

Striver: Stream Runtime Verification for Real-Time Event-Streams ^{*}

Felipe Gorostiaga and César Sánchez

IMDEA Software Institute, Spain
{felipe.gorostiaga,cesar.sanchez}@imdea.org

Abstract. We study the problem of monitoring rich properties of real-time event streams, and propose a solution based on Stream Runtime Verification (SRV), where observations are described as output streams of data computed from input streams of data. SRV allows a clean separation between the temporal dependencies among incoming events, and the concrete operations that are performed during the monitoring. SRV specification languages typically assume that all streams share a global synchronous clock and input events arrive in a synchronous manner. In this paper we generalize the time assumption to cover real-time event streams, but keep the essential explicit time dependencies present in synchronous SRV languages. We introduce **Striver**, which shares with SRV the simplicity, and the separation between the timing reasoning and the data domain. **Striver** is a general language that allows to express other real-time monitoring languages. We show in this paper translations from other formalisms for (piece-wise constant) signals and timed event streams. Finally, we report an empirical evaluation of an implementation of **Striver**.

1 Introduction

Runtime verification (RV) is a lightweight formal method that studies the problem of whether a single trace from the system under analysis satisfies a formal specification. From the point of view of coverage, static verification must consider all possible executions of the system while RV only considers the traces observed. In this manner, RV sacrifices completeness but offers a readily applicable formal method that can be combined with testing or debugging. See [16, 20] for surveys on RV, and the recent book [2]. Early specification languages proposed in RV were based on temporal logics [17, 12, 7], regular expressions [25], timed regular expressions [3], rules [5], or rewriting [24].

Stream runtime verification (SRV), pioneered by Lola [11] defines monitors by declaring the dependencies between output streams (results) and input streams (observations). The main idea of SRV is that the same sequence of operations

^{*} This research has been partially supported by: the EU H2020 project Elastest (num. 731535), by the Spanish MINECO Project “RISCO (TIN2015-71819-P)” and by the EU ICT COST Action IC1402 ARVI (*Runtime Verification beyond Monitoring*).

performed during the monitoring of a temporal logic formula can be followed to compute statistics of the input trace, if the data type and the operations are changed. The generalization of the outcome of the monitoring process to richer verdict values brings runtime verification closer to monitoring and data stream-processing. See [22, 15, 8] for further works on SRV. Temporal testers [23] were later proposed as a monitoring technique for LTL based on Boolean streams. SRV was initially conceived for monitoring synchronous systems, where all observations proceed in cycles. In this paper we present a specification formalism for timed asynchronous observations, where streams are sequences of timed events, not necessarily happening at the same time in all input streams, but where all time-stamps are totally ordered according to a global clock (following the timed asynchronous model of distributed systems [10]). The formalism that we propose in this paper targets the outline, non-intrusive monitoring (see [18] for definitions), where the model of time is that of timed asynchronous distributed systems. Our target application is the monitoring and testing of cloud systems and multi-core hardware monitoring, where this assumption is reasonable.

Related work The work [19] presents an asynchronous evaluation engine for a simple event stream language for timed events, based on a collection of language constructs that compute aggregations. This language does not allow explicit time references and offsets. Moreover, recursion is not permitted and all recursive computations are encapsulated implicitly in the language constructs. A successor work of [19] is TeSSLa [9] which allows recursion and offers a smaller collection of language constructs. Still, TeSSLa precludes explicit offset dependencies, and the target application domain is hardware based monitoring. We sketched that **Striver** subsumes TeSSLa. Another similar work is RTLola [14], which also aims to extend SRV from the synchronous domain to timed streams. However, in RTLola defined streams are computed at predefined periodic instants of time, collecting aggregations between these predefined instants using language constructs. In this manner, the output streams in RTLola are *isochronous*¹, while in **Striver** defined streams are computed at the specific real-time instants where they are required, resulting in a completely *asynchronous* SRV system (in the sense that streams can tick at arbitrary time points). **Striver** can be used as a low level language to compile TeSSLa, RTLola and similar specifications.

The rest of the paper is organized as follows. Section 2 describes the **Striver** specification language. Section 3 presents a trace-length independent online algorithm. Section 4 shows some extensions of **Striver**. Section 5 reports on an empirical evaluation and Section 6 concludes the paper.

2 The **Striver** Specification Language

In this section we introduce **Striver**, a specification language that allows defining efficiently monitorable specifications [11], those for which all streams can be

¹ We borrow this term from telecommunications and signal processing where an isochronous signal is one in which events happen at regular intervals.

resolved immediately. We show in Section 3 an online monitoring algorithm and prove that this algorithm is also trace length independent.

2.1 Preliminaries

The main idea behind SRV is to separate two concerns: the temporal dependencies and the data manipulated, for which we use data domains.

Data Domains. We use many-sorted first order logic to describe data domains. A simple theory, *Booleans*, has only one sort², *Bool*, two constants **true** and **false**, binary functions \wedge and \vee , unary function \neg , etc. A more sophisticated signature is *Naturals* that consists of two sorts (*Nat* and *Bool*), with constant symbols $0, 1, 2, \dots$ of sort *Nat*, binary symbols $+, *$, etc (of sort $\text{Nat} \times \text{Nat} \rightarrow \text{Nat}$) as well as predicates $<, \leq$, etc of sort $\text{Nat} \times \text{Nat} \rightarrow \text{Bool}$, with their usual interpretation. All theories have equality and are typically (e.g. *Naturals*, *Booleans*, *Queues*, *Stacks*, etc) equipped with a ternary symbol **if · then · else**. In the case of *Naturals*, the **if · then · else** symbol has sort $\text{Bool} \times \text{Nat} \times \text{Nat} \rightarrow \text{Nat}$.

Our theories are interpreted, so each sort S is associated with a domain D_S (a concrete set of values), and each function symbol \mathbf{f} is interpreted as a total computable function f , with the given arity and that produces values of the domain of the result given elements of the arguments' domains. For simplicity, we omit the sort S from D_S .

