

RoboChart: a State-Machine Notation for Modelling and Verification of Mobile and Autonomous Robots

Alvaro Miyazawa¹, Pedro Ribeiro², Wei Li³,
Ana Cavalcanti⁴, Jon Timmis⁵, and Jim Woodcock⁶

¹alvaro.miyazawa@york.ac.uk

²pedro.ribeiro@york.ac.uk

³wei.li@york.ac.uk

⁴ana.cavalcanti@york.ac.uk

⁵jon.timmis@york.ac.uk

⁶jim.woodcock@york.ac.uk

^{1,2,4,6}Department of Computer Science, University of York, York, YO10 5GH, UK

^{3,5}Department of Electronics, University of York, York, YO10 5DD, UK

June 10, 2016

Abstract

Autonomous and mobile robots are becoming ubiquitous. From domestic robotic vacuum cleaners to driverless cars, such robots interact with their environment and humans, leading to potential safety hazards. We propose a three-pronged solution to the problem of safety of mobile and autonomous robots: (1) domain-specific modelling with a formal underpinning; (2) automatic generation of sound simulations; and (3) verification based on model checking and theorem proving. Here, we report on a UML-like notation called RoboChart, designed specifically for modelling autonomous and mobile robots, and including timed and probabilistic primitives. We discuss a denotational semantics for a core subset of RoboChart, an approach for the development of sound simulations, and an implementation of RoboChart and its formal semantics as an Eclipse plugin supported by the CSP model checker FDR.

1 Introduction

The current practice of programming mobile and autonomous robots does not reflect the modern outlook of their applications. Such practice is often based on standard state machines, without formal semantics, to describe the robot controller only, with time and probabilistic properties discussed in natural language. In the design stage, the state machine guides the development of a simulation, but no rigorous connection between them is established.

In this report, we present a state-machine based notation, called RoboChart, for the specification and design of robotic systems. Besides state machines, RoboChart includes elements to organise specifications, fostering reuse and taming complexity. These extra constructs embed the notions of robotic platforms and their controllers; communication between controllers can be synchronous or asynchronous. The state-machine notation is fully specified, including an action language and constructs to specify timing and probabilistic properties. Operations used in a state machine can be taken from a domain-specific API or defined by other state machines; communication between state machines inside a controller is synchronous. Operations can be given pre and postconditions.

The time primitives of RoboChart allow time budgets and deadlines to be specified for operations and events directly as part of a state machine. Constraints can be specified in association with the relative-time elapsed since the occurrence of events or the entering of states. Our time primitives are inspired by constructs of timed automata [5] and Timed CSP [31].

RoboChart also includes probabilistic transitions suggested for UML state machines in [18], with a semantics based on Markov Decision Processes. Probabilistic state machines are used in many robotic applications, such as adaptive foraging and swarming behaviour [22]. Probability in [18] is defined over actions, and the resulting state machines can be analysed using PRISM, a well-established model checker for probabilistic automata [21].

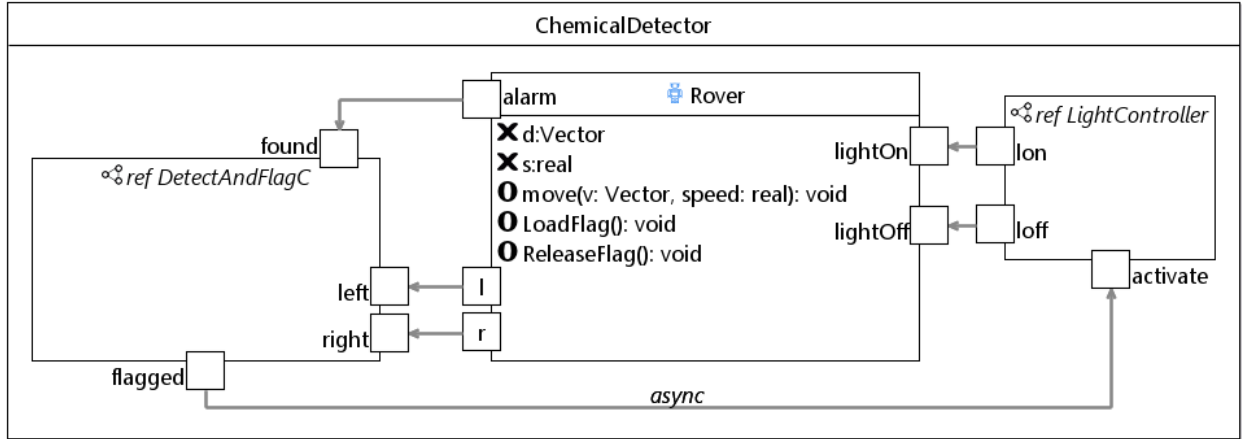


Figure 1: Chemical Detector

UML [25] state machines are popular. RoboChart, however, is customised to make it suitable for verification and automatic generation of simulations.

In this paper, we formalise the semantics of the core constructs of RoboChart using CSP [29]. Importantly, CSP is a front end for a mathematical model that supports model-checking and theorem proving, namely, Hoare and He’s Unifying Theories of Programming [17] (UTP). Use of CSP enables model checking with FDR [14]. On the other hand, the underlying UTP model makes our core semantics adequate for extension to deal with time [34] and probability [37].

RoboChart and its semantics can also be used for the generation of sound simulations that can shed light on the actual behaviours of the robots within various configurations and environment. We describe a general architecture for simulations that can be automatically generated from RoboChart models.

Finally, we describe RoboTool, which provides support for modelling using the RoboChart graphical notation (and its optional textual counterpart), and for verification. Specifically, it provides a code generator that produces CSP specifications defined by the RoboChart semantics for use with FDR.

Section 2 describes RoboChart models. Section 3 gives an overview of their semantics in CSP. Section 4 describes our simulation approach. Section 5 describes RoboTool. Finally, Section 6 reviews related works, and Section 7 concludes with a summary of the results and future work.

2 RoboChart: notation, metamodel, and examples

In this section, we first describe the core features of a RoboChart model (Section 2.1). To illustrate the concepts, we present the model of a robot for chemical detection based on that in [16]¹. In our example, the robot employs a random walk and, upon detection of a chemical source, it turns on a light and drops a flag. Sections 2.2 and 2.3 describe the features to define time and probabilistic properties. Finally, Section 2.4 describes the RoboChart metamodel.

2.1 RoboChart diagrams

A robotic system is specified in RoboChart by a module, where a robotic platform is connected to one or more controllers. A robotic platform is characterised by variables, events, and operations representing in-built facilities of the hardware. For our example, the module **ChemicalDetector** is shown in Figure 1, where we have a robotic platform **Rover** and controllers **DetectAndFlagC** and **LightC**.

Rover declares a number of events via named boxes on its border. The events **lightOn** and **lightOff** are used to request that the in-built light is turned on or off. The events **l** and **r** represent two sensors, one to the left and one to the right. They allow the rover to detect whether there is a wall on either side. The event **alarm** represents an in-built sensor for the position of a chemical source.

The operation **move(v,s)** declared in **Rover** takes a **Vector v** as parameter; it moves the rover in the direction defined by **v** with the speed **s**. In RoboChart, we can define given types (uninterpreted sets), record types, and other structured types. The primitive types include numbers and strings.

