



# Runtime Verification for Timed Event Streams with Partial Information

Martin Leucker<sup>1(✉)</sup>, César Sánchez<sup>2(✉)</sup>, Torben Scheffel<sup>1(✉)</sup>,  
Malte Schmitz<sup>1(✉)</sup>, and Daniel Thoma<sup>1(✉)</sup>

<sup>1</sup> University of Lübeck, Lübeck, Germany

{leucker,scheffel,schmitz,thoma}@isp.uni-luebeck.de

<sup>2</sup> IMDEA Software Institute, Madrid, Spain

cesar.sanchez@imdea.org

**Abstract.** Runtime Verification (RV) studies how to analyze execution traces of a system under observation. Stream Runtime Verification (SRV) applies stream transformations to obtain information from observed traces. Incomplete traces with information missing in gaps pose a common challenge when applying RV and SRV techniques to real-world systems as RV approaches typically require the complete trace without missing parts. This paper presents a solution to perform SRV on incomplete traces based on abstraction. We use TeSSLa as specification language for non-synchronized timed event streams and define abstract event streams representing the set of all possible traces that could have occurred during gaps in the input trace. We show how to translate a TeSSLa specification to its abstract counterpart that can propagate gaps through the transformation of the input streams and thus generate sound outputs even if the input streams contain gaps and events with imprecise values. The solution has been implemented as a set of macros for the original TeSSLa and an empirical evaluation shows the feasibility of the approach.

## 1 Introduction

Runtime verification (RV) is a dynamic formal method for software system reliability. RV studies how to analyze and evaluate traces against formal specifications and how to obtain program traces from the system under observation, e.g., through software instrumentation or utilization of processors' embedded trace units. Since RV only inspects one execution trace of the system, it is often regarded to be a readily applicable but incomplete approach, that combines formal verification with testing and debugging.

Most early RV languages were based on logics common in static verification, like LTL [21], past LTL adapted for finite paths [4, 11, 18], regular expressions [22] or timed regular expressions [2]. For these logics, the monitoring problem consists on computing a Boolean verdict indicating whether the trace fulfills the


---

This work was funded in part by the Madrid Regional Government under project *S2018/TCS-4339 (BLOQUES-CM)*, by EU H2020 projects 731535 *Elastest* and 732016 *COEMS*, by Spanish National Project *BOSCO (PGC2018-102210-B-100)* and by the BMBF project *ARAMiS II* with funding ID 01 IS 16025.

© The Author(s) 2019

B. Finkbeiner and L. Mariani (Eds.): RV 2019, LNCS 11757, pp. 273–291, 2019.

[https://doi.org/10.1007/978-3-030-32079-9\\_16](https://doi.org/10.1007/978-3-030-32079-9_16)



**Fig. 1.** Example trace for a typical SRV specification (left) with two input streams *values* (with numeric values) and *resets* (with no internal value). The intention of the specification is to accumulate in the output stream *sum* all values since the last reset. The intermediate stream *cond* is derived from the input streams indicating if *reset* has currently the most recent event, and thus the sum should be reset to 0. If the input streams contain gaps (dotted regions on the right) some information can no longer be computed, but after a reset event the computation recovers from the data loss during the gap.  $\top$  denotes events with unknown data.

specification. In contrast to static analysis, however, considering only a single concrete trace enables the application of more complex analyses: Stream Runtime Verification (SRV) [6, 7, 10] uses stream transformations to derive additional streams as verdicts from the input streams. Using SRV one can still check if the input stream is conformant with a specification, but additionally verify streams in terms of their events’ data: streams in SRV can store data from richer domains than Booleans, including numerical values or user defined data-types, so SRV languages can extract quantitative values and express quantitative properties like “*compute the average retransmission time*” or “*compute the longest duration of a function*”. SRV cleanly separates the temporal dependencies that the stream transformation algorithms follow from the concrete operations to be performed on the data, which are specific to each data-type. As an example for SRV consider the trace diagram on the left of Fig. 1. We consider non-synchronized event streams, i.e., sequences of events with increasing timestamps and values from a data domain. Using non-synchronized event streams one can represent events arriving on different streams with different frequencies in a compact way with little computation overhead because there is no need to process additional synchronization events in the stream-transformation process. In this paper we use the TeSSLa specification language [7], an SRV language for non-synchronized, timed event streams. TeSSLa has been defined to be general enough to allow for a natural translation from other common SRV formalisms, e.g., Lola [10] and Striver [16]. Therefore, our results carry over to these languages as well.

Since RV is performed on traces obtained from the system under test in the deployed environment, it is a common practical problem for RV techniques that the traces do not cover the entire run of the system. However, most of the previous RV approaches require the trace to be available without any interruptions in order to obtain a verdict, because this knowledge is assumed in the semantics of the specification logics. Especially in the case of interrupted traces with some data losses applying previous RV techniques can be very challenging. Unfortunately those traces occur very often in practical testing and debugging

scenarios, e.g., due to interrupted experiments, buffer overflows, network errors or any other temporary problem with the trace retrieval.

In this paper we present a solution to the problem of evaluating traces with imprecise values and even interrupted traces. Our only assumption is that we have exact knowledge of the imprecision of the trace in the following sense: (1) for events with imprecise values we know the range of values and (2) for data losses we know when we stop getting information and when the trace becomes reliable again. We call such a sequence of uncertainty a *gap* in the trace. Our solution automatically propagates gaps and imprecisions, and allows to obtain sound verdicts even in the case of missing information in the input trace.

Figure 1 on the right displays a case where the input stream *values* has a long gap in the middle. It is not possible to determine the events in the output stream *sum* during that gap, because we do not even know if and how many events might have happened during that gap. Thus, the intermediate stream *cond* and the output stream *sum* simply copy that gap representing any possible combination of events that might occur. The first event after the gap is the one with the value 3 on *values*. Because no reset happened after the end of the gap, we would add 3 to the latest event's value on *sum*, but the gap is the latest on *sum*. Thus, we only know that this input event on *values* causes an event on *sum* independently of what might have happened during the gap, but the value of that event completely depends on possible events occurring during the gap. After the next event on *reset* the values of the following events on *sum* are independent of any previous events. The monitor can fully recover from the missing information during the gap and can again produce events with precise values.

In order to realize this propagation of gaps through all the steps of the stream-transformation we need to represent all potentially infinitely many concrete traces (time is dense and values are for arbitrary domains) that might have happened during gaps and imprecise events. An intuitive approach would be a symbolic representation in terms of constraint formulas to describe the set of all possible streams. These formulas would then be updated while evaluating the input trace. While such a symbolic execution might work for shorter traces, the representation can grow quickly with each input event. Consequently the computational cost could grow prohibitively with the trace length for many input traces. Instead, in this paper we introduce a framework based on abstraction [8,9]. We use abstraction in two ways:

- (1) Streams are lifted from concrete domains of data to abstract domains to model possible sets of values. For example, in our solution a stream can store intervals as abstract numerical values.
- (2) We define the notion of abstract traces, which extend timed streams with the capabilities of representing gaps. Intuitively, an abstract trace over-approximates the sets of concrete traces that can be obtained by filling the gaps with all possible concrete events.

Our approach allows for both gaps in the input streams as well as events carrying imprecise values. Such imprecise values can be modelled by abstract domains,

e.g., intervals of real numbers. Since we rely on abstraction, we can avoid false negatives and false positives in the usual sense: concrete verdicts are guaranteed to hold and imprecise verdicts are clearly distinguished from concrete verdicts. The achievable precision depends on the specification and the input trace.

