Code generation from UML/MARTE/OCL Environment Models to Support Automated System Testing of Real-Time Embedded Software[1]

MUHAMMAD ZOHAIB IQBAL[a,b], ANDREA ARCURI [a], LIONEL BRIAND [a, c]

[a] *Simula Research Laboratory, P.O. Box 134, Lysaker, Norway*

[b] *Department of Informatics, University of Oslo, Norway*

[c] *SnT Centre for Security, Reliability, and Trust, University of Luxembourg, Luxembourg*

{zohaib, arcuri, briand}@simula.no, lionel.briand@uni.lu

Abstract. Given the challenges of testing at the system level, only a fully automated approach can really scale up to industrial real-time embedded systems (RTES). Our goal is to provide a practical approach to the model-based testing of RTES by allowing system testers, who are often not familiar with the system's design but are application domain experts, to model the system environment in such a way as to enable its black-box test automation. Environment models can support the automation of three tasks: the code generation of an environment simulator to enable testing on the development platform or without involving actual hardware, the selection of test cases, and the evaluation of their expected results (oracles). From a practical standpoint—and such considerations are crucial for industrial adoption—environment modeling should be based on modeling standards (1) that are at an adequate level of abstraction, (2) that software engineers are familiar with, and (3) that are well supported by commercial or open source tools. In this paper, we propose a precise environment modeling methodology fitting these requirements and discuss how these models can be used to generate environment simulators. The environment models are expressed using UML/MARTE and OCL, which are international standards for real-time systems and constraint modeling. The presented techniques are evaluated on a set of three artificial problems and on two industrial RTES.

Key Words and Phrases: Environment modeling, environment simulation, automated testing, model-based testing, real-time embedded systems, search based software engineering

1. Introduction

Real-time embedded systems (RTES) are widely used in many different domains, as for example from integrated control systems to consumer electronics. Already 98% of computing devices are embedded in nature and it is estimated that, by the year 2020, there will be over 40 billion embedded

[1] This paper is an extension of a conference paper "Environment Modeling with UML/MARTE to Support Black-Box System Testing for Real-Time Embedded Systems: Methodology and Industrial Case Studies" published in proceedings of ACM/IEEE International Conference on Model Driven Engineering Languages and Systems (MODELS), 2010 [1]

computing devices worldwide [2]. Testing these systems such that they are functionally correct and do not lead their environment into critical states (e.g., unsafe) is vital. RTES environments typically comprise a number of physical components (e.g., sensors and actuators) and possibly other RTES systems (e.g., in systems of systems). Typically, there is a large number and variety of stimuli to the RTES with different patterns of arrival times. These characteristics make the testing of RTES challenging and increase the need for automated, systematic testing strategies.

Because RTES are developed for diverse domains presenting different constraints (e.g., different timing, safety, security requirements), different testing approaches are required to handle the varying set of characteristics required by these domains [3]. Our main target RTES in this paper are soft-real time systems with time deadlines in the order of hundreds of milliseconds with an acceptable jitter of a few milliseconds in response time. Our testing approach (black-box system level testing) not only encompasses functional correctness of the system under test (SUT), but also enable to focus testing on particularly critical aspects of the RTES, i.e., potentially hazardous situations.

Typically, large scale testing of RTES software in real environments and on actual deployment platforms is not a viable option. It would be expensive, the consequences of failures might be catastrophic (e.g., in safety critical systems), and the number of variations in the environment that can be exercised within a reasonable time frame are small. Moreover, some of the environment components might not be available at the time of testing, since hardware and software components are typically developed concurrently. To test RTES software in this kind of situations, a common strategy is to develop a simulator for these environment components.

When testing RTES, the simulation of three concepts (or their combinations) is typically considered: the SUT, its hardware platform, and the environment with which the SUT interacts. Depending on the goal of testing, different combination of these three concepts can be simulated [3]: (i) at early stages of the development process, a typical approach is to model and simulate the SUT, its hardware and its environment to ensure that the specifications of the SUT do not violate the environment assumptions; (ii) the embedded software is tested on the development platform with a simulated environment to ensure that the developed software works according to the environment assumptions and can handle possible environment failures. This is done with either an adapter for the hardware platform that forwards the signals from the SUT to the simulated environment or a simulation of the hardware platform; (iii) another level of simulation is when the actual software is deployed on the hardware platform (or part of the platform, e.g., only the processor) and testing is done with a simulated environment.

The focus of this paper is on the second type (ii) of modeling and simulation in which the actual SUT is used, the environment is simulated, the hardware platform is simulated or bypassed through an adapter communicating with the environment simulator. In our experience of working with two industrial organizations, which were developing RTES for different domains (seismic acquisition

systems and automated bottle recycling machines), this form of testing was highly critical as it enabled early verification of the RTES.

To address the above objective, in this paper, we propose an automated methodology for RTES based on environment behavioral models developed using software modeling standards: Unified Modeling Language (UML) [4], UML Profile for Modeling and Analysis of Real-time and Embedded Systems (MARTE) [5], and Object Constraint Language (OCL) [6]. The main contributions of this paper include an environment modeling methodology and an approach to generate a simulator of the environment from the environment model in a way to enable the automated testing of industrial RTES. As further discussed below, our focus is to devise a *practical* approach in a *system testing* context, and we evaluate both the modeling methodology and simulation generation on two industrial case studies.

Environment models describe both the structural and behavioral properties of the environment. Given an appropriate level of detail, defined by our methodology, they enable the automatic generation of the environment simulator. These models can also be used to generate automated test oracles, which are typically modeled as "error states" that should never be reached by the environment during the execution of a test case. Moreover, the models can further be used to automatically select test cases and sophisticated heuristics are used to automatically do so from the models without any intervention of the tester. To summarize, the only required artifacts to be developed by testers is the environment model and the rest of the process is expected to be fully automated. Incidentally, by using this automated Model-Based Testing (MBT) technology, one of our industrial partners was able to find new critical faults in their RTES.

To support environment modeling in a practical fashion, we have selected standard and widely accepted notation for modeling software systems, the UML and its standard extensions. We use the MARTE [5] extensions for modeling real-time features and OCL for specifying constraints. We have also provided lightweight extensions to UML to ease its use in our context. As we will discuss later, environment modeling is not a new concept. But, most of the approaches use non-standardized notations or grammars for modeling, which makes them difficult to apply from a practical standpoint. Modeling the environment of industrial RTES systems using a combination of UML, MARTE, and OCL has not been addressed in the literature. By using the proposed methodology, the software testers (who are primarily software engineers) can model the environment with a notation that they are familiar with, using commercial or open source tools, and at a level of precision required to support automated MBT. The importance of relying on standards for modeling was confirmed on the two industrial case studies across entirely different domains.

Although code generation from models has been widely studied, the context of black-box RTES system testing poses specific challenges and problems that are not fully discussed and addressed in the literature. For this purpose we present extensions to the *state pattern* [7] specifically aimed at

enabling environment simulation for system testing and define rules for transforming environment models to Java code (the simulator).

To summarize, the fundamental motivation here is that system testers, in many industry sectors, are usually application domain experts but have little or no knowledge of the system design and implementation. Our approach is therefore black-box and does not require the RTES itself to be modeled. It only requires its environment to be modeled at the right level of abstraction and in such a way as to enable effective test automation. The reliance on software modeling standards offers significant advantages, such as the possibility of using (1) different commercial and open source modeling tools (e.g., IBM Rational Software Architect (RSA)[2], Papyrus[3], or Enterprise Architect[4]), (2) notations that many software engineers—including system testers—might already be familiar with and that can be used to also model the SUT, and (3) existing analysis tools (e.g., [8]) that can take such models as input.

The paper is organized as follows: Section 2 shed light on the practical motivations and aspects of the work presented in this paper, setting the context to better justify our approach; Section 3 discusses the related work. Section 4 presents the motivating example that we use throughout the paper to explain various concepts. Section 5 discusses the proposed environment modeling methodology. Section 6 goes into the details of the most important decisions regarding the transformation of models to simulation code, whereas Section 7 presents the case studies. Section 8 discusses the limitation of the proposed work and finally, Section 9 concludes the paper.

2. Practical Aspects

The work discussed in this paper was motivated by the problems faced and practices followed by two industrial organizations that we worked with, namely WesternGeco AS, Norway and Tomra AS, Norway. These two organizations were developing RTES for two different domains; WesternGeco was developing a seismic acquisition system and Tomra was developing automated recycle machines. Both the RTES were developed to run in an environment that enforces time deadlines in the order of hundreds of milliseconds with an acceptable jitter of a few milliseconds in response time. In one of the organizations, testing the SUT on the development platform with a simulated environment was considered to be mandatory before deploying the software on the operational hardware. To achieve this, software engineers were writing application specific simulators directly in Java. Test cases for system level testing were written by hand by the software test engineers and were executed on the SUT with the environment simulator. The research presented in this paper was strongly driven by our investigation of the practical needs of our industry partners which, based on our experience, are shared by many others in numerous industry sectors. Our understanding of these needs is presented in the remainder of this section.

[2] Webpage: http://www.ibm.com/developerworks/rational/products/rsa/, date last accessed: 05/02/2012
[3] Webpage: http://www.papyrusuml.org/, date last accessed: 05/02/2012
[4] Webpage: http://www.sparxsystems.com.au/, date last accessed: 05/02/2012

Manually writing an environment simulator using a programming language (e.g., Java or C) appeared to pose a number of issues, the main one being that software engineers have to develop such simulator at a low-level of abstraction while simultaneously focusing on the logic of the simulator, complex programming constructs (e.g., multiple threads, handling timers), and the handling of test case configurations (when the simulator is used for testing). Making this problem even more acute, over the course of the RTES development, these simulators frequently change due to changes in the specifications of the hardware components.

Typically, modeling and simulation (M&S) approaches focus on simulating hardware components, execution platforms, and natural phenomena in the RTES environment using various simulation tools, such as DEVS [9] and Modelica [10]. These M&S tools support precise simulations of both discrete and continuous system behaviors and are typically based on mathematical models. However, in our context, such M&S tools are not practical, for a number of reasons: (i) *Software engineers*, who are typically in charge of system testing at this level, are often not familiar with such simulation languages. To enable technology transfer in industrial practice, it would be more convenient for them to develop or generate the simulator using a language that they are familiar with, as for example the languages used to program and model the SUT; (ii) These simulation tools do not support automated environment-based testing of RTES software. A number of features must be modeled to enable this kind of testing. For example, the models need to provide information for the automated generation of oracles (to verify whether test cases pass or fail). Furthermore, the simulator needs to interact with a test harness to get appropriate values for various non-deterministic events. The exact occurrences of such events in the environment cannot be determined. These events may follow different probability distributions (e.g., probability of failure of a sensor) or may occur at any time within a given time interval (e.g., a gate at a railroad intersection may take from 5 – 7 seconds to close); (iii) Another issue is that in simulation languages such as ModelicaML [11] and DEVS [12], since they were developed for a different purpose, there are limitations regarding the interactions of the simulator with the production code of the RTES (e.g., handling of operating system resources, such as inter-process communication with the production code over TCP/UDP). Such an interaction is a requirement for the type of testing we deal with in this paper, since the environment simulator has to interact with the *actual RTES production code* to receive stimuli and to send responses. In dealing with such interactions, we do not want any constraint regarding the programming language in which the RTES is written.

The modeling methodology presented here provides an automated model-based approach to derive environment simulators, test cases, and test oracle, taking into consideration all the practical aspects described above, which are common place in many industrial environments. The only major input required are the environment models describing the structure and behavior of the environment. Since the intended users are software engineers, we chose standard software modeling languages for environment modeling with the aim to make the modeling methodology as simple as possible. This

paper discusses the methodology for modeling the environment based on the selected modeling standards and a specific profile and furthermore describes the process of simulation generation from those environment models. It does not, however, address in detail test generation and test oracles.

3. Related Work

In this section we discuss the related literature in the areas of (1) modeling and simulation for RTES Testing, (2) environment modeling and environment model-based testing of RTES, and finally (3) code generation approaches from UML state machines and class diagrams.

3.1. Modeling & Simulation for RTES Testing

As discussed earlier, based on the goals of testing of RTES, the SUT, the hardware platform, the environment, or their combinations can be modeled and simulated.

At the early stages of the development process for RTES, a typical approach is to model and simulate the SUT, its hardware and its environment. The aim is to ensure that the model of the SUT complies with the requirement specifications and does not violate the environment and hardware assumptions. This approach is sometimes also referred as "model-in-the-loop" simulation [3, 13, 14].

Another level of simulation for testing is when the actual executable software is deployed on the real hardware platform (e.g., electronic control unit) and their combination is tested with a simulated environment (e.g., with the simulation of plant model [3]). This approach is generally referred to as hardware-in-the-loop testing [15, 16]. Typically, a prototype of the hardware platform is used at this stage. A variation to hardware-in-the-loop testing is the case where only the actual processor is used during testing and rest of the hardware and environment are simulated. This variation is referred to as processor-in-the-loop testing [17].

Before the hardware or the processor is available, the embedded software can also be tested on the development platform (e.g., Linux or Windows-based machine) with a simulated environment and hardware platform. This is typically done to ensure that the developed software works according to the environment assumptions and behave appropriately in hazardous or abnormal situations. This is mostly referred to as software-in-the-loop simulation [3, 13].

Existing modeling and simulation languages have been developed and are widely used for the first three types of simulations. In these cases the environment simulation needs to interact with the actual hardware or its simulation. In such cases, precise simulation of both discrete and continuous phenomena is required and is typically based on mathematical models.

In this paper, we target a slight variation to the typical software-in-the-loop simulation. We only model and simulate the environment and use an adapter for the hardware platform that forwards the signals from the SUT to the simulated environment. Our research problem definition is motivated primarily by the practical needs of our industrial partners (Section 2) but it is expected to be relevant in many other industrial environments developing similar RTES. Other approaches that do testing with a similar focus are discussed next.

3.2. Environment Modeling and Environment Model-based Testing

In this subsection, we will discuss various environment modeling and model-based testing approaches discussed in the literature. Kishi and Noda [18] present an approach for modeling the environment of an embedded system using an aspect-oriented modeling technique. Karsai *et al*. [19] propose a new language for modeling the environment of an embedded system. Choi *et al*. [20] use annotated UML class and sequence diagrams for modeling and simulation of environment. Kreiner *et al*. [21] present a process to develop environment models for simulation of automatic logistic systems and its environment. Axelsson [22] evaluates how UML can be used to model real-time features and provides extension to UML for modeling of real-time systems and their environments. Gomaa [23] discusses the use of a context diagram for modeling the relationship between an RTES and its external entities. Friedentahl *et al*. use the concept of SysML block diagram and activity diagrams to represent the system and its interfaces with environment components [24].

There are a few works reported in literature that discuss testing of RTES based on its environment and without considering the hardware platform or its simulation. Auguston *et al*. [25] discuss the development of environment behavioral models using Attributed Event Grammar for testing of RTES. The behavioral models contain details about the interactions with the SUT and possible hazardous situations in the environment. These models are then traversed to obtain various test scenarios. The approach is applied on a simulation of the RTES specifications. Heisel *et al*. [26] propose the use of a requirement model and an environment model using UML state machines along with the model of the SUT for testing. Adjir *et al*. [27] discuss a technique for testing RTES based on a model of the system and intended assumptions in the environment in Labeled Prioritized Timed Petri Nets. Larsen *et al*. [28] propose an approach for online testing of RTES based on time automata to model the SUT and environmental constraints. Bousquet *et al*. [29] present an approach for testing of synchronous reactive software by representing the environmental constraints using temporal logic.

To summarize, there are approaches in literature that deal with modeling the environment of a system for various purposes. Most of these approaches are only limited to modeling the static structure of the environment, as they do not focus on test automation. The approaches that deal with modeling of behavioral aspects either use notations with which software engineers are not familiar, or provide extensions for environment modeling that in a non-standard way.

All environment modeling approaches aimed at supporting testing in literature, except the one by Heisel *et al*. [26], use non-standard languages for modeling. This work, however, models the concepts of probabilities and time using non-standard notations, without relying on UML extension mechanisms. Furthermore, most of the works on environment model-based testing, model both the SUT and the environment, which does not fit our purpose: black-box, system testing. Moreover, none of these works simulate the behavior of the environment. To be able to execute different possible behaviors of the environment based on its interaction with SUT for system testing, a simulator is essential. Generating a random set of scenarios from the environment models, as done by Auguston *et*

al., [25] provides a limited coverage of the environment model. Last but not least, none of the relevant works assess their environmental methodology on an actual RTES system, which we believe is a requirement to assess the credibility and applicability of any MBT approach.

3.3. Code Generation from UML Classes and State Machines

There is no reported approach in the literature that generates RTES environment simulators from UML models or their extensions. As we discussed in Section 3.2, though modeling and simulation languages and environments have been proposed, they do not fit our purpose of black-box system testing (Section 2).

Generating code from state machines is not a new problem. Even though a number of works are reported in the literature (e.g., [30], [31]) and a number of tools are available that generate code from state machines (e.g., SmartState [32], IBM Rhapsody [33]), no technology was developed having the required features for black-box testing based on environment models, such as non-determinism (a common feature of the environment) and test-specific behavior (e.g., modeling illegal or unsafe environment states). The use of standards is also an important requirement for our modeling methodology, as discussed earlier. The modeling standards that we selected for our methodology are supported by a wide range of tools and support is available for training. The existing code generation tools and techniques discussed in the literature are focused on generating system code and not environment simulators. They are not applicable to our purpose because of several reasons. Following, we discuss the most important ones:

1) The models need to capture specific states of failures in the environment components (the "failure states", e.g., a sensor break down). The environment simulators then must be able to simulate these situations, which occurrences can be controlled by the test cases. Similarly, the models need to capture information of situations in the environment that should never happen if the SUT is correct (which we call the "error states"). Such information is required to derive test oracles and to guide testing strategies.
2) Since we simulate the environment for testing, the generated code is strongly coupled with the test harness, which is not the case with existing simulators.
3) The generated simulator includes code that collects execution information to guide testing strategies. Such information is used in heuristics to generate test cases that are effective to reveal faults. The heuristics that we use are specific to our modeling and testing methodology, existing tools for code generation do not provide this support.
4) Extensions to existing tools are not feasible in most cases as some of them are proprietary and others are not based on model-driven engineering standards, such as the Eclipse Modeling Framework.