We will use *stream variables* with an associated sort, but from the point of view of the theories, these stream variables are atoms. As usual, given a set of sorted atoms A and a theory, the set of terms is the smallest set containing A and closed under the use of function symbols in the theory as constructors (respecting sorts).

We consider a special *time* domain \mathbb{T} , whose interpretation is a (possibly infinite, possibly dense) set with a total order and a minimal element 0 , and a binary addition symbol $+$. Examples of time domains are \mathbb{R}_0^+ , \mathbb{Q}_0^+ and \mathbb{N}_0 with their usual order. Given $t_a, t_b \in \mathbb{T}$ we use $[t_a, t_b] = \{t \in \mathbb{T} \mid t_a \leq t \leq t_b\}$, and also (t_a, t_b) , $[t_a, t_b)$ and $(t_a, t_b]$ with the usual meaning. We say that a set of time points $S \subseteq \mathbb{T}$ does not contain bounded infinite subsets, whenever for every $t_a, t_b \in \mathbb{T}$, the set $S \cap [t_a, t_b]$ is finite, in which case we say that S is a non-Zeno set.

We extend every domain D into D^\perp that includes two special fresh symbols \perp_{notick}^D and \perp_{outside}^D . These new symbols allow capturing when a stream does not generate an event, and when the time offset falls off the beginning and the end of the trace.

Streams. Monitors observe sequences of events as inputs, where each event is time-stamped and contains a data value from its domain.

² We use sort and type interchangeably in the rest of the paper.

Definition 1 (Event stream). An event stream of sort D is a partial function $\eta : \mathbb{T} \rightarrow D$ such that $\text{dom}(\eta)$ does not contain bounded infinite subsets, where $\text{dom}(\eta)$ is the subset of \mathbb{T} where η is defined.

The set $\text{dom}(\eta)$ is called the set of *event points* of η . An event stream η can be naturally represented as a *timed word*: $s_\eta = (t_0, \eta(t_0))(t_1, \eta(t_1)) \cdots \in (\text{dom}(\eta) \times D)^*$, or as an ω -timed word $s_\eta = (t_0, \eta(t_0))(t_1, \eta(t_1)) \cdots \in (\text{dom}(\eta) \times D)^\omega$ for infinite streams, such that:

- (1) s_η is ordered by time ($t_i < t_{i+1}$); and
- (2) for every $t_a, t_b \in \mathbb{T}$ the set $\{(t, d) \in s_\eta \mid t \in [t_a, t_b]\}$ is finite.

The set of all event streams over D is denoted by \mathcal{E}_D .

We introduce some notation for event streams. The functions $\text{prev}_<$ and prev_\leq with type $\mathcal{E}_D \times \mathbb{T} \rightarrow \mathbb{T}^\perp$ are defined as follows. Note that the functions can return a value in \mathbb{T}^\perp because sup can return $\perp_{\text{outside}}^\mathbb{T}$ when the stream has no event in the interval provided.

$$\begin{aligned} \text{prev}_<(\sigma, t) &\stackrel{\text{def}}{=} \text{sup}(\text{dom}(\sigma) \cap [0, t)) & \text{sup}(S) &\stackrel{\text{def}}{=} \begin{cases} \max(S) & \text{if } S \neq \emptyset \\ \perp_{\text{outside}}^\mathbb{T} & \text{otherwise} \end{cases} \\ \text{prev}_\leq(\sigma, t) &\stackrel{\text{def}}{=} \text{sup}(\text{dom}(\sigma) \cap [0, t]) \end{aligned}$$

Essentially, given a stream σ and a time instant $t \in \mathbb{T}$, the expression $\text{prev}_<(\sigma, t)$ provides the nearest time instant in the past of t at which σ is defined. Similarly, $\text{prev}_\leq(\sigma, t)$ returns t if $t \in \text{dom}(\sigma)$, otherwise it behaves as $\text{prev}_<$.

Synchronous SRV In synchronous SRV, specifications are given by associating every output stream variable y with a defining equation that, once the input streams are known, associates y to an output stream. For example:

```
define bool always_p := p /\ always_p[-1,true]
define int   count_p := (count_p[-1,0]) + if p then 1 else 0
```

defines two output streams: `always_p`, which calculates whether Boolean input stream p was true at every point in the past (that is, $\Box p$) and `count_p`, which counts the number of times p was true in the past. Offset expressions like `count_p[-1,0]` allow referring to streams in a different position (in this case in the previous position) with a default value when there is no previous position (the beginning of the trace). In this paper we introduce a similar formalism for timed event streams. Our goal is to provide a simple language with few constructs including explicit references to the previous position at which some stream contains an event, contrary to other stream languages like TeSSLa [9] and RTLola [14] which preclude to reason about real-time instants. We say that Striver is an *explicit time* SRV formalism.

2.2 Syntax of Striver

A Striver specification describes the relation between input event-streams and output event-streams, where an input stream is a sequence of observations from the system under analysis.

The key idea in **Striver** is to associate each defined stream variable with:

- a *ticking expression* that captures when the stream may contain an event;
- a *value expression* that defines the value contained in the event.

Note that in synchronous SRV, only a value expression is necessary because every stream has a value at every clock tick.

Formally, a **Striver** specification $\varphi : \langle I, O, V, T \rangle$ consists of input stream variables $I = \{x_1, \dots, x_n\}$, output stream variables $O = \{y_1, \dots, y_m\}$, a collection of ticking expressions $T = \{T_1, \dots, T_m\}$ and a collection of value expressions $V = \{V_1, \dots, V_m\}$. For output variable y , T_y captures when stream y ticks and V_y what the value is when y ticks. All input and output streams are associated with a sort. It is sometimes convenient to partition output streams into proper outputs and intermediate streams, that are introduced only to simplify specifications.

In practice, it is very useful that T_y defines an over-approximation of the set of instants at which y ticks, and then allow the value expression to evaluate to \perp_{notick}^D . The stream associated with y does not contain an event at t if V_y evaluates to \perp_{notick}^D at t , even if t is in T_y . For example, if one wishes y to filter out events from a given stream x it is simple to define in T_y that y ticks whenever x does, and delegate to V_y to decide whether an event is relevant or should be filtered out.

