

Chapter 7

Experimenting with Adaptation in Smart Cyber-Physical Systems: A Model Problem and Testbed



Vladimir Matena, Tomas Bures, Ilias Gerostathopoulos, and Petr Hnetynka

Abstract The chapter focuses on experimentation with adaptation in the field of smart cyber-physical systems (sCPS). In particular, it provides a model problem that features a coordination of autonomous cleaning robots. The model problem is accompanied with a testbed which allows the execution of the model problem along with custom adaptation logic. The testbed can be executed as a simulation of multiple robots running or deployed on an actual TurtleBot robot. Both the simulated and actual deployment environment are based on the same software stack. The offered simulation is precise timing-, bandwidth-, and mobility-aware and brings together a ROS-based Stage simulation of a swarm of robots and OMNeT++-based simulation of 802.15.4 wireless network, while the actual deployment is based on the TurtleBot robotic platform. The adaptation business logic is based on the DEECo component model and points to specific places, where the user code can be easily plugged in.

7.1 Introduction

Smart cyber-physical systems (sCPS) are distributed and decentralized systems that closely cooperate with their physical environment by sensing and actuating [9]. A characteristic feature of sCPS is that they exhibit a high level of “intelligence” in terms of opportunistic cooperation, dynamic self-organization, self-healing, and self-adaptation [6]. As such, sCPS are regarded as vital for building applications for smart mobility, smart energy grids, ambient assisted living, smart cities, etc.

V. Matena (✉) · T. Bures · P. Hnetynka
Faculty of Mathematics and Physics, Charles University in Prague, Prague, Czech Republic
e-mail: matena@d3s.mff.cuni.cz; bures@d3s.mff.cuni.cz; hnetynka@d3s.mff.cuni.cz

I. Gerostathopoulos
Faculty of Mathematics and Physics, Charles University in Prague, Prague, Czech Republic
Fakultät für Informatik, Technische Universität München, Munich, Germany
e-mail: iliassg@d3s.mff.cuni.cz

Software engineering of sCPS is largely an open challenge, as sCPS combine autonomous decentralized cooperative behavior, with concerns of real-time, limited communication, dependability, etc. The lack of software engineering support also applies to self-adaptation [7], which is a central feature of sCPS, crucial for coping with the uncertain environments in which sCPS operate.

While there is a large body of knowledge for experimenting with adaptation in the context of enterprise services and other traditional software systems, there is rather a vacuum in terms of knowledge and especially tools for experimenting with adaptation in sCPS. This is in our view because sCPS combine multiple relatively distinct disciplines (real-time, control, networking, agents, learning, data-analysis, etc.) [4]. This consequently requires engineering approaches and tools for sCPS to build synergies between the disciplines and support the mutual interplay of the concerns.

In this chapter,¹ we partially address the problem of development of self-adaptive sCPS by providing a model problem and testbed for experimenting with, comparing, and developing new adaptation solutions pertinent to sCPS.

In particular, the model problem and testbed provide challenges in coordination of autonomous robots with the interplay of concerns of (a) realistic communication (i.e., communication limited by bandwidth and subject to latencies), (b) real-time control, and (c) decentralized operation.

To enable fast prototyping, the testbed abstracts robots as autonomous components (implemented in Java) and allows describing robot communication via dynamic collaboration groups. It also points to specific places in the code where adaptation logic can be plugged in and provides metrics for evaluating the plugged-in adaptation. Thus, together, the model problem and the testbed provide a concrete ready-to-use benchmark for experiments in the relatively new field of sCPS.

Either the implementation can be executed as a simulation, or it can be directly deployed to actual robots (currently, the implementation out of the box supports the TurtleBot robots²).

The chapter is organized as follows. Section 7.2 describes the model problem in detail. Section 7.3 presents the testbed from both the user perspective and also implementation point of view. Section 7.4 describes a sample adaptation we have used for evaluating the testbed and further discusses lessons learned and limitations. Section 7.5 briefly details the structure of the provided testbed (detailed instructions are packaged together with the testbed). Section 7.6 discusses related work, while Sect. 7.7 concludes the chapter by summarizing the contributions.

¹It is based on material included in a SEAMS 2016 publication by the same authors [5].

²<http://www.turtlebot.com/>

7.2 Model Problem

The model problem provided by our testbed is the “Autonomous Cleaning Robots Coordination” (ACRC) problem. In ACRC, a number of cleaning robots is deployed in in-door space consisting of corridors and multiple office rooms (see Fig. 7.1).

Every robot is equipped with a camera which provides depth information. The robots use the cameras to observe obstacles (other robots, walls, etc.) and for navigation, by means of Adaptive Monte Carlo Localization (AMCL). Robots are equipped with a map of the place that they are supposed to clean. This map is used in the AMCL-based navigation, which works by comparing a depth scan with the map.

Robots are capable of limited communication using an IEEE 802.15.4 transceiver (with approx. 10 m direct visibility range), which allows building mobile ad hoc networks. This means that robots can exchange data only when they are close to one another. Robots can extend the communication range by acting as proxies that rebroadcast messages further. Generally, however, no global communication can be assumed as situations when no proxy is close enough or too much interference exists are rather often.

The basic software of the robots is formed by the Robot Operating System³ (ROS), which is the de facto standard set of libraries and services for building open-source robotic platforms.




Fig. 7.1 A visualization of the model problem

³<http://wiki.ros.org/>

7.2.1 *Operation and Adaptation Challenges*

Each robot is initially given its own set of places it is supposed to visit and clean. In the naïve solution, which can be considered as the baseline, robots act completely independently of one another (i.e., they do not communicate nor coordinate) and visit places on their list in the given order.

