitive *Agent*. Fig. 2 illustrates this subsystem and outlines the abstract FOL contract of each component to be verified.

Given the camera input and the size of the square grid to be explored,  $n$ , the specification of the *Detection* component guarantees that it detects an obstacle at a particular coordinate if an obstacle actually exists (in the physical or simulated environment) at that point. Note that we assume the existence

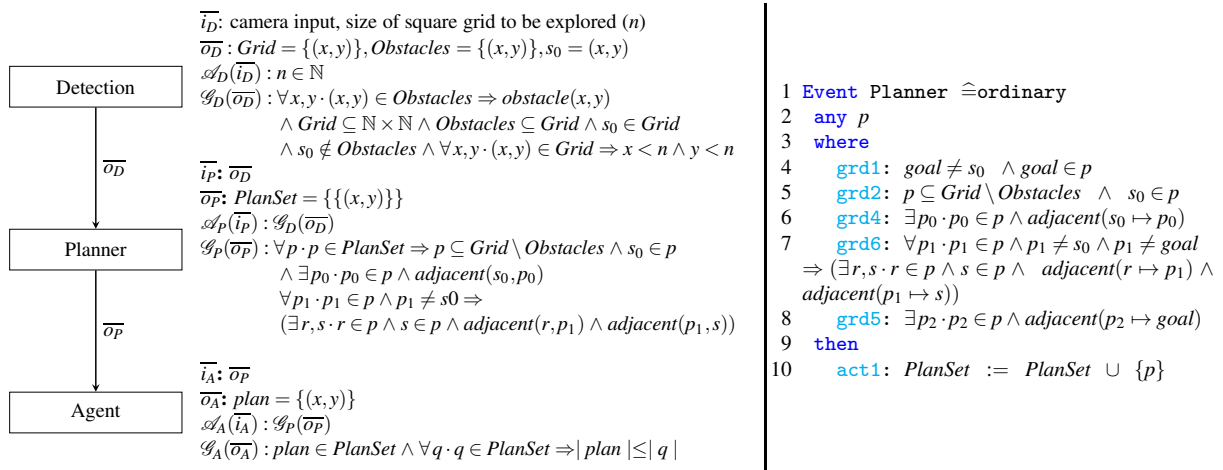


Figure 2: An overview of the robotic system to be verified. On the far left, each rectangle represents a component, and each arrow represents data flow between components. To the right of each individual component, we summarise the  $\bar{i}$ s,  $\bar{o}$ s, and their respective FOL contract. On the far right, we provide an Event-B event specification corresponding to the *Planner*. Here, ‘ $\mapsto$ ’ denotes tuples.

of the  $obstacle(x,y)$  function to represent this physical or simulated observation. Further, the *Detection* component outputs the rover’s current location, denoted by  $s_0$ , and guarantees that  $s_0$  is in the grid and does not coincide with an obstacle.

The *Planner* component produces a *PlanSet* whose elements are sets of coordinates capturing obstacle-free plans under the assumption that the *Detection* component has produced a set of *Obstacles* that refers to points that exist and contain an obstacle. The guarantee also ensures that each of these plans (sets of coordinates) in *PlanSet* can be transformed into a sequence of coordinates where each is adjacent to the previous. Note that we use sets here rather than sequences because sequences may contain duplicate entries whereas sets cannot, and so our specification rules out the case where a plan loops.

The specification of the cognitive *Agent* component states that, under the assumption that all potential plans in the *PlanSet* are obstacle-free, then it chooses a *plan* from the *PlanSet* and that this plan is, in fact, the shortest one available. Note that we use  $|plan|$  to denote the length of *plan*.

The far right of Fig. 2 contains an Event-B event specification [1] of the *Planner* component. Notice that the high-level FOL contract that we have written is captured in the Event-B specification as guards on lines 5–7. The FOL contract has been refined and the Event-B specification has a notion of a goal that it must plan paths towards. This has resulted in the extra guards on lines 1 and 8. This is an excerpt from a larger *Planner* specification that we have verified via theorem proving in the *Rodin* Platform [2].

In this illustrative example, we formally verified the *Planner* contracts using Event-B and the *Agent* contracts using the Gwendolen agent programming language [11]. The *Detection* component would be verified via standard software testing practices and simulation-based testing since it may contain a learning component which cannot be formally verified.

In §3.1, we presented rules about how we chain contracts together. In this example, two applications of PR1 would allow us to derive the following property:

$$\forall \bar{i}_D, \bar{o}_A \cdot \mathcal{A}_C(\bar{i}_D) \Rightarrow \diamond \mathcal{G}_A(\bar{o}_A)$$

meaning that if the *Detection* component receives valid input then the *Agent* will eventually return the shortest path from the start to the goal position in the grid.

We have shown how FOL can be used as a unifying language for the high-level Assume-Guarantee

specification of components in a modular robotic system. A top-down approach to software engineering of such a robotic system would begin with these FOL contracts. Then, each contract would be further refined [23], ideally to a more detailed formal specification and, eventually to its concrete implementation.

