

PROMISE: High-Level Mission Specification for Multiple Robots

Sergio García
Chalmers | University of Gothenburg
Gothenburg, Sweden
sergio.garcia@gu.se

Patrizio Pelliccione
Chalmers | University of Gothenburg
Gothenburg, Sweden
University of L'Aquila
L'Aquila, Italy
patrizio.pelliccione@gu.se

Claudio Menghi
University of Luxembourg
Luxembourg, Luxembourg
claudio.menghi@uni.lu

Thorsten Berger
Chalmers | University of Gothenburg
Gothenburg, Sweden
thorsten.berger@gu.se

Tomas Bures
Charles University
Prague, Czech Republic
bures@d3s.mff.cuni.cz

ABSTRACT

Service robots, a type of robots that perform useful tasks for humans, are foreseen to be broadly used in the near future in both social and industrial scenarios. Those robots will be required to operate in dynamic environments, collaborating among them or with users. Specifying the list of requested tasks to be achieved by a robotic team is far from being trivial. Therefore, mission specification languages and tools need to be expressive enough to allow the specification of complex missions (e.g., detailing recovery actions), while being reachable by domain experts who might not be knowledgeable of programming languages. To support domain experts, we developed PROMISE, a Domain-Specific Language that allows mission specification for multiple robots in a user-friendly, yet rigorous manner. PROMISE is built as an Eclipse plugin that provides a textual and a graphical interface for mission specification. Our tool is in turn integrated into a software framework, which provides functionalities as: (1) automatic generation from specification, (2) sending of missions to the robotic team; and (3) interpretation and management of missions during execution time. PROMISE and its framework implementation have been validated through simulation and real-world experiments with four different robotic models.

Video: <https://youtu.be/RMtqwY2GOIQ>

ACM Reference Format:

Sergio García, Patrizio Pelliccione, Claudio Menghi, Thorsten Berger, and Tomas Bures. 2020. PROMISE: High-Level Mission Specification for Multiple Robots. In *42nd International Conference on Software Engineering Companion (ICSE '20 Companion)*, May 23–29, 2020, Seoul, Republic of Korea. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3377812.3382143>

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).
ICSE '20 Companion, May 23–29, 2020, Seoul, Republic of Korea
© 2020 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-7122-3/20/05.
<https://doi.org/10.1145/3377812.3382143>

1 INTRODUCTION

It is envisioned the increase of investment and inclusion of service robots in several market sectors (e.g., logistics, medical, or farming) in the next few years.¹ As the number of deployed service robots increases, humans will likely have to interact with robots in their everyday life. Imagine for instance a near-future social or industrial environment (e.g., a hotel, a hospital, or a warehouse) where teams of service robots are deployed to achieve tasks collaboratively. Those tasks will be configured and specified by domain experts and end-users who may not be knowledgeable on programming languages, according to the H2020 Robotics Multi-Annual Roadmap (MAR).² According to this roadmap, this will require: (1) “highly abstracted mission definition [...] algorithms for interaction and operation with untrained users.” (2) and that robots “need to become intuitively integrated with human operators.” In fact, these requirements are confirmed by two recent studies [8, 9], whose authors claim that precisely specifying missions and transforming them for automatic processing are one of the main challenges in robotics software engineering. Therefore, the current state of robot mission specification methods and tools must become friendlier in their usage to the user. We aim to promote the simplicity of our tool while keeping its expressiveness. At the same time, we strive to enable a rigorous specification that describes precisely and unambiguously the mission the robot must perform. As explained in a recent study [7], existing solutions fail at providing a trade-off of those qualities.

In this paper, we present a framework that supports users in mission specification for multi-robot applications. The framework integrates PROMISE (simPle ROBot MIssion SpEcification), a language realized as a Domain-Specific Language (DSL). The DSL permits the specification of complex missions (i.e., complex behaviors, including recovery actions) while being accessible for users without programming skills. PROMISE builds upon a catalog of mission specification Linear Temporal Logic (LTL)-based patterns [9], ensuring the

¹www.ifr.org

²<https://eu-robotics.net>

correctness of the semantics and therefore the rigor of the tool. The framework also integrates software components for the automatic mission generation, sending, and interpretation and management. We detail the framework implementation, workflow, and evaluation of PROMISE through experimentation.