Due to the complexity of the environment and the deficiencies in the ROS stack (which we consider as a black-box component that is given and one has to live with), the naïve solution gives rise to multiple problems:

- A robot has only an approximation of its position and orientation. Often, especially when other robots are present nearby, the AMCL localization fails as the depth scans (which include other robots) cannot be matched with the known map. As a result, the robot navigation becomes very imprecise and sometimes, when in dense traffic, fails completely, and the robot stops.
- The navigation module in a robot sometimes fails to find a route to the destination because other robots moving by obstruct it. As the result the robot stops.
- Due to physical space constraints, robots often get to a deadlock situation – e.g., when one robot wants to enter the office and another wants to exit it. The result is again that the robots stop to avoid collision. (Note that this is a different situation to the previous point, where the failure to find a way is only transient. In this case, however, it persists until the deadlock is explicitly solved.)

Generally, each of these problems can be solved by pointing the robot to the right direction. However, it practically turns out to be quite difficult to (1) distinguish the cause of the problem and (2) to know where to navigate the robot to recover it from the failed state.

Though these problems could be targeted by modifying ROS, our experience with extending and customizing ROS shows that a more practically viable solution is to regard ROS as a black-box and build an adaptation layer over it. As such, the robotic scenario constitutes an excellent case for adaptation. (Of course, this is by no way a criticism of ROS, which itself is the most comprehensive open-source solution for robotics. It is more an acknowledgment of the complexity inherently connected with developing systems that perform in and interact with real environments.)

To remediate the deficiencies of the baseline solution, the adaptation layer has generally free access to the robot navigation. In particular, it can obtain estimates of the position and can sense whether the robot moves. Based on this, it can:

- manipulate the queue of locations to be visited (destinations),
- pause the robot and command the robot to move to any place on the map.

Additionally, the adaptation layer on one robot may communicate with the adaptation layers of other robots to realize more complex adaptation strategies via cooperation.

The adaptation however comes with another set of problems once we try not only to recover the robots from failures and deadlocks but also optimize the overall

performance of the system. Clearly, by reordering the locations to be visited and by transferring the responsibility of cleaning a place from one robot to another, the system can highly optimize itself. Theoretically, it can even get to a point when no collisions happen because robots exchange their destinations in such a way that they do not interfere. This is however subject to multiple problems, which can be regarded as additional adaptation challenges:

- The uncertainty in location makes planning not completely reliable.
- Communication range is limited, which means that robots in different rooms cannot communicate directly but only through proxies (if present), which have to be located in the corridor close to the office entrances.
- The communication is subject to latencies and unreliability (due to interference) which makes it impossible for a robot to have an up-to-date knowledge of the global state of the system and disallows strong synchronization among robots.

7.2.2 Solution Comparison Dimensions

Having the adaptation logic in place, various metrics can be considered for evaluation and comparison of different adaptation strategies (solutions to ACRC). We list below metrics which we found useful in our experiments with ROS-controlled robots. Note that since ACRC contains random elements and non-determinism, the evaluation of a solution requires multiple simulation runs of ACRC and statistical evaluation (e.g., by statistical testing of sample means or quantiles).

Time to complete all the tasks (i.e., visit and clean all locations assigned to the robots at the start) can be regarded as the basic metric when we assume that the evaluated solution is able to make the robots complete all their tasks. Our experience showed that this is more difficult than it appears to be. For evaluating partial successes, we thus suggest the following metrics.

Number of cleaning tasks that were completed This covers situations when time limit for completion expires or when the system itself realizes that certain locations cannot be cleaned – e.g., if a robot gets stuck in a room entrance and any attempts to move the robot out of the way fail.

Total running time till system completes or gives up This can be used as a metric complementary to the above one, to reward solutions which possess the ability to recognize that certain problems cannot be solved. It can serve to resolve ties in case two solutions are statistically similar (e.g., a statistical test cannot reject the hypothesis of the two solutions that have the same average number of cleaning tasks completed).

7.3 Testbed

The provided testbed allows for easy experimentation with adaptation techniques and algorithms for the ACRC problem. The model problem is implemented on top of ROS, and both the adaptation logic and the adapted system are specified using DEECo, which is a component model for sCPS. Details on the technical architecture are given in Sect. 7.3.5.

The testbed offers two modes of deployment and execution. The first mode is a simulation of a swarm of TurtleBots solving the ACRC problem. This mode is primarily suitable for early stages of development and/or for conducting quantitative measurements. The second mode allows for actual deployment using real TurtleBots.

The *simulation mode* is implemented on top of the Stage simulator, which is tightly integrated with ROS. The Stage simulator is capable of simulating robot physics, movement, laser scan sensing, and odometry readings. Currently the included Stage is configured to simulate TurtleBots only, but it is possible to change robot shape, movement model, and sensors to match different robots as well.

The Stage simulator is further extended by a custom integration of the OMNeT++ network simulation into ROS – called ROSOMNeT++⁴ – which enables sending and receiving IEEE 802.15.4 packets using ROS facilities.

For the *actual deployment*, it is necessary to equip each TurtleBot with an onboard computer and wireless network interface. Regarding the onboard computer, any average contemporary machine is suitable (we tested it with the Intel P9600 CPU and 6 GB of RAM, and such a configuration was completely sufficient). For the wireless network interface, an external microcontroller with an IEEE 802.15.4 module is expected. The testbed has been tested with and is prepared for the STM32F4⁵ board equipped with the extension board⁶ and the BEE click⁷ module. All of them are off-the-shelf components.