After reproducing the semantics of the basic TeSSLa operators in Sect. 2, we introduce abstract semantics of the existing basic operators of TeSSLa in Sect. 3. Using these abstract TeSSLa operators, we can take a TeSSLa specification on streams and replace every TeSSLa operator with its abstract counterpart and derive an abstraction of the specification on abstract event streams. We show that the abstract specification is a sound abstraction of the concrete specification, i.e., every concrete verdict generated by the original specification on a set  $S$  of possible input traces is represented by the abstract verdict applied to an abstraction of  $S$ . We further show that the abstract TeSSLa operators are a perfect abstraction of their concrete counterparts, i.e., that applying the concrete operator on all individual elements of  $S$  doesn't get you more accurate results. Finally, we show that an abstract TeSSLa specification can be implemented using the existing TeSSLa basic operators by representing an abstract event stream as multiple concrete event streams carrying information about the events and the gaps. Since the perfect accuracy of the individual abstract TeSSLa operators does not guarantee perfect accuracy of their compositions, we discuss the accuracy of composed abstract TeSSLa specifications in Sect. 4. Next we present in Sect. 5 an advanced use-case where we apply abstract TeSSLa to streams over a complex data domain of unbounded queues, which are used to compute the average of all events that happened in the sliding window of the last five time units. Section 6 evaluates the overhead and the accuracy of the presented abstractions on representative example specifications and corresponding input traces with gaps. An extended preprint version of this paper is available as [19].

*Related Work.* SRV was pioneered by LOLA [10,13,14]. TeSSLa [7] generalises to asynchronous streams the original idea of LOLA of recursive equations over stream transformations. Its design is influenced by formalisms like stream programming languages [5,15,17] and functional reactive programming [12]. Other approaches to handle data and time constraints include Quantitative Regular Expressions QRE [1] and Signal Temporal Logic [20].

While ubiquitous in practice, the problem of gaps in an observation trace has not been studied extensively. To the best of our knowledge, abstraction techniques have not been applied to the evaluation of stream-based specifications. However, approaches to handle the absence of events or ordering information have been presented for MTL [3] and past-time LTL [24]. State estimation based on Markov models has been applied to replace absent information by a probabilistic estimation [23]. The concept of abstract interpretation used throughout this paper has been introduced in [8].

## 2 The TeSSLa Specification Language

A *time domain* is a totally ordered semi-ring  $(\mathbb{T}, 0, 1, +, \cdot, \leq)$  that is positive, i.e.,  $\forall t \in \mathbb{T} 0 \leq t$ . We extend the order on time domains to the set  $\mathbb{T}_\infty = \mathbb{T} \cup \{\infty\}$  with  $\forall t \in \mathbb{T} t < \infty$ . Given a time domain  $\mathbb{T}$ , an *event stream* over a data domain  $\mathbb{D}$  is a finite or infinite sequence  $s = t_0 d_0 t_1 \dots \in \mathcal{S}_\mathbb{D} = (\mathbb{T} \cdot \mathbb{D})^\omega \cup (\mathbb{T} \cdot \mathbb{D})^* \cdot (\mathbb{T}_\infty \cup \mathbb{T} \cdot \mathbb{D}_\perp)$  where  $\mathbb{D}_\perp := \mathbb{D} \cup \{\perp\}$  and  $t_i < t_{i+1}$  for all  $i$  with  $0 < i + 1 < |s|$  ( $|s|$  is  $\infty$  for infinite sequences). An infinite event stream is an infinite sequence of timestamps and data values representing the stream's events. A finite event stream is a finite sequence of timestamped events up to a certain timestamp that indicates the progress of the stream. A stream can end with:

- a timestamp without a data value that denotes progress up to but not including that timestamp,
- a timestamp followed by  $\perp$  (or a data value) which denotes progress up to and including that timestamp (and an event at that timestamp),
- $\infty$ , which indicates that no additional events will ever arrive on this stream.

We refer to these cases as *exclusive*, *inclusive* and *infinite progress*, resp.

Streams  $s \in \mathcal{S}_\mathbb{D}$  can be seen as functions  $s : \mathbb{T} \rightarrow \mathbb{D} \cup \{\perp, ?\}$  such that  $s(t)$  is a value  $d$  if  $s$  has an event with value  $d$  at time  $t$  or  $\perp$  if there is no event at time  $t$ . For timestamps after the progress of the stream  $s(t)$  is  $?$ . Formally,  $s(t) = d$  if  $s$  contains  $td$ ,  $s(t) = \perp$  if  $s$  does not contain  $t$ , but contains a  $t' > t$  or  $s$  ends in  $t_\perp$ , and  $s(t) = ?$  otherwise. We use  $\text{ticks}(s)$  for the set  $\{t \in \mathbb{T} \mid s(t) \in \mathbb{D}\}$  of timestamps where  $s$  has events. A stream  $s$  is a *prefix* of stream  $r$  if  $\forall t \in \mathbb{T} s(t) \in \{r(t), ?\}$ . We use the unit type  $\mathbb{U} = \{\square\}$  for streams carrying only the single value  $\square$ . A TeSSLa specification consists of a collection of stream variables and possibly recursive equations over these variables using the operators **nil**, **unit**, **time**, **lift**, **last** and **delay**. The semantics of recursive equations is given as the least fixed-point of the equations seen as a function of the stream variables and fixed input streams. See [7] for more details.

- **nil** =  $\infty \in \mathcal{S}_\emptyset$  is the stream without any events and infinite progress.
- **unit** =  $0 \square \infty \in \mathcal{S}_\mathbb{U}$  is the stream with a single unit event at timestamp zero and infinite progress.
- **time** :  $\mathcal{S}_\mathbb{D} \rightarrow \mathcal{S}_\mathbb{T}$ , **time**( $s$ ) :=  $z$  maps the event's values to their timestamps:  $z(t) = t$  if  $t \in \text{ticks}(s)$  and  $z(t) = s(t)$  otherwise.
- **lift** :  $(\mathbb{D}_\perp \times \dots \times \mathbb{D}_\perp \rightarrow \mathbb{D}_\perp) \rightarrow (\mathcal{S}_{\mathbb{D}_1} \times \dots \times \mathcal{S}_{\mathbb{D}_n} \rightarrow \mathcal{S}_\mathbb{D})$ , **lift**( $f$ )( $s_1, \dots, s_n$ ) :=  $z$  lifts a function  $f$  on values to a function on streams by applying  $f$  to the stream's values for every timestamp. The function  $f$  must not generate new events, i.e., must fulfill  $f(\perp, \dots, \perp) = \perp$ .

$$z(t) = \begin{cases} f(s_1(t), \dots, s_n(t)) & \text{if } s_1(t) \neq ?, \dots, s_n(t) \neq ? \\ ? & \text{otherwise} \end{cases}$$

- **last** :  $\mathcal{S}_\mathbb{D} \times \mathcal{S}_{\mathbb{D}'}$   $\rightarrow \mathcal{S}_\mathbb{D}$ , **last**( $v, r$ ) :=  $z$  takes a stream  $v$  of values and a stream  $r$  of triggers. It outputs an event with the previous value on  $v$  for every event

on  $r$ .

$$z(t) = \begin{cases} d & t \in \text{ticks}(r) \text{ and } \exists_{t' < t} \text{isLast}(t, t', v, d) \\ \perp & r(t) = \perp \text{ and } \text{defined}(z, t), \text{ or } \forall_{t' < t} v(t') = \perp \\ ? & \text{otherwise} \end{cases}$$

$\text{isLast}(t, t', v, d) \stackrel{\text{def}}{=} v(t') = d \wedge \forall_{t'' | t' < t'' < t} v(t'') = \perp$  holds if  $t'd$  is the last event on  $v$  until  $t$ , and  $\text{defined}(z, t) \stackrel{\text{def}}{=} \forall_{t' < t} z(t') \neq ?$  holds if  $z$  is defined until  $t$  (exclusive).

