TeSSLa: Runtime Verification of Non-synchronized Real-Time Streams*

Martin Leucker
ISP, Univ. of Lübeck, Germany
leucker@isp.uni-luebeck.de

César Sánchez
IMDEA Software Inst., Spain
cesar.sanchez@imdea.org

Torben Scheffel
ISP, Univ. of Lübeck, Germany
scheffel@isp.uni-luebeck.de

Malte Schmitz
ISP, Univ. of Lübeck, Germany
schmitz@isp.uni-luebeck.de

Alexander Schramm
IMDEA Software Inst., Spain
alexander.schramm@imdea.org

ABSTRACT

We present TeSSLa, a specification language based on stream runtime verification, designed for monitoring a specific class of real-time signals. Our monitors can observe concurrent systems with a shared clock, but where each component reports observations as signals that arrive to the monitor at different speeds and with different and varying latencies. The signals and streams that TeSSLa supports (including inputs and final verdicts) are not restricted to be Booleans but can be data from richer domains, including integers and reals with arithmetic operations and aggregations. Consequently, TeSSLa can be used both for checking logical properties, and for computing statistics and general numeric temporal metrics (and properties on these richer metrics). We present an online evaluation algorithm for TeSSLa specifications and show a formal proof of the correctness of concurrent implementations of the evaluation algorithm. Finally, we report an empirical evaluation of a highly concurrent Erlang implementation of the monitoring algorithm.

CCS CONCEPTS

• **Theory of computation** → **Logic and verification**; *Modal and temporal logics*; *Online algorithms*; *Semantics and reasoning*;

KEYWORDS

Runtime Verification, Online Monitoring, Stream Processing

ACM Reference Format:

Martin Leucker, César Sánchez, Torben Scheffel, Malte Schmitz, and Alexander Schramm. 2018. TeSSLa: Runtime Verification of Non-synchronized Real-Time Streams. In *Proceedings of SAC 2018: Symposium on Applied Computing (SAC 2018)*. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3167132.3167338>

*This work is supported in part by EU COST Action IC1402 “ArVi”, the BMBF projects ARAMIS II with funding ID 01 IS 16025 and CONIRAS with funding ID 01 IS 13029, the EU H2020 project COEMS under num. 732016, the EU H2020 project Elastest under num. 731535 and by Spanish MINECO Project “RISCO (TIN2015-71819-P)”.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SAC 2018, April 9–13, 2018, Pau, France

© 2018 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-5191-1/18/04.

<https://doi.org/10.1145/3167132.3167338>

1 INTRODUCTION

Runtime verification (RV) is an applied formal technique for verifying, analyzing and supporting software reliability. In contrast to static verification, in RV only one trace of the system under scrutiny is considered. Thus, RV sacrifices completeness guarantees to obtain an immediately applicable and formal extension of testing and debugging. A central problem in runtime verification is how to generate monitors from formal specifications (see [17, 23] for RV surveys).

In this paper we study how to perform runtime verification on concurrent systems that have a shared global clock but whose concurrent components emit events to the monitor at different speeds and with different delays. This assumption is common, for example, when observing embedded systems or when observing execution traces of software running on multi-core processors. At the low-level software analysis, the signals that these systems emit are real-time signals that remain constant between two observations, also known as *piece-wise constant* signals.

We are interested in *online* monitoring performed while the system is running (as opposed to offline monitoring through post-mortem analysis of dumped traces). Our target application also requires *non-intrusive* monitoring, meaning that the monitoring activity cannot perturb the execution of the system under observation. We use some hardware capabilities to obtain run-time information while the concurrent system executes. This information is dispatched to an external monitoring infrastructure that executes *outline* (as opposed to inlining the monitors within the system itself). See [22] for a definition and classification of these RV concepts.

In a nutshell, the goals of this paper are (1) to study how to describe sophisticated properties of continuous piece-wise constant signals, and (2) to efficiently monitor these properties against systems where each component is an event source that can emit events at different speeds and with different latencies. We say that these systems emit “non-synchronized in-order streams”. In our setting it is relevant to distinguish between the *system time* and *monitor time*. System time refers to the moments at which events are produced within the observed system. These instants are captured by the synchronized global clock that is used to time-stamp events. Monitor time refers to the instants at which events arrive at the monitor and when these events are processed in order to produce verdicts.

Event streams from hardware processors come at very high speeds, which imposes the additional requirement of crafting highly efficient monitoring implementations. We explore here how to

exploit parallelism and distributed implementations using multi-core platforms, while still formally guaranteeing the correctness of the monitors.

We propose here the specification language TeSSLa to achieve these goals. TeSSLa stands for Temporal Stream-based Specification Language. TeSSLa is based on Stream Runtime Verification (SRV) and has already been used for creating monitors in FPGA hardware in [10] without providing a concrete definition or further theoretical background. Early specification languages for RV were based on their counterparts in static verification, typically logics like LTL [25] or past LTL adapted for finite paths [5, 12, 18]. Similar formalisms proposed are based on regular expressions [31], timed regular expressions [2], rule based languages [3], or rewriting [28]. Stream runtime verification, pioneered by the tool LOLA [9], is an alternative to define monitors using streams. In SRV one describes the dependencies between input streams of values (observable events from the system under analysis) and defined streams (alarms, errors and output diagnosis information). These dependencies can relate the current value of a depending stream with the values of the same or other streams at the present moment, in past instants (like in past temporal formulas), or in future instants. In SRV there is a clean separation between the evaluation algorithms—that exploit the explicit dependencies between streams—and the data manipulation—expressed by each individual operation. SRV allows to generalize well-known evaluation algorithms from runtime verification to perform collections of numeric statistics from input traces.

SRV resembles synchronous languages [8]—like Esterel [6], Lustre [16] or Signal [14]—but these systems are causal because their intention is to describe systems and not observations, while SRV removes the causality assumption allowing to refer to future values. Another related area is Functional Reactive Programming (FRP) (see [13]), where reactive behaviors are defined using functional programs as building blocks to express reactions. As with synchronous languages, FRP is a programming paradigm and not a monitoring specification language, so future dependencies are not allowed in FRP. On the other hand SRV, was initially conceived for monitoring synchronous systems. See [7, 15, 26] for further developments on SRV. The semantics of temporal logics can also be defined using declarative dependencies between streams of values. For example, temporal testers [27] defined these dependencies for LTL. Likewise, the semantics of Signal Temporal Logic (STL) [11, 24] is defined in terms of the relation between a defined signal and the signals for its sub-expressions, based on Metric Interval Temporal Logic [1].