Running example. We use an example inspired by a milestone experiment from a European project (<http://www.co4robots.eu/>) with which our DSL’s development is involved to illustrate different aspects throughout the paper. The example represents two robots collaboratively working in an industrial warehouse. A mobile platform $r1$ patrols a set of locations (assembly stations). Users may use the robot to deliver items to other stations and can stop the patrolling using a specific gesture. $r1$ continuously evaluates if station 2 ($l2$) has assembled a product (which is recognizable by the robot). If $r1$ detects such a product, it requests help to $r2$. The mobile manipulator $r2$ that has been waiting in an idle state now moves to $l2$ and loads the object on top of $r1$. If the loading task fails, $r2$ communicates the failure through its speakers.

Users. We now categorize the users of our framework.

Domain expert. User characterized as a worker from one of the scenarios previously described referring to the MAR. This user might be knowledgeable of robotics but not of programming languages. Their main interaction with PROMISE is to specify missions (as the one shown in Fig. 1) by instantiating actions that were previously implemented by *developers*.

Developer. User with programming knowledge in charge of developing new robotic skills and actions to the tool. Referring to the running example, a possible use case for the developer is to develop and implement robotic actions, as the ones that permit the robots to recognize specific gestures from a human.

Maintainer. User in charge of performing maintenance services over the whole framework. We envision a use case for the maintainer: to assure the correct functioning of the framework and its implementation with other solutions that may provide robotic functionalities as planning or motion control.

2 RELATED WORK

Petri Nets [13] and Statecharts [1] are solutions for mission specification. As opposed to our goal, these solutions rely on an imperative, step-by-step description of a robot’s mission. DSLs are a great approach to solve domain-specific problems [11], such as mission specification in robotics. Various DSLs have been proposed for modeling robotic systems and their missions [10]. However, most of the effort has been made for the industrial robots domain, a field where users are often specialized and where robots work in controlled environments not populated by humans. On the contrary, human-populated environments are typically dynamic and unconstrained.

Many robotic companies have developed mission specification mechanisms—e.g., Choregraphe (<https://community.ald.softbankrobotics.com>)—which are made available for their robots as IDEs and frameworks [10]. Yet, those mechanisms are often platform- and robot-dependent. Among the existing mission specification proposals that are not developed or are constrained to a specific robot, many require programming

skills from the user [5, 12]. Contrarily, our goal is to promote the simplicity of our language to make it accessible by end-users and domain experts. Some works also aimed at providing user-friendly solutions, like the one proposed by Bozhinoski et al. [2]. Their tool allows non-technical users to define high-level specifications of missions for teams of multi-copters. However, the user is required to precisely detail the robots’ behavior while we aim to provide a declarative specification where the user needs only to specify a high-level goal instead of the intermediate steps required to achieve it.

Finally, Doherty et al. [4] propose a high-level mission-specification language for multi-robot applications that supports abstract yet complex and declarative mission specification that uses operators. However, their catalog of composition operators is limited and their application targets multi-copter vehicles, while we focus on ground service robots.

3 PROMISE

PROMISE allows the definition of *global missions* to be achieved by a team of robots and its manual decomposition into *local missions*, i.e., robot-specific missions [6]. Our language builds upon a recently proposed catalog of mission specification LTL-based patterns [9]. These patterns are used as atomic tasks that may be composed utilizing a set of operators we proposed in our previous work [7] and which are inspired by behavior tree operators [3]. We conceived two types of operators: (1) composition operators allow the definition of complex missions by combining other operators and managing events that may occur in the environment; (2) delegate operators make a robot perform a specific task instantiated with parameters as locations or actions. In turn, each task translates to an LTL formula—e.g., from the running example, the task *Sequenced patrolling* $l1, l2, l3$ performed by $r1$ translates to $\mathcal{G} \left(\mathcal{F} (l1 \wedge \mathcal{F} (l2 \wedge \mathcal{F} l3)) \right)$. This specific task makes the robot visit infinitely often a set of locations one after the other. These formulas are used by an underlying planner, which uses a model of the environment and of the robots to compute a plan—a set of low-level actions that, if performed, allows the achievement of the given mission.

