

KompicsTesting - Unit Testing Event Streams

IFEANYI W. UBAH, LARS KROLL, ALEXANDRU A. ORMENIŞAN, SEIF HARIDI
 Department of Software and Computer Systems, School of ICT
 KTH Royal Institute of Technology, Stockholm, Sweden

Abstract

In this paper we present *KompicsTesting*, a framework for unit testing components in the *Kompics* component model. Components in *Kompics* are event-driven entities which communicate asynchronously solely by message passing. Similar to actors in the actor model, they do not share their internal state in message-passing, making them less prone to errors, compared to other models of concurrency using shared state. However, they are neither immune to simpler logical and specification errors nor errors such as dataraces that stem from nondeterminism. As a result, there exists a need for tools that enable rapid and iterative development and testing of message passing components in general, in a manner similar to the *xUnit* frameworks for functions and modular segments code. These frameworks work in an imperative manner, ill suited for testing message-passing components given that the behavior of such components are encoded in the streams of messages that they send and receive. In this work, we present a theoretical framework for describing and verifying the behavior of message-passing components, independent of the model and framework implementation, in a manner similar to describing a stream of characters using regular expressions. We show how this approach can be used to perform both black box and white box testing of components and illustrate its feasibility through the design and implementation a prototype based on this approach, *KompicsTesting*.

Contents

1	Introduction	2
1.1	The <i>Kompics</i> Component Model	3
1.2	Formal Languages and Grammars	5
1.3	Regular Grammars and Expressions	6
1.4	Finite Automata	7
1.5	Context-free Grammars	9
2	Unit Testing in Message Passing Systems	9
2.1	A Test Specification for Executions	11
2.2	Regular Operations on Execution Sets	11
2.3	Blocks and Executions	13
2.4	Ordering an Execution	15
2.5	Specifying Constraints and Requirements on Executions	17
2.6	Request-Response Patterns	20
2.7	White Box Testing and Dependencies	21
3	<i>KompicsTesting</i> - A Framework Implementation	23
3.1	The Alphabet of Executions	23

2	<i>I. W. Ubah et al.</i>	
3.2	The Proxy Component	24
3.3	Scheduling	25
3.4	A DSL for Building Specifications	25
3.5	Creating and Executing a Test Case	26
3.6	Building an Automaton	27
3.7	Expecting Events	27
3.8	Request-Response	28
3.9	Creating Blocks	29
3.10	Conditional Statements	29
3.11	Interactive Testing	29
3.12	Ambiguous Test Specifications	30
3.13	Accepting Sequences and Termination	31
4	Related Work	32
5	Conclusions	33
	References	34

1 Introduction

Companies spend a large amount of resources in a quest to detect and fix errors in their software systems. Such errors are usually introduced into the systems either at development time, due to mistakes made by programmers, or during system requirements specification, and may untimely present themselves, resulting in failures with possibly severe consequences. The chances of introducing errors rapidly increase with the complexity of the system and as more complex systems are developed to take prominent roles in our lives, it becomes increasingly important that as many errors as possible are found and fixed as early as possible in the development process.

Message-passing systems are inherently complex and this complexity is amplified when processes work independent of each other in a distributed setting, due to the inherent asynchrony and possibility of partial failure where subsets of processes may fail at any time during execution.

One way of finding and detecting errors in a system is by testing, that is the execution of a system with the sole purpose of finding errors. This can be performed on various levels depending on how much of, and at what granularity the system is to be exercised. At the lowest level is unit testing (IEEE, 2010), focusing on the verification of the smallest testable entities in the system, usually a single function, class or module in object oriented programming languages, or an actor, process, or component in a message-passing system.

However, when unit testing an entity in a message passing system (which we refer to such an entity simply as a component from hereon) however, the task of specifying meaningful tests become convoluted if at all feasible. Most tools available for unit testing only allow the tester to describe tests in a low-level manner, for example, using mechanisms for sequentially making assertions on the return values of functions. While this way of testing may prove effective for the intended languages or frameworks, they are not sufficient for verifying the behavior of a component as this can not be done properly without examining the sequences of messages sent and received by the component.

Our approach enables the tester to describe the expected behavior of the component as a stream or sequence of messages and their directions. For example, a tester may deem the behavior of a component under test (CUT) to be correct if the component responds with a message m_2 on receiving a message m_1 and subsequently responds with one or more messages m_4 on receiving m_3 . This could be specified as $(in(m_1), out(m_2), in(m_3), out(m_4), out(m_4)^*)$ where in and out denote incoming and outgoing messages from the perspective of the CUT and $*$ is the Kleene closure operator denoting zero or more occurrences of the message $out(m_4)$. Such a specification of sequences form a language. If we assign each message and direction pair the values e_1, e_2, e_3, e_4 respectively and call each such pair an event, then the language of the specification is thus described by the regular expression $e_1e_2e_3e_4^*$ and the alphabet of the described language are the event symbols $\{e_1, e_2, e_3, e_4\}$.

Such a specification can subsequently be converted to an automaton by a framework implementation and used to verify the CUT by running an instance and verifying that the events observed conform to the specification. Thus, if the automaton ends up in an accepting state when all events have been observed, the test case passes.

This paper presents a theoretical framework for unit testing message-passing systems using this approach as well as KompicsTesting, an implementation of the framework for the Kompics component model. Section 1.1 presents the Kompics component model while sections 1.2 to 1.5 give a brief background on formal languages and finite automata needed to follow the paper. We then describe the aforementioned theoretical framework in section 2, followed by a description of the design and implementation of KompicsTesting in section 3. This is followed by a discussion of related work (section 4) and concluding remarks (section 5).

1.1 The Kompics Component Model

Kompics is a component based, message-passing model for building distributed systems. Components in Kompics are event-driven entities that communicate by exchanging messages, in the form of events, with each other. Events are simply data-carrying objects in the system. Components provide communication interfaces via bidirectional ports and are connected to each other via channels binding any two ports. The following sections describe the key primitives and concepts used in the Kompics model; more details can be found in (Arad, 2013).

1.1.1 Ports

Ports in Kompics embody the interface between a component and its environment. They are *bidirectional* entities through which events are sent to and received from a component. Through ports, Kompics provides a type system for the events exchanged in the system. Unlike systems like Akka (Wyatt, 2013) and Erlang (Armstrong, 2003) where any message may be addressed to an actor or component, ports define which events may go in and out of a component.

Each port has two directions which we label *positive* and *negative*, as well as a *port type* which declares a specific set of event types that are allowed to pass through it in each

direction. We denote a port $\alpha = (p_\alpha, n_\alpha)$ where p_α is the positive direction and n_α is its negative direction, and say that a port α allows an event e in some direction d_α if its port type declares the type of e in direction d_α .

As communication interfaces, we can think of a port $\alpha = (p_\alpha, n_\alpha)$ as a service interface and associate requests with its negative side n_α and responses with its positive side p_α . Thus a component that produces service α , declares α so that requests events are received, incoming from n_α while responses events are outgoing from p_α . We also say that the component declaring α in this manner *provides* the port. Conversely, a component that consumes service α , declares α so that responses are incoming from p_α while requests are outgoing from n_α . We say that the component *requires* the port α .

We again, label the two sides of a port from the perspective of a particular component declaring the port by saying that the side of the port emitting incoming events to the component is the *inside* port while the side emitting outgoing events is the *outside* port. Note that for a component that provides the port α , the inside and outside ports are n_α and p_α respectively while for the component that requires α , they are conversely p_α and n_α respectively

Finally, we say a component *triggers* an event on a port, if it sends an event through that port — this is typically done from the inside, going out, except when interfacing component code with traditional thread-based libraries.

1.1.2 Channels

Channels create connections between two components via their declared ports. They can be thought of as bidirectional communication pipes that carry events from one port to another. Connections are only possible for any two ports of the same port type that have opposite directions, that is a channel can be created between two ports α and β by connecting p_α to n_β or n_α to p_β . Events are forwarded through channels in first in, first out (FIFO) order, that is they are delivered to the destination port in the order that they were triggered at the source port. An event triggered on a port is broadcast on all channels connected to that port, thus Kompics does not provide a mechanism for addressing events to specific components in the system. On arrival of events at their destination ports, they are queued up on that port until the component that declared the port is scheduled to execute that particular event on that port.

1.1.3 Event Handlers

An event handler, or *handler*, in Kompics is a user-defined function of a component. A handler accepts events of a particular type and any of its subtypes (in the strongly typed programming language sense). Handlers are registered on ports and as functions are executed whenever the component receives an event on an event that is accepted by any registered handlers. We say that a registered handler for a given port is *subscribed* to the port.

Kompics guarantees the sequential execution of the handlers of a single component, preventing the need for internal state synchronization between handlers. Handlers belonging to different components on the other hand, may be scheduled and executed concurrently.

1.1.4 Components

Components in Kompics are reactive entities which communicate asynchronously with each other by exchanging messages. Similar to actor based systems, a component has some internal state associated with it as well as an event queue at its declared ports. It also has a set of event handlers as described, *subscribed* on its declared ports, which are executed whenever some accepted event is received on that port. Handlers are the means by which a component updates its internal state.

A component can be encapsulated within another component, using parent-child relationships that form a component heirarchy, with a single *Main* component that is initially started at runtime. The relationship between a component and its child components enables a flexible architecture for managing system complexity as well as the delegation of configuration of a component to the parent. For example, on creating and starting a component, the sub-components or child components are recursively created and started. It is the job of such a parent component to bootstrap its child components by say, setting up their communication channels.

1.2 Formal Languages and Grammars

A formal language (Mateescu & Salomaa, 1997) is a set of sequences of symbols formed together according to some specified rule. A language has an alphabet, which is the set of allowed symbols that may be used to form valid sequences of that language. As an example, the set $\{ab, ac\}$ is a language containing exactly two sequences, the alphabet is the set of *characters* $\{a, b, c\}$.

It is rarely feasible to describe a language as we just did by listing out all sequences that are members of that language, nor is it particularly desirable to verbally describe all possible languages. Instead, one may simply give a *formation rule* that describes the intended language. This rule may then be used to either generate sequences that belongs to the described language or check that a given sequence belongs to the language. Such a rule may be a grammar (Chomsky, 1956) or an automaton (Hopcroft & Ullman, 1979) and several of these exist for a variety of classes of languages. We discuss finite automata as well as regular and context-free grammars.