The original state pattern, discussed in [7], provided a design pattern to implement state-driven behavior but did not address a number of important features present in UML 2.x state machines, e.g., concurrency, time events, change events, and actions. A number of extensions for the pattern have

been discussed over time to handle missing features (e.g., [34, 35]). Most of these extensions are focused on increasing the understanding and usability of the code obtained using the state pattern to support programmers (e.g., [35]) and are not very useful in our context where the code is automatically generated from the models. Chin and Millstein [36] propose an extension to the state pattern for handling state behavior in inherited sub-classes. Holt et al. [37] propose an extension for handling state and transition actions and report its application on an industrial case. Palfinger [38] provides an extension to the state pattern that allows extension of object's behavior at runtime by using a mapper class. In our approach, this was not applicable as we do not require the addition of new behavior during the execution of environment components. The work in [39] extends the state pattern to supports hierarchical state machine and time events. We handle time in a similar way, i.e., by using a separate timer class that calls `timeout()` on the context object in case of a timeout. The approach in [39] does not handle parallel regions and change events, does not provide details on handling actions, and does not provide support for the test-related features required by our approach. Overall, none of the extensions of the state pattern in the literature completely meets the needs for RTES environment simulation to support system testing.

We could have used some optimizations to improve ease of understanding and modification, and cleaner code generation, but these would not have had a large impact as the generated source code is not visible to the end-user and is only provided as an executable archive. Furthermore, there are optimizations in the literature to improve the performance of the generated code (e.g., minimizing the number of running threads to avoid overheads due to context switching). However, in our framework (as we have explained in details throughout the paper) we do not need to optimize performance. The generated simulators are used only for testing purposes, and each test case runs on a different process, lasting from a few seconds to a few minutes (depending on the RTES). As long as the environment simulators can behave as expected (e.g., providing the right stimuli at the right time), this would be sufficient for our testing purposes. This was the case for all our case studies.

3.4. Summary

To summarize, this paper differs from existing works in several of the following ways: (1) It provides an environment modeling methodology based on international software engineering modeling standards (UML 2.x, MARTE, OCL) that is dedicated to black-box, RTES system testing. The targeted RTES have complex environments and have soft-real time constraints in the order of hundreds of milliseconds pertaining to the response time of the SUT and operations of the environment. This is the first work focusing on such methodology, allowing the modeling of important concepts for testing such as modeling non-determinism and oracle information, while relying only on light-weight, standard extensions of UML (i.e., by defining a UML profile); (2) It provides an approach to generate simulators, based on environment models developed using the proposed environment modeling methodology, addressing the specific needs of black-box system

testing; (3) Regarding simulator generation, the paper provides extensions to the state pattern that handle time-related features and various UML 2.x features that were not previously discussed in the literature, including handling of change events and concurrency; (4) Unlike most of the works reported in the literature, this paper assesses the proposed methodology and simulator generation on two industrial RTES, which we believe is a requirement to assess the applicability of any test automation approach.

4. Motivating Example

To motivate and explain our modeling methodology and simulator generation strategy for RTES, we take as example a subset of one of our industrial case studies. Note that we have sanitized the information due to confidentiality restrictions. This example is an Automated Bottle Recycling System developed by Tomra AS, Norway. It is representative of the type of RTES we are targeting in this paper: it has soft-real time constraints in the order of hundreds of milliseconds pertaining to the response time of the SUT and operations of the environment (e.g., the sorting of item should be done within a couple of minutes after an item is inserted).

The portion of the case study considered in this paper is focused on the important functionality of sorting the recycled items to their proper storage locations (or destinations). Users insert the items to be recycled inside the front-end of the recycling system, called the *Reverse Vending Machine (RVM)*. The items can be of three different types for the subset we are discussing: plastic bottles, cans, or glass bottles. The *RVM* forwards the items to the *Sorter*, which is a sorting arm (we only consider a simplified backroom with a single sorting arm). On its way from an *RVM* to *Sorter*, an item can be lost if it is not detected in time or if it falls from the moving belt. The *Sorter* can move in three directions (each leading to a specific destination) and its movement is controlled by a *Sorting Controller*.

The *Sorting Controller* is the system under test in our case study. The *Sorting Controller* receives information of the type of the item inserted from the *RVM* and when it is supposed to reach the *Sorter*. The *Sorting Controller* is responsible for moving the *Sorter* in a position that leads the items to their appropriate destinations. There can be different destinations based on the type of items. Plastic bottles and cans are placed in their appropriate bins, whereas the glass bottles are placed in the crates. The *Sorting Controller* should prevent certain erroneous situations from happening. For the subset of the case study discussed in this paper, we consider two such situations: (i) when an item is not correctly sorted and it goes to a wrong destination (for example, a plastic bottle going into a bin of cans) (ii) when an item reaches the *Sorter* while it is still moving.

5. Environment Modeling Methodology

If environment models are to be used for testing RTES, they should not only be sufficiently detailed, but should also be easy to understand and modify as the environment and RTES evolve. To handle the complexity of realistic RTES environments, the modeling language should have provision

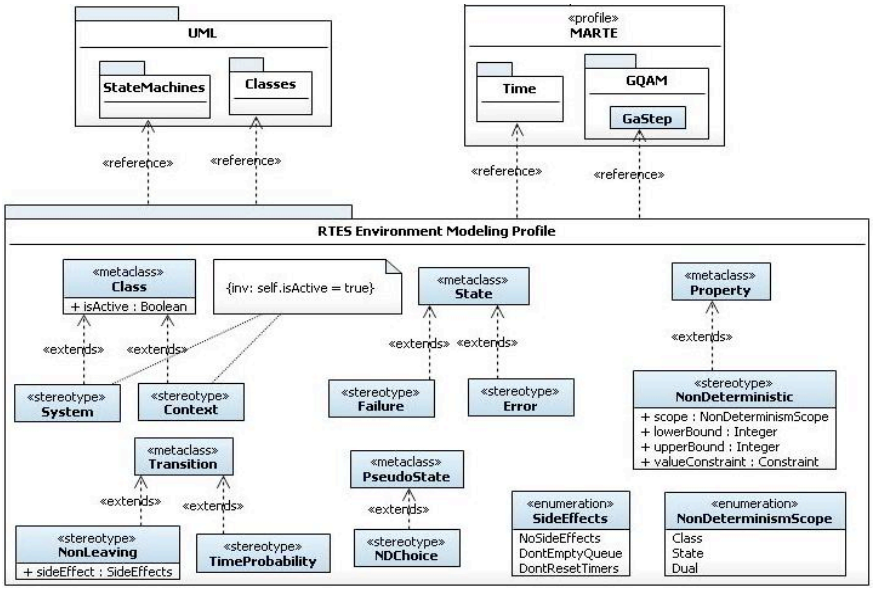

Fig. 0. RTES Environment Modeling Profile

for modeling at various levels of abstraction. The modeling language should also be well-defined for the tools to analyze the models and for the humans to accurately understand them. The language should also provide features (or allow possible extensions) for modeling real world concepts, real-time features, and other concepts, such as non-determinism, required by the environment components. The UML, MARTE profile, and the OCL together fulfill the important requirements of an environment modeling language.

Even though we are using the same notations to model the environment that are used for modeling software systems, it is important to note that the methodology for environment modeling is significantly different from system modeling. While modeling our industrial cases, we abstracted the functional details of the environment components to such an extent that only the details visible to the SUT were included. An environment of a RTES typically features a number of non-deterministic events (e.g., breakdown of a sensor), which must be modeled. Such events are not common when modeling the internal behavior of a system.

To model RTES environments, we have developed a profile that provides support for modeling various concepts central to our methodology and highlights the subset of UML/MARTE that is required for such modeling. For testing the system based on its environment, the behavioral details of the environment are as important as its structural details. Structural details of the RTES environment are important to understand the overall composition of the environment (e.g., number and configuration of sensors/actuators), the characteristics of various components, and their relationships. We choose to model these details in the form of a Domain Model developed using UML class diagrams annotated with our defined profile. The behavioral details of environment components are required to specify the dynamic aspects of the environment, for example, to determine the possible environment states, before and after its interactions with the SUT, and to specify the possible

interactions between the SUT and its environment. For behavioral details, we used the UML State Machines augmented with the MARTE profile and our defined profile.

In the kind of testing this paper addresses, the focus is on the interactions of the RTES with the components in its environment, i.e., what are the possible inputs/outputs to/from the RTES from/to these components at any given point in time? How does the RTES behave in abnormal situations, such as a hardware failure in any of the environment components? A test case for a RTES would typically consist of a sequence of actions from the user(s), signals from/to sensors/actuators, and possibly hardware component breakdowns. This would correspond, in our context, to *non-deterministic events* that can happen during the environment simulations.

In the following sub-sections, we discuss the environment modeling profile that we developed followed by the guidelines for modeling domain and behavioral models that were developed based on our experience of modeling two large-scale industrial RTES - a marine seismic acquisition system and an automated bottle recycling system.

5.1. Environment Modeling Profile

Our goal was to model the environment based, to the extent possible, on the standard UML and its existing extensions. We applied the standard notations and based on our needs for those case studies, where required, we provided light weight extensions to UML. In this section we will discuss the subsets of UML and MARTE that we used and the lightweight extensions that we have provided for environment modeling. From a practical standpoint, it was important to identify these subsets for the methodology, since the UML and MARTE standards are very large and most organizations would be reluctant to adopt such large notations.

Developing UML profiles is a way to provide lightweight extensions to UML that do not conflict with its original semantics. To model an RTES environment, generate its simulator, test cases, and obtain test oracle from these models, we need more specific notations than what the standard UML provides. We provided extensions to the standard UML class diagram and state machine notations in the form of a profile. The profile also resolved various semantic variation points left open by the standard (discussed later in Section 6.5) to address our specific needs. Fig. 1 depicts a profile diagram for our proposed RTES environment modeling profile. The profile defines a set of stereotypes for modeling our methodology specific features on UML classes and state machines. It also shows the subset of MARTE that the profile is using, i.e., the Time package and the concept of *GaStep* from the Generic Quantitative Analysis Modeling (GQAM) package. The Time package allows the software engineers to model various time related features, such as timed events and action durations [5]. This small subset of UML and MARTE was sufficient for modeling our two industrial case studies for the purpose of automated black-box testing.

5.2. Domain Modeling

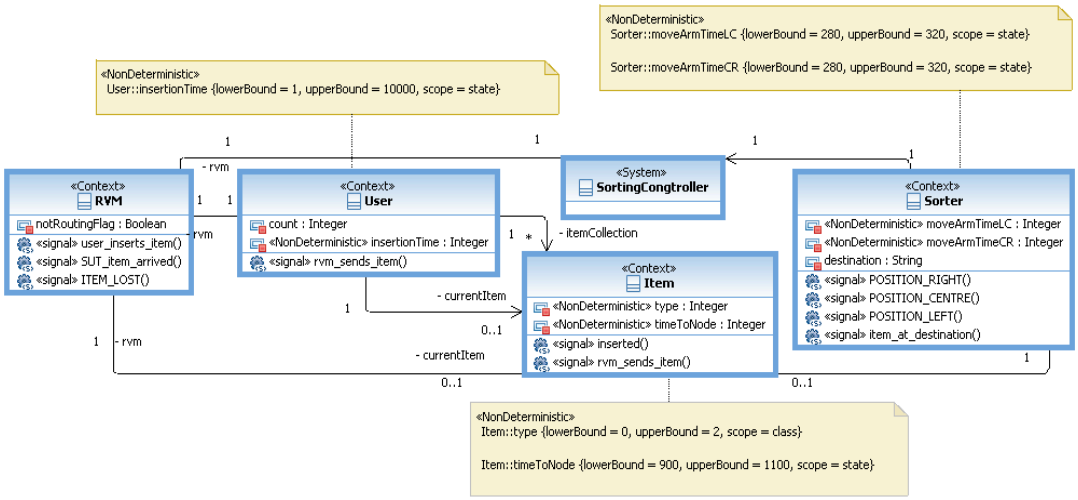

Fig. 0. Domain Model of Sorting Machine Case study

Our environment modeling methodology for system testing requires the modeler to create an environment domain model that captures relevant structural details of the environment including the various components of the environment, their cardinalities, characteristics, and relationships. The domain model is developed using the UML class diagram notation.

The various components modeled in the domain model together form the overall environment of the SUT. This means that all these components (their instances) will run in parallel with each other. The domain model represents various possible forms that the environment of RTES can take. Each component in the domain model can have a number of instances in the RTES environment. The information about the number of possible instances of a component in the environment is modeled as cardinalities on the associations between different components in the domain model. Therefore, the domain model can be used to obtain a number of potential configurations of the environment. To restrict the possible forms an environment of an RTES can take, OCL constraints can be specified. These constraints can for example be used to restrict the possible combinations of environment components or to restrict the possible values of attributes.

The domain model for the Sorting Machine case study is shown in Fig. 2. *Sorting Controller* in the domain model is the SUT and the components *RVM*, *User*, *Sorter* and *Item* are the environment components. All the environment components are considered to be active objects, i.e., having their own thread of execution, and communicate with each other through signals. Each environment component in the domain model can have multiple instances. For example, in the domain model, shown in Fig. 2, Item is represented as one environment component, but during simulation it can have multiple instances. The number of instances to be created, which we refer to as an *environment configuration*, is determined based on the cardinality of relationships, i.e., in this case the cardinality of the association between *User* and *Item* with the role name *itemCollection* and the OCL constraints restricting the possible combination of environment components. In the motivating example we have restricted the possible number of items a user can enter to be less than 100. This is shown as an OCL constraint in Fig. 3. A valid *environment configuration* for this example is a single *RVM,* a single

Sorter and a *User* with three *Items*. A test case in our context is a combination of a setting of the simulator for the *non-determinism* in the environment models (e.g., a specific time at which a sensor stops working), which we call *simulation configuration* and an *environment configuration*. A test case uses these settings for a particular simulation run. During testing, the selected test strategy decides the way these configurations for a test case are generated by the *Test Framework* (e.g., random values for *simulation* and *environment configurations* when using random testing).

Note that the domain model that we develop is different from the ones commonly discussed in literature (e.g., [40]). The components represented as classes in the environment domain model will not necessarily relate to software classes. They may correspond to systems, users and concepts related to various natural phenomena. Domain modeling here is not a starting point for software analysis. The identification of components in the domain model, their properties, and their relationships is also different from what is commonly done for software analysis. Following, we further discuss various guidelines for modeling the structural details of a RTES environment.

5.2.1. Environment Components to be Included

Initially, all the environment components that are directly interacting with the SUT are included in the domain model. Then, each of these components is further refined to a level where we are certain to cover the important details for simulating the environment needed to test the SUT. If at any time the behavior of an environment component is getting too complex, when possible, we can decompose the component and divide its behavior into multiple concurrent state machines. This is especially useful if a component can be divided into components that are similar to existing components, so that we can specialize existing state machines. The environment components in the domain are stereotyped with «Context». The environment components are modeled as active objects and can communicate with each other and the SUT through signals.

5.2.2. Relationships to be Included

All those associations representing the physical or logical relationships among various environment components, or that were needed for components to communicate, should be included. A number of components in the environment might be similar to each other (e.g., various types of sensors). It is useful to relate these components (and their behavior) using the generalization/specialization relationship for simplifying the model, as our experience shows that such domain models get highly complex. For example, in the sorting machine case study, we modeled the association of the *SortingController* with the *Sorter*, which is controlled by the board.

5.2.3. Properties to be Included

From all properties that may characterize environment components, it is important to include only

```
context User inv:
self.itemCollection→size() > 0 and self.itemCollection→size() < 100
```

Fig. 0. An example OCL constraint

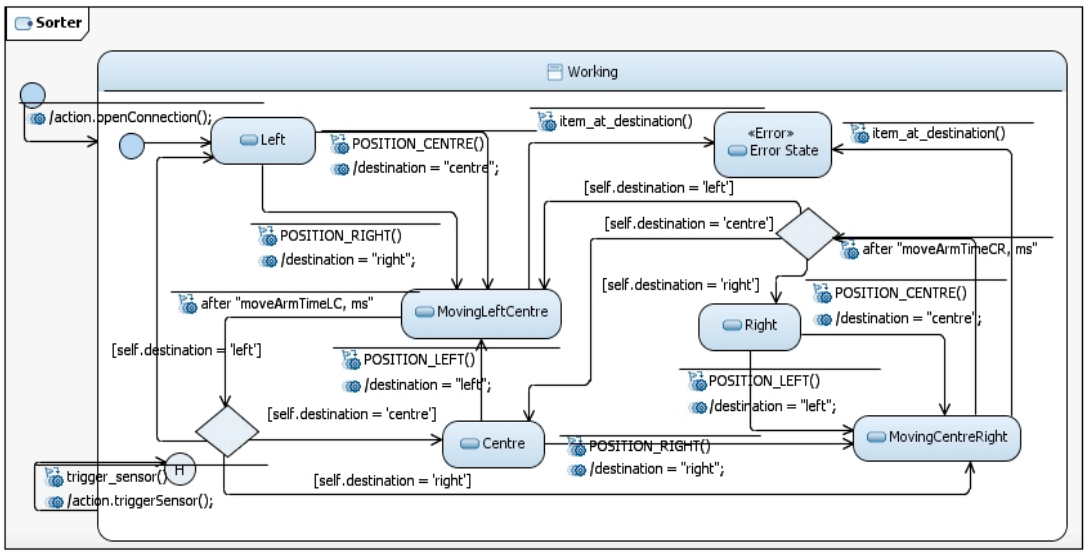

Fig. 0. State machine for the Sorter component

those properties that are visible to the SUT (or have an impact on a component that is visible to the SUT). These may include attributes that have a relationship to the inputs of the SUT, that constrain the behavior of a component with respect to the SUT, or that contribute to the state invariant of a component that is relevant to the SUT. For example, in Fig. 2, the attribute *type* of *Item* is used by the *SortingController* to move the *Sorter* in appropriate position before the items arrive.

