

Automated Testing of Distributed and Heterogeneous Systems Based on UML Sequence Diagrams

Bruno Lima^{1,2}(✉) and João Pascoal Faria^{1,2}

¹ INESC TEC, FEUP campus, Rua Dr. Roberto Frias, s/n, 4200-465 Porto, Portugal

² Faculty of Engineering, University of Porto, Rua Dr. Roberto Frias, s/n, 4200-465 Porto, Portugal
{bruno.lima, jpf}@fe.up.pt

Abstract. The growing dependence of our society on increasingly complex software systems makes software testing ever more important and challenging. In many domains, several independent systems, forming a distributed and heterogeneous system of systems, are involved in the provisioning of end-to-end services to users. However, existing test automation techniques provide little tool support for properly testing such systems. Hence, we propose an approach and toolset architecture for automating the testing of end-to-end services in distributed and heterogeneous systems, comprising a visual modeling environment, a test execution engine, and a distributed test monitoring and control infrastructure. The only manual activity required is the description of the participants and behavior of the services under test with UML sequence diagrams, which are translated to extended Petri nets for efficient test input generation and test output checking at runtime. A real world example from the Ambient Assisted Living domain illustrates the approach.

Keywords: Software testing · Distributed systems · UML sequence diagrams · Heterogeneous systems · Systems of systems

1 Introduction

Due to the increasing ubiquity, complexity, criticality and need for assurance of software based systems [2], testing is a fundamental lifecycle activity, with a huge economic impact if not performed adequately [21].

In a growing number of domains, the provision of services to end users depends on the correct functioning of large and complex systems of systems [5]. A system of systems consists of a set of small independent systems that together form a new system, combining hardware components and software systems. Systems of systems are in most cases distributed and heterogeneous, involving mobile and cloud-based platforms.

Testing these distributed and heterogeneous systems is particularly important and challenging. Some of the challenges are: the difficulty to test the system

as a whole due to the number and diversity of individual components; the difficulty to coordinate and synchronize the test participants and interactions, due to the distributed nature of the system; the difficulty to test the components individually, because of the dependencies on other components.

An example of a distributed and heterogeneous system is the Ambient Assisted Living (AAL) ecosystem that was prototyped in the context of the nationwide AAL4ALL project [1]. The AAL4ALL ecosystem comprises a set of interoperable AAL products and services (sensors, actuators, mobile and web-based applications and services, middleware components, etc.), produced by different manufacturers using different technologies and communication protocols (web services, message queues, etc.). To assure interoperability and the integrity of the ecosystem, it was developed and piloted a testing and certification methodology [6], to be applied on candidate components of the ecosystem. A major problem faced during test implementation and execution was related with test automation, due to the diversity of component types and communication interfaces, the distributed nature of the system, and the lack of support tools. Similar difficulties have been reported in other domains, such as the railway domain [22]. In fact, we found in the literature limited tool support for automating the whole process of specification-based testing of distributed and heterogeneous systems.

Hence, the main objective of this paper is to propose an approach and a toolset architecture to automate the whole process of model-based testing of distributed and heterogeneous systems in a seamless way, with a focus on integration testing, but supporting also unit (component) and system testing. As compared to existing approaches, the proposed approach and architecture provide significant benefits regarding efficiency and effectiveness: the only manual activity required from the tester is the creation (with tool support) of partial behavioral models of the system under test (SUT), using feature-rich industry standard notations (UML 2 sequence diagrams), together with model-to-implementation mapping information, being all the needed runtime test components provided by the toolset for different platforms and technologies; the ability to test not only the interactions of the SUT with the environment, but also the interactions among components of the SUT, following an adaptive test generation and execution strategy, to improve fault detection and localization and cope with non-determinism in the specification or the SUT.

The rest of the paper is organized as follows: Sect. 2 describes the state of the art. Section 3 presents an overview of the proposed approach and test process. Section 4 introduces the toolset architecture. Section 5 summarizes the novelties and benefits of the proposed approach. Section 6 concludes the paper and points out future work. A running example from the AAL domain is used to illustrate the approach presented.

2 State of the Art

The highest level of test automation is achieved by automating both test generation and test execution, but the approaches for automating test generation

and automating test execution are in most cases orthogonal. Hence, we analyze in separate subsections approaches for automatic test generation (from models or specifications) and automatic test execution that have the potential to be applied for distributed and heterogeneous systems.

2.1 Model-Based Test Generation

Model-based testing (MBT) techniques and tools have attracted increasing interest from academia and industry [24], because of their potential to increase the effectiveness and efficiency of the test process, by means of the automatic generation of test cases (test sequences, input test data, and expected outputs) from behavioral models of the system under test (SUT).

However, MBT approaches found in the literature suffer from several limitations [4]. The most common limitation is the lack of integrated support for the whole test process. This is a big obstacle for the adoption of these approaches by industry, because of the effort required to create or adapt tools to implement some parts of the test process.

Another common problem with existing MBT approaches is the difficulty to avoid the explosion of the number of test cases generated. In recent MBT approaches [8, 18], researchers try to overcome the test case explosion problem by the usage of behavioral models focusing on specific scenarios, i.e., by following a scenario-based testing approach instead of a state-based testing approach. Being a feature-rich industry standard, UML 2 sequence diagrams (SDs) are particularly well suited for supporting scenario-based MBT approaches. With the features introduced in UML 2, parameterized SDs can be used to model both simple and complex behavioral scenarios, with control flow variants, temporal constraints, and conformance control operators. UML SDs are also well suited for modeling the interactions that occur between the components and actors of a distributed system.