We start by defining the Kleene closure, or closure, on a set of symbols. Let Σ be a finite set of symbols. The closure Σ^* on Σ is the set of sequences that can formed taking any number of symbols in Σ , with repetitions allowed. For example if $\Sigma = \{a, b\}$, the closure $\Sigma^* = \{a, b, aa, ab, aba, \dots\}$. The closure on any set always includes the empty symbol ϵ since one may choose to take zero number of symbols.

A grammar G is formally described as a four-tuple (N, Σ, P, S) where N is a set of *nonterminal* symbols, Σ is the alphabet containing *terminal* symbols, P a set of *productions* and $S \in N$ a start symbol (Hopcroft & Ullman, 1979). A production $p \in P$ is of the form $(\Sigma \cup N)^* N (\Sigma \cup N)^* \rightarrow (\Sigma \cup N)^*$ where $*$ is the closure operator — hence the right hand side may be empty. Thus a production may have at least one nonterminal symbol and any number of terminal and nonterminal symbols on the left-hand side of the arrow while the right-hand side may contain any number of terminals and nonterminals.

1	$S \rightarrow ABc$
2	$S \rightarrow Bcd$
3	$A \rightarrow a$
4	$B \rightarrow b$

Listing 1. A Simple Grammar.

If the grammar G generates exactly the set of sequences that make up some language L , then we say that the grammar describes L and write $L(G) = L$. To generate a sequence in L using G , we start with the start symbol S and expand using the right-hand side of any production of S , replacing the expanded symbol with its right-hand side. We repeatedly do this until no nonterminal symbols are left in the expanded sequence.

As an example, consider the simple grammar G whose productions are shown in listing 1. Nonterminal symbols are denoted with capital letters while terminal symbols are denoted by small letters. So $N = \{A, B\}$ and alphabet $\Sigma = \{a, b, c, d\}$, S is the start symbol. We begin with the start symbol and expand the sequence using the first production. Thus we replace S with the sequence ABc . Next, we replace the next nonterminal A with its right-hand side a leaving the new sequence abc . Finally, we do the same for B to retrieve the final sequence containing no terminals abc . Note that we could have started by expanding S using the second production which would have resulted in the string bcd .

1.3 Regular Grammars and Expressions

Regular grammars restrict the types of productions they contain. Here, the left-hand side must contain exactly one nonterminal symbol while the right-hand side may contain a single terminal symbol followed by a nonterminal symbol, or a nonterminal symbol followed by a terminal symbol. The right-hand side may also be empty. A language described by a regular grammar is called a *regular* language. We can forego the use of regular grammars in place of an equivalent *regular expressions* (Kleene, 1951) notation for describing regular languages in a more concise manner.

A single regular expression describes a language that is regular. The basis for regular expressions are the symbols of the alphabet Σ , such that if $a \in \Sigma$ then $L(a) = \{a\}$, that is the language described by a terminal symbol of an alphabet consists of the single symbol. Regular expressions can be built from smaller regular expressions using the following operations on that are closed on regular languages (Hopcroft & Ullman, 1979).

Union The union of two languages L and M denoted $L \cup M$ is the language containing all sequences that are either in L or in M or both. For example, if $L = \{ab\}$ and $M = \{ab, cd\}$, then the union $L \cup M$ is the language $\{ab, cd\}$. Hence if R and S are regular expressions, then $R|S$ is also a regular expression describing the language $L(R) \cup L(S)$.

Concatenation The concatenation of two language L and M denoted LM is the language containing all sequences that are formed by taking a sequence in L and appending a sequence in M to it. Using the previous example, the concatenation of both languages $LM = \{abcd\}$. Hence if R and S are regular expressions, then RS is also a regular expression describing the language $L(R)L(S)$.




Fig. 1. A DFA recognizing the sequence “cat”.

Kleene Closure The closure on a language L , denoted L^* is the language containing all sequences formed by taking any number of sequences, with repetitions allowed, from L and concatenating them. Using the previous example, the closure L^* on L is the language $\{ab, abab, ababab, \dots\}$. Hence if R is a regular expression R^* is also a regular expression.

1.4 Finite Automata

A finite state automaton (from hereon automaton) is an abstract machine made up of states and transitions which are labelled with an input symbol and connect any two states (Hopcroft & Ullman, 1979). An automaton can be in any such state at a given time, changing states by following transitions and consuming the input associated with that transition. An automaton can either be a *deterministic* finite automaton (DFA) or a *nondeterministic* finite automaton (NFA). A DFA has exactly one destination state from any given state and input symbol while an NFA may have several possible destination states. As a result, an NFA may be in several states at once.

1.4.1 Deterministic Finite Automata

Formally, a DFA M is defined as the five-tuple $(Q, \Sigma, \delta, q_0, F)$ where Q is a finite set of states, Σ is the set of input symbols, δ is a transition function that takes an input symbol $i \in \Sigma$ and a state $q_i \in Q$ and maps them to a new state $q_n \in Q$, q_0 is a start state which is the initial state of the automaton, $F \subseteq Q$ is a set of accepting states (Hopcroft & Ullman, 1979).

Instead of explicitly listing out the mappings of the transition function δ of a finite automaton, we can visually describe it using a state diagram as shown in figure 1. Here, $Q = \{q_0, q_1, q_2, q_3\}$, $\Sigma = \{a, c, t\}$, $F = \{q_3\}$. Transitions are arrows labelled by an input symbol. The start state q_0 has no incoming transitions while final states, here q_3 , are depicted by a double-bordered circle.

A DFA can be used to recognize a sequence of symbols. Given a sequence, it can be simulated such that at every given state, it consumes the next symbol and follows a transition labeled with that symbol to the next state. When the sequence is completely consumed, if the DFA is in an accepting state, then we say that it *recognizes* or *accepts* the sequence otherwise we say that it *rejects* the sequence. The language $L(M)$ of a DFA M is the set of all sequences recognized by M , we also say that M describes $L(M)$.

As an example, listing 1 shows a DFA that recognizes the sequence “cat”. It starts at state q_0 and consumes the first symbol ‘c’, from q_0 it follows the transition labeled with this symbol to q_1 , then it consumes the next symbol ‘a’ and so on until no more symbols are left. Since it ends up in state q_3 , we say that it recognizes the input sequence “cat”.

Note that we have not specified transitions for every other symbol at each state, and the DFA must have exactly one valid transition for each state and input pair. We only show those transitions that lead to an accepting state and imply that all other non-specified transitions lead to an implicit error state such that the automaton rejects the sequence.

Finally we note that, like for regular grammars and regular expressions, the language $L(M)$ of a DFA M is also a regular language (Hopcroft & Ullman, 1979). Thus if a language can be described by a regular grammar, then it can also be described by a DFA. In fact, every DFA can be converted into an equivalent regular expression and vice-versa.

1.4.2 Nondeterministic Finite Automata

As mentioned previously, a NFA can be in several states at the same time since it may have multiple transitions for a given input symbol and state pair. The automaton is said to guess its next state by following all possible transitions.

Formally, an NFA M is defined as the five-tuple $(Q, \Sigma, \delta, q_0, F)$, where Q is a finite set of states, Σ is a finite set of input symbols, $q_0 \in Q$ is the start state, $F \subseteq Q$ is a set of accepting states and δ is a transition function that takes a state $q_i \in Q$ and an input symbol $i \in \Sigma$ and returns a *subset* of Q (Hopcroft & Ullman, 1979). Thus, the transition function δ alone marks the only difference between an NFA and a DFA. NFAs describe the same class of languages as DFAs, in fact, any NFA can be converted to an equivalent DFA using the subset construction (Rabin & Scott, 1959). Conversely, note that every DFA is an NFA such that the transition function always returns a singleton set.

Finally, an NFA with epsilon transitions, or ϵ -NFA is an extended NFA with the single additional ability to allow a transition without consuming any input string. Thus an ϵ -NFA can have unlabeled transitions, called ϵ -transitions, allowing it to make spontaneous a transition to the next state. In a state diagram, such labels are marked with the special ϵ symbol only as a visual convenience, the symbol does not belong to the input alphabet of the automaton. ϵ -NFAs do not extend the class of languages defined by NFAs or DFAs or regular grammars, they all define regular languages. In fact, any given ϵ -NFA can be converted to an equivalent *NFA*. However, they offer more illustrative and programming convenience and will be used later in this paper.

1.4.3 Concatenating Finite Automata

Since the described finite automata all define the same class of language as regular grammar and expressions do, the union, concatenation and closure operations can also be performed on any such automaton. Here we describe the construction of an automaton MN from the concatenation of any two automata M and N as presented in (Hopcroft & Ullman, 1979). This operation is implied throughout this paper and as such is presented here for reference. A more specific description for the union and closure operations on automata is illustrated in sections 2.2.2 and 2.3.2

To construct an automaton $MN = (Q, \Sigma, \delta, q_0, F)$ from two automata M and N where $M = (Q_M, \Sigma_M, \delta_M, q_{0M}, F_M)$ and $N = (Q_N, \Sigma_N, \delta_N, q_{0N}, F_N)$, We set $Q = Q_M \cup Q_N$, next we set the start state of the first automaton q_{0M} as the start state q_0 of MN and the accepting states F_N of the second automaton as the accepting states F of MN . The transition function δ uses the mappings from δ_M and δ_N with additional ϵ -transitions from each final state $q \in F_M$ in the first automaton to the start state q_{0N} of the second automaton.

The idea here is that the first part of recognized sequence is delegated to the automaton M and once this part is recognized, an ϵ -transition takes the automaton MN to the start




Fig. 2. A DFA recognizing the sequence “ac”.




Fig. 3. A DFA recognizing the sequence “dc”.

state q_{0N} of N , where the second part of the sequence takes the automaton to an accepting state.

An example of this is shown in figure 4, as an automaton constructed from the automata in figure 2 and 3.