By using the profile, it is also possible to leave the decision of selecting exact values for properties of the environment components till the time of testing (where it is decided by the *simulation configurations*). This concept is modeled by assigning «NonDeterministic» to the properties of environment components. This stereotype has three properties: an *upper bound*, a *lower bound*, a *valueConstraint*, and a *scope*. The upper and lower bound specify the possible range of values that an integer property of an environment component can take during simulation. This is provided to ease the modeling of time events' bounds. Alternatively, an OCL constraint can be provided as a *valueConstraint* that restricts the possible values that an environment property can take. This constraint can, for example, be used to restrict a string property to certain specific values.

As shown in Fig. 1, the scope property can have three possible values: *Class*, *State*, or *Dual*. If the value is set to *Class*, the properties of the environment component instances are initialized with a value obtained from *simulation configuration* only once when the instances are created. If the value is set to *State*, the values are obtained whenever there is a state change in an instance. If the value of *scope* is set to *Dual*, then a value is obtained for this environment component's property from the *simulation configuration* when an instance is created and the property is reassigned a value when there is a state change in the instance. For example, in Fig. 2, the property *type* of *Item* is a non-deterministic variable with the scope *Class* and its value is initialized based on a simulation configuration when an instance of *Item* is created.

5.2.4. Modeling the SUT

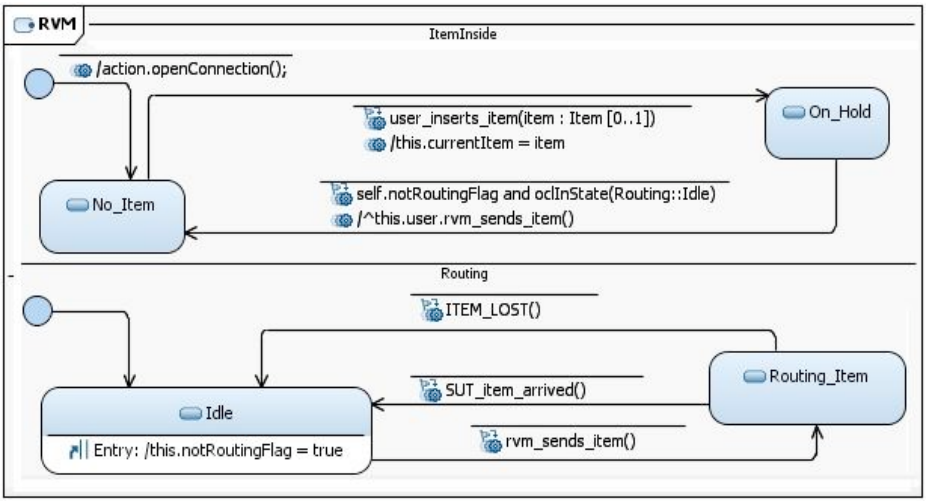

Fig. 0. State machine of the RVM Component

It is important to include the SUT in the environment domain model, so that its relationship with the other environment components can be specified. It is also useful to include the details of signal receptions by the SUT from other environment components. The SUT is stereotyped as «System». The stereotype was used initially by Gomaa [23] to refer the system in a context diagram. Since, our goal is *software-in-the-loop* testing, the SUT modeled in the domain model represents the SUT and its execution platform as a single component.

5.3. Behavior Modeling

For each environment component in the domain model that has a behavior affecting the SUT, our methodology requires to create a state machine representing this behavior. The state machine captures such behavior at the level of abstraction that is visible to the SUT. The state machines are developed using UML 2.x state machine notation and concepts, MARTE real-time extensions, and our profile to assist in modeling the environmental aspects of RTES. The MARTE profile is used to model the features related to time and a form of non-determinism. As discussed earlier, during simulation, the instances of the environment components run in parallel to form the environment of the RTES. They can send signals to each other and to the SUT. We can also view the environment as having one state machine with orthogonal regions, one for each component. Fig. 4 - Fig. 7 show the state machines of the four environment components of our motivating example. Note that the diagrams are developed in IBM RSA, which adds some additional symbols to the triggers and effects in the state machines. A change event is not shown with a *when* keyword as for example in the transition from *On_Hold* to *No_Item* in the *ItemInside* region of *RVM* state machine shown in Fig. 5. All the guards in the state machines have the corresponding environment components as the context for OCL constraints. Following, we discuss the details of the methodological guidelines for modeling behavior of the environment components.

5.3.1. Identifying Stateful Components

Components whose states either affect the SUT or are affected by the SUT should be modeled with state machines.Overall, the environment should be modeled in a way that enables, after the

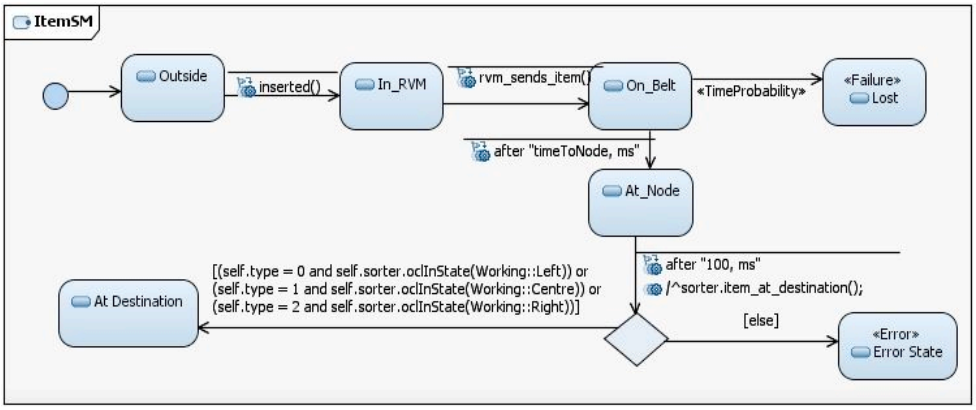

Fig. 0. State machine of *Item* environment component

initialization and provision of *simulation* and *environment configurations*, the full simulation of the interactions with the SUT. All the environment components shown in Fig. 2 are stateful components of the sorting machine case study. For example, the *Sorter* component was modeled as stateful since it receives signals from the *SortingController* and reacts differently based on its current state.

5.3.2. States to be Included

It is important to determine the right level of abstraction for a component state machine. If we want to precisely model the behavior of an environment component, this might lead to a large number of states. We are, however, only interested in state changes that have an impact on the SUT. A single state in an environment model state machine may correspond to a large number of concrete or physical states. For example, in the sorting machine, the states of *Item* that we modeled were all related to its movement through the sorting machine whereas its other possible states were not of interest as an environment component of the *SortingController*.

A state in a UML state machine can be a simple state, a composite state (i.e., containing substates) or it can be a submachine state. UML state machines can also have multiple orthogonal regions. The concept of orthogonal regions is particularly useful in environment modeling as one environment component can in reality be composed of multiple sub-components. For example, *RVM* in our motivating example is composed of two sub-components: an item feeder that handles item insertion and a conveyer that is responsible for routing the items. From the perspective of the SUT (*Sorting Controller*) it is not important to distinguish these two components as it sees *RVM* as a single component. For the *RVM,* to completely simulate the behavior visible to the *Sorting Controller*, it must manage the movement of items on the conveyer in parallel to handling items in the feeder. From the *RVM* point of view, functionality must be provided for both of these components, conveyer and feeder. Therefore this information is modeled as two orthogonal regions of the *RVM* (named *ItemInside* and *Routing*) in the state machine shown in Fig. 5. In addition, according to our modeling methodology, failure behavior of a component that is independent of its nominal behavior can also be modeled as a separate orthogonal region.

5.3.3. Modeling Users in the Environment

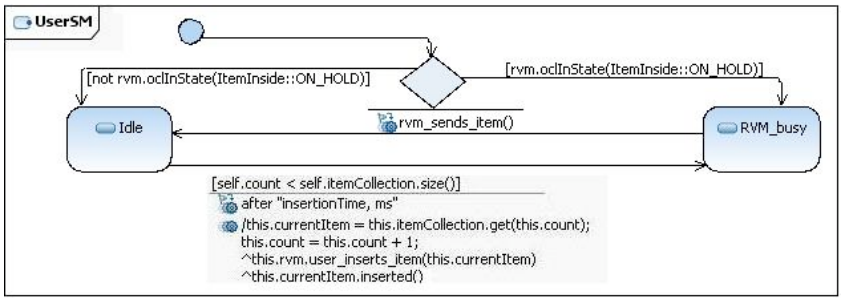

Fig. 0. State machine of User

Generally, for software system modeling, users are only modeled as sources of inputs and data. The behaviors of users with respect to the system are mostly not considered. In the environment modeling methodology, it is useful to model the behavior of users in the environment to have a control over the inputs/outputs of the various components or the SUT. If a user participates in multiple roles, it is useful to model each role a user plays as a separate component. In the motivating example, we modeled the persons who enter the items for sorting as a *User* environment component (state machine shown in Fig. 7). In certain cases it can be interesting to model both the expected and unexpected behavior of users using the proposed methodology. Overall, the behavior of a user in an RTES environment is modeled using the same notations as any other environment component.

5.3.4. Modeling Events

When using UML 2.x state machines for environment modeling, only three types of events are required to be modeled: signal events, time events, and change events. Call events are not required since the components in the environment represent active objects and communicate asynchronously. OCL is used to model guards on transitions and conditions in the change events. For example, Fig. 4 shows the state machine of the *Sorter* component. As discussed earlier, a *Sorter* can be at three different positions. This is represented by the three states, *Left*, *Centre*, and *Right*. Movement between these states is represented by the outgoing transitions from these states to the two movement related states: *MovingLeftCentre* and *MovingCentreRight*. For the *Sorter* to move from *Left* to *Center* it needs to transition first from *Left* to *MovingLeftCentre*, which is triggered on receiving a signal event *POSITION_CENTRE()* from the SUT (*Sorting Controller*). A transition from *MovingLeftCentre* to *Centre* state is triggered by the time event *after "movingArmTimeLC, ms"*, where *movingArmTimeLC* is the name of a non-deterministic property of *Sorter* and *ms* is the unit of time, milliseconds. This transition is only triggered if the guard on the transition, written in OCL (*self.destination = "centre"*), is true. An example of a change event can be seen in the state machine of the *RVM* component (Fig. 5) in the *ItemInside* region on a transition from *On_Hold* to *No_Item*. The transition has an effect `^this.user.rvm_sends_item()` written in Java, which we chose as the action language as further discussed later.

The MARTE *TimedEvent* concept is used to model all timeout transitions, so that it is possible for them to explicitly specify a clock (if needed). Each environment component may have its own clock

```
1. import simula.embt.commons.*;
2. public class SortingMachineActionCode implements ExternalCode
3. {
4.      private int port;
5.      private String address;
6.      private Connection con;
7.      private IActiveObject sorter;
8.      public SortingMachineActionCode(Object[] args)
9.      {
10.             port = (Integer) args[0];
11.             address = (String) args[1];
12.     }
13.     public void openConnection() throws Exception
14.     {
15.             con = TCPConnection.openConnection(address, port, 1000);
16.             String msg = con.readMessage();
17.             sorter.receiveSignal(msg, null);
18.     }
19.     public void triggerSensor() throws Exception
20.     {       con.sendMessage(Signals.TRIGGER_SENSOR);        }
21.     @Override
22.     public void startExecution(IActiveObject ao)
23.     {       sorter= ao;                                     }
24.     @Override
25.     public void stopExecution() throws Exception
26.     {       con.stopExecution();                            }
27. }
```

Fig. 0 Excerpt of ExternalActionCode for the Sorter component

or multiple components may share the same clock for absolute timing. The clocks are modeled using the MARTE's concept of clocks. If no clock is specified (as in the case of motivating example), then by default the notion of time is considered to be according to the physical time. Specifying a threshold time for an action execution or for a component to remain in a state is possible using the MARTE *TimedProcessing* concept. This is also a useful concept and can be used, for example, to model the behavior of an environment component when the RTES expects a response from it within a time threshold.

The proposed environment modeling profile allows the modeler to apply three stereotypes to transitions in the state machines: «NonLeaving», «TimeProbability», and «gaStep» (defined by MARTE profile). Following, we discuss the stereotype «NonLeaving», whereas the other two stereotypes are related to non-determinism and will be discussed later (Section 5.3.7).

By default whenever a component transitions from one state to another (i.e., transitions that are not internal), its event queue is emptied. To control the effect of a transition on the event queue and timers of the context component, we defined the stereotype «NonLeaving». In other words, this stereotype is a way to give more control to the modeler over the internal handling of the queues and timers. The stereotype has a property called *sideEffect*, which can have three possible values: (1) *NoSideEffects*, to denote that the transition should have no side effects on the source state. A transition with this stereotype will result in no alteration of the event queue and the various timers in the environment component; (2) *DoNotEmptyQueue*, which will result in no alteration of the event queue, but will reset the timers; (3) *DontResetTimers* where the queue is emptied, but does not reset the timers.

5.3.5. Modeling Actions & Action Durations

In our methodology, we chose Java as the action language for writing actions. The decision to choose Java as the action language at the model level is due to the current lack of tool support for the UML action language (ALF) [41] at the time of our tool development. Moreover, the testers of the SUT are expected to be more familiar with Java (consistent with our experience of applying the approach in two industrial contexts), rather than with a newly accepted standard language.

In the environment models, actions can be written in two places. Simple actions can be written inside the models, e.g., in the *RVM* state machine (Fig. 5), the simple assignment action of the transition from *ItemInside::No_Item* to *ItemInside::On_Hold* (i.e., `this.currentItem = item`) is placed directly as an effect. Relatively complex actions and communication related details are written in a separate source file and are referred to as the *external action code*. Calls to methods of external action code classes are simply made by using the keyword *action*: *action.<method name>*. For example, in Fig. 5 the effect of the initial transition in the *ItemInside* region is in fact making a call to `openConnection()` in the action class corresponding to *RVM* (see line # 13 in Fig. 8).

External action code is the code that is to be written manually by the tester in a separate source code file, to communicate with the SUT and compute complex effects (e.g., action code computing continuous physical effects could also be generated by modeling and simulation tools, such as Modelica [10]). An example for the type of external action code is signals transmitted to the SUT over a UDP/TCP communication layer.

An excerpt of the external action code for the *Sorter* component is shown in Fig. 8 (line # 13 and line # 19). The action code for the two messages sent to the action object in the state machine of *Sorter* (Fig. 4) - `openConnection()` and `triggerSensor()` – can also be seen in the excerpt shown. Class `TCPConnection` is part of the communication library that we used. The method `triggerSensor()` simply forwards the signal to the SUT over the TCP connection. The decision to keep such action code separate was made to avoid cluttering the models with unnecessary details and to allow developers to write this code in a familiar programming tool. It was also important to keep the communication related information separate to avoid changing the models in case of changes in the communication mechanism. For example, if we want to change the communication from TCP to UDP, the only change will be in the external action code classes. Given a communication layer, even if the simulator is generated in Java, there is no particular restriction on the programming language in which the SUT is implemented.

To provide a mapping between the environment components and corresponding classes containing the external action code, an *External Code Mapping* file is provided by the modeler. Fig. 9 shows an excerpt of this file for the sorting machine case study. The file contains the mapping details of external action code and environment components for the two components (*Sorter* and *RVM*) that are

tomra.embt.env.Sorter	tomra.embt.action.SortingMachineActionCode
tomra.embt.env.RVM	tomra.embt.action.SortingMachineActionCode

Fig. 0 Excerpt of External Code Mapping File for the Sorting Machine case study

communicating with the SUT. For the other two classes, no action class was required. This information could have also been placed in the models, but we decided to keep things related to external action code separate from the models, so that it can be changed without affecting the models.

Specifying a time threshold for an action execution or for a component to remain in a state is possible using the MARTE *TimedProcessing* concept. This is also a useful concept and can be used, for example, to model the behavior of an environment component when the RTES expects a response from it within a time threshold. Though in our case studies we did not face a situation where we needed to model action durations, the methodology supports this feature.

5.3.6. Modeling Error & Failure states

Two of the important features that are modeled in the state machines of environment components are the *Error* and *Failure* states. Failure states represent possible failures in the environment of SUT, e.g., hardware failure in the components. These states are required to test the robustness of the SUT when confronted to failures in the environment components. The failure states are modeled with the «Failure» stereotype. Failures that are independent of any specific aspect of an environment component's behavior (e.g., a hardware failure that can occur in any component state) are typically modeled using separate parallel regions within the state machine of the environment component.

Error states are the states of the environment that can only be reached due to faulty behavior of the SUT. These states are conditions that should never happen in the environment, otherwise indicating that the SUT is faulty. For example, a *Sorter* should never receive an *item* while it is moving and when there are no simulated failures in the hardware of the environment components, all items should always be delivered to the correct destinations based on their types. It is the responsibility of the *Sorting Controller* (SUT) to make the *Sorter* reach the appropriate position before an *item* reaches it. Otherwise, this would mean that there is a fault in the implementation of the *Sorting Controller*. This behavior of *Sorter* is modeled in the state machine shown in Fig. 4 as an error state, which is labeled with «Error». Error states are key oracle information that is used during the test execution of the SUT. By modeling erroneous situations as states, the methodology allows modeling of erroneous situations due to violation of temporal constraint (modeled as time transitions leading to error states), due to illegal change in the state of the environment (modeled as transitions leading to error states triggered by change events), and due to erroneous signal receptions (modeled as a transitions leading to error states triggered by signal events).

5.3.7. Modeling Non-Determinism

Non-determinism is a particularly important concept for environment modeling and is one of the fundamental differences between models for system modeling and models for environment modeling. Following we discuss different types of non-determinism that we have modeled for our case studies.

For a number of RTES environment components, specifying the exact values for timeout transitions is not always possible. To model their behavior in a realistic way, it is often more

appropriate to specify a range of values for a possible timeout, rather than an exact value. Moreover, the behavior of humans interacting with the RTES is by definition non-deterministic. For modeling these behaviors, the modeler can add an attribute in the environment component and label it with the stereotype «NonDeterministic». The stereotype allows the modeler to provide an upper and lower bound values by directly specifying them in the properties or by using an OCL constraint to restrict the possible set of values of the attribute. This attribute can then be used as a parameter of a time event. In the state machine of the *User* (Fig. 7), the transition between the state *Idle* and *RVM_Busy* is modeling the behavior that this transition is non-deterministic and that the user can insert the next item with a delay ranging from 1 to 10,000 milliseconds. The information constraining the values is provided in the domain model by applying the stereotype «NonDeterministic» on the *insertionTime* attribute (see Fig. 2). The attribute is then used as a parameter to the time event on the transition. The actual value (between the range specified) to be used during simulation is obtained from the *Simulation Configuration*. Since the *scope* property of the stereotype is set to *State*, the value for this attribute will be obtained from the *Simulation Configuration* every time the User enters *Idle* state, i.e., every time a new item is inserted.