Here we extend SRV to real-time piece-wise constant signals, and study how to deal with the non-synchronized arrival of events. All previous approaches to SRV assume synchronous sampling and synchronous arrivals of events in all input streams. It is theoretically feasible, in some cases, to reduce the setting in this paper to the synchronous SRV, for example by assuming that all samples are made at instants multiple of a minimum delay, and executing the specification synchronously after every delay. However, the fast arrival of events would render such an approach impractical due to the large number of processing steps that would be required.

STL has also been used to create monitors on FPGAs [20] and for monitoring in different application areas (see for example [21, 30]). However, the assumptions of STL on the signals is different,

because the goal of STL is to analyze arbitrary continuous signals and not necessarily changes from digital circuits with accurate clocks. Sampling ratios and sampling instants are important issues in STL, while our signals are accurately represented by the stream of events at the changing points of the signal. In timed regular expressions (TRE) [2] the signals are also assumed to be piece-wise constant. However, our framework can handle much richer data domains than TREs and STL¹. TREs have been combined with STL [29] to get the advantages of both domains but again the signals analyzed are not necessarily piece-wise constant. Consequently, the results are approximate and sampling becomes, again, an important issue.

Contributions. In summary, the contributions of this paper are:

- (1) TeSSLa, an SRV-based specification language for real-time piece-wise constant signals. The syntax and semantics of TeSSLa are presented in Section 2, including the numerous core and library functions.
- (2) A method for the systematic generation of parallel and asynchronous online monitors for software monitoring TeSSLa specifications. These monitors handle the non-synchronized arrival of events from different input stream sources. In Section 3 we introduce a computational model for asynchronous concurrent monitors which allows to implement an online evaluation algorithm for TeSSLa specifications. Section 4 describes a prototype implementation developed in Erlang and an early empirical evaluation.

Finally, Section 5 concludes.

2 SYNTAX AND SEMANTICS OF TESSLA

We introduce in this section the real-time specification language TeSSLa². We first present some preliminaries, and then introduce the syntax and semantics of TeSSLa.

2.1 Preliminaries

We use two types of stream models as underlying formalism: piece-wise constant signals and event streams. We use $\mathbb{T}$ for the time domain (which can be $\mathbb{N}$, $\mathbb{Q}$, $\mathbb{R}$, etc), and D for the collection of data domains (Booleans, integers, reals, etc). Values from these data domains model observations and the output produced by the monitors.

Definition 2.1 (Event stream). An event stream is a partial function $\eta : \mathbb{T} \rightarrow D$ such that $E(\eta) := \{t \in \mathbb{T} \mid \eta(t) \text{ is defined}\}$ does not contain bounded infinite subsets. The set of all event streams is denoted by $\mathcal{E}_D$.

The set $E(\eta)$ is called the set of *event points* of η . When η is not defined at a time point t , that is $t \in \mathbb{T} \setminus E(\eta)$, we write $\eta(t) = \perp$. We use $\top$ as the “unit” value (the singleton domain). A finite event stream η can be naturally represented as a timed word, that is, a sequence $s_\eta = (t_0, \eta(t_0))(t_1, \eta(t_1)) \cdots \in (E(\eta) \times D)^*$ ordered by time ($t_i < t_{i+1}$) that contains a D value at all event points.

¹In the synchronous non-real-time case, [7] contains a thorough theoretical comparison of SRV versus temporal logics, regular expressions, etc. A similar comparison for real-time piece-wise signals is out of the scope of this paper.

²TeSSLa is available at www.isp.uni-luebeck.de/tessla

The second type of stream model that we consider is piece-wise constant signals, which have a value at every point in time. These signals change value only at a discrete set of positions, and remain constant between two change points.

Definition 2.2 (Signal). A *signal* is a total function $\sigma : \mathbb{T} \rightarrow D$ such that the set of *change points*

$$\Delta(\sigma) := \{t \in \mathbb{T} \mid \nexists t' < t : \forall t'' . t' < t'' < t : \sigma(t) = \sigma(t'')\}$$

does not contain bounded infinite subsets. The set of all signals is denoted by $\mathcal{S}_D$.

Every piece-wise constant signal can be exactly represented by an event stream that contains the change points of the signal as events, and whose value is the value of the signal after the change point. Hence, one can convert signals into event streams and vice-versa. Note that while in STL sampling provides an approximation of fully continuous signals, in TeSSLa event streams represent piece-wise constant signals with perfect accuracy.

Example 2.3. Consider the following streams e , s and $e2$, where e and $e2$ are interpreted as event streams and s as a signal.

The signal s has been created from e by using the value of the last event on e as value, with a default value 0. In turn, stream $e2$ is defined as the changes in value of s . When converting an event stream into a signal, only events that represent actual changes are generated. ■

We define next the syntax and semantics of TeSSLa.

2.2 Syntax of TeSSLa

We begin with an example to illustrate a simple TeSSLa specification. Specifications are written by defining streams in terms of other streams (and ultimately in terms of input streams). Streams marked as **out** are the verdict of the monitor and will be reported to the user.

Example 2.4. Consider the following TeSSLa specification:

```

in e: Events<Unit>
in s: Signal<Int>
define comp := eventCount(e) > s
define allowed := within(-1, 1, filter(e, comp))
define ok := implies(s > 5, allowed)
out ok

```

The first two lines define two input streams, e (an event stream without values) and s (a signal of integers). The Boolean signal $comp$ is *true* if the number of events of e (denoted by $eventCount(e)$) is greater than the current value of s , and *false* otherwise. The Boolean signal $allowed$ is *true* when there is an event that has not been filtered out from e in the interval $[-1, +1]$ around the current instant. The function $filter$ eliminates an event if the Boolean signal as the second parameter is *false*. Finally, the Boolean signal ok is *false* whenever s is greater than 5 and $allowed$ is *false*. Consider the input shown in the box below.