Expressions. We fix a set of stream variables $Z = I \cup O$. Apart from ticking expressions and value expressions, offset expressions (used inside value expressions) allow defining temporal dependencies between ticking instants.

– *Ticking Expressions:*

$$\alpha := \{c\} \mid v.\text{ticks} \mid \alpha \cup \alpha \mid \text{delay } w$$

where $c \in \mathbb{T}$ is a time constant, v is an arbitrary stream variable, and w is a stream variable of type \mathbb{T}_ϵ , and \cup is used for the union of sets of ticks. The type \mathbb{T}_ϵ is defined as $\mathbb{T}_\epsilon = \{t \mid t \geq \epsilon\}$ for a given $\epsilon > 0$. This restriction on the argument of **delay** guarantees that the ticking instants are non-zero if all their inputs are non-zero (see Section 3).

– *Offset Expressions,* which allow fetching previous events from streams:

$$\tau_x ::= x < \sim \tau' \mid x \ll \tau' \quad \tau' ::= \mathbf{t} \mid \tau_z \text{ for } z \in Z$$

Offset expressions have sort \mathbb{T} . Here, \mathbf{t} represents the current value of the clock. The intended meaning of $x \ll \tau'$ is to refer to the previous instant strictly in the past of τ' where x ticks (or \perp_{outside}^D if there is not such an instant). The expression $x < \sim \tau'$ also considers the present as a candidate.

– *Value Expressions,* which give the value of a defined stream at a given ticking point candidate:

$$E ::= d \mid x(\tau_x) \mid \mathbf{f}(E_1, \dots, E_k) \mid \mathbf{t} \mid \tau_x \mid \text{outside}_D \mid \text{notick}_D$$

where d is a constant of type D , $x \in Z$ is a stream variable of type D and f is a function symbol of return type D . Note that in $x(\tau_x)$ the value of stream x is fetched at an offset expression indexed by x , which captures the ticking points of x and guarantees the existence of an event. Expressions \mathbf{t} and τ_x build expressions of sort \mathbb{T} . The two additional constants outside_D and notick_D allow to reason about accessing the end of the streams, or not generating an event at ticking candidate instant.

We also use the following syntactic sugar:

$$\begin{aligned} x(\sim e) &\stackrel{\text{def}}{=} x(x < \sim e) & x(\sim e, d) &\stackrel{\text{def}}{=} \text{if } (x < \sim e) == \text{outside} \text{ then } d \text{ else } x(\sim e) \\ x(< e) &\stackrel{\text{def}}{=} x(x << e) & x(< e, d) &\stackrel{\text{def}}{=} \text{if } (x << e) == \text{outside} \text{ then } d \text{ else } x(< e) \end{aligned}$$

Essentially, $x(\sim \mathbf{t})$ provides the value of x at the previous ticking instant of x (including the present) and $x(< \mathbf{t})$ is similar but not including the present. Also, $x(< \mathbf{t}, d)$ is the analogous to $x[-1, d]$ in synchronous SRV allowing to fetch the value in the previous event in stream x , or d if there is not such previous event.

Example 1. Consider two input event streams: `sale` that represents sales of a certain product, and `arrival` which represents the arrivals to the store:

```
input int sale, int arrival
ticks stock := sale.ticks U arrival.ticks
define int stock := stock(<t,0) +
    (if isticking(arrival) then arrival(~t) else 0) -
    (if isticking(sale ) then sale(~t ) else 0)
```

where `isticking(sale)` is defined as `(sale<~t)==t`. Note that `stock` is defined to tick when either `sale` or `arrival` (or both) tick. \square

Example 2. To illustrate the use of `delay` consider the following specification:

```
ticks clock := {0} U delay clock
define Time_eps clock := 1sec
```

The stream `clock` emits an event every second since time 0. \square

2.3 Semantics

As common in SRV, the semantics is defined denotationally first. This semantics establishes whether a given input and a given output satisfy the specification, which is defined in terms of *valuations*. Given a set of variables Z , a valuation σ is a map that associates every x in Z of sort D with an event stream from \mathcal{E}_D . Given a valuation σ we define the result of evaluating an expression for σ . We define three evaluation maps $\llbracket \cdot \rrbracket_\sigma$, $\llbracket \cdot \rrbracket_\sigma$, $\llbracket \cdot \rrbracket_\sigma$ depending on the type of the expression³:

³ we use colors to better distinguish between semantic maps

- *Ticking Expressions.* The semantic map $\llbracket \cdot \rrbracket_\sigma$ assigns a set of time instants to each ticking expression as follows:

$$\begin{aligned} \llbracket \{c\} \rrbracket_\sigma &\stackrel{\text{def}}{=} \{c\} \\ \llbracket v.\text{ticks} \rrbracket_\sigma &\stackrel{\text{def}}{=} \text{dom}(\sigma_v) \\ \llbracket a \cup b \rrbracket_\sigma &\stackrel{\text{def}}{=} \llbracket a \rrbracket_\sigma \cup \llbracket b \rrbracket_\sigma \\ \llbracket \text{delay}(w) \rrbracket_\sigma &\stackrel{\text{def}}{=} \{t' \mid \text{there is a } t \in \text{dom}(\sigma_w) \text{ such that } t + \sigma_w(t) = t' \\ &\quad \text{and } \text{dom}(\sigma_w) \cap (t, t') = \emptyset \} \end{aligned}$$

- *Offset Expressions.* For offset expressions $\llbracket \cdot \rrbracket_\sigma$ provides, given a time instant t , another time instant:

$$\begin{aligned} \llbracket t \rrbracket_\sigma(t) &\stackrel{\text{def}}{=} t \\ \llbracket x \llcorner e \rrbracket_\sigma(t) &\stackrel{\text{def}}{=} \begin{cases} \perp_{\text{outside}}^\mathbb{T} & \text{if } \llbracket e \rrbracket_\sigma(t) = \perp_{\text{outside}}^\mathbb{T} \\ \text{prev}_{<}(\sigma_x, \llbracket e \rrbracket_\sigma(t)) & \text{otherwise} \end{cases} \\ \llbracket x \llcorner \sim e \rrbracket_\sigma(t) &\stackrel{\text{def}}{=} \begin{cases} \perp_{\text{outside}}^\mathbb{T} & \text{if } \llbracket e \rrbracket_\sigma(t) = \perp_{\text{outside}}^\mathbb{T} \\ \text{prev}_{\leq}(\sigma_x, \llbracket e \rrbracket_\sigma(t)) & \text{otherwise} \end{cases} \end{aligned}$$