The definitions of **move** and **Vector** are part of the RoboChart API. This API is a collection of types, events, and operations organised in packages by function (moving, sensing, and so on), and, for each function,

¹<http://tinyurl.com/hdaws7o>

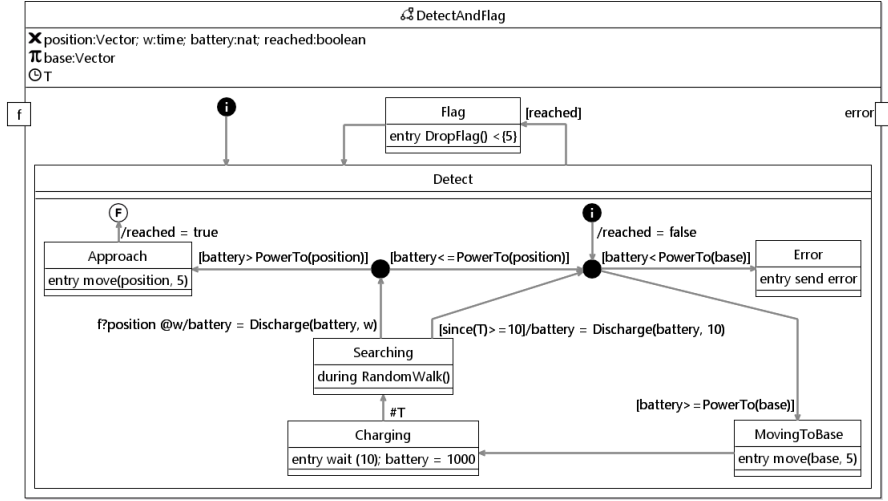


Figure 4: Timed Chemical Detector

tectAndFlag; flagged to done in DropFlag(); and left and right to the events of the same name in RandomWalk(). The definition of RandomWalk() is global, since this is an operation relevant for many applications. So, DetectAndFlagC contains a reference to this state machine, which we present in Section 2.3.

RoboChart state machines are standard, but restricted and with a well defined semantics. They can have composed states, junctions, and entry, during, exit, and transition actions defined using a well defined action language. Features of UML state machines [24] deemed not essential for robotics are not included, resulting in a streamlined semantics. Figure 3 illustrates our notation.

The state-machine DetectAndFlag models the robot roaming behaviour. Two local variables store the position where the chemical is detected and whether that position has been reached. The transition into the initial composed state Detect and its initial state Searching sets reached to false. Its during action is a call to RandomWalk(); this causes the robot to search indefinitely. If the event f, communicating the position of a chemical source, is raised, the machine moves to the state Approach. Its entry action move moves the robot towards the detected source. Once this operation finishes, the machine transitions to its final state, when the variable reached is set to true. This enables the transition from the state Detect to Flag, whose entry action is the operation DropFlag(). The state machine that defines it, also in Figure 3, has just junctions. The transitions defines that the operations LoadFlag and ReleaseFlag are used to select and release a flag next to the chemical source. When DropFlag completes, in DetectAndFlag, the transition from Flag to Detect is taken and the robot restarts the search.

2.2 Time primitives

RoboChart operations take zero time, and enabled transitions take place as soon as they can be triggered. So, time constraints need to be explicitly defined. The time budget t for an operation call can be specified by sequentially composing it with the primitive wait(t), that waits for t time units. A deadline of d time units for a statement S is specified by $S < \{d\}$. Deadlines can also be set on events.

Clocks provide a way to associate the instant in time $\#T$ in which a transition is triggered with subsequent conditions on this instant. The primitive since(T) yields the time elapsed since the most recent time instant $\#T$. A similar primitive sinceEntry(s) yields the time elapsed since entering state s . Finally, the transition primitive $e@t$ records the time elapsed between the moment the event e is available and when the transition it guards is triggered.

To illustrate the time primitives we extend DetectAndFlag. In this version, shown in Figure 4, the additional states Charging, MovingToBase and Error model a robot that can charge its battery on a base station, whose location is fixed by a constant base. The robot can perform different actions depending on how much charge is available, as recorded by the variable battery. We assume that the robot cannot measure this quantity directly, but instead can estimate battery drain using time and an estimate of how much charge is consumed over time.

The state-machine starts at a junction, with outgoing transitions guarded by conditions on the variable battery: if battery is lower than the charge required to reach the base, as estimated by the operation PowerTo (omitted here), then the state-machine transitions to the state Error, otherwise it transitions to MovingToBase. In Error, the event error is sent (to LightC), which turns on the light to indicate that there is a problem. When

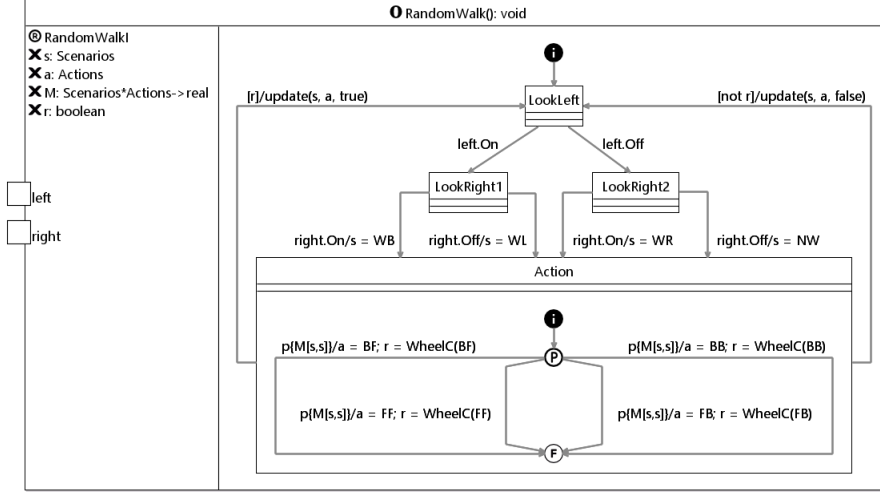


Figure 5: Random walk State-machine

the robot reaches the **base**, the state machine transitions to **Charging**, whose entry action models the charging activity: the primitive `wait(10)` models the fact that it takes 10 time units to charge the battery to its full capacity of 1000 mAh.

Once the battery is charged, a transition leads to the state **Searching**. There are now two possibilities: either the robot finds a chemical source, as signalled by the event `f`, or, if it takes more than 10 time units, then the walk is interrupted and another attempt is made to return to the **base**. The transition to the initial junction updates the estimated **battery** charge using the operation **Discharge** (whose definition is omitted here). To constrain the time allowed for the random walk, we associate with the incoming transition to **Searching** a timed instant `#T`, which is related to a modelling clock `T`. The transition from **Searching** to the initial junction is guarded by a condition on the time primitive `since(T)`.

In case the robot finds a chemical source during the walk, signalled by `f`, we record in the variable `w` the amount of time elapsed since `f` was available, in this case immediately upon entering **Searching**, and its occurrence. This is used by **Discharge** to calculate a new estimate for the **battery** charge based on the time `w`. Moreover, the transition to **Approach** is guarded by a condition that requires that there is enough **battery** charge to reach the **position**, as estimated by **PowerTo**, otherwise the robot attempts to move back to the **base**. Finally, when the robot finishes approaching the chemical source, it transitions to the state **Flag**, where there is a deadline imposed on the operation **DropFlag** to terminate within 5 time units, so that the robot can quickly leave the place.

2.3 Probability primitives

RoboChart has just a single additional construct to cope with modelling of probabilistic behaviour, taken from [18]. It is the P-node: a junction, inscribed with the letter P, and a number of outgoing transitions. Each outgoing transition is labelled with an expression $p\{e\}$ that denotes a probability, together with an optional action. The probabilities for the outgoing edges must sum to 1.0.

In Figure 5, we describe the operation `RandomWalk()`. This state machine has a local variable `s`, which can take values WB, to indicate that there is a wall on both sides (that is, the rover is facing a corner), WL, if there is just a wall to the left, WR, for a wall to the right, or NW, if there is no adjacent wall. Another variable `a` takes as values actions corresponding to actuation of the rover's driving wheels: forwards or backwards. FF takes the rover forward; FB turns right; BF turns left; and BB reverses. Finally, the variable `M` is a matrix: $M[s,a]$ shows the probability of the action `a` occurring in the scenario `s`.

The transition into the initial state **LookLeft** initializes `M` to give probability of $\frac{1}{4}$ to all actions in all scenarios using `init()`. The omitted operation `init()` is in the required interface **RandomWalkI**, which also includes two other operations. **WheelC(a)** controls the wheels to take the input action `a` and has a boolean result that indicates whether the action completes without hitting anything or not. Finally, `update(s,a,b)` updates `M` to record that in scenario `s` the action `a` is good or bad, as indicated by the boolean `b`. If `a` is good, then its probability in `s` is increased, at the expense of the other actions. If it was bad, then its probability is decreased, at the benefit of the others. The precise changes in probabilities can be varied to produce different behaviours.