When $allowed$ is *true*, so will be ok . The signal ok will also be *true* as long as s is lower than 6. When s becomes 7, not enough events have happened on e and then $comp$ is *false*. Consequently, no event is left through the filter and $allowed$ is *false* too. But because s is greater than 5, ok becomes *false*. When s is set back to 6 and more events on e have happened, $allowed$ becomes *true* again. ■

The basic syntax of a TeSSLa specification *spec* is

```

spec ::= define name[: stype] := texpr | out name |
       in name: stype | spec spec
texpr ::= expr[: type]
expr ::= name | literal | name (texpr(, texpr)*)
type ::= btype | stype
stype ::= Signal<btype> | Events<btype>

```

A *name* is a nonempty string. Basic types *btype* cover typical types found in programming and verification like **Int**, **Float**, **String** or **Bool**. One of the main contributions of SRV is to generalize existing monitoring algorithms for logics (that produce Boolean verdicts) to algorithms that compute values from richer domains. The production **in** introduces input stream variables, and **define** introduces defined stream variables (also called output variables). Given a specification φ we use I for the set of input variables and O for the set of output variables, and write $\varphi(I, O)$. For example, in Example 2 above, $I = \{e, s\}$ and $O = \{comp, allowed, ok\}$. The marker **out** is used to denote those output variables that are the result of the specification and will be reported to the user. Each defined variable x is associated with a *defining equation* E_x given by the expression on the right hand side of the $:=$ symbol. Literals *literal* denote explicit values of basic types such as integers $-1, 0, 1, 2, \dots$, floating point numbers $0.1, -3.141593$ or strings "foo", "bar" (enclosed in double quotes).

We expand the syntax of basic TeSSLa by adding builtin functions, user defined macros and timing functions.

```

name ::= defName | timingFun | builtinFun | macro
timingFun ::= delay | shift | within

```

A *defName* is simply a name of a previously defined stream or constant. Timing functions allow to describe timing dependencies between streams. The function *delay* delays the values of a signal (or events of an event stream) by a certain amount of time. The function *shift* shifts the values of an event stream one unit into the future, that is, the first event becomes the second event, etc. The function *within* defines a signal which is *true* as long as some event of the given stream exists within the specified interval.

Macros are user defined functions identified by the construct **fun**. Macros can be expanded at compile time using their definition on a purely syntactical level because macros are not recursive.

Example 2.5. An example of a macro has already been used in the Example 2.4 because *implies* is not a builtin function, but the following macro:

```
fun implies(x, y) := or(not(x), y)
```

The full expressivity of TeSSLa is obtained using a set of *builtin functions*. Before describing the builtin functions, we define the temporal core of TeSSLa (which allows to define real time relations between streams) and the general semantics of TeSSLa.

2.3 Semantics

Semantics of Timing Functions. There are three timing functions *delay*, *shift* and *within*. The function *delay* is overloaded for signals

$$\begin{aligned} \text{delay} : \mathcal{S}_D \times \mathbb{T} \times D &\rightarrow \mathcal{S}_D \\ \text{delay}(\sigma, d, v)(t) &= \begin{cases} \sigma(t - d) & \text{if } t - d \geq 0 \\ v & \text{otherwise} \end{cases} \end{aligned}$$

and for event streams

$$\begin{aligned} \text{delay} : \mathcal{E}_D \times \mathbb{T} &\rightarrow \mathcal{E}_D \\ \text{delay}(\eta, d)(t) &= \begin{cases} \eta(t - d) & \text{if } t - d \geq 0 \\ \perp & \text{otherwise} \end{cases} \end{aligned}$$

The function *delay* delays a signal or an event stream by a given amount of time units. Since signals must always carry a value, a value v is provided as default. For event streams, the occurrence of each event is delayed by the indicated amount of time.

The *shift* function receives an event stream and produces the event stream that results from moving the value of each event to the next event. We use the following notation. Let

$$s_\eta = (t_0, \eta(t_0))(t_1, \eta(t_1))(t_2, \eta(t_2)) \dots,$$

we use $s_\eta^\rightarrow$ for the stream $(t_1, \eta(t_0))(t_2, \eta(t_1)) \dots$. The signature and interpretation of *shift* is:

$$\begin{aligned} \text{shift} : \mathcal{E}_D &\rightarrow \mathcal{E}_D \\ \text{shift}(s_\eta) &= \begin{cases} \varepsilon & \text{if } s_\eta = (t_0, \eta(t_0)) \text{ or } s_\eta = \varepsilon \\ s_\eta^\rightarrow & \text{otherwise} \end{cases} \end{aligned}$$

The last timing function is *within*, which already appeared in Example 2.4. It produces a Boolean valued signal that captures whether there is an event within a timing window:

$$\begin{aligned} \text{within} : \mathbb{T} \times \mathbb{T} \times \mathcal{E}_D &\rightarrow \mathcal{S}_\mathbb{B} \\ \text{within}(a, b, \eta)(t) &= \begin{cases} \text{true} & \text{if } E(\eta) \cap [t + a, t + b] \neq \emptyset \\ \text{false} & \text{otherwise} \end{cases} \end{aligned}$$

Semantics of TeSSLa. We define the semantics of TeSSLa in terms of evaluation models, as commonly done in SRV [9]. The intended meaning of TeSSLa specifications is to define output signals and event streams from input signals and event streams.

Consider a TeSSLa specification over input variables I and defined variables O . A *valuation* of a signal variable x of type D is an element of $\mathcal{S}_D$. Similarly, a valuation of a stream variable y of type D is an element of $\mathcal{E}_D$. We extend valuations to sets of variables in the usual way. If σ_I and σ_O are valuations of sets of variables I and

Figure 1: The dependency graph for the spec in Example 2. Inputs are shown in brown, constants in blue, outputs in green, computation nodes in red and some possible merges of computation nodes in dashed red.

O with $I \cap O = \emptyset$ then we use $\sigma_I \cup \sigma_O$ for the valuation of $I \cup O$ that coincides with σ_I on I and σ_O on O .