1.5 Context-free Grammars

Context-free grammars (Autebert *et al.*, 1997) (CFG), place less restrictions on their production forms than regular grammars, the left-hand side of every production must contain a single nonterminal symbol. They can thus be used to describe a larger class of languages than regular grammars. A language described by a CFG is called a context-free language. A typical application of CFGs is to describe programming languages, as those form a subset of the context-free languages. We will also make use of this convenience, when describing domain specific languages (DSLs) later.

2 Unit Testing in Message Passing Systems

Consider that the basic entity for building message-passing systems, usually, is an actor or a component based on the particular systems. Regardless of the particular system involved, each such message-passing entity, which we refer to generically as a component, can be considered as a state machine. Its execution can be described by the performance of a sequence of atomic events, each of which serve as a transition label in the state machine. These events can be categorized as input, output and internal events (Lynch & Tuttle, 1989). *Internal events* such as the scheduling and execution of event handlers occur internally within the component while input and output events, further labelled as *external actions*, correspond to the receipt of incoming events and emission of outgoing events respectively by the component. External events for a component occur at the interface between the component and its environment.

In accordance with the definition of unit testing (IEEE, 2010), we would like to assert assumptions against the behavior of some component under test (CUT) and its interaction with the environment — the goal being to increase our confidence in the component’s implementation. In the context of unit testing message-passing components, we adopt black-



Fig. 4. Concatenating automata in figures 2 with 3 recognizing the sequence “acdc”.

box testing (IEEE, 2010) as testing that ignores the internal implementation of the CUT and uses general knowledge of input and output events to verify the behavior of the CUT and white-box testing (IEEE, 2010) as testing that uses internal implementation knowledge of the CUT to verify its functionality, for example, by running assertions on its internal state. To thus verify the behavior of the CUT against a desired protocol, the tester must provide some test specification containing a description of the testable expected behavior, while the test framework, on the other hand, takes on the responsibility of executing an instance of the CUT and verifying its execution against the provided specification. This raises the question of what the contents of such a specification should be and how it assists the test framework in the verification process at runtime.

2.0.1 Executions

We define an *execution sequence* (or *execution* if this is clear from the context) of a component c to be the sequence of external events performed by a single instance of c when run in a given environment. Such an environment comprises c , along with other peer components as well as all messages sent between these peers. In other words, an execution may be considered as a linear arrangement of input and output actions, similar to a string of characters in formal languages, that models the behavior of a component.

2.0.2 Black-Box Testing as Well-formedness of Executions

The act of black box-testing a component becomes equivalent to verifying the correspondence of actual external events of the CUT at runtime against the expected behavior of the CUT. This means that a tester describes the set of sequences, that are deemed to be correct for a given test case, as the specification, while the framework matches the actual occurrence of events, obtained by running an instance of the CUT, against this specification. This problem becomes equivalent to the decidability of the well-formedness of a word in the field of formal languages — that is, whether or not a given word (sequence of symbols) belongs to a language. A task very well suited for finite automata in the case of regular languages as discussed in section 1.2.

Given that an execution is composed of a finite sequence of events, the set of all such events forms an alphabet, over which a language can be described. The tester thus specifies a language while the word to be tested for well-formedness is the execution sequence observed by the test framework at runtime by listening on events incoming to and outgoing from the component. The idea of black box-unit testing thus can be summarized as the act of writing a specification describing a language L of all acceptable executions for that particular test case, while the test framework verifies the runtime execution of the CUT against L .

2.0.3 Formal Grammar for Specifying Executions

Our approach uses a mechanism for specifying execution sequences as regular languages, allowing the utilization of regular expression techniques for writing tests. As a result, a test specification may be converted into a finite automaton by the testing framework

<pre> 1 Body → expect Event⁺ 2 Event → <i>e</i> </pre>
--

Listing 2. Concatenation of executions

such that the test passes if the automaton ends in an accepting state. This allows for the runtime execution of the CUT while observed events are being matched by the automaton simultaneously. An added advantage of this is that it enables the tester to use likely familiar concepts from regular expressions when writing tests.

2.1 A Test Specification for Executions

With the aim of producing a framework for performing automated and repeatable testing, we desire the test specification to be unambiguous as well as machine executable. For this reason, a test specification can be considered as a program, written in some language, not to be confused with the language of execution sequences it describes. The test framework implementation then acts as a machine executing this program, responsible for verifying the correspondence of the CUT to the specification.

An advantage of using a specialized language is that it enables the specification of instructions for additionally controlling the behavior of the test framework, resulting in more concise and flexible test specifications. For brevity, we assume that environment setup activities such as the actual creation of CUT instances and other participating components and their dependencies are handled prior to program execution.

The following sections introduce the theoretical framework using syntactic constructs that generate a test specification. The formation rule used for illustration is a context-free grammar that describes a block-structured language for specifications. The CFG allows us to succinctly describe most of the constructs, independent of any particular message-passing system. We show how these constructs can be mapped to automata and where applicable, how they correspond to their regular expression variant. The discussions also include how they can be used to write tests for the vast majority of desired behaviors in both deterministic and non-deterministic environments.

2.2 Regular Operations on Execution Sets

In this section, we illustrate the variants of regular expression operations offered by our framework for describing events comprising execution sequences and building larger executions from smaller ones. In particular, we illustrate the union and concatenation of the languages described by a specification while the Kleene closure operation is explained in Section 2.3.2 when the concept of blocks have been introduced.

2.2.1 Concatenation of Execution Sets

Listing 2 shows the production for concatenating a sequence of single events using the **expect** keyword. The symbol ⁺ means one or more occurrences. Statements that appear on the right hand side of the *Body* non-terminal describe a unique language over the set of



Fig. 5. An automaton recognizing a single event e .

Fig. 6. Automaton for recognizing the concatenation of events.

Body \rightarrow **either** $Body^+$ **or** $Body^+$ **end**

Listing 3. Union of executions

event alphabet, that is, they describe a unique set of execution sequences. As with the symbols of regular expressions, a single event e matches itself and the language $L(e)$ described by e consists only of itself — that is $L(e) = \{e\}$. In terms of the automaton created, at runtime, the statement “**expect** e ” would cause the automaton to move from the start to next and final state only when e has occurred — thus matching e . Any other event observed at this state leads to a failed test case. Figure 5 shows a deterministic finite automaton DFA M_e recognizing a single event e .

For concatenation of sets of executions, we also use the definition of concatenation of regular languages as defined in 1.3. The *expect* keyword allows us to specify a list of execution sets to be concatenated. As an example, the statement **expect** $e_1 e_2 e_3$ describes the language $L = \{e_1 e_2 e_3\}$ formed by concatenating the three languages $L(e_1) = \{e_1\}$, $L(e_2) = \{e_2\}$ and $L(e_3) = \{e_3\}$.

Generally, given a statement $S = \text{“expect } e_1 e_2 e_3 \dots e_n\text{”}$, where each a_i describes a unique language $L(a_i)$, a DFA M_S recognizing $L(S)$ is created by concatenating each automaton M_{a_i} that recognizes $L(a_i)$ sequentially. Figure 6 shows a DFA M_S constructed in this manner.

2.2.2 Union of Executions Sets

Listing 3 extends the grammar in Listing 2 with constructs for creating the union of execution sets via the **either-or** conditional statement. As each *Body* nonterminal describes a unique language, the conditional statement contains two independent languages from the *either* and *or* branches which are combined via the union operation on regular languages to form the described language of the conditional statement.

Conditional statements allow the tester to describe possible paths within the state space of the component, exactly one of which would be traversed depending on the observed event at runtime. These are convenient in situations where there are several, alternative and possibly equivalent paths outgoing from a certain state of the CUT. In some cases, it may be inconvenient or difficult to reproduce a test environment that consistently guides the framework through any one of the alternate paths, while in other cases, the paths may be supplied to increase the robustness of the test case — the test case may be designed so that a random path is chosen at each iteration.




Fig. 7. NFA for conditional statement: “**either expect $e_1 e_2$ or expect $e_1 e_3$ end**”.

```
Exec → repeat n? Hdr* body Body* end
```

Listing 4. Production for generating an execution.

With the *Body* nonterminal, any statements are allowed inside the *either* and *or* branch of a conditional statement, including other conditional statements.

Given a conditional statement S with statements A and B as specified by its *either* and *or* branches respectively, the language $L(S)$ described by S is defined to be the union of the sets described by both branches, that is $L(S) = L(A) \cup L(B)$. Consequently, an automaton M_S recognizing $L(S)$ would be equivalent to the automaton $M_{A \cup B}$ that recognizes $L(A) \cup L(B)$.

We construct an NFA for this purpose by combining sub-automata M_A and M_B , for $L(A)$ and $L(B)$ respectively, alongside each other as described in (Hopcroft & Ullman, 1979). All states and transitions of the sub-automata remain throughout the construction. A new start state $q_{A \cup B}$ of M_S is created by merging the start states q_A and q_B of M_A and M_B respectively. This new state will have the same outgoing transitions as the combined start states, allowing the NFA to reach the final states of either sub-automaton when verifying events at runtime. The set of final states of the NFA is the union of the final states of both sub-automata since the NFA accepts any execution that is accepted by either sub-automaton.

As an example, consider the statement $S =$ “**either expect $e_1 e_2$ or expect $e_1 e_3$ end**”. Figure 7 shows the constructed NFA M_S for this statement. The start state $q_{A \cup B}$ contains the same transitions outgoing from the start states q_{A0} and q_{B0} of the sub-automata while the final states q_{A2} and q_{B2} of the NFA are those of the sub-automata. The constructed NFA can be subsequently be converted to a DFA or simulated at runtime using a set of states that keep track of the possible states of the automaton.

2.3 Blocks and Executions

As mentioned, our illustrative specification language is a block-structured language. This means that the language allows for the creation of blocks as well as nested blocks, enabling the tester to group a sub-sequence of events into a single unit.