Using the basic operators we can now derive the following utility functions:

- ▷  $\text{const}(c)(a) := \text{lift}(f_c)(a)$  with  $f_c(d) := c$ . This function maps the values of all events of the input stream  $a$  to a constant value  $c$ . Using  $\text{const}$  we can lift constants into streams representing a constant signal with this value, e.g.,  $\text{true} := \text{const}(\text{true})(\text{unit})$  or  $\text{zero} := \text{const}(0)(\text{unit})$ .
- ▷  $\text{merge}(x, y) := \text{lift}(f)(x, y)$  with  $f(a \neq \perp, b) = a$  and  $f(\perp, b) = b$ , which combines events from two streams, prioritizing the first stream.

Event streams in TeSSLa can also be interpreted as a *continuous signals*. Using **last** one can query the last known value of an event stream  $s$  and interpret the events on  $s$  as points where a piece-wise constant signal changes its value. By combining the **last** and **lift** operators, we can realize:

- ▷ *signal lift* for total functions  $f : \mathbb{D} \times \mathbb{D}' \rightarrow \mathbb{D}''$  as  $\text{slift}(f)(x, y) := \text{lift}(g_f)(x', y')$  with  $x' := \text{merge}(x, \text{last}(x, y))$  and  $y' := \text{merge}(y, \text{last}(y, x))$ , as well as  $g_f(a \neq \perp, b \neq \perp) := f(a, b)$ ,  $g_f(\perp, b) := \perp$ , and  $g_f(a, \perp) := \perp$ .

*Example 1.* We can now specify the stream transformations shown on the left in Fig. 1 in TeSSLa. Let  $\text{resets} \in \mathcal{S}_{\mathbb{U}}$  and  $\text{values} \in \mathcal{S}_{\mathbb{Z}}$  be two external input event streams. We then derive  $\text{cond} \in \mathcal{S}_{\mathbb{B}}$  and  $\text{lst}, \text{sum} \in \mathcal{S}_{\mathbb{Z}}$  as follows:

$$\begin{aligned} \text{cond} &= \text{slift}(\leq)(\text{time}(\text{resets}), \text{time}(\text{values})) & f : \mathbb{B} \times \mathbb{Z} \times \mathbb{Z} \rightarrow \mathbb{Z} \text{ with} \\ \text{lst} &= \text{merge}(\text{last}(\text{sum}, \text{values}), \text{zero}) & f(c, l, v) = \begin{cases} 0 & \text{if } c = \text{true} \\ l + v & \text{otherwise} \end{cases} \\ \text{sum} &= \text{slift}(f)(\text{cond}, \text{lst}, \text{values}) \end{aligned}$$

Using the operators described above one can only derive streams with timestamps that are already present in the input streams. To derive streams with events at computed timestamps one can use the **delay** operator, which is described in [7].

### 3 Abstract TeSSLa

*Preliminaries.* Given two partial orders  $(A, \preceq)$  and  $(B, \preceq)$ , a *Galois Connection* is a pair of monotone functions  $\alpha : A \rightarrow B$  and  $\gamma : B \rightarrow A$  such that, for all  $a \in A$  and  $b \in B$ ,  $\alpha(a) \preceq b$  if and only if  $a \preceq \gamma(b)$ . Let  $(A, \preceq)$  be a partial

order,  $f : A \rightarrow A$  a monotone function and  $\gamma : B \rightarrow A$  a function. The function  $f^\# : B \rightarrow B$  is an *abstraction* of  $f$  whenever, for all  $b \in B$ ,  $f(\gamma(b)) \preceq \gamma(f^\#(b))$ . If  $(\alpha, \gamma)$  is a Galois Connection between  $A$  and  $B$ , the function  $f^\# : B \rightarrow B$  such that  $f^\#(b) := \alpha(f(\gamma(b)))$  is a *perfect abstraction* of  $f$ .

In this section we define the abstract counterparts of the TeSSLa operators, listed in Sect. 2. A *data abstraction* of a data domain  $\mathbb{D}$  is an abstract domain  $\mathbb{D}^\#$  with an element  $\top \in \mathbb{D}^\#$  and an associated concretisation function  $\gamma : \mathbb{D}^\# \rightarrow 2^{\mathbb{D}}$  with  $\gamma(\top) = \mathbb{D}$ . The abstract value  $\top$  represents any possible value from the data domain and can be used to model an event with known timestamp but unknown value. A *gap* is a segment of an abstract event stream that represents all combinations of events that could possibly occur in that segment (both in terms of timestamps and values). Hence an abstract event stream consists of an event stream over a data abstraction and an associated set of known timestamps:

**Definition 1 (Abstract Event Stream).** *Given a time domain  $\mathbb{T}$ , an abstract event stream over a data domain  $\mathbb{D}$  is a pair  $(s, \Delta)$  with  $s \in \mathcal{S}_{\mathbb{D}}^\#$  and  $\Delta \subseteq \mathbb{T}$  such that  $\Delta$  can be represented as union of intervals whose (inclusive or exclusive) boundaries are indicated by events in an event stream. Further, we require  $s(t) \neq \perp \Rightarrow t \in \Delta$ . The set of all abstract event streams over  $\mathbb{D}$  is denoted as  $\mathcal{P}_{\mathbb{D}}$ . The concretisation function  $\gamma : \mathcal{P}_{\mathbb{D}} \rightarrow 2^{\mathcal{S}^{\mathbb{D}}}$  is defined as*

$$\gamma((s, \Delta)) = \{s' \mid \forall t \in \text{ticks}(s) s(t) \in \gamma(s'(t)) \wedge \forall t \in \Delta \setminus \text{ticks}(s) s(t) = s'(t)\}$$

If the data abstraction is defined in terms of a Galois Connection a refinement ordering and abstraction function can be obtained. The refinement ordering  $(\mathcal{P}_{\mathbb{D}}, \preceq)$  is defined as  $(s_1, \Delta_1) \preceq (s_2, \Delta_2)$  iff  $\Delta_1 \supseteq \Delta_2$  and  $\forall t \in \text{ticks}(s_2) s_1(t) \preceq s_2(t) \wedge \forall t \in \Delta_2 \setminus \text{ticks}(s_2) s_1(t) = s_2(t)$ . The abstraction function  $\alpha : 2^{\mathcal{S}^{\mathbb{D}}} \rightarrow \mathcal{P}_{\mathbb{D}}$  is defined as  $\alpha(S) = \sup\{(s, \mathbb{T}) \mid s \in S\}$ . Note, if the data abstraction is defined in terms of a Galois Connection,  $(\alpha, \gamma)$  is a Galois Connection between  $2^{\mathcal{S}^{\mathbb{D}}}$  and  $\mathcal{P}_{\mathbb{D}}$ .

An abstract event stream  $s = (s', \Delta) \in \mathcal{P}_{\mathbb{D}}$  can also be seen as a function  $s : \mathbb{T} \rightarrow \mathbb{D}^\# \cup \{?, \perp, \smile\}$  with  $s(t) = s'(t)$  if  $t \in \Delta$  and  $s(t) = \smile$  otherwise. A particular point  $t$  of an abstract event stream  $s$  can be either (a) directly at an event ( $s(t) \in \mathbb{D}$ ), (b) in a gap ( $s(t) = \smile$ ), (c) in a gapless segment without an event at  $t$  ( $s(t) = \perp$ ), or (d) after the known end of the stream ( $s(t) = ?$ ).