Let $\llbracket I \rrbracket$ be the value of a literal I , which is an element of its corresponding domain. Also, given a function name f we use $\llbracket f \rrbracket$ for the mathematical function that gives an interpretation of f (that is, a map from elements of the domain to an element of the co-domain). Given a valuation σ for each of the variables $I \cup O$ of a specification $\varphi(I, O)$, we can extend the valuation to expressions E over variables I and O , written $\llbracket E, \sigma \rrbracket$, recursively as follows:

- *variable name* ($E = \text{name}$):

$$\llbracket \text{name}, \sigma \rrbracket = \sigma(\text{name});$$

- *literal* ($E = I$):

$$\llbracket I, \sigma \rrbracket = \llbracket I \rrbracket;$$

- *function application* ($E = f(e_1, \dots, e_n)$):

$$\llbracket E, \sigma \rrbracket = \llbracket f \rrbracket(\llbracket e_1, \sigma \rrbracket, \dots, \llbracket e_n, \sigma \rrbracket)$$

An *evaluation model* of a specification $\varphi(I, O)$ is a valuation σ for variables I and O for every $x \in O$ with defining equation E_x : $\llbracket x, \sigma \rrbracket = \llbracket E_x, \sigma \rrbracket$. Informally, a valuation σ is an evaluation model whenever, for every defined variable x , the value that results when evaluating x and when evaluating its defining expression E_x coincide. We say that a specification $\varphi(I, O)$ is *well-defined* whenever for every valuation σ_I of input variables I there is a unique valuation σ_O of output variables O such that $\sigma_I \cup \sigma_O$ is an evaluation model of φ .

Non-recursive specifications. In order to guarantee that every specification is well-defined, we restrict TeSSLa specifications such that no variable x can depend on itself. More formally, given a specification $\varphi(I, O)$ we say that a variable x directly depends on a variable y if y appears in the defining equation E_x , and we write $x \rightarrow y$. We say that x depends on y if $x \rightarrow^+ y$ (where $\rightarrow^+$ is the transitive closure of $\rightarrow$). The dependency relation $x \rightarrow^+ y$ gives a necessary condition for y to affect in any way the value of x . The dependency graph has variables as nodes and the dependency relation as edges. Note that input variables and constants are leafs

in the dependency graph. The dependency graph of legal TeSSLa specifications must be non-recursive (i.e. for every x , $x \rightarrow^+ x$), which is easily checkable at compile time. If this is the case, the dependency graph is a DAG and a reverse topological order gives an evaluation order to compute the unique evaluation model. If all variables y preceding x have been assigned a valuation (the only one for which $\llbracket y \rrbracket = \llbracket E_y \rrbracket$) then $\llbracket E_x \rrbracket$ can be evaluated, which is the only possible choice for x .

Hence, this restriction guarantees that all TeSSLa specifications are well-defined. Figure 1 shows the dependency graph of the specification from Example 2.4.

Note also that if one merges a node n and the nodes n directly depends on, and replaces the function of n with the composition of the functions of the merged nodes, the resulting graph is still a DAG, and the streams computed will be the same. For example in Figure 1 nodes $>$ and *eventCount* could be merged. Such a node is called *computation node* or *node* for short.

TeSSLa Library of Builtin Functions. There are five types of functions in TeSSLa, apart from logical functions: arithmetic functions, aggregations, stream manipulators, timing functions (explained above) and temporal property functions. Tables 1 and 2 show a representative set of the functions provided by TeSSLa and the semantics of some of these functions.

Simple arithmetic functions provide capabilities for performing arithmetic operations on streams. In general, these functions take a set of signals as input and output another signal. Examples include basic arithmetic operations like *add*, *mul*, etc. More complex calculation functions in TeSSLa are *aggregations*, which take event streams as input and output a signal. Examples are *sum* that computes the sum of all events that happened on an event stream, and *eventCount* that counts the events. Another important aggregation function is *mrval*: $\mathcal{E}_D \times D \rightarrow \mathcal{S}_D$ which converts an event stream into a signal that receives the most recent value in the event stream (or a default value of type D provided as second argument).

Sampling functions convert a signal into an event stream. The function *changeOf*: $\mathcal{S}_D \rightarrow \mathcal{E}_D$ returns an event stream with an event at the point in time at which the signal changes. The function *sample*: $\mathcal{S}_D \times \mathcal{E}_D \rightarrow \mathcal{E}_D$ samples a signal by an event stream and returns an event stream with the values obtained from the signal. *Stream manipulators* allow to process event streams. Examples include a *filter* operator which allows to delete events and *merge* which fuses two event streams.

The *monitor* function delivers a scope for specifying properties in temporal logics like LTL or SALT [4] with classical or three-valued semantics [5]. Input properties are then compiled into a state machine which can be executed using Boolean signals as propositions and an event stream to step the monitor every time an event happens on this stream (to model discrete inputs).

3 ONLINE EVALUATION OF EFFICIENTLY MONITORABLE SPECIFICATIONS

The semantics provided in the previous section is denotational in the sense that it associates for each complete input valuation, the unique output valuation. We develop now an operational semantics for online monitoring TeSSLa specifications. In this section we *restrict TeSSLa specifications to refer to present and past values only*. These

specifications are known as efficiently monitorable [9], and satisfy that the values of an output variable x can immediately be resolved to their unique possible values (by evaluating E_x on the variables lower in the dependency graph) when a new input is processed. This fact leads to an online algorithm, which is implemented in the evaluation engines presented in this section. TeSSLa specifications are compiled into a single monitor that receives multiple inputs from the system under observation (each input is called a source and is associated with an input stream of the TeSSLa specification). Recall that each source can send events at different speeds and with different delays.

We represent behaviors of monitors as transducers on timed finite words, whose input is the stream of events arriving from any source. Let A be an input alphabet and B be an output alphabet (which correspond to the domains of input and output streams). A timed input letter is an element of $(A \times \mathbb{T})$ and a timed output letter is an element of $(B \times \mathbb{T})$. Given a timed letter a we use $t(a)$ to denote its time component. We use *source*(a) to represent the source of an input letter. We reserve the special symbol $\$$ (not in A or B) to denote the end of a timed word $a_0 \dots a_n\$$. We use Σ for $(A \times \mathbb{T})$ and Γ for $(B \times \mathbb{T})$, and ϵ for the empty word. Given a word w we use $L(w)$ for the letters occurring in w and $pos(a)$ for the position of a in w . The time-stamps of letters model the *system time*, while the position of a letter in the word models the *monitoring time*. We will show that every execution of a TeSSLa specification generates the exact same output if each input stream is the same even if the streams are non-synchronized, as long as these streams are in-order.

Definition 3.1 (In-order & Synchronized Inputs). We say that an input word w is

- *in-order* whenever for every $a, b \in L(w)$ if $pos(a) < pos(b)$ and $source(a) = source(b)$ then $t(a) \leq t(b)$.
- *synchronized* whenever for every $a, b \in L(w)$ if $pos(a) < pos(b)$ then $t(a) \leq t(b)$.