PROMISE provides a graphical and a textual syntax, each of which is supported by a dedicated editor. The graphical syntax maps mission specification concerns to graphical elements that can be understood by users not knowledgeable on programming languages. Using this syntax, operators and tasks can be just *dragged&dropped* from a palette and interconnected accordingly or specified textually.

Our DSL defines missions that can react to external events, permitted by the semantics of certain operators. Robots may synchronize among them through message exchanging—one robot performs a sending message action and another “reads” it by reacting to a predefined event. This method is however rather rigid, we plan to conceive a new operator to define synchronization capabilities among robots, which will enhance the flexibility of PROMISE.

Figure 1 shows the graphical definition of a mission, the circled numbers were added to improve readability. Node ①

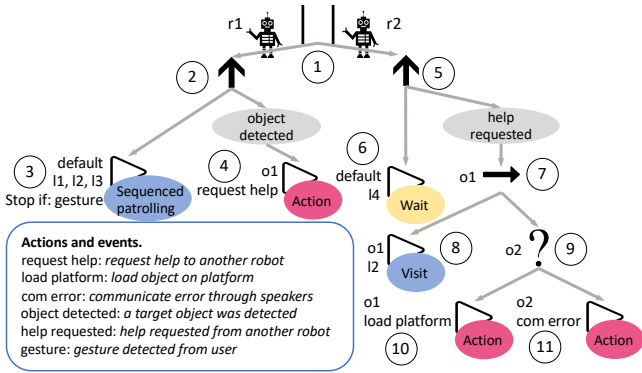


Figure 1: Mission specification using the graphical syntax (running example).

decomposes the global mission into two local missions, (2) for $r1$ and (5) for $r2$. Note that the ordering of every operator is textually represented in its label (i.e., *default* and o_i). Node (2) represents an operator *event handler*, which takes as input a set of operators and events. Operators and events are associated with other operators in a parent-child structure through edges—represented in Fig. 1 with gray arrows. When executed, the operator starts by performing its default mission, that is, the operator labeled as *default*. The execution of the event handler will not end until the execution of the default mission is finished (either succeeding or failing). In the example, the default mission of $r1$ is to sequentially patrol $l1$, $l2$, and $l3$ (3). PROMISE also allows the definition of events that if detected, will stop the execution of a task. A configured delegate operator will display in its label the string “Stop if: *event*”. In this case, if $r1$ detects a specific gesture from a user it will stop its patrolling (3), which otherwise would go on forever. The event handler is also associated with behaviors that are triggered by events (represented as gray ellipses). Thus, if *object detected* is sensed by $r1$ during its default mission (3), $r1$ will perform the action *request help* (4). This specific action sends a message requesting help to a specific robot ($r2$). The default mission of an event handler is resumed as soon as the execution of one of its event-triggered behaviors is finished. An excerpt from the mission specification using the textual syntax is shown in Fig. 2, and it corresponds to the instantiation of the event handler (2) and its children.

Meanwhile, $r2$ waits by default in $l4$ (5 and 6), and if during this idle state the robot receives a request of help, it will perform a set of tasks in a sequential manner (7), starting from moving to $l2$ (8). Then, the operator *fallback* (9) starts by executing its first child (10), i.e., load the mobile platform $r1$ with a target object. If this task fails (e.g., the object falls from the robot’s gripper), the operator *fallback* dictates that the following child (11) is executed, and so on. With (11), the robot must communicate that a failure has occurred and then resume its waiting state in $l4$.

```

eventHandler(
  default(
    delegate (SequencedPatrolling locations l1, l2, l3 stoppingEvents finish))
  except found_object (
    delegate (SimpleAction actions request_help)))

```

Figure 2: Textual syntax excerpt.




Figure 3: PROMISE and its software framework.