ROS already contains modules, which serve as drivers for TurtleBot, and we have developed extensions to support the IEEE 802.15.4 network. In particular, we have developed two projects. The beeclickarm⁸ provides the firmware and Java interface for the used microcontroller, while the beeclickarmROS⁹ exports features of the beeclickarm as ROS topics and services.

The detailed instructions about hardware installation and deployment are available in the testbed's README file.¹⁰

⁴<https://github.com/d3scomp/ROSOMNeT>

⁵<http://www.st.com/stm32f4>

⁶<http://www.mikroe.com/stm32/stm32f4-discovery-shield>

⁷<http://www.mikroe.com/click/bee>

⁸<https://github.com/d3scomp/beeclickarm/tree/robot-additions>

⁹<https://github.com/d3scomp/beeclickarmROS>

¹⁰<https://github.com/d3scomp/deeco-adaptation-testbed>

The testbed seamlessly supports both deployment modes; the simulated and actual devices are accessible via the same ROS interface, and thus no changes at user code are required when switching between the deployment modes.

7.3.1 Modeling Concepts for Decentralized Coordination

The robots' behavior is developed using DEECo [2], which is a component model and framework for developing complex sCPS. DEECo is based on concepts of *ensemble-based component systems* (EBCS) (designed primarily in the scope of the EU FP7 ASCENS project¹¹). In EBCS, a system is modeled as a set of dynamic cooperation groups of software components – *ensembles*. DEECo itself is an abstract component model; however it comes with two implementations – one in Java¹² (JDEECo) and one in C++¹³ (CDEECo). In the testbed, we use JDEECo as we found Java easier for prototyping the components.

A component in DEECo is represented by its data (*knowledge* in EBCS) and its tasks (*processes* in EBCS). Figure 7.2 shows a code skeleton of the baseline implementation of the robot component in JDEECo. JDEECo-specific constructs are expressed using an internal domain-specific language (DSL) defined via Java annotations. A component is defined as a plain Java class annotated with the `@Component` annotation. Component's knowledge is defined as Java class fields (lines 3–10 in Fig. 7.2). Knowledge that is not supposed to be shared with other components via ensembles (as described below) is marked as `@Local`. Component's fields are manipulated by the component's processes (e.g., lines 13–33). Processes are defined as, respectively, annotated static Java class methods. Processes are either periodically executed or event-triggered (i.e., commonly as a reaction to knowledge change). This is determined by the annotation attached to the process.

Typically, processes involve sensing, computation, mutation of the component knowledge fields, and actuating. The signature of the process defines which knowledge fields are read/written (as in/out/in-out). Technically, the processes are scheduled by JDEECo runtime, which also takes care of thread-safe retrieval of component's knowledge to be used by a process and storing of the process results back in component's knowledge.

Figure 7.2 lists the processes defined in the baseline implementation of the cleaner robot as provided by ACRC. These are (i) setting the next destination, (ii) reading the position, (iii) reporting the status, and (iv) controlling the movement of the robot.

Communication between components is in DEECo modeled by *ensembles*. An ensemble dynamically determines which components are in the communication

¹¹<http://ascens-ist.eu/>

¹²<http://github.com/d3scomp/JDEECo>

¹³<http://github.com/d3scomp/CDEECo>

```

1. @Component
2. public class CleanerRobot {
3.     public String id;
4.     public Position destination;
5.     public Position position;
6.     public State state;
7.     @Local public Long blockedCounter;
8.     @Local public Long noPosChangeCounter;
9.     @Local public Position oldPosition;
10.    @Local public List<Position> route;
11.    ...
12.
13.    @Process
14.    @PeriodicScheduling(period = 500)
15.    public static void setDestination(
16.        @InOut("destination")
17.        ParamHolder<Position> destination,...)
18.    {...}
19.
20.    @Process
21.    @PeriodicScheduling(period = 100)
22.    public static void sense( @Out("position")
23.        ParamHolder<Position> position,
24.        @In("positioning") Positioning positioning)
25.    {...}
26.
27.    @Process
28.    @PeriodicScheduling(period = 1000)
29.    public static void reportStatus( @In("id")
30.        String id, ...)
31.    {...}
32.
33.    @Process
34.    @PeriodicScheduling(period = 2000)
35.    public static void driveRobot(
36.        @In("position") Position pos,
37.        @In("positioning") Positioning positioning,
38.        @In("destination") Position destination,
39.        @InOut("curDestination")
40.        ParamHolder<Position> curDestination,...)
41.    {...}
42. }

```

Fig. 7.2 Model of ACRC baseline in JDEEC0

group via a membership condition. Topologically, an ensemble in JDEEC0 is a star featuring one *coordinator* and multiple *members*. The communication within ensembles is implicit, i.e., the ensemble defines an exchange method, which performs knowledge exchange among components grouped in the ensemble (i.e., copying data from a knowledge field of one component to a knowledge field of another component).




Fig. 7.3 Adaptation architecture

The baseline implementation does not involve any ensembles. However, ensembles are to be exploited for decentralized coordination of adaptation across several robots. This is demonstrated in Fig. 7.12, where an ensemble for location exchange is given. It is established between robots which are close to each other, and both of them are stuck. The ensembles are defined again as plain Java classes with annotations. The membership and exchange methods are periodically executed (with prescribed period – line 2), and their parameters specify the read/written knowledge fields of particular components (prefixes *coord-* and *mbr-* are used to identify coordinator and member role, respectively).