In **LookLeft**, the rover decides how to move based on the events `left` and `right` of the wall sensors. These events are used to record the position of the walls in `s`. The action that follows depends on the history of

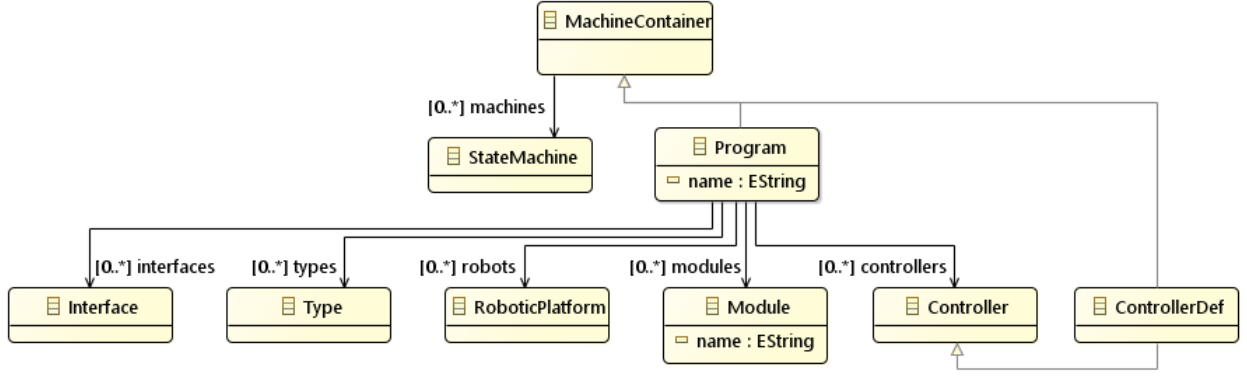


Figure 6: Metamodel of RoboChart models

successful actions recorded in M . The probabilistic choice defined by the P-node selects an action a with probability $M[s,a]$ and executes it using *WheelC*. The fact that alternatives in the probabilistic choice sum to 1.0 is enforced through a row invariant in M . Once *WheelC* finishes, a transition out of the state *Action* back to *LookLeft* is taken, when the matrix M is updated using the operation *update*.

2.4 RoboChart metamodel

As illustrated above, RoboChart models are structured using elements shown in the metamodel in Figure 6, namely, modules, robotic platforms, controllers, state machines, interfaces, and types. Modules give a complete account of a robotic system. They define robotic platforms or include references to platforms defined elsewhere to indicate the robots available. Modules associate their robotic platforms with particular controllers and state machines to specify the behaviours of the robots. State machines can be directly associated to robotic platforms, but, when the behaviour is complex and is specified by multiple (potentially interacting) state machines, controllers can be used.

A module comprises a number of module nodes, which can be controllers, robotic platforms and state machines, and connections between the nodes that establish the relationship between platforms and their specified behaviours.

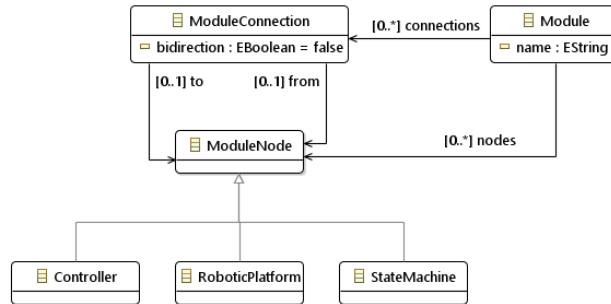


Figure 7: Metamodel of RoboChart modules

A controller encapsulates state machines that can communicate with each other and the external environment through synchronous events. In this way, simpler behaviours can be coordinated to model a complex controller. Robotic platforms, controllers and state machines share features such as variables, operations, events, and provided and required interfaces.

The metamodel of RoboChart state machines is similar to that of UML state machines. Features that have been removed are parallel regions, history junctions, and interlevel transitions. Whilst the state machines are designed with sequential controllers (which may be in parallel with other controllers) in mind, there is space for parallelism in the execution of during actions (see Section 3). When state machines are used to specify the implementation of operations, they declare an operation signature, and may also specify pre and postconditions.

Expressions and statements include time primitives, such as *wait*. Triggers (for state transitions) include probability as well as time recording and deadline primitives (for example, *#T*). In addition, the metamodel for state machines includes P-nodes. These are simple variations of the usual metamodel.