There can be situations in which the modeler wants to restrict that a non-deterministic value is either only obtained once for each instance (i.e., assigned at the time of instantiation) or is obtained at the time of instance creation and every state change. These restrictions can be modeled by setting the property *scope* of the stereotype «NonDeterministic» to *Class* or *Dual* respectively. For example, The attribute *type* for an *Item* (see Fig. 2) on the basis of which the *Item* is sorted is modeled as «NonDeterministic» with *scope* set to *Class* and values constrained between 0 and 2 representing different types of items. This means that for each instance of *Item*, the attribute *type* is given a value by a *simulation configuration* when an instance of *Item* is created.

Another important form of non-determinism is to assign probabilities to the transitions of state machines. In an RTES environment, we sometimes only know the probability of a component to go into a particular state over time and we are not sure about the exact occurrence of such conditions. For example, we can say that the probability of a car engine to overheat after running continuously for 10 hours is 0.05, but we cannot be certain about the exact instance in time when this situation will happen. We can model this in the engine state machine with a transition going from *Normal Temperature* state to *Overheated* state, during an interval of 10 hours, with probability of 0.05.

For modeling these scenarios, we can assign a probability on the transitions using the property *prob* of the MARTE *GaStep* concept. Whenever a timeout transition has the *gaStep* stereotype applied with a non-zero value of *prob*, the combination will be comprehended as the probability of taking the transition over time of the test case execution. In the sorting machine case study, a *Sorter* can get stuck in a position (e.g., because of a bottle blocking it or the arm jamming) for example with a probability 0.02 in a minute if it is not moving and a higher probability when it is moving. The

sending of non-deterministic signals can also be modeled using this type of transitions, by placing them in the actions of such transitions.

If the goal is *validation*, for example based on reliability estimation, then these probability values can be used as a sort of *operational profile* of the SUT [42]. On the other hand, if the goal of testing is the *verification* of the SUT, then the actual values of these probabilities are not important (*Test Framework* decides if an event happens, as long as its probability is not zero). For example, if the goal was validation, the above discussed scenario of a *Sorter* getting stuck could have been modeled with the *gaStep* stereotype to provide an exact value, range, or a probability distribution of occurrence of this failure.

For verification purposes, typically we only require modeling of such situations without specifying an exact probability value or distribution and leave the decision of exact value to the *Test Framework*. The stereotype «TimeProbability» on a transition is used to model such a non-deterministic trigger, whose occurrence is decided by the *Test Framework* and obtained from the *simulation configuration*. Such a transition is very useful to represent failures in environment components. For example, in the Sorting machine case study, an item can fall of the belt and be lost at any time while it is moving inside the machine. This is modeled as a transition from *On_Belt* state to a failure state named *Lost* in the *Item* state machine shown in Fig. 6.

Another type of probability that we modeled in our case studies is for the situations where one event can lead to multiple possible scenarios, but all of them are mutually exclusive. For example, we might want to represent the fact that during the communication with the SUT there is a chance that signals are received with or without distortion. For modeling such scenarios in UML state machines, we can use choice nodes. We provided a stereotype «NDChoice» that can be applied on choice nodes, where each outgoing transition has the same probability. The decision of taking one of the outgoing transitions from such a choice node is made at the time of execution by the *Test Framework*.

If the modeler wants to provide precise probabilities for such scenarios, she can assign the MARTE *gaStep* stereotype to each of the multiple possible outgoing transitions. The example of communication with the SUT can be modeled by having two transitions going out of the environment component state on receiving of a signal, one labeled with a probability that the signal was corrupted and the other with the probability that the signal was fine. As mentioned earlier, modeling the distribution of event arrivals and timeout transitions can be useful for validation purposes, but is out of the scope of this paper, since our goal is verification of the SUT.

6. Simulator Generation

The environment models, comprising a domain model (UML class diagram) and behavioral models (UML state machines), are converted into a Java-based simulator using model to text transformations. The transformations are based on an extension of the state pattern [7], which is a well-known way of implementing state machines. The transformations proposed here are defined to

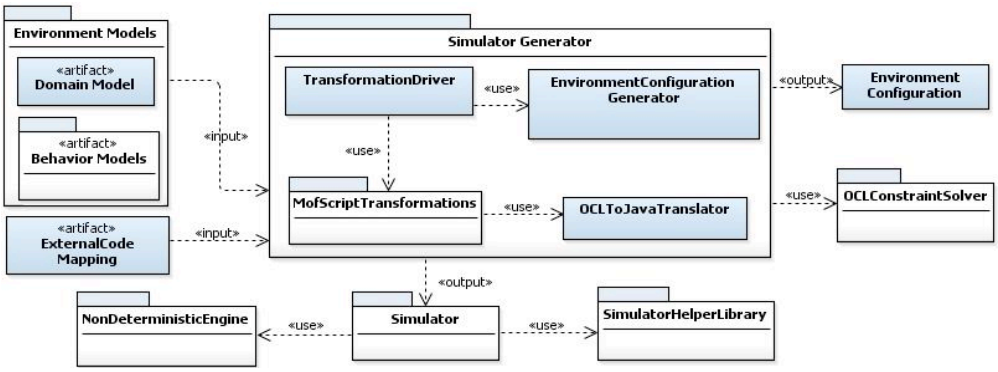

Fig. 0 Architecture Diagram of Simulation Framework

address the specific requirements for environment simulation and RTES system testing. In this section, we first provide an overview of the overall simulation framework. Then we discuss our extended state pattern followed by a discussion on detailed transformation rules for domain and behavior models to simulator code, thus providing a more thorough description of the pattern.

6.1. Simulation Framework

Fig. 10 shows the architecture of the *Simulation Framework*. Components marked with the stereotype «artifact» represent the artifacts that are provided by the software testers to use the framework. The two inputs include the *Environment Models* that are developed according to the methodology discussed in Section 5 and the *External Code Mapping* file that provides a mapping between the external action code and environment models (as discussed in Section 5.3.5)

The package named `Simulator Generator` contains the core components required for simulator generation. The sub-package `GeneratorDrivers` contains driver classes provided with the framework that are responsible for configuring and running the model transformations. The `Model Transformations` package contains the transformations we wrote in MOFscript [43] to translate the environment models to Java classes representing the environment simulator. Class `EnvironmentConfigurationGenerator` is responsible for generating an *environment configuration* representing one possible setting of the environment. The class `OCLToJavaTranslator` is used by the MOFScript transformations to translate the OCL expressions in the model representing guards and change events to their Java equivalent. More details on how the simulator generator handles change events are provided later in Section 6.4.1. The components inside the Simulator Generator package generate a set of classes in Java corresponding to the environment models given as input. This is represented as a `Simulator` package in Fig. 10. The generated simulator is statically linked to classes from two packages: the *Simulator Helper Library (SHL)* and the *Non-Deterministic Engine* which we discuss below.

The *Simulator Helper Library* is developed to support a number of features required by the generated simulator. The library is independent of the case studies and hence is developed as a separate library. The library contains generic features required by active objects (message queue, event handling mechanism, etc.), time related functionalities (including features for handling clocks

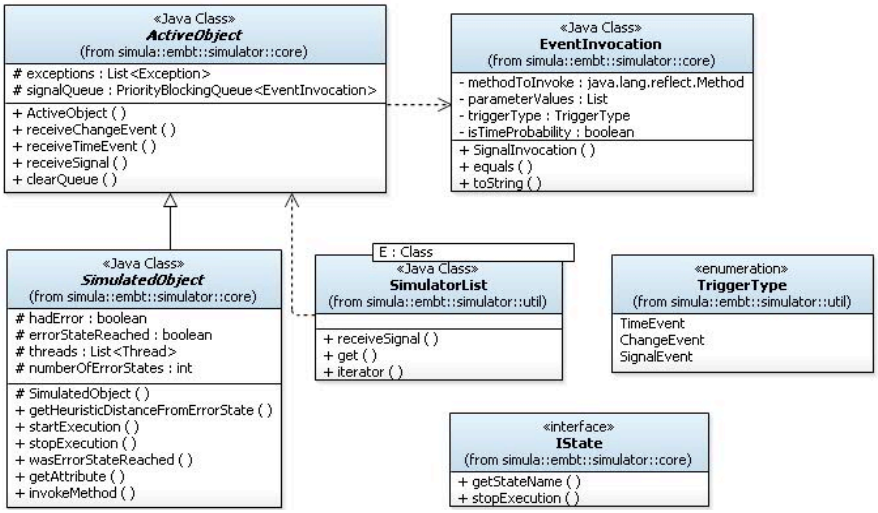

Fig. 0. Important classes in the *SimulatorHelperLibrary*

and timed events), collection classes (providing facility for sending broadcast signals to all elements in the collection), and support for implementing the defined extension of the state pattern, as discussed further in Section 6. The core package of the library is shown in Fig. 11. The class `ActiveObject` represents the UML concept of an *active object*. The class provides an event queue for the active objects in the form of a `java.util.PriorityBlockingQueue`. The queue is a blocking queue and enables setting the priority of elements in the queue. The queue holds instances of class `EventInvocation`. An `EventInvocation` instance represents an event that is ready to be executed and is placed in the event queue of an `ActiveObject`. `EventInvocation` has an attribute `methodToInvoke` that is of type `Method` from the Java reflections package and contains the method of the `ActiveObject` to be invoked along with its parameter types, an attribute `parameterValues` containing the values for the parameters of the method to be invoked, an attribute `triggerType` representing the type of the trigger (signal event, change event, or time event), and an attribute `isTimeProbable` that indicates whether the method to be invoked represents a time probability trigger. The class `ActiveObject` is an abstract class containing the generic behavior of an active object. The behavior provided by the class is used both for the environment components (extending the further generalized class `SimulatedObject`) and implementation of parallel regions (since each region has its own thread of execution and an event queue). The interface `IState` is implemented by all the classes representing UML states. The class `SimulatorList` is used to implement relationships having a multiplicity greater than 1. The class provides facility to broadcast events to all the elements it contains at any given time. For example `SimulatorList` is used to implement the relationship (*itemCollection*, Fig. 2) between *Item* and *User* for the Sorting Machine case study, so that signals to items can be easily broadcasted.

The *Non-Deterministic Engine* is responsible to provide a link between the simulator and various *simulation configurations* produced by the *Test Framework*. The *Non-Deterministic Engine* is called by the simulator each time a non-deterministic occurrence needs to be produced, which in turn queries

```
public Item(int id){
        super(1);
        this.instanceId = id;
        type = (Integer)TestCaseHandler.getTestCase().
                getNextNonDeterministicValue(instanceId*3 + 0);
}
```

Fig. 0. Code snippet showing call to Non-Deterministic Engine

the current *simulation configuration* and returns the value generated by the *Test Framework* corresponding to the non-deterministic occurrence. This is handled by assigning a unique id to each non-deterministic occurrence during the entire simulation, based on the formula: <NDID> = <INSTANCE_ID> * <MAX_ND_COUNT> + <LOCAL_ND_ID>, where NDID is the non-deterministic occurrence id to be calculated, INSTANCE_ID is the unique id assigned to instances of environment components, MAX_ND_COUNT is the maximum number of non-deterministic occurrences that any component in the domain model can have, and LOCAL_ND_ID is the unique id for the occurrence for each environment component. For example, in the *Item* component state machine (Fig. 6), the transition from *On_Belt* to *Lost* is stereotyped as «TimeProbability», which is a non-deterministic event. The actual value for the time when this transition is to be taken is obtained by the simulator through the non-deterministic engine. As discussed later in Section 6.4.3, non-determinism can be of multiple types. An excerpt of generated code in Fig. 12 shows a call to the *Non-deterministic Engine* in order to initialize the value of the attribute *type* of the *Item* environment component. The statement `instanceId * 3 + 0` will result in a unique id for this non-deterministic occurrence during simulation.

6.2. An Extended State Pattern for Environment Simulation

The original state pattern, discussed in [7], provides a design pattern to implement state-driven behavior in an object-oriented programming language. The idea of the pattern was to provide a clean way to implement state-based behavior and make it easy to add new states or transitions by confining the code related to states in separate classes and by providing a mechanism to change the state class at runtime. The original state pattern did not, however, specify a number of important features present in UML 2.x state machines, such as concurrency, time events, change events, and effects. A number of works have provided extensions to the basic state-pattern for various purposes. We discuss these extensions in the related work section and explain why these extensions do not entirely match our needs. In this section, we describe how we extend the basic state pattern for various features required by our environment modeling methodology.

Fig. 13 shows the meta-model of the proposed extensions to the state pattern. The meta-model is included here for ease of understanding. The actual transformation is from UML models directly to Java code (without an explicit target meta-model). The abstract meta-classes `Active`, `IState` and `SimulatedObject` represent the classes with the same names in the *Simulator Helper Library*. The core package of *Simulator Helper Library* is shown in Fig. 11. These classes were required for the environment simulation and are not defined in the original state pattern.

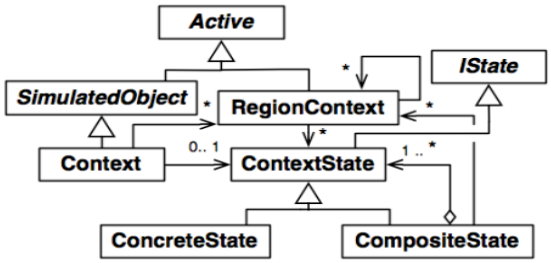

Fig. 0. Extended state pattern meta-model

The concept of `Active` is similar to the notion of an active object in UML. Instances of this class hold an event queue and run as a separate thread. All state classes extend the `IState` class. The class is implemented as a Java interface in *Simulator Helper Library*. `SimulatedObject` holds a list of threads that have been created so far for the instance of an environment component and whether or not the simulation has resulted in reaching an error state. The classes `Context`, `ContextState` (called *State* in state pattern), and `ConcreteState` perform similar roles as in the state pattern. The class `Context` corresponds to a «Context» component of the environment. It holds a reference to all the states and to the current state of the environment component instance and forwards the incoming events to the current state object. A `ContextState` can be a `ConcreteState` or a `CompositeState`. The `ConcreteState` class represents the simple states where actual implementations of triggers are defined. According to the modeling methodology, all events that are not explicitly defined on a state are ignored (Section 6.4). To implement this behavior, the `ContextState` class provides an empty implementation of all the signals defined for the `Context` class. Since all the concrete state classes extend the `ContextState` class, if a signal is not accepted in a state, then its implementation is not provided in the corresponding `ConcreteState` class. This results in executing the empty implementation and the event is ignored. Since the original state pattern does not handle parallel regions or composite states, we have added the `CompositeState` and `RegionContext` classes for this purpose. An instance of `CompositeState` class is created for each composite state in the state machine. A `CompositeState` holds substates that can be either a `ConcreteState` or `CompositeState` (implemented as composite pattern [7] in the meta-model). Instances of `RegionContext` are generated for each parallel region in the state machines. All the states within a region are created as instances of `ContextState`. A `RegionContext` can have further links with other `RegionContext` objects in case of further parallel regions within a region.

Further extensions to the state pattern are related to handling the required simulation details for RTES system testing. These extensions include the support of non-determinism that is a common feature in RTES environments. We also provide support for change events, time events modeled using the MARTE profile, and handling actions written in Java. We also provide a generic way to develop and integrate a communication layer in the generated code for communication with the SUT (via the

external action code as discussed earlier in Section 5.3.5). The generated code also supports the generation of code for error and failure states and various heuristics necessary to apply search-based testing techniques (e.g., for calculation of branch distances [44]).

The overall event processing in the simulator for active environment components is based on run to completion processing as defined by UML semantics. The active object waits on the event queue, performs a dequeue operation, processes the received events, and waits on the event queue again if the queue is empty.

6.3. Transformation of the Domain Model

The domain model is used to obtain information regarding various environment components, their relationships and possible cardinalities of associations, attributes, non-deterministic attributes, signals, and signal receptions. This information is used throughout the transformation process and is also included in the generated simulator classes. Since the transformation rules are based on an extension of the state pattern, a number of Java classes are generated for every stateful component in the domain model, e.g., *RVM, User*. As discussed earlier, every environment component is translated into a Java class which is an instance of the `Context` meta-class. A list of the important auto-generated methods in such context classes along with their descriptions is provided in Table 1 and a list of generated attributes is given in Table A.1 of Appendix A. The various methods listed in the table are further discussed in Section 6.4 (since most of them relate to the behavior models). Every context class instance for each environment component will be assigned a unique id, called `instanceId`, during simulation. The `instanceId` is decided by the *environment configuration* and is passed to the constructor of context classes when the instances are created, as shown in Table 1. Each environment component holds a reference to instance of a state object that represents the current state of the component. A method `oclInState(stateName)` is provided for every component that returns true if the component is in a state with name equal to `stateName`. The method corresponds to the `oclInState` method defined by the OCL specifications and is called during evaluation of various OCL expressions that need to check the state of a component. Whenever a component

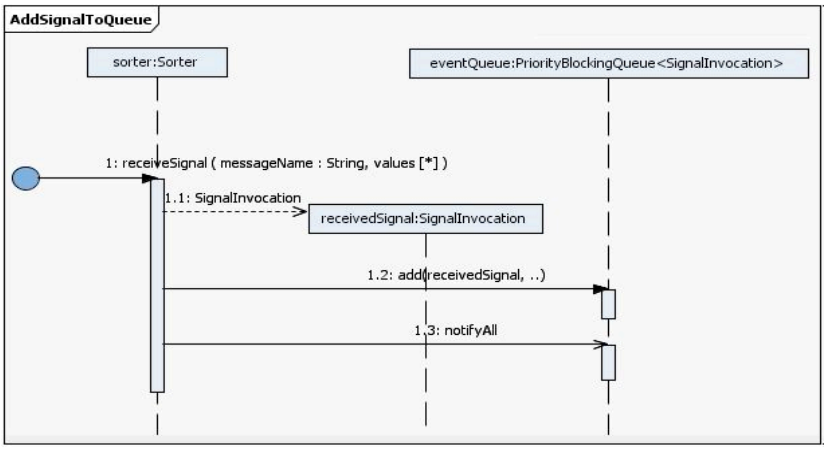

Fig. 0 Sequence of message calls on receiving a signal

changes its state, the method `changeState(fromState, toState)` is called, which updates the state object referring the current state. The method `startExecution()` is called when the simulation is started and `stopExecution()` is called when the simulation is stopped. These methods call methods with the same name in the *external action code* (shown in line # 22 and line # 29 in Fig. 8). Code related to acquiring or releasing resources can be placed in these methods. Effects defined on transitions where either the constructor of the environment component is the trigger or the transitions are initial transitions (i.e., their source is the initial pseudo-state) are implemented in the `startExecution()` method. For example in the *Sorter* state machine (Fig. 4), the effect of initial transition (`action.openConnection()`is implemented in the `startExecution()` method of context class `Sorter`.