- *Value Expressions.* Finally, value expressions are evaluated into event streams of the appropriate type. For a given instant t :

$$\begin{aligned} \llbracket d \rrbracket_\sigma(t) &\stackrel{\text{def}}{=} d \\ \llbracket x(e) \rrbracket_\sigma(t) &\stackrel{\text{def}}{=} \begin{cases} \perp_{\text{outside}}^D & \text{if } \llbracket e \rrbracket_\sigma(t) = \perp_{\text{outside}}^\mathbb{T} \\ v & \text{if } \llbracket e \rrbracket_\sigma(t) = t' \text{ and } \sigma_x(t') = v \end{cases} \\ \llbracket f(E_1, \dots, E_k) \rrbracket_\sigma(t) &\stackrel{\text{def}}{=} f(\llbracket E_1 \rrbracket_\sigma(t), \dots, \llbracket E_k \rrbracket_\sigma(t)) \\ \llbracket t \rrbracket_\sigma(t) &\stackrel{\text{def}}{=} t \\ \llbracket \tau_x \rrbracket_\sigma(t) &\stackrel{\text{def}}{=} \llbracket \tau_x \rrbracket_\sigma(t) \\ \llbracket \text{outside}_D \rrbracket_\sigma(t) &\stackrel{\text{def}}{=} \perp_{\text{outside}}^D \\ \llbracket \text{notick}_D \rrbracket_\sigma(t) &\stackrel{\text{def}}{=} \perp_{\text{notick}}^D \end{aligned}$$

Note that $\llbracket x(e) \rrbracket_\sigma$ includes the possibility that (1) the expression cannot be evaluated because the time instant given by $\llbracket e \rrbracket_\sigma(t)$ is outside the boundaries of domain of the stream and (2) the expression is not defined because the stream does not tick at t . It is easy to see that the cases for $\llbracket x(e) \rrbracket_\sigma$ are exhaustive because $\llbracket e \rrbracket_\sigma(t)$ guarantees that $\sigma_x(t')$ is defined.

For example, consider the following stream $(1.0, 17), (2.5, 21), (3.5, 12)$ for variable `sale` from Example 1. Then

$$\llbracket \text{sale}(\sim t) \rrbracket_\sigma(3.1) = \llbracket \text{sale}(\text{sale} \llcorner \sim t) \rrbracket_\sigma(3.1) = \llbracket \text{sale} \rrbracket_\sigma(2.5) = 21$$

Definition 2 (Evaluation Model). Given a valuation σ of variables $I \cup O$ the evaluation of the equations for stream $y \in O$ is:

$$\llbracket T_y, V_y \rrbracket_\sigma \stackrel{\text{def}}{=} \{(t, d) \mid t \in \llbracket T_y \rrbracket_\sigma \text{ and } d = \llbracket V_y \rrbracket_\sigma(t) \text{ and } d \neq \perp_{\text{notick}}^D\}$$

An evaluation model is a valuation σ such that for every $y \in O$: $\sigma_y = \llbracket T_y, V_y \rrbracket_\sigma$.

The goal of a **Striver** specification is to define a monitor, that intuitively should be a computable function from input streams into output streams. The following definition captures whether a specification indeed corresponds to such a function.

Definition 3 (Well-defined). *A specification φ is well-defined if for all σ_I , there is a unique σ_O , such that $\sigma_I \cup \sigma_O$ is an evaluation model of φ .*

As with synchronous SRV, specifications can be ill-defined. For example, the following specification (`define bool a:= not a`) admits no evaluation model, and (`define bool a:= a`) admits many evaluation models. Additionally, a specification is efficiently monitorable if the output at time t only depends on the input at time t , which enable the incremental computation of the output stream.

Definition 4 (Efficiently monitorable). *A well-defined specification φ is efficiently monitorable whenever for every two input σ_I and σ'_I with evaluation models σ_O and σ'_O , and for every time t , if $\sigma_I|_t = \sigma'_I|_t$ then $\sigma_O|_t = \sigma'_O|_t$.*

2.4 Well-formedness

The condition of well-definedness is a semantic condition, which is not easy to check for a given specification (undecidable for expressive enough domains). We present here a syntactic condition, called well-formedness, that is easy to check on input specifications and guarantees that specifications are well-defined. Most specifications encountered in practice are well-formed.

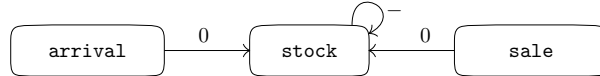
We first define a subset of the offset expressions, called the *Present* subset, as the smallest subset that contains \mathfrak{t} and such that if $e \in \text{Present}$ then $(x <^{\sim} e) \in \text{Present}$. We say that an output stream variable y directly depends on a stream variable x (and we write $x \rightarrow y$) if x appears in T_y or V_y . We say that y has a present direct dependency on x (and write $x \xrightarrow{0} y$) if $x \rightarrow y$ and either

- $x.\text{ticks}$ appears in T_y , or
- $(x <^{\sim} e)$ appears in V_y and $e \in \text{Present}$.

A direct dependency captures whether in order to compute a value of a stream variable y at position t , it is necessary to know the value of stream variable x up to t . If $x \rightarrow y$ but $x \not\xrightarrow{0} y$ we say that y directly depends on x in the past (and we write $x \xrightarrow{-} y$).