Example 3.2. Consider two input sources, one receives $(T, 0)(T, 3)$ and another receives $(F, 1)(F, 6)$. The following two inputs

- $w_1 : (T, 0)(F, 1)(T, 3)(F, 6)$ and
- $w_2 : (T, 0)(F, 1)(F, 6)(T, 3)$

are in-order inputs for these sources. However, w_1 is synchronized but w_2 is not, because in w_2 $(T, 3)$ is received after $(F, 6)$. The input w_2 is in-order because the input sources are different for the letters received in reverse order of time-stamps. ■

3.1 Evaluation Engines

We introduce *evaluation engines* to define the operational semantics and reason about the correctness of different implementations of TeSSLa online monitors. Consider a dependency graph G for a given specification $\phi(I, O)$, and let N be the collection of nodes of G . Nodes will be implemented by separate execution entities, possibly executing concurrently and asynchronously. Nodes that read some input stream are called *sources nodes*. Nodes communicate by sending timed letters, which we call (internal) *events*. Events are sent along the reversed edges of G . Evaluation engines equip each node with one queue for each of the node's inputs, which stores the non-processed events in that input. Queues support the standard operation for lists to get the head, the tail or to append to the end of the queue.

Arithmetics		Aggregations	Sampling & Filter
$add : S_N \times S_N \rightarrow S_N$	$max : S_N \times S_N \rightarrow S_N$	$sma : E_N \times N \rightarrow E_N$	$timestamps : E_D \rightarrow E_T$
$sub : S_N \times S_N \rightarrow S_N$	$min : S_N \times S_N \rightarrow S_N$	$maximum : E_N \times N \rightarrow S_N$	$changeOf : S_D \rightarrow E_{\{T\}}$
$mul : S_N \times S_N \rightarrow S_N$	$abs : S_N \rightarrow S_N$	$maximum : S_N \rightarrow S_N$	$ifThen : E_{D_1} \times S_{D_2} \rightarrow E_{D_2}$
$div : S_N \times S_{N^+} \rightarrow S_N$	$abs : E_N \rightarrow E_N$	$minimum : E_N \times N \rightarrow S_N$	$sample : S_{D_1} \times E_{D_2} \rightarrow E_{D_1}$
$gt : S_N \times S_N \rightarrow S_B$	$and : S_B \times S_B \rightarrow S_B$	$minimum : S_N \rightarrow S_N$	$filter : E_D \times S_B \rightarrow E_D$
$geq : S_N \times S_N \rightarrow S_B$	$or : S_B \times S_B \rightarrow S_B$	$sum : E_N \rightarrow S_N$	$ifThenElse : S_B \times S_D \times S_D \rightarrow S_D$
$leq : S_N \times S_N \rightarrow S_B$	$not : S_B \rightarrow S_B$	$eventCount : E_{D_1} \times E_{D_2} \rightarrow S_N$	$occursAny : E_{D_1} \times E_{D_2} \rightarrow E_{\{T\}}$
$eq : S_D \times S_D \rightarrow S_B$	$neg : E_B \rightarrow E_B$	$mrsv : E_D \times D \rightarrow S_D$	$occursAll : E_{D_1} \times E_{D_2} \rightarrow E_{\{T\}}$
$N \in \{N, Z, R\}, N^+ = N \setminus \{0\}$			$merge : E_D \times E_D \rightarrow E_D$

Table 1: A representative collection of the set of builtin functions.

$add(\sigma_1, \sigma_2)(t) = \sigma_1(t) + \sigma_2(t)$ $div(\sigma_1, \sigma_2)(t) = \frac{\sigma_1(t)}{\sigma_2(t)}$ $ge(\sigma_1, \sigma_2)(t) = \sigma_1(t) > \sigma_2(t)$ $max(\sigma_1, \sigma_2)(t) = \max\{\sigma_1(t), \sigma_2(t)\}$ $not(\sigma)(t) = \neg\sigma(t)$ $neg(\eta)(t) = \begin{cases} \neg\eta(t) & \text{if } t \in E(\eta) \\ \perp & \text{otherwise} \end{cases}$ $timestamps(\eta)(t) = \begin{cases} t & \text{if } t \in E(\eta) \\ \perp & \text{otherwise} \end{cases}$ $mrsv(\eta, d)(t) = \begin{cases} \eta(\max E(\eta) \cap [0, t]) & \text{if } E(\eta) \cap [0, t] \neq \emptyset \\ d & \text{otherwise} \end{cases}$	$changeOf(\sigma)(t) = \begin{cases} \top & \text{if } t \in \Delta(\sigma) \\ \perp & \text{otherwise} \end{cases}$ $sample(\sigma, \eta) = ifThen(\eta, \sigma)$ $filter(\eta, \sigma)(t) = \begin{cases} \eta(t) & \text{if } t \in E(\eta) \text{ and } \sigma(t) = true \\ \perp & \text{otherwise} \end{cases}$ $merge(\eta_1, \eta_2)(t) = \begin{cases} \eta_1(t) & \text{if } t \in E(\eta_1) \\ \eta_2(t) & \text{if } t \in E(\eta_2) \setminus E(\eta_1) \\ \perp & \text{otherwise} \end{cases}$ $occursAny(\eta_1, \eta_2)(t) = \begin{cases} \top & \text{if } t \in E(\eta_1) \cup E(\eta_2) \\ \perp & \text{otherwise} \end{cases}$
--	---

Table 2: Semantics for some of the representative builtin functions.

Formally, an *evaluation engine* is a transition system $E : \langle S, T, s_0 \rangle$, where the set of states S consists of the internal state of each node n , together with the state of each queue. In the initial state $s_0 \in S$ all queues are empty and all internal states of the nodes are set to their initial values.

An evaluation engine can be fed with input events from the input word w , which are placed into the input queues of the corresponding source nodes. During execution the evaluation engine can produce output events in the streams marked as output streams. A *transition* $\tau \in T$ involves the execution of exactly one node. A transition is *enabled* if there are events present in every input queue of the node. *Firing* a transition follows the operational semantics of the TeSSLa operation associated with the node, and consumes at least one event from some of the node's input queues, producing events into the output queues and updating the internal state of the node. The events produced are pushed to the corresponding input queues of the nodes directly depending on the executing node. We add the special transition $\lambda \in T$ for the empty transition where no event is consumed. We use $apply : S \times T \rightarrow S$ for the application of a transition to a state of the evaluation engine. The function $node : T \setminus \{\lambda\} \rightarrow N$ returns the node involved in a transition. A run is obtained by the repeated application of transitions.