The architectural view of the adaptation is depicted in Fig. 7.3, which shows the split into “adapted” and “adaptation” layers. The adapted layer consists of robot drivers, ROS modules, and business logic implemented as DEECo processes. The adaptation layer is implemented by DEECo constructs. In case of local adaptations, which do not take multiple robots into account, the adaptation is implemented as a DEECo process. In more advanced cases when the adaptation layers of multiple robots need to cooperate, a DEECo ensemble is used to implement the adaptation logic.

7.3.2 Setup

The testbed models the ACRC problem via DEECo component model (in detail in Sect. 7.3.1). In particular, it represents each robot as an instance of the *CleanerRobot* component and provides its baseline behavior (i.e., the base-level subsystem [13])

```

1. ROSSimulation ros = new ROSSimulation(SIM_SRV_ADDRESS,
2.     11311, SIM_SRV_ADDRESS, "corridor", 0.02, 100);
3. DEECoSimulation realm = new DEECoSimulation(ros.getTimer());

4. realm.addPlugin(Network.class);
5. realm.addPlugin(DefaultKnowledgePublisher.class);
6. realm.addPlugin(KnowledgeInsertingStrategy.class);
7. realm.addPlugin(BeeClick.class);

8. for (int i = 0; i < NUM_ROBOTS; ++i) {
9.     final String name = "Collector" + i;
10. List<Position> garbage = generateGarbage(NUM_GARBAGE);
11.     Positioning positioning = new Positioning();
12.     DEECoNode robot = realm.createNode(i, positioning,
13.         ros.createROSServices(colors[i]), positionPlug[i]);
14.     robot.deployComponent(new CleanerRobot(name, positioning,
15.         ros.getTimer(), garbage,...));
16.     robot.deployEnsemble(DestAdoptionEnsemble.class);
17.     robot.deployEnsemble(AdoptedDestRemoveEnsemble.class);
18. }

19. realm.start(600_000);

```

Fig. 7.4 Deploying multiple robots in simulation

in Java. The testbed provides a well-defined place in the component where the adaptation logic is to be plugged in (i.e., the reflective subsystem). Technically, this is done by introducing additional periodic processes to the Collector Robot component and additional ensemble specifications (e.g., see Sect. 7.4.1). There is no difference between the setup for the simulator deployment and the actual TurtleBot deployment; only the initialization and launching differ as follows.

Simulation setup In Fig. 7.4 the code responsible for initializing the simulation is shown. Lines 1–3 establish DEECo simulation using ROS, lines 4–7 load DEECo plug-ins shared by all robots, the loop on lines 8–18 deploys robots, and finally line 19 runs simulation for 600 s. The simulation is configured by the number of robots, their initial positions, and the map of the environment. The testbed comes with one map that comprises a corridor and two offices. Custom maps can be provided as PNG files similar to the one shown in Fig. 7.1.

Actual TurtleBot setup Instead of deploying all the robots at once (as in the case of simulation), the deployment code depicted in Fig. 7.5 deploys a single cleaning robot component on a single actual robot. Thus, it is necessary to run the code on each individual robot. Lines 1–4 are responsible for establishing the DEECo system using wall timer and actual robot ROS interface, lines 6–11 deploy a DEECo node with all required plug-ins, lines 13–20 are responsible for deployment of the cleaning robot component and ensembles, and finally line 22 runs the system and blocks forever as the real deployment has no time limit.

```
1. WallTimeTimer wallTimer = new WallTimeTimer();
2. RosServices rosServices = new RosServices(
3.     System.getenv("ROS_MASTER_URI"),
4.     InetAddress.getLocalHost().getHostName());
5.
6. DEECoNode node = new DEECoNode(ROBOT_ID, wallTimer,
7.     new Network(),
8.     new BeeClick(),
9.     new DefaultKnowledgePublisher(),
10.    new KnowledgeInsertingStrategy(),
11.    rosServices, positionPlugin);
12.
13. final String name = "Collector" + ROBOT_ID;
14. garbage = generateGarbage();
15. Positioning positioning = new Positioning();
16. node.deployComponent(new CleanerRobot(name, positioning,
17.    wallTimer, garbage,...));
18.
19. node.deployEnsemble(DestinationAdoptionEnsemble.class);
20. node.deployEnsemble(AdoptedDestinationRemoveEnsemble.class);
21.
22. wallTimer.start();
```

Fig. 7.5 Single actual robot deployment

Further guidelines on deployment as well as the whole source code of the testbed can be found on GitHub.¹⁴

7.3.3 Debugging

The testbed enables usage of several debugging tools that can be used to observe the system in order to deploy adaptation techniques as well as debug existing adaptation code. The testbed is using ROS topics to control the simulated robot using standard messages described by ROS. Thus it is possible to use ROS tools to inspect and visualize messages in the system at no extra cost. These can be used to obtain a robot-centric view of the system (as displayed in Fig. 7.6) and thus realize what is wrong at a local level. In the following paragraphs, the most important tools are briefly described.

Stage visualization The primary output of the testbed is the direct visualization of the scenario shown in Fig. 7.1 (the visualizer itself comes with the Stage robot simulator – Sect. 7.3.5). Via it, the user can observe the movement of the robots in real time. Black lines represent walls and other obstacles impenetrable for the

¹⁴<https://github.com/d3scomp/deeco-adaptation-testbed>



Fig. 7.6 Robots' perception of the environment

robots (i.e., the map provided to the testbed). The colored dots represent the robots as located in the simulated system. Thus the output of the Stage visualization is global view of the systems' ground truth data.