Definition 5 (Dependency Graph). *The dependency graph of a specification φ is a graph (V, E) where $V = I \cup O$ and $E = V \times V \times \{\xrightarrow{0}, \xrightarrow{-}\}$.*

The dependency graph of Example 1 is:



The following definition captures whether an output stream variable cannot depend on itself at the present moment.

Definition 6 (Well-Formed Specifications). A specification φ is well-formed if every closed path in its dependency graph contains a past dependency edge $\overrightarrow{}$.

Closed paths in the dependency graph correspond to dependencies between a stream and itself in the specification φ . These closed paths do not create problems if the path corresponds to accessing the strict past of the stream. Note that if one removes $\overrightarrow{}$ edges from the dependency graph of a well-formed specification, the resulting graph is necessarily a DAG. In other words $\xrightarrow{0}^*$ is irreflexive. The following lemma formally captures the information that is sufficient to determine the value of a given stream at a given time instant.

Lemma 1. Let y be an output stream variable of a specification φ , σ, σ' be two evaluation models of φ , such that, for time instant t :

- (i) For every variable x , $\sigma_x(t') = \sigma'_x(t')$ for every $t' < t$, and
- (ii) For every x , such that $x \xrightarrow{0}^* y$, $\sigma_x(t') = \sigma'_x(t')$ for every $t' \leq t$

Then $\sigma_y(t) = \sigma'_y(t)$.

The proof proceeds by structural induction on expressions, with the observation that only values in the past are necessary, as in conditions (i) and (ii). We are now ready to show that well-formed specifications cannot have two different evaluation models.

Theorem 1. Every well-formed Striver specification is well-defined.

The proof proceeds by showing that for well-formed specifications two evaluation models must be equal. This is shown by induction on the events in the traces to prove that the i -th event must be identical. Lemma 1 guarantees that induction can be applied.

3 Operational semantics

The semantics of Striver specifications introduced in the previous section are denotational in the sense that these semantics associate for a given input stream valuation exactly one output stream valuation, but does not provide a procedure to compute the output streams, let alone do it incrementally. We provide in this section an operational semantics that computes the output incrementally. We fix a specification φ with dependency graph G and we let $G^=$ be its pruned dependency graph (obtained from G by removing $\xrightarrow{0}$ edges). We also fix $<$ to be an arbitrary total order between stream variables that is a reverse topological order of $G^=$.

We first present an online monitoring algorithm that stores the full history computed so far for every output stream variable. Later we will provide bounds on the portion of the history that needs to be remembered by the monitor, showing that only a bounded number of events needs to be recorded, and that this bound depends only on the size of the specification (number of streams)

Algorithm 1 MONITOR: Online Monitor

```
1: procedure MONITOR
2:    $H_s \leftarrow \langle \rangle$  for every  $s$ 
3:    $t_q \leftarrow -\infty$ 
4:   loop ▷ Step
5:      $t_q \leftarrow \min_{s \in O} \{t \mid t = \text{VOTE}(H, T_s, t_q)\}$ 
6:     if  $t_q = \infty$  then break
7:     for  $s$  in  $G^\#$  following  $<$  do
8:       if  $t_q \in \llbracket T_s \rrbracket_H$  then
9:          $v \leftarrow \llbracket V_s \rrbracket_H(t_q)$ 
10:        if  $v \neq \perp_{\text{notick}}^D$  then
11:           $H_s \leftarrow H_s ++ (t_q, v)$  ▷ Updates history  $H$ 
12:           $\text{emit}(t_q, v, s)$ 
13:        end for
14:    end loop
```

and not on the length of trace. This modified algorithm is a trace-length independent monitor for efficiently monitorable **Striver** specifications. The algorithm maintains the following state (H, t_q) :

- **History:** H is a finite event stream one for each output stream variable. We use H_y for the event stream prefix for stream variable y .
- **Quiescence time:** t_q is the time up to which all output streams have been computed.

The monitor runs a main loop, calculating first the next relevant time t_q for the monitoring evaluation and then computing all outputs (if any) for time t_q . We show that no event exists in any stream in the interval between two consecutive quiescence time instants. We assume that at time t , the next event for every input stream is available to the monitor, even though knowing that there is no event up-to some t_q is sufficient.

The core observation follows from Lemma 1, which limits the information that is necessary to compute whether stream y at instant t contains an event (t, d) . All this information is contained in H , so we write $\llbracket T_y \rrbracket_H$ and $\llbracket V_y \rrbracket_H$ to remark that only H is needed to compute $\llbracket T_y \rrbracket_\sigma$ and $\llbracket V_y \rrbracket_\sigma$.

The main algorithm, MONITOR, is shown in Algorithm 1. Lines 2 and 3 set the history and initial quiescence time. The main loop continues until no more events can be generated. Line 5 computes the next quiescence time, by taking the minimum instant after the last quiescence time at which some output stream may tick. A stream y “votes” (see Algorithm 2) for the next possible instant at which its ticking equation T_y can possibly contain a value. Consequently, if no input stream votes for an earlier time it is guaranteed that no ticking equation will contain a value t lower than the lowest vote. Note that recursive calls at line 28 terminate because the graph $G^\#$ is acyclic (recall that the specification is well-formed).

Algorithm 2 VOTE: Compute the next ticking instant

```
15: function VOTE( $H, expr, t$ )
16:   switch  $expr$  do
17:     case  $\text{delay}(s)$ 
18:       if  $(t' + v) > t$  (where  $(t', v) = \text{last}(H_s)$ ) then return  $t' + v$ 
19:       else return  $\infty$ 
20:     case  $\{c\}$ 
21:       if  $c > t$  then return  $c$ 
22:       else return  $\infty$ 
23:     case  $a \cup b$ 
24:       return  $\min(\text{VOTE}(H, a, t), \text{VOTE}(H, b, t))$ 
25:     case  $y.\text{ticks}$  with  $y \in O$ 
26:       return  $\text{VOTE}(H, T_y, t)$ 
27:     case  $s.\text{ticks}$  with  $s \in I$ 
28:       return  $\text{succ}_>(\sigma_s, t_q)$ 
```

The algorithm follows a topological order over the $G^=$, so the information about the past required in Lemma 1 is contained in H . The following result shows that, assuming that σ_I is non-zero, the output is also non-zero. Hence, for every instant t , the algorithm eventually reaches $t_q > t$ in a finite number of executions of the main loop.