4 SOFTWARE FRAMEWORK

The software framework (Fig. 3) into which PROMISE is integrated consists of four elements distributed into two units. The unit *Central station* provides an interface to the user (typically deployed into a computer) and integrates elements (a), (b), and (c).

The unit *Robot* represents what is deployed and executed at run-time within each robot. Our framework builds upon an existing robotic platform and architecture, SERA [6]. This platform’s current implementation relies on ROS³ and provides a set of functionalities, including motion control, collision avoidance, image recognition, self-localization, and planning. Nevertheless, PROMISE is developed as a standalone tool and could be integrated with various robotic tools and platforms.

In Fig. 3, the component (a) encapsulates the language and DSL of PROMISE, realized as a plugin for Eclipse, using Xtext⁴ for the textual interface and Sirius⁵ for the graphical one. The compiler (component (b)) contains a script that automatically generates the local missions to be sent to each robot. This component also generates a description of the specified mission using natural English, which might help users while specifying missions to evaluate whether the description corresponds to what they wanted to express.

The intermediate language (component (c)) describes the set of tasks to be performed by each robot. The used software platform currently implements an LTL-based planner, so PROMISE’s intermediate language is composed of a set of LTL formulae with the addition of the used operators’ semantics. The intermediate language permits decoupling the mission specification from the robotic platform and the development of interpreter tools [7]. Finally, an interpreter ((d)) is deployed within each robot. This interpreter receives the robot-specific local mission and communicates the tasks to the local planner appropriately. In this way, PROMISE becomes robot-agnostic since only the robot-specific interfaces of the interpreter with the lower-level components of the platform must be adapted when using a new robot.

5 PROMISE IN ACTION

In this section, we present the workflow of PROMISE. The seamless integration of Sirius with Xtext allows the mission specification using the previously explained syntaxes in parallel. Mission specification begins by using a wizard that supports users creating the basic structure of the mission, i.e., names of the used robots, locations, events, and actions. Once finished with the wizard, the user may continue using

³<https://www.ros.org>

⁴<https://www.eclipse.org/Xtext/>

⁵<https://www.eclipse.org/sirius/>

the textual syntax (writing the mission specification in a Java-like fashion), or using the graphical syntax dragging and dropping elements from the provided palette, or using both seamlessly. We put special emphasis on the development of the graphical syntax, since our goal is to support users and domain experts without deep programming skills. This syntax displays helpful messages in the labels of the operators and interconnecting edges. For instance, labels display the order of children operators and their configurations (i.e., actions or locations set as properties to them). The labels also warn the user when operators have not been configured or remain unlinked. All these mission specification mechanisms are encapsulated in component ④ of Fig. 3.

Whenever the user saves a mission specification, the compiler (⑤) automatically generates two files for each robot. One file encodes the set of LTL formulae that conform each local mission as an intermediate language (③). The second file contains a description of the local mission in natural English. This description proved to be useful for users who want to evaluate whether their concerns have been properly specified [7].

Once the user is satisfied with the specification they shall proceed to send it to the target robots as services (the IP of each robot must be specified) using a provided Java script. Deployed into each robot there is an interpreter (④), which receives the local mission and parses it. Finally, the interpreter uses the ROS environment to feed the underlying planner with tasks as requested. For instance, the interpreter of *r1* firstly feeds its planner with the task *patrol l1, l2, and l3*. If the robot detects a recognizable object, this task is stopped and a new task is sent to the planner (in this case, *perform action request help*). Once the task is finished, the interpreter requests the planner to resume the stopped default mission. The interpreter displays useful information for the user through a terminal. Examples of such information are the mission received, the task which is being executed by the robot, and detected events and whether they trigger any task of the local mission.

6 EXPERIMENTATION

We evaluate our DSL and implementation framework through experimentation, and with this aim, we integrated them into an existing platform (SERA) that provides different robotic functionalities (e.g., planning, motion control, self-localization). The experimentation serves us to evaluate the workflow described in Sec. 5 and the components shown in Fig. 3. We provide videos of our experiments with different robots and environments in a dedicated website.⁶

We also replicated missions which are available on literature, specifically from the RoboCup@Home'18 competition.⁷ We tested the restaurant-management scenario in simulation and in real environments. For the experimentation we used a Turtlebot2 robot in the offices of the University of Gothenburg and a TIAGo robot in the facilities of PAL Robotics (a robotics manufacturer). The repository with the implementation of our DSL (<https://github.com/SergioGarG/>

PROMISE_implementation) contains also the specification of two other missions from that edition of RoboCup@Home: the *dishwasher challenge* and the *tour guide*.