Logging ROS messages As mentioned above robot control is ROS based on sending messages. Fortunately those can be printed to command-line or a file for later processing by plotting or using statistical tools. ROS defines how different datatypes are represented as text, thus printing robot location requires no extra output formatting as shown in Fig. 7.7.

Robot Visualizer (RViz)¹⁵ This tool provides a convenient way to observe a robot's view of the environment by displaying ROS messages in a 2D or 3D. Most of the messages used in the system are directly understood by RViz, so that having data visualized is as easy as choosing the correct data source.

The real power of RViz is the visualization of the data from real robot. Figure 7.8 shows RViz visualization of data from the TurtleBot deployed in a real environment. The background is a static map of the environment which is used for long-range planning. The colored rectangle around the robot is a local map capturing temporary obstacles detected by distance scanner such as chairs and persons. Below the 3D model of the robot, a cloud of green arrows used by AMCL to guess robot location

¹⁵<http://wiki.ros.org/rviz>

Fig. 7.7 Printing ROS location message

```

1. $ rostopic echo /robot_0/amcl_pose
2. header:
3.   seq: 78
4.   stamp:
5.     secs: 91
6.     nsecs: 900000000
7.   frame_id: /map
8. pose:
9.   pose:
10.    position:
11.     x: 13.901734935
12.     y: 11.8805620695
13.     z: 0.0
14.    orientation:
15.     x: 0.0
16.     y: 0.0
17.     z: 0.955308313164
18.     w: 0.29561127651

```




Fig. 7.8 RViz using data from real robot

is visible. Finally a depth image captured by onboard camera is displayed in 3D in order to help guess how guessed location matches reality.

rqt_plot¹⁶ Working on top of ROS messages, an `rqt_plot` tool can generate plots of various messages in real-time and store them for later use. The output of this tool is depicted in Fig. 7.9.

¹⁶http://wiki.ros.org/rqt_plot

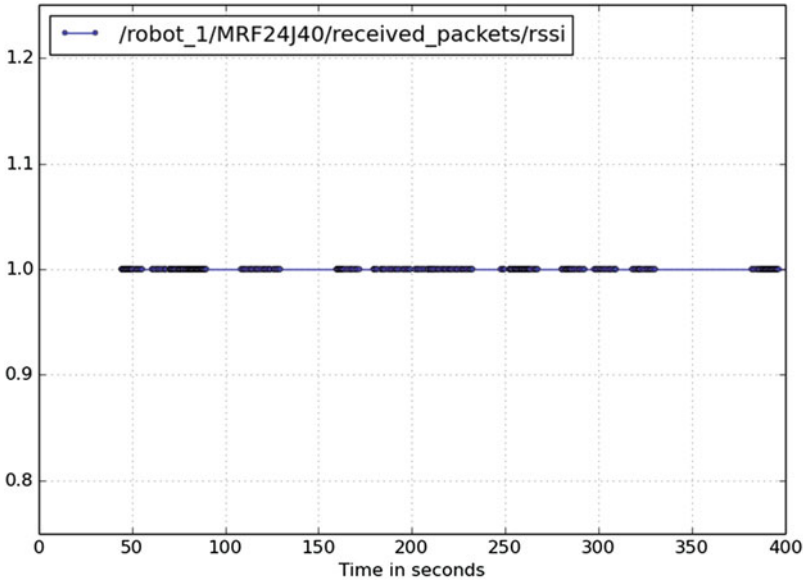


Fig. 7.9 rqt_plot showing MANET packet arrivals

ROS Bags¹⁷ ROS has an ability to record all messages in the system into a file, which can be used for offline analysis. All the aforementioned tools using ROS messages can work on top of replayed messages recorded during simulation or actual system execution. For instance, it is possible to visualize trajectories of the robots and replay the visualization over and over.


This feature is important for recording simulation runs as it saves time needed to execute the same simulation repeatedly. It is even more important for the actual deployments as it is in fact impossible to execute a scenario repeatedly with the exactly same results.

Eclipse debugger¹⁸ As the testbed and all the adaptation code are written in Java, it is possible to run the testbed as a Java application directly from Eclipse IDE and thus use all debugging features of Eclipse. This is possible for both the simulation and actual run. The limitation here (stemming from the soft real-time nature of ROS) is that ROS continues running even if jDEECo and the adaptation logic are paused by debugging. However, this is typically not a problem due to the fact that jDEECo controls ROS essentially only by setting robot waypoints. As such, if jDEECo is paused, the robot only continues to its next waypoint or stops sooner if there seems to be an obstacle preventing its move and then it waits for the adaptation logic to instruct it further.

¹⁷<http://wiki.ros.org/rosbag>

¹⁸<https://eclipse.org/>

Fig. 7.10 Boxplots of results from 10 experiment runs



7.3.4 Obtaining Results

The testbed comes with a script which computes statistics of the evaluation from the logs collected in multiple simulation runs. It generates boxplots of the results for the last two metrics defined in Sect. 7.2.2 (as in Fig. 7.10).