Lemma 2. *MONITOR generates non-zero output for a given non-zero input.*

The proof proceeds by contradiction assuming a t with non-zero output, and the minimum output stream in $G^=$ that has a non-zero output, and then showing that there must be a non-zero output for $t - \epsilon$. This can be applied $\frac{t}{\epsilon}$ times to conclude that there is non-zero output before 0 which is a contradiction.

We finally show that the output of MONITOR is an evaluation model. We use $H_s^i(\sigma_I)$ for the history of events H_s after the i -th execution of the loop body, and $H_s^*(\sigma_I)$ for the sequence of events generated after a continuous execution of the monitor. Note that $H_s^*(\sigma_I)$ can be a finite sequence of events (if the input is bounded and no repetition is introduced in the specification using **delay**) or an infinite sequence of events. In the first case, the vote is eventually ∞ and the monitoring algorithm halts.

Theorem 2. *Let σ_I be an input event stream, and let σ_O consist of $\sigma_x = H_x^*(\sigma_I)$ for every output stream x . Then (σ_I, σ_O) is an evaluation model of φ .*

The proof proceeds by induction on the number of rounds in the loop, showing that the output is an evaluation model up-to the quiescence time. Putting together Theorem 2, Lemma 1 and Lemma 2 we obtain the following result.

Corollary 1. *Let φ be a well-formed specification, σ_I a non-zero input stream and H^* the result of MONITOR. Then, H^* is the only evaluation model for input σ_I , and H^* is non-zero.*

Trace Length Independent Monitoring The algorithm MONITOR shown above computes incrementally the only possible evaluation model for a given input stream, but this algorithm stores the whole prefix H_y for every output stream variable y . We show now a modification of the algorithm that is trace length independent, based on flattening the specification. A specification is *flat* if every occurrence of an offset expression in every T_y is either $x(<\sim t)$ or $x(<<t)$. In other words, there can be no nested term of the form $x(<\sim (y<\sim t))$ or $x(<\sim (y<<t))$ or $x(<<(y<\sim t))$ or $x(<<(y<<t))$. We first show that every specification can be transformed into a flat specification. The flattening applies incrementally the following steps to every nested term $x(E(y<<t))$, where E is an arbitrary offset term:

1. introduce a fresh stream s with equations $T_s = y.ticks$ and $V_s = x(E(t))$
2. replace every occurrence of $x(E(y<<t))$ by $s(<t)$.

Example 3. Consider the following specification of a continuous integration process in software engineering. The intended meaning is to report in `faulty` those commits to a repository that fail the unit tests.

```
input commit_id commits, unit push, bool tests
ticks faulty := tests.ticks
define commit_id faulty := if tests(~t) then notick
                                else commits(<push<<t)
```

After applying the flattening process the specification becomes:

```
define commit_id faulty := if tests(~t) then notick else s(<t)
ticks s := push.ticks
define commit_id s := commits(<t)
```

Here, `s` stores the `commit_id` of the last commit at the point of a `push`, which is precisely the information to report at the time of a `faulty` commit. \square

Lemma 3. *Let φ be a specification. There is an equivalent flat specification φ' that is linear in the size of φ .*

Now, let φ' be the flat specification obtained from φ and let y be an output stream variable. Consider the cases for offset sub-expressions in the computation of $\llbracket V_y \rrbracket_H(t)$ in line 9 of MONITOR:

- $s<\sim t$: the evaluation fetches the value H_s at time t (if s ticks at t) or at the previous ticking time (if s does not tick at t).
- $s<<t$: the evaluation fetches the value H_s at the previous ticking time of s .

In either case, only the last two elements of H_s are needed. The similar argument can be made to compute T_y because only the last event of s is needed for `delay(s)`. Hence, to evaluate MONITOR on flat specifications, the algorithm only needs to maintain the last two elements in the history for every output stream variable to compute the next value of every value and ticking equation.

Theorem 3. *Every flat specification φ can be monitored online with linear memory in the size of the specification and independently of the length of the trace. Moreover, every step can be computed in linear time on the size φ .*

4 Extensions and Comparison

We first sketch how to define the most complex operator⁴ of TeSSLa: $x = \text{delay}\langle s_0, s_1 \rangle$, which creates an event stream x which will tick at an instant t if there is an event (t', v) in s_0 such that $t' + v = t$ and also $\text{dom}(s_1) \cap (t', t) = \emptyset$. TeSSLa does not handle explicit time and offsets but builds specifications from building blocks like *delay*. Given inputs s_0 and s_1 the Striver specification is:

```

ticks aux := s0.ticks U s1.ticks
define Time_eps aux := if isticking(s1) then infty
    else if aux(<t, infty) = infty || aux(<t) + aux<<t <= t
        then s0(~t) else notick
ticks x := delay x_aux
define unit x := ()

```

We now present three extensions to the basic Striver introduced previously.

Accessing successors. The first extension allows accessing future events, via the dual of the offset operators $x >^{\sim} e$ and $x \gg e$, and the syntactic sugar to access the successor value $x(e>)$, $x(e^{\sim})$, $x(e, d>)$ and $x(e, d^{\sim})$. As for Lola, well-formedness can be guaranteed as long as all strongly connected components in the dependency graph contain only $\xrightarrow{-}$ and $\xrightarrow{0}$ edges, or only $\xrightarrow{+}$ and $\xrightarrow{0}$ edges, and additionally, there is no cycle with only $\xrightarrow{0}$ edges. For example, this guarantees that there is no cyclic dependency, as every stream either depends on itself in the future or in the past (or none at all).

All Delays. This allows defining tick sets that consider all delays. The ticking expressions are extended with an operator `delayall` with the following semantics:

$$\llbracket \text{delayall}(w) \rrbracket_{\sigma} \stackrel{\text{def}}{=} \{t' \mid \text{there is a } t \in \text{dom}(\sigma_w) \text{ such that } t + \sigma_w(t) = t'\}$$

This extension requires only to change VOTE to accommodate for a set of possible pending delays and not just a single delay. In general, this cannot be implemented in finite memory for arbitrary event rates and delays, but MONITOR works directly for the online monitoring this construct.