Definition 3.3 (Run). A run of an evaluation engine

$$r = (\lambda, s_0)(\tau_1, s_1)(\tau_2, s_2) \dots \in (T \times S)^*$$

is a sequence of transitions and states such that for every $i > 1$, $node(\tau_i)$ is enabled at state s_{i-1} and $apply(s_{i-1}, \tau_i) = s_i$.

Given an input $w \in \Sigma^*$ and a run we get the output of the evaluation engine by concatenating all the output events produced in the run.

It is possible that more than one node is enabled in a given state. A *scheduler* chooses a transition to fire among the enabled transitions. While output events produced by any given node are ordered, outputs produced by different nodes may not be ordered, so concatenating output streams does not necessarily lead to a timed order sequence.

LEMMA 3.4 (ALL RUNS ARE FINITE). *Let E be an evaluation engine and $w \in \Sigma^*$ be an input. All runs $r \in (T \times S)^*$ of E on w are of finite length.*

3.1.1 Output Completeness. An event carries information about its occurrence, but there is so far no means to convey information about the absence of an event. We introduce extra events, called *progress events* whose only purpose is to inform nodes downstream about absences of events. A node n of an evaluation engine is called *output complete* if whenever n fires it produces at least one event in its output queue, either a real event or a progress event. In order to guarantee that all queues are emptied after a run finishes we extend the operational semantics of all TeSSLa operators to be output complete, while still implement the intended function. To see the importance of output completeness, consider a node n that is not output complete. This node could just consume all inputs without producing any events in its output (e.g a filter whose condition is always false). Every other node m depending directly on n will never be enabled, because the input queue of m coming from n will always be empty. If m has other input queues filled with some events, these queues will not be emptied at the end of the run. Lemma 3.4 guarantees that the run will terminate (because no node is enabled) but not necessarily that all queues are empty.

Progress events indicate that an input has progressed up to the time-stamp in the progress event, with no value change. It is easy to see that with output complete building blocks all nodes consume at least one input and all nodes generate exactly one output.

3.2 Timed Transducers

In order to prove properties of evaluation engines, in particular that for any given input the scheduler does not affect the events emitted in the output, we introduce a theory of timed asynchronous transducers. The main result is Theorem 3.12 which states that all runs of an engine are observationally equivalent, independently of the scheduler.

A “classical” synchronous transducer is an element of $(\Sigma \times \Gamma)^*$. However, we wish to model asynchronous transducers because we want to decouple the rate of arrival at input sources from the internal execution of the evaluation engine. A *timed transducer* F is $F \subset \Sigma^* \times \Gamma^*$. Our timed transducers will relate every input to some output (possibly ε). A timed transducer F is *complete* if for all $w \in \Sigma^*$ there is some $v \in \Gamma^*$ such that $(w, v) \in F$.

Definition 3.5. A timed transducer F is *strictly deterministic* if for all $w \in \Sigma^*$, and for all $v, v' \in \Gamma^*$, if $(w, v) \in F$ and $(w, v') \in F$, then $v = v'$.

As an example consider a transducer that delays every input letter by one time instant. Such a transducer is complete and strictly deterministic and would translate $(a, 0)(b, 1)(c, 2)$ into $(a, 1)(b, 2)(c, 3)$. However, strict determinism is too fine grained for our purposes, because we want to allow output letters to be produced out of order. We introduce a *timed reordering* function $timed : \Gamma^* \rightarrow \Gamma^*$ which reorders a word according to the time-stamps:

$$timed(b_0 \dots b_{i-1} b_i b_{i+1} \dots b_n) = b_i \cdot timed(b_0 \dots b_{i-1} b_{i+1} \dots b_n) \\ \text{if } t(b_i) < t(b_j) \text{ for all } j \neq i$$

Since words are finite, this recursive definition is well-defined. The following notion of asynchronous determinism captures more precisely the deterministic nature of asynchronous evaluation engines.

Definition 3.6 (Asynchronous determinism). A timed transducer F is called asynchronous deterministic if for all $w \in \Sigma^*$ and for all $v, v' \in \Gamma^*$ with $(w, v) \in F$ and $(w, v') \in F$, $timed(v) = timed(v')$.

Asynchronous determinism allows non-deterministic transducers to produce different outputs for the same input as long as the outputs are identical up-to reordering. Another important notion is asynchronous causality. We use $w \sqsubseteq w'$ to denote that w is a prefix of w' .

Definition 3.7 (Asynchronous causality). A timed transducer F is called asynchronous causal if for all $w, w' \in \Sigma^*$ with $w \sqsubseteq w'$ and $v, v' \in \Gamma^*$ with $(w, v) \in F$ and $(w', v') \in F$, $timed(v) \sqsubseteq timed(v')$.

Finally, we introduce observational equivalence between transducers.

Definition 3.8 (Observational Equivalence). Let F and G be two timed transducers over the same input and output alphabets, and let $w \in \Sigma^*$. We say that F and G are observational equivalent, and we write $F \equiv_O G$ whenever for all $v, u \in \Gamma^*$ with $(w, v) \in F$ and $(w, u) \in G$, $timed(v) = timed(u)$.

It is easy to see that observational equivalence is an equivalence relation for asynchronous deterministic transducers, because the definition of $\equiv_O$ is symmetric and transitive, and if F is asynchronous deterministic then $F \equiv_O F$.

3.3 Correctness

Before we give the main result we need some preliminary lemmas.

LEMMA 3.9 (PERSISTENCE OF ENABLEDNESS). Consider a run $r \in (T \times S)^*$. If a node n is enabled in a state s_i of r , then n stays enabled until it gets scheduled. In particular, the run contains a transition which involves n .

We say that a node n is *independent* of a node m if n is not a descendant of m in the dependency graph. Similarly, let $\tau_1, \tau_2 \in T$ be two transitions. We say that τ_1 is independent of τ_2 whenever $node(\tau_1)$ is independent of $node(\tau_2)$.