A block is the topmost construct of a test specification. As shown in Listing 4, it generates an execution which may be made up of smaller sub-executions using nested blocks. Here, n is a positive number while $?$ means that it is optional. Every specified event belongs to a single block. As with a lot of block-structured programming languages, the benefits

of this are manifold — a scope can be invoked throughout a block, allowing declarations of constraints and requirements on events that only affect the block’s sub-sequence at test runtime. For example, a tester may want to specify that some particular event must be present within a block’s sequence. For example, consider the execution $e_1e_2e_3e_4$, the tester may declare that the sub-sequence e_2e_3 delimit a nested block so that e_1 and e_4 continue to be associated with the outer block. Now within the newly declared block, the tester may want to declare a requirement that some event e_0 must occur at some point within the block — hence the actual described sequence is $e_1e_0e_2e_3e_4|e_1e_2e_0e_3e_4|e_1e_2e_3e_0e_4$. This is further explained in Section 2.5.2

The production in Listing 4 splits a block into a *header* and a *body* identified by the Hdr and Body nonterminals. The Hdr non-terminal generates constraints and requirements on the events observed within the entire block as well as nested blocks while the Body as mentioned generates executions belonging to the block.

The language of an entire block is the set of execution sequences described by applying the specified constraints and requirements declared in the header, to the execution sequences described by the body. The entire block is declared within keywords *repeat* and *end* with an optional positive integer n . This is used to generate a block that consecutively matches a sub-sequence at runtime, the specified number of times or in the case without any number, a block that matches the Kleene closure on the defined language.

2.3.1 Matching Repeated Executions

The block construct, when declared with a positive integer n , instructs the framework at runtime to match an sub-sequence against its described execution sequence n times. This mechanism facilitates concise test specifications, allowing the tester specify a repeating sequence of events only once. For example, if the tester expects an execution of the form “ $abcabcabc \dots$ ” where a, b, c are events, the sequence may only be declared once as “**repeat** n **body expect** $a \ b \ c$ ” where n is the number of expected occurrences of the subsequence “ abc ”.

Suppose that a DFA M_s recognizes the language of a block s without a *repeat* operation invoked. What we know about the actual language S described by the block, with the *repeat* operation invoked, is that it is a language formed by concatenating s with itself n times. An execution in S is formed by taking any n executions in s and concatenating them. As a result, a DFA M_S recognizing $L(S)$ can be constructed using n copies of the automaton M_s by concatenating them together.

Consider as an example, the statement $S = \text{“repeat } 2 \text{ body expect } e_1 \ e_2 \ \text{end”}$. $L(S) = \{e_1e_2e_1e_2\}$. The language described by the block B before the *repeat* operation is invoked is $L(B) = \{e_1e_2\}$. However, the sequence is expected twice, resulting in the final language described by the automaton $M_B M_B$ where M_B is an automaton recognizing $L(B)$.

2.3.2 Matching the Kleene Closure of an Execution Set

The *repeat* statement, when declared without a positive integer, instructs the framework to match the Kleene closure on the described language. The framework matches zero or more occurrences of observed executions that belong to the described language. Analogous to the




Fig. 8. ε -NFA for “**repeat body expect $e_1 e_2$ end**”.

Kleene closure on a language, we denote such a statement S^* and define its language $L(S^*)$ as the set of executions that can be formed by concatenating any number of executions described by S , the associated block. For example, given a block B describing the set of executions $L(B) = \{e_1, e_2\}$, $L(B^*)$ describes the set consisting of all executions containing only the actions e_1 and e_2 — i.e. $\{e_1, e_2, e_1 e_2, e_2 e_2, \dots\}$. Since the closure of a language matches zero or more occurrences, the language $L(S^*)$ always includes an empty execution (containing no events) regardless of the execution described by S .

The addition of the closure operation into our test specification necessarily introduces nondeterminism when verifying the occurrence of events at runtime using an automaton. Consider the language $L(S^*)$ of some block S . An equivalent automaton M_{S^*} at runtime, that matches some initial execution $E_0 \in L(S)$ must transition to a next state that implies the current state of the CUT, without having access to future events which might be yet to occur. In such a case, the automaton must correctly guess between two options — a transition to a final state signalling that it is done matching executions in $L(S)$, or a transition to a next state that expects to match another execution $E_1 \in L(S)$.

In accordance to the construction of a finite automaton for the Kleene closure of a language (Hopcroft & Ullman, 1979), we build a nondeterministic finite automaton with epsilon transitions (ε -NFA) as M_{S^*} to recognize the closure on an execution set described by some block S . To do this, we start with the automaton M_S that recognizes $L(S)$ and transform M_S into M_{S^*} by introducing two types of ε -transitions corresponding to the choices to be made by the automaton. We form an ε -transition from the start state q_0 of M_S to every final state of M_S allowing the automaton to go directly to the final state when it guesses that all executions have been verified. We also form ε -transitions from each final state q_n back to q_0 , allowing any number of execution sequences described by S to be verified by the automaton.

As an example, figure 8 shows an ε -NFA M_{S^*} recognizing the language described by the specification **repeat body expect $e_1 e_2$ end**. Here, the framework transitions directly from the q_0 to q_2 if no actions are observed while the the path from q_0 to q_2 via q_1 is traversed n times where n is the number of consecutive $e_1 e_2$ executions that occur at test runtime.

Such an automaton could be directly converted to a DFA using the subset construction or simulated directly by keeping a set of states that the automaton could possibly be in as facilitated by the eclosure mechanism (Hopcroft & Ullman, 1979).

2.4 Ordering an Execution

Accurately predicting the order of events that make up an execution becomes problematic in message-passing systems with multiple components, due to the inherent asynchrony

```
Event → e | unordered e+ end
```

Listing 5. Generating an unordered execution.

involved. This is especially problematic in asynchronous distributed environments where there are no upper bounds on computation and message transmission time (Raynal & Singhal, 1996; Lamport, 1978). However, this problem becomes manageable in the scope of unit testing since the only matched events are local to a single component in the system. Inevitably, the class of verifiable scenarios by the framework are limited to those of local properties of the CUT. For example, scenarios verifying global properties of algorithms, which likely involve asserting properties across several components, are not specifiable. Nonetheless, it sufficiently serves the purpose of unit testing.

In the scope of unit testing of a single component in a message passing system, we consider inherent problems of nondeterminism such as the unpredictable scheduling of components, lack of upper bounds on computation steps and message delays as it pertains to the events observed locally at the CUT’s interface. For example, a set M of outgoing requests from the CUT to a set of peer components may expect all incoming responses in a set N . The order in which the requests are sent or responses arrive at the CUT may not be accurately predictable when writing the tests. Hence a need to explicitly specify a sequence of unordered actions for nondeterministic events in an execution.

Listing 5 generates an unordered sequence of events which are specified between the *unordered* and a matching *end* keyword. Since the order of the events do not matter, the language described by the statement consists of all permutations of the originally specified sequence.

As an example, the following statement “**expect** $e_1 e_2$ **unordered** $e_3 e_4$ **end** e_5 ” instructs the framework to fail the test case at runtime on observing, say, the execution “ $e_2 e_1 e_3 e_4 e_5$ ” but not on executions “ $e_1 e_2 e_3 e_4 e_5$ ” and “ $e_1 e_2 e_4 e_3 e_5$ ”.

2.4.1 Constructing Automata for Unordered Executions

Consider the statement $S = \text{“expect unordered } e_1 e_2 \text{ end”}$ describing the language $L(S) = \{e_1 e_2, e_2 e_1\}$. Figure 9 shows an automaton M_S recognizing an equivalent language. Generally, the statement “**unordered** $e_1 e_2 \dots e_n$ **end**” for n actions e_1 to e_n describes a language containing $n!$ executions — each, a permutation of the originally specified sequence. As a result, a DFA M describing an equivalent language must accept exactly all $n!$ possible sequences.

One way to build M is to construct a path for each possible sequence from the start to the end states of the automaton. Each path contains exactly n transitions and $n + 1$ states and each state within a path is used to remember which events have been matched and which are pending. Such a state can be thought of as being associated with a bit string of length n representing the set of specified events e_1 to e_n such that the i th bit is 1 if the event e_i has been matched in the execution and 0 otherwise.

Consider the start state q_0 of M . No events have been matched at this state so all bits are set to 0 - that is, the bit string associated with this state is “000...0”. Now suppose at runtime that the event e_n occurs first. Then the automaton transitions to the next state q_δ




Fig. 9. DFA for “**unordered** $e_1 e_2$ **end**”; key $\{e_1, e_2\}$.




Fig. 10. EFSM for unordered executions.

$Hdr \rightarrow \mathbf{allow} e^+ \mid \mathbf{disallow} e^+ \mid \mathbf{drop} e^+ \mid \mathbf{blockExpect} e^+$
--



Listing 6. Generating restrictions on observed executions

associated with the bit string “000...1” with only the n th bit set to 1. A transition from q_δ on action e_1 leads to the next state $q_\gamma = “100\dots 1”$, and so on with the final state of the automaton $q_\phi = “111\dots 1”$ signalling that all specified events have been matched.

This mechanism generates a total of $n!$ states where n is the number of specified events, making it impractical for even modest values of n . In implementation however, the need for extra states is easily avoided since unlike with DFAs, the task of remembering the matched actions can be accomplished programmatically by the framework. Figure 10 shows a more practical scheme using an EFSM - an extended finite state machine (Alagar & Periyasamy, 2011). The self-transition labelled α represents the set of events specified by the *unordered* statement. The automaton uses exactly 2 states regardless of the size of n and enables the transition from the start to the end state only when the trigger condition, that all bits in the string have been set to 1, has been met.

2.5 Specifying Constraints and Requirements on Executions

Having a programmable test specification allows the possibility of placing flexible restrictions on the events observed within the execution. Listing 6 shows the right hand of the *header* nonterminal Hdr that is used to generate such restrictions in a test specification. The semantics of these constructs assume that, unlike events specified using the *expect* keyword, it is not possible to predict an exact instance in the test execution where the constraints must be satisfied. As a result, they apply to an entire block and any nested blocks in the specification. For this reason, they can only be specified within the *header* of a block.

Fig. 11. Automaton recognizing “**repeat 1 body expect $e_1 e_2$ end**”.Fig. 12. Extending Figure 11 with headers “**allow $e_3 e_4$ drop e_5** ”.