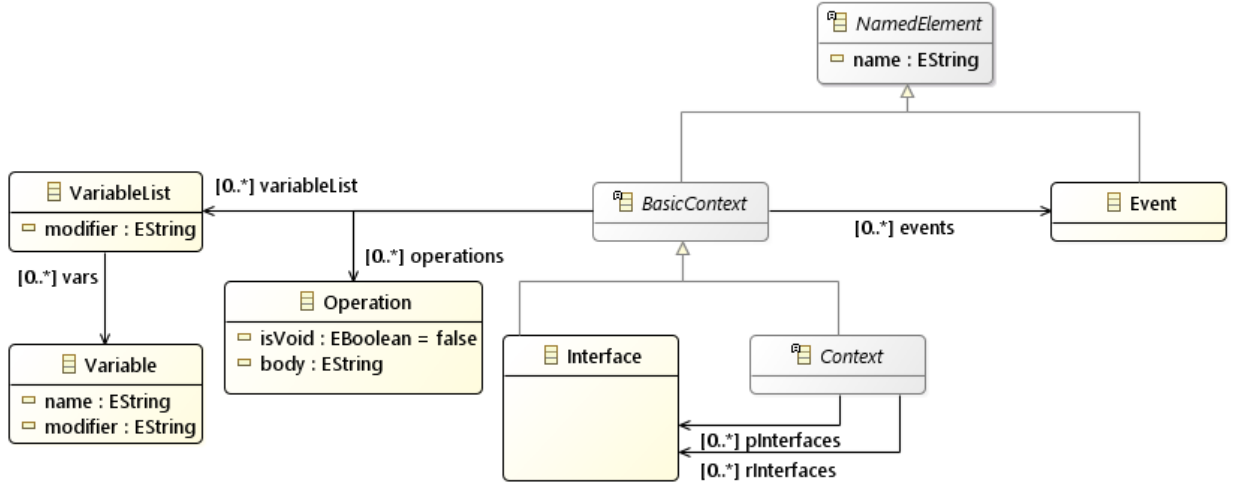


Figure 8: Metamodel of RoboChart constructs

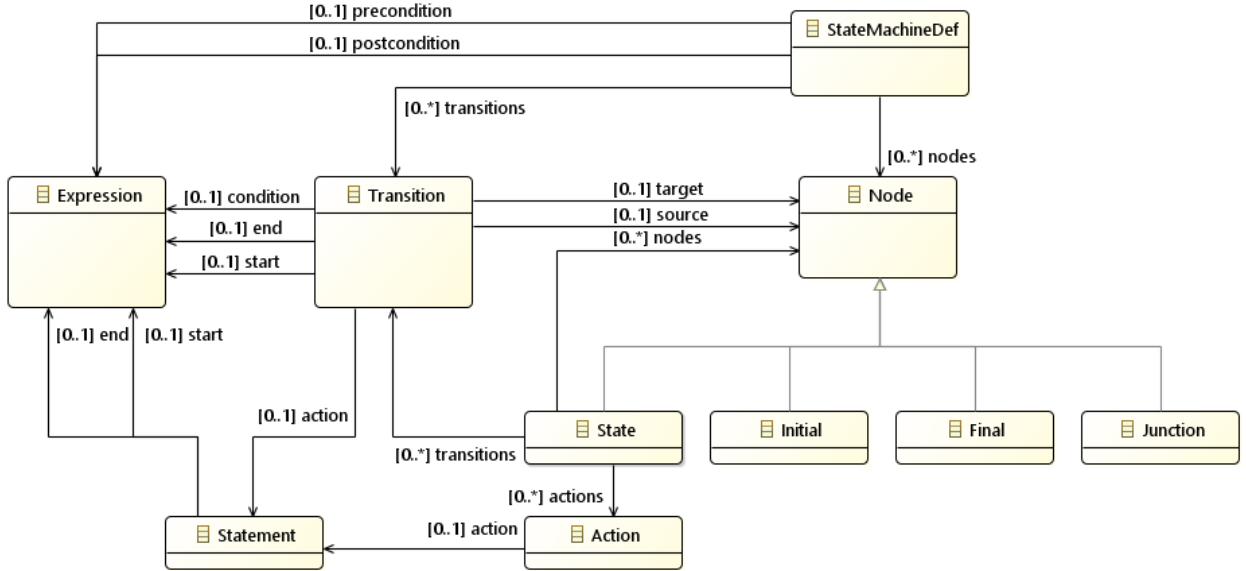


Figure 9: Metamodel of RoboChart state machines

2.5 API operations

In this section, we describe the API in more details. The purpose of the API is to provide a library of generic operations that are useful across a variety of robotic applications and platforms. Some operations are primitive, while others are composite. The composite operations can be implemented by a set of primitive operations. The API is organized by type of robots (e.g. wheeled robots, flying robots, and so on) and equipment (e.g. camera, infrared sensors). The operations also accept inputs such as the speed of the robot.

For example, in the operation `MoveForward`, the linear speed of the robot is set if it is positive and the angular speed will be set to 0. Other operations related with the movement of a 2D robot are shown in Figure 10.

Some operations such as `MoveUp`, `MoveDown`, `Roll/Pitch` are included in the API of flying robots. For humanoid robots, the specific operations include `Lean`, `StandUp` and `FallDown`.

For a composite operation such as `ObstacleAvoidance`, the robot needs to decide how to move depending on the position of the obstacle. `ObstacleAvoidance` can be realized using different sensors such as camera or infrared sensors. Note that although we define the operations in the API in an abstract way, the implementation of a particular operation is related with one robotic platform.

The API can be extended in the tool.

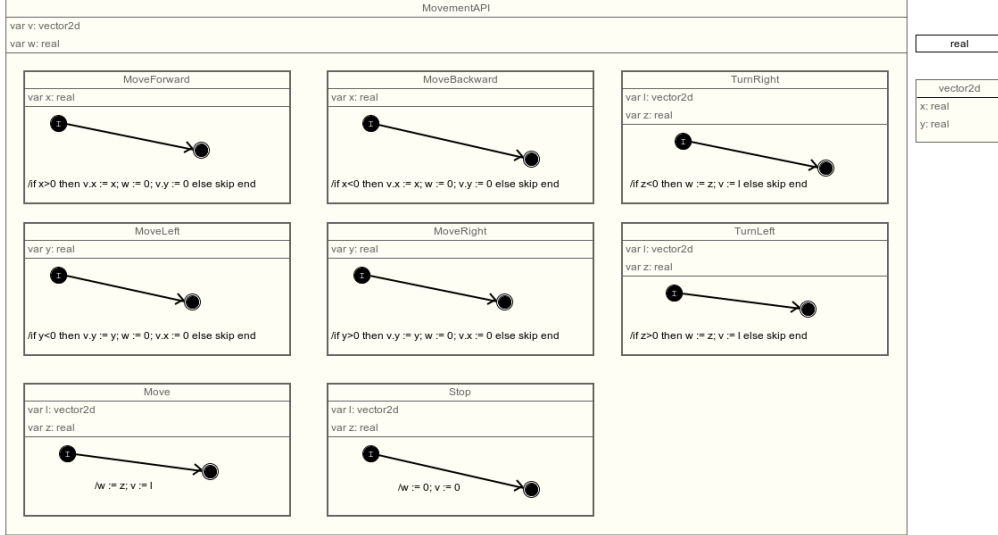


Figure 10: An example of API defining the movement of a robot moving in a 2D environment.

3 Behavioural semantics and verification

To support the verification of models and the generation of sound simulations, we give RoboChart a formal semantics. In this section, we discuss the semantics of the core RoboChart notation (Section 3.1) and its use for verification (Section 3.2). As already mentioned, our formalisation relies on the UTP framework, and on its expressiveness to cater for time and probability in an integrated way. On the other hand, for the core notation, CSP provides the ideal level of abstraction and support for validation via model checking.

3.1 CSP semantics

The semantics of a module is given by the parallel composition of the processes that define the semantics of its controllers interacting according to the connections in the module. We show below the semantics of *ChemicalDetector*, in terms of processes for the controllers *LightC* and *DetectAndFlagC*.

$$\begin{aligned}
 \text{ChemicalDetector} &= (\text{Controllers} \\
 &\quad || \{ \text{activate}, \text{flagged} \} \\
 &\quad \text{Buffer}[\text{flagged}, \text{activate}/\text{in}, \text{out}] \setminus || \{ \text{activate}, \text{flagged} \} \\
 \text{Controllers} &= \text{LightC}[\text{lightOn}, \text{lightOff}/\text{lon}, \text{loff}] \\
 &\quad || \\
 &\quad \text{DetectAndFlagC}[l, r, \text{alarm}/\text{left}, \text{right}, \text{found}]
 \end{aligned}$$

Synchronous interactions such as the one between events *lon* and *lightOn* are modelled by renaming the events of the controller to those of the robotic platform. On the other hand, asynchronous interaction such as that between *activate* and *flagged* is realised via a buffer that runs in parallel ($||$) with the controllers, synchronising on these events. The occurrences of *flagged* are inputs to the buffer, whose outputs are *activate* events. These events are hidden (\setminus). The visible interactions of the system are the events of the robotic platform only.

The semantics of a controller is the parallel composition of the semantics of its main state machines, not including those that define operations, again, interacting according to their connections. This semantics is similar to that of a module, except that the components are the processes that model state machines.

A state machine specifies a sequential control flow hierarchically. Its semantics is the parallel composition of a process *States*, which models each of the states, with a process *InitialTransitions*, which describes the transition to the initial state. Below, we show the semantics of *DetectAndFlag*.

$$\begin{aligned}
 \text{DetectAndFlag} &= \\
 &\quad (\text{InitialTransitions} || \text{EnterExitChannels} || \text{States}) \setminus (\Sigma \setminus \text{Events}) \\
 \text{InitialTransitions} &= \\
 &\quad \text{enter.DetectAndFlag.Detect} \rightarrow \text{entered.DetectAndFlag.Detect} \rightarrow \text{SKIP}
 \end{aligned}$$

The CSP channels used to model the control flow of a state activation and deactivation are *enter*, *entered*, *exit* and *exited*. Events using these channels model the beginning and end of the actions of entering and

exiting a state. Each of them takes two parameters: the state that requested the action to start and the target of the request. For instance, in *InitialTransitions*, the state machine itself requests the activation of the state *Detect*.

A process like *Detect* or *Flag* that models a state *s* does so in a compositional way, capturing only information about *s* itself, irrespective of the context (parent state or state machine) where it occurs. Such a process first offers events *enter* to request and then *entered* to acknowledge entry to *s*, and then offers a choice of events that trigger its transitions, including those of its substates, and any other transitions that might be available for its parents (which, if taken, also lead to an exit of *s*). When any of the transition events is chosen, then the *exit* and *exited* events to request and acknowledge exit are offered.

We note that the state is not aware of the transitions of its parents and, therefore, accepts any of them. In the definition of the process for a parent state or for the state machine, the available choices are further restricted. Only the transitions associated with the ancestor states are feasible.

At the top level, *States* composes in parallel restricted versions of processes for the top states; in our example, versions *DetectR* and *FlagR* of *Detect* and *Flag*. Because there are no parent states, the transition events enabled are those for their own transitions (that triggered by *f*, the final transition, and that with condition reached for *Detect*, and just the silent transition for *Flag*).