In the literature it can be found some test automation approaches based on UML SDs, but those approaches has some limitations for the testing of distributed and heterogeneous systems, namely regarding the support for features specific to those systems, such as parallelism, concurrency and time constraints.

Of particular relevance in the context of this paper is the UML Checker toolset developed in recent work of the authors [7, 8], with several advantages over other approaches, namely regarding the level of support of UML 2 features. The toolset supports the conformance testing of standalone object-oriented applications against test scenarios specified by means of so called test-ready SDs. Test-ready SDs are first translated to extended Petri Nets for efficient incremental conformance checking, with a limited support for parallelism and concurrency. Besides external interactions with users and client applications, internal interactions between objects in the system are also monitored using Aspect-Oriented Programming (AOP) techniques [16], and checked against the ones specified in the model. The testing of distributed systems is not supported, but some of the techniques developed have the potential to be reused for the modeling and

testing of interactions between components in a distributed system, instead of interactions between objects in a standalone application.

Other examples of test automation approaches based on UML SDs are the SCENTOR tool, targeting e-business EJB applications [27], the MDA-based approach of [14], and the IBM Rational Rhapsody TestConductor AddOn [13], targeting real-time embedded applications. A comparison of the strengths and weaknesses of these approaches can be found in [8]. The main limitations of these approaches are the limited support for the new features of UML 2 SDs and the limited support for testing internal interactions (besides the interactions with the environment).

2.2 Test Execution Frameworks

Regarding test concretization and execution for distributed systems, we found in the literature several frameworks that can be adapted and integrated for building a comprehensive test automation solution.

The Software Testing Automation Framework (STAF) [20] is an open source, multi-platform, multi-language framework designed around the idea of reusable components, called services (such as process invocation, resource management, logging, and monitoring). STAF removes the tedium of building an automation infrastructure, thus enabling the tester to focus on building an automation solution. The STAF framework provides the foundation upon which to build higher level solutions, and provides a pluggable approach supported across a large variety of platforms and languages.

Torens and Ebrecht [22] proposed the RemoteTest framework as a solution for the testing of distributed systems and their interfaces. In this framework, the individual system components are integrated into a virtual environment that emulates the adjacent modules of the system. The interface details are thereby abstracted by the framework and there is no special interface knowledge necessary by the tester. In addition to the decoupling of components and interface abstraction, the RemoteTest framework facilitates the testing of distributed systems with flexible mechanisms to write test scripts and an architecture that can be easily adapted to different systems.

Zhang et al. [28] developed a runtime monitoring tool called FiLM that can monitor the execution of distributed applications against LTL specifications on finite traces. Implemented within the online predicate checking infrastructure D³S [17], FiLM models the execution of distributed applications as a trace of consistent global snapshots with global timestamps, and it employs finite automata constructed from Labelled transition systems (LTL) specifications to evaluate the trace of distributed systems.

Camini et al. [3] proposed DiCE, an approach that continuously and automatically explores the system behavior, to check whether the system deviates from its desired behavior. At a highlevel DiCE (i) creates a snapshot consisting of lightweight node checkpoints, (ii) orchestrates the exploration of relevant system behaviors across the snapshot by subjecting system nodes to many possible inputs that exercise node actions, and (iii) checks for violations of properties

that capture the desired system behavior. DiCE starts exploring from current system state, and operates alongside the deployed system but in isolation from it. In this way, testing can account for the current code, state and configuration of the system. DiCE reuses existing protocol messages to the extent possible for interoperability and ease of deployment.

One difficulty in testing distributed systems is that their distributed nature imposes theoretical limitations on the conformance faults that can be detected by the test components, depending on the test architecture used [11, 12]. Hierons [11] devised a hybrid framework for solving the problem that exist in many systems that interact with their environment at distributed interfaces without the possibility in some cases to place synchronised local testers at the ports of the SUT. Before this framework existed only two main approaches to test this type of systems: having independent local testers [23] or a single centralised tester that interacts asynchronously with the SUT. The author proved that the hybrid framework is more powerful than the distributed and centralised approaches.

2.3 Synthesis

Although we didn't find in the literature an integrated approach for fully automating the testing of distributed and heterogeneous systems, the concepts used by each can be harnessed in the development of an architecture that can be fully supported by tools, so that all the testing process can be automated.

3 Approach and Process

Our main objective is the development of an approach and a toolset to automate the whole process of model-based testing of distributed and heterogeneous systems in a seamless way, with a focus on integration testing, but supporting also unit (component) and system testing. The only manual activity (to be performed with tool support) should be the creation of the input model of the SUT.

To that end, our approach is based on the following main ideas:

- the adoption of different ‘frontend’ and ‘backend’ modeling notations, with an automatic translation of the input behavioral models created by the user in an accessible ‘frontend’ notation (using industry standards such as UML [19]), to a formal ‘backend’ notation amenable for incremental execution at runtime (such as extended Petri Nets as in our previous work for object-oriented systems [8]);
- the adoption of an online and adaptive test strategy, where the next test input depends on the sequence of events that has been observed so far and the resulting execution state of the formal backend model, to allow for non-determinism in either the specification or the SUT [11];
- the automatic mapping of test results (coverage and errors) to the ‘frontend’ modeling layer.