We evaluated our language simulating two other missions. The first is used as a running example in our previous work [7], and the second is represented in Fig. 1. Both comprise two robots performing collaborative and reactive tasks. It is worth to mention that we validated the simplicity and expressiveness of PROMISE and its implementation through two user studies (more details are available in our previous publication [7]).

7 CONCLUSION

In this paper, we present a software framework that supports end-users and domain experts in multi-robot mission specification and execution. The framework integrates different software components, including a language realized as a DSL, a compiler, and an interpreter. We strive to keep the balance between simplicity and expressiveness in our language. We report the usage of the framework and its components by listing the possible stakeholders that may use it and illustrating use cases. We also introduce a running example, which is specified using PROMISE's graphical syntax. Finally, we show real examples of mission specification and execution in real and simulated environments using PROMISE and its framework.

ACKNOWLEDGMENTS

Research supported by the EU H2020 Research and Innovation Progs. under GA No. 731869 and GA No. 694277 and the Centre of Excellence on Connected, Geo-Localized and Cybersecure Vehicle (EX-Emerge) funded by the Italian Government under CIPE resolution No. 70/2017 (Aug. 7, 2017).

REFERENCES

- [1] J. Bohren and S. Cousins. 2010. The SMACH high-level executive [ROS news]. *IEEE Robotics & Automation Magazine* (2010).
- [2] D. Bozhinoski, D. Di Ruscio, I. Malavolta, P. Pelliccione, and M. Tivoli. 2015. FLYAQ: Enabling Non-expert Users to Specify and Generate Missions of Autonomous Multicopters. In *ASE'15*.
- [3] Michele Colledanchise and Petter Ögren. 2018. *Behavior Trees in Robotics and AI: An Introduction*. CRC Press.
- [4] P. Doherty, F. Heintz, and J. Kvarnström. 2013. High-level mission specification and planning for collaborative unmanned aircraft systems using delegation. *Unmanned Systems* (2013).
- [5] P. Doherty, F. Heintz, and D. Landén. 2010. A distributed task specification language for mixed-initiative delegation. In *PRIMA'10*.
- [6] S. García, C. Menghi, P. Pelliccione, T. Berger, and R. Wohlrab. 2018. An Architecture for Decentralized, Collaborative, and Autonomous Robots. In *ICSA'18*.
- [7] S. García, P. Pelliccione, C. Menghi, T. Berger, and T. Bures. 2019. High-level mission specification for multiple robots. In *SLE'19*.
- [8] S. Maoz and J. Ringert. 2018. On the software engineering challenges of applying reactive synthesis to robotics. In *RoSE'18*.
- [9] C. Menghi, C. Tsiganos, P. Pelliccione, C. Ghezzi, and T. Berger. 2019. Specification Patterns for Robotic Missions. *TSE* (2019).
- [10] A. Nordmann, N. Hochgeschwender, and S. Wrede. 2014. A survey on domain-specific languages in robotics. In *SIMPAR'14*.
- [11] Douglas C. Schmidt. 2006. Guest Editor's Introduction: Model-Driven Engineering. *Computer* 39, 2 (2006).
- [12] DC Silva, PH Abreu, LP Reis, and E Oliveira. 2014. Development of a Flexible Language for Mission Description for Multi-robot Missions. *Information Sciences* (2014).
- [13] V. Ziparo, L. Iocchi, D. Nardi, P. F. Palamara, and H. Costelha. 2008. Petri net plans: a formal model for representation and execution of multi-robot plans. In *AAMAS'08*.

⁶<https://sites.google.com/view/promise-dsl/home>

⁷<http://www.robocupathome.org/rules>