LEMMA 3.10 (EXCHANGE OF INDEPENDENT TRANSITIONS). Let τ be independent of τ' , then the following holds:

If

$$r_1 = (\lambda, s_0) \dots (\tau_{i-1}, s_{i-1})(\tau, s)(\tau', s'')(\tau_{i+1}, s_{i+1}) \dots (\tau_l, s_l)$$

is a run then

$$r_2 = (\lambda, s_0) \dots (\tau_{i-1}, s_{i-1})(\tau', s')(\tau, s'')(\tau_{i+1}, s_{i+1}) \dots (\tau_l, s_l)$$

is a run.

Definition 3.11 (Distance between Runs). Let $r, r' \in (T \times S)^*$ be two runs of an engine for the same input. Let p be the common prefix between r and r' and let τ be the transition taken in r after p (that is, the first different transition). We define the *distance* between r and r' as $\delta(r, r') = (|r| - |p|, j)$ where j is the position in r' at which τ is taken after p .

Note that this is well-defined because two runs for the same input: (1) are of the same length, because it takes exactly the same number of transitions to empty all the queues in all cases and (2) contain exactly the same transitions, but possibly in different order. For any run $r \in (T \times S)^*$ we get $\delta(r, r) = (0, 0)$.

Let $r = r_0 \dots r_n$ be a run of an engine for a given input w . The output of r is $output(r) = o_1 \dots o_n$ where o_i is the output produced by taking τ_i at s_{i-1} (or ε otherwise). The timed transducer defined by an engine E is:

$$\{(w, v) \in (\Sigma^* \times \Gamma^*) \mid \text{there is a run } r \text{ of } E \text{ on } w \text{ with } output(r) = v\}.$$

We are now ready to prove the main result.

THEOREM 3.12. Let E be an engine, $w \in \Sigma^*$ an input and r and r' two runs of E on w . Then $timed(output(r)) = timed(output(r'))$.

The proof of Theorem 3.12 proceeds by induction on the distance δ between runs defined above proving that all runs are equivalent using Lemma 3.10.

COROLLARY 3.13. All evaluation engines are asynchronous deterministic and asynchronous causal.

Theorem 3.12 allows to conclude that the scheduler has no impact on the output of an engine, up-to timed reorderings of the output. Hence, one can take a very deterministic scheduler to reason about the outcome, because every other scheduler is guaranteed to

Figure 2: Benchmarks. The black solid plot indicates the total time and the blue dashed line indicates the relative time per event or node, resp.

produce an observationally equivalent output (for the same input). For example, a simple reverse topological order allows to easily prove that the resulting deterministic transducer corresponds to the intended denotational semantics of the corresponding TeSSLa specification.

4 IMPLEMENTATION AND EVALUATION

We report here an empirical evaluation of an implementation of the TeSSLa evaluation engine³. Our implementation consists of two parts. First, a *compiler* translates a TeSSLa specification into a dependency graph (and performs type checking, macro expansion, and type inference for the defined streams, and also checks that the specification is non-recursive). Then, the *evaluation engine*, written in Elixir, takes an input trace and the dependency graph generated by the compiler and produces an output trace. Elixir, built on top of Erlang, is based on the actor model [19] which allows to deploy code over multiple cores or even distributed systems. Our implementation maps nodes to actors. Theorem 3.12 guarantees that the outcome of an execution is independent of the concrete execution of the Erlang scheduler. Additionally, we also implemented auxiliary tools to use TeSSLa for the runtime verification of C programs based on a software instrumentation. To evaluate the performance of our implementation we created several artificial benchmarks. We measured the execution time in relation to the number of processor cores, the length of the input trace and the size of the specification. All benchmarks were performed on the same machine with up to 16 cores and 48GB of RAM. The results displayed in Figure 2 are the average of 20 runs.

Number of cores. For this benchmark we created a TeSSLa specification with 16 computation nodes, consisting of a chain of *abs* functions. The input trace contains exactly 10,000 events. The results show that the execution time decreases drastically with an increase on the number of cores available, suggesting that our implementation is able to make an effective use of parallelism even though the dependency graph is completely linear. Our asynchronous implementation is able to exploit parallelism automatically in a pipeline fashion.

Number of events in the input. To study the dependency of the execution time with the input length we modified the specification

with different input sizes. We also display (with a dash line) the average time per event. With an increase on the input length, the static overhead is amortized quickly and the length of the trace becomes less relevant for the average event time. Consequently, the relative time shows a decay as more events are added.

Number of nodes in the specification. The execution time also grows linearly in the size of the specification. For this benchmark we used input traces of 1,000 events and increased the number of computation nodes in the specification by adding more calls to *abs* to the composition chain. Again, the relative time per event shows a decay as more nodes are added, because the static overhead becomes less relevant.

5 CONCLUSION

We presented TeSSLa, a stream-based runtime verification language for non-synchronized streams of piece-wise constant real-time signals. We defined the operational semantics of TeSSLa in terms of evaluation engines. We defined a version of timed transducers to prove that all evaluations of a TeSSLa specification produce the same output up-to timed reordering, independently of the scheduler. This result enables different evaluation engines, including asynchronous evaluation engines based on actors—which allow to exploit multi-core parallelism—and evaluation engines implemented in FPGAs which enable the utilization of massive hardware parallelism. We report in this paper an implementation of an evaluation engine written in Elixir/Erlang.

For simplicity, our timed transducers operate on finite words, but all definitions and results can be extended to infinite words under the assumption of a fair scheduler. Similarly, the interleaving semantics of the evaluation engine can be extended to true concurrency between independent nodes. We are currently working on an extension of the TeSSLa language that allows to define the time dependencies between streams explicitly (in a controlled way) and a library of builtin functions on top of the resulting core language.

ACKNOWLEDGMENTS

We thank Jannis Harder and Sebastian Hungerecker for their work on TeSSLa and its compiler.