Windows. The last extension allows implementing computations over precise windows, like “*count the number of events in every window of one second*”. This cannot be described in TeSSLa [9], which is limited to finite memory monitors, or in RTLola [14] because this specification is not isochronous. Note that this property cannot be monitored by splitting the time in intervals of one second and counting the events in each of the intervals obtained (as in RTLola) as this approach misses the case of counting the events in part of one window and the remaining time in the adjacent window. The main idea of this extension is to

⁴ Due to limitations a full comparison against TeSSLa and STL is not presented here.

enrich time expressions with a tag, in such a way that every tick carries an additional value (we called this extension *dependent time*). Then, `delay` and `delayall` are enriched with the ability to use tagged time streams, with the caveat that the `U` combinator must now indicate how to combine tags. Consider the following example with input `int s`:

```
ticks wcount := (const 1 s) U delayall (const (-1,5) s)
define int wcount t aux := wcount(<t,0) + aux
```

The stream `wcount` must only be computed when a new event arrives in `s` (adding 1) or when an event leaves the window (subtracting 1), which is monitored with a constant number of operations per event, but requires storing a number of events that depends on the event rate.

The Signal Temporal Logic STL [21, 1]—when interpreted over piecewise-constant signals—is subsumed by `Striver`. First note that event streams have a dual interpretation as piece-wise constant signals, where the signal only changes at the point where events are produced. The translation to `Striver` opens the door to a quantitative computation of STL by enriching the data types of expressions and verdicts. We show the operator $x\mathcal{U}_{[0,b]}y$:

```
ticks v := x U y U delayall -b x U delayall -b y
define bool v t := if y(~t,false) then true else
  if !x(~t,false) then false else
  let t' := yT(t~) in
  if t'==outside || t' > t+b then false else t' <= xF(t~)
```

5 Empirical Evaluation

We report an empirical evaluation of a prototype sequential `Striver` implementation, written in the Go programming language⁵. We measure the memory usage and time per event for two collections of specifications. The first collection, from Example 1, computes the stocks of p independent products. These specifications contain a number of streams proportional to p , where each defining equation is of constant size. The second collection computes the average of the last k sales of a fixed product, via streams that tick at the selling instants and compute the sum of the last k sales (see the appendix for the concrete specs). The resulting specifications has depth proportional to k . We instantiate k and p from 10 to 500 and run each resulting specification with a set of generated input traces. We run the experiments on a virtual machine on top of an Intel Xeon at 3GHz with 32GB of RAM, and measure the average memory usage (using the OS) and the number of events processed per second.

In the first experiment, we run the synthesized monitors with traces of varying length (top two plots in Figure 1). The results illustrate that the memory needed to monitor each specification is independent of the length of the trace (the curves are roughly constant). Also, the ratio of events processed is independent of the length of the trace. In the second experiment, we fix a trace of

⁵ `Striver` is available at <http://github.com/imdea-software/striver>

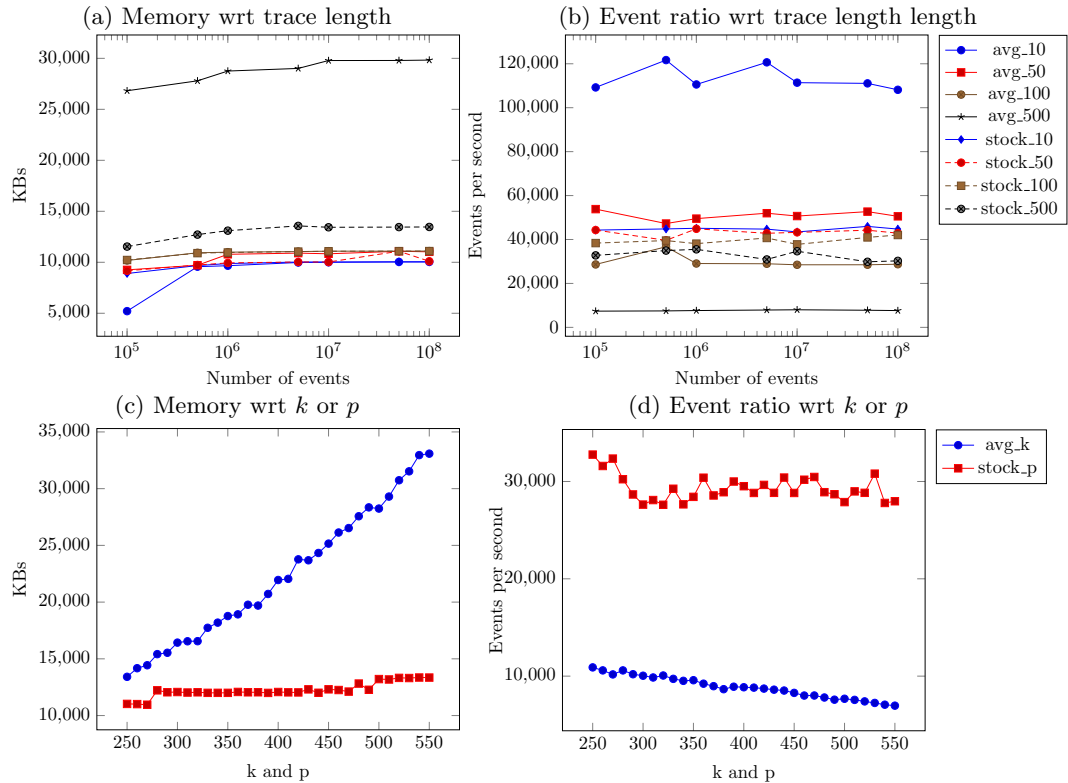


Fig. 1. Empirical evaluation

1 million events and run the specifications with k and p ranging from 250 to 550. The results (lower diagrams) indicate that the memory needed to monitor `stock_p` is independent of the number of products while the memory needed to monitor each `avg_k` specification grows linearly with k . Recall that theoretically all specifications can be monitored with memory linearly on the size of the specification.