7.3.5 Technical Architecture

Figure 7.11 shows the architecture of the testbed. Technically, it is a merger of four main existing modules. The contribution of the testbed lies in properly configuring them and bridging them by glue and synchronization code. The modules are:

- ROS Core – this module provides publish/subscribe middleware for robotic systems and the basic software of the robot. In particular, it implements the AMCL localization, navigation, and low-level movement control of the robot. The messaging system is used to interconnect robot basic software as well as to connect remaining modules described later.
- OMNeT++¹⁹ – it is a network simulator. It runs independently of ROS. We have implemented a bridge between ROS Core publish/subscribe mechanism

¹⁹<http://omnetpp.org/>

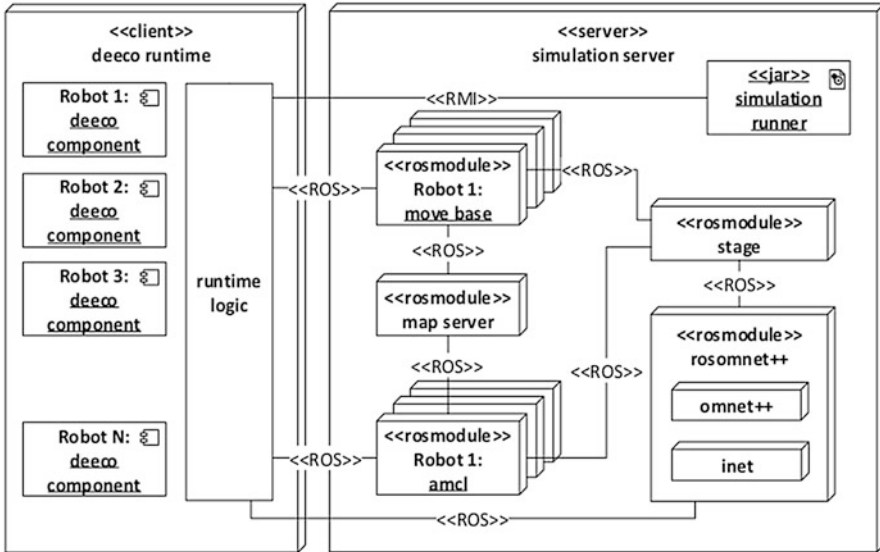


Fig. 7.11 Testbed deployment diagram

and OMNeT++, which exposes the MANET transceiver as a ROS topic. This allows modules connected to ROS to communicate. OMNeT++ simulates the latency, physical range, and interference of the communication based on robots' positions.

- Stage²⁰ – it is a robot simulator, which controls the simulation. It connects to ROS Core and simulates sensors and actuators of the robot given the simulated robot position and the map of the environment. Robots sensors and actuators are exported as ROS topics. The interface of simulated robot is the same as the interface of the real TurtleBot. The only difference is the usage of namespaces which enable deployment of multiple simulated robots into one ROS system.
- JDEECO – it provides the component abstraction and concepts for decentralized coordination as described in Sect. 7.3.1. It abstracts ROS topics on location and navigation and exposes them to DEECO components to allow for adaptation. It further exploits the ROS topic on MANET-based communication (backed by OMNeT++) to implement inter-component communication via ensembles. JDEECO again runs independently of ROS and is synchronized with it by a bridge that we have developed as part of the testbed.

²⁰<http://playerstage.sourceforge.net/>

7.4 Evaluation

7.4.1 Example Adaptation Logic

We complement the model problem specification and the testbed with an example adaptation logic as part of the model problem. It provides a comprehensive example of the modeling concepts (described in Sect. 7.3.1) and also serves as evaluation of the testbed to perform simulation of physical, mobility, networking, and coordination concerns.

In the example adaptation, we tackle the problems described in Sect. 7.2.1 in the following way:

1. We introduce a process (on each robot), which periodically detects the situation when a robot is stuck. This is done by checking whether the robot is moving and whether the robot has a destination set. The robot that is not moving and wants to move is considered stuck.
2. If a robot is detected to be stuck, we select a random location from its queue of destinations and set it as its current one. This resets the navigation module in the robot and typically gets the robot to move. We monitor the outcome via the process described in (1) and repeat if no visible outcome is detected.
3. If another robot is stuck in close proximity (up to 1.5 m), we establish an ensemble with it. Within the ensemble, one robot adopts the current destination of the other robot and vice versa. This solves the (deadlock) situations when two robots meet in the office entrance and cannot proceed.

Strategy (3) is illustrated in Fig. 7.12. Ensemble membership is defined on lines 7–13, and destination adoption is defined on lines 20–24.

7.4.2 Lessons Learned and Limitations of the Testbed

The experience with development of the testbed on top of ROS led us to several observations, which we believe are of general interest. We thus share them here.

Generally, a relatively big surprise was the overall immaturity of the frameworks. This most likely stems from the fact that ROS is primarily used as a platform for controlling a single robot at real time. Though it has very flexible architecture, which allows running multiple robots within a single ROS system and allows connecting different environment simulators (e.g., Stage), the practice shows that these setups work out of the box only for trivial examples. Deploying multiple robots without careful configuration of the environment would make ROS or the Stage simulator crash. Similar story applies for OMNeT++, which is a mature and production-ready network simulator used in many applications. Nevertheless, when it comes to complex exercising of the mobile ad hoc network, the simulator again becomes very fragile, and without careful configuration and patching, it crashes for no obvious reason. From this perspective, we believe that even without the DEECo abstraction

```

1. @Ensemble
2. @PeriodicScheduling(period = 3000)
3. public class DestinationAdoptionEnsemble {
4.     double MAX_DIST_M = 3.0;
5.     long BACKOFF_MS = 10000;
6.
7.     @Membership
8.     bool membership(@In("coord.id") String
coordId, ...) {
9.         return coordState == Block &&
10.             mbrState == Block &&
11.             !coordId.equals(mbrId) &&
12.             coordPos.distTo(mbrPos) < MAX_DIST_M;
13.     }
14.
15.     @KnowledgeExchange
16.     void exchng(@InOut("mbr.destination") destination, ...) {
17.         if (now-lastAdoption.value) < BACKOFF_MS) {
18.             return;
19.         }
20.         mbrAdoptedDestinations.value.add(coordDestination);
21.         mbrDestination.value = coordDestination;
22.         mbrRoute.value.add(coordDestination);
23.         mbrBlockCnt.value = 01;
24.         lastAdoption.value = now;
25.     }
26. }

```

Fig. 7.12 Excerpt from example ACRC adaptation strategy

layer, the pre-configured testbed we provide can save a couple of months of painful debugging.