Relationships are implemented following the standard class diagram conversion rules [45]. Signal receptions are translated into Java methods in the context class. The associations with an environment component at the navigable end, with a multiplicity above one, are implemented using a `SimulatorList` collection from the *Simulator Helper Library* (Fig. 11). If, in the action code, a signal is sent to a role name having multiplicity more than one, then it is sent to all the elements of the collection. The attributes of classes that are stereotyped as «NonDeterministic» (e.g., *moveArmTimeLC* in *Sorter*) are used to generate an output file, called `NonDeterministicOccurrences` that contains the range of values specified in the model for these attributes. This file is used by the *Test Framework* to identify the domain of valid test cases. Initial values for the *environment configuration* are randomly generated based on the OCL constraints defined on the attributes.

Table 1. Automatically generated methods in instances of the Context meta-class

Method Name	Description
«Constructor» (int instanceId)	The constructor is passed with the unique instanceId for this instance during the simulation. Actions of constructor are included in startExecution().
startExecution	This method is called at the start of execution and all the initialization of attributes, regions, and states is done here. Threads for concurrent regions are started in this method.
stopExecution	This method is called at the end of execution. Various threads are stopped, and resources are released in this method. For example in the sorter case, a call to action code to close the opened sockets is sent from this method.
evaluateChangeEvents	This method is called from setter methods to evaluate whether any of the change events is affected by the current change.
executeChangeEvent (condition)	This method is called when the condition of a change event has been satisfied. The call is forwarded to executeChangeEvent of the current state object.
Setter methods	Setter methods are generated for every attribute and association of the environment components.
Getter methods	Getter methods are generated for every attribute and association of the environment components.
oclInState(stateName)	Evaluates whether the component is in the specified state. The Semantics is similar to OCL method oclInState, except that the parameter stateName is of type String.
timeout	Called whenever a timeout has occurred (i.e., a timer of a time event has expired). The method calls the timeout method in the current state object.
Signal methods	These methods are generated for every signal that is accepted. The method forwards the call to the method implementing the signal in the current state.

State getters	A getter method is generated for every ContextState class
executeCompletionEvent	This method is executed when the entry actions of the current state object have been executed. The call is forwarded to current state object.
changeState(fromState: IState, toState:IState)	The fromState object refers to the source state and the toState refers to the target state. onStateExit()in fromState and onStateEntry()in toState are called. Current state is changed to toState.

6.4. Transformation of Behavioral Models

In this section, we will discuss some of the important rules for transformation of environment behavioral models (i.e., UML/MARTE state machines). As discussed earlier, depending on the type of the states in state machines of environment components (i.e., simple, orthogonal, and composite states), instances of corresponding sub-classes of meta-class `ContextState` are generated. Various methods that are generated for the instances of meta-class `ContextState` and its subclasses are discussed in following sections. A summarized list is provided in Table A.2 in Appendix A.

6.4.1. Event Handling

As discussed earlier (in Section 5.3.4) the environment modeling methodology allows three types of events to be modeled: signal events, change events, and time events. Since environment components represent active objects, each of them contains an event queue that holds the events that have been dispatched, but have not been executed yet. An environment component instance is considered to be busy when it is executing behavior corresponding to an event. When an event is triggered on a busy instance, the event is kept in the event queue and is processed after the current execution is finished. Following we discuss how each of these events are translated into simulator code.

Handling Signals The rules for handling signals are defined according to UML semantics. Here we discuss how they are realized using the proposed extensions of the state pattern. Sending of a signal from one component to another is done in the generated code by calling a `receiveSignal` method of the target context instance, which extends the *Active* class in *Simulator Helper Library* where `receiveSignal` is defined. The method places the received signal in the queue as an `EventInvocation`, which represents the method to be invoked and the parameters. This behavior is shown as a sequence diagram in Fig. 14. When the context object is ready to process the signal, then the method of the `EventInvocation` is invoked on the context object using Java Reflection API. Whenever an instance exits a state, its event queue is emptied (except for «TimeProbability» events, discussed later).

Signals to a SUT are sent through the action code. All the signals towards the SUT are first forwarded to a corresponding method (with the same signature) of the action code. For the reasons discussed earlier (in Section 5.3.5), the low level details for sending the signal to the SUT over a communication medium are written in the external action code manually by the developer. For example, in the state machine of *Sorter* shown in Fig. 4, as a result of receiving a signal

`trigger_sensor()`, the *Sorter* sends a signal to the SUT. This is modeled as a signal to the action class `action.triggerSensor().` The implementation of the action class corresponding to the *Sorter* is shown in Fig. 8. The action class contains `triggerSensor()` (line # 19 in Fig. 8) which implementation is to forward this signal to the SUT over the TCP connection (implemented as `sendMessage(..)`). Similarly, it is the responsibility of an action code developer to define a mechanism by which the signals are received from the SUT and are forwarded to the context objects. For this purpose, every action class has an access to its corresponding context object. For example in the external action code of the *Sorter* shown in Fig. 8, an object of type *Sorter* is passed as an `IActiveObject` to the method `startExecution()` (line # 22 in Fig. 8). This method is called at the start of environment simulation and is a way to allow the action code writer to provide initializations required at the start of execution (e.g., opening of sockets). As shown in the implementation of the action code in the figure, the reference of the instance is kept in a local variable of the class (line # 23), which is then used to send messages to the *Sorter* component (see line # 19 in Fig. 8) by invoking message `sorter.receiveSignal().` In the state machine of *Sorter* shown in Fig. 4, the signals `position_left()`, `position_centre()`, and `position_right()` are sent by the SUT to the *Sorter* and are handled in the above mentioned way.

Handling Change Events. A special mechanism was implemented for handling change events. In the generated code, setter methods are generated for the attributes of environment components. All action code statements that are assigning values to attributes are converted to corresponding setter calls of state machines. Within the code of the setter method of every attribute there is a call to `evaluateChangeEvents()` in the context object. The method forwards the call to the current state object. Within the state object, `evaluateChangeEvents()` evaluates whether the change in the attribute value has an impact on any of the possible change events. If this is the case, then the corresponding condition which was evaluated to true is returned by the method. In the context object's setter method, if the condition returned is not null, then a call to `executeChangeEvent()` with the condition as parameter is placed in the event queue of the context object. This mechanism is similar to handling signals in the queue. The only difference is that the change events have a higher priority than the signal events (reasons discussed later in Section 6.5.3). The mechanism was adopted in order to execute change events asynchronously for active objects. As an example, the behavior of what happens when a setter method is called for the `notRoutingFlag` attribute of the *RVM* component is shown in Fig. 15. A change event depending on the value of this attribute is shown in Fig. 5 (i.e., `self.notRoutingFlag` and `oclInState(Routing::Idle)`). When the `executeChangeEvent` method is executed, it forwards the call to the current state object, which in turn evaluates the condition again and executes the transition corresponding to the change event. If the condition is not satisfied, then nothing happens and the event will be considered as lost.

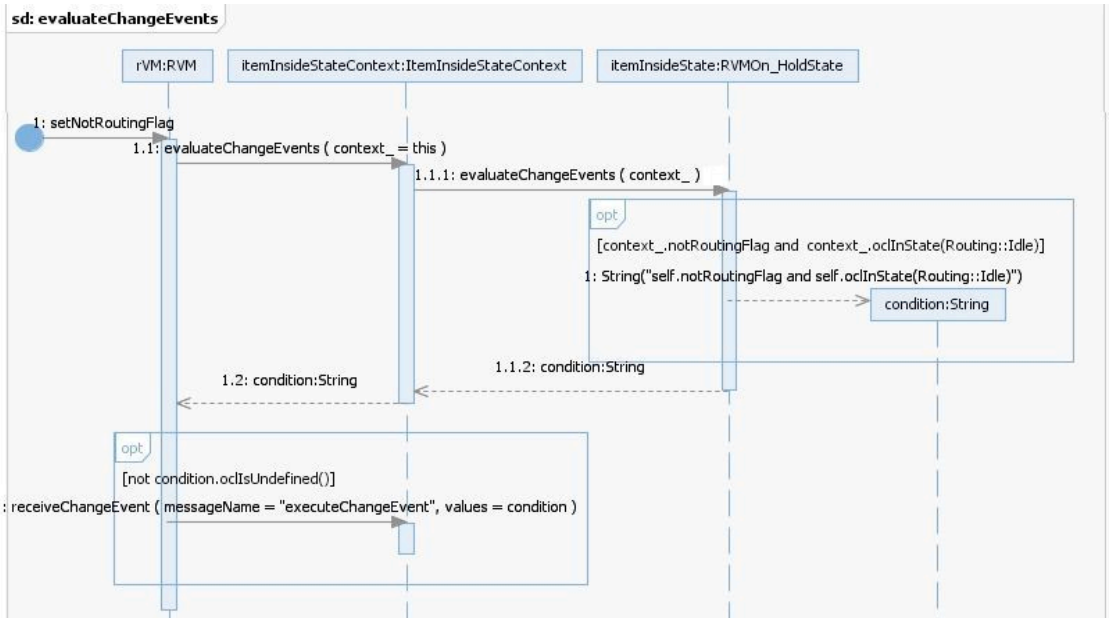

Fig. 0 Sequence diagram showing the behavior for evaluating change events

Handling Time Events. Another non-trivial transformation is for the time events in state machines. We have created a *TimeService Library*, as part of *Simulator Helper Library* that is responsible for managing time-related operations (e.g., clock handling, event scheduling). In every state class that has an outgoing transition with time trigger(s), an object of type `TimeService` is added, which is responsible for handling time-related operations. For each time event, a corresponding `TimeInstance` object is included that is initialized to the value of the time event. A method `afterT`<i> is also included for every time event, where <i> is the unique index associated with each time event. For such state classes, a `timeout()` method is also implemented, which has the logic of forwarding the call to the correct `afterT<i>` method.

The time events are scheduled according to their corresponding clock. If no specific clock is associated, then they are scheduled on first entry into the state class and are reset on every next entry into the state. Non-deterministic time events (where times of occurrence are determined by the simulation configuration) are handled as discussed in a specific section below (Section 6.4.3). When a timer associated with a scheduled time event expires, a signal `timeout()` is sent to the context object with the information of the time instance. The context object forwards the call to the *timeout* method of its current state.

6.4.2. Handling Hierarchical State machines

We have already discussed how a simple state is handled by the code generator (in Section 6.2). Since a submachine state is semantically equivalent to a composite state (*p. 566* [4]), both of them are treated in the same way for code generation. In the remainder of the section, we describe in detail how the code generator handles orthogonal regions and composite states.

To translate parallel regions into code, we introduced the concept of a `RegionContext` class (Fig. 13). The class acts as the state pattern's context class for various states in that region (i.e., it

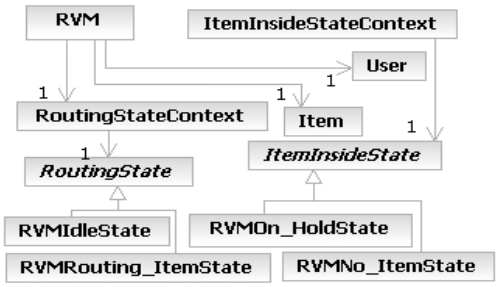

Fig. 0. Generated code structure for RVM environment component

holds the current state object and forwards the messages to it). The `Context` class in these cases has a reference to this `RegionContext` class. Each `RegionContext` class extends the *Active* class and thus holds a queue of events. This was important for simulation in order to allow each region for separate processing and execution of events in its queue. Fig. 16 shows the static structure of the code (without operations and attributes) produced by the code generator for the *RVM* state machine shown in Fig. 5. Classes name `ItemInsideStateContext` and `RoutingStateContext` are the two `RegionContext` classes corresponding to the two regions of the state machine. Both these classes have an association to their corresponding `ContextState` object, `ItemInsideState` and `RoutingState`, respectively. These `ContextState` classes are specialized by the `ConcreteState` classes, for example, in the case of *RVM*, `RVMIdleState` and `RVMRouting_ItemState` are instances of meta-class `ConcreteState` (for the states *Idle* and *Routing_Item* in Fig. 5 respectively) and specialize the `RoutingState` class, which is an instance of the `ContextState` meta-class.

For every composite state, at least two classes are added to the state hierarchy. If a composite state does not contain parallel regions, then an instance of `CompositeState` class is added and, for all the sub-states of this composite state, instances of `ConcreteState` are added. If there are parallel regions, then an instance of a `RegionContext` class is added and the `CompositeState` class has an association with it (Fig. 13). The rest of the handling is similar to other states as defined by the UML semantics (e.g., when a sub-state is entered, the entry actions of composite states are executed first followed by the entry actions of the sub-state).

6.4.3. Handling Non-Determinism

Non-determinism can be of five types in the environment models, as discussed in Section 5.3.7. Following we discuss how each of the five types of non-determinism are handled for simulating the environment. Note that, as discussed earlier in Section 6.1, a unique id is assigned to each non-deterministic occurrence during the entire simulation, which is used by the *Non-deterministic Engine* to select an appropriate value for this occurrence from the *simulation configuration*.

The first form of non-deterministic occurrence is due to a trigger accessing a class attribute that has a «NonDeterministic» stereotype. On entering a state, when one of the outgoing transitions contains a trigger with such an attribute, the generated code passes the unique id of the non-

deterministic occurrence to the *Non-deterministic Engine* and obtains an appropriate attribute value from the *simulation configuration*. Handling of time events accessing such variables was discussed earlier in Section 6.4.1.

The second form of non-determinism is when a variable of an environment component (modeled by assigning a stereotype «NonDeterministic») needs to be initialized at the time of instance creation during simulation. This information will be implemented by having the component constructor query for a value from the *Non-deterministic Engine*. For example, for the domain model in Fig. 2, the code that initializes the value for the attribute `type` of the `Item` component is shown in Fig. 12 .

The third form of non-determinism is the explicit representation of the probability to take a transition, which is specified by the «gaStep» MARTE stereotype. During simulation whenever the code corresponding to such transitions is reached, based on the probability it is decided whether or not to take the transition.

The fourth type of non-determinism can be due to the «TimeProbability» stereotype on a transition. The time value specifying the delay in taking the transition is obtained from the *Non-deterministic Engine* when the state is entered for the first time or when the instance is reentering the state after this transition has been executed.

The fifth type of non-determinism is possible when we have a choice node with the «NDChoice» stereotype. During simulation, whenever the code corresponding to such a choice node is reached, the option of which branch to select is obtained from the *Non-deterministic Engine* (again, by passing the id of this occurrence).

6.4.4. Handling Oracle Information

As discussed earlier (in Section 5.3), error states represent the states of the environment that are reached due to a faulty implementation of the SUT. For example, in the Sorting machine case study, there are two possible errors: the items are not sorted correctly according to their type or an item reaches the Sorter while it is moving. The first error scenario is modeled in the state machine of Item (Fig. 6) and the second scenario is modeled in the state machine of the Sorter (Fig. 4). For the purpose of verification, the testing that we performed in [44] aimed at reaching the error states of the SUT. The oracle consists in checking that any environment component instances do not traverse any of the error states during test execution. Each error state for each instance of the environment components is assigned a unique id during the simulation and the search heuristics (to help the generation of test cases, discussed in Section 6.7.1) use this id to report information relevant to the selected test heuristic, e.g., the distance from the error state for search-based testing.

During the simulation, a number of components in the environment might fail, but with a correct SUT, the environment should never enter in an error state, although the SUT would likely operate with degraded functionalities. In terms of code generation, the execution is stopped once an error state

is reached, a complete log showing the execution trace is generated, and a *JUnit* test case is generated corresponding to the values used in the simulation in order to enable the re-execution of the test case.

6.4.5. Handling Guards and Actions

Since the environment models are translated to Java code, the OCL guards on the models also need to be translated to Java. For search-based testing we need to evaluate the OCL guards and the corresponding branch distance [44]. Therefore, the OCL guards are translated to their Java equivalent along with instrumentation code to support testing heuristics at run time [44].

As mentioned earlier (in Section 5.3.5), we have used Java as an action language. Complex action code and the code related to the communication with the SUT (e.g., handling UDP/TCP sockets, writing to the file system) are written in a separate action code file developed as part of the modeling activity as discussed in Section 5.3.5. Recall that the mapping between the Context class and its action code class is provided in the *External Code Mapping* (Section 5.3.5). An object of this class is accessed in the models by using the keyword *action* as shown in Fig. 4 in the transition action of the initial transition. The actions on the transitions are placed inside the body of the corresponding events in the instance of `ConcreteState` class corresponding to the state from which the transition is outgoing (e.g., the action discussed above is placed in the method for the signal `user_inserts_item()` in the class `RVMNo_ItemState`). Internal state activities (entry, do, and exit) are handled as defined by the UML semantics. Their implementation is provided in the `onStateEntry()` method of the state classes. On completion of the do activity, `executeCompletionEvent()` in the state class is called.