## 4 Confidence in Verification

When linking the results from multiple verification techniques a key question becomes how using these different techniques affects our confidence in the verification of the whole system. One might think that a formal proof of correctness corresponds to a higher level of confidence than simple testing methods (especially over unbounded environments). However, formal verification is usually only feasible on an abstraction of the system whereas testing can be carried out on the implemented code. Therefore, it is our view that we achieve higher levels of confidence in verification when multiple verification methods have been employed for each component in the system [30].

Component	Testing	Simulation-Based Testing	Formal Methods
<i>Detection</i>	✓	✗	✗
<i>Planner</i>	✓	✓	✓
<i>Cognitive Agent</i>	✓	✓	✓

Table 1: Verification techniques applied to each component.

We have broadly partitioned current verification techniques into three categories: testing, simulation-based testing and formal methods. We have determined which of these techniques might be employed for each component in our example as shown in Table 1. Examining how this metric can be calculated for more complex systems with loops is a future direction for this work.

## 5 Conclusions and Future Work

We have presented an initial description of a framework for compositional verification of modular systems using FOL as a unifying language. FOL contracts are used to guide the verification techniques applied to each component. We have briefly illustrated this by verifying and refining the FOL contract for the *Planner* component using Event-B in our example of a remote inspection rover. Furthermore, we introduce the notion of confidence in verification techniques and provide a broad categorisation.

Future work includes assessing how well our approach scales to industrial-sized robotic systems development. Also, ensuring that our approach is usable by developers and understandable by certification organisations is crucial to its use and success. This needs tool support for writing and reasoning about the FOL contracts, with similar functionality to an IDE, integrated robotic development other tools. We will also demonstrate the approach with a wider variety of verification techniques. Our illustrative example employs an event-based style of communication, however, we intend to explore how to extend our approach to stream-based communications as future work.

Finally, our proposed method could be used to generate runtime monitors to check that each component’s contract holds during execution. If a contract is violated, the monitor could flag this to the robot’s operator (if one exists) or trigger a mitigating action (if the robot is autonomous). This could augment the verification of components that can only be verified at a low level of confidence, or aid in fault finding when diagnosing failures.

## References

- [1] Jean-Raymond Abrial (2010): *Modeling in Event-B*. Cambridge University Press, doi:10.1017/CB09781139195881.
- [2] Jean-Raymond Abrial, Michael Butler, Stefan Hallerstede, Thai Son Hoang, Farhad Mehta & Laurent Voisin (2010): *Rodin: an open toolset for modelling and reasoning in Event-B*. *International Journal on Software Tools for Technology Transfer* 12(6), pp. 447–466, doi:10.1007/s10009-010-0145-y.
- [3] Yves Bertot & Pierre Castéran (2013): *Interactive theorem proving and program development: Coq'Art: the calculus of inductive constructions*. Springer, doi:10.1007/978-3-662-07964-5.
- [4] Robert Bogue (2011): *Robots in the nuclear industry: a review of technologies and applications*. *Industrial Robot: An International Journal* 38(2), pp. 113–118, doi:10.1108/01439911111106327.
- [5] Manfred Broy (2018): *A logical approach to systems engineering artifacts: semantic relationships and dependencies beyond traceability – from requirements to functional and architectural views*. *Software and System Modeling* 17(2), pp. 365–393, doi:10.1007/s10270-017-0619-4.
- [6] Manfred Broy (2018): *Theory and methodology of assumption/commitment based system interface specification and architectural contracts*. *Formal Methods in System Design* 52(1), pp. 33–87, doi:10.1007/s10703-017-0304-9.
- [7] Adrien Champion, Arie Gurfinkel, Temesghen Kahsai & Cesare Tinelli (2016): *CoCoSpec: A mode-aware contract language for reactive systems*. In: *International Conference on Software Engineering and Formal Methods, LNCS 9763*, Springer, pp. 347–366, doi:10.1007/978-3-319-41591-8\_24.
- [8] Alessandro Cimatti, Michele Dorigatti & Stefano Tonetta (2013): *OCRA: A tool for checking the refinement of temporal contracts*. In: *International Conference on Automated Software Engineering (ASE), IEEE*, pp. 702–705, doi:10.1109/ASE.2013.6693137.
- [9] Edmund M Clarke, Orna Grumberg & Doron Peled (1999): *Model checking*. MIT press.
- [10] Darren Cofer, Andrew Gacek, Steven Miller, Michael W Whalen, Brian LaValley & Lui Sha (2012): *Compositional verification of architectural models*. In: *NASA Formal Methods Symposium, LNCS 7226*, Springer, pp. 126–140, doi:10.1007/978-3-642-13464-7\_5.
- [11] Louise A. Dennis, Michael Fisher, Matthew P. Webster & Rafael H. Bordini (2012): *Model checking agent programming languages*. *Automated Software Engineering* 19(1), pp. 5–63, doi:10.1007/s10515-011-0088-x.
- [12] Ankush Desai, Shaz Qadeer & Sanjit A. Seshia (2018): *Programming Safe Robotics Systems: Challenges and Advances*. In Tiziana Margaria & Bernhard Steffen, editors: *Leveraging Applications of Formal Methods, Verification and Validation. Verification*, Springer International Publishing, Cham, pp. 103–119, doi:10.1007/978-3-030-03421-4\_8.
- [13] Edsger W Dijkstra (1975): *Guarded commands, nondeterminacy and formal derivation of programs*. *Communications of the ACM* 18(8), pp. 453–457, doi:10.1145/360933.360975.
- [14] Marie Farrell, Matt Luckcuck & Michael Fisher (2018): *Robotics and Integrated Formal Methods: Necessity meets Opportunity*. In: *Integrated Formal Methods, LNCS 11023*, Springer, pp. 161–171, doi:10.1007/978-3-319-98938-9\_10.
- [15] Marie Farrell, Rosemary Monahan & James F Power (2017): *Specification Clones: An Empirical Study of the Structure of Event-B Specifications*. In: *International Conference on Software Engineering and Formal Methods, LNCS 10469*, Springer, pp. 152–167, doi:10.1007/978-3-319-66197-1\_10.
- [16] Helen F. Hastie, Katrin Solveig Lohan, Mike J. Chantler, David A. Robb, Subramanian Ramamoorthy, Ronald P. A. Petrick, Sethu Vijayakumar & David Lane (2018): *The ORCA Hub: Explainable Offshore Robotics through Intelligent Interfaces*. CoRR abs/1803.02100. Available at <http://arxiv.org/abs/1803.02100>.
- [17] C. A. R. Hoare (1969): *An axiomatic basis for computer programming*. *Communications of the ACM* 12(10), pp. 576–580, doi:10.1145/363235.363259.