Another class of problems comes from the fact that though ROS has been used in simulations, it is not a discrete event simulator. It consists of a number of modules, which just run in wall-clock time. This means that (1) the simulation is non-deterministic and (2), if extra care is not taken, the system crashes because the simulator, ROS, OMNeT++, and DEECo are not synchronized. We have solved this problem by introducing explicit synchronization at critical places, but still one has to keep in mind that this solution does not result in fully deterministic simulations.

Surprisingly enough, our experience with developing the sample adaptation logic has shown that the wall-clock-timed simulation has certain advantages over a standard off-line discrete-event simulation. Since the system is live (and behaves as if the robots were moving in real time), one can watch the system as it runs, inspect the laser scans, etc. Additionally, it is possible to modify the system while it is running – e.g., a robot can be dragged by mouse to another location. While this is not important in classical batch simulations which focus on statistical comparison of different algorithms, it is very useful in debugging and especially in prototyping (which in fact is one of the primary goals of our testbed and the reason why we equipped the testbed with DEECo abstractions).

7.5 Testbed Structure

The complete testbed is available at http://d3s.mff.cuni.cz/projects/components_and_services/deeco/files/deeco-adaptation-testbed.zip. It contains the source code of the testbed, together with installation and usage instructions. Moreover, a pre-configured virtual machine image is included in order to enable rapid hands-on experience without the hassle of installing tons of libraries.

7.6 Related Work

In this section, we briefly review three model problems/exemplars that have been contributed to the self-adaptive systems community repository [8]. This is an ongoing effort to provide benchmarks to evaluate new techniques against the state of the art, a popular strategy in other communities such as performance engineering (DaCapo suite [1], SPEC benchmarks [11], Cloud Efficiency Metric [10]).

The automated traffic routing problem (ATRP) [14] is a model problem that can be used as benchmark for the evaluation of different self-adaptation mechanisms. ATRP features cars traveling on a map. Each car has a specific starting point, a specific destination, and a specific starting time. Each car has the goal to reach its destination by traversing the map while respecting the speed limits on the streets. Problems arise due to conflicts between individual goals (e.g., all cars select the same street resulting in traffic congestion on the street), traffic accidents, and road closures. ATRP can have solutions that are fundamentally different ranging from centralized to completely decentralized ones and generating answers that are optimal or suboptimal. These solutions can be compared w.r.t. dimensions such as scalability, answer quality, robustness to sensor uncertainty, etc.

To evaluate and compare ATRP solutions, ADASIM has been proposed [14]. ADASIM is a Java-based discrete-event simulator that simulates the execution of an ATRP solution on an ATRP instance. It provides configuration files for specifying the problem instance, built-in routing algorithms, traffic delay functions, filters for introducing measurement uncertainty, and Java interfaces that can be implemented to specify an ATRP solution. An event logging and analysis infrastructure is also provided. In summary, ATRP provides a vehicle suitable for experimentation with different self-adaptation strategies that try to resolve conflicts between goals of individual agents, prioritize between nonfunctional properties, and provide robustness to faults. Similarly, our model problem and testbed stand as a benchmark for self-adaptation mechanisms but focus more on run-time uncertainty and unreliable communication and coordination in sCPS.

Znn.com [3] is another model problem for self-adaptation. Znn.com is a news service that serves multimedia content to its customers. It is realized by a classical N-tier client-server architecture. The objective of Znn.com is to serve content while optimizing operating costs and keeping the response time bounded. Self-adaptation

capabilities are needed in order for Znn.com to react to spikes on user load or other external changes while satisfying its objectives. In such cases, Znn.com can choose from a limited number of predesigned adaptation decisions, e.g., switching the content served from multimedia to textual or incrementing the server pool size. While Znn.com is primarily suitable for comparing self-optimization solutions, our model problem and exemplar are more suitable for comparing solutions that focus on self-healing and survivability (robot unblocking, deadlock resolution) properties of sCPS.

Tele Assistance System (TAS) [12] is an exemplar for self-adaptation in the area of service-based systems. TAS aims to aid patient suffering from chronic conditions via tele-assistance. It is realized by wearable sensors measuring vital parameters and three remote services for data analysis, medication delivery, and ambulance dispatching in case of emergency. TAS comes with a number of generic adaptation scenarios. Each scenario consists of the type of uncertainty that warrants self-adaptation (e.g., service failure), appropriate self-adaptation actions (e.g., switching to equivalent service), and type of quality attributes (QoS) impacted (e.g., cost). For measuring the satisfaction level of each QoS, respective metrics are specified. A reference implementation of TAS [12] provides a convenient way of comparing self-adaptation solutions w.r.t. user-specified requirements in user-specified settings (instances of the generic TAS scenarios) by simulating them and analyzing the results with built-in graphical tools. While an excellent representative exemplar for self-adaptive systems, TAS focuses specifically on service-based systems, not CPS.