6.5. Various Design Decisions and Their Rationale

Following we discuss various design decisions and their rationale for the code generation approach. Decisions corresponding to the UML semantic variation points that had to be resolved for simulator generation are also discussed.

6.5.1. Object Concurrency Model

We used the Active object model [4] to handle the concept of a concurrent object. In our case, most of the environment components are considered as active objects. This is because they operate independently in the RTES environment and can communicate asynchronously with each other and the SUT. These objects have their own thread of execution and receive asynchronous messages that are handled using an event queue. For example, *Sorter* shown in Fig. 2 is implemented as an active object and executes independently from the other environment components, such as *RVM* and *Item*.

An active object simulating an environment component can have multiple internal threads associated with it. These threads correspond to the parallel regions of the state machines. In our motivating example, *RVM* (Fig. 5) has two internal threads, each for a parallel region (*ItemInside*, *Routing*). This was required to simulate the behavior of an *RVM* when routing and handling item

insertion at the same time. In other cases, the parallel regions were used for modeling component failures that could happen at any time independently of the current state of the component.

6.5.2. Time Semantics

Typically, in simulation approaches, the aim is to simulate and analyze the behavior of a system or environment before it is actually built. For the type of simulator that we have developed, the aim is to simulate the environment in order to test the RTES in diverse situations, without involving actual hardware or people. The SUT in our case is always the actual executable production code and is seen as a black-box. We have no control over the SUT behavior and its definition of time. Therefore, there is no point in simulating the SUT clock, unlike the case in typical simulation approaches such as SystemC[5]. From its point of view, the SUT interacts with the simulator as it would interact with the actual environment. The time it takes for SUT to process a message will be same in both cases. The notion of time for the environment is therefore based on the software implementation of the physical time. In our case, we used the implementation provided by Java time semantics which are based on the CPU clock.

A typical issue in using the CPU clock is the jitter that might be introduced because of computation overhead on the processor (e.g., garbage collector, other operating system processes). This jitter can range up to a few milliseconds. Fortunately, for the type of environments for the systems that we tested, as in many embedded applications, a delay of few milliseconds was not a major issue as the time events were generally in the magnitude of seconds (for example the time of a bottle traveling on a conveyor). To be on the safe side, we explicitly executed the Java garbage collector before and after the simulation. Moreover, for the experiments that we conducted, the garbage collector never executed during simulation and we never faced synchronization issues due to jitter. For the type of industrial systems that we were focusing on (see Section 2), using Java was sufficient. For the environments which have hard time constraints from the system, for magnitudes of milliseconds or less, this can be a critical problem. A possible solution to address this problem is to use real-time Java virtual machines (e.g. Sun Java Real-Time System [46]) running over a real-time operating system (e.g. SUSE Linux Enterprise Real Time Extension [47]), which will result in nanosecond level accuracy. The code that our tool generates is in theory compatible to run with real-time Java virtual machines since this is one of the requirements of such virtual machines. Though, the practical implications of doing this for testing hard real time systems still remain to be investigated.

6.5.3. Execution Semantics and Order of Events in Queue

The state machines of environment components are implemented using the run-to-completion semantics as specified by the standard UML [4]. For handling asynchronous messages, active objects need to implement event queues. We have used the `PriorityBlockingQueue` Java class to implement these queues. They are priority queues and hold various events during the life cycle of an

[5] Webpage: http://www.systemc.org/, date last accessed: 04/01/2012

environment component. The time events in the queue are assigned the highest priority, change events the second highest, and the signal events the lowest priority. Time events have the highest priority since the behavior that they trigger is explicitly related to time, so it should be executed as close to the event occurrence time as possible. Change events have higher priority than the signal events since it is important to execute the corresponding behavior at the earliest opportunity once the change condition turns to true, as it may become false again.

6.5.4. Default Entry & Handling Conflicting Triggers

Another UML semantic variation point is the decision about the default behavior of state machines when there is no explicit initial pseudo state defined in an enclosed region (e.g., in sub states). Our environment modeling methodology requires the modeler to put at least one initial pseudo state in every enclosed region. A region without an explicit initial transition will be considered ill-formed.

There can be situations during simulation when one outstanding event satisfies multiple triggers in an environment component. This issue is left as a semantic variation point in UML [4]. For our methodology, when such a case arises, one of the satisfied behaviors is selected and triggered at random. Our environment modeling methodology recommends avoiding such situations as they indicate imprecise or incomplete environment models.

6.5.5. Event not satisfying any Trigger

The behavior in the cases when the occurring events do not specify triggers on active states are left as semantic variation points in UML. In our case, all the events that do not satisfy any trigger are simply ignored. The modeling methodology requires the modeler to only model those triggers that have a significant behavior, e.g., they have a corresponding effect. If accepting a signal coming from the SUT in some state represents a faulty behavior (i.e., that signal should not have been sent), then it should be modeled with a transition leading to an «Error» state. For example, in the sorting machine, a *Sorter* should never accept a signal `item_at_destination()` when it is in *MovingLeftCentre*, and so we modeled this with a transition from the moving state to an error state, as shown in Fig. 4.

6.5.6. Event Evaluation Time

In UML semantics, the time it should take for a component to dispatch an event after it was received is not defined and is left as a semantic variation point to be decided by specific methodologies. In our methodology, this time is dependent on the event queue size of the receiver, the priority of the event, and the time to enqueue and dequeue the events before consumptions.

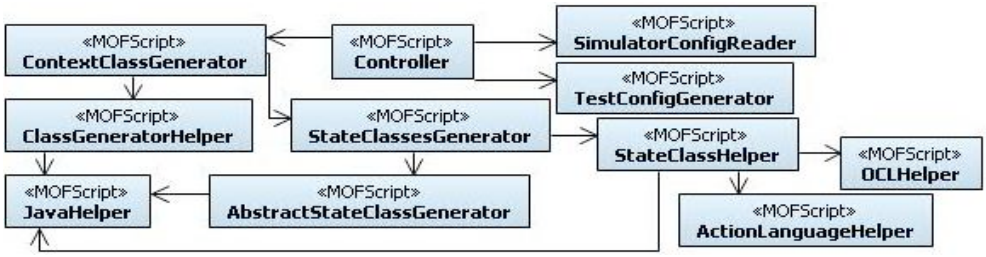

Fig. 0. MOFscript Transformations used for generating simulators

For the specific case of time events, as discussed earlier (Section 6.4.1), when a time event is received by the environment component, it is placed in the event queue. If there are more time events ready to be dispatched in the queue, then the event received may never be executed. If a time event already at the front of the queue is executed, it will result in a time transition and hence, the queue will be emptied from all its events, unless for specific cases already mentioned in Section 5.3.4 (e.g., for transitions with «NonLeaving»). If the time event belongs to a parallel region (or a substate of a parallel region) then it will again be placed in the region's queue (see Section 6.4.2). Thus, the dispatch time of a time event depends on the time events that are already in the queue(s) and the time to enqueue and dequeue in the queue(s), and the number of internal orthogonal regions. Overall this time will only be within a few milliseconds, which is not a major obstacle for the type of system testing that we deal with and for many embedded applications (as discussed in Section 6.5.2).

6.5.7. Signal Transmission

As discussed earlier, signals are transmitted from the source to the target by calling the target's `receiveSignal()` method with the signal parameters. The `receiveSignal` method creates an instance of `EventInvocation and` puts it in the event queue of the target component. The sender and receiver will be running on the same process since all the environment components run on a single process. This decision was made because intra-process communication is easier and faster than inter-process communication. Since SUT is a black-box in our testing approach, the SUT runs on a separate process (or even a separate machine) and the communication between environment components and SUT is handled through the external action code.

6.6. Automation

We implemented the rules mentioned earlier in Section 6 by using MOFScript model to text transformations [43]. Fig. 17 shows various MOFScript transformations that we developed for transformation from environment models to Java code. These transformations are contained in the package named `MOFScript Transformations` shown in Section 6.1. Stereotype «MOFScript» denotes the MOFscript m2t files containing the transformation rules.

The control of transformations is handled by `Controller`. The `ContextClassGenerator` transformation is responsible for transforming the domain model with the help of `ClassHelper`. `ContextClassGenerator` also calls `StateClassesGenerator`, which is responsible for transforming the state machine of an environment components to Java code with the help of

`StateHelper`. `JavaHelper` contains the rules specific to Java and are used by various other transformations; `OCLHelper` uses the `OCLToJavaTranslator` class discussed in Section 6.1 to convert OCL expressions to their equivalent Java code. `ActionLanguageHelper` is responsible for modifying the action language code in the models according to the generated code structure. `SimulatorConfigReader` reads the configuration file given as input and the `TestConfigGenerator` generates a configuration file that is read by the test engine.

The tool requires models exported as standard EMF format for UML (.uml file). This is a widely accepted format for interchanging UML models and is supported by a number of modeling tools, including Rational Software Architect and Papyrus. Our tool reads the .uml file and generates a simulator corresponding to the models. The software engineer then needs to provide a test driver that specifies the testing strategies to be used and that initializes the SUT and the generated environment simulator.

6.7. Interaction with Test Framework

The *Test Framework* queries the simulator to obtain information required to generate the *Simulation Configuration* (e.g., number of non-deterministic variables and their value domain and type). The *Test Framework* then generates valid *Simulation Configurations* based on the testing strategy in use (e.g., at random for random testing) and uses the *Test Driver* to run the SUT with the environment simulator. The test framework uses a set of heuristics based on the previous test case executions (e.g., rewarding test case diversity and Genetic Algorithms) to choose new test cases to run and evaluate (for details see [7]).

The goal of the *Test Framework* is to find a *simulation configuration* for which, once executed with the simulator, an error state is reached (if any fault is present in the SUT). Once such a configuration is found, it automatically generates a *JUnit* test case. In [7], we show how the models developed using the methodology described above can be used for automated system testing of RTES. A test case is used to define two important components of the simulation: (1) the configuration of the environment, e.g., number of sensors/actuators and their initialization; (2) the non-deterministic events in the simulation, e.g., variance in time-related events such as physical movements of hardware components, occurrence and type of hardware failures and actions of the user(s). In [7], we investigate three strategies to automate these choices: Random Testing, Adaptive Random Testing and Search-Based Testing (using a Genetic Algorithm). The results of the experiments (on three artificial problems and one large industrial RTES) showed that environment model-based testing is able to automatically find faults in all the considered case studies. In particular, previously uncaught critical faults were automatically found in the industrial RTES [44].

Test data generation can be reformulated as a *search problem* [48, 49], in which for example the goal can be to find test data for which failures are triggered (if any fault is present in the SUT). To achieve such goal, it is important to have a heuristic to evaluate how good test data are, even if they

do not trigger any failure. For example, test data that lead to execute most of the code of the SUT would a priori be more useful than test data for which the computation finishes very quickly after executing few lines of code. A *search algorithm* would use such information to focus on areas of the search space (i.e., test data) that are more likely to contain test data for which the SUT fails. Since it is not feasible to evaluate and run all possible test data, a search algorithm has to focus only on some promising areas. Such type of test data generation is referred to as *search-based software testing* [48], and there are many search algorithms and *fitness functions* (i.e., functions to evaluate how good the test data are). For details regarding how to use search algorithms to system testing of RTES, see [44].

However, to use such search algorithms, it is important to obtain information on the execution of test cases after they are run (e.g., which parts of the SUT code they execute). Such information would be used to compute the fitness function. Which type of information is needed depends on the selected testing strategy.

Typically, in white-box testing, when information regarding source code execution is needed for the heuristics, the code of the SUT needs to be *instrumented*. Instrumentation means that the code is augmented with probes to collect execution data (e.g., to check which branches of "if" statements are executed). Since our system testing approach is based on environment models where the SUT is treated as a black-box, the probes are inserted only in the code of the environment simulator. The models contain valuable information to guide search-based testing, and such information would be lost or difficult to reverse-engineer after the simulator is generated. For practical reasons, the probes are automatically inserted when the code is generated from the models. This is another advantage of using simulators generated from models rather than manually coding a simulator, as manually inserting those probes would be likely tedious and time consuming.

6.7.1. Search Heuristics

Following we discuss four types of instrumentation to collect test case execution information used in fitness functions in the context of environment-based system testing for RTES [44], namely *approach level*, *branch distance*, *time distance*, and *risky states*. Because the goal of such type of testing is to find simulation configurations for which error states are reached, we collect information regarding *how close* the execution was from reaching any of these error states.

The *approach level* is a common heuristic [48] in white-box, search-based software testing, where executions that get close to the testing target (e.g., branches in branch coverage) are rewarded. When a test case is run, several states in the environment state machines are reached, while others are not. The approach level calculates the *minimum* number of transitions required to reach any error state from any visited state of the environment component. To obtain this value, we consider the state machine as a graph and perform a breadth first search on each state to obtain the minimum distance (in number of transitions) to reach the error states. This calculation is done only once, when the simulator code is generated, and then hard-coded directly in the simulator code to ease fitness computations during simulation. When a test case is executed, the approach level for all reachable

error states is calculated and reported. For example, consider the environment component *Sorter* (Fig. 4): the distance for the *Sorting::Working::Error* state from *Sorter::Working::Left* state is two whereas such distance is one from *Sorter::Working::MovingCentreRight*. If both these two non-error states are reached during test case execution, then the approach level would be the minimum value among those two values (i.e, 1).

The second piece of information used in the fitness functions is the *branch distance*. To reach an error state, it is necessary to follow some specific paths in the state machine. A path would be a sequence of state transitions, driven by triggers. However, state transitions often have *guards* (e.g., logical predicates expressed in OCL), which need to be satisfied (i.e., their predicates need to be evaluated to true) to take such transitions. The predicates in these guards depend on variables, which values cannot be directly manipulated [49] by the *Test Framework* and depend on the entire test case execution carried out so far until the guard is evaluated. Some guards can be difficult to satisfy (i.e., only few simulation configurations lead to it) and, because the variables in the guards cannot be directly manipulated, it is not possible to use external *constraint solvers* to satisfy them. The *branch distance* is a heuristic to reward simulation configurations that brings the guards closer to satisfaction. Consider for example the guard "x==0", and two test cases for which "x=1" and "x=100". None of the test cases satisfy that guard but the case "x=1" is heuristically closer. The branch distance is a common and effective heuristic in search-based software testing for structural coverage [48]. In previous work, we have developed a search based constraint solver, in which we extended the branch distance functions for white-box testing to support all the constructs of OCL constraints [49]. In this paper, when we generate Java code to represent the OCL guards, we instrument such predicates to calculate their branch distance each time they are evaluated. For the details of how these branch distances are calculated, see [49].

The third type of information used in the fitness functions is the *time distance*. In some cases, a transition is taken only after a timeout, and this type of transitions can appear on the paths that lead to the error states. For example, assume that, in a particular state, the environment expects a signal from the SUT within one second.

If such a signal is not received, then an error state is reached. In a state machine, this would be modeled as a transition to an error state with trigger *after (1, s)*, whereas receiving the signal from the SUT would trigger a transition toward another state. The *time distance* calculates how much longer it would have taken to get a given time trigger fired. Taking the example above, if we receive the signal after one millisecond, it would be worse than receiving the signal after 900 milliseconds, although in neither of the cases the error state is reached. Each time there is a time transition, during code generation such transition is instrumented to calculate its time distance.

The fourth type of information used in the fitness functions to guide the search is about risky states. The states that have a direct transition to error states are considered to be *risky* states. For the search, this information is important as these are the closest states to the error states. For example, in

the *Sorter* component, the state *Sorter::Working::MovingCentreRight* is a risky state. How often a risky state has been reached, and for how long the environment was in such risky states, can be used by the search algorithms to reward test cases that keep the environment in these risky states as long as possible.

Table 2. A simulation configuration for the sorting machine

Environment Component	Instance Id	Nondeterministic Occurrence Id	Related Property	Value
User	0	0	insertionTime	500
RVM	1	- None -	- None -	- N/A -
Sorter	2	6	moveArmTimeLC	292
Sorter	2	7	moveArmTimeCR	303
Item	3	9	type	1
Item	3	10	timeToNode	1062
Item	3	11	«TimeProbability»	0, 0
Item	4	12	type	0
Item	4	13	timeToNode	970
Item	4	14	«TimeProbability»	1, 300
Item	5	15	type	0
Item	5	16	timeToNode	1011
Item	5	17	«TimeProbability»	0, 0

6.7.2. Simulation Configuration

The *Simulation Configuration* is generated by the *Test Framework*. During the simulation, the simulator queries the *simulation configuration* to obtain the values of non-deterministic occurrences, e.g., exact time in time event. For this purpose each non-deterministic occurrence is assigned a unique id during the simulation. Notice that, once a configuration is defined, the simulation becomes deterministic. In other words, executing again the simulator environment with the same *simulation configuration* should result in the same behavior. However, this latter point is not strictly correct, because a simulation would still be affected by non-deterministic components such as the thread scheduler and other operating system resources. Fortunately, this is not a serious problem for the type of system level testing done here where, for most environments, variances of few milliseconds in the interactions between the environment and SUT are simply negligible as they have no impact on the resulting states of the environment and SUT. When this is not the case, as further discussed in Section 6.5.2, the modeling methodology and code generated are still valid but a real-time operating system and Java RT would need to be used. Therefore, for all practical purposes, a test case is uniquely characterized by a simulation configuration.

As an example consider Table 2 that shows a random generated simulation configuration. The simulation configuration shown is based on an *environment configuration* having 1 *RVM* instance, 1 *Sorter* instance, 1 *User* instance, and 3 *Item* instances. This *environment configuration* results in a total of twelve non-deterministic occurrences, 2 for *Sorter*, 1 for *User*, and 3 for each of the three instances of *Item*. Each instance has a unique id during the simulation ("Instance Id"). Each of the non-deterministic occurrences has a unique id during simulation as shown by the "Nondeterministic

Occurrence Id" column. The "Related Property" column in the table shows properties of the environment components that are related to the non-deterministic occurrences. When a non-deterministic occurrence is based on a transition with «TimeProbability», the stereotype is mentioned in the column. The values for these occurrences selected by the *Test Framework* are shown in the column labeled "Value". In the case of «TimeProbability», each value pair specifies the choice of value as 1 or 0 referring to whether or not to take the transition and the time in milliseconds at which the transition is to be triggered (irrelevant when the transition is not to be taken). Other values in the column are assigned by the test-engine based on the ranges (e.g., the upper and lower bounds) of «NonDeterministic» environment component properties. The ranges are shown in the domain model as stereotype properties of the environment component properties (Fig. 2). For example the lower and upper bounds for *moveArmTimeLC* are 280 and 320. The value in the *simulation configuration* decided by the *Test Framework* is 292.