- [18] C. A. R. Hoare (1978): *Communicating sequential processes*. *Communications of the ACM* 21(8), pp. 666–677, doi:10.1145/359576.359585.
- [19] Cliff B. Jones (1983): *Tentative Steps Toward a Development Method for Interfering Programs*. *ACM Transactions on Programming Languages and Systems* 5(4), pp. 596–619, doi:10.1145/69575.69577.
- [20] Matt Luckcuck, Marie Farrell, Louise A. Dennis, Clare Dixon & Michael Fisher (2019): *Formal Specification and Verification of Autonomous Robotic Systems: A Survey*. *ACM Comput. Surv.* 52(5), pp. 1–41, doi:10.1145/3342355.
- [21] Christoph Luckeneder & Hermann Kaindl (2018): *Systematic top-down design of cyber-physical models with integrated validation and formal verification*. In: *International Conference on Software Engineering*, pp. 274–275, doi:10.1145/3183440.3194967.
- [22] Claudio Menghi, Christos Tsigkanos, Patrizio Pelliccione, Carlo Ghezzi & Thorsten Berger (2019): *Specification patterns for robotic missions*. *IEEE Transactions on Software Engineering*, doi:10.1109/TSE.2019.2945329.
- [23] Carroll Morgan, Ken Robinson & Paul Gardiner (1988): *On the Refinement Calculus*. Springer, doi:10.1007/978-1-4471-3273-8.
- [24] Peter D Mosses (2004): *CASL reference manual: The complete documentation of the common algebraic specification language*. Springer, doi:10.1007/b96103.
- [25] A. Pnueli (1977): *The Temporal Logic of Programs*. In: *18th Symposium on the Foundations of Computer Science*, IEEE, pp. 46–57, doi:10.1109/SFCS.1977.32.
- [26] Azamat Shakhimardanov, Nico Hochgeschwender & Gerhard Kraetzschmar (2010): *Component models in robotics software*. In: *Workshop on Performance Metrics for Intelligent Systems*, ACM, pp. 82–87, doi:10.1145/2377576.2377592.
- [27] Stefano Spellini, Michele Lora, Franco Fummi & Sudipta Chattopadhyay (2019): *Compositional Design of Multi-Robot Systems Control Software on ROS*. *ACM Trans. Embed. Comput. Syst.* 18(5s), doi:10.1145/3358197.
- [28] J Michael Spivey (1988): *Understanding Z: a specification language and its formal semantics*. Cambridge University Press.
- [29] Lucas Wagner, David Greve & Andrew Gacek (2017): *SIMPAL: A Compositional Reasoning Framework for Imperative Programs*. In: *Proceedings of the 24th ACM SIGSOFT International SPIN Symposium on Model Checking of Software*, SPIN 2017, Association for Computing Machinery, New York, NY, USA, p. 90–93, doi:10.1145/3092282.3092290.
- [30] Matt Webster, David Western, Dejanira Araiza-Illan, Clare Dixon, Kerstin Eder, Michael Fisher & Anthony G Pipe (2020): *A corroborative approach to verification and validation of human–robot teams*. *The International Journal of Robotics Research* 39(1), pp. 73–99, doi:10.1177/0278364919883338.
- [31] Brian H. Wilcox (1992): *Robotic vehicles for planetary exploration*. *Applied Intelligence* 2(2), pp. 181–193, doi:10.1007/BF00058762.