States =

$$DetectR \parallel \{ \{ enter.x.y, entered.x.y, exit.x.y, exited.x.y \mid x, y : \{ Detect, Flag \} \} \}$$

$$FlagR$$

$$\setminus \{ \{ enter.x.y, exit.x.y, exited.x.y \mid x, y : \{ Detect, Flag \} \}$$

The processes synchronise on their common activation and deactivation events. For instance, the event *enter.Detect.Flag* is in the synchronisation set, indicating that *Detect* may request the activation of *Flag*, but *enter.DetectAndFlag.Flag* is not in the synchronisation set, because it involves the state machine and *Flag*, but not *Detect*. In other words, synchronisation establishes the flow of activation and deactivation of the states at the same level.

In the parallelism that defines a state machine, like *DetectAndFlag* above, the possibility of entering a state different from the initial state is blocked. To recall, *States* offers the possibility of entering (*Detect* or *Flag*), and later exiting the states. All these events are included in the set *EnterExitChannels*. So, *InitialTransitions* defines exactly which state is entered. The machine terminates if it reaches a final state, but it does not require that state to exit. So, an event like *exit.SM.S*, where *SM* is the state-machine identifier is not possible. Instead, if a machine terminates, a special event *terminated* is raised.

Finally, we hide all CSP events (set Σ), except those in the set *Events*, which represent events in the RoboChart model itself. These are the only events visible in the semantics, although, for verification, we may hide fewer events.

The CSP semantic models just described can be automatically generated for a RoboChart model using our RoboTool described in Section 5. The complete model for our example can be found in Section 3.3

3.2 Verification

The automatically generated semantics outlined above, can be used for analysis and verification using the FDR model-checker. We can, for example, establish determinism, and absence of divergence and deadlock. In our example, the *DetectAndFlag* state machine is very simple. We are encouraged, however, because, although to give a compositional semantics, we use a lot of extra events and parallelism, the use of the compression functions *diamond* and *sbisim* (that preserve the semantics of the processes) radically optimises the analysis reducing the number of states from thousands to just four as expected.

A second example of the kind analysis we can perform is the verification of assumptions about the hardware controlled by the state machine. For instance, the rover can carry a limited amount of flags. We have modelled this mechanism as a CSP process and used a refinement to verify whether the controller satisfies this restriction of the hardware. As probably already observed, it does not, and must be refined. In such a verification, FDR produces a counterexample that pinpoints a scenario that break this assumption. This information can be used to guide the redesign of the state machine.

As a final example, we can perform reachability analysis using FDR. For each state *s* of a state machine, we can define a process by renaming events of the form *entered.x.s* to *hasEntered.s*, and hiding all events except *hasEntered.s*. We can then verify that this process contains the trace $\langle hasEntered.s \rangle$, and so *s* is entered at least once. We have applied this technique to a modified version of our example where the variable *reached* is not updated. In this case, the verification correctly points out that the state *Flag* is unreachable.

No doubt, further evaluation is necessary to gauge the effectiveness of model checking RoboChart semantic models. We are, however, encouraged by the results with compression, which are related to the structure of the models, and have the possibility of exploring theorem proving using the UTP theory for CSP.

3.3 Complete Semantics of *DetectAndFlag*

In this section, we present and explain the complete model of the state-machine *DetectAndFlag* shown in Figure 3. As briefly described in Section 3 this model is a denotational semantics of RoboChart state-machines specified in CSP where parallelism is used mainly to conjoin requirements. The behaviour of each state is specified as a CSP process, and these processes are composed in parallel to formalise the behaviours of composite state and the state-machine itself.

```

nat = {1,2,3}

nametype ID = Seq(Char)

FINAL = "_final_"

channel terminate

channel enter, entered: ID.ID
channel exit, exited: ID.ID

transparent chase
external prioritise
transparent diamond
transparent sbisim

sbdia(P) = sbisim(diamond(P))

channel not_deadlocked

channel internal__: ID

nametype Vector = (nat,nat)
nametype boolean = {True,False}

channel event_f: ID.Vector
channel f: Vector

InterruptsBut(ids) = {| internal__.x, event_f.x | x <- diff(TIDS,ids) |}
Interrupts(id) = {| internal__.id, event_f.id |}
InterruptsOf(ids) = {| internal__.x, event_f.x | x <- ids |}

nametype IDS = {"DetectAndFlag", "Detect", "Searching", "Approach", "Flag"}
nametype TIDS = {"t1", "t2", "t3", "t4"}

t1 = "t1"
t2 = "t2"
t3 = "t3"
t4 = "t4"

channel set_position, get_position: Vector
channel set_reached, get_reached: boolean

channel cmove, walk, drop

DropFlag = drop -> SKIP
RandomWalk = walk -> SKIP
move(x,y) = cmove -> SKIP

Flag = enter?x: diff(IDS,{"Flag"})!"Flag" -> (
  (
    (DropFlag; entered!x!"Flag" -> SKIP);
    (STOP/\(

```

```

        internal...t1 -> enter!"Flag"!"Detect" -> entered!"Flag"!"Detect" -> SKIP
        [] ([ e: InterruptsBut({t1}) @ e -> exit?x:diff(IDS,{"Flag"})!"Flag" ->
            (exited!x!"Flag" -> SKIP))
    ))
); Flag

Searching = enter?x:diff(IDS,{"Searching"})!"Searching" -> (
    (
        (entered!x!"Searching" -> SKIP);
        (RandomWalk;STOP/\(
            event.f.t2?position -> set_position!position ->
                enter!"Searching"!"Approach" ->
                entered!"Searching"!"Approach" -> SKIP
            [] ([ e: InterruptsBut({t2}) @ e ->
                exit?x:diff(IDS,{"Searching"})!"Searching" ->
                (exited!x!"Searching" -> SKIP))
        ))
    )
); Searching

Approach = enter?x:diff(IDS,{"Approach"})!"Approach" -> (
    (
        (get_position?p -> move(p,5); entered!x!"Approach" -> SKIP);
        (STOP/\(
            internal...t3 -> set_reached!True -> Final
            [] ([ e: InterruptsBut({t3}) @ e ->
                exit?x:diff(IDS,{"Approach"})!"Approach" ->
                (exited!x!"Approach" -> SKIP))
        ))
    )
); Approach

Final = ([ e: InterruptsBut({t3}) @ e ->
    exit?x:diff(IDS,{FINAL})!FINAL ->
    (exited!x!FINAL -> SKIP))

DetectAux = enter?x:diff(IDS,{"Detect"})!"Detect" -> (
    set_reached!False -> enter!"Detect"!"Searching" ->
    entered!"Detect"!"Searching" -> entered!x!"Detect" -> SKIP
);
(STOP/\(
    internal...t4 -> exit!"Detect"?y:{"Searching","Approach", FINAL} ->
    exited!"Detect"!y -> enter!"Detect"!"Flag" ->
    entered!"Detect"!"Flag" -> SKIP
    []
    ([ e: InterruptsBut({t2,t3,t4}) @ e ->
        exit?x:diff(IDS,{"Detect"})!"Detect" ->
        exit!"Detect"?y:{"Searching","Approach", FINAL} ->
        exited!"Detect"!y -> exited!x!"Detect" -> SKIP)
)); DetectAux

SearchingR = Searching
[] diff(union(Interrupts(t2),Interrupts(t3)),{|event.f.t2|})||
SKIP

ApproachR = Approach
[] diff(union(Interrupts(t2),Interrupts(t3)),{|internal...t3|})||
SKIP