6.7.3. OCL Constraint Solver

```
1.  public class Test_Sorter {
2.   @Test
3.   public void testCase(){
4.      ProblemData pd = new sorter.embt.SorterProblemData();
5.      TestCaseRunner runner = new TestCaseRunner();
6.      runner.init(pd);
7.      TestCase tc = getTestCaseData(pd);
8.      Environment env = pd.getEnvironment();
9.      try{    runner.runTestCase(tc);       }
10.     catch(Exception e){    fail(e.toString());    }
11.     assertEquals(false, env.hadError());
12.  }
13.  protected TestCase getTestCaseData(){
14.     RingTestCase tc = new RingTestCase();
15.     tc.setVariable(17, 0, 0);
16.     tc.setVariable(16, 1011);
17.     tc.setVariable(15, 0);
18.     tc.setVariable(14, 1, 300);
19.     tc.setVariable(13, 970);
20.     tc.setVariable(12, 0);
21.     tc.setVariable(11, 0, 0);
22.     tc.setVariable(10, 1062);
23.     tc.setVariable(9, 1);
24.     tc.setVariable(7, 303);
25.     tc.setVariable(6, 292);
26.     tc.setVariable(0, 500);
27.     return tc;
28.  }
29. }
```

Fig. 0 An auto generated JUnit test case for Sorting Machine case study

According to the modeling methodology, the domain model captures the different forms the SUT environment can take (see Section 5.2). For a given test execution, we need to select one possible *environment configuration*. We use an OCL constraint solver to generate a possible configuration fro7m the domain model. This consists in selecting appropriate initial values for the association multiplicities and attributes that are not labeled with the «NonDeterministic» stereotype, as the latter are determined by the simulation configuration. Fig. 3 shows an example of OCL constraint on the *User* component of the domain model of the Sorting machine case study (Fig. 2). The constraint specifies that a *User* instance is always associated with a non-empty collection of items. To obtain an appropriate value of an instance according to such constraints, we have developed a search-based OCL constraint solver [49], since current OCL solvers in the literature do not scale up to the complexity of real constraints found in industrial systems. The generated simulator code calls this solver to generate values for which the OCL constraints are satisfied.

6.7.4. Test Driver & JUnit Test Case

A Test Driver needs to be written by the tester, which is used to start the execution of the simulator and SUT, and to stop them after a timeout (a set of predefined libraries are provided to help the tester in this task). In the case studies for this paper, the environment simulator and the SUT are run on different processes on the same machine.

Whenever the *Test Framework* leads the simulation to an error state, this is due to a faulty implementation of the SUT. The specific *simulation configuration* of the simulator at that time is embedded in a *JUnit test case* and a source file representing the test case is generated by the *Test Framework*. This is done so that the *simulation configuration* is saved and can be executed later for debugging purposes. The *JUnit test case* calls the *Test Driver* based on a simulation configuration.

Assert statements are automatically added to check whether any error state has been reached during the simulation. Once generated, these *JUnit test cases* do not need the *Test Framework* for their execution. Fig. 18 shows an auto generated test case for the Sorting Machine case study. The test case is based on the simulation configuration shown in Table 2. The class `ProblemData` used in the JUnit test case holds information of various non-deterministic occurrences.

The method `getTestCaseData()` creates and returns an object of class named `TestCase` based on a simulation configuration decided by the *Test Framework*. The implementation of the method shows the various values being assigned to non-deterministic occurrences. Class `ProblemData` used in the JUnit test case (see line number 9 in Fig. 18) is supposed to hold information related to various settings of the *Test Framework*, such as the *environment configuration* and the total time for simulation. Objects of class `Environment` represent *environment configurations* for the simulator.

7. Case Study

In this section, we discuss the case study we conducted to evaluate the proposed modeling methodology and simulator generation.

7.1. Case Study Design

The objective of the case study is to evaluate whether (1) the transformation rules are sufficient to convert environment models of different complexity levels, and belonging to various domains, to simulator code, (2) the automated generation of simulators is likely to significantly reduce development effort, (3) the generated simulators enables the detection of failures in RTES system testing, (4) the transformations implemented are correct, and (5) the proposed methodology and profile are sufficient for modeling environments of RTES for the type of testing we are interested in.

We selected five different RTES as part of our study. Two of the cases were industrial RTES. One industrial RTES, Industrial Case A (IC-A) is the sorting machine system that we have discussed as a motivating example throughout the paper[6]. As mentioned earlier, in this paper, we are only considering a subset of the case study focusing on the sorting functionality, having four environment components and an average of five states per component.

The second industrial system (Industrial Case B) is a marine seismic acquisition system, which has five environment components with an average of 12 states per component. The aim was to select two industrial systems that belong to different domains, with different functionalities, to study diverse environment models. We also developed three artificial problems of varying complexity that also belong to different domains. Two of the artificial problems (AP1 and AP2) were inspired by one of our industrial case studies and deal with RTES interacting with multiple sensors in different situations. The third artificial problem (AP3) is inspired by a train control gate system discussed in

[6] Notice that in [1] we considered the entire sorting machine case study. In this paper, we only discuss the subset of the case study that we used for testing and simulator generation. Therefore, the data presented in this paper is not exactly the same as in [1].

Table 3. Environment Models for Artificial Problems and Industrial Cases

Case	EC	States		
		Simple	Orth	Comp
AP1	1	3	0	0
AP2	1	6	2	0
AP3	2	9	0	0
IC-A	4	20	2	1
IC-B	3	23	3	1

*EC=Environment components, Orth=Orthogonal, Comp=Composite, Trans=Transition, NDV = Non

[50]. The RTES for these artificial problems were developed in Java. Note that during the simulator execution a number of instances are generated for each environment component. For example, in industrial case B (IC-B), tens of instances were created and ran in parallel for simulating the environment. For the industrial case (IC-A), a hardware interface layer was not separated from rest of the code while the RTES was being developed (as it was for IC-B). As a refactoring task to improve the testing processes, the software engineers are currently working on separating out the code that deals with hardware components and providing a standard mechanism of communication for the SUT. Since the adapter was not yet available at the time of writing, we manually developed the portions of SUT that are related to the subset of the case being used in this paper. However, this has *no effect* on the code generation discussed in this paper, as the environment models would be *exactly* the same for the actual SUT. Table 3 shows the statistics of various modeling elements used in the five cases.

As mentioned earlier, the aim of the case study is to evaluate five aspects: completeness of the transformation rules, effect on model-based simulation generation on development effort, effectiveness of the approach in test automation, the correctness of transformations, and the completeness of profile and methodology. We define them below and briefly mention how they will be assessed, before going into more details in the next section. *Completeness of transformation rules* refers to whether the transformation rules that we wrote are sufficient for generating a simulator from environment models developed using our methodology. Evaluating the completeness of model transformations is still an open research question [51]. Note that in this paper, we have only discussed the most important rules for simulator generation. As further discussed below, completeness of the transformation rules is evaluated by checking that, in all five case studies, there were no elements in the environment models that were ignored or could not be handled by the simulation code generator.

Regarding the *effect on development time*, we evaluate whether the automated generation of the simulator is likely to help significantly reduce the effort required by the developers to perform system level testing. To gain insight into this question, we compare the source code to be manually developed with the size of the models required for automated simulator generation.

Effectiveness in test automation refers to the effectiveness of the *Test Framework* in detecting failures when it uses an automatically generated simulator. We evaluated the fault detection effectiveness for the simulators generated for all the five cases with various testing strategies. On one

hand the generated simulator should not prevent test cases that should fail to do so and, on the other hand, it should not trigger spurious failures, the latter being related to *correctness* discussed below.

Correctness of the transformation rules is about how correct are the transformation rules in generating simulators that behave as expected according to the environment models. Evaluating the correctness of model transformations is still an open research question [51] and no standard mechanism or appropriate tool is available for testing them. To evaluate the correctness of our transformations, we focused on both the structure (the code is what it is supposed to be) and behavior (the code works as it is supposed to work) of the generated code. We evaluated the structural correctness of the code keeping in mind two properties: *syntactic correctness* (the generated simulators are syntactically correct, i.e., no compilation errors) and *semantic correctness* (the generated code is what it should be according to the extended state pattern). To evaluate that the behavior of the generated simulators is correct, we manually developed test cases keeping in mind three properties: (1) behavior of the generated simulators conform to what is specified in their environment models, (2) the simulators correctly report the information required to guide test heuristics, and (3) the simulators correctly detect and report failures in SUT.

Completeness of profile and methodology refers to whether the proposed modeling methodology and the identified subset of UML/MARTE along with the proposed profile is sufficient for modeling the environment of RTES for the automated black-box system level testing that we are interested in. This includes the ability to generate simulators from the environment models, automated generation of test cases, and use of the models as oracles.

To evaluate completeness and correctness of the transformation rules we also created several test models. The models were not linked to any SUT and were only developed to evaluate the transformation rules and were developed in order to cover different sets of modeling elements according to our environment modeling profile. In total the test models comprised of 50 components, with each component having on average of 4 states and 12 transitions and covered various state machine and class diagram constructs and all the modeling features defined by the environment modeling methodology.

7.2. Case study procedure

This section describes how the data was collected for the five evaluation criteria.

7.2.1. Completeness of the Transformation Rules

To evaluate completeness of the transformation rules, we generated simulators for the five cases, each having different level of complexity and modeling different concepts. We also generated simulators for the different test models that covered different sets of modeling elements. The aim was to see whether the transformations completely handle code generation from that diverse set of models.

7.2.2. Effect on Development Effort

To evaluate this we generated simulators for the five cases and obtained size and complexity information about the generated code. When compared to the size of the models, such data can help assess the potential amount of effort saved by generating simulators from models rather than developing the simulators manually. Of course, we can only provide qualitative insights because the quantitative results depend on the skills of developers with respect to modeling and coding. Table 3 summarized the details of the environment models that were developed to generate simulators for the five cases.

7.2.3. Effectiveness in Test Automation

To assess this, we manually seeded non-trivial faults in the three artificial problems and industrial case A (one fault for each SUT). We could have rather seeded those faults in a systematic way, for example by using a mutation testing [52] tool. We did not follow such procedure because the SUTs are highly multi-threaded and use a high number of network features (e.g., opening and reading/writing from TCP sockets), which could be a problem for current mutation testing tools. Furthermore, our testing is taking place at the system level, and though small modifications made by a mutation testing tool might be representative of faults at the unit level, it is unlikely to be the case at the system level for RTES. For the SUT of Case B, a previously uncaught critical fault was found with our *Test Framework* [44] and used to assess the effectiveness of the simulator for test automation. We ran the simulators generated for the five cases (i.e., the three artificial problems and two industrial cases) with various testing strategies to evaluate whether the test cases that were expected to fail did and whether no other test cases failed.

7.2.4. Correctness of Transformations

To evaluate the transformation rules based on the five correctness properties, we adopted a procedure that is summarized in Table 4. The structure of the generated code can be considered to be syntactically correct if it has no compilation errors for various types of models. To achieve this we created a number of test models that contained different set of modeling elements for state machines and class diagrams, generated simulators for them, and used the Java compiler to compile the generated simulator code. For semantic correctness, we inspected the generated code for the developed test models to see if the code was conforming to the extended state pattern.

To evaluate the effectiveness of simulators to detect failures in SUTs, we created two versions of SUTs for the three artificial problems and industrial case A: one version was a bug free version and for the other one we seeded a fault manually (same as we did for evaluating effectiveness in test automation). We created two test cases for each artificial problem and industrial case A, one test case was supposed to detect and report the failure corresponding to the seeded fault. The other test case was not expected to detect the fault. We ran the two test cases on the two versions of the SUT and observed their behavior. Our assumption was that a faulty simulator might lead to falsely reporting a bug in the correct SUT or prevent the triggering of an expected failure in the faulty SUT.

The above strategy was iteratively applied to check correctness and achieve a stable version of the tool. At any rate, testing cannot prove the absence of defects, although it increases our confidence.

7.2.5. Completeness of the Modeling Methodology and Profile

To evaluate whether the modeling methodology and profile are sufficient for simulator generation we generated simulators for the five cases according to the transformation rules presented in the paper. To evaluate the completeness of the methodology and profile with respect to testing we generated test cases with different test strategies and evaluated them based on the information of oracle obtained from the models.

7.3. Results

Following, we present the results of the case study for each evaluation criterion.

7.3.1. Completeness of the Transformation Rules

As far as generating simulators from environment models for the five cases and test models presented above, the transformation rules are complete. These test models along with the three artificial problems and two industrial cases covered all the modeling elements defined in the methodology. The MOFScript transformations developed were able to generate Java code for all of the UML/MARTE/OCL model constructs used in the case study artifacts and the test models.

7.3.2. Effect on Development Effort

When using our methodology, the only significant effort required by the software engineers is to create environment models for the environment of RTES. Once developed, these models are used for generating the simulator, executable test cases, and automated oracles. Table 3 shows relevant size data for the various environmental models and Table 5 summarize the details for corresponding generated simulators for the five cases (three artificial problems and two industrial case studies) in terms of the total number of generated classes, number of methods, threads, and lines of code. The first three rows of the tables show data about the three artificial problems (AP) and the last two rows about the two industrial cases. Note that, even though the number of components in each case is small, during the simulator execution a number of instances are generated for each environment component. As discussed earlier the number of instances to be generated is decided based on the OCL constraints on the domain model and is specified in an *environment configuration*. For example, in industrial case B, as shown in Table 3, the total number of environment components is three, but for one of the components, tens of instances were created for simulating the environment.

One of the reasons for the large number of generated classes, methods, and lines of code is the use of the state pattern, which requires a separate class for each state. For example, industrial case A has only four environment components, but because of 23 simple, 3 orthogonal, and 1 composite state over 5000 lines of simulator code were generated with 35 classes and 386 methods. Even if the code were manually written, we would expect developers to follow a similar pattern (extension of state pattern) in order to facilitate changes, which would have resulted in a similar number of lines of code.

This conjecture is supported by the fact that, in one of the industrial case studies, an existing, manually-written simulator was of similar complexity. The simulator is a complex, multithreaded application and various components run in parallel. If the simulator were to be manually written, the developers would have had to resolve synchronization issues related to concurrency. For example, in IC-B, the generated simulator included 20 threads to handle active objects and various timers. With a large number of possible instances running during an environment simulation (as tens of instances for IC-B) the overall behavior of the environment is quite complex. In our approach, the simulator generator takes care of these issues and the software engineers only have to develop the environment models, which are expressed at a higher level of abstraction than the source code.

Table 4. Summary of procedure to evaluate the correctness of transformations

Property / *Artifacts*	**Syntactic correctness**	**Semantic correctness**	**Conformance to models**	**Heuristics Reporting**	**Failure Detection**
Input	Test models, AP1 – AP3, IC-A, IC-B	Test models	Test models, models for AP1, AP2, AP3, & IC-A	Test models, models for AP1, AP2, AP3, & IC-A	Models for AP1, AP2, AP3, & IC-A
Execution Procedure	Manual drivers	Manual inspection	Manual test cases	Manual test cases	Manual test cases
SUT	None	None	None for test models, stubs for artificial problems & IC-A	None for test models, stubs for artificial problems & IC-A	Buggy & Correct versions for AP1, AP2, AP3, & IC-A
Oracle	Java compiler	Checking compliance with extended state pattern	Comparison with source models	Manually added in the test cases based on source models	Failure reporting mechanism during simulation

Table 5. Details of Generated Simulators

Case	**Classes**	**Methods**	**LOC**	**Threads**	**Manually written LOC**
AP-1	8	67	975	3	123
AP-2	13	133	1871	6	79
AP-3	16	174	2396	8	137
IC-A	35	386	5545	12	181
IC-B	37	573	9209	20	360

The statistics about the generated code are only used to provide an estimation of the complexity of the simulators. The reason is that, the objective of our tool was on generating simulators that are correct and are usable for environment-based system level testing by utilizing reasonable level of computing resources. The generated simulator is given to the end-user as an executable archive so we did not focus on optimizing the code for better understanding or cleaner code generation. Therefore there is room for further optimizations to reduce the number of lines of code, classes, and methods.

The column in Table 5 labeled ‘Manually written lines of code’ show the lines of code that the developers had to write by hand. These were mostly written for external action code dealing with communication. It is worth noting that, even if the simulator were manually developed, this communication-related code would have to be written in any case and would have been very similar

to the code written using our methodology (as we have experienced in one of the industrial case studies). Therefore, the effort required to develop such communication layers is not something specific to our approach, rather it is required whenever environment-based simulations are run.

Overall, the automated generation of the simulator code can be expected to save significant effort to the developers. Though there is a considerable effort involved in developing environment models, given the amount and complexity of the source code generated, it is expected to be less than the effort required for manually developing and maintaining environment simulator code with concurrency and complex synchronization issues. However, to ascertain this claim with confidence, controlled empirical studies in industrial contexts are required.

7.3.3. Effectiveness in Test Automation

As discussed earlier, to evaluate the effectiveness of the generated simulator to help in test automation, we ran the simulator with various testing strategies on all the five cases, i.e., the three artificial problems and two industrial cases.

Overall, the testing framework was able to trigger system failures corresponding to all the seeded faults for these cases. For IC-B, we ran the testing framework for three different testing strategies: Random Testing, Adaptive Random Testing, and Genetic Algorithms and we were able to find a critical fault in the production code. The detailed results for the experiment conducted on the three artificial problems and this industrial case are presented in [44]. Taken together, the results of these experiments increased our confidence that the generated simulators are effective in detecting faults in the SUT when used in combination with various test automation strategies.