Figure 1 depicts the main activities and artifacts of the proposed test process based on the above ideas. The main activities are described in the next subsections and illustrated with a running example.

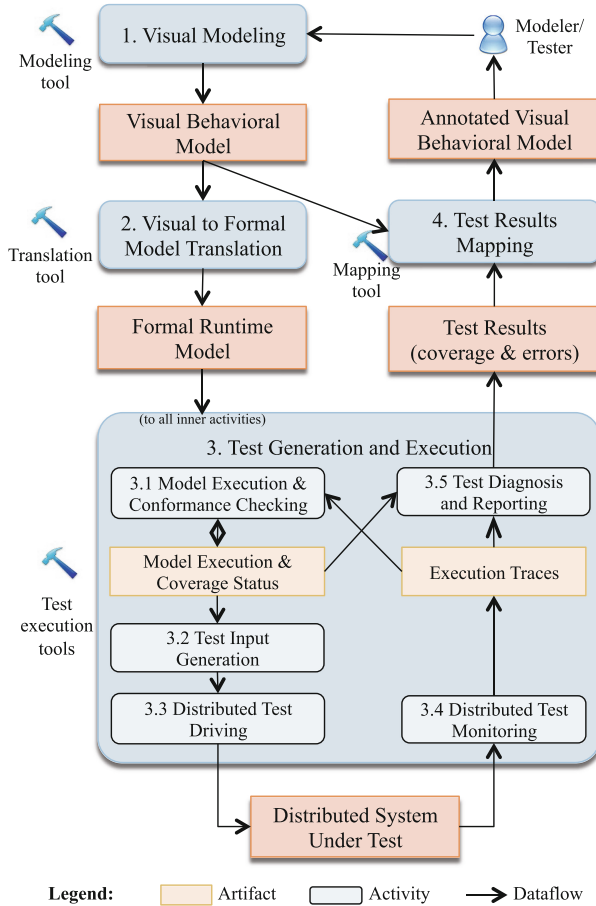


Fig. 1. Dataflow view of the proposed test process.

3.1 Visual Modeling

The behavioral model is created using an appropriate UML profile [10, 19] and an existing modeling tool. We advocate the usage of UML 2 SDs, with a few restrictions and extensions, because they are well suited for describing and visualizing the interactions that occur between the components and actors of a distributed system. UML deployment diagrams can also be used to describe the distributed structure of the SUT. Mapping information between the model and the implementation, needed for test execution (such as the actual location of each component under test), may also be attached to the model with tagged values.

To illustrate the approach, we use a real world example from the AAL4ALL project, related with a fall detection and alert service. As illustrated in Fig. 2, this service involves the interaction between different heterogeneous components

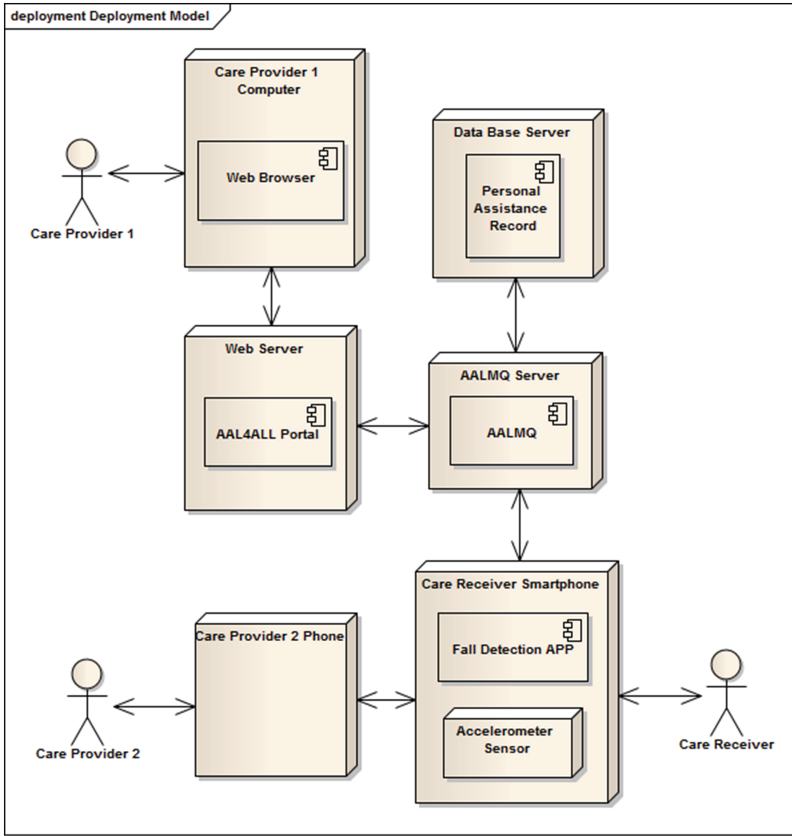


Fig. 2. UML deployment diagram of a fall detection scenario.

running in different hardware nodes in different physical locations, as well as three users.