2.5.1 Constraints on Blocks

Within a block, some events may be *disallowed* by the tester. The presence of such events at any point when the framework executes a statement in the block is undesirable so that it is unnecessary to continue running the test case - the test case should terminate immediately with a failure. In some other cases, the occurrence of certain events in an execution is not necessary to validate the test case. In fact, such events may not even be observed in multiple executions of the same test case. In other words, these events are not required for a successful test case but are *allowed* if they occur. Finally, the messages associated with some events may be *dropped* on occurrence. This means that such messages, if outgoing from the CUT, should not be forwarded to recipients and if incoming, should not be delivered to the CUT. Such a mechanism is particularly useful when writing test logic that verify edge cases and error conditions by enabling the tester guide the CUT into the vulnerable state. However, it does require knowledge of events that should or should not be handled by the CUT in order to enter such a state be available to the tester.

As is typical of block-structured languages, constraints are only valid within the scope of the block where it was declared as well as its nested blocks. However, a constraint C_1 on an event e in block B_1 can be *shadowed* by redeclaring a new constraint C_2 on e in nested block B_2 of B_1 . This means that C_2 is valid if e is observed within B_2 and C_1 is valid if e is observed within B_1 but not B_2 . Conflicts occur when constraints on the same events are declared within the same block. In such cases, only the last constraint declared is enforced.

As much as these statements within block headers describe the execution sequence, they also control the behavior of the framework. For example the *disallow* statement may be interpreted as an instruction to the framework to fail the testcase if any of the specified events appear within the sequence while *allow* and *drop* instruct the implementation whether or not to forward the specified events if they do appear in the sequence. In other words, from the perspective of building an automaton that recognizes an execution, actions specified by these constructs do not necessarily drive the automaton closer to a final state.

Consider the specification $S = \text{“repeat 1 body expect } e_1 e_2 \text{ end”}$, an equivalent automaton M_S shown in Figure 11 and the specification $C = \text{“repeat 1 allow } e_3 e_4 \text{ drop } e_5 \text{ body expect } e_1 e_2 \text{ end”}$ having an equivalent block body to S . To build the automaton M_C shown in Figure 12 recognizing $L(C)$, we incorporate *allow* and *drop* statements by adding a self transition on the set of specified events, to each state of the automaton for the entire block body including nested blocks (q_0, q_1). In the case of M_C , this implies that states recognizing events e_1 and e_2 become annotated

a self transition on events (e_3, e_4, e_5) as specified by the block header. No distinction is made between *allow* and *drop* statements in the automaton, the actual difference is in the behavior of the framework implementation (whether or not it forwards the event).

Since events declared with the *disallow* keyword causes the testcase to fail on occurrence, an equivalent event from the automaton’s perspective would be to include a transition leading to an error state on the specified actions — such transitions are implicit as mentioned in section 1.4.1. Just as with *allow* and *disallow* statements, such a transition would be included on each state within the block.

2.5.2 Requirements on Blocks

Events preceded by the *blockExpect* keyword as shown in Listing 6 are required the occur at runtime at some point within the block. This places a requirement on the framework to successfully exit the block not only when all statements have been executed but also that all such required events have been verified. For example, given the specification “**repeat blockExpect** e_0 **body expect** $e_1 e_2$ **end**. The test framework does not exit successfully until events e_1 and e_2 have been matched in order but also the event e_0 has been observed as well. Hence, the actual language of this specification is $\{e_0e_1e_2, e_1e_0e_2, e_1e_2e_0\}$ as e_0 may occur at any position in the sequence while the relative positions of all other events are respected. Since such events are not specified in any order relative to other events within the block, this construct is able to describe a wider range of non-deterministic scenarios than *unordered* statements.

To incorporate the *blockExpect* statement into an automaton, note that any sequence described by this statement must be interleaved with the sequences described by the rest of the block. Since the statement itself describes its own language (all permutations of the specified execution), we can distinguish this language from that described by the rest of the block.

We illustrate the construction process with an example using the specification $S = \text{“repeat 1 blockExpect } e_3 e_4 \text{ body expect } e_1 e_2 \text{ end”}$. The block header declaration $H = \text{“blockExpect } e_3 e_4 \text{”}$ describes the language $L(H) = \{e_3e_4, e_4e_3\}$, while the body B describes the language $L(B) = \{e_1e_2\}$. Let h be the number of events specified by H and b the number of events specified by B . In this case, $h = b = 2$. The semantics of the *blockExpect* produces the language $L(S) = \{s \mid s \text{ is a permutation of “} e_1e_2e_3e_4 \text{”} \wedge e_1 \text{ appears before } e_2\}$. An automaton M describing $L(S)$ must thus recognize only sequences containing exactly the events e_1 through e_4 such that the order imposed by the entire block is preserved.

We start with the automaton M_B recognizing $L(B)$. Each state in M_B remembers the next expected event as specified in the block body B . To incorporate H , note that for each state in M_B , M must remember what actions from H have already occurred. This can be done using a bit string of length h where the i^{th} bit represents whether or not the i^{th} event specified by H has occurred. Thus each state in M is annotated with a bit string, as well as the next expected action in B , leading to 2^h states for each event e in B as each state for e remembers a unique possible configuration of the bit string.

To define the transition function δ for the automaton M , consider each state as a pair (e, s) where e is the next expected event in B and s is the bit string associated with that




Fig. 13. DFA for “**repeat 1 blockExpect e_3 body expect $e_1 e_2$ end**”; key {nextaction in body, bitstring of blockExpect}.




Fig. 14. EFSM for Figure 13.

state. Valid transitions from a state q are only available for an unverified event e_n such that $e_n \in I_B$ or $e_n \in I_H$ - where I_B and I_H are the sets of events declared by the body B and header H respectively. For a given state $q = \{e_i, s\}$, If $e_n \in I_B$, then the next state is the state $q_\gamma = \{e_{i+1}, s\}$ else if $e_n \in I_H$ such that e_n equals the k^{th} event specified by H , then the next state is $q_\gamma = \{e_i, s_\gamma\}$ where s_γ is a copy of s with the k^{th} bit set to 1. The final state $q_\phi = \{\emptyset, s_\phi\}$ has all bits in s_ϕ set to 1 and e empty since no events are expected.

Figure 13 shows an automaton constructed using this technique for the specification “**repeat 1 blockExpect e_3 body expect $e_1 e_2$ end**” This technique creates a $(b + 1)2^h$ states for the constructed automaton. Just as with the automaton constructed in Section 2.4, an implementation similar to an extended finite state machine can accomplish the task using $b + 2$ states. Such an automaton uses $b + 1$ states to represent the body of the specification with the logic for remembering previously verified events in H delegated to the framework as opposed to the automaton itself. As a result, we include an additional *sink* state before the final state that transitions to the final state only when all actions in H have been verified. Figure 14 shows such an EFSM with the transition from the sink state q_2 to the final state q_3 enabled only when all required events specified by the *blockExpect* have been verified, in this case e_3 .

2.6 Request-Response Patterns

In some cases a tester may expect a set of matching incoming messages (responses) in response to a set of outgoing messages (requests). A variation of which can be accomplished using a specification similar to “**expect unordered $rq_1 rq_2 \dots rq_n rs_1 rs_2 \dots rs_n$ end**” where rs_i is the matching response for request rq_i . The requests and responses are expected in any order, accounting for the non-determinism in response order. In typical request-response scenarios, there exists a way of uniquely identifying the response to a request, for example

using a unique id i shared by both rq_i and rs_i . However, the example specification assumes knowledge of this id and this may not always be predictable at test specification time. It is not unusual for such ids to be randomly generated by the CUT or some other artefact in the system. Thus we need a mechanism for matching responses for requests with randomly generated ids.

We note that the language described by this problem is not regular as it requires some context in order to determine whether or not a response matches some previous request. Since we do not have a unique identity of the requests at specification time to begin with, we need some way to set a checkpoint when the automaton starts expecting requests and subsequent responses. The tester describes some property of the requests, like the class name in for example Java, as well as a mapping function $\alpha : E \times E \rightarrow \mathbb{B}$ where $\mathbb{B} = \{true, false\}$ for that request, that takes a request and a response events and returns true if the provided request matches the response. Such a function may simply compare the ids of the messages associated with both events. The framework implementation on reaching the defined checkpoint, keeps a set of observed requests and for each response received, uses the matching function to determine if it matches some requests, advancing the automaton as needed.

2.7 White Box Testing and Dependencies

2.7.1 Testing with external dependencies

Normally, the CUT communicates with other peer components at runtime in its deployed environment in order to complete its computations. These peers thus become dependencies when testing and must be present or simulated in order to correctly exercise the CUT. There are several reasons for decoupling the CUT from its dependencies when unit testing. Tests themselves becomes less complex to write when the logic does not additionally need to coordinate peer components in order to reproduce the necessary conditions for the test. Also, testing with actual dependencies may also cause computations of other components to cross the boundaries of the unit test, for example, when a peer happens to be a component located on a different network and network messages are not desirable in the tests. More important is the certainty that the source of a test failure does not originate external to the CUT — a failure in some peer component may cause the test case to inaccurately verify the CUT.

When such situations are undesirable, a solution is usually to replace these dependencies with mocks (Mackinnon *et al.*, 2000) — deterministic and usually less complex versions of a dependency offering the minimal features that facilitate the test execution. This requires no new constructs with respect to our proposed framework as no distinction is made between a mock and a real version of a component in the system. Both parties are simply components interacting with the CUT, allowing the tester to decide to mock on a per-component basis. Note that it may be simpler to forego the mock and use the actual dependency in some cases after all.

The second option offered by the framework is to send events interactively into the test environment. This is done using the the *trigger* construct within the body of a block as shown in Listing 7. It can thus be used to mimic some peer component's participation in

```
Body → trigger  $e^+$ 
```

Listing 7. Interactively producing events

```
Body → inspect  $\alpha$ 
```

Listing 8. Inspecting Internal State

the test scenario by sending messages to the CUT instance at runtime as if they were sent by the peer itself. This option produces more flexible but potentially verbose test specifications diffused with operations. A *trigger* statement prefixes a list of events, instructing the test framework to send the messages associated with those events to the CUT. In fact, this operation needs not be limited to CUT events. Specified events may be destined for any peer component in the test environment as long as the framework is able to provide some mechanism for identifying the recipient. Thus the framework is free to include any additional parameters associated with the events used in a *trigger* statement. Consequently, unlike the other operations, we associate the *trigger* of a single event as the act of sending a message to any component in the system. Events are triggered in their specified order.