³Tools available at www.isp.uni-luebeck.de/tessla

REFERENCES

- [1] Rajeev Alur, Tomás Feder, and Thomas A. Henzinger. 1996. The benefits of relaxing punctuality. *J. ACM* (1996).
- [2] Eugene Asarin, Paul Caspi, and Oded Maler. 2002. Timed regular expressions. *J. ACM* 49, 2 (2002), 172–206.
- [3] Howard Barringer, Allen Goldberg, Klaus Havelund, and Koushik Sen. 2004. Rule-Based Runtime Verification. In *Proc. of VMCAI’04 (LNCS 2937)*. Springer, 44–57.
- [4] Andreas Bauer and Martin Leucker. 2011. The Theory and Practice of SALT. In *NASA Formal Methods (NFM)*. Springer, 13–40.
- [5] Andreas Bauer, Martin Leucker, and Chrisitan Schallhart. 2011. Runtime Verification for LTL and TLTL. *ACM T. Softw. Eng. Meth.* 20, 4 (2011), 14.
- [6] Gérard Berry. 2000. *Proof, language, and interaction: essays in honour of Robin Milner*. MIT Press, Chapter The foundations of Esterel, 425–454.
- [7] Laura Bozelli and César Sánchez. 2014. Foundations of Boolean Stream Runtime Verification. In *In Proc. RV’14 (LNCS)*, Vol. 8734. Springer, 64–79.
- [8] Paul Caspi and Marc Pouzet. 1996. Synchronous Kahn Networks. In *Proc. of ICFP’96*. ACM Press, 226–238.
- [9] Ben D’Angelo, Sriram Sankaranarayanan, César Sánchez, Will Robinson, Bernd Finkbeiner, Henny B. Sipma, Sandeep Mehrotra, and Zohar Manna. 2005. LOLA: Runtime Monitoring of Synchronous Systems. In *Proc. of TIME’05*. IEEE, 166–174.
- [10] Normann Decker, Philip Gottschling, Christian Hochberger, Martin Leucker, Torben Scheffel, Malte Schmitz, and Alexander Weiss. 2017. Rapidly Adjustable Non-Intrusive Online Monitoring for Multi-core Systems. In *20th Brazilian Symposium on Formal Methods (SBMF 2017)*. Springer.
- [11] Alexandre Donzé, Oded Maler, Ezio Bartocci, Dejan Nickovic, Radu Grosu, and Scott A. Smolka. 2012. On Temporal Logic and Signal Processing. In *In Proc. of ATVA’12 (LNCS)*, Vol. 7561. Springer, 92–106.
- [12] Cindy Eisner, Dana Fisman, John Havlicek, Yoad Lustig, Anthony McIsaac, and David Van Campenhout. 2003. Reasoning with Temporal Logic on Truncated Paths. In *Proc. of CAV’03 (LNCS 2725)*, Vol. 2725. Springer, 27–39.
- [13] Conal Eliot and Paul Hudak. 1997. Functional Reactive Animation. In *Proc. of ICFP’07*. ACM, 163–173.
- [14] Thierry Gautier, Paul Le Guernic, and Lic Besnard. 1987. SIGNAL: A declarative language for synchronous programming of real-time systems. In *Proc. of FPCA’87 (LNCS 274)*. Springer, 257–277.
- [15] Alwyn E. Goodloe and Lee Pike. 2010. *Monitoring distributed real-time systems: A survey and future directions*. Technical Report. NASA Langley Research Center.
- [16] Nicolas Halbwachs, Paul Caspi, D. Pilaud, and J.A. Plaice. 1987. Lustre: a declarative language for programming synchronous systems. In *Proc. of POPL’87*. ACM Press, 178–188.
- [17] Klaus Havelund and Allen Goldberg. 2005. Verify your runs. In *Proc. of VSTTE’05 (LNCS 4171)*. Springer, 374–383.
- [18] Klaus Havelund and Grigore Rou. 2002. Synthesizing Monitors for Safety Properties. In *Proc. of TACAS’02 (LNCS 2280)*. Springer, 342–356.
- [19] Carl Hewitt, Peter Bishop, and Richard Steiger. 1973. A Universal Modular ACTOR Formalism for Artificial Intelligence. *IJCAI* (1973), 235–245.
- [20] Stefan Jaksic, Ezio Bartocci, Radu Grosu, Reinhard Kloibhofer, Thang Nguyen, and Dejan Nickovic. 2015. From signal temporal logic to FPGA monitors. In *Proc. of MEMOCODE 2015*. 218–227.
- [21] Stefan Jaksic, Ezio Bartocci, Radu Grosu, and Dejan Nickovic. 2016. Quantitative Monitoring of STL with Edit Distance. In *Proc. of RV’16 (LNCS)*, Vol. 10012. 201–218.
- [22] Martin Leucker. 2011. Teaching Runtime Verification. In *Proc. of RV’11 (LNCS)*. Springer, 34–48.
- [23] Martin Leucker and Christian Schallhart. 2009. A Brief Account of Runtime Verification. *J. Logic Algebr. Progr.* 78, 5 (2009), 293–303.
- [24] Oded Maler and Dejan Nickovic. 2004. Monitoring Temporal Properties of Continuous Signals. In *FTRFT*. 152–166.
- [25] Zohar Manna and Amir Pnueli. 1995. *Temporal Verification of Reactive Systems: Safety*. Springer, New York.
- [26] Lee Pike, Alwyn Goodloe, Robin Morisset, and Sebastian Niller. 2010. Copilot: A Hard Real-Time Runtime Monitor. In *Proc. of RV’10 (LNCS 6418)*. Springer.
- [27] Amir Pnueli and Aleksandr Zaks. 2006. PSL Model Checking and Run-Time Verification Via Testers. In *Proc. of FM’06 (LNCS 4085)*. Springer, 573–586.
- [28] Grigore Rou and Klaus Havelund. 2005. Rewriting-Based Techniques for Runtime Verification. *Automated Software Engineering* 12, 2 (2005), 151–197.
- [29] Konstantin Selyunin, Stefan Jaksic, Thang Nguyen, Christian Reidl, Udo Hafner, Ezio Bartocci, Dejan Nickovic, and Radu Grosu. 2017. Runtime Monitoring with Recovery of the SENT Communication Protocol. In *Proc. of CAV’17 (LNCS)*, Vol. 10426. Springer, 336–355.
- [30] Konstantin Selyunin, Thang Nguyen, Ezio Bartocci, and Radu Grosu. 2016. Applying Runtime Monitoring for Automotive Electronic Development. In *Proc. of RV’16 (LNCS)*, Vol. 10012. 462–469.
- [31] Koushik Sen and Grigore Rou. 2003. Generating Optimal Monitors for Extended Regular Expressions. *ENTCS* 89, 2 (2003), 226–245.