A behavioral model for a typical fall detection scenario is shown in Fig. 3. In this scenario, a care receiver has a smartphone that has installed a fall detection application. When this person falls, the application detects the fall using the smartphone's accelerometer and provides the user a message which indicates that it has detected a drop giving the possibility for the user to confirm whether he/she needs help. If the user responds that he/she does not need help (the fall was slight, or it was just the smartphone that fell to the ground), the application does not perform any action; however, if the user confirms that needs help or does not respond within 5 s (useful if the person became unconscious due to the fall), the application raises two actions in parallel. On the one hand, it makes a call to a previously clearcut number to contact a health care provider (in this case can be a formal or informal caregiver); on the other hand, it sends the fall occurrence for a Personal Assistance Record database and sends a message to a

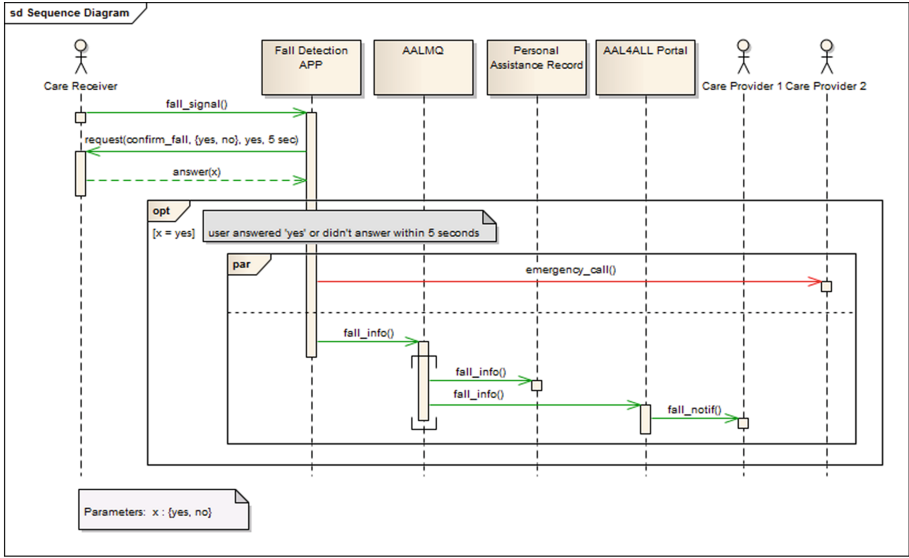


Fig. 3. UML sequence diagram representing the interactions of the fall detection scenario. The diagram is already painted after a failed test execution in which the fall detection application didn't send an emergency call.

portal that is used by a caregiver (e.g. a doctor or nurse) that is responsible for monitoring this care receiver. The last two actions are performed through a central component of the ecosystem called AALMQ (AAL Message Queue), which allows incoming messages to be forwarded to multiple subscribers, according to the publish-subscribe pattern [9]. To facilitate the representation of a request for input from the user with a timeout and a default response, we use the special syntax `request(confirm_fall, {yes, no}, yes, 5 sec)`, where the first argument identifies the message, the second argument is the set of valid answers, the third is the default answer in case of timeout, and the last argument is the timeout time.

3.2 Visual to Formal Model Translation

For the formal runtime model, we advocate the usage of *Event-Driven Colored Petri Nets* – a sort of extended Petri Nets proposed in our previous work for testing object-oriented systems [8], with the addition of time constraints as found in Timed Petri Nets. We call the resulting Petri Nets *Timed Event-Driven Colored Petri Nets*, or TEDCPN for short. Petri Nets are well suited for describing in a rigorous and machine processable way the behavior of distributed and concurrent systems, usually requiring fewer places than the number of states of equivalent finite state machines. Translation rules from UML 2 SDs to Event-Driven Colored Petri Nets have been defined in [8]. Rules for translating time and duration constraints in SDs to time constraints in the resulting Petri Net can also be defined.




Fig. 4. TEDCPN derived from the SD of Fig. 3. The net is marked in a final state of a failed test execution in which the fall detection application didn't send an emergency call.

Figure 4 shows the TEDCPN derived from the SD of Fig. 3, according to the rules described in [8] and additional rules for translating time constraints.

The generated TEDCPN is partitioned into a set of fragments corresponding to the participants in the source SD. Each fragment describes the behavior local to each participant and the communication with other participants via boundary places.

Transitions may be optionally labeled with an *event*, a *guard* (with braces) and a *time interval* (with square brackets). Events correspond to the sending or receiving of messages in the source SD. Guards correspond to the conditions of conditional interaction fragments in the source SD. Time intervals correspond to duration and time constraints in the source SD. A transition can only fire when there is at least one token in each input place, the event (if defined) has occurred, the guard (if defined) holds, and the time elapsed since the transition became enabled (i.e., since there is a token in each input place) lies within the time interval (if defined).

Incoming and outgoing arcs of a transition may be labeled with a pattern matching expression describing the value (token) to be taken from the source place or put in the target place, respectively, being 1 the default. For example, in Fig. 4 the transition labeled “?answer(x)” has an input arc labeled “x”, where “x” represents a local variable of the transition. The transition can only fire if the value of the token in the source place is the same as the value of the argument of the event. Then, the value of “x” is placed in the target place.

For testing purposes, the events in the runtime model are marked as *observable* (default) or *controllable*. Controllable events (underlined) are to be injected by the test harness (playing the role of a test driver, simulating an actor) when the corresponding transition becomes enabled. Controllable events correspond to the sending of messages from actors in the source SD. All other events are observable, i.e., they are to be monitored by the test harness. For example, when the TEDCPN of the example starts execution (i.e., a token is put in the start place), the initial unlabeled transition is executed and a token is placed in the initial place of each fragment. At that point, the only transition enabled is the one labeled with the “!fall_signal” controllable event, so the test harness will inject that event (simulating the user) and test execution proceeds.