In terms of implementation, we can annotate our automaton with special *active* states that represent *trigger* operations. These states do not transition on any events but simply instruct the framework to perform an action (send a message) on entry and immediately transition to the next state. As an example, the statement **expect** e_1 **trigger** e_2 e_3 **expect** e_4 causes the test framework to send the messages associated with events e_2 and e_3 in between matching events e_1 and e_4 .

2.7.2 White-box Testing

In addition to black-box testing by verifying external events using the constructs explained so far, Listing 8 shows the production for generating an *inspect* statement. The statement allows the tester to request an assertion over the internal state of the CUT at any point within the test execution. This keyword prefixes a predicate function α that takes as an argument the CUT and returns a boolean depending on whether the internal assertions run against the CUT succeeded or not. On executing an *inspect* statement, the framework pauses the execution of the test case and passes the instance of the CUT to the supplied predicate and resumes the test case when the predicate returns. In the case of a unsuccessful inspection, where the returned value is false, the test case fails.

The idea of an inspection is that the tester puts the CUT into some known state by reproducing a set of conditions in the test environment and subsequently verifying the effects of internal events if any, performed by the CUT. Producing and optionally verifying the required conditions to inspect the component internals is done using the same constructs for black-box testing such as *trigger* and *expect*.

A requirement is placed on the framework to offer some mechanism for performing inspections with the ability to pause the CUT's execution at runtime while its internal state is being inspected. In view of components being scheduled nondeterministically, the framework must also enforce the happened-before relation (Lamport, 1978) between all

preceding actions, both internal and external, and the inspect operation. As the components of message-passing systems are usually associated with a queue, messages to be handled by the CUT as a result of previous events may still be queued at the time of inspection. Inspecting the CUT under such an incomplete environment results in an inaccurate verification of the CUT. The framework must thus provide a transparent mechanism for ensuring that handlers for all events that happen-before the inspection at runtime have been executed and completed by the CUT before performing the inspection.

The *inspect* statement can be incorporated into an automaton using the same technique as described in Section 2.7.1 where an active state is associated with each *inspect* statement, instructing the framework to perform the inspection when it reaches that state.

3 KompicsTesting - A Framework Implementation

In this section we describe an implementation of the presented theoretical framework on Kompics. KompicsTesting is implemented on the reference implementation of Kompics (Arad, 2008) which is written in Java. However, there are other implementations of Kompics in Python and Scala which are not considered in this paper. Thus we refer to the reference implementation simply as Kompics in this section. We start with a description of the event symbols used to describe executions, followed by the architectural and design decisions that enable the implementation verify the described executions and enforce the operational semantics of the framework. Then we illustrate the DSL of the framework used for describing these executions according to the context-free grammar presented in the theoretical framework.

3.1 The Alphabet of Executions

As the main goal of KompicsTesting is to enable the tester describe a set of executions as a language, this section highlights the alphabet of the language from which executions are formed. The theoretical framework described assumes the notion of an event and declares the set of all events to be the alphabet of the language. Here, we make this idea concrete for our implementation.

Let E be the set containing all possible events observable in the environment of a test case in question, P the set of ports provided or required by all peer components in the environment including the CUT, and a set of literals $D = \{in, out\}$ denoting the direction of events relative to the CUT. As mentioned in Section 1.1, Kompics provides a form of type system for events. We thus define our alphabet Σ , the set of valid symbols for constructing executions, as the relation on the set $E \times P \times D$ consisting of triples (e, p, d) where $e \in E, p \in P, d \in D$ and the port type of p allows events of type e in direction d . For example $(e, p, in) \in \Sigma$ if and only if the CUT provides port p and the port type of p allows e in the negative direction, or it requires p and the port type of p allows e in the positive direction.

The direction literals belonging to D are Java enums while event and port literals are provided as references to the test framework implementation.

3.2 The Proxy Component

The theoretical framework assumes that events incoming to and outgoing from the CUT are not only observable by an implementation, but in operations like allow and drop, intercepted by the framework, since the CUT or peer components may be required not to receive such events.

KompicsTesting uses a special component as the *Main* component in the test case called the *proxy*, responsible for intercepting and observing events to and from the CUT. This is enabled by Kompics' component hierarchy. As the first component to be created and bootstrapped by the Kompics runtime, the proxy acts as a composite component, in charge of creating the CUT and other peers as children components.

Internal Port Implementation in Kompics: As bidirectional interfaces, a port has both a negative and a positive side. In the Kompics reference implementation, a port declaration (provided or required) α by a component is internally represented by a pair of port references (p, n) where p is the positive port and n is the negative port representing the two sides of α . Both ports p and n are associated with an *owner* which is either the component C that declares (provides or requires) the port or the parent component P of C . On receiving an event on a port, the port's owner is scheduled for execution of its handlers. For a port $\alpha = (p, n)$ provided by a component C , n is the inside port and owned by C while p is the outside port owned by parent P . Conversely, for a port $\alpha = (p, n)$ required by a C , p is the inside port and owned by C while n is the outside port and owned by the parent P .

Establishing a parent-child relationship between the proxy and all other components enables a sufficient degree of control over the events that go into and out of the child components. Every peer component's port declaration assigns the outside port to the proxy. Consequently, the proxy can subscribe handlers to these ports and listen for events on these ports. This, coupled with the fact that creating channels between any ports in the test environment must be done through the test framework, enables the proxy to completely isolate the CUT from other components, intercept events at its interfaces and only forward those which have not been explicitly disabled in the specification.

Consider a provided port $\alpha = (p_\alpha, n_\alpha)$ by the CUT and a required port of the same type $\beta = (p_\beta, n_\beta)$ by some peer component. Say we connect both outside ports via a channel c from p_α to n_β as would be done in a normal environment. Handlers subscribed by the CUT for α are done on n_α and execute when an incoming event is received while outgoing events triggered by the CUT must go through p_α and c before reaching the peer component. If we subscribe intercepting handlers h_i to n_α and h_o to p_α , we are able to observe both incoming and outgoing events respectively on port α . However, since the handlers get executed at the same instance as those subscribed by the CUT, we are not able to optionally forward events as required by the framework, the CUT would have already received the event by the time our handlers execute.

The solution is to connect the proxy component directly to any peers that would normally be connected to the component while the CUT has no direct connections to any other components. Due to Kompics' channel broadcast mechanism, peers are unaware of the source of received events and recipients of triggered events. Thus by mirroring ports declared by

the CUT on the proxy component, peers may be directly connected to the proxy while the CUT remains isolated. In this case, a mirror port $\delta = (p_\delta, n_\delta)$ of port α would be created with c connected from p_δ to n_β . Incoming events to are thus intercepted by subscribing h_i to n_δ while h_o remains subscribed to p_α . Handler h_o keeps track of the connected channels and components and manually forwards outgoing events when required.

3.3 Scheduling

Scheduling decisions must also be made in order to enforce the semantics of the framework. So far we have described the process of intercepting and optionally forwarding events but not how to determine if an intercepted event should be forwarded. KompicsTesting runs the automaton on a single control thread while other components except for the proxy are scheduled on a thread pool. The proxy component uses a calling thread scheduler, executing its handlers immediately on the same thread of the component triggering the event on one of its ports. This enables exactly those events that are going in and out of the CUT to be placed in an event queue as soon as they are triggered on any of the proxy's ports. Since these components are scheduled independent from the automaton, the automaton can progress at its own pace and consume events from the queue when it needs to transition to some next state. For example, a trigger state may add a new event to the queue while an expect state consumes an event from it. Only when an event is consumed from a queue can it be *handled*, that is forwarded to the recipient(s), as determined by the current state's transitions.

3.4 A DSL for Building Specifications

The simplicity of the presented context-free grammar allows us to implement our test specification scheme using a simple Java domain-specific modelling Language DSL, based on the builder pattern (Gamma *et al.*, 1993). It allows the tester to describe languages over our alphabet of events by calling API methods. The DSL uses a similar formation rule to the CFG, mapping keywords to API methods which can then be consumed by the tester as if they were generating a specification using the CFG. An advantage of implementing a DSL directly in Java allows the tester to supply terminals of the CFG such as events, ports and predicate functions directly to the framework as references.

Generally, the API methods are named after an equivalent keyword in the CFG and where appropriate and take as arguments a representation of the literals, if any, as defined in the CFG. For example, the *inspect* statement requiring a predicate literal is implemented as a method **inspect**(Function<Component, Boolean>); while the *end* keyword is implemented as a method declared without any parameters.

Listing 10 shows how the specification in Listing 9 would be written in KompicsTesting. To avoid long-winding method calls, statements such as *trigger* and *expect* that prefix a sequence of literals are implemented as methods that take a single literal. As a result, specifying multiple such literals in a sequence are done by calling the methods multiple times in succession, as shown using multiple calls to the *expect* and *allow* methods. The *trigger* method only takes an event and a port reference but not direction. Since ports in

```

1 repeat 10
2   allow e3 e4
3 body
4   expect e1 e2
5   trigger e5
6 end

```

Listing 9. Generating a specification from context-free grammar.

```

1 testcontext.body()
2   .repeat(10)
3     .allow(e_3, p_3, in)
4     .allow(e_4, p_4, in)
5   .body()
6     .expect(e_1, p_1, in)
7     .expect(e_2, p_2, out)
8     .trigger(e_5, p_5)
9   .end();

```

Listing 10. Converting Listing 9 to a DSL.

```

1 TestContext<CUT> tc = TestContext.create(CUT.class);
2 Component cut = tc.getComponentUnderTest();
3 Component peer = tc.create(PeerComponent.class);
4 tc.connect(cut.getPositive(Message.class),
5           peer.getNegative(Message.class));
6 //...
7 tc.check();

```

Listing 11. Creating and Executing a Test Case

Kompics are implemented as bidirectional port references, the direction is inferred from the supplied port reference.

3.5 Creating and Executing a Test Case

In KompicsTesting, the first step of testing a component is by creating an instance of the *TestContext* class by providing a component definition of the CUT to the test framework. Listing 11 illustrates the process of creating and running a test case. Line 1 creates a test context for the component under test having the component definition class *CUT*. Line 2 retrieves a reference to the created CUT instance while Line 3 creates a peer component whose required port of type *Message* is subsequently connected with the provided port of the CUT. Once all such setup activities have been carried out, methods for describing the execution of the CUT may be called (Line 6) and finally the **boolean** `check()` method is called, running the test case and returning true if the observed execution is a member of the language described by the test case (the created automaton ends up in an accepting state) and false otherwise.