We denote  $\mathbb{D}^\#_\perp \stackrel{\text{def}}{=} \mathbb{D}^\# \cup \{\perp, \smile\}$ . If  $\mathbb{D}^\#$  is a data abstraction of a data domain  $\mathbb{D}$  with an associated concretisation function  $\gamma$ , then  $\mathbb{D}^\#_\perp$  is a data abstraction of  $\mathbb{D}_\perp$  with an associated concretisation function  $\gamma_\perp : \mathbb{D}^\#_\perp \rightarrow 2^{\mathbb{D} \cup \{\perp\}}$  with

$$\gamma_\perp(d) = \begin{cases} \perp & \text{if } d = \perp \\ \mathbb{D} \cup \{\perp\} & \text{if } d = \smile \\ \gamma(d) & \text{if } d \in \mathbb{D}^\# \end{cases}$$

The above diagram shows a possible data abstraction  $\mathbb{B}^\#$  of  $\mathbb{B}$  and the corresponding data abstraction  $\mathbb{B}^\#_\perp$ . Using the functional representation of an abstract event stream we can now define the abstract counterparts of the TeSSLa operators:


- ▶  $\mathbf{nil}^\# = (\infty, \top) \in \mathcal{P}_\emptyset$  is the empty abstract stream without any gaps.
- ▶  $\mathbf{unit}^\# = (0 \square \infty, \top) \in \mathcal{P}_\cup$  is the abstract stream without any gaps and a single event at timestamp 0.
- ▶  $\mathbf{time}^\# : \mathcal{P}_\mathbb{D} \rightarrow \mathcal{P}_\top$ ,  $\mathbf{time}^\#(s) := z$  is equivalent to its concrete counterpart; only the data domain is extended:  $z(t) = t$  if  $t \in \mathit{ticks}(s)$  and  $z(t) = s(t)$  otherwise.
- ▶  $\mathbf{lift}^\# : (\mathbb{D}_1^\# \times \dots \times \mathbb{D}_n^\# \rightarrow \mathbb{D}_\perp^\#) \rightarrow (\mathcal{P}_{\mathbb{D}_1} \times \dots \times \mathcal{P}_{\mathbb{D}_n} \rightarrow \mathcal{P}_\mathbb{D})$ ,  $\mathbf{lift}^\#(f^\#)(s_1, \dots, s_n) := z$  can be defined similarly to its concrete counterpart, because the abstract function  $f^\#$  takes care of the gaps:

$$z(t) = \begin{cases} f^\#(s_1(t), \dots, s_n(t)) & \text{if } s_1 \neq ?, \dots, s_n \neq ? \\ ? & \text{otherwise} \end{cases}$$

The operator  $\mathbf{lift}^\#$  is restricted to those functions  $f^\#$  that are an abstraction of functions  $f$  that can be used in  $\mathbf{lift}$ , that is,  $f(\perp, \dots, \perp) = \perp$ . Using the abstract lift we can derive the abstract counterparts of  $\mathbf{const}$  and  $\mathbf{merge}$ :

- ▷  $\mathbf{const}^\#(c)(a) := \mathbf{lift}^\#(f_c)(a)$  with  $f_c(d) := c$  if  $d \neq \smile$  and  $f_c(\smile) := \smile$  otherwise maps all events' values to a constant while preserving the gaps. Using  $\mathbf{const}^\#$  we can define constant signals without any gaps, e.g.,  $\mathbf{true}^\# := \mathbf{const}^\#(\mathbf{true})(\mathbf{unit}^\#)$  or  $\mathbf{zero}^\# := \mathbf{const}^\#(0)(\mathbf{unit}^\#)$ .
- ▷  $\mathbf{merge}^\#(x, y) := \mathbf{lift}^\#(f)(x, y)$  with  $f(a \notin \{\smile, \perp\}, b) = a$ ,  $f(\perp, b) = b$ ,  $f(\smile, b \in \{\smile, \perp\}) = \smile$ , and  $f(\smile, b \notin \{\smile, \perp\}) = \top$ .

The diagram on the right shows an example trace merging the events of the streams  $x$  and  $y$ . The symbol  $\circ$  indicates a point-wise gap. Note how an event on the first stream takes precedence over a gap on the second stream, but not the other way round, similarly to how events from the first stream are prioritized if both streams have an event at the same timestamp.



- ▶  $\mathbf{last}^\# : \mathcal{P}_{\mathbb{D}_1} \times \mathcal{P}_{\mathbb{D}_2} \rightarrow \mathcal{P}_{\mathbb{D}_1}$ ,  $\mathbf{last}^\#(v, r) := z$  has three major extensions over its concrete counterpart:


- (1)  $\top$  is added as an output in case an event on  $r$  occurs and there were events on the stream  $v$  of values but all followed by a gap.
- (2)  $\smile$  is outputted for all gaps on the stream  $r$  of trigger events if there have been events on the stream  $v$  of values.
- (3)  $\smile$  can also be output if an event occurs on  $r$  and no event occurred on  $v$  before except for a gap.

The parts similar to the concrete operator are typeset in gray:

$$z(t) = \begin{cases} d & t \in \mathit{ticks}(r) \wedge \exists_{t' < t} \mathit{isLast}(t, t', v, d) \wedge \forall_{t'' | t' < t'' < t} v(t'') \neq \smile \\ \top & t \in \mathit{ticks}(r) \wedge \exists_{t' < t} \mathit{isLast}(t, t', v, \smile) \wedge \exists_{t'' | t' < t'' < t} v(t'') = \smile \\ \perp & r(t) = \perp \wedge \mathit{defined}(z, t) \vee \forall_{t' < t} v(t') = \perp \\ \smile & \mathit{defined}(z, t) \wedge r(t) = \smile \wedge \exists_{t' < t} v(t') \neq \perp \\ \smile & \mathit{defined}(z, t) \wedge t \in \mathit{ticks}(r) \wedge \forall_{t' < t'} \notin \mathit{ticks}(v) \wedge \exists_{t' < t} v(t') = \smile \\ ? & \text{otherwise} \end{cases} \quad (1)$$




The trace diagram on the right shows an example trace covering most edge cases of the abstract last. The output stream  $z$  is a point-wise gap if triggered after initial gaps (3);  $z$  is  $\top$  if triggered after non-initial gaps (1);  $z$  is an event if triggered after a gapless sequence (d); and  $z$  inherits all gaps from the stream of trigger events (2).



We can now combine the **last**<sup>#</sup> and the **lift**<sup>#</sup> operators to realize:

- ▷ *abstract signal lift* for total functions  $f : \mathbb{D} \times \mathbb{D}' \rightarrow \mathbb{D}''$  as  $\text{slift}^\#(f)(x, y) := \text{lift}^\#(g_f)(x', y')$  with  $x' := \text{merge}^\#(x, \text{last}^\#(x, y))$  and  $y' := \text{merge}^\#(y, \text{last}^\#(y, x))$ , as well as  $g_f(a \notin \{\smile, \perp\}, b \notin \{\smile, \perp\}) = f(a, b)$ ,  $g_f(\perp, b) = g_f(a, \perp) = \perp$ ,  $g_f(\smile, \smile) = \smile$ , and  $g_f(\smile, b \notin \{\smile, \perp\}) = g_f(a \notin \{\smile, \perp\}, \smile) = \smile$ .