This mechanism provides a unified framework with monitoring, testing and simulation capabilities. In one extreme case, all events in the model may be marked as observable, in which case the test system acts as a runtime monitoring and verification system. In the other extreme case, all events in the model may be marked as controllable, in which case the test system acts as a simulation system. This also allows the usage of the same model with different markings of observable and controllable events for integration and unit testing.

3.3 Test Generation and Execution

Test Generation. Using the UML 2 interaction operators, a single SD, and hence the TEDCPN derived from it, may describe multiple control flow variants, that require multiple test cases for being properly exercised.

In the running example, from the reading of the set of interactions represented in Fig. 3, one easily realizes that there are three test paths to be exercised (with at least one test case for each test path). The first test path (TP1) is the case where the care receiver responds negatively to the application and the application doesn’t trigger any action. The second test path (TP2) is the situation where the user confirms to the application that he/she needs help and after that the application triggers the actions. The last test path (TP3) corresponds to the situation where the user doesn’t answer within the defined time limit and the application triggers the remaining actions automatically. If one wants also to exercise the boundary values of allowed response time (close to 0 and close to 5 seconds), then two test cases can be considered for each of the test paths TP1 and TP2, resulting in a total of 5 test cases.

Equivalently, in order to exercise all nodes, edges and boundary values in the TEDCPN, several test cases are needed. In the example, one could exercise the two outgoing paths after the “?conf_fall” event, the two possible values of variable “x” in the “!answer(x)” event, and the two boundary values of the “[0, 5 sec]” interval, in a total of 5 test cases.

In general, the required test cases can be generated using an offline strategy (with separate generation and execution phases) or an online test strategy (with intermixed generation and execution phases) [25]. In an offline strategy, the test cases are determined by a static analysis of the model, assuming the SUT behaves deterministically. But that is not often the case, so we prefer an online,

adaptive, strategy, in which the next test action is decided based on the current execution state. Whenever multiple alternatives can be taken by the test harness in an execution state, the test harness must choose one of the alternatives and keep track of unexplored alternatives (i.e., model coverage information) to be exercised in subsequent test repetitions.

Test Execution. Test execution involves the simultaneous execution of: (i) the set of components under test (CUTs); (ii) the formal runtime model (TEDCPN), dictating the possible test inputs and the expected outputs from the CUTs in each step of test execution; (iii) a local test component for each CUT, running in the same node of the CUT, able to perform the roles of test driver (i.e., send test inputs to the CUT, simulating an actor) and test monitor (i.e., monitor all the messages sent or received by the CUT).

The collection of monitored events (message sending and receiving events) forms an *execution trace*. Testing succeeds if the observed execution trace *conforms* to the formal behavioral model, in the sense that it belongs to the (possibly infinite) set of valid traces defined by the model.

Conformance checking is performed *incrementally* as follows: (i) initially, the execution of the TEDCPN is started by placing a token in the start place and firing transitions until a quiescent state is reached (a state where no transition can fire); (ii) each time a quiescent state is reached having an enabled transition labeled with a controllable event, the test harness itself generates the event (i.e., the message specified in the event is sent to the target CUT by the appropriate test driver) and the execution status of the TEDCPN is advanced to a new quiescent state; (iii) each time an observable event is monitored (by a test monitor), the execution state of the TEDCPN is advanced until a new quiescent state is reached; (iv) the two previous steps are repeated until the final state of the TEDCPN is reached (i.e., a token is placed in the final place), in which case test execution succeeds, or until a state is reached in which there is no controllable event enabled and no observable event has been monitored for a defined wait time, in which case test execution fails. The latter situation is illustrated in Fig. 4. Depending on the conformance semantics chosen, the observation of an unexpected event may also be considered a conformance error.

To minimize communication overheads, the TEDCPN can itself be executed in a *distributed* fashion, by executing each fragment of the ‘global’ TEDCPN (describing the behavior local to one participant and the communication with other participants via boundary places) by a local test component. Communication between the distributed test components is only needed when tokens have to be exchanged via boundary places.

When a final (success or failure) state is reached, the *Test Diagnosis and Reporting* activity is responsible to analyze the execution state of the TEDCPN and the collected execution trace, and produce meaningful error information.

Model coverage information is also collected during test execution, to guide the selection of test inputs and the decision about when to stop test execution, as follows: when it is reached a quiescent state of the TEDCPN with multiple

controllable events enabled leading to different execution paths, the test harness shall generate an event that leads to a previously unexplored path; when a final state of the TEDCPN is reached, test execution is restarted if there are still unexplored (but reachable) paths.

3.4 Test Results Mapping

At the end of test execution it is important to reflect the test results back in the visual behavioral model created by the user. As an example, the marking shown in the net of Fig. 4 corresponds to the final state of a failed test execution in which the Fall Detection App didn't send an emergency call. By a simple analysis of this final state (and traceability information between the source SD and the TEDCPN), it is possible to point out to the tester which messages in the source SD were covered and what was the cause of test failure (missing "emergency call" message), as shown in Fig. 3.

4 Toolset Architecture

Figure 5 depicts a layered architecture of a toolset for supporting the test process described in the previous section, promoting reuse and extensibility.