The component definition is used by the framework to create an instance of the CUT for test execution, a reference to which is accessible to the tester throughout the specification.

```
1 tc.  
2   body()  
3     .expect(e_0, p_0, in)  
4     .expect(e_1, p_1, out)  
5     .unordered()  
6       .expect(e_2, p_2, in)  
7       .expect(e_3, p_3, out)  
8     .end()
```

Listing 12. Expecting Events

3.6 Building an Automaton

From the perspective of the tester, the DSL simply provides a means for describing the expected execution sequence. Once the *check* method is called, the test framework takes control of the program. It builds up an automaton using the techniques presented for each method called and concatenates them accordingly. The final automaton is an ϵ -NFA, subsequently simulated by tracking a set of states that the automaton may possibly be in at anytime during test execution. Simulation of the constructed automaton is done in parallel with the execution of the CUT. This enables the CUT to produce and handle events while the test framework matches these events and advances the automaton when required.

3.7 Expecting Events

As mentioned in section 3.1, event symbols are of the form (e, p, d) where e is an event reference, p is a port reference and d is a direction enum. Consequently, this triple must be specified to methods that add events to the described sequence. Listing 12 shows an example of expecting both ordered and unordered events. Line 3 and 4 expects an incoming e_0 and an outgoing e_1 event in that specified order, creating two states with these events as transitions. Line 5 through 8 expects both incoming e_2 and outgoing e_3 in any order, creating a single state, 2.4, while the framework transitions to the final state only when both events have occurred.

3.7.1 Matching Events

As KompicsTesting is written in Java, we would like to make sure that events are correctly matched according to the intentions of the tester, that is, we need some mechanism for defining the equality of events. For example, in line 3 of listing 12, we say that we expect to match the event $E_0 = (e_0, p_0, in)$ and having observed some event $E_i = (e_i, p_i, d)$ at runtime, would like to verify that E_0 is in fact E_i . Since directions and ports are internal to Kompics and the test framework, we can define our notion of equality transparently, albeit sensibly to the user. These are simply compared using the `equals` method of `Object`. In other words, we need some user-defined notion of equality of the event references specified in executions. Instead of forcing the tester to implement equals methods for each class with an instance that appears as an expected event, a comparator may be registered with the framework during test specification, for any class that may appear as an event, representing the users definition of equality for objects of that class.

```

1 tc
2   .setComparator(E.class, comparatorE)
3   .body()
4     .expect(e_0, p_0, in)
5     .expect(E.class, predicate, p_i, in)

```

Listing 13. Matching Events with Comparators and Predicates

Listing 13 shows how this is accomplished via the `setComparator` method. The method can appear only within the initial header (the implicitly defined block header which will be explained in 3.9) of a test specification. Say an event (e_i, p_0, in) is received at the state with the transition labelled with the event specified in line 3, and that both e_0 and the e_i belong to some class C which is assignable to the class E of the registered comparator in line 2, then this comparator is used to check for equality. If no such comparator is registered for a class of events, then the framework defaults to using the `equals` method of the specified event. Thus testers have the option of either providing a comparator for events or defining the equality of the event classes themselves.

Line 5 uses an explicit matching function for a class of events. The function provided is a predicate, in this case for class E , that returns true if the event argument specified is matched and false otherwise. This method provides allows the tester reuse the same predicate for describing a wider range of events instead of explicitly providing equivalent references of the expected events. Consider the case of observing an event (e_i, p_i, in) at some state q with the transition described in Line 5, and the class of e_i is assignable to E , the specified class of the predicate. Then the predicate is used to match e_i . Note that the behavior is not defined for a state with multiple transitions having predicates with the same signature as labels. Since such predicates could potentially match events of the same class so there is no way of correctly selecting an appropriate one. We thus declare such cases as a symptom of an ambiguous test specification and leave their behavior undefined.

3.8 Request-Response

KompicsTesting implements a variation of the more general request-response pattern described in section 2.6. Here the tester provides a function $\alpha : E \rightarrow E$ mapping a given request to a response, which is subsequently triggered on some specified port. Listing 14 illustrates this pattern. Line 2 marks the starting point of the pattern while line 5 marks the end of the pattern. The `expect` methods take in the class of the expected request, the outgoing port of the request (note that the direction is not specified since this is implicit as an outgoing request message), a second port where the provided response should be triggered and finally the mapping function that provides such a response.

The matching function is called with a received event and returns a response if the event is matched otherwise null. Internally, the entire scheme is implemented as a single state that transitions to the next state only when all requests have been matched and the corresponding responses have been triggered.

```

1 tc.body()
2   .expectWithMapper()
3     .expect(R1.class, p_1, p_2, m1)
4     .expect(R2.class, p_2, p_1, m2)
5   .end()

```

Listing 14. Request Response in KompicsTesting

```

1 tc.body()
2   .repeat(3)
3   .body()
4     .trigger(e_0, p_0)
5   .end()
6   .repeat()
7   .body()
8     .expect(e_1, p_1, d_1)
9   .end()

```

Listing 15. Creating Blocks in KompicsTesting

3.9 Creating Blocks

Blocks declared within the *repeat* and *end* keywords are represented by their method variants. Listing 15 shows the pattern for creating blocks within KompicsTesting. As every event must belong to a block, an entire test case is always run in a block that is repeated exactly once. This block is not explicitly created by the tester, in the sense that the corresponding *repeat* and *end* methods are not called. The *body* method however, must still be called before using non-header methods as shown in line 1. Two blocks are explicitly created by the tester; the first spans Line 2 through 5 and triggers some e_0 on the specified port 3 times, while the second block creates a Kleene closure block that matches as many incoming e_1 events as possible on port p_1 . Blocks with a fixed number of iterations are implemented using a loop counter that is decremented for each traversal of the state graph representing the block. The graph is exited when the counter reaches zero. This avoids the need to create a large number of states for large sizes of n .

3.10 Conditional Statements

Listing 16 shows the pattern for creating a conditional statement in KompicsTesting. In this example, lines 3 and 4 comprise the *either* branch and expects two events while lines 6 through 10 comprise the *or* branch of the statement, expecting any number of the event sequence described in the body of the repeat block. Similar to the Kleene closure block equivalent, this statement creates a NFA. In this example, the start state contains a transitions on $E = (e_1, p_1, d_1)$ to both the states representing statements in lines 4 and 9, from where on the automaton diverges.

3.11 Interactive Testing

Listing 17 shows a demonstration of the *trigger* and *inspect* methods for a specification. This example interleaves the *trigger* with the *inspect* methods that assert any changes to the component state made due to the previous *trigger*. Here the reference supplied as an

```

1 tc.body()
2   .either()
3     .expect(e_1, p_1, d_1)
4     .expect(e_2, p_2, d_2)
5   .or()
6     .repeat()
7     .body()
8       .expect(e_3, p_3, d_3)
9       .expect(e_4, p_4, d_4)
10    .end()
11  .end()

```

Listing 16. Conditional Statements

```

1 tc.body()
2   .trigger(e_0, p_0)
3   .inspect(predicate)
4   .trigger(e_1, p_1)
5   .inspect(predicate)

```

Listing 17. Interactive Tests

argument to the `inspect` method is an instance of a `Function<C, Boolean>` where `<C extends ComponentDefinition>` represents the type of the component definition of the CUT.

As the *trigger* operation is used to simulate components sending events across channels, a triggered event in `KompicsTesting` is simply published on any connected channels by triggering on the given port.

The *inspect* operation requires that all previous events that happened-before the operation must have been handled by the CUT before the inspection can commence. To accomplish this, on entry to the `inspect` state, `KompicsTesting` verifies that the work count, the amount of pending events and events currently being handled, is zero otherwise it waits until this becomes true, signalling that all events have been handled by the CUT.

3.12 Ambiguous Test Specifications

The introduction of interactive operations like `trigger` and `inspect`, while useful tools for testing, come with caveats when writing test cases with inherent nondeterminism. Consider a Kleene closure that does not describe any sequence as shown in 18 and a conditional statement with both branches beginning with interactive operations as shown in 19. In the first case there is no way of knowing how many times the `trigger` operation should be performed since the block does not define any clear entry or exit conditions. The second case does not provide a way to accurately choose between the *either* and *or* branches since they both do not have entry conditions. Now consider listing 20 without a clear entry condition for the *either* branch but an entry condition for the *or* branch. From the start state, the automaton may consult the event queue and only if it does not consume the event (e_1, p_1, d_1) successfully does it traverse the *either* branch.

The testers intent can not be inferred by the framework from specifications such as those in listings 18 and 19 and in fact, the intent of the first case is questionable to begin

```

1 tc.body()
2   .repeat()
3     .trigger(e_0, p_0)
4   .body()
5   .end()

```

Listing 18. A Kleene block that does not describe any execution

```

1 tc.body()
2   .either()
3     .trigger(e_0, p_0)
4   .or()
5     .trigger(e_1, p_1)
6   .end()

```

Listing 19. A Conditional statement with no clear entry conditions

```

1 tc.body()
2   .either()
3     .trigger(e_0, p_0)
4   .or()
5     .expect(e_1, p_1, d_1)
6   .end()

```

Listing 20. A Conditional statement with an entry condition for the *or* branch

with. Such specifications are deemed ambiguous and the behavior of the framework is left undefined. Listing 20 on the other hand is not ambiguous.

3.13 Accepting Sequences and Termination

A test case is successful when the current set of states contains a final state and no other events are received as that would take the automaton a next state. Since the processing of events and simulation of the automaton is done at real-time with the execution of components, it may sometimes become problematic to predict when no new events are expected or when the test case should terminate. The unpredictable scheduling and asynchronous behavior of components are the reasons for this.