```

```

DetectSubStates = (
  SearchingR
  [|{|enter.x.y, entered.x.y, exit.x.y, exited.x.y |
    x <- {"Searching","Approach",FINAL},
    y <- {"Searching","Approach",FINAL}}|}]
  ApproachR)
)\{|enter.x.y, exit.x.y, exited.x.y |
  x <- {"Searching","Approach"},
  y <- {"Searching","Approach"}}|}]

Detect =
(
  DetectAux
  [|
    union(diff(InterruptsBut({}),{|event_f.t2 ,internal_.t3|}),
      {|enter.y.x, entered.y.x, exit.y.x, exited.y.x |
        x <- {"Searching","Approach",FINAL},
        y <- diff(IDS,{"Searching","Approach",FINAL})|})
  |]
  DetectSubStates
)\{|enter."Detect".x, exit."Detect".x, exited."Detect".x |
  x <- {"Searching","Approach",FINAL}}|}]

Machine = enter."DetectAndFlag"."Detect" -> entered."DetectAndFlag"."Detect" -> SKIP
  [|
    {|enter.x.y,entered.x.y,exit.x.y,exited.x.y|
      x<-diff(IDS,{"Detect","Flag"}),
      y<-{"Detect","Flag"}}|}]
  [|]
  MachineSubStates

DetectR = Detect
  [| diff(InterruptsOf({t1,t2,t3,t4}),
    {|internal_.t4, event_f.t2, internal_.t3|})|]
  SKIP

FlagR = Flag
  [| diff(InterruptsOf({t1,t2,t3,t4}),{|internal_.t1|})|]
  SKIP

MachineSubStates = (
  DetectR
  [|{|enter.x.y, entered.x.y, exit.x.y, exited.x.y |
    x <- {"Detect","Flag"},
    y <- {"Detect","Flag"}}|}]
  FlagR
)\{|enter.x.y, exit.x.y, exited.x.y |
  x <- {"Detect","Flag"}, y <- {"Detect","Flag"}}|}]

DetectAndFlagForReachability =
(((
  Machine\{|enter.x.y,exit.x.y,exited.x.y|x<-IDS,y<-{"Detect","Flag"}}|}
  [|{|get_position,get_reached,set_position,set_reached,internal_.t4|}]
  Memory((1,1),false)
)\[event_f.x <- f | x <- TIDS])
\{|get_position,get_reached,set_position,set_reached,internal_.x|x<-TIDS|}]

DetectAndFlag = DetectAndFlagForReachability\{|entered.x.y|x<-IDS,y<-IDS|}]

```

```

Memory(position, reached) =
  get_position!position -> Memory(position, reached)
  []
  get_reached!reached -> Memory(position, reached)
  []
  set_position?x -> Memory(x, reached)
  []
  set_reached?x -> Memory(position, x)
  []
  (reached)&internal...t4 -> Memory(position, reached)

```

The properties discussed in Section 3.2 can be verified using the following assertions.

The limitation of the DropFlag mechanism that only two flags can be dropped during the execution of the robots is specified by the abstract models *AbstractSpecification1* and *AbstractSpecification2*. These processes model the behaviour where any events can happen any number of times, except for the event *drop* that can only happen twice; the first process is used to verify the property under the traces model, whilst the second is used for verification in the failures-divergences model.

```

assert sbdia(DetectAndFlag) :[deterministic [FD]]
assert sbdia(DetectAndFlag) :[deadlock-free [FD]]
assert sbdia(DetectAndFlag) :[divergence-free [FD]]

```

```
Trace = walk -> f.(2,2) -> cmove -> drop -> Trace
```

```

AbstractSpecification1 =
  Run({|f, cmove, walk|});
  drop -> Run({|f, cmove, walk|}); drop -> Run({|f, cmove, walk|})
AbstractSpecification2 =
  Chaos({|f, cmove, walk|});
  drop -> Chaos({|f, cmove, walk|}); drop -> Chaos({|f, cmove, walk|})

```

```

Run(events) = [] e: events @ e -> Run(events) [] SKIP
Chaos(events) = |~| e: events @ e -> Chaos(events) |~| SKIP

```

```

assert sbdia(AbstractSpecification1) [T= sbdia(DetectAndFlag)
assert sbdia(AbstractSpecification2) [FD= sbdia(DetectAndFlag)

```

```
channel hasEntered: IDS
```

```

DetectAndFlagReachability(s) =
  (sbdia(DetectAndFlagForReachability)[[entered.x.y<-hasEntered.y|x<-IDS,y<-IDS]])
  \{|f, walk, cmove, drop, hasEntered.x|x<-diff(IDS,{s})|}

```

```

assert DetectAndFlagReachability("Approach") :[has trace]: <hasEntered."Approach">
assert DetectAndFlagReachability("Searching") :[has trace]: <hasEntered."Searching">
assert DetectAndFlagReachability("Detect") :[has trace]: <hasEntered."Detect">
assert DetectAndFlagReachability("Flag") :[has trace]: <hasEntered."Flag">

```

```
assert DetectAndFlag :[deterministic [FD]]
```

4 Simulation

Typically, there is no rigorous connection between a design and its simulation. Here, we outline our approach to simulation of RoboChart controllers. Simulations are platform independent and can be automatically generated for use with a wide variety of simulation tools. Examples are V-rep [19] and Webots [23], which are integrated development tools that can simulate different kinds of robots, simulation scenarios and sophisticated physics engines. It is also possible to directly import the generated code into a real robot.

We use object-oriented simulations. This paradigm is adopted in many simulators [19, 23, 28] and supports a direct mapping of RoboChart constructs (controllers, state machines, and so on) to simulations in a well




Figure 11: RoboChart simulations

RoboChart	Simulation
states	attribute of enumerated type
variable	attribute
events	attribute of enumerated type
actions	methods

Table 1: RoboChart state-machine classes

structured way.

A module defines a complete simulation, although additional information about the environment is needed. We adopt the model-view-controller (MVC) pattern in the design of simulations, where, here, the terms model and controller are used in an entirely different way from that adopted in RoboChart. Figure 11 maps the RoboChart constructs to an MVC architecture.

The model (M) component contains a simulation of the environment and of the RoboChart controller. We can generate a simulation of the RoboChart controller, which, together with a simulation of the environment defines the M component. The controller component (C) implements the robotic platform. It provides the variables, events, and operations defined in the RoboChart robotic platform. Finally, the view component (V) defines the interface of the simulation.

With this approach, we have simulations that can take advantage of different simulations of robotic platforms and can be visualised in different ways, without requiring changes to the M component. For example, the environment model in the M component can be implemented with a variety of different physics engines, for example, Bullet² or Delta3D³. Likewise, for the V component, we may choose to view the robot with different visualization tools in 3D, or 2D or in fact choose only a textual representation of the platform to allow for debugging.

The simulation of a controller is the simulation of its main state machine. Each main state machine generates a class in the simulation. Table 1 defines how constructs of a state machine are mapped to elements of this class.

The variables, events, and states are defined as attributes of the class. The actions (entry, during and exit) that the robot executes in each state give rise to methods. We note that, even if an operation is specified by a state machine, that state machine also generates a method.

The operations of the API have default implementations that assume that the preconditions hold. This allows more efficient implementations. We plan to include the preconditions in the simulation code to support verification.

In a simulation, the motion of the robot is updated in a cycle-based way (that is, in control steps). Therefore, the time in a simulation is discrete and the length of the control step is the smallest granularity for a clock. If the controller is implemented in a sequential way, we also need to assume that the length of the control step is small enough in order not to miss any event.

For our example, a class `DetectAndFlag` contains attributes `position` and `reached` for the variables of the state machine, and `done`, `f`, `right`, and `left` for its events and those of the operations it uses. The methods for the operations of the state machine are `DropFlag()`, `RandomWalk()`, and `move(position,speed)`. A main method runs the state machine: checks the incoming events, makes transitions, and executes the actions in each state.

5 Tool support

Support for RoboChart modelling is provided by RoboTool⁴, a set of Eclipse [1] plugins implemented using the Xtext [3] and Sirius [2] frameworks. Xtext, using the definition of a grammar, automatically generates plugins that implement a parser and a linker, and provides mechanisms for the implementation of scope providers, validators, type checkers, code generators, and so on. Sirius supports the definition of graphical representation for metamodels, and produces an Eclipse plugin for construction and visualisation of instances of the metamodel.