At the bottom layer in Fig. 5, the SUT is composed by a set of components under test (CUT), executing potentially in different nodes [19]. The CUT interact with each other (usually asynchronously) and with the environment (users or external systems) through well defined interfaces at defined interaction points or ports [10,11].

The three layers of the toolset are described in the following sections.

4.1 Visual Modeling Environment

At the top layer, we have a visual modeling environment, where the tester can create a *visual behavioral model* of the SUT, invoke test generation and execution, and visualize test results and coverage information back in the model.

This layer also includes a *translation tool* to automatically translate the visual behavioral models created by the user into the formal notation accepted by the test execution manager in the next layer, and a *mapping tool* to translate back the test results (coverage and error information) to annotations in the visual model.

The model transformations can be implemented using existing MDA technologies and tools [26].

4.2 Test Execution Engine

At the next layer, the test execution engine is the core engine of the toolset. It comprises a *model execution & conformance checking engine*, responsible for incrementally checking the conformance of observed execution traces in the SUT

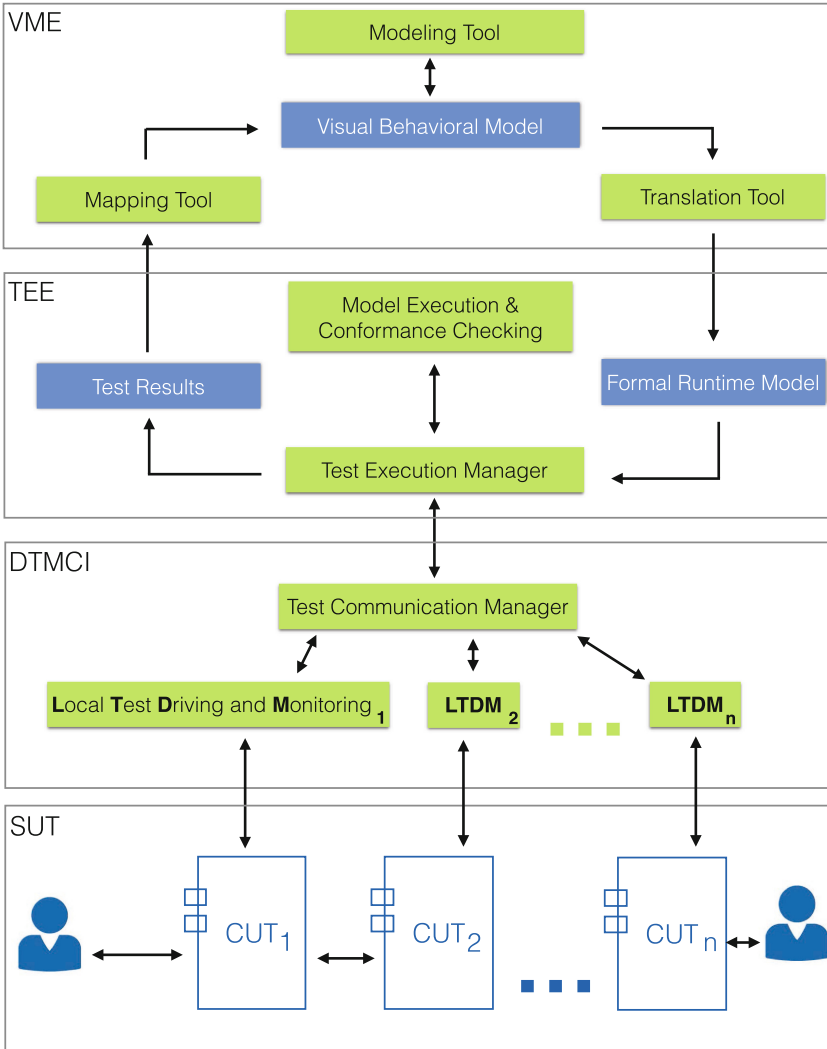


Fig. 5. Toolset architecture.

against the *formal runtime model* derived from the previous layer, and a *test execution manager*, responsible for initiating test execution (using the services of the next layer), forward execution events (received from the next layer) to the model execution & conformance checking engine, decide next actions to be performed by the local test driving and monitoring components in the next layer of the system, and produce *test results* and diagnosis information for the layer above.

The model execution & conformance checking engine can be implemented by adapting existing Petri net engines, such as CPN Tools [15].

4.3 Distributed Test Monitoring and Control Infrastructure

We adopt a hybrid test monitoring approach as proposed in [11], combining a centralized ‘tester’ and a local ‘tester’ at each port (component interaction point) of the SUT, that was shown to lead to more effective testing than a purely centralized approach (where a centralized tester interacts asynchronously with the ports of the SUT) or a purely distributed approach (where multiple independent distributed testers interact synchronously with the ports of the SUT).

Hence, the Distributed Test Monitoring and Control Infrastructure comprises a set of *local test driving and monitoring* (LTDM) components, each communicating (possibly synchronously) with a component under test (CUT), performing the roles of test monitor, driver and stub; and a *test communication manager* (TCM) component, that (asynchronously) dispatches control orders (coming from the previous layer) to the LTDMs and aggregates monitoring information from the LTDMs (to be passed to the previous layer).