7.7 Summary

Responding to the pressing need for model problems and testbeds to evaluate the research ideas in the area of self-adaptive smart cyber-physical systems (sCPS), we have presented ACRC, a model problem in the realm of sCPS that lends itself to a number of self-adaptation techniques that increase its self-healing, survivability, and self-optimization properties. It facilitates the process of trying out and comparing self-adaptation solutions to this problem. Our pre-configured testbed provides a starting point for experimental research. Moreover, the experiments can be easily extended to actual robots as the simulation shares interface with off-the-shelf robotic platform. We hope that ACRC will help increase the awareness of the yet-to-be-addressed challenges in the exciting field of self-adaptive sCPS and drive further advances in the field.

Acknowledgements The work on this paper has been partially supported by the Charles University Grant Agency project No. 391115 and Charles University institutional funding SVV-2016-260331. This work is part of the TUM Living Lab Connected Mobility (TUM LLCM) project and has been funded by the Bayerisches Staatsministerium für Wirtschaft und Medien, Energie und Technologie (StMWi).

References

1. Blackburn, S.M., Garner, R., Hoffmann, C., Khang, A.M., McKinley, K.S., Bentzur, R., Diwan, A., Feinberg, D., Frampton, D., Guyer, S.Z., Hirzel, M., Hosking, A., Jump, M., Lee, H., Moss, J.E.B., Phansalkar, A., Stefanović, D., VanDrunen, T., von Dincklage, D., Wiedermann, B.: The dacapo benchmarks: java benchmarking development and analysis. *SIGPLAN Not.* **41**(10), 169–190 (2006). <https://doi.org/10.1145/1167515.1167488>
2. Bures, T., Gerostathopoulos, I., Hnetyinka, P., Keznikl, J., Kit, M., Plasil, F.: Deeco: an ensemble-based component system. In: *Proceedings of the 16th International ACM Sigsoft Symposium on Component-Based Software Engineering, CBSE'13*, pp. 81–90. ACM, New York (2013). <https://doi.org/10.1145/2465449.2465462>
3. Cheng, S.W., Schmerl, B.: Znn model problem. <http://self-adaptive.org/exemplars/model-problem-znn-com>. Accessed 22 Jan 2016
4. Kim, K., Kumar, P.R.: Cyber-physical systems: a perspective at the centennial. *Proc. IEEE* **100**(Centennial-Issue), 1287–1308 (2012). <https://doi.org/10.1109/JPROC.2012.2189792>
5. Matena, V., Bures, T., Gerostathopoulos, I., Hnetyinka, P.: Model problem and testbed for experiments with adaptation in smart cyber-physical systems. In: *Proceedings of the 11th International Symposium on Software Engineering for Adaptive and Self-Managing Systems, SEAMS'16*, pp. 82–88. ACM, New York (2016). <https://doi.org/10.1145/2897053.2897065>
6. NIST: Cyber physical systems: situation analysis of current trends, technologies, and challenges. In: *NIST CPS Workshop* (2012). http://events.energetics.com/NIST-CPSWorkshop/pdfs/CPS_Situation_Analysis.pdf
7. Salehie, M., Tahvildari, L.: Self-adaptive software: landscape and research challenges. *ACM Trans. Auton. Adapt. Syst.* **4**(2), 14:1–14:42 (2009). <https://doi.org/10.1145/1516533.1516538>
8. Self-adaptive systems community repository. <http://self-adaptive.org/exemplars>. Accessed 22 Jan 2016
9. Sha, L., Gopalakrishnan, S., Liu, X., Wang, Q.: Cyber-physical systems: a new frontier. In: Singhal, M., Serugendo, G.D.M., Tsai, J.J.P., Lee, W., Römer, K., Tseng, Y., Hsiao, H.C.W. (eds.) *IEEE International Conference on Sensor Networks, Ubiquitous, and Trustworthy Computing (SUTC 2008)*, Taichung, 11–13 June 2008, pp. 1–9. IEEE Computer Society (2008). <https://doi.org/10.1109/SUTC.2008.85>
10. Shtern, M., Smit, M., Simmons, B., Litoiu, M.: A runtime cloud efficiency software quality metric. In: *Companion Proceedings of the 36th International Conference on Software Engineering, ICSE Companion 2014*, pp. 416–419. ACM, New York (2014). <https://doi.org/10.1145/2591062.2591127>
11. Spec benchmarks. <http://www.spec.org/benchmarks.html>. Accessed 22 Jan 2016
12. Weyns, D., Calinescu, R.: Tele assistance: a self-adaptive service-based system exemplar. In: *Proceedings of the 10th International Symposium on Software Engineering for Adaptive and Self-Managing Systems, SEAMS'15*, pp. 88–92. IEEE Press, Piscataway (2015). <http://dl.acm.org/citation.cfm?id=2821357.2821373>
13. Weyns, D., Malek, S., Andersson, J.: Forms: a formal reference model for self-adaptation. In: *Proceedings of the 7th International Conference on Autonomic Computing, ICAC'10*, pp. 205–214. ACM, New York (2010). <https://doi.org/10.1145/1809049.1809078>
14. Wuttke, J., Brun, Y., Gorla, A., Ramaswamy, J.: Traffic routing for evaluating self-adaptation. In: Müller, H.A., Baresi, L. (eds.) *7th International Symposium on Software Engineering for Adaptive and Self-Managing Systems, SEAMS, Zurich*, 4–5 June 2012, pp. 27–32. IEEE Computer Society (2012). <https://doi.org/10.1109/SEAMS.2012.6224388>