7.3.4. Correctness of the Transformation Rules

On the stable release of the transformations, we did not find any compilation errors for the test models. Inspecting the code generated revealed that the code was generated according to the extended state pattern. For conformance correctness, the transitions in the models were triggered correctly in the code and all the events that were not defined were ignored (as defined in the methodology). The heuristics reported also matched the desired results, except that the time distance had a jitter of 2 – 4 milliseconds due to the possible noise in timers. To evaluate failure reporting, for the sixteen test cases that we executed, the four that were supposed to trigger the failure in the buggy SUT reported the failure and the rest did not report any failure. This increases our confidence that the generated simulators report failures correctly.

7.3.5. Completeness of the Modeling Methodology and Profile

For all the five cases, we were able to model the RTES environments with the subset of UML and MARTE that we identified and the lightweight extensions that we proposed. The models were sufficient to generate simulators that could be used to support large-scale test automation. Results in Section 7.3.1 support this claim. In one of our industrial case study, using random testing and the search-based testing, combined with using the environment model to identify error states (oracle),

new critical faults were detected. For the other cases, as discussed in 7.3.3, we were able to detect various seeded faults using the generated test cases.

For both the industrial case studies, the number of components identified at the time of domain modeling was larger than what was finally required. During successive revisions and based on insight obtained through behavioral modeling, some components turned out to be unnecessary and were removed from the domain model. One practical challenge is that it was not easy in practice to identify the right level of abstraction to model the behavior of environment components. Sub-machines were widely used to incrementally refine the behavioral models until the right level of detail was achieved to simulate the behavior of component from the viewpoint of the SUT.

8. Limitations

The focus of the work presented here was only on RTES with soft time deadlines in the order of hundreds of milliseconds, with an accepted jitter of a few milliseconds. However, the modeling methodology proposed in this paper is independent of this limitation and can be used to model systems with stricter deadlines. This limitation is due to the choice of Java as a target language for simulation. The choice of Java was made based on the needs of our industry partners but may obviously not be appropriate in other environments. To use the approach in the presence of stricter deadlines, a possible option is to use a virtual machine supporting Real Time Java (Java RT) (e.g., [46]) on a real-time operating system (e.g., [47]). We have not currently evaluated the practical implications of using Java RT, but the specifications claim that the standard java code is completely portable to a Java RT machine, which provides a precision in the order of nanoseconds.

One of the limitations of the proposed simulator generator is that we still cannot be completely sure about its correctness. As discussed earlier, to test the simulator generator, we wrote a number of test cases by hand. We also ran the generated simulator with the testing engine to test faulty RTES (artificial problems and industrial cases) and were able to trigger failures on test cases revealing seeded faults without triggering failures that were unwarranted. This increased our confidence in the correctness of the transformation rules. The evaluation of such transformations is still an open research question [51] and no suitable tool for testing model to text transformation is available yet.

Based on the experiments that we ran [44] to test different types of RTES (artificial problems and industrial cases), and as discussed above, the generator simulator seems effective in supporting test automation for fault detection. Because this is very time consuming, we however did so only on five RTES. One important question is whether our simulation rules are complete enough to simulate the environment of *any* RTES. To address this possible limitation, the industrial cases and artificial problems that we selected were from diverse domains. One case was of an automated bottle recycling system and the other was a marine seismic acquisition system. Two of the artificial problems that we developed depicted the common scenarios of RTES interactions with sensors in the environment. The third artificial problem was selected from the domain of train control systems. The diversity of the

domains of these RTES, which environment was simulated, increased our confidence in the completeness of the transformation rules and the simulation framework for the RTES developed using our modeling methodology.

We evaluated correctness and completeness of the transformation rules by generating simulators for several test models, three artificial problems, and two industrial cases. Though it cannot be guaranteed that the implemented transformation rules are complete and correct, note that this does not affect the general validity and applicability of the simulation generation approach based on environment modeling using extensions of UML. Such rules will be refined and augmented over time.

9. Conclusion

Black-box system testing of Real-time Embedded Systems (RTES) on their development platforms is required to verify the correctness of these systems without involving the deployed hardware and other physical components of their environments. This approach typically involves simulations of the behavior of environment components in a way that is transparent to the RTES. Such a strategy allows early and fully automated system testing, even when the hardware is not yet available. It is also helpful in situations where testing RTES for critical failures in their actual environments is either not feasible, too costly, or might have catastrophic consequences.

This paper reported on a model-driven automated approach for such black-box system testing strategy based on environment simulation. We purposefully took a practical angle and our approach does not require software engineers to use additional, specific notations for simulation and testing purposes, but only involve slight extensions of existing software modeling standards and a specific modeling methodology. This paper focuses on environment modeling and rules for simulator generation to enable automated black-box system testing and only briefly discusses the test generation strategies, which are reported elsewhere.

As mentioned above, to facilitate its adoption, the methodology is based on standards: UML, MARTE profile and OCL for modeling the structure, behavior, and constraints of the environment. We, and this is part of our methodology, made a conscious effort to minimize the notation subset used from these standards. Our modeling methodology entails the use of constructs (e.g., non-determinism, error states, and failure states), which are essential to enable fully automated system testing (i.e., choice, execution and evaluation of the test cases). We modeled the environment of three artificial problems and two industrial RTES in order to investigate whether our methodology and the notation subsets selected were sufficient to fully address the need for automated software testing. Our experience showed that this was the case.

Based on a careful analysis of the literature, we concluded that none of the existing code generation approaches in the literature supports the constructs required to support the testing of RTES through environment simulation. We implemented the code generation rules for the simulator using model-to-text transformations with MOFScript, thus producing a set of Java classes. Our empirical

evaluation based on our five case studies shows that the developed rules are sufficient and that they are correct as far as fault detection is concerned. The automated simulator generation is expected to save a significant amount of effort, although controlled empirical studies in industrial contexts will be necessary to support such a claim with increased confidence. By using our environment models and the generated simulators, it was possible to automatically find new, critical faults in one of the industrial case studies using fully automated, large scale random and search-based testing.

Acknowledgements.

The work presented in this paper was supported by Norwegian Research Council and was produced as part of the ITEA 2 VERDE project. We are thankful to Christine Husa, Tor Sjøwall, John Roger Johansen, Erling Marhussen, Dag Kristensen, and Anders Emil Olsen, all from Tomra and Bjørn Nordmoen, Petter Eide, Tore Andre Nilsen, and Sofia Wegger from WesternGeco for their technical support.

Appendix

The appendix provides a description of auto-generated attributes for instances of Context meta-class in Table A.1 and a description of auto-generated methods for instances of `ContextState` meta-class of the extended state pattern in Table A.2.

Table A.1. Auto generated attributes in the Context class

Attribute Name	Description
instanceId: int	Refers to a unique id for every instance in the simulation
action: ? extends ExternalCode	The reference contains an object of action class associated to the context class
states:IState[*]	An array of possible states the environment component can be in
stateContext[*]:Region	The array has object(s) when the class has an orthogonal state machine associated with it.
currentState:IState [0.. 1]	Refers to the state object that the context object is currently in.

Table A.2. Auto generated methods in the State class

Method Name	Description
Constructor	Generally empty unless the state has outgoing time events. In that case *TimeService* object is initialized
evaluateChangeEvents	Returns the condition corresponding to any of the change triggers of the out-going transition from the state that is satisfied
executeChangeEvents	The condition that is satisfied is compared and actions corresponding to the change event that the condition relates to are executed.
executeCompletionEvent	If the only outgoing transition from the state is without a trigger, then that tran-sition is taken (in this case changeState method is executed)
getStateName	Returns the name of the current state
onStateEntry	Timers corresponding to time events are initialized. For non-deterministic time events, a value for the timer is obtained from test-engine. Timers associated with «TimeProbability» transition are only initialized/reset if this is the first entry into the state after instance creation or the transition has been taken. Entry activity is executed, do activity is executed in a parallel thread, and completion event is triggered.
onStateExit	Exit activity of the state is executed
stopExecution	Stops the current thread
afterT<i> methods	An *afterT*<*i*> method is generated for every time event and is called by the timeout method if the corresponding time event is triggered. <*i*> is the automatically assigned id of the time event.

Signal methods	A signal method is generated for every signal that is defined to be accepted for the state. Action code corresponding to the transition is placed in this method along with a call to changeState method
timeout()	Calls the *afterT<i>* method whose time event has been triggered

References

[1] M. Z. Iqbal, A. Arcuri, and L. Briand, "Environment Modeling with UML/MARTE to Support Black-Box System Testing for Real-Time Embedded Systems: Methodology and Industrial Case Studies," in *Model Driven Engineering Languages and Systems*. Springer Berlin / Heidelberg, 2010, pp. 286-300.

[2] Artemis. (2011, June 13, 2011). *Artemis Joint Undertaking - The public private partnership for R & D Embedded Systems*. Available: http://artemis-ju.eu/embedded_systems

[3] B. M. Broekman and E. Notenboom, *Testing Embedded Software*: Addison-Wesley Co., Inc., 2003.

[4] OMG, "Unified Modeling Language Superstructure, Version 2.3, http://www.omg.org/spec/UML/2.3/," ed, 2010.

[5] OMG, "Modeling and Analysis of Real-time and Embedded systems (MARTE), Version 1.0, http://www.omg.org/spec/MARTE/1.0/," ed, 2009.

[6] OMG, "Object Constraint Language Specification, Version 2.2, http://www.omg.org/spec/OCL/2.2/," ed: Object Management Group Inc., 2010.

[7] E. Gamma, R. Helm, R. Johnson, and J. Vlissides, *Design patterns: elements of reusable object-oriented software*, 1995.

[8] Cheddar. (2011). Available: http://beru.univ-brest.fr/~singhoff/cheddar/

[9] G. A. Wainer, *Discrete-Event Modeling and Simulation: A Practitioner's Approach*: CRC, 2009.

[10] P. Fritzson and V. Engelson, "Modelica - A Unified Object-Oriented Language for System Modeling and Simulation," in *ECOOP'98 — Object-Oriented Programming*. Springer Berlin / Heidelberg, 1998, p. 67.

[11] W. Schamai, P. Fritzson, C. Paredis, and A. Pop, "Towards Unified System Modeling and Simulation with ModelicaML: Modeling of Executable Behavior Using Graphical Notations," presented at the 7th International Modelica Conference, Como, Italy, 2009.

[12] J. Mooney and H. Sarjoughian, "A framework for executable UML models," presented at the Proceedings of the 2009 Spring Simulation Multiconference, San Diego, California, 2009.

[13] P. M. Kruse, J. Wegener, and S. Wappler, "A highly configurable test system for evolutionary black-box testing of embedded systems," presented at the Proceedings of the 11th Annual conference on Genetic and evolutionary computation, Montreal, Canada, 2009.

[14] F. Lindlar, A. Windisch, and J. Wegener, "Integrating Model-Based Testing with Evolutionary Functional Testing," presented at the Proceedings of the 2010 Third International Conference on Software Testing, Verification, and Validation Workshops, 2010.

[15] F. Lindlar and A. Windisch, "A Search-Based Approach to Functional Hardware-in-the-Loop Testing," presented at the Proceedings of the 2nd International Symposium on Search Based Software Engineering, 2010.

[16] M. Short and M. J. Pont, "Assessment of high-integrity embedded automotive control systems using hardware in the loop simulation," *J. Syst. Softw.,* vol. 81, pp. 1163-1183, 2008.

[17] G. Francis, R. Burgos, P. Rodriguez, F. Wang, D. Boroyevich, R. Liu, and A. Monti, "Virtual Prototyping of Universal Control Architecture Systems by means of Processor in the Loop Technology " presented at the Twenty Second Annual IEEE Applied Power Electronics Conference, APEC 2007 - Twenty Second Annual IEEE 2007

[18] T. Kishi and N. Noda, "Aspect-oriented Context Modeling for Embedded Systems," presented at the Workshop on Early Aspects: Aspect-Oriented Requirements Engineering and Architecture Design, 2004.

[19] G. Karsai, S. Neema, and D. Sharp, "Model-driven architecture for embedded software: A synopsis and an example," *Science of Computer Programming,* vol. 73, pp. 26-38, 2008.

[20] K. S. Choi, S. C. Jung, H. J. Kim, D. H. Bae, and D. H. Lee, "UML-based Modeling and Simulation Method for Mission-Critical Real-Time Embedded System Development," presented at the IASTED International Conference Proceedings, 2006.

[21] C. Kreiner, C. Steger, and R. Weiss, "Improvement of Control Software for Automatic Logistic Systems Using Executable Environment Models," presented at the EUROMICRO '98: Proceedings of the 24th Conference on EUROMICRO, 1998.

[22] J. Axelsson, "Unified Modeling of Real-Time Control Systems and Their Physical Environments Using UML," presented at the Eighth Annual IEEE International Conference and Workshop on the Engineering of Computer Based Systems (ECBS '01), 2001.

[23] H. Gomaa, *Designing Concurrent, Distributed And Real-Time Applications With UML*: Addison-Wesley Educational Publishers Inc, 2000.

[24] S. Friedenthal, A. Moore, and R. Steiner, *A Practical Guide to SysML: The Systems Modeling Language*: Elsevier, 2008.

[25] M. Auguston, M. J. B, and M. Shing, "Environment behavior models for automation of testing and assessment of system safety," *Information and Software Technology,* vol. 48, pp. 971-980, 2006.

[26] M. Heisel, D. Hatebur, T. Santen, and D. Seifert, "Testing Against Requirements Using UML Environment Models," in *Fachgruppentreffen Requirements Engineering und Test, Analyse & Verifikation*, 2008, pp. 28-31.

[27] N. Adjir, P. Saqui-Sannes, and K. M. Rahmouni, "Testing Real-Time Systems Using TINA," in *Testing of Software and Communication Systems*. Lecture Notes in Computer Science, Springer Berlin / Heidelberg, 2009.

[28] K. G. Larsen, M. Mikucionis, and B. Nielsen, "Online Testing of Real-time Systems Using Uppaal," in *Formal Approaches to Software Testing*. Lecture Notes in Computer Science, Springer Berlin / Heidelberg, 2005.

[29] L. Du Bousquet, F. Ouabdesselam, J. L. Richier, and N. Zuanon, "Lutess: a specification-driven testing environment for synchronous software," presented at the ICSE '99: Proceedings of the 21st International Conference on Software Engineering, Los Angeles, California, United States, 1999.
[30] R. Pilitowski and A. Dereziñska, "Code Generation and Execution Framework for UML 2.0 Classes and State Machines," in *Innovations and Advanced Techniques in Computer and Information Sciences and Engineering*. Springer Netherlands, 2007, pp. 421-427.
[31] F. Chauvel and J.-M. Jézéquel, "Code Generation from UML Models with Semantic Variation Points," in *Model Driven Engineering Languages and Systems*. Springer Berlin / Heidelberg, 2005, pp. 54-68.
[32] SmartState. (2011). *SmartState - UML statemachine code generation tool*. Available: http://www.smartstatestudio.com/
[33] IBM. (2011). *IBM Rational Rhapsody* Available: http://www.ibm.com/software/awdtools/rhapsody/
[34] M. Samek, *Practical UML statecharts in C/C++: event-driven programming for embedded systems*: Newnes, 2009.
[35] L. Ferreira and C. Rubira, "The reflective state pattern," presented at the Proceedings of the Pattern Languages of Program Design, Monticello, Illinois-USA, 1998.
[36] B. Chin and T. Millstein, "An Extensible State Machine Pattern for Interactive Applications," in *ECOOP 2008 – Object-Oriented Programming*. Springer Berlin / Heidelberg, 2008, pp. 566-591.
[37] N. Holt, B. Anda, K. Asskildt, L. Briand, J. Endresen, and S. Frøystein, "Experiences with Precise State Modeling in an Industrial Safety Critical System," presented at the Critical Systems Development Using Modeling Lanuguages, CSDUML'06, 2006.
[38] G. Palfinger, "State Action Mapper," presented at the 4th Pattern Languages of Programming (PLoP), USA, 1997.
[39] I. A. Niaz and J. Tanaka, "An object-oriented approach to generate Java code from UML Statecharts," *International Journal of Computer & Information Science,* vol. 6, pp. 83-98, 2005.
[40] C. Larman, *Applying UML and Patterns: An Introduction to Object-Oriented Analysis and Design and the Unified Process*: Prentice Hall PTR Upper Saddle River, NJ, USA, 2001.
[41] OMG, "Concrete Syntax for UML Action Language (Action Language for Foundational UML - ALF), Version 1.0 - Beta 1, http://www.omg.org/spec/ALF/," ed, 2010.
[42] J. D. Musa, "The operational profile in software reliability engineering: an overview," presented at the Third International Symposium on Software Reliability Engineering, 1992.
[43] J. Oldevik, "MOFScript user guide," *Version 0.6 (MOFScript v 1.1. 11),* 2006.
[44] A. Arcuri, M. Iqbal, and L. Briand, "Black-Box System Testing of Real-Time Embedded Systems Using Random and Search-Based Testing," in *Testing Software and Systems*. Springer Berlin / Heidelberg, 2010, pp. 95-110.
[45] B. Bruegge and A. Dutoit, *Object-Oriented Software Engineering Using UML, Patterns, and Java*: Prentice Hall Press, 2009.
[46] *Sun Java Real-Time System*. Available: http://java.sun.com/javase/technologies/realtime/index.jsp, last accessed 09/02/2012
[47] *SUSE Linux Enterprise Real Time Extension*. Available: http://www.novell.com/products/realtime/, last accessed: 09/02/2012
[48] P. McMinn, "Search based software test data generation: a survey," *Software Testing, Verification and Reliability,* vol. 14, pp. 105-156, 2004.
[49] S. Ali, M. Z. Iqbal, A. Arcuri, and L. Briand, "A Search-based OCL Constraint Solver for Model-based Test Data Generation," presented at the 11th International Conference on Quality Software, 2011.
[50] M. Zheng, V. Alagar, and O. Ormandjieva, "Automated generation of test suites from formal specifications of real-time reactive systems," *The Journal of Systems & Software,* vol. 81, pp. 286-304, 2008.
[51] C. Fiorentini, A. Momigliano, M. Ornaghi, and I. Poernomo, "A constructive approach to testing model transformations," in *Theory and Practice of Model Transformations*, 2010, pp. 77-92.
[52] J. Andrews, L. Briand, Y. Labiche, and A. Namin, "Using mutation analysis for assessing and comparing testing coverage criteria," *IEEE Transactions on Software Engineering,* vol. 32, pp. 608-624, 2006.