During test execution, the TEDCPN may be executed in a centralized or a distributed mode, depending on the processing capabilities that can be put in the LTDM components. In centralized mode, the LTDM components just monitor all observable events of interest and send them to the central TEM; they also inject controllable events when requested from the central TEM. In distributed mode, a copy of each fragment (up to boundary places) is sent to the respective LTDM component for local execution. When there is the need to send a token to a boundary place, the LTDM sends the token to the central TEM, which subsequently dispatches it to the consumer LTDM. Because of possible delays in the communication of tokens through boundary places, the LTDM components must be prepared to tentatively accept observable events before receiving enabling tokens in boundary places.

This infrastructure may be implemented by adapting and extending existing test frameworks for distributed systems, such as the ones described in Sect. 2.2.

Different LTDM components have to be implemented for different platforms and technologies under test, such as WCF (Windows Communication Foundation), Java EE (Java Platform, Enterprise Edition), Android, etc. However, a LTDM component implemented for a given technology may be reused without change to monitor and control any CUT that uses that technology. For example, in our previous work for automating the scenario-based testing of standalone applications written in Java, we developed a runtime test library able to trace and manipulate the execution of any Java application, using AOP (aspect-oriented programming) instrumentation techniques with load-time weaving. In the case of a distributed Java application, we would need to deploy a copy of that library (or, more precisely, a modified library, to handle communication) together with each Java component under test. In the case of a distributed system implemented using other technologies (with different technologies for different components in case of heterogeneous systems), similar test monitoring components suitable for the technologies involved will have to be deployed.

5 Synthesis of Novelties and Benefits

As compared to existing approaches (see Sect. 2), the approach proposed in this paper provides the following novelties and benefits.

Our approach provides a higher level of automation of the testing process because all phases of the test process are supported in an integrated fashion. The only manual activity needed is the development in a user friendly notation of the model required as input for automatic test case generation and execution; there is no need to develop test components specific for each SUT.

This approach also provides a higher fault detection capability. The use of a hybrid test architecture allows the detection of a higher number of errors as compared to purely distributed or centralized architectures. Interactions between components in the SUT are also monitored and checked against the specification, besides the interactions of the SUT with the environment. To facilitate fault diagnosis, it is used an incremental conformance checking algorithm allowing to capture the execution state of the SUT as soon as a failure occurs. Because of the support for temporal constraints, timing faults can also be detected. Our approach has the ability to test non-deterministic SUT behaviors, using an online, adaptive, test generation strategy.

The proposed approach provides easier support for multiple test levels because the same input model can be used to perform tests at different levels (unit, integration, and system testing), simply by changing the selection of observable and controllable events in the input model. A scenario-oriented approach simplifies the level of detail required in the input models.

With this approach the test execution process is more efficient. With a distributed conformance checking algorithm, communication overheads during test execution are minimized and the usage of a state-oriented runtime model allows a more efficient model execution and conformance checking.

6 Conclusions

In this paper, it was presented a novel approach and process for automated scenario-based testing of distributed and heterogeneous systems. It was also presented the architecture of a toolset able to support and automate the proposed test process. Based in a multilayer architecture and using a hybrid test monitoring approach combining a centralized ‘tester’ and a local ‘tester’ this toolset promotes reuse and extensibility. In the approach proposed, the tester interacts with a visual modeling front-end to describe key behavioral scenarios of the SUT using UML sequence diagrams, invoke test generation and execution, and visualize test results and coverage information back in the model using a color scheme (see Fig. 3). Internally, the visual modeling notation is converted to a formal notation amenable for runtime interpretation (see Fig. 4) in the back-end. A distributed test monitoring and control infrastructure is responsible for interacting with the components of the SUT, under the roles of test driver, monitor and stub. At the core of the toolset, a test execution engine coordinates

test execution and checks the conformance of the observed execution trace with the expectations derived from the visual model. For better understanding the approach and toolset architecture proposed, a real world example from the AAL domain was presented along the paper.

As future work we will implement a toolset following the architecture (represented in Fig. 5) and working principles presented in this paper, taking advantage of previous work for automating the integration testing of standalone object-oriented systems. To experimentally assess the benefits of the approach and toolset, industrial level case studies will be conducted, with at least one in the AAL domain.

With such a toolset, we expect to significantly reduce the cost of testing distributed and heterogeneous systems, from the standpoint of time, resources and expertise required, as compared to existing approaches.