*Example 2.* By replacing every TeSSLa operator in Example 1 with their abstract counterparts and applying it to the abstract input streams  $\text{values} \in \mathcal{P}_{\mathbb{Z}}$  and  $\text{resets} \in \mathcal{P}_{\mathbb{U}}$ , we derive the abstract stream  $\text{cond} \in \mathcal{P}_{\mathbb{B}}$  and the recursively derived abstract stream  $\text{sum} \in \mathcal{P}_{\mathbb{Z}}$ : After the large gap on  $\text{values}$ , the  $\text{sum}$  stream eventually recovers completely. The first reset after the point-wise gap does not lead to full recovery, because at that point the last event on  $\text{values}$  cannot be accessed, because of the prior gap. The next reset falls into the gap, so again  $\text{cond}$  cannot be evaluated. In a similar fashion one can define an abstract **delay**<sup>#</sup> operator as counterpart of the concrete **delay**. See [19] for details.



Following from the definitions of the abstract TeSSLa operators we get:

**Theorem 1.** *Every abstract TeSSLa operator is an abstraction of its concrete counterpart.*

Theorem 1 implies that abstract TeSSLa operators are sound in the following way. Let  $o$  be a concrete TeSSLa operator with the abstract counterpart  $o^\#$  and let  $s \in \mathcal{P}_{\mathbb{D}}$  be an abstract event stream with a concretization function  $\gamma$ . Then,  $o(\gamma(s)) \preceq \gamma(o^\#(s))$ . Since abstract interpretation is compositional we can directly follow from the above theorem:

**Corollary 1.** *If a concrete TeSSLa specification  $\varphi$  is transformed into a specification  $\psi$  by replacing every concrete operator in  $\varphi$  with its abstract counterpart, then  $\psi$  is an abstraction of  $\varphi$ .*

Theorem 1 guarantees that applying abstract TeSSLa operators to the abstract event stream is still sound regarding the underlying set of possible concrete event streams. However, we have established no result so far about the accuracy of the abstract TeSSLa operators. The abstraction returning only the completely unknown stream ( $\Delta = \emptyset$ ) is sound but useless. The following theorem states, that our abstract TeSSLa operators are optimal in terms of accuracy.

Using a perfect abstraction guarantees the abstract TeSSLa operators preserve as much information as can possibly be encoded in the resulting abstract event streams.

**Theorem 2.** *Every abstract TeSSLa operator is a perfect abstraction of its concrete counterpart.*

Given a concrete TeSSLa operator  $o$  and its abstract counterpart  $o^\#$ , and any abstract event stream  $s \in \mathcal{P}_{\mathbb{D}}$  with the Galois Connection  $(\alpha, \gamma)$  between  $2^{\mathcal{S}^{\mathbb{D}}}$  and  $\mathcal{P}_{\mathbb{D}}$  one can show that  $o^\#(s) = \alpha(o(\gamma(s)))$ . Applying the abstract operator on the abstract event stream is as good as applying the concrete operator on every possible event stream represented by the abstract event stream. Thus  $o^\#$  is a perfect abstraction of  $o$ . (The detailed proof can be found in [19].) Note that we assume that  $f^\#$  is a perfect abstraction of  $f$  to conclude that  $\mathbf{lift}^\#(f^\#)$  is a perfect abstraction of  $\mathbf{lift}(f)$ .

In Corollary 1 we have shown that a specification  $\psi$  (generated by replacing the concrete TeSSLa operator in  $\varphi$  with their abstract counterparts) is an abstraction of  $\varphi$ . Note that  $\psi$  is in general not a perfect abstraction of  $\varphi$ . We study some special cases of perfect abstractions of compositional specifications in Sect. 4.

The next result states that the abstract operators can be defined in terms of concrete TeSSLa operators. Realizing the abstract operators in TeSSLa does not require an enhancement in the expressivity of TeSSLa.

**Theorem 3.** *The semantics of the abstract TeSSLa operators can be encoded in TeSSLa using only the concrete operators.*


*Proof.* One can observe that the abstract TeSSLa operators are monotone and future independent (the output stream up to  $t$  only depends on the input streams up to  $t$ .) As shown in [7], TeSSLa can express every such function.  $\square$

### 3.1 Fixpoint Calculations Ensuring Well-Formedness

A concrete TeSSLa specification consists of stream variables and possibly recursive equations applying concrete TeSSLa operators to the stream variables. Theorem 1 and Corollary 1 guarantee that a concrete TeSSLa specification can be transformed into an abstract TeSSLa specification, which is able to handle gaps in the input streams. Additionally, Theorem 3 states that the abstract TeSSLa operators can be implemented using concrete TeSSLa operators. Combining these two results, one can transform a given concrete specification  $\varphi$  into a corresponding specification  $\psi$ , which realizes the abstract TeSSLa semantics of the operators in  $\varphi$ , but only uses concrete TeSSLa operators.

However, using the realization of the abstract TeSSLa operators in TeSSLa adds additional cyclic dependencies in  $\psi$  between the stream variables. A TeSSLa specification is well-formed if every cycle of its dependency graph contains at least one edge guarded by a last (or a delay) operator, which is required to guarantee the existence of a unique fixed-point and hence computability (see [7]).

Consider the trace diagram on the right showing  $\mathbf{last}^\#(v, r)$ . If  $v$  is used in a recursive manner, i.e.,  $v$  is defined in terms of  $\mathbf{last}^\#(v, r)$ , then the first event on  $v$  could start a gap on  $\mathbf{last}^\#(v, r)$  that could start a gap on  $v$  at the same timestamp. As a result  $v$  has an unguarded cyclic dependency and hence the specification is not well-formed. To overcome this issue one can split up the value and gap calculation sequentially, reintroducing guards in the cyclic dependency:



**Definition 2 (Unrolled Abstract Last).** We define two variants of the abstract last,  $\mathbf{last}^\#_\perp$  and  $\mathbf{last}^\#_\smile$  as follows. Let  $z = \mathbf{last}^\#(v, r)$ , then  $\mathbf{last}^\#_\perp(v, r) := z_\perp$  and  $\mathbf{last}^\#_\smile(v, r, d) := z_\smile$ .

$$z_\perp(t) = \begin{cases} z(t) & \text{if } z(t) \neq \smile \\ \perp & \text{otherwise} \end{cases} \quad z_\smile(t) = \begin{cases} d(t) & \text{if } t \in \text{ticks}(d) \\ \smile & \text{if } t \notin \text{ticks}(d) \wedge z(t) = \smile \\ \perp & \text{otherwise} \end{cases}$$

Function  $\mathbf{last}^\#_\perp$  executes a normal calculation of the events, in the same way an abstract last would do, but neglecting gaps and outputting  $\perp$  as long as there is no event. Function  $\mathbf{last}^\#_\smile$  takes a third input stream and outputs its events directly, but calculates gaps correctly as  $\mathbf{last}^\#$  would do.

Since the trigger input of a  $\mathbf{last}$  operator cannot be recursive in a well-formed specification, a recursive equation using one last has the form  $x = \mathbf{last}^\#(v, r)$  and  $v = f(x, \mathbf{c})$ , where  $\mathbf{c}$  is a vector of streams not involved in the recursion and  $f$  does not introduce further last (or delay) operators. Now, this equation system can be rewritten in the following equivalent form:

$$x' = \mathbf{last}^\#_\perp(v, r) \quad v' = f(x', \mathbf{c}) \quad x = \mathbf{last}^\#_\smile(v', r, x') \quad v = f(x, \mathbf{c})$$

This pattern can be repeated if multiple recursive abstract lasts are used and can also be applied in a similar fashion to mutually recursive equations and the delay operator.