The RoboChart metamodel presented in Section 2.4 has been implemented using the *Eclipse Modeling Framework* (EMF), and used to generate a textual editor using Xtext and a graphical editor using Sirius. Figure 12 shows these editors with the model of the module `ChemicalDetector`.

Table 2 shows the existing and planned features of RoboChart and how far they have been implemented. Both the textual and graphical editors have been implemented for the core and timed languages, and partially

²<http://bulletphysics.org/wordpress/>

³<http://delta3d.org/>

⁴<https://www.cs.york.ac.uk/circus/RoboCalc-tools/>

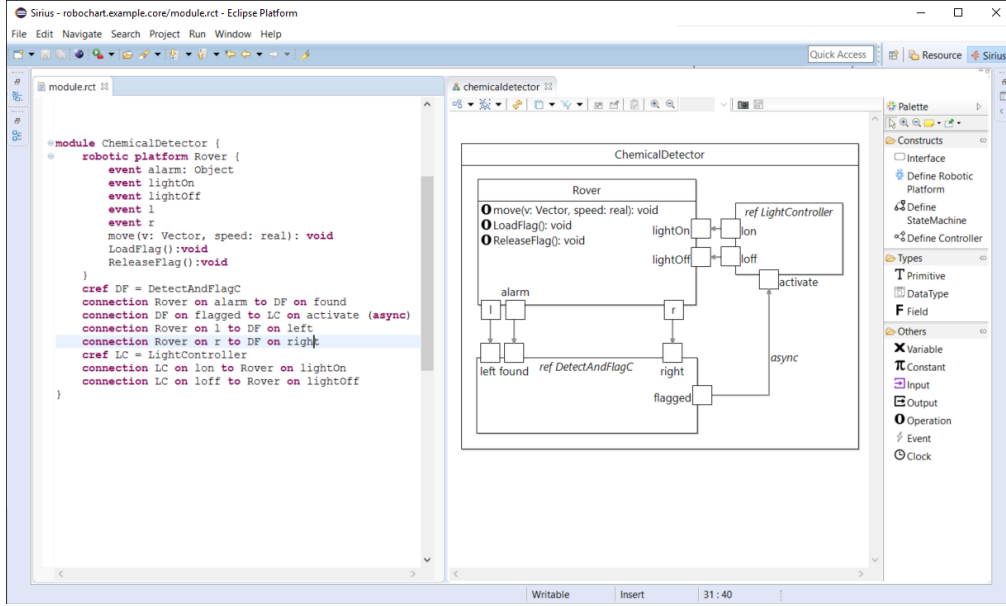


Figure 12: RoboChart textual and graphical editors

	Textual	Graphical	Validation	Semantics
Core Language	done	partial	partial	partial
Time	done	done		
Probability	partial	partial		

Table 2: Features of RoboTool

for the probabilistic language. Validation and semantics have been addressed at the level of the textual language so far. In addition to standard features implemented out-of-the-box by Xtext, we have implemented validation rules that verify well-formedness conditions such as the existence of an initial node in a composite state, as well as the semantics of RoboChart as a code-generator.

The code generator produces CSP specifications that can be animated and analysed using FDR. It is possible to check state machines and controllers for deadlock-freedom and nondeterminism. More interestingly, it is possible to check for refinement between any two state-machines or controllers. To optimise the analysis, we make extensive use of compression functions available in FDR3, including the *chase* function, which, in general, does not preserve semantics. This is adequate only for deterministic state machines and controllers. Due to the compositional nature of our semantics, we can use this function selectively to the deterministic components of the specification.

6 Related work

Another graphical domain-specific language for robotics is presented in [10]. It also aims to support design modelling and automatic generation of platform-independent code. It was defined as a UML profile. Model-based engineering of robotic systems is also advocated in [30], where a component-based framework that uses UML to develop robotics software is presented. In contrast, RoboChart is a small language, with a well defined semantics to support sound generation of formal models as well as simulations. It is a generic domain-specific language for designing and verifying the robot controllers.

There are a multitude of models for UML state machines. Kuske *et al.* [20] gave an integrated semantics for UML class, object, and state-machine diagrams using graph transformation. Rasch and Wehrheim [27] presented integrated semantics in CSP for (extended) class diagrams and state machines. Davies and Crichton [9] described CSP models for UML class, object, statechart, sequence and collaboration diagrams. Broy *et al.* [6] presented one of the first foundational semantics for a subset of UML2. Similarly, our semantics gives a precise characterisation of state machines and is close to [27] and [9] in our use of CSP. Our state machines, however, do not include components like history junctions and inter-level transitions to enable a compositional semantics.

UML has a simple notion of time. Its profile UML-MARTE [33] supports logical, discrete and continuous time through the notion of clocks. Complex constraints may be specified using CCSL (Clock Constraint Specification Language). Specification of time budgets and deadlines is focused on particular instances of

behaviour specified through sequence and time diagrams. It is not possible to define timed constraints in terms of transitions and states.

UML-RT [32], an extension to UML, includes the notion of capsules, which encapsulate state machines; communication between capsules takes place through ports, whose valid communications are defined by protocols. A timing protocol can act as a timer by raising timeouts in response to the passage of a certain amount of time. It is not obvious how timed constraints, such as deadlines, can be specified directly on state machines beyond informal annotations.

Timed automata [5] caters for timed models using synchronous continuous-time clocks. Interesting properties can be checked using the model checker UPPAAL. Modelling the execution time of operations, and more complex constraints, requires UPPAAL patterns consisting of additional states and appropriate state invariants. Our aim is to provide a rich language suitable for directly capturing timed aspects of interest in robotic controllers.

Calendar automata [11] have been used to model time-triggered architecture systems. Support for model-checking is available using SAL. Calendar automata adopts a strict interleaving between time evolving till the next calendar entry, and events taking place. This model, strictly less expressive than timed automata, is not adequate for modelling scenarios where events do not necessarily alternate with time passing, and where transitions may be nondeterministic.

In [26] a semantics is given for a subset of UML-RT in *Circus* without considering time. An extension to UML-RT is considered in [4] with semantics given in terms of CSP+T [38], an extension of CSP that supports the timing of events. Inspired by the constructs of CSP+T, in [4] annotations are added for recording the occurrence time of events and constraining the occurrence time of other subsequent events. Although timed primitives such as *since* bear a resemblance, we have a richer set of primitives inspired by timed automata and Timed CSP [31].

7 Conclusions

We have presented here, RoboChart, a new notation for modelling of robots. It is based on UML state machines, but includes the notions of robotic platform and controller, synchronous and asynchronous communications, an API of operations common to autonomous and mobile robots, a well defined action language, pre and postconditions, and time and probability primitives.

We have described a semantics for the core constructs of RoboChart. It uses CSP, but we envisage its extension to use *Circus* [7], a process algebra that combines Z [35] and CSP, and includes time constructs [34]. We can already reason about untimed and non-probabilistic robotic systems [13]. Use of *Circus* and its UTP foundation will enable use of theorem proving as well as model checking. Work on probability is available in the UTP [37], but we will pursue an encoding of Markov decision processes in the UTP.

An approach for writing object-oriented simulations of RoboChart diagrams has also been defined. Automatic generation of simulations is possible and part of our future work. Verification of correctness of simulations will use the object-oriented version of *Circus* [8], with a semantics given by the UTP theory in [36].

Finally, we have presented RoboTool, a user friendly tool for modelling and verification of core RoboChart diagrams. Extensions of RoboTool will support the timed and probabilistic primitives, and automatic generation of simulations.

RoboChart itself misses support for modelling the environment and the robotic platforms in model detail. It is also in our plans to take inspiration from hybrid automata [15] to extend the notation, and from the UTP model of continuous variables [12] to define the semantics.

Acknowledgements This work is funded by the EPSRC grant EP/M025756/1. No new primary data was created during this study.

References

- [1] Eclipse. <http://www.eclipse.org/>. Accessed: 2016-05-25.
- [2] Sirius. <https://www.eclipse.org/sirius/>. Accessed: 2016-05-25.
- [3] Xtext. <https://eclipse.org/Xtext/>. Accessed: 2016-05-25.
- [4] K. B. Akhlagi, M. I. C. Tunon, J. A. H. Terriza, and L. E. M. Morales. A methodological approach to the formal specification of real-time systems by transformation of UML-RT design models. *Science of Computer Programming*, 65(1):41–56, 2007.