References

1. AAL4ALL: Ambient Assisted Living For All (2015). <http://www.aal4all.org>
2. Boehm, B.: Some future software engineering opportunities and challenges. In: Nanz, S. (ed.) *The Future of Software Engineering*, pp. 1–32. Springer, Heidelberg (2011). http://dx.doi.org/10.1007/978-3-642-15187-3_1
3. Canini, M., Jovanović, V., Venzano, D., Novaković, D., Kostić, D.: Online testing of federated and heterogeneous distributed systems. *SIGCOMM Comput. Commun. Rev.* **41**(4), 434–435 (2011). <http://doi.acm.org/10.1145/2043164.2018507>
4. Dias Neto, A.C., Subramanyan, R., Vieira, M., Travassos, G.H.: A survey on model-based testing approaches: a systematic review. In: *Proceedings of the 1st ACM International Workshop on Empirical Assessment of Software Engineering Languages and Technologies: Held in Conjunction with the 22nd IEEE/ACM International Conference on Automated Software Engineering (ASE) 2007, WEASELTech 2007*, pp. 31–36. ACM, New York (2007). <http://doi.acm.org/10.1145/1353673.1353681>
5. DoD: *Systems Engineering Guide for Systems of Systems*. Technical report, Office of the Deputy Under Secretary of Defense for Acquisition and Technology, Systems and Software Engineering Version 1.0 (2008)
6. Faria, J.P., Lima, B., Sousa, T.B., Martins, A.: A testing and certification methodology for an open ambient-assisted living ecosystem. *Int. J. E-Health Med. Commun. (IJEHMC)* **5**(4), 90–107 (2014)
7. Faria, J.: A toolset for conformance testing against UML sequence diagrams (2014). <https://blogs.fe.up.pt/sdbt/>
8. Faria, J., Paiva, A.: A toolset for conformance testing against UML sequence diagrams based on event-driven colored Petri nets. *Int. J. Softw. Tools Technol. Transf.*, pp. 1–20 (2014). <http://dx.doi.org/10.1007/s10009-014-0354-x>
9. Gamma, E., Helm, R., Johnson, R., Vlissides, J.: *Design Patterns: Elements of Reusable Object-Oriented Software*. Pearson Education, Upper Saddle River (1994)
10. Gross, H.G.: *Component-Based Software Testing with UML*. Springer, Heidelberg (2005)
11. Hierons, R.M.: Combining centralised and distributed testing. *ACM Trans. Softw. Eng. Methodol.* **24**(1), 5:1–5:29 (2014). <http://doi.acm.org/10.1145/2661296>

12. Hierons, R.M., Merayo, M.G., Núñez, M.: Scenarios-based testing of systems with distributed ports. *Softw. Pract. Experience* **41**(10), 999–1026 (2011). <http://dx.doi.org/10.1002/spe.1062>
13. IBM: IBM® Rational® Rhapsody® Automatic Test Conductor Add On User Guide, v2.5.2 (2013)
14. Javed, A., Strooper, P., Watson, G.: Automated generation of test cases using model-driven architecture. In: *Second International Workshop on Automation of Software Test, 2007, AST 2007*, p. 3, May 2007
15. Jensen, K., Kristensen, L., Wells, L.: Coloured Petri Nets and CPN tools for modelling and validation of concurrent systems. *Int. J. Softw. Tools Technol. Transf.* **9**(3–4), 213–254 (2007). <http://dx.doi.org/10.1007/s10009-007-0038-x>
16. Kiczales, G., Lamping, J., Mendhekar, A., Maeda, C., Lopes, C., Loingtier, J.M., Irwin, J.: Aspect-oriented programming. In: Akşit, M., Matsuoka, S. (eds.) *ECOOP 1997. LNCS*, vol. 1241, pp. 220–242. Springer, Heidelberg (1997). <http://dx.doi.org/10.1007/BFb0053381>
17. Liu, X., Guo, Z., Wang, X., Chen, F., Lian, X., Tang, J., Wu, M., Kaashoek, M.F., Zhang, Z.: D3S: debugging deployed distributed systems. In: *NSDI*, vol. 8, pp. 423–437 (2008)
18. Moreira, R.M., Paiva, A.C.: PBGT tool: an integrated modeling and testing environment for pattern-based GUI testing. In: *Proceedings of the 29th ACM/IEEE International Conference on Automated Software Engineering, ASE 2014*, pp. 863–866. ACM, New York (2014). <http://doi.acm.org/10.1145/2642937.2648618>
19. OMG: *OMG Unified Modeling Language™ (OMG UML), Superstructure*. Technical report, Object Management Group (2011)
20. STAF: *Software Testing Automation Framework (STAF)* (2014). <http://staf.sourceforge.net/>
21. Tassef, G.: *The Economic Impacts of Inadequate Infrastructure for Software Testing*. Technical report, National Institute of Standards and Technology (2002)
22. Torens, C., Ebrecht, L.: RemoteTest: a framework for testing distributed systems. In: *2010 Fifth International Conference on Software Engineering Advances (ICSEA)*, pp. 441–446 August 2010
23. Ulrich, A., König, H.: Architectures for testing distributed systems. In: Csopaki, G., Dibuz, S., Tarnay, K. (eds.) *Testing of Communicating Systems. IFIP – The International Federation for Information Processing*, vol. 21, pp. 93–108. Springer, US (1999). http://dx.doi.org/10.1007/978-0-387-35567-2_7
24. Utting, M., Legeard, B.: *Practical Model-Based Testing: A Tools Approach*. Morgan Kaufmann Publishers Inc., San Francisco (2007)
25. Utting, M., Pretschner, A., Legeard, B.: A taxonomy of model-based testing approaches. *Softw. Test. Verification Reliab.* **22**(5), 297–312 (2012). <http://dx.doi.org/10.1002/stvr.456>
26. Völter, M., Stahl, T., Bettin, J., Haase, A., Helsen, S.: *Model-Driven Software Development: Technology, Engineering, Management*. Wiley, Chichester (2013)
27. Wittevrongel, J., Maurer, F.: SCENTOR: scenario-based testing of e-business applications. In: *Proceedings of the Tenth IEEE International Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises, 2001, WET ICE 2001*, pp. 41–46 (2001)
28. Zhang, F., Qi, Z., Guan, H., Liu, X., Yang, M., Zhang, Z.: FiLM: a runtime monitoring tool for distributed systems. In: *Third IEEE International Conference on Secure Software Integration and Reliability Improvement, 2009, SSIRI 2009*, pp. 40–46, July 2009