Listing 21 illustrates both cases. The tester initially expects to receive any number of the event specified at line 4. However the automaton may prematurely consult an empty event queue and either immediately fail with an end-of-stream error, or try to wait, potentially indefinitely, for the event. Lines 3 through 6 on the other hand expect any number of the event specified on line 5 for successful termination. Here, the automaton may choose to transition to an accepting state if the event queue is empty but this behavior would not be correct if there are some events still in transit. To resolve this, KompicsTesting includes a default timeout, configurable by the tester that specifies how long the framework should wait for an event to occur. The responsibility of enforcing a correct verification of the component in such cases is thus placed on the tester.

```
1 tc.body()
2   .expect(e_0, p_0, d_0)
3   .repeat()
4   .body()
5     .expect(e_1, p_1, d_1)
6   .end()
```

Listing 21. A test without timeout for events

4 Related Work

This paper presents an approach for verifying the correctness of message passing systems. Another strategy for accomplishing the same goal is through the use of model checking, a formal method used to verify the conformance of a given model of a system to its intended specification, containing properties which the system must satisfy. Model checking methods use state-space exploration (Clarke & Schlingloff, 2001), and can instill more confidence into a system implementation by automatically enumerating and exhaustively exploring the state-space of the system in order to find errors. If all paths in the state-space have successfully been executed, then the system is said to be correct. However, generating the state-space of a system can be an expensive process and in the presence of concurrency as is typical of message passing systems, it is not uncommon for a exponentially larger or infinite number of states to be required in order to perform properly verify the system due to the number of possible interleavings in the initial state space, a problem known as state-space explosion. Techniques such as partial-order reduction (Maiya *et al.*, 2016), (Bokor *et al.*, 2011) and dynamic partial-order reduction exist to (Lauterburg *et al.*, 2010) try to mitigate this problem by analysing opportunities for reducing the number of explored states but do not always succeed managing the problem.

Previously, model checking methods have been applied to testing of message passing systems. Systems like Basset (Lauterburg *et al.*, 2009), Jacco (Zakeriyan *et al.*, 2015) and McErlang (Fredlund & Svensson, 2007) have been developed for actor systems. Basset provides a model checker for actor systems that compile to Java bytecode, using Java PathFinder JPF (Visser *et al.*, 2003), an environment for performing model checking, verification and analysis of Java programs. In order to efficiently explore a Java actor system, it avoids exploring low-level and library code by concentrating on message scheduling related errors by providing a transparent interface to the actor application being tested while replacing the underlying library code with equivalent code that is easier to explore. It uses dynamic partial-order reduction as an optimization technique to prune the number of explored states. However, its message scheduling assumptions leads to false negatives in generating explorable states leading to large state spaces of programs. Its reliance on the JPF platform also leads to inefficient test execution time. Jacco, improves on the Basset system by implementing a new scheduling approach and technique for reducing the state space of the actor system in order to reduce the time taken for model checking. Coupled with an architecture that abandons the reliance on the JPF platform, it shows better performance in both execution time and generated state space of programs compared to Basset. McErlang is a model checker for Erlang programs. Being developed in Erlang itself, it facilitates rapid prototyping and testing of Erlang programs and supports a substantial part of the programming language including the distributed parts. It allows the

user to encode correctness properties of the system as automata specified in Erlang while the model checker checks the system against the specified properties. However, it uses state-space exploration for systematic checking of the system against the specification and suffers from the state-explosion problem especially in the face of nondeterminism.

Our approach presents a platform independent framework and foregoes the use state-space generation and exploratory techniques, depending solely on the tester to specify the particular state-space to be exercised. Thus it does not run into problems like state-explosion that can make verification infeasible. However, the possibility of offering full test coverage of the component depends solely on the tester. Our approach also focuses solely on verifying a single component in the system as opposed to system-wide verification as in the mentioned tools. As a consequence, a tester is unable to verify global properties of the system using a tool based on our approach.

The Akka TestKit tool (Akka, n.d.) on the other hand provides a platform for performing unit and integration testing on actor systems based on the Akka framework at various level of granularity. In terms of unit testing, it allows the tester perform synchronous tests on specific units of code without involving the actor model, in addition to performing unit tests on a higher level of abstraction using single actors as units. In the latter case, a tester is able to send messages to an actor under test and subsequently verify responses in a synchronous manner. In between message exchanges, the tester is able to access the actors internal state in order to perform white-box testing similarly to our approach. It also allows the user test that an incoming sequences of events is processed correctly, even in the face of nondeterminism, leading to reordering of events. However, such tests can only be performed using multiple actors to generate the stream of events. Unlike our approach that allows an interactive mechanism for generating streams of events. Since it does not use an approach of describing the entire sequence, verifying outgoing events is done manually by the tester, for example setting up a recipient component that asserts the outgoing events.

5 Conclusions

In this paper we presented a language based approach to unit testing message passing systems as well as a prototype implementation of this approach for the reference implementation of the Kompics component framework. In the first part of the paper, we presented the rationale for treating sequences of events incoming to, and outgoing from a messaging component as a language and how such a language could be used to efficiently verify the behavior of the component using blackbox and white box based testing techniques. This was done by presenting a theoretical framework for testing components by describing its expected sequence of events and interactively introducing events into this sequence in order to produce stimuli during test execution. We showed how a tester may describe deterministic and nondeterministic sequences as well as how an implementing framework may cope with the feasibility of verifying nondeterministic sequences.

The second part of the paper presented KompicsTesting as a prototype implementation of the theoretical framework. We discussed the design choices and internal implementation of the prototype and illustrated the provided DSL for describing sequences. We showed how the DSL may be used to specify a wide variety of sequences modelling the behavior of components in Kompics.

References

- Akka. *Testing actor systems*.
- Alagar, V. S., & Periyasamy, K. (2011). *Extended finite state machine*. London: Springer London. Pages 105–128.
- Arad, Ci Cosmin Ionel. (2013). *Programming Model and Protocols for Reconfigurable Distributed Systems*. Ph.D. thesis, KTH - Royal Institute of Technology, Stockholm.
- Arad, Cosmin. (2008). *Kompics project*.
- Armstrong, Joe. (2003). Making reliable distributed systems in the presence of software errors. 295.
- Autebert, Jean-Michel, Berstel, Jean, & Boasson, Luc. (1997). Context-free languages and push-down automata. *Pages 111–174 of: Handbook of formal languages*. Springer.
- Bokor, Peter, Kinder, Johannes, Serafini, Marco, & Suri, Neeraj. (2011). Supporting domain-specific state space reductions through local partial-order reduction. *Pages 113–122 of: Proceedings of the 2011 26th IEEE/ACM international conference on automated software engineering*. ASE '11. Washington, DC, USA: IEEE Computer Society.
- Chomsky, Noam. (1956). Three models for the description of language. *Ire transactions on information theory*, **2**(3), 113–124.
- Clarke, Edmund M., & Schlingloff, Bernd-Holger. (2001). *Handbook of automated reasoning*. Amsterdam, The Netherlands, The Netherlands: Elsevier Science Publishers B. V.
- Fredlund, Lars-Åke, & Svensson, Hans. (2007). Mcerlang: a model checker for a distributed functional programming language. *Pages 125–136 of: Acm sigplan notices*, vol. 42. ACM.
- Gamma, Erich, Helm, Richard, Johnson, Ralph, & Vlissides, John. (1993). *Design patterns: Abstraction and reuse of object-oriented design*. Berlin, Heidelberg: Springer Berlin Heidelberg. Pages 406–431.
- Hopcroft, John E., & Ullman, Jeffrey D. (1979). *Introduction to automata theory, languages, and computation*. Addison-Wesley Publishing Company, Inc.
- IEEE. (2010). *ISO/IEC/IEEE 24765: Systems and Software Engineering - Vocabulary*. Tech. rept.
- Kleene, Stephen Cole. (1951). *Representation of events in nerve nets and finite automata*. Tech. rept. DTIC Document.
- Lamport, Leslie. (1978). Time, clocks, and the ordering of events in a distributed system. *Commun. acm*, **21**(7), 558–565.
- Lauterburg, Steven, Dotta, Mirco, Marinov, Darko, & Agha, Gul. (2009). A framework for state-space exploration of java-based actor programs. *Pages 468–479 of: Proceedings of the 2009 IEEE/ACM international conference on automated software engineering*. IEEE Computer Society.
- Lauterburg, Steven, Karmani, Rajesh K., Marinov, Darko, & Agha, Gul. (2010). Evaluating ordering heuristics for dynamic partial-order reduction techniques. *Pages 308–322 of: Proceedings of the 13th international conference on fundamental approaches to software engineering*. FASE'10. Berlin, Heidelberg: Springer-Verlag.
- Lynch, Nancy A., & Tuttle, Mark R. (1989). An introduction to input/output automata. *Cwi quarterly*, **2**, 219–246.
- Mackinnon, Tim, Freeman, Steve, & Craig, Philip. (2000). Endo-testing: Unit testing with mock objects. *extreme programming and flexible processes in software engineering - xp2000*.
- Maiya, Pallavi, Gupta, Rahul, Kanade, Aditya, & Majumdar, Rupak. (2016). Partial order reduction for event-driven multi-threaded programs. *Pages 680–697 of: Proceedings of the 22nd international conference on tools and algorithms for the construction and analysis of systems - volume 9636*. New York, NY, USA: Springer-Verlag New York, Inc.
- Mateescu, Alexandru, & Salomaa, Arto. (1997). *Handbook of formal languages*, vol. 1. New York, NY, USA: Springer-Verlag New York, Inc.

- Rabin, M. O., & Scott, D. (1959). Finite automata and their decision problems. *Ibm j. res. dev.*, **3**(2), 114–125.
- Raynal, Michel, & Singhal, Mukesh. (1996). Logical time: Capturing causality in distributed systems. *Computer*, **29**(2), 49–56.
- Visser, Willem, Havelund, Klaus, Brat, Guillaume, Park, SeungJoon, & Lerda, Flavio. (2003). Model checking programs. *Automated software engineering*, **10**(2), 203–232.
- Wyatt, Derek. (2013). *Akka Concurrency*. USA: Artima Incorporation.
- Zakeriyan, Arvin, Khamespanah, Ehsan, Sirjani, Marjan, & Khosravi, Ramtin. (2015). Jacco: More efficient model checking toolset for java actor programs. *Pages 37–44 of: Proceedings of the 5th international workshop on programming based on actors, agents, and decentralized control. AGERE! 2015*. New York, NY, USA: ACM.