## 4 Perfection of Compositional Specifications


A concrete TeSSLa specification  $\varphi$  can be transformed into an abstract TeSSLa specification  $\psi$  by replacing the concrete operators with their abstract counterparts. For two functions  $f$  and  $g$  with corresponding abstractions  $f^\#$  and  $g^\#$  the function composition  $f^\# \circ g^\#$  is an abstraction of  $f \circ g$ . Unfortunately, even if  $f^\#$  and  $g^\#$  are perfect abstractions,  $f^\# \circ g^\#$  is not necessarily a perfect abstraction. Hence,  $\psi$  needs not be a perfect abstraction of  $\varphi$ . In this section we discuss the perfection of two common compositional TeSSLa operators: (1) the  $\mathbf{slift}^\#$  defined in Sect. 3 is a composition of  $\mathbf{last}^\#$  in  $\mathbf{lift}^\#$ , which realizes signal semantics; (2)  $\mathbf{last}^\#(\mathbf{time}^\#(v), r)$ , which is a common pattern used when comparing timestamps.

The  $\text{slift}^\#$  is defined as the  $\text{lift}^\#$  applied to the synchronized versions  $x'$  and  $y'$  of the input streams  $x$  and  $y$ . The input stream  $x$  is synchronized with  $y$  by keeping the original events of  $x$  and reproducing the last known value of  $x$  for every timestamp with an event on  $y$ , but not on  $x$ .

**Theorem 4.** *If  $f^\#$  is a perfect abstraction of  $f$  then  $\text{slift}(f^\#)^\#$  is a perfect abstraction of  $\text{slift}(f)$ .*

*Proof.* Since  $\text{slift}^\#$  is defined on abstract event streams we need to consider gaps. The stream  $x'$  does not have any gap or event until the first gap or event on  $x$ . After the first gap or event on  $x$  the synchronized stream  $x'$  contains a gap or event at every timestamp where  $x$  or  $y$  contain a gap or event. Because  $\text{slift}^\#$  is symmetric in terms of the event pattern the same holds for  $y'$ . By definition,  $\text{slift}^\#(f^\#)(x, y) = z$  contains an event or gap iff  $x'$  and  $y'$  contain an event or gap, because  $f$  is a total function. The output stream  $z$  contains an event iff  $x'$  and  $y'$  contain events. The events values are ensured to be as precise as possible, because  $f^\#$  is a perfect abstraction of  $f$ . □

TeSSLa allows arbitrary computations on the timestamps of events using the **time** operator. The specification  $z = \text{time}(v)$  derives a stream  $z$  from  $v$  by replacing all event's values in  $v$  with the event's timestamps. The stream vari-



able  $z$  can now be used in any computation expressible in TeSSLa. Hence, TeSSLa does not distinguish between timestamps and other values, and consequently abstract TeSSLa specifications cannot make use of the monotonicity of time. As an example consider the trace diagram on the right. The stream  $\text{last}^\#(\text{time}^\#(v), r)$  is derived from  $v$  by composing  $\text{time}^\#$  and  $\text{last}^\#$ . Since  $\text{time}^\#$  changes the events values with their timestamps, the  $\text{last}^\#$  does not know any longer that we are interested in the last timestamp of  $v$  and can only produce an event with the value  $\top$  representing all possible values. To overcome this issue we define  $\text{lastTime}(v, r) := \text{last}(\text{time}(v), r)$  and provide a direct abstraction, which allows a special treatment of timestamps.

**Definition 3 (Time Aware Abstract Last).** *Let  $y = \text{last}^\#(\text{time}^\#(v), r)$ , we define  $\text{lastTime}^\# : \mathcal{P}_{\mathbb{D}} \times \mathcal{P}_{\mathbb{D}'} \rightarrow \mathcal{P}_{2^{\mathbb{T}}}$ ,  $\text{lastTime}^\#(v, r) := z$  as  $z(t) = [a, b]$  if  $y(t) = \top$  with  $a = \inf\{t' < t \mid \forall v' < t' < t v(t'') \neq \cup\}$  and  $b = \max\{t' < t \mid t' \in \text{ticks}(v)\}$  and  $z(t) = y(t)$  otherwise.*


Now the following result holds (the proof can be found in [19]).

**Theorem 5.**  $\text{lastTime}^\#$  is a perfect abstraction of  $\text{lastTime}$ .

A similar problem occurs if  $\text{slift}^\#$  is used to compare event's timestamps. In Example 2 the stream  $\text{cond}$  derived by comparing the timestamps of values and resets has two events with the unknown data value  $\top$  because of prior gaps on values. Since the  $\text{slift}^\#$  is defined in terms of  $\text{lift}^\#$  and  $\text{last}^\#$  we can

define the function  $\text{sliftTime}^\#(f^\#)(x, y)$  as an abstraction for the special case  $\text{sliftTime}(f)(x, y) = \text{slift}(f)(\mathbf{time}(x), \mathbf{time}(y))$  by using  $\text{lastTime}^\#$  instead of  $\text{last}^\#$  and ensuring that  $f^\#$  uses interval arithmetics to abstract  $f$ . Note that  $\text{sliftTime}^\#(f^\#)$  is a perfect abstraction of  $\text{sliftTime}(f)$ .

*Example 3.* To illustrate the perfect abstraction  $\text{sliftTime}^\#$  we update the definition of `cond` in Example 2 as follows:  $\text{cond} = \text{sliftTime}(\leq)(\text{resets}, \text{values})$ . The events drawn in red now have concrete values instead of  $\top$  as in Example 2.



### 5 Abstractions for Sliding Windows

In this section we demonstrate how to apply the techniques presented in this paper to specifications with richer data domains. In particular, we show now a TeSSLa specification that uses a queue to compute the average load of a processor in the last five time units. The moving window is realized using a queue storing all events that happened in the time window. The stream load  $\in \mathcal{S}_{\mathbb{R}}$  contains an event every time the input load changes:

$$\begin{aligned} \text{stripped} &= \text{slift}(\text{remOlder}_5)(\mathbf{time}(\text{load}), \text{merge}(\text{last}(\text{queue}, \text{load}), \langle \rangle)) \\ \text{queue} &= \text{lift}(\text{enq})(\mathbf{time}(\text{load}), \text{load}, \text{stripped}) \\ \text{avg} &= \text{lift}(\text{int})(\text{queue}, \mathbf{time}(\text{load})) \\ \text{int}(q, u) &= \text{fold}(f, q, 0, u) \quad f(a, b, v, \text{acc}) = \text{acc} + v \cdot (b - a) / 5 \end{aligned}$$

The queue operation `enq` adds elements to the queue, while `remOlder5` removes elements with a timestamp older than five time units. The function `int` accumulates all values in the queue weighted by the length of the corresponding signal piece. The queue operation `fold` is used to fold the function  $f$  over all elements from the queue with the initial accumulator 0 until the timestamp  $u$ . Hence  $f$  is called for every element in the queue with the timestamps  $a$  and  $b$ , the element’s value  $v$  and the accumulator. Consequently, the specification adds elements to the queue, removes the expired elements and accumulates the remaining values. Using our approach we replace every operator with its abstract counterpart and represent abstract queues appropriately such that also queues with partly unknown entries can be modeled. By doing this we obtain a specification that is able to handle gaps in the input stream, as illustrated in Fig. 2.