6 Conclusion and Future Work

We have introduced `Striver`, a specification language with explicit time and offset reference for the stream runtime verification of timed event streams. We have presented a trace-length independent online monitoring algorithm for the efficiently monitorable fragment. Future work includes the extension of the language with parametrization, (like in `QEA` [4], `MFOTL` [6] and `Lola2.0` [13]), to dynamically instantiate monitors for observed data items. We are also studying offline evaluation algorithms, and algorithms that tolerate deviations in the time-stamps and asynchronous arrival of events from the different input streams.

References

1. *Lectures on Runtime Verification*, volume 10457 of *LNCS*, chapter Specification-Based Monitoring of Cyber-Physical Systems: A Survey on Theory, Tools and Applications, pages 135–175. Springer, 2018.
2. *Lectures on Runtime Verification - Introductory and Advanced Topics*, volume 10457 of *LNCS*. Springer, 2018.
3. Eugene Asarin, Paul Caspi, and Oded Maler. Timed regular expressions. *J. ACM*, 49(2):172–206, 2002.
4. Howard Barringer, Yliès Falcone, Klaus Havelund, Giles Reger, and David Rydeheard. Quantified event automata: Towards expressive and efficient runtime monitors. In *Proc of the 18th Int'l Symp. on Formal Methods (FM'12)*, volume 7436 of *LNCS*, pages 68–84. Springer, 2012.
5. Howard Barringer, Allen Goldberg, Klaus Havelund, and Koushik Sen. Rule-based runtime verification. In *Proc. of VMCAI'04*, LNCS 2937, pages 44–57. Springer, 2004.
6. David A. Basin, Felix Klaedtke, Samuel Mller, and Eugen Zalinescu. Monitoring metric first-order temporal properties. *Journal of the ACM*, 62(2), 2015.
7. Andreas Bauer, Martin Leucker, and Chrisitan Schallhart. Runtime verification for LTL and TLTL. *ACM T. Softw. Eng. Meth.*, 20(4):14, 2011.
8. Laura Bozelli and César Sánchez. Foundations of Boolean stream runtime verification. In *In Proc. RV'14*, volume 8734 of *LNCS*, pages 64–79. Springer, 2014.
9. Lukas Convent, Sebastian Hungerecker, Martin Leucker, Torben Scheffel, Malte Schmitz, and Daniel Thoma. TeSSLa: Temporal stream-based specification language. In *Proc. of the 21st. Brazilian Symp. on Formal Methods (SBMF'18)*, volume 11254 of *LNCS*. Springer, 2018.
10. Flaviu Cristian and Christof Fetzer. The timed asynchronous distributed system model. *IEEE Transactions on Parallel and Distributed Systems*, 10(6):642–657, 1999.
11. Ben D'Angelo, Sriram Sankaranarayanan, César Sánchez, Will Robinson, Bernd Finkbeiner, Henny B. Sipma, Sandeep Mehrotra, and Zohar Manna. LOLA: runtime monitoring of synchronous systems. In *Proc. of TIME'05*, pages 166–174. IEEE, 2005.
12. Cindy Eisner, Dana Fisman, John Havlicek, Yoad Lustig, Anthony McIsaac, and David Van Campenhout. Reasoning with temporal logic on truncated paths. In *Proc. of CAV'03*, volume 2725 of *LNCS 2725*, pages 27–39. Springer, 2003.
13. Peter Faymonville, Bernd Finkbeiner, Sebastian Schirmer, and Hazem Torfah. A stream-based specification language for network monitoring. In *Proc. of the 16th Int'l Conf. on Runtime Verification (RV'16)*, volume 10012 of *LNCS*, pages 152–168. Springer, 2016.
14. Peter Faymonville, Bernd Finkbeiner, Maximilian Schwenger, and Hazem Torfah. Real-time stream-based monitoring. *CoRR*, abs/1711.03829, 2017.
15. Alwyn E. Goodloe and Lee Pike. Monitoring distributed real-time systems: A survey and future directions. Technical report, NASA Langley Research Center, 2010.
16. Klaus Havelund and Allen Goldberg. Verify your runs. In *Proc. of VSTTE'05*, LNCS 4171, pages 374–383. Springer, 2005.
17. Klaus Havelund and Grigore Roşu. Synthesizing monitors for safety properties. In *Proc. of TACAS'02*, LNCS 2280, pages 342–356. Springer, 2002.

18. Martin Leucker. Teaching runtime verification. In *Proc. of RV'11*, number 7186 in LNCS, pages 34–48. Springer, 2011.
19. Martin Leucker, César Sánchez, Torben Scheffel, Malte Schmitz, and Alexander Schramm. TeSSLa: Runtime verification of non-synchronized real-time streams. In *Proc. of the 33rd Symposium on Applied Computing (SAC'18)*. ACM, 2018.
20. Martin Leucker and Christian Schallhart. A brief account of runtime verification. *J. Logic Algebr. Progr.*, 78(5):293–303, 2009.
21. Oded Maler and Dejan Nickovic. Monitoring temporal properties of continuous signals. In *Proc. of FORMATS/FTRTFT 2004*, volume 3253 of LNCS, pages 152–166. Springer, 2004.
22. Lee Pike, Alwyn Goodloe, Robin Morisset, and Sebastian Niller. Copilot: A hard real-time runtime monitor. In *Proc. of RV'10*, LNCS 6418. Springer, 2010.
23. Amir Pnueli and Aleksandr Zaks. PSL model checking and run-time verification via testers. In *Proc. of FM'06*, LNCS 4085, pages 573–586. Springer, 2006.
24. Grigore Roşu and Klaus Havelund. Rewriting-based techniques for runtime verification. *Automated Software Engineering*, 12(2):151–197, 2005.
25. Koushik Sen and Grigore Roşu. Generating optimal monitors for extended regular expressions. *ENTCS*, 89(2):226–245, 2003.