- [5] R. Alur and D. L. Dill. A theory of timed automata. *Theoretical Computer Science*, 126(2):183–235, 1994.
- [6] M. Broy, M. V. Cengarle, and B. Rumpe. Semantics of UML – Towards a System Model for UML: The State Machine Model. Technical Report TUM-I0711, Institut für Informatik, Technische Universität München, February 2007.
- [7] A. L. C. Cavalcanti, A. C. A. Sampaio, and J. C. P. Woodcock. A Refinement Strategy for *Circus*. *Formal Aspects of Computing*, 15(2 - 3):146–181, 2003.
- [8] A. L. C. Cavalcanti, A. C. A. Sampaio, and J. C. P. Woodcock. Unifying Classes and Processes. *Software and System Modelling*, 4(3):277–296, 2005.
- [9] J. Davies and C. Crichton. Concurrency and Refinement in the Unified Modeling Language. *Formal Aspects of Computing*, 15(2-3):118–145, 2003.
- [10] S. Dhouib, S. Kchir, S. Stinckwich, T. Ziadi, and M. Ziane. *Simulation, Modeling, and Programming for Autonomous Robots*, chapter RobotML, a Domain-Specific Language to Design, Simulate and Deploy Robotic Applications, pages 149–160. Springer, 2012.
- [11] B. Dutertre and M. Sorea. *Formal Techniques, Modelling and Analysis of Timed and Fault-Tolerant Systems: Joint International Conferences on Formal Modeling and Analysis of Timed Systems*, chapter Modeling and Verification of a Fault-Tolerant Real-Time Startup Protocol Using Calendar Automata, pages 199–214. Springer, 2004.
- [12] S. Foster, B. Thiele, A. L. C. Cavalcanti, and J. C. P. Woodcock. Towards a UTP semantics for Modelica. In *Unifying Theories of Programming*, Lecture Notes in Computer Science. Springer, 2016.
- [13] S. Foster, F. Zeyda, and J. C. P. Woodcock. Isabelle/UTP: A Mechanised Theory Engineering Framework. In D. Naumann, editor, *Unifying Theories of Programming*, volume 8963 of *Lecture Notes in Computer Science*, pages 21–41. Springer, 2015.
- [14] T. Gibson-Robinson, P. Armstrong, A. Boulgakov, and A. W. Roscoe. FDR3 A Modern Refinement Checker for CSP. In *Tools and Algorithms for the Construction and Analysis of Systems*, pages 187–201, 2014.
- [15] T. A. Henzinger. The theory of hybrid automata. In *11th Annual IEEE Symposium on Logic in Computer Science*, pages 278–292, 1996.
- [16] J. A. Hilder, N. D. L. Owens, M. J. Neal, P. J. Hickey, S. N. Cairns, D. P. A. Kilgour, J. Timmis, and A. M. Tyrrell. Chemical detection using the receptor density algorithm. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 42(6):1730–1741, 2012.
- [17] C. A. R. Hoare and He Jifeng. *Unifying Theories of Programming*. Prentice-Hall, 1998.
- [18] D. N. Jansen, H. Hermanns, and J.-P. Katoen. A Probabilistic Extension of UML Statecharts. In W. Damm and E.-R. Olderog, editors, *Formal Techniques in Real-Time and Fault-Tolerant Systems*, volume 2469 of *Lecture Notes in Computer Science*, pages 355–374. Springer, 2002.
- [19] N. Koenig and H. Andrew. Design and use paradigms for gazebo, an open-source multi-robot simulator. In *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems*, volume 3, pages 2149–2154. IEEE, 2004.
- [20] S. Kuske, M. Gogolla, R. Kollmann, and H.-J. Kreowski. An Integrated Semantics for UML Class, Object and State Diagrams Based on Graph Transformation. In M. Butler, L. Petre, and K. SereKaisa, editors, *Integrated Formal Methods*, volume 2335 of *Lecture Notes in Computer Science*, pages 11–28. Springer, 2002.
- [21] M. Kwiatkowska, G. Norman, and D. Parker. Probabilistic symbolic model checking with PRISM: a hybrid approach. *International Journal on Software Tools for Technology Transfer*, 6(2):128–142, 2004.
- [22] Wenguo Liu and Alan F. T. Winfield. Modeling and optimization of adaptive foraging in swarm robotic systems. *The International Journal of Robotics Research*, 29(14):1743–1760, 2010.
- [23] O. Michel. Webots: Professional mobile robot simulation. *International Journal of Advanced Robotic Systems*, 1(1):39–42, 2004.

- [24] Object Management Group. OMG Unified Modeling Language (OMG UML), superstructure, version 2.4.1. Technical report, OMG, 2011.
- [25] Object Management Group. *OMG Unified Modeling Language*, March 2015.
- [26] R. Ramos, A. C. A. Sampaio, and A. C. Mota. A Semantics for UML-RT Active Classes via Mapping into *Circus*. In *Formal Methods for Open Object-based Distributed Systems*, volume 3535 of *Lecture Notes in Computer Science*, pages 99–114, 2005.
- [27] H. Rasch and H. Wehrheim. Checking consistency in UML diagrams: Classes and state machines. In E. Najm, U. Nestmann, and P. Stevens, editors, *Formal Methods for Open Object-Based Distributed Systems*, volume 2884 of *Lecture Notes in Computer Science*, pages 229–243. Springer, 2003.
- [28] E. Rohmer, S. P. N. Singh, and M. Freese. V-rep: A versatile and scalable robot simulation framework. In *IEEE International Conference on Intelligent Robots and Systems*, volume 1, pages 1321–1326. IEEE, 2013.
- [29] A. W. Roscoe. *Understanding Concurrent Systems*. Texts in Computer Science. Springer, 2011.
- [30] C. Schlegel, T. Hassler, A. Lotz, and A. Steck. Robotic software systems: From code-driven to model-driven designs. In *14th International Conference on Advanced Robotics*, pages 1–8. IEEE, 2009.
- [31] S. Schneider. *Concurrent and Real-time Systems: The CSP Approach*. Wiley, 2000.
- [32] B. Selic. Using UML for modeling complex real-time systems. In F. Mueller and A. Bestavros, editors, *Languages, Compilers, and Tools for Embedded Systems*, volume 1474 of *Lecture Notes in Computer Science*, pages 250–260. Springer, 1998.
- [33] B. Selic and S. Grard. *Modeling and Analysis of Real-Time and Embedded Systems with UML and MARTE: Developing Cyber-Physical Systems*. Morgan Kaufmann Publishers Inc., 2013.
- [34] A. Sherif, A. L. C. Cavalcanti, J. He, and A. C. A. Sampaio. A process algebraic framework for specification and validation of real-time systems. *Formal Aspects of Computing*, 22(2):153–191, 2010.
- [35] J. C. P. Woodcock and J. Davies. *Using Z—Specification, Refinement, and Proof*. Prentice-Hall, 1996.
- [36] F. Zeyda, T. L. V. L. Santos, A. L. C. Cavalcanti, and A. C. A. Sampaio. A modular theory of object orientation in higher-order UTP. In *Formal Methods*, volume 8442 of *Lecture Notes in Computer Science*, pages 627–642. Springer, 2014.
- [37] H. Zhu, J. W. Sanders, He Jifeng, and S. Qin. Denotational Semantics for a Probabilistic Timed Shared-Variable Language. In B. Wolff, M.-C. Gaudel, and A. Feliachi, editors, *Unifying Theories of Programming*, volume 7681 of *Lecture Notes in Computer Science*, pages 224–247. Springer, 2013.
- [38] J. J. Zic. Time-constrained Buffer Specifications in CSP + T and Timed CSP. *ACM Transactions on Programming Languages and Systems*, 16(6):1661–1674, 1994.