We can extend the example such that the queue only holds a predefined maximum number of events (to guarantee a finite state implementation). When removing events we represent these as unknown entries in the abstract queues. The abstract  $\text{fold}^\#$  is capable of computing the interval of possible average loads for queues with unknown elements anyhow.

Note that the average load is only updated for every event on the input stream. Using a delay operator, we can set a timeout whenever an element leaves




Fig. 2. Example trace of the abstract queue specification.

the sliding window in the abstract setting. The element is removed from the queue at that timeout and the new value of the queue is updated with the remaining elements. Formal definitions of the queue functions as well as the complete specifications are available online<sup>1</sup>.

## 6 Implementation and Empirical Evaluation

As discussed in Sect. 3.1 the abstract TeSSLa operators can be implemented using only the existing concrete TeSSLa operators. We implemented the abstract TeSSLa operators as macros specified in the TeSSLa language itself such that the existing TeSSLa engine presented in [7] can handle abstract TeSSLa specifications. An abstract event stream  $(s, \Delta) \in \mathcal{P}_{\mathbb{D}}$  can be represented as two TeSSLa streams  $s \in \mathcal{S}_{\mathbb{D}^{\#}}$  and  $s_d \in \mathcal{S}_X$ , where  $X$  contains the following six possible changes of  $\Delta$ : inclusive start, exclusive start, inclusive end, exclusive end, point-wise gap and point-wise event in a gap. Using this encoding it is sufficient to look up the latest  $s_d(t')$  with  $t' \leq t$  to decide whether  $t \in \Delta$ . While this encoding already allows a decent implementation of abstract TeSSLa we go one step further and assume a finite time domain with a limited precision, e.g., 64 bit integers or floats. Under this assumption there is always a known smallest relative timestamp  $\varepsilon$ . Hence, we can use the encoding  $s_d \in \mathcal{S}_{\mathbb{B}}$  where an event  $s_d(t) = \text{true}$  encodes a start inclusive and  $s_d(t) = \text{false}$  an end exclusive. This encoding captures the most common cases and simplifies the implementation of union and intersection on  $\Delta$  enormously since they can now be realized as  $\text{slift}(\vee)$  and  $\text{slift}(\wedge)$ , resp. The other possible switches at timestamp  $t$  can be represented as follows:  $s_d(t + \varepsilon) = \text{true}$  encodes an exclusive start,  $s_d(t + \varepsilon) = \text{false}$  encodes an inclusive end,  $s_d(t) = \text{true}$  and  $s_d(t + \varepsilon) = \text{false}$  encodes a point-wise event in a gap, and  $s_d(t) = \text{false}$  and  $s_d(t + \varepsilon) = \text{true}$  encodes a point-wise gap. Using this encoding the abstract TeSSLa operators do not need to handle these additional cases explicitly.

Furthermore, assuming the smallest relative timestamp  $\varepsilon$ , we can avoid the need to perform the unrolling defined in Definition 2 by delaying the second part of the computation to the next possible timestamp  $t + \varepsilon$ .


<sup>1</sup> <http://tessla-a.isp.uni-luebeck.de/>.

As a final efficiency improvement we simplified **last**<sup>#</sup> before the first event on the stream of values, which are not relevant in practice. The abstract operator and hence abstract specifications are of course still a sound abstraction of their concrete counterparts, but due to over-abstractions no longer a perfect one during this initial event-less phase of the stream of values.


The implementation in form of a macro library for the existing TeSSLa engine is available together with all the examples and scripts used in the following empirical evaluation and can be experimented with in a web IDE (see Footnote 1).

In the following empirical evaluation we measure the accuracy of the abstractions presented in this paper. An abstract event stream represents input data with some sequences of data loss, where we do not know if any events might have been occurred or what their values have been. Applying an abstract TeSSLa specification to such an input stream takes these gaps into account and provides output streams that in turn contain sequences of gaps and sequences containing concrete events. To evaluate the accuracy of this procedure we compare the output of an abstract TeSSLa specification with the best possible output.

Let  $r \in \mathcal{P}_{\mathbb{D}}$  be an abstract event stream. We obtain the set  $R$  of all possible input streams containing all possible variants that might have happened during gaps in  $r$  by applying the concretization function  $\gamma$  on the abstract input stream. Now we can apply the concrete TeSSLa specification  $\varphi$  to all streams in  $R$  and get the set  $S$  of concrete output streams. On the other hand we apply the abstract TeSSLa specification  $\varphi^\#$  directly to  $r$  and get the abstract output stream  $s$ . Now  $S$  is the set of all possible output streams and  $\gamma(s)$  is the set of output streams defined by the abstract TeSSLa specification. The diagram on the right depicts this comparison process.



To compare  $\gamma(s)$  and  $S$  in a quantitative way we define the *ignorance measure*  $\iota : 2^{S_{\mathbb{D}}} \rightarrow \mathbb{I} = [0, 1]$  scoring the ambiguity of such a set of streams, i.e., how similar the different streams in the set are. Events in non-synchronized streams might not have corresponding events at the same timestamp on the other streams. Hence we refer to the signal semantics of event streams where the events represent the changes of a piece-wise constant signal. As depicted on the right with three event streams over the finite data domain  $\{0, 1, 2\}$ , we score timestamps based on how many event streams have the same value with respect to the signal semantics at that timestamp. These scores are then integrated and normalized throughout the length of the streams. See [19] for the technical details. Using this ignorance measure we can now compute the optimal ignorance  $i := \iota(S) \in \mathbb{I}$  and the ignorance  $k := \iota(\gamma(s)) \in \mathbb{I}$  of the streams produced by the abstract TeSSLa specification.



For the evaluation we took several example specifications and corresponding input traces representing different use-cases of TeSSLa and compared the opti-

mal ignorance with the ignorance of abstract TeSSLa. Note that computing the optimal ignorance requires to derive all possible variants of events that might have happened during gaps, which are in general infinitely many and in the special case of only point-wise gaps still exponentially many. Hence this can only be done on rather short traces with only a few point-wise gaps. As a measure for the overhead imposed by using the abstraction compared to the concrete TeSSLa specification we use the computation depth, i.e., the depth of the dependency graph of the computation nodes of the specifications. While runtimes are highly depending on implementation details of the used TeSSLa engines, the computation depth is a good indicator for the computational overhead in terms of how many concrete TeSSLa operators are needed to realize the abstract TeSSLa specification. Figure 3 shows the empirical results.




Fig. 3. Empirical results.

The first three examples represent the class of common, simple TeSSLa specifications without complex interdependencies and no generation of additional events with **delay**: Reset-count counts between reset events; reset-sum sums up events between reset events; and filter-example filters events occurring in a certain timing-pattern. For these common specifications the overhead is small and the abstraction is perfectly accurate. The burst example checks if events appear according to a complex pattern. In the abstraction we loose accuracy because the starting point of a burst is not accessible by **last**<sup>#</sup> after a gap. A similar problem occurs in the queue example where we use a complex data domain to develop a queue along an event stream. If **last**<sup>#</sup> produces  $\top$  after a gap all information about the queue before the gap is lost. For variable-period the abstraction is not perfectly accurate, because the **delay** is used to generate events periodically depending on an external input. This gets even worse for the self-updating queue where complex computations are performed depending on events generated by a **delay**. Surprisingly, the finite-queue is again perfectly accurate, because the size of the queue is limited in a way that eliminates the inaccuracy of the abstraction in this particular example.



## 7 Conclusion

By replacing the basic operators of TeSSLa with abstract counterparts, we obtained a framework where properties and analyses can be specified with respect to complete traces and automatically evaluated for partially known traces. We have shown that these abstract operators can be encoded in TeSSLa, allowing existing evaluation engines to be reused. This is particularly useful as TeSSLa comprises a very small core language suitable for implementation in soft- as well as hardware. Using the example of sliding windows, we demonstrated how complex data structures like queues can be abstracted. Using finite abstractions, our approach even facilitates using complex data structures when only limited memory is available. Evaluating the abstract specification typically only increases the computational cost by a constant factor. In particular, if a concrete specification can be monitored in linear time (in the size of the trace) its abstract counterpart can be as well. Finally, we illustrated the practical feasibility of our approach by an empirical evaluation using the freely available TeSSLa engine.

## References

1. Alur, R., Fisman, D., Raghothaman, M.: Regular programming for quantitative properties of data streams. In: Thiemann, P. (ed.) ESOP 2016. LNCS, vol. 9632, pp. 15–40. Springer, Heidelberg (2016). [https://doi.org/10.1007/978-3-662-49498-1\\_2](https://doi.org/10.1007/978-3-662-49498-1_2)
2. Asarin, E., Caspi, P., Maler, O.: Timed regular expressions. *J. ACM* **49**(2), 172–206 (2002)
3. Basin, D.A., Klaedtke, F., Zalinescu, E.: Failure-aware runtime verification of distributed systems. In: FSTTCS. LIPIcs, vol. 45, pp. 590–603. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2015)
4. Bauer, A., Leucker, M., Schallhart, C.: Runtime verification for LTL and TLTL. *ACM T. Softw. Eng. Meth.* **20**(4), 14 (2011)
5. Berry, G.: The foundations of Esterel. In: Plotkin, G., Stirling, C., Tofte, M. (eds.) *Proof, Language, and Interaction: Essays in Honour of Robin Milner*, pp. 425–454. MIT Press, Cambridge (2000)
6. Bozzelli, L., Sánchez, C.: Foundations of Boolean stream runtime verification. In: Bonakdarpour, B., Smolka, S.A. (eds.) RV 2014. LNCS, vol. 8734, pp. 64–79. Springer, Cham (2014). [https://doi.org/10.1007/978-3-319-11164-3\\_6](https://doi.org/10.1007/978-3-319-11164-3_6)
7. Convent, L., Hungerecker, S., Leucker, M., Scheffel, T., Schmitz, M., Thoma, D.: TeSSLa: temporal stream-based specification language. In: Massoni, T., Mousavi, M.R. (eds.) SBMF 2018. LNCS, vol. 11254, pp. 144–162. Springer, Cham (2018). [https://doi.org/10.1007/978-3-030-03044-5\\_10](https://doi.org/10.1007/978-3-030-03044-5_10)
8. Cousot, P., Cousot, R.: Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In: POPL, pp. 238–252. ACM Press (1977)
9. Cousot, P., Cousot, R.: Abstract interpretation frameworks. *J. Log. Comput.* **2**(4), 511–547 (1992)
10. D’Angelo, B., et al.: LOLA: runtime monitoring of synchronous systems. In: TIME, pp. 166–174. IEEE (2005)

11. Eisner, C., Fisman, D., Havlicek, J., Lustig, Y., McIsaac, A., Van Campenhout, D.: Reasoning with temporal logic on truncated paths. In: Hunt, W.A., Somenzi, F. (eds.) CAV 2003. LNCS, vol. 2725, pp. 27–39. Springer, Heidelberg (2003). [https://doi.org/10.1007/978-3-540-45069-6\\_3](https://doi.org/10.1007/978-3-540-45069-6_3)
12. Eliot, C., Hudak, P.: Functional reactive animation. In: Proceedings of ICFP '07, pp. 163–173. ACM (1997)
13. Faymonville, P., Finkbeiner, B., Schirmer, S., Torfah, H.: A stream-based specification language for network monitoring. In: Falcone, Y., Sánchez, C. (eds.) RV 2016. LNCS, vol. 10012, pp. 152–168. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-46982-9\\_10](https://doi.org/10.1007/978-3-319-46982-9_10)
14. Faymonville, P., et al.: StreamLAB: stream-based monitoring of cyber-physical systems. In: Dillig, I., Tasiran, S. (eds.) CAV 2019. LNCS, vol. 11561, pp. 421–431. Springer, Cham (2019). [https://doi.org/10.1007/978-3-030-25540-4\\_24](https://doi.org/10.1007/978-3-030-25540-4_24)
15. Gautier, T., Le Guernic, P., Besnard, L.: SIGNAL: a declarative language for synchronous programming of real-time systems. In: Kahn, G. (ed.) FPCA 1987. LNCS, vol. 274, pp. 257–277. Springer, Heidelberg (1987). [https://doi.org/10.1007/3-540-18317-5\\_15](https://doi.org/10.1007/3-540-18317-5_15)
16. Gorostiaga, F., Sánchez, C.: Striver: stream runtime verification for real-time event-streams. In: Colombo, C., Leucker, M. (eds.) RV 2018. LNCS, vol. 11237, pp. 282–298. Springer, Cham (2018). [https://doi.org/10.1007/978-3-030-03769-7\\_16](https://doi.org/10.1007/978-3-030-03769-7_16)
17. Halbwachs, N., Caspi, P., Pilaud, D., Plaice, J.: LUSTRE: a declarative language for programming synchronous systems. In: Proceedings of POPL 1987, pp. 178–188. ACM Press (1987)
18. Havelund, K., Roşu, G.: Synthesizing monitors for safety properties. In: Katoen, J.-P., Stevens, P. (eds.) TACAS 2002. LNCS, vol. 2280, pp. 342–356. Springer, Heidelberg (2002). [https://doi.org/10.1007/3-540-46002-0\\_24](https://doi.org/10.1007/3-540-46002-0_24)
19. Leucker, M., Sanchez, C., Scheffel, T., Schmitz, M., Thoma, D.: Runtime verification for timed event streams with partial information. [arXiv:1907.07761](https://arxiv.org/abs/1907.07761) (2019). <https://arxiv.org/abs/1907.07761>
20. Maler, O., Nickovic, D.: Monitoring temporal properties of continuous signals. In: Lakhnech, Y., Yovine, S. (eds.) FORMATS/FTRTFT -2004. LNCS, vol. 3253, pp. 152–166. Springer, Heidelberg (2004). [https://doi.org/10.1007/978-3-540-30206-3\\_12](https://doi.org/10.1007/978-3-540-30206-3_12)
21. Manna, Z., Pnueli, A.: Temporal Verification of Reactive Systems: Safety. Springer, New York (1995). <https://doi.org/10.1007/978-1-4612-4222-2>
22. Sen, K., Roşu, G.: Generating optimal monitors for extended regular expressions. ENTCS **89**(2), 226–245 (2003)
23. Stoller, S.D., et al.: Runtime verification with state estimation. In: Khurshid, S., Sen, K. (eds.) RV 2011. LNCS, vol. 7186, pp. 193–207. Springer, Heidelberg (2012). [https://doi.org/10.1007/978-3-642-29860-8\\_15](https://doi.org/10.1007/978-3-642-29860-8_15)
24. Wang, S., Ayoub, A., Sokolsky, O., Lee, I.: Runtime verification of traces under recording uncertainty. In: Khurshid, S., Sen, K. (eds.) RV 2011. LNCS, vol. 7186, pp. 442–456. Springer, Heidelberg (2012). [https://doi.org/10.1007/978-3-642-29860-8\\_35](https://doi.org/10.1007/978-3-642-29